



Micromachined 2D Transducers and Phantoms for 3D Super-Resolution Ultrasound Imaging

Ommen, Martin Lind

Publication date:
2020

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Ommen, M. L. (2020). *Micromachined 2D Transducers and Phantoms for 3D Super-Resolution Ultrasound Imaging*. DTU Health Technology.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

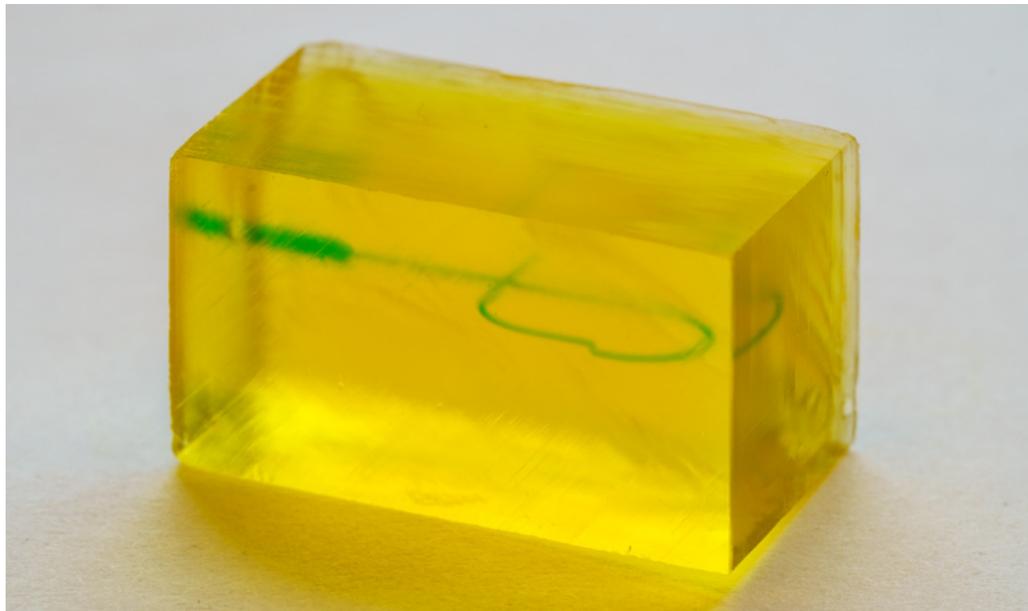
TECHNICAL UNIVERSITY OF DENMARK

PH.D. THESIS

MICROMACHINED 2D TRANSDUCERS AND
PHANTOMS FOR 3D SUPER-RESOLUTION
ULTRASOUND IMAGING

Author:
Martin Lind Ommen

Supervisors:
Prof. Erik V. Thomsen
Prof. Jørgen A. Jensen
Prof. Niels B. Larsen



14th March 2020

Kgs. Lyngby, Denmark

Cover image: Photograph of a 3D printed 200 μm diameter flow phantom. Water containing blue fruit dye was pumped through the channel to increase the contrast between the phantom polymer and the channel. The channel enters from the end, passes down through the phantom and loops around, passing itself at 90° with a separation of 108 μm . The phantom was designed to demonstrate 3D super-resolution ultrasound.

Technical University of Denmark
Department of Health Technology
Ørsteds Plads 345C
2800 Kgs. Lyngby
DENMARK
Author e-mail: marlio@dtu.dk

Preface

This PhD thesis has been submitted to the Department of Health Technology at the Technical University of Denmark in partial fulfillment of the requirements for acquiring the PhD degree. The research providing the foundation for the thesis has been conducted over a period of three years from March 15, 2017, to March 14, 2020. It has been carried out at the Department of Health Technology at the Technical University of Denmark (DTU) under supervision of Professor Erik Vilain Thomsen and co-supervised by Professor. Jørgen Arendt Jensen, and Professor Niels Bent Larsen.

The project combines the competences at each of the three research groups through CMUT development with Erik Vilain Thomsen, 3D printing of ultrasound phantoms with Niels Bent Larsen, and ultrasound experimentation with Jørgen Arendt Jensen.

Martin Lind Ommen
Kgs. Lyngby, March 2020

Abstract

Super-resolution ultrasound imaging (SRUS) is a new ultrasound imaging technique which aims to visualise the smallest branches of the vascular system, namely sub-100 μm arterioles and venules and 5-9 μm capillaries. The technique breaks the conventional diffraction limited resolution through super-localisation of micro-bubble contrast agents injected into the vascular system. The method has been demonstrated to obtain resolutions of only a few tens of micrometres which should be compared to 500 μm of conventional diffraction limited ultrasound system. The goal of this project has been to develop tools to improve the SRUS techniques and transfer them from 2D to 3D imaging through development of capacitive micro-machined ultrasonic transducer (CMUT) fabrication processes and 3D printed phantom fabrication for improved validation.

A continuous goal in ultrasound transducer fabrication is to create larger transducer arrays for increased field of views (FOVs), combined with larger operating frequencies for increased resolution. The increased size of arrays means that even the smallest sample contamination might ruin the few devices available. The fabrication process optimisation presented in the thesis is about fusion bonding. Fusion bonding conducted directly in hand without a wafer bonder has been shown to provide a wafer bond of comparable quality to fusion bonding performed in dedicated wafer bonders. Handbonding allows for forming the bond directly after cleaning the wafers, minimizing the risk of particle contamination, therefore improving the processing yield.

To properly develop the SRUS techniques, suitable phantom structures need to be made. A stereolithography (SLA) 3D printing solution for fabrication of ultrasound phantoms is presented in the thesis. Conventional phantom fabrication methods consist of tubes suspended in water which can be perfused by micro-bubble-containing liquids. However, these methods are incapable of providing feature control on the scale required for SRUS, and suffer from limited three-dimensional feature placement capabilities.

The printed structures are hydrogels, water-containing polymer networks, printed with a voxel size of $10.8 \mu\text{m} \times 10.8 \mu\text{m} \times 20 \mu\text{m}$. The acoustic and structural properties, as well as potential ways to manipulate them are presented. The phantoms have an average speed of sound of 1577 m/s and an average density of 1.045 g/ml. The printed phantoms swell approximately 2.6% post printing, making compensation of design features necessary when using the phantoms as reference structures.

A new type of phantom was developed based on printed cavities which have been shown to reflect sound. By keeping the cavities smaller than the imaging wavelength, they can be used as stable point targets for repeated imaging. Design optimisation of the scatterers has been conducted in terms of actual printed size and reflected intensity, modelling different sizes, shapes and local overexposure schemes. The “Single pixel” scatterers, printed with a single voxel wide local overexposure around each cavity, yielded the highest reflected intensity.

Calibration of a 3D SRUS pipeline imaged with a row-column addressed (RCA) probe using a scatterer phantom containing eight randomly placed scatterers showed high accuracy of the pipeline. The localisation precision was found to be smaller than 27.6 μm in all directions, which is less than 1/18th of the imaging wavelength used in the experiment. The high precision allowed for detection of distortion in the beamforming on a micrometre scale. This would not have been possible to discover using conventional tube phantom setups.

A series of flow phantoms were created to perform well controlled SRUS experiments with micro-bubbles. A fiducial marker grid layout was presented, which allow for easy alignment of the ultrasound probe to the phantom features. A flow phantom was created to demonstrate super-localisation of micro-bubbles in 3D using a RCA array. The localisation precision was estimated using the flow phantom, evaluated based on the micro-bubble distributions in the flow channel, with the estimates being in line with the precision estimates based on the scatterer phantom. Finally, flow phantoms for demonstrating true resolution of the SRUS pipelines were developed, utilizing the 3D fabrication freedom of the 3D printing technique.

The results illustrate the great obtainable achievements with a high resolution 3D printing phantom fabrication method, but only scratches the surface of the potential solutions that the phantom printing method provides. The printing method allows for three-dimensional freedom of design and an unparalleled control of phantom feature placement and feature size control.

Resumé (Danish)

Super-resolution ultralydsbilleddannelse (SRUS) er en ny ultralydsteknik hvor målet er at afbilde de mindste forgreninger af det vaskulære system, herunder sub-100 μm arterioler og venoler samt 5-9 μm kapillærer. Teknikken bryder den konventionelle diffraktionsbegrænsede opløsningsevne ved super-lokalisering af mikroboble kontrast agenter injiceret i det vaskulære system. Metoden har vist opløsningsevner på få snesevis af mikrometer, hvilket skal sammenlignes med 500 μm for konventionelle ultralydssystemer. Projektets mål har været at udvikle værktøjer til forbedring af SRUS teknikkerne samt at overføre dem fra 2D til 3D billeddannelse ved udvikling af capacitive mikrofabrikerede ultralydstransducer (CMUT) fabrikationsprocesser samt 3D printede fantomer for forbedret validering.

Et fortsat mål i ultralydstransducerfabrikation er at lave større transducerarrays for øget field of view (FOV), kombineret med højere operationsfrekvens for øget opløsningsevne. Den øgede størrelse af arraysne betyder at selv de mindste prøveforureninger vil kunne ødelægge de få enheder der produceres. Fabrikationsprocesoptimeringen der præsenteres i afhandlingen omhandler fusion bonding. Det vises at fusion bonding lavet direkte i hånden uden en wafer bonder resulterer i et bond af sammenlignelig kvalitet som havde den været lavet i en dedikeret wafer bonder. Håndbondning tillader at bondet kan laves direkte efter rens af waferne, hvilket mindsker risikoen for partikelkontaminering og dermed øger procesyieldet.

For ordentlig udvikling af SRUS teknikkerne er det nødvendigt at have egnede fantomstrukturer. En stereolitografi 3D printer løsning til fabrikation af ultralydsfantomer præsenteres i afhandlingen. Konventionelle fantom fabrikationsmetoder består af slanger nedsænket i vand, hvori væsker med mikrobobler kan skylles igennem. Disse metoder tilbyder imidlertid ikke kontrol over placering af strukturer på et tilstrækkeligt niveau i forhold til SRUS, og er begrænsede med hensyn til placering af strukturer i det tredimensionelle rum.

Printene kaldes hydrogeler og er vandholdige polymernetværk, printet med en voxelstørrelse på $10.8 \mu\text{m} \times 10.8 \mu\text{m} \times 20 \mu\text{m}$. De akustiske og strukturelle egenskaber og potentielle metoder de kan manipuleres præsenteres. Fantomerne har en gennemsnitlig lydshastighed på 1577 m/s og en gennemsnitlig densitet på 1.045 g/ml. De printede fantomer kvæller ca. 2.6% efter printning, hvilket nødvendiggør kompensering af designfeatures når fantomerne skal bruges som referencestrukturer.

En ny type fantom blev udviklet baseret på printede kaviteter hvilke reflekterer lyd. Ved at holde kaviteterne mindre end billeddannelsesbølgelængden fungerer de som stabile punktkilder der kan afbildes kontinuerligt. Designoptimering af kaviteterne blev foretaget med henblik på den faktisk printede størrelse og den reflekterede lydintensitets afhængighed af kavitetsstørrelse, form og det lokale eksponeringsmønster. "Single pixel" kaviteter, som printes med en enkelt pixel bred ramme af overeksponering, giver den største refleksion af kaviteterne.

Kalibrering af en 3D SRUS pipeline foretaget med et fantom bestående af otte kaviteter, afbildet med en row-column addressed (RCA) probe, demonstrerede høj nøjagtighed af pipelinen. Lokaliseringspræcisionen blev estimeret mindre end 27.6 μm i alle retninger, hvilket er mindre end 1/18 af den anvendte billeddannelsesbølgelængde. Den høje præcision muliggjorde detektering af forvrængning i beamformningen, hvilket ville være umuligt med konventionelle slangefantomer.

En serie af flowfantomer blev lavet til gennemførsel af velkontrollerede SRUS eksperimenter med mikrobobler. Et fiducial marker design mønster præsenteres som tillader let alignment mellem ultralydsprobe og fantomstrukturer. Et flowfantom blev lavet til demonstration af superlokalisering af mikrobobler i 3D ved hjælp af et RCA array. Lokaliseringspræcisionen blev estimeret ved brug af flowfantomet ud fra mikrobobledistributionen i flowkanalen, med estimerer i overensstemmelse med de foregående estimerer baseret på kavitetsfantomet. Slutteligt præsenteres udviklingen af flowfantomer til demonstration af reel opløsningsevne af SRUS pipelines der udnytter 3D fabriktionsfriheden af 3D printmetoden.

Resultaterne viser det store potentiale med fabrikation af højopløsnings 3D print fantomer, men ridser kun lige overfladen af alle de potentielle muligheder som fabrikationsmetoden tilbyder. Printmetoden tillader tredimensionel designfrihed og en kontrol over størrelse og placering af strukturer uden sidestykke.

Acknowledgements

During the last three years, I have had the pleasure of being able to do research within a very interesting field full of lots of possibilities. I have been able to undertake a lot of exciting projects, and I am very much aware that I have a lot of people to thank for getting these opportunities and achieving the results.

First of all, my three supervisors Niels Bent Larsen, Jørgen Arendt Jensen, and Erik Vilain Thomsen. There is no doubt that the project in its entirety has only been possible because of the combined different research interest of you. You have all been very involved and interested in my progress throughout the last three years. I have no doubt that the combined efforts, and at times discussions and disagreements between all of us, has pushed the project further than any of us would have been able to do by ourselves. I am very thankful that you have always been interested in taking the time to assist me whenever requested.

I want to give a special thank you to my main supervisor Erik Vilain Thomsen. It is wonderful to work in a research group in which there is a clear atmosphere of trust, faith and freedom to seek the solutions that I choose for myself, while always having the opportunity to get support and assistance when I might be at a loss. I am very thankful that this thesis does not mark the end of that, but that I am able to continue working with you.

Aside from my supervisors, I have a lot of wonderful co-workers with whom I have interacted with on a daily basis. These are primarily the PhDs and Post Docs working in the groups of my supervisors. Together with you, I have been able to utilise my results in broader and new perspectives. You have provided a wonderful work environment, both in terms of professional sparring, but certainly for so much more as well. I am thankful for having had the opportunity to collaborate with you all.

There are three colleagues that I want to mention explicitly.

First of all, Rujing Zhang, who originally developed the 3D printer as well as the printer processes. It is unlikely to find another person with as much knowledge about any system and all of the smaller quirks. Your constant willingness to contribute, be that intellectually towards new solutions, or practically, assisting in unfamiliar printing processes has been wonderful.

Second, Mikkel Schou, who remains the central figure at CFU for development of imaging sequences and ultrasound experimentation. Your assistance in discussions of experimental and phantom designs has provided another perspective and practical experience that I could not have hoped to be able to benefit from.

Lastly, Andreas Havreland, who has been my office mate for the last three years. Your incredible experience, knowledge, and ease with which you are able to gain an overview of new techniques and research fields is admirable and a privilege to have been allowed to exploit over the last three years.

Finally, and certainly not the least, my daughter, Maggie and my wife, Ditte. Observing a small person in a lab coat, during the times when I have brought her to assist in the lab, full of amazement over the “yellow weird things my dad is making” is indescribable. The enthusiasm has changed from “well, my dad just makes some really boring yellow things” to “please, can I come with you to the lab again?” (thankfully in that order), and it has been an absolute pleasure to show her a tiny bit of what is without a doubt a very weird world to her as she has grown from 2 to 5 years old in the span of this project. There is no question that my project and work over the last three years would not have been as successful, or as enjoyable, without the love and support of my wife. It is too easy to forget your indulgence during the stressful periods of the project and your constant great belief in me, both of which are invaluable to me.

Sincerely,

Martin Lind Ommen

Publications	ix
Presentations	xii
List of abbreviations	xv
Overall Introduction	3
1 Thesis content	3
1.1 Central topic of this thesis	3
1.1.1 Working hypotheses for this Ph.D. project	3
1.2 Motivation	4
1.2.1 Medical imaging techniques	5
1.3 Ultrasound as a research field	6
1.3.1 Typical ultrasound research workflow	8
1.4 Thesis outline	10
2 Ultrasound	13
2.1 Basic ultrasound physics	13
2.2 Medical ultrasound	17
2.2.1 Super-resolution ultrasound imaging (SRUS)	17
2.2.2 Theoretical microfluidics for simple geometries	20
2.3 2D imaging versus 3D imaging	22
2.3.1 Row-column addressed arrays	25
2.4 Ultrasound phantoms	26
2.4.1 Phantoms for super-resolution ultrasound imaging	27
2.4.2 3D printing a new type of ultrasound phantom	28
3 Ultrasound transducers	31
3.1 Conventional ultrasound transducers	31
3.2 Capacitive micro-machined ultrasonic transducers (CMUTs)	32
3.2.1 Basic CMUT physics	33
3.2.2 Conventional CMUT fabrication methods	34

I	CMUT process optimisation	39
4	Hand-bonded CMUTs	41
4.1	Motivation	41
4.1.1	Fusion bonding	42
4.2	Materials and methods	44
4.2.1	Experimental design	44
4.2.2	Material choice - silicon nitride plates	44
4.2.3	Fabrication of test devices	50
4.3	Results	52
4.3.1	Hypothesis for bond interface diffusion	54
4.3.2	Deflection test with a silicon plate	56
4.3.3	Bond interface leak rate test	56
4.4	Chapter summary	58
II	3D printed phantoms	59
5	Introduction to 3D printing of phantoms	61
5.1	3D printing overview	61
5.1.1	Stereolithography (SLA)	63
5.2	Custom built 3D printing system	64
5.2.1	The 3D printer	65
5.2.2	Resin composition	65
5.3	Chapter summary	69
6	Hydrogel material characterisation	71
6.1	Hydrogel structural properties	71
6.2	Hydrogel swelling	72
6.2.1	Swelling uniformity	72
6.2.2	Swelling at different printing doses	78
6.3	Hydrogel density	80
6.4	Acoustic characterisation	81
6.4.1	Speed of sound	81
6.4.2	Sound attenuation	83
6.4.3	Acoustic impedance	86
6.5	Discussion	88
6.6	Chapter summary	89
7	Calibration phantoms for SRUS	91
7.1	A new type of phantom for SRUS validation	91
7.2	Micro-engineering of the 3D printed scatterers	92
7.2.1	Concept description	92
7.2.2	Experimental setup	98
7.2.3	Printed scatterer size	99
7.2.4	Dose manipulation for increased scattering intensity	106
7.2.5	Scatterer separation distance	112
7.3	Cavity scatterer micro-phantoms for validation of SRUS in 3D	112
7.3.1	Methods	114
7.3.2	Results	116
7.4	Discussion	122
7.5	Chapter summary	125

8	Flow phantoms for SRUS	127
8.1	General flow phantom considerations	127
8.2	Flow phantom for 2D SRUS	129
8.2.1	Phantom description	130
8.2.2	2D SRUS results	130
8.3	MATLAB phantom generation	133
8.4	Single channel phantom for 2D and 3D super-localisation	136
8.4.1	Phantom description	136
8.4.2	Fiducial marker layout	137
8.4.3	3D super-localisation results	138
8.5	Looping flow phantom for 3D SRUS	141
8.5.1	Phantom description	141
8.5.2	Looping flow phantom results	141
8.6	Flow phantom optimisation	145
8.6.1	New fiducial marker layout	145
8.6.2	Decreasing the flow channel size	145
8.7	Phantom concepts for future exploration	148
8.7.1	Optimisation of channel separation	148
8.7.2	Different flow velocities along different axes in a single phantom	149
8.7.3	Branching channel systems to quantify local print variability	149
8.8	Chapter summary	150
III	Overall conclusion and outlook	151
9	Conclusion	153
9.1	Outlook	155
	Bibliography	157
IV	Appendix	167
A	Published papers	169
A.1	Paper A - BCB polymer based row-column addressed CMUT	169
A.2	Paper B - 3D Printed Flow Phantoms With Fiducial Markers for Super-Resolution Ultrasound Imaging	174
A.3	Paper C - Wafer Level Characterization of Row-Column Addressed CMUT Arrays	179
A.4	Paper D - Ultrasound Multiple Point Target Detection and Localization using Deep Learning	184
A.5	Paper E - History and Latest Advances in Flow Estimation Technology: From 1-D in 2-D to 3-D in 4-D	189
A.6	Paper F - 3-D Super Resolution Imaging using a 62+62 Elements Row-Column Array	200
A.7	Paper G - 3D Printed Calibration Micro-phantoms for Validation of Super-Resolution Ultrasound Imaging	205
A.8	Paper H - Three-Dimensional Super Resolution Imaging using a Row-Column Array	210
B	Papers under review	223
B.1	Paper I - 3D Printed Calibration Micro-Phantoms for Validation of Super-Resolution Ultrasound Imaging	223
B.2	Paper J - Detection and Localization of Ultrasound Scatterers Using Convolutional Neural Networks	235

C	Papers in preparation	249
C.1	Paper K - Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary?	249
D	Presented posters	257
D.1	Poster 1 - 3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging	257
D.2	Poster 2 - Reduced cavity pressure in fusion bonded devices: is a wafer bonder necessary? (no)	259
E	Statistical modelling	261
E.1	Optical characterisation of scatterer sizes	261
E.2	Scattering strength	262
F	Process optimisation and analysis scripts	269
F.1	Thin film thermal processing time	269
F.2	Film thickness map	270
F.3	Characterisations of furnaces	274
G	Phantom generation scripts	277
G.1	Flow channel scripts	277
G.1.1	Cylindrical channels	277
G.1.2	Square channels	283
H	Experimental setups	287
H.1	Hand-bonding recess	287
H.2	Pressure chamber	287
H.3	3D printed holder for optical measurements of swelling	287
I	Additional phantom designs	291
I.1	Scatterer phantom for neural network testing	291
J	3D printing of hydrogels	293
J.1	True PEGDA concentration	293
J.2	Optical microscope images of scatterers	293
J.3	Additional process optimisation and printing issues	293
J.3.1	Stress-induced bending of hydrogel samples	293
J.3.2	Ghost image	294
J.3.3	Minimum scatterer separation	294
J.3.4	Scratches in the film	294

List of publications

Published papers:

Paper A - BCB polymer based row-column addressed CMUT

*Andreas Spandet Havreland, **Martin Lind Ommen**, Chantal Silvestre, Mathias Engholm, Jørgen Arendt Jensen and Erik Vilain Thomsen*

Proceedings of 2017 IEEE International Ultrasonics Symposium (IUS)

Paper B - 3D Printed Flow Phantoms With Fiducial Markers for Super-Resolution Ultrasound Imaging

***Martin Lind Ommen**, Mikkel Schou, Rujing Zhang, Carlos Armando Villagómez Hoyos, Jørgen Arendt Jensen, Niels Bent Larsen and Erik Vilain Thomsen*

Proceedings of 2018 IEEE International Ultrasonics Symposium (IUS)

Paper C - Wafer Level Characterization of Row-Column Addressed CMUT Arrays

*Erik Vilain Thomsen, Kitty Steenberg, Magnus Galsgård Petersen, Mads Weile, Andreas Havreland, **Martin Lind Ommen**, Rune Sixten Grass, and Mathias Engholm*

Proceedings of 2019 IEEE International Ultrasonics Symposium (IUS)

Paper D - Ultrasound Multiple Point Target Detection and Localization using Deep Learning

*Jihwan Youn, **Martin Lind Ommen**, Matthias Bo Stuart, Erik Vilain Thomsen, Niels Bent Larsen, Jørgen Arendt Jensen*

Proceedings of 2019 IEEE International Ultrasonics Symposium (IUS)

Paper E - History and Latest Advances in Flow Estimation Technology: From 1-D in 2-D to 3-D in 4-D

*Jørgen Arendt Jensen, Svetoslav Ivanov Nikolov, Kristoffer Lindskov Hansen, Matthias Bo Stuart, Carlos Armando Villagomez Hoyos, Mikkel Schou, **Martin Lind Ommen** Sigrid Husebø Øygaard, Lasse Thurmann Jørgensen and Marie Sand Traberg*

Proceedings of 2019 IEEE International Ultrasonics Symposium (IUS)

Paper F - 3-D Super Resolution Imaging using a 62+62 Elements Row-Column Array

*Jørgen Arendt Jensen, Mikkel Schou, **Martin Lind Ommen**, Sigrid Husebø Øygaard, Thomas Sams, Matthias Bo Stuart, Erik Vilain Thomsen, Niels Bent Larsen, Christopher Beers, and Borislav Gueorguiev Tomov*

Proceedings of 2019 IEEE International Ultrasonics Symposium (IUS)

Paper G - 3D Printed Calibration Micro-phantoms for Validation of Super-Resolution Ultrasound Imaging

Martin LindOmmen, Mikkel Schou, Christopher Beers, Jørgen Arendt Jensen, Niels Bent Larsen, and Erik Vilain Thomsen,

Proceedings of 2019 IEEE International Ultrasonics Symposium (IUS)

Paper H - Three-Dimensional Super Resolution Imaging using a Row-Column Array

Jørgen Arendt Jensen, Martin Lind Ommen, Sigrid Husebø Øygaard, Mikkel Schou, Thomas Sams, Matthias Bo Stuart, Christopher Beers, Erik Vilain Thomsen, Niels Bent Larsen, and Borislav Gueorguiev Tomov

IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, Volume 67, Issue 3, page 583-546

Papers under review:

Paper I - 3D Printed Calibration Micro-Phantoms for Validation of Super-Resolution Ultrasound Imaging

Martin Lind Ommen, Mikkel Schou, Christopher Beers, Jørgen Arendt Jensen, Niels Bent Larsen, and Erik Vilain Thomsen

IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control

Paper J - Detection and Localization of Ultrasound Scatterers Using Convolutional Neural Networks

Jihwan Youn, Martin Lind Ommen, Matthias Bo Stuart, Erik Vilain Thomsen, Niels Bent Larsen, Jørgen Arendt Jensen

IEEE Transactions on Medical Imaging

Papers in preparation:

Paper K - Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary?

Martin Lind Ommen and Erik Vilain Thomsen

Contributions to publications

Published papers

Paper A - BCB polymer based row-column addressed CMUT

Contributed with the BCB process development to increase the functional properties of BCB. Participated in design, and planning and design of CMUTs, and participated in internal review of the article.

Paper B - 3D Printed Flow Phantoms With Fiducial Markers for Super-Resolution Ultrasound Imaging

Main author of the article. Conducted development of fiducial marker concept and phantoms for optical characterisation, ultrasound experimentation using the phantoms and analysis of the

results.

Paper C - Wafer Level Characterization of Row-Column Addressed CMUT Arrays

Contributed with development of the measurement schemes for RCA arrays. Conducted initial experiments for the concepts and participated in internal review of the article.

Paper D - Ultrasound Multiple Point Target Detection and Localization using Deep Learning

Contributed with design development of phantoms for the experiments in collaboration with the main author. Assisted on the experiments to ensure the correct phantom behaviour and the correct use of the phantoms. Contributed with writing of the sections regarding the phantoms as well as internal review of the article.

Paper E - History and Latest Advances in Flow Estimation Technology: From 1-D in 2-D to 3-D in 4-D

Contributed with micro-phantom design and concepts included in the article. Participated in the experiments leading to the presented phantom based results. Participated in internal review of the article.

Paper F - 3-D Super Resolution Imaging using a 62+62 Elements Row-Column Array

Contributed with micro-phantom design and concepts included in the article. Participated in the presented experiments in the article, and analysis of some of the results. Contributed with writing of the sections regarding the phantoms as well as internal review of the article.

Paper G - 3D Printed Calibration Micro-phantoms for Validation of Super-Resolution Ultrasound Imaging

Main author of the article. Conducted development of the scatterer phantom designs and concept presented, ultrasound experimentation and analysis of the results.

Paper H - Three-Dimensional Super Resolution Imaging using a Row-Column Array

Contributed with micro-phantom design and concepts included in the article. Participated in the presented experiments in the article, and analysis of some of the results. Contributed with writing of the sections regarding the phantoms as well as internal review of the article.

Papers under review

Paper I - 3D Printed Calibration Micro-Phantoms for Validation of Super-Resolution Ultrasound Imaging

Main author of the article. Conducted development of the scatterer phantom designs and concept presented for optical and ultrasound characterisation, all experimentation and analysis of the results.

Paper J - Detection and Localization of Ultrasound Scatterers Using Convolutional Neural Networks

Contributed with design development of phantoms for the experiments in collaboration with the main author. Assisted on the experiments to ensure the correct phantom behaviour and the correct use of the phantoms. Contributed with writing of the sections regarding the phantoms as

well as internal review of the article.

Papers in preparation

Paper K - Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary?

Main author of the article. Conducted device and experimental design. Conducted all experimentation, including the development of a pressure chamber for leak rate analysis.

List of presentations

Oral presentations:

MUT:

MUT 2017 - Process optimisation of BCB-polymer for use in CMUTs

MUT 2018 - Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary?

IUS:

IUS 2019 - 3D Printed Calibration Micro-phantoms for Validation of Super-Resolution Ultrasound Imaging

Poster presentations:

MNE:

MNE 2018 - Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary? (No)

IUS:

IUS 2018 - 3D Printed Flow Phantoms With Fiducial Markers for Super-Resolution Ultrasound Imaging

List of abbreviations

.png portable network graphics

ABS acrylonitrile butadiene styrene

AG agarose

AM amplitude modulation

AM additive manufacturing

ASIC application-specific integrated circuit

Au gold

BCB benzocyclobutene

BHF buffered hydrofluoric acid

BIC bayesian information criterion

C carbon

CAD computer aided design

CCD charge coupled device

CEUS contrast enhanced ultrasound

CMUT capacitive micro-machined ultrasonic transducer

Cr chromium

CT x-ray computed tomography

DED direct energy deposition

DKD diabetic kidney disease

DMD digital micromirror device

F fluorine

- FBAR** thin-film bulk acoustic resonator
- FDA** Food and Drug Administration
- FDM** fused deposition modelling
- FOV** field of view
- FPALM** fluorescence photoactivated localisation microscopy
- FPM** fully populated matrix
- FWHM** full width at half maximum
- H** hydrogen
- HEPA** high efficiency particulate arrestance
- ICR** inter quartile range
- IUS** international ultrasonics symposium
- LAP** lithium phenyl-2,4,6-trimethylbenzoylphosphinate
- LED** light emitting diode
- LOCOS** local oxidation of silicon
- LOM** laminated object manufacturing
- LPCVD** low pressure chemical vapour deposition
- MEMS** micro electro-mechanical system
- MRI** magnetic resonance imaging
- MSE** mean square error
- MW** molecular weight
- NDE** non-destructive evaluation
- O** oxygen
- PAA** polyacrylamide
- PALM** photoactivated localisation microscopy
- Pb** lead
- PDMS** polydimethylsiloxane
- PEGDA** poly(ethylene glycol) diacrylate
- PET** positron emission tomography
- PI** pulse inversion
- PL** photoluminescence
- PLA** polylactic acid

- PMMA** poly(methyl methacrylate)
- PSF** point spread function
- PZT** piezoelectric ceramic
- Q-Q** quantile-quantile
- QY** quinoline yellow
- RCA** row-column addressed
- ROI** region of interest
- SA** synthetic aperture
- SARUS** synthetic aperture real-time ultrasound system
- SAW** surface acoustic wave
- SFF** solid freeform fabrication
- Si** silicon
- Si₃N₄** silicon nitride
- SiO₂** silicon dioxide
- SLA** stereolithography
- SLA** stereolithography apparatus
- SLM** selective laser melting
- SLS** selective laser sintering
- SNR** signal to noise level
- SOI** silicon on insulator
- SPM** sparsely populated matrix
- SRUS** super-resolution ultrasound imaging
- STORM** stochastic optical reconstruction microscopy
- SVD** singular value decomposition
- TEM** transmission electron microscopy
- TGC** time-gain-compensation
- TOBE** top orthogonal to bottom electrode
- ULM** ultrasound localisation microscopy

Overall introduction

1.1 Central topic of this thesis

Super-resolution ultrasound imaging (SRUS) has emerged during the last decade, and has been shown to be able to visualise micro-vascular structures with unprecedented resolution for non-invasive medical imaging modalities, being able to resolve structures only separated by a few tens of micrometres. An imaging resolution that high will allow physicians to follow the development of the vascular network on a microscopic scale, to evaluate the state of vascularisation, or monitor the vascular response to treatment methods. A number of new challenges follow with the technique which offers increased resolution. So far, most SRUS has been conducted using 2D imaging probes, i.e. vertical image “sheets”, since 3D volume imaging probes are not yet widely available, while the structures to be imaged are inherently three-dimensional objects. Transducer development will be tailored towards 3D imaging with higher resolution performance than before, and in doing so move into unexplored fabrication parameter spaces, likely with specific challenges to follow. Furthermore, the obtained increase in imaging resolution also places new demands for imaging test phantoms for calibration and validation of the techniques, with currently used phantom fabrication techniques not providing a feature resolution on par with the imaging technology nor the three dimensional structure freedom to match the desired transducer imaging capabilities.

1.1.1 Working hypotheses for this Ph.D. project

The underlying goal of this Ph.D. project has been to develop tools to expand and improve the SRUS techniques in the ultrasound research field, in part by transferring them from 2D to 3D imaging. The development should push the capabilities of the techniques to achieve greater imaging resolutions than had been possible before, and to reveal the intricate three-dimensional details of the vascular networks. The fundamental understanding has been that this can be attained through software and hardware improvements and developments, ideally with synergetically improved outcomes. This project has been conducted with two central hypotheses addressing current shortcomings in imaging hardware and imaging validation.

The first is that the development of the SRUS techniques will benefit from improved transducer performance, tailored specifically towards the SRUS techniques.

The second is that the developed ultrasound techniques should be tested in controlled settings, for which no ideal test phantoms existed. 3D printing of hydrogel phantoms will allow for unprecedented accuracy and precision, which in turn will allow for validation of imaging techniques with higher accuracy and precision.

The first thesis half concerning the development of ultrasound transducers assumes a firm understanding of silicon micro-fabrication techniques. This involves knowledge of common micro-fabrication techniques such as UV-lithography, thin film deposition and growth, wet chemical etching, as well as dry plasma etching.

The second thesis half concerning the development of ultrasound phantoms assumes a solid understanding of ultrasound imaging, in particular typical validation methods of the imaging techniques, as well as an understanding of 3D printing techniques. The underlying physical principles of ultrasound is essential to understand and select the important properties for ultrasound phantom materials, as well as the phantom design techniques.

1.2 Motivation

The living organism is a complex combination of organ-, mechanical- and neural systems, all working in (more or less) unison in the human being. Central for all of these systems is the dependence on the vascular system consisting of billions of vessels distributed throughout the body, facilitating the transport of vital nutrients, hormones and gasses beyond the possible distribution lengths determined by nutrient diffusion [1]. The diameter of the vessels range from a few centimetres in the aorta, down to the smallest vessels in the human body, i.e. arterioles and venules with sub-100 μm dimensions and capillaries of 5-9 μm diameters [2]. The smallest capillaries are on average approximately 25% smaller in diameter than the red blood cells, the oxygen carrying erythrocytes shaped as bi-concave discs, forcing them to bend to pass the capillaries. It has been hypothesised that this enhances the transfer of oxygen to the tissue [3].

Living tissue continuously adapts to changes in external stresses, or internal requirements. As the tissue adapts, the local nutrient requirements will likely change as well. The fundamental example is following mitosis of cells, where the increased number of cells will require an increasing amount of nutrients. This of course happens during normal growth, where the vascular system develops and expands in unison with the growth of the tissue to match the increased requirements.

The formation of new vessels from pre-existing vessels is called angiogenesis, and is a normal and vital process. It happens throughout our lives as we grow, but also as the needs of the body changes for instance at particular stages of life, with pregnancy being an example. The properly functioning vascular system is crucial for the well-being, living organism, and has conversely also been shown to be directly linked to certain disease processes in cancer, diabetes, Alzheimer's and Parkinson's disease [4, 5]. Although angiogenesis is a common process in the body, it may indicate issues in the body, for instance if the proliferation of cells continues uncontrollably, such as seen in cancer development [6]. But angiogenesis happens regardless of the proliferation being malignant or benign [7].

The correlation between cancer and vascularisation can be used as a tool in the treatment of cancer patients. Common cancer treatment plans rely on administration of a medical drug, after which it is in some cases necessary to wait for multiple months before it is possible to observe whether the chosen treatment method has had an effect. Only at this point is it possible to switch to a different treatment plan. Unfortunately, the treatments themselves are often straining on the patients, which combined with the time required for each treatment method limits how many different methods can be applied. Fundamentally, the reason for the long treatment blocks is that the rate of tumour size decrease is fairly small, therefore making it difficult to observe changes. However, a change in the tumour growth will also influence the local capillaries. Thus, if this change could be detected on the micrometre scale, much quicker intervention might become possible.

Micro-vascular changes in the kidneys are also associated with diabetes, in diabetic kidney disease (DKD). One in three diabetic patients experience DKD, and there is unfortunately no cure. Once DKD becomes apparent in the clinic, the renal damage is already quite severe, often with poor prognosis for the patients and high medical costs [8]. If instead the diagnosis could be made based on changes to the micro-vasculature, treatments could be administered earlier, likely improving the prognosis.

1.2.1 Medical imaging techniques

Many different imaging techniques have been developed for non-invasive visualization of the internal structures of the body. Some of the commonly known methods are X-ray, x-ray computed tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET) and ultrasound. These techniques each have their own benefits and intended applications, resulting in most hospitals requiring equipment for all of these imaging modalities.

X-rays are high energy electromagnetic radiation with wavelengths approximately ranging from 0.01 nm to 10 nm. The radiation attenuation depends on the propagation media, and are for air and most tissue types quite similar, but significantly larger for bone, lead (Pb) and other metals. By placing a source between an x-ray source and a plate or digital sensor capable of absorbing the remaining radiation, shadow images of the object can be formed, showing differences in absorption. The imaging method is very fast and is ideal for imaging the bone structure, or fractures. X-ray setups are often quite large and therefore typically stationary pieces of equipment. Furthermore, the high energy photons also ionise atoms in the body, which may over time lead to cancer, which is why physicians leave the room, to areas shielded from stray radiation when the images are acquired. X-ray images are necessarily projection images of the object. CT is a modification in which a series of x-ray line images are taken and processed to form a cross-sectional image instead. CT images are made in large pieces of equipment, by placing the object in the centre of a tube, and having the x-ray source and sensor array rotate around the object within the tube wall. This creates a series of projection images from many different angles, from which the cross-sectional structure can be reverse engineered. 3D volumes can be constructed of the data by translating the object through the CT machine, thereby translating the imaged slice through the object. The method is more sensitive to different types of tissue, and can be used to image internal organs as well, and not only the bone structure. Since it uses x-rays, it also ionises atoms in the body, and the equipment is also very large.

MRI utilizes the magnetic spin of protons in the atoms of the body. A large magnetic field is applied to the imaged object, aligning the net spin of the protons in the body. Primarily the water molecules of the body which are affected. A smaller magnetic field is applied, oscillating at the resonance frequency of the spins, thereby making significant perturbations to the spin orientation. Depending on the tissue type, re-alignment to the main magnetic field once the perturbing magnetic field is removed takes a varying amount of time. Thereby, the different tissue types can be identified, and imaged. MRI also creates cross-sectional image slices. However, as opposed to CT, the orientation of the plane can be adjusted by application of gradients to the perturbing magnetic fields, allowing for full freedom of slice selecting. The slice can also be translated through the object to create 3D data. Functional imaging can be done in the brain, since blood oxygenation in the brain is closely linked to neural activity [9], resulting in a higher blood flow. There is no ionizing radiation making the imaging method safe. However, the equipment is large, very loud, and data acquisition takes a very long time, requiring the patient to stay still for extended periods, often upwards of an hour.

PET is a nuclear medicine functional imaging modality, used to observe metabolic processes in the body. Biologically active molecules bound with radioactive tracers are injected into the patients blood stream. The active molecule is chosen to highlight the biological function of interest. Once the radioactive particle decays, it will emit a pair of gamma rays. When these are detected, the spatial point of decay can be determined, and an image can be formed of the positions of the radioactive particles. These are located in three dimensions to create 3D volume data. A wide range of molecules can be chosen to image only selected biological functions. The gamma rays are however also ionising.

Ultrasound will be discussed in more detail in the next chapter, but in short, it is a non-ionising imaging method using non-audible high frequency sound waves, which for imaging purposes are harmless. Images are constructed by measuring the time from sound is transmitted until it is reflected from an object and is detected by the ultrasound probe again. Advanced transducer designs allow for 3D imaging as well. Sound attenuates significantly in tissue, thereby limiting the depth of imaging. Contrast agents can be used to allow for functional imaging. The ultrasound

Table 1.1: Comparison of the properties of common medical imaging modalities. *Ultrasound resolution is greatly improved through the super-resolution ultrasound imaging technique.

	X-ray	CT	MRI	PET	Ultrasound
Image resolution	≈ 0.5 mm	≈ 0.5 mm	≈ 1 mm	≈ 1 mm	≈ 0.5 mm*
Image orientation	Projection	Perpendicular cross-section	Free orientation	Perpendicular cross-section	Vertical slice
Field of view	Large	Large	Large	Large	≈ 20 cm deep
Acquisition time	Short	Medium	Long	Long	Short
2D/3D imaging	2D	2D/3D	2D/3D	3D	2D+3D
Functional imaging	No	Yes	Yes	Yes	Yes
Ionising radiation	Yes	Yes	No	Yes	No
Mobile equipment	No	No	No	No	Yes
Real time imaging	No	No	No	No	Yes

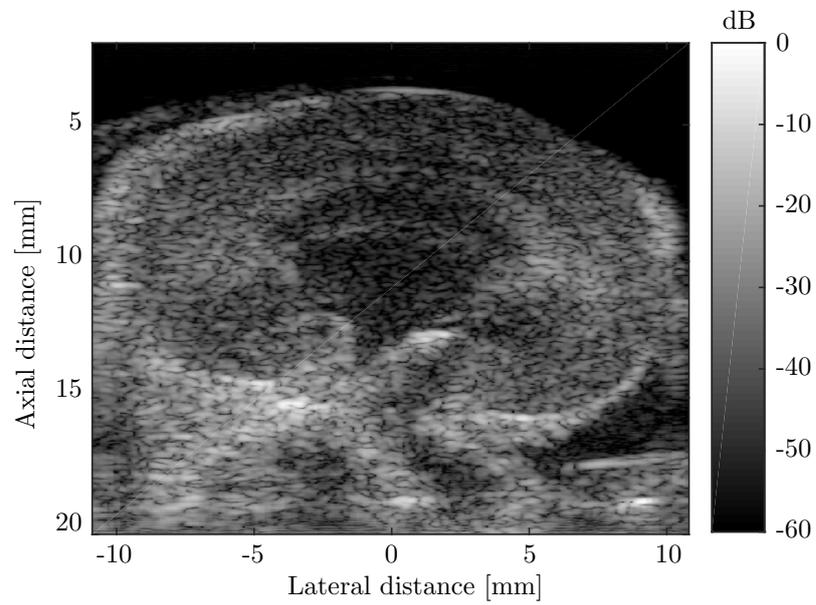
scanning equipment is in most cases fairly small, often made mobile to allow for bedside scanning of the patients. As opposed to all of the other imaging methods, ultrasound creates images in real time, with the potential of hundreds of images per second. A comparison of the imaging modalities can be seen in Table 1.1.

Returning to the imaging of micro-vasculature, most often the micro-vascular parameters are investigated on stained biopsy samples using optical microscopy, due to the small size of the vessels. The reason for this is that high resolution in imaging typically comes at the expense of field of view. High spatial resolution optical methods exist such as two-photon imaging [10] or optical coherence tomography [11], which have very limited field of view (FOV) and penetration depth. The extraction of tissue samples also introduces the issue that the sample preparation method might influence the structure of the sample. It will often be relevant to investigate the full context of the affected organ, requiring much larger FOVs, using non-invasive imaging methods. Perfusion imaging versions do exist for the previously mentioned modalities of MRI [12], CT [13], PET [14], and ultrasound [15], but the resolution of all of these systems are limited to millimetre sized features, and will not be able to resolve the minute changes of the micro-vasculature.

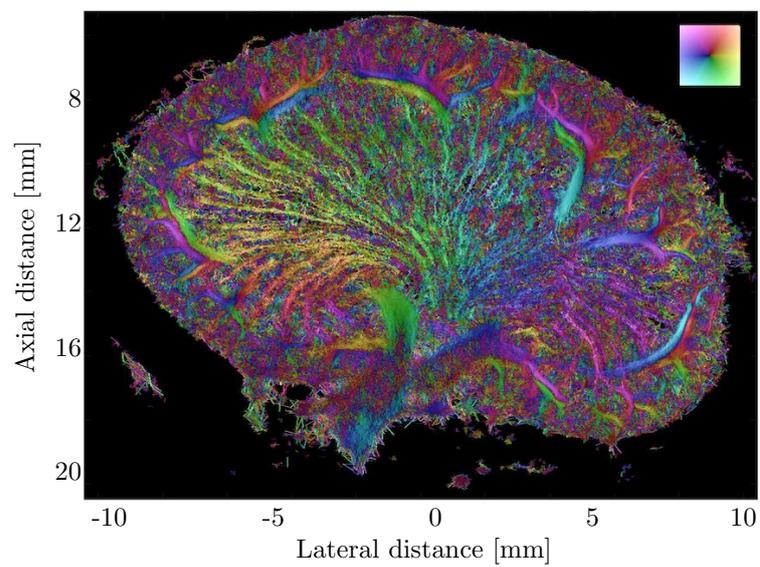
However, recently a development in the ultrasound field has emerged, called super-resolution ultrasound imaging (SRUS), that aims to break the resolution limit set by the imaging wavelength, specifically to create detailed images of the micro-vasculature. The technique uses regular “low” resolution B-mode images such as Figure 1.1(a) of a rat kidney, to create high resolution images of the vasculature such as Figure 1.1(b). The colours in the latter indicate the direction of the flow in the vessels, as indicated by the colour wheel. The SRUS technique is the cornerstone of the work conducted during this Ph.D. project and is described in Section 2.2.1.

1.3 Ultrasound as a research field

The ultrasound field of research is large and active, with more than 2000 attendees at the main international ultrasound conference, the IEEE international ultrasonics symposium (IUS). The various activities in the research field have been separated into five main groups as defined by the IEEE IUS, illustrated in Figure 1.2. The five groups are: “Medical ultrasonics”; “Sensors, NDE & Industrial application”, with NDE being non-destructive evaluation; “Physical Acoustics”; “Micro-acoustics SAW, FBAR, MEMS”, with SAW being surface acoustic waves, FBAR being thin-film bulk acoustic resonators, and MEMS being micro electro-mechanical systems; and “Transducers and transducer materials”. Each of these main groups span several subtopics, of which a few are noted in the following. “Medical ultrasonics” contains all the topics which are related to medical investigations. This spans beamforming algorithms [16], contrast agents and



(a) Ultrasound B-mode image of rat kidney



(b) Super-resolution image of rat kidney

Figure 1.1: Two rat kidneys imaged with ultrasound. (a) is a regular B-mode image showing the typical resolution. (b) is a SRUS image. The colours indicate the direction of flow in the micro-vasculature, according to colour wheel in the top right corner.

development of these [17], cardiac imaging and elasticity [18], photo-acoustics for medical purposes [19], various flow imaging techniques [20] and SRUS. “Sensors, NDE & Industrial application” contains topics of non-destructive evaluation [21] and material change monitoring [22]. “Physical Acoustics” contains the topics concerned with the physical phenomena associated with ultrasound such as particle separation using ultrasound. “Micro-acoustics SAW, FBAR, MEMS” contains the various different types of resonator devices, such as SAW, SAW an SAW. “Transducers and transducer materials” contains all the topics associated with development of transducers and new transducer schemes, for instance for the 3D imaging row-column addressed scheme [23], material development of ultrasound transducer crystals, and development of capacitive micro-machined ultrasonic transducers (CMUTs). Figure 1.2 illustrates how the ultrasound research field encompasses many different areas of expertise and research, from material science, to electronics, to micro-fabrication, as well as medicine.

The activities carried out in this project address different sub-areas. This is illustrated by the red circles in Figure 1.2, highlighting the nodes in which the activities belong. They are distributed across the two main groups of “Medical Ultrasonics” and “Transducers & Transducer Materials”, which is why these are presented in more detail in the figure. This helps to illustrate how the activities within this project from the general point of view might be considered isolated, uncorrelated topics. However, under the right circumstances they are a set of individual pieces which can be fitted together in the greater puzzle of the ultrasound research workflow, working towards greater multidisciplinary achievements.

1.3.1 Typical ultrasound research workflow

Most research fields are subject to the same overall workflow switching between development of new ideas and techniques, and testing and validation. Depending on the results, this might result in necessary changes to the original concepts, which will then again need testing and validation before eventually finishing up with a robust solution. This is no different for medical ultrasound. A generalised workflow for ultrasound is illustrated in Figure 1.3, and will be discussed in the following.

The end goal from a pure research standpoint, marked by the green ellipse, is to develop new techniques in research labs which can eventually be tested and used in the clinic. It is often also in the clinic that the conceptual ideas for new imaging techniques are formed, perhaps based on a specific use case, requiring improved imaging performance or new imaging features. However, before a new imaging concept and/or transducer design can be accepted for clinical use, it will need to go through the same generalised steps. Many concepts will be partly based in hardware and partly based in software. The first step is to formulate the imaging concept of transducer design, a). Next up is a general development block, b1) and b2), for respectively the transducer development and the imaging algorithms. Practical challenges or limitations in fabrication of the transducers might already at this point make it necessary to adapt the original concept, thereby also influencing the imaging development. The same goes for the imaging development, which might reveal necessary restrictions or changes to parameters that influence the original concept and therefore the transducer development. Next up is a testing phase. The developed imaging algorithms can be tested in simulations using simulated data, c), to reveal expected performances. Following testing in simulations, the imaging concepts can be validated in practical experiments, d). In the experimental validation phase it is of course important to mimic the intended clinical use as best as possible, to reveal potential issues before moving to the clinic. Although c) and d) both are testing phases, a new imaging concept will always need experimental validation before moving to clinical testing; it is unlikely that one could move directly to the clinic based on simulated results. At all stages up through the testing phase, the results might lead to adjustments affecting all of the other stages.

The workflow illustrates the multidisciplinary activities required to make a final imaging solution suitable for use in the clinic. At that point, a new set of clinical testing phases will need to be conducted to comply with regulatory requirements. All of the activities are facilitated in the multidisciplinary project which this Ph.D. project has been conducted as a part of. The red

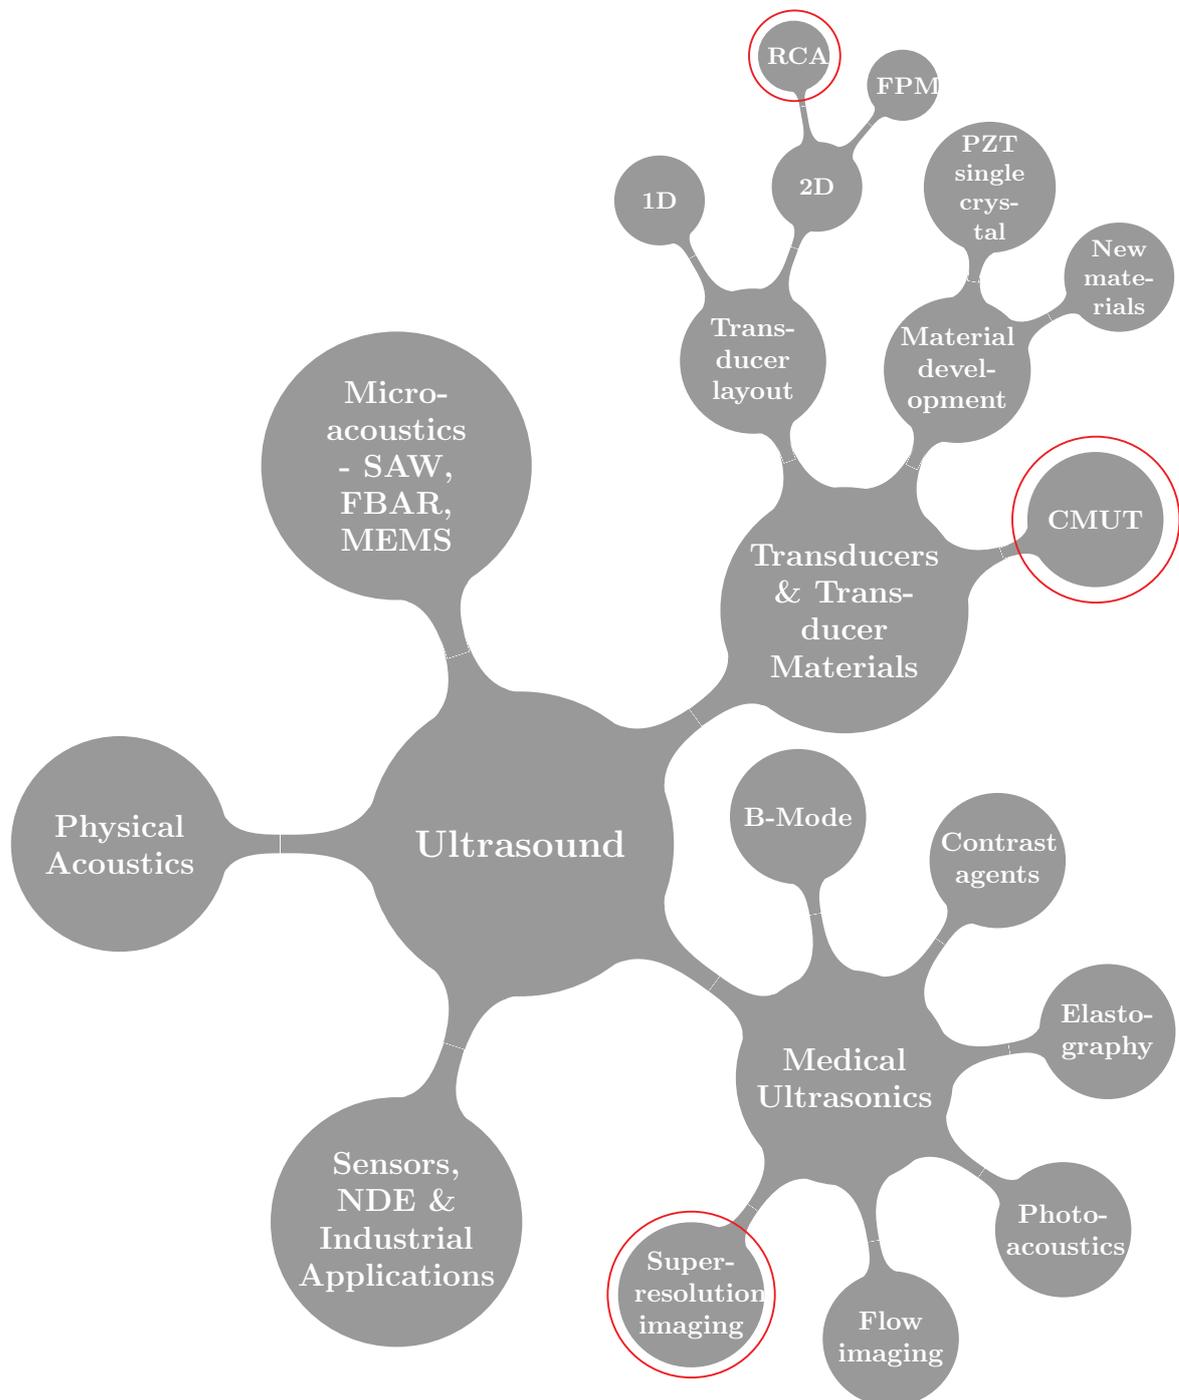


Figure 1.2: Graphical overview of the different topics in the ultrasound research field. The main groups are based on the official separation of the field according to the IEEE IUS. The red circles indicate the activities conducted during this project, highlighting how the activities are separated into isolated branches. The graphic focuses on the branches of the activities in this Ph.D. project.

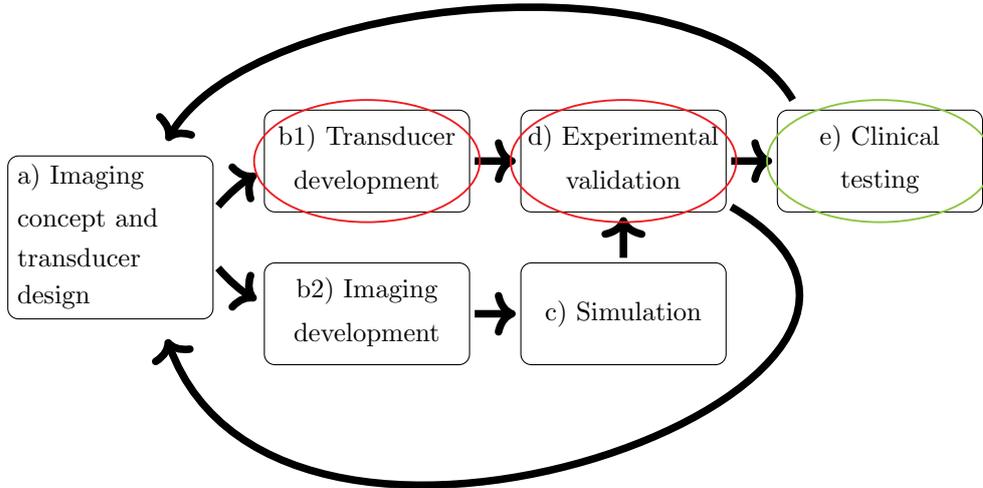


Figure 1.3: Generalised workflow for the development of a new ultrasound technique for the medical ultrasound industry.

circles illustrates the blocks which have been worked on directly during this Ph.D. project, namely development of transducer fabrication methods in b1) and ultrasound phantom development for experimental validation in d). These activities have been conducted in collaboration with the other project participants, involved in the other activities of the workflow.

1.4 Thesis outline

Given the two topics of the Ph.D. project, the main developments presented in this thesis is separated into two parts: Capacitive micro-machined ultrasonic transducer (CMUT) process optimisation and 3D printed phantoms. In addition, two introductory chapters provide a fundamental broad overview of ultrasound and transducer fabrication. Finally, the thesis is ended by a collective conclusion of the presented accomplishments, as well as an outlook. A brief description of the content of each chapter is given in the following.

Chapter 2 - Ultrasound. Firstly, an introduction to the physics of ultrasound is presented, including some basic understanding of which material parameters provides contrast in an image. Next, an overview of the SRUS technique is presented as well as the microfluidic behaviour in micrometre sized channels. Then, a discussion of some differences between 2D and 3D ultrasound imaging is presented, highlighting some differing properties which have been directly used for phantom fabrication. Finally, the use of ultrasound phantoms is described, both with a general overview, but also focused on how phantoms have been used in the SRUS field, and in the work presented in this thesis.

Chapter 3 - Transducer fabrication. An overview of transducer functioning and fabrication is provided, focusing on CMUT devices.

Part I - CMUT process optimisation

Chapter 4 - Hand-bonded CMUTs. The work with fusion bonded CMUT devices is presented. The results show that it is possible to create fusion bonded CMUTs directly in hand of similar quality to CMUTs bonding in dedicated wafer bonders, while minimising the risk of

particle contamination.

Part II - 3D printed phantoms

Chapter 5 - Introduction to 3D printing of phantoms. Firstly, an overview of some of the 3D printing methods is given, before an in depth description of the stereolithography (SLA) method which has been utilised is presented. The details of the specific printer and liquid resin which are used is provided.

Chapter 6 - Hydrogel material characterisation. A study of the swelling of the printed structures is provided to see whether the expansion is isotropic, and whether it depends on the printing exposure time. Analyses of the density, speed of sound, sound attenuation and acoustic impedance is also presented.

Chapter 7 - Calibration phantoms for SRUS. A new type of scatterer phantom based on fixated printed cavities is presented. An analysis of how the scatterer size and ultrasound scattering intensity varies as different sizes, shapes, and dosing schemes are applied is presented. The scatterer phantom is then used to determine the accuracy and precision of a 3D SRUS pipeline.

Chapter 8 - Flow phantoms for SRUS. A series of flow phantom designs and ultrasound experiments using them is presented. A fiducial marker design for proper alignment to the phantom features is presented, and 3D super-localisation of micro-bubbles is demonstrated in a single channel flow phantoms. Additional phantom designs for demonstration of SRUS in 3D are presented.

Part III - Overall conclusion and outlook

Chapter 9 - Conclusion. Conclusion on both parts of the project. An outlook is provided, on all of the possibilities left for exploring following directly from the presented results presented in this PhD thesis.

In this chapter, a fundamental introduction to the physics of ultrasound is given, as well as an overview of medical ultrasound, with a special focus on an introduction to super-resolution ultrasound which is the central topic of this Ph.D. project. The fundamental changes of doing 2D vs 3D imaging is outlined, with details provided on the specific method applied in this thesis. Finally, an overview of how ultrasound phantoms are used, as well as the challenges which need to be handled and the approach that has been utilized in this thesis. This chapter is in part based on Paper B, Paper G and Paper I.

2.1 Basic ultrasound physics

Sound is propagating vibrations of particles through either solid, liquid or gaseous media. Without sound or any other disturbance, the particles in the medium would be evenly distributed in space. Sound is periodic disturbances in the average position of the particles. No net particle displacement is happening, but small oscillations of the particles around their average position result in small local changes of the particle density, and thereby the pressure. Sound consists of longitudinal waves meaning the oscillations are happening in the same direction as the travelling wave. When the frequencies of the oscillations are larger than 20 kHz, which is the auditory limit of humans, the sound waves are classified as ultrasound. In medical ultrasound, the used frequencies, denoted f , typically ranges from 1 MHz to 15 MHz [24], with 3 MHz being a frequently used centre frequency in the clinic. The acoustic wavelength of an ultrasound wave is given as

$$\lambda = \frac{c}{f}, \quad (2.1)$$

where c is the velocity of the propagating wave, or speed of sound, and can be expressed as [25]

$$c = \sqrt{\frac{1}{\rho_0 \kappa}}, \quad (2.2)$$

with ρ_0 being the average density of the medium, and κ being the adiabatic compressibility of that medium. In isotropic solids, the speed of sound can be expressed as

$$c = \sqrt{\frac{Y}{\rho_0}}, \quad (2.3)$$

Table 2.1: Acoustic properties of selected media, including several types of human tissue. Data in the first section is from Jensen, 2013 [25]. Salt water speed of sound is from Kleis and Sanchez [26] and the data is recreated in Section 7.4.

Medium	Density ρ , [kg/m ³]	Speed of sound c , [m/s]	Characteristic acoustic impedance Z , [kg/(m ² s)]
Air	1.2	333	0.4×10^3
Blood	1.06×10^3	1566	1.66×10^6
Bone	$(1.38 - 1.81) \times 10^3$	2070 – 5350	$(3.75 - 7.38) \times 10^6$
Brain	1.03×10^3	1505 – 1612	$(1.55 - 1.66) \times 10^6$
Fat	0.92×10^3	1446	1.33×10^6
Kidney	1.04×10^3	1567	1.62×10^6
Lung	0.40×10^3	650	0.26×10^6
Liver	1.06×10^3	1566	1.66×10^6
Muscle	1.07×10^3	1542 – 1626	$(1.65 - 1.74) \times 10^6$
Spleen	1.06×10^3	1566	1.66×10^6
Distilled water	1.0×10^3	1480	1.48×10^6
Salt water (8.95%, 19 °C)	1.07×10^3	1581	1.69×10^6
Salt water (13.25%, 19 °C)	1.10×10^3	1635	1.80×10^6
Aluminium	2.70×10^3	6420	17.3×10^6

where Y is Youngs modulus. These expressions show that the speed of sound increases with the stiffness of the material, and decreases with density. For anisotropic materials, the speed of sound will be different along different directions, and can be calculated using the stiffness tensor. The speed of sound in selected media can be seen in Table 2.1.

Using Equation (2.1) along with the data in Table 2.1, it can be seen that the wavelength in distilled water at 3 MHz will be $\approx 500 \mu\text{m}$, with the speed of sound of most human tissues only deviating slightly from that value. This is one of the reasons for many *ex vivo* experiments being conducted in water.

The speed of sound of a medium is a critical parameter for creating ultrasound images. The fundamental component in ultrasound imaging is a device capable of both transmitting and receiving ultrasound signals. Ultrasound transducers are devices which are capable of transducing an electrical signal into an acoustic signal, and vice versa. Ultrasound images can be created of objects in the imaged medium, which reflects sound back to the transducer, through time-of-flight evaluation of the ultrasound signal. Knowing the speed of sound in the medium, the distance to the sound reflecting object, d , can be calculated based on the time from signal transmission until it is received again by the transducer as

$$d = \frac{tc}{2}, \quad (2.4)$$

with t being the time between the signal was transmitted until it was received, c being the speed of sound of the medium, and the factor of 2 entering since the sound travelled the distance to the object twice, going back and forth. Thereby, the time of flight measured as the electrical signal can be converted to images of how far objects are from the transducer surface.

Sound only scatters or reflects when it encounters a medium of different acoustic properties. The characteristic acoustic impedance, Z , can be expressed as [25]

$$Z = c\rho, \quad (2.5)$$

where c is the speed of sound in the propagation medium and ρ is the density of the propagation medium. The unit of acoustic impedance is also called Rayl, with $1 \text{ Rayl} = 1 \text{ kg}/(\text{m}^2 \text{ s})$. Values

Table 2.2: Attenuation in different types of human tissue. Data is from Jensen, 2013 [25], assembled from [28].

Tissue	Attenuation [dB/(MHz cm)]
Liver	0.6 – 0.9
Kidney	0.8 – 1.0
Spleen	0.5 – 1.0
Fat	1.0 – 2.0
Blood	0.17 – 0.24
Plasma	0.01
Bone	16.0 – 23.0

for densities and acoustic impedances for typical tissues encountered in ultrasound are also shown in Table 2.1. The differences in the characteristic acoustic impedances at an interface between dissimilar objects results in partial reflection of sound, making the interfaces between media of different acoustic impedances visible in B-mode images. For normal incidence the intensity reflection coefficient is [27]

$$R = \left(\frac{Z_2 - Z_1}{Z_1 + Z_2} \right)^2, \quad (2.6)$$

with Z_1 being the acoustic impedance of the first medium, and Z_2 being the acoustic impedance of the second medium. It can be seen that the listed acoustic impedances are very similar, with the exception of air, bone, and aluminium. This is the reason why it is practically impossible to image the lungs or through bone, and why fractures which have been corrected using metal pins and bolts can be difficult to image using ultrasound. Practically all of the energy is reflected back at these tissue or metal interfaces, resulting in what simply appears as shadows behind them.

The ultrasound intensity is also attenuated when sound propagates through media. Attenuation is often presented as scaling linearly with frequency, typically being measured in dB/(MHz cm). A few examples of attenuation in different human tissue is presented in Table 2.2. To compensate for the attenuation loss, a time-gain-compensation (TGC) is most often applied to the received electrical signal to apply a varying amplification to the signal depending on how much of the imaged medium the ultrasound wave has propagated through, with longer propagation requiring additional amplification. The exact profile of the TGC will vary depending on the imaged medium, as well as the imaging parameters such as ultrasound frequency.

The description of time of flight measurements using a single transducer allows to determine the distance from the reflecting objects to the transducer surface. However, no unambiguous description of the sizes and shapes of the objects in the medium can be determined. Today, most equipment and imaging techniques in use in medical ultrasound are based on 2D ultrasound imaging. However, the ultrasound probes are assembled as arrays of individual transducers, also called transducer elements as illustrated in Figure 2.1(a). Current ultrasound probes will consist of more than 100 elements, making a typical array a few centimetres long (azimuth) and less than a centimetre wide (elevation). Each element can transmit and receive ultrasound and are electrically interfaced individually. Thereby, all elements can be used simultaneously. This makes it possible to acquire full 2D images very quickly, as well as to focus the ultrasound beam by applying a delay profile to the electrical transmit signals across the transducer array along the azimuth direction, such that the ultrasound fields from the individual elements interferes constructively at the focus point, d , in the imaged medium. This is illustrated in Figure 2.1(b). By tilting the delay profile, by not having the profile symmetric as it has been illustrated, the ultrasound beam can be steered along the azimuth direction in the medium, to focus off-axis.

The received electrical signals undergoes a series of signal processing before being presented as the conventional B-mode images most often seen on ultrasound scanners. Some of the typical steps are the beamforming, often a delay-and-sum beamformer, Hilbert transformation of the data,

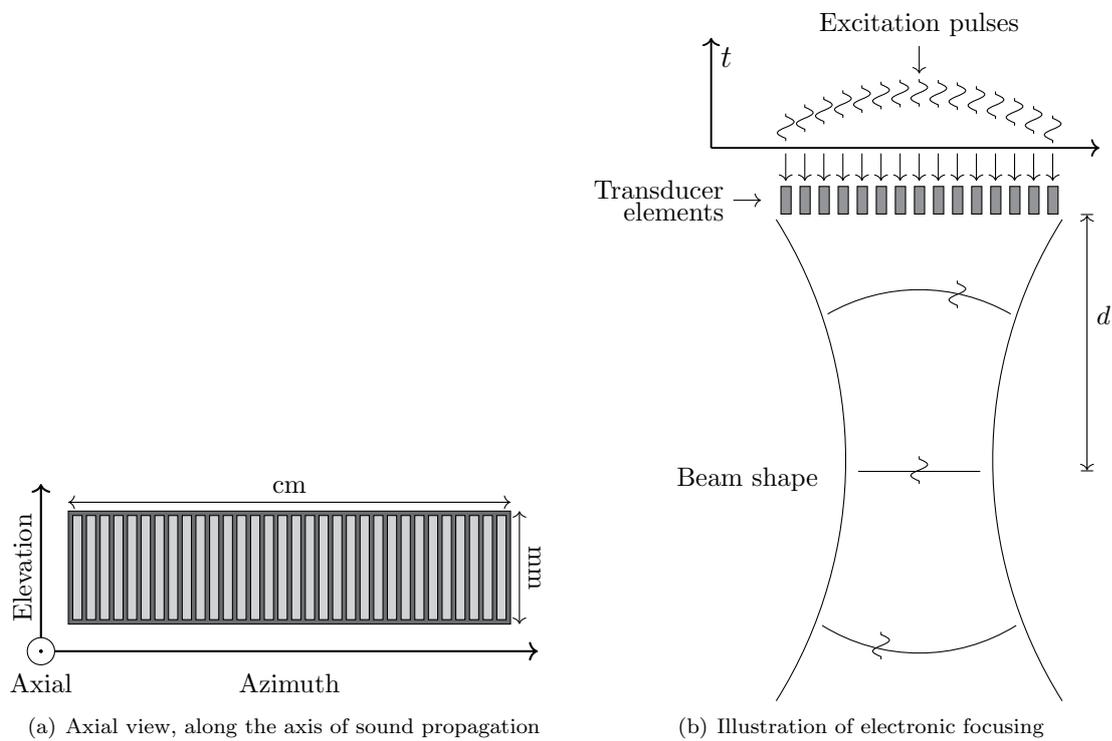


Figure 2.1: Sketch of a 1D transducer array. (a) shows the element layout of the 1D array. (b) illustrates the application of electronic delay profiles across the transducer elements along azimuth to a linear ultrasound array to focus the ultrasound field at the distance d from the array surface.

of which the absolute values are taken. Due to wide dynamic range, the data is often presented on a logarithmic scale. The final images are only a representation of the imaged object. Given that signal processing and direct generation of the B-mode images have not been the topic of this Ph.D., no further details of this will be given here. The process is described in [24].

2.2 Medical ultrasound

A few examples of the topics within medical ultrasound was presented in Section 1.3. Overall, medical ultrasound encapsulates every topic which can be applied for medical purposes, be that diagnosis or therapy. Ultrasound has been used as a medical diagnostics tool since the 1950s [29, 30, 31]. Since then, the techniques and applications have expanded drastically. The mentioned topics of contrast enhanced ultrasound (CEUS), cardiac imaging, photo-acoustics, and flow imaging techniques are only a few of the numerous topics. New techniques are continuously developed, to illuminate topics which would have previously been impossible to investigate. Resolution has been a topic of constant interest, pushing probe developers to create ultrasound transducers with higher operating frequencies. Recently, an alternative approach for improving imaging resolution called super-resolution ultrasound imaging (SRUS) was developed.

2.2.1 Super-resolution ultrasound imaging (SRUS)

Although this project has not been concerned with development of the super-resolution ultrasound imaging (SRUS) algorithms, it is still important with a proper fundamental overview of how the SRUS techniques work, to obtain a better understanding of the requirements for these components.

Obtaining a “superior” resolution in ultrasound imaging is not new, but something that has been discussed since at least 1979 [32]. During the years the proposed method for obtaining that increased resolution has changed, but the goal has remained to increase the resolution of the imaging system beyond that set by the imaging frequency and the classic diffraction limit, and be able to separate the signals from objects closer than the diffraction limit. The diffraction limit is fundamentally a wave phenomenon, and the objective of breaking this limit was therefore not limited to the ultrasound field. In 2006, a number of techniques for optical microscope were presented, namely fluorescence photoactivated localisation microscopy (FPALM), photoactivated localisation microscopy (PALM), and stochastic optical reconstruction microscopy (STORM) [33, 34, 35]. The techniques rely on sequential imaging of fluorescent sources, with photoactivation of only a subset of the fluorescent particles in the field of view at any given instance. This allowed for imaging isolated particles, and localising them with extreme precision. The resolution of the accumulated localisations would end up being in the tens of nanometres.

Only a few years later, the principle of FPALM was transferred to ultrasound as ultrasound localisation microscopy (ULM) [36], with the fluorescent particles having been replaced by ultrasound contrast agents, and the optical cameras by ultrasound scanning systems, with the result that micro-bubbles could be super-localised with micrometer precision.

The overall workflow of typical SRUS algorithms, and the one applied in this work, is illustrated in Figure 2.2. It can overall be summarised in 6 steps: a) acquisition of a series of ultrasound images, typically both contrast enhanced images (a_1) with B-mode images interleaved (a_2); b) data processing, such as noise reduction; c) filtering of the data; d) motion compensation based on the acquired B-mode images; e) super-localisation of the micro-bubble positions; f) tracking of micro-bubbles across images to create flow tracks, and eliminate false detections. The steps are discussed in slightly more detail in the following.

All ULM techniques in the literature rely on the infusion of micro-bubble contrast agents into the medium, which in the clinical setting is the human vasculature. The approved contrast agent SonoVue [37] is often used, however, other contrast agents are being used and developed. Common for these contrast agents are that they remain in the vascular network, meaning indirect imaging of the vascular networks can be obtained. The contrast agents are typically $< 10 \mu\text{m}$ in diameter [37], but will appear much larger in regular B-mode images due to the diffraction limit and the

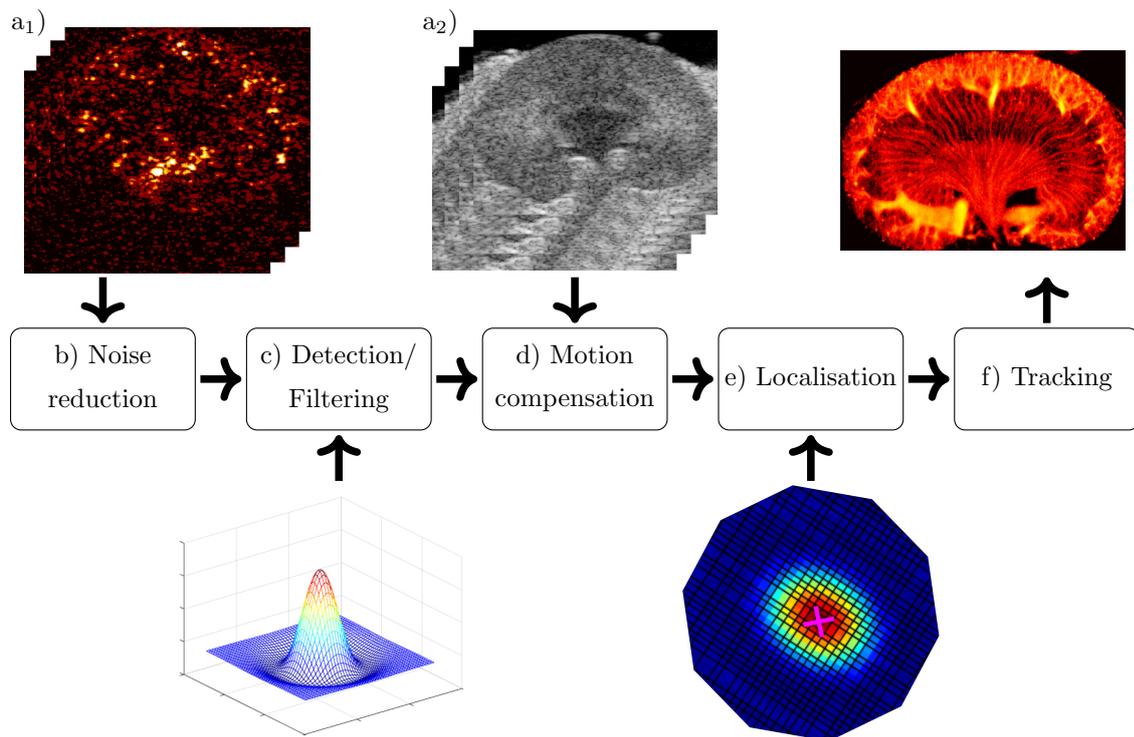


Figure 2.2: Illustration of the typical SRUS workflow. A series of ultrasound images, a), containing contrast agents are processed first by noise reduction, b), filtering based on knowledge of the micro-bubble ultrasound response, c), motion compensation, d), super-localisation of the individual micro-bubbles, e), with finally tracking of the individual micro-bubbles over time to form maps of vascular networks as well as flow velocities, f).

associated point spread function (PSF) of the scanning system. These contrast agents provide high contrast when imaged, often in contrast enhancing imaging sequences, such as pulse inversion (PI) [38] or amplitude modulation (AM) [39], combinations of the two [40], and singular value decomposition (SVD) [41]. Since the localisation algorithms are limited by the diffraction limited PSF of the imaging system, the concentration of the contrast agent needs to be low enough that the individual bubbles can be spatially separated. A series of ultrasound images are acquired of the micro-bubbles. These individual images are illustrated as stacks in Figure 2.2 as a_1) and a_2). Each image will likely need some form of noise reduction, b). Next the images are filtered based on the knowledge of the micro-bubble ultrasound response, c), in order to detect the micro-bubbles, and find isolated micro-bubbles while discarding signals from overlapping PSFs stemming from the micro-bubbles. Contrast agents have been used in ultrasound for a long time in CEUS [15], and many of the techniques for extracting the micro-bubble signals in CEUS can be used directly. The acquisition time will vary significantly depending on the specific technique used, as well as the geometry of the imaged structure. Imaging schemes vary from conventional frame rates of a few 100 of frames per second to ultrafast imaging schemes, which will directly influence how much time is required to obtain a set number of images. At the same time, the diameter of the vessels being imaged influences the flow velocity, with smaller vessels experiencing slower velocities, which will require additional time for the vascular branch to be mapped by a sufficient amount of micro-bubbles. This results in typical acquisition times of several minutes [42]. During the acquisition, motion artefacts might be introduced to the images, either due to subject movement or operator movement. This motion can be compensated by finding the relative motion from image to image, perhaps with varying amount of compensation for different parts of the image, d). Once the motion has been compensated, the micro-bubbles are localised by determining the position of the centroid of their signal on a much finer grid than that of the ultrasound scanning system, through fitting of the known micro-bubble response and interpolation, e). The super-localised positions of the micro-bubbles can then be accumulated across the acquired images, to create a single super-resolution image of the micro-bubble positions, or the individual bubbles can be tracked from image to image, to create velocity maps of the flow in the micro-vasculature, f). Through that framework one can create SRUS images such as the image of a rat kidney seen in Figure 2.3. The brightness in the image indicates localised bubble counts, with bright yellow indicating many counts, red indicating fewer counts and black indicating no counts.

The SRUS field of research has seen tremendous development, with some notable achievements being the super-resolved micro-vasculature of a rat brain [42] resolving 9 μm micro-vessels, and the super-resolved micro-vasculature of a mouse ear [44] resolving 19 μm micro-vessels. Here,

Within the last three years, other techniques have been introduced which aim to separate micro-bubbles closer than the diffraction limit through deep learning or convolutional neural network based methods [45, 46, 47]. The implementation of these would relax the requirement for low concentrations of the contrast agents, which would therefore also significantly decrease acquisition times.

A few definitions are worth discussing. Resolution in any field refers to how close two features can be placed and still be distinguished using the imaging modality. For SRUS the features to be distinguished will be vessels that are very close to each other. However, as was described, it is a requirement that the concentration of micro-bubbles is low enough that the signals from the individual micro-bubble in one frame do not overlap. Thereby, the resolution during imaging is still diffraction limited; The superior resolution is only achieved once many localisations, acquired at different points in time, are added together. In the end, you end up with an image containing localisations which are closer than the diffraction limit. Thus only at the final step of the workflow is the super-resolution achieved.

The localisation precision is the precision with which one can localise a single scatterer repeatedly, and could also be called the measurement uncertainty. This could be a single isolated micro-bubble in a tube without flow, or another somehow fixated sub-wavelength object. Ideally continued localisation of the fixated scatterer should provide the same exact localisation. However, noise in the scanning system will influence the localised position [48]. The distribution of the

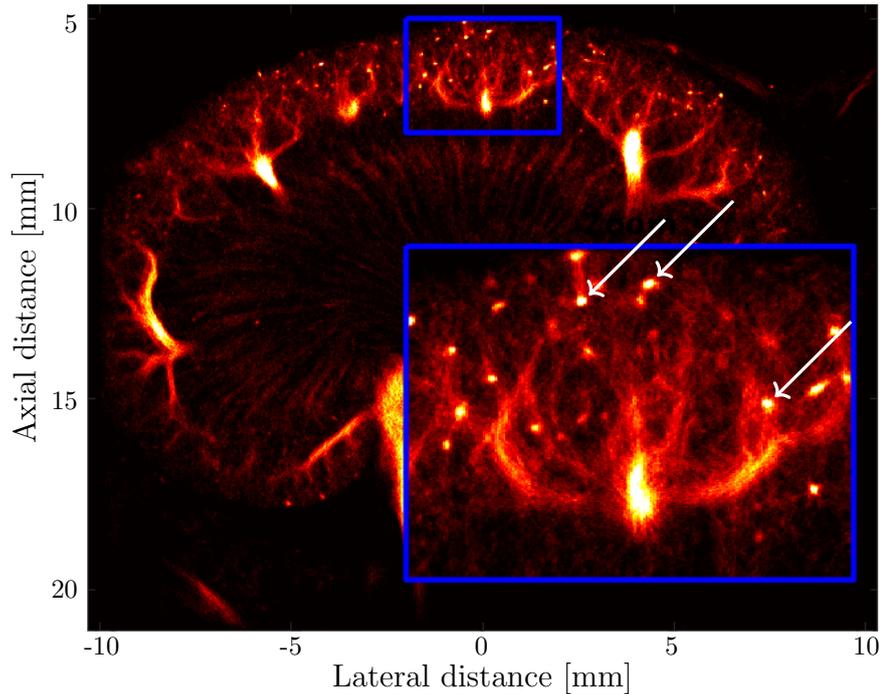


Figure 2.3: SRUS image of a rat kidney. The insert shows a magnified view of the framed section of the renal cortex. The white arrows mark three examples of local regions with many bubble localisations. There are many of these small bright regions distributed across the renal cortex. Image from [43].

localisations will define the localisation precision.

It is important to note that resolution and localisation precision are not the same. This is in part because super-resolution is inherently an indirect measurement method. Images are created based on the micro-bubble positions, however, it is not actually the micro-bubble positions which are of interest; The images are used to infer the structure of the vessels in which the micro-bubbles are placed. The localisation precision thus is a measure for how precisely the micro-bubbles can be located, whereas the resolution would be how close two actual vessels can be, and still be separated by the imaging method. While not the same, the two concepts are inherently linked. Assuming one has an infinitely narrow straight channel which micro-bubbles can flow through, the localisation precision will define how wide this channel will appear in the super-resolution image. The distribution of the localisations will necessarily define how close another infinitely narrow straight channel can be placed, before the distributions from each channel starts to overlap, and they would become indistinguishable. Thereby, the localisation precision will indirectly indicate the systems resolution. It could be argued that there might be interference of signal when two channels are actually being placed closely together, which would not be seen in an isolated localisation distribution. However, given that the typical SRUS schemes require low concentrations of micro-bubbles, this interference is avoided per definition. This is not to say that it is irrelevant to show and demonstrate the resolution directly, but it is important to understand that the high resolution is only achieved through averaging over time.

This section is meant as a brief overview of the concepts in SRUS. For additional details, see the review paper on SRUS published just a few months ago [49].

2.2.2 Theoretical microfluidics for simple geometries

A large and active field within ultrasound research relates to flow estimation, and different methods for doing so. Within the human body, the flow of interest is primarily the flow in the circulatory

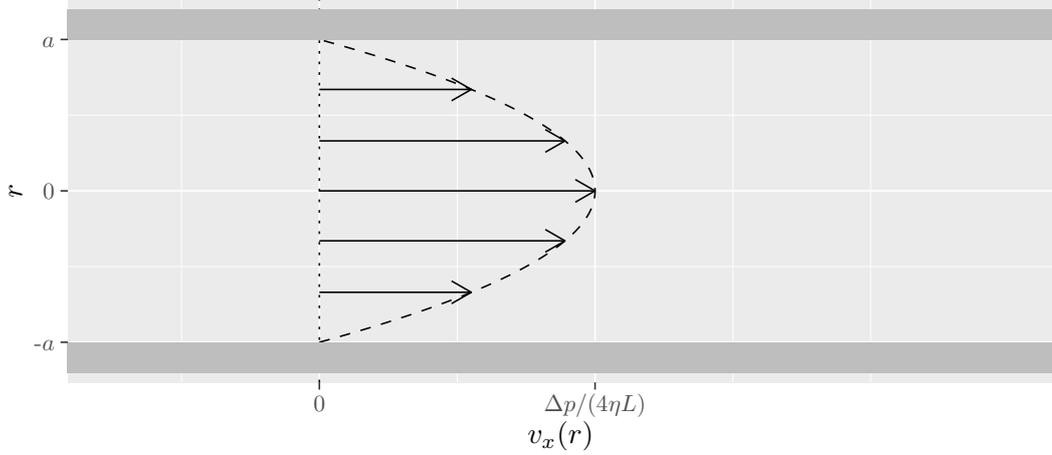


Figure 2.4: Flow velocity profile solution to the Navier-Stokes equation for a circular cross-sectional channel. This is popularly referred to as “Parabolic flow” profile.

system. One reason for the many activities and different subtopics of ultrasound flow estimation is that the flow in the circulatory system happens on a lot of different length scales, in many different specific geometrical configurations, with some medical diseases causing perturbations to the local flow conditions. Thereby, imaging of the vascular system can visualize disease progression.

One of the main goals of SRUS is to determine the flow velocities in the microscopic sections of the vasculature. It will be beneficial to have a fundamental understanding of microfluidics to explain the flow behaviour in properly dimensioned flow-channels, both for SRUS in general, but also for designing and fabricating micro-phantoms for SRUS. Flow is described by the Navier-Stokes equation, for which analytical solutions exist for very specific idealised cases. The analytical solutions to the Navier-Stokes equation for pressure driven flow in tubes are called Hagen-Poiseuille flows and assumes an infinitely long, translational invariant channel. Even though these are idealised conditions, the solutions provide general insight into the flow phenomena.

The no-slip boundary condition is employed when solving the equation, which specifies that the flow velocity at the boundary of the channel follows the boundary itself. Thus, for a fixated channel system, fluid at the channel boundaries will be stationary. The fluid is driven through the channel system by imposing a pressure difference between the channel ends. For a circular cross-section channel oriented along the x -axis, the solution to the Navier-Stokes equation for the velocity component along the channel is [50]

$$v_x(y, z) = \frac{\Delta p}{4\eta L} (a^2 - y^2 - z^2), \quad (2.7)$$

where Δp is the pressure difference, η is the fluid viscosity, L is the channel length, a is the channel radius, and x and y are the coordinates in the cross-sectional plane. $y^2 + z^2$ can be replaced by r to illustrate the radial symmetry. Figure 2.4 shows the radial solution. The flow profile is also known as a “parabolic flow” profile, showing the largest velocity in the centre of the channel and zero velocity at the boundaries. The volume flow rate, Q , in the circular channel is

$$Q = \frac{\pi a^4}{8\eta L} \Delta p = \frac{1}{R_{\text{hyd}}} \Delta p, \quad (2.8)$$

where R_{hyd} is the hydraulic resistance. This is called the Hagen-Poiseuille law, which is analogous to Ohm’s law, with Δp corresponding to voltage, R_{hyd} corresponding to resistance, and Q corresponding to current.

Analytical solutions exist for other special cases of channel profiles, such as ellipses, which is only a slight modification of the circular cross-section, infinite parallel plates, equilateral triangular

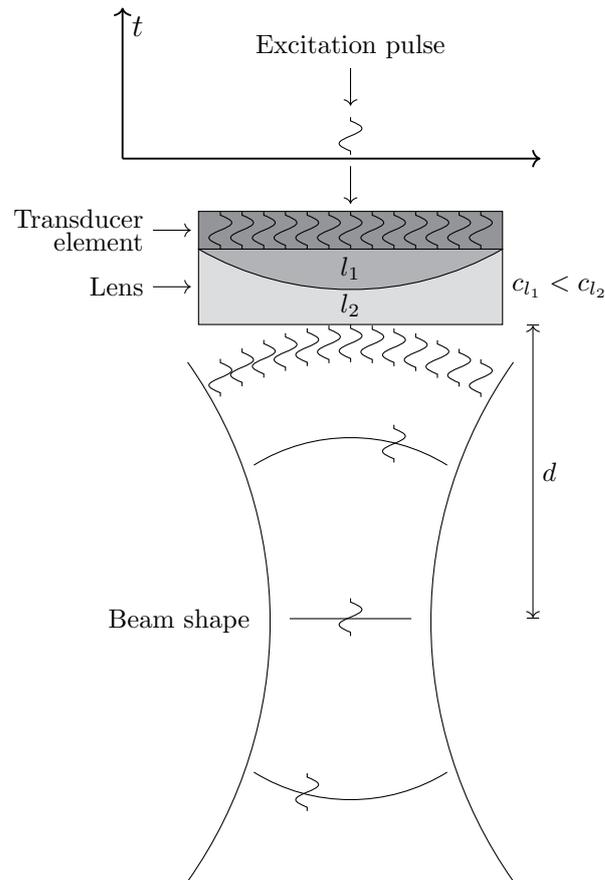


Figure 2.5: Side view of a 1D array, showing the length of an individual transducer element. Since it is only a single element, the excitation pulse is distributed without delay along the length of the element. A lens is applied on top of the element, consisting of two materials of different speeds of sound. As the sound propagates through the lens, a delay profile is effectively applied directly to the acoustic signal, focusing the sound waves at the distance d .

cross-sections, rectangular cross-sections, and shape perturbations of those shapes [50]. However, the geometry of small vessels will predominantly be fairly symmetrical, with flow profile behaviour somewhat similar to that of the circular cross-section.

2.3 2D imaging versus 3D imaging

In Section 2.1, it was illustrated how the individual elements of a 1D transducer array can be used to focus the ultrasound waves. However, since the elements are only distributed in one direction, electronic focusing and steering of the beam can only be done along that direction; it is not possible to make electronic focusing along the orthogonal direction, which is called the elevation direction, since there is only a single long element along this direction. This means that by default, the energy of the ultrasound field will be dispersed as a circular wave along the elevation direction. Thereby, the received signal is effectively integrated along the entire elevation direction resulting in it once again becoming impossible to determine the location of objects in the elevation direction. To avoid this issue, an acoustic lens is applied on top of the array, and moulded to obtain a convex or concave surface depending on the acoustic properties of the lens material. This results in practise to a fixed delay profile, which is applied to the acoustic signal instead of the electric signal. This is illustrated in Figure 2.5. However, this means that the elevation focus of a 2D ultrasound probe

cannot be modified. It will be focused at a specific depth, tailored to the intended use of the probe. At the focus distance, the elevation width is typically $2 - 5 \lambda$ [24], which corresponds to approximately 1 mm – 2.5 mm for a 3 MHz probe. This is a great improvement over simply dispersing the energy cylindrically, and is suitable for most ultrasound imaging applications. It should be noted though, that anywhere but at the focusing distance the width of the elevation plane will be larger. The only way the elevation focus can be moved, is by physical movement of the transducer.

For SRUS, where the goal is to image the smallest vessels in the human body, i.e. arterioles and venules with sub-100 μm dimensions and capillaries of 5-9 μm diameters [2], integration of signal across a region of a few millimetres can be a significant problem. Figure 2.3 of the rat kidney illustrates this potential problem. The image is generated based on cumulated micro-bubble localisations. In the renal cortex, the outermost layer of the kidney, a number of high intensity, bright yellow spots can be seen, of which three have been marked with white arrows. Anatomically, there should not be larger vessels out there, which could otherwise explain a larger bubble-count. A probable explanation could be that the vessels in the renal cortex also travels perpendicular to the imaged plane. Thereby, micro-bubbles can be tracked and added across the width of the elevation plane, erroneously resulting in spots of higher intensity. Most SRUS development has been done using 1D transducer arrays for 2D imaging.

A method to achieve a narrower elevation plane, is to apply dynamic focusing along that direction as well. This could be achieved by modifying a regular 1D array for 2D imaging, illustrated in Figure 2.6(a) with 32 elements, by separating the elements along the elevation direction as well, as illustrated in Figure 2.6(b) with 32×8 elements. If all individual elements can be activated individually, delay profiles can be expanded for 2D arrays, to not only achieve focusing in 3D, but to acquire 3D data, to reconstruct full 3D B-mode volumes. With the exception of special use cases where a rectangular transducer footprint would be an advantage, for instance when imaging through the space between the ribs, 2D transducer footprints are most often quadratic, as illustrated in Figure 2.6(c), with 32×32 elements. This is both because it will allow for imaging of a larger volume, but also because the imaging resolution of the probe improves with an increasing number of elements [51]. Thus, if there is a different number of elements along elevation and azimuth, the probe will have different resolutions along these two axes, which is typically not ideal.

Such arrays are called fully populated matrix (FPM) arrays, and have been used widely in research [24, 52, 53], and to a lesser extent in clinical settings. The size of the array is denoted by the total number of elements in the array, as the product of the number of rows and columns, $N \times M$, or $N \times N$ for a square matrix array. The key argument for using such complex arrays, is that the human body is in fact a 3D object, and not just 2D slices. Thus, it will be impossible to apprehend the full complexity by only observing 2D slices. However, the number of required connections also increase drastically. While a 1D array with N elements would require N electrical channels, a comparative array, supposedly of the same resolution along both directions would require N^2 electrical channels. The illustrated example in Figure 2.6 is a matrix probe which has $32 \times 32 = 1024$ channels. The illustrated matrix size effectively corresponds to that of a state of the art 2D phased array matrix probe by Vermon S.A. (Tours, France), with a footprint of roughly $1 \times 1 \text{ cm}^2$ [53]. Such a high number of channels results in a very large cable from the transducer to the scanner, an equal number of connections on the scanner, and drastically increased complexity for all components of the ultrasound scanning system, essentially making FPM arrays larger than the current 32×32 elements as well as scanners suitable for using them at the very least impractical if not impossible to fabricate and use in practice.

Alternatives such as sparsely populated matrix (SPM) arrays have been demonstrated, were only a smaller part of the full matrix array channels are used, with the used elements being distributed across the probe surface [54, 55, 56, 57]. Another alternative is the row-column addressing scheme, which is the method that has been used for 3D imaging in this thesis.

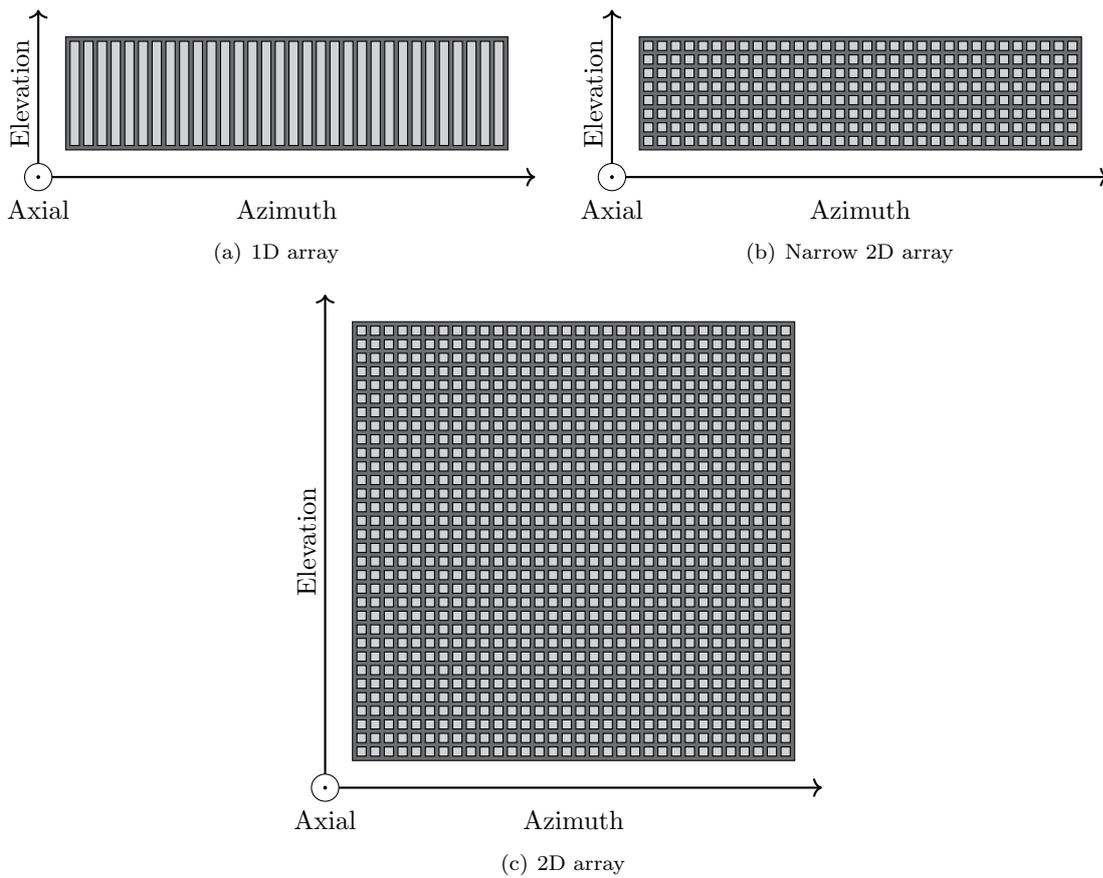


Figure 2.6: Illustration of 1D and 2D array designs. The light grey areas are the transducer elements. (a) is a 1D array for 2D imaging, with 32 elements. To allow for focusing along the elevation direction the elements would need to be split similarly along the other direction as shown in (b), with 8×32 elements. Most often, arrays are made quadratic as in (c), illustrated with 32×32 elements.

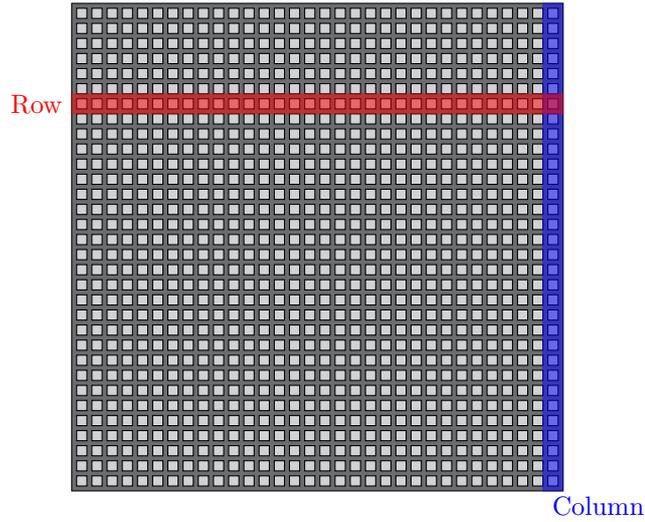


Figure 2.7: Sketch of the row-column addressing scheme on a 32 + 32 row-column addressed array. The matrix elements are contacted collectively along rows on the top (red), and along columns on the bottom (blue).

2.3.1 Row-column addressed arrays

The row-column addressing scheme was first presented theoretically in 2003 [58]. The concept is by some research groups referred to as top orthogonal to bottom electrode (TOBE) [59]. Row-column addressed (RCA) arrays, are matrix arrays where the matrix elements are electrically connected along the rows on the top of the elements, and along the columns on the bottom of the elements. The addressing scheme is illustrated in Figure 2.7. The size of the array is denoted by the sum of the number of rows and columns, $N + M$, or $N + N$ for a square array. The first experimental results were presented in 2006, based on a 64 + 64 element RCA probe [60].

By only addressing the rows and the columns respectively, the total number of channels required for connection reduces to $2N$ instead of N^2 for the FPM array, a factor of $N/2$ less. The decrease in number of channels comes at the expense of increasing complexity in the ultrasound transmission sequences, reception of the signals and subsequent beamforming of the signal to construct the imaged 3D volumes. On the other hand, 2D RCA arrays can have significantly larger footprints than the FPM arrays, while using the same or a smaller amount of channels. For a FPM array with $N \times N$ channels which would have a side length of N when normalised to the element pitch, the side length of the comparative RCA array would be $N^2/2$, again normalised to the element pitch, assuming the same element pitch for both. Thereby, the fraction of the area of the RCA array to that of the FPM array will be

$$\frac{\left(\frac{N^2}{2}\right)^2}{N^2} = \frac{N^2}{4}. \quad (2.9)$$

By addressing the array by the rows and columns, it effectively works as two orthogonal 1D arrays. The imaging scheme is illustrated in Figure 2.8. One of the 1D arrays are used as a transmit array, focusing the ultrasound wave in one direction. The other 1D array is used as the receive array, thereby being able to focus in the orthogonal direction. The transmit and receive focus combines to focus at a point in the volume, thereby allowing for 3D volumetric imaging. More details can be found in [61]. Advanced imaging schemes have been developed, making it possible to obtain real time 3D images and flow estimation in 3D, also called vector flow [63].

While the receive focus can be changed electronically by imposing delays to the signals obtained on each element, the transmit focus will be fixed for each emission. Thereby, it is possible to obtain high resolution along the receive direction, but lower resolution along the transmit direction. The

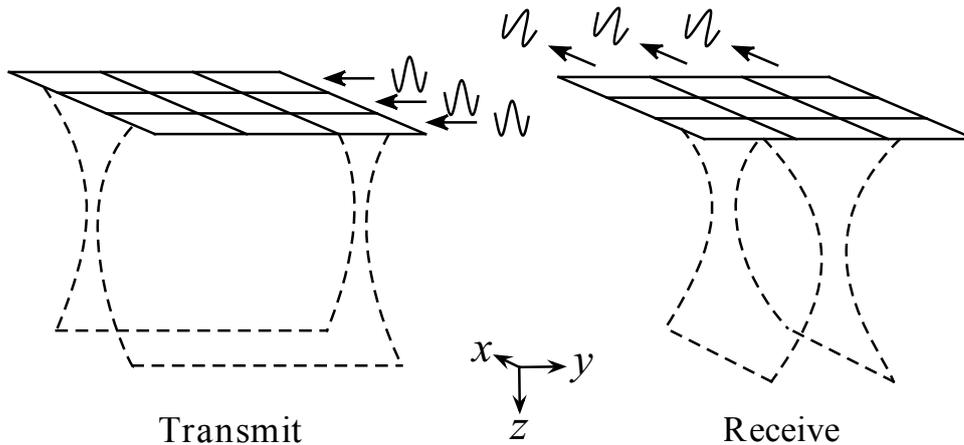


Figure 2.8: Row-column imaging scheme. The 2D array is addressed as two orthogonal 1D arrays. One 1D array is used to transmit, focusing along one direction. The other 1D array is used to receive, focusing along the orthogonal direction. The two array foci in combination results in focusing on a point in the volume, thereby allowing for 3D imaging. Illustration from [62].

transmit resolution can be increased by creating multiple images with different transmit foci and summing the results. The transmit direction resolution will then vary depending on the number of transmit events in the final image, making volume rate and resolution a compromise. Thereby, the resolution in the two lateral directions will not necessarily be the same.

A direct comparison between an RCA probe and a FPM probe both with a channel count of 256, corresponding to $N = 128$ and $N = 16$ respectively, was presented in a simulation study in 2013, showing that the detail resolution could be more than doubled when using the same number of electrical channels [51]. The companion experimental study used a 32×32 FPM probe to compare the FPM performance with the performance of a $32 + 32$ probe, by addressing the same FPM probe as a RCA probe, in other words, keeping the same value of N for both probes [53]. The experiment showed comparative results between both schemes when imaging a wire phantom and comparing the full width at half maximum (FWHM) of the imaged wire, and a poorer performance of the RCA addressing scheme, when determining the cystic resolution at -20 dB. Thus, somewhere between using a factor of $N/2$ less channels for a RCA probe, and using the same exact number of channels, there is a threshold, where the performance of the RCA addressing scheme surpasses the FPM scheme, while using significantly fewer channels.

This section illustrates how different probes have different benefits and limitations. While the work conducted in this thesis has not focused directly on development of the imaging schemes, being aware of these limitations is important. In some cases, utilization the apparent imaging limitations, such as the elevation focus in 2D imaging, can be used as an advantage, for instance when designing phantoms.

2.4 Ultrasound phantoms

When new techniques or transducer designs have been developed and perhaps shown to work in simulations, they need to be tested in practice. In medical imaging, the end goal of the techniques are to image structures in humans. However, before the techniques can be applied on humans in the clinic, they need to clear a lot of regulatory requirements. Most often, the techniques will be tested on animals first, assuming that the features observable in the animals are representative of the human counterpart. However, animal testing is also a quite tedious process, being time consuming, and often expensive as well, without even discussing the general controversy with using animals as test objects. Therefore, prior to doing animal testing, the methods should be tested in another controlled setting. This is where phantoms enter.

Phantoms will vary in complexity, but the finest purpose of a phantom is to provide a well known foundation for what to expect from the experiment. That might involve knowledge of the acoustic properties of the phantom, knowledge of the geometrical layout such as channel dimensions, or knowledge of the relative positioning of features, depending on the specific application. That foundation will often not be available when doing animal or human testing. As the ultrasound field has been developed over time, so has the requirements for phantoms.

Phantoms for medical ultrasound are made to mimic the properties of the structures to be recreated. Slight modifications in fabrication material will often result in changes to the acoustic properties. Therefore, Cafarelli et al. have conducted studies of how the speed of sound and acoustic impedance changes when varying the amount of agarose (AG), polyacrylamide (PAA) and polydimethylsiloxane (PDMS) in phantom structures [64].

With flow estimation currently being an active ultrasound field of research, flow phantoms are presently also being developed. Nguyen et al. developed a phantom mimicking the geometry of the carotid bulb to validate vector flow pressure estimation techniques, and compare them against the intra-vascular pressure catheter which is the clinical reference standard [65]. Yiu and Yu developed a spiral tube phantom allowing for simultaneous flow along all directions for 2D imaging [66]. A 3D equivalent was presented at the IUS 2018 conference, of a “Helical toroid” flow phantom, essentially capable of providing flow in most directions along three dimensions. The work on this has seemingly not been published.

The latest major change in the ultrasound field has been the development of the SRUS field, which once again places new requirements on phantoms.

2.4.1 Phantoms for super-resolution ultrasound imaging

With the goal of SRUS being to image the smallest vascular structures in tissue by breaking the conventional diffraction limit of ultrasound, a phantom for SRUS should ideally replicate the structure and dimensionality of the smallest vessels, which form three-dimensional networks with dimensions smaller than 100 μm . Given that the documented resolution of the SRUS algorithms is a few tens of micrometres, the precision, the accuracy, and the repeatability of the phantom fabrication method should be at a similar level or better.

The phantom studies that have been published have typically consisted of channels defined in various ways. Viessmann et al. [67] and Christensen-Jeffries et al. [40] employed tube phantoms of 3 mm and 200 μm diameters, respectively, to validate their SRUS algorithms. Both of these are significantly larger than the vessels of interest. Desailly et al. presented a phantom study in which the channel dimensions were reduced to $40 \times 80 \mu\text{m}^2$ by utilizing the high resolution of UV lithography on PDMS [68]. While the dimensions of the channel in the latter article was approaching the scale of capillaries, the ability to expand the phantom types to three dimensions are limited in all cases. Harput et al. [57] used a 2D sparse array to do 3D SRUS on a tube phantom consisting of two tubes twisted around each other. This method ensures that the separation between the lumina is defined by the tube wall thickness. This provides a test structure which contains 3D features. However, the twist might deform the tubes and affect the lumina. Furthermore, this does not provide good control of the absolute positioning of the tubes within the imaging field of view. A completely different approach uses the vasculature of chicken embryos, which is optically visible [69, 70, 71, 72]. In this case, the imaged structure is in fact arterioles, venules and capillaries, meaning the scale and complexity is ideal. Since the vascular networks are optically visible one can obtain high resolution optical images of the vasculature for comparison with the vascular maps based on the ultrasound data. However, it is impossible to obtain a three-dimensional representation of the vascular network using commonly available optical microscopes. This is not a limitation of the chicken embryo model itself, since this will feature complex three-dimensional structures. However, the characterisation of those networks is very complex, and not possible to do using regular optical microscopes. Optical mapping of the structures could be performed with other more complex methods such as optical coherence tomography [11, 73], but this has not been utilized so far in the literature.

All of the above mentioned methods are channel based, and thereby meant to provide an outer limit for the positions of the micro-bubbles which are tracked. But that leaves the inherent problem that it is not possible to control the position of the micro-bubbles within the tubes or vessels, and therefore, the source of the signal will not be precisely known. On top of that, it might be difficult to control the actual position of the tubes of a phantom, to obtain a straight tube segment with precisely known positions all along the length of the tube for instance.

2.4.2 3D printing a new type of ultrasound phantom

3D printing of phantoms is a promising new approach, which is not subject to the mentioned limitations. It provides complete three-dimensional flexibility in fabrication and can replicate features in the sub-100 μm range as demonstrated by Jacquet et al., who recently demonstrated 3D printed phantoms for ultrasound [74], supposedly not with SRUS in mind. The phantoms contained highly scattering solid features as small as $30 \times 50 \mu\text{m}^2$ in cross section, demonstrating the exciting potential for point spread function evaluation provided by the method, as well as other possibilities for phantom features and uses.

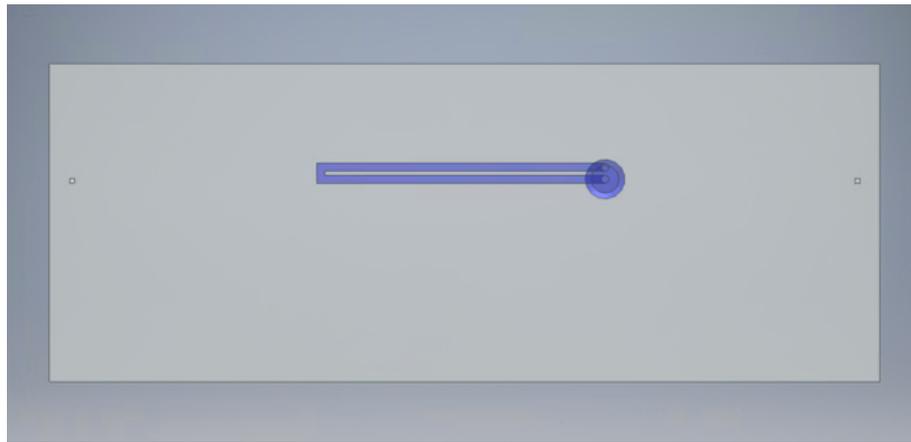
We have been working with a different type of 3D printing, namely SLA. The method is used to print hydrogels, a soft material with acoustic properties similar to tissue. The printer system had in an unrelated research project previously been shown to allow printing of channel systems with cross-sections as small as $100 \times 100 \mu\text{m}^2$ [75]. The printer system and the SLA technique is described in detail in Section 5.1 and Section 5.2.

Flow phantoms for micro-bubble tracking have been made in this project, utilizing the capability of a 3D printing solution to print channel systems in three dimensions. Although not a 3D imaging representation, Figure 2.9 shows a cross-sectional view of a phantom and a SRUS image of micro-bubble localisations in the same phantom. The phantom contains a single channel which returns back on top of itself, allowing for detection of micro-bubble flow in both directions in a single phantom. The 3D printing solution will in principle allow for arbitrarily complex 3D structures or channel networks to be made, only limited by the actual obtainable feature sizes. It would not be possible to create even the 2D channel example illustrated here with the previously mentioned phantom types, simply because of the lack of dimensional freedom. However, in this case the cross-sectional side length of the channel is 200 μm , similar to some of the previous phantom examples, meaning it is still too large to be used as a perfect replication of capillaries. On top of this, it of course suffers from the same limitation of indirect measurements of the vessel structure, and inability to control the exact position of the micro-bubbles within the tubes.

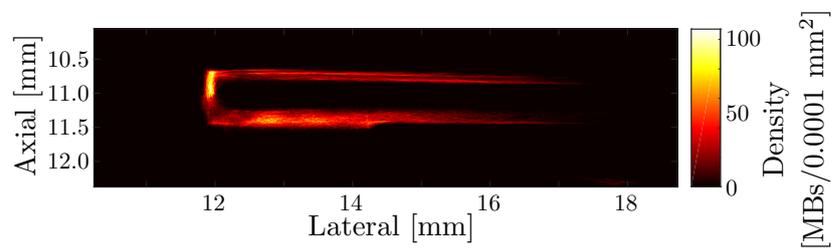
A fundamentally different type of phantom has also been developed in this project. Instead of aiming to replicate the channel systems, the goal was to directly create the scattering source, somewhat similarly to the solid encapsulation presented by Jacquet et al. [74]. If the scattering source can be printed smaller than the imaging wavelength, it will appear as a point target, similarly to the micro-bubbles used in SRUS, although it might not be possible to obtain the same size of the scatterers as that of the micro-bubbles.

In this work, the scatterers are not made by solid encapsulation, but primarily by printing small cavities as illustrated in Figure 2.10. In Figure 2.10(a) the yellow region is the region to be printed, and the black region is the cavity region. The small yellow squares represent the individual voxels in the print. The sketch is scaled similarly to the optical microscope image in Figure 2.10(b), where the voxel grid is optically visible in the printed structure. The voxel size of the printer system is $(x,y,z) = (10.8, 10.8, 20.0) \mu\text{m}$, defining the grid on which the printed features are placed, and therefore also the fundamental limit of the printer accuracy and precision. The scatterers are both in the illustration and in the optical image designed to be 12 voxels wide, and can be placed precisely on the voxel grid using the printer system, for complete control of where the scatterers are placed. Figure 2.10(c) shows a B-mode image of a phantom containing a row of scatterers, with the white arrows marking the visible scatterers. These structures will be stable in time, enabling repeated imaging, in direct contrast to small channels and micro-bubbles.

The different types of phantoms made by 3D printing and the results obtainable with them are described in Part II of the thesis.

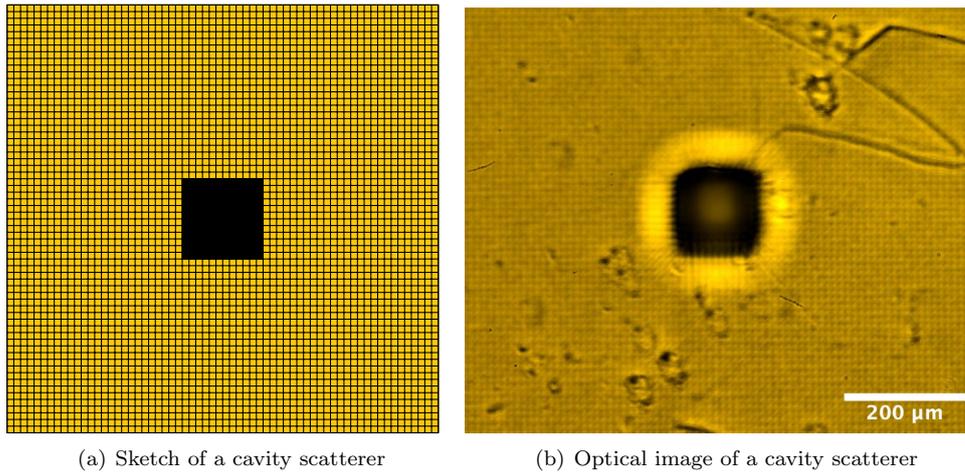


(a) Cross-section of 3D model of channel phantom



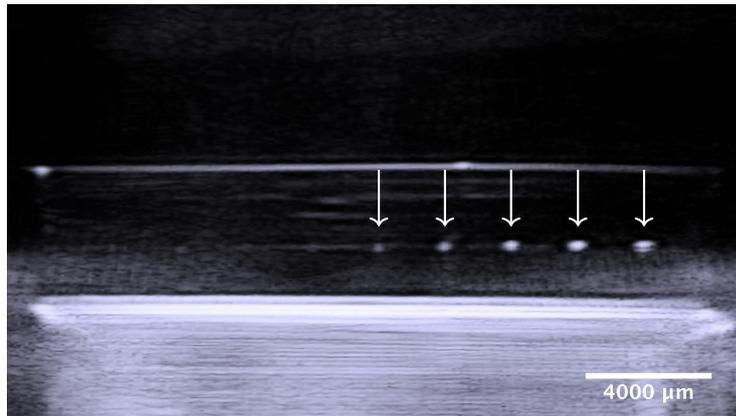
(b) 2D SRUS image of channel phantom

Figure 2.9: A channel phantom bending back on top of itself, to detect micro-bubble flow simultaneously in to opposite directions. (a) shows a cross-section of the phantom model in Autodesk Inventor. The large circle is an inlet channel perpendicular to the cross-plane, and the red regions show the flow channel. (b) shows an SRUS image obtained using the phantom. The colours indicate the micro-bubble counts. Channel side length was 200 μm , and the vertical separation is 108 μm .



(a) Sketch of a cavity scatterer

(b) Optical image of a cavity scatterer



(c) Ultrasound image of a row of cavity scatterers. Image is a slight modification from Paper B.

Figure 2.10: The cavity scatterer concept. (a) is a sketch of the cavity, with black indicating an unexposed region, and yellow indicating printed hydrogel. The small squares frame the voxels, and are matched in size to (b), a microscope image of an actual printed scatterer placed at the top of a print. (c) is an ultrasound image of a row of cavities, which have been marked by white arrows.

In this chapter, an introduction to ultrasound transducers is given. Conventional transducers are briefly outlined, before a slightly more detailed description of capacitive micro-machined ultrasonic transducers (CMUTs) is given, which have been used in this project, both in terms of how they work, and how they are typically fabricated.

3.1 Conventional ultrasound transducers

A transducer is a device which facilitates the conversion of energy from one form to another. For ultrasound transducers, energy is taken from the electrical domain, in the form of electrical signals, which are transformed to the mechanical domain, in which deformations of material can periodically compress the surrounding air, forming sound waves. In many cases, the transduction will work in both directions, i.e. both from the electrical to the mechanical domain, and from the mechanical to the electrical domain, allowing the device to function both as a transmitter and a receiver, in which case the device is called a transceiver.

Multiple different types of transducers have been made for ultrasound equipment. Common for the devices in medical ultrasound imaging is that the arrays are built as demonstrated in Figure 3.1 which is a zoomed in version of the array sketch from Section 2.1 in Figure 2.1 with some additional descriptive details showing the transducer element pitch, element width, and kerf separating the elements.

Most ultrasound transducers since the start of the field in 1950 have been based on piezoelectric ceramic (PZT). Strictly speaking, PZT refers to lead zirconate titanate, which is the most commonly used material. However, many other piezoelectric materials have been used and are still actively being developed, with fundamentally the same overall function, namely piezoelectricity. Therefore, they will in the following all be referred to as PZT.

Piezoelectric materials are materials which experience a change in potential across the material when a stress is applied to it, and conversely also experience mechanical deformation when a potential is applied across the material. Specifically, the material will contract and expand as the potential across it is changed. This is exactly the transceiving principle utilized for ultrasound probes. A slab of PZT is covered by one electrode on the front and another on the back of the material. When an oscillating potential is applied across the two electrodes, the material will start expanding and contracting as a response to the potential, thereby compressing the surrounding air, and generating sound waves. Oscillation frequencies in megahertz range will make it suitable for medical ultrasound. When ultrasound waves impinges on the piezoelectric material, it will deform

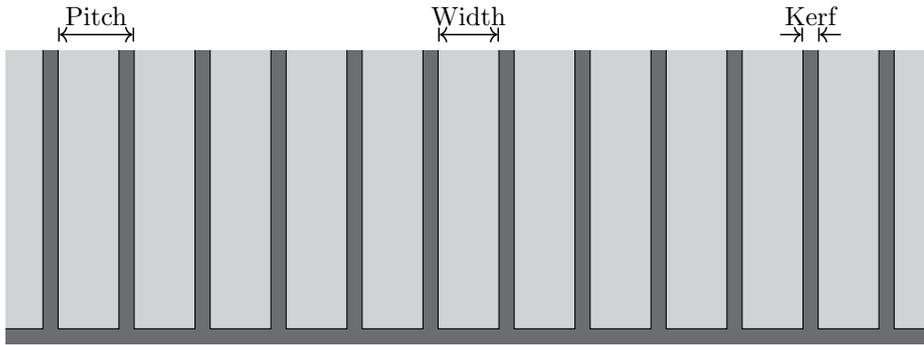


Figure 3.1: Zoom of 1D array sketch in Figure 2.1(a), in which the dark grey region marks the outline of the array, and the light grey rectangles are the transducer elements. The common geometrical terms element pitch, element width and kerf have been marked.

in response to the ultrasound waves, resulting in electric potential changes which can be measured and converted to ultrasound images.

The dimensions of the PZT crystal defines the resonance frequency of the transducer according to

$$f_c = \frac{c_{\text{crystal}}}{2t_{\text{crystal}}}, \quad (3.1)$$

with c_{crystal} being the speed of sound in the used PZT crystal, and t_{crystal} being the thickness of the crystal. Thus, for increasing frequencies, a decreasing crystal thickness is required. For 15 MHz probes, the required thickness of a lead zirconate titanate is approximately 100 μm . It becomes increasingly more difficult to achieve sufficient uniformity for thinner layers of PZTs, which is one of the apparent limitations of PZTs.

The elements of PZTs are defined by sawing out the PZT, thereby allowing for individual electrical contacting of the sub-crystals. Typical saw blades are between 15 μm and 40 μm wide, and directly define the width of the kerf. It has been shown that imaging improvements can be gained by designing the element pitch as a half wavelength, $\lambda/2$ [76, 77]. Thereby, the width of the elements also scale with the imaging frequency. For a frequency of 15 MHz, the element pitch would need to be approximately 100 μm , of which the kerf would remove 15-40%. This percentage will of course increase for increasing frequencies as the size of the kerf does not change, while the pitch does. This also decreases the emitted pressure, since the transducer elements decrease in size.

However, there are alternatives to PZT transducers, which do not in the same way suffer from these fabrication limitations, with the one used in this work being CMUTs.

3.2 Capacitive micro-machined ultrasonic transducers (CMUTs)

Capacitive micro-machined ultrasonic transducers (CMUTs) were invented in the early 1990s and first presented in 1994 by Haller and Khuri-Yakub [78]. The fundamental transducing unit is called a cell, and is a drum-like structure, illustrated in Figure 3.2. It consists of a rigid substrate on which the CMUT cavity is defined, often in another support material, on top of which is a thin plate facilitating the transduction by being set into vibration.

CMUTs are typically fabricated using the well established silicon micro-fabrication techniques such as UV-lithography which easily allows for lateral definition of features down to 1 μm , and thin film growth or deposition for height control of layers with nanometre precision, both of which are significant improvements on the techniques used when fabricating PZT transducers.

Over the years, many different fabrication methods have been used to fabricate CMUTs with many different specific purposes, for instance sensing, microphones, and medical imaging. Recently,

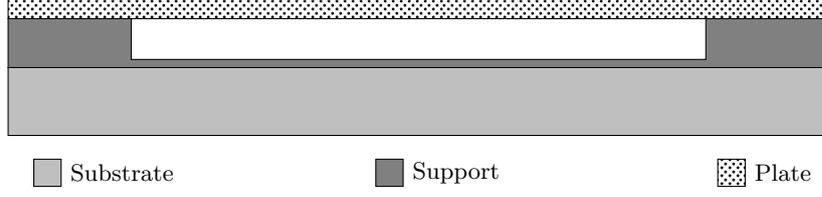


Figure 3.2: Illustration of the cross-section of the fundamental unit of a CMUT transducer array called a cell, the CMUT. It consists of a rigid substrate, a support structure defining the cavity, and a plate on top.

Butterfly Network, Inc. (Guilford, CT, USA) brought a CMUT array transducer to market called Butterfly iQ, which is both Food and Drug Administration (FDA) approved and CE marked. It is a 2D array, designed to be connected to a phone or tablet for imaging directly on that mobile device. It is inexpensive, and designed with versatility and ease of use as the main selling points. Due to the CMUT technology, it can be used and addressed in many different common transducer schemes to emulate linear arrays, curved arrays, and phased arrays. It is however limited in its image quality capabilities, still leaving room for conventional ultrasound imaging systems.

3.2.1 Basic CMUT physics

CMUTs are fundamentally plate capacitors. Referring back to Figure 3.2, the plate is used as one electrode and the substrate as the other, either directly if the plate and substrate are electrically conductive, or indirectly by deposition of metal electrodes on top of them. The capacitance of a parallel plate capacitor is

$$C = \frac{\epsilon_0 \epsilon_r A}{l}, \quad (3.2)$$

where ϵ_0 is the vacuum permittivity, ϵ_r is the relative permittivity of the material in between the plates, A is the area of the plates, and l is the distance between the plates. An array of CMUT cells is not just a simple parallel plate capacitor for a number of reasons though. First of all, part of the region between the plates is a rigid structure, the support, and part is vacuum or air. Thereby, the dielectric properties are not constant in the region between the plates. Second, as a potential is applied to the two electrodes of the CMUT, the accumulating charges on the plates will attract each other. Since part of the space between them are simply a gap, they plate, being much thinner than the substrate, will be pulled down towards the substrate. Thus the distance between the plates is not constant a constant, and will also influence the capacitance.

Most often, the CMUTs are shaped as circles. Solving the plate equation, it can be shown that the plate deflection of an isotropic circular plate is

$$w(r) = w_0 \left(1 - \left(\frac{r}{a} \right)^2 \right)^2, \quad (3.3)$$

where $w(r)$ is the plate deflection depending on the radial position from the centre of the cavity, a is the cell radius, and w_0 is the maximum deflection at the centre given by

$$w_0 = \frac{p a^4}{64D}, \quad (3.4)$$

where p is the load applied to the top plate, and D is the flexural rigidity of a plate given as

$$D = \frac{E h^3}{12(1 - \nu^2)}, \quad (3.5)$$

where E is Young's modulus of the plate material, h is the thickness of the plate, and ν is Poisson's ratio of the plate material. This expression for the plate deflection can be used as input to the

expression for the capacitance in Equation (3.2), with the separation l being replaced by $g - w(r)$, where g is the distance between the plates at rest.

For isotropic materials, the equations above explain the behaviour very well. However, it has been shown that for an-isotropic materials, such as silicon (Si) which is a commonly used plate material, an error of up to 10% will be obtained when using the equations for isotropic plates, compared to finite element modeling [79]. Other corrections will need to be applied if using plate materials with built in stress. For a plate with a tensile stress, the centre deflection can be shown to be [80]

$$w_0 = \frac{p a}{2} \sqrt{\frac{C D}{N_t^3}} \left(\frac{1 - I_0\left(\sqrt{\frac{N_t}{C D}} a\right)}{I_1\left(\sqrt{\frac{N_t}{C D}} a\right)} \right) + \frac{p a^2}{4 N_t}, \quad (3.6)$$

where C is a constant based on the solution to the plate equation, D is the flexural rigidity of the plate, I_n is the modified Bessel function of first kind, and $N_t = \sigma h$ is the stress resultant, where σ is the planar biaxial stress in the plate. It should be noted that the load applied to the top plate is potentially a combination of both a pressure difference across the plate, and the attracting force between charges in the case of an applied potential.

The frequency of a CMUT is again defined by the cell geometry. For an isotropic circular plate, the eigenfrequency can be shown to be [81]

$$\omega_0 \approx \sqrt{\frac{80}{9} \frac{E}{\rho(1-\nu^2)} \frac{h}{a^2}} \quad (3.7)$$

where ρ is the density. It can be seen that everything except h and a are material constants. The eigenfrequency can thus be tuned by modifying the plate thickness and the cell radius. For CMUTs using a Si plate of $\approx 3 \mu\text{m}$, the radius will be $\approx 25 \mu\text{m}$ to be in the range for medical ultrasound, easily within the capabilities of UV-lithography. Note that the lateral dimensions of the CMUT cells are smaller than a common element pitch, on the order of a few hundreds of micrometres, as well as the element length which is on the order of millimetres. To obtain the rectangular layout defined for the transducer arrays, multiple CMUT cells are placed next to each other, and connected in parallel, to operate in unison as a single large element.

The ultrasound transduction principle of CMUTs is illustrated in Figure 3.3. The CMUT is biased with a DC voltage, pulling the plate down to the operating position, as this increases the efficiency [82]. When an AC potential is applied on the electrodes on top of the DC bias, the plate start oscillating around the operating position, illustrated in Figure 3.3(a). The plate will move up and down in tune with the AC potential, compressing the air, creating sound waves. When in turn a sound wave impinges on the plate, the pressure wave will move the plate, changing the capacitance of the CMUT, which can be measured as a potential shift across the electrodes, illustrated in Figure 3.3(b).

UV-lithography is also used to define the kerf between elements, once again allowing for designing the kerf to be only a few micrometres wide. This illustrates how CMUTs are a promising candidate for ultrasound transducer fabrication going into the future, as the fabrication issues which are to be expected for PZT transducers are not an issue with CMUTs.

This has only been a brief overview of the CMUT physics, jumping directly to the relevant results necessary in this thesis. Although Part I is directly concerned with CMUTs and CMUT fabrication, the focus in this work has directly been on fabrication process optimisation, for which the above description is sufficient. For more details, see [81].

3.2.2 Conventional CMUT fabrication methods

So far, the CMUT has only been described conceptually, consisting of a cavity on top of which a plate is suspended. However, this can be fabricated in a number of different ways. In the following, the most common methods, which are sketched in Figure 3.4, will be described.

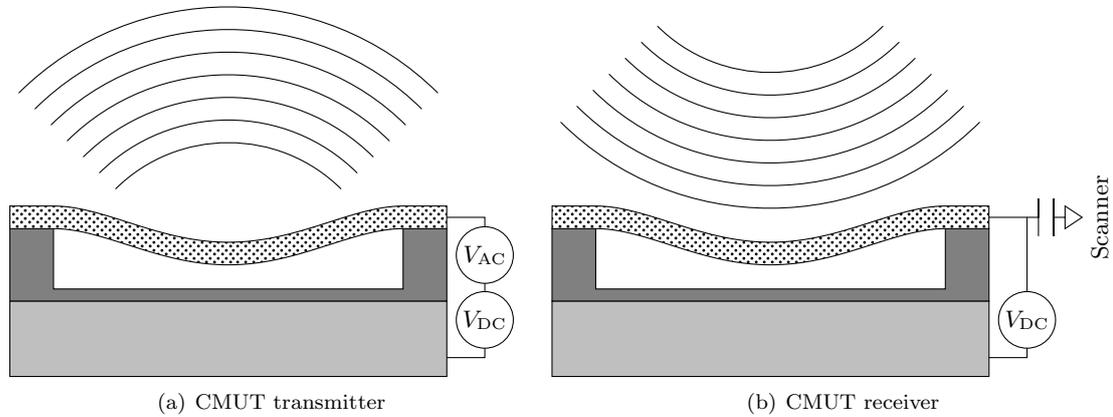


Figure 3.3: Illustrations of the transceiving capabilities of CMUTs. A DC bias is applied, pulling down the plate. (a) illustrates how an alternating potential applied on top makes the plate vibrate according to the signal, emitting sound waves. (b) illustrates that an incoming sound signal will make the plate oscillate, which can be read out electrically on a scanner system to be processed into ultrasound images.

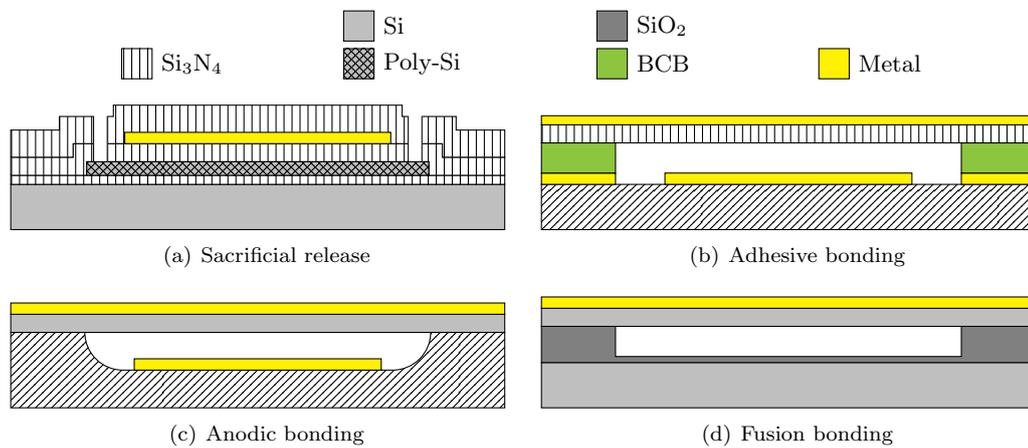


Figure 3.4: Sketches of the cavity structures in the most common CMUT fabrication methods. (a) is one of the main groups of fabrication methods, sacrificial release, in which the cavity is defined by a buried material, often poly-silicon, which is etched away late in the process. The other main group is wafer bonding, in which the CMUT plate is bonded on top of a pre-fabricated cavity. (b) adhesive bonding, (c) anodic bonding, and (d) fusion bonding, are all different wafer bonding methods. Sketches are not to scale since cavities typically have very high aspect ratios, being a few 100 nm tall, and $\approx 100 \mu\text{m}$ wide.

The CMUT fabrication methods can be split into two main groups: Sacrificial release and wafer bonding. The main difference of the two groups is the way in which the cavity is defined.

Sacrificial release

An example of a sacrificial release CMUT structure is illustrated by the cross-section seen in Figure 3.4(a). The shown structure could be fabricated bottom up, starting with a Si wafer which serves as the bottom electrode of the CMUT. On top of that a silicon nitride (Si_3N_4) layer is deposited for insulation of the bottom electrode. Next, a layer of a sacrificial material is deposited, which is often either poly-silicon or chromium (Cr). This layer will eventually define the cavities, and the thickness should therefore match the desired cavity height. After deposition, it should be patterned into the desired shape and layout of the CMUT cells. Then, a conformal deposition of additional Si_3N_4 is made to bury the sacrificial layer in the structure. This part of the layer on top of the sacrificial metal structure will form the bottom of the CMUT plate, which serves as isolation of the top electrode. On top of that, another metal, perhaps gold (Au), is deposited and patterned, forming the top electrode of the CMUT. This might be covered by another conformal Si_3N_4 layer to define the final plate thickness, which will then be a stack of Si_3N_4 -Au- Si_3N_4 , resulting in the structure illustrated in Figure 3.4(a). In the illustration, vertical channels have been dry etched on both sides of the cavity, to access the sacrificial metal layer. This can then be removed by wet etching, releasing the plate and forming the cavity. Variations of the process illustrated here have been demonstrated in a lot of publications [83, 84, 85, 86], including the very first CMUT publication by Haller and Khuri-Yakub [78]. Many variations of different complexity have been developed, in which spaces for the different metal layers are etched out in the Si_3N_4 [87]. This allows for more robust fabrication, and a reduced parasitic capacitance, which is a term for the undesired capacitances present in CMUT structures.

Wafer bonding

The wafer bonded CMUT structures are in general made by defining the cavities as recesses in one wafer, and sealing them off by bonding another wafer on top of the recesses. This is the method that has been applied during this Ph.D. project. There are many different ways in which the bonds can be formed. The three methods, adhesive bonding, anodic bonding, and fusion bonding, that have been applied during this Ph.D. project are described briefly in the following, of which the work on adhesive bonding and fusion bonding is included in the thesis.

An example of an adhesive bonding CMUT structure is illustrated by the cross-section seen in Figure 3.4(b). There are multiple variations of adhesive bonding, but the figure illustrates the general concept. First, a substrate, either conductive itself [88], or alternatively isolating with a metal bottom electrode defined on top (illustrated) [89], is coated with an adhesive material. The CMUT cavity is defined directly in the adhesive material. Next, a plate is bonded on top. The bonding is finalised by thermal treatment of the wafer stack, to cure the adhesive material, ensuring a high bond strength between the layers in the stack. The top plate is likely transferred on a handle wafer, which will need to be removed by etching, before the top elements can be defined by patterning and etching. The plate material should either be conductive enough to form the top electrode itself, or another metal layer should be added on top (illustrated) [89, 90]. Finally, access to the bottom electrodes can then be made. Often, an additional layer is added to the stack prior to bonding for insulation of the electrodes, to avoid short-circuiting the device in the case the plate is pulled down to the substrate during operation. If the plate material is isolating, no additional layer is needed. SU-8 or benzocyclobutene (BCB) are commonly used adhesives [90, 91, 92], with BCB having been used during this Ph.D. project. The choice of top plate material might put additional requirements on the adhesive material. If the top plate is conducting, the adhesive will be required to isolate the two electrodes from each other. According to the manufacturer, the breakdown field in BCB should be 0.53 V/nm. Alternatively, isolating plate materials, such as Si_3N_4 can be used [93, 89], in which case the plate itself will also contribute to the electrode isolation. An argument for using an isolating substrate is that it will decrease the electric cross-talk between elements,

which is an unwanted effect of coupling of electrical signals down through the substrate wafer [94]. The adhesive wafer bonding method is very forgiving towards particle contamination and surface roughness, resulting in a robust bonding mechanism. Furthermore, the fabrication process is fairly simple and short, meaning it could be a good solution for rapid prototyping. The temperatures used in this process are lower than 300 °C, making the process CMOS compatible [93], allowing for integration of the CMUTs directly on application-specific integrated circuits (ASICs) for on-chip signal processing [95, 96]. The focus of the work with adhesive bonding during this Ph.D. project has been to improve the dielectric properties of the BCB through curing schemes. The breakdown field of BCB stated by the manufacturer seemed promising, and would be sufficient for the designs of interest. However, our testing showed very poor performance with typical breakdown fields as low as 0.08 V/nm [82]. The thermal curing process of the BCB layers was optimized through fractional factorial experimental designs, resulting in systematic improvements. Even so, the improvements were too small to be viable. The solution chosen was to utilize an isolating plate material, which would be able to carry the applied potentials. This was presented in Paper A. The activities in adhesive bonding has since been stopped, and will not be discussed in further details in this thesis.

An example of an anodic bonding CMUT structure is illustrated by the cross-section seen in Figure 3.4(c). Once again, this is quite a simple process, which is also based on an insulating substrate, which would remove the effect of electric cross-talk through the substrate wafer. In this case, the substrate wafer is a borosilicate glass wafer, which contains around 81% silicon dioxide (SiO_2), 13% B_2O_3 , 4% $\text{Na}_2\text{O}/\text{K}_2\text{O}$, and 2% Al_2O_3 [97]. The cavities are etched directly into the glass wafer, and metal contacts are deposited and defined inside the etched recesses. At elevated temperatures (often less than 400 °C), the ions in the glass become mobile, and will redistribute in the glass when a potential is applied across a wafer stack of a borosilicate wafer and a silicon wafer. Positive Na^+ ions are depleted from the region near the bonding interface, polarising the glass, and creating a large electric field right at the interface which effectively pulls the wafers together. The required potential for a strong bond depends on the exact structures and layers in the wafer stack. Often, a voltage ramp, starting around 300 V and incrementally increasing in a few steps to upwards of 800 V is used. Subsequently, the top handle wafer can be removed by etching, the top elements can be defined, and access to the bottom electrodes can be made. Anodic bonding has been demonstrated with wafers having a surface roughness up to 50 nm [98]. The final anodic bond is so strong that the substrates will break before the bond fails [99]. The temperatures used for this process are again low enough that they are CMOS compatible. For this structure, the glass substrate is the insulator between the top and bottom electrodes. Typical breakdown fields in the glass are 0.92 V/nm. CMUT fabrication based on the anodic bonding process was initiated during this project, but has since been continued by other Ph.D. and master students working in the group. Therefore, it will not be discussed further in this thesis.

An example of an fusion bonding CMUT structure is illustrated by the cross-section seen in Figure 3.4(d). For fusion bonding, a substrate Si wafer is first oxidised. The oxide is used to define the cavities, either by etching parts of the oxide away directly, in which case a subsequent oxidation is typically carried out to form an oxide in the bottom of the cavity to isolate the electrodes (illustrated), or by local oxidation of silicon (LOCOS) [100]. The breakdown field in the SiO_2 available has been measured to be ≈ 0.74 V/nm. The substrate wafer and the to-be-bonded silicon on insulator (SOI) wafer [100] or Si_3N_4 wafer [101] are then cleaned before being placed in contact, at which point a relatively weak pre-bond is formed between the wafers. The bond strength is then increased by annealing at 1100 °C. Next, the handle layer of the SOI wafer is removed, releasing the plate. Then the top electrodes can be defined, and access to the bottom electrodes can be made. While the process appears very brief, some of the variants of cavity definitions require several thermal processes [100, 101], making fusion bonding a time consuming process in comparison to the other wafer bonding techniques. It is also very sensitive to particles in comparison to the others, and requires a surface roughness of less than 1 nm [102]. The fusion bonding method is also prone to the substrate coupling issue. However, the structure is also typically very stable in operation over time, with little to no charging effects [82]. A variant of fusion bonding, carried out without the use of any equipment has been investigated during this

Table 3.1: Comparison of the properties of common wafer bonding techniques. Adapted from [82].

	BCB	Anodic	Fusion
Breakdown voltage	0.53 V/nm*	0.92 V/nm [†]	0.74 V/nm
Particle sensitivity	Low	Low	High
Surface roughness tolerance	High	Medium	Low
Highest processing temperature	≈ 250 °C	≈ 350 °C	≈ 1100 °C
Fabrication time	Days	Days	Weeks
Successful practical experience	Limited	Ongoing	Yes

* according to manufacturer
[†] from [103]

project, and is presented in Chapter 4, in which the fusion bonding process is also described in more detail.

A summarising comparison between the techniques can be seen in Table 3.1, adapted from [82]. Consulting the table, it might be surprising that the fusion bonding method is the one applied and presented in this thesis. While fusion bonding might appear to be lacking in most of the presented properties, the major benefit is the accumulated knowledge of this process during the previous years within the research group and the known and documented good properties of the resulting structures. This is illustrated by the last row. The properties of the other processes are of a more hypothetical nature, particularly at the time of the CMUT fabrication activities in this PhD project, based on material properties but not practical CMUT results. The ongoing activities within the group have since provided promising results for the anodic bonding method, due to continued process improvement.

Part I

CMUT process optimisation

This chapter presents the work conducted on hand-bonded CMUTs, which is bonding of two wafers without the use of a dedicated wafer bonder. An analysis is given of what to expect of the gas content in the resulting devices. The chapter is in part based on Paper C, presented work at MUT 2018 and MNE 2018 - Poster 2, and the unpublished manuscript Paper K.

4.1 Motivation

As described in Section 3.2.2, many different capacitive micro-machined ultrasonic transducer (CMUT) fabrication schemes are used in the research field, with wafer bonding being the one applied in this work. Fusion bonding has been the work horse for almost a decade now in the MEMS-AppliedSensors group, but during the last few years, a considerable amount of work has been put into experiments involving the alternative bonding techniques adhesive bonding and anodic bonding. Over the years, the group has acquired many pieces of equipment, a significant example being the Cascade 12K Summit semi-automatic wafer prober [104]. The combination of these pieces of equipment means that it becomes possible to analyse a greater number of fabricated devices, allowing for broader analyses of parameters and assessment of wafer level variation of those parameters, instead of single array, even single element, evaluation, which tends more towards proof-of-concept demonstrations.

The other side of the coin is the necessity to be able to deliver a suitable volume of CMUT arrays. In combination with the desire stated in Section 2.3.1 for the fabricated RCA arrays to become larger for increased ultrasound resolution, the projected RCA array designs would reach a size where only a single array would fit on a 4" Si wafer. As arrays become larger, the fabrication process becomes increasingly more sensitive to particle contamination simply due to the array footprint. One way to mitigate that issue is process on larger Si wafers. Combined with the tendency for SOI wafer manufacturers to shut down production of 4" SOI wafer production due to the majority of their industry clients only using larger wafers, meant that a transfer of our wafer bonding process to suit a 6" production line will be necessary. However, currently 4" Si wafers are the largest wafer size the majority of cleanroom equipment available to us today will allow. With adjustments to the processes, most fabrication steps could be transferred to a 6" Si wafer process, with the unfortunate, critical exception of the wafer bonder.

The solution we pursued was to bond directly in hand, without the use of any wafer bonder, referred to as 'Hand-bonded' in the following. This also means that these results do not only provide a solution to a local issue of equipment availability. In some cases you might not even need a wafer bonder, which could otherwise cost several millions of Euro; you might be able to fusion bond wafers

together, without having to purchase additional equipment. These results were presented at MUT 2018, with great interest due to some of the participants now being able to move into other CMUT fabrication methods that they had considered impossible from an equipment availability standpoint. Hand-bonding allows the bond to be formed under high efficiency particulate arrestance (HEPA) filters directly after cleaning of the wafers, without needing any form of transportation of the cleaned wafers, for a decreased risk of particle contamination of the bond interface. The outcome of this study has been that all of our bonding processes have been changed, with the purpose of decreasing the particle contamination.

4.1.1 Fusion bonding

Wafer bonding is a common processing technique for combining multiple wafers into a single structure. It can be used to stack structures which would otherwise not be possible to combine by epitaxy or film deposition. Examples of wafer bonding applications are in the production of SOI wafers and in the formation of sealed cavities, such as in pressure sensors [105] or in CMUTs [106, 100].

The original fusion bonding method was first described in the literature in 1985 – 1986 [107, 108]. Fusion bonding, also called direct bonding, is typically considered a three step process: pre-treatment of the wafers; pre-bonding of the wafers; annealing of the bond.

The pre-bond is formed when the surfaces of two wafers are placed in contact. If the surfaces are sufficiently smooth, the bond will form immediately at the point where the two wafers are placed in contact, and spread as a wave from the point. Si wafers which have been exposed to the oxygen in the air will form a native oxide on the surface, effectively covering the surface in hydrophilic silanol groups, $-\text{SiOH}$. These groups will form hydrogen bonds either to the oxygen atom of the silanol group on the opposing wafer, or to water molecules which are trapped in the interface [109]. The wafer bond is at this point strong enough that force needs to be applied in order to separate the wafers from each other again. If the wafers are contacted simultaneously near the edge of the wafers, the bonding wave front will move inwards, and likely trap air in voids. If instead the wafers are contacted with a small amount of force in the centre of the wafers, the bond front will move outwards, effectively pushing air radially out, significantly decreasing the risk of voids due to air pockets. Subsequent to the conclusion of the work on hand-bonding, a device which would only apply the pre-bond contact force in the middle of the contacted wafers was designed and milled in aluminium. The model can be seen in Appendix H.1.

The bonding method is extremely sensitive to the surface roughness, requiring a roughness of less than 1 nm or 0.5 nm for conventional wafers thicknesses of 500 μm [102, 110]. Due to the remaining surface roughness, the wafers will still only be locally in contact, with many unbonded micro-areas. It has been shown however, that for extremely thin wafers, with thicknesses ranging from 2 – 200 μm , fusion bonding can be achieved even when the surface roughness is 10 – 50 nm, likely due to the flexibility of the thin wafers [111].

For the same reason, the method is also extremely sensitive to particles. That is the reason for the pre-treatment step, which is a cleaning step, either performed as a plasma clean or a wet chemical cleaning, often an RCA clean, which here is short for Radio Corporation of America, the corporation at which it was invented [112]. The cleaning process consists of two main cleaning solutions, one which removes organic films, particles, and some metals, the other which removes heavy metals, alkalis, and metal hydroxides, combined with intermittent buffered hydrofluoric acid (BHF) etching to remove chemically grown SiO_2 layers.

The final step is the annealing of the interface bond, gradually increasing the bonding strength over time. Between 100°C and 200°C, covalent bonds between the silanol groups are formed, creating Si-O-Si bonds at the interface [113]. Beyond this temperature, no changes are observed in terms of the bonding strength. The bond-strength is limited by the actual contact area of the wafers, and thus by the unbonded micro-areas. At 800°C the native SiO_2 layers become viscous enough that they start to fill the locally unbonded micro-areas, making a more complete bonding across the interface, therefore also increasing the bonding strength. At 1100°C, the viscous flow of the SiO_2 will complete the bonding [109]. The new structure is found to be practically equivalent

to a single bulk silicon substrate [108]. Transmission electron microscopy (TEM) images of the interfaces showed an epitaxial-like lattice continuity, apart from local dislocations, which might be explained by smaller misalignments between the crystal orientations of the two wafers, combined with gradual rearrangement of atoms near the interface at the high temperatures [108].

Fusion bonding of hydrophobic surfaces can also be achieved, for instance if a SiO_2 layer is removed by BHF. The Si surface will then mainly be terminated by hydrogen (H) and fewer fluorine (F) atoms. The F terminated regions are particularly reactive to water, which would form silanol groups again and make the surface hydrophilic again. When two such wafers are contacted, van der Waals bonds form between the H on the opposite wafers. Furthermore, the few F terminated regions might form hydrogen bonds to other formed silanol groups. Due to the interwafer bonding stemming primarily from van der Waals bonds, the pre-bond strength is significantly lower for a hydrophobic surface, than for a hydrophilic surface [113]. Annealing of the structures does not increase the bonding strength before 400°C , at which point the H atoms desorb from the wafer surfaces, and diffuse into the silicon substrates, into unbonded micro-areas, or out along the bonding interface, while Si-Si bonds are formed instead. Note that the H atom is so small that it is capable of diffusing into or through structures which would for other gasses be impermeable. The structures reach the cohesive strength of bulk Si at about 700°C . The bonding of hydrophobic silicon surfaces results in a clean interface with no boundary SiO_2 layer. On the other hand, the temperatures are not high enough that the Si atoms rearrange significantly, as is the case for SiO_2 which would become viscous and fill out the micro-gaps. As a consequence, the interface of hydrophobic wafers contains many nanometre sized voids.

In the original work by Shimbo et al. [108], two hydrophilic, mirror-polished silicon wafers were fusion bonded by being brought in contact in a clean environment at room temperature. The wafer stacks were then heated to 1000°C , which completed the bonding process. Since the initial introduction, the overall processing has not changed; it still consists of a cleaning and/or surface activation process; a pre-bond in which the wafers to be bonded are placed in contact, typically with an applied pressure on the wafer stack; and a subsequent high-temperature bond-anneal. However, sophisticated wafer bonders have been made, which are typically used during the pre-bond. Such pieces of equipment allow for more advanced bonding methods, including alignment bonding when structures on the top and bottom wafers should be aligned. They can also provide control of the bonding conditions, in terms of the bonding temperature, the pressure applied on the wafer stack during the pre-bond, and the pre-bond atmosphere. Additionally, it has been shown that it is possible to fusion bond other materials together, such as SiO_2 to silicon [114], Si_3N_4 to Si_3N_4 [115] or in fact any combination of the three materials. The potential to obtain a reduced cavity pressure, by bonding in a vacuum, is in some cases used as an integral part of the functional device [105, 106, 100]. The three step process is generally considered collectively. However, only limited investigations have been conducted of the intermediate state of the bonded structure, for instance after the pre-bond and before the bond-anneal. Typically, the studies have focused on how the bond strength increases with higher annealing temperatures and/or longer annealing times [109, 116], with only a few considerations of how the bond-anneal influences the resulting atmosphere inside fabricated cavities in a device. Harendt et al. [117] investigated whether bonding to wafers structured with cavities would change the bonding strength, time and number of voids, compared to full wafers, but found overall similar results. They analysed the resulting gas content in the cavities by mass spectrometry, and found that it depended on the annealing temperature, and argued that their findings of lower water content subsequent to higher annealing temperatures showed that the water and oxygen were oxidizing the silicon surfaces. No investigations on how the pre-bond atmosphere influences the final atmosphere inside fabricated cavities in a device.

The success criteria for the bonding process would be that the bond strength is high, and that the final bond is air-tight. The following sections will show that it is possible to obtain a reduced cavity pressure, without bonding in a vacuum. The results not only show that the bond anneal is the essential process in determining the final cavity pressure, they also indicate that even though the pre-bond is conducted in a vacuum chamber, the resulting cavity pressure is not a vacuum. Essentially, this means that unless alignment bonding or elevated pre-bond temperatures

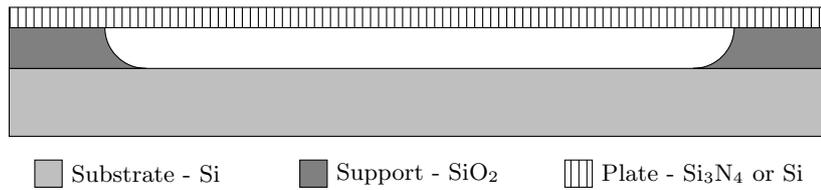


Figure 4.1: Cross section of a cavity device consisting of a substrate wafer, a support structure defining the cavities, and a plate which is bonded on top.

are required, a wafer bonder will not be necessary for obtaining reduced pressures in fusion bonded cavities.

4.2 Materials and methods

4.2.1 Experimental design

Given the success criteria of a strong bond, and an air-tight interface, test structures were made to determine any differences in the cavity pressure after the bonding process. Simple wafer bonded cavity test structures enable indirect determination of the cavity pressure, by measurement of the deflection of a plate suspended over the cavities in an ambient environment. A cross section of such a device can be seen in Figure 4.1. The cavity would be defined in a SiO_2 thin film on a Si wafer.

The centre deflection of an isotropic circular plate, w_0 , can be expressed according to [118]

$$w_0 = \frac{3}{16} \frac{(1 - \nu^2)^2 a^4}{E h^3} \Delta p, \quad (4.1)$$

where ν is Poisson's ratio, Δp is the pressure difference across the plate, a is the radius of the plate, E is Young's modulus, and h is the plate thickness. Hence, any difference in plate deflection between the devices is proportional to the difference in cross-plate differential pressure, and thus, to the cavity pressure as well.

To test the effect of the pre-bond environment on the resulting cavity pressure, four bonding conditions were compared: three formed inside a wafer bonder and one formed directly by hand (hand-bonded). In the wafer bonder, the atmospheric environment was changed between 2×10^{-4} mbar (Vacuum), atmospheric air at 1 bar, and argon at 1 bar. Assuming a perfect seal of the cavities, the three different atmospheres should result in different cavity pressures for the final fabricated devices. For the devices bonded in a vacuum, the cavity pressure should be 0 bar, and $\Delta p = 1$ bar when the ambient pressure is 1 bar. For the devices bonded in 1 bar of argon $\Delta p = 0$ bar, as the argon atmosphere is inert and should remain intact. For the devices bonded in air $\Delta p \approx 0.2$, since air is composed of 78% nitrogen, 21% oxygen and 1% argon, of which the 21% oxygen will be consumed in oxidation of any silicon surfaces of the cavities during the high-temperature bond-anneal. The oxygen consumption has previously been described in [119, 117]. Finally, for the devices bonded directly in hand $\Delta p \approx 0.2$ as the atmospheric environment is the same as that of the air devices. As the plate deflection is linear in pressure, these differences in Δp should correspond directly to the relative differences in plate deflections. The maximum deflection is expected for the Vacuum devices, whereas the Air and Hand-bond devices would only deflect one fifth of the Vacuum devices, and the Argon devices should not deflect at all. These expectations are illustrated in Figure 4.2.

4.2.2 Material choice - silicon nitride plates

There are a number of ways to fabricate fusion bonded cavities, as fusion bonding can be made with the combination of any two substrates with a surface of either silicon, SiO_2 , or Si_3N_4 . For

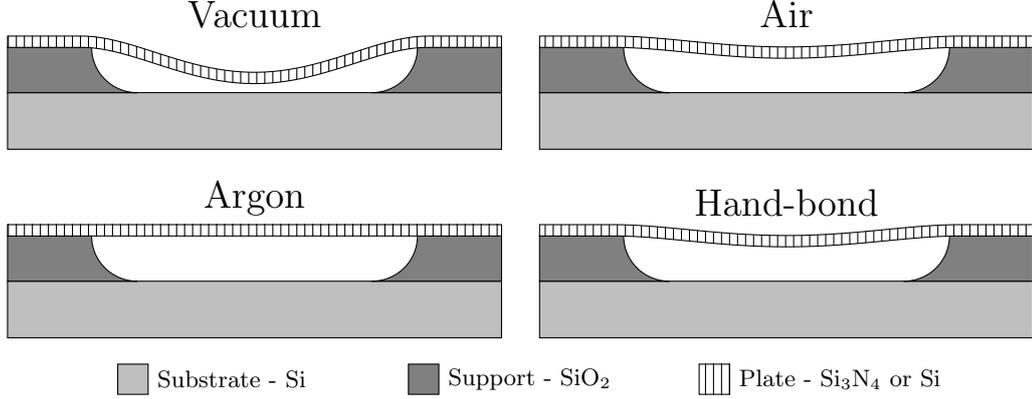


Figure 4.2: Schematic of the expected deflections for the four different bonding conditions.

ease of fabrication, the cavities are etched in a SiO_2 layer, which is grown on a silicon wafer. This provides control of the cavity depth when using a selective wet etchant, due to essentially an etch stop once the etchant reaches the silicon below the SiO_2 . The plate can then be fabricated using either silicon, SiO_2 , or Si_3N_4 . To obtain a device layer of a few μm as would be required for the chosen design when using silicon as the plate material, commercial SOI wafers are typically thinned by chemical mechanical polishing. This processing unfortunately results in a thickness variation between 300 nm and 500 nm, providing large variations in deflection as Equation (4.1) scales with h^3 . Therefore, two alternative types of plates were considered. Firstly, a SiO_2 thin film plate grown on a silicon wafer. SiO_2 would provide control of the layer thickness and uniformity, but also introduce built-in compressive stress when grown on a silicon substrate. This could result in buckling of the plate with a direct influence on the deflection measurements. Secondly, a Si_3N_4 thin film plate deposited on a silicon wafer. Low pressure chemical vapour deposition (LPCVD) of Si_3N_4 layers can, similarly to growth of SiO_2 films, provide good control of thickness and uniformity compared to the SOI wafers. There will be a built-in stress, but as it is tensile, it will not result in buckling, and is therefore acceptable.

To be able to compare the deflections of the different bonded structures illustrated in Figure 4.2, the variability of the thickness and the stress in the Si_3N_4 films is a critical parameter. As described in the introductory Chapter 3.2.1, Engholm et al. [80] showed that both thickness and stress will influence how much a plate deflects, describing the centre deflection, w_0 , of a plate with a built in tensile stress as

$$w_0 = \frac{\Delta p a}{2} \sqrt{\frac{C D}{N_t^3}} \left(\frac{1 - I_0\left(\sqrt{\frac{N_t}{C D}} a\right)}{I_1\left(\sqrt{\frac{N_t}{C D}} a\right)} \right) + \frac{\Delta p a^2}{4 N_t}, \quad (4.2)$$

where C is a constant based on the solution to the plate equation, D is the flexural rigidity of the plate, I_n is the modified Bessel function of first kind, $N_t = \sigma h$ is the stress resultant, where σ is the planar biaxial stress in the plate. p has been exchanged with a Δp since for this experiment, we are only considering a cross-plate pressure difference as the load. Since it is not possible to change the bonding atmosphere or method of pressure application locally on a single wafer, the comparison of the four bonding conditions will necessarily need to be between devices fabricated on separate wafers. Therefore, it is essential that the inter-wafer variability in thickness and stress of the plate is not so large, that it makes distinguishing between the expected differences in deflection across the different wafers impossible.

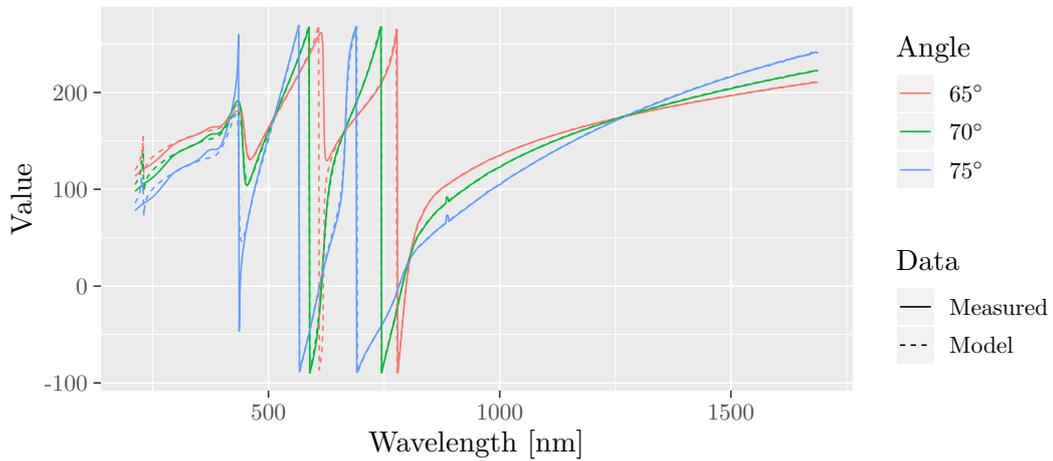
The first part is to be able to determine the processing parameters such as processing temperature and time, such that the desired film thickness is obtained. The theoretical foundation of thin film deposition and growth based on the underlying physics has been described in the literature

many times [120, 121], with models such as the Deal-Grove model for thermal oxidation, and linear deposition models for LPCVD deposition processes. On top of that, many additions or corrections to these models have been proposed [122]. Based on these, simulation software packages have been made, to allow for simulation of complex process flows, with the desired parameters [123]. However, while the theoretical models provides a great insight into the physical principles, it might be difficult to transfer those parameters to a deposition furnace, or vice versa, due to a number of inevitable experimental deviations from theory. In practice, once the furnace is in stable operation, it might be better to use data from a processing log, to predict the necessary processing parameters. A program was built to determine the necessary processing parameters based on a desired outcome. The details of the script are discussed in Appendix F.1. The script takes the film thickness as an input, and provides the processing time required to obtain that thickness, along with the residual error of the statistical model based on the log data from the LPCVD Si₃N₄ furnace used. The data is based on measurements of a new wafer centrally placed in the quartz boat during every process. Thus the residual error marks the variability of the film thicknesses from deposition to deposition. The issue with using processing logs is that the furnace might change behaviour over time, meaning not all data is necessarily equally representative of what to expect from the furnace when it is actually going to be used. Therefore, it is possible to input exactly how many of the log entries to include, resulting in only that amount of the newest entries being analysed. Another issue is the human factor, with users occasionally inputting bad data in the process log. This is tested by iteratively analysing the residuals of the model, and discarding outlier values which do not fit the distribution of outliers sufficiently well, before remodelling the remaining data. For the furnace used, the residual error of the mode is ≈ 7 nm, which from experience is found to mirror the expected outcome film thickness well.

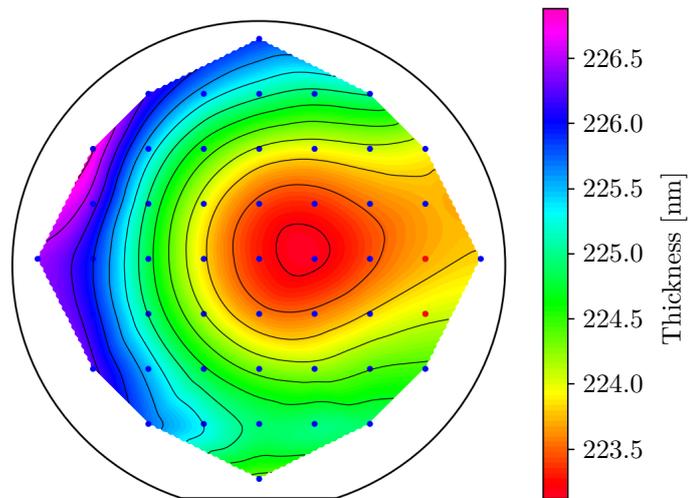
To determine the variability across a full quartz boat of Si wafers, 15 Si wafers were placed in the LPCVD Si₃N₄ furnace. After deposition, the film thicknesses were measured. 49 measurements distributed across the wafer surface were taken on each wafer using a M2000XI-210 ellipsometer (J.A. Woollam Co., Inc., Nebraska). An example spectrum can be seen in Figure 4.3(a). The solid lines are the measured data, the dashed lines are the theoretical spectrum of a thin film and the different colours are the data from different angles of incident. The ellipsometer determines the thickness of the film by comparing the reflected light spectrum to the corresponding spectrum based on a theoretical model, for a given thin film and angle of incidence. It then modifies the thickness of the theoretical model, in search of a minimum in the mean square error (MSE) between the measured spectrum and the theoretical spectrum. The determined film thickness is thus the thickness of the model when the MSE reaches a minimum.

An example of a thickness wafer map based on 49 measurements can be seen in Figure 4.3(b). The ellipsometer also provides a map of the MSE values. These values can be used to estimate the correctness of the corresponding thickness estimate. The absolute MSE values will vary depending on the thickness and type of film being measured. Therefore, it is not possible to set a general MSE value as a threshold for outlier detection. However, if the film is perfectly uniform, it would be expected that the MSE values would all be very similar. By assuming that the small variations in MSE follow a normal distribution, outlier detection can be based on extreme MSE values. The details of the outlier detection script can be seen in Appendix F.2. The measurement positions are marked by dots in Figure 4.3(b). The red dots mark outlier measurements and have been discarded in the image and analysis. The contours and gradients are determined from interpolation between the valid thickness measurements.

Figure 4.4 shows how the LPCVD Si₃N₄ film thickness varies across a full quartz boat of silicon wafers after a single batch process. Each box in the box-plot represents the individual wafers in the quartz boat, and consists of the 49 thickness measurements of a wafer map similar to Figure 4.3(b), thereby showing the intra-wafer variability. Similar characterisations of other furnaces can be seen in Appendix F.3. The line in the middle of the box is the median value, and the lower and upper edge of the boxes correspond to the 25th and 75th percentile of the data respectively, with outliers marked as dots. Outliers are defined as measurements further than 1.5 times the inter quartile range (ICR) away from the nearest box edge, where ICR being the distance between the 25th and



(a) Ellipsometer model and measurement



(b) Measured thickness wafer map

Figure 4.3: The thickness measurements from the ellipsometer is based on a satisfying correlation between the measured spectrum (solid lines) and the corresponding model (dashed lines) of a thin film as specified in the text (a). The measurements can then be collected into a wafer map (b). The dots mark the measurement positions. Blue dots mark used points, red dots mark outlier points. The underlying contours are based on interpolations of the blue dots.

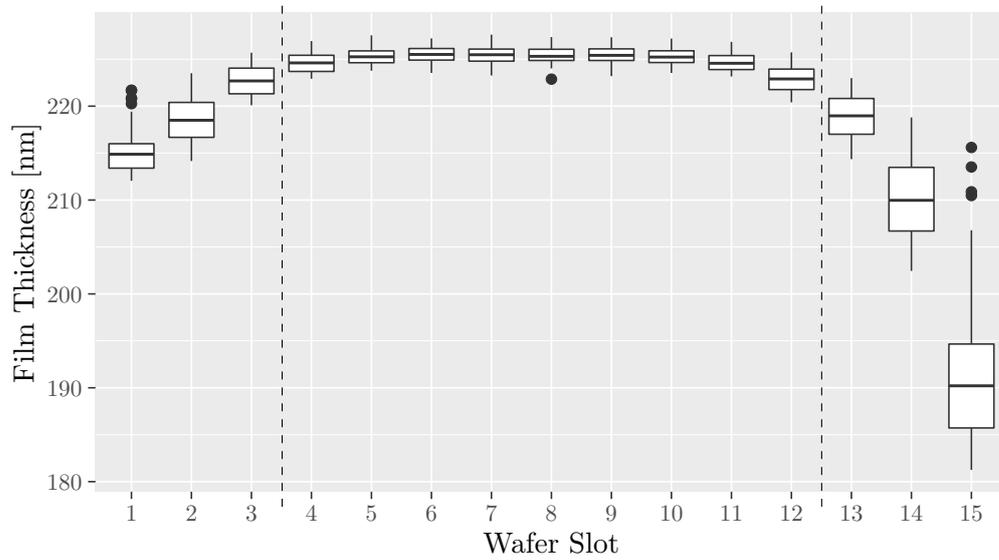


Figure 4.4: Box-plot of the film thickness distribution across a furnace boat of Si_3N_4 films. The dots mark outlier thickness values. The film thickness is $225.0 \text{ nm} \pm 0.8 \text{ nm}$ between the dashed lines.

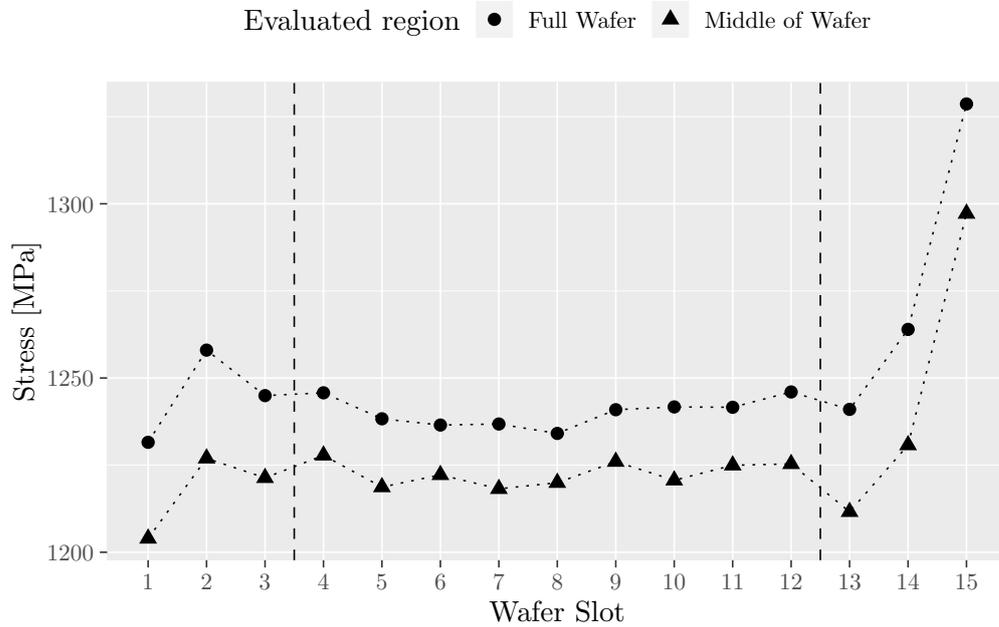


Figure 4.5: Stress distribution across a furnace boat of Si_3N_4 films. The dashed lines indicate the same region of the wafer boat selected in Figure 4.4. The stress measured near the middle of the wafer within that region is $1223 \text{ MPa} \pm 4 \text{ MPa}$.

the 75th percentile. The stress of the same wafers has been calculated from the wafer curvature, through the Stoney equation, and is shown in Figure 4.5, in two separate data ranges, either measured across the full wafer diameter, or measured only across the central part of the wafer, spanning the region in which the devices to be tested are placed.

Using Equation (4.2), it is possible to determine the expected deflections of the four different types of devices, and more importantly, the smallest expected difference between the devices. The expected differential pressure across the plates for the four different bonding conditions were 1 bar (Vacuum), 0.2 bar (Air), 0 bar (Argon) and 0.2 bar (Hand-bond), which means that the smallest difference in differential pressure between any two devices would be 0.2 bar. Thus, for the devices with a cavity radius of 32 μm , the smallest difference in deflection for the average values of plate thickness and plate stress can be calculated using Equation (4.2) to be

$$\Delta w_{0,\min} = 17.5 \text{ nm.}$$

The inter-wafer variability of the data in Figure 4.4 and Figure 4.5 can be used to estimate how large the expected variations in deflection are. Assuming the sources of variation are independent and random, the propagation of error in the centre plate deflections, δw_0 , can be estimated using Equation (4.2) to calculate

$$\delta w_0 = \sqrt{\left(\frac{\partial w_0}{\partial h} \delta h\right)^2 + \left(\frac{\partial w_0}{\partial \sigma} \delta \sigma\right)^2}, \quad (4.3)$$

where δh is the uncertainty in the plate thickness, and $\delta \sigma$ is the uncertainty in the stress, and it is assumed that h and σ are the only varying parameters [124].

This expression can also be used to see how much variation in plate thickness and stress will be tolerable while still providing less than 17.5 nm of variation in the resulting plate deflection. However, given that the plate deflection uncertainty stems from the contributions from both thickness uncertainty and stress uncertainty, which are assumed independent, it is difficult to provide reasonable limits for the combination of the parameters. However, the limiting scenarios can be calculated, by observing how large the uncertainty of one of the parameters can be while the uncertainty of the other is set to 0. The largest pressure difference will result in the largest deflections and therefore also the largest uncertainty estimates. Thus, the plate deflection uncertainty estimates are calculated for $\Delta p = 1$ bar. Setting the resulting plate deflection uncertainty to 17.5 nm and the stress uncertainty to 0 MPa, the tolerable plate thickness uncertainty will be 43.7 nm. Similarly, setting the resulting plate deflection uncertainty to 17.5 nm and the plate thickness uncertainty to 0 nm, the tolerable stress uncertainty will be 249 MPa. Combining the uncertainties will of course mean that the actual tolerable uncertainties will be smaller. However, this provides a sense of the tolerable scale of the uncertainties.

Referring to Figure 4.4 and Figure 4.5, it is clear that apart from the outermost quartz boat positions, all thickness measurements and stress calculations are presumably within acceptable limits. However, the three last wafers at each end of the boat shows significantly larger variation than the central part of the boat. To minimize the variability, only the wafers in the slots between the dashed lines were used for the experiment. This allows for two wafers to be fabricated with each of the bonding conditions (in addition to a required single furnace processing test wafer). Using the standard deviations of the measurements between the dashed lines in Figure 4.4 and Figure 4.5 as estimates for the uncertainty, the actual expected uncertainty in the plate deflection can be calculated. The average thickness between the dashed lines is $225.0 \text{ nm} \pm 0.8 \text{ nm}$ and the average stress is $1223 \text{ MPa} \pm 4 \text{ MPa}$. In both cases, the uncertainty is the standard deviation. For the largest pressure difference ($\Delta p = 1$ bar), the uncertainty in plate deflection will be

$$\delta w_0 = 0.4 \text{ nm.} \quad (4.4)$$

This difference from the processing uncertainties is much smaller than the expected difference between devices. Consequently, the expected deflection differences due to different cavity pressures should be distinguishable when choosing Si_3N_4 as the plate material.

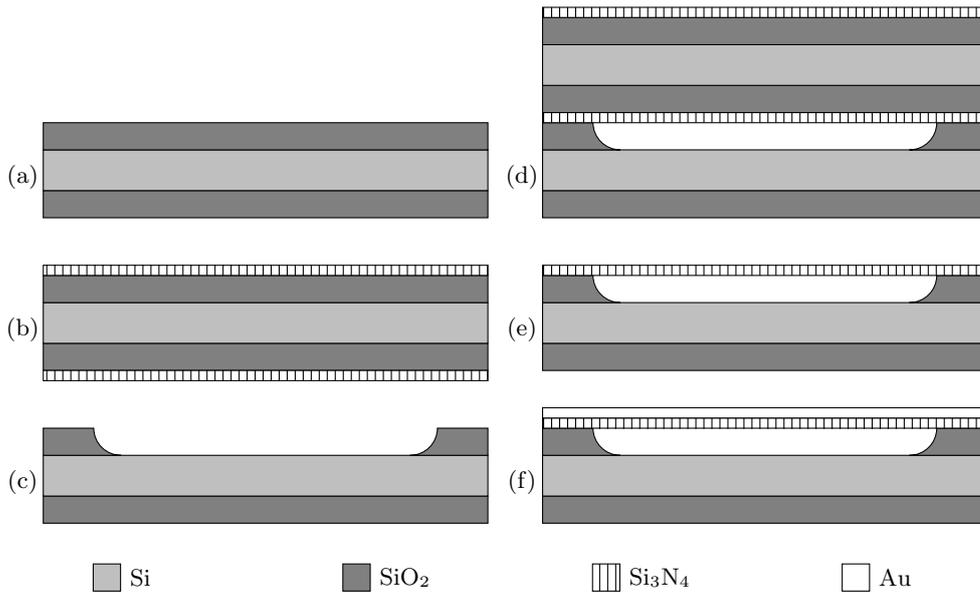


Figure 4.6: Process flow for the fabricated test devices.

4.2.3 Fabrication of test devices

An illustration of the process flow can be seen in Figure 4.6. A $405 \text{ nm} \pm 0.5 \text{ nm}$ SiO_2 layer was grown on a batch of both single side polished four inch silicon (100) wafers and double side polished four inch silicon (100) wafers in a dry thermal oxidation process at $1100 \text{ }^\circ\text{C}$ (a). The single side polished wafers were used as substrates, in which the cavities were to be etched. The double side polished wafers were used as support substrates for the plate layers. The plate wafers were transferred directly to an LPCVD furnace for deposition of a $226 \text{ nm} \pm 0.8 \text{ nm}$ Si_3N_4 layer (b). After the deposition, the plate wafers were transferred to an oxidation furnace for oxidation of the Si_3N_4 layer, which has been shown to improve the bonding strength between SiO_2 and Si_3N_4 [115]. The plate wafers were left inside the furnace until needed for bonding to minimize particle contamination. The circular cavities on the substrate wafers were defined in a lithographic process with a radius of $a = 32 \text{ }\mu\text{m}$. They were then etched in a wet BHF etch to define the 405 nm deep cavities (c). The substrate wafers were RCA cleaned to remove particles, directly after which the substrate and plate wafers were fusion bonded together, under the four different bonding conditions (d). A Süss SB6 wafer bonder (Garching, Germany) was used to bond the non-hand-bond devices. The Vacuum devices were bonded at a pressure of $2 \times 10^{-4} \text{ mbar}$, the Air devices were bonded without pumping down the chamber, and the Argon devices were bonded in an argon atmosphere at a pressure of 1 bar . All of the devices bonded in the wafer bonder had a 600 mbar pressure applied on the wafer stack during the pre-bond. All bonds were made at room temperature. After the pre-bond, all bonded structures were annealed at $1100 \text{ }^\circ\text{C}$ in 1 bar of N_2 for 3 hours. The bonding interfaces were then characterized by infrared reflectance measurements using the infrared photoluminescence system Accent RPM2000 Compound Semiconductor Photoluminescence System (Nanometrics, Massachusetts) to check for voids. The system maps a wafer by emitting infrared light of a given wavelength. The light is then reflected on the sample, back towards a photo-sensor, in which the reflected intensity is converted to an electrical signal. The reflected intensity is then represented by the voltage generated by the sensor in millivolts. The light is incident normal to the wafer surface, and based on complex analysis of light interaction at interfaces, light will be reflected at each material interface, and the power reflectance will be given

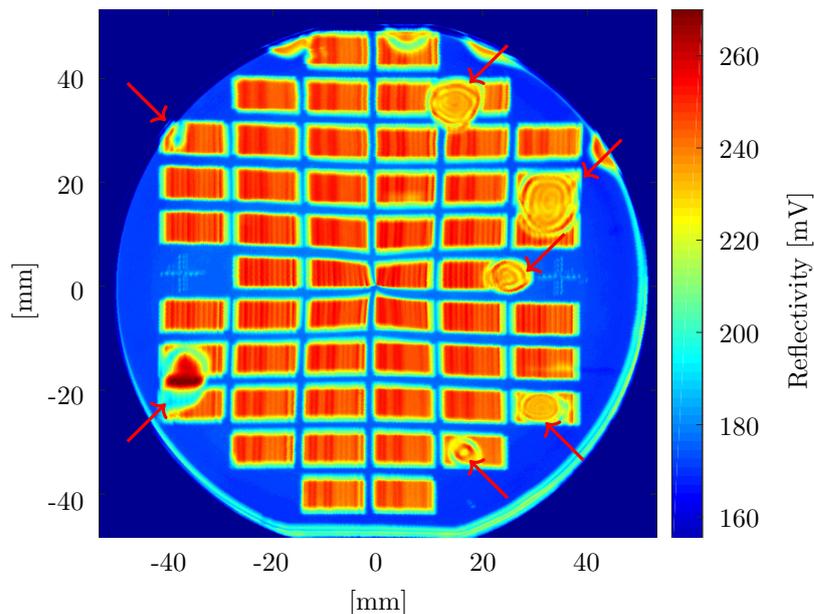


Figure 4.7: Infrared reflectance map of one of the device wafers bonded directly in hand. The red rectangles are CMUT arrays. The arrows point toward interface voids. Most of the interface is void-free.

as

$$R = \left(\frac{n_1 - n_2}{n_1 + n_2} \right)^2, \quad (4.5)$$

where n_1 and n_2 are the refractive indices of the materials of the interface. Thus, the larger the contrast in refractive index, the larger the reflectivity. The bonded interface is between silicon, which has a refractive index of ~ 3.45 , and SiO_2 , which has a refractive index of ~ 1.45 . In the case that air is present in the interface micro-cavities, the refractive index will be ~ 1 , which gives a larger contrast to the silicon, and therefore a larger reflection. As a consequence, voids will be clearly visible in a photoluminescence (PL)-map.

An infrared reflectance map for one of the hand-bonded wafers can be seen in Figure 4.7. A few voids can be seen, marked by red arrows, but most of the interface has been properly bonded. The pincushion distortion of the data, seen towards the centre of the wafer as well as horizontally along some of the outer arrays, is expected to be an artefact of the scanning method, which consists of scanning a laser on the surface, while rotating the wafer at a controlled angular velocity. The handle layer of the top wafer was etched away using a sequential combination of dry etching and wet KOH etching to release the nitride plates (e). Finally, a layer of gold was sputtered on top of the wafer to increase the reflectivity of the surface for the subsequent analysis (f). The layer thickness was not determined directly, but a similar process has previously shown a layer thickness of approximately 10 nm. Being much thinner than the thickness of the nitride plate with built in tensile stress, the gold layer is not expected to modify the plate deflection significantly.

The PL-map can provide an indication of the bond strength, in the sense that if the PL-map does not appear uniform or with larger issues of voids, the bond-strength will likely not be high. However, a better indication is step (e) in the process flow, the dry etching of the handle layer, since this is a rough process. In the case of poorer bonds, the top wafer will often detach at the bond interface early in the etching process. Thus, the fact that it was possible to etch the hand-bonded wafers down using this method indicates a strong bond. Whether it is as strong as the devices bonded in a wafer bonder is not clear, but it is strong enough for CMUT processing.

The wafer layout is illustrated in Figure 4.8. (a) shows the CMUT array layout, with all the arrays in red, and contact pads in black. The zoom-in in (b) shows how each array consists of

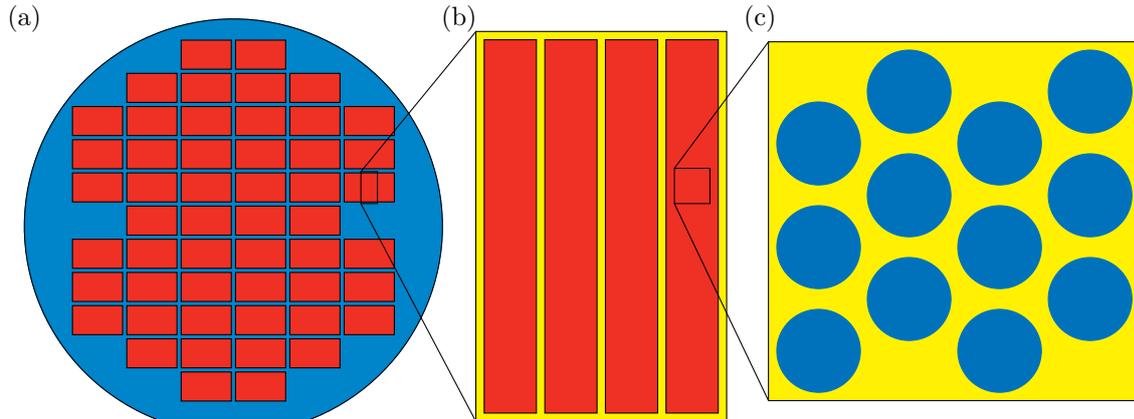


Figure 4.8: Overall device layout. (a) shows a top down view of the CMUT array (red) layout on the wafer (blue). (b) shows how each array consists of multiple elements (red). (c) shows how each element contains several CMUT cells (blue).

several elements, marked by red. Finally, the zoom-in in (c) shows how each element consists of many circular CMUT cells, marked in blue. The colours of (a) and (b) have been chosen to match the PL-map in Figure 4.7, and the colours of (c) are chosen to match Figure 4.9 and similar images.

4.3 Results

Deflection measurements

The plates on eight different CMUT arrays on each wafer were measured using the Sensorfar PLu Neox Optical Profiler (Sensorfar, Terrassa, Barcelona) to determine their deflections. An example of such an optical height distribution is shown in Figure 4.9 for a Vacuum devices. The blue regions are the deflecting plates. The corresponding histogram can be seen in Figure 4.10. These histograms are very distinct, and can be used to determine a systematic estimate of the deflection of each measurement. The highest probability density, P_{\max} , correlates with the area between the cavities, the yellow region in Figure 4.9. This peak in the histogram can be used to offset the data to align all measurements to the same reference point. For decreasing values (larger deflections) the density initially decreases rapidly, reaching a local minimum, P_{\min} , between -40 nm and -80 nm in the case of Figure 4.10 before increasing again slightly and finally dropping to zero. This non-monotonic behaviour means it is not possible to set a lower density threshold and use that to find the maximum deflection value, as it could result in the deflection value corresponding to P_{\min} . Also, choosing the lowest probability value increases the susceptibility to data outliers. By locating the first bin in the histogram with a value larger than the tenth quantile of the density data, and choosing the deflection corresponding to this as the deflection, it is possible to systematically determine the deflection of the plates near P_0 , while avoiding the risks listed previously.

Figure 4.11 shows a comparison of the height distributions of the four different types of bond conditions. All plates of the test structures deflect significantly and almost the same amount. The deflection data can be seen in Figure 4.12. It should be noted that the magnitude of deflections of all devices is large, regardless of the bonding conditions. This is a remarkable result, as the Argon devices were not expected to deflect at all. However, it seems that there are two groups, namely the devices bonded in the wafer bonder which all deflect ≈ 110 nm, and the hand-bonded devices which all deflect ≈ 60 nm. The Air devices and Hand-bond devices are directly comparable in terms of the pre-bond atmosphere, but the Air devices which were bonded in a wafer bonder deflect significantly more. These measurements indicate that whether the devices were pre-bonded

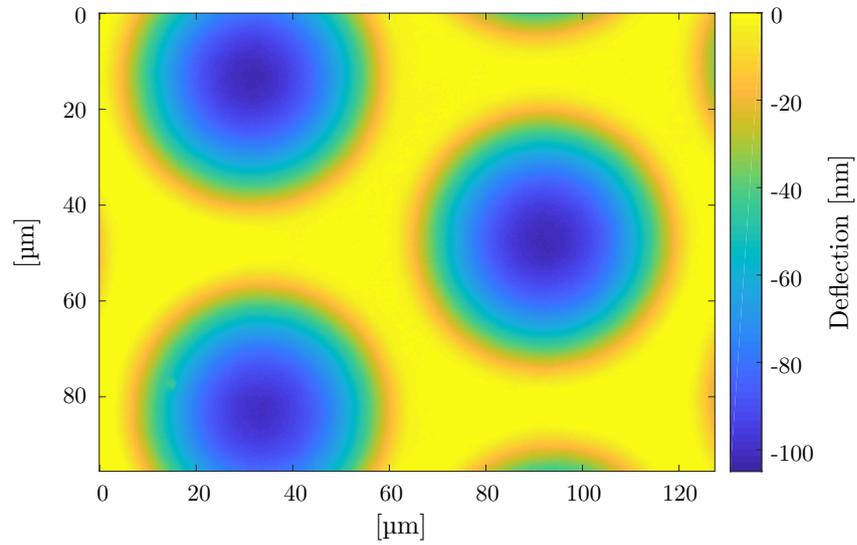


Figure 4.9: Optical profile of one of the Vacuum devices. The blue regions are deflecting plates.

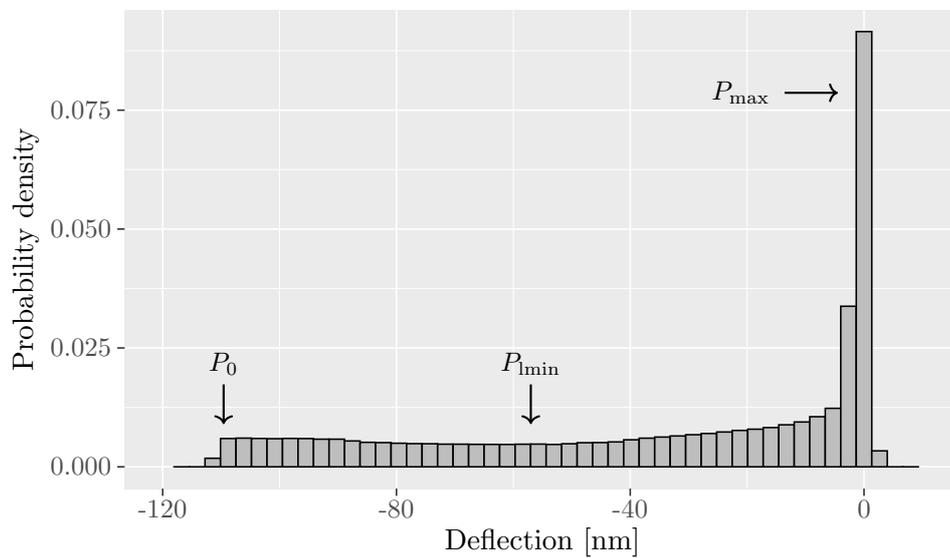


Figure 4.10: Histogram of the deflection data shown in Figure 4.9. The distinct behaviour, with points of interest marked, allows for automatic detection of the maximum deflection. P_0 marks the maximum (negative) deflection, P_{\max} corresponds to the top surface, seen as yellow in Figure 4.9, and P_{\min} is a local minimum in the histogram.

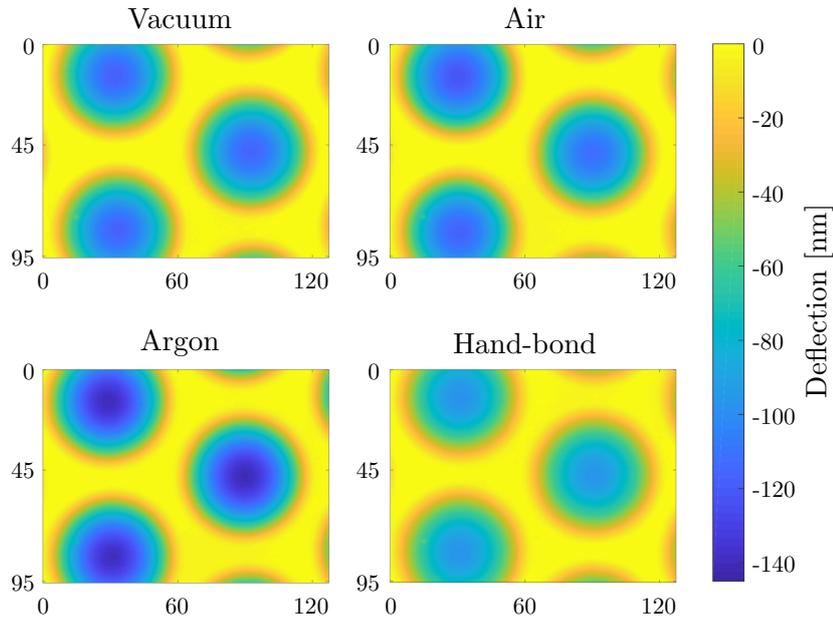


Figure 4.11: Comparisons of typical optical profiles of the four different bond conditions. The lateral dimensions are all in μm .

in a vacuum, in air, or in argon, is not critical for the final cavity pressure, and that the cavity pressures end up being similar regardless, but whether the pre-bond is done in a wafer bonder or in hand will have an effect. Although the deflection of the Hand-bond devices is lower than that of the other devices, they still deflect more than the expected fifth of the Vacuum devices. Finally, the magnitude of the intra-wafer variation on most of the devices as well as the inter-wafer variation, does not correlate with the uncertainty estimate presented in Section 4.2.2, which must mean that there is a source of variability not accounted for. However, it should not have an effect on the Argon devices, since there should be no pressure difference across the plate according to the hypothesis.

4.3.1 Hypothesis for bond interface diffusion

The expectations for the experiments presented in Section 4.2.1 were based on the cavities being sealed during the pre-bond. However, if the bond-interfaces are not leak tight after the pre-bond, a gas exchange between the cavities and the external environment can occur. During the bond-anneal, the temperature is increased to $1100\text{ }^{\circ}\text{C}$ in 1 bar of nitrogen. According to the ideal gas law, the high temperature will increase the pressure inside of the cavities by a factor of about 4.5. As illustrated in Figure 4.13, any pressure gradient will be able to drive gas diffusion between the cavities and the external environment, potentially equilibrating the pressures. For the Vacuum devices, the pressure inside the cavities is initially 0 bar, while the external pressure in the furnace is 1 bar. Therefore, gas will diffuse into the cavities during annealing. For the Air, Argon and Hand-bond devices, the pressure inside the cavities will initially be around 4 bar, and gas will diffuse out of the cavities during the anneal.

Once the bond-anneal is finished, and the cavities are sealed, the cavity pressure will be reduced by the same factor of about 4.5 when the temperature is returned to room temperature. This explains the large deflection of the Argon devices, when the hypothesis of a sealed cavity predicts no deflection at all. It also explains why most wafers deflect the same amount, due to the pressure inside the cavities having been equilibrated to the same value.

Both the intra-wafer and inter-wafer deflection variations which are seen could potentially be

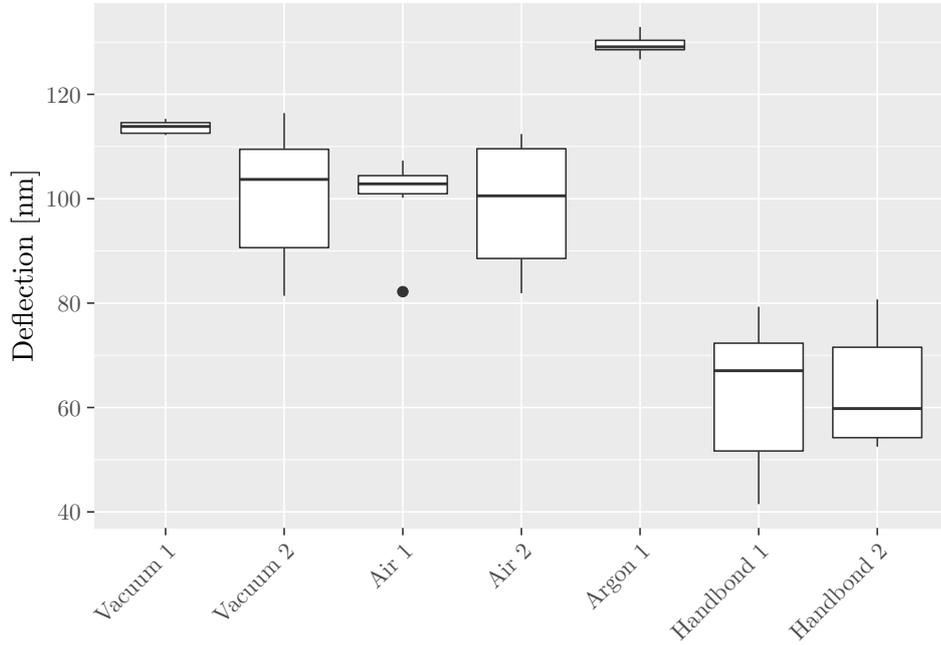


Figure 4.12: Box-plot showing the deflection measurements of each of the devices. Each box consists of measurements of eight different cavities distributed across a wafer.

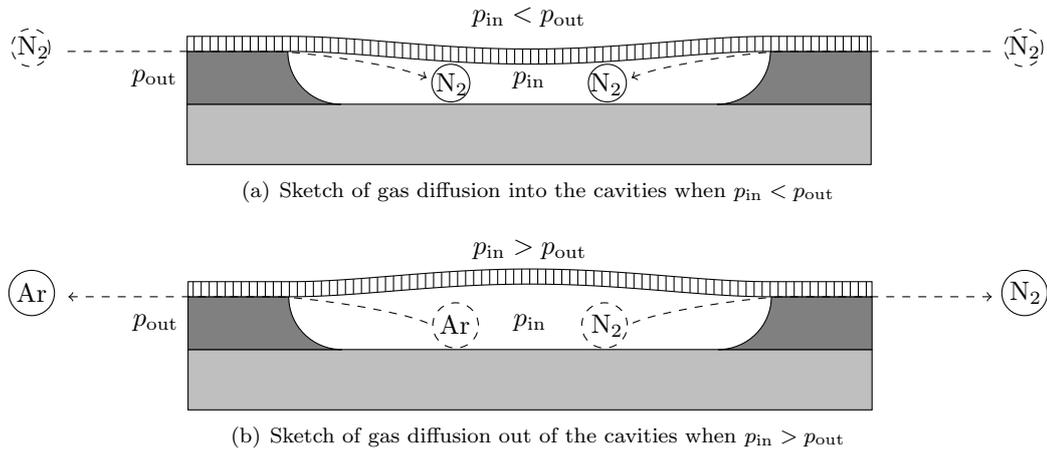


Figure 4.13: Sketches of gas diffusion during the annealing process, determined by the direction of the pressure gradient between the cavities and the ambient environment. (a) represents the Vacuum devices, where the pressure initially is larger outside of the cavities. (b) represents the Air, Argon and Hand-bond devices, where the pressure initially is larger inside the cavities. Note that the deflecting plates are used to illustrate the pressure variation. During the annealing process, the handle layer of the top wafer is still attached, meaning practically no deflection will occur.

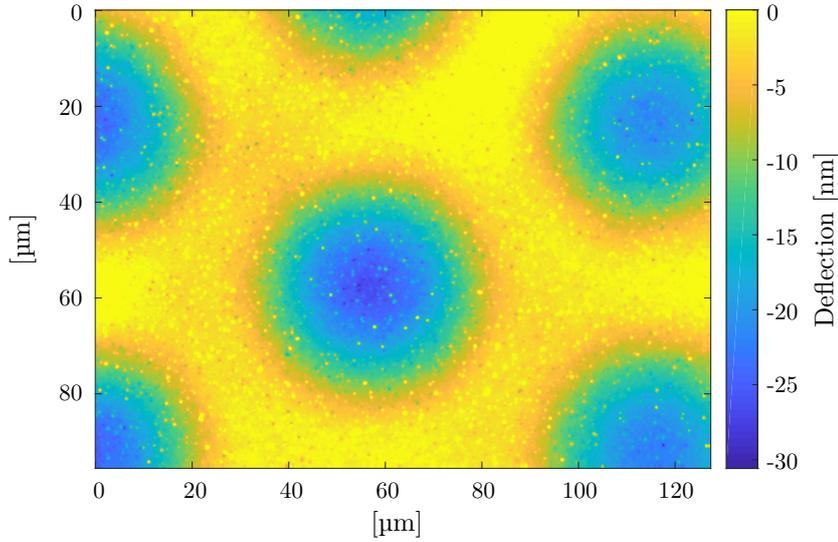


Figure 4.14: Deflection measurement of a hand-bonded cavity device fabricated with a silicon top plate.

explained by the sealing of the cavities being obtained at different points in time for the different cavities. If the gas diffusion has not managed to equilibrate the pressure, a sealing of the cavity would result in an off-equilibrium static pressure in the cavities. This could suggest that the pre-bond of the Hand-bond devices was stronger than that of the other devices.

Considering the experimental differences to Harendt et al. [117], these findings add new knowledge of the actual processes happening at high temperature annealing. Harendt et al. found water and oxygen content in the cavities based on low annealing temperatures, but none in those annealed at higher temperatures. They attributed the lack of water and oxygen to oxidation of exposed interface silicon at higher temperatures, however, the results presented here indicate that the lack of water content should be explained by a high interface diffusion rate of gases between the cavities and the furnace. At lower annealing temperatures, the diffusion rate will be lower as well, and the results by Harendt et al. indicate that it in this case will be too low for a complete exchange before the bonding interface is sealed.

4.3.2 Deflection test with a silicon plate

In order to test whether the out-diffusion of gas was a unique effect of using Si_3N_4 as the plate material, another set of wafers were fabricated, using SOI wafers as the top plate. Apart from the change of plate material and the metal being used as an etch mask, the process flow was the same as described in Section 4.2.3. The thickness of the SOI device layer was $3 \mu\text{m} \pm 0.5 \mu\text{m}$, as specified by the manufacturer. A height distribution of the surface after fabrication of a silicon plate device which was hand-bonded can be seen in Figure 4.14. The grainy appearance is not noise in the data, but roughness of the surface due to partial etching of the metal.

By using Equation (4.1), the centre deflection is expected to be $w_0 \approx 6 \text{ nm}$ when $\Delta p = 0.2 \text{ bar}$. To obtain the deflection of $w_0 \approx 30 \text{ nm}$ seen in the figure, it would be required that $\Delta p \approx 1 \text{ bar}$. This is similar to the Si_3N_4 devices, and shows that it is plausible that the same out-diffusion is obtained when bonding silicon to SiO_2 , as when bonding Si_3N_4 to SiO_2 .

4.3.3 Bond interface leak rate test

Deflection measurements do not reveal differences in leak rates of the bonding interfaces for the four different types of devices. In particular, such measurements do not indicate whether there is a

Table 4.1: Deflection measurements before and after 7 hours at 2 bar helium. The Before value has been measured five months after the bonding process. The Vacuum2 wafer was broken in half.

	Deflection after five months [nm]	Deflection after helium exposure [nm]	Deflection difference [nm]
Vacuum1	-113.2	-115.3	-2.1
Vacuum2	-134.4	+10.1	144.5
Air1	-105.4	-107.3	-1.9
Air2	-102.4	-103.5	-1.1
Argon1	-129.4	-128.4	1.0
Hand-bond1	-42.4	-43.2	-0.8
Hand-bond2	-45.6	-48.6	-3.0

difference in leak rate between the devices bonded in a wafer bonder and those bonded directly in hand. They only document a static situation in a very narrow period of time. A simple method for investigating the leak rates is to measure the deflection after long periods of time. Table 4.1 shows a collection of measurements conducted at different points in time. The first data column consists of deflection measurements conducted five months after the device fabrication. Considering the large variation in deflections observed for each device type in Figure 4.12, it is imperative that the exact same cavities are compared. Unfortunately, the deflection measurements conducted directly after fabrication were not logged with exact cell placement, leaving direct comparison to the deflections immediately after processing impossible. However, the values after five months align well with those presented in Figure 4.12, suggesting that the leak rate must be small. Another method of assessing the leak rate would be by helium leak testing. Helium is commonly used for testing of leak rates due to its small atomic size enabling faster propagation through porous structures, than would be the case for e.g. the molecules of atmospheric air.

A simple pressure chamber was designed to contain a single 4" wafer. Details of the pressure chamber can be seen in Appendix H.2. Each device was exposed to 2 bar of helium for 7 hours, and the deflections were measured within 30 minutes. Table 4.1 also lists the comparative measurements after the helium test, as well as the difference to the measurement before.

Positive differences in deflection means the cavity pressure has decreased, and that some amount of helium has reached the device cavities, as would be the expectation. A number of interesting points can be made from these measurements. Firstly, that most differences are small, suggesting a low permeability of the bonding interface. Secondly, that the Vacuum2 device behaves significantly different than the others. The measurements are highlighted in red font in the table. This particular wafer had broken in half during the latter half of processing, certainly setting it apart from the others, and likely influencing the bonding interface. However, the wafer had a large deflection after five months, indicating very little leak rate at that point. Only the exposure to helium changed the deflection significantly, demonstrating that a difference can actually be detected simply by changing the gas to helium. In addition, cavities next to the one investigated deflected similarly to the measurement conducted five months after bonding, showing that this is a very local effect, simply due to the breakage. Thirdly, most values are negative, corresponding to fewer gas particles being in the cavities after helium bombardment, which does not make physical sense. More likely, small changes in the exact region which is investigated before and after helium testing, influence the deflection estimate more than the leakage flow of helium. Thus, the listed differences in pressure are indicators of the uncertainty of each individual measurement. As a consequence, it does not make sense to calculate individual leak rates for the devices. However, by estimating that the true change in deflection is no larger than the measurement uncertainty, estimated to be the largest deflection difference in Table 4.1 namely 3 nm, it is possible to calculate an upper limit of the leak

rate. The simplest estimate of the leak rate, will be calculated as

$$L = \frac{dp}{dt}V, \quad (4.6)$$

where V is the cavity volume. Using Equation (4.2), a 3 nm deflection difference is found to correspond to a 34 mbar pressure difference. Combining this with the geometrical parameters of the cavity, namely a radius of $a = 32 \mu\text{m}$ and a height of $h = 405 \text{ nm}$, the upper estimate for the leak rate is becomes

$$L \leq 1.76 \times 10^{-15} \frac{\text{mbar L}}{\text{s}}. \quad (4.7)$$

This order of magnitude agrees with the literature [125], albeit having been determined through a different method.

4.4 Chapter summary

The presented results provide an overview of the resulting cavity pressure after fusion bonding. The fabricated test structures consisting of etched SiO_2 cavities, with Si_3N_4 plates fusion bonded on top, enabled measurements of the deflection of the nitride plates as an indirect measure of the cavity pressure. Four sets of bonding conditions were used, three in a wafer bonder in atmospheres of vacuum, air and argon, and the last set was bonded directly in hand in atmospheric conditions. Qualitative arguments and observations of the plate deflections over time and after helium testing revealed a maximum leak rate of the fusion bonded structures of $1.76 \times 10^{-15} \text{ mbar L s}^{-1}$. Comparison of the test devices revealed similar deflections for all devices. The same phenomenon was observed for devices fabricated with silicon plates, showing that this is not an isolated feature of using Si_3N_4 plates. This has led to the conclusion, that the initial pre-bond of the wafers did not provide a leak-tight bond-interface. Instead, gas is able to diffuse from and to the cavities during the subsequent bond-anneal, until reaching an equilibrium pressure between the cavities and the surrounding atmosphere, prior to the cavities being sealed. The cavities reaching an equilibrium pressure explains why the different bonding conditions resulted in similar plate deflections, and reveals that the bonding conditions do not influence the final cavity pressure, and even bonding in vacuum does not ensure a vacuum cavity. Thus, whether the pre-bond is made in a wafer bonder or directly in hand does not matter. Therefore, if you have no need for alignment bonding or other advanced techniques, you might not need to acquire a wafer bonder. You can even achieve a reduced cavity pressure without it.

The results also has merit even if you already have a wafer bonder. Hand-bonding allows you to place the wafers to-be-bonded together at any point in time. If you have proper quality control of your equipment, your wafers will likely never be cleaner than they were right after a cleaning process. A cleaning process will likely be conducted under a HEPA filter, ensuring a very clean flow of air. Directly after cleaning of the wafers, the pre-bond can be made by hand-bonding the wafers. Subsequent to that, they will not be susceptible for further particle contamination of the bonding interface. They can then be transferred to a wafer bonder for a systematic control of the applied pressure to the wafer stack during the wafer bonder pre-bond. This would be possible for both fusion bonding and anodic bonding, and is in fact the procedure which has been adapted and is used today, based on the presented results.

Part II

3D printed phantoms

Introduction to 3D printing of phantoms

In this chapter, a broad perspective of 3D printing is provided, before a more detailed description of the used 3D printing method stereolithography (SLA) is given. Next, a description of the specific printer that has been used and the printing solution is provided. The content is in part based on Paper B, Paper G, and Paper I.

5.1 3D printing overview

3D printing has seen tremendous development during the last decade. 3D printing is also referred to as additive manufacturing (AM) or solid freeform fabrication (SFF), and is a collective term for a large number of manufacturing techniques capable of creating three-dimensional components in a layer by layer fashion. The methods were originally developed for rapid prototyping, allowing for testing many smaller variations of components, which might otherwise require a lot of work on handmade moulds and casting [126]. 3D printing makes it possible to test many different smaller iterations of a product in a very short time, allowing for much better optimisation of components. Instead of hand-crafting the development models, a digital 3D model can be created in a computer aided design (CAD) program, which can be directly interpreted by the printer software, and created in the desired material. 3D printing is being applied increasingly more in production, where it allows for fairly inexpensive customisation of the product to meet specific customer demands. In the medical field, 3D printing makes it possible to create varying tissue replicas based directly on CT images [127] or patient-specific models from CT or MRI images [128].

The first commercial 3D printing system was a stereolithography (SLA) printing system developed by 3D Systems in 1986 [128]. Since then, a lot of different methods have been invented. These allow for printing various materials, such as metals, polymers, ceramics and concrete, for very different applications and scales. Over time, the capabilities and robustness of the printer systems have increased dramatically, while the cost of the printer systems has decreased drastically, to the point that a 3D printer could almost be considered a household appliance today. The surge in popularity is mirrored in the literature, as seen in Figure 5.1, showing the development in publications relating to 3D printing as the light grey bars. The reason for the drastic increase around 2013 is that a number of patents related to 3D printing processes expired [129]. That allowed for the development of many new 3D printing techniques, which lead to an enormous increase in activity in the research field.

Most people associate 3D printing with the extrusion based 3D printing technique. This is also called fused deposition modelling (FDM). In this printing method, a filament of a thermoplastic polymer, often polylactic acid (PLA) or acrylonitrile butadiene styrene (ABS), is used to print

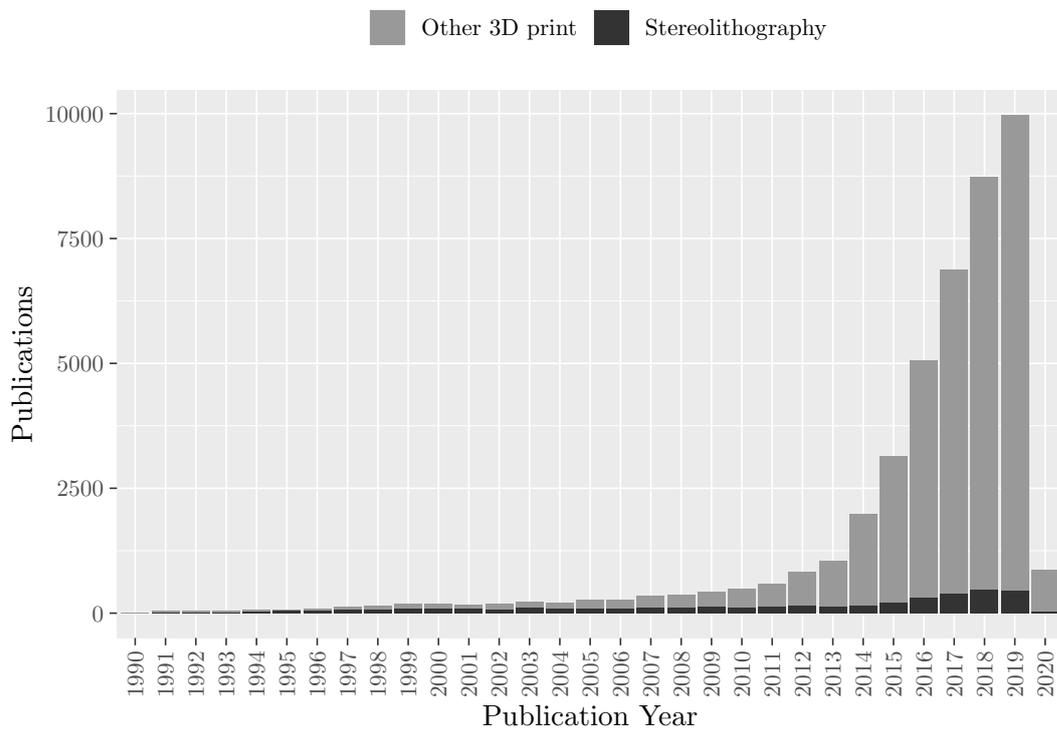


Figure 5.1: Publications with 3D printing per year on Web of Science. Search criteria were: Stereolithography - Topic “stereolithography”; Other 3D print - Topic: “3D Print*” OR “additive manufacturing” OR “solid freeform fabrication”, with “stereolithography” entries subtracted, for the combined height to represent the full “3D printing” results. Both datasets were collected on webofknowledge.com on February 21, 2020.

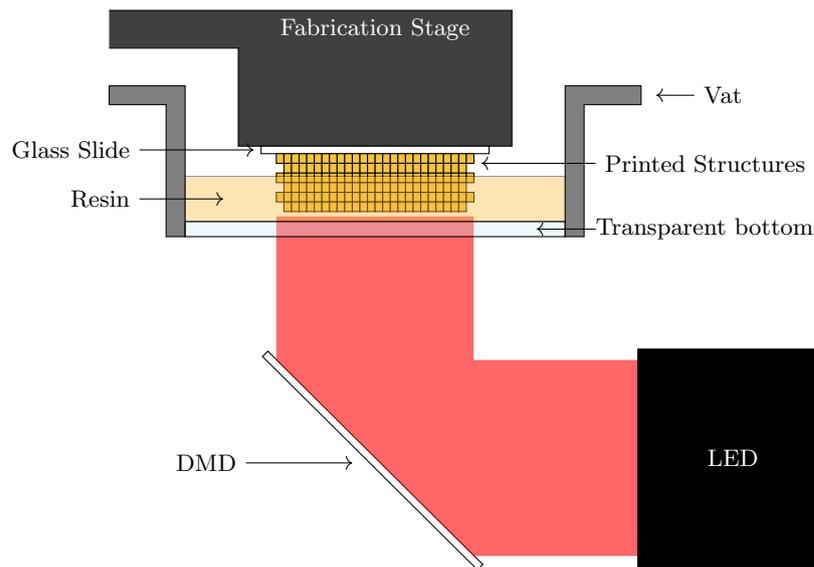


Figure 5.2: Sketch of a stereolithography setup. Light from an LED illuminates a DMD, which reflects the light in the desired pattern through the transparent printer vat bottom. The illuminated resin in the vat will then start cross-linking. The initial layer of the printed structures is cross-linked to a glass slide mounted on the movable fabrication stage.

the desired structure. The filaments is extruded at a nozzle which heats the thermoplastic to a semi-liquid state, in which it can be extruded onto a platform, or previously printed structures, by translation of the nozzle. Then the extruded structure will cool down and solidify. Such systems often have a position accuracy of a few hundreds of micrometres. However, that is only one of many different 3D printing methods. The 3D printing methods are typically separated in the following main methods: fused deposition modelling (FDM), powder bed fusion with subgroups of selective laser sintering (SLS) and selective laser melting (SLM), inkjet printing, stereolithography (SLA), direct energy deposition (DED) with many subgroups, and laminated object manufacturing (LOM), with each having their own benefits depending on the task at hand [129, 130].

SLA is the 3D printing method which has been utilised in this project. The black bars in Figure 5.1 shows the number of publications on stereolithography per year. The remaining publications in 3D printing, shown as the light grey bars, stems from other 3D printing techniques.

5.1.1 Stereolithography (SLA)

Stereolithography (SLA), originally stereolithography apparatus, uses a liquid resin hardened to the shape of the desired pattern through local illumination by a light source, in a layer by layer process to produce the designed 3D object. Figure 5.2 illustrates the bottom-up SLA method. Light from a light emitting diode (LED) illuminates a digital micromirror device (DMD), which reflects the light in the desired pattern through the transparent bottom of a vat. The printer vat contains the resin. The structures are printed on a glass slide which is mounted on the fabrication stage. The fabrication stage is lowered into the liquid resin, until a short distance from the vat bottom. The distance between the glass slide and the vat bottom limits the thickness of the first printed layer. Upon illumination, the liquid resin will start cross-linking until reaching the ‘gel point’ at which the resin solidifies. The illumination system allows for local exposure of the polymer to enable printing of hollow structures. After a layer has been exposed, the fabrication stage is moved a specified distance away from the bottom of the vat, thereby defining the next layer thickness. This is repeated until all layers of the object have been printed.

The design to be printed can be created in any 3D modelling software. However, the MATLAB code controlling the 3D printer requires that the 3D model is sliced into separate layer files. This

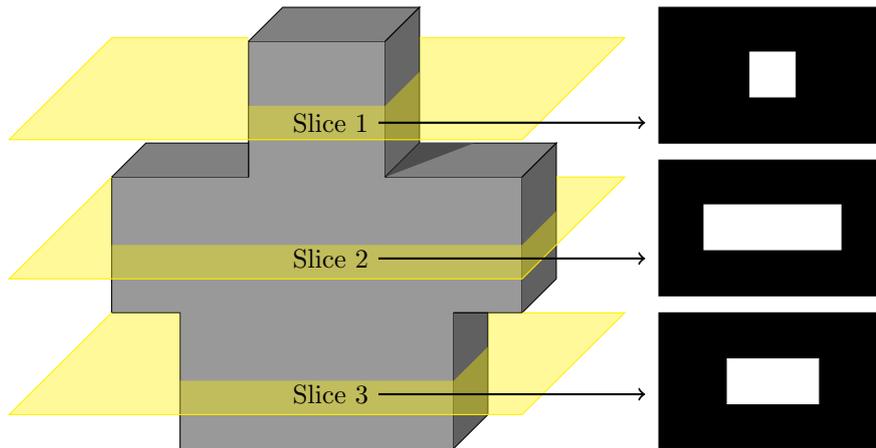


Figure 5.3: 3D model slicing sketch. The slicing software takes a 3D model, as the one on the left, and decomposes it into slices, which are exported as .png files. The figures on the right are examples of slices corresponding to the yellow planes.

can be done by the open source slicing software Slic3r (www.slic3r.org), which converts the model to a set of portable network graphics (.png) files. In addition, an accompanying build list is generated, coordinating the order and exposure time of each layer. Figure 5.3 shows the slicing concept, where a 3D model, on the left, is sliced into separate .pngs, on the right, at each of the yellow planes.

Using a slicer is a very convenient way of obtaining a set of print files in most cases. However, it does not allow for controlling the vertical placement of the slices. This means that the features of the 3D model may end up being split into multiple layers, if they do not align precisely with the slices. The same is valid if the size of the features to be printed do not match with the layer height of the printer, for instance if you want to print a 30 μm tall feature with a 20 μm layer height. In many cases it is not obvious what the results will be.

A way to mitigate that issue is to generate the .png files and the build list manually. This provides full control of the individual slices. As structures become smaller, this level of control will become increasingly more important. As the goal is to push towards the limit of the smallest features attainable, control of the printing structures is of high importance. Therefore, MATLAB scripts were made to create all slice .pngs and build lists manually.

SLA has many different applications, and the printer systems will often be very different depending on the application. As is the case with medical imaging systems, there is a general inverse correlation between printer field of view, and printer resolution. Systems capable of printing in dimensions of more than half a metre are available, for instance at Protolabs (MN, USA), a company offering a 3D printing service, with SLA systems capable of printing as large as 736 mm \times 635 mm \times 533 mm objects with a resolution of 0.254 mm, or 127 mm \times 127 mm \times 63.5 mm with a resolution of 0.0508 mm. The SLA system used in this project is designed to achieve an even higher resolution at the cost of a more limited field of view.

5.2 Custom built 3D printing system

The SLA printer built at DTU was originally developed for printing micro-containers for cell culture chips with nutrient diffusion open 3D micro-channels for vascular networks [131, 132]. This would provide a 3D *in vitro* alternative to animal models which would capture both the structural and dynamic complexity of the *in vivo* counterpart, while being less expensive, time-consuming and controversial [133, 134]. The field is called “organs-on-chips”, where microfluidics are integrated with cell culturing [135, 136, 137, 138]. The goal is to be able to culture cells in an environment which is sufficiently similar to the *in vivo* counterpart. For this application, nutrients are transported to the cells via vascular-like channel systems, which is distributed in all three dimensions.

Within that field, the smallest features printed in a hydrogel so far is $18\ \mu\text{m} \times 20\ \mu\text{m}$ [139], however perfusability of those structures were not presented.

Although the intended use and optimisation of the printer was for a completely different field of research, the resulting hydrogel properties proved to be suitable for ultrasound experimentation as well, due to the acoustic properties being very similar to many types of tissue.

5.2.1 The 3D printer

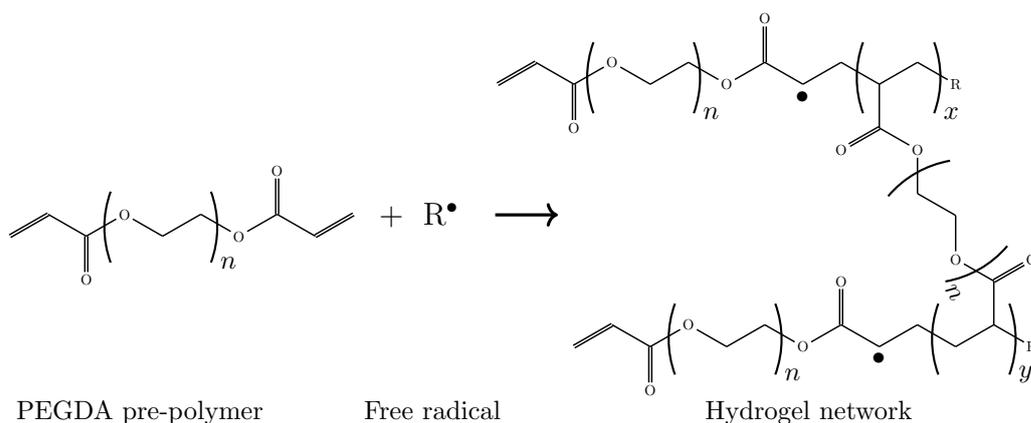
As previously described, the SLA printer projects a full image of the current layer of the 3D model at a time. A 365 nm LED light source was used, emitting light at an intensity of $20\ \text{mW}/\text{cm}^2$ measured at the printing position. The light is expanded and focused through a series of optical components. Each layer image is a one-to-one projection of a digital image generated on a DMD (DLP9500UV, Texas Instruments, TX; part of a V-9501 UV SuperSpeed Digital Light Processing module, Vialux) with a centre-to-centre pixel spacing of $10.8\ \mu\text{m}$ in both lateral dimensions. Thus, there will inherently be a physical mapping of the targeted phantom design layers onto a square grid of $10.8\ \mu\text{m}$ spacing. The DMD consists of 1920 by 1080 micro-mirrors, resulting in a printer field of view of $20.736\ \text{mm}$ by $11.664\ \text{mm}$. The layer images of the phantom shapes were created to match the DMD pixel pitch using a MATLAB (MathWorks, MA) script with a layer thickness of $20\ \mu\text{m}$. The phantoms were printed on $22 \times 22 \times 0.40\ \text{mm}^3$ cover glasses (MEN-ZDA022022A4E0, Menzel Gläser, DE) pretreated with (3-glycidioxypropyl)trimethoxysilane (440167, Sigma-Aldrich) to enhance the adhesion to the printed poly(ethylene glycol) diacrylate (PEGDA). The illumination system homogeneity has been tested by illuminating a charge coupled device (CCD) to determine the light intensities across the printer FOV. The intensity map showed a not perfectly centred, radially decaying intensity. The map was then used to correct the system, by modifying the illumination time on an individual micro-mirrors basis, to match the resulting doses across the printer FOV.

The resulting printed structures are not in equilibrium with water directly after printing, but will swell slightly when subsequently transferred to water. Previous work showed that after four hours, the printed structure reaches its equilibrium swelling [131]. It is important to note that the part of the print which is fixated on the cover glass is not free to swell. This will initially induce stress in the print, and bending of the phantom, as the phantom layers will be gradually more free to swell the further they are away from the cover glass. It is important to remove the print from the cover glass post printing, as the stress induced bending might be frozen into the printed structure if the print remains on the cover glass for too long. This can be seen in Appendix J.3.1.

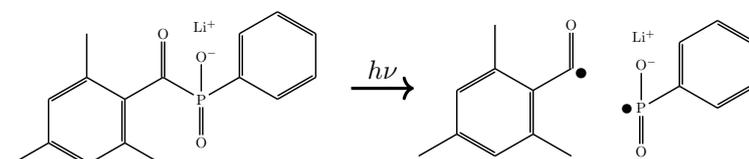
5.2.2 Resin composition

The printing resin consists of three parts: an aqueous pre-polymer solution, a photo-initiator, and a photo-absorber. The pre-polymer will polymerise to form a solid when locally initiated by the light activated photo-initiator. The choice of pre-polymer, photo-initiator, and photo-absorber, along with their respective concentrations is an optimization problem which is in part dictated by the choice of printer components, as well as the requirements for the printed structures. This optimization was conducted previously in [131], with the purpose of printing cell culture chips.

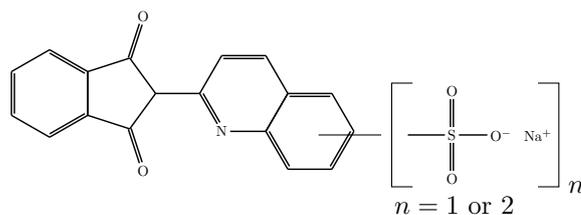
The pre-polymer of the resin is poly(ethylene glycol) diacrylate (PEGDA). The chemical structure of the monomer is sketched on the left in Figure 5.4(a), in which the central back bone is a OC_2H_4 region repeating n times, and an acrylate group at both ends. PEGDA can be acquired in different molecular weights (MWs). The one used during this project has a MW of $700\ \text{g}/\text{mol}$, a density of $1.12\ \text{g}/\text{ml}$, and a melting point between 12°C and 17°C , with the range of temperatures likely being a consequence of PEGDA having a distribution of MWs, in this case with an average of $700\ \text{g}/\text{mol}$. Based on the constituent atomic weights, it can be calculated that on average $n = 13$, making the actual chemical five times longer on average when the backbone is stretched out, compared to the illustrated structure with $n = 1$. However, all the single-bonded carbon (C) atoms are sp^3 hybridized, meaning there is free rotation about the C-C bond, as there is about the



(a) PEGDA - poly(ethylene glycol) diacrylate



(b) LAP - lithium phenyl-2,4,6-trimethylbenzoylphosphinate



(c) QY - quinoline yellow

Figure 5.4: Sketch of the chemical structure of the lithium phenyl-2,4,6-trimethylbenzoylphosphinate photo-initiator molecule (b) and the quinoline yellow photo-absorber (c).

C-O bonds. As a consequence it is entropically favourable for the monomer to curl up in a more spherical structure.

The photo-initiator used is lithium phenyl-2,4,6-trimethylbenzoylphosphinate (LAP). The chemical structure is sketched in Figure 5.4(b). When activated by light at a suitable wavelength, the molecule splits to create two molecules, each with a free radical, seen as the black dots to the right in Figure 5.4(b). The reactivity of the two free radicals is different, with the phosphinate-containing group being more reactive. The initiator is then able to bind to either of the acrylate groups, by breaking the double bond at the end. The initiator will bind to the outermost C atom of the acrylate group, breaking the double bond at the end of the PEGDA molecule. In doing so, the other C of the double bond becomes a new free radical. The initiator always binds to the outermost C atom as the other carbon atom forms a more stable free radical due to more carbon substituents. The new free radical is seen as the black dots in the polymer network to the right in Figure 5.4(a). This free radical can then break the double bond of another acrylate group in the solution, essentially starting a chain reaction. This will continue until terminated by another radical binding to the free radical of the acrylate group, thereby stopping the chain reaction. The terminating radical may either be located on an activated photo-initiator molecule or on another growing polymer chain. For the polymerisation to be characterised as a polymer network, it will

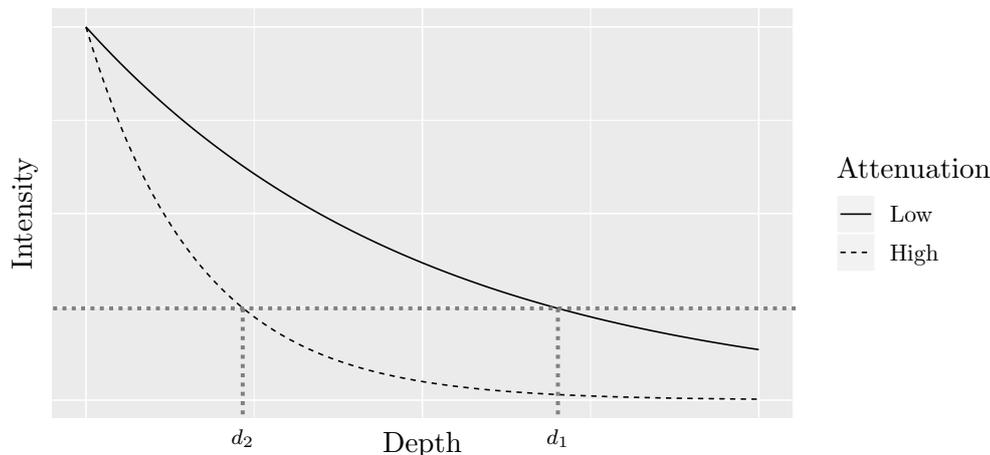


Figure 5.5: Light intensity against depth for two (arbitrary) different attenuation coefficients. d_1 and d_2 mark the depths, at which the light dose (intensity times exposure duration) has decreased to the threshold of resin solidification, marked by the grey horizontal dotted line.

require that the monomers can be cross-linked in at least three sites each; any less, and it would simply be a single polymer string. In the case of the PEGDA network, each monomer contains two double bonds which can each result in cross-linking to two other monomers, for a total of four cross-linking sites per monomer.

The PEG groups, marked by the parentheses in Figure 5.4(a) are hydrophilic groups, which attract and bind water. As the printing resin is aqueous, water is present during the printing process and will be bound in the structure immediately, making the printed structure a hydrogel. The prints will contain ≈ 75 wt% water thereby making it resemble many types of tissue in terms of the water content [140].

The photo-absorber used is quinoline yellow (QY). The chemical structure is sketched in Figure 5.4(c), where it can be seen that it contains one or two sulfonate groups, which might be at different positions of the main C molecule skeleton. To understand the reasoning for using it, one needs to understand that light in a medium is attenuated according to Lambert-Beers law

$$I = I_0 e^{-\mu d}, \quad (5.1)$$

where I_0 is the initial intensity, μ is the attenuation coefficient, and d is the depth at which the intensity, I , is measured. Figure 5.5 shows two examples of the exponential decay of the light intensity in two different media with different μ . The pre-polymer reaction initiated by the photo-initiator will need a certain dose (intensity times the exposure duration) of light to solidify (reach its gel point). Thus, illumination of the polymer will only cross-link the polymer until a certain depth. In Figure 5.5 this is marked by d_1 and d_2 for the two curves corresponding to different levels of attenuation. While the interlayer movement of the stage sets a lower limit on the vertical resolution, the light dose may be sufficiently large to induce further solidification in previously exposed layers, as illustrated in Figure 5.6(a).

The aqueous solution of pre-polymer and photo-initiator has little attenuation, so light will propagate far into the resin before being absorbed. Addition of a highly absorbent photo-absorber allows to limit the propagation depth, thereby setting the depth resolution. A light absorber will modify μ of the solution according to

$$\mu = \epsilon c, \quad (5.2)$$

where ϵ is the extinction coefficient, and c the concentration of the photo-absorber. Thus, μ can be modified by the type and concentration of the added photo-absorber from a slowly attenuating

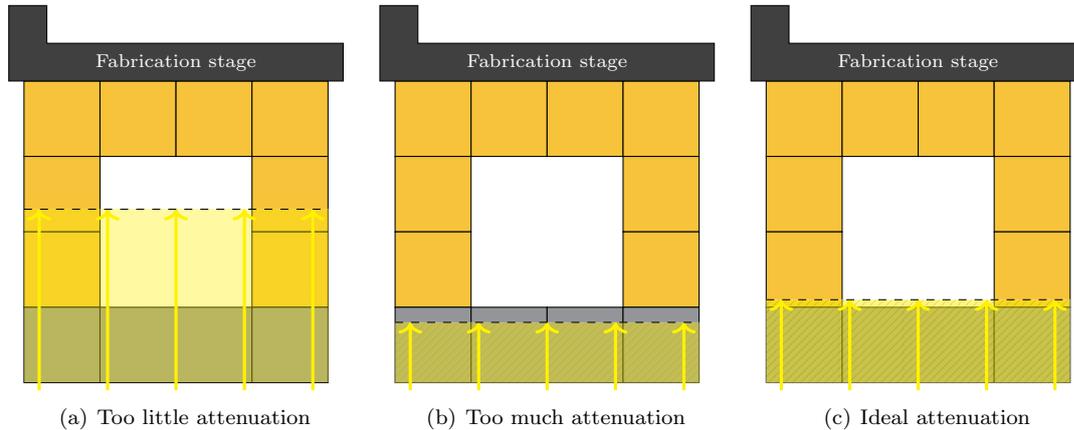


Figure 5.6: Sketches of the cross-linking depth for different levels of attenuation. The yellow squares are previously exposed voxels, and the grey squares are the voxels which are to be exposed in the new layer. The dashed lines mark the depth of the threshold dose. (a) has too little attenuation, and previous layers are re-exposed. (b) has too much attenuation, and the newly exposed layer is unable to cross-link with the structures in the previous layer. (c) has sufficient attenuation, with only a minimum overlap in the exposed region to the previous cross-linked structures.

medium such as water, exemplified by the solid curve and d_1 in Figure 5.5, to a higher attenuating medium exemplified by the dashed curve and d_2 . If μ becomes too large, the polymer will not cross-link sufficiently deep to chemically bond to the overlaying structures in the previous layer, as illustrated in Figure 5.6(b). Thus, the choice of photo-absorber and concentration must be matched to get a slight overlap between the newly exposed regions, and the previously printed structures, as shown in Figure 5.6(c).

The optimized concentrations of each component were established in previous work [131]: ≈ 200 mg/ml, or $\approx 20\%$ (w/v) of PEGDA, 5 mg/ml, or 0.5% (w/v) of LAP, and 12 mg/ml, or 1.2% (w/v) of QY, all dissolved in water. The photo-initiator LAP is chosen for its water solubility and its absorption spectrum which matches the light source used. The photo-absorber QY is also chosen for its water solubility in addition to its high extinction coefficient at the wavelength used. Absorption spectra for both can be seen in [131]. These concentrations have been optimized for using a layer exposure time of three seconds which is sufficient to facilitate cross-linking of the pre-polymers in the resin to previously printed layers. At this layer exposure time, the overlap between printed layers is approximately 10 μm . However, the printed structures are not completely saturated in terms of cross-linking at this layer exposure time, and a change in the exposure will modify the amount of cross binding.

The resulting printed structures have a Young's modulus of ≈ 1 MPa. This is low compared to prints in PLA and ABS through FDM, which both have a Young's moduli on the order of hundreds of MPa [141, 142], and the hydrogel phantoms are therefore in this context considered consisting of soft materials.

The structure of the printed polymer network is fairly open. This is a utilized property given the intended use of the network as a diffusion open structure, capable of delivering nutrients to cells. This is also visible to the naked eye over time after printing. Right after the print is finished, the hydrogel contains a significant amount of QY, making the print yellow as seen in Figure 5.7(a). The water is clear because it has just been exchanged. The sample is 2 mm thick. After 30 minutes, a significant amount of QY has diffused out through the hydrogel polymer network into the surrounding water, making the water yellow as seen in Figure 5.7(b). After having exchanged the water a few times over a few days, the QY has effectively been washed out of the hydrogel sample, as seen in Figure 5.7(c). Although not visible, the extra PEGDA monomers left inside the print, for instance in a printed cavity, are expected to diffuse out in a similar manner. As

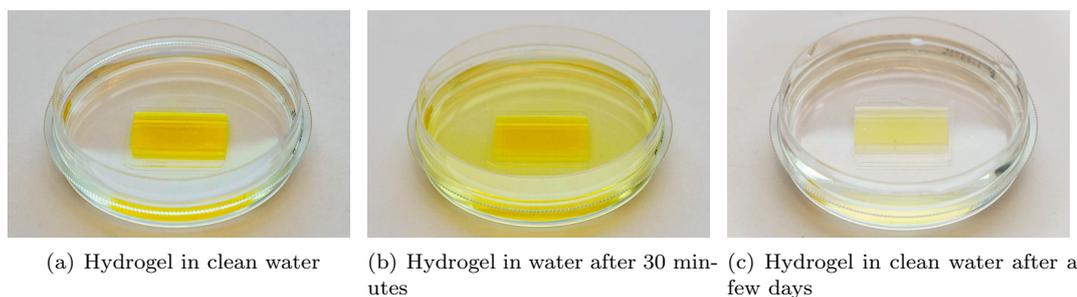


Figure 5.7: QY diffuses out through the hydrogel network over time. (a) shows how the yellow QY is initially trapped in the hydrogel sample right after the print. The water is clear because it has just been exchanged. (b) shows the same hydrogel 30 minutes later. A significant amount of QY has diffused out of the hydrogel sample into the water. (c) shows the same hydrogel sample after the water has been exchanged multiple times over a couple of days. Most of the QY is now gone.

a consequence of the folding nature of organic molecules into spherical structures, the radius of a molecule can generally be assumed to scale with the cube root of the molecular weight of the molecule, corresponding weight being proportional to volume which in turn scales with the radius cubed [143]. The molecular weight of QY is between 352 g/mol and 432 g/mol, depending on whether it has one or two sulfonate groups attached. So the size difference between the PEGDA monomer and QY will not be large. In fact, the radius of the PEGDA molecule will only be a factor of 1.26 or 1.17 larger than the QY molecules with one and two sulfonate groups respectively.

The printed structures need to be stored in water. If left in air, the water in the hydrogel will start evaporating. Due to the bottom-up method of printing, the printed layers will gradually be pulled out of the printer resin as more layers are added. For very tall prints, the samples will eventually dry out during printing, as the printed layers which are exposed to air will dehydrate. The consequence might be bending of the prints, resulting in misalignment of the subsequent layers. This effectively sets a limit as to how tall prints can be made. The limit depends on the used layer exposure time, since a larger exposure time will result in a longer total print time. The prints made in this project are both significantly larger than the typical prints being made with the printer, with some of the prints also having significantly larger doses. The majority of the prints during this project were 11.66 mm tall with a three second layer exposure time utilizing the full print area, and printed consistently without failure. Other lateral geometries printed just as tall with small openings between print regions have been made with interlayer exposure times as high as 23 seconds. For the largest exposure, the prints would occasionally start failing.

5.3 Chapter summary

In this chapter, the stereolithography 3D printing technique has been presented, along with the printer components, and the optimised resin formula. The printed structures are hydrogels with a water content of $\approx 75\%$, which is desirable for ultrasound imaging, since that water content matches that of many different tissues in the human body. It was also demonstrated how the hydrogel structures are diffusion open to water, and even remaining unpolymerised resin components, which means that the unpolymerised remains left inside printed cavities, will be washed out by water over time. It has been shown in this project that prints can systematically be printed 11.66 mm tall. This is not the limit, but suitable for use in SRUS.

In the chapters to follow, the PEGDA based hydrogel is characterised specifically focusing on properties relevant for ultrasound experimentation, before some examples of optimisation and general uses of the hydrogel phantoms are presented. In addition, a number of other observed features and issues with the hydrogel printing are presented in Appendix J.3.

Hydrogel material characterisation

This chapter describes the material properties of the printed PEGDA hydrogels. With the goal being to mimic vascular structures in terms of scale and complexity to develop ultrasound techniques with a supposed resolution of only a few micrometers, it would be critical to know and compensate for any systematic geometrical deviations in the printed structures relative to the designed structures. Furthermore, it is important to have an understanding of the acoustic properties when using the phantoms for ultrasound experimentation. This chapter is in part based on Paper B, Paper G, and Paper I.

6.1 Hydrogel structural properties

Whenever a new phantom type or fabrication method is introduced, it is important that it is characterised in detail in terms of the relevant parameters. For SRUS where the purpose is to mimic vascular structures much smaller than the resolution of conventional ultrasound techniques the precision and accuracy are arguably some of the most important parameters. Given the intended use for the SLA printer, the resin component and printer optimization had focused on the bio-compatibility and network porosity of the hydrogel, as well as the obtainable resolution of features. Typical sizes of the cell culture models are close to 5 mm, which means that accuracy and precision across the entire printer field of view had not previously been investigated. Conversely, conventional ultrasound phantoms are typically several centimetres in all dimensions, so the printed phantoms would in general employ the full printer field of view, obtaining lateral dimensions of $\approx 21 \times 12 \text{ mm}^2$.

In addition to the printer precision and accuracy, the acoustic parameters of the materials used are always very important when working with ultrasound phantoms. The early experiments we did showed us empirically that it is possible to manipulate the acoustic impedance of the printed hydrogel material. Figure 6.1 shows a B-mode image of three scatterers of different sizes, printed using different parameters. The two white arrows point to cavity markers, which simply consists of a region of unexposed, unpolymerised voxels. The red arrow points to a solid marker. The solid marker was created by providing additional layer exposure time to an already exposed hydrogel region, in this case an additional 20 seconds. The left and right markers were designed to be $240 \times 240 \times 240 \text{ }\mu\text{m}^3$, and the central marker was designed to be $400 \times 400 \times 400 \text{ }\mu\text{m}^3$. As presented in Section 5.2.2 the unpolymerised resin in the printed cavities will be washed out over time and replaced by water since the hydrogel is diffusion open to water and PEGDA monomers and QY. The phantom was imaged at 15 MHz using a BK Medical "Hockey Stick" X18L5s probe and a BK 5000 experimental scanner. The high frequency leads to high resolution of the B-mode

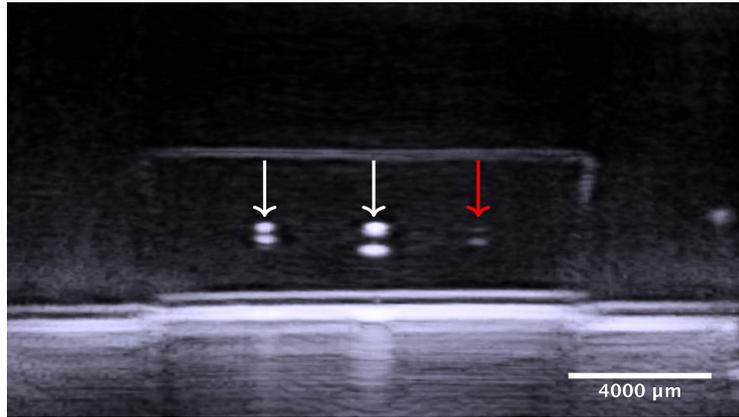


Figure 6.1: B-mode image of scatterers created with different methods of different sizes. The two white arrows point to cavity markers. The red arrow points to a solid marker.

image, which results in each scatterer being represented by two reflections, one at the top of the scatterer, and one at the bottom. This will be discussed further in Chapter 7. The solid marker is separable from the background, although with a 28 dB lower intensity compared to the hollow markers.

For the majority of phantoms printed including the one imaged in Figure 6.1, the interlayer exposure in the bulk of the phantom was set to the default optimised exposure of three seconds, in the following referred to as the base exposure. Thus the material properties should be constant across most of a phantom. While the three second exposure time is sufficient to facilitate cross-linking of the pre-polymers in the resin to previously printed layers, the printed structures are not completely saturated in terms of cross-linking at this exposure time, and a change in the dose will modify the amount of cross-linking. The solid marker in Figure 6.1 shows that the additional dose leads to a change in acoustic impedance. Therefore it would be interesting to test the effect of changing the dose, if any, on the acoustic impedance. However, the effect on structural properties, such as the printer precision and accuracy might also be affected by changing the dose.

6.2 Hydrogel swelling

The hydrogel material was known to swell post-printing, but the exact amount of swelling was unknown and might depend on the exposure dose.

6.2.1 Swelling uniformity

First of all, it is important to know whether the swelling is isotropic, or whether it is different along the different print-axes. This could potentially be an effect of the anisotropic voxel.

Experimental procedure

The phantom swelling was originally investigated during development of SRUS calibration phantoms. The use of these phantoms are presented in Chapter 7. The foundation for the work presented there is a phantom containing eight randomly placed scatterers. The designed outer dimensions of the phantom is $20.736 \times 11.664 \times 11.660 \text{ mm}^3$ with each scatterer being $205 \times 205 \times 200 \text{ μm}^3$. The designed layout is shown in Figure 6.2, in which the blue points represents the randomly placed scatterers. Separate droplines lead from the points out along the y -axis and along the z -axis respectively, and are included to aid the 3D perception of the scatterer placements. The

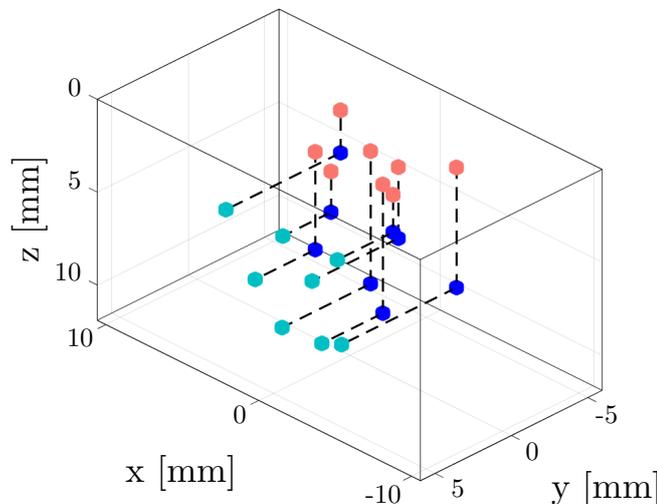
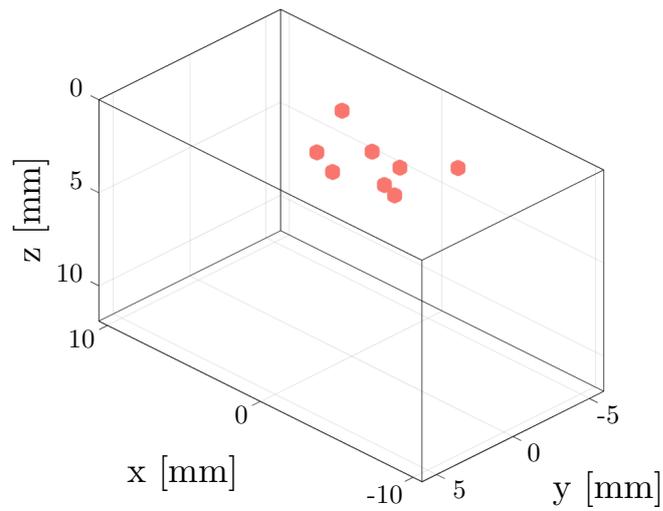


Figure 6.2: The designed layout of the scatterers within the $\sim 20.7 \times 11.7 \times 11.7 \text{ mm}^3$ phantom. The blue points are the randomly placed scatterers. The droplines are included to aid the 3D perception of the scatterer placement. The red points mark the scatterer positions collapsed into the x - y -plane near the top surface, and the turquoise points mark the scatterer positions collapsed into the x - z -plane near the side surface.

droplines end up 1 mm from the respective surfaces in the collapsed x - y -plane version (red) and the collapsed x - z -plane version (turquoise) of the scatterers.

The accuracy of the phantom fabrication method should be verified by another characterisation method before being used to calibrate ultrasound techniques. Although the printer specifications have been presented, they only specify the lower limit of the attainable feature sizes and accuracies. Furthermore, the phantom expansion due to post-printing swelling needs to be determined to compensate the designed feature sizes before using the phantom as a calibration tool. Optical characterisation using an optical microscope can be used to locate phantom features with high precision. Unfortunately, the printed hydrogel scatters light, rendering it impossible to use the phantom containing the eight randomly placed scatterers, since these are placed too far inside the phantom. Instead, the same coordinates were used to make two new phantoms, in which the coordinates were collapsed either into the x - y -plane and placed near the top of the phantom, as seen as the red points in Figure 6.2 and in the overviews in Figure 6.3, or into the x - z -plane and placed near the side of the phantom, as seen as the turquoise points in Figure 6.2 and in the overviews in Figure 6.4. The scatterers were placed 1 mm from the surfaces in both cases. By placing them near the surfaces, the light scattering is minimised and the scatterers become clearly visible in the optical microscope. Each scatterer was physically moved into a defined centre point in the optical field of view using an X-Y microscope stage with integrated linear encoders for accurate readout of the actual position. This procedure circumvents possible measurement errors due to distortions in the optical components. The measurements were performed using a Zeiss LSM 700 upright microscope with a Zeiss 130x85 PIEZO stage having a positioning reproducibility of $\pm 0.6 \mu\text{m}$. The positioning accuracy of the procedure was assessed by repeatedly locating the same scatterer. The position was found with a standard deviation of $1.3 \mu\text{m}$ along both the x -axis and the y -axis ($n = 50$).

To estimate the swelling, the design distances of the phantom model can be correlated to the optically measured distances. The analysis procedure is sketched in Figure 6.5. The distance between all scatterers can be determined from the individual scatterer positions, and the correlation



(a) Phantom layout

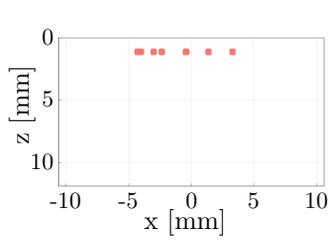
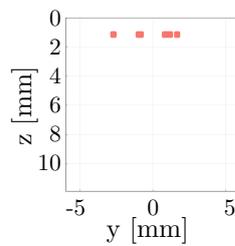
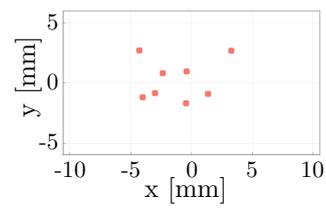
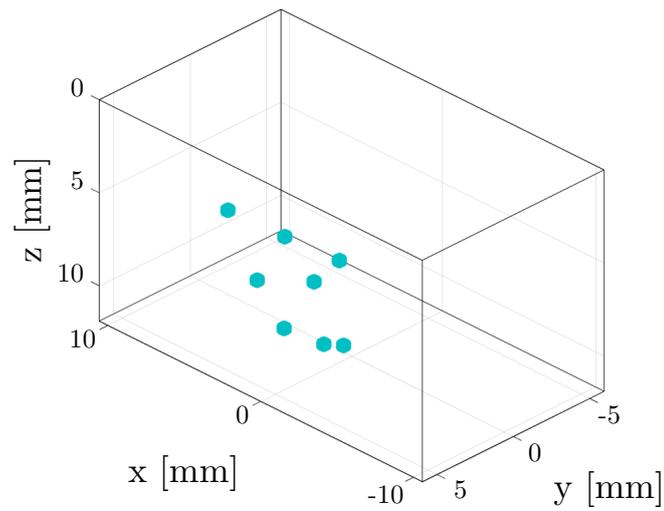
(b) x - z -plane(c) y - z -plane(d) x - y -plane

Figure 6.3: Scatterer phantom 1 for validation in optical microscope. The scatterer positions in this phantom is based on the the blue scatterers in Figure 6.2 being collapsed into the x - y -plane 1 mm from the top surface of the phantom.



(a) Phantom layout

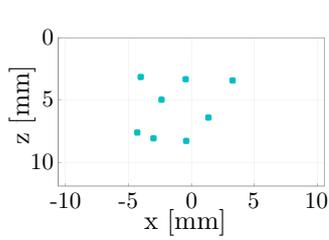
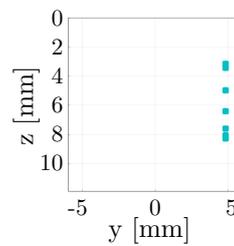
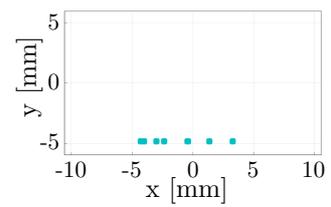
(b) x - z -plane(c) y - z -plane(d) x - y -plane

Figure 6.4: Scatterer phantom 1 for validation in optical microscope. The scatterer positions in this phantom is based on the the blue scatterers in Figure 6.2 being collapsed into the x - z -plane 1 mm from the side surface of the phantom.

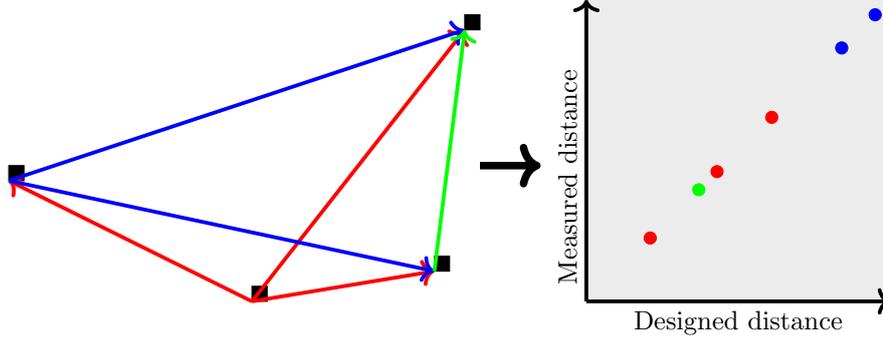


Figure 6.5: Sketch of the procedure to determine the accuracy of the printed phantoms. The black squares represent printed scatterers. The distances between scatterers are determined, and a correlation between the measured distances and the designed distances is made. The slope of the correlation will be the expansion factor post-printing.

Table 6.1: Summary of the variables and their data types used in the optical correlation analysis.

Predictors	Sample values	Variable type	Description
Optical distance [mm]	4.041, 1.950,..., 8.189	Numerical values	The response variable, measured by optical microscope
Design distance [mm]	3.973, 1.927,..., 8.087	Numerical values	The designed distance between points
Plane	XY, XZ	Fixed factor	The cross-plane investigated
Phantom	1, 2, 3, 4	Random factor	The phantom group

should be linear. The slope of the correlation is the factor by which the printed structure has expanded relative to the design. If the printed structures are a perfect replication of the design, the correlation will be a linear relationship with a slope of 1.

Results

Two replicates of each of the two projected-scatterer phantoms for optical validation were made using the base exposure. Each scatterer was located using the optical microscope and the translation stage coordinates of each scatterer was determined. Subsequently, the scatterer coordinates were used to determine the distance between the scatterers. The correlation between the optically measured distances and the designed distances can be seen in Figure 6.6. In addition to analysing the direct correlation between measured distances and design distances, it was also investigated whether there was any difference between the two sets of cross-planes (x - y and x - z), which could potentially be explained by the anisotropic voxels. The different phantoms were also modelled as a random factor, to test and compensate for print-to-print variability. The combination of fixed and random factors makes the fitted model a linear mixed effects model. Such a model can be analysed using the `lmerTest` package [144] in R [145]. A summary of the data types and the factors included in the analysis can be seen in Table 6.1. The initial mixed effects model is given as

$$Y_i = \mu + \alpha(\text{Plane}_i) + (\beta_1 + \beta_2(\text{Plane}_i)) x_{\text{design},i} + c(\text{Phantom}_i) + \epsilon_i, \quad (6.1)$$

where Y_i is the optically measured distances, μ is the overall intercept, $\alpha(\text{Plane}_i)$ is an intercept addition due to the Plane factor, β_1 is the average slope of the model, $\beta_2(\text{Plane}_i)$ is a plane dependent correction to the slope, $c(\text{Phantom}_i) \sim N(0, \sigma_{\text{Phantom}}^2)$ is a random offset from phantom to phantom, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution

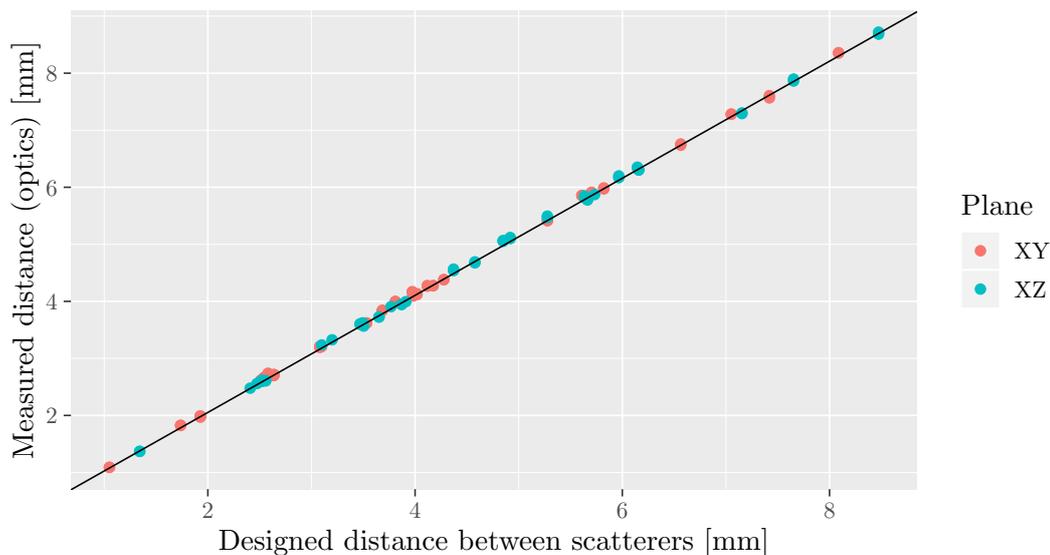


Figure 6.6: Correlation between the distance between the designed scatterer positions and the distances measured using an optical microscope. The black line is the final reduced model seen in Eq. (6.2).

Table 6.2: Model parameter estimates of the final reduced model including confidence intervals of correlation between optical measurements and design distances.

Predictors	Estimate	2.5%	97.5%	<i>p</i> -value
Fixed effects				
Intercept [mm]	0.023	0.005	0.042	0.0136
β_1 (slope)	1.026	1.022	1.030	<0.0001
Residual error				
σ [μm]	36.6			

with mean μ and standard deviation σ , all for the i th response. All $c(\text{Phantom}_i)$'s and ϵ_i 's are independent.

The model reduction was conducted by removing only a single term at a time, based on a 5% level of significance. Neither the random effect of the individual phantoms ($c(\text{Phantom}_i)$), nor the Plane dependent intercept addition ($\alpha(\text{Plane}_i)$), nor the Plane dependent slope ($\beta_2(\text{Plane}_i)$) were significant at 5%. Thereby the model reduction converged at the final model

$$Y_i = \mu + \beta_1 \cdot x_{\text{design},i} + \epsilon_i. \quad (6.2)$$

The model coefficients and confidence intervals of the reduced model can be seen in Table 6.2.

The analysis shows that the phantom swelling is isotropic, since there was no effect of the Plane factor. There was no significant difference between the four test phantoms, indicating good print repeatability. The parameter estimate of β_1 indicates that the phantom expands by approximately 2.6% along all dimensions. The residual standard error of the model is 36.6 μm . Model diagnostics showed that the residuals appeared to be normally distributed. Thereby, the model is a good describer for the phantom expansion. The overall good correlation of all points to the straight line indicate that the expansion is uniform and isotropic in the investigated region of the print area. The analysis showed a significant intercept of 23 μm , which was unexpected. Given that the intercept lies outside of the data range of interest, it has not been analysed any further. It is worth noting that the confidence interval for the intercept varies from less than a single voxel width, to four voxel widths.

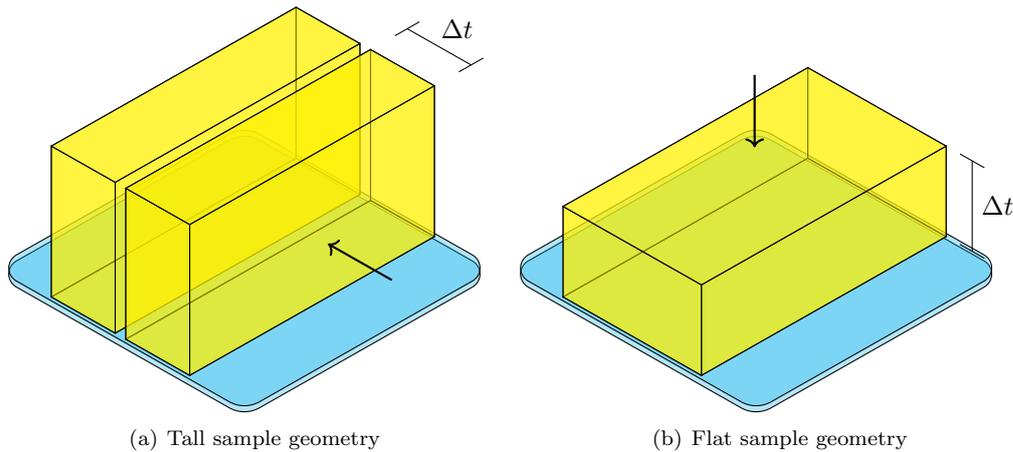


Figure 6.7: Sketch of the test sample geometries. The geometry shown in (a) was used for testing the hydrogel swelling, density, cross-linking density, and the speed of sound. The geometry in (b) was used for testing the acoustic attenuation. The arrows indicate the intended direction of sound, relative to the sample geometries.

6.2.2 Swelling at different printing doses

Experimental procedure

In the following, it is assumed that the swelling is isotropic regardless of the dose. To test whether the expansion changed with dose, a number of samples were made with interlayer exposure times from 2 seconds to 23 seconds, which spans the full range of exposure times that have been used throughout this project. It should be noted that this range of exposure times far exceeds the interlayer exposure times that had previously been used for cell culture chips, which was typically only in the range of 2 seconds to 8 seconds. Initially, the sample geometry in Figure 6.7(a) was used, primarily to increase the number of test samples per print. The arrow shows the intended sound propagation direction. The samples were all designed to be 5 mm wide. With the 2.6% expansion for a 3 second interlayer exposure, the sample width would be expected to become 5.13 mm. Each sample was then mounted in a custom 3D printed hydrogel holder allowing it to be submerged in water during optical measurement. The 3D printed holder design can be seen in Appendix H.3. The final dimensions of the print were then determined using the same Zeiss LSM 700 upright microscope with a Zeiss 130x85 PIEZO stage as for the optical characterisation of the scatterer phantom. This time, each edge of the phantom was moved into a defined centre point in the optical field of view using the X-Y stage, and the position was read out.

Results

A total of 25 hydrogel samples were 3D printed, using twelve different doses, spanning the layer exposure time range of 2 seconds to 23 seconds. In addition to the exposure time range specified, an interlayer exposure time of one second was also used, however, this was not sufficient to facilitate interlayer cross-binding, rendering it impossible to make one-second-samples for testing. An image of a 3 second layer exposure time sample can be seen in Figure 6.8(a). Two or three samples were made using each dose. The resulting measured sample thicknesses and corresponding swelling percentages are presented in Figure 6.9. The dashed lines in the figure represents the measured swelling based on the data from Figure 6.6 in the previous section. The data shows that the swelling is not constant with dose, and increases beyond 6% in the layer exposure time range from 9 seconds to 17 seconds, before decreasing again for higher exposure times. The data from the previous section aligns well with these data points.

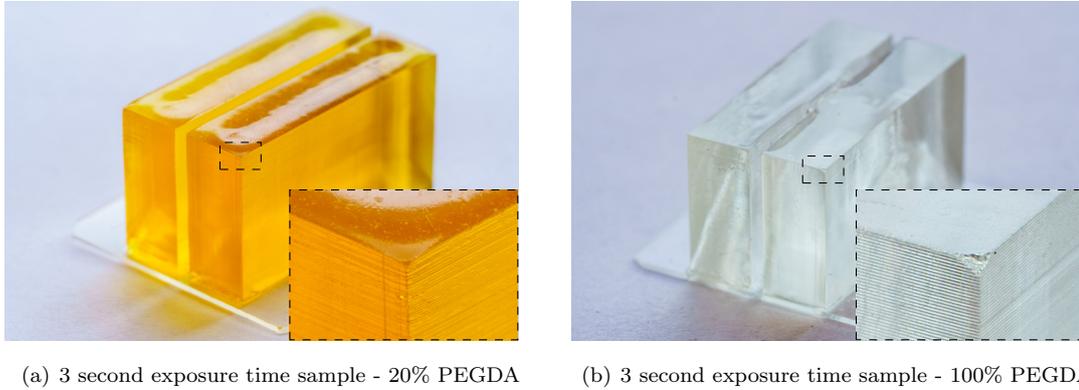


Figure 6.8: Images of test samples. (a) is a 20% PEGDA sample and (b) is a 100% PEGDA sample. The difference in colour stems from different types of absorber used. Note that the surfaces on the sides appear more rough than those on the top.

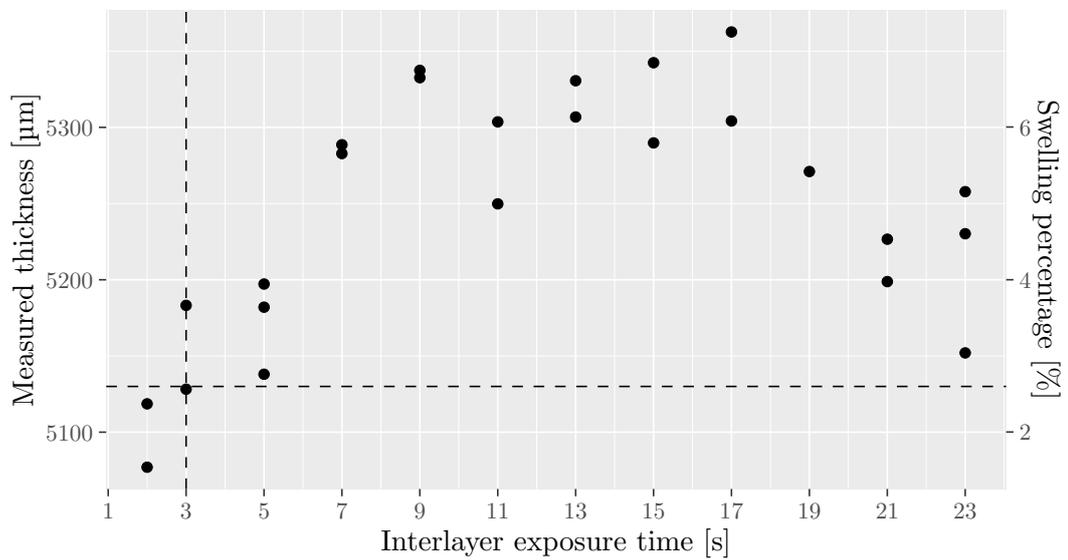


Figure 6.9: Measured thickness and swelling percentage of hydrogel samples designed to be 5000 µm against interlayer exposure time. The dashed lines represent the used dose and determined swelling from Section 6.2.1.

6.3 Hydrogel density

Experimental procedure

The amount of cross-linking might have an effect on the density of the material. The density was determined using the Archimedes principle, which states that

$$\rho_{\text{sample}} = \frac{w_{\text{sample,dry}}}{w_{\text{sample,dry}} - w_{\text{sample,submerged}}} \rho_{\text{fluid}}, \quad (6.3)$$

where ρ_{sample} is the density of the sample, $w_{\text{sample,dry}}$ is the weight of the sample when measured in air, $w_{\text{sample,submerged}}$ is the apparent weight of the sample when submerged in a liquid, and ρ_{fluid} is the density of the fluid which the sample was submerged into. The same hydrogel samples which were used to determine the swelling at different doses can be used to determine the density.

A Mettler Toledo XS105 DualRange scale (Mettler Toledo, Columbus, OH, USA) equipped with a ML-DNY-43 density kit (Mettler Toledo, Columbus, OH, USA) was used to determine $w_{\text{sample,dry}}$ and $w_{\text{sample,submerged}}$. The scale has a built in function to determine the density of a sample using the Archimedes principle. Unfortunately, when the hydrogel samples are exposed to air, the water in the sample will slowly start to evaporate. The rate of evaporation was high enough that the scale was not consistently able to stabilise due to the loss in weight, rendering it impossible to use the built in program. Instead, it was simply used as a regular scale, since there was a clear difference in the rate of evaporation, to the initial stabilisation of the scale when the sample was placed on the scale. Thereby, it was possible to manually read of the weights, both for the dry weights and the submerged apparent weights.

Aside from the 25 samples used previously, a different print solution was made, which would allow for direct printing of a 100% PEGDA structure, to be able to determine the PEGDA content in the hydrogels which were printed from the 20% printed PEGDA resin. The samples were printed using the same design as the 25 samples, shown in Figure 6.7(a). Since neither the photo-initiator LAP nor the photo-absorber QY are soluble directly in PEGDA, the photo-initiator Irgacure 819 at 4 mg/mL and the photo-absorber Avobenzene at 1 mg/mL were used instead. An image of the printed structure can be seen in Figure 6.8(b). Due to the different absorber, the print is completely clear. This combination of absorber and initiator is more reactive than the 20% PEGDA resin, revealing some printer system artefacts, namely the ‘‘Ghost image’’, which is an additional systematic offset of the printed pattern with a very low dose. This is also seen to a small degree for the sample in Figure 6.8(b), in which there is no longer any separation between the two parts in the middle of the sample. A more severe case can be seen in Appendix J.3.2. The effect is simply that more of the PEGDA was polymerised than the 3D model design had dictated. If the density of the 100% samples vary with dose, the samples affected by the ‘‘ghost image’’ effect might be poor representatives of a true printed 100% PEGDA density. Comparison of the densities of severe ‘‘ghost image’’ samples and less severe ‘‘ghost image’’ samples will reveal whether the densities are affected.

Results

The same 25 hydrogel samples were used to determine the density of the hydrogel with different doses. The calculated densities can be seen in Figure 6.10. For low doses the density appears to increase slightly before decreasing again for medium to high doses. The overall average density is 1.045 g/ml, marked as the dashed line.

Three 100% PEGDA samples were printed. The sample shown in Appendix J.3.2 was severely affected by the ‘‘ghost image’’ effect, while the other two samples only were slightly affected. The average density and standard deviation was (1.186 ± 0.001) g/ml. The small standard deviation indicates that the density of the 100% samples was unaffected by the ‘‘ghost image’’ effect. The right y -axis in Figure 6.10 shows the percentage of PEGDA in the printed hydrogel samples, from comparison of the densities to that of the average 100% printed PEGDA sample. The results show that the printed structures have a slightly larger PEGDA content than the liquid resin, which was a

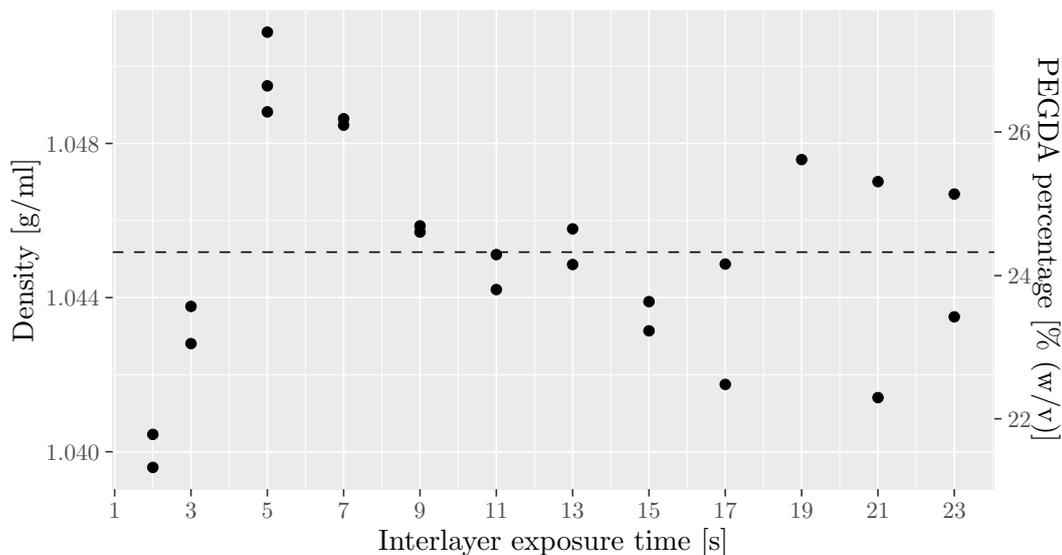


Figure 6.10: Calculated density against interlayer exposure time. The right axis shows the (w/v) percentage of PEGDA in the printed hydrogel samples, determined from comparison to a 100% printed PEGDA sample. The dashed line marks the average density.

20% (w/v) PEGDA solution. Although observations indicate that the true starting solution might actually be 21% PEGDA (see Appendix J.1), the results show a systematically higher concentration of PEGDA. These higher concentrations of PEGDA likely means that the printing resin is slightly depleted in PEGDA content during printing, potentially meaning that the PEGDA percentage in a final print could differ from the bottom of the sample to the top. This will in particular be the case for layer exposure times right around 5 seconds and 7 seconds. This indicates it might be important to use fairly large resin volumes when printing to minimise this effect.

6.4 Acoustic characterisation

The acoustic parameters are critical when using the hydrogel samples for ultrasound imaging. In this section the acoustic properties will be evaluated.

6.4.1 Speed of sound

Methods

To determine the acoustical properties at different doses, a custom 3D printed pulse-echo measurement setup was created. The 3D model was made in Autodesk inventor and can be seen in figure Figure 6.11.

At each end, the holes are exactly dimensioned to fit the round ultrasound transducers used. When mounted, the transducers face towards each other. The alignment of the printed model was tested by twisting the probes within the holes to make small adjustments away from the model alignment. In all cases, the signal received decreased in amplitude, suggesting that the printed model is aligned well. The structure in the middle allows for mounting of poly(methyl methacrylate) (PMMA) sample holders, centring the samples along the ultrasound propagation path. This design can be used both for transmission measurements, by using one transducer as a transmitter and the other as a receiver, or for pulse-echo measurements, using a single transducer as both transmitter and receiver. The setup was used to determine the speed of sound and the attenuation of the samples. The PMMA holders were designed with holes in front of the hydrogel

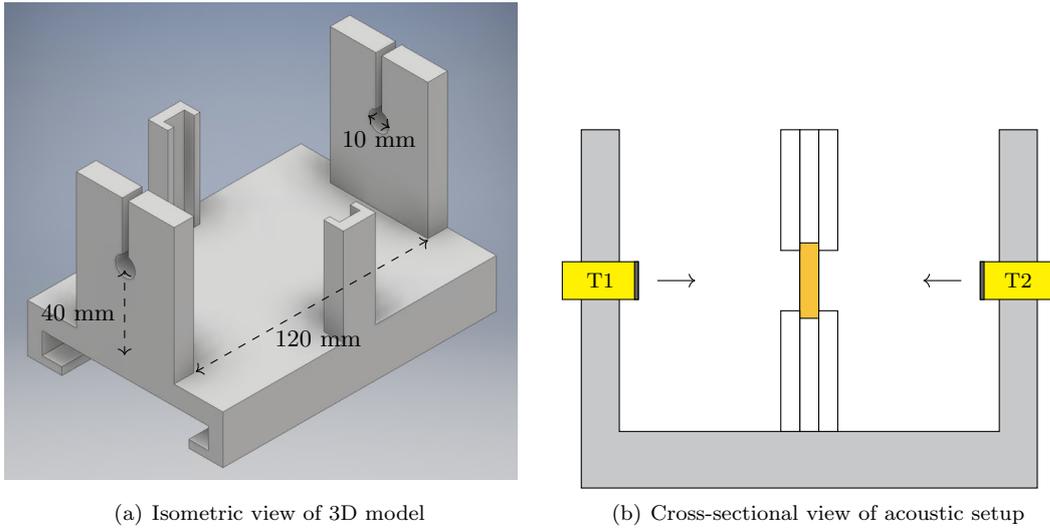


Figure 6.11: The acoustic setup. (a) shows an isometric view of the 3D model in Autodesk inventor with dimensions. (b) shows a cross-sectional view of the setup. The holes at each end are exactly dimensioned to fit the round ultrasound transducers used, marked T1 and T2. When mounted, the transducers face towards each other. The structure seen at midway in the isometric view allows for mounting of PMMA sample holders seen as the white rectangles. Hydrogel samples can be mounted in the middle of the PMMA holders, centring the samples along the ultrasound propagation path. The PMMA samples are kept together by a bolt and nut. This design can be used both for transmission measurements, by using T1 as a transmitter and T2 as a receiver, or for pulse-echo measurements, using T1 as both transmitter and receiver.

mounting position, to allow the ultrasound beam to propagate through the hydrogel sample without the beam being transmitted through the PMMA holder. The size of the holes were confirmed to be large enough to not influence the measurements, by inserting the PMMA holders in the test setup without any samples present, and comparing the signal to that when no holder is present in between. No change in the received signal was detected, which confirmed that the part of the ultrasound signal which is received by the transducer is unaffected by the PMMA holder.

For the investigation of the speed of sound, the pulse-echo configuration of the measurement setup was employed, utilizing only one of the transducers. Some of the ultrasound intensity will be reflected at each edge of the hydrogel sample. The time delay between these two reflections can be determined by cross-correlation of the received signal with itself and subsequent interpolation. The time delay between the reflections can thereby be determined and the speed of sound can be calculated as

$$c = \frac{2d}{t} \quad (6.4)$$

where c is the speed of sound, d is the thickness of the sample, t is the time delay, and the factor of 2 enters since the sound from the second reflection travelled back and forth.

Results

The same 25 hydrogel samples which were used to determine the swelling were used for determining the speeds of sound against interlayer exposure time. Thereby, the sample thickness d was known accurately for each sample. Two examples of pulse-echo signals from two different samples, each signal having been averaged over 32 measurements, can be seen in Figure 6.12. One is from a 2 second interlayer exposure time sample and one from a 5 second interlayer exposure time sample. The x -axis shows time from the signal was transmitted, until it was received. The main elements

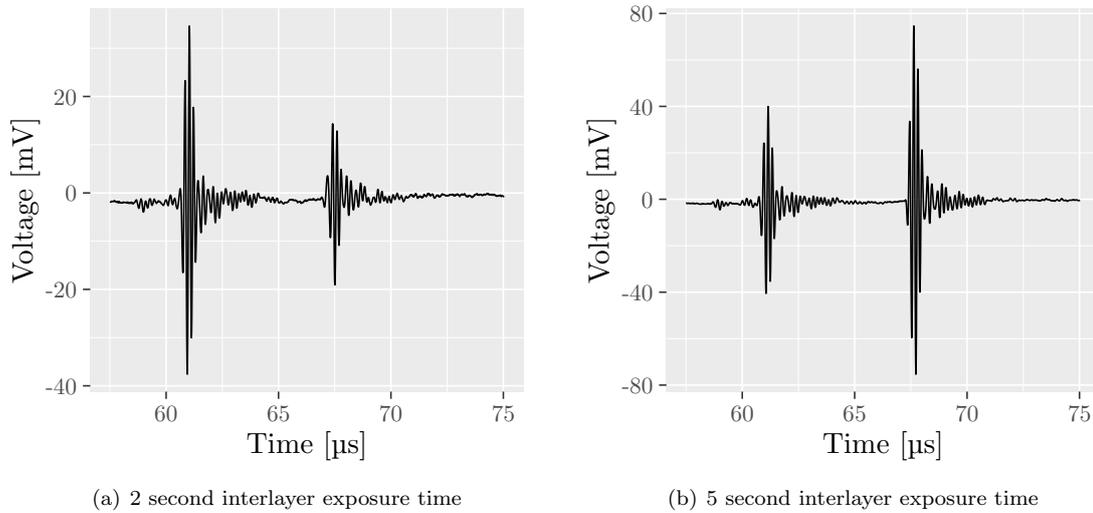


Figure 6.12: Pulse-echo measurements showing the reflections at the front and the back interface of the samples. (a) shows the expected behaviour of a larger reflection at the front surface and a smaller at the back surface, consistent with losses due to the front reflection and attenuation through the sample. (b) shows an example of how some samples showed the opposite behaviour, with a larger reflection at the back interface. The x -axis shows time from the signal was transmitted, until it was received.

of the received signals stems from the two reflecting interfaces seen as two isolated oscillations in both graphs. In Figure 6.12(a), the reflection to the left, stemming from the front interface, is larger than the reflection to the right from the back interface. Due to sound attenuation in the hydrogel material, this is the expected behaviour. The opposite is seen in Figure 6.12(b), where the reflection from the back interface is larger than that from the front. This indicates that the surfaces are not identical. Since the time delay between the reflections is determined through cross-correlation of the signal with itself, and the waveforms are largely unaffected, the time delay estimates are not affected by the difference in amplitudes. The calculated speeds of sound for the samples can be seen in Figure 6.13. The dotted line marks the mean speed of sound of all samples, 1561 m/s. The dashed line marks the mean speed of sound excluding the 13 second, 21 second, and 23 second samples, 1577 m/s. Apart from the 13 second samples which appear as outliers, the speed of sound of the samples with exposure times from 2 seconds to 19 seconds seem unaffected by the interlayer exposure time. For the highest interlayer exposure times of 21 seconds and 23 seconds, the speed of sound appears to decrease to between ≈ 1515 m/s to ≈ 1545 m/s. Overall, these speeds of sound correspond very well to the typical speeds of sound found in tissue shown in the introductory Chapter 2 Table 2.1.

During experimentation it was noticed that the highest interlayer exposure time samples were quite brittle, with some 21 second and 23 second samples de-laminating along the printed layers during handling.

6.4.2 Sound attenuation

Method

By close examination of the images of the samples in Figure 6.8, the sample surfaces on the sides appear more rough than the top surfaces. Analysis of the acoustic data, some of which was presented in Figure 6.12, showed that the roughness on the sides is not necessarily consistent from print to print as the amplitude of the reflections varied between prints. This was also the case for samples printed at the same dose. The data also showed that the surfaces might not be identical, since the reflection from the back surface of the samples was sometimes larger than that at the

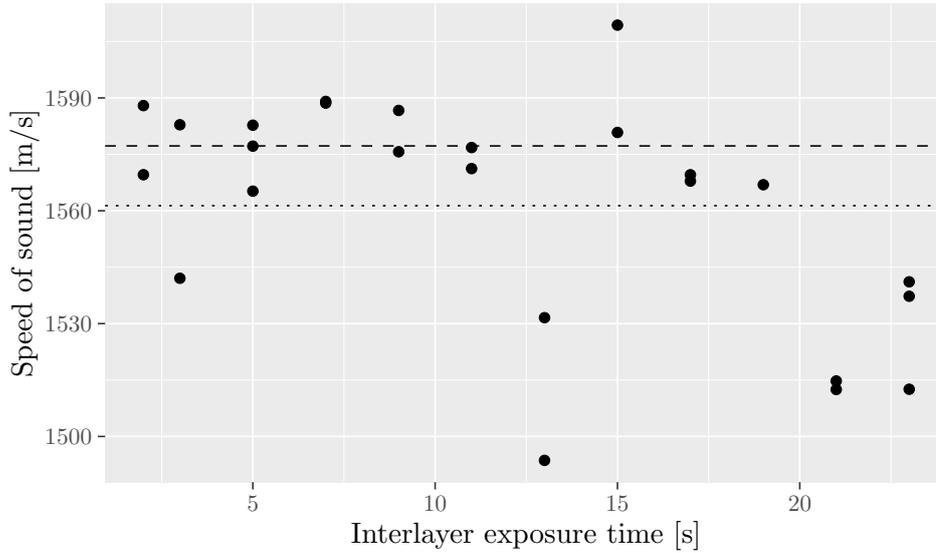


Figure 6.13: Calculated speed of sound against interlayer exposure time. All values are compensated for the measured thickness from Section 6.2.2.

front. While this was not a problem with the applied method to determine the speed of sound, the estimation of attenuation is based on the assumption that the surfaces of the samples are perfectly identical across all different thickness samples. Therefore, due to the roughness of the side sample surfaces, these samples would not be usable for attenuation measurements. Instead, the sample geometry seen in Figure 6.7(b) was used, in which the sound propagation direction is intended to be vertical. The benefit is that the top and bottom surfaces of the printed structures are physically defined in the printing process, with one being defined by the bottom of the vat, and the other defined by the cover glass surface. Therefore, these surfaces are expected to be more uniform, and suitable for measurements of attenuation.

To determine the attenuation, the acoustic measurement setup was used in the transmission configuration, with one transducer transmitting and one transducer receiving. If a pure hydrogel sample should be mounted, it would need to be removed from the glass slide it was printed on, which is impossible to do without modifying the bottom surface of the sample. Any deviation from a perfect flat sample would influence the attenuation measurements similarly to the native roughness on the side surfaces and would therefore be a problem. Instead, four samples of different thicknesses at each of the investigated doses were printed. The used thicknesses were 3 mm, 4 mm, 5 mm, and 6 mm. The print interlayer exposure times do not match exactly with the previous samples, since this print geometry was more time consuming from a manufacturing point of view, resulting in prioritising of the exposure time selection. However, the chosen interlayer exposure times span roughly the same range, namely from 3 seconds to 20 seconds. Higher doses were tested, but the different geometry utilizing the full print area unfortunately only allowed for interlayer exposure times up to 20 seconds to be printed, with higher exposure times resulting in insufficient interlayer cross-binding, resulting in systematic print failure, similarly to the high exposure time samples of the narrow sample geometry used for speed of sound testing.

The successful samples were mounted such that the ultrasound is transmitted through the hydrogel and the glass slide before it is received by the second transducer. By keeping the samples on the glass slide, the interfaces are kept equal between samples. Since the samples are of different thicknesses, the amplitude of the transmitted signal can be fitted across the four samples for each frequency, which eliminates the effect of the glass slide and surface reflections, with the fitted effect being the sound attenuation. This is done in the frequency domain. The attenuation spectra can subsequently be fitted as a power-law model as $\mu = a \cdot f^b$, where f is the ultrasound frequency in

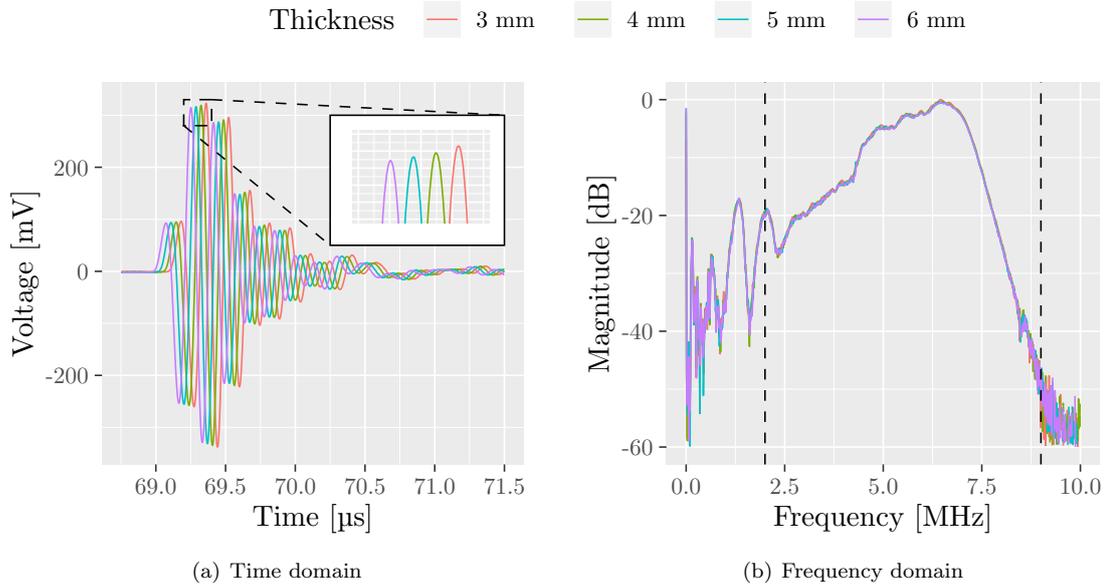


Figure 6.14: Received signal after ultrasound transmission through the hydrogel samples for four different samples printed with an interlayer exposure time of 3 seconds. (a) shows the time domain signal. The x -axis shows time from the signal was transmitted, until it was received. The small inset shows how the amplitude decays, and a time shift is introduced for increasing sample thicknesses. (b) shows the frequency domain of the same data. The attenuation fitted across the four thicknesses in the frequency domain.

megahertz, a is the attenuation coefficient at 1 MHz and b describes the degree of non-linearity of the dependence on frequency [146].

The probing signal was unipolar wideband ultrasound pulse, transmitted and received by two identical ultrasound transducers.

Results

The transmission signals for the four different thickness samples of the 3 second interlayer exposure time can be seen in Figure 6.14. The twofold effect of samples of different thicknesses can be seen in the time domain in Figure 6.14(a). First of all, thicker samples attenuate more of the signal, which can be seen from the decreasing oscillation amplitude. Second of all, since the speed of sound in the hydrogels is larger than that of water, the signals from the thicker samples also reach the receiving transducer faster, seen as the offset of the curves. The frequency spectrum of each of the transmission signals is seen in Figure 6.14(b). All spectra are very similar, essentially showing the frequency response of the system. For frequencies above ≈ 9 MHz, the signal is reduced to the transducer noise floor, and the irregular pattern at low frequencies are likely system artefacts. Therefore only the region between the dashed lines are used for analysis. Careful examination of the frequency response shows that very small differences are seen for the different thickness samples. The best example is that the slight oscillations of the red curve, which is the thinnest sample, are only seen on the top of the group of curves, whereas the purple curve, which is the thickest sample, is trending below for all frequencies. The magnitude for the different samples are fitted at each frequency against the sample thickness to obtain the attenuation spectrum normalised to material thickness. This can be seen in Figure 6.15, to which the power-law model was fitted, seen as the red curve. The black dots show the estimates of the attenuation across the four samples for each frequency, with the error bars showing the standard error of the attenuation estimate. The power-law model was weighed by the standard errors squared. Figure 6.15(a) shows the full obtained spectrum. The values at low and high suffer from large variation, and as a consequence

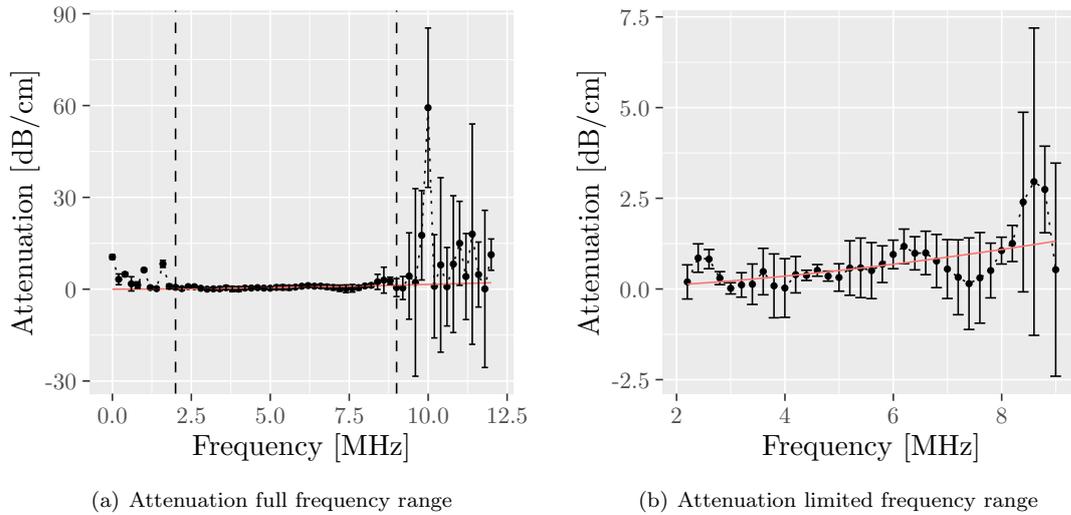


Figure 6.15: Spectra of the fitted attenuation for the different thickness samples printed with an interlayer exposure time of 7 seconds. The error bars are the standard error of the fitted values. The red curve is the power-law fit to the data between 2 MHz and 9 MHz. (a) shows the full frequency range of the data. The standard error and variation in the data are both large outside of 2 MHz to 9 MHz. (b) shows only the fitted range.

also large standard errors of the attenuation estimates. This is the reason for only fitting the power-law model between 2 MHz and 9 MHz. Figure 6.15(b) shows the same data and fit, in the fitted range, in which the overall increasing tendency can be seen.

The similar attenuation spectra for all interlayer exposure times fitted across the different sample thicknesses can be seen in Figure 6.16. Unfortunately, as is particularly evident for the zoom-in in Figure 6.16(b), the power-law behaviour is not seen for all samples in general. The reason for this is unclear, but a possible reason could be that the top surfaces of the prints are not sufficiently uniform from sample to sample as expected. Despite this, the attenuation, when normalised by frequency is similar to the values in the tissue types presented in Section 2.1, spanning 0.6 dB/(MHz cm) to 2.0 dB/(MHz cm). Regardless, the result to not invoke confidence in the experimental method, and therefore the power-law fit parameters are not included. The small differences seen in the magnitudes in Figure 6.14(b) suggest that overall larger samples, and larger differences in thicknesses would be ideal. Alternatively, the small differences might be averaged out by using multiple groups of samples at each dose. These solutions can be tested in the future.

6.4.3 Acoustic impedance

The acoustic impedance can be calculated according to Equation (2.5) as the product between the density and the speed of sound. Figure 6.17 shows the acoustic impedance based on the measured density and speed of sound for all 25 samples. Given that both the speed of sound and the density were fairly constant with interlayer exposure time, the acoustic impedance is as well. The dotted line marks the average acoustic impedance for all samples, 1.634×10^6 kg/(m² s). The dashed line marks the mean acoustic impedance excluding the 13 second, 21 second and 23 second samples, 1.649×10^6 kg/(m² s). Those samples do not seem to follow the tendency, with the 13 second samples appearing as outliers.

Based on the calculated acoustic impedance, the intensity reflection coefficient when the phan-

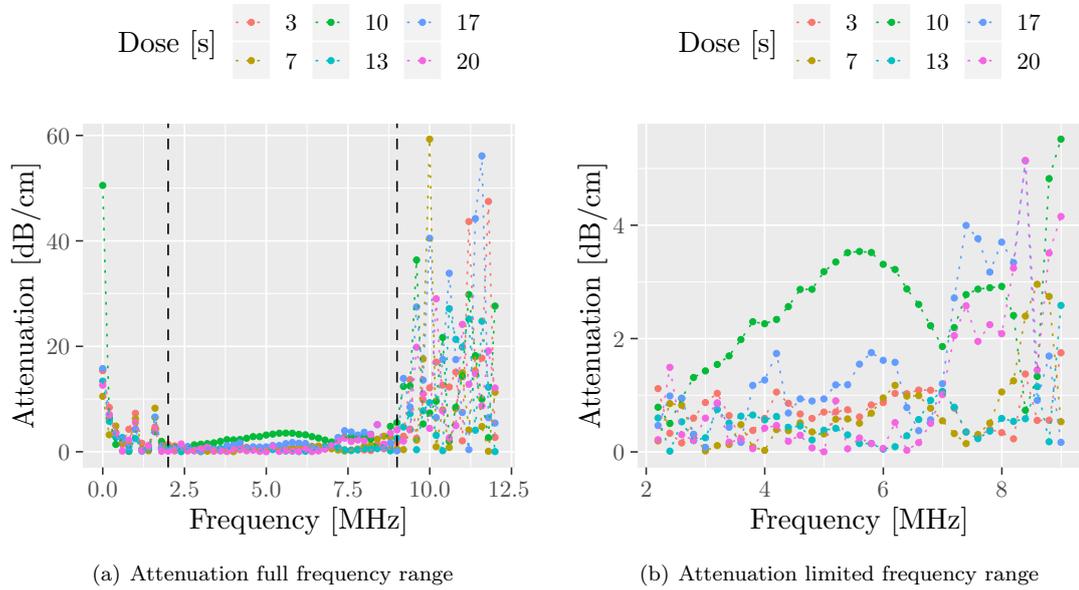


Figure 6.16: Spectra of the fitted attenuation for the different thickness samples for all the doses applied. The error bars are the standard error of the fitted values. (a) shows the full frequency range of the data. The standard error and variation in the data are both large outside of 2 MHz to 9 MHz. (b) shows only the range between 2 MHz and 9 MHz.

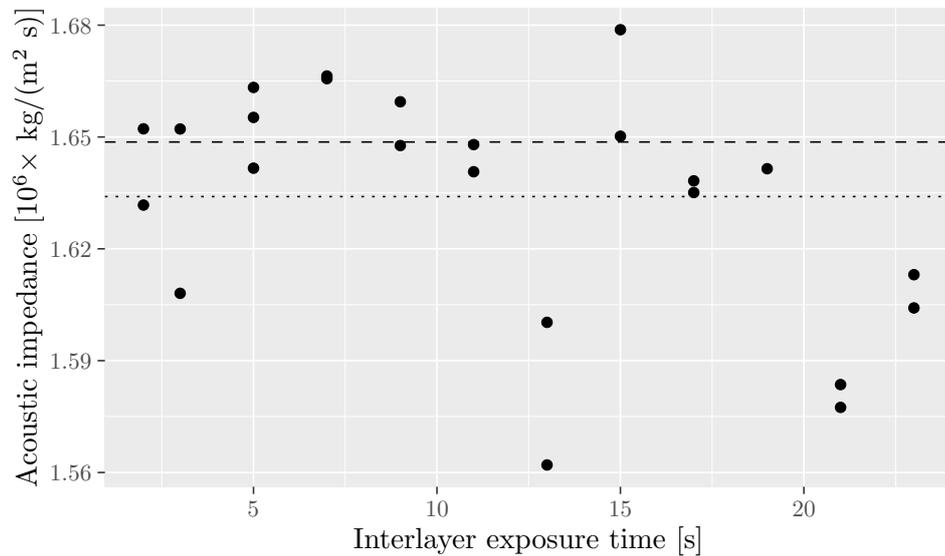


Figure 6.17: Calculated acoustic impedance against interlayer exposure time. The dotted line shows the mean acoustic impedance across all samples. The dashed line shows the mean acoustic impedance excluding the 13 second, 21 second and 23 second samples.

tom is submerged in distilled water becomes

$$R = \left(\frac{Z_{\text{hydrogel}} - Z_{\text{Water}}}{Z_{\text{hydrogel}} + Z_{\text{Water}}} \right)^2 = \left(\frac{1.649 \times 10^6 - 1.48 \times 10^6}{1.649 \times 10^6 + 1.48 \times 10^6} \right)^2 = 0.00291. \quad (6.5)$$

Since the reflection coefficient is so small, the amplitude of the transmitted sound waves will be virtually unaffected, only decreasing by 0.3%.

6.5 Discussion

For the initial optical characterisation of the eight scatterer phantoms, a 36.6 μm residual error was found for the correlation between the designed distances and those measured using an optical microscope. This is significantly larger than the position repeatability claimed by the microscope stage manufacturer and the experimentally validated position repeatability which was tested. A possible explanation might be that the experiment to determine the position repeatability was made by locating the same scatterer multiple times. On the other hand, the correlation in Figure 6.6 was made localising many different scatterers. Local distortion of the printed structures might make the scatterer shapes slightly unequal, resulting in localisation of comparative features (for instance a specific corner) more difficult between scatterers, than when locating the same feature on the same scatterer. It should be noted that the model diagnostics showed that the residuals appeared to be normally distributed, indicating that the model is a good describer for the phantom expansion.

It is remarkable that although a significant difference in swelling is observed, it seemingly has no effect on the speed of sound in the material or on the acoustic impedance. None of the results presented within this chapter explain where the change in acoustic impedance, that was demonstrated by the scatterer reflections in the phantom in the beginning of the chapter, comes from. There is one big difference between the samples tested in this chapter and the printed scatterers. The samples in this section are large prints with uniform properties, which for instance means that the entire print can expand similarly. If instead it is only a small local region which receives a higher dose within a phantom, it will not be able to expand freely if the two doses do not result in equal swelling. The empirically observed difference might thus be a consequence of local stress around the high dose region. This hypothesis is unfortunately seemingly impossible to test. The large uniform sample design was chosen as the precision of the estimates of swelling, density, and speed of sound, all increase with size. Furthermore, in order to cross-correlate ultrasound reflections for speed of sound measurements, received signals from the two reflections need to be separated. Thus, the sample will need to be larger than the probing pulse length, which was a problem for samples smaller than 3 mm. As will be presented in Section 7, local dose changes are in some cases visible in optical microscopes. However, they are not as clear as cavities, and will therefore be impossible to see within the hydrogel samples due to light scattering in the hydrogel. If the high dose areas were moved close to the surfaces as was done with the scatterers when testing the swelling, the stress effects would likely change drastically as it would be released to the surface, making the investigation meaningless. Ultrasound probing of uniform samples while they were being exposed to an external strain would show whether this actually has an effect. However, it is unclear how one could do that in practice.

The sound attenuation study showed that a different approach will need to be taken. It is likely that it is not possible to obtain perfectly similar reflections from each sample. In that case, the only way to mitigate that challenge will be to make multiple prints and average out the variation through statistical analysis. It would also be beneficial to print larger samples. However, even when a proper understanding of the material attenuation has been obtained, the variation in reflection at the surfaces will still be the same, making predictions of the acoustic behaviour difficult.

6.6 Chapter summary

To use a printed structure as a phantom will require confidence in knowledge of the exact location of the printed features. The hydrogel post-printing swelling was investigated, and it was found that it is isotropic, expanding by $\approx 2.6\%$ in all directions when printed with an interlayer exposure of 3 seconds. The swelling changes with the dose to more than 6% for exposure times between 9 and 17 seconds. The hydrogel density increases slightly with dose. Although the printer resin contains 20% PEGDA, the printed structures contain between 21% and 27%, which means the printer resin will be slightly depleted during printing. The speed of sound appears unaffected by the dose, with an average speed of sound of 1577 m/s for interlayer exposure times between 2 seconds and 19 seconds. For higher exposure times it decreases to ≈ 1530 m/s. Experiments to determine the sound attenuation was conducted, however, the print surfaces appear to vary too much in roughness, resulting in a lack of confidence in the data.

Calibration phantoms for SRUS

This chapter describes the work towards an alternative from the conventional flow channel phantoms for SRUS validation - fixation of sub-wavelength scatterers. The fixated scatterers started simply as fiducial markers for flow phantoms, in order to be able to align the phantom to the ultrasound probe with micrometer precision. However, it quickly became apparent that they could be utilised directly for a new type of phantom. The content of this chapter is in part based on Paper B, Paper D, Paper G and Paper I.

7.1 A new type of phantom for SRUS validation

Conventional phantoms used for SRUS, as those described in Section 2.4.1, consists of tubes which are supposed to mimic the micro-vasculature of tissue. The methods work as intended, providing outer boundaries for micro-bubbles. However, most phantom types do not provide a true three-dimensional structure, and none of them provide any control of positioning of the micro-bubbles with the precision that the SRUS techniques are supposed to provide.

When we made our first 3D printed flow phantoms, which was shown right at the start of this thesis in Section 2.4.2 and which will be discussed further in Chapter 8, we realised the importance of being able to align the ultrasound probe to the micro-channels. That motivated the experimentation with fiducial markers, and inclusion of them in subsequent phantoms, which was mentioned in Section 6.1, and some of which was shown in Figure 6.1. Aligning a 2D image plane in 3D requires recognizable and visible structures. The structures should be able to be fixated in the phantom, be able to be imaged repeatedly to align the probe to the phantoms with high precision. It was realised that these same properties would in themselves be ideal for a phantom for characterisation of SRUS pipelines. That led to the development of a phantom containing what had previously been considered fiducial markers, but now without any accompanying micro-channel. These scatterers would be fixated in the phantom and stable over time, in direct contrast to conventional flow phantoms. It would be possible to conduct multiple experiment separated in time with the same phantom, and end up with the exact same results, which would be impossible if one was using micro-bubbles.

The basis for the phantom was the cavity scatterer, since the reflection from a solid scatterer was much smaller, as shown in Section 6.1. The cavity scatterer concept is illustrated again in Figure 7.1 for convenience. In Figure 7.1(a) the yellow region is the illuminated region, ending up as hydrogel, and the black region is the region which receives no dose, with no cross-linked polymers. The small yellow squares represent the individual voxels. The sketch is scaled similarly

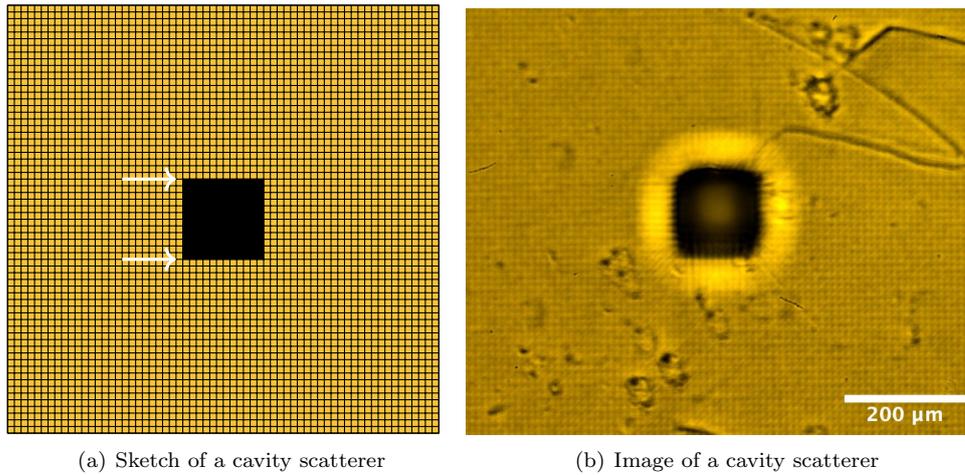


Figure 7.1: The cavity scatterer concept. (a) is a sketch of the cavity, with black indicating an unexposed region, and yellow indicating printed hydrogel. The small squares frame the voxels, and are matched in size to (b), a microscope image of an actual printed scatterer placed at the top of a print. Both scatterers are designed to be 12 voxels wide. The white arrows mark depths where the acoustic impedance is expected to change significantly.

to the image in Figure 7.1(b), where the voxel grid is also visible. The scatterers are in both cases designed to be 12 voxels wide.

Since ultrasound is reflected at interfaces between media of different acoustical impedances, such a scatterer will actually reflect the ultrasound twice as indicated by the white arrows in the figure: first at the top interface between the hydrogel material and the cavity, and then at the bottom interface when the sound propagates from the cavity into the hydrogel again. This is why two reflections were observed for each cavity in Figure 6.1. In that experiment, the imaging frequency was 15 MHz, corresponding to a wavelength of $\approx 100 \mu\text{m}$. If using an imaging frequency of 3 MHz instead, the wavelength becomes $\approx 500 \mu\text{m}$, and the scatterer will end up appearing as a single point target.

Comparing Figure 7.1(a) and Figure 7.1(b) it can be seen that the actual printed structure is not necessarily a perfect replica of the design. In this case, the printed cavity is actually slightly larger than the design. A good understanding of what to expect when printing is crucial, making it important to determine the correlation between the designed dimensions of scatterers, and the actual printed ones.

7.2 Micro-engineering of the 3D printed scatterers

7.2.1 Concept description

The printer system uses a DMD with a micro-mirror size of $10.8 \mu\text{m}$ by $10.8 \mu\text{m}$. That means that the ultimate resolution of the system is $10.8 \mu\text{m}$. However, it is not possible to obtain $10.8 \mu\text{m}$ features in practice. The optical components have been implemented such that the light is parallel and illuminates a 1:1 replication of the DMD in the vat. However, once it reaches the vat, the light is transmitted through the glass bottom of the vat, and through the liquid resin itself. Once the resin starts polymerising, it will scatter the light slightly, resulting in feature broadening. This means that the printed features sizes do not necessarily match perfectly with the design feature sizes. And the extent of this discrepancy will likely change based on the implemented dose. The intuitive effect will be that a cavity printed with a high dose will become smaller than it was designed to be, as illustrated in Figure 7.2, and conversely, a cavity printed with a low dose might not have sufficient cross-linking at the edges, resulting in the cavity becoming larger than it was

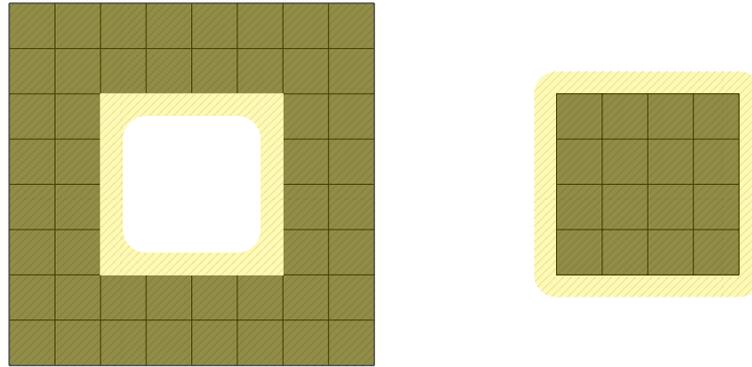


Figure 7.2: Sketch of feature broadening. The grey region is the desired pattern for illumination, and the squares represent the $10.8\ \mu\text{m}$ by $10.8\ \mu\text{m}$ DMD micro-mirrors. The yellow is the unintended exposure due to light scattering. This illustrates how a designed cavity will end up becoming smaller than intended due to light scattering in the transparent vat bottom and the liquid resin, and that a local region of illumination will result in widening of the printed feature.

designed to be.

As the side length of the scatterers are decreased, this will eventually result in designed cavities being illuminated and closed in the actual print due to feature broadening. To have a better understanding of the capabilities of the 3D printer system, it is important to investigate how the printed features deviate from the design. A correlation between the design and printed features sizes can then be made, revealing the practical resolution limit of the printer system, and allowing for compensation of design features to realise the desired feature sizes. Results from our other work in Paper D [47], based on 2D ultrasound imaging have demonstrated reflecting scatterers designed to be $75\ \mu\text{m}$ by $75\ \mu\text{m}$. In this work, we are looking for the lower limit of scatterer sizes, both directly in terms of actual printed sizes, but also in terms of the reflected intensity, by investigating designed side lengths in the range from $32.4\ \mu\text{m}$ to $129.6\ \mu\text{m}$.

As was shown empirically in Section 6.1, overexposing a region of hydrogel will also modify its acoustic impedance. However, the analyses of the acoustic properties presented in Section 6.4 provided no clear indication of what that change in acoustic impedance stems from. The empirical evidence suggests that there might be other effects involved when it is only a local region which experiences the dose change, perhaps local stress effects due to local differences in swelling. Building on this it was considered whether it would be possible to utilize the overexposure in combination with the cavity scatterers.

Figure 7.3 illustrates the implemented dosing schemes. Figure 7.3(a) shows the regular scatterer, which is simply a non-illuminated region surrounded by hydrogel printed using only the base dose. Figure 7.3(b) implements a single voxel wide overexposure at the edge of the scatterer. Figure 7.3(c) implements a dose gradient to increase the cross-binding of the hydrogel to a maximum at the edge of the hollow cavity. The Gradient dose scheme is based on the “Black Silicon” concept from the silicon micro-fabrication industry [147, 148, 149]. By gradually increasing the dose, the acoustic impedance could be expected to gradually change similarly, which with a perfect gradient would result in no reflection from the gradient region.

These three dosing schemes are implemented in all the tested phantoms. The base dose corresponds to a 3 second interlayer exposure time, and the maximum interlayer exposure time is 23 seconds. For the gradient, the interlayer exposure time is decreased from 23 seconds at the edge of the cavity by one second for each voxel, reaching the base interlayer exposure time after 20 voxels.

Signal interference for sub-wavelength features

The potential interfaces of the scatterers with changing acoustic impedance which would reflect sound are marked by arrows in Figure 7.3. Given that the scatterers in general are designed to be

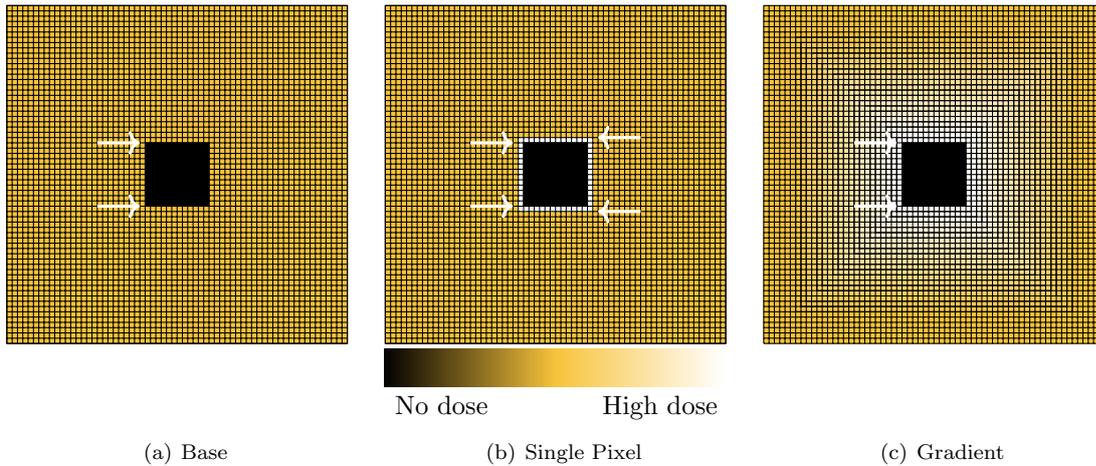


Figure 7.3: Sketch of the different dosing schemes applied in this work. In (a), only the base dose is used around the cavity. In (b), only a single voxel wide increased dose is used. In (c), a dose gradient gradually increases the dose from the base dose, up to a maximum at the edge of the cavity. The arrows mark depths where the acoustic impedance is expected to change significantly.

of sub-wavelength sizes, the two interfaces should not be separable due to the diffraction limit of the scanner system. The sub-wavelength separation does however mean that the reflected signal from the two interfaces might interfere with each other, and depending on the separation, the phase shift between the two waves will vary. Figure 7.4 illustrates the case of two reflecting interfaces, separated by 1 voxel, 4 voxels or 12 voxels. The interface dictates whether interference might be almost completely constructive, as for small separations, illustrated by a separation of 1 voxel and 4 voxels, or almost completely destructive, as for 12 voxels. The illustrations are based on a 3 MHz single period sine pulse, propagating in a hydrogel with $c = 1577$ m/s. The top x -axis is spatial separation referring to the interface separation, and the bottom x -axis is a temporal axis referring to the waveforms. The two x -axes correlate 1:2 due to the compensation factor of 2 when converting temporal reflection signals to spatial positions in ultrasound, but are aligned to show the correlation of the scattering source and the reflected waveform.

Losses are not considered in these illustrations. The reflection coefficient at an interface between water and hydrogel was shown to be only $\approx 0.3\%$, which therefore does not change the transmitted amplitude significantly. The distances between the interfaces are only of a few micrometer, whereas the expected magnitude of attenuation is only a few dB per millimeter, and will therefore also only have a minor influence on the amplitude of the transmitted signal. Thus, the waveform transmitted at the first interface will have an almost unaltered amplitude, therefore also resulting in a similar amplitude reflection at the second interface, making the illustration reasonably representative. The waveforms are simplified to ideal sinusoidal oscillations to illustrate the interference. In reality, the transmitted waveforms have also been modified by the transducer impulse response, which in general makes it less symmetric, therefore also less likely to have any part of the resulting waveform be completely removed by destructive interference.

For all the dosing schemes, the transitions from hydrogel to the water containing cavity and back again will result in reflections. In the case of a 12 voxel wide scatterer, the 12 voxel example is a reasonable representation. For both the dose gradient and the single pixel overexposure, the largest cross-linking density will be attained at the edges of the cavity. However, for the single pixel overexposure, there will also be a sharp difference in the amount of cross-linking between the base dose region and the increased dose, across only a single voxel in the design. This is illustrated by the 1 voxel example, showing that the signals from the two interfaces will add up to a larger amplitude than either of the individual reflections. Of course, due to feature widening, the actual separation between the interfaces might be closer to that shown for four voxels. However, even

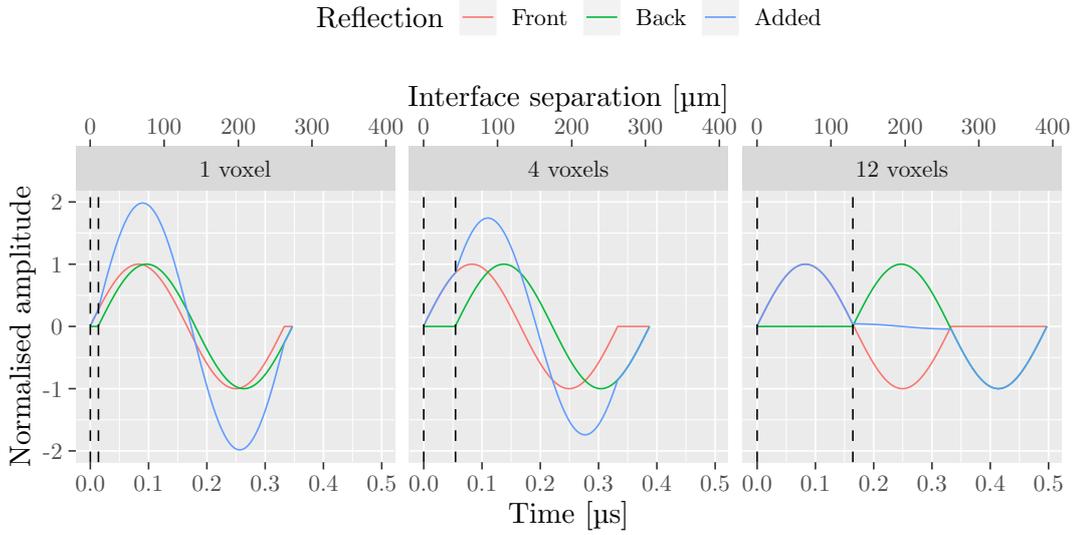


Figure 7.4: Illustration of interference between reflected waveforms at two interfaces of a scatterer. The interface dictates whether interference might be almost completely constructive, as for small separations, illustrated by a separation of 1 voxel and 4 voxels, or almost completely destructive, as for 12 voxels. The illustrations are based on a 3 MHz single period sine pulse, propagating in a hydrogel with $c = 1577$ m/s. The top x -axis is spatial separation referring to the interface separation, and the bottom x -axis is a temporal axis referring to the waveforms. The two x -axes correlate 1:2 due to the compensation factor of 2 when converting temporal reflection signals to spatial positions in ultrasound, but are aligned to show the correlation of the scattering source and the reflected waveform.

then, the waveforms constructively interfere for a resulting reflected waveform of larger amplitude than the individual reflections.

Of course, the two reflections due to the single voxel overexposure is present at both sides of the cavity, which means the resulting waveform interference might look as in Figure 7.5. The same parameters were used for the ultrasound waveform. The interference essentially becomes a mixture of the cases presented in figure 7.4. “1 voxel/12 voxels” illustrates interference according to the design structure with only a single voxel of overexposure for a 12 voxel wide scatterer. “4 voxels/10 voxels” is an example more representative of the actual printed scatterer, as will be presented in the following. The analysis shows that a 12 voxel Single Pixel scatterer on average will be printed approximately 10 voxels wide due to feature widening, and the overexposure region optically appears to be widened to approximately 4 voxels.

Test phantom design

In order to test the effect of the dosing schemes on the printed cavity size, and the resulting reflected intensity, the phantom layout seen in Figure 7.6 was designed. Each black spot is a cavity, and the black region in the top left corner is used for orientation purposes. The phantom is split laterally into three regions, one for each of the dosing schemes. Each of the three regions are separated into four groups of ten scatterers, with each group containing ten different sized scatterers. The shape of the scatterers is changed between being square and circular. The sidelength or diameter was changed from three voxels to twelve voxels within each group of ten scatterers. In order to avoid pairing the shapes with specific scatterer sizes, the scatterer positions were permuted by one column from one group of ten scatterers to the next. Six phantom configurations were made in total, one for each of the six different permutations of the three dosing scheme positions, for a total of 720 scatterers. For the ultrasound intensity experiments, the scatterers were positioned in the middle of the phantom, and each scatterer was printed 3 mm long to take advantage of integration of signal

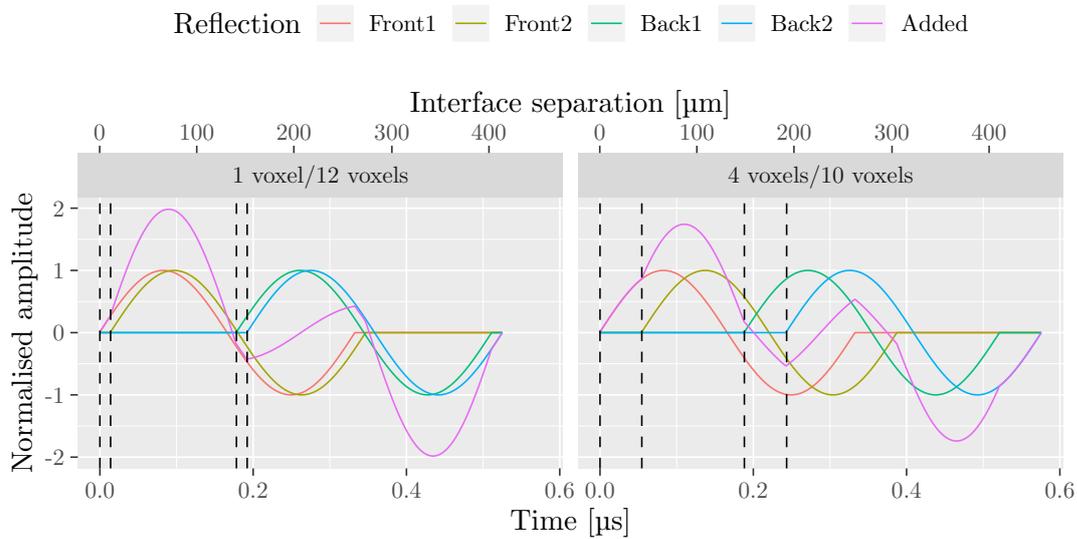


Figure 7.5: Illustration of interference between reflected waveforms at four interfaces of a Single Pixel scatterer. The interference pattern becomes a combination of the cases presented in Figure 7.4. “1 voxel/12 voxel” illustrates the design separation of only a single pixel of overexposure at each side of a 12 voxel wide scatterer. In the following sections, it is shown that the overexposure region appears to widen to approximately 4 voxels, and the cavity narrows to approximately 10 voxels, illustrated by the “4 voxels/10 voxels” curves. The illustrations are based on a 3 MHz single period sine pulse, propagating in a hydrogel with $c = 1577$ m/s. The top x -axis is spatial separation referring to the interface separation, and the bottom x -axis is a temporal axis referring to the waveforms. The two x -axes correlate 1:2 due to the compensation factor of 2 when converting temporal reflection signals to spatial positions in ultrasound, but are aligned to show the correlation of the scattering source and the reflected waveform.

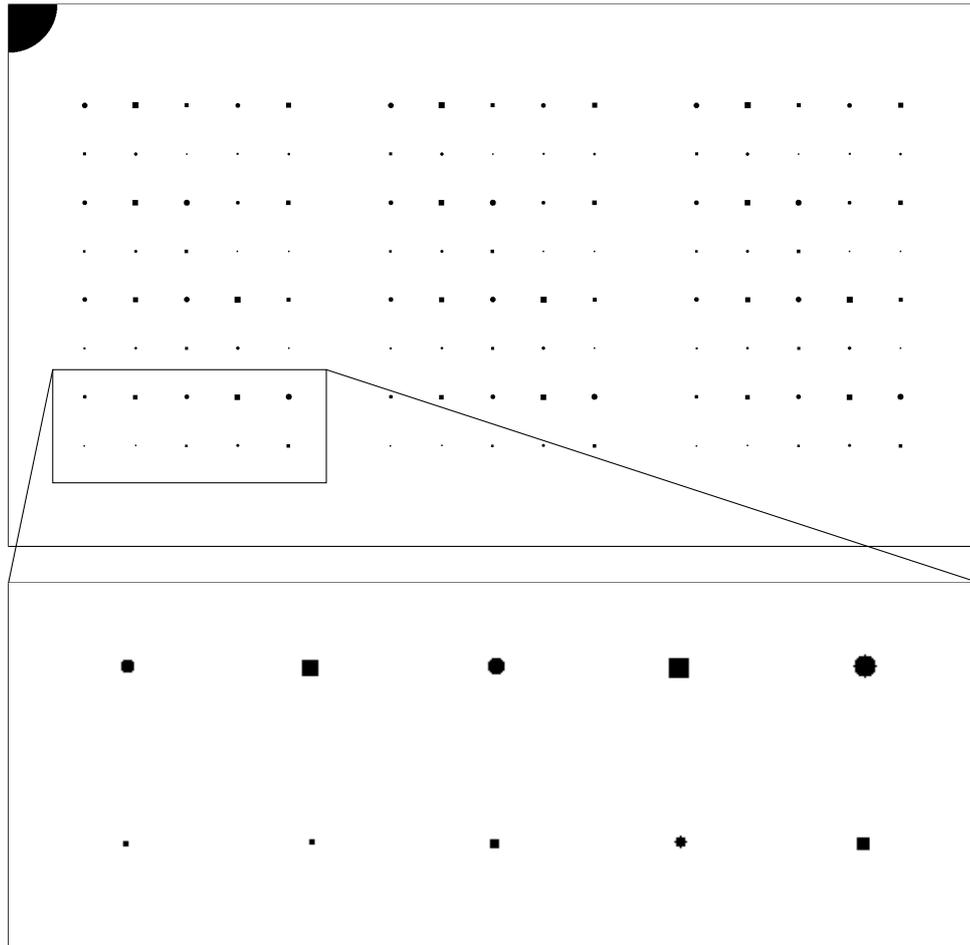


Figure 7.6: The designed phantom layout. The three groups which are separated laterally, are the three different dosing schemes. The zoom-in shows how each subsection of two rows contains ten different sizes, with the shape changing between square and circular.

across the elevation direction. However, as noted in [150], the hydrogel scatters light, rendering it impossible to image printed features in the middle of a phantom. Instead, six phantoms with the same scatterer configurations were printed, with the scatterers being placed near the surface of the phantom. This allows for investigating the dimensions of the scatterers using optical microscopes.

These phantom designs allow for analysing the effect of many different factors on the actual printed scatterer size and the resulting reflected intensity.

For the optical microscope investigation of the printed scatterer size, it is first of all possible to test how the printed size changes with the design size, the shape, and the dosing schemes, and any interactions between these factors. Permutation of both the dosing scheme and the scatterer position ensures that they are not confounded with the lateral (column) position of the scatterers in the phantom. By replicating and permuting the scatterer sizes in the four groups within each dosing scheme, the same is true for the vertical (row) position of the scatterers in the phantom. Thereby, it becomes possible to test whether the printer itself contributes with an effect on the scatterer size, depending on the position in the phantom (row, column, both linear and second order interactions), and remove this effect from the analysis of the dependence on design size, shape and dosing scheme. Finally, since each dose scheme is printed in the same position in two separate phantoms, it is possible to test whether there is any random variation between the printed phantoms, which would represent the printer variability, once again, with the purpose of removing this effect from the analysis of the dependence on design size, shape and dosing scheme.

For the ultrasound intensity experiment, the same layout is used, and all of the same effects can be tested. In addition, by measuring each phantom twice, with the phantom being flipped 180 degrees in one of the measurement, it is possible to investigate whether the imaging system provides a uniform ultrasound field, or whether this has a linear dependence of the lateral position in the field of view of the ultrasound probe. It is also possible to test for a quadratic effect on the lateral position. However, this will be confounded with the printed position. Thus, in principle it will not be possible to know whether any quadratic effect seen is due to the printer system or the ultrasound system. However, in combination with the optical analysis of the sizes, it might still be possible to get an indication of an effect.

The number of factors to include and test results in an large model with a lot of factors. It is important to note that it is primarily the main factors: the correlation to the design size, the shape, and the dosing scheme, and their respective interactions. Inclusion of all of the other factors in the analysis allows us to remove the effects of these factors, to determine the true underlying effects of the main factors, as well as get some insights into the printing and measurement systems.

The scatterer size range is set to suit 2D imaging, in which the scatterers can be elongated in the elevation direction, for summation of signal. A similar experiment could be carried out for 3D imaging characterisation. This would likely require that the scatterer size range is increased, since it would no longer be possible to benefit for the poor elevation focus in 2D imaging.

7.2.2 Experimental setup

Optical characterisation setup

A Zeiss Axioskop 40 optical microscope equipped with a 10x magnifying lens was used for optical characterisation of the side lengths and diameters of the phantoms. The phantoms were submerged in MQ water to avoid water evaporation from the hydrogel phantoms during inspection. An IDS UI-3280CP-C-HQ camera on the microscope was used to acquire the images of the individual scatterers. The size of the images were 2456 x 2054 pixels, with a sensor resolution of 0.3423 $\mu\text{m}/\text{pixel}$ at the used magnification, corresponding to a total field of view of 840 μm x 703 μm . The images were captured using $\mu\text{Manager}$ [151] and subsequently analysed in Fiji [152].

Ultrasound experiment setup

A new experimental setup was built for all the following SRUS experimentation. When aiming to measure position changes on the order of a few micrometers, vibrations of the measurement setup

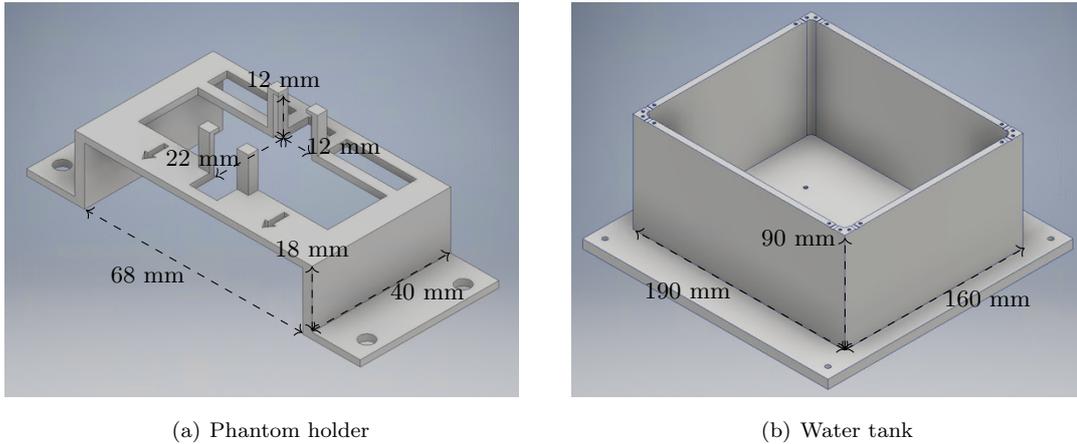


Figure 7.7: 3D models of the (a) 3D printed phantom holder fitted for the $\approx 21 \times 12 \text{ mm}^2$ hydrogel samples, and (b) water tank of the newly developed SRUS experimental setup, which the phantom holder fits into.

Figure 7.8: 3D models of the experimental setup. (a) is the 3D printed phantom holder, which has been fitted to hold the the $\approx 21 \times 12 \text{ mm}^2$ hydrogel samples in the central position. (b) is the water tank of the newly developed SRUS experimental setup, which the phantom holder fits into.

or inadequate fixation of the phantom to the measurement stage will be detrimental. Therefore, a holder system consisting of phantom holder designed for high precision mounting in a water tank was designed. The 3D printed holder 3D model can be seen in Figure 7.7(a). It was fitted to the phantom dimensions enabling mounting of the phantoms on top of an absorbing polyurethane rubber sheet (Sorbothane, Inc., Kent, Ohio, USA). It only holds the phantom by each corner. This serves two purposes: It allows for design freedom since the needle from the flow controller can be inserted almost anywhere on the phantom and the cut-outs minimise ultrasound signals from the holder itself. The arrows allows for systematic mounting of the phantom in the holder, as well as the holder on the remaining system. The water tank was milled in aluminium, and the 3D model can be seen in Figure 7.7(b). Small recesses designed for mounting tubes from the flow controller were milled out at the top of the water tank walls. Thereby, if too much translation or rotation of the water is accidentally carried out, the pulling of the tubes will not pull at the phantoms, which could destroy them, but will only pull at the aluminium water tank. The water tank was mounted on a 8MR190-2-28 rotation stage (0.01° resolution) combined with a 8MTF-75LS05 x - y translation stage ($0.31 \mu\text{m}$ resolution) (Standa, Vilnius, Lithuania). To minimize the effect of vibrations, everything was mounted on a Newport PG Series floating optical table (Irvine, California). A sketch of the combined setup can be seen in Figure 7.9. For this experiment, a BK Medical "Hockey Stick" X18L5s probe was used with a BK 5000 experimental scanner to acquire the 2D ultrasound images. All images were obtained with an imaging frequency of 15 MHz, and the field of view was set to rectilinear imaging, imaging only that which is directly below the transducer footprint. The B-mode images used for analysis were averaged over 50 acquired ultrasound B-mode frames.

7.2.3 Printed scatterer size

Figure 7.10 shows three microscope images of square scatterers designed to have a side length of 12 voxels, or $129.6 \mu\text{m}$, one of each of the three dosing schemes: a) Base, b) Gradient, and c) Single Pixel. More microscope images are included in Appendix J.2. The yellow regions are the hydrogel, and the black are the scatterer cavities. The images are scaled equally. The square pattern seen in the hydrogel regions are the individual voxels. The regions of different dose are

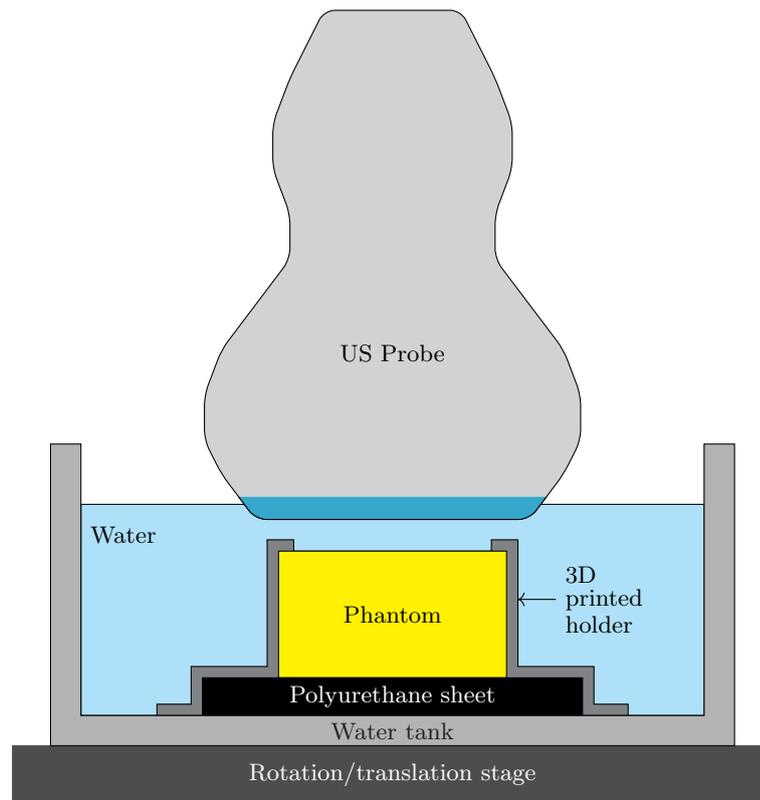


Figure 7.9: Sketch of the experimental ultrasound setup.

visibly distinguishable, with the 20 voxel wide dose gradient in Figure 7.10(c) particularly clear, seen as a large square surrounding the actual scatterer. It should be noticed that the bright yellow regions are not in general correlated to the amount of cross-linking since Figure 7.10(a) shows a bright yellow frame around the cavity, without there being any additional dose, which looks similar to the bright yellow frame of the Single Pixel cavity. Furthermore, there is a wide apparently uniform frame for the Gradient cavity, which could be falsely interpreted as the dose only changing right at the cavity, and 20 voxels away from it, whereas the dose has actually been changed gradually from the outer edge to the cavity. The optical effects might be a stress related artefact.

The side lengths and diameters were measured for all scatterers across the six phantoms. The measured side lengths and diameters of all scatterers are presented in Figure 7.11, plotted against the designed size in terms of voxels. The discrete grouping of the data along the x -axis is due to the designed size not being completely free, but limited to an integer number of voxels. The colours group the data into the different dose schemes, and the shapes group the data into the two different cross-sectional shapes. The solid and dashed lines are the statistical models of the square and the circular cross sectional scatterers respectively, for the different dosing schemes. With all of the factors of the phantom design in mind, a lot of information is compounded into the plot. Even the distribution of the different groups of data points are difficult to discern, due to the number of data points. The following figures offer alternative perspectives on the same data of the isolated effect of the main factors.

Figure 7.12(a) shows a box-plot of the measured side length against the two different scatterer shapes. The centre line in each box is the median value, and the lower and upper edge of the boxes correspond to the 25th and 75th percentile of the data respectively. Each box contains the data points across all sizes for the two shapes, resulting in the wide extend of the boxes. Even so, the median offset between the two shapes indicates that the square scatterers are generally larger than

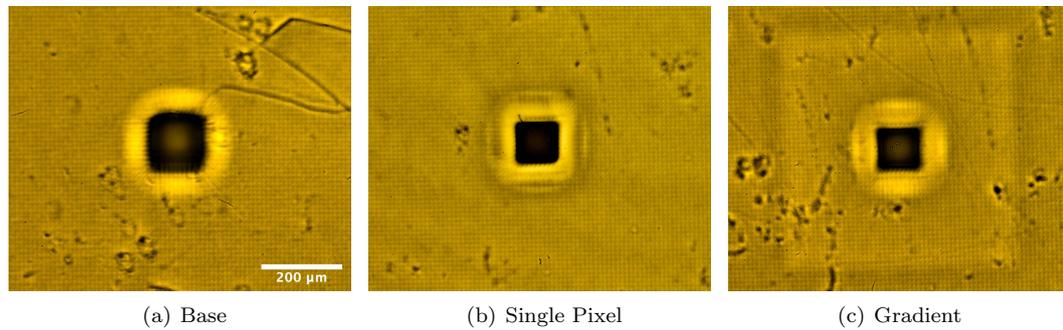


Figure 7.10: Optical microscope images of individual scatterers. The yellow regions are printed hydrogel. The scatterers were designed to be $129.6 \mu\text{m}$ wide. (a) is Base dosing scheme, (b) is Single pixel, and (c) is Gradient. The scale bar is common for all images. Additional images, with examples of all sizes, shapes and dose schemes can be seen in Appendix J.2.

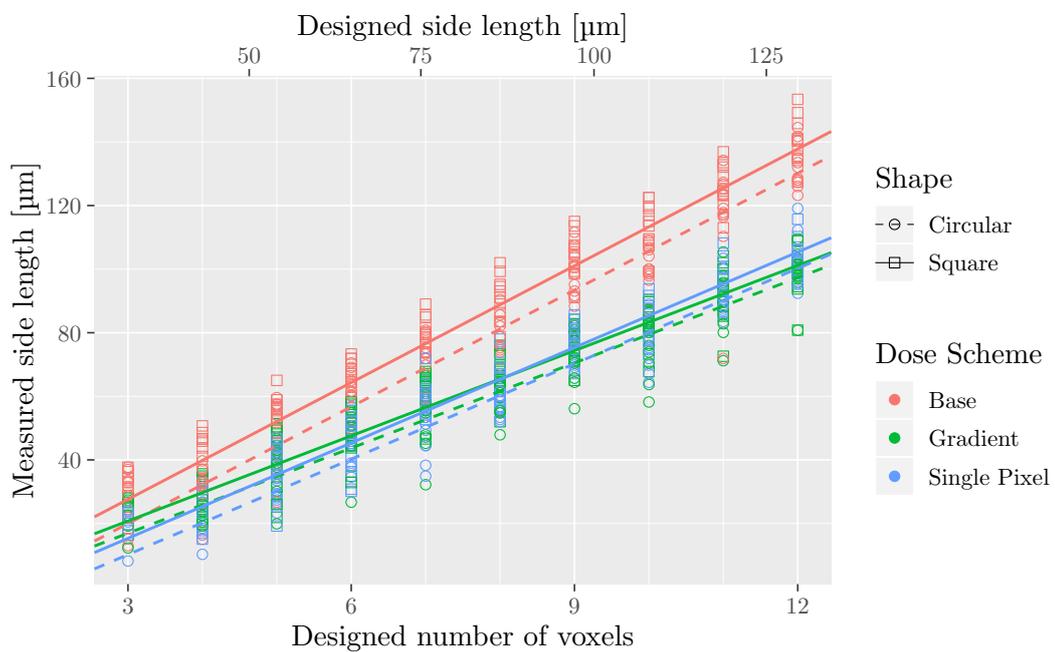


Figure 7.11: Measured side length or diameter against the designed scatterer size. The colours group the data into the different dose schemes, and the symbols group the data into the two different cross sectional shapes. The solid and dashed lines are linear fits to the square and the circular cross sectional scatterers respectively.

Table 7.1: Summary of the variables and their data types used in the following analyses. The variables in the top block are tested both in the scatterer size and scatterer intensity analysis, while the factor in the bottom block is only applicable in the scatterer intensity analysis. The parenthesis under variable type indicates whether the factor provides insight about the general scatterer size correlation, the printer uniformity, or the ultrasound uniformity.

	Sample values	Variable type	Description
Measured size [mm]	28.8, 39.7, ..., 153.4	Numerical values	The response variable, optically measured
Design size [voxels]	3, 4, ..., 12	Numerical values (<i>Phantom</i>)	The designed size of the scatterer
Shape	Square, Circular	Fixed factor (<i>Phantom</i>)	The three different dosing schemes
Dosing Scheme	Base, Single Pixel, Gradient	Fixed factor (<i>Phantom</i>)	The three different dosing schemes
Row	1, 2, ..., 8	Fixed factor (<i>Printer</i>)	Row position in the scatterer grid
PrintColumn	-8, -7, ..., 8	Fixed factor (<i>Printer</i>)	Column position in grid relative to the printer DMD
Phantom	1, 2, ..., 6	Random factor (<i>Printer</i>)	The six different printed phantoms
Phantom:Flip	1:0, 1:180, ..., 6:180	Random factor (<i>Ultrasound</i>)	Factor checking for random variation between the B-mode images
ImageColumn	-8, -7, ..., 8	Fixed factor (<i>Ultrasound</i>)	Column position in grid relative to the ultrasound probe

the circular scatterer.

Figure 7.12(b) shows a box-plot of the measured side length against the three different scatterer dose schemes. Again, the content of each box is compounded across other factors, such as the scatterer size, and scatterer shapes. The range of the Base dose scheme is larger than the others, with a noticeable offset of the median value to that of the other dosing schemes.

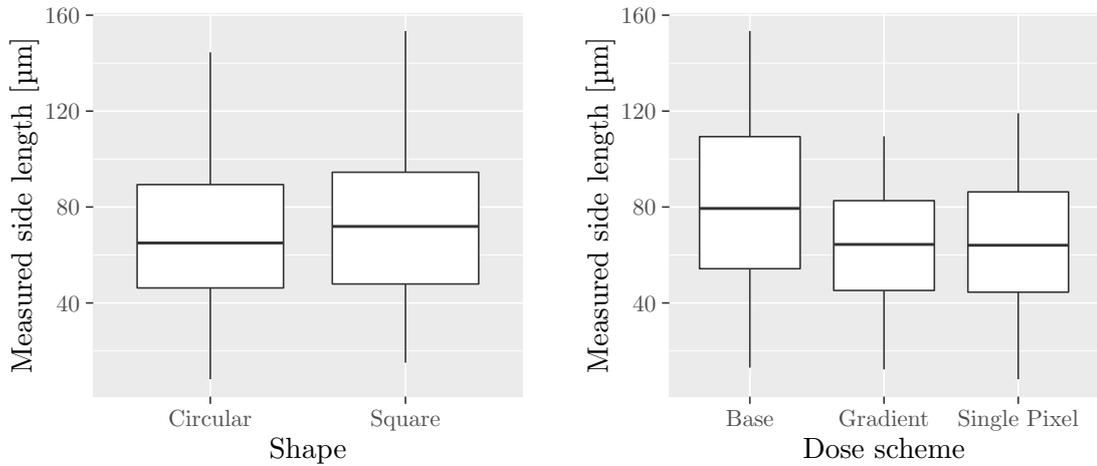
Figure 7.13 shows box-plots of the measured side length or diameter against the designed scatterer size, separated into the different dose schemes and shapes. Dots mark outlier values, with outliers being defined as measurements further than 1.5 times the ICR away from the nearest box edge, with ICR being the distance between the 25th and the 75th percentile. All plots are scaled equally to allow for easy comparison. The box plots provide a good overview of the point distributions within each group. It can be seen that there is a different slope of correlation between the data groups for different dose schemes, and slight apparent offsets depending on the shapes.

The plots show that the actual printed size of the scatterers printed using only the Base dose are generally larger than those printed with a Single Pixel overexposure, or with the Gradient overexposure.

Scatterer size statistical analysis

In Section 7.2.1, the factors which could be analysed due to the chosen phantom layout were listed. A summary of the data types and the factors included in the analysis can be seen in the top block in Table 7.1. Under variable type is also written whether the factor provides insight about the general scatterer size correlation or the printer uniformity. Interactions between factors, and quadratic effect of position in the scatterer grid were also investigated.

The full model of all the factors and interactions investigated can be seen in Appendix E.1, along with model diagnostics and more summarizing plots. The combination of fixed and random factors makes the fitted model a linear mixed effects model. Such a model can be analysed using the



(a) Measured side length or diameter against shape

(b) Measured side length or diameter against dose scheme

Figure 7.12: Measured side length against (a) against the scatterer shape and (b) against the dose scheme.

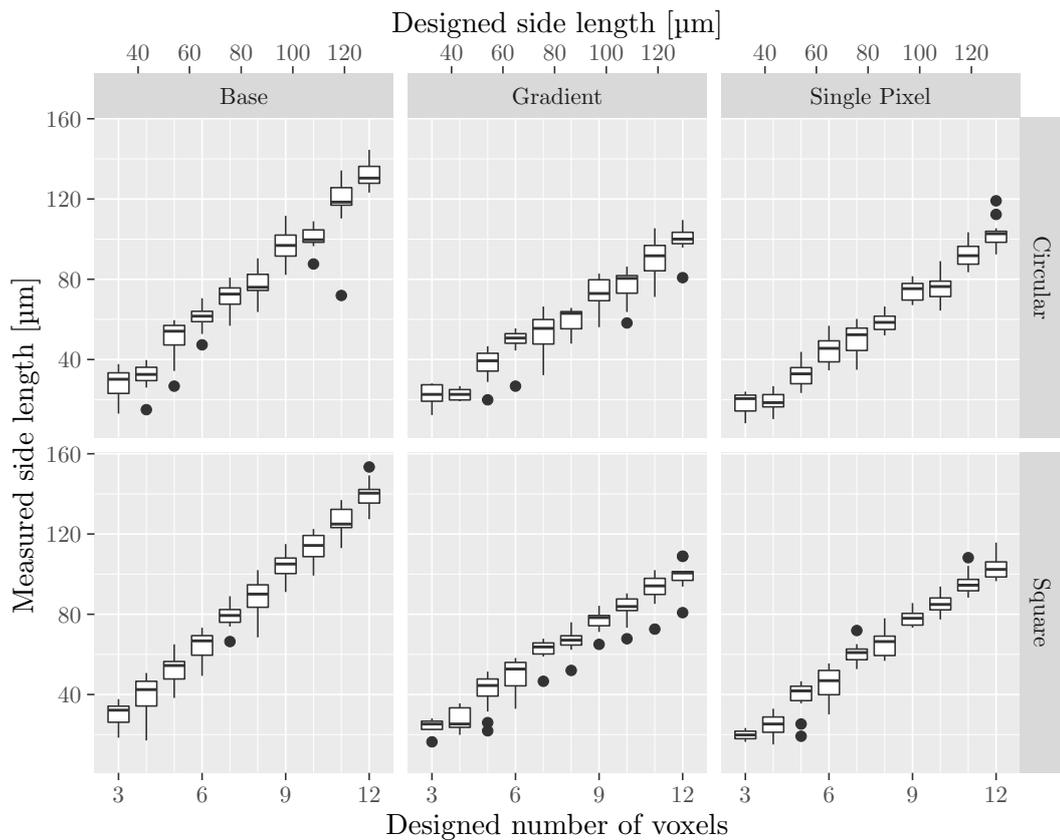


Figure 7.13: Measured side length or diameter against the designed scatterer size. The data is separated into the dose schemes and shapes for a better overview of the data distributions. Dots mark outliers.

Table 7.2: Model coefficients for the reduced model of the printed scatterer size, with confidence intervals and p -values.

Predictors	Estimate	2.5%	97.5%	p -value
Fixed effects				
μ [μm]	-16.8	-20.4	-13.1	<0.0001
α_1 (Square) [μm]	7.6	6.2	9.1	<0.0001
α_2 (Gradient) [μm]	6.9	3.6	10.2	<0.0001
α_2 (Single Pixel) [μm]	-3.0	-6.4	0.3	0.0804
α_3 (Square:Gradient) [μm]	-3.8	-5.8	-1.7	0.0004
α_3 (Square:Single Pixel) [μm]	-2.5	-4.6	-0.4	0.0189
β_1	12.2	12.0	12.5	<0.0001
β_3 (Gradient)	-3.3	-3.7	-2.9	<0.0001
β_3 (Single Pixel)	-2.2	-2.6	-1.9	<0.0001
γ_1 (Row $_i$)	-1.1	-1.3	-0.9	<0.0001
γ_2 (PrintColumn $_i$)	0.60	0.51	0.68	<0.0001
γ_4 (PrintColumn $_i^2$)	0.04	0.02	0.06	<0.0001
Random effects				
σ_{Phantom} [μm]	3.5	1.9	6.5	
σ [μm]	5.6	5.3	5.9	

lmerTest package [144] in R [145]. The factors of the model were reduced according to minimization of the bayesian information criterion (BIC).

The final reduced model takes the form

$$\begin{aligned}
Y_i = & \mu + \alpha_1(\text{Shape}_i) + \alpha_2(\text{DoseScheme}_i) \\
& + \alpha_3(\text{Shape:DoseScheme}_i) \\
& + (\beta_1 + \beta_3(\text{DoseScheme}_i)) x_{\text{design},i} \\
& + \gamma_1(\text{Row}_i) + \gamma_2(\text{PrintColumn}_i) \\
& + \gamma_4(\text{PrintColumn}_i^2) \\
& + d(\text{Phantom}_i) + \epsilon_i,
\end{aligned} \tag{7.1}$$

where Y_i is the measured printed side length or diameter, μ is the overall intercept, $\alpha_1(\text{Shape}_i)$ is an intercept addition due to the Shape factor, $\alpha_2(\text{DoseScheme}_i)$ is an intercept addition due to the DoseScheme factor, $\alpha_3(\text{Shape:DoseScheme}_i)$ is an intercept addition due to the interaction between the Shape and the DoseScheme factor, β_1 is the overall slope of the model for correlation with the design number of voxels, $\beta_3(\text{DoseScheme}_i)$ is a correction to the slope depending on the DoseScheme factor, $\gamma_1(\text{Row}_i)$ is a slope addition due to the Row factor, $\gamma_2(\text{PrintColumn}_i)$ is a slope addition due to the PrintColumn factor, $\gamma_4(\text{PrintColumn}_i^2)$ is a quadratic addition due to the PrintColumn factor, $d(\text{Phantom}_i) \sim N(0, \sigma_{\text{Phantom}}^2)$ is a random offset from phantom to phantom, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All $d(\text{Phantom}_i)$'s and ϵ_i 's are independent. No significant effect was found of slope dependence on the shape (β_2), interaction between shape and dose scheme (β_4), quadratic effects on the row position (γ_3), or interaction between the row and column position (γ_5). The model coefficient estimates of the reduced model along with their confidence interval and p -value can be seen in Table 7.2.

The model coefficients are discussed in the following, separated into two sections depending on whether the coefficient relates to the general scatterer size correlation or the printer uniformity.

Scatterer predictors

The negative value of μ is an indication of the feature broadening discussed in Section 7.2.1, showing that the actual printed size goes to zero even before the designed size does. This was also directly evident in the microscope images of the smallest scatterers, in particular for the Gradient and Single pixel dosing schemes, for which it was in some cases not possible to measure the scatterer size for the three voxel and four voxel wide scatterers due to apparent closure of the printed structure. α_1 indicates that square scatterers on average become 7.6 μm larger than the circular scatterers. For the Gradient and Single pixel dosing schemes, the difference is only about half of that though, due to smaller negative corrections given by the α_3 values. The α_2 values indicate that the Gradient scatterers are 6.9 μm larger, and Single pixel scatterers are 3 μm smaller. However, it should be noted that these offsets refer to the intercept at a voxel count of 0, and is in part countered by the differences in the slope corrections. β_1 indicates that for each additional voxel of 10.8 μm in the design, the actual printed scatterer increases by 12.2 μm . β_3 indicates that the slopes are reduced by 3.3 μm and 2.2 μm respectively for the Gradient and the Single Pixel dose schemes compared to the Base dose. It should be noted that the β coefficients directly show that the model is not generally valid, but only valid in the investigated region, potentially able to be extrapolated slightly. For small scatterers it makes sense that the correlation is not 1:1. However, it would be expected that the scatterers eventually become large enough that β_1 should correspond to the pixel pitch plus the hydrogel swelling, i.e. $\approx 10.8 \times 1.03 = 11.1 \mu\text{m}/\text{voxel}$.

The scatterer coefficients presented have been modelled together with the printer related predictors, and have therefore been compensated for these.

Printer related predictors

γ_1 indicates there is a -1.1 μm linear difference from row to row, corresponding to an 8.8 μm difference from one side of the phantom to the other. Similarly, γ_2 indicates there is a 0.60 linear difference from column to column, corresponding to a 10.2 μm difference from one side of the phantom to the other of the phantom. Finally, on top of that γ_4 indicates there is a 0.04 quadratic difference between the columns, corresponding to a 2.56 μm difference from the centre to the edge of the phantom. Based on the α_3 values, showing that a larger dose will make scatterers smaller, a smaller dose might increase the scatterer size. Thus, this outwards increasing scatterer size might indicate that the dose illumination system dose compensation map is not perfect, but still decaying outwards. σ_{Phantom} shows that the random variation in resulting sizes from phantom to phantom has a standard deviation of 3.5 μm . The residual error of the model was 5.6 μm .

Model summary

The model explains the average trend of printed scatterer sizes after compensation for the printer inhomogeneities. The analysis revealed that scatterer shape, dose scheme as well as interaction between the two had an effect on the resulting scatterer size. The slope of the correlation between the designed scatterer size and the actual resulting scatterer size also depends on the dose scheme. Furthermore, variation based on the position within the printer FOV was documented, with systematic changes based on the row position, the column position, a quadratic dependence on the column position having been determined.

The Base dose scheme in general provides larger scatterers than Single pixel or Gradient scatterers. This is a consequence of less feature widening. For the three square scatterers in the optical images in Figure 7.10 which were all designed to be 12 voxels, or 129.6 μm , the model shows that the average similarly sized Base dose scatterer will be 137.8 μm , a Single Pixel dose scatterer will be 105.4 μm , and a Gradient dose scatterer will be 101.2 μm .

The residual error σ was 5.6 μm , slightly more than half of a voxel. This might be a combination of actual print variability and operator error when determining the side lengths. The scatterer size varied from phantom to phantom with a random variation of 3.5 μm , about a third of a voxel.

It is worth noting that the model curves in Figure 7.11 seems to roughly meet at the same value for low voxel counts. Due to the discrete square voxel grid, the actual shape of the circular

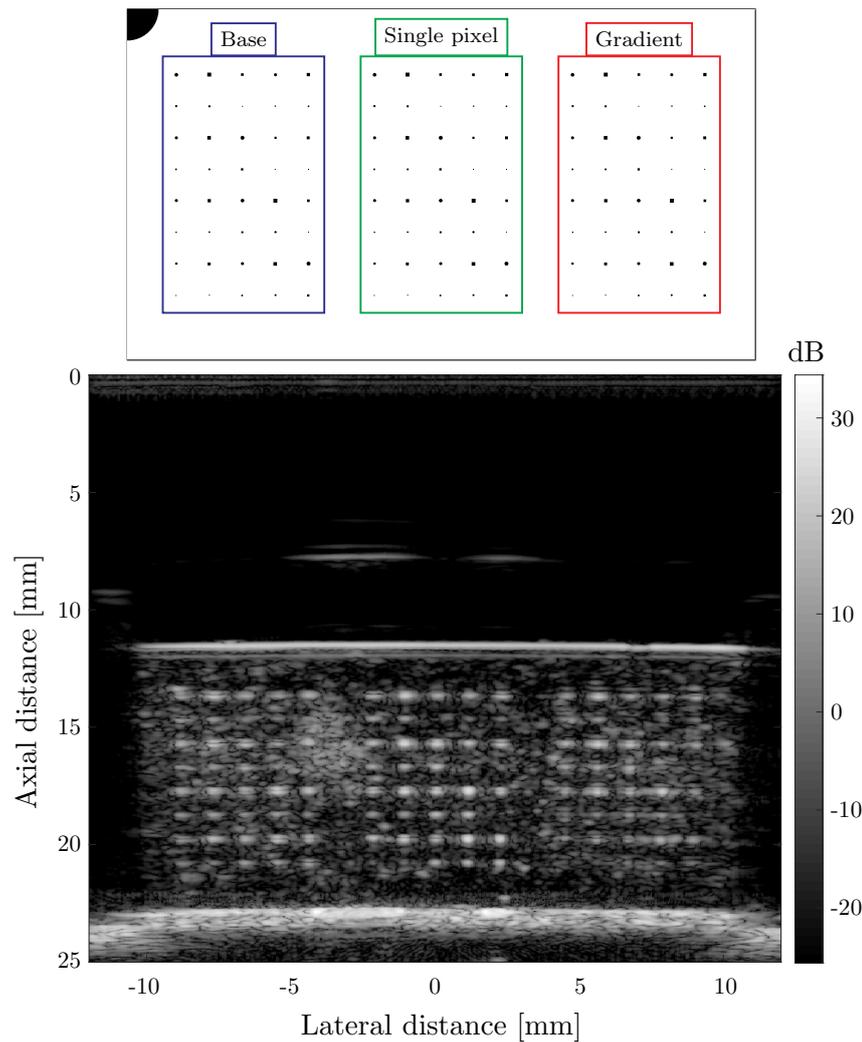


Figure 7.14: B-mode image of one of the six hydrogel phantoms containing scatterers. The image is averaged across 50 frames. The dosing scheme is “Base” in the left region, “Single pixel” in the middle, and “Gradient” to the right. The dynamic range is 60 dB. The intensities have been normalised to the background scattering signal in the phantom.

scatterers are of course only approximations, and for low voxel counts they become increasingly more similar to the square scatterer shapes, with the three voxel “circular” scatterer actually being a 3 by 3 square scatterer. Thus, it makes sense that the starting point for low voxel counts is the same.

7.2.4 Dose manipulation for increased scattering intensity

The six scatterer phantoms were imaged using the BK Medical “Hockey Stick” X18L5s probe with a BK 5000 experimental scanner. A B-mode image of one of the six hydrogel phantoms can be seen in Figure 7.14. A sketch of the layout is placed above the B-mode image, in which the three dose schemes in the used phantom are shown. The dose scheme positions have been perturbed in the other five phantoms. The dynamic range of the image has been set to 60 dB. Each phantom was also imaged rotated by 180° for a total of 12 B-mode images, all of which were averaged over 50 frames.

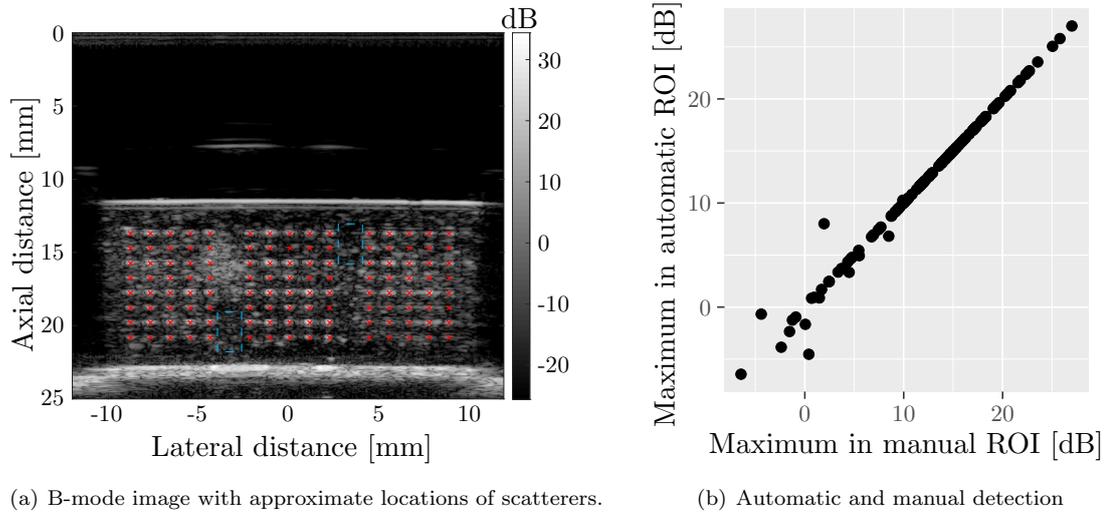


Figure 7.15: Automatic peak detection validation. The scatterer intensities were determined through automatic peak value detection, by feeding the expected centre coordinates of the scatterers, marked by red crosses in (a) the the peak-finding function. The blue dashed rectangles outline scatterer free regions which were used to determine the background intensity level for normalisation of the data. (b) shows the correlation between the automatic detection, and manual detected values for each scatterer. Only few outliers are present, all for low intensities.

The intensity of each scatterer was found using a ultrasound peak-finding function. A region of interest (ROI) is provided as input to the function, in which it searches for the highest scattering intensity. It then extracts the maximum intensity, centre-coordinates, FWHM along the two axes, along with a number of additional parameters. The procedure was automated by feeding the peak-finding function the coordinates of the scatterers, and the size of the ROI. The centre coordinates can be seen marked in Figure 7.15(a). To check the procedure, all scatterers on a single phantom was measured manually, and the correlation between the manual and automatic detections were tested. This is seen in Figure 7.15(b) where the automatic detected values are plotted against the manual detected values. Ideally, the same exact values would be found, and all points would fall on a straight line. Only a few outliers are present, all of them being for low intensities, which is not likely to influence the analysis significantly due to the total number of observations.

The data has been normalised to the background intensity in the phantoms. The background intensity was estimated by searching for peaks in the two regions marked by blue dashed rectangles in Figure 7.15(a). This was done in all 12 B-mode images, with the average value used for normalisation. In the presented B-mode image, the two regions avoid the high intensity background on the left. However, the same regions have been used for all images, so the high intensity region is also included when analysing the background level of the same phantom when flipped by 180°. Thus, in all images of the scatterer phantom and the following analysis, an intensity value of 0 dB corresponds to the average background peak intensity. The analysis of the scattering intensity will thus describe the scattering intensity above the phantom noise level.

Figure 7.16 shows an overview of all detected scatterer intensities across all six phantoms in the two configurations, plotted against the designed size in terms of voxels. The discrete grouping of the data along the x -axis is again due to the designed size not being completely free, but limited to an integer number of voxels. The colours group the data into the different dose schemes, and the shapes group the data into the two different cross-sectional shapes. The solid and dashed lines are the statistical models of the square and the circular cross sectional scatterers respectively, for the different dosing schemes. It appears that the highest intensity scatterers are predominantly square, and the smallest intensities are predominantly circular for each scatterer size.

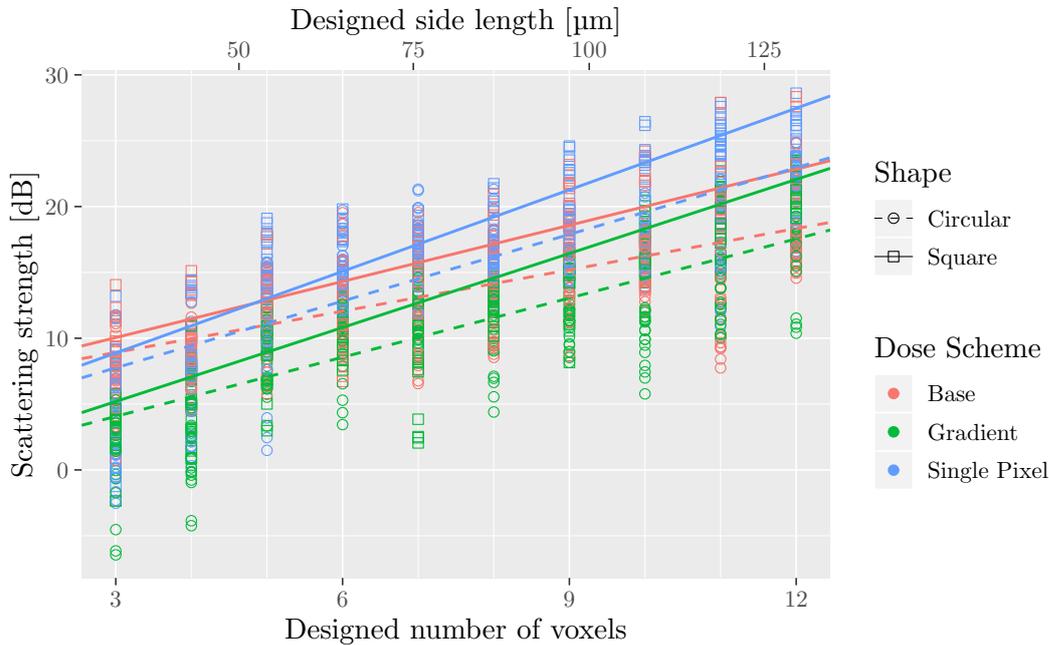


Figure 7.16: Measured scatterer intensity against the designed scatterer size. The colours group the data into the different dose schemes, and the symbols group the data into the two different cross sectional shapes. The solid and dashed lines are linear fits to the square and the circular cross sectional scatterers respectively.

Figure 7.17(a) shows a box-plot of the measured scatterer intensities against the two different scatterer shapes, while compounding all other effects. The median offset between the two shapes indicates that the square scatterers generally reflect more sound than the circular scatterer.

Figure 7.17(b) shows a box-plot of the measured scatterer intensities against the three different scatterer dose schemes. Again, the content of each box is compounded across other factors, such as the scatterer size, and scatterer shapes. The median offset indicates slight differences, with Single pixel scatterers providing the most intensity, Base scatterers a little bit less intensity, and Gradient scatterers the least.

Figure 7.18 shows box-plots of the measured side length or diameter against the designed scatterer size, separated into the different dose schemes and shapes. All plots are scaled equally to allow for easy comparison. However, while the intensity distributions become more clear, the overall tendencies are difficult to isolate. It is noticeable that the Gradient and Single pixel box-plots seem to bend around 6 voxels, with different slopes on both sides. The quickly decaying intensities for low voxel counts might be due to feature broadening and closing of the scatterers due to the additional doses of these dose schemes, as was observed in the optical microscope images. That will explain why the same trend is not seen for the Base dose scheme.

Scatterer intensity statistical analysis

In addition to the factors which were tested in the scatterer size experiment, imaging of the same phantom rotated by 180° allows for testing the ultrasound field homogeneity as well. The summary of the data types and the factors are similar to those of the scatterer size analysis, and are included in Table 7.1. For this experiment, the Imaging column and the random interaction between the phantom and the mounting orientation in the bottom of the table are included as well. Interactions between factors, and quadratic effect of position in the scatterer grid relative to the printer or the ultrasound probes were also investigated.

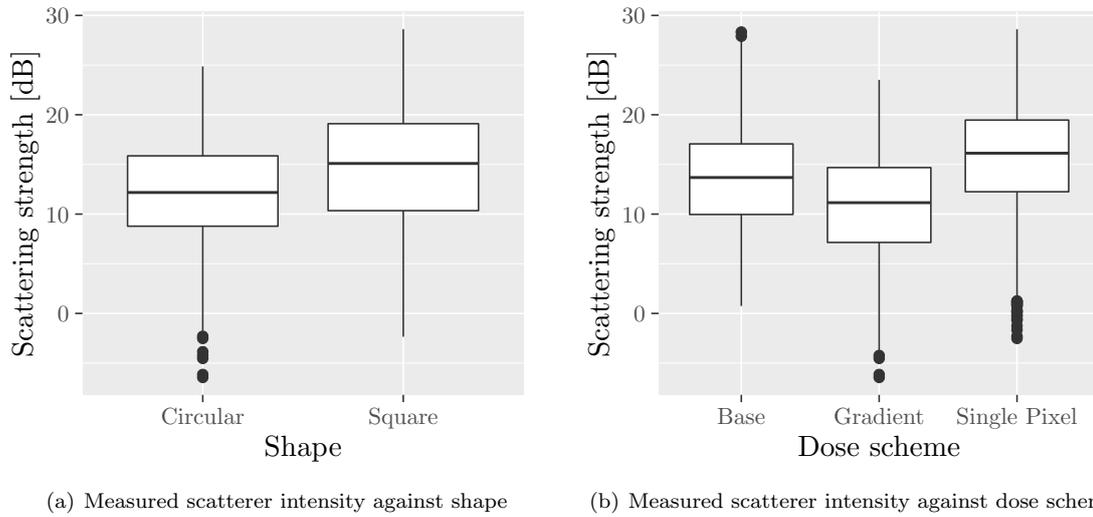


Figure 7.17: Measured scatterer intensity (a) against the scatterer shape and (b) against the dose scheme.

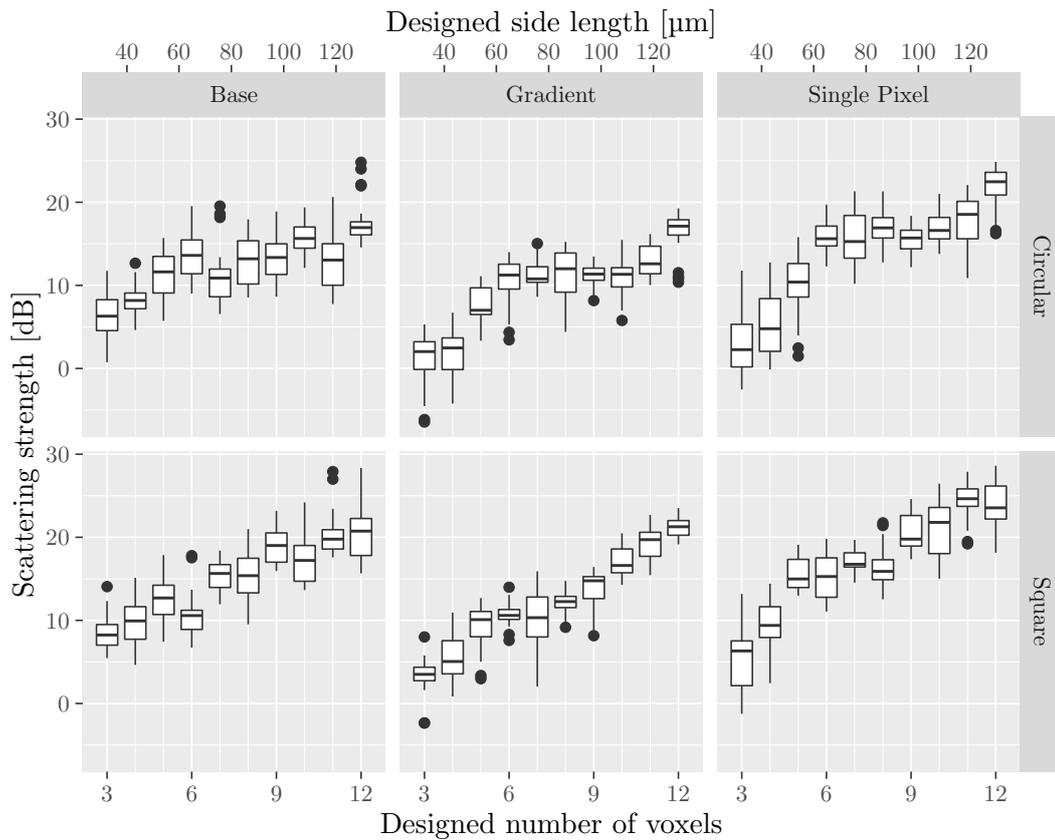


Figure 7.18: Measured scattering intensity against the designed scatterer size. The data is separated into the dose schemes and shapes for a better overview of the data distributions. Dots mark outliers.

The full model of all the factors and interactions investigated can be seen in Appendix E.2, along with model diagnostics and more summarizing plots. The combination of fixed and random factors makes the fitted model a linear mixed effects model. Such a model can be analysed using the `lmerTest` package [144] in R [145]. The factors of the model were reduced according to minimization of the BIC.

The final reduced model takes the form

$$\begin{aligned}
Y_i = & \mu + \alpha_2(\text{DoseScheme}_i) \\
& + [\beta_1 + \beta_2(\text{Shape}_i) \\
& + \beta_3(\text{DoseScheme}_i)] x_{\text{design},i} \\
& + \gamma_1(\text{Row}_i) + \gamma_2(\text{PrintColumn}_i) \\
& + \gamma_3(\text{Row}_i^2) + \gamma_4(\text{Column}_i^2) \\
& + d(\text{Phantom}_i) + \epsilon_i,
\end{aligned} \tag{7.2}$$

where Y_i is the measured printed side length or diameter, μ is the overall intercept, $\alpha_2(\text{DoseScheme}_i)$ is an intercept addition due to the DoseScheme factor, β_1 is the overall slope of the model for correlation with the design number of voxels, $\beta_2(\text{Shape}_i)$ is a correction to the slope depending on the DoseScheme factor, $\beta_3(\text{DoseScheme}_i)$ is a correction to the slope depending on the DoseScheme factor, $\gamma_1(\text{Row}_i)$ is a slope addition due to the Row factor, $\gamma_2(\text{PrintColumn}_i)$ is a slope addition due to the PrintColumn factor, $\gamma_3(\text{Row}_i^2)$ is a quadratic addition due to the Row factor, $\gamma_4(\text{Column}_i^2)$ is a quadratic addition due to the Column factor, $d(\text{Phantom}_i) \sim N(0, \sigma_{\text{Phantom}}^2)$ is a random offset from phantom to phantom, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All $d(\text{Phantom}_i)$'s and ϵ_i 's are independent. The model was plotted in Figure 7.16 as the solid and dashed lines.

No significant random effect was found between the B-mode images, showing that the data range in the images were similar. There was also no intercept dependence on the shape (α_1), or the interaction between shape and dose scheme (α_3). The scatterer column position within the ultrasound field was not significant either, showing good uniformity across the probe FOV. There was no effect of slope dependence on the interaction between the Shape and the dose scheme (β_4), nor any interaction between the print column and the row (γ_5). The model coefficient estimates of the reduced model along with their confidence interval and p -value can be seen in Table 7.3.

Scatterer predictors

The small positive value of μ is an indication that even the smallest scatterers are more intense than the average background, with an overall average 5.79 dB above. The μ intercept corresponds to the overall intercept, of the Base dose scatterers. The alpha coefficients are in general difficult to conclude on, since most effects are offset by the slopes of the individual curves. α_2 indicates that the intercepts are 6.21 dB and -3.10 dB lower for the Gradient and Single Pixel dose schemes respectively, compared to the base dose. For Single pixel, this is not the tendency observable in the box-plot of the intensity split into the three dose schemes in Figure 7.17(b). However, observing Figure 7.18 instead, this is clearly due to the sharp drop-off in intensity for scatterers smaller than 6 voxels. This is likely in turn correlated to the Gradient and Single pixel scatterers being smaller as a consequence of feature widening due to the large additional doses. β_1 shows that the reflected intensity increases by 1.05 dB for each additional voxel as the voxel size is increased. This slope corresponds to a Base dose circular scatterer. β_2 increases the intensity by 0.38 dB per additional voxel. Similarly, β_3 shows that if the dose scheme is changed to either Gradient or Single pixel, the intensity increases by 0.38 dB and 0.64 dB respectively. For larger scatterers, this undoes the negative intercept additions from the α_2 coefficients, rendering a complete overview of the correlations quite convoluted.

The scatterer coefficients presented have been modelled together with the printer and ultrasound related predictors, and have therefore been compensated for these.

Table 7.3: Model coefficients for the reduced model of the scatterer intensity, with confidence intervals and p -values.

Predictors	Estimate	2.5%	97.5%	p -value
Fixed effects				
μ [dB]	5.79	4.83	6.74	<0.0001
α_2 (Gradient) [dB]	-6.21	-7.19	-5.23	<0.0001
α_2 (Single Pixel) [dB]	-3.10	-4.08	-2.12	<0.0001
β_1	1.05	0.96	1.13	<0.0001
β_2 (Square)	0.38	0.34	0.41	<0.0001
β_3 (Gradient)	0.45	0.33	0.57	<0.0001
β_3 (Single Pixel)	0.64	0.52	0.77	<0.0001
γ_1 (Row $_i$)	-0.35	-0.41	-0.28	<0.0001
γ_2 (PrintColumn $_i$)	0.09	0.07	0.12	<0.0001
γ_3 (Row $_i^2$)	-0.16	-0.19	-0.13	<0.0001
γ_4 (Column $_i^2$)	-0.02	-0.03	-0.01	<0.0001
Random effects				
σ_{Phantom} [dB]	0.76	0.41	1.43	
σ [dB]	2.78	2.67	2.87	

Printer and scanner related predictors

γ_1 indicates that the intensity increases by 0.35 dB per row down into the phantom. The sign of the coefficient is the same as that in the scatterer analysis, meaning it might simply be a matter of the scatterers being printed smaller, therefore reflecting less sound. However, the total size change from top to bottom should be less than a voxel according to the scatterer size model, which would therefore result in approximately 1 dB according to β_1 . However, a coefficient of 0.35 dB should result in a difference in almost 2.5 dB across 8 rows. The decreasing tendency is also opposite to regular attenuation of signal with depth, and might also indicate a too high TGC in this experiment. A likely explanation would be a combination of the two. γ_2 shows that there is a difference of 0.09 dB per column correlated with printed orientation, i.e. the slope changes when the phantom is flipped 180°. The change from one side of the phantom to the other is 1.44 dB. Referring to the scatterer size experiment, the scatterer size changed by approximately 1 voxel from one side to the other, thus the change in intensity is in good agreement with the change in scatterer size. γ_3 indicates a negative quadratic effect in depth. No similar effect was found for the scatterer sizes, and there is no obvious physical argument as for why that is. γ_4 indicates a negative quadratic effect of the scattering intensity with column position of -0.02, corresponding to -1.28 dB from centre to the edge of the phantom. The tendency is opposite to that of the scatterer size analysis, and would therefore not be expected to be a consequence of the scatterer size. However, the quadratic column factor represented both the printer column and the ultrasound column, since these would be indistinguishable. Therefore, the quadratic effect of the column position likely an effect of ultrasound energy loss the further out laterally in the FOV the scatterer is placed.

Model summary

The model explains the average trend of scatterer intensity after compensation for the printer inhomogeneities. The analysis showed that the dose scheme had a direct effect on the intensity. The correlation slope between the designed scatterer size and the intensity also depends on both the shape of the scatterer and the dose scheme. Variation based on the scatterer position within the printer FOV was also documented, with systematic changes based on the row, printed column, and quadratic effect of the column position.

Overall, the Single pixel dose scheme combined with the square shape provides the largest reflected intensity.

The residual error σ was 2.78 dB. The intensity varied systematically from phantom to phantom with a random variation of 0.76 dB.

7.2.5 Scatterer separation distance

New machine learning and neural network schemes are being developed for SRUS to be able to detect individual micro-bubbles with partially overlapping PSFs. This is the topic worked toward in Paper D and Paper J. For those techniques, the scatterer size will be even more important. To demonstrate the neural network approach, a phantom containing a scatterer array of 10 by 10 scatterers placed with a lateral separation of 518 μm and an axial separation of 342 μm was created. Each scatterer was designed to be 7 by 7 voxels in cross-section, exposed with the Single pixel dose scheme. Based on the scatterer size statistical model, this results in 55.3 μm by 55.3 μm printed scatterer. The design can be seen in Appendix I.1. The neural network detected the two reflections of each scatterer, needing additional training to combine them to a single scatterer localisation.

In addition to decreasing the scatterer size for these techniques, it will be important to be able to place scatterers close to each other. Once the distance becomes too short, the separating hydrogel pattern might become fragile and break. To test the minimum separation between scatterers, the phantom design seen in Figure 7.19 was developed. Figure 7.19(a) shows the phantom design, Figure 7.19(b) shows the base exposure cross-sectional pattern and Figure 7.19(c) shows the Single pixel overexposure cross-sectional pattern. The scatterers were placed at the top surface of the phantom to allow for simple optical characterisation. The scatterers were designed to be 7 by 7 voxels, placed in small arrays of 4 by 4 scatterers. The separation within each group was varied from 1 voxel to 10 voxels.

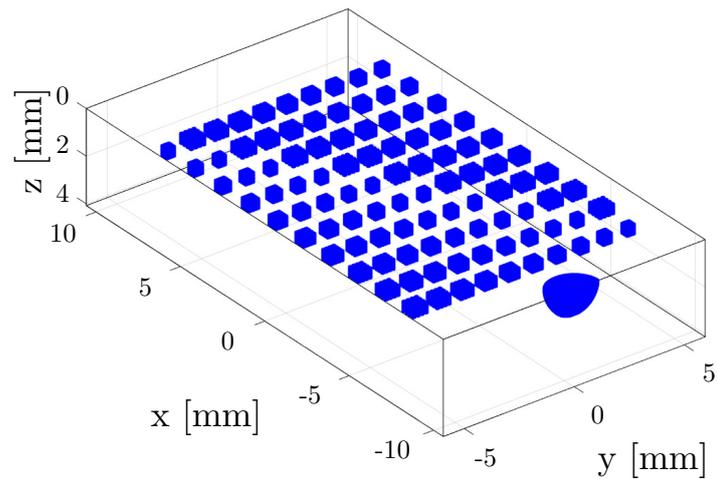
The Zeiss Axioskop 40 optical microscope equipped with a 5x magnifying lens was used for optical characterisation. Figure 7.20 shows two groups of scatterers, one designed to be separated by 10 voxels, and one designed to be separated by 1 voxel. The scatterer groups were printed without failure for all separation distances. Thus, phantoms can be created with scatterers only separated by a single voxel.

It should be noticed that although the group of scatterers in Figure 7.20(b) were designed to be separated by only a single voxel, the actual printed separation is approximately 3 voxels wide, due to the feature widening from the Single pixel overexposure. The shapes of the scatterers in the different groups were not consistent. Another group of scatterers separated by 1 voxel can be seen in Appendix J.3.3.

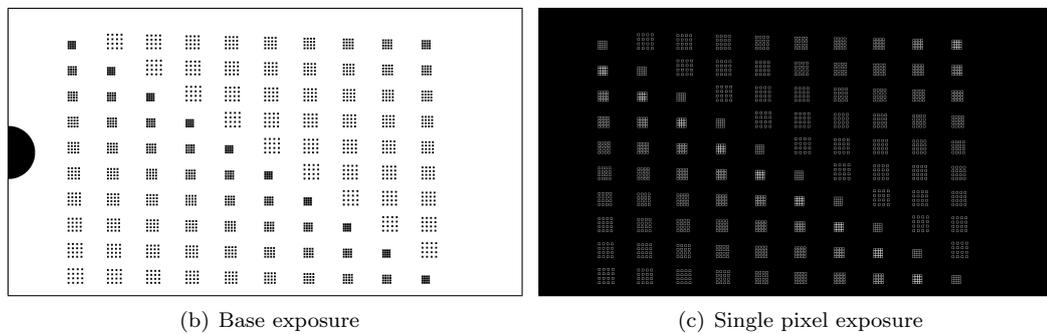
In conventional B-mode imaging with sub-wavelength sized scatterers, the detected scatterer position will be in the centre of the scatterer, between the reflections at the front of the scatterer and at the back. For two 7 voxels wide scatterers separated by a single voxel, the centre separation will be 8 voxels, or 86.4 μm . The only way to place scatterers closer than that will be to decrease the scatterer size further.

7.3 Cavity scatterer micro-phantoms for validation of SRUS in 3D

Most publications on SRUS have been based on 2D imaging. The reason for that is primarily lack of availability of 3D imaging probes and scanning equipment capable of handling the increased amount of data. But the issue with using 2D imaging equipment is, that the vessel structures to be imaged in the end are inherently three-dimensional. 2D SRUS is just that: SRUS along two directions, with the caveat that the received signals have been summed across the elevation plane, which at its focus is $2 - 5\lambda$, and therefore by no means super-resolved in the elevation direction. As mentioned previously, the only way to fix that problem is to do 3D imaging, with which it would be possible to focus along the elevation plane, thereby enabling SRUS in 3D. For this experiment a 3D RCA probe was used.



(a) Phantom layout



(b) Base exposure

(c) Single pixel exposure

Figure 7.19: Cross-sections of the phantom design for testing the minimum separation between scatterers. The scatterers were designed to be 7 by 7 voxels, placed in small arrays of 4 by 4 scatterers. The separation within each group was varied from 1 voxel to 10 voxels. (a) shows the phantom design, (b) shows the base exposure cross-sectional pattern, and (c) shows the Single pixel overexposure cross-sectional pattern. White pixels are illuminated.

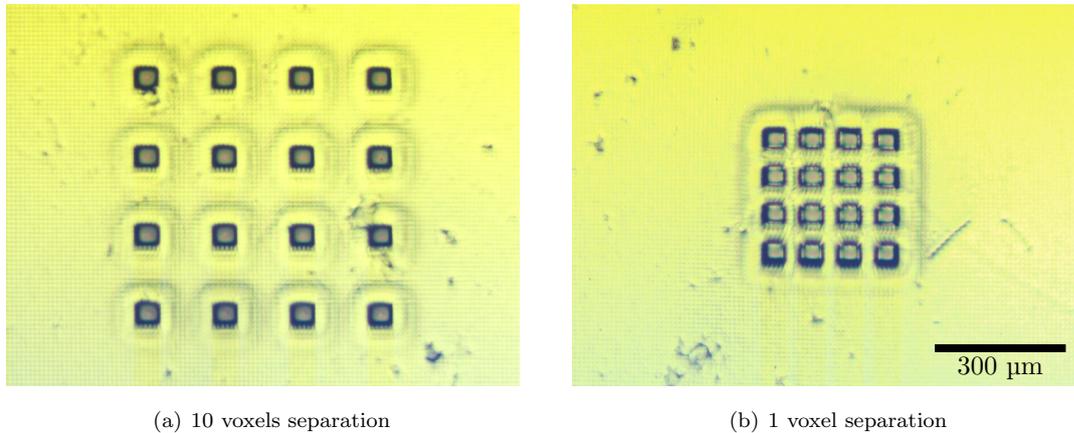


Figure 7.20: Images of scatterer separation phantom groups. (a) shows a group of scatterers separated by 10 voxels. (b) shows a group of scatterers separated by 1 voxel. Each group of scatterers were printed systematically without failure. Images are scaled equally.

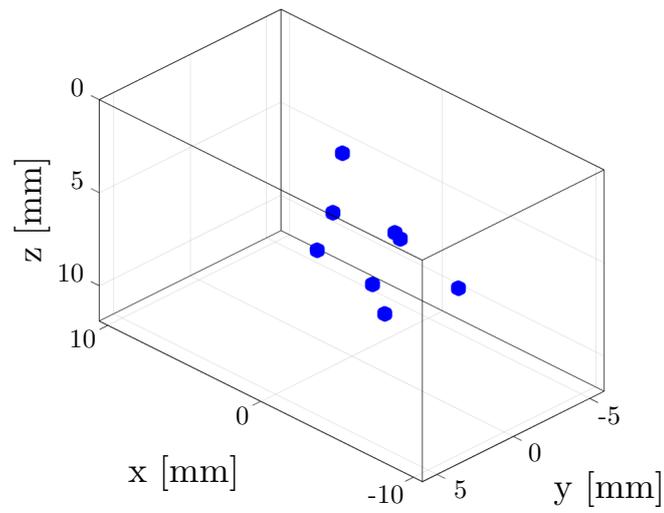
7.3.1 Methods

Experimental procedure

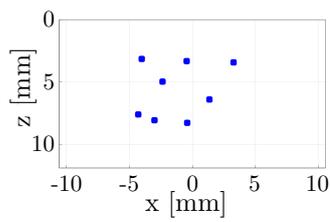
The foundation for the experiment was the eight scatterer phantom presented in Chapter 6 Section 6.2 in Figure 6.2, where it was used to create collapsed scatterer versions for optical characterisation. The actual phantom used for ultrasound experimentation is shown in Figure 7.21. The outer dimensions of the phantom were $21.1 \times 11.9 \times 11.9 \text{ mm}^3$, after correction for the expansion for a three second base dose, with each scatterer having been designed to be $205 \times 200 \times 205 \text{ }\mu\text{m}^3$. While the printing setup allows for printing significantly smaller scatterers as demonstrated in the previous sections, it was necessary with an increased size to obtain reflections with intensities larger than background scattering due to unavoidable small random print artefacts in the phantom. The scatterers will function as point targets in regular B-mode volumes, when the imaging wavelength is larger than the scatterer size, in this case for any frequency below 6 MHz. They were placed with a minimum separation distance of 3 mm, which will eliminate overlapping signals for any frequency above 0.5 MHz. Since this experiment was conducted prior to the investigation of the effect of the dosing schemes, the scatterers were printed using the Gradient scheme, due to the belief that this would be better. Extrapolating the model of the scatterer size presented in the previous section, the average actual size of a 19 voxel wide square scatterer printed with the Gradient scheme would be expected to be $163 \times 200 \times 163 \text{ }\mu\text{m}^3$, with the second dimension being unaffected, since the gradient is not applied in this direction.

With the print swelling factor determined, the true distances between the scatterers in the 3D version of the scatterer phantom will be known, and can be used to compare against those found by ultrasound. The phantom was translated relative to the ultrasound probe using the translation stage along a single axis; in the first experiment along the x -axis, and in the second experiment along the y -axis. The inter-volume stage movement in both experiments was $12.5 \text{ }\mu\text{m}$, corresponding to a 2 mm/s velocity acquired at a volume rate of 160 Hz. This speed corresponds to common flow velocities in small vessels. By moving the phantom in between volume acquisitions, any differences depending on the phantom placement within the field of view of the transducer will be included in the analysis, instead of simply testing the SRUS pipeline parameters locally within the transducer field of view.

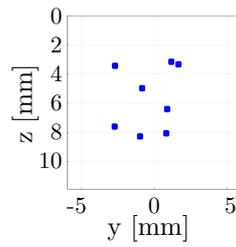
The experimental setup shown in Figure 7.9 was also used for this experiment. The imaging probe was a prototype 62 + 62 elements 3 MHz PZT, RCA array [61]. The probe was connected to the experimental synthetic aperture real-time ultrasound system (SARUS) [153], which is capable



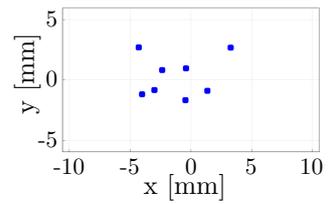
(a) Phantom layout



(b) x - z -plane



(c) y - z -plane



(d) x - y -plane

Figure 7.21: Scatterer phantom for ultrasound experimentation.

of storing channel data for offline processing. A single frame is a summation of 32 defocused emissions using a synthetic aperture (SA) imaging approach [154]. Rows were transmitting and columns were receiving, thereby resulting in 62 channels in receive per emission. The phantom was stationary while a frame was being measured to avoid intra-frame motion artefacts. In total 2×640 volumetric frames were acquired over the 2×640 positions. The volumetric frames were then passed to the SRUS pipeline.

Super-resolution pipeline

The SRUS pipeline which was used has been described in detail in [154]. It is briefly summarised in the following. The super resolution pipeline consists of three steps. The first is SA beamforming. Each imaged volume spans a volume of $14.86 \times 14.86 \times 7.43 \text{ mm}^3$, corresponding to $61 \times 61 \times 243$ voxels. Each high resolution volume was a summation of 32 volumes beamformed from 32 emissions, using a specialised beamformer [155] implemented on a GPU [156]. The volume was dynamically focused in receive (F-number of 1.5) and synthetically in transmit (F-number of 1), with an optimized sequence for SA B-mode. This was done for all 2×640 frames. In the next step, a stationary echo filter was applied to remove stationary tissue. In a micro-bubble experiment, this would remove the signal stemming from the tissue as it is stationary, leaving only the micro-bubble signal. However, since the entire phantom was translated between each frame in this experiment, the stationary echo filter would have no effect on the results. The final step is to determine the points scatterer positions based on local maxima. Sub-pixel positioning is obtained by interpolating the peak location using a second order polynomial in all three dimensions. The 3D coordinates $\{x_p, y_p, z_p\}$ of the detected points is then provided as the output from the third stage. Tracks of the individual scatterers can then be formed by collecting spatially similar coordinates across all imaged frames. The pipeline was implemented in MATLAB, and was processed offline [154].

7.3.2 Results

Scatterer localisation

Figure 7.22 shows three selected cross planes of a B-mode volume. The coloured dots mark the localised positions of the scatterers detected in one of the 640 volumes. The example cross planes have been chosen such that they all contain the scatterer marked by a blue dot. The $x - z$ cross plane, Figure 7.22(c), also contains an additional scatterer, marked in red. The selected volume contains a total of five scatterers, with the remaining scatterers not visible within the selected cross-planes. The large reflection at $x \approx 3.5 \text{ mm}$ and $z \approx 4 \text{ mm}$ does not correlate with any of the designed scatterer positions, and likely stems from a print artefact.

The localised positions of the 3D printed scatterers, accumulated over the 640 volumes, can be seen in Figure 7.23. The colours group the tracked points of the individual scatterers, while the black tracks illustrate the expected tracks based on the design coordinates. The latter are included for visual confirmation that the localisations are indeed the designed scatterers. It is recommended to always include such a comparison to confirm that the localisations indeed correspond to the features of the designed phantom. Drop lines are included to aid the 3D perception. The horizontal field of view in the figures have been limited to the measured data tracks, removing parts of the black tracks. The actual cross-sectional field of view of the probe is $14.86 \times 14.86 \text{ mm}^2$.

Although eight scatterers were printed, not all were found in the two experiments: seven scatterers were correctly localised for the movement along the x -axis (Figure 7.23 a)) and five scatterers were correctly localised for the movement along the y -axis (Figure 7.23 b)). In addition, the track length varies from 81 localisations to 633 localisations, across the 640 volumes. Two additional tracks, which did not align with the design coordinates, have been omitted from the images and the analyses. It is expected that these tracks stem from print artefacts, resulting in unintended cavities in the phantom, which therefore reflect the ultrasound similarly as the designed scatterers. They aligned well with the reflection seen in Figure 7.22(c) at $x \approx 3.5 \text{ mm}$

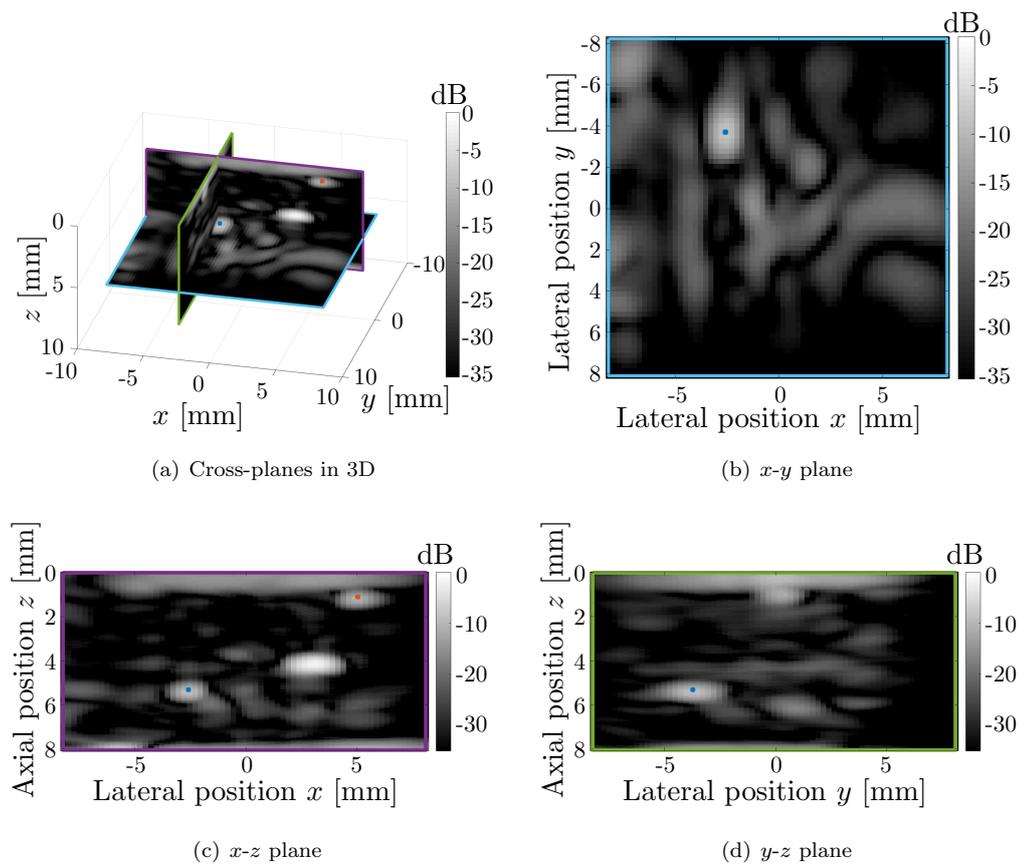


Figure 7.22: (a) B-mode volume containing scatterers. Three cross planes of the B-mode volume are shown, (b) $x-y$, (c) $x-z$, and (d) $y-z$. The super-localised positions of the scatterers are marked by coloured dots.

Table 7.4: Summary of the variables and their data types used in the ultrasound correlation analysis.

Predictors	Sample values	Variable type	Description
Ultrasound distance [mm]	8.717, 3.730,..., 6.279	Numerical values	Distance between points calculated through SRUS pipeline
Compensated design distance [mm]	8.719, 3.811,..., 6.384	Numerical values	Compensated designed distance between points
Motion	X, Y	Fixed factor	The axis of translation

and $z \approx 4$. While these print artefacts would also be fixed in position, and be moved along the same trajectory as the designed scatterers, the print artefact geometry is not known. If a print artefact is significantly larger than the imaging wavelength, localisation of the centroid might be ambiguous, and therefore, these tracks were omitted from the analysis.

Super-resolution accuracy

The SRUS pipeline accuracy was investigated in a similar manner to the optical validation, by comparing the known distances between the designed points to the measured distances between points from the ultrasound experiments. There are two main differences to the optical experiment: The scatterers are now positioned not in collapsed planes but in 3D, visualised as the blue points in Figure 6.2, and the design distances are compensated for the expansions according to the results in Table 6.2 before analysing the correlation between the designed distances and those calculated from the ultrasound data. After the compensation, the correlation should be a straight line with a slope of 1, in the case of perfect correlation. Since there are two sets of experiments, one for each direction of motion of the translation stage, the variables of the analysis are the compensated design distances, the measured ultrasound distances, and a factor separating the data into the x - and y -motion, all summarised in Table 7.4. In this experiment, the entire beamformed volume has been assumed to have a speed of sound equal to that in pure water, 1480 m/s.

As was mentioned in Section 7.3.2 and shown in Figure 7.23, an unequal number of scatterers were localised by the SRUS pipeline in the two experiments, and the tracks were of unequal length. This means there will be more data for the x -direction of motion, resulting in an unbalanced dataset from a statistical point of view. In addition, our analysis of the variation in the data showed that the data was heteroscedastic. Modelling the correlation of the raw distances between points might be heavily biased toward certain parts of the data simply due to the large number of samples. Instead, a weighted least squares analysis of the distance distributions was conducted. This was performed by modelling the mean distance between each point across all measurements, with each mean value being weighted by the variance of the measurements contributing to that mean. The correlation between the compensated design distances and the mean of the distances calculated by the SRUS pipeline is shown in Figure 7.24.

The initial linear model is given as

$$Y_i = \mu + \alpha(\text{Motion}_i) + (\beta_1 + \beta_2(\text{Motion}_i)) x_{\text{design},i} + \epsilon_i, \quad (7.3)$$

where Y_i is the mean of the distance between points calculated from the SRUS pipeline output, μ is the overall intercept, $\alpha(\text{Motion}_i)$ is an intercept addition due to the Motion factor, β_1 is the average slope of the model, $\beta_2(\text{Motion}_i)$ is a Motion dependent correction to the slope, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All ϵ_i 's are independent.

The model reduction was conducted by removing only a single term at a time, based on a 5% level of significance. Neither the overall intercept (μ), nor the direction of motion dependent addition to the intercept ($\alpha(\text{Motion}_i)$), nor the direction of motion dependent correction to the slope

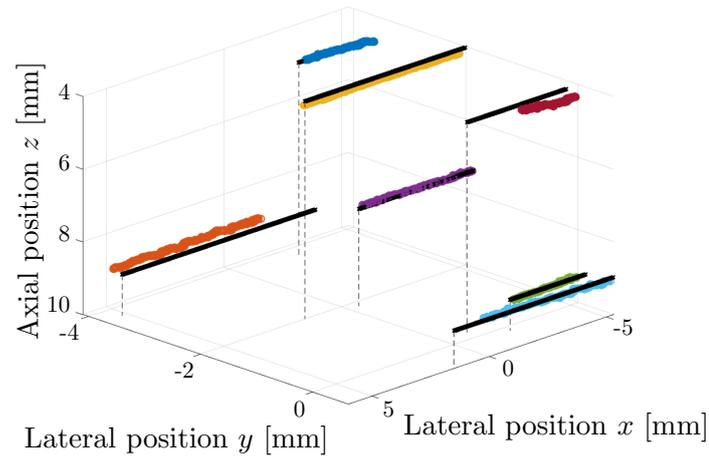
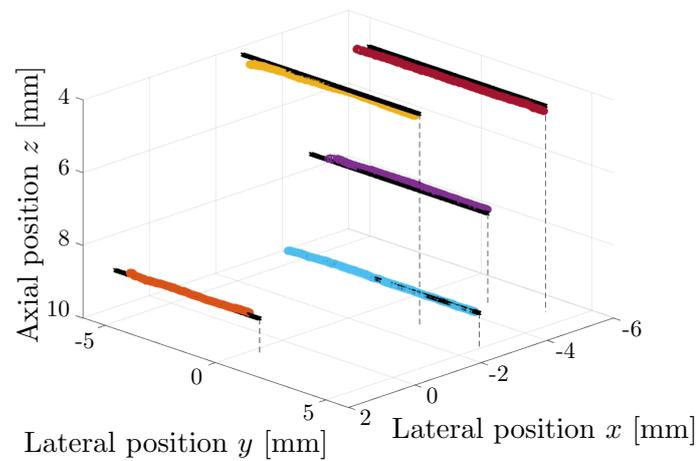
(a) Motion along x (b) Motion along y

Figure 7.23: Cumulated localized scatterers acquired over 640 volumes. The phantom was translated in two separate experiments, along the transducer x -axis (a), and along the transducer y -axis (b). The black tracks illustrate the expected tracks based on the design coordinates. Drop-lines end on the $z=10$ mm plane, and are included to aid the 3D perception.

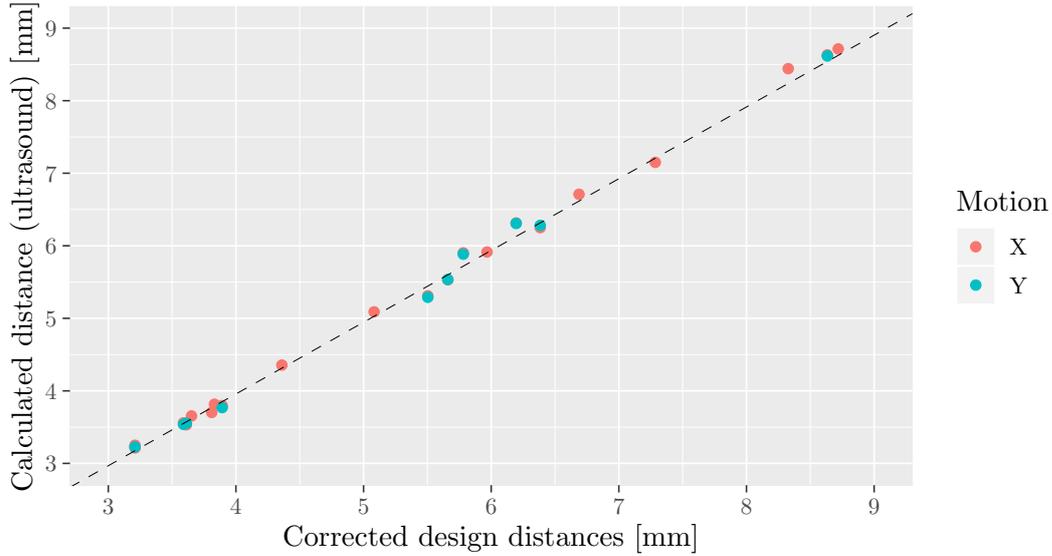


Figure 7.24: Correlation between the compensated design distances and the mean of the distances calculated by the SRUS pipeline. The line represent the final reduced model seen in Eq. (7.4).

Table 7.5: Model parameter estimates of the final reduced model including confidence intervals of correlation between ultrasound distances and compensated design distances.

	Estimate	2.5%	97.5%	p -value
β_1 (slope)	0.989	0.982	0.996	<0.0001

$(\beta_2(\text{Motion}_i))$ were significant at 5%, and were therefore removed. Thereby the model reduction converged at the final model

$$Y_i = \beta_1 \cdot x_{\text{design},i} + \epsilon_i. \quad (7.4)$$

The model coefficient and confidence interval of the reduced model are presented in Table 7.5. The analysis showed no dependence of the direction of motion, nor any intercept of the correlation. The modelled average behaviour of the fitted line has a slope of 0.989, close, yet not equal, to a perfect correlation with a slope of 1. Based on the heteroscedastic assumption of the data, a direct estimate of the residual standard error is not meaningful.

Super-resolution precision

The same ultrasound data was used to estimate the SRUS pipeline precision. The precision was estimated by investigating the variation of the individual localisations relative to the trajectories of the translated scatterers. The tracks with motion along the x -direction were used to estimate the precision in y . The tracks with motion along the y -direction were used to estimate the precision in x . Both datasets were used to estimate the precision in z . To visualise the variation, the mean x -, y - and z -coordinate were subtracted from each individual track, to centre the tracks around the transducer coordinate-system origin. This is illustrated in Figure 7.25, where two cross-planes (x - y and x - z) are shown for the tracks with motion along the x -axis, which corresponds to the tracks in Figure 7.23(a).

The colour of the points represent the tracks of the different design points, and are matched to those of the tracks in Figure 7.23(a). The movement was uni-axial along the translation stage x -axis. However, slight misalignment between the ultrasound transducer and the translation stage have resulted in the localisation tracks not being perfectly aligned to the transducer axes. This

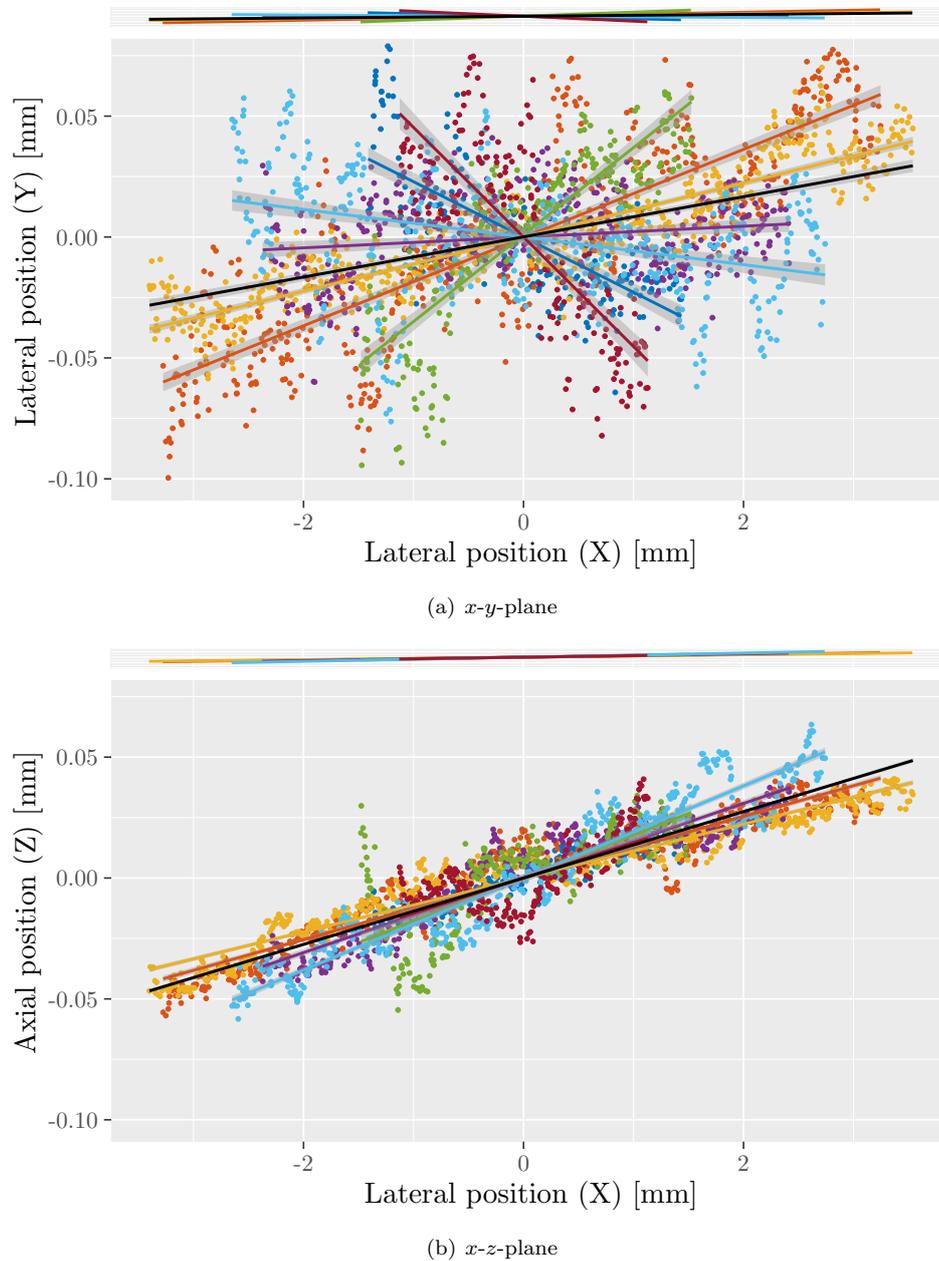


Figure 7.25: Cross-planes of the tracks with motion along the x -axis, offset to be centred around the coordinate system origin. The black lines show the average trajectory, while the coloured lines are linear fits to the individual trajectories of the different scatterers. The scaling is equal across (a) and (b), but the main plots do not have equally scaled axes. The narrow graphs on top in both figures show the same linear fits, with equally scaled axes.

Table 7.6: Estimated precision for the super-resolution algorithm.

	Average trajectory	Individual trajectories
$\tilde{\sigma}_x$ [μm]	17.7	17.3
$\tilde{\sigma}_y$ [μm]	27.6	19.3
$\tilde{\sigma}_z$ [μm]	9.5	8.7

can be observed in Figure 7.25, in which the black line is the average trajectory of all tracks in the dataset. It should be noted however, that the axes are not equally scaled in the main plots, but only in the inserts of the same data shown on top of each plot. The misalignment angle is 0.49° in the x - y plane, and 0.79° in the x - z plane. This misalignment should be compensated for when determining the variation of the tracks. The scatterers are fixed in the phantom and have been moved collectively by the translation stage. Then all tracks should have moved in the same direction, and the average trajectory of the tracks would be a good estimate of that. An estimate of the precision could be determined as the variation relative to the average trajectory. The precision along all three dimensions based on the variability relative to the average trajectory is displayed in Table 7.6 (“Average trajectory”).

However, the coloured lines indicate that the tracks are in fact not parallel, but at small angles to each other. It is fairly small angles relative to the average trajectory, with the largest angle in any plane being 3.1° . This indicates that there is an error somewhere in the SRUS pipeline, and that determining the precision relative to the average trajectory might be misleading. As an alternative, the estimate of the precision could be determined relative to the individual trajectories of the tracks. The precision along all three dimensions based on the variability relative to the individual trajectories is displayed in Table 7.6 (“Individual trajectories”). However, given that the tracks should have been parallel, this latter estimate of the precision might also be misleading. It is expected that the two presented estimates of the precision are limiting cases, and that the true precision of the SRUS pipeline will lie somewhere in between.

7.4 Discussion

The scatterer size analysis provided insight into the actual printed sizes for the different dosing schemes and shapes. The overall trend is that dose schemes involving longer exposure times results in smaller scatterers, perfectly in line with the feature widening concept. The residual error of the model was $5.6 \mu\text{m}$, about half the size of a voxel. This difference might in part be due to measurement uncertainty. The model provides a clearer overview of what to expect when printing features of a certain size. The smallest designed scatterers were not consistently visible for the dose schemes involving increased exposure times. In many cases, the scatterers designed to be three voxels appeared completely closed.

The scatterer intensity analysis showed that on average all scatterer sizes reflect sound larger than the average background intensity. The largest reflections are obtained from the square Single pixel scatterers. The background intensity of the phantoms is so far not controllable, and varies a lot even across a single phantom. It is difficult to provide an exact minimum scatterer size which will provide a sufficient signal guaranteed. This is to a lesser extend due to the residual error of the model of slightly less than 3 dB. However, depending on the intended application, the required intensity will vary. If only a single scatterer is needed, it will need to reflect more strongly to stand out against the background, than a pattern of scatterers would. Methods for decreasing the background noise from the bulk of the phantom should be investigated. Some of the unintended structures observed in ultrasound might originate from issues in the printer system, one of which is illustrated in Appendix J.3.4. Incidentally, this section also shows how precisely the printer vat can be positioned from print to print.

It should be noted that the scatterer intensity analysis was conducted specifically for 2D imag-

ing, utilizing the integration of signal across the elevation focus. A similar analysis conducted for 3D imaging would be equally relevant. In this case the size range of scatterers will be different since it will no longer be possible to integrate the signal across the elevation focus. The 3D scatterer phantom presented was made with 200 μm scatterers, which provided sufficient signal. However, the size had not been optimized, and was not based on the Single pixel dose scheme. The actual limit for 3D scatterer phantoms is still unknown.

The new phantom concept introduced has successfully been used for SRUS pipeline characterisation. The presented results illustrate that it is possible to obtain estimates for precision and accuracy, using these specialised phantoms. The obtained precision is an improvement of at least a factor of 18 compared to the ultrasound wavelength. It is particularly worth noting that although there are some questions regarding how to interpret the estimates of precision, even the worst obtained estimates for precision are comparable to the size of the smallest vessels in tissue. Thereby it is clear that the used method is suitable for resolving features at the size of the smallest vessels in tissue in three dimensions, and the stability of the phantom features allows for documentation of this.

The high positioning control has allowed for the detection of distortion in the SRUS pipeline, through the non-parallel tracks, which would not have been possible using conventional phantoms. The tracks should have been parallel given that the scatterers are fixated in the phantom, and that they have only been moved collectively using the translation stage. The distortion is the reason for the discrepancy between the precision estimates. However, it was quite small with an angular distortion of at most 3.1° . A possible explanation could be that the experiment has been conducted assuming a speed of sound of 1480 m/s in the entire beamformed volume. This was chosen, since the phantom was submerged in water, and the phantom itself consists of $\approx 75\%$ water. However, the speed of sound of the phantom has been measured to be ≈ 1580 m/s, which will lead to distortion. One way to match the speed of sound of the water to the speed of sound in the phantom could be to add salt to the water. Figure 7.26 is recreated from [26]. The data shows how increasing the salinity of the water will increase the speed of sound. At 19 $^\circ\text{C}$, an 8.95% salinity will result in a speed of sound of ≈ 1581 m/s, which practically matches the measured speed of sound for the base layer exposure time of three seconds shown in Section 6.4.1 exactly. It should be noted that since the hydrogel is diffusion open to water, the salt water would also be absorbed in the phantom, likely changing its speed of sound as well. The exact concentration needed would need to be tested.

An alternative or additional explanation could be that the ultrasound system has both a spatially dependent sensitivity and a spatially dependent point spread function, which changes in shape and intensity. This would not only explain the non-parallel tracks, but could also explain the difference in the number of tracks detected in the two ultrasound experiments, and that the eighth scatterer was not localised in either experiment. A consequence of a spatially dependent point spread function could be that full calibration of a SRUS pipeline should perhaps be performed with local parameter estimates throughout the field of view of the probe instead of globally, as presented here. Thus the properties of a SRUS pipeline would then be given by accuracy and precision estimates, both as functions of the x , y , and z coordinates. This might even be necessary, illustrated by the results in this paper, as proper thresholding can become difficult to implement globally in the field of view.

The presented phantom concept could be expanded to investigate other aspects of super-resolution algorithms and systems, such as resolvability and separability. The separation of 3 mm in the scatterer phantom experiment was chosen to ensure no overlap between the reflected signals from the individual scatterers, thereby mimicking how many SRUS pipelines work today. The resolution that can be expected from an SRUS pipeline will be given by the variability of the positions, presented here as the σ values in Table 7.6. This is an indirect measure of resolution as it does not directly show feature separability. However, only slight modifications to the experiment would need to be made to show resolution directly. After having tracked the scatterers in the phantom when it was translated along a single direction, the phantom could be offset a sub-wavelength distance along a perpendicular direction, before being translated back again, parallel to the original

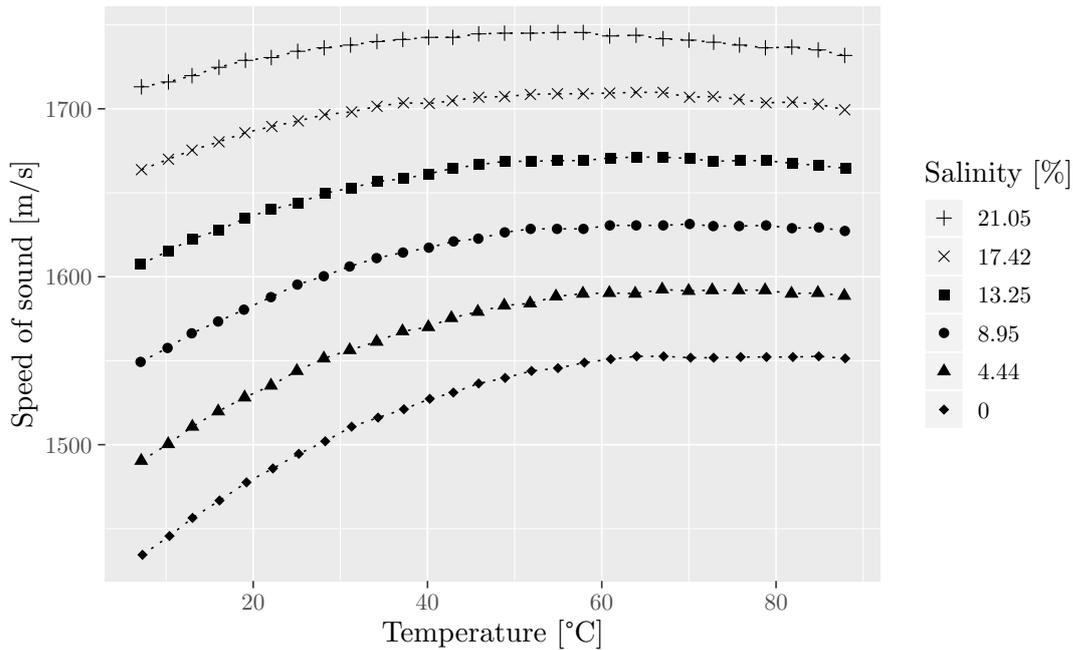


Figure 7.26: Speed of sound against temperature for different salinities. Graph is recreated from [26].

track. This would exactly mimic the functioning of current diffraction limited SRUS algorithms schemes. By changing the offset distance, the exact distance when the tracks become separable could be determined, demonstrating the pipeline resolution.

The 3 mm separation of scatterers in the phantom is not the limit, and phantoms could be developed with scatterers placed much closer as demonstrated in Section 7.2.5. This could be done to tune algorithms to be able to separate signals from partially overlapping reflections. This has been the objective of the work in Paper D and Paper J. In these papers, neural network scatterer localisation was presented, in part through experiments using phantoms with arrays of scatterers. To demonstrate the neural network detector, the separation was not sub-wavelength, but was created larger for initial demonstration. The lateral separation was $518 \mu\text{m}$ and the axial separation was $342 \mu\text{m}$. The scatterers in this experiment were according to the scatterer size model $55.3 \mu\text{m}$ by $55.3 \mu\text{m}$. As demonstrated, the scatterers can be separated by only 8 voxels, or $\approx 86.4 \mu\text{m}$. This opens up for a whole new set of experiments to optimize the neural network scatterer detection. An alternative method for decreasing the scatterer separation would be to use the two reflections from the front and the back of each scatterer as individual targets instead of training the neural network to recognize it as a single scatterer. However, initial testing has shown that signals from the two reflections are different, with one going from hydrogel to water, and one from water to hydrogel, which will therefore require another approach for training of the neural network.

The precision, accuracy and repeatability of the 3D printed phantoms would be incredibly difficult to achieve, if not impossible, using the traditional types of tube phantoms or chicken embryos. Yet, it still provides the opportunity of creating complex three-dimensional phantom features, providing the opportunity for full volumetric characterisation of an ultrasound system, which is not offered by any other phantom fabrication method available today.

The presented phantoms illustrates an alternative solution for SRUS pipeline calibration to regular tube phantoms. However, this does not mean that it is irrelevant to create phantoms, which allow flow of micro-bubbles to be tracked. These will be discussed in the next chapter.

7.5 Chapter summary

A new phantom concept for SRUS was presented, utilizing fixated scatterers in the phantom instead of micro-channels and micro-bubbles for a temporally stable reference structure. Three different scatterer concepts and two different shapes were analysed to determine how the printed size differed, and what the influence on the final reflected intensity was. It was found that square scatterers reflect higher intensity than circular scatterers, and that overexposing a 1 voxel wide frame at the edge of the the scatterers increase the reflected intensity the most. This is likely explained by multiple sub-wavelength reflections constructively interfering for increased reflected intensity. Scatterers can be placed as close as a single voxel from each other, providing a good test foundation for imaging methods capable of sub-wavelength scatterer separation. A scatterer phantom containing eight scatterers was created for evaluating the precision and accuracy of a 3D SRUS pipeline using a RCA array. Analysis of the data showed a good correlation between designed distances between the scatterers and the distances calculated based on the SRUS results, with a slope of correlation being 0.989, close to a perfect correlation slope of 1. Based on the same data, the precision of the SRUS pipeline was found to be between the two limiting estimates of $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.7 \mu\text{m}, 27.6 \mu\text{m}, 9.5 \mu\text{m})$ and $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.3 \mu\text{m}, 19.3 \mu\text{m}, 8.7 \mu\text{m})$, with the worst precision estimates being about 1/18th of the wavelength of $500 \mu\text{m}$ used in the experiment. The two sets of precision estimates stems from distortion in the beamforming, on a micrometre scale. This would not have been possible to discover using conventional tube phantom setups.

Flow phantoms for SRUS

This chapter describes the work on flow phantoms for SRUS. The chapter provides insight into the practical experience that has been obtained from each phantom iteration. First, some general considerations for making flow phantom structures are presented, before three different phantoms and the experiments conducted with them are presented. Finally, a number of phantom designs which have been created but not yet used are discussed. These designs are meant to further the controlled testing of SRUS. The content is in part based on Paper B, Paper F and Paper H.

8.1 General flow phantom considerations

The 3D printing method presented in this thesis is a solution which provides unparalleled control of feature placement in three dimensions across the entire phantom. It does however not allow for completely free control of feature dimensions or placement, as the structures need to be placed on the voxel grid. This means that by default, the features to be printed, whether being cavities, channels, or solid objects, will end up being built by $10.8 \times 10.8 \times 20 \mu\text{m}^3$ building blocks. This has a few consequences. First of all, the channel size cannot by default be controlled on a continuous scale. It is only possible to vary the size in steps of $10.8 \mu\text{m}$ or $20 \mu\text{m}$. Furthermore, whether it is possible to create a completely symmetric channel will depend on the orientation of that channel within the phantom. Since the voxel dimensions are equal along x and y , vertical channels can be printed completely symmetrical. However, horizontal channels will be subject to a print resolution of $10.8 \mu\text{m}$ in one direction and $20 \mu\text{m}$ in another direction. Channels at angles different than the main axes will also inevitably end up with dimensions differing from the main axes. This is illustrated in Figure 8.1, where channels designed to be $30 \mu\text{m}$ in diameter are placed along different directions of the phantom, and the actual dimensions are rounded off to fit the voxel grid. It should also be noted that while the printer and the print solution can theoretically be optimised to print the exact design dimensions in x and y , this will by definition of the printing method not be possible in z since a sufficient overlap between layers is needed. Thereby, dimensions in z will always be smaller than the design. Where all previous phantom sketches have shown the structures from the top on a 2D isotropic voxel grid, Figure 8.1 illustrates the channels from the side of the phantom, on the 2D anisotropic voxel grid. The side walls of channels which are not along the main axes will also inherently be “stair-like” as illustrated. As seen, although the channels were designed to be $30 \mu\text{m}$ in diameter, they end up being $32.4 \mu\text{m}$ if printed along the z -axis, and $40 \mu\text{m}$ if printed along the x -axis. If the channels are large, these difference might all be negligible, but for smaller channels this can be a problem as illustrated. At the very least, it is something to be aware of if the exact dimensions of the channel system is important.

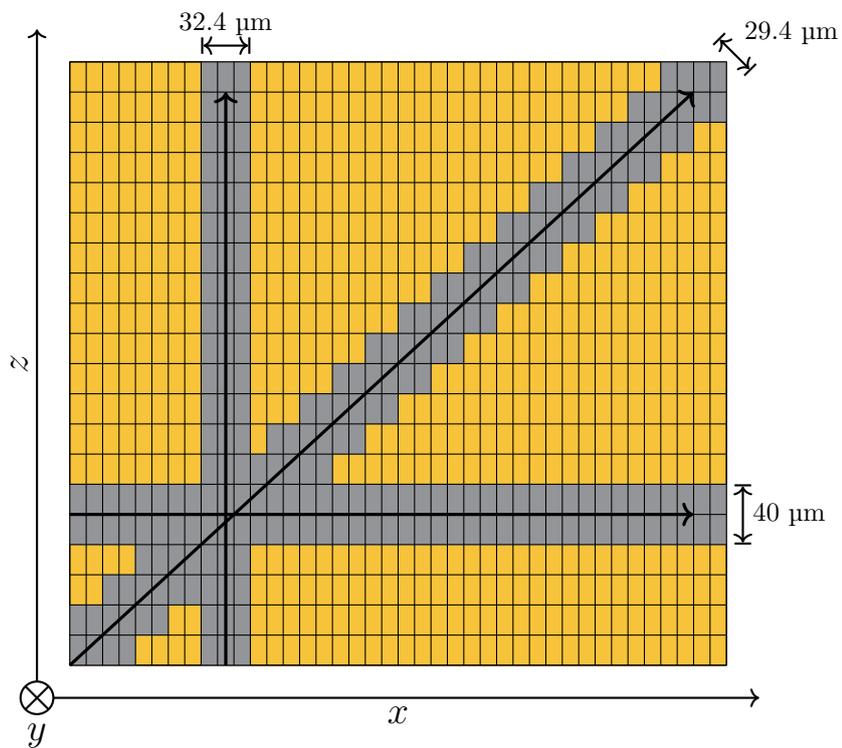


Figure 8.1: Channels designed to be 30 μm in diameter along different axes in the phantom. The anisotropic voxel grid means the channel dimensions will be rounded to fit the grid, and are in this case changing between the different axes.

The channel dimensions will be important if one wants to validate the algorithm velocity estimates. For such experiments, the flow is typically controlled by a flow controller, which can provide a stable volume flow rate through the system. For a volume flow rate of Q , the average flow velocity \bar{v} will be

$$\bar{v} = \frac{Q}{A} = \frac{Q}{\pi a^2}, \quad (8.1)$$

in a channel with cross-sectional area A and radius a . Thus, a change to the radius of the channel will also result in a change in the average flow velocity. Thereby, if the goal is to verify that the algorithm is capable of estimating the correct flow velocity, it is important to know the geometry of the channel.

To demonstrate the superior resolution of a SRUS algorithm with flow phantoms, it will be necessary to have two channels placed closer together than the diffraction limit of the imaging system. This has also previously been done in the literature [40, 157]. In order to obtain flow in two channels simultaneously, the volume flow from the flow controller can be split into multiple channels, which can then be imaged. However, unless the channels which the flow is split into are geometrically exactly the same, the volume flow rate will not be divided evenly between the two channels, due to different hydraulic resistance in the channels [50]. The Hagen-Poiseuille law, repeated here for convenience, expresses that

$$\Delta p = R_{\text{hyd}} Q, \quad (8.2)$$

where Δp is the pressure drop across the channel system, R_{hyd} is the hydraulic resistance and Q once again is the volume flow rate. For a circular straight channel with radius a , the hydraulic resistance is

$$R_{\text{hyd}} = \frac{8}{\pi} \eta L \frac{1}{a^4}, \quad (8.3)$$

where η is the dynamic viscosity, and L is the channel length. When a channel is split into multiple channels, the resulting channels are effectively placed in parallel, with the same pressure drop applied across all channels. Thereby, if one channel is smaller than the others, the volume flow will not be distributed evenly between the channels, which will influence the flow velocities. This property can be used as a benefit. If multiple different sized channels are placed simultaneously within the field of view, the performance of the velocity estimator can be evaluated across a range of velocities simultaneously. However, if the difference in geometry is unintended, for instance due to variability in tube dimensions or due to a local narrowing of one channel, one would likely conclude that the velocity estimator performs differently in different parts of the imaged volume, and it might be difficult to verify the accuracy of the estimator.

One way to mitigate this problem would be to only use a single channel, looping it around to return close to itself. In this way, the volume flow rate is necessarily the same everywhere in the channel segments. Even if there is a local narrowing of the channel, the flow velocity will only be affected directly at the narrowing, but will obtain the expected value everywhere else. Thus, the velocity magnitude at any point in the flow channel should be the same and can be compared throughout the phantom.

With the purpose of the phantoms being to provide the most predictable outcome, all flow phantoms have been designed with only a single channel, which is then bent in different ways to obtain the desired outcomes. Changing the flow trajectory like this is very simple when the phantoms are 3D printed.

8.2 Flow phantom for 2D SRUS

The first 3D printed flow phantom was an ambitious attempt at creating a unique flow phantom capable of demonstrating 2D SRUS with a geometry unattainable by conventional phantom fabrication methods.

8.2.1 Phantom description

The first flow phantom created was designed for 2D SRUS. It consisted of two 7 mm long square channel segments with side lengths designed to be 200 μm before mapping to the voxel grid, with the channels being separated by 100 μm . The phantom was modelled in Autodesk Inventor. Figure 8.2(a) shows an isometric view of the 3D model, and Figure 8.2(b) shows the x - z plane of the model. The channel is highlighted in blue to emphasise it. A larger inlet can be seen on the left of the model, to mount a tube needle from the flow controller. The diameter in this and all needle sections in the following is 700 μm to make a tight seal with a 800 μm outer diameter hypodermic needle. The flow should enter the bottom channel first, before exiting through the top channel. The printed phantom can be seen in Figure 8.2(c). Water containing blue fruit dye has been pumped through the channel to create optical contrast between the hydrogel and the channel. The phantom was printed with an exposure time of three seconds for each layer.

8.2.2 2D SRUS results

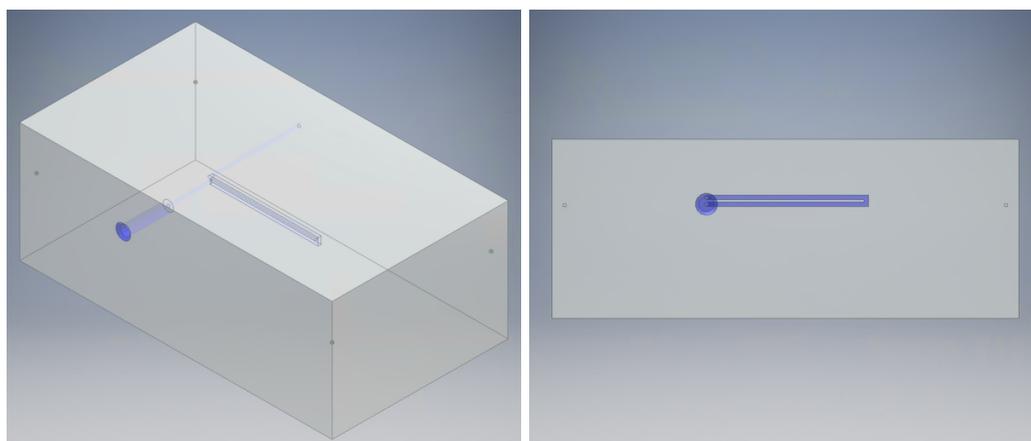
The phantom was used for some of the first controlled SRUS experimentation and provided a lot of valuable experience. SonoVue micro-bubbles were used, in a 1:50 dilution of the standard solution, infused into the system at 2 $\mu\text{L/s}$. The data was subsequently analysed using the SRUS pipeline developed by collaborators at CFU at DTU. The resulting data can be seen in Figure 8.3, in which Figure 8.3(a) shows the accumulated density of the detected micro-bubbles across the imaged frames, Figure 8.3(b) shows the direction of flow, indicated by the inserted colour wheel, and Figure 8.3(c) shows the velocity magnitude.

The obtained B-mode video of the data combined with the three images, provide some interesting insights. First of all, Figure 8.3(b) shows that the bubbles enter from the right in the bottom channel, connects vertically upwards to the top channel, and exits to the right of the image again, as would be expected. Figure 8.3(a) show the density of micro-bubbles in the image. With a constant flow, and a lot of bubbles being detected, it would be expected that one might see a density corresponding to the parabolic flow profile across the channel due to the velocity profile itself, but otherwise constant through the length of the channel. However, that is not the case. Furthermore, there is a high number of localisations in the top left corner of the channel. The video showed a lot of micro-bubbles flowing through the channel at a very high velocity. This is mirrored in Figure 8.3(c), with some micro-bubbles having a flow velocity faster than 60 mm/s, and the majority having a flow velocity around 20 mm/s. As previously stated, typical flow velocities in the smallest vessels are around 2 mm/s, meaning the applied flow velocities are not particularly representative. Furthermore, considering that the length of the two horizontal channel segments is 7 mm, a micro-bubble travelling at 20 mm/s or 60 mm/s would only be present in the field of view for 0.75 second or 0.25 second, respectively. This will decrease the number of frames in which the same micro-bubble can be located and tracked. This might also explain the apparent errors in the flow velocity figures, with some micro-bubbles being illustrated as travelling across the boundary between the channels. This shows how important flow control is, and that it will be a good idea to decrease the flow velocities, not only to match values which are biologically relevant, but also for more reliable performance.

The B-mode video showed many micro-bubbles being very close together in many frames. This likely means that many of the micro-bubbles would not have contributed to the SRUS images, but would have been discarded due to partial overlap of their PSFs. With a more suitable concentration of micro-bubbles, the resulting SRUS might have contained even more tracks.

The results also indicated the need for fiducial markers. It appears in Figure 8.3(a) that there is an overall tendency of higher density of micro-bubbles to the left in the image compared to the right. This might be due to misalignment between the phantom channels and the ultrasound probe. Thereby, it became apparent that an alignment system would be required for future experiments, which resulted in development of fiducial markers to be included in the phantoms.

Another practical issue with all phantoms is that the printed features are in general so small that they are difficult to see with the naked eye. This makes mounting of the phantoms in the correct

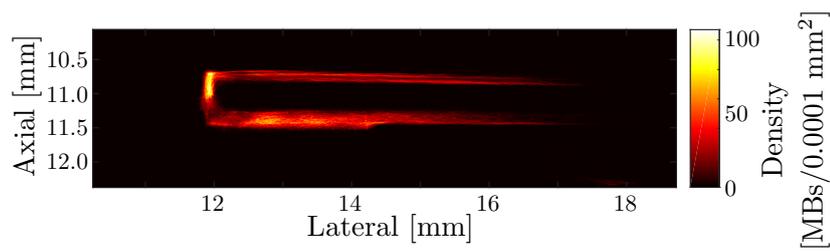


(a) Isometric view of 3D model of phantom in Autodesk Inventor (b) x - z cross-sectional view of 3D model of phantom in Autodesk Inventor

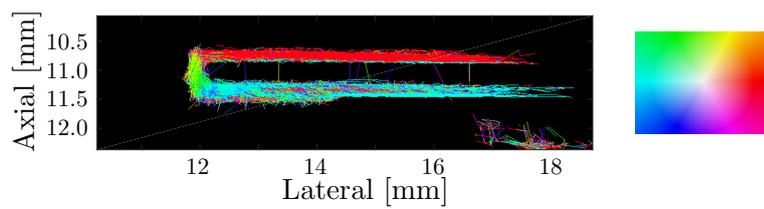


(c) Image of the printed phantom

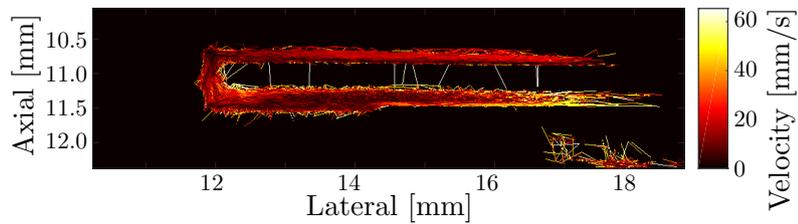
Figure 8.2: The 2D SRUS flow phantom. (a) Isometric view of the 3D model of the flow phantom for 2D SRUS created in Autodesk Inventor. (b) Cross-sectional view of the phantom. The channel surfaces have been highlighted in blue. (c) Image of the printed phantom. Water containing blue fruit dye was pumped through the channel to increase contrast between the hydrogel and the channel.



(a) Accumulated micro-bubble density



(b) Flow velocity direction



(c) Flow velocity magnitude

Figure 8.3: SRUS results using the 2D flow phantom. (a) shows the accumulated density of the detected micro-bubbles across the imaged frames. (b) shows the direction of flow, indicated by the inserted colour wheel. Flow enters from the right in the bottom channel, connects vertically upwards to the top channel, and exits to the right of the image again. (c) shows the velocity magnitude.

orientation a difficult task. Therefore, mounting marks have been included in all subsequently printed phantoms. The mark is seen as a quarter sphere placed in the middle of one of the top edges of the phantom, for instance in Figure 8.6(b). This placement allows for unambiguous mounting of the phantoms.

This phantom design was the last one created in Autodesk Inventor. The lack of control when doing automatic slicing of a 3D model and therefore uncertainty in the actual printed phantoms, meant that the requirements for precision of placement of fiducial markers and dimensioning of flow channels, would be unattainable. All other phantom models, including those presented in previous chapters, were created directly in MATLAB.

8.3 MATLAB phantom generation

Each phantom model was created as a matrix with each matrix element representing a voxel in the printed phantom. Using the full printing area, and matching the height of the phantom to the width of the phantom, the matrix dimensions become $1920 \times 1080 \times 583$. Thereby, any design created would by default be matched to the voxel grid, with each layer in the matrix representing a single slice.

Inclusion of fiducial markers is quite simple, as only a list of the marker coordinates matched to the voxel grid is needed. A number of surrounding voxels is then just marked, to obtain the desired size of the marker. As the markers are typically smaller than the imaging wavelength, the fiducial marker coordinates should be placed in the centre of each scatterer. In practice, the matrix elements corresponding to the fiducial markers are marked by first finding the centre voxel corresponding to the fiducial coordinate, and subsequently marking the number of voxels to both sides of the centre voxel, which adds up to the total number of voxels. However, mathematically, it is not just a matter of dividing the total number of voxels in two, rounding off, and highlighting that number of voxels on both sides, as one will end up with the scatterers becoming larger than intended. A small correction needs to be made regarding the method of rounding off the number of voxels. The issue is illustrated in Figure 8.4(a). The centre voxels are marked by white crosses. White arrows mark the distance corresponding to half the number of voxels of the marker size to each side of the centre voxel, in this example case 2.5 and 3 to each side. Using the regular “round” function will highlight more voxels than intended. The exact number will vary, but is always larger than intended. The solution consists of two steps: Making a slight offset of 0.1 to the location of the centre voxel, illustrated by the location of the small vertical white lines. The rounding is then always done towards the centre voxel on both sides, using “ceil” for the lower value and “floor” for the higher value in MATLAB. The combination of the two steps will result in the correct number of voxels being highlighted regardless of whether the desired number of voxels n in the fiducial marker is an odd number, illustrated with 5 voxels, or an even number, illustrated with 6 voxels. It is of course not possible to have a true centre voxel in a fiducial marker consisting of an even number of voxels. In this case, a correction to the centre coordinate of half a voxel will need to be made. Note that the rounding is done towards the centre of the voxels. The described rounding correction has been applied in all of the phantom designs presented.

Special functions (included in Appendix G.1.1) were created for defining the flow channels. The scripts requires the phantom matrix, channel diameter, start- and end coordinate of the phantom, and the radius of curvature if the phantom is to bend. The scripts were designed to align the main sections of the channels along the printer x -axis, y -axis, and z -axis, which so far has suited our needs. Furthermore, in the case that the start- and end coordinates along all three axes, the channel section order will always be first x then y then z . This might seem limiting. However, if the channel is split into segments only containing bends along two axes at a time, any channel system can be made. This is illustrated in Figure 8.4(b) for a channel with two segments, one along x and one along y . There are two possible configurations as shown in green and blue. Changing which end of the channel is considered the “start coordinate” and which is considered the “end coordinate” will result in the two different configurations. Based on the coordinates and the radius of curvature, the centre line of the channel will be defined and mapped to the voxel grid. Then

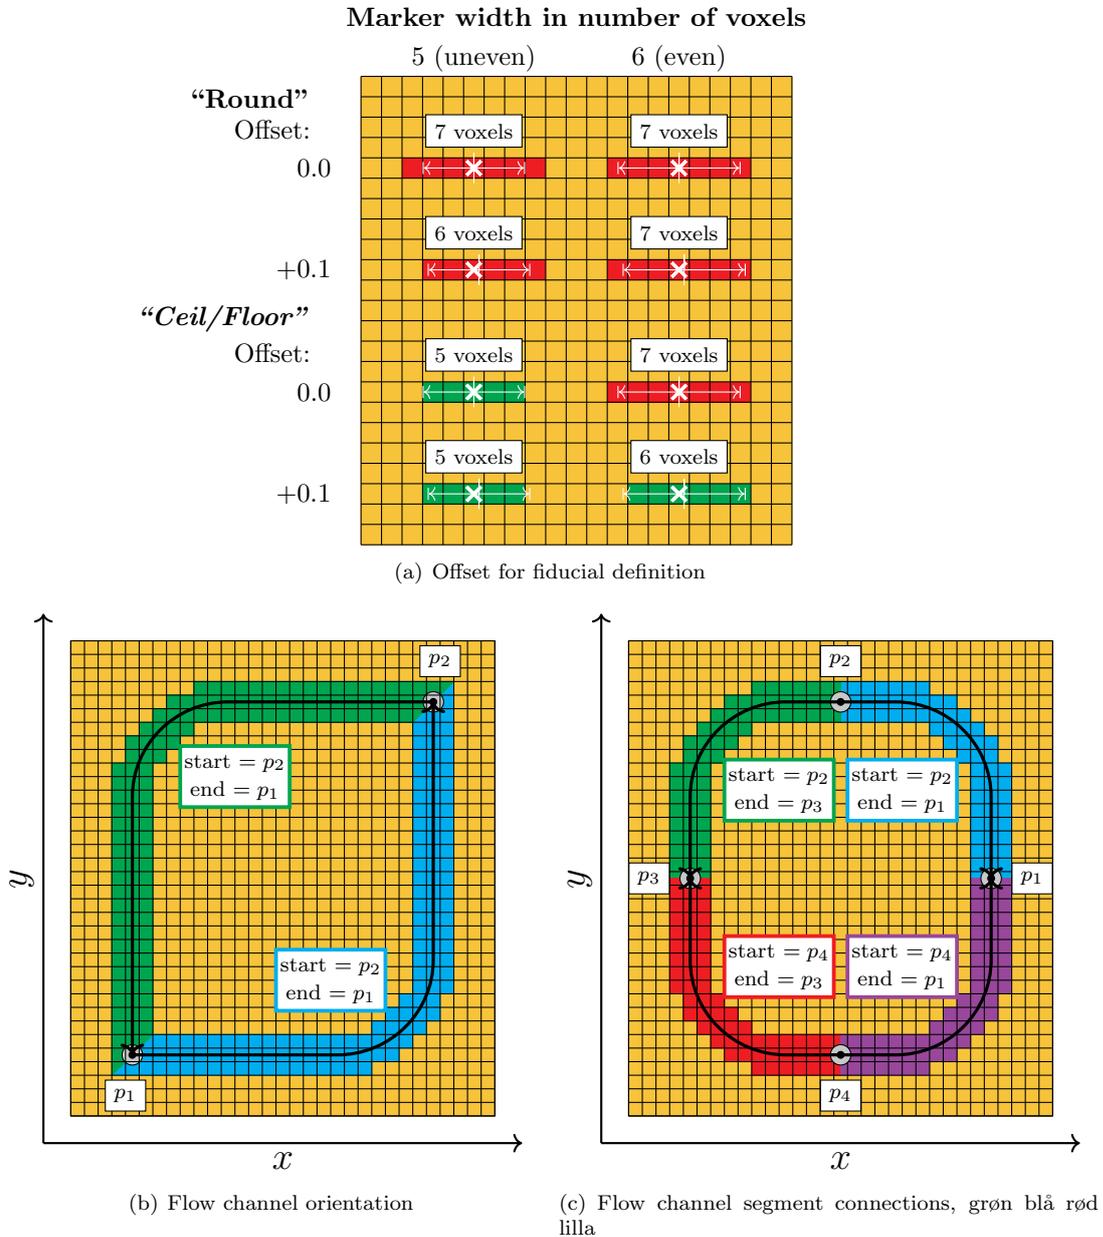


Figure 8.4: MATLAB phantom design concept sketches. (a) illustrates how a small offset and special rounding needs to be implemented when the fiducial marker sizes are defined. (b) shows how the channel function, which is limited to orient the channel segments first along x then y then z can create any channel configuration. For a channel with two sections, one along x and one along y , there are two possible configurations as shown in blue and green. Changing which end of the channel is considered the “start coordinate” and which is considered the “end coordinate” will result in the two different configurations. (c) illustrates how channel segments should be connected at straight segments to round all corners. The arrows in (b) and (c) also mark the start and end coordinates.

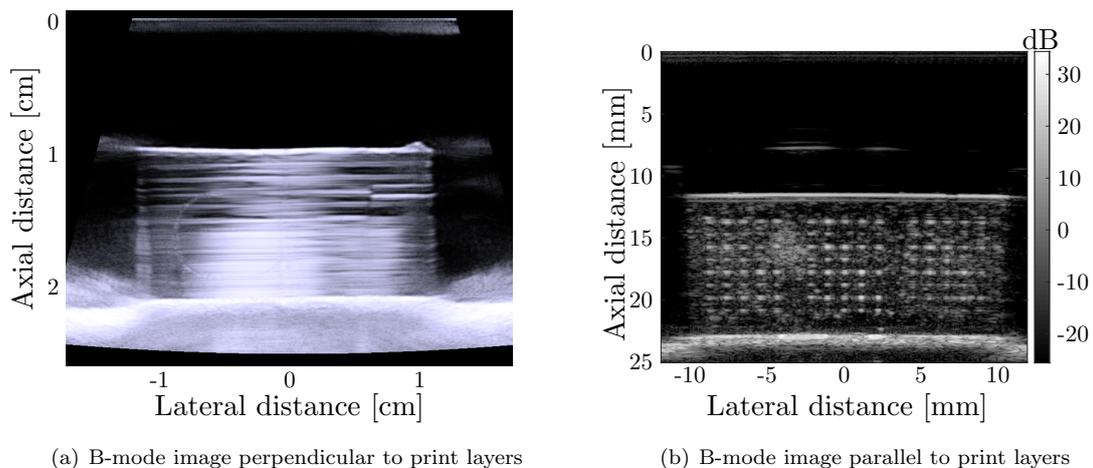


Figure 8.5: (a) B-mode image of a channel phantom. The horizontal lines are reflections at interfaces between layers in the phantom. To avoid this effect, all phantoms were printed on the side. Thereby, the axial orientation in images becomes the printer y -axis, between which the uniformity is significantly higher, therefore not resulting in systematic lines in the images. The difference can be seen in (b) for the previously shown scatterer phantom.

a sphere with a diameter matching the desired channel diameter is slid along this path, marking all of the voxels which form the channel. By splitting the channel into multiple segments, the diameter can be changed locally, for instance to accommodate a larger inlet for the needle from the flow controller. In all subsequent descriptions of flow phantoms, the used centreline is also shown.

Although not presented in the thesis, the same approach has been applied to create square channel systems. The only necessary change was that a cylinder is slid along the channel path instead of a sphere. When the cylinder is dragged along the path, the proper channel width will be obtained at the bends as well. It is important to consider in which orientation the bend is, as the cylinder will need to be rotated on the side if the channel bends from either x or y into z , or vice versa. This script can be seen in Appendix G.1.2.

Given that the channel function only takes a single start- and end coordinate, it will have no knowledge of the orientation of previous channel segments. This means that if the start- and end coordinates actually refer to corners in the channel system, it will not be possible to apply a radius of curvature to these corners. This can be seen in Figure 8.4(b), where the start and end coordinates are marked by grey disks. If instead, the start- and end coordinates mark positions along straight segments of the channels, all corners can be rounded, as illustrated in Figure 8.4(c).

Fairly simple adjustments to the script could be implemented which would allow for any channel orientations to be made. It would only require that the bending radii of curvature were directly included in the channel path, by which a cylindrical channel system of arbitrary orientations could be made very easily by sliding the sphere along this new arbitrary path.

Figure 8.5(a) shows a B-mode image of a printed phantom, in which the phantom is mounted in the same orientation as it was printed. The horizontal lines are reflections at interfaces between layers in the phantom. To avoid this effect, all phantoms were printed on the side. By flipping the phantoms, the axial orientation in images becomes the printer y -axis, between which the uniformity is significantly higher, therefore not resulting in systematic lines in the images. The scatterer phantom was printed on the side, and is included in Figure 8.5(b) for comparison. It can be seen that the background structure is significantly different. The background pattern is constant in the images, and it will therefore always be desirable for it to be of as small intensity as possible. For flow phantoms, in which the flow is constantly moving, it can quite easily be removed through stationary echo cancellation, leaving only the moving elements in the images. This will be demonstrated in Section 8.5.2. However, it is critical to minimise the effect of the pattern for

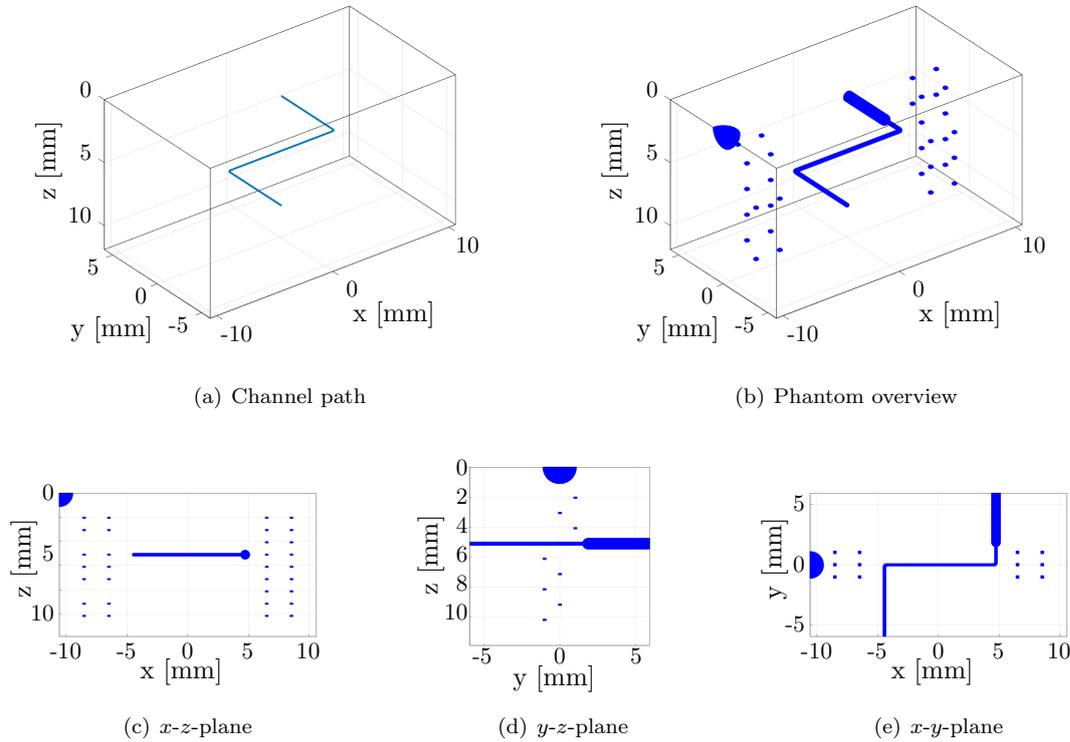


Figure 8.6: Design of the single channel phantom for 2D and 3D imaging.

the scatterer phantoms.

Through the presented design methods and functions, accurate documentation of the location of fiducial markers as well as the centre line in the flow channels can be provided.

8.4 Single channel phantom for 2D and 3D super-localisation

To get a better basis for doing SRUS experiments, optimise alignment and demonstrate SRUS in 3D, an even simpler phantom was created, consisting of a single central straight channel. The phantom had multiple purposes: Characterisation and optimization of the acoustic micro-bubble response, which will not be presented in this thesis; Simple channel foundation for 3D SRUS using a RCA probe.

8.4.1 Phantom description

The phantom design can be seen in Figure 8.6. It consists of a 200 μm cylindrical channel only in a single plane, with a 5.8 mm long inlet with a wide section for needle insertion. The channel then bends 90° with a 200 μm radius of curvature, continuing in a 7 mm straight segment, before bending 90° with a 200 μm radius of curvature into the 5.8 mm long outlet channel. For 2D imaging, the 7 mm channel segment would be used, whereas practically the entire channel system can be included in the FOV of an RCA array for 3D imaging.

The printed phantom can be seen in Figure 8.7. Water containing blue fruit dye has been pumped through the channel to create optical contrast between the hydrogel and the channel. The phantom including fiducial markers was printed with a layer exposure time of three seconds.

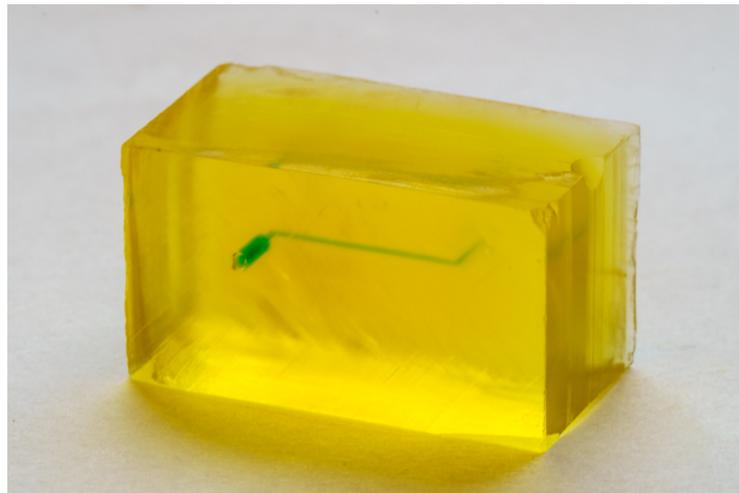


Figure 8.7: Image of the single channel phantom. Water containing blue fruit dye was pumped through the channel to increase contrast between the hydrogel and the channel.

8.4.2 Fiducial marker layout

Fiducial markers are conceptually meant as smaller structures which can be imaged and aligned to. For the experimental setups used, the alignment can be done along the x -axis, the y -axis, and rotation about the z -axis. If the phantom consists of a single straight channel, a simple design could be to place one scatterer in extension of the channel at both ends. For best rotational alignment, the markers should be placed as far from each other as possible, within the probe FOV. However, in practice it is not that simple, and the fiducial marker layouts have been iteratively improved throughout our experimentation. Looking at the B-mode images presented so far in the thesis, for instance Figure 8.5(b), it can be seen that there are a lot of background phantom inhomogeneities observable as dots of varying intensity. Finding a single scatterer at the end of a channel, on a background of dots of varying intensities is like looking for a needle in a haystack. However, the human brain is excellent at identifying patterns. This is also well demonstrated in Figure 8.5(b), in which finding the regular pattern of scatterers on top of the noisy background is quite easy. Therefore, it was decided that the fiducial markers should be arranged in specified patterns. Ideally, these patterns would make it very easy to identify left and right in the phantom, allow for alignment to a specified plane, and provide visual assistance on how to correct the misalignments.

The developed fiducial marker design is illustrated in Figure 8.8. Figure 8.8(a) shows the x - z -plane of the design. Figure 8.8(b) shows the y - z -plane of the design. The patterns are not symmetric from left to right. The dashed lines in Figure 8.8(b) represent the elevation focus of a 2D imaging probe. The green fiducial marker columns should be aligned to the flow channel of interest. The separation between columns in the elevation direction should be large enough that only the centre column of markers in the y - z -plane are within the elevation focus, as illustrated in the figure. Thereby, misalignment in any direction will allow an extra column of scatterers to be within the elevation focus, and additional markers will therefore appear either at the top or the bottom of the design. If both sides of the phantom show only extra markers at the top, or at the bottom, the misalignment is purely translational, and can be corrected using the design both in terms of the direction of the corrections as well as the amount of correction, based on the knowledge of how far the marker columns are separated. If one side shows markers at the bottom, and the other side shows markers at the top, the misalignment is also rotational. In the implemented design, marker columns are separated by 2 mm in the x - z -plane, by 1 mm in the y - z -plane, and by 1 mm vertically.

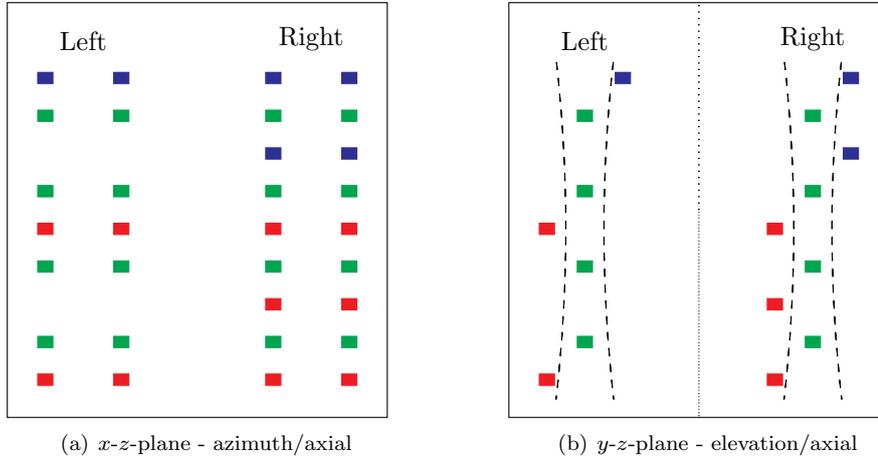


Figure 8.8: Fiducial marker layout concept. (a) shows the x - z -plane of the design. The patterns are not symmetric from left to right. (b) shows the y - z -plane of the design. The dashed lines in represent the elevation focus of a 2D imaging probe. The green fiducial marker columns should be aligned to the flow channel of interest. The separation between columns in the elevation direction should be large enough that only the centre, green, markers are within the elevation focus.

8.4.3 3D super-localisation results

Ultrasound data was acquired over 28 seconds using a prototype 62+62 elements 3 MHz PZT RCA array [61]. The probe was connected to the experimental synthetic aperture real-time ultrasound system (SARUS) [153]. The maximum volume rate of the imaging sequence was 156 Hz at a pulse repetition frequency of 10 kHz. For the experiment, the volume rate was lowered to ≈ 14 Hz for a total of 400 volumes. The sequence is described in detail in Paper H. SonoVue micro-bubbles were used, in a 1:10 dilution of the standard solution, infused into the channel system at 1.61 $\mu\text{L/s}$, which resulted in a peak velocity of 102.4 mm/s. The data was subsequently analysed using the SRUS algorithm developed by my collaborators at DTU.

The super-localised positions of the micro-bubbles in the phantom can be seen as the blue dots in Figure 8.9. The used experimental parameters reflect that the results from the initial 2D double channel phantom had not been analysed in depth at the time of the experiment. The results in this experiment thus suffer from some of the same problems. A micro-bubble with a velocity of 102.4 mm/s will only be present inside the phantom for ≈ 0.18 seconds. With the utilized volume rate, a single micro-bubble will only be captured in three consecutive volumes before exiting the phantom again. Clearly, an attempt at tracking the individual micro-bubbles and imaging the flow velocities based on that would once again result in images where micro-bubbles would apparently travel through the hydrogel, and not only follow the flow channel, similarly to Figure 8.3(c). However, the actual super-localised micro-bubble positions presented in Figure 8.9 show that the micro-bubbles are in fact only ever localised in positions which correlate with the phantom design.

At each of the bends in the flow channel, the inferred diameter of the channel based on the micro-bubble localisations appear wider than in the straight segments. There has been no evidence suggesting that this is a true effect in the printed phantom; it could be an artefact in the SRUS pipeline. This should be investigated further.

Given that the phantom only has a single straight channel and no other channel segments placed closer than the diffraction limit set by the scanning system, it is not possible to demonstrate the resolution capabilities of a SRUS pipeline using this phantom. However, as noted in the introductory Section 2.2.1, the localisation precision and the resolution of the system are closely linked, and thus, the localisation precision will hint at the resolution. By isolating the straight segments of the inlet, central channel, and outlet, individually, as demonstrated by the blue crosses in Figure 8.10(a), it is possible to make an indirect estimate of the localisation precision. The localisation precision

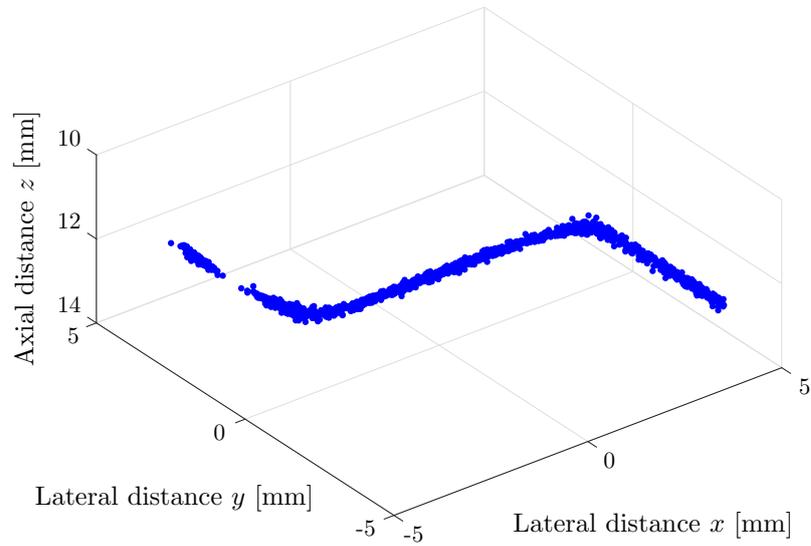


Figure 8.9: Super-localised micro-bubble positions in the single channel phantom.

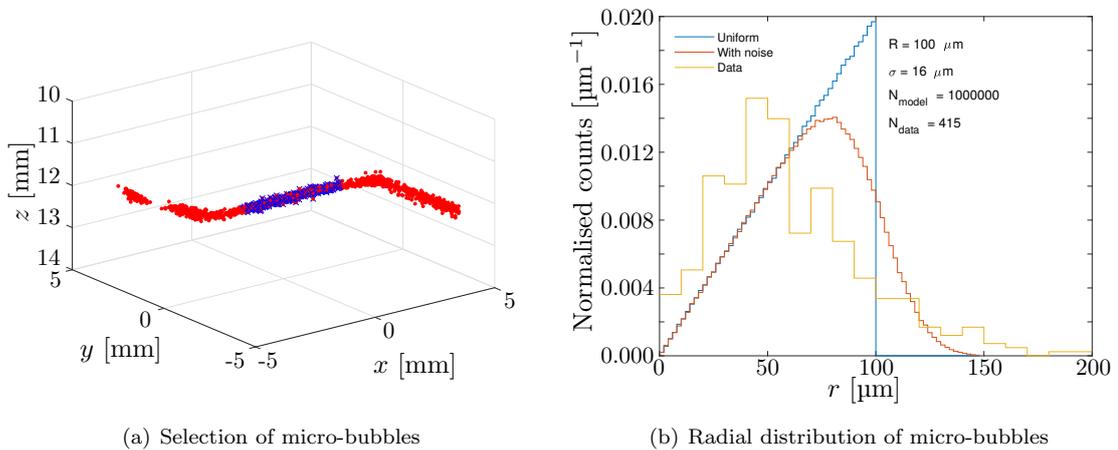


Figure 8.10: Procedure for estimating the localisation precision from flow channel data. (a) highlights a subset of the micro-bubble localisations located along a single straight segment. (b) shows the distribution of the localisations radially from the centreline of the used subset. The blue curve illustrates the radial distribution with uniform micro-bubble distribution in the channel. The red curve shows the expected distribution when measurement uncertainty is accounted for, showing how some micro-bubbles might appear to be located outside of the channel boundary. The yellow curve shows the actual radial distribution from the selected subset.

stems from the noise in the scanning system. Due to that noise, a micro-bubble right at the edge of the channel might erroneously be localised outside of the channel, with the localisation precision determining how far the position estimate can be offset. A straight line is fitted to the data in each segment, which therefore represents the centreline of the channel. The distance from each localised micro-bubble to the centre is calculated. The procedure effectively takes the data from 3D, to a 2D cross-sectional projection, into a 1D radial distribution. The phantom was mounted roughly along the transducer axes, but calculating distances relative to the fitted line instead of the transducer axes will remove a potential effect from misalignment of the phantom to the ultrasound probe. If one assumes that the micro-bubble localisation are uniformly distributed across the cross-sectional area of the channel, the radial distribution of micro-bubbles will be a straight line from zero, out to the radius of the channel, as illustrated by the blue curve in Figure 8.10(b). This is because each infinitesimal area increment scales as the circumference, $2\pi r$, and is therefore linear in r . Assuming the localisation precision is normal distributed with the same variance in each dimension, the radial distribution of all micro-bubbles in the segment will follow the distribution

$$f(r) = 2\pi r \int_{|\vec{r}_t| < R} \frac{1}{\pi R^2} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{|\vec{r} - \vec{r}_t|^2}{2\sigma^2}\right) d^2r_t, \quad (8.4)$$

where r is radial position, R is the radius of the tube, and σ is the standard deviation. The integral is a convolution of a constant density ($1/(\pi R^2)$) with a two-dimensional Gaussian. The non-analytical integral (8.4) is estimated in a Monte-Carlo calculation and is a Rayleigh distribution convolved with a uniform disk distribution of radius $R = 100 \mu\text{m}$. The factor $2\pi r$ is the Jacobian needed to convert from Cartesian to cylindrical coordinates. The convolved theoretical micro-bubble distribution is shown as the red curve in Figure 8.10(b). After convolution, a certain fraction of the micro-bubble positions will be outside of the channel boundary, illustrated as the vertical blue line, where the uniformly distributed micro-bubble curve goes to zero. This fraction will be given by the localisation precision of the actual measured data, or the standard deviation σ of the theoretical distribution. Thus, when the data has been collected, the fraction of micro-bubbles localised outside the boundary of the tube can be calculated, and the localisation precision can be reverse engineered by simulation, by finding the standard deviation of a 2D Gaussian distribution which results in the same fraction of localisations outside of the channel boundary. The radial distribution of the measured micro-bubble localisation is shown as the yellow curve in Figure 8.10(b). For the channel with radial distributions in the $y - z$ plane, the fraction of localisations outside of the channel boundary was found to be 13%, corresponding to a localisation precision of $16.5 \mu\text{m}$. For the channel with radial distributions in the $x - z$ plane, the fraction of localisations outside of the channel boundary was found to be 18%, corresponding to a localisation precision of $23 \mu\text{m}$.

This model is based on a number of assumptions, of which some are known not to be met. First, it can be problematic to use a model as a predictor when there are so large differences in the bin sizes between the simulated data and the measured data. The reason for this is simply due to the availability of data, with a total micro-bubble count of 415. Of course more data would always be beneficial when basing calculations on data distributions. Second, reducing this to a radial problem is problematic, since the ultrasound imaging system is not expected to have equal precision along the different axes: Ultrasound systems provide higher precision in the axial direction compared to the lateral direction, and the distributions in both channel segments are mixtures of the axial direction with one of the lateral directions, which consequently means the localisation precision estimates are as well. Thus it is reasonable to expect that the radial estimates will overestimate the true axial precision and underestimate the true lateral precision. Third, the foundation is a uniform distribution of micro-bubbles in the channel cross-sectional area. Due to the no-slip boundary condition of the fluid in the micro-channel, the flow velocity at the edge of the channel is zero, and consequently, there will not be any micro-bubbles there. The velocity profile increases towards the centre of the micro-channel, and the micro-bubble distribution would be expected to gradually increase from zero. When the distribution is then normalised with circumference for a radial distribution, the result would be a curve somewhat similar to the red curve in Figure 8.10(b),

but with the tail reaching zero at the channel boundary. Finally, the model does not really seem to fit the data that well. The data is skewed significantly towards the centre. This is actually quite well in line with the previous point of the wrong foundational distribution. None of these issues are present for the scatterer calibration phantom presented in the previous chapter.

Considering all of these erroneous assumptions, it is surprising that the localisation precision estimates based on this model are quite well in line with those determined using the scatterer phantom. In fact, both estimates obtained from the flow phantom split the lateral precision estimate of $17.3 - 27.6 \mu\text{m}$ and axial precision estimate of $8.7 - 9.5 \mu\text{m}$ found using the scatterer phantom, which would be expected. Thus, the method functions well as a first order estimate of the localisation precision.

8.5 Looping flow phantom for 3D SRUS

The previous phantom demonstrated how localisation precision in 3D can be estimated directly from the micro-bubble localisations, which will hint at the system resolution. However, instead of inferring the resolving power of the SRUS pipeline, direct demonstration of it would be preferable. That is the purpose of the next design.

8.5.1 Phantom description

The phantom design can be seen in Figure 8.11. It consists of a $200 \mu\text{m}$ cylindrical channel which loops around and passes $108 \mu\text{m}$ above itself at a 90° angle. The radius of curvature at the corners of the loop is 2.7 mm to ensure that the bend itself will be visible in B-mode, while the separated crossing will not. The phantom has been designed to allow for inclusion of the crossing channels, the entire channel loop, as well as the fiducial markers within the FOV of the RCA probe used in previous experiments. The vertical bend only has a radius of curvature of $200 \mu\text{m}$. This was the simple implementation based on how the channel generation works, but the ideal displacement would be done gradually, for a smooth transition with a small disturbance of the flow. However, the important feature of the phantom is the crossing of the channels, and the vertical bend was therefore placed at the back of the loop to allow the flow to re-stabilize before the channel crossing.

The printed phantom can be seen in Figure 8.12. Water containing blue fruit dye has been pumped through the channel to create optical contrast between the hydrogel and the channel. The phantom was printed with a layer exposure time of three seconds, while the Single pixel dosing scheme was applied to the fiducial markers.

A new fiducial marker layout was made for this phantom. The overall layout of the fiducial markers are the same as for the single channel phantom. However, as seen in Figure 8.11(e) the two pairs of markers are not placed along a channel, or even in the same plane. This was done to avoid any signal interference between that from the markers and micro-bubbles in the channel. Instead, the groups are placed on each side of the channel crossing along the x -axis, offset along the y -axis. It has been designed strictly for 3D imaging, with the placement of the fiducial markers not allowing to align a 2D imaging probe to a plane. Thus, if 2D cross-sectional images are desired, alignment will have to be done directly to the channel, which is not ideal. This should be changed in the future.

8.5.2 Looping flow phantom results

The phantom has been used for illustrating a number of different features, and has illuminated some of the issues which have been stated previously. As mentioned it was designed to demonstrate the 3D resolution capabilities of the 3D SRUS pipeline. Building on the previous experience the volume flow velocity was decreased to $0.06 \mu\text{L/s}$ and the micro-bubble concentration decreased to a 1:100 dilution compared to the standard solution. The prototype $62 + 62$ elements 3 MHz PZT, RCA array was used again, connected to the experimental SARUS. Figure 8.13 shows a 3D B-mode volume of the channel crossing and the full loop. The image illustrates that the two channels are not

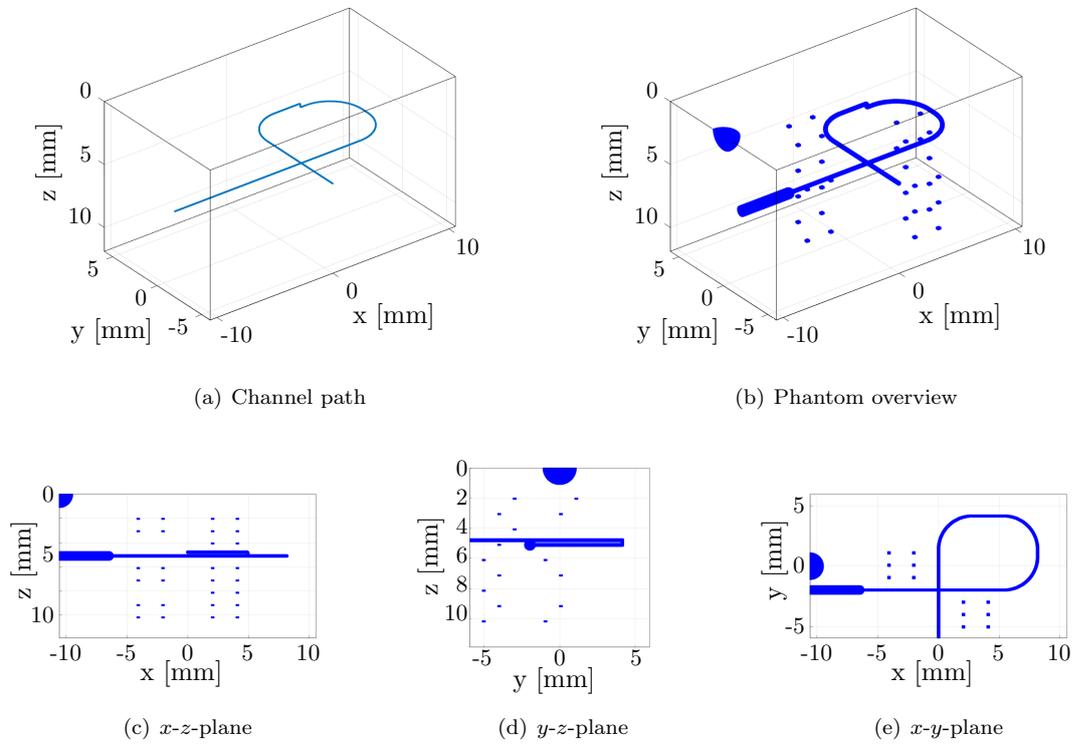


Figure 8.11: Design of the looping flow phantom for 3D SRUS.

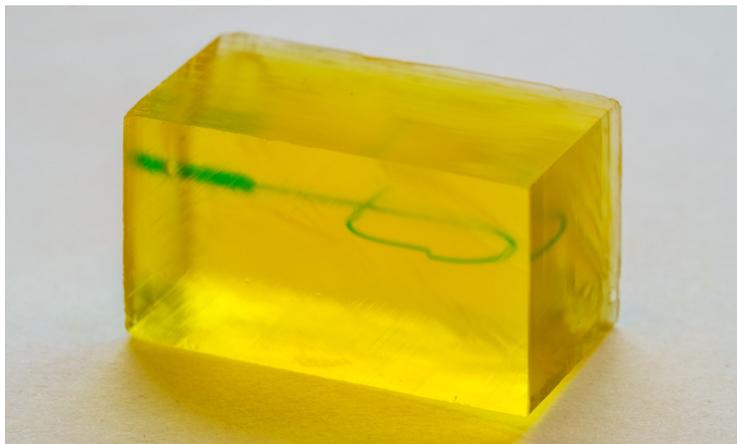


Figure 8.12: Image of the looping channel phantom. Water containing blue fruit dye was pumped through the channel to increase the contrast between the hydrogel and the channel.

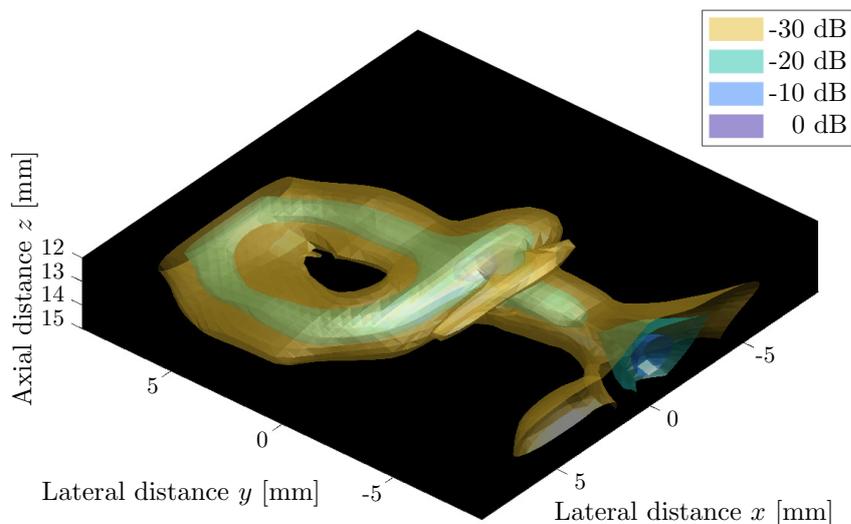


Figure 8.13: B-mode volume of the looping phantom acquired with a 3 MHz RCA array. The iso-surfaces mark intensities at 0 dB, -10 dB, -20 dB and -30 dB.

distinguishable in regular B-mode images, separated by only $108\ \mu\text{m}$. Although the micro-bubble concentration was decreased, no micro-bubbles were detected in the images, despite the fact that they had been visualised in the phantom in control 2D B-mode images using 2D imaging probes. The hypothesis is that the pressure from the RCA array is too high, bursting the bubbles, therefore resulting in no micro-bubble detections. The micro-bubble debris is expected to be providing the wide signal seen in the B-mode volume. It should be noted that the image is not a summation of volumes across the acquisition time, but a single volume. Micro-bubbles bursting might also explain why it was possible to detect micro-bubbles in the single straight channel phantom in the previous section, even though a high concentration of micro-bubbles was used. If most of the bubbles in the solution are bursting but a few remain, the resulting concentration might end up being sufficiently low, and suitable for SRUS. The experiment has not yet been repeated with a higher micro-bubble concentration.

In an ongoing series of experiments, the phantom is used to illustrate the benefits obtained when doing super-resolution in 3D in terms of the elevation resolvability. By placing the channel loop vertically, and aligning a 2D transducer to the channel segments, the entire channel system can be captured within the width of the elevation focus. This is illustrated in Figure 8.14. Figure 8.14(a) shows the B-mode image of the phantom. The vertical configuration places the phantom perpendicular to the layer orientation seen as the horizontal lines again. Figure 8.14(b) shows a contrast enhanced image taken interleaved with the B-mode image, illustrating the isolation potential when imaging non-linear objects. In this case, the micro-bubble signal within the channel is clearly visible.

Another method for isolating the flow signal would be to apply a stationary echo filter. Such filters can be applied in various forms, but the simplest method will be to consider only the difference between consecutive B-mode images. In this way, only features which have moved between frames will be visible. Figure 8.15(a) shows an example of the difference between two consecutive images. Even though each B-mode image appears similar to Figure 8.14(a), the difference between two consecutive frames reveals significant movement due to the flow, beneath the static background. These difference images can then be quadrated to obtain the energy in each image, and integrated over time to reveal the areas of significant movement in the images, as seen in Figure 8.15(b). The channel shape is clearly identifiable with two hot regions shown in the vertical segments of channels. The exact reason for this is unknown, but is clearly observable when watching the video

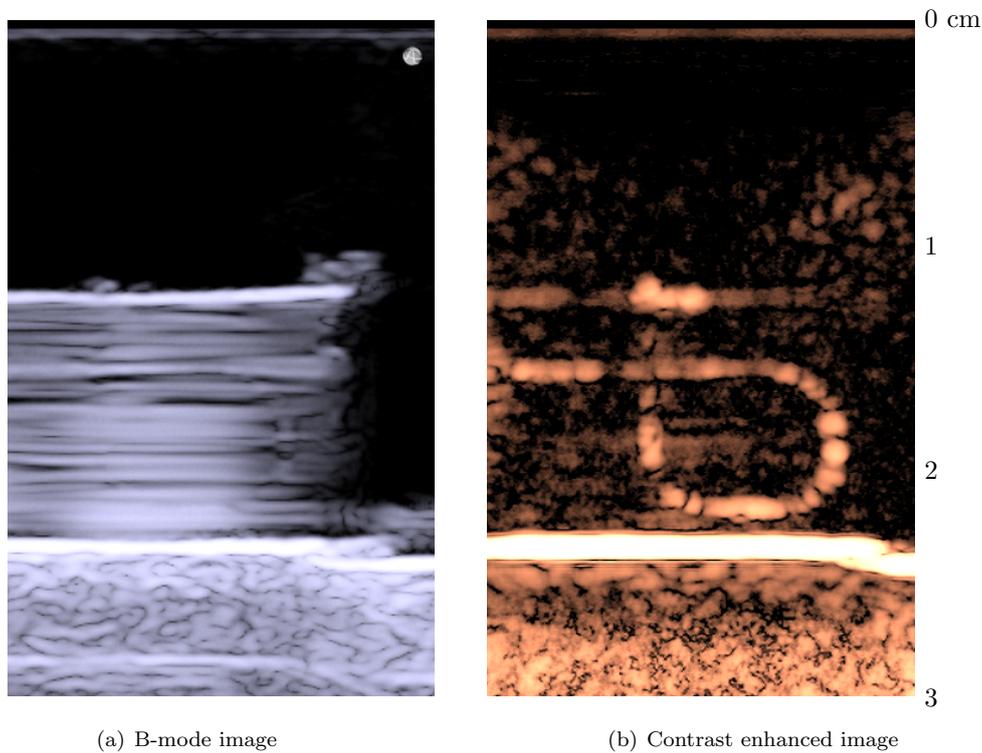


Figure 8.14: 2D ultrasound images of the looping flow phantom, with the channel system placed vertically. (a) shows a B-mode image of the phantom. The phantom orientation results in only the horizontal print layers being the only visible features. (b) shows a contrast enhanced image taken interleaved with the B-mode image.

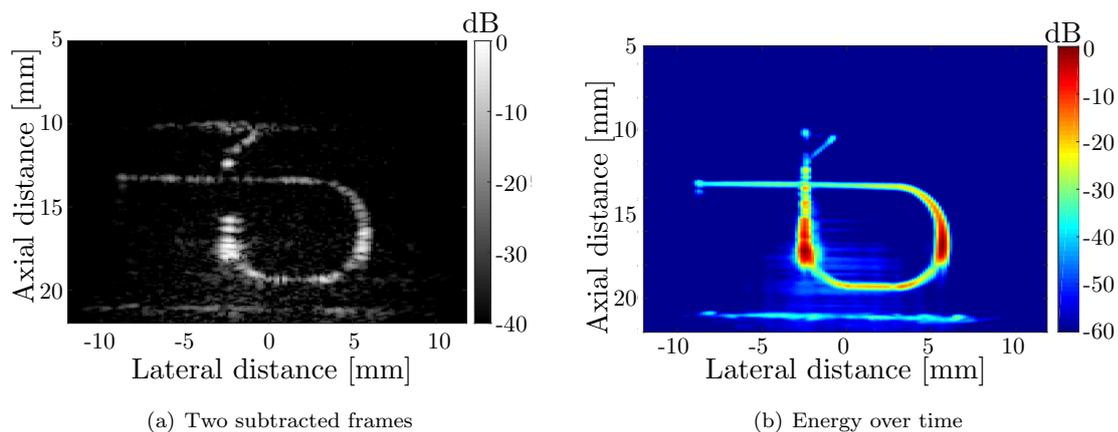


Figure 8.15: Resulting images after application of a stationary echo filter on the B-mode images in the series presented in Figure 8.14(a). (a) Subtraction of consecutive frames will result in only the features which have moved between frames being visible, in this case, the micro-bubbles in the flow channel. (b) shows the integrated energy over time of the filtered B-mode images. A tear near the outlet can be seen in which the channel branches into two segments.

of the subtracted frames. The effect might stem from the opposing forces acting on the micro-bubbles such as the fluid velocity, buoyancy, and the ultrasound pressure field pushing down the micro-bubbles. It is also notable that the vertical parts of the channels are printed in the x - y -plane, on an isotropic voxel grid, whereas the horizontal channel segments are printed on a anisotropic voxel grid. Thus, there might be an actual difference in channel dimensions, disturbing the flow.

8.6 Flow phantom optimisation

8.6.1 New fiducial marker layout

A practical issue with using 3D imaging probes, is that the data rate is so large, that there is typically no ultrasound live view, and therefore no easy feedback on alignment of the probe to the phantom. Instead, one will have to set everything up, make rough alignments, record a volume and beamform the data. Then the alignment can be checked, and altered if needed. To ease with alignment, the experimental setup includes 3D printed probe holders, which are all designed to mount the centreline of the ultrasound probe in the same position. In that way, if one probe is aligned all other probes will be aligned as well. Thus, if the fiducial markers are aligned along a plane, a 2D transducer can be used for live feedback on alignment to that plane. Any other probe exchanged with the used one will then be aligned similarly. This greatly simplifies alignment of 3D imaging probes. Inclusion of another set of fiducial markers rotated 90° around the centre vertical axis allows for aligning the probe along x and y . Figure 8.16 shows an updated version of the straight channel phantom, with offset fiducial markers, of the same overall pattern used in the original straight channel phantom, only now placed in the centre of the phantom and with two sets of orthogonal alignment marks. The new design was implemented to have the needle inlet further away from the central channel to be imaged, in this case the orthogonal channel section, potentially to move it completely out of the probe FOV. The needle itself is made of metal, which in general have high acoustic impedances, and will therefore provide large reflections in ultrasound imaging. These reflections could potentially influence the image of structures surrounding it.

Practical use of the original fiducial marker layout has shown that even at 1 mm of lateral separation of the fiducial marker columns, the elevation focus of the used 2D probes is still too wide for good unambiguous alignment to the fiducial marker structure. Updated versions will be separated by 1.5 mm.

8.6.2 Decreasing the flow channel size

The presented phantoms have been made with varying channel feature complexity, mainly to highlight different features of SRUS. However, they have all been fabricated with safe dimensions for the channel which were known to work without problems based on previous experiments. The fabrication method had not been characterised regarding the channel dimensions. It could be argued that this follows directly from the analysis of the scatterer sizes presented in Section 7.2.3, which was also one of the reasons why the circular cross-section was included in the design. However, observing a cavity appearing to be open does not mean that a channel of the same size will be perfusable. It might depend on the localisation of the channel, which was also illustrated by the small systematic variations seen in the printed scatterer size depending on the position of the scatterer within the print area. Furthermore, even a small print error at a single point along the channel might result in blocking of the channel. The chance of actually locating such a print error when investigating a channel by slicing the hydrogel for cross-sectional optical imaging is slim at best. Additionally there is significant risk of actually contaminating the channel directly when slicing it. The only way to test for perfusability is to perfuse the printed samples.

A preliminary investigation was made using the phantom design seen in Figure 8.17. Each channel starts with a $702\ \mu\text{m}$ (65 voxels) diameter needle inlet, which narrows to $205\ \mu\text{m}$ (19 voxels), before entering the main segment which is designed to be from $32.4\ \mu\text{m}$ to $108\ \mu\text{m}$ (3 to

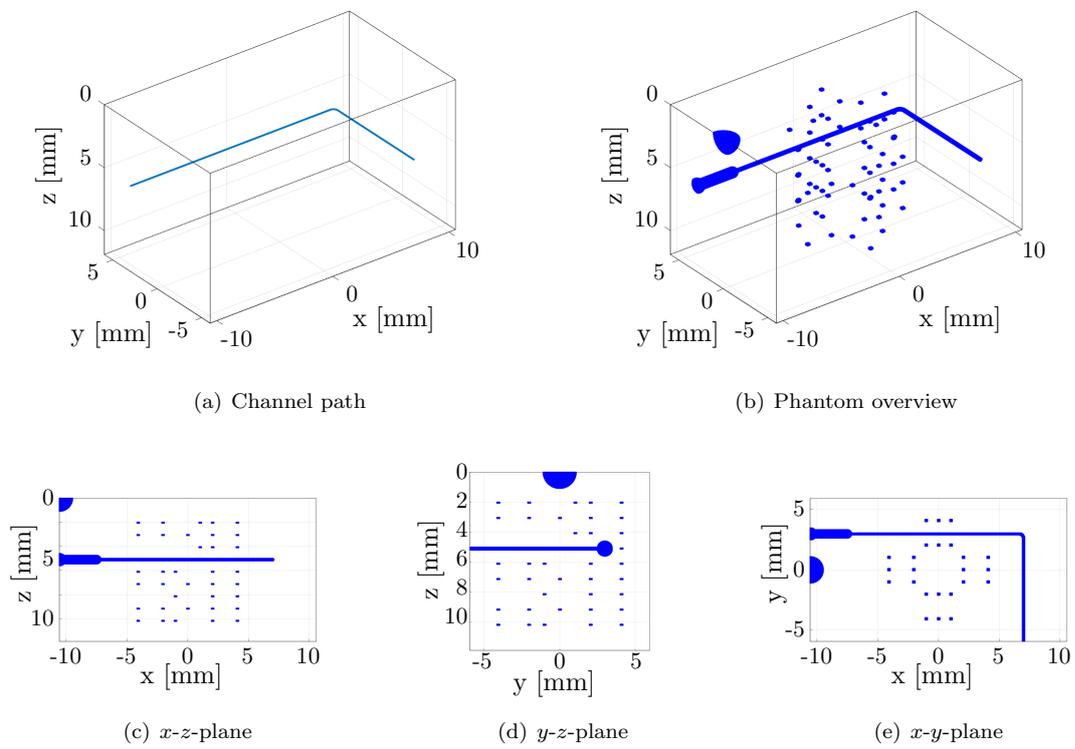


Figure 8.16: Updated version of the single channel phantom. In this case, the needle inlet is moved to the end of the phantom to remove it from the probe FOV. The fiducial markers are located near the centre of the phantom. Knowing the offset to the flow channels, proper alignment can easily be done.

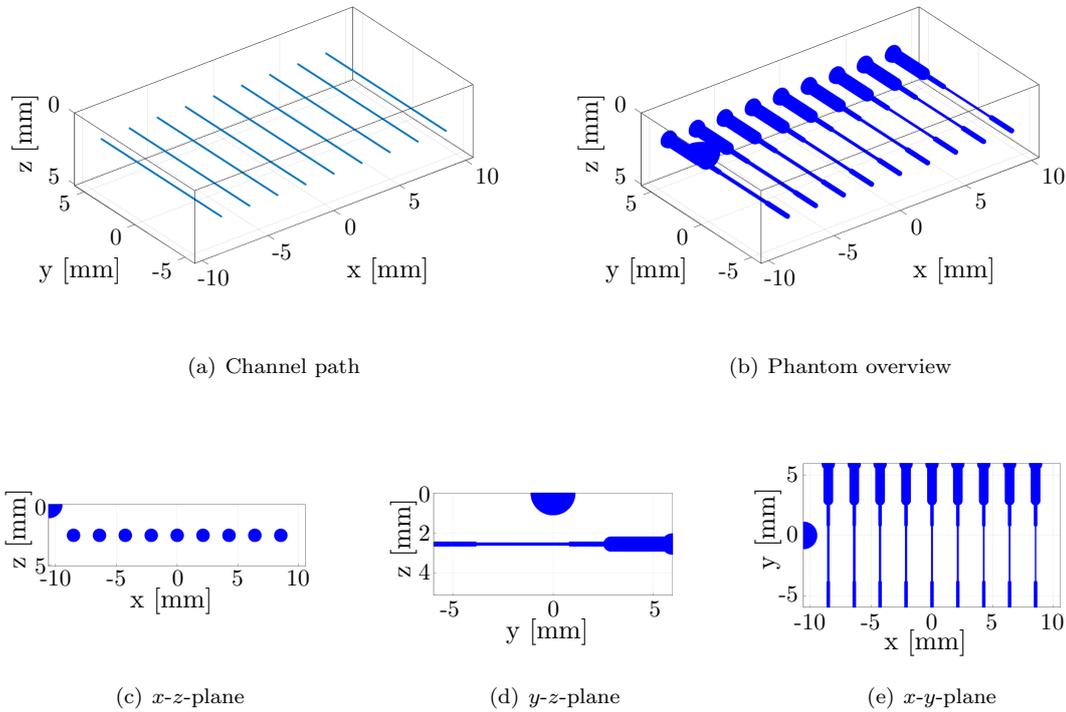


Figure 8.17: Perfusion test phantom. The channels are all cylindrical, and are designed to be from 3 to 10 voxels in diameter.

10 voxels), before going out into another 205 μm (19 voxels) segment. The phantoms were printed using a layer exposure time of 3 seconds.

An image of a phantom submerged in water during perfusion of the middle channel can be seen in Figure 8.18. The channels were perfused with water containing blue fruit dye for visual confirmation of perfusability. Due to the yellow colour of the phantom, and likely QY mixing with the perfusion solution, the dye appears green instead of blue. The wider 205 μm segments were included since the water with blue fruit dye was not visible to the naked eye in the smallest of the channels. By making a wider inlet before and after the narrow channel, the dye would be visible in these sections. Alternatively, perfusability could have been confirmed simply by observing the blue dye at the outlet of the channels. The channel perfused in the image was designed to be 86.4 μm in diameter. The blue dye can be seen in the water at the outlet of the perfused channel.

As the design was printed on the side, it was possible to fabricate two samples per print. Four samples were made in total in two prints. The results were inconclusive. In both of the two prints, one phantom was perfusable down to a design diameter of 86.4 μm , which according to the model presented in the scatterer size analysis will actually be printed as 80.8 μm . The other phantoms were only perfusable at the largest or next largest channels, 129.6 μm and 118.8 μm . Whether this difference stems from which side of the print area the phantom was printed on is unknown since the print side was not logged, but it could indeed be a consequence of imperfect printer uniformity. However, the optical investigation of the scatterer size in Section 7.2.3 did not suggest variations this large.

A more systematic study could be conducted in the future, taking into account multiple factors as in the scatterer size analysis, including the print position. Dosing schemes around the channels could also be applied, either through more controlled narrowing of channels by local increase of the dose, or by decreasing the risk of closing down the channels through lowering of the dose locally

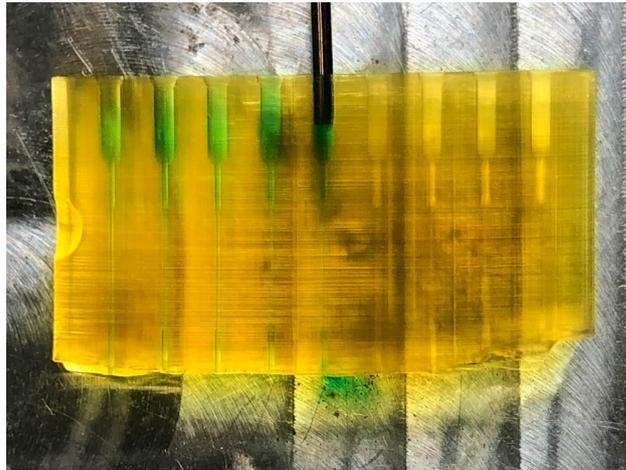


Figure 8.18: Image of perfusion test sample. The channels were perfused with water containing blue fruit dye for visual confirmation of perfusability. Due to the yellow colour of the phantom, and likely QY mixing with the perfusion solution, the dye appears green instead of blue. The channel perfused in the image was designed to be $86.4 \mu\text{m}$ in diameter.

around the channels. Thereby, the channel diameters might actually be controlled on a continuous scale, instead of being limited by the voxel dimensions.

8.7 Phantom concepts for future exploration

A number of additional phantom concepts have already been designed, but have not yet been utilized fully. These are designed both to determine the limit of the phantom fabrication method, but also to demonstrate resolution capabilities of SRUS algorithms in three dimensions.

8.7.1 Optimisation of channel separation

The minimum channel separation distance will set the limit for the SRUS resolution testing. The currently demonstrated $108 \mu\text{m}$ of separation is already well below the diffraction limit of most ultrasound systems, but at the same time not near the limit of SRUS algorithms. The following designs have been considered to test the printer capabilities, but have not yet been tested. It should be noted that in both cases the failure points likely will vary with channel diameter and volume flow rate.

Channels running in parallel

Figure 8.19 shows a design concept for testing how close flow channels can be printed when placed parallel to each other. A single channel bends multiple times passing parallel close by itself. Initial separation is set to $108 \mu\text{m}$, and decreased by 1 voxel for each bend. Note that only the separations between channels are scaled to the voxel grid in the illustration. The channel would need to be larger. Supposedly, when the separation between channel segments becomes too small at a certain pressure, the separation will fail. By using blue fruit dye again, it will be possible to determine the point of failure optically. The channel diameter should be kept large enough that the dyed water is visible to the naked eye. Unless debris at the failure ends up blocking the flow channel, the liquid can bypass the remaining loops to escape to the outlet. Thereby, the exact point of failure can be determined for parallel channels.

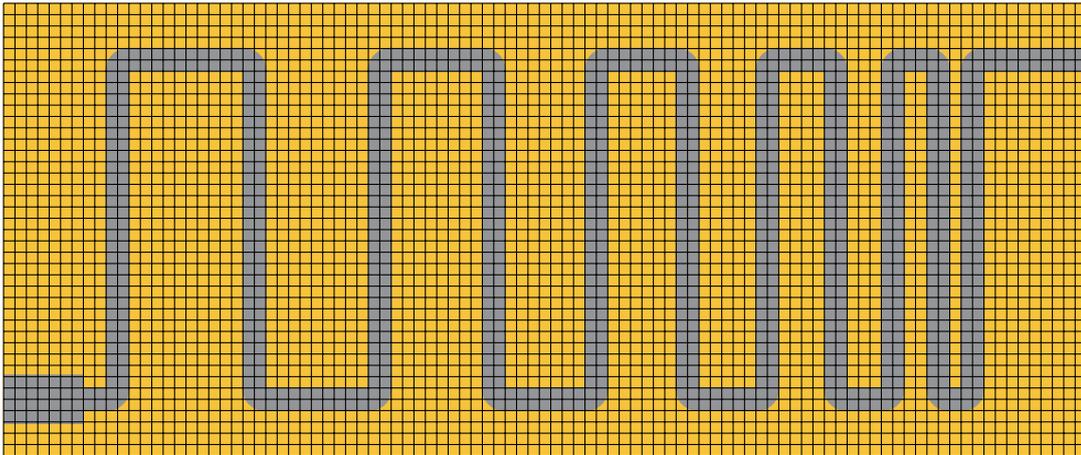


Figure 8.19: Parallel channels phantom concept. A single channel bends multiple times passing parallel close by itself. Initial separation is set to 108 μm , and decreased by 1 voxel for each bend. Note that only the separations between channels are scaled to the voxel grid. The channel would need to be larger.

Channels crossing perpendicularly

Figure 8.20 shows a design concept for testing how close flow channels can be printed when placed perpendicular to each other. A single channel is looped around and displaced vertically by initially only a single voxel of separation before passing over itself. The channel passes itself multiple times, each time increasing the separation by one voxel. The crossings are seen in (a), and the vertical channel separation is seen in (b). Given that the region in which the channels are close to each other is small compared to the parallel channel design concept, it is expected that the channel separation will fail at a smaller separation. The order of vertical separation is critical in this design. By placing the smallest separation furthest away from the inlet, the flow will only bypass smaller separations than that at the failure point. With the order reversed, the first crossing from the inlet would be the one with smallest separation. If that fails, flow would bypass all larger bends on the way to the outlet, making it impossible to observe the actual minimum separation distance. Once again, this assumes that debris at the point of failure does not end up blocking the channel.

8.7.2 Different flow velocities along different axes in a single phantom

Changes to the channel diameter will modify the hydraulic resistance and thereby the volume flow rate, and as a consequence the flow velocity. If only a single channel is used, the volume flow rate will necessarily be the same all the way through the channel system. By intentionally modifying the diameter of selected segments of the channel, the flow velocities in the different segments would be known, which could be used for velocity estimator validation. By having the channel bend in different directions, a single phantom with a single channel could be bent to be oriented in all directions at different positions in the phantom, with locally different well controlled flow velocities.

8.7.3 Branching channel systems to quantify local print variability

Once the flow estimator has been properly validated, phantom fabrication can start getting even more creative, utilizing the three-dimensional freedom of the printing method to the fullest. Flow phantoms can be designed with features mimicking actual vasculature. Designs could contain all of the elements previously avoided, for instance branching of channels oriented in many different directions. With the SRUS pipeline validated at this level, it could be utilised to illuminate the print variability of the channel phantoms.

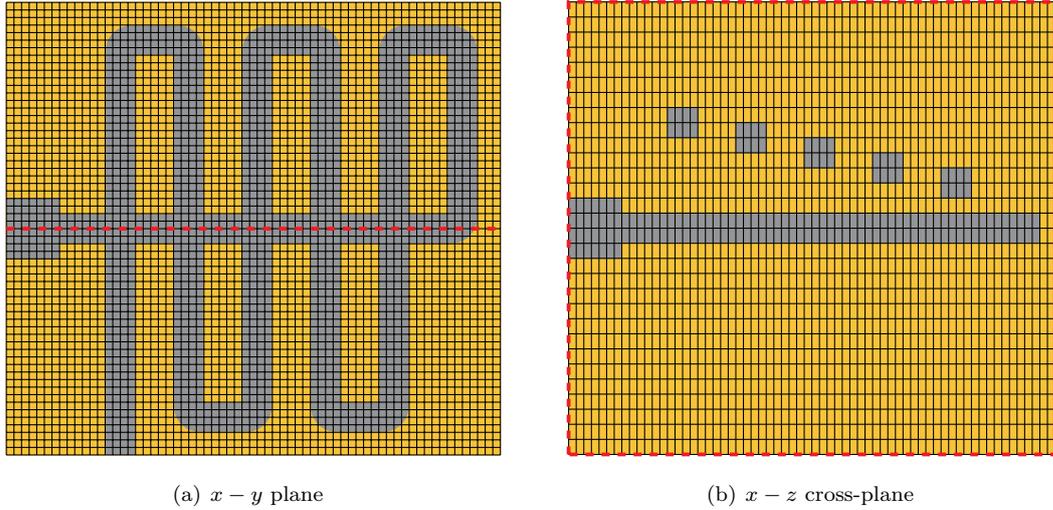


Figure 8.20: Channels crossing perpendicularly phantom concept. A single channel is looped around and elevated by initially only a single voxel of separation before passing over itself. The channel passes itself multiple times, each time increasing the separation by one voxel. The crossing is seen in (a), and the channel separation is seen in (b). Note that only the vertical separations between channels are scaled to the voxel grid. The channel would need to be larger.

8.8 Chapter summary

A series of flow phantom designs with 200 μm diameter channels have been presented, providing insight into the progressive development of the phantom and fiducial marker designs utilized. The design methods in MATLAB was described, demonstrating how to obtain most control possible over phantom dimensions and feature placement. A phantom containing two channels with a separation of 108 μm was presented to demonstrate SRUS in 2D. This sparked the development of fiducial marker patterns for precise alignment of ultrasound probes to the flow channels. A fiducial marker pattern was presented, which provides visual feedback on which corrections to alignment are needed. A single channel flow phantom was designed for demonstration of super-localisation of micro-bubbles using a 3D RCA probe. The micro-bubble localisations were used to determine SRUS precision estimates of 16.5 μm in the $y - z$ plane and 23 μm in the $x - z$ plane, both in line with the precision estimates determined using the scatterer phantom in Section 7.3.2. A looping flow phantom was designed to demonstrate 3D SRUS pipeline resolution capabilities. Initial experimentation utilizing this phantom design was presented, however, 3D SRUS experimentation has not yet been conducted. Initial experiments have shown that it is possible to 3D print perfusable channels as small as 80.8 μm .

Part III

Overall conclusion and outlook

The underlying goal of this Ph.D. project has been to develop tools to expand and improve the super-resolution ultrasound imaging (SRUS) techniques in the ultrasound research field, in part by transferring them from 2D to 3D imaging. The foundation has been that the improvements will be found both in software and hardware development, which in turn needs validation. This has led to the solutions presented in the two main parts of this thesis: capacitive micro-machined ultrasonic transducer (CMUT) process development, and 3D printed phantom fabrication.

CMUT process development

A critical processing step for fusion bonded CMUTs is the bonding step, which is highly vulnerable to particle contamination. The results presented in the thesis show that it is possible to obtain fusion bonds of comparable quality when conducting the pre-bonding process directly in hand, referred to as handbonding, instead of in a dedicated wafer bonder. A CMUT device structure using silicon nitride (Si_3N_4) plates were used to indirectly determine the resulting cavity pressure by measuring the resulting plate deflection under ambient conditions, as this is proportional to the cross-plate pressure difference. Handbonded devices were compared to devices bonded in a wafer bonder in vacuum, argon and air. No difference in plate deflection was found between the devices bonded in different atmospheres, indicating gases trapped in cavities formed through fusion bonding will diffuse out through the bonding interface during the 1100°C bond anneal step. Thereby, the gas content inside the cavity devices will depend on the annealing atmosphere: If the anneal is conducted in an N_2 atmosphere, the cavities will end up containing N_2 , regardless of what the initial gaseous content was. The only way to obtain a vacuum cavity is to anneal inside a vacuum.

The study has changed the bonding procedures of the research group, not only for fusion bonding, but also other types of wafer bonding. Due to the particle sensitivity of the bonding processes, a wafer cleaning process is always conducted as a final step prior to wafer bonding. A by-product of handbonding is that it is not necessary to transfer the wafers from the wafer cleaning station to the wafer bonder before bonding, with inevitable additional particle exposure. The handbond can be conducted right after wafer cleaning, minimising the risk of particle contamination. If controlled bonding parameters are necessary, the wafer stack can then subsequently be transferred to the wafer bonder, but now without the risk of additional particle contamination.

3D printed phantoms

In this thesis, ultrasound phantoms have been created using a stereolithography (SLA) 3D printing method. The method is a layer-by-layer printing process in which a liquid resin is poly-

merised into a water-containing polymer network called hydrogels, with acoustic parameters similar to those found in tissue. The printer system provides a voxel size of $10.8 \mu\text{m} \times 10.8 \mu\text{m} \times 20 \mu\text{m}$.

The properties of the phantoms have been characterised. Empirical results showed that changes in the acoustic properties could be attained by varying the interlayer exposure time of the printed structures locally within a phantom. However, characterisation of the hydrogel samples revealed a practically constant speed of sound of 1577 m/s for exposure times from 2 seconds to 19 seconds with a slight decrease to ≈ 1530 m/s for exposure times up to 23 seconds. The density was also found to be practically constant for most exposure times, with an average density of 1.045 g/ml. As a consequence, there should be no change in acoustic impedance, meaning there is no clear explanation for the empirical observation that the phantom properties do change. The phantom swells post printing, and the amount of swelling changes with the exposure time, from a little less than 2.6% for the most utilized exposure time of 3 seconds, to more than 6% for doses in the range of 9 seconds to 17 seconds. The swelling was shown to be isotropic. By using local exposures, this difference in swelling will result in stress effects locally in the phantom which might explain the change in acoustic impedance. However, this is not easily testable.

A new type of phantom was also developed for SRUS. One of the shortcomings of tube phantoms and micro-bubbles is that it is not possible to control the position of the micro-bubbles within the tubes. It is only possible to create an outer boundary. Thereby, the source of the detected signals are not completely known. It was shown that fixated scatterers could be printed by printing hollow cavities. By making them smaller than the imaging wavelength, they can be used as point targets. Statistical models were made to correlate the designed sizes of scatterers, ranging from 3 to 12 voxels, to the resulting printed scatterer size, also modelling the effect of using square or circular scatterer cross-sections, and printing them with different dose schemes with local overexposure around the scatterers. The statistical models provide not only an overview of these scatterer related factors, but also models inhomogeneities in the printer field of view (FOV) which can then be compensated. The model showed a significant size dependence on the dose schemes, with scatterers employing local overexposure around the cavity becoming smaller than those printed using only the base exposure time. A similar study was conducted on the reflected intensity, modelling all of the same factors, including compensation for inhomogeneities in the ultrasound imaging field. It was shown that the scatterers printed with overexposure in a single voxel wide frame around the cavity provided the largest intensities, and should thus be used for scatterer phantoms and fiducial markers.

A scatterer phantom containing eight randomly placed scatterers was used for determining the accuracy and the localisation precision of a 3D SRUS pipeline using a row-column addressed (RCA) array. With the scatterer coordinates compensated for the swelling, the SRUS pipeline could be validated. Good correlation between the distances between the scatterers based on the super-localised positions of the scatterers, and the corresponding design distances compensated for swelling was found. The slope of correlation was 0.989, close to a perfect correlation slope of 1. The true precision is expected to be between the two limiting estimates of $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.7 \mu\text{m}, 27.6 \mu\text{m}, 9.5 \mu\text{m})$ and $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.3 \mu\text{m}, 19.3 \mu\text{m}, 8.7 \mu\text{m})$, with the worst precision estimates being about 1/18th of the wavelength of $500 \mu\text{m}$ used in the experiment. The two sets of precision estimates stems from distortion in the beamforming, on a micrometre scale. This would not have been possible to discover using conventional tube phantom setups.

A series of flow phantoms were created to perform well controlled SRUS with micro-bubbles. The initial flow phantom results illustrated the fundamental need for fiducial markers to align the ultrasound probe to the phantom features. A fiducial marker layout was presented, which allow for easy alignment. The layout is based on a grid of scatterers for simpler identification of orientation. Furthermore, additional fiducial markers were included to provide visual feedback on how to correct alignment.

Another flow phantom was created to demonstrate super-localisation of micro-bubbles in 3D using a RCA array. An alternative method for estimation of the localisation precision was demonstrated, through consideration of the distributions of micro-bubbles. The localisation precision

estimates were $16.5\ \mu\text{m}$ in the $y - z$ plane and $23\ \mu\text{m}$ in the $x - z$ plane, both in line with the precision estimates determined using the scatterer phantom.

A looping flow phantom was designed to demonstrate 3D SRUS pipeline resolution capabilities. The design concept and intended use was described, including initial experiments utilizing the phantom, however, 3D SRUS experimentation has not yet been conducted.

The results illustrate the great obtainable achievements with a high resolution 3D printing phantom fabrication method, but only scratches the surface of the potential for future solutions which the phantom printing method provides. The printing method allows for three-dimensional freedom when designing phantoms, and an unparalleled control of phantom feature placement and feature size control.

9.1 Outlook

Many conclusive results and elements of transducer and phantom development were presented, showing off some of the capabilities of the technologies. While these results are directly usable, each of them opens up a series of new possibilities and areas to explore.

CMUT process development

The handbonding method has already been implemented as a standardised intermediate bonding step, conducted right after wafer cleaning. Although pressure is still applied to the pre-bonded structures for a more controlled pre-bond condition, this could potentially be exchanged with a weight placed on the wafer stack. This would allow for the transfer of the CMUT processing to a 6" wafer process, likely increasing the device throughput, which would provide the potential for increased characterisation of performance of the finished fusion bonded devices. A number of other processing steps still need transferring to a 6" wafer process.

3D printed phantoms

The 3D printing method has also proven a lot of use cases already. The implemented scatterer phantoms illustrate the intended concepts of SRUS pipeline characterisation using static and repeatable scatterer targets. However, a lot of modifications can be made to illuminate other SRUS features. Direct investigations into exactly how much separation between scatterers is needed for a properly functioning SRUS pipeline could be implemented. Sub-wavelength placement of scatterers might also open up for alternative SRUS techniques capable of separating the scatterers. In relation to the neural network sub-wavelength scatterer detection system, phantoms with scatterers placed much closer will be needed. These might be fixated in grids as used so far or in random patterns with a varying number of scatterers to better mimic the *in vivo* micro-bubble situation. The neural network scatterer detection has a high enough resolution that each scatterer is recognised with two scattering interfaces, one at the front, and one at the back of the scatterer. These might also be utilized as individual scattering targets, for increased flexibility in scatterer placement.

The scatterer sizes and signal intensities have only been characterised for 2D imaging, but should also be optimised for 3D imaging. To improve the signal to noise level (SNR) in the scatterer phantoms further, lowering of the background scattering intensity should also be investigated.

The scatterer phantoms illustrated high sensitivity to small distortions in the beamforming. It is likely a consequence of not perfectly matched speeds of sound between the phantom hydrogel and the water. Addition of salt to the water will increase the speed of sound towards that of the hydrogel. However, given that the hydrogel is a diffusion open polymer network, the salt water would also enter the hydrogel material, which is likely to change the hydrogel speed of sound as well. This, and alternative methods of matching of the speeds of sound should be tested.

A number of flow phantoms have been presented in the thesis, mainly using channel dimension known to provide perfusable channels. An initial study of decreasing the channel diameter was presented showing that it is possible to perfuse $80.8\ \mu\text{m}$ channels. The study should be done more

systematically, as the results showed high variability, potentially due to the channel placement within the printer FOV. Furthermore, modifications to the exposure schemes could result in even smaller channels being perfusable.

In addition to optimization of the channel diameter, the possible channel separation should also be tested, as this would be directly usable to demonstrate the resolution of the SRUS pipelines. However, at some point the channel separation might become so small that the flow channel separation fails due to the flow pressure. Two designs of phantoms were presented to investigate this, one for parallel flow channels and one for perpendicular flow channels.

The presented phantoms primarily investigated super-localisation of scatterers, as inexperience with using micro-bubbles resulted in practical flow velocities not suitable for velocity estimation. Optimisation of the flow velocities and micro-bubble concentrations could be combined with flow phantoms tailored for illustrating SRUS flow velocities. Phantoms with different diameter channel segments, oriented in many different directions could demonstrate simultaneous velocity estimation in full imaged volumes, which would not be possible utilizing any other phantom methods.

Finally, the 3D printing method fundamentally provides complete freedom in structure design. This should eventually be utilized for creating phantoms with channel systems mimicking real vasculature in terms of approximate scale and complexity.

Bibliography

- [1] M. K. Pugsley and R. Tabrizchi, "The vascular system: An overview of structure and function," *Journal of Pharmacological and Toxicological Methods*, vol. 44, no. 2, pp. 333–340, 2000.
- [2] H. C. H. Ko, B. K. Milthorpe, and C. D. McFarland, "Engineering thick tissues - the vascularisation problem," *Eur. Cells Mater.*, vol. 14, pp. 1–19, 2007.
- [3] G. K. Snyder and B. A. Sheafor, "Red blood cells: Centerpiece in the evolution of the vertebrate circulatory system," *American Zoologist*, vol. 39, 1999.
- [4] D. B. Stanimirovic and A. Friedman, "Pathophysiology of the neurovascular unit: Disease cause or consequence," *Journal of Cerebral Blood Flow and Metabolism*, vol. 32, no. 7, pp. 1207–1221, 2012.
- [5] B. V. Zlokovic, "Neurovascular pathways to neurodegeneration in alzheimer's disease and other disorders," *Nature Reviews Neuroscience*, vol. 12, no. 12, pp. 723–738, 2011.
- [6] J. Folkman, "Angiogenesis," *Annual Review of Medicine*, vol. 57, pp. 1–18, 2006.
- [7] D. Hanahan and R. A. Weinberg, "Hallmarks of cancer: The next generation," *Cell*, vol. 144, no. 5, pp. 646–674, 2011.
- [8] M. Narres, H. Claessen, S. Droste, T. Kvitkina, M. Koch, O. Kuss, and A. Icks, "The incidence of end-stage renal disease in the diabetic (compared to the non-diabetic) population: A systematic review," *Plos One*, p. e0147329, 2017.
- [9] S. A. Huettel, A. W. Song, and G. McCarthy, *Functional magnetic resonance imaging*, 2nd ed. Sinauer Associates, Inc, 2009.
- [10] C. Soeller and M. B. Cannell, "Examination of the transverse tubular system in living cardiac rat myocytes by 2-photon microscopy and digital image-processing techniques," *Circulation Research*, vol. 84, no. 3, pp. 266–275, 1999.
- [11] Y. Jia, O. Tan, J. Tokayer, B. Potsaid, Y. Wang, J. J. Liu, M. F. Kraus, H. Subhash, J. Fujimoto, J. Hornegger, and D. Huang, "Split-spectrum amplitude-decorrelation angiography with optical coherence tomography," *Optics Express*, vol. 20, no. 4, pp. 4710–4725, 2012.
- [12] D. S. Williams, J. A. Detre, and A. A. P. K. J. S. Leigh, "Magnetic resonance imaging of perfusion using spin inversion of arterial water," *Proc. Natl. Acad. Sci. USA*, 1992.

- [13] K. A. Miles, "Measurement of tissue perfusion by dynamic computed-tomography," *British Journal of Radiology*, vol. 64, no. 761, pp. 409–412, 1991.
- [14] R. Nakazato, D. S. Berman, E. Alexanderson, and P. Slomka, "Myocardial perfusion imaging with PET," *Imaging in Medicine*, vol. 5, no. 1, pp. 35–46, 2013.
- [15] D. Cosgrove and N. Lassau, "Imaging of perfusion using ultrasound," *Eur. J. Nucl. Med. Mol. Imaging*, pp. s65–s85, 2010.
- [16] G. Montaldo, M. Tanter, J. Bercoff, N. Benezek, and M. Fink, "Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 56, no. 3, pp. 489–506, March 2009.
- [17] F. Dong, J. Zhang, K. Wang, Z. Liu, J. Guo, and J. Zhang, "Cold plasma gas loaded microbubbles as a novel ultrasound contrast agent," *Nanoscale*, vol. 11, p. 1123–1130, 2019.
- [18] A. Caenen, M. Pemot, M. Peirlinck, L. Mertens, and P. Segers, "Analyzing the shear wave mechanics in cardiac shear wave elastography using finite element simulations," *Proceedings of 2018 IEEE International Ultrasonics Symposium (IUS)*, vol. 2018, p. 8579698, 2018.
- [19] L. V. Wang and S. Hu, "Photoacoustic tomography: In vivo imaging from organelles to organs," *Science*, vol. 335, no. 6075, pp. 1458–1462, 2012.
- [20] N. F. R. M. J. B. O. M. B. N. Simon Holbek, Kristoffer Lindskov Hansen and J. A. Jensen, "Real-time 2-d phased array vector flow imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 65, no. 7, pp. 1205–1213, 2018.
- [21] Z. Qiu, H. Fang, R. O'Leary, A. Gachagan, and A. Moldovan, "Broadband piezocrystal transducer array for non-destructive evaluation imaging applications," *Proceedings of 2018 IEEE International Ultrasonics Symposium (IUS)*, vol. 2018, p. 8579855, 2018.
- [22] J. E. Carlson, E. Andr n, M. Mohamad, and J. v. d. Beek, "Monitoring changes in mechanical properties of rock bolts using a low-power coded-excitation scheme," *Proceedings of 2018 IEEE International Ultrasonics Symposium (IUS)*, vol. 2018, p. 8580038, 2018.
- [23] M. Engholm, H. Bouzari, T. L. Christiansen, C. Beers, J. P. Bagge, L. N. Moesner, S. E. Diederichsen, M. B. Stuart, J. A. Jensen, and E. V. Thomsen, "Probe development of CMUT and PZT row-column-addressed 2-D arrays," *Sens. Actuators A: Phys.*, vol. 273, pp. 121–133, 2018.
- [24] T. L. Szabo, *Diagnostic ultrasound imaging: Inside out*, 2nd ed. Elsevier (Oxford, UK), 2014.
- [25] J. A. Jensen, *Estimation of Blood Velocities Using Ultrasound: A Signal Processing Approach*, 3rd ed. Department of Electrical Engineering, 2013.
- [26] S. J. Kleis and L. A. Sanchez, "Dependence of speed of sound on salinity and temperature in concentrated NaCl solutions," *Solar Energy*, vol. 45, no. 4, pp. 201–206, 1990.
- [27] L. M. Brekhovskikh, *Waves in Layered Media*. New York: Academic Press, 1960.
- [28] M. J. Haney and W. D. O'Brien, "Temperature dependency of ultrasonic propagation properties in biological materials," in *Tissue Characterization with Ultrasound*. Boca Raton, Fla.: CRC Press, 1986.
- [29] J. J. Wild, "The use of ultrasonic pulses for the measurement of biologic tissues and the detection of tissue density changes," *Surgery*, vol. 27, pp. 183–188, 1950.
- [30] D. H. Howry and W. R. Bliss, "Ultrasonic visualization of soft tissue structures of the body," *J. Lab. Clin. Med.*, vol. 40, pp. 579–592, 1952.

- [31] I. Edler and C. H. Hertz, "The use of ultrasonic reflectoscope for the continuous recording of the movement of heart walls," *Kungl. Fysiogr. Sällskap. i Lund Föhandl.*, vol. 24, pp. 40–58, 1954.
- [32] O. Ikeda, T. Sato, and K. Suzuki, "Super-resolution imaging system using waves with a limited frequency bandwidth," *J. Acoust. Soc. Am.*, vol. 65, no. 1, pp. 75–81, 1979.
- [33] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, "Imaging intracellular fluorescent proteins at nanometer resolution," *Science*, vol. 313, no. 5793, pp. 1642–1645, 2006.
- [34] S. T. Hess, T. P. K. Girirajan, and M. D. Mason, "Ultra-high resolution imaging by fluorescence photoactivation localization microscopy," *Biophysical Journal*, vol. 91, no. 11, pp. 4258–4272, 2006.
- [35] M. J. Rust, M. Bates, and X. Zhuang, "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)," *Nature methods*, vol. 3, no. 10, pp. 793–795, 2006.
- [36] O. Couture, B. Besson, G. Montaldo, M. Fink, and M. Tanter, "Microbubble ultrasound super-localization imaging (MUSLI)," in *Proc. IEEE Ultrason. Symp.*, 2011, pp. 1285–1287.
- [37] M. Schneider, "Characteristics of SonoVue," *Echocardiography*, vol. 16, no. s1, pp. 743–746, 1999.
- [38] D. H. Simpson, C. T. Chin, and P. N. Burns, "Pulse inversion Doppler: a new method for detecting nonlinear echoes from microbubble contrast agents," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 46, no. 2, pp. 372–382, 1999.
- [39] G. A. Brock-Fischer, "Means for increasing sensitivity in non-linear ultrasound imaging systems," *Journal of the Acoustical Society of America*, vol. 101, no. 6, p. 3240, 1997.
- [40] K. Christensen-Jeffries, J. Brown, P. Aljabar, M. Tang, C. Dunsby, and R. J. Eckersley, "3-D in vitro acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 64, no. 10, pp. 1478–1486, 2017.
- [41] A. Bar-Zion, C. Tremblay-Darveau, O. Solomon, D. Adam, and Y. C. Eldar, "Fast vascular ultrasound imaging with enhanced spatial resolution and background rejection," *IEEE Trans. Med. Imag.*, vol. 36, no. 1, pp. 169–180, 2016.
- [42] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, pp. 499–502, November 2015.
- [43] J. A. Jensen, S. B. Andersen, C. A. V. Hoyos, K. L. Hansen, C. M. Sørensen, and M. B. Nielsen, "Tissue motion estimation and correction in super resolution imaging," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [44] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433–440, February 2015.
- [45] A. Bar-Zion, O. Solomon, C. Tremblay-Darveau, D. Adam, and Y. C. Eldar, "Sushi: Sparsity-based ultrasound super-resolution hemodynamic imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 12, pp. 2365–2380, 2018.
- [46] R. J. G. van Sloun, O. Solomon, M. Bruce, Z. Z. Khaing, H. Wijkstra, Y. C. Eldar, and M. Mischi, "Super-resolution ultrasound localization microscopy through deep learning," [arXiv:1804.07661v2 \[eess.SP\]](https://arxiv.org/abs/1804.07661v2), 2018.

- [47] J. Youn, M. L. Ommen, M. B. Stuart, E. V. Thomsen, N. B. Larsen, and J. A. Jensen, "Ultrasound multiple point target detection and localization using deep learning," in *Proc. IEEE Ultrason. Symp.*, 2019.
- [48] Y. Desailly, J. Pierre, O. Couture, and M. Tanter, "Resolution limits of ultrafast ultrasound localization microscopy," *Phys. Med. Biol.*, vol. 60, no. 22, pp. 8723–8740, 2015.
- [49] K. Christensen-Jeffries, O. Couture, P. A. Dayton, Y. C. Eldar, K. Hynynen, F. Kiessling, M. O'Reilly, G. F. Pinton, G. Schmitz, M. X. Tang, M. Tanter, and R. J. G. van Sloun, "Super-resolution ultrasound imaging," *Ultrasound in Med. & Biol.*, 2020.
- [50] H. Bruus, *Theoretical Microfluidics*, 1st ed. Oxford University Press, 2008.
- [51] M. F. Rasmussen and J. A. Jensen, "3D ultrasound imaging performance of a row-column addressed 2D array transducer: a simulation study," in *Proc. SPIE Med. Imag.*, 2013, pp. 1–11, 86750C.
- [52] E. D. Light, R. E. Davidsen, J. Fiering, T. A. Hruschka, and S. Smith, "Progress in two-dimensional arrays for real-time volumetric imaging," *Ultrasound Imaging*, vol. 20, no. 1, pp. 1–15, 1998.
- [53] M. F. Rasmussen and J. A. Jensen, "3-D ultrasound imaging performance of a row-column addressed 2-D array transducer: A measurement study," in *Proc. IEEE Ultrason. Symp.*, July 2013, pp. 1460–1463.
- [54] J. W. Choe, O. Oralkan, and P. T. Khuri-Yakub, "Design optimization for a 2-D sparse transducer array for 3-D ultrasound imaging," *Proceedings of 2010 IEEE International Ultrasonics Symposium (IUS)*, vol. 2010, pp. 1928–1931, 2010.
- [55] B. Diarra, M. Robini, P. Tortoli, C. Cachard, and H. Liebgott, "Design of optimal 2-d nongrid sparse arrays for medical ultrasound," *IEEE Transactions on biomedical engineering*, vol. 60, no. 11, pp. 3093–3102, 2013.
- [56] A. Austeng and S. Holm, "Sparse 2-d arrays for 3-d phased array imaging - design methods," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 49, no. 8, pp. 1073–1086, 2002.
- [57] S. Harput, K. Christensen-Jeffries, J. Brown, J. Zhu, G. Zhang, C. H. Leow, M. Toulemonde, A. Ramalli, E. Boni, P. Tortoli, R. J. Eckersley, C. Dunsby, and M. Tang, "3-D super-resolution ultrasound (SR-US) imaging using a 2-D sparse array with high volumetric imaging rate," *arXiv preprint*, p. 1902.02631v1, 2019.
- [58] C. E. Morton and G. R. Lockwood, "Theoretical assessment of a crossed electrode 2-D array for 3-D imaging," in *Proc. IEEE Ultrason. Symp.*, 2003, pp. 968–971.
- [59] R. K. W. Chee, A. Sampaleanu, D. Rishi, and R. J. Zemp, "Top orthogonal to bottom electrode (TOBE) 2-D CMUT arrays for 3-D photoacoustic imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 8, pp. 1393–1395, 2014.
- [60] C. H. Seo and J. T. Yen, "64 x 64 2-D array transducer with row-column addressing," in *Proc. IEEE Ultrason. Symp.*, vol. 1, 2006, pp. 74–77.
- [61] H. Bouzari, M. Engholm, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Imaging performance for two row-column arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 7, pp. 1209–1221, 2019.
- [62] T. L. Christiansen, "Micromachined ultrasonic transducers for 3-D imaging," Ph.D. dissertation, Technical University of Denmark, 2015.

- [63] S. Holbek, T. L. Christiansen, M. B. Stuart, C. Beers, E. V. Thomsen, and J. A. Jensen, "3-D vector flow estimation with row-column addressed arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 11, pp. 1799–1814, 2016.
- [64] A. Cafarelli, A. Verbeni, A. Poliziani, P. Dario, A. Menciassi, and L. Ricotti, "Tuning acoustic and mechanical properties of materials for ultrasound phantoms and smart substrates for cell cultures," *Acta Biomaterialia*, vol. 49, 2017.
- [65] T.-Q. Nguyen, M. S. Traberg, J. B. Olesen, P. H. Møller-Sørensen, R. Moshavegh, L. Lönn, J. A. Jensen, M. B. Nielsen, and K. L. Hansen, "Pressure difference estimation in carotid bulbs using vector flow imaging - a phantom study," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [66] B. Y. S. Yiu and A. C. H. Yu, "Spiral flow phantom for ultrasound flow imaging experimentation," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 64, no. 12, pp. 1840–1848, 2017.
- [67] O. M. Viessmann, R. J. Eckersley, K. Christensen-Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, pp. 6447–6458, 2013.
- [68] Y. Desailly, O. Couture, M. Fink, and M. Tanter, "Sono-activated ultrasound localization microscopy," *Appl. Phys. Lett.*, vol. 103, no. 17, p. 174107, 2013.
- [69] C. Huang, M. R. Lowerison, F. Lucien, P. Gong, D. Wang, P. Song, and S. Chen, "Noninvasive contrast-free 3D evaluation of tumor angiogenesis with ultrasensitive ultrasound microvessel imaging," *Scientific Reports*, vol. 9, no. 1, p. 4907, 2019.
- [70] B. D. Lindsey, S. E. Shelton, K. H. Martin, K. A. Ozgun, J. D. Rojas, F. S. Foster, and P. A. Dayton, "High resolution ultrasound superharmonic perfusion imaging: In vivo feasibility and quantification of dynamic contrast-enhanced acoustic angiography," *Annals of Biomedical Engineering*, vol. 45, no. 4, pp. 939–948, 2017.
- [71] M. A. Schellpfeffer and G. L. Kolesari, "Microbubble contrast imaging of the cardiovascular system of the chick embryo," *Ultrasound in Med. & Biol.*, vol. 38, 2012.
- [72] T. Faez, I. Skachkov, M. Versluis, K. Kooiman, and N. de Jong, "In vivo characterization of ultrasound contrast agents: Microbubble spectroscopy in a chicken embryo," *Ultrasound in Med. & Biol.*, vol. 38, no. 9, pp. 1608–1617, 2012.
- [73] G. Liu, W. Qi, L. Yu, and Z. Chen, "Optimized doppler optical coherence tomography for choroidal capillary vasculature imaging," in *Proc. SPIE*, vol. 7889, 2011, pp. 1–10.
- [74] J. R. Jacquet, F. Ossant, F. Levassort, and J. M. Grégoire, "3-D-printed phantom fabricated by photopolymer jetting technology for high-frequency ultrasound imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 6, pp. 1048–1055, 2018.
- [75] R. Zhang, "Stereolithographic hydrogel printing of 3D culture chips with biofunctionalized complex 3d perfusion networks," *Lab Chip*, vol. 17, no. 24, pp. 4273–4282, 2017.
- [76] J. Jensen, M. B. Stuart, and J. A. Jensen, "Optimized plane wave imaging for fast and high-quality ultrasound imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 11, pp. 1922–1934, 2016.
- [77] P. Wilcox and J. Zhang, "Quantification of the effect of array element pitch on imaging performance," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2018.

- [78] M. I. Haller and B. T. Khuri-Yakub, "A surface micromachined electrostatic ultrasonic air transducer," in *Proc. IEEE Ultrason. Symp.*, vol. 2, 1994, pp. 1241–1244.
- [79] M. F. la Cour, T. L. Christiansen, J. A. Jensen, and E. V. Thomsen, "Modelling of CMUTs with anisotropic plates," in *Proc. IEEE Ultrason. Symp.*, 2012, pp. 588–591.
- [80] M. Engholm, T. Pedersen, and E. V. Thomsen, "Modeling of plates with multiple anisotropic layers and residual stress," *Sensors and Actuators A: Physical*, pp. 70–79, 2016.
- [81] S. E. Diederichsen, "Micromachined integrated transducers for ultrasound imaging," Ph.D. dissertation, The Technical University of Denmark, 2019.
- [82] A. Havreland, "Micromachined integrated 2d transducers for ultrasound imaging," Ph.D. dissertation, The Technical University of Denmark, 2019.
- [83] D. H. L. T. O. Ahrens, A. Buhrdorf and J. Binder, "Fabrication of gap-optimized CMUT," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 49, no. 9, pp. 1321–1329, 2002.
- [84] A. S. Ergun, Y. Huang, X. Zhuang, Ö. Oralkan, G. G. Yaralioglu, and B. T. Khuri-Yakub, "Capacitive micromachined ultrasonic transducers: Fabrication technology," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 52, no. 12, pp. 2242–2258, 2005.
- [85] E. Cianci, V. Foglietti, G. Caliano, and M. Pappalardo, "Micromachined capacitive ultrasonic transducers fabricated using silicon on insulator wafers," *Microelectron. Eng.*, vol. 61, pp. 1025–1029, 2002.
- [86] K. Brenner, A. S. Ergun, K. Firouzi, M. F. Rasmussen, Q. Stedman, and B. Khuri-Yakub, "Review: Advances in capacitive micromachined ultrasonic transducers," *Micromachines*, vol. 10, no. 152, pp. 1–27, Year.
- [87] R. B. Roy, O. Farhanieh, A. S. Ergün, and A. Bozkurt, "Fabrication of high-efficiency cmuts with reduced parasitics using embedded metallic layers," *IEEE Sensors Journal*, vol. 17, no. 13, p. 7929385, 2017.
- [88] A. Bakhtazad, R. Manwar, and S. Chowdhury, "Cavity formation in bonded silicon wafers using partially cured dry etch bisbenzocyclobutene (BCB)," in *Circuits and Systems (LAS-CAS)*, 2014, pp. 1–4.
- [89] A. S. Havreland, M. L. Ommen, C. Silvestre, M. Engholm, J. A. Jensen, and E. V. Thomsen, "Bcb polymer based row-column addressed cmut," in *Proc. IEEE Ultrason. Symp.*, 2017, pp. 1–4.
- [90] A. Bakhtazad, R. Manwar, and S. Chowdhury, "Fabrication and characterization of sealed cavities realized by adhesive wafer bonding with dry etched cyclotene," *Microsyst Technol.*, vol. 21, p. 2435–2442, 2015.
- [91] X. Zhang, O. Adelegan, F. Y. Yamaner, and O. Oralkan, "An optically transparent capacitive micromachined ultrasonic transducer (cmut) fabricated using su-8 or bcb adhesive wafer," in *Proc. IEEE Ultrason. Symp.*, 2017, pp. 1–4.
- [92] R. Manwar, T. Simpson, A. Bakhtazad, and S. Chowdhury, "Fabrication and characterization of a high frequency and high coupling coefficient cmut array," *Microsyst. Technol.*, vol. 23, no. 10, pp. 4965–4977, 2017.
- [93] Z. Li, L. L. P. Wong, A. I. H. Chen, S. Na, J. Sun, and J. T. W. Yeow, "Fabrication of capacitive micromachined ultrasonic transducers based on adhesive wafer bonding technique," *J. Micromech. Microeng.*, vol. 26, no. 11, p. 115019, 2016.

- [94] M. Engholm, H. Bouzari, J. A. Jensen, and E. V. Thomsen, "Capacitive substrate coupling of row-column-addressed 2-d cmut arrays," in *Proc. IEEE Ultrason. Symp.*, 2016, pp. 1–4.
- [95] J. Song, S. Jung, Y. Kim, K. Cho, B. Kim, and S. Lee, "Reconfigurable 2d cmut-asic arrays for 3d ultrasound image," *Progress in Biomedical Optics and Imaging - Proceedings of SPIE*, vol. 8320, no. 1, p. 83201A, 2012.
- [96] P. L. Muntal, S. E. Diederichsen, I. H. H. Jørgensen, J. A. Jensen, and E. V. Thomsen, "Transmitting performance evaluation of ASICs for CMUT-based portable ultrasound scanners," *Proceedings of 2017 IEEE International Ultrasonics Symposium (IUS)*, vol. 2017, pp. 1–4, 2017.
- [97] Abrisa Technologies, *Why Use Borofloat® Borosilicate Glass?*, 2020.
- [98] G. Wallis and D. I. Pomerantz, "Field assisted glass-metal sealing," *J. Appl. Phys.*, vol. 40, no. 10, pp. 3946–3949, 1969.
- [99] K. M. Knowles and A. T. J. V. Helvoort, "Anodic bonding," *International Materials Reviews*, vol. 51, no. 5, p. 273–311, 2006.
- [100] S. E. Diederichsen, J. M. F. Hansen, M. Engholm, J. A. Jensen, and E. V. Thomsen, "Output pressure and pulse-echo characteristics of cmuts as function of plate dimensions," in *Proc. IEEE Ultrason. Symp.*, 2017, pp. 1–4.
- [101] M. H. J. T. L. A. E. V. T. Mathias J.G. Mølgaard*, Milan Laustsen, "Combined colorimetric and gravimetric cmut sensor for detection of benzyl methyl ketone," *Sensors and Actuators B: Chemical*, vol. 275, pp. 483–489, 2018.
- [102] M. A. Schmidt, "Wafer-to-wafer bonding for microstructure formation," *Proceedings of the IEEE*, vol. 86, no. 8, pp. 1575–1585, 1998.
- [103] J. Vermeer, "The impulse breakdown strength of pyrex glass," *Physica*, vol. 20, no. 4, pp. 313–326, 1954.
- [104] E. V. Thomsen, K. Steenberg, M. G. Petersen, M. A. Weile, A. Havreland, M. L. Ommen, R. S. Grass, and M. Engholm, "Wafer level characterization of row-column addressed CMUT arrays," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [105] G. Fragiaco, T. Pedersen, O. Hansen, and E. V. Thomsen, "Intrinsic low hysteresis touch mode capacitive pressure sensor," *IEEE Sens. J.*, vol. 2010, pp. 2279–2282, 2011.
- [106] Y. Huang, A. S. Ergun, E. Hægström, M. H. Badi, and B. T. Khuri-Yakub, "Fabricating capacitive micromachined ultrasonic transducers with wafer-bonding technology," *J. Microelectromech. Syst.*, vol. 12, no. 2, pp. 128–137, 2003.
- [107] J. B. Lasky, S. R. Stiffler, F. R. White, and J. R. Abernathey, "Silicon-on-insulator (SOI) by bonding and etch-back," *International Electron Devices Meeting. Technical Digest*, vol. 85, pp. 684–687, 1985.
- [108] M. Shimbo, K. Furukawa, K. Fukuda, and K. Tanzawa, "Silicon-to-silicon direct bonding method," *Journal of Applied Physics*, vol. 60, no. 8, pp. 2987–2989, 1986.
- [109] W. P. Maszara, "Silicon-on-insulator by wafer bonding: a review," *J. Electrochem. Soc.*, vol. 138, no. 1, pp. 341–347, 1991.
- [110] K. K. Park, H. J. Lee, M. Kupnik, and B. T. Khuri-Yakub, "Fabrication of capacitive micromachined ultrasonic transducers via local oxidation and direct wafer bonding," *J. Microelectromech. S.*, vol. 20, no. 1, pp. 95–103, 2011.

- [111] M. Beggans, K. Farmer, J. Federici, T. G. J. Digges, S. Garofalini, and D. Hensley, "Bondability and surface roughness of ultra-thin single crystal silicon wafers," *Proceedings of the Fourth International Symposium on Semiconductor Wafer Bonding: Science, Technology, and Applications*, pp. 64–71, 1998.
- [112] W. Kern and D. A. Puotinen, "Cleaning solutions based on hydrogen peroxide for use in silicon semiconductor technology," *RCA Review*, vol. 31, no. 2, p. 187, 1970.
- [113] A. Plössl and G. Kräuter, "Wafer direct bonding: tailoring adhesion between brittle materials," *Materials Science and Engineering: R: Reports*, vol. 25, no. 1-2, pp. 1–88, 1999.
- [114] A. Berthold, B. Jakoby, and M. J. Vellekoop, "Wafer-to-wafer fusion bonding of oxidized silicon to silicon at low temperatures," *Sens. Actuator A Phys.*, vol. 68, no. 1, pp. 410–413, 1998.
- [115] K. Reck, C. Østergaard, E. V. Thomsen, and O. Hansen, "Fusion bonding of silicon nitride surfaces," *J. Micromech. Microeng.*, vol. 21, no. 12, pp. 1–5, 2011.
- [116] B. K. Ju, M. H. Oh, and K. H. Tehah, "Interfacial oxide growth and filling-up behaviour of the micro-gap in silicon fusion bonding processes," *Journal of materials science*, 1993.
- [117] C. Harendt, B. Hofflinger, H. Graf, and E. Penteker, "Silicon direct bonding for sensor applications: Characterization of the bond quality," *Sens. Actuator A Phys.*, vol. 25, no. 25, pp. 87–92, 1991.
- [118] S. Timoshenko and S. Woinowsky-Krieger, *Theory of Plates and Shells*, 2nd ed. McGraw-Hill College, jun 1959.
- [119] J. B. Lasky, "Wafer bonding for silicon-on-insulator technologies," *Appl. Phys. Lett.*, vol. 48, no. 1, pp. 78–80, 1986.
- [120] B. E. Deal and A. S. Grove, "General relationship for the thermal oxidation of silicon," *J. Appl. Phys.*, vol. 36, no. 12, pp. 3770–3778, 1965.
- [121] G. S. May and S. M. Sze, *Fundamentals of semiconductor fabrication*. John Wiley and Sons, Inc., 2004.
- [122] J. D. Plummer, M. D. Deal, and P. B. Griffin, *Silicon VLSI technology: Fundamentals, practice and modeling*. Prentice Hall, 2000.
- [123] S. I. Inc., *ATHENA user's manual*, SILVACO International Inc., Santa Clara, California, January 2016.
- [124] J. R. Taylor, *An introduction to error analysis: the study of uncertainties in physical measurements*. University Science Books, 1997.
- [125] K. Schjolberg-Henriksen, N. Malik, A. Sandvand, G. Kittilsland, and S. T. Moe, "Leak rates and residual gas pressure in cavities sealed by metal thermo-compression bonding and silicon direct bonding," *ECS Trans.*, vol. 64, no. 5, pp. 305–314, 2014.
- [126] D. T. Pham and R. S. Gault, "A comparison of rapid prototyping technologies," *International Journal of Machine Tools and Manufacture*, vol. 38, no. 10-11, pp. 1257–1287, 1998.
- [127] J. W. Stansbury and M. J. Idacavage, "3d printing with polymers: Challenges among expanding options and opportunities," *Dental Materials*, vol. 32, pp. 54–64, 2016.
- [128] F. P. W. Melchels, J. Feijen, and D. W. Grijpma, "A review on stereolithography and its applications in biomedical engineering," *Biomaterials*, vol. 31, pp. 6121–6130, 2010.

- [129] T. D. Ngo, A. Kashani, G. Imbalzano, K. T. Q. Nguyen, and D. Hui, "Additive manufacturing (3D printing): A review of materials, methods, applications and challenges," *Composites Part B*, vol. 143, pp. 172–196, 2018.
- [130] X. Wang, M. Jiang, Z. Zhou, J. Gou, and D. Hui, "3d printing of polymer matrix composites: A review and prospective," *Composites Part B*, vol. 110, pp. 442–458, 2017.
- [131] R. Zhang, "Stereolithographic hydrogel printing of 3D microfluidic cell culture chips," Ph.D. dissertation, The Technical University of Denmark, October 2017.
- [132] X. Zhang, O. J. Adelegan, F. Y. Yamaner, and O. Oralkan, "An optically transparent air-coupled capacitive micromachined ultrasonic transducer (CMUT) fabricated using adhesive bonding," in *Proc. IEEE Ultrason. Symp.*, 2017, pp. 1–4.
- [133] J. Jang, H. Yi, and D. Cho, "3d printed tissue models: Present and future," *Acs Biomaterials Science and Engineering*, vol. 2, no. 10, pp. 1722–1731, 2016.
- [134] R. Edmondson, J. J. Broglie, A. F. Adcock, and L. Yang, "Three-dimensional cell culture systems and their applications in drug discovery and cell-based biosensors," *Assay and Drug Development Technologies*, vol. 12, no. 4, pp. 207–218, 2014.
- [135] A. Hasan, A. Paul, A. Memic, and A. Khademhosseini, "A multilayered microfluidic blood vessel-like structure," *Biomed Microdevices*, vol. 17, no. 88, pp. 1–13, 2015.
- [136] B. N. Johnson, K. Z. Lancaster, I. B. Hogue, F. Meng, Y. L. Kong, L. W. Enquist, and M. C. McAlpine, "3d printed nervous system on a chip," *Lab chip*, vol. 16, no. 8, pp. 1393–1400, 2016.
- [137] C. Huang, P. Lee, P. Chen, and T. Liu, "Design and implementation of a smartphone-based portable ultrasound pulsed-wave Doppler device for blood flow measurement," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 59, no. 1, pp. 182–188, 2012.
- [138] W. H. Kim, J. M. Chang, C. Kim, J. Park, Y. Yoo, W. K. Moon, N. Cho, and B. Choi, "Synthetic aperture imaging in breast ultrasound: a preliminary clinical study," *Academic Radiology*, vol. 19, pp. 923–929, 2012.
- [139] H. Gong, B. P. Bickham, A. T. Woolley, and G. P. Nordin, "Custom 3D printer and resin for 18 μm \times 20 μm microfluidic flow channels," *Lab Chip*, vol. 17, no. 17, 2017.
- [140] R. M. Forbes, A. R. Cooper, and H. H. Mitchell, "The composition of the adult human body as determined by chemical analysis," *Journal of Biological Chemistry*, vol. 203, no. 1, pp. 359–366, 1953.
- [141] B. M. Tymrak, M. Kreiger, and J. M. Pearce, "Mechanical properties of components fabricated with open-source 3-d printers under realistic environmental conditions," *Materials and Design*, vol. 58, pp. 242–246, 2014.
- [142] A. Rodríguez-Panes, J. Claver, and A. M. Camacho, "The influence of manufacturing parameters on the mechanical behaviour of pla and abs pieces manufactured by fdm: A comparative analysis," *Materials*, vol. 11, no. 8, pp. 1–21, 2018.
- [143] M. B. Jackson, *Molecular and Cellular Biophysics*. Cambridge University Press, 2006.
- [144] A. Kuznetsova, P. B. Brockhoff, and R. H. B. Christensen, "lmerTest package: Tests in linear mixed effects models," *Journal of Statistical Software*, vol. 82, no. 13, pp. 1–26, 2017.
- [145] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2018. [Online]. Available: <https://www.R-project.org/>

- [146] R. A. Crescenti, J. C. Bamber, M. Partridge, N. L. Bush, and S. Webb, “Characterization of the ultrasonic attenuation coefficient and its frequency dependence in a polymer gel dosimeter,” *Phys. Med. Biol.*, vol. 52, pp. 6747–6759, 2007.
- [147] R. S. Davidsen, J. Ormstrup, M. L. Ommen, P. E. Larsen, M. S. Schmidt, A. Boisen, Ø. Nordseth, and O. Hansen, “Angle resolved characterization of nanostructured and conventionally textured silicon solar cells,” *Sol. Energy Mater. Sol. Cells*, vol. 140, pp. 134–140, 2015.
- [148] R. S. Davidsen, H. Li, A. To, X. Wang, A. Han, J. An, J. Colwell, C. Chan, A. Wenham, and M. S. Schmidt, “Black silicon laser-doped selective emitter solar cell with 18.1vol. 144,” pp. 740–747, 2016.
- [149] B. Iandolo, N. Sánchez, P. Adriana, R. S. Davidsen, and O. Hansen, “Black silicon with ultra-low surface recombination velocity fabricated by inductively coupled power plasma,” *Phys. Status Solidi RRL*, vol. 13, no. 2, p. 1800477, 2019.
- [150] M. L. Ommen, M. Schou, R. Zhang, C. A. V. Hoyos, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, “3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging,” in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [151] A. D. Edelstein, M. A. Tsuchida, N. Amodaj, H. Pinkard, R. D. Vale, and N. Stuurman, “Advanced methods of microscope control using µmanager software,” *J Biol Methods*, vol. 1, no. 10, pp. 1–10, 2014.
- [152] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J. Tinevez, D. J. White, V. Hartenstein, K. Eliceiri, P. Tomancak, and A. Cardona, “Fiji: an open-source platform for biological-image analysis,” *Nat. Methods*, vol. 9, no. 7, pp. 676–682, 2012.
- [153] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, “SARUS: A synthetic aperture real-time ultrasound system,” *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.
- [154] J. A. Jensen, M. L. Ommen, S. H. Øygaard, M. Schou, T. Sams, M. B. Stuart, C. Beers, E. V. Thomsen, N. B. Larsen, and B. G. Tomov, “Three-dimensional super resolution imaging using a row-column array,” *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, pp. 1–9, 2019.
- [155] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, “3-D imaging using row-column-addressed arrays with integrated apodization — Part I: Apodization design and line element beamforming,” *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 947–958, 2015.
- [156] M. B. Stuart, M. Schou, and J. A. Jensen, “Row-column beamforming with dynamic apodizations on a GPU,” in *Proc. SPIE Med. Imag.*, 2019, pp. 1–7, paper number 10955-20.
- [157] S. Harput, K. Christensen-Jeffries, A. Ramalli, J. Brown, J. Zhu, G. Zhang, C. H. Leow, M. Toulemonde, E. Boni, P. Tortoli, R. J. Eckersley, C. Dunsby, and M. Tang, “3-D super-resolution ultrasound (SR-US) imaging with a 2-D sparse array,” *arXiv preprint*, p. 1902.01608v1, 2019.

Part IV
Appendix

APPENDIX **A**

Published papers

A.1 Paper A - BCB polymer based row-column addressed CMUT

BCB polymer based row-column addressed CMUT

Andreas Spandet Havreland*, Martin Lind Ommen*, Chantal Silvestre*, Mathias Engholm*,
Jørgen Arendt Jensen† and Erik Vilain Thomsen*

* Department of Micro and Nanotechnology, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

† Center for Fast Ultrasound Imaging, Department of Electrical Engineering,
Technical University of Denmark, DK-2800 Lyngby, Denmark

Abstract—This paper presents an inexpensive, low temperature and rapid fabrication method for capacitive micromachined ultrasonic transducers (CMUT). The fabrication utilizes the bonding and dielectric properties of the photosensitive polymer Benzocyclobutene (BCB). A BCB based row-column addressed CMUT with integrated apodization has been fabricated and characterized with initial impedance measurement. Furthermore, two linear BCB CMUT arrays have been fabricated with different bottom electrode designs and characterized acoustically. All the fabricated arrays have a center frequency of 2.5 MHz when immersed into water and a pull-in voltage of 75 V. Stability tests have showed a stable coupling coefficient of approximately 0.1 during 10 hours of biased operation. Acoustic measurements, with a hydrophone positioned 1 cm from the CMUTs, have showed a peak-to-peak pressure of 14 kPa.

I. INTRODUCTION

Ultrasound transducers have for decades been based on piezoelectric technology, which today is a highly optimized technology. A promising alternative technology is the Capacitive Micromachined Ultrasonic Transducer (CMUT), which has the advantages of a wider bandwidth, and less acoustic impedance mismatch between transducer and tissue compared to a piezo transducer among others. The CMUT technology has furthermore higher design flexibility and allows for more closely packed active elements, thus, enabling the possibility for high frequency transducers. CMUTs are made by conventional MEMS (microelectromechanical systems) fabrication techniques. A CMUT transducer relies on having thousands of tightly sealed cavities, which can be obtained by either bonding or sacrificial release [1], [2]. If a bonding method is used the bonding process becomes one of the more critical steps in the fabrication process. Several methods have been demonstrated in the literature with different wafer-bonding techniques, where fusion/direct bonding and anodic bonding are among the more common methods [1], [3].

A third option is to use Benzocyclobutene (BCB) as a spacer and adhesive bonding material. CMUTs fabricated by adhesive wafer bonding using BCB have been demonstrated by [4]–[7]. BCB, from the Dow Chemical Company’s Cyclotene 4000 series, is a spin and spray coatable photosensitive polymer and at the same time being an excellent bonding material, with a reported mean bonding strength up to 35Jm^{-2} [8]. Furthermore, BCB exhibits a high chemical resistance when fully cured, which becomes an apparent advantage in the presented fabrication process. A BCB based CMUT fabrication has the advantage of being less sensitive to parti-

cles during the bonding process compared to fusion bonding. Hence, transducer arrays with large footprints such as a Row-Column Addressed (RCA) CMUT arrays can benefit from this fabrication technique. A RCA array is a 2-D transducer configuration for 3D ultrasound imaging [9]. Compared with conventional 2-D arrays RCA significantly reduces the required number of connections from N^2 to $2N$, with N being the number of channels.

In the presented CMUT design, BCB is used as a spacer defining the gap in the CMUT. However, BCB is inferior to the typical CMUT spacers SiO_2 and Si_3N_4 in terms of breakdown voltage, Young’s modulus, and thickness uniformity. High quality SiO_2 and Si_3N_4 thin films have a breakdown field of approximately 1Vnm^{-1} [10] whereas the breakdown field of BCB is 0.53Vnm^{-1} according to the data sheet. Hence, to make the breakdown voltage of BCB competitive towards SiO_2 new designs must be incorporated.

This paper addresses an inexpensive and fast BCB fabrication, which is a suitable rapid prototyping platform for different CMUT designs, and in the long term potentially a candidate for ultrasound imaging. This prototyping platform has, however, some challenges with hermetically sealing the cavities, less thin film uniformity, and softer clamping conditions. These effects lead to increased squeeze flow damping, variations in pull-in voltage across the wafer, and presumably more mechanical cross talk between CMUT cells.

II. FABRICATION AND ARRAY DESIGN

The BCB CMUT fabrication is entirely based on MEMS fabrication techniques. The fabrication is a three mask process schematically shown in Fig 1. A quartz wafer substrate is chosen in order to reduce the number of required masks by enabling backside alignment after the bonding process. In addition, the dielectric properties of the quartz substrate have the benefit of lowering substrate coupling [11].

Step 1 is a lithography step with a negative tone resist (AZ nLOF 2020) followed by 400 nm aluminum deposition and lift-off in Microposit™ remover 1165. CYCLOTENE resin 4022-25 BCB from Dow Chemical Company is, in step 2, spin coated on top of the structured bottom electrode followed by a 60°C bake for 90 s. The BCB is afterwards exposed (@ 365 nm) for 3.2 s with an intensity of 13mWcm^{-1} resulting in a dose of 41.6mJcm^{-2} . A post exposure bake at 60°C for 90 s is done subsequently. Development is carried out in two beakers using DS3000 from Dow Chemical Company.

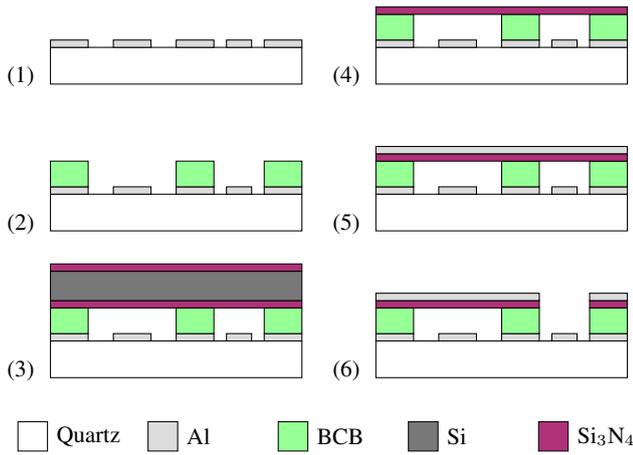


Fig. 1. BCB CMUT fabrication process. Step 1: Lithography, metallization and lift-off. Step 2: BCB lithography on top of patterned metal. Step 3: A double side polished wafer with silicon nitride is bonded to the BCB. Step 4: Top nitride layer and the silicon are etched away in a dry etch and potassium hydroxide (KOH) respectively. Step 5: Metal deposition. Step 6: Lithography followed by an etch through metal and Silicon. Figure is not to scale.

The development time and temperature in the first beaker are approximately 1 min at 30 °C, next in the second beaker for 2 min at room tempered DS3000 which stops the development. A final 90 °C bake for 1 min is performed after development resulting in an approximate final BCB thickness of 450 nm. Notice that the BCB is on top of aluminum pads. The purpose of these pads are to planarize the surface, and thereby achieve the best possible condition for the spin coating process. The planarization pads are not electrically connected in the final device. The plate of the CMUT is made of silicon nitride, and can be fabricated in a cleanroom with a Low Pressure Chemical Vapour Deposition (LPCVD) nitride furnace. The top-wafer consist of a 350 μm double side polished wafer with a 350 nm low stress LPCVD silicon nitride on both sides. To lower the stress-induced curvature across the wafer, the nitride is kept on both sides during bonding in step 3. The nitride top-wafer is bonded to the patterned BCB in a CNI v2.0 desktop nanoimprint tool from NIL Technology. During bonding the temperature is first ramped to 125 °C and kept for 15 min and subsequently raised to 240 °C and kept constant for 1 h. This step is a combined bonding and curing step.

A nitride plate serves two purposes. First, it lowers the fabrication cost compared to a SOI based method, and second the dielectric properties of the nitride increases the overall breakdown voltage and allows the plate to go into pull-in without short circuiting. The top nitride layer and the silicon are etched away in step 4. The nitride is removed in a fluorine plasma and the silicon is etched in potassium hydroxide (KOH). The silicon has a thickness of 350 μm and requires approximately 4.5 h of etching in 28wt.% KOH at 80 °C. The bonded wafers are etched without any backside or edge protection and the BCB will therefore be directly exposed to KOH at the edges. The transparent quartz substrate

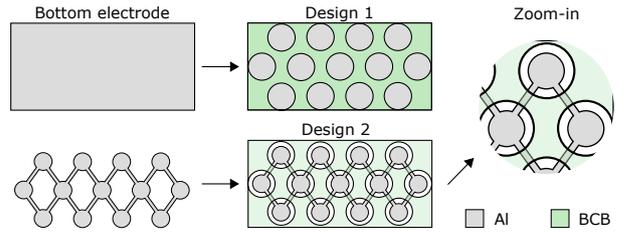


Fig. 2. A top view of two different array designs. Design 1 has a full bottom electrode. Design 2 has a structured bottom electrode and the overlapping region with BCB is significantly reduced (see the zoom-in figure). The planarization pads are for simplicity not included in this sketch.

enables visual inspection of potential damages of the CMUT structures, and it has by visual inspection been observed that KOH etches BCB at the edges with a rate of approximate 1 mm h⁻¹. Hence, a safety margin of 1 cm should be sufficient for this process.

In step 5, 400 nm aluminum is deposited on top of the remaining nitride device layer and followed by a lithography step, which later defines the top electrode. In step 6, aluminum is first etched in a wet solution of H₂O:H₃PO₄ in the volume ratio 1:2. Finally, the nitride is removed in a dry etch process to access the bottom electrode.

The highest temperature during fabrication is 240 °C, which makes this process CMOS compatible. The low process temperature allows for a metal bottom electrode, which can be structured. Two different bottom electrodes designs have been fabricated and are shown schematically in Fig 2. Design 1 has a full bottom electrode. It has the disadvantage of increased parasitic capacitance, however, this design reduces the required masks from three to two. Design 2, inspired by [4], has a structured bottom electrode, which serves two purposes. First, it provides a reduction in parasitic capacitance and second less BCB will be exposed to a high electric field. The BCB on top of the connecting wires must be able to withstand the direct electric field between top and bottom electrodes. However, as opposed to Design 1, it is only in a very small area where BCB should withstand the direct electric field. Hence, the probability of trapping particles or other breakdown lowering defects is reduced.

Often, CMUT transducers for medical imaging are designed to be operated with a DC bias voltage of approximately 200 V [12], [13]. The presented CMUT is a rapid prototype of a BCB CMUT, and the nitride plate is too thin with respect to the mask layout, resulting in a lowered pull-in voltage and resonant frequency. The pull-in voltage is measured to be approximately 75 V. The resonance frequency in air is 4 MHz and reduces to 2.5 MHz when immersed into water.

A BCB RCA CMUT with integrated apodization [14], [15] has been fabricated along with two linear arrays. The RCA CMUT array has a footprint of 2.6×2.6 cm², and design 2 bottom electrode configuration. The footprint of the two linear arrays are 5.4×25.4 mm², where one of them is design 1 and the other is design 2.

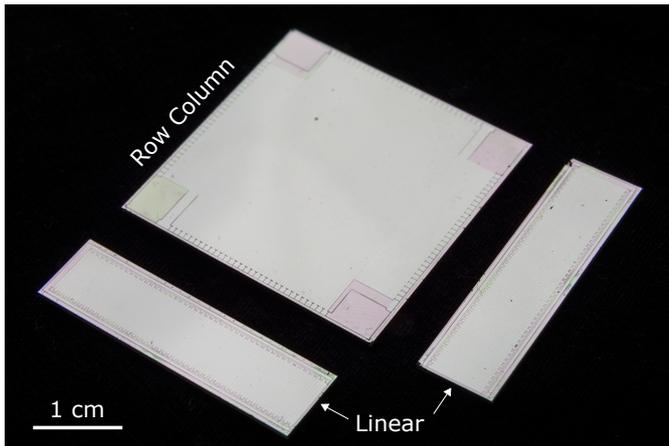


Fig. 3. Image of the presented BCB CMUTs. 62+62 BCB Row-Column addressed CMUT array with integrated apodization and two 92-element linear BCB CMUT arrays with different bottom electrode configurations.

TABLE I
CMUT ARRAY PARAMETERS

Parameter	Value	
	1D	2D
Array	1D	2D
Number of elements	94	62+62
Element width	270 μm	270 μm
Element length	5.4 mm	26 mm
Element pitch	270 μm	270 μm
Cell diameter	60 μm	
Electrode diameter (design 2)	40 μm	
Cell to cell distance	7.5 μm	
BCB cavity height	45 nm	
Nitride plate thickness	350 nm	

The final transducers are shown in Fig 3. The 1-D linear arrays consist of 94 elements each with 288 CMUT cells and the RCA array consists of 62 rows and 62 columns each having 1254 CMUT cells with apodization included. All array parameters can be found in Table I.

III. RESULTS

For this BCB fabrication technique to be of interest the performance has to be constant in time. Hence, operational stability is a key parameter, and a Row-Column element with integrated apodization has been stability tested by impedance measurements during 10 hours of biased operation. The measurements were performed using an Agilent 4294A Precision Impedance Analyzer with a varying DC voltage and an AC voltage of 50 mV. The stability results can be found in Fig 4. The figure contains four plots where a) is the applied voltage sequence. The voltage steps of 0 V, 65 V and -65 V are applied to the CMUT in a sequence such all transition combinations are present. 65 V corresponds to $\sim 85\%$ of the pull-in voltage. The capacitance and the coupling coefficient κ^2 are depicted in b) and c). The capacitance and coupling coefficient have the same qualitative behavior. An increasing tendency over time is observed for both biased polarities. The underlying reason for this increasing behavior is not yet understood. However, it is

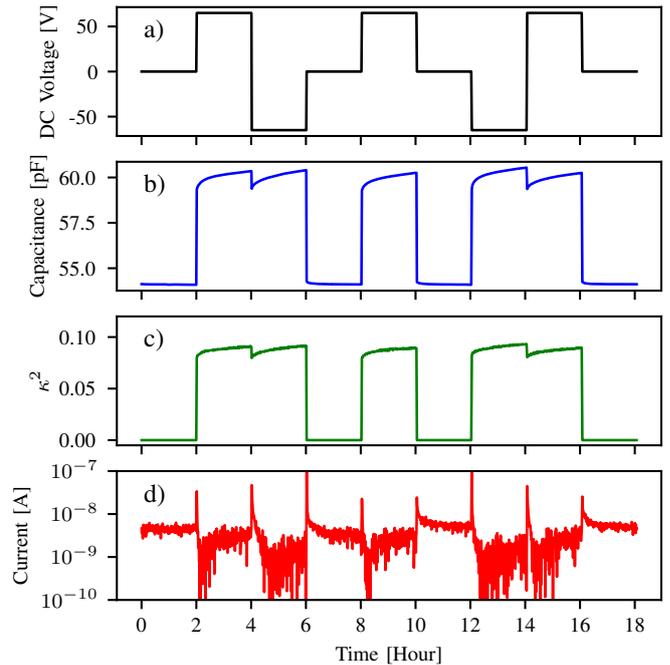


Fig. 4. Stability measurements during 10 hours of operation on a Row test element having a design 2 bottom electrode configuration. a) The applied DC voltage sequence. b) The measured capacitance. c) The fitted coupling coefficient. A slow increasing tendency during 2 hours of operation is observed in both plots. d) The measured current. Notice none of the measured parameters show any polarity depended characteristics.

not believed to be dielectric charging, since these parameters would then be expected to decay in time in one of the biased polarities or both. The coupling coefficient is estimated by fitting a lumped element model to the impedance spectra. The magnitude of the coupling coefficient is approximately 0.1 during operation, and is furthermore observed to be equal for both polarities.

To verify the absence of dielectric charging the leak current through the CMUT is shown in d). The current in a polymer depends on several parameters and has multiple regimes in its IV characteristic [16]. The measured current stabilizes in the nA range for both polarities, and does not show any signs of dielectric charging.

The two 1-D linear arrays have been acoustically characterized. Hydrophone measurements were performed in water, using a calibrated ONDA HGL-0400 hydrophone positioned 1 cm from the CMUT. Prior to the measurement, PDMS has been coated on top of the arrays to ensure electrical isolation of the elements during operation. A 2.5 MHz two-cycle sine pulse was used to excite the CMUTs with a 70 V DC bias and an AC peak-to-peak voltage of 10 V. The acoustic results are shown in Fig 5. The two upper plots show the hydrophone response and the lower plot shows the Fourier spectrum of the two signals. The hydrophone responses are an average of 30 measurements of the same element during a time period of 3 min. The averaged hydrophone responses show a small

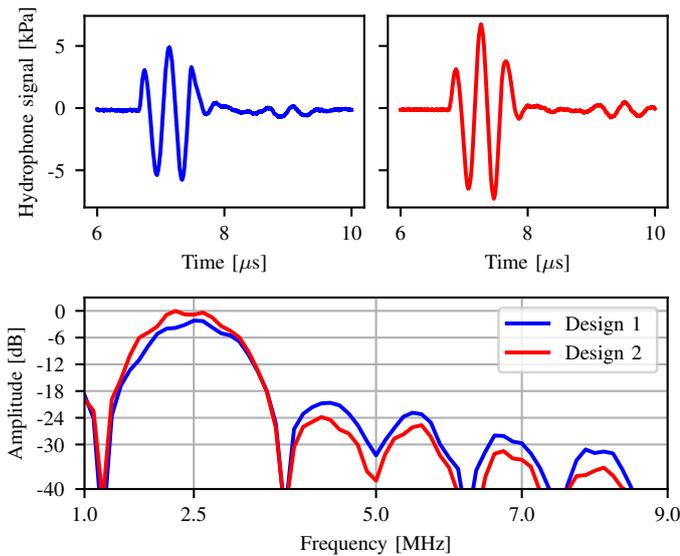


Fig. 5. Acoustic measurements of linear CMUT elements. Upper figures: Hydrophone response for the two array designs. To reduce noise 30 signals have been averaged. The smaller signal at $9\mu\text{s}$ corresponds to the reflection from the PDMS. Lower figure: Fourier spectrum for a two period sine pulse with a DC bias voltage of 70 V and a peak-to-peak AC voltage of 10 V.

amplitude increase for Design 2. The peak-to-peak pressure is measured to be 10 kPa and 14 kPa for design 1 and 2, respectively, which might suggest the amplitude is design depend. The current output pressure is too low and noisy to be used for medical imaging, and further optimization will be needed to increase the output pressure. In the Fourier spectrum it is seen that the side lobes from the higher harmonics starts in the range between -18 dB and -24 dB, thus, not nearly as good as other CMUT in the literature, where a -40 dB suppression of the second and third harmonic can be found [17]. In conclusion structuring the bottom electrode does not only reduce the parasitic capacitance and thereby increase the sensitivity in receive, but it might also have the advantage of an increased peak-to-peak pressure. However, despite improvements in the design, more optimization is needed to make BCB CMUTs competitive towards other CMUT fabrication techniques and ultimately the piezoelectric technology.

IV. CONCLUSION

An inexpensive and rapid BCB fabrication process for 1-D and 2-D CMUT array is proposed. Functional 2.5 MHz BCB CMUTs have been fabricated and demonstrated. Two linear arrays with different bottom electrode configuration have been characterized one with a uniform bottom electrode and another with a structured bottom electrode. The presented BCB CMUTs have been stability tested during 10 hours of operation, and a coupling coefficient κ^2 of approximately 0.1 has been observed. BCB based 1-D linear elements have been acoustically characterized with a 2-cycle sine pulse, where the peak-to-peak pressure, at 1 cm, was measured to be 14 kPa.

V. ACKNOWLEDGEMENT

We would like to thank the Danish Innovation Fund and BK Ultrasound for funding this research, and the Otto Mønsted fund for financial support of travel expenses.

REFERENCES

- [1] Y. Huang, A. Sanli Ergun, E. Hægström, M. H. Badi, and B. T. Khuri-Yakub, "Fabricating capacitive micromachined ultrasonic transducers with wafer-bonding technology," *Journal of Microelectromechanical Systems*, vol. 12, no. 2, pp. 128–137, 2003.
- [2] A. S. Ergun, G. G. Yaralioglu, and B. T. Khuri-Yakub, "Capacitive micromachined ultrasonic transducers: Theory and technology," *Journal of Aerospace Engineering*, vol. 16, pp. 76–84, 2003.
- [3] S. Olcum, K. Ouz, M. N. enlik, F. Y. Yamaner, A. Bozkurt, A. Atalar, and H. Köymen, "Wafer bonded capacitive micromachined underwater transducers," *Proceedings - IEEE Ultrasonics Symposium*, p. 5441699, 2009.
- [4] Z. Li, A. I. Chen, L. L. Wong, S. Na, and J. T. Yeow, "Fabrication of polymer-based wafer-bonded capacitive micromachined ultrasonic transducers," in *2015 IEEE International Ultrasonics Symposium, IUS 2015*, 2015.
- [5] Z. Li, L. L. P. Wong, A. I. H. Chen, S. Na, J. Sun, and J. T. W. Yeow, "Fabrication of capacitive micromachined ultrasonic transducers based on adhesive wafer bonding technique," *Journal of Micromechanics and Microengineering*, vol. 26, no. 11, p. 115019, 2016.
- [6] R. Manwar, T. Simpson, A. Bakhtazad, and S. Chowdhury, "Fabrication and characterization of a high frequency and high coupling coefficient CMUT array," *Microsystem Technologies*, pp. 1–13, 2016.
- [7] R. Manwar and S. Chowdhury, "Experimental analysis of bisbenzocyclobutene bonded capacitive micromachined ultrasonic transducers," *Sensors*, vol. 16, no. 7, p. 959, 2016.
- [8] F. Forsberg, F. Saharil, T. Haraldsson, N. Roxhed, G. Stemme, W. van der Wijngaart, and F. Niklaus, "A comparative study of the bonding energy in adhesive wafer bonding," *Journal of Micromechanics and Microengineering*, vol. 23, no. 8, p. 85019, 2013.
- [9] C. C. E. Morton and G. R. G. Lockwood, "Theoretical assessment of a crossed electrode 2-D array for 3-D imaging," *IEEE Symposium on Ultrasonics*, 2003, vol. 1, no. c, pp. 968–971, 2003.
- [10] C. M. Osburn and E. J. Weitzman, "Electrical conduction and dielectric breakdown in silicon dioxide films on silicon," *Journal of The Electrochemical Society*, vol. 119, no. 5, p. 603, 1972.
- [11] M. Engholm, H. Bouzari, J. A. Jensen, and E. V. Thomsen, "Capacitive substrate coupling of row-column-addressed 2-D CMUT arrays," *Proceedings of 2016 IEEE International Ultrasonics Symposium*, vol. 2016-, p. 7728384, 2016.
- [12] M. Engholm, T. L. Christiansen, C. Beers, J. P. Bagge, L. N. Moesner, H. Bouzari, A. Lei, M. Berkheimer, M. B. Stuart, J. A. Jensen, and E. V. Thomsen, "A hand-held row-column addressed CMUT probe with integrated electronics for volumetric imaging," *Proceedings of 2015 IEEE International Ultrasonics Symposium*, p. 7329132, 2015.
- [13] A. S. Savoia, G. Caliano, and M. Pappalardo, "A CMUT probe for medical ultrasonography: From microfabrication to system integration," *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, vol. 59, no. 6, pp. 1127–1138, 2012.
- [14] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, "3-D imaging using row-column-addressed arrays with integrated apodization. Part I: Apodization design and line element beamforming," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 62, no. 5, pp. 947–958, 2015.
- [15] T. L. Christiansen, M. F. Rasmussen, J. P. Bagge, L. N. Moesner, J. A. Jensen, and E. V. Thomsen, "3-D imaging using row-column-addressed arrays with integrated apodization. Part II: Transducer fabrication and experimental results," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 62, no. 5, pp. 959–971, 2015.
- [16] L. A. Dissado and J. C. Fothergill, *Electrical degradation and breakdown in polymers*. London: Peter Peregrinus, 1992.
- [17] A. Lei, S. E. Diederichsen, S. M. Hansen, M. B. Stuart, J. P. Bagge, J. A. Jensen, and E. V. Thomsen, "Elimination of second-harmonics in CMUTs using square pulse excitation," *Proceedings of 2016 IEEE International Ultrasonics Symposium*, vol. 2016-, p. 7728824, 2016.

A.2 Paper B - 3D Printed Flow Phantoms With Fiducial Markers for Super-Resolution Ultrasound Imaging

3D Printed Flow Phantoms with Fiducial Markers for Super-Resolution Ultrasound Imaging

*Martin Lind Ommen*¹, *Mikkel Schou*², *Rujing Zhang*¹, *Carlos Armando Villagómez Hoyos*²,
*Jørgen Arendt Jensen*², *Niels Bent Larsen*¹, and *Erik Vilain Thomsen*¹

¹Department of Micro and Nanotechnology, Technical University of Denmark, Kgs. Lyngby, Denmark

²Center for Fast Ultrasound Imaging, Technical University of Denmark, Kgs. Lyngby, Denmark

Abstract—The improved resolution provided by ultrasound super-resolution imaging (SRI) sets new demands on the fabrication of phantoms for the validation and verification of the technique. Phantoms should resemble tissue and replicate the 3D nature of tissue vasculature at the microvascular scale. This paper presents a potential method for creating complex 3D phantoms, via 3D printing of water-filled polymer networks. By using a custom-built stereolithographic printer, projected light of the desired patterns converts an aqueous poly(ethylene glycol) diacrylate (PEGDA) solution into a hydrogel, a material capable of containing 75 wt% of water. Due to the hydrogel mainly consisting of water, it will, from an acoustical point of view, respond very similar to tissue. A method for printing cavities as small as $(100\ \mu\text{m})^3$ is demonstrated, and a 3D printed flow phantom containing channels with cross sections of $(200\ \mu\text{m})^2$ is presented. The designed structures are geometrically manufactured with a 2% increase in dimensions. The potential for further reduction of the flow phantom channels size, makes 3D printing a promising method for obtaining microvascular-like structures.

Index Terms—3D printing, stereolithography, phantom, hydrogel, microvasculature, resolution, ultrasound

I. INTRODUCTION

Living tissue continuously adapts to changes in external stresses, or internal requirements. A local increase in cell activity will require an increase in nutrients and oxygen to the local tissue. One reason for increased cell activity could be cell proliferation due to cancer. The increased demands are met through increased vascularisation of the tissue [1]. This correlation between cancer and vascularisation can be used in the diagnosis of cancer patients. Once carcinogenic tissue has been detected, it will be possible to closely follow the vasculature response to any cancer treatments, with the potential to react fast to treatment responses.

Ultrasound super-resolution imaging (SRI) has within the last few years been introduced as a non-invasive method for imaging of the vasculature [2], [3]. The characteristic signal from contrast bubbles can be used to form images of higher resolution than would be possible due to the diffraction limit of regular ultrasound methods. To verify and optimise imaging- and tracking algorithms, it is necessary to use tissue mimicking phantoms, which not only recapitulate the acoustic properties of tissue, but also the scale and dimensionality of interest. The latter two points have proven difficult to engineer. At the introduction of the ultrasound SRI field, a few articles were published, demonstrating the method principles using measurements of phantoms. In one case, the channel

dimensions were decreased to $40 \times 80\ \mu\text{m}^2$ [4] by defining the channel using UV lithography. However, in practice this imposes a limit of the channels being in a single plane. Other examples of phantoms consisting of tubes have been presented in [5] (3 mm diameter) and [6] (200 μm diameter). In both cases this is larger than the vessels of interest, which are less than 20 μm in diameter.

3D printing of polymers is a promising new technique, which allows the flexibility of fabricating complex 3D structures, as well as printing of very small features in the sub-100 μm range. Alignment of the ultrasound probe to the phantom becomes increasingly difficult as phantom features become smaller. Inclusion of fiducial markers within the phantoms can greatly reduce this problem.

We present a method for 3D printing of phantoms, by stereolithography. Our hypothesis is that by 3D printing phantoms using stereolithography, it is possible to create geometrically stable 3D phantoms to a precision within 5% of the designs, with acoustical properties similar to tissue, containing feature sizes relevant to SRI in the sub-100 μm range, print flow channels in 3D space and define fiducial markers with the purpose of probe alignment to the phantom features. Multiple iterations of fiducial markers have been made with a twofold purpose. Firstly, to investigate the type of markers made, and the influence on the contrast, and secondly, to investigate how small the markers can be made while still providing sufficient contrast.

II. MATERIALS AND METHODS

A. Stereolithography

Stereolithography uses a liquid resin hardened to the shape of the desired pattern through local illumination by a light source, in a layer by layer process to produce the designed 3D object. Fig. 1 illustrates the method employed. The printer vat contains the resin. The bottom of the vat is transparent to enable transmission of light to the resin. A glass slide is mounted on the fabrication stage, which is lowered into the resin, until a short distance from the vat bottom. The exact distance sets a limit on the thickness of the first printed layer. Upon illumination, the liquid resin will start crosslinking until reaching the ‘gel point’ at which the resin solidifies. The illumination system allows for local exposure of the polymer to enable printing of hollow structures. The fabrication stage is moved a specified distance away from the bottom of the vat once a layer has been exposed, thereby defining the next

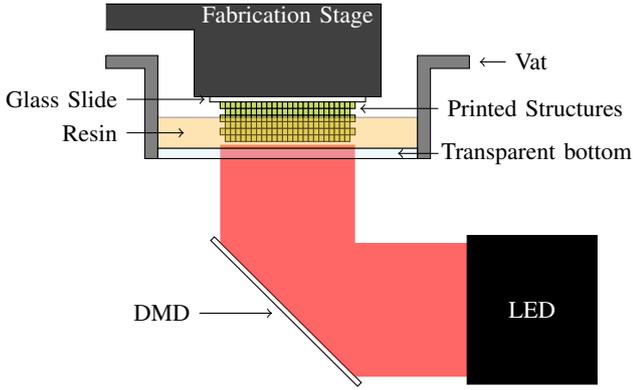


Figure 1. Sketch of the stereolithography setup. Light from an LED illuminates a digital mirror device (DMD), which reflects the light in the desired pattern through the transparent printer vat bottom. The illuminated resin in the vat will then start crosslinking. The initial layer of the printed structures is crosslinked to a glass slide mounted on the movable fabrication stage.

layer thickness. This is repeated until all layers of the object have been printed.

B. Custom Built Stereolithographic Printer

The printing setup has previously been described in [7]. Briefly, the custom-built high-resolution printer is based on a 1-to-1 projection of light reflected off a digital mirror device (DMD, 10.6 μm pixel pitch, DLP9500UV, Texas Instruments; part of a V-9501 UV SuperSpeed Digital Light Processing module, Vialux) onto the transparent bottom of a vat containing the resin. The printing process is controlled by a custom written MATLAB (MathWorks) code, that synchronises digital mask exposure on the DMD, with light exposure using a 365 nm high power LED (LZ1-00UV00, Ledengin), and fabrication stage movement via a linear stage (LNR50S, Thor Labs). The power density at the vat bottom was 18 mW cm^{-2} as measured using a UV power meter (S130VC + PM100D, Thorlabs). The stage is moved in 20 μm steps, matched to the light absorber concentration, which determines the vertical printing resolution. This results in an ultimate printing resolution of $10 \times 10 \times 20 \mu\text{m}^3$.

C. Resin

The resin used for printing consists of 3 parts: an aqueous pre-polymer solution, a photo-initiator and a photo-absorber. The pre-polymer will polymerise to form a solid when locally initiated by the light activated photo-initiator. Light in a medium is attenuated according to Lambert-Beers law

$$I = I_0 e^{-\mu d}, \quad (1)$$

where I_0 is the initial intensity, μ is the attenuation coefficient, and d is the depth at which the intensity, I , is measured. Fig. 2 shows two examples of the exponential decay of the light intensity in two different media with different μ . The pre-polymer reaction initiated by the photo-initiator will need a certain dose (intensity times the exposure duration) of light to solidify (reach its gel point). Thus, illumination of the polymer will only crosslink the polymer until a certain depth. In Fig. 2 this is marked by d_1 and d_2 for the two

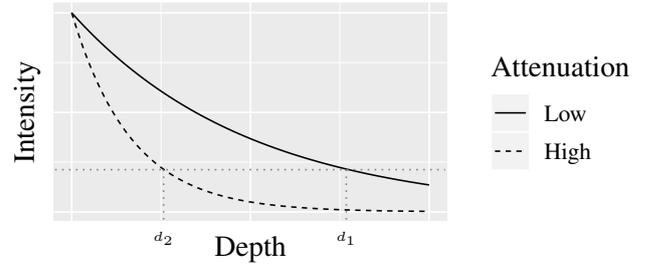


Figure 2. Light intensity against depth for two (arbitrary) different attenuation coefficients. d_1 and d_2 mark the depths, at which the light dose (intensity times exposure duration) has decreased to the threshold of resin solidification, marked by the grey horizontal dotted line.

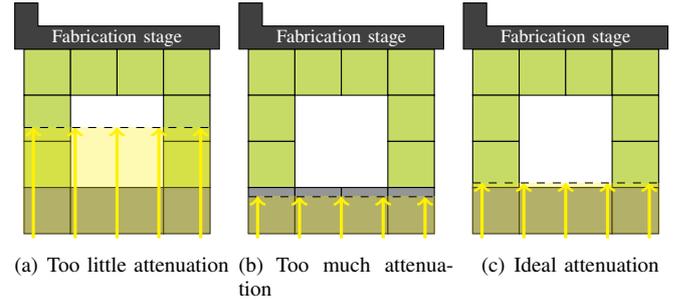


Figure 3. Sketches of the crosslinking depth for different levels of attenuation. The green squares are previously exposed voxels, and the grey squares are the voxels which are to be exposed in the new layer. The dashed lines mark the depth of the threshold dose. (a) has too little attenuation, and previous layers are re-exposed. (b) has too much attenuation, and the newly exposed layer is unable to crosslink with the structures in the previous layer. (c) has sufficient attenuation, with only a minimum overlap in the exposed region to the previous crosslinked structures.

curves corresponding to different levels of attenuation. While the interlayer movement of the stage sets a lower limit on the vertical resolution, the light dose may be sufficiently large to induce further solidification in previously exposed layers, as illustrated in Fig. 3(a). The aqueous solution of pre-polymer and photo-initiator has little attenuation, so light will propagate far into the resin before being absorbed. Addition of a highly absorbent photo-absorber allows to limit the propagation depth, thereby setting the depth resolution. A light absorber will modify μ of the solution according to (2),

$$\mu = \epsilon c, \quad (2)$$

where ϵ is the extinction coefficient, and c the concentration of the photo-absorber. Thus, μ can be modified by the type and concentration the added photo-absorber from a slowly attenuating medium such as water, exemplified by the solid curve and d_1 in Fig. 2, to a higher attenuating medium exemplified by the dashed curve and d_2 . If μ becomes too large, the polymer will not crosslink sufficiently deep to chemically bond to the overlying structures in the previous layer, as illustrated in Fig. 3(b). Thus, the choice of photo-absorber and concentration must be matched to get a slight overlap between the newly exposed regions, and the previously printed structures, as shown in Fig. 3(c).

The pre-polymer used is poly(ethylene glycol) diacrylate (PEGDA) 700 g/mol at 200 mg/mL. This polymer will contain 75 wt% water when converted to a hydrogel, thereby making it resemble tissue, both in terms of the water content [8] and, as a consequence, also in terms of the acousti-

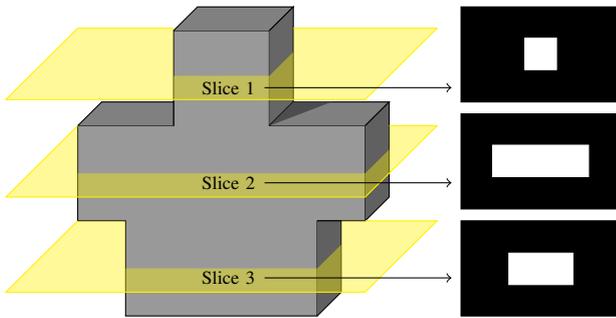


Figure 4. 3D model slicing. The slicing software takes a 3D model, as the one on the left, and decomposes it into slices, which are exported as .png files. The figures on the right are examples of slices corresponding to the yellow planes.

cal properties. The photo-initiator is lithium phenyl-2,4,6-trimethylbenzoylphosphinate (LAP) at 5 mg/mL, and is chosen for its water solubility and its absorption spectrum, which matches the light source used. The photo-absorber is quinoline yellow (QY) at a concentration of 12 mg/mL, also chosen for its water solubility and high extinction coefficient at the wavelength used.

D. 3D Design Creation

The design to be printed can be created in any 3D software. However, the 3D printer MATLAB code requires that the 3D model is sliced into separate layers. This can be done by the open source slicing software Slic3r (www.slic3r.org), which converts the model to a set of portable network graphics (.png) files. Additionally, the slicing software produces a build list, detailing the order and exposure time of each individual layer. Fig. 4 shows the slicing concept, where a 3D model, on the left, is sliced into separate .pngs, on the right, at each of the yellow planes.

When using the slicing software, the exact vertical placement of the slice is not controlled. The features of the 3D model may end up being split into multiple layers, if they do not align precisely with the slices. On the other hand, if the designs are simple, the .png files and the build list can be made manually, giving full control of the individual slices. As structures become smaller, this level of control will become increasingly more important. For the investigation of fiducial marker creation, all models have been made manually, by defining the individual .pngs in a MATLAB script.

E. Fiducial Marker Designs

Two types of markers were tested: hollow markers and solid markers. Ultrasound scattering is a function of perturbations in density and propagation velocity [9]. At the given polymer crosslinking density, the speed of sound within the material is practically constant. However, as more polymer polymerises, the crosslinking density will increase, which in turn increases scattering. Thus, when hollow markers are made, the scattering will increase as the density changes from that of the polymer, to that of the water contained within the hollow marker. When solid markers are made, the scattering increases due to a local increase in the polymer density, formed by overexposure of that region.

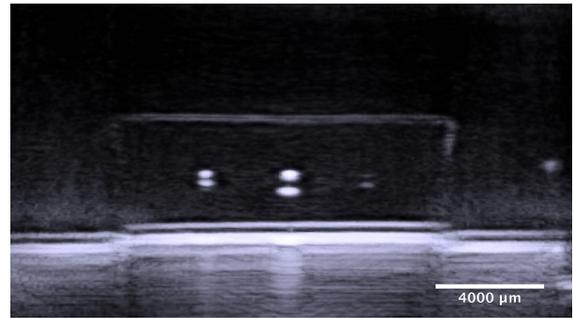


Figure 5. Ultrasound image showing the signal of two types of fiducial markers: Hollow (left and centre) and solid (right). The hollow markers have significantly larger contrast than the solid markers. Each marker shows two contrast regions. These correspond to the top and bottom of the cavity or solid region.

III. RESULTS

A. Type of Markers

A 3D printed phantom containing both solid and hollow markers was made to investigate the difference in contrast from the two types of markers. All ultrasound images presented have been scanned using a BK Medical "Hockey Stick" X18L5s probe and a BK 5000 scanner. A cross-sectional image showing both types of markers can be seen in Fig. 5. The left and centre markers are hollow markers, and the right marker is a solid marker. The left and right markers are both $(240 \mu\text{m})^3$, and the centre marker is $(400 \mu\text{m})^3$. Each marker can be seen as two high contrast spots on top of each other. The reason is that the propagation media density changes twice for each marker, namely at the top of the marker (the transition from polymer to unexposed resin - or higher density polymer for the solid markers), and at the bottom of the marker (back to the polymer from the unexposed polymer solution - or high-density polymer for the solid markers). Each of these interfaces will result in scattering of sound, and therefore contrast in the image. The markers on the left and the right are of the same size, and therefore suitable for comparison. The reflected intensity from the solid marker is measured to be 28 dB lower than the reflected intensity from the hollow marker.

B. Size of Markers

In an iterative process, the size of the markers was sought optimised. Fig. 6 shows a cross-sectional image of a single line of markers of decreasing sizes from $(140 \mu\text{m})^3$ on the far right and decreasing by $10 \mu\text{m}$ to $(60 \mu\text{m})^3$ on the far left, for a total of 9 markers across the phantom. Five markers are clearly visible and useful as point spread functions, with the smallest being $(100 \mu\text{m})^3$. Arguably, a few more of the smaller markers can be seen, barely above the noise level. The reason may be the limit of the chosen design method, but it could also be probe alignment, the choice of focus point, or a sub-optimal time gain compensation.

C. Manufacturing Accuracy

The polymer networks, and thereby the phantoms, will swell in size after printing due to water uptake. This means that

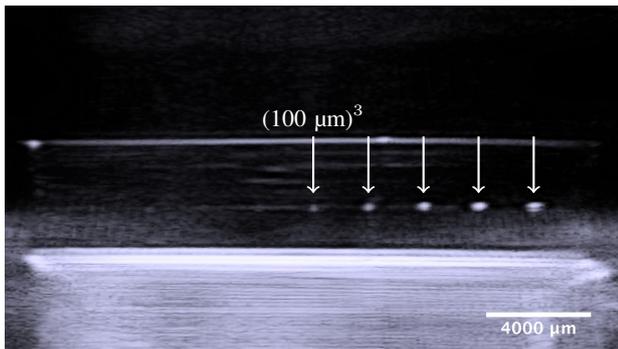


Figure 6. Ultrasound image showing the signal from hollow markers of decreasing design sizes from $(140 \mu\text{m})^3$ on the far right and decreasing by $10 \mu\text{m}$ in all dimensions to $(60 \mu\text{m})^3$ on the far left. The clearly visible markers are indicated by arrows.



Figure 7. Optical microscope image of two square fiducial markers.

the designed structures will need to account for this, by incorporating a scaling factor corresponding to the increase. This difference is quantified by optical microscope- and ultrasound images. Fig. 7 shows an optical microscope image of two hollow fiducial markers. The designed pitch between all markers was $2074 \mu\text{m}$. By measuring the pitch in optical microscope images similar to Fig. 7, the distance was found to be $2115 \pm 9 \mu\text{m}$ (an average deviation from the design of 2%). In a similar investigation of the ultrasound image in Fig. 6, the pitch was found to be $2078 \pm 71 \mu\text{m}$ (an average deviation of 0.2%).

D. Flow Channel Phantom

Fig. 8 shows an ultrasound SRI image of micro-bubble flow through a channel with a cross section of $(200 \mu\text{m})^2$ of a 3D printed phantom without fiducial markers. Liquid enters the lower channel to the right and leaves the upper channel to the right. The inserted colour wheel indicates the flow direction of the micro-bubbles. Fewer detected bubbles are observed passing through the top channel. Possible reasons could be channel misalignment or that the contrast bubbles were destroyed before reaching the end of the channel. Fig. 8 demonstrates that the channel is perfusable, and that the bubbles flow from the entrance in the lower right, up into the top channel, and out through the top right, as expected. This preliminary flow phantom demonstrates the 3D capabilities of printing by stereolithography. Using the same printing setup, channels with a cross section of $(100 \mu\text{m})^2$ have been demonstrated for other purposes than ultrasound [7],

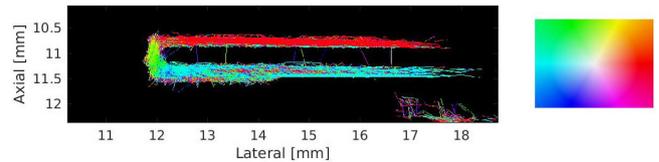


Figure 8. Super-resolution image of micro-bubbles through a 3D printed channel. The colour wheel indicates the micro-bubble flow direction.

and in-house testing have demonstrated $(50 \mu\text{m})^2$ perfusable channels.

IV. DISCUSSION AND CONCLUSION

We have presented a custom-built stereolithographic 3D printer capable of printing hydrogels, which are capable of storing 75 wt% of water, thereby resembling tissue in terms of the acoustic properties, making them well suited for the fabrication of ultrasound phantoms.

Fiducial markers have been printed, with hollow markers providing 28 dB more signal than solid markers. Markers of only $(100 \mu\text{m})^3$ have been shown to provide a well defined point spread function. A 3D printed flow phantom has been demonstrated with $(200 \mu\text{m})^2$ cross sectional channels. The designed structures have been printed with a final 2% increase in dimensions.

The preliminary findings presented within the previous sections, highlight the potential for stereolithographic printing of phantoms, and furthermore indicates that the fiducial markers obtained in this work are not demonstrating the ultimate resolution limit. Optimised printing schemes will be explored in an effort to obtain the highest possible resolution of the printing system.

REFERENCES

- [1] D. Hanahan and R. A. Weinberg. Hallmarks of cancer: The next generation. *Cell*, 144(5):646–674, 2011.
- [2] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter. Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging. *Nature*, 527(7579):499–, 2015.
- [3] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley. In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles. *IEEE Trans. Med. Imag.*, 34(2):433–440, 2015.
- [4] Y. Desailly, O. Couture, M. Fink, and M. Tanter. Sono-activated ultrasound localization microscopy. *Appl. Phys. Lett.*, 103(17):174107, 2013.
- [5] O. M. Viessmann, R. J. Eckersley, K. Christensen-Jeffries, M. X. Tang, and C. Dunsby. Acoustic super-resolution with ultrasound and microbubbles. *Phys. Med. Biol.*, 58(18):6447–6458, 2013.
- [6] K. Christensen-Jeffries, J. Brown, P. Aljabar, M. Tang, C. Dunsby, and R. J. Eckersley. 3-D in vitro acoustic super-resolution and super-resolved velocity mapping using microbubbles. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 64(10):1478–1486, 2017.
- [7] R. Zhang and N. B. Larsen. Stereolithographic hydrogel printing of 3D culture chips with biofunctionalized complex 3D perfusion networks. *Lab Chip*, 17(24):4273–4282, 2017.
- [8] R. M. Forbes, A. R. Cooper, and H. H. Mitchell. The composition of the adult human body as determined by chemical analysis. *J. Biol. Chem.*, 203(1):359–366, 1953.
- [9] J. A. Jensen. A model for the propagation and scattering of ultrasound in tissue. *J. Acoust. Soc. Am.*, 89(1):182–190, 1991.

A.3. *PAPER C - WAFER LEVEL CHARACTERIZATION OF ROW-COLUMN ADDRESSED CMUT ARRAYS*

A.3 Paper C - Wafer Level Characterization of Row-Column Addressed CMUT Arrays

Wafer Level Characterization of Row-Column Addressed CMUT Arrays

Erik. V. Thomsen, Kitty Steenberg, Magnus G. Petersen, Mads Weile, Andreas Havreland, Martin L. Ommen, Rune S. Grass, and Mathias Engholm
Department of Health Technology, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

Abstract—This paper presents a measurement methodology for wafer level characterization of row-column addressed (RCA) capacitive micromachined ultrasonic transducers (CMUT). Characterization of a 62+62 element RCA CMUT is presented. To facilitate wafer level electrical characterization measurements between adjacent electrodes can be used to characterize the device. This allows for determination of the individual element capacitance. Current-voltage measurements between adjacent top or bottom electrodes provides valuable information about process yield.

I. INTRODUCTION

Row-Column Addressed (RCA) capacitive micromachined ultrasonic transducers (CMUTs) are an attractive alternative to fully-populated matrix arrays, as they provide volumetric imaging with a greatly reduced number of electrical connections [1], [2]. Furthermore, they can be mass-produced on wafer scale using silicon process technology which allows for tight control of the dimensions. RCA CMUTs have been fabricated using various fabrication techniques including fusion bonding, anodic bonding and surface micromachining. Several devices have been presented including a 32+32 RCA CMUT chip based on a fusion bonding process with silicon nitride [3], a RCA CMUT chip fabricated using a surface micromachining process [4], a 32+32 anodically bonded RCA CMUT chip [5], a BCB based RCA CMUT chip [6], a 62+62 fusion bonded RCA CMUT probe [7], and a 120+120 element RCA CMUT probe fabricated using sacrificial release microfabrication [8]. As the resolution of the RCA transducer scales linearly with the number of elements the current trend is to increase the number of elements and fabricate large area chips.

In this work we present wafer level characterization of a local oxidation of silicon (LOCOS, originally used by [9] for CMUTs) based 62+62 element RCA CMUT transducer as shown in Fig. 1. Manual electrical characterization of a transducer having 124 element becomes a tedious job and automated tests need to be implemented on wafer scale, i.e. electrical wafer level test can be used to ensure that only fully functional arrays are used for probe manufacturing and to provide valuable information during process optimization.

Conventional linear arrays are easily characterized electrically using the two terminals on the device, e.g. the capacitance of an element can be measured directly using the two terminals of the element. However, RCA arrays cannot be measured in the same way as e.g. a measurement of the capacitance of a row-element will require that all columns are

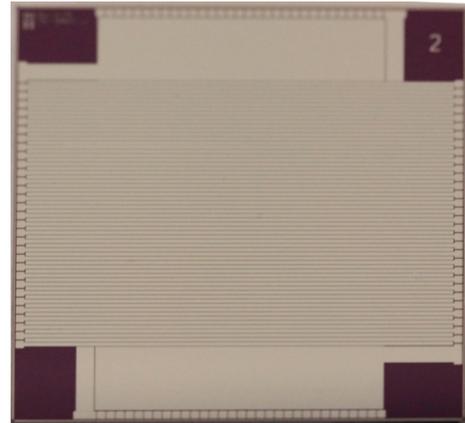


Fig. 1. Image of the 62+62 RCA CMUT chip

shorted and vice versa. This is not easily done on commercial wafer probers where micro manipulators are used to place probe needles or a probe card in fixed positions. During wafer level characterization the probe needles remain fixed and the wafer is automatically moved to perform the measurements needed. The objective of this work is therefore to present a methodology for characterization of RCA CMUT arrays with the aim of providing an automated and non-destructive wafer level test to determine if an array is fully functional or not.

The article is organized as follows: First, the measurement methodology is described. Then, the experimental details of the wafer level electrical characterization is given. Finally, the article ends with conclusions.

II. CMUT CHARACTERIZATION METHODOLOGY

The characterization methodology is illustrated in Fig. 2. Two fixed probes, $P1$ and $P2$, are used to connect two neighboring top electrodes (e.g. rows $R2$ and $R3$) to the measurement equipment.

In a typical situation a range of electrical measurements are performed:

- 1) Current voltage (IV) measurements are made to determine if adjacent top or bottom electrodes are short circuited due to e.g. problems during fabrication.
- 2) Impedance frequency (Z-f) measurements are performed to determine the frequency for capacitance voltage (CV) measurements and investigate the characteristics of the transmission line.

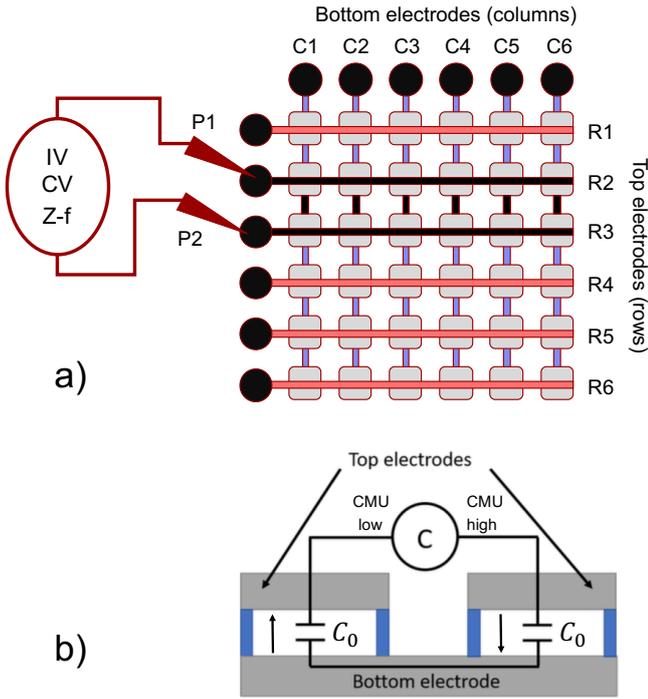


Fig. 2. a) Illustration of the measurement of the capacitance between two neighboring top electrodes, $R2$ and $R3$ using two probes, $P1$ and $P2$. b) The bottom electrodes, $C1$ to $C6$, effectively connect the CMUT cell capacitors in $R2$ in series with the cell capacitors in $R3$. The measured capacitance is therefore $C = C_0/2$. The arrows show the direction of the electric field in the CMUT cavities. The electric field points in opposite directions.

3) CV measurements are used to verify the electro mechanical behavior of the elements which is seen as a parabolic CV curve.

Once the measurement is completed the wafer is moved to allow the fixed probes to connect to e.g. $R3$ and $R4$ and this procedure is then repeated until all elements have been measured. When all rows have been characterized the procedure is repeated on the columns. This allows to determine the electrical characteristics of the array and identify defective elements.

During impedance measurements, the bottom electrodes, columns $C1$ to $C6$, effectively connect the CMUT cell capacitors in $R2$ in series with the cell capacitors in $R3$ as illustrated in Fig. 2b). Ignoring electrode resistance [10] and substrate coupling [11] the measured capacitance, C , is therefore

$$C = \frac{C_0}{2}, \quad (1)$$

where C_0 is the capacitance of a single row element. In this way the element capacitance can be estimated on wafer scale.

The electrical measurements allow to determine if the elements of the array perform as expected and process errors can be detected as abnormal electrical behavior indicating a faulty element. For example, if two neighboring rows are not completely separated during element etching, as illustrated on Fig. 3, the IV measurement will show that the elements are

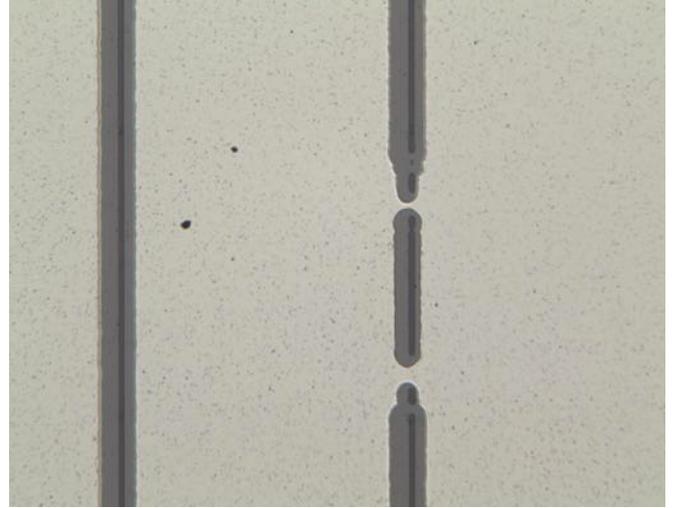


Fig. 3. Short circuit between two adjacent top electrodes (rows).

short circuited and impedance measurements, used to calculate the capacitance, will show both a small resistance and phase angle. Likewise, the capacitance between the short circuited elements and a neighboring element will represent a series coupling of $2C_0$ and C_0 , i.e. the measured capacitance is (again ignoring substrate coupling) $C = 2C_0/3$. In general, the capacitance between m short circuited elements and n short circuited elements will be

$$C = \frac{mn}{m+n}C_0. \quad (2)$$

Thus, also the capacitance between two neighboring elements can reveal structural defects.

III. EXPERIMENTAL

The methodology for wafer level characterization of CMUTs relies on a semi-automatic wafer prober and IV, CV and Z-f are the basic tests performed. These measurements were performed on a 62+62 RCA CMUT array from the same wafer as the array described in [12], however, a chip with lithographic errors possessing dielectric charging was selected to demonstrate the possibilities of wafer level test. A Cascade 12K Summit semi-automatic wafer prober and a KEYSIGHT B1500A Semiconductor Device Parameter Analyzer equipped with a B1520A multi frequency capacitance measurement unit (CMU) and six source measurement units (SMU) for IV measurements were used. The KEYSIGHT B1500A is connected to the probe manipulators using a SMU-CMU unify unit (SCUU) which allows to automatically switch between current-voltage and capacitance measurements. The measurement system is also equipped with a guard switch unit (GSWU) which is used for an accurate impedance measurement by connecting the guard lines between CMU high and low near the CMUT. The electrical measurements were performed between neighboring bottom or top electrodes.

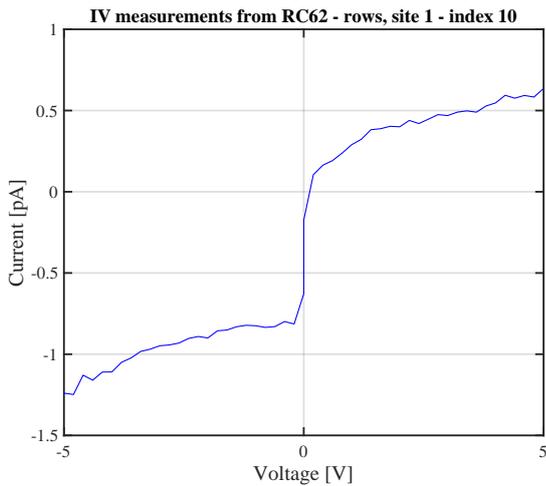


Fig. 4. Typical current-voltage measurement between adjacent row elements show a leakage current in the pA range.

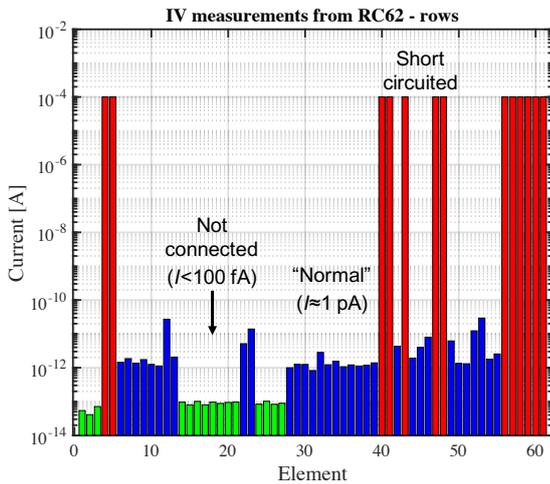


Fig. 5. Plot of the maximum current for each adjacent row elements determined from IV measurements as shown on Fig. 4. The blue bars show normal functioning elements. The red bars indicate elements that are shorted, and the green bars show elements that are not connected.

A. IV characterization

Fig. 4 shows the IV curve measured between two adjacent row elements on the 62+62 RCA CMUT array. The currents is as expected very low, in the pA range, corresponding to a large electrical resistance of around $10 \text{ T}\Omega$. Such IV measurements were performed between all adjacent elements in the array and the maximum current was extracted and the result is shown on Fig. 4. The plot reveals three categories of elements. The blue bars show normal functioning elements where the current between adjacent rows is around 1 pA. The red bars indicate elements that are shorted leading to a high current. The reason for the short circuited elements were found to be errors in the

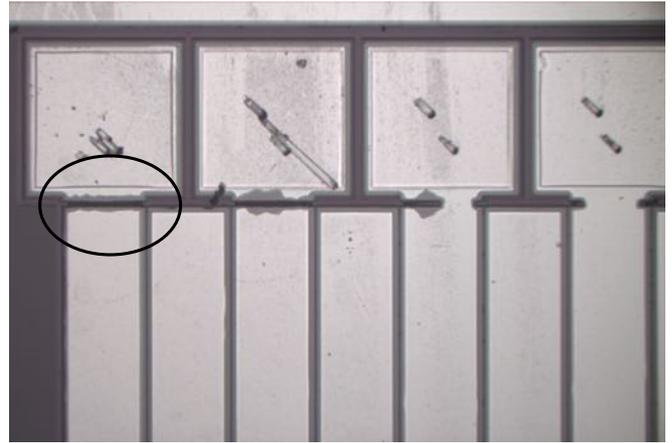


Fig. 6. Bond pad and row element not connected due to errors in the lithographic process.

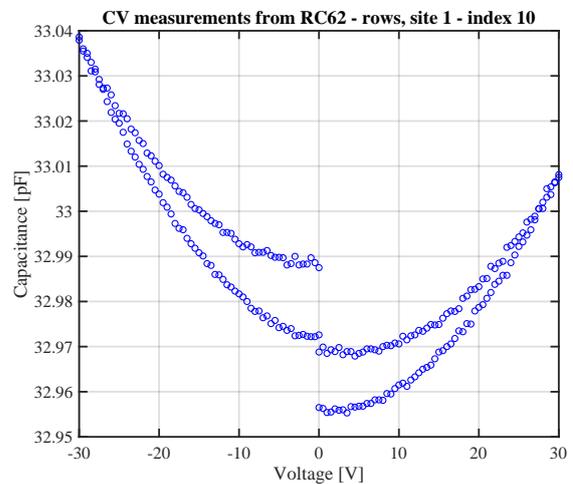


Fig. 7. Capacitance-voltage measurement between adjacent row elements. The CV curve reveals dielectric charging.

definition of the top electrodes as show on Fig. 3. Finally, the green bars show elements that are not connected leading to a very low current on the order of fA. The reason for this error was in all cases found to be an over etch of the aluminum electrode close to the bonding pad leaving the element without connection to the bond pad as illustrated in Fig. 6. Thus, IV measurements are very well suited for process control and can be performed on both top (row) and bottom (columns) electrodes.

B. CV characterization

Fig. 7 shows a typical CV curve as measured on the RCA CMUT chip shown on Fig. 1. The CV curve is parabolic as expected for a CMUT where the applied voltage decreases the gap in the CMUT cells leading to an increasing capacitance. It is noted that the CV curve has hysteresis indicating dielectric charging. Measurements on linear test elements has shown

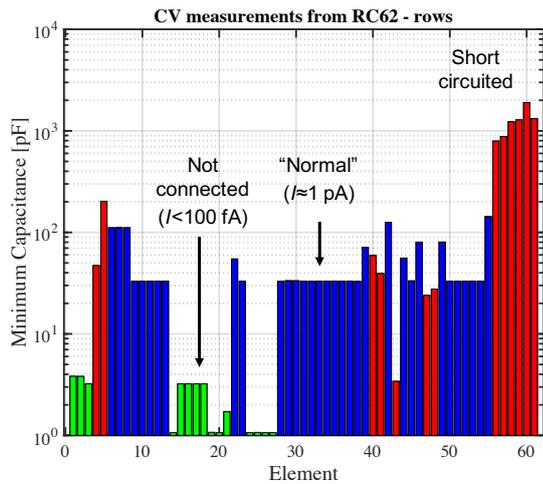


Fig. 8. Plot of the minimum capacitance for each adjacent row elements determined from CV measurements as shown on Fig. 7. The blue bars show normal functioning elements having a capacitance around 30 pF. The red bars indicate elements that are shorted so the capacitance cannot be determined. The green bars show elements that are not connected where the capacitance is around 1 pF.

that the device is electrically stable for one bias polarity. That the charging behavior is observed for *both* voltage polarities is because that measuring between adjacent electrodes means that the electric field in the CMUT cavities of the two elements always points in opposite directions as illustrated in Fig. 2. CV measurements were performed between all adjacent elements in the array and the minimum in capacitance was extracted and the result is shown on Fig. 7. The capacitance of the normal functioning elements is around 33 pF. The capacitance measured between elements which are shorted to neighboring elements is higher as the area of the elements are larger. CV measurements can also be performed between adjacent bottom electrodes.

IV. CONCLUSION

This paper presented a measurement methodology for wafer level characterization of RCA CMUTs. To facilitate automated wafer level tests, electrical measurements were performed between adjacent top (rows) or bottom electrodes (columns). IV measurements between fully functioning elements showed a maximum current in the pA range. It was found that short circuited elements and elements that are not connected to the bond pads can be identified electrically. The reason for the errors was found to be related to the etching process used for definition of the top electrodes. CV measurements allowed to determine the element capacitance as this is approximately twice the measured capacitance. When CV measurements are performed between adjacent electrodes the electric field has opposite directions in the two elements and if dielectric charging is present it will show up in the CV curve for both polarities even if the device is stable in one voltage polarity. In conclusion, wafer level characterization of RCA CMUT

devices is a valuable tool for selecting the best performing chips and to assist in process optimization.

V. ACKNOWLEDGEMENT

This work was financially supported by grant 7050-00004B from Innovation Fund Denmark, and from BK Medical, Herlev, Denmark.

REFERENCES

- [1] C. H. Seo and J. T. Yen, "A 256 x 256 2-D array transducer with row-column addressing for 3-D rectilinear imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 56, no. 4, pp. 837–847, April 2009.
- [2] C. E. M. Démoré, A. Joyce, K. Wall, and G. Lockwood, "Real-time volume imaging using a crossed electrode array," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 56, no. 6, pp. 1252–1261, 2009.
- [3] A. S. Logan, L. L. P. Wong, and J. T. W. Yeow, "2-D CMUT wafer bonded imaging arrays with a row-column addressing scheme," in *Proc. IEEE Ultrason. Symp.*, sep 2009, pp. 984–987.
- [4] A. Sampaleanu, P. Zhang, A. Kshirsagar, W. Moussa, and R. Zemp, "Top-orthogonal-to-bottom-electrode (TOBE) CMUT arrays for 3-D ultrasound imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 2, pp. 266–276, 2014.
- [5] A. Zeshan, X. Zhang, Ö. Oralkan, and F. Y. Yamaner, "2D CMUT array based ultrasonic micromanipulation platform," in *Proc. IEEE Ultrason. Symp.*, 2016, pp. 1–4.
- [6] A. S. Havreland, M. L. Ommen, C. Silvestre, M. Engholm, J. A. Jensen, and E. V. Thomsen, "BCB polymer based row-column addressed CMUT," in *Proc. IEEE Ultrason. Symp.*, 2017, pp. 1–4.
- [7] M. Engholm, H. Bouzari, T. L. Christiansen, C. Beers, J. P. Bagge, L. N. Moesner, S. E. Diederichsen, M. B. Stuart, J. A. Jensen, and E. V. Thomsen, "Probe development of CMUT and PZT row-column-addressed 2-D arrays," *Sens. Actuators A: Phys.*, vol. 273, pp. 121–133, 2018.
- [8] A. S. Savoia, B. Mauti, L. Fanni, G. Caliano, E. Boni, P. Mattesini, M. Scaringella, and P. Tortoli, "A 120+120-element crisscross CMUT probe's with real-time switchable electronic and fresnel focusing capabilities," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [9] K. K. Park, H. J. Lee, M. Kupnik, Ö. Oralkan, and B. T. Khuri-Yakub, "Fabricating capacitive micromachined ultrasonic transducers with direct wafer-bonding and LOCOS technology," in *IEEE 21st Int. Conf. Micro Electro Mech. Syst.*, 2008, pp. 339–342.
- [10] A. Havreland, M. Engholm, B. Tomov, J. Jensen, O. Hansen, and E. Thomsen, "CMUT electrode resistance design: Modelling and experimental verification by a row-column array," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, pp. 1110–1118, 2019.
- [11] M. Engholm, H. Bouzari, J. A. Jensen, and E. V. Thomsen, "Capacitive substrate coupling of row-column-addressed 2-D CMUT arrays," in *Proc. IEEE Ultrason. Symp.*, 2016, pp. 1–4.
- [12] T. L. Christiansen, M. F. Rasmussen, J. P. Bagge, L. N. Moesner, J. A. Jensen, and E. V. Thomsen, "3-D imaging using row-column-addressed arrays with integrated apodization — part II: Transducer fabrication and experimental results," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 959–971, 2015.

A.4 Paper D - Ultrasound Multiple Point Target Detection and Localization using Deep Learning

Ultrasound Multiple Point Target Detection and Localization using Deep Learning

Jihwan Youn, Martin Lind Ommen, Matthias Bo Stuart, Erik Vilain Thomsen,
Niels Bent Larsen, Jørgen Arendt Jensen

Department of Health Technology, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark

Abstract—Super-resolution imaging (SRI) can achieve sub-wavelength resolution by detecting and tracking intravenously injected microbubbles (MBs) over time. However, current SRI is limited by long data acquisition times since the MB detection still relies on diffraction-limited conventional ultrasound images. This limits the number of detectable MBs in a fixed time duration. In this work, we propose a deep learning-based method for detecting and localizing high-density multiple point targets from radio frequency (RF) channel data. A Convolutional Neural Network (CNN) was trained to return confidence maps given RF channel data, and the positions of point targets were estimated from the confidence maps. RF channel data for training and evaluation were simulated in Field II by placing point targets randomly in the region of interest and transmitting three steered plane waves. The trained CNN achieved a precision and recall of 0.999 and 0.960 on a simulated test dataset. The localization errors after excluding outliers were within $\pm 46 \mu\text{m}$ and $\pm 27 \mu\text{m}$ in the lateral and axial directions. A scatterer phantom was 3-D printed and imaged by the Synthetic Aperture Real-time Ultrasound System (SARUS). On measured data, a precision and recall of 0.976 and 0.998 were achieved, and the localization errors after excluding outliers were within $\pm 101 \mu\text{m}$ and $\pm 75 \mu\text{m}$ in the lateral and axial directions. We expect that this method can be extended to highly concentrated microbubble (MB) detection in order to accelerate SRI.

I. INTRODUCTION

Super-resolution imaging (SRI), often referred to as ultrasound localization microscopy (ULM), has demonstrated that it is possible to surpass the diffraction limit of conventional ultrasound imaging. Microvessels laying closer than a half-wavelength apart have been resolved by deploying microbubbles (MBs) as a contrast agent and using SRI [1]–[5]. The centroids of individual MBs can be easily found as MB echoes are much stronger than surrounding tissues when insonified, and their sizes are much smaller than a wavelength. Sub-wavelength imaging is achieved by accumulating the detected MB positions over time, revealing the fine structure of the microvasculature.

The MB detection in SRI, however, is still diffraction-limited because it is performed in conventional ultrasound images which are commonly formed by delay-and-sum (DAS) beamforming [6]. For accurate and reliable detection and localization, the MBs need to be more than a wavelength apart to avoid the overlaps of MB point spread functions (PSFs). Diluted concentrations of MBs are commonly used to satisfy this criteria as the behavior of MBs is hard to control. The number of detectable MBs, therefore, is constrained and this

leads to very long data acquisition times in order to map the entire microvasculature.

In this work, we propose a deep learning-based method for detecting and localizing multiple ultrasound point targets. The method especially aims to identify high-density point targets whose PSFs are overlapping, by feeding radio frequency (RF) channel data directly as input. A fully convolutional neural network (CNN) was designed to return 2-D confidence maps given RF channel data. The pixel values of the confidence maps correspond to the confidence of point targets existing in the pixels. The point target positions were extracted from the confidence maps by identifying local maxima. The CNN was trained and evaluated using simulated RF channel data. To further investigate the method on measured data, a phantom experiment was performed using a 3-D printed PEGDA 700 g/mol hydrogel phantom [7].

II. METHOD

A. Simulated Dataset

1) *RF channel data*: The Field II ultrasound simulation program [8], [9] was used to simulate RF channel data for generating a training and a test datasets. The datasets were composed of a certain number of frames. One frame was created by transmitting three steered plane waves after placing 100 point targets randomly within a region of $6.4 \times 6.4 \text{ mm}^2$ (an average target density of 2.44 mm^{-2}) where the center was 18 mm away from a transducer. The transducer was modeled after a commercial 192-element linear array, and the measured impulse response [10], [11] was applied to make the RF data as close to real measured data as possible. The parameters used in simulation are listed in Table I.

The simulated raw RF data were not beamformed but delayed, based on the time-of-flight calculated by

$$\tau_i(x, z) = \left(\sqrt{(x - x_i)^2 + z^2} + z \right) / c \quad (1)$$

where τ_i is the time-of-flight of the i -th transmission, (x, z) is the point, x_i is the center of the i -th transmission aperture, and c is the speed of sound. The delayed RF data were then sampled to have the same number of samples as that of confidence maps along the axial direction. The size of resulting RF data for one frame was $256 \times 64 \times 3$.

TABLE I
RF CHANNEL DATA SIMULATION PARAMETERS

Category	Parameter	Value
Transducer	Center frequency	5.2 MHz
	Pitch	0.20 mm
	Element width	0.18 mm
	Element height	6 mm
	Number of elements	192
Imaging	Number of TX elements	32
	Number of RX elements	64
	Steering angles	$-15^\circ, 0^\circ, 15^\circ$
Environment	Speed of sound	1480 m/s
	Field II sampling frequency	120 MHz
	RF data sampling frequency	29.6 MHz
Scatterer	Number of scatterers	100
	Lateral position range	$(-3.2, 3.2)$ mm
	Axial position range	$(14.8, 21.2)$ mm

2) *Confidence Map*: Non-overlapping Gaussian confidence maps were used as labels for training CNNs. Initially, binary confidence maps were created, where pixel values of one indicated a point target and the remaining pixel values were zero. A 21×21 Gaussian filter with a standard deviation of six was then applied at each point target position in the binary confidence maps. The filter values from the targets will be overlapped when some targets are closer than a half of the filter size in the confidence maps. In that case, the maximum value at each pixel location was taken. This maintained local maxima at target positions as opposed to the overlapping PSFs of DAS beamforming, and enabled the CNN to resolve targets closer than the diffraction limit.

The pixel size of the confidence maps was set to $25 \mu\text{m}$, and the image size of them became 256×256 , given the pixel size and the region of interest.

B. Convolutional Neural Network

1) *Network Architecture*: The proposed CNN is adapted from U-Net [12] which has an encoder-decoder structure. The feature maps are downsampled while the number of feature maps increases in the encoding path. Then, the feature maps are upsampled to their original size while the number of feature maps decreases in the decoding path. U-Net has a large receptive field, an effective input size that is covered by a convolution operation in an unit, for the sake of this structure. This is beneficial because a partial view of RF data is not enough to determine point target existence.

A detailed CNN architecture is illustrated in Fig. 1. Convolution and rectified linear unit (ReLU) layers in U-Net were replaced with pre-activation residual units (Fig. 1a) [13]. The pre-activation residual units ease optimization problem by introducing shortcuts, thereby improving performance. The proposed CNN (Fig. 1e) mainly consisted of four *down-blocks* (Fig. 1b), one *conv-block* (Fig. 1c), and four *up-blocks* (Fig. 1d). The skip-connections in U-Net was removed since it hindered the training. Instead, CoordConv [14] was added to transfer spatial information over convolution layers. Dropout [15] was attached after the shortcut in residual blocks for regularization. For pooling and unpooling, strided convolution and

pixel shuffle [16] were chosen, respectively. Leaky rectified linear units (Leaky ReLU) [17] were applied as non-linear activation to avoid dying ReLU problem causing nonactivated units.

2) *Training Details*: The CNN was trained by minimizing the mean squared error (MSE) between true confidence maps and CNN outputs. The training dataset consisted of a total of 10,240 frames. The kernel weights were initialized with orthogonal initialization [18] and optimized with ADAM [19] by setting $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-7}$. The initial learning rate was 10^{-4} and it was halved at every 100 epoch while limiting the minimum learning rate to 10^{-6} . The number of epochs was 600 and the mini-batch size was 32.

C. 3-D Printed Scatterer Phantom

A PEGDA 700 g/mol hydrogel scatterer phantom [7] was 3-D printed to investigate the proposed method on measured data. The phantom contained water-filled cavities which acted as scatterers. A total of 100 scatterers were placed on a 10×10 grid with a spacing of $518 \mu\text{m}$ in the lateral direction and $342 \mu\text{m}$ in the axial direction, as illustrated in Fig. 2.

The 3-D printed phantom was scanned by the Synthetic Aperture Real-time Ultrasound System (SARUS) [20] to acquire RF channel data. The same imaging scheme and transducer described in Table I were used. The phantom was placed on a motion stage and scanned at different positions by moving the motion stage at a step of $50 \mu\text{m}$ in the lateral direction. A total of 33 frames were obtained.

III. RESULTS

A. Simulation Experiment

The trained CNN was initially evaluated on a simulated test dataset. It was simulated in the same way as the training dataset in Field II, and consisted of 3,840 frames. In Fig. 3, the result of applying the CNN method to a test frame is compared with simply using the conventional DAS beamforming on the same frame. The CNN method was able to identify highly concentrated point targets while the DAS beamforming failed due to the overlapping PSFs. Full width at half maximum (FWHM) of the DAS beamforming at a depth of 18 mm was $387 \mu\text{m}$ (1.36λ) in the lateral direction and $140 \mu\text{m}$ (0.49λ) in the axial direction.

The CNN's capability to detect and localize point targets were quantitatively evaluated. Detection was measured by precision and recall that are defined by

$$\text{Precision} = \frac{TP}{TP + FP} \quad (2)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3)$$

where TP is the number of true positives, FP is the number of false positives, and FN is the number of false negatives. The positive and negative detections were determined by comparing estimated target positions with true target positions based on their pair-wise distances. The CNN method achieved

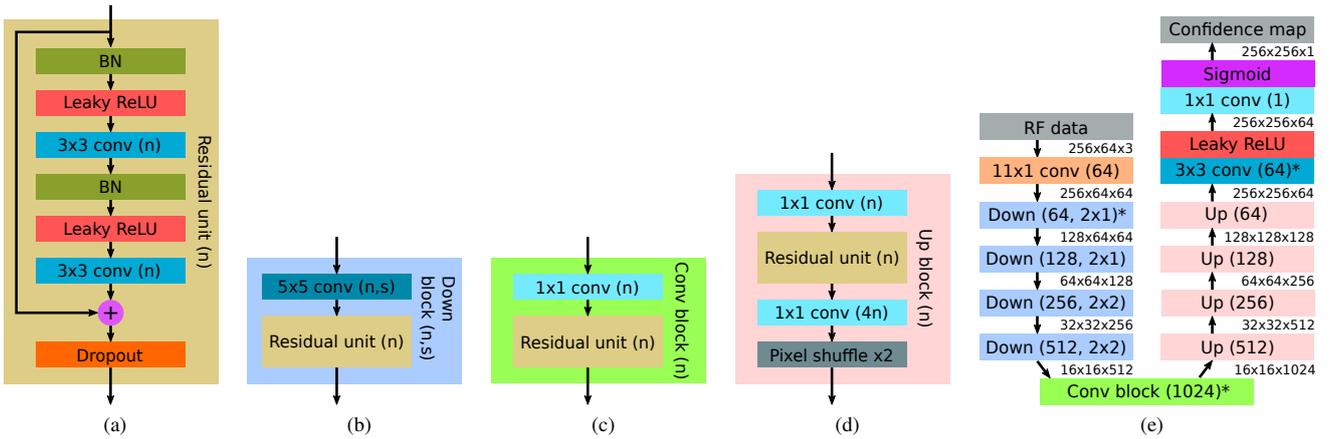


Fig. 1. The proposed CNN architecture and its components. (a) residual unit, (b) down-block, (c) conv-block, (d) up-block, and (e) the network overview. n and s in the parenthesis are the number of kernels and stride. The asterisk in (e) indicates that its first convolution in the block is CoordConv. The three numbers between blocks in (e) represent feature map size in the order of height, width, and the number of feature maps.

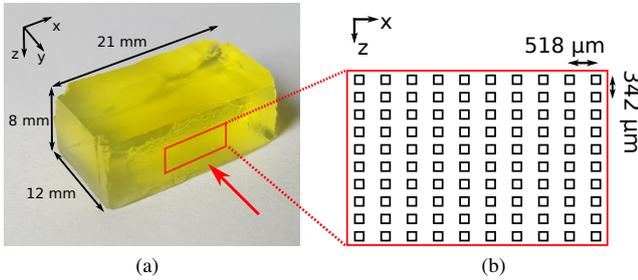


Fig. 2. Fabricated 3-D scatterer phantom: (a) photograph of the phantom and (b) 100 scatterers placed in a 10×10 grid.

a precision and recall of 0.999 and 0.960, while DAS beamforming achieved a precision and recall of 0.986 and 0.756.

Localization uncertainties in the lateral and axial position were calculated using the positive detections, and is illustrated using a box-and-whisker plot in Fig. 4a. The bottom and top edges of the blue box indicate the 25th (q_1) and 75th percentiles (q_3) and the center red edge indicates the median. The vertically extended line from the box (whisker) indicates the range of inliers which are smaller than $q_3 + 1.5 \times (q_3 - q_1)$ and greater than $q_1 - 1.5 \times (q_3 - q_1)$. The inliers were within $\pm 46 \mu\text{m}$ (0.16λ) in the lateral direction and $\pm 27 \mu\text{m}$ (0.09λ) in the axial direction.

B. Phantom Experiment

The CNN trained for the simulation experiment was not effective on the measured data because the scatterers in the phantom are not infinitesimally small point targets. The ultrasound beam is actually scattered twice at each scatterer in the phantom. Therefore, the RF data in the training dataset were simulated a second time by modeling a target using two points. In addition, the first scattering was phase reversed since the acoustic impedance is higher in the phantom than in the water inside the targets.

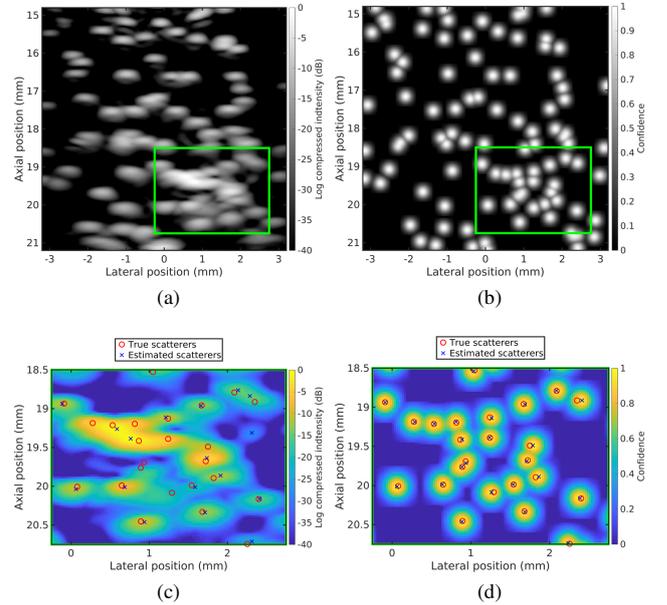


Fig. 3. Comparison of point target detection between DAS beamforming and CNN on a simulated test data using three steered plane wave transmissions. (a) DAS beamformed B-mode image, (b) confidence map returned from CNN, (c) true and estimated scatterer positions in the green square region of (a), and (d) true and estimated scatterer positions in the green square region of (b).

A new CNN was trained using the modified training dataset, and it successfully identified scatterers from the measured data as shown in Fig. 5. The achieved precision and recall were 0.976 and 0.998. The inliers were within $\pm 101 \mu\text{m}$ (0.33λ) in the lateral direction and $\pm 75 \mu\text{m}$ (0.25λ) in the axial direction, as illustrated in Fig. 4b.

IV. CONCLUSION

A CNN-based ultrasound multiple point target detection and localization method was demonstrated. The CNN was trained

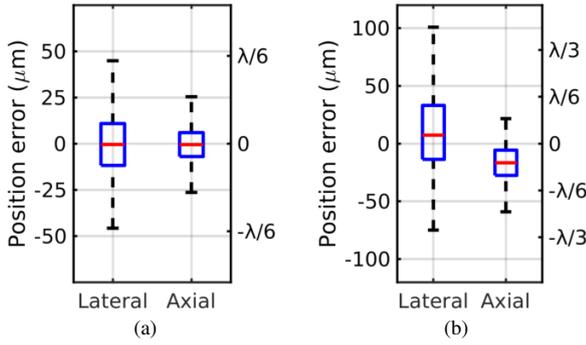


Fig. 4. Localization uncertainty in the lateral and axial direction measured (a) on the simulated test dataset and (b) on the measured phantom data.

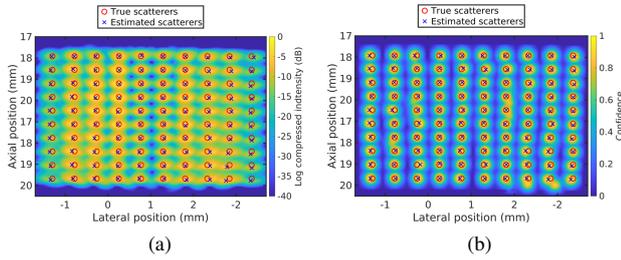


Fig. 5. Comparison of scatterer detection between DAS beamforming and CNN on phantom data using three steered plane wave transmissions. (a) DAS beamformed B-mode image and (b) confidence map returned from CNN with true and estimated scatterer positions

to learn a mapping from RF channel data to non-overlapping Gaussian confidence maps, and point target positions were estimated from the confidence maps by identifying local maxima. The non-overlapping Gaussian confidence maps were introduced to relax the sparsity of binary confidence maps while maintaining local maxima as target positions. The CNN method resolved point targets closer than the diffraction limit, whereas DAS beamforming failed as shown in Fig. 3.

It is also shown that the CNN method is applicable to real-world data, as well as simulated data, through the phantom experiment. It is notable that the training was performed solely using simulated data because it is nearly impossible to obtain a large number of measurements with ground truth for these kinds of work. It was also imperative to employ the measured impulse response and model targets following realistic physical modeling in the simulation.

We expect that this method can be extended to MB detection and potentially shorten the data acquisition time of SRI by detecting a greater number of MBs in a shorter amount of time.

ACKNOWLEDGMENT

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan V Volta GPU used for this research.

REFERENCES

- [1] O. Couture, B. Besson, G. Montaldo, M. Fink, and M. Tanter, "Microbubble ultrasound super-localization imaging (MUSLI)," in *Proc. IEEE Ultrason. Symp.*, 2011, pp. 1285–1287.
- [2] O. M. Viessmann, R. J. Eckersley, K. C. Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, pp. 6447–6458, 2013.
- [3] M. A. O'Reilly and K. Hynynen, "A super-resolution ultrasound method for brain vascular mapping," *Med. Phys.*, vol. 40, no. 11, pp. 110701–7, 2013.
- [4] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, pp. 499–502, November 2015.
- [5] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433–440, February 2015.
- [6] F. L. Thurstone and O. T. von Ramm, "A new ultrasound imaging technique employing two-dimensional electronic beam steering," in *Acoustical Holography*, P. S. Green, Ed., vol. 5. New York: Plenum Press, 1974, pp. 249–259.
- [7] M. L. Ommen, M. Schou, R. Zhang, C. A. V. Hoyos, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, "3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [8] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 39, no. 2, pp. 262–267, 1992.
- [9] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comp.*, vol. 10th Nordic-Baltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1, pp. 351–353, 1996.
- [10] —, "Safety assessment of advanced imaging sequences, II: Simulations," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 1, pp. 120–127, 2016.
- [11] B. G. Tomov, S. E. Diederichsen, E. V. Thomsen, and J. A. Jensen, "Characterization of medical ultrasound transducers," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [12] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Eur. Conf. Computer Vision*, 2016, pp. 630–645.
- [14] R. Liu, J. Lehman, P. Molino, F. P. Such, E. Frank, A. Sergeev, and J. Yosinski, "An intriguing failing of convolutional neural networks and the coordconv solution," in *Neural Information Processing Systems*, 2018, pp. 9605–9616.
- [15] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, pp. 1929–1958, 2014.
- [16] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2016, pp. 1874–1883.
- [17] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *ICML Workshop on Deep Learning for Audio, Speech, and Language Processing*, 2013.
- [18] A. M. Saxe, J. L. McClelland, and S. Ganguli, "Exact solutions to the nonlinear dynamics of learning in deep linear neural networks," *arXiv:1312.6120v3 [cs.NE]*, 2013.
- [19] D. Kingma and L. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980 [cs.LG]*, 2015.
- [20] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, "SARUS: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.

A.5. PAPER E - HISTORY AND LATEST ADVANCES IN FLOW ESTIMATION TECHNOLOGY: FROM 1-D TO 4-D

A.5 Paper E - History and Latest Advances in Flow Estimation Technology: From 1-D to 2-D to 3-D to 4-D

History and Latest Advances in Flow Estimation Technology: From 1-D in 2-D to 3-D in 4-D

Jørgen Arendt Jensen¹, Svetoslav Ivanov Nikolov², Kristoffer Lindskov Hansen⁴, Matthias Bo Stuart¹, Carlos A. Villagomez Hoyos², Mikkel Schou¹, Martin Lind Ommen¹, Sigrid Husebø Øygaard¹, Lasse Thumann Jørgensen¹, Marie Sand Traberg¹, Tin-Quoc Nguyen⁴, Erik Vilain Thomsen¹, Niels Bent Larsen¹, Christopher Beers³, Borislav Gueorguiev Tomov¹, and Michael Bachmann Nielsen⁴

¹Department of Health Technology, Technical University of Denmark, Lyngby, Denmark

²BK Medical, Herlev, Denmark, ³BK Medical, State College, PA 16803, USA,

⁴Department of Diagnostic Radiology, Rigshospitalet, Denmark

Abstract—Ultrasound imaging of flow has seen a tremendous development over the last sixty years from 1-D spectral displays to color flow mapping and the latest Vector Flow Imaging (VFI). The paper gives an overview of the development from current commercial vector flow systems to the latest advances in fast 4-D volumetric visualizations. It includes a description of the radical break with the current sequential data acquisition by the introduction of synthetic aperture imaging, where the whole region of interest is insonified using either spherical or plane waves also known as ultrafast imaging. This makes it possible to track flow continuously in all directions at frame rates of thousands of images per second. The latest research translates this to full volumetric imaging by employing matrix arrays and row-column arrays for full 3-D vector velocity estimation at all spatial points visualized at very high volume rates (4-D).

I. INTRODUCTION

The measurement of blood flow has undergone a tremendous development since the first system devised by Satomura in Japan in 1957 and 1959 [1, 2] more than sixty years ago. The continuous wave system could detect heart wall movements and flow patterns in peripheral arteries. Pulsed systems developed by Baker [3] and Wells [4] could display the spectral content of the flow signals at one depth in the vessel. These early 1-D systems forms the basis for the spectral Doppler systems of today, which are used for investigating and quantifying flow everywhere in the human circulation. Even the continuous wave systems are still in use in cardiology, where the velocities can be too high to measure for a pulsed system. Both yield quantitative estimates but can only measure at a single spatial location.

This limit was lifted by the Color Flow Mapping (CFM) system developed by Kasai et al [5, 6], where an auto-correlation estimator can estimate the velocity from only 8 to 16 emissions, thereby making it possible to acquire and display axial velocity images. This introduced the second most important innovation in velocity estimation, which is implemented in all commercial scanners for flow imaging of the vessels and the heart. The estimator has been investigated and improved in numerous papers using e.g. both RF averaging [7, 8] and cross-correlation [9, 10].

Although these systems are widely used in the clinic, and a whole range of diagnostic measures are routinely used, they also have a number of drawbacks and technical problems. Most importantly, only the axial velocity component is estimated. This is often compensated for by finding the beam-to-flow angle using the B-mode image, but it is inherently unreliable as the angle can vary over the cardiac cycle, and the flow is not necessarily parallel to the vessel wall. Often the beam-to-flow angle can be difficult to keep below 60°, and even a modest error of 5° can here lead to 20-30% errors in the estimated velocities. In many cases the axial velocity is actually the smallest component for e.g. peripheral vessels, and the lateral component is more important. The problem is addressed by the 2-D Vector Flow Imaging (VFI) systems presented in Section II, which also describes how more accurate measures of flow and turbulence can be attained in Section II-B.

A second problem is that CFM systems are limited in their frame rate by the sequential data acquisition due to the speed of sound [11, 12]. Eight to sixteen emissions must be acquired in multiple directions to yield an image, and the precision of the velocity estimates is limited by the number of emissions in the same direction. It is, thus, not possible to have both a large imaging region (large depth), fast frame rates, and precise estimates at the same time. Further, it is often difficult to detect flow in both the systolic and diastolic phase. The limits number of lines making low velocity estimation difficult, if aliasing should be avoided at the same time. These problems are addressed in Section III with the introduction of Synthetic Aperture (SA) systems, which radically breaks the trade-off between frame rate and precision [12]. It opens a whole range of new possibilities for flow imaging, where both slow and fast velocities can be estimated from the same data with a very high precision.

The third problem is that current systems only show flow in a 2-D image. Recently, 3-D volumetric imaging has been introduced, and these systems can show CFM images in a volume. Even though parallel beamforming is employed, it is still difficult to attain decent frame rates for real-time

cardiac imaging, and often the scanners have to resort to ECG gated sequences to stitch the volume together from multiple acquisitions. A further problem is the use of matrix array probes. Attaining a high resolution and contrast in ultrasound images require 64 to 128 transducer elements along the imaging plane, and for 3-D volumetric imaging matrix probes have to be used. These should ideally have at least 4,000 to 16,000 elements making them prohibitively expensive to develop and costly to use. Current state-of-the-art probes have more than 9,000 elements, which is still too low to attain a state-of-the-art image quality. Further, the velocity estimation is still only in the axial direction and not in full 3-D. These problems are addressed in Section IV, which shows how the latest research in Row-Column (RC) matrix probes potentially can be a solution to the problems of fast 4-D imaging with display of the full 3-D velocity vector in all points in the volume in real time.

II. 2-D VECTOR FLOW IMAGING

It was early realized that only estimating the axial velocity component was not sufficient to give a complete picture of the complex human blood flow. Fox [13] suggested the first system with two crossing beams to enable estimation of the lateral velocity component from triangulation. This has later been investigated and optimized by a number of authors [14, 15]. A second approach developed by Trahey et al [16] used speckle tracking, where a small search region was correlated to a larger image region. The velocity could then be found for both components.

A. Transverse oscillation

The first approach to make it into commercial scanners was the Transverse Oscillation (TO) method developed by Jensen, Munk, and Anderson [17, 18]. Axial velocity estimators rely on the sinusoidal signal emitted, and the velocity is estimated by correlating multiple emissions in the same direction. The motion between emissions is then found through either an autocorrelation using the phase shift or a cross-correlation for the time shift [19]. The idea in TO is to introduce an oscillation transverse to the ultrasound beam and then find the lateral displacement. A Fourier relationship exists between the transducer's aperture sensitivity and the lateral far-field sensitivity [17, 20, 21]. Introducing two peaks in the receive apodization therefore generates a lateral oscillation, where the frequency is determined by the separation of the two peaks. A dedicated estimator was developed for separately estimating the axial and lateral velocity components [22]. The method was implemented on BK Medical scanners (Herlev, Denmark) and FDA approved in 2012 [23]. It made it possible for the first time to visualize the complex flow in the body in real-time, and vortices in e.g. the bulbous of the carotid artery could be seen as shown in Fig. 1. The approach has been implemented on linear [17, 22], convex [24], and phased array probes [25] and can also be used for finding the spectrum of the transverse velocity [26].

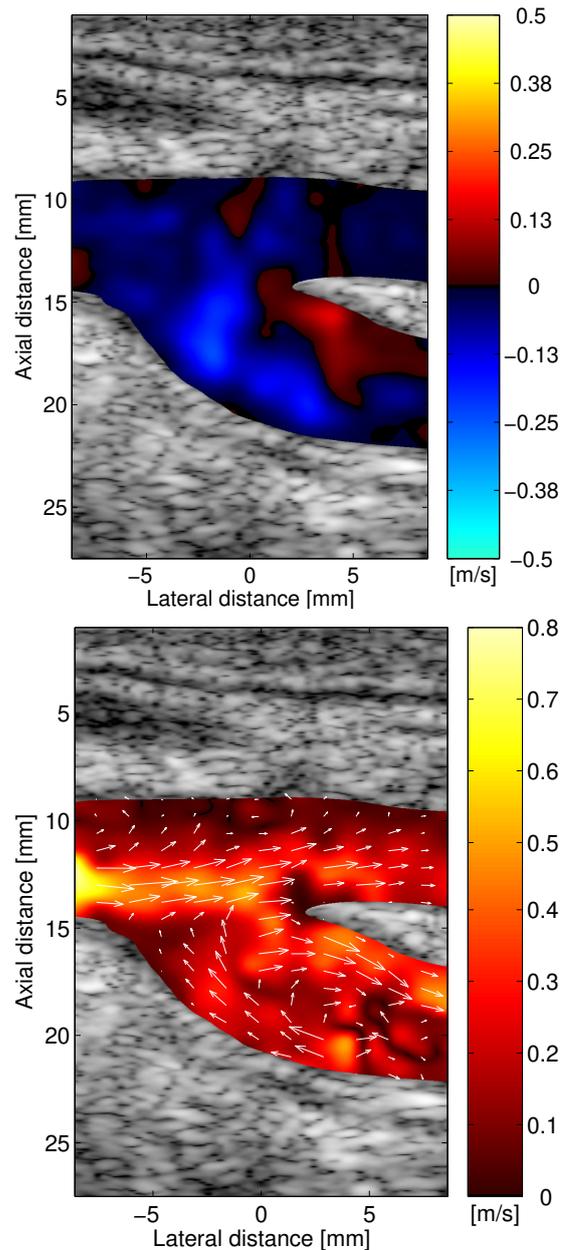


Fig. 1. In-vivo images of flow in the carotid bifurcation right before peak systole. The top images shows the CFM image for the axial velocity, and the bottom image shows the VFI using TO. A vortex is seen in the carotid bulb, and the velocity estimates are more consistent with what is found in the carotid artery (from [27]).

An example of flow in the aorta is shown in Fig. 2 for a short-axis view. The direction and velocity magnitude of the blood flow are displayed as colored pixels defined by the 2-D color bar with arrows superimposed for showing direction and magnitude. The short-axis view shows the rotation of the flow, which is nearly always found during the cardiac cycle, and the image demonstrates that the velocity can be estimated for all directions [28].

A range of studies have been conducted using the BK

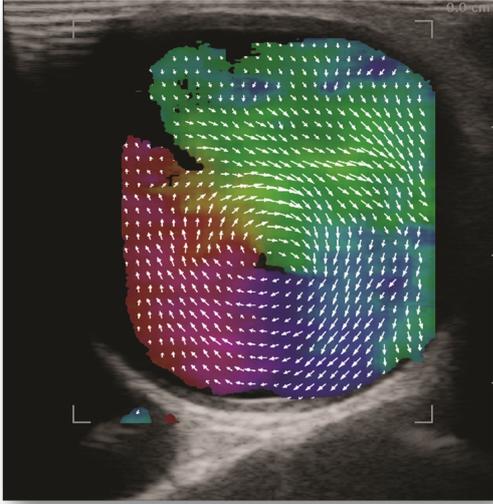


Fig. 2. Vector velocity imaging of blood flow in the ascending aorta in a short-axis view. The colors and arrows indicate velocity direction and magnitude (modified from [28]).

implementation. This includes investigating volume flow in arteriovenous fistulas [29], intraoperative cardiac examinations [30], flow in the aorta [28], flow in the ascending aorta for normal, stenotic and replaced aortic valves [31], and transthoracic VFI examination of newborns and infants with congenital heart defects [32]. Other groups have also investigated VFI and compared it to e.g. spectral velocity methods [33].

Vector flow is now also implemented on systems from Mindray and Toshiba, and a comprehensive review of all the developed methods can be found in [11], which also lists the comprehensive literature in the field for a range of different methods and clinical investigations.

B. Quantitative Measurements in VFI

Currently, quantification of velocities is obtained by using the axial velocity component from spectral velocity estimates, as the measurements are more precise than CFM results due to the continuous acquisition in one direction. The measurements have to be corrected for the beam-to-flow angle, and variations in this can lead to a serious bias. A 5° error at a 60° beam-to-flow angle can lead to a 20% error in the velocity. VFI can automatically compensate such errors and can also handle that the beam-to-flow angle varies over the cardiac cycle. An example of quantitative VFI measurements is shown in Fig. 3, where both the mean value and the standard deviation (SD) can be estimated by measuring over several cardiac cycles [34].

Many other quantities can be derived from VFI data including flow complexity for revealing disturbed and turbulent flow [31, 35], volume flow [36], and pressure gradients [37]. In the last example, the pressure gradients are estimated by solving a simplified version of the Navier-Stokes equation with the VFI estimates as input. An example of this is shown in Fig. 4, where the top image shows the trajectory for the pressure gradient calculation, and the lower graph shows the mean

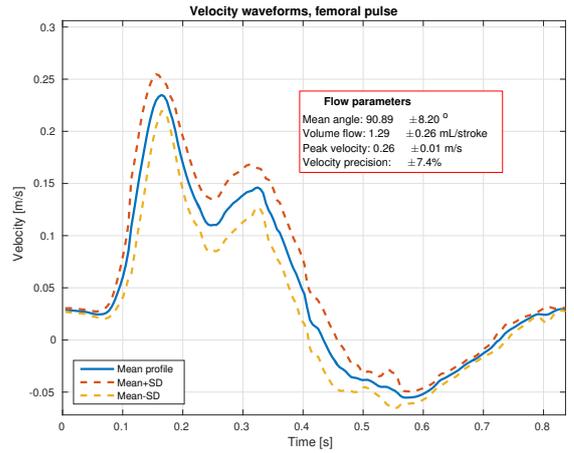


Fig. 3. Quantitative velocity measurements from a carotid phantom using a linear array probe with a directional TO velocity estimator (from [34]). Several cardiac cycles are automatically aligned and the mean value and SD are estimated from the 10 cycles for both the beam-to-flow angle and various velocity measures.

pressure gradient and its SD found from 11 cardiac cycles. The pressure gradient can be retrospectively found from the 10 seconds of data for any trajectory within the vector flow imaging region with a precision of 19%. A large improvement compared to a pressure catheter, which had a relative SD of 786% [38].

III. SYNTHETIC APERTURE FLOW IMAGING

A major problem in conventional flow imaging is the sequential data acquisition, which limits the frame rate and the amount of data available for velocity estimation [12, 41]. This limits the penetration depth, the maximum detectable velocity, and the precision of the estimates. A break with this paradigm is to employ SA imaging as shown in Fig. 5, where the region of interest is broadly insonified by using spherical or plane waves. The scattered signal is then received on part or all of the elements, and a full Low Resolution Image (LRI) can be generated. Combining LRIs from a number of emissions then yields a High Resolution Image (HRI) dynamically focused in both transmit and receive. The focusing is performed by summing the waves in phase, and for spherical emissions the focusing times are calculated as:

$$t_{i,j} = \frac{|\vec{r}_i - \vec{r}_p|}{c} + \frac{|\vec{r}_j - \vec{r}_p|}{c}, \quad (1)$$

where \vec{r}_i is the origin of emission i , \vec{r}_p is the location of the imaging point, and \vec{r}_j is the position of the receiving element j . The high resolution image is then made by:

$$y(\vec{r}_p) = \sum_{i=1}^{N_i} \sum_{j=1}^{N_j} a(\vec{r}_i, \vec{r}_p, \vec{r}_j) r(t_{i,j}), \quad (2)$$

where N_i is the number of transmissions and N_j the number of receiving elements. Here $a()$ is the apodization function or relative weight between emissions and between receiving elements, which is often calculated from the F-number in transmit and receive. The same calculations are performed for

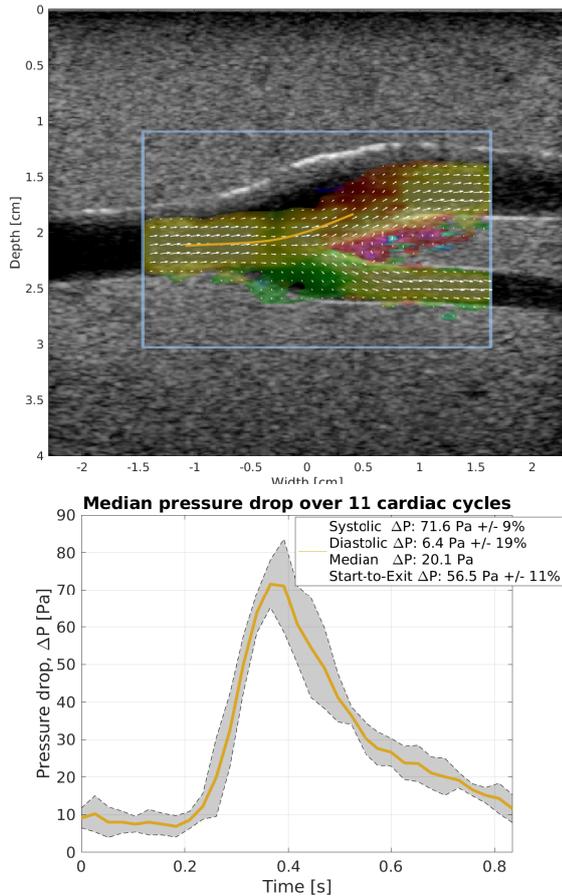


Fig. 4. Estimated pressure gradient from a carotid artery phantom. The top image shows the VFI and the trajectory for finding the pressure gradient (orange line). The lower graph shows the estimated pressure gradient from 11 cardiac cycles including the relative SD.

plane wave imaging with a replacement of the transmit delay (the first term in (1)) with the corresponding equation for a plane wave. This is often called ultrafast imaging [42], but the imaging scheme is really the same for both types of waves. The only difference is the calculation of the transmit delay in the beamforming, and we will, therefore, call both schemes for SA imaging in this paper. Creating SA images decouples frame rate from the number of lines in the image, and the frame rate is only determined by the number of emissions.

It might be counter-intuitive that such images acquired over multiple emissions can be used for velocity estimation, as the investigated object is moving between emissions and, thus, cannot be summed coherently. The initial idea for SA flow imaging is illustrated in Fig. 6, where a short sequence is used for SA imaging [43–45]. The emissions are shown on the top and the LRIs beneath. The bottom row shows the HRIs when the different LRIs are combined. A single scatterer moving towards the probe is investigated. The LRIs are not summed in phase, and HRI $H^{(n-3)}$ is different from $H^{(n-2)}$, but equal to $H^{(n-1)}$ apart from the shift in position, where n is the emission number. The basic idea is that the HRIs are highly correlated, if their emission sequences are the same. They may not be

summed fully in phase, but the defocusing from motion is the same for all HRIs, as the emission sequence is the same. They can, therefore, be correlated to find the velocity.

This might seem like a small detail, but it has major implications for flow imaging. Firstly, imaging is continuous, and data are available everywhere in the imaging region for all time. It is, thus, possible to average the correlation functions over as long time as the flow can be considered stationary [46]. Also, flow can be followed in any direction, as data is available for the whole imaging region, and beamforming can be made in all directions. Any echo canceling filter can be used without detrimental initialization effects, making it much easier to separate out flow from tissue [47–50].

An example of the benefits from SA flow imaging can be seen in Fig. 7, which shows a velocity magnitude image acquired using an 8 emissions SA sequence [51]. The data have been beamformed along the flow direction and the velocity estimated by cross-correlating these directional lines for 16 HRIs, which yields the velocity magnitude. No post processing has been employed on the image, and only the raw estimates are shown. The relative standard deviation to the peak velocity is 0.3% for very precise quantitative data, ideal for the quantification described in Section II-B. Data can be beamformed in any direction, making it also possible to estimate transverse flow [51]. Methods for estimating the correct beam-to-flow angle have also been developed [52, 53].

The current state-of-the-art in SA flow imaging is shown in Fig. 8, where the flow in the carotid bifurcation is measured on a healthy volunteer [53]. Here, a five emissions sequence was used, and it can potentially yield more than 3000 frames per second. Images at three different time points in the cardiac cycle are shown at the top. The bottom graph shows the velocity magnitude estimated in the white circle in graph c). The evolution on the vortex in the carotid bulb can be studied in detail using such ultrafast imaging.

A major issue in these images is the very large amount of data and the significant number of calculations to conduct for creating real time imaging. The current trend is to employ fast GPUs to perform the beamforming and this can often approach real time imaging [54–57]. Another approach is to reduce the amount of data and thereby the calculation load. Dual stage beamforming has been developed to reduce the sampled data to one channel, and the processing demand is thereby also reduced proportionally. It was demonstrated in [58] that very fast SA VFI could be attained by this approach using TO and dual stage beamforming, and the processing could be performed in real time on a Tablet [59].

A. Fast Flow

One problem in SA imaging has been the reduction of the detectable peak velocity. For SA flow imaging the data has to be acquired over N_e emissions, and the effective pulse repetition frequency $f_{prf,eff}$ is equal to f_{prf}/N_e . The maximum detectable velocity v_{max} in velocity estimation is generally proportional to $\lambda f_{prf,eff} = v_{max}$, which is reduced by a factor N_e compared to traditional flow imaging. There is,

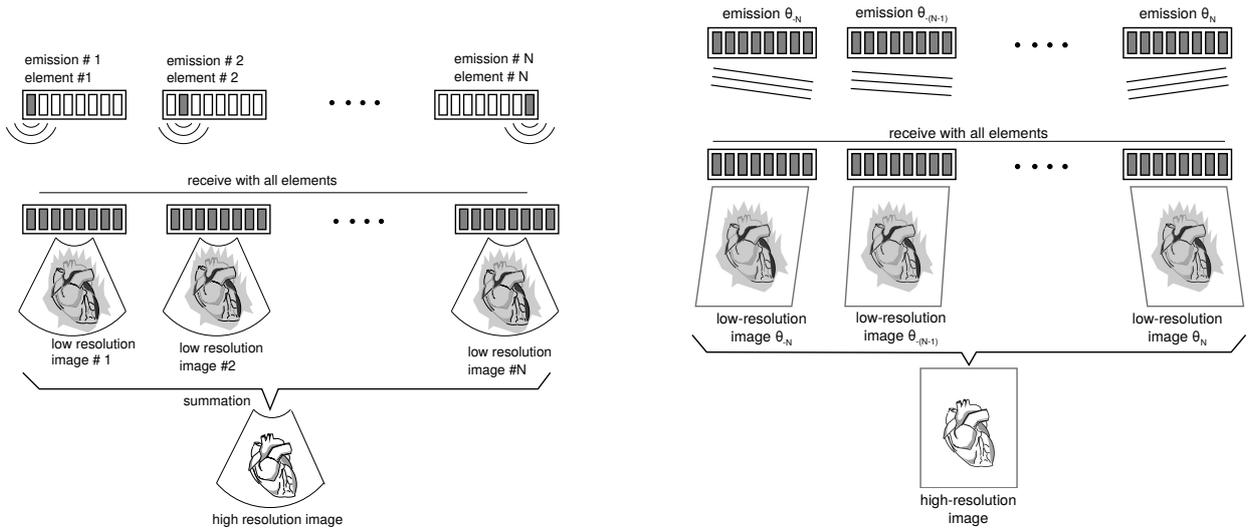


Fig. 5. Principle of SA imaging (left figure from [39]). The spherical emissions are shown on the top row with reception on all elements in the middle row. Plane wave imaging is shown in the right figure (from [40]).

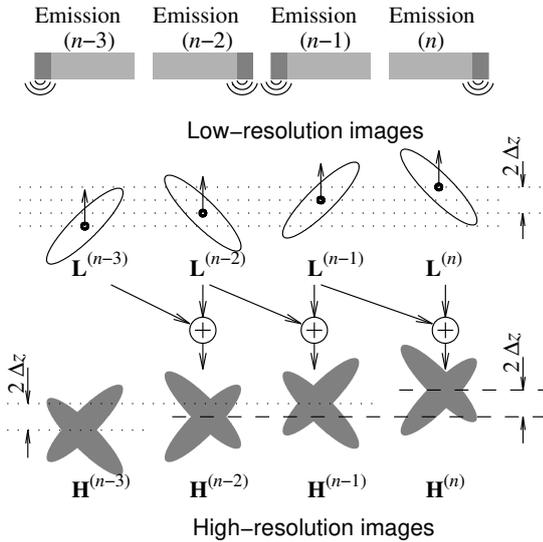


Fig. 6. Principle of SA flow imaging (from [43]).

thus, a compromise between sequence length and v_{max} . Often a longer sequence is preferred to enhance contrast and this reduces v_{max} . A possible solution is to use single emissions like in [61–64], but this reduces contrast and makes it difficult to estimate flow in small vessels.

The problem has recently been solved by introducing interleaved sequences, where an emission is repeated as shown in Fig. 9. The beamformed HRIs are then only temporally separated by $1/f_{prf}$ and not $1/f_{prf,eff}$, and v_{max} is increased by a factor N_e . Combined with a cross-correlation estimator made it possible to estimate velocities above 5 m/s for imaging down to 15 cm [60, 65], and it is also possible to further increase the limit by using directional beamforming as in Fig. 7.

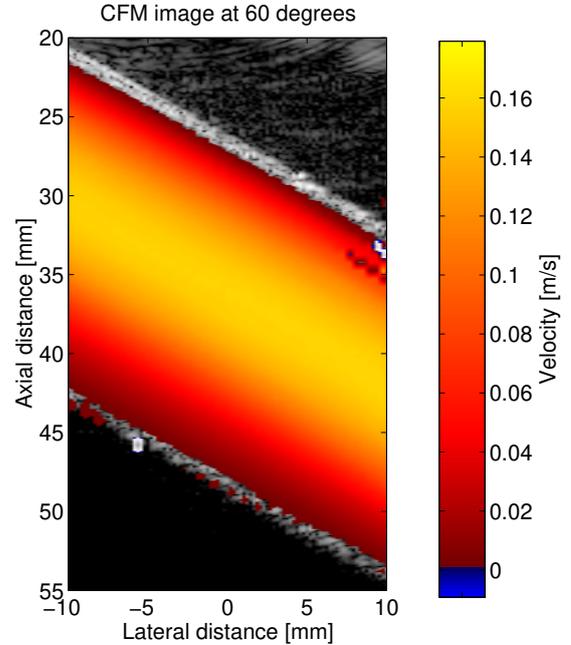


Fig. 7. Velocity magnitude velocity image acquired using SA flow imaging and directional beamforming (from [51]).

B. Slow Flow

A major advantage of continuous imaging is the possibility of using advanced echo canceling filters to separate flow from tissue. This is especially important for low velocities, and SA imaging has created major breakthroughs in studying slow flow in e.g. the rat brain as shown in Fig. 10 and the kidney [66, 67]. In particular the employment of Singular Value Decomposition (SVD) echo canceling methods has benefited low velocity imaging and introduced a whole new range of possibilities [47, 50, 68].

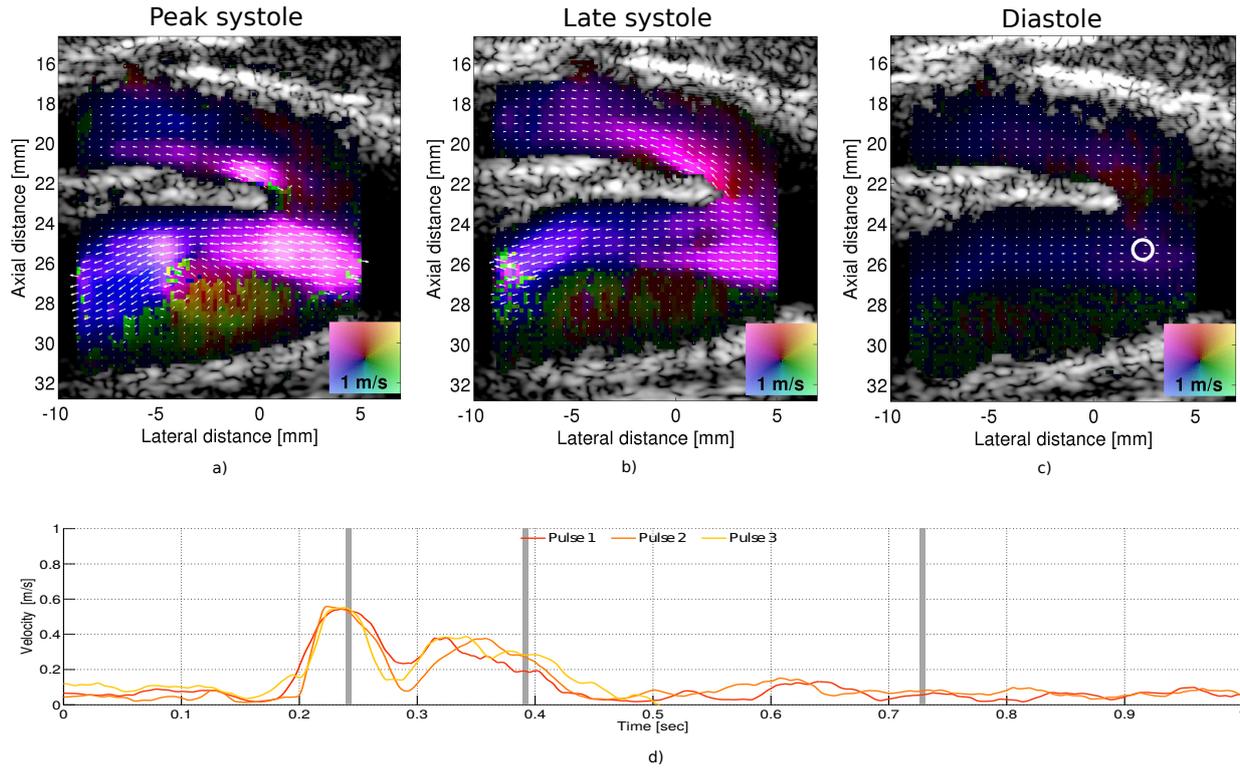


Fig. 8. VFI acquired using a SA flow sequence and directional beamforming (from [53]). Images at three different time points in the cardiac cycle is shown on the top, and the measured velocity magnitude over time for three cardiac cycles are shown in the bottom figure.

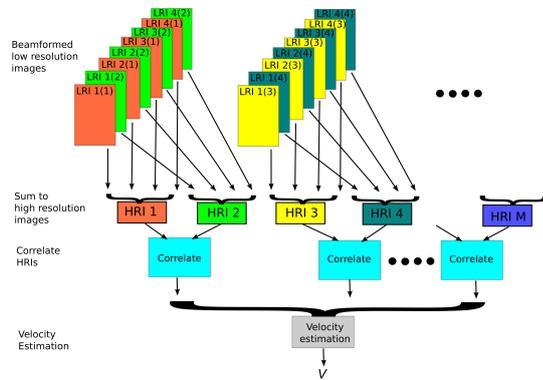


Fig. 9. Inter-leaved SA sequence where LRIs are repeated to minimize the distance between HRIs. The same colored LRIs are summed to yield one HRI. The effective $f_{prf,eff}$ is equal to the highest possible value (f_{prf}) due to the inter-leaving. Correlations in the blue boxes yield the same correlation function, which are then averaged to improve precision (from [60]).

IV. FROM 2-D TO 4-D

The ultimate goal for VFI is to yield a full 3-D volumetric image at a high frame rate (4-D) with the full velocity vector determined for all three velocity components (3-D). This could be called 3-D VFI in 4-D. SA imaging can be used for this using matrix probes, where the emitted waves can be steered in all directions to insonify the whole volume continuously. The TO approach has been modified to estimate all three velocity components [71, 72]. A 1024 elements Vermon matrix probe

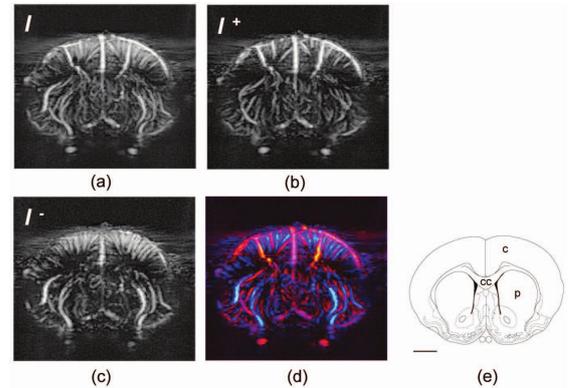


Fig. 10. Directional power Doppler. (a) Initial μ Doppler image. (b) Positive part of the Doppler power spectrum I^+ quantifying the volume of blood flowing up. (c) Negative part of the Doppler power spectrum I^- quantifying the volume of blood flowing down. (d) Color-coded μ Doppler image: in each pixel, the positive part is colored on a red range of intensities and the negative part on a blue range of intensities. (e) Anatomy of the brain slice (bregma + 1.0 mm). Main structures: cortex (denoted c), corpus callosum (cc) and caudate putamen (p). Scale bar: 2 mm. (from [67]).

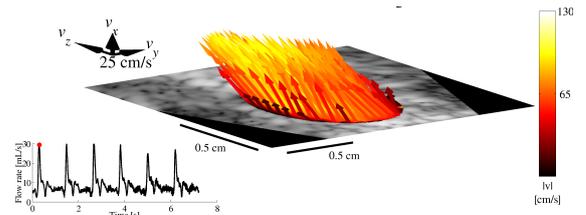


Fig. 11. Three-dimensional vector flow from the common carotid artery of a volunteer during peak systole using a 3-D TO estimator (modified from [69]).

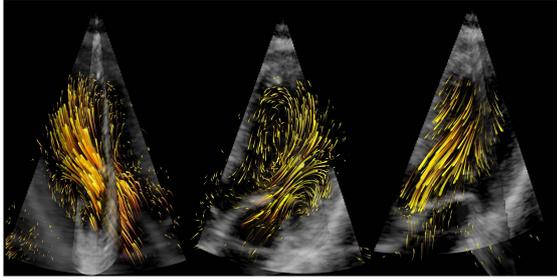


Fig. 12. Three-dimensional cardiac VFI rendering of the flow path lines at three time points in the cardiac cycle corresponding to diastole, diastasis, and systole. (from [70]).

[73] was used with the SARUS research scanner [74]. In-vivo imaging of ten volunteers was conducted on the carotid artery in [69] as shown in Fig. 11, and the volume flow could be determined with a SD of 5.7%. 3-D VFI has also been conducted in the heart using a modified GE Vivid E95 ultrasound scanner (GE Vingmed, Horten, Norway) using a GE 4V-D matrix array transducer for full volumetric coverage of the left ventricle at 50 volumes/second utilizing ECG-gating [70]. An example of these measurements is shown in Fig. 12.

One major problem is, however, the amount of elements needed. Both examples above use more than 1000 transducer elements, with probe foot-prints that are small, thus, impeding focusing. Good focusing in 2-D demands larger probes with 128 to 192 elements to maintain a low F-number for all imaging depths. Translating this to 3-D yields $192^2 = 36,864$ independent elements, which is impossible to connect through a cable to the scanner. A possible solution is to use a sparse array or electronic beamforming in the handle. This still restricts the number of elements to around 9,000 for roughly 100 elements on each side of the array. Low F-number focusing is therefore very difficult and expensive to attain in 3-D imaging, and compromises have to be made in both the imaging schemes and beamforming.

A novel solution to this problem is to employ Row-Column Arrays (RCAs), where rows and columns are independently addressed [75–78]. The number of interconnects is then transformed from N^2 to $2N$, thus reducing it by a factor of $N/2$. This makes very large arrays possible, and much lower F-numbers can be maintained for larger depths. A further advantage of the large array size is the increased penetration depth. This again can be used for increasing the center frequency of the probe and thereby resolution. Arrays with only 64+64 elements at 3 MHz have attained a decent volumetric image quality and a penetration down to 30 cm for SA imaging sequences [79].

The RCAs can be combined with all the methods presented here, and, thereby, attain the previously mentioned advantages. Three-dimensional VFI was presented for a line and a plane in [81] and for a volume [80, 81] using a 64+64 RC array, and the TO approach adapted to 3-D VFI as shown in Fig. 13.

Recently, a SA RCA imaging sequence has also been developed using an interleaved sequence for fast imaging, high detectable velocities, and continuous data available in the full

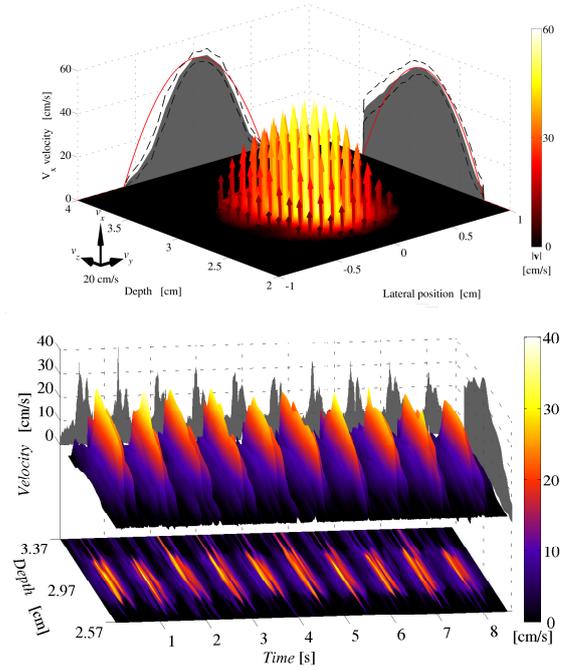


Fig. 13. Cross sectional mean 3-D vector flow averaged over 100 frames acquired using a 62+62 element RCA. The magnitude and direction of the flow is depicted by the length and the color of the arrows on the top figure. The shaded gray areas represent the projection of the flow in the respective direction with their standard deviations (dotted line). The theoretical flow profiles are illustrated by the red lines. The bottom figure shows an M-mode of the out-of-plane v_x velocity component measured for a pulsating carotid flow waveform (modified from [80]).

volume [82]. Results from simulated flow with components in all directions are shown in Fig. 14, where the vessel is rotated $45^\circ = \beta$ compared to the probe, and the beam-to-flow angles α are 90° , 75° , and 60° . All velocity components can be estimated with a bias less than -6.2% and an SD below 4.5% for situations. An example of 3-D vector flow in 4-D is shown in Fig. 15, which was measured on pulsating flow in a bifurcation phantom using the 62+62 RCA, SARUS and the SA sequence. It is possible to obtain new VFI estimates of all components and a B-mode image after 56 emissions, which yields 275 volumes/second for imaging down to 5 cm. This demonstrates that quantitative 3-D VFI can be attained in a full volume at high volume rates (4-D) using only 62 receive channels.

The continuous data for SA RC imaging can also be employed for estimating low velocities using the methods described in Section III-B. Another example is to use super resolution imaging with RC arrays and ultrasound contrast agents. An example of this is shown in Fig. 16 for flow in a micro-phantom. The 3 MHz 62+62 array was used together with a 32+32 emission SA pulse inversion sequence. The full volume was beamformed continuously, and the envelope signal was processed in a 3-D super resolution pipeline for bubble detection and presentation. A precision of roughly $20 \mu\text{m}$ was attained in all three coordinates in the full volume [83].

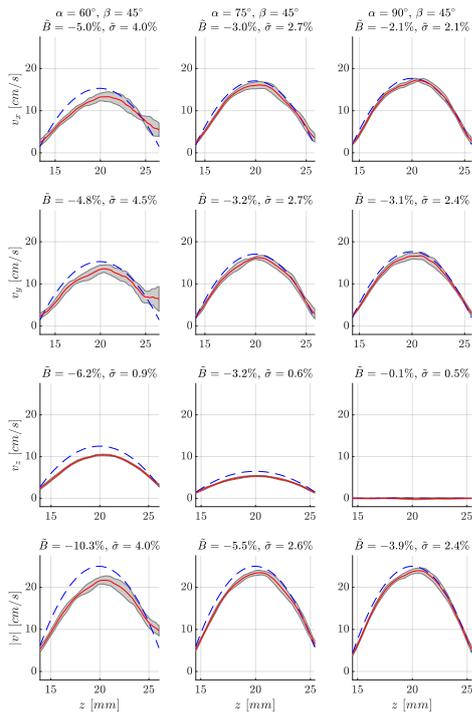


Fig. 14. Simulated velocity profiles for the SA RC flow sequence. The vessel is rotated $45^\circ = \beta$ compared to the probe, and the beam-to-flow angles α are 90° , 75° , and 60° . The true profiles are shown as dashed blue lines, the mean profiles are red, and the gray backgrounds show ± 1 SD.

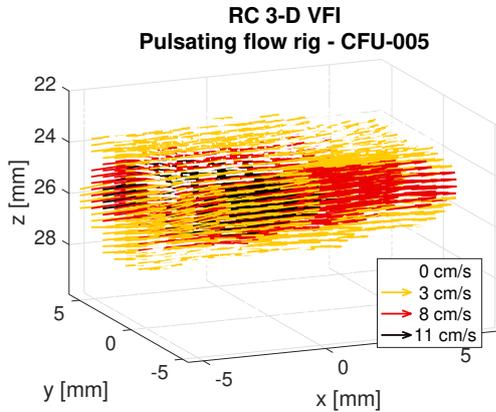


Fig. 15. Three-dimensional RC VFI from pulsating flow in a carotid artery phantom.

V. CHALLENGES AND OPPORTUNITIES

Flow imaging has progressed in the last sixty years from simple 1-D measurements to the potential of revealing the full 3-D velocity vector in a full volume in real time at very high volumetric frame rates. The development has included new imaging schemes, new estimators and progress in making advanced arrays for both 2-D and 3-D imaging.

Many challenges still lie ahead. Larger 2-D probes should be developed to fully exploit the potential of RCA SA imaging. The field-of-view should also be expanded by employing e.g.

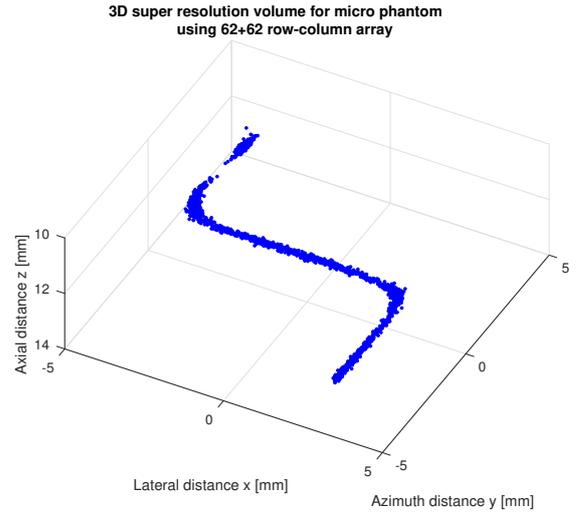


Fig. 16. Three-dimensional super resolution of of 3-D printed micro-phantom with a $200 \mu\text{m}$ diameter channel. The blue dots each indicate a single detected bubble (from [83]).

lenses on the RC array as investigated in [79]. Much research is also needed for developing imaging schemes for such arrays using sparse sets of interleaved emissions to yield the fastest imaging with the fewest emissions for an optimal contrast and resolution. The years of development has also shown that new estimators can increase precision at the same time as the number of calculations is reduced by using TO estimators. This is quite a significant point, when real time flow estimation has to be conducted in a large volume at high frame rates. Echo canceling has been an object of intense research, and the new SVD based methods are very promising for separating flow from tissue, especially when employed on the new ultrafast SA sequences.

Implementation of the processing of the data from the probes is also a problem. The data rates from RC probes are comparable to the rates already processed in commercial consoles, but the output rate is higher since a full volume has to be made. Often, several volumes have to be made from the same data at a rate of f_{prf} for flow imaging, when 3-D VFI is made.

The large amount of 3-D data being made available at fast rates is a challenge to visualize and understand in the clinic, and new display methods have to be developed in collaboration with clinicians. It is especially important to keep in mind, what is usable in the clinic, and what can improve work flow and diagnostic reliability. The further development of quantitative measures can be an avenue for improving diagnostic information. Volume flow, peak velocities, and pressure gradients might be beneficial, and their precision can be directly deduced from the data for showing diagnostic reliability.

ACKNOWLEDGMENT

This work was financially supported by grant 82-2014-4 from the Danish National Advanced Technology Foundation, by grant 7050-0004B from Innovation Fund Denmark, and from BK Medical, Herlev, Denmark.

REFERENCES

- [1] S. Satomura, "Ultrasonic Doppler method for the inspection of cardiac functions," *J. Acoust. Soc. Am.*, vol. 29, pp. 1181–1185, 1957.
- [2] —, "Study of the flow patterns in peripheral arteries by ultrasonics," *J. Acoust. Soc. Jap.*, vol. 15, pp. 151–158, 1959.
- [3] D. W. Baker, "Pulsed ultrasonic Doppler blood-flow sensing," *IEEE Trans. Son. Ultrason.*, vol. SU-17, no. 3, pp. 170–185, 1970.
- [4] P. N. T. Wells, "A range gated ultrasonic Doppler system," *Med. Biol. Eng.*, vol. 7, pp. 641–652, 1969.
- [5] K. Namekawa, C. Kasai, M. Tsukamoto, and A. Koyano, "Realtime bloodflow imaging system utilizing autocorrelation techniques," in *Ultrasound '82*, R. Lerski and P. Morley, Eds. New York: Pergamon Press, 1982, pp. 203–208.
- [6] C. Kasai, K. Namekawa, A. Koyano, and R. Omoto, "Real-Time Two-Dimensional Blood Flow Imaging using an Autocorrelation Technique," *IEEE Trans. Son. Ultrason.*, vol. 32, no. 3, pp. 458–463, 1985.
- [7] T. Loupas, J. T. Powers, and R. W. Gill, "An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the Doppler equation by means of a two-dimensional autocorrelation approach," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 42, pp. 672–688, 1995.
- [8] T. Loupas, R. B. Peterson, and R. W. Gill, "Experimental evaluation of velocity and power estimation for blood flow imaging, by means of a two-dimensional autocorrelation approach," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 42, pp. 689–699, 1995.
- [9] P. M. Embree and W. D. O'Brien, "The accurate ultrasonic measurement of volume flow of blood by time-domain correlation," in *Proc. IEEE Ultrason. Symp.*, 1985, pp. 963–966.
- [10] O. Bonnefous and P. Pesqué, "Time domain formulation of pulse-Doppler ultrasound and blood velocity estimation by cross correlation," *Ultrason. Imaging*, vol. 8, pp. 73–85, 1986.
- [11] J. A. Jensen, S. I. Nikolov, A. Yu, and D. Garcia, "Ultrasound vector flow imaging I: Sequential systems," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 11, pp. 1704–1721, 2016.
- [12] —, "Ultrasound vector flow imaging II: Parallel systems," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 11, pp. 1722–1732, 2016.
- [13] M. D. Fox, "Multiple crossed-beam ultrasound Doppler velocimetry," *IEEE Trans. Son. Ultrason.*, vol. SU-25, pp. 281–286, 1978.
- [14] B. Dunmire, K. W. Beach, K.-H. Labs., M. Plett, and D. E. Strandness, "Cross-beam vector Doppler ultrasound for angle independent velocity measurements," *Ultrasound Med. Biol.*, vol. 26, pp. 1213–1235, 2000.
- [15] B. Dunmire and K. W. Beach, "A brief history of vector doppler," in *Proc. SPIE Med. Imag.*, vol. 4325, Feb. 2001, pp. 200–214.
- [16] G. E. Trahey, J. W. Allison, and O. T. von Ramm, "Angle independent ultrasonic detection of blood flow," *IEEE Trans. Biomed. Eng.*, vol. BME-34, no. 12, pp. 965–967, 1987.
- [17] J. A. Jensen and P. Munk, "A new method for estimation of velocity vectors," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 45, no. 3, pp. 837–851, 1998.
- [18] M. E. Anderson, "Multi-dimensional velocity estimation with ultrasound using spatial quadrature," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 45, pp. 852–861, 1998.
- [19] J. A. Jensen, *Estimation of Blood Velocities Using Ultrasound: A Signal Processing Approach*. New York: Cambridge University Press, 1996.
- [20] J. W. Goodman, *Introduction to Fourier optics*, 2nd ed. New York: McGraw Hill Inc., 1996.
- [21] J. Udesen and J. A. Jensen, "Investigation of transverse oscillation method," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 53, pp. 959–971, 2006.
- [22] J. A. Jensen, "A new estimator for vector velocity estimation," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 48, no. 4, pp. 886–894, 2001.
- [23] P. M. Hansen, M. M. Pedersen, K. L. Hansen, M. B. Nielsen, and J. A. Jensen, "Demonstration of a vector velocity technique," *Ultraschall in Med.*, vol. 32, pp. 213–215, 2011.
- [24] J. A. Jensen, A. H. Brandt, and M. B. Nielsen, "Convex array vector velocity imaging using transverse oscillation and its optimization," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 12, pp. 2043–2053, 2015.
- [25] M. J. Pihl, J. Marcher, and J. A. Jensen, "Phased-array vector velocity estimation using transverse oscillations," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 59, no. 12, pp. 2662–2675, 2012.
- [26] J. A. Jensen, "Transverse spectral velocity estimation," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 11, pp. 1815–1823, 2014.
- [27] J. Udesen, M. B. Nielsen, K. R. Nielsen, and J. A. Jensen, "Examples of in-vivo blood vector velocity estimation," *Ultrasound Med. Biol.*, vol. 33, pp. 541–548, 2007.
- [28] K. L. Hansen, H. Møller-Sørensen, J. Kjaergaard, J. T. Lund, M. M. Pedersen, T. Lange, J. A. Jensen, and M. B. Nielsen, "Analysis of systolic backflow and secondary helical blood flow in the ascending aorta using vector flow imaging," *Ultrasound Med. Biol.*, vol. 42, no. 4, pp. 899–908, 2016.
- [29] P. M. Hansen, J. B. Olesen, M. J. Pihl, T. Lange, S. Heerwagen, M. M. Pedersen, M. Rix, L. Lönn, J. A. Jensen, and M. B. Nielsen, "Volume flow in arteriovenous fistulas using vector velocity ultrasound," *Ultrasound Med. Biol.*, vol. 40, no. 11, pp. 2707–2714, 2014.
- [30] K. L. Hansen, M. M. Pedersen, H. Møller-Sørensen, J. Kjaergaard, J. C. Nilsson, J. T. Lund, J. A. Jensen, and M. B. Nielsen, "Intraoperative cardiac ultrasound examination using vector flow imaging," *Ultrason. Imaging*, vol. 35, no. 4, pp. 318–332, Oct 2013.
- [31] K. L. Hansen, H. Møller-Sørensen, J. Kjaergaard, M. B. Jensen, J. T. Lund, M. M. Pedersen, T. Lange, J. A. Jensen, and M. B. Nielsen, "Intraoperative vector flow imaging using ultrasound of the ascending aorta among 40 patients with normal, stenotic and replaced aortic valves," *Ultrasound Med. Biol.*, vol. 42, no. 10, pp. 2414–2422, 2016.
- [32] K. L. Hansen, K. Juul, H. Møller-Sørensen, J. C. Nilsson, J. A. Jensen, and M. B. Nielsen, "Pediatric transthoracic cardiac vector flow imaging - A preliminary pictorial study," *Ultrasound International Open*, vol. 5, no. 1, pp. E20–E26, 2019.
- [33] P. Tortoli, M. Lenge, D. Righi, G. Ciuti, H. Liebgott, and S. Ricci, "Comparison of carotid artery blood velocity measurements by vector and standard Doppler approaches," *Ultrasound Med. Biol.*, vol. 41, no. 5, pp. 1354–1362, 2015.
- [34] J. A. Jensen, "Directional transverse oscillation vector flow estimation," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 64, no. 8, pp. 1194–1204, 2017.
- [35] K. L. Hansen, H. Møller-Sørensen, J. Kjaergaard, M. B. Jensen, J. A. Jensen, and M. B. Nielsen, "Aortic valve stenosis increases helical flow and flow complexity: A study of intra-operative cardiac vector flow imaging," *Ultrasound Med. Biol.*, vol. 43, no. 8, pp. 1607–1617, 2017.
- [36] J. Jensen, J. B. Olesen, P. M. Hansen, M. B. Nielsen, and J. A. Jensen, "Accuracy and sources of error for an angle independent volume flow estimator," *Proc. IEEE Ultrason. Symp.*, pp. 1714–1717, 2014.
- [37] J. B. Olesen, C. A. Villagómez Hoyos, N. D. Møller, C. Ewertsen, K. L. Hansen, M. B. Nielsen, B. Bech, L. Lönn, M. S. Traberg, and J. A. Jensen, "Non-invasive estimation of pressure changes using 2-D vector velocity ultrasound: An experimental study with in-vivo examples," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 5, pp. 709–719, 2018.
- [38] T.-Q. Nguyen, M. S. Traberg, J. B. Olesen, P. H. Møller-Sørensen, R. Moshavegh, L. Lönn, J. A. Jensen, M. B. Nielsen, and K. L. Hansen, "Pressure difference estimation in carotid bulbs using vector flow imaging - a phantom study," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [39] S. I. Nikolov, "Synthetic aperture tissue and flow ultrasound imaging," Ph.D. dissertation, Ørsted•DTU, Technical University of Denmark, 2800, Lyngby, Denmark, 2001.
- [40] J. Jensen, "Fast plane wave imaging," Ph.D. dissertation, DTU Elektro, Technical University of Denmark, 2017.
- [41] J. A. Jensen, S. Nikolov, K. L. Gammelmark, and M. H. Pedersen, "Synthetic aperture ultrasound imaging," *Ultrasonics*, vol. 44, pp. e5–e15, 2006.
- [42] M. Tanter and M. Fink, "Ultrafast imaging in biomedical ultrasound," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 1, pp. 102–119, January 2014.
- [43] S. I. Nikolov and J. A. Jensen, "In-vivo synthetic aperture flow imaging in medical ultrasound," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 50, no. 7, pp. 848–856, 2003.
- [44] —, "Velocity estimation using synthetic aperture imaging," in *Proc. IEEE Ultrason. Symp.*, 2001, pp. 1409–1412.
- [45] J. A. Jensen, S. I. Nikolov, T. Misaridis, and K. L. Gammelmark, "Equipment and methods for synthetic aperture anatomic and flow imaging," in *Proc. IEEE Ultrason. Symp.*, 2002, pp. 1555–1564.
- [46] N. Oddershede and J. A. Jensen, "Effects influencing focusing in synthetic aperture vector flow imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 54, no. 9, pp. 1811–1825, 2007.
- [47] L. A. F. Ledoux, P. J. Brands, and A. P. G. Hoeks, "Reduction of the clutter component in Doppler ultrasound signals based on singular value

- decomposition: A simulation study," *Ultrasound Imaging*, pp. 1–18, 1997.
- [48] L. Løvstakken, S. Bjærum, and H. Torp, "Optimal velocity estimation in ultrasound color flow imaging in presence of clutter," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 54, no. 3, pp. 539–549, March 2007.
- [49] A. C. H. Yu and L. Løvstakken, "Eigen-based clutter filter design for ultrasound color flow imaging: a review," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 57, no. 5, pp. 1096–1111, 2010.
- [50] C. Demene, T. Defieux, M. Pernot, B.-F. Osmanski, V. Biran, J.-L. Gennisson, L.-A. Sieu, A. Bergel, S. Franqui, J.-M. Correas, I. Cohen, O. Baud, and M. Tanter, "Spatiotemporal clutter filtering of ultrafast ultrasound data highly increases Doppler and fUltrasound sensitivity," *IEEE Trans. Med. Imag.*, vol. 34, no. 11, pp. 2271–2285, 2015.
- [51] J. A. Jensen and S. I. Nikolov, "Directional synthetic aperture flow imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 51, pp. 1107–1118, 2004.
- [52] J. A. Jensen and N. Oddershede, "Estimation of velocity vectors in synthetic aperture ultrasound imaging," *IEEE Trans. Med. Imag.*, vol. 25, pp. 1637–1644, 2006.
- [53] C. A. Villagomez-Hoyos, M. B. Stuart, K. L. Hansen, M. B. Nielsen, and J. A. Jensen, "Accurate angle estimator for high frame rate 2-D vector flow imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 6, pp. 842–853, 2016.
- [54] B. Y. S. Yiu, I. K. H. Tsang, and A. C. H. Yu, "GPU-based beamformer: Fast realization of plane wave compounding and synthetic aperture imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 58, no. 7, pp. 1698–1705, 2011.
- [55] J. M. Hansen, D. Schaa, and J. A. Jensen, "Synthetic aperture beam-formation using the GPU," in *Proc. IEEE Ultrason. Symp.*, 2011, pp. 373–376.
- [56] M. Lewandowski, M. Walczak, B. Witek, P. Kulesza, and K. Siewlewicz, "Modular & scalable ultrasound platform with GPU processing," *Proc. IEEE Ultrason. Symp.*, pp. 1–4, 2012.
- [57] M. B. Stuart, M. Schou, and J. A. Jensen, "Row-column beamforming with dynamic apodizations on a GPU," in *Proc. SPIE Med. Imag.*, 2019, pp. 1–7, paper number 10955-20.
- [58] T. di Ianni, K. L. Hansen, C. A. V. Hoyos, R. Moshavegh, M. B. Nielsen, and J. A. Jensen, "Portable vector flow imaging compared with spectral doppler ultrasonography," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 3, pp. 453–462, 2018.
- [59] T. Di Ianni, T. Kjeldsen, C. Villagomez-Hoyos, J. Mosegaard, and J. A. Jensen, "Real-time implementation of synthetic aperture vector flow imaging in a consumer-level tablet," in *Proc. IEEE Ultrason. Symp.*, 2017, pp. 1–4.
- [60] J. A. Jensen, "Estimation of high velocities in synthetic aperture imaging: I: Theory," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 6, pp. 1024–1031, 2019.
- [61] M. Tanter, J. Bercoff, L. Sandrin, and M. Fink, "Ultrafast compound imaging for 2-D motion vector estimation: application to transient elastography," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 49, pp. 1363–1374, 2002.
- [62] J. Udesen, F. Gran, K. L. Hansen, J. A. Jensen, C. Thomsen, and M. B. Nielsen, "High frame-rate blood vector velocity imaging using plane waves: Simulations and preliminary experiments," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 55, no. 8, pp. 1729–1743, 2008.
- [63] K. L. Hansen, J. Udesen, F. Gran, J. A. Jensen, and M. B. Nielsen, "In-vivo examples of flow patterns with the fast vector velocity ultrasound method," *Ultraschall in Med.*, vol. 30, pp. 471–476, 2009.
- [64] S. Fadnes, S. A. Nyrnes, H. Torp, and L. Lovstakken, "Shunt flow evaluation in congenital heart disease based on two-dimensional speckle tracking," *Ultrasound Med. Biol.*, vol. 40, no. 10, pp. 2379–2391, 2014.
- [65] J. A. Jensen, "Estimation of high velocities in synthetic aperture imaging: II: Experimental investigation," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 6, pp. 1032–1038, 2019.
- [66] E. Mace, G. Montaldo, I. Cohen, M. Baulac, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain," *Nature methods*, vol. 8, no. 8, pp. 662–664, 2011.
- [67] E. Mace, G. Montaldo, B. Osmanski, I. Cohen, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain: Theory and basic principles," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 3, pp. 492–506, 2013.
- [68] J. Baranger, B. Arnal, F. Perren, O. Baud, M. Tanter, and C. Demene, "Adaptive spatiotemporal SVD clutter filtering for ultrafast Doppler imaging using similarity of spatial singular vectors," *IEEE Trans. Med. Imag.*, vol. 37, no. 7, pp. 1574–1586, July 2018.
- [69] S. Holbek, K. H. Lindskov, H. Bouzari, C. Ewertsen, M. B. Stuart, C. Thomsen, M. B. Nielsen, and J. A. Jensen, "Common carotid artery flow measured by 3-D ultrasonic VFI and validated with MRI," *Ultrasound Med. Biol.*, vol. 43, no. 10, pp. 2213–2220, 2017.
- [70] M. S. Wigen, S. Fadnes, A. Rodriguez-Molares, T. Bjåstad, M. Eriksen, K. H. Stensæth, A. Støylen, and L. Løvstakken, "4-D intracardiac ultrasound vector flow imaging-reasibility and comparison to phase-contrast MRI," *IEEE Trans. Med. Imag.*, vol. 37, no. 12, pp. 2619–2629, December 2018.
- [71] M. J. Pihl and J. A. Jensen, "A transverse oscillation approach for estimation of three-dimensional velocity vectors. Part I: Concept and simulation study," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, pp. 1599–1607, 2014.
- [72] M. J. Pihl, M. B. Stuart, B. G. Tomov, M. F. Rasmussen, and J. A. Jensen, "A transverse oscillation approach for estimation of three-dimensional velocity vectors. Part II: Experimental validation," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 51, no. 10, pp. 1608–1618, 2014.
- [73] L. Ratsimandresy, P. Mauchamp, D. Dinet, N. Felix, and R. Dufait, "A 3 MHz two dimensional array based on piezocomposite for medical imaging," in *Proc. IEEE Ultrason. Symp.*, 2002, pp. 1265–1268.
- [74] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, "SARUS: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.
- [75] C. E. Morton and G. R. Lockwood, "Theoretical assessment of a crossed electrode 2-D array for 3-D imaging," in *Proc. IEEE Ultrason. Symp.*, 2003, pp. 968–971.
- [76] N. M. Daher and J. T. Yen, "2-D array for 3-D ultrasound imaging using synthetic aperture techniques," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 53, no. 5, pp. 912–924, 2006.
- [77] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, "3-D imaging using row-column-addressed arrays with integrated apodization — Part I: Apodization design and line element beamforming," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 947–958, 2015.
- [78] T. L. Christiansen, M. F. Rasmussen, J. P. Bagge, L. N. Moesner, J. A. Jensen, and E. V. Thomsen, "3-D imaging using row-column-addressed arrays with integrated apodization — part II: Transducer fabrication and experimental results," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 959–971, 2015.
- [79] H. Bouzari, M. Engholm, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Imaging performance for two row-column arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 7, pp. 1209–1221, 2019.
- [80] S. Holbek, T. L. Christiansen, M. B. Stuart, C. Beers, E. V. Thomsen, and J. A. Jensen, "3-D vector flow estimation with row-column addressed arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 11, pp. 1799–1814, 2016.
- [81] S. Holbek, T. L. Christiansen, M. F. Rasmussen, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "3-D vector velocity estimation with row-column addressed arrays," in *Proc. IEEE Ultrason. Symp.*, 2015, pp. 1–4.
- [82] M. Schou, L. T. Jørgensen, M. B. Stuart, M. S. Traberg, B. G. Tomov, and J. A. Jensen, "Full volumetric 3-D vector flow imaging using a 62+62 row-column array," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [83] J. A. Jensen, M. L. Ommen, S. H. Øygaard, M. Schou, T. Sams, M. B. Stuart, C. Beers, E. V. Thomsen, N. B. Larsen, , and B. G. Tomov, "Three-dimensional super resolution imaging using a row-column array," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, p. Submitted, 2019.

A.6 Paper F - 3-D Super Resolution Imaging using a 62+62 Elements Row-Column Array

3-D Super Resolution Imaging using a 62+62 Elements Row-Column Array

Jørgen Arendt Jensen¹, Mikkel Schou¹, Martin Lind Ommen¹, Sigrid Husebø Øygard¹,
Thomas Sams¹, Matthias Bo Stuart¹, Erik Vilain Thomsen¹,
Niels Bent Larsen¹, Christopher Beers² and Borislav Gueorguiev Tomov¹

¹Department of Health Technology,

Technical University of Denmark, DK-2800 Lyngby, Denmark

² BK Medical, 401 Science Park Road, State College, PA 16803, USA

Abstract—Current 2-D Super Resolution (SR) imaging is limited by the slice thickness determined by the elevation focus. The fixed, geometric elevation focus is often poor due to its high F-number. SR images are, thus, a summation of vessels across the elevation plane without the possibility to track scatterers in 3-D for full visualization. 3-D SR imaging has been obtained by translating the probe, but this does not remove the elevation summation. Full 3-D can be acquired using 2-D matrix probes, but the equipment is expensive, and the amount of data is excessive, when channel data are acquired over thousands of elements for minutes. This paper demonstrates that full volumetric SRI can be attained using a 62+62 channels Row-Column (RC) probe with a high frame rate and with μm precision. Data were acquired by a 3 MHz 62+62 PZT RC probe with $\lambda/2$ pitch connected to the SARUS scanner. A synthetic aperture scan sequence with 32 positive and 32 negative emissions was employed for pulse inversion (PI) imaging with an MI of 0.3. The pulse repetition frequency was 10 kHz for a 156 Hz volume rate. A PEGDA 700 g/mol based hydrogel flow-microphantom was 3-D printed by stereo-lithography. It contains a single cylindrical 200 μm diameter channel placed 3 mm from the top surface of the phantom. After a 5.8 mm long inlet, the channel bends 90° into a 7 mm long central region before bending 90° again into the 5.8 mm outlet. The flow channel was infused at 1.61 $\mu\text{L/s}$ with Sonovue (Bracco) in a 1:10 dilution. The received RF signals from the 62 row elements were beamformed with PI to yield a full volume of 15 x 15 x 15 mm³. The interpolated 3-D positions of the bubbles were estimated after local maximum detection. The reconstructed 3-D SR volume clearly shows the 200 μm channel shape with a high resolution in all three dimensions. The center line for the channel was found by fitting a line to all bubble positions, and their radial position calculated. The observed fraction of bubbles falling outside the channel was used for estimating the location precision. The precision was 16.5 μm in the $y-z$ plane and 23.0 μm in the $x-z$ plane. The point spread function had a size of 0.58 x 1.05 x 0.31 mm³, so the interrogated volume was 15,700 times smaller than for normal volumetric B-mode imaging. This demonstrates that full 3-D SRI can be attained with just 62 receive channels. The SA sequence has a low MI, but attains a large measured penetration depth of 14 cm in a tissue mimicking phantom, due to the large RC probe size. The 156 Hz volume rate also makes it possible to track high velocities in 3-D in the volume.

I. INTRODUCTION

Super resolution (SR) imaging has recently been introduced in ultrasound. The method is based on tracking the centroid of contrast agent bubbles and thereby paint an image of the

micro-vasculature [1–6]. Very high resolutions can be attained from the non-linear estimation of the target’s centroid positions, and reports of image resolution in the 10 μm range have been given [7]. The results presented are predominantly in 2-D, and the high resolution is only attained in the image plane. The out-of-plane resolution is determined by the elevation focus, which often is poor due to the fixed-focus lens. The F-number (imaging depth divided by the probe width) is usually 2 to 5 giving an ideal resolution of $2\lambda - 5\lambda$ at the focal depth and worse away from it (λ is the wavelength given by c/f_0 , where c is the speed of sound and f_0 is transducer center frequency). The images are acquired over several seconds to minutes generating Gbytes of data. Currently, most SRI is conducted using 1-D array probes due to the large amount of data generated, and that few scanners are capable of full 3-D imaging.

Visualization of 3-D SR volumes has been performed by several groups using mechanically translated linear array probes [6, 8, 9], but such a setup does not make it possible to estimate the out-of-plane location. SR has also been made using two orthogonal probes for 3-D localization in a line [10], and mechanical scanning is needed to cover a full volume. A matrix probe is, thus, needed for avoiding mechanical scanning.

Currently, the largest research scanners have 1024 channels [11, 12], and they generate around 20-50 Gbytes/s of data for 3 MHz probes, only making short acquisitions possible and precluding the use of high-frequency probes. They can handle 2-D arrays with $32 \times 32 = 1024 = N^2$ elements, which have been fabricated with $\lambda/2$ pitch. This makes them suitable for phased array imaging, but severely limits their focusing ability due to their small size and hence high F-numbers.

The problem can be somewhat alleviated by using sparse arrays, and Harput et al. [13] recently used a 512 elements sparse 2-D array based on a spiral pattern to acquire full 3-D SR imaging. Two 256 channels research scanners [14] were used for scanning of 200 μm cellulose tubes with a final localization precision of 18 μm . The main drawback of this approach is the many transducer channels needed to avoid grating lobes and the corresponding large amounts of data generated per second. Further, the probe is quite small

(\varnothing 10.4 mm), as it has to be nearly fully populated to avoid side and grating lobes, limiting the possible F-numbers.

This paper describes a 3-D SR method based on a Row-Column (RC) array with only 62+62 elements. The approach is implemented using a prototype RC array, and the imaging is conducted using the SARUS research scanner [11]. Its precision is investigated using a 3-D printed micro-phantom and is estimated from the located bubbles in the phantom.

II. METHODS

A. Data acquisition and beamforming

A prototype 3 MHz PZT RC array with 62 rows and 62 columns was used for the data acquisition [15]. It contains amplifiers in the handle and was fabricated with edge apodization to reduce ghost echoes after the main point spread function (PSF) [16]. The probe has $\lambda/2$ pitch to avoid grating lobes. It was connected to the SARUS scanner [11], which acquired full RF data for all the receiving channels.

A synthetic aperture, pulse inversion sequence was used for imaging. Transmissions were conducted using the rows, and data were received on all 62 columns. The virtual line sources emitted cylindrical waves [17] in a sequence with 32 positive emissions and 32 negative emission to make pulse inversion imaging possible. The transmit F-number was -1 using 32 Hanning apodized active elements, with the virtual source placed behind the array.

B. 3-D micro-phantom

A flow micro-phantom is fabricated for validating the approach by 3-D printing of a PEGDA 700 g/mol hydrogel using stereo-lithography, as described in [18]. The phantom measures $21.1 \times 8.16 \times 11.9$ mm³, and the voxel size of the printer is $(\Delta x, \Delta y, \Delta z) = 10.8 \times 10.8 \times 20$ μ m³. The flow micro-phantom contains a single cylindrical 100 μ m radius channel placed 3 mm from the top surface of the phantom. After a 5.8 mm long inlet, the channel bends 90° into a 7 mm long central region before bending 90° again into the 5.8 mm outlet. The flow channel is infused at 1.61 μ L/s with SonoVue (Bracco, Milano, Italy) in a 1:10 dilution, giving a peak velocity of 102.4 mm/s.

C. Processing pipeline

The beamformed volumes are processed in Matlab using our 3-D SR processing pipeline consisting of three steps. The first is to beamform the stored RF data from the SARUS scanner using the beamforming strategy described by Rasmussen et al. [16, 17] implemented in Matlab and running on an Nvidia GeForce GTX 1050 Ti (Nvidia, Santa Clara, CA, USA) GPU [19]. For the flow micro-phantom the second harmonic signal is employed, and a filter matched to the second harmonic is employed on all the received signals. The GPU beamformer was used for making the focusing of the full volumes for all emissions with an F-number of 1.5 in transmit and 1 in receive with a dynamic Hanning apodization weighting the elements. The volumes with a size of $\pm 15\lambda$ in both the x and y directions were beamformed with a line density of $\lambda/2$ covering the

full depth of the phantom. The sampling density in the z direction is $\lambda/16$. All emissions are added to generate the high resolution volume (HRV), and the positive and negative emissions HRVs are added to enhance the bubble signals.

The second step is to subtract the stationary background signal. The mean value of twenty volumes is found and subtracted from all the 400 volumes acquired. The envelope of the HRV is then found using a Hilbert transform and log compressed to a 40 dB dynamic range in relation to the data in the volume for finding locations.

The bubble locations can either be found from calculation of the centroid of local maxima, or the peak locations can be interpolated to increase the location accuracy. Experimentation with the data showed that the interpolation scheme is the most stable and accurate method, and this is the one used in this paper.

The third stage finds bubble locations by interpolating the peak position by fitting a second order polynomial to the data and then finding its interpolated maximum position x_i , as:

$$x_i = i - \frac{0.5(d(i+1, j, k) - d(i-1, j, k))}{d(i+1, j, k) - 2d(i, j, k) + d(i-1, j, k)}, \quad (1)$$

where i, j, k are the indices of the maximum and d is the envelope data for the volume. This is conducted in all three coordinates x_i, y_j, z_k with similar equations for an increased resolution in all three directions.

D. Statistical evaluation

The detected bubble locations are randomly distributed in the flow micro-phantom tube due to noise in the localization estimation, and some of them will appear to be located outside the phantom wall. The distribution of positions found can therefore yield an estimate of the localization precision. An estimate of the $y-z$ and $x-z$ precision can be obtained from the two straight segments of the 200 μ m channel phantom. In the straight segments a line is fitted to the data and considered an estimate of the center of the channel, and the distance from each bubble to the center is calculated. Assuming the measurement uncertainty in each dimension is normal distributed, the radial distribution of all bubbles in the segment will follow the distribution

$$f(r) = 2\pi r \int_{|\vec{r}_i| < R} \frac{1}{\pi R^2} \frac{1}{2\pi\sigma^2} \exp\left(-\frac{|\vec{r} - \vec{r}_i|^2}{2\sigma^2}\right) d^2 r_i, \quad (2)$$

where r is radial position, R is the radius of the tube, and σ is the standard deviation of the uncertainty. The integral is a convolution of a constant density ($1/(\pi R^2)$) with a two-dimensional Gaussian. The non-analytical integral (2) is estimated in a Monte-Carlo calculation and is a Rayleigh distribution convolved with a uniform disk distribution of radius $R = 100$ μ m. The fraction of bubbles estimated to fall outside the tube can then be translated into an estimate for the standard deviation σ (localization precision).

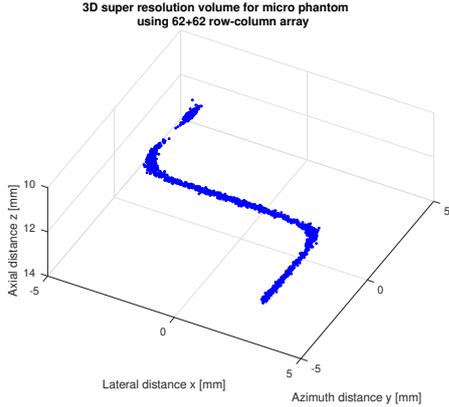


Fig. 1. Visualization of the 3-D phantom after detection of bubble locations each indicated by a blue dot.

III. RESULTS

Initially the penetration depth for the scheme is measured. It gives a penetration depth of 14 cm (0 dB signal-to-noise ratio) when using a tissue mimicking phantom with an attenuation of 0.5 dB/[MHz cm]. The SA imaging sequence and array were also simulated in Field II [20, 21] and yielded a PSF with a size of $(1.17\lambda \times 2.12\lambda \times 0.63\lambda)$ at 15 mm.

The resulting 3-D SR image is shown in Fig. 1, where each blue dot indicates the identification of a bubble. The full geometry of the phantom can be seen with the inlet and outlet and the detected bubbles seem confined to the tube.

The localization in the $y-z$ has been investigated by selecting the bubble only moving in the x direction as is shown in the top graph in Fig. 2, where blue crosses are the selected bubbles and red dots indicates all localized bubbles. Lines have then been fitted to the center of all the locations as shown in Fig. 3, so the distance from the tube center to the bubble locations can be found. The radial positions are then found and shown in Fig. 4. Bubbles inside the tube are marked by a cross and bubbles outside are marked by a red circle with a blue cross.

The fraction of bubbles outside the tube, as shown in Fig. 5, is then an indication of the precision of the bubble localization as described in Section II-D. The fraction is in this case 13.0%, which translates to a precision of $16.5 \mu\text{m}$. The fraction is 18.2% in the $x-z$ plane translating to a precision of $23.0 \mu\text{m}$. The simulated point spread function of the imaging setup at this depth is $0.58 \times 1.05 \times 0.31 \text{ mm}^3 (x, y, z)$, which corresponds to an interrogated volume of 0.189 mm^3 . Assuming the precision in all three coordinates is $23.0 \mu\text{m}$ gives a volume of $12,167 \mu\text{m}^3$, which is 15,700 times smaller than for the PSF limited system.

IV. DISCUSSION AND CONCLUSION

A 3-D SR measurement scheme and processing pipeline have been presented. The approach uses a 62+62 elements RC probe, where only rows are used for emission and columns for reception. The scheme employs two times 32 emissions for pulse inversion imaging attaining a volume rate of 240

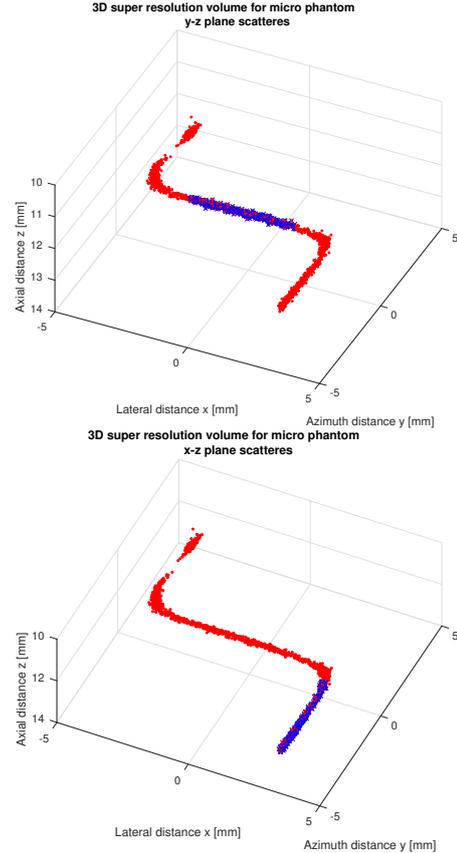


Fig. 2. Identification of bubbles only flowing in the x direction (top) and only in the y direction in the outlet (bottom). Red dots show all the bubble locations and blue crosses are the selected bubbles.

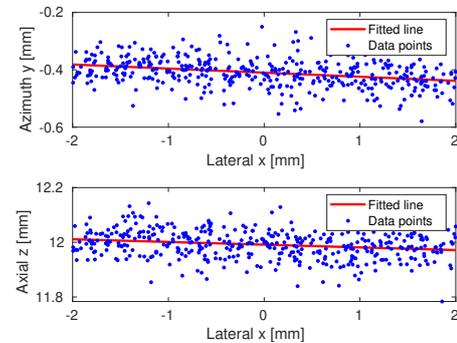


Fig. 3. Fitted center line for the bubble locations with movement in the x direction.

Hz down to 5 cm or 85 Hz down to 14 cm, which is the penetration depth of the imaging scheme.

The major advantage is that the volume is focused in all three directions including the elevation direction, which yields a resolution of $(1.17\lambda \times 2.12\lambda \times 0.63\lambda)$ at a depth of 15 mm. This was attained for a modest 62 elements, which both reduces the amount of data from the probe by a factor of 8 compared to previous 3-D SRI [13] as well as the beamforming time for a probe with 4 times the area of a 1024 elements 2-D matrix probe.

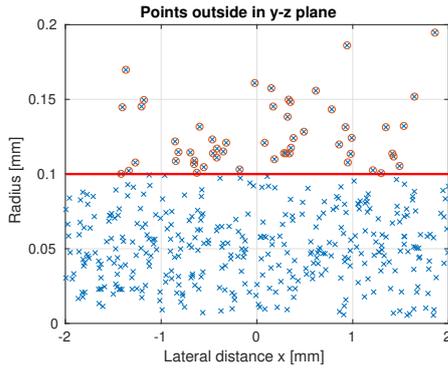


Fig. 4. Radius of the bubble position relative to the fitted center line in the $y-z$ plane. Blue crosses marks the locations of all bubbles and red circles mark bubbles outside the vessel boundary.

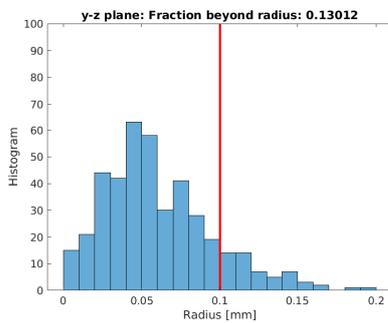


Fig. 5. Histogram for the bubble radius. The red line indicates the $100\ \mu\text{m}$ radius of the tube and the fraction of bubbles outside is 13.0%.

The processing pipeline yielded a precision of $16.5\ \mu\text{m}$ in the $y-z$ plane and $23.0\ \mu\text{m}$ in the $x-z$ plane, which is 15,700 times smaller in volume than for the PSF limited RC system. The number of active elements is three times less than what is available in commercial ultrasound consoles, and the beamforming can be attained in near real-time, when employing a state-of-the-art GPU card [22].

ACKNOWLEDGEMENT

This work was financially supported by grant 7050-00004B from Innovation Fund Denmark, and from BK Medical, Herlev, Denmark.

REFERENCES

- [1] O. Couture, B. Besson, G. Montaldo, M. Fink, and M. Tanter, "Microbubble ultrasound super-localization imaging (MUSLI)," in *Proc. IEEE Ultrason. Symp.*, 2011, pp. 1285–1287.
- [2] O. M. Viessmann, R. J. Eckersley, K. C. Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, pp. 6447–6458, 2013.
- [3] M. Siepmann, G. Schmitz, J. Bzyl, M. Palmowski, and F. Kiessling, "Imaging tumor vascularity by tracing single microbubbles," *Proc. IEEE Ultrason. Symp.*, pp. 6293–6297, 1906–1908, 2011.
- [4] M. A. O'Reilly and K. Hynynen, "A super-resolution ultrasound method for brain vascular mapping," *Med. Phys.*, vol. 40, no. 11, pp. 110701–7, 2013.
- [5] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433–440, February 2015.

- [6] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, pp. 499–502, November 2015.
- [7] O. Couture, V. Hingot, B. Heiles, P. Muleki-Seya, and M. Tanter, "Ultrasound localization microscopy and super-resolution: A state of the art," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 8, pp. 1304–1320, 2018.
- [8] J. Zhu, E. M. Rowland, S. Harput, K. Riemer, C. H. Leow, B. Clark, K. Cox, A. Lim, K. Christensen-Jeffries, G. Zhang, J. Brown, C. Dunsby, R. J. Eckersley, P. D. Weinberg, and M.-X. Tang, "3D super-resolution US imaging of rabbit lymph node vasculature in vivo by using microbubbles," *Radiology*, vol. 291, no. 3, pp. 642–650, 2019.
- [9] F. Lin, S. E. Shelton, D. Espindola, J. D. Rojas, G. Pinton, and P. A. Dayton, "3-D ultrasound localization microscopy for identifying microvascular morphology features of tumor angiogenesis at a resolution beyond the diffraction limit of conventional ultrasound," *Theranostics*, vol. 7, no. 1, pp. 196–204, 2017.
- [10] K. Christensen-Jeffries, S. Harput, J. Brown, P. N. T. Wells, P. Aljabar, C. Dunsby, M. Tang, and R. J. Eckersley, "Microbubble axial localization errors in ultrasound super-resolution imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 64, no. 11, pp. 1644–1654, 2017.
- [11] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, "SARUS: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.
- [12] L. Petrusca, F. Varray, R. Souchon, A. Bernard, J. Y. Chapelon, H. Liebgott, W. A. N'Djin, and M. Viallon, "Fast volumetric ultrasound B-mode and Doppler imaging with a new high-channels density platform for advanced 4D cardiac imaging/therapy," *Applied Sciences (Switzerland)*, vol. 8, no. 2, pp. 200:1–15, 2018.
- [13] S. Harput, K. Christensen-Jeffries, A. Ramalli, J. Brown, J. Zhu, G. Zhang, C. H. Leow, M. Toulemonde, E. Boni, P. Tortoli, R. J. Eckersley, C. Dunsby, and M. Tang, "3-D super-resolution ultrasound (SR-US) imaging with a 2-D sparse array," *arXiv preprint*, p. 1902.01608v1, 2019.
- [14] E. Boni, L. Bassi, A. Dallai, F. Guidi, V. Meacci, A. Ramalli, S. Ricci, and P. Tortoli, "ULA-OP 256: A 256-channel open scanner for development and real-time implementation of new ultrasound methods," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 10, pp. 1488–1495, 2016.
- [15] M. Engholm, H. Bouzari, T. L. Christiansen, C. Beers, J. P. Bagge, L. N. Moesner, S. E. Diederichsen, M. B. Stuart, J. A. Jensen, and E. V. Thomsen, "Probe development of CMUT and PZT row-column-addressed 2-D arrays," *Sens. Actuators A: Phys.*, vol. 273, pp. 121–133, 2018.
- [16] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, "3-D imaging using row-column-addressed arrays with integrated apodization — Part I: Apodization design and line element beamforming," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 947–958, 2015.
- [17] H. Bouzari, M. Engholm, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Imaging performance for two row-column arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 7, pp. 1209–1221, 2019.
- [18] M. L. Ommen, M. Schou, R. Zhang, C. A. V. Hoyos, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, "3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [19] M. B. Stuart, M. Schou, and J. A. Jensen, "Row-column beamforming with dynamic apodizations on a GPU," in *Proc. SPIE Med. Imag.*, 2019, pp. 1–7, paper number 10955-20.
- [20] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 39, no. 2, pp. 262–267, 1992.
- [21] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comp.*, vol. 10th Nordic-Baltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1, pp. 351–353, 1996.
- [22] M. B. Stuart, P. M. Jensen, J. T. R. Olsen, A. B. Kristensen, M. Schou, B. Dammann, H. H. B. Sørensen, and J. A. Jensen, "Fast GPU-beamforming of row-column addressed probe data," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.

A.7. PAPER G - 3D PRINTED CALIBRATION MICRO-PHANTOMS FOR VALIDATION OF SUPER-RESOLU

**A.7 Paper G - 3D Printed Calibration Micro-phantoms for
Validation of Super-Resolution Ultrasound Imaging**

3D Printed Calibration Micro-phantoms for Validation of Super-Resolution Ultrasound Imaging

*Martin Lind Ommen*¹, *Mikkel Schou*¹, *Christopher Beers*², *Jørgen Arendt Jensen*¹, *Niels Bent Larsen*¹, and *Erik Vilain Thomsen*¹

¹Department of Health Technology, Technical University of Denmark, Kgs. Lyngby, Denmark

²BK Medical, State College, Pennsylvania, USA

Abstract—This study evaluates the use of 3D printed phantoms for super-resolution ultrasound imaging (SRI) algorithm calibration. Stereolithography is used for printing calibration phantoms containing eight randomly placed scatterers of nominal size $205\ \mu\text{m} \times 205\ \mu\text{m} \times 200\ \mu\text{m}$. The purpose is to provide a stable reference for validating new ultrasonic imaging techniques such as SRI. SRI algorithm calibration is demonstrated by imaging a phantom using a $\lambda/2$ pitch 3 MHz 62+62 row-column addressed (RCA) ultrasound probe. As the imaging wavelength is larger than the dimensions of the scatterers, they will appear as single point spread functions in the generated volumes. The scatterers are placed with a minimum separation of 3 mm to avoid overlap of the point spread functions of the scatterers. 640 volumes containing the phantom features are generated, with an intervolumetric uniaxial movement of $12.5\ \mu\text{m}$, emulating a flow velocity of 2 mm/s at a volume frequency of 160 Hz. A super-resolution pipeline is applied to the obtained volumes to localise the positions of the scatterers and track them across the 640 volumes. The standard deviation of the variation in the scatterer positions along each track is used as an estimate of the precision of the super-resolution algorithm, and was found to be between the two limiting estimates of $(x, y, z) = (17.7, 27.6, 9.5)\ \mu\text{m}$ and $(x, y, z) = (17.3, 19.3, 8.7)\ \mu\text{m}$. In conclusion, this study demonstrates the use of 3D printed phantoms for determining the precision of super-resolution algorithms.

Index Terms—3D printing, stereolithography, phantom, hydrogel, calibration, resolution, ultrasound

I. INTRODUCTION

Super-resolution ultrasound imaging (SRI) has emerged over the last few years as a non-invasive method for imaging of the smallest sections of the vasculature [1], [2], [3]. This is enabled by the use of micrometer sized bubbles acting as a contrast agent and is believed to be capable of passing vessels as small as $10\ \mu\text{m}$. Tracking of individual contrast bubbles over a period of time can reveal images of the vasculature. Structures well below the diffraction limit of conventional ultrasound have been shown. However, a fundamental problem is to validate the spacial accuracy of these new techniques. Biological structures are often very complex in geometry with liquid flow and tissue motion as added complications. This means that the ground truth of the imaging system accuracy can never be obtained on biological samples. Instead, it is attained

using phantoms designed for ultrasound. With the introduction of SRI, the requirements for phantoms focus on the scale of features, to utilize the increased resolution, while still being able to replicate the dimensionality of vascular networks. Since the introduction of SRI development, the algorithm principles have been demonstrated in fairly few phantom studies. The phantoms have all contained channels with varying dimensions, fabricated using different methods. In one case, the channel dimensions were decreased to $40 \times 80\ \mu\text{m}^2$ [4] by utilizing the high resolution of silicon micro-fabrication methods. In practice, this enforces a limit of the channels at most being in a few planes, and true 3D replication of structures will be very limited. Other examples of phantoms consisting of tubes have been presented by Viessmann et al. (3 mm diameter) [5] and by Christensen-Jeffries et al. (200 μm diameter) [6]. These structures are larger than the vessels of interest, which would be smaller than $100\ \mu\text{m}$ in diameter. One of the most recent super-resolution phantom studies was made using a 2D sparse array transducer imaging two cellulose tubes arranged in a double helix pattern to create a three-dimensional phantom structure [7]. A common element of the mentioned phantom studies is that they are channel based and meant to provide a known outer limit for the localization of the tracked microbubbles. However, that does not allow for exact control of the microbubbles' positions within the channels. Therefore the spatial position of the imaging source is still not fully known.

3D printing is a promising new method for creating ultrasound phantoms. It allows for the flexibility of fabricating complex 3-dimensional structures, as well as printing of very small features in the sub- $100\ \mu\text{m}$ range [8]. Recently, the first 3D printed phantoms for ultrasound were demonstrated [9], albeit seemingly not with SRI vascular replication in mind. The phantoms contained highly scattering solid features as small as $30 \times 50\ \mu\text{m}^2$ in cross section, demonstrating the exciting potential for point spread function evaluation provided by the method. We recently demonstrated 3D printing by stereolithography to obtain small cavities and channels in a hydrogel, which is suitable for SRI ultrasound [10]. It was shown that small cavities in the 3D printed hydrogel will scatter sound, and therefore be visible in ultrasound images. Thus, in a similar manner to solid encapsulation 3D printing by Jacquet et al. [9], it is possible

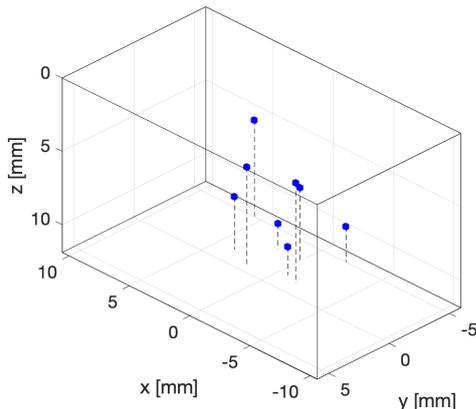


Figure 1. The designed layout of the scatterers within the $\sim 21.1 \times 11.9 \times 11.9 \text{ mm}^3$ phantom. The drop-lines are included to aid 3D perception, and end on the $z = 10 \text{ mm}$ plane.

to fixate small scatterers within the hydrogel. These structures will be stable in time, and can be imaged repeatedly, in direct contrast to microbubbles in small channels. In this paper, we demonstrate through 3D SRI how these stable structures can be used to determine the precision of the SRI hardware and algorithms. This method can potentially be used to demonstrate local variations in the precision of the algorithms based on the scatterer’s position within the localization field of view.

II. MATERIALS AND METHODS

A. Fabrication of the phantoms

Calibration phantoms were fabricated by stereolithographic 3D printing of an aqueous solution of poly(ethylene glycol) diacrylate (PEGDA, M_n 700 g/mol, 455008, Sigma-Aldrich) to form a hydrogel solid. Stereolithography is based on printing a stack of individual thin layers of materials, calling for prior digital slicing of the targeted 3D design into separate layer designs matching the printing system. The method and components have previously been presented in more detail [10]. To determine the precision of the SRI algorithms, a phantom containing eight randomly placed scatterers was printed. The outer dimensions of the phantom was $\sim 21.1 \times 11.9 \times 11.9 \text{ mm}^3$ with each scatterer being $205 \times 205 \times 200 \text{ }\mu\text{m}^3$. While the printing setup allows for printing significantly smaller scatterers, it is necessary with an increased size in order to obtain a reflection with intensity larger than that of bulk noise scatterers in the phantom. The scatterers will appear as single point spread functions in regular B-mode volumes when the imaging wavelength is larger than the scatterer size, in this case for any frequency below $\approx 6 \text{ MHz}$. They were placed with a minimum separation distance of 3 mm which will eliminate overlapping signals for any frequency above 0.5 MHz. The designed layout is shown in Fig. 1.

B. Experimental setup and procedure

A custom built stage was designed to mount the 3D printed phantom. A 3D printed frame fitted to the phantom was used to mount it on top of an absorbing polyurethane rubber sheet (Sorbothane, Inc., Kent, Ohio, USA) on the stage, which can then be submerged in water. Subsequently, the stage was mounted on a 8MR190-2-28 rotation stage (0.01° resolution) combined with a 8MTF-75LS05 x - y translation stage ($0.31 \text{ }\mu\text{m}$ resolution) (Standa, Vilnius, Lithuania), which in turn was mounted on a Newport PG Series floating optical table (Irvine, California) for stability.

Two separate experiments were conducted. The translation stage was used to move the phantom relative to the ultrasound probe along a single axis; in the first experiment along the x -axis, and in the second experiment along the y -axis. In both experiments, the inter-volume stage movement was $12.5 \text{ }\mu\text{m}$, which emulates a 2 mm/s velocity at a 160 Hz volume rate.

The imaging probe was a prototype $62 + 62$ elements 3 MHz piezo-electric, row-column addressed (RCA) transducer [11]. The probe was connected to the experimental scanner SARUS [12], a system capable of storing all channel data to be processed offline. A single frame consisted of 32 defocused emissions using a synthetic aperture (SA) imaging approach. Rows were transmitting and columns receiving, resulting in 62 channels in receive per emission. The data was passed to the SRI pipeline, described in Section II-C. The phantom was stationary when a frame was measured to avoid intra-frame motion artefacts. In total 640 volumetric frames were acquired over 640 positions. The volumetric frames were then passed to the SRI pipeline [13].

C. Super-resolution Pipeline

The super resolution pipeline has three steps. The first step is the SA beamforming. A single frame consists of a volume with dimensions of $14.86 \times 14.86 \times 7.43 \text{ mm}^3$, corresponding to $61 \times 61 \times 243$ voxels. The volume is created by beamforming all 32 emissions in a frame with a specialized beamformer [14] implemented on a GPU [15]. The 32 beamformed volumes were then summed to reveal a single high resolution volume. The volume was dynamically focused in receive (F-number of 1.5) and synthetically in transmit (F-number of 1), with an optimized sequence for SA B-mode. All 640 frames were beamformed and a stationary echo filter was applied to remove stationary tissue. However since the entire phantom is moving in this experiment, this step had no effect on the results. The third step is to locate the point scatterers from local maxima. The peak location is interpolated using a second order polynomial in all three dimensions to attain sub-pixel positioning. The output from this stage is a series of 3-D coordinates $\{x_p, y_p, z_p\}$ for all detected points. The coordinates can then be collected across all imaged frames to form tracks. The pipeline was implemented in MATLAB, and processing was performed offline [13].

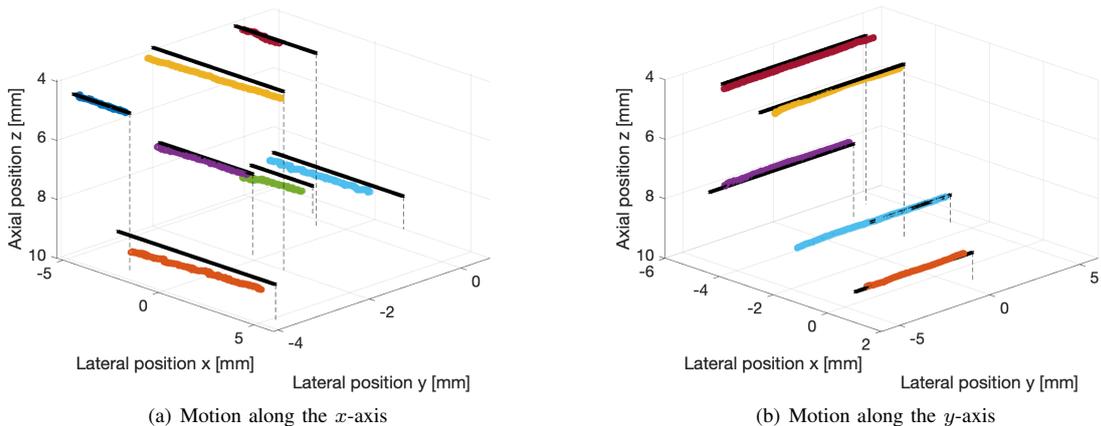


Figure 2. Cumulated localized scatterers. Both data sets are acquired over 640 volumes. In (a) the phantom was translated along the transducer x -axis. In (b) the phantom was translated along the transducer y -axis. The black tracks mark the expected linear movements based on the printed scatterer coordinates. The drop-lines are included to aid 3D perception and end on the $z = 10$ mm plane.

III. RESULTS

A. Scatterer localisation

Fig. 2 shows the cumulated localized positions of the 3D printed scatterers after acquisition of the 640 volumes. The coloured tracks mark the super-localized tracks of scatterers, while the black tracks mark the expected tracks based on the design coordinates. The horizontal field of view of the plots has been limited to the data tracks. The actual horizontal field of view of the probe is 14.86×14.86 mm². For the movement along the x -axis (Fig. 2(a)), seven of the eight printed scatterers have been correctly localized, and for the movement along the y -axis (Fig. 2(b)), five of the eight printed scatterers have been correctly localized. While the same phantom was used, differences in the point spread function of the probe based on the position within the imaged volume could mean that slight differences in the starting position of the phantom could result in one experiment correctly detecting more tracks than the other. This can also be the reason for the eighth scatterer not being detected in either experiment. Furthermore, two additional tracks which were not associated with the printed scatterers, have been omitted from the plots and analyses. A possible explanation for these localizations could be print errors, resulting in additional cavities in the phantom which act as extra scatterers. While they would be expected to follow the same trajectories as the designed scatterers, the exact shape of them is unknown. If a printing error is significantly larger than the wavelength itself, localisation of the centroid might be ambiguous. Therefore, the analysis has been limited to the designed scatterers.

B. Precision

The precision of the SRI algorithm can be determined as the localisation variation around the expected trajectories of the

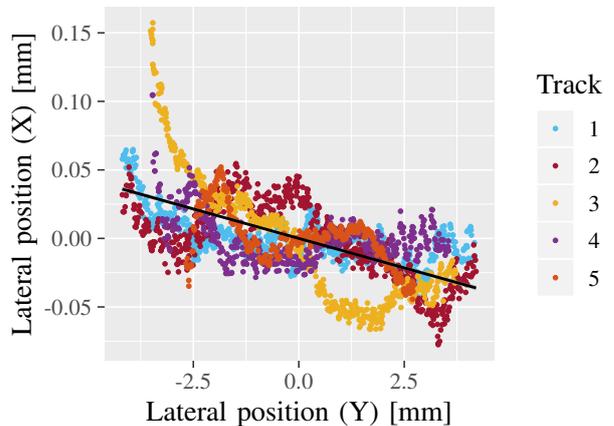


Figure 3. $y - x$ cross-plane for the tracks with motion along the y -axis. Note that the axes are not equally scaled, and the misalignment of the track movement to the transducer axes is small.

tracks. While the movement of the tracks should be solely along the x - and y -axes respectively, misalignment between the translation stage and the transducer axes could result in a slight offset of the axes. This can be seen in Fig. 3 for the tracks with motion along the y -axis, in which the $y - x$ cross-plane is shown.

The colour of the points indicate tracks from different scatterers matched to the corresponding tracks in Fig. 2(b), and the black line is the least squares fit to the localized positions, representing the average trajectory of the points, which in this case is not perfectly along the y -axis. It should be noted that the axes are not equally scaled, so the misalignment is small. A logical argument can be made for fitting to the data from all tracks combined, instead of the individual tracks, since the scatterers are fixed and are translated together, which therefore provides a good approximation for the average trajectory. However, close

Table I
ESTIMATED PRECISION FOR THE SUPER-RESOLUTION ALGORITHM.

	Average trajectory	Individual trajectory
x [μm]	17.7	17.3
y [μm]	27.6	19.3
z [μm]	9.5	8.7

examination of Fig. 3, shows that although the individual tracks roughly follow the black lines, they are not perfectly parallel, but at a small angle to each other, with the tracks crossing the average trajectory at different angles. The same tendency is observed for the x -motion tracks. This can indicate distortion in the beamforming.

The standard deviation of the residuals is used as an estimate for the precision of the SRI algorithm. The tracks with movement along the x -axis are used to estimate the precision in y and z , and the tracks with movement along the y -axis are used to estimate the precision in x and z . One of the five tracks for the y -dataset was omitted due to it not following a normal distribution, for an unknown reason. Since the tracks are not parallel, using the average trajectory would make the deviations from different tracks represent different normal distributions, and therefore not be a single characteristic property of the pipeline. Table I shows two sets of the estimated precision of the SRI pipeline averaged across all tracks: one set uses the average trajectory to determine the deviations; the other uses the individual trajectories.

IV. DISCUSSION AND CONCLUSION

We have presented a 3D printed phantom for SRI algorithm calibration. The phantom contains fixated scatterers measuring $205 \times 205 \times 200 \mu\text{m}^3$. By using fixated scatterers instead of micro-bubbles in a phantom-tube, the scatterers will be stable in time, and can therefore be used for calibration of the SRI algorithms. The phantom was imaged using a $\lambda/2$ pitch 3 MHz 62 + 62 PZT row-column addressed (RCA) probe. The beamformed volumes were processed in an SRI pipeline, tracking the scatterers as they were moved with a translation stage. The localised positions were supposed to be detected along parallel trajectories due to the movement being induced by a translation stage. However, they were found not to be parallel, which likely indicates distortion of the beamforming. Therefore, the correct estimate of the precision is likely in between the estimate using the common trajectory $(x, y, z) = (17.7, 27.6, 9.5) \mu\text{m}$, and the individual trajectories $(x, y, z) = (17.3, 19.3, 8.7) \mu\text{m}$. These results demonstrate the novel use of 3D printed phantoms for calibration of SRI algorithms.

ACKNOWLEDGMENT

This work was financially supported by grant 7050-00004B from Innovation Fund Denmark, and from BK Medical, Herlev,

Denmark.

REFERENCES

- [1] K. Christensen-Jeffries, R. J. Browning, M. X. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imaging*, vol. 34, no. 2, p. 6908009, 2015.
- [2] M. A. O'Reilly and K. Hynynen, "A super-resolution ultrasound method for brain vascular mapping," *Med. Phys.*, vol. 40, no. 11, pp. 1–7, 2013.
- [3] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, no. 7579, pp. 499–502, 2015.
- [4] Y. Desailly, O. Couture, M. Fink, and M. Tanter, "Sono-activated ultrasound localization microscopy," *Appl. Phys. Lett.*, vol. 103, no. 17, p. 174107, 2013.
- [5] O. M. Viessmann, R. J. Eckersley, K. Christensen-Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, no. 18, pp. 6447–6458, 2013.
- [6] K. Christensen-Jeffries, J. Brown, P. Aljabar, M. Tang, C. Dunsby, and R. J. Eckersley, "3-D in Vitro Acoustic Super-Resolution and Super-Resolved Velocity Mapping Using Microbubbles," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 64, no. 10, pp. 1478–1486, 2017.
- [7] S. Harput, K. Christensen-Jeffries, J. Brown, J. Zhu, G. Zhang, C. H. Leow, M. Toulemonde, A. Ramalli, E. Boni, P. Tortoli, R. J. Eckersley, C. Dunsby, and M. X. Tang, "3-D Super-Resolution Ultrasound Imaging Using a 2-D Sparse Array with High Volumetric Imaging Rate," *IEEE Int. Ultrason. Symp. IUS*, vol. 2018-October, pp. 23–26, 2018.
- [8] H. Gong, B. P. Bickham, A. T. Woolley, and G. P. Nordin, "Custom 3D printer and resin for $18 \mu\text{m} \times 20 \mu\text{m}$ microfluidic flow channels," *Lab Chip*, vol. 17, no. 17, pp. 2899–2909, 2017. [Online]. Available: <http://dx.doi.org/10.1039/C7LC00644F>
- [9] J. R. Jacquet, F. Ossant, F. Levassort, and J. M. Grégoire, "3-D-Printed Phantom Fabricated by Photopolymer Jetting Technology for High-Frequency Ultrasound Imaging," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 65, no. 6, pp. 1048–1055, 2018.
- [10] M. L. Ommen, M. Schou, R. Zhang, C. A. Villagomez Hoyos, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, "3D Printed Flow Phantoms with Fiducial Markers for Super-Resolution Ultrasound Imaging," *IEEE Int. Ultrason. Symp. IUS*, vol. 2018-October, p. 8580217, 2018.
- [11] H. Bouzari, M. Engholm, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Imaging Performance for Two Row-Column Arrays," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 66, no. 7, pp. 1209–1221, 2019.
- [12] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domstien, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, "Sarus: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 60, no. 9, pp. 6 587 394, 1838–1852, 2013.
- [13] J. A. Jensen, M. L. Ommen, S. H. Øygaard, M. Schou, T. Sams, M. B. Stuart, C. Beers, E. V. Thomsen, N. B. Larsen, and B. G. Tomov, "Three-dimensional super resolution imaging using a row-column array," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, p. Accepted for publication, 2019.
- [14] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, "3-d imaging using row-column-addressed arrays with integrated apodization. part i: Apodization design and line element beamforming," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 62, no. 5, pp. 947–958, 2015.
- [15] M. B. Stuart, M. Schou, and J. A. Jensen, "Row-column beamforming with dynamic apodizations on a gpu," *Proceedings of Spie*, vol. 10955, p. 109550Q, 2019.

A.8 Paper H - Three-Dimensional Super Resolution Imaging using a Row-Column Array

Three-Dimensional Super Resolution Imaging using a Row-Column Array

Jørgen Arendt Jensen¹, Martin Lind Ommen¹, Sigrid Husebø Øygaard¹, Mikkel Schou¹,
Thomas Sams¹, Matthias Bo Stuart¹, Christopher Beers², Erik Vilain Thomsen¹,
Niels Bent Larsen¹ and Borislav Gueorguiev Tomov¹

¹Department of Health Technology,

Technical University of Denmark, DK-2800 Lyngby, Denmark

² BK Medical, 401 Science Park Road, State College, PA 16803, USA

Abstract—A 3-D super resolution (SR) pipeline based on data from a Row-Column (RC) array is presented. The 3 MHz RC array contains 62 rows and 62 columns with a half wavelength pitch. A Synthetic Aperture (SA) pulse inversion sequence with 32 positive and 32 negative row emissions are used for acquiring volumetric data using the SARUS research ultrasound scanner. Data received on the 62 columns are beamformed on a GPU for a maximum volume rate of 156 Hz, when the pulse repetition frequency is 10 kHz. Simulated and 3-D printed point and flow micro-phantoms are used for investigating the approach. The flow micro-phantom contains a 100 μm radius tube injected with the contrast agent SonoVue. The 3-D processing pipeline uses the volumetric envelope data to find the bubble's positions from their interpolated maximum signal and yields a high resolution in all three coordinates. For the point micro-phantom the standard deviation on the position is (20.7, 19.8, 9.1) μm (x, y, z). The precision estimated for the flow phantom is below 23 μm in all three coordinates, making it possible to locate structures on the order of a capillary in all three dimensions. The RC imaging sequence's point spread function has a size of $0.58 \times 1.05 \times 0.31 \text{ mm}^3$ ($1.17\lambda \times 2.12\lambda \times 0.63\lambda$), so the possible volume resolution is 28,900 times smaller than for SA RC B-mode imaging.

I. INTRODUCTION

Ultrasound super resolution imaging (SRI) was introduced by a number of groups for increasing the resolution of ultrasound imaging beyond the diffraction limit [1–6]. The approach is based on injection of a diluted ultrasound contrast agent to enable tracking of individual bubbles. The centroids of the bubble signals are calculated, and their tracks are determined and displayed to show an image of the vasculature. This can reveal the micro vasculature down to vessel sizes of 10 μm [7]. The images are acquired over several seconds to minutes generating Gbytes of data. Currently most SRI is conducted using 1-D array probes due to the large

amount of data, and that few scanners are capable of full 3-D imaging. The 2-D SR images therefore have a high resolution in the imaging plane, but localization in the elevation direction is not possible. 2-D SRI therefore displays a summation of vessels in the elevation plane.

Visualization of 3-D SR volumes has been performed by several groups using mechanically translated linear array probes [6, 8, 9], but such a setup does not make it possible to estimate the out-of-plane location. SR has also been made using two orthogonal probes for 3-D localization in a line [10], and mechanical scanning is needed to cover a full volume. A matrix probe is, thus, needed for avoiding mechanical scanning.

Currently, the largest research scanners have 1024 channels [11, 12], and they generate around 20–50 Gbytes/s of data for 3 MHz probes, only making short acquisitions possible and precluding the use of high-frequency probes. They can handle 2-D arrays with $32 \times 32 = 1024 = N^2$ elements, which have been fabricated with $\lambda/2$ pitch (λ is the wavelength given by c/f_0 , where c is the speed of sound and f_0 is transducer center frequency). This makes them suitable for phased array imaging, but severely limits their focusing ability due to their small size and hence high F-numbers (imaging depth divided by the probe width).

The problem can be somewhat alleviated by using sparse arrays, and Harput et al. [13] recently used a 512 elements sparse 2-D array based on a spiral pattern to acquire full 3-D SR imaging. Two 256 channels research scanners [14] were used for scanning of 200 μm cellulose tubes with a final localization precision of 18 μm . The main drawback of this approach is the many transducer channels needed to avoid grating lobes and the corresponding large amounts of data generated per second. Further, the probe is quite small (\varnothing 10.4 mm), as it has to be nearly fully populated to avoid side and grating lobes, limiting the possible F-numbers.

An new approach is therefore needed for 3-D SR volumetric imaging. One possibility for reducing the number of elements by a factor of $N/2$ is the employment of Row-Column (RC) arrays as introduced by Morton and Lockwood [15], and later investigated by a number of groups [16–20]. Here, the array is addressed by either its rows or columns, and imaging can be conducted using synthetic aperture (SA) imaging schemes [21] for both a high resolution, deep penetration depth, and high volume rate. Furthermore, RC SA imaging schemes can have a low mechanical index due to the emission of cylindrical waves, making them ideally suited for contrast agent imaging. The RC arrays can be made large without having an excessive amount of elements, making it possible to both have low F-numbers for high resolutions and still have modest data rates from the arrays.

This paper presents a 3-D SR imaging method using a prototype 62 + 62 RC array [22] connected to the SARUS research scanner [11]. Two 3-D printed micro phantoms are used for validating the approach along with simulation of a point phantom. The precision of the pipeline is revealed from these simulations and measurements.

II. METHODS

This Section describes the various methods used in the 3-D SR pipeline including the imaging scheme, processing pipeline, and statistical evaluation.

A. Imaging sequence and processing

The imaging sequence was optimized for a 62+62 RC PZT 3 MHz experimental probe with dimensions given in Table I. The probe includes a mechanical apodization at each end of the elements to reduce edge element artifacts as described in previous publications on the probe [19–22]. The volumetric RC SA imaging scheme consists of 32 virtual focus lines using 32 active elements per emissions. An F-number of -1 was used for emitting de-focused line sources with the focal point placed behind the probe surface and with a Hanning weighting to reduce side-lobes. The 32 different virtual lines were placed to generate a sliding aperture imaging sequence across the rows. Transmission were only made with the rows and reception was made with the column elements, resulting in 62 signals to be stored per emission. Pulse inversion imaging was conducted by emitting two 2 cycle sinusoidal 3 MHz waves, one positive and one negative, for each virtual line source. The imaging sequence was implemented on the the SARUS experimental scanner [11], with a transmit sampling frequency of 70 MHz. The

TABLE I
RC 62 + 62 PZT PROBE DIMENSIONS.

Parameters	62+62 RC
Number of elements	62+62
Center frequency f_0	3 MHz
Wavelength λ	513 μm
Kerf	25 μm
Pitch	270 μm ($\approx \lambda/2$)
Apodization region length	4.05 mm
Element length	24.84 mm
Total Active Surface area	282.2 mm ²

receive sampling frequency was 23.7 MHz to preserve the second harmonic component in the signal.

Each emission was beamformed using a MATLAB based GPU beamformer [23] to generate a low resolution volume (LRV). The 32 different LRVs were summed to reveal a high resolution volume (HRV). A simplified schematic of the sequence can be seen in Fig. 1. The positive +LRV(1:32) and negative -LRV(1:32) beamformed emissions were summed to reveal a second harmonic HRV using pulse inversion. This data was then passed to the SRI processing pipeline described in Section II-B. The pulse repetition frequency (f_{prf}) was 10 kHz, and a pause of 10 ms was inserted between volumes to reduce the memory usage and extend the duration of the acquisition. The Mechanical Index (MI) of the sequence was determined to be 0.2 at 12 mm from the probe surface, which is the location of the micro-phantom flow channel. The actual MI in the phantom is probably slightly lower due to the attenuation in the phantom.

Three sets of measurements were performed using a precision translation stage. The RC probe was mounted on a Newport PG Series floating optical table (Irvine, California) for stability with the micro-phantoms mounted on a 8MR190-2-28 rotation stage combined with a 8MTF-75LS05 $x-y$ translation stage (Standa, Vilnius, Lithuania). These were used to align both micro-phantoms with the imaging axis, and to generate translation in the $x-y$ plane used for the validation described in Section III-B.

B. Processing pipeline

The processing pipeline consists of several stages. The first step is to beamform the stored RF data from the SARUS scanner using the beamforming strategy described by Rasmussen et al. [19, 21] implemented in Matlab and running on an Nvidia GeForce GTX 1050 Ti (Nvidia, Santa Clara, CA, USA) [23]. A volume with a size of $\pm 15\lambda$ in both the x and y directions are beamformed with a line density of $\lambda/2$ and covering the depth of the phantom. The sampling density in the z direction is

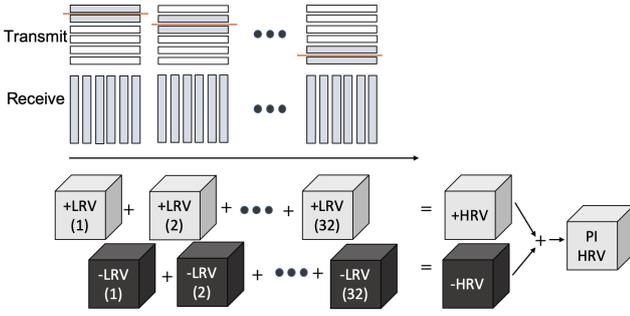


Fig. 1. The transmitting row elements and their translation across the aperture is shown in the top figure along with the receiving column elements. The time sequence of the positive and negative emissions and their combination is shown in the bottom figure.

$\lambda/16$. A matched filter is applied on the received signals. It is designed using the measured impulse response of the probe to match the first harmonic signal found in the linearly simulated data. The same filter is also used for PSF phantom. The positive and negative emissions are then subtracted to increase the signal-to-noise ratio (SNR). For the flow micro-phantom the second harmonic signal is employed, the filter is matched to this, and the two emissions are added. The full LRV is beamformed for all emissions with an F-number of 1.5 in transmit and 1 in receive with a dynamic Hanning apodization weighting the elements, and all emissions are added to generate the HRV. The mean value of the first 20 HRVs are averaged and subtracted from the processed HRVs to remove stationary objects in the processing. The envelope of the HRV is then found using a Hilbert transform and log compressed to a 40 dB dynamic range in relation to the data in the volume for finding locations.

The peak location can either be found from calculation of the centroid of a global maximum, or the peak location can be interpolated to increase the location accuracy. Experimentation with the data has shown that the interpolation scheme is the most stable and accurate method, and this is the one used in this paper.

The second stage finds bubble locations by interpolating the peak position by fitting a second order polynomial to the data and then finding its interpolated maximum position x_i , as:

$$x_i = i - \frac{0.5(d(i+1, j, k) - d(i-1, j, k))}{d(i+1, j, k) - 2d(i, j, k) + d(i-1, j, k)}, \quad (1)$$

where i, j, k are the indices of the maximum and d is the envelope data for the volume. This is conducted in all three coordinates x_i, y_j, z_k with similar equations for an increased resolution in all three directions.

The third step used only on the point micro-phantom finds contiguous tracks of target locations. A target in a

first HRV is used as a reference point, and the adjacent HRV is searched to find a detected target location lying within a radius of $r = v_s/f_r$ from the reference, where f_r is the volume rate and v_s is the maximum search velocity, where $v_s = 10$ mm/s was used. The track is terminated, if no target is found, and the whole track is discarded, if it does not contain more than 200 contiguous locations. No tracks were formed for the micro-flow phantom due to the high velocity employed, and all bubble locations in all images are shown in Section III-C.

C. Simulations and measurement phantoms

The method is evaluated using both simulations and measurements from two 3-D printed micro-phantoms, which are all described in this Section. The penetration depth is also determined from measurements on a tissue mimicking phantom with a 0.5 dB/[MHz cm] attenuation.

1) *Simulation of 3-D SRI system:* The SA RC sequence has been simulated using Field IIpro [24–26] to generate reference data, where the positions of the scatterers are known in the volume. The phantom contains a number of point targets located at a depth of 5, 15, and 25 mm at the center axis of the probe. It is used for determining the point spread function (PSF) of the imaging method.

2) *Fabrication of micro-phantoms:* Two micro-phantoms have been made and used for validating the approach. Both have been fabricated by 3-D printing of a PEGDA 700 g/mol hydrogel using stereo-lithography. The phantoms measure $21.1 \times 8.16 \times 11.9$ mm³, and the voxel size of the printer is $(\Delta x, \Delta y, \Delta z) = 10.8 \times 10.8 \times 20$ μ m³. More information about the fabrication process can be found in [27].

The first point phantom contains eight markers with a size of $10.8 \times 10.8 \times 20 = 233$ μ m³. The marker sizes are in all dimensions smaller than the imaging wavelength of 500 μ m for the RC probe used, resulting in markers appearing as single targets in the B-mode volume. The markers are positioned with a minimum distance of 3 mm to ensure a clear separation of the reflected signals. The phantom is moved relative to the ultrasound probe using the $x-y$ translation stage in two experiments, one along x and one along y . An inter-volume movement of 12.5 μ m is used to emulate a constant velocity of 1.95 mm/s at 156 Hz. After each movement the positions of the markers are determined and tracks for the targets are made.

The second flow micro-phantom contains a single cylindrical 100 μ m radius channel placed 3 mm from the top surface of the phantom. After a 5.8 mm long

inlet, the channel bends 90° into a 7 mm long central region before bending 90° again into the 5.8 mm outlet. The flow channel was infused at 1.61 $\mu\text{L/s}$ with SonoVue (Bracco, Milano, Italy) in a 1:10 dilution giving a peak velocity of 102.4 mm/s.

D. Statistical evaluation

The bubble locations are randomly distributed in the flow micro-phantom tube due to noise in the localization estimation, and some of them will appear to be located outside the phantom wall. The distribution of positions found can then yield an estimate of the localization precision. An estimate of the $y-z$ and $x-z$ precision can be obtained from the two straight segments of the 200 μm channel phantom. In the straight segments a line is fitted to the data and considered an estimate of the center of the channel, and the distance from each bubble to the center is calculated. Assuming the measurement uncertainty in each dimension is normal distributed, the radial distribution of all bubbles in the segment will follow the distribution

$$f(r) = 2\pi r \int_{|\vec{r}_i| < R} \frac{1}{\pi R^2} \frac{1}{2\pi\sigma^2} \exp\left(\frac{-|\vec{r} - \vec{r}_i|^2}{2\sigma^2}\right) d^2r_i \quad (2)$$

where r is radial position, R is the radius of the tube, and σ is the standard deviation of the uncertainty. The integral is a convolution of a constant density ($1/(\pi R^2)$) with a two-dimensional Gaussian. The non-analytical integral (2) is estimated in a Monte-Carlo calculation and is a Rayleigh distribution convolved with a uniform disk distribution of radius $R = 100 \mu\text{m}$. The factor $2\pi r$ is the Jacobian needed to convert from Cartesian to cylindrical coordinates. The fraction of bubbles estimated to fall outside the tube can then be translated into an estimate for the standard deviation σ (localization precision), as is performed in Section III-C1

III. RESULT

A. Imaging performance

The performance of the imaging scheme has both been simulated and measured. The response from several point scatterers were simulated using Field II, and the FWHM was determined for the first harmonic signal to be ($\text{FWHM}_x, \text{FWHM}_y, \text{FWHM}_z = 0.58 \times 1.05 \times 0.31 \text{ mm} = (1.17\lambda \times 2.12\lambda \times 0.63\lambda)$) at a depth of 15 mm. The receive F-number is 1, so the PSF is close to the theoretical limit of λ . The transmit F-number is 1.5 and, thus, gives a slightly wider PSF in the y -direction along emissions.

The penetration depth of the scheme was determined using a uniformly scattering phantom model 571 (Danish

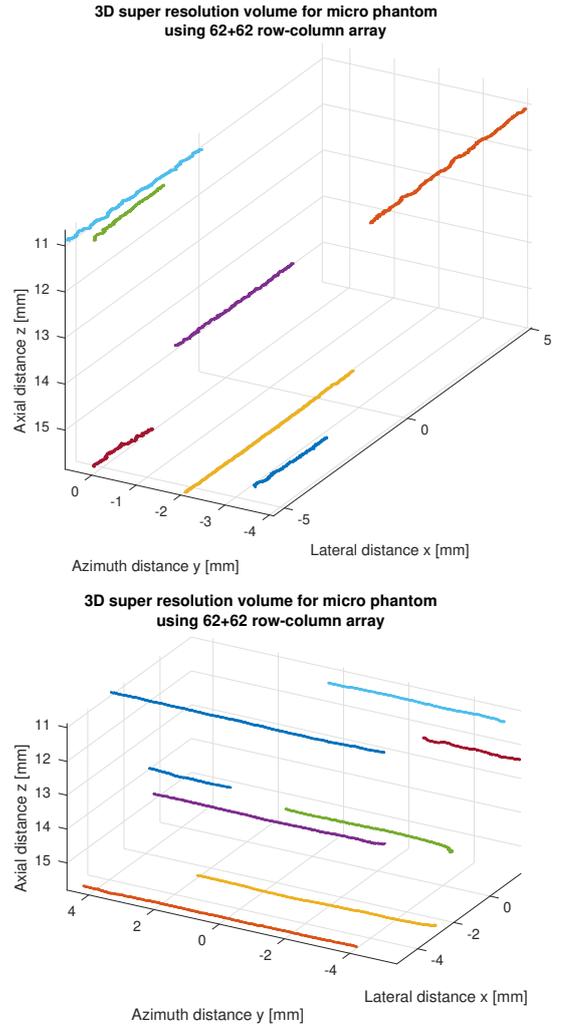


Fig. 2. Tracks estimated from mechanical translation of the PSF micro phantom where the colors indicate detected positions. The top graph is for translation in the x -direction and bottom for translation in y .

Phantom Service, DK-3600 Frederikssund, Denmark) with a speed-of-sound of 1540 m/s and a uniform attenuation of 0.5 dB/[MHz cm]. Determining the SNR from ten independent measurements gave a penetration depth of 14 cm (SNR=0 dB).

B. Validation in point micro-phantom

Fig. 2 shows the cumulative localized positions of 3-D printed markers within the micro phantom acquired over 640 beamformed volumes at the emulated speed of 1.95 mm/s. The top figure shows movement in the x -direction and in the y -direction at the bottom. Seven markers have been detected and are shown as colored points. The eighth marker was too weak to be detected. Lines are fitted to the positions using a least squares fit, as shown in Fig. 3 for one of the tracks. The deviations

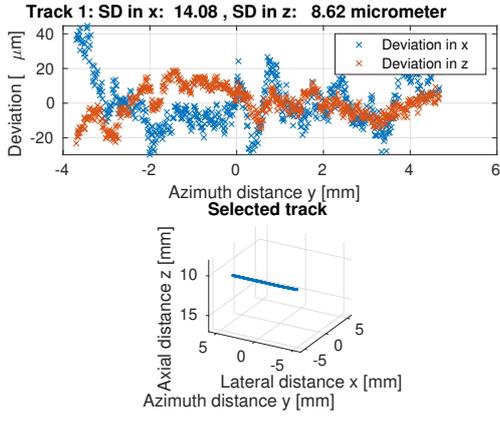


Fig. 3. Deviations calculated for one of the tracks, when a line has been fitted to the data. The bottom graph shows the estimated target locations, and the top graph shows the deviations in x and z , when the line has been subtracted from the target position.

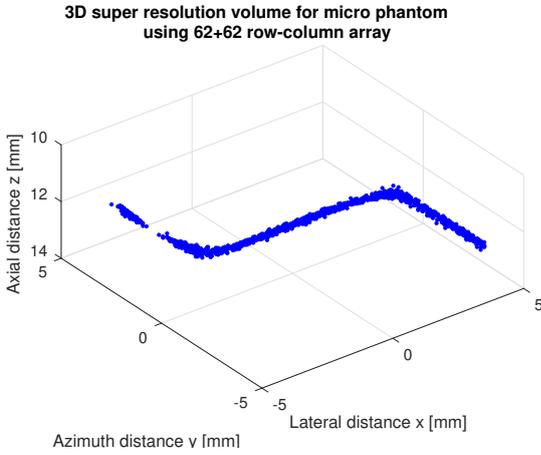


Fig. 4. View of the $200\ \mu\text{m}$ channel phantom. Each blue dot represents a detected contrast agent bubble. See accompanying video for a 3-D view of the phantom.

from the fitted line are calculated and the standard deviations is estimated to $(\sigma_x, \sigma_y, \sigma_z) = (20.7, 19.8, 9.1)\ \mu\text{m}$, when taking the average across all tracks.

C. 3-D SRI imaging

The measured data from the flow micro-phantom acquired from 400 frames of the SA imaging sequence has been processed by the SR pipeline, including beam-forming and detection of bubble locations, using the interpolation scheme in (1). A 3-D view of the detected bubbles is shown in Fig. 4, where each blue dot is a detected bubble. The geometry of the phantom can clearly be seen.

1) *Precision of bubble locations:* Bubbles in the central part of the phantom ($-2\ \text{mm} < x < 2\ \text{mm}$) have been selected for estimating the localization precision in

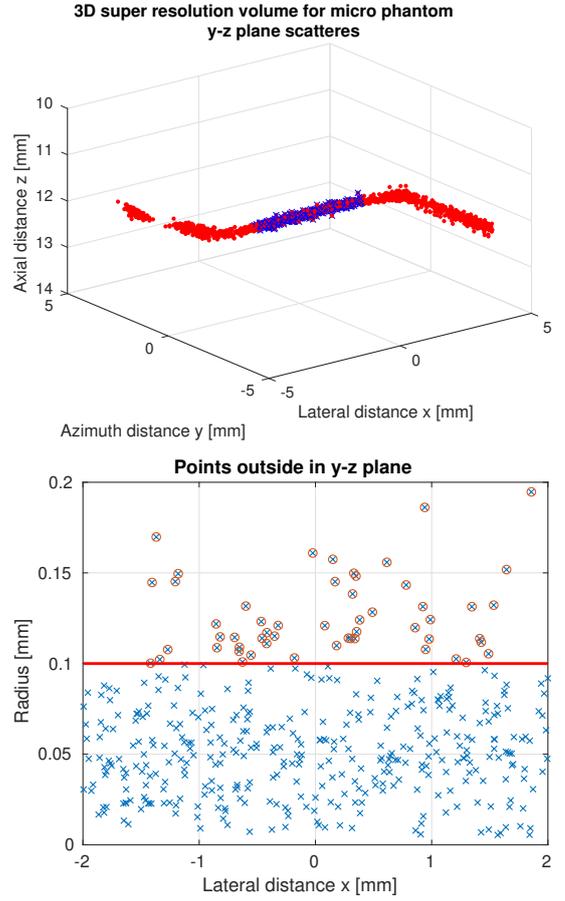


Fig. 5. Selected bubbles in the $y-z$ plane (blue crosses) for estimating precision (top graph), and the calculated radial position of the bubbles in the vessel (bottom graph). Blue crosses indicate bubbles inside the vessel and red circles indicates outside.

the $y-z$ plane as shown on the top in Fig. 5, where blue crosses indicate bubbles used for this estimation. Center lines for all selected bubbles are estimated with a least square fit as shown in Fig. 6. The channels center depth is at $12.0\ \text{mm}$ from the probe, and the channel is slightly rotated in the $x-y$ plane ($57\ \mu\text{m}$ tilt of $x-y$ line in the top graph). These lines are used for calculating the radial positions of the bubbles in the vessel, as shown on the bottom figure in Fig. 5. Here, a blue cross indicates bubbles inside the vessel, and a red circle indicates bubbles outside of the vessel boundary, shown as the solid red line ($r = 100\ \mu\text{m}$). The same graph for the $x-z$ plane is shown in Fig. 7, where the outlet part of the vessel has been employed for finding the precision ($-5.5\ \text{mm} < y < -1\ \text{mm}$).

The fraction of bubbles estimated to fall outside the tube can then be translated into an estimate for the standard deviation as described in Section II-D. The fitted distribution for the bubble locations in the $y-z$ plane is shown in Fig. 8. For the $y-z$ plane 13% of the

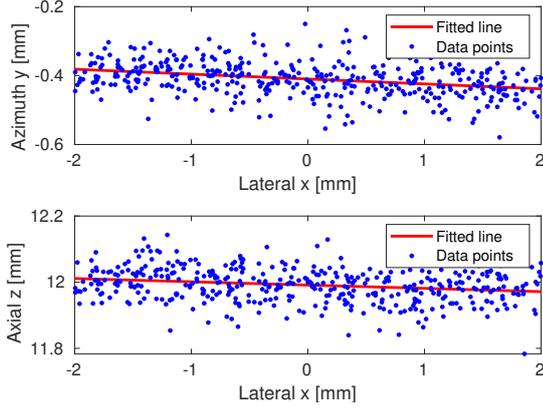


Fig. 6. Fitted line for calculating the center position of the vessel in the phantom in the $y-z$ plane.

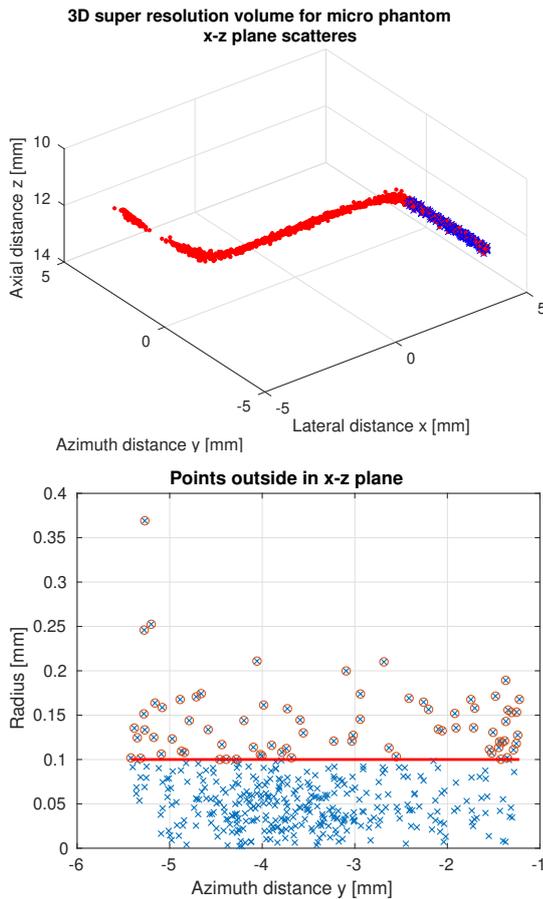


Fig. 7. Selected bubbles in the $x-z$ plane (blue dots) for estimating precision.

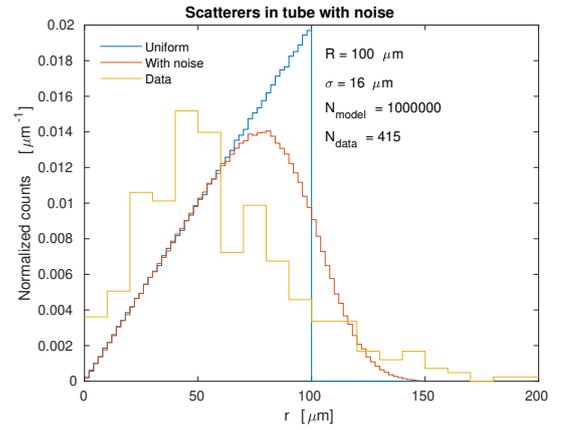


Fig. 8. Fitted distribution for the bubble locations in the $y-z$ plane.

bubbles fall outside the tube, which leads to a position uncertainty of $16.5 \mu\text{m}$. Similarly, for the $x-z$ plane 18% of the bubbles are estimated to fall outside the tube, which leads to a position uncertainty of $23 \mu\text{m}$.

IV. DISCUSSION

A method for 3-D super resolution imaging has been developed based on a RC array and a pulse-inversion SA imaging sequence using 32 positive and negative emissions. A full volume is, thus, created in 64 emissions for a possible volume rate of 156 Hz at $f_{prf} = 10 \text{ kHz}$, and the modest number of emissions makes it possible to have a 100 Hz volume rate down to a depth of 12 cm. The 3 MHz array's penetration depth is 14 cm due to its low frequency and fairly large size ($31\lambda \times 31\lambda$). Only 62 elements were employed during receive, making it possible to implement the approach on a standard ultrasound console with the advantage that a limited amount of data is generated. A $\lambda/8$ sampling density on the receiving elements can be employed and will continuously generate 2.9 Gbytes/s, which is well within reach of modern ultrasound research scanners [11, 14, 28–30]. This is significantly less than for a fully populated array, where a 32×32 array yields 49 Gbytes/s for an array with one fourth the area of the RC probe used here.

The attained precision of the schemes was investigated using both a point micro-phantom and a flow micro-phantom with a $200 \mu\text{m}$ diameter tube. The point phantom yielded a localization precision of $(20.7, 19.8, 9.1) \mu\text{m}$ in the x, y, z coordinates. The flow micro-phantom yielded an estimated radial precision of $16.5 \mu\text{m}$ in the $y-z$ plane, and $23 \mu\text{m}$ in the $x-z$ plane. Assuming the coordinate precisions are independent, the radial precision would be $15.4 \mu\text{m}$ in the $y-z$ plane, and $16.0 \mu\text{m}$ in the $x-z$ plane, when using the estimated precisions from the point phantoms. The $10.8 \times 10.8 \times 20$

μm^3 voxel size of the printer will give rise to tube-wall fluctuations, with an increase in precision, so the estimated precision for the two phantoms are therefore similar.

The precision should be compared to the emitted wavelength of $500 \mu\text{m}$, and an improved localization of a factor of at least 20 times is attained in all three coordinates. The measured PSF has a size of $0.58 \times 1.05 \times 0.31 \text{ mm}^3$ making it in theory possible to interrogate a volume 28,900 times smaller than a PSF limited system.

The main advantage of a 3-D system compared to the current 2-D systems is the increased resolution in the elevation plane. Current 2-D SR displays images averaged across the elevation plane thickness, which can often be $5\text{-}15 \lambda$ away from the elevation focus. The resolution is, thus, improved by a factor 100-300 times compared to a 1D probe, even though the number of elements is 3 times lower than a 192-elements 1D probe.

Several factors can be improved in the current setup and should be incorporated into a clinically useful 3-D SR imaging scheme. Currently, no motion correction is conducted, but the SA imaging scheme makes it possible to beamform a full volume at more than 100 Hz. This is sufficient to employ speckle tracking [31] in 3-D to yield and compensate for the motion as described for SA flow imaging [32, 33]. Although many schemes use very high frame rates with thousands images per second [6, 9], it has been shown that a conventional linear array scan with frame rates at 54 Hz can yield excellent super resolution images with both motion estimation [34] and quantification of flow [35]. The 154 Hz volume rate should, thus, be sufficient for in-vivo imaging.

The fairly high velocity of 102.4 mm/s in the flow micro-phantom is used to prevent clogging of the phantom. This currently prevents the formation of tracks in the SR pipeline as is done for the PSF micro phantom, but further experiments should be conducted to lower the velocity, and maybe introduced a phantom with less sharp bends to prevent clogging. No efforts have been made to reduce false detections in the SR pipeline. Forming long tracks can significantly reduce the number of false detections, and this could potentially improve the precision of the location estimates.

The numbers of bubbles used here was sparse to make isolation easier. The acquisition length could be reduced, if more detections could be made per second. Methods for increasing the number of bubble detections have been the topic of a number of articles [36–38]. Such approaches can also be employed here, as full RF data are acquired and can be processed using more advanced schemes.

The RC array can also be improved. The current

array is a 5-years old prototype PZT array with only 62 elements. The array has 8 non-functioning row elements and is slightly curved with a deviation around 0.1λ from a flat surface. This introduces phase errors and impedes image contrast. Other more advanced focusing schemes, like matched filter focusing, could also be used for increasing contrast [21, 39]. It is also possible to optimize the emission sequence for contrast agent enhancement, where amplitude modulation potentially could be used [40–42], and it could also be possible to optimize the imaging sequence with fewer emissions for yielding less data and higher volume rates [43]. Adding more elements to the probe can also increase resolution and thereby reduce acquisition time, as more bubbles can be separated. Early investigations have been made for a $192 + 192 \text{ RC}$ array and showed an increased resolution proportional to the F-number and wavelength [44]. Such arrays can directly be used on modern ultrasound consoles with few modifications in the beamforming.

The approach can fairly easily be translated to clinical use by modifying our current 2-D super resolution pipeline to include searches and localizations in 3-D [34, 35]. The motion correction schemes developed for 2-D imaging and needed for in-vivo imaging can then also be applied [34]. The main clinical applications could be superficial structures, where the F-number in beamforming can be kept low. The penetration depth is 14 cm for this array and imaging scheme, which is beneficial for larger organs like the liver. The bubble density would have to be reduced for reliable detection, and the imaging region will only be within the rectilinear area of the probe. This can potentially be alleviated by using a lens in front of the array [45].

A first in-vivo target would be to scan a rat kidney as performed in [34, 35]. The acquisition time was between 1 and 10 minutes for a 2-D image, with 1 minute giving an overall rough view of the vasculature and 10 minutes giving precise quantitative data for the blood flow. We predict that the same scan times can be kept with the method presented here for a full volume, and maybe with a shorter time when employing more advanced SRI [36–38].

The high resolution will also give some future challenges. The RC array can image a volume of $31\lambda \times 31\lambda \times 280\lambda$, which, with a voxel size of $(10\mu\text{m})^3$, would yield 33 GVoxels. This might give some challenges in the display of such data.

V. CONCLUSION

A method for 3-D SRI has been investigated, where a $62+62 \text{ RC}$ array was employed. A detection precision

better than $23\ \mu\text{m}$ was attained in all three coordinates for both the SonoVue contrast agent flowing in a micro-phantom and the point micro-phantom. The precision was obtained using 1/8 of the elements employed in previous 3-D SRI, which reduces both the storage and processing demands by a factor of eight. The approach yielded an increase in volumetric resolution by a factor of more than 28,900 with a possible penetration depth down to 14 cm and corresponding increase in the amount of volumetric data to 10-40 GVoxels. Potentially a volume of $16 \times 16 \times 140\ \text{mm}^3$ can be resolved with a voxel size of $(10\ \mu\text{m})^3$.

ACKNOWLEDGMENT

This work was financially supported by grant 82-2014-4 from the Danish National Advanced Technology Foundation, by grant 7050-00004B from Innovation Fund Denmark, and from BK Medical, Herlev, Denmark.

REFERENCES

- [1] O. Couture, B. Besson, G. Montaldo, M. Fink, and M. Tanter, "Microbubble ultrasound super-localization imaging (MUSLI)," in *Proc. IEEE Ultrason. Symp.*, 2011, pp. 1285–1287.
- [2] O. M. Viessmann, R. J. Eckersley, K. C. Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, pp. 6447–6458, 2013.
- [3] M. Siepmann, G. Schmitz, J. Bzyl, M. Palmowski, and F. Kiessling, "Imaging tumor vascularity by tracing single microbubbles," *Proc. IEEE Ultrason. Symp.*, pp. 6293–6297, 1906–1908, 2011.
- [4] M. A. O'Reilly and K. Hynynen, "A super-resolution ultrasound method for brain vascular mapping," *Med. Phys.*, vol. 40, no. 11, pp. 110701–7, 2013.
- [5] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433–440, February 2015.
- [6] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, pp. 499–502, November 2015.
- [7] O. Couture, V. Hingot, B. Heiles, P. Muleki-Seya, and M. Tanter, "Ultrasound localization microscopy and super-resolution: A state of the art," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 8, pp. 1304–1320, 2018.
- [8] J. Zhu, E. M. Rowland, S. Harput, K. Riemer, C. H. Leow, B. Clark, K. Cox, A. Lim, K. Christensen-Jeffries, G. Zhang, J. Brown, C. Dunsby, R. J. Eckersley, P. D. Weinberg, and M.-X. Tang, "3D super-resolution US imaging of rabbit lymph node vasculature in vivo by using microbubbles," *Radiology*, vol. 291, no. 3, pp. 642–650, 2019.
- [9] F. Lin, S. E. Shelton, D. Espindola, J. D. Rojas, G. Pinton, and P. A. Dayton, "3-D ultrasound localization microscopy for identifying microvascular morphology features of tumor angiogenesis at a resolution beyond the diffraction limit of conventional ultrasound," *Theranostics*, vol. 7, no. 1, pp. 196–204, 2017.
- [10] K. Christensen-Jeffries, S. Harput, J. Brown, P. N. T. Wells, P. Aljabar, C. Dunsby, M. Tang, and R. J. Eckersley, "Microbubble axial localization errors in ultrasound super-resolution imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 64, no. 11, pp. 1644–1654, 2017.
- [11] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, "SARUS: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.
- [12] L. Petrusca, F. Varray, R. Souchon, A. Bernard, J. Y. Chapelon, H. Liebgott, W. A. N'Djin, and M. Viallon, "Fast volumetric ultrasound B-mode and Doppler imaging with a new high-channels density platform for advanced 4D cardiac imaging/therapy," *Applied Sciences (Switzerland)*, vol. 8, no. 2, pp. 200:1–15, 2018.
- [13] S. Harput, K. Christensen-Jeffries, A. Ramalli, J. Brown, J. Zhu, G. Zhang, C. H. Leow, M. Toulemonde, E. Boni, P. Tortoli, R. J. Eckersley, C. Dunsby, and M. Tang, "3-D super-resolution ultrasound (SR-US) imaging with a 2-D sparse array," *arXiv preprint*, p. 1902.01608v1, 2019.
- [14] E. Boni, L. Bassi, A. Dallai, F. Guidi, V. Meacci, A. Ramalli, S. Ricci, and P. Tortoli, "ULA-OP 256: A 256-channel open scanner for development and real-time implementation of new ultrasound methods," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 10, pp. 1488–1495, 2016.
- [15] C. E. Morton and G. R. Lockwood, "Theoretical assessment of a crossed electrode 2-D array for 3-D imaging," in *Proc. IEEE Ultrason. Symp.*, 2003, pp. 968–971.
- [16] C. E. M. Démoré, A. Joyce, K. Wall, and G. Lockwood, "Real-time volume imaging using a crossed electrode array," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 56, no. 6, pp. 1252–1261, 2009.
- [17] J. T. Yen, C. H. Seo, S. I. Awad, and J. S. Jeong, "A dual-layer transducer array for 3-D rectilinear imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 56, no. 1, pp. 204–212, 2009.
- [18] A. Sampaleanu, P. Zhang, A. Kshirsagar, W. Moussa, and R. Zemp, "Top-orthogonal-to-bottom-electrode (TOBE) CMUT arrays for 3-D ultrasound imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 2, pp. 266–276, 2014.
- [19] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, "3-D imaging using row-column-addressed arrays with integrated apodization — Part I: Apodization design and line element beamforming," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 947–958, 2015.
- [20] T. L. Christiansen, M. F. Rasmussen, J. P. Bagge, L. N. Moesner, J. A. Jensen, and E. V. Thomsen, "3-D imaging using row-column-addressed arrays with integrated apodization — part II: Transducer fabrication and experimental results," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 959–971, 2015.
- [21] H. Bouzari, M. Engholm, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Imaging performance for two row-column arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 7, pp. 1209–1221, 2019.
- [22] M. Engholm, H. Bouzari, T. L. Christiansen, C. Beers, J. P. Bagge, L. N. Moesner, S. E. Diederichsen, M. B. Stuart, J. A. Jensen, and E. V. Thomsen, "Probe development of CMUT and PZT row-column-addressed 2-D arrays," *Sens. Actuators A: Phys.*, vol. 273, pp. 121–133, 2018.
- [23] M. B. Stuart, M. Schou, and J. A. Jensen, "Row-column beamforming with dynamic apodizations on a GPU," in *Proc. SPIE Med. Imag.*, 2019, pp. 1–7, paper number 10955-20.
- [24] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 39, no. 2, pp. 262–267, 1992.

- [25] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comp.*, vol. 10th Nordic-Baltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1, pp. 351–353, 1996.
- [26] —, "A multi-threaded version of Field II," in *Proc. IEEE Ultrason. Symp.* IEEE, 2014, pp. 2229–2232.
- [27] M. L. Ommen, M. Schou, R. Zhang, C. A. V. Hoyos, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, "3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [28] J. A. Jensen, O. Holm, L. J. Jensen, H. Bendsen, S. I. Nikolov, B. G. Tomov, P. Munk, M. Hansen, K. Salomonsen, J. Hansen, K. Gormsen, H. M. Pedersen, and K. L. Gammelmark, "Ultrasound research scanner for real-time synthetic aperture image acquisition," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 52, no. 5, pp. 881–891, May 2005.
- [29] P. Tortoli and J. A. Jensen, "Introduction to the special issue on novel equipment for ultrasound research," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 53, no. 10, pp. 1705–1706, Oct. 2006.
- [30] E. Boni, A. C. H. Yu, S. Freear, J. A. Jensen, and P. Tortoli, "Ultrasound open platforms for next-generation imaging technique development," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 7, pp. 1078–1092, 2018.
- [31] G. E. Trahey, J. W. Allison, and O. T. von Ramm, "Angle independent ultrasonic detection of blood flow," *IEEE Trans. Biomed. Eng.*, vol. BME-34, no. 12, pp. 965–967, 1987.
- [32] S. I. Nikolov and J. A. Jensen, "In-vivo synthetic aperture flow imaging in medical ultrasound," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 50, no. 7, pp. 848–856, 2003.
- [33] K. L. Gammelmark and J. A. Jensen, "2-D tissue motion compensation of synthetic transmit aperture images," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 4, pp. 594–610, April 2014.
- [34] J. A. Jensen, S. I. Nikolov, K. L. Hansen, M. B. Stuart, C. A. V. Hoyos, M. Schou, M. L. Ommen, S. H. Øygaard, L. T. Jørgensen, M. S. Traberg, T.-Q. Nguyen, E. V. Thomsen, N. B. Larsen, C. Beers, and et al., "History and latest advances in flow estimation technology: From 1-D in 2-D to 3-D in 4-D," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [35] S. B. Andersen, C. A. V. Hoyos, I. Taghavi, F. Gran, K. L. Hansen, C. M. Sørensen, and J. A. J. M. B. Nielsen, "Super-resolution ultrasound imaging of rat kidneys before and after ischemia-reperfusion," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [36] E. Kanoulas, M. Butler, C. Rowley, V. Voulgaridou, K. Diamantis, W. C. Duncan, A. Mcneilly, M. Averkiou, H. Wijkstra, M. Mischi, R. S. Wilson, W. Lu, and V. Sboros, "Super-resolution contrast-enhanced ultrasound methodology for the identification of in vivo vascular dynamics in 2d," *Investigative Radiology*, vol. 54, no. 8, pp. 500–516, 2019.
- [37] A. Bar-Zion, O. Solomon, C. Tremblay-Darveau, D. Adam, and Y. C. Eldar, "Sushi: Sparsity-based ultrasound super-resolution hemodynamic imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 12, pp. 2365–2380, 2018.
- [38] O. Solomon, R. J. van Sloun, H. Wijkstra, M. Mischi, and Y. C. Eldar, "Exploiting flow dynamics for super-resolution in contrast-enhanced ultrasound," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, p. Early Access, 2019.
- [39] J. T. Yen, "Beamforming of sound from two-dimensional arrays using spatial matched filters," *J. Acoust. Soc. Am.*, vol. 134, no. 5, pp. 3697–3704, 2013.
- [40] V. Mor-Avi, E. G. Caiani, K. A. Collins, C. E. Korcarz, J. E. Bednarz, and R. M. Lang, "Combined assessment of myocardial perfusion and regional left ventricular function by analysis of contrast-enhanced power modulation images," *Circulation*, vol. 104, no. 3, pp. 352–357, 2001.
- [41] R. J. Eckersley, C. T. Chin, and P. N. Burns, "Optimising phase and amplitude modulation schemes for imaging microbubble contrast agents at low acoustic power," *Ultrasound Med. Biol.*, vol. 31, no. 2, pp. 213–219, 2005.
- [42] C. Tremblay-Darveau, R. Williams, L. Milot, M. Bruce, and P. N. Burns, "Visualizing the tumor microvasculature with a nonlinear plane-wave Doppler imaging scheme based on amplitude modulation," *IEEE Trans. Med. Imag.*, vol. 35, no. 2, pp. 699–709, 2016.
- [43] R. Moshavegh, J. Jensen, C. A. Villagomez-Hoyos, M. B. Stuart, M. C. Hemmsen, and J. A. Jensen, "Optimization of synthetic aperture image quality," in *Proc. SPIE Med. Imag.*, vol. 9790, 2016, pp. 97900Z–97900Z–9.
- [44] M. Schou, A. S. Havreland, M. Engholm, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Design of a novel zig-zag 192+192 row column addressed transducer: A simulation study," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [45] H. Bouzari, M. Engholm, C. Beers, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Curvilinear 3-D imaging using row-column addressed 2-D arrays with a diverging lens: Phantom study," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 7, pp. 1182–1192, 2018.



Jørgen Arendt Jensen (M'93-SM'02-

F'12) received the MSc degree in 1985, the Ph.D. degree in 1989, and the Dr.Techn. degree all from the university in 1996. Since 1993, he has been a Full Professor of Biomedical Signal Processing with the Department of Health Technology, Technical University of Denmark. He has been the founder and head of the Center for Fast Ultrasound Imaging since its inauguration in 1998. CFU has contributed with innovations in transverse oscillation vector flow imaging, synthetic aperture flow imaging in 2-D and 3-D, ultrasound simulation, research scanners, and row-column probes and beamforming. He has published more than 500 journal and conference papers on signal processing and medical ultrasound and the book Estimation of Blood Velocities Using Ultrasound (Cambridge Univ. Press), 1996. He is also the developer and maintainer of the Field II simulation program. He has been a visiting scientist at Duke University, Stanford University, and the University of Illinois at Urbana-Champaign. He was founder and head of the Biomedical Engineering group from 2007 to 2010. In 2003, he was one of the founders of the biomedical engineering program in Medicine and Technology, which is a joint degree program between the Technical University of Denmark and the Faculty of Health and Medical Sciences at the University of Copenhagen. The degree is one of the most sought-after engineering degrees in Denmark. He was chairman of the study board from 2003 to 2010 and Adjunct Professor with the University of Copenhagen from 2005 to 2010. He has given a number of short courses on simulation, synthetic aperture imaging, and flow estimation at international scientific conferences and teaches biomedical signal processing and

medical imaging at the Technical University of Denmark. His research is centered around simulation of ultrasound imaging, synthetic aperture imaging, vector blood flow estimation, 3-D and super resolution imaging, row-column probes, and construction of ultrasound research systems. He has educated 41 PhD students and currently advises 17 PhD students. Dr. Jensen has given more than 60 invited talks at international meetings and received several awards for his research, most recently the Grand Solutions Prize from the Danish Minister of Science, the order of the Dannebrog by her Majesty the Queen of Denmark, and the Rayleigh award from the UFFC Society in the field of Ultrasonics in 2019.



Martin Lind Ommen received the M.Sc. degree in Physics and Nanotechnology from the Technical University of Denmark, Lyngby, Denmark, in 2017. He is currently a Ph.D. student with the Biomedical Engineering Section, Department of Health Technology at the Technical University of Denmark. His current research interests include micro- and nanofabrication in general, in particular the fabrication of capacitive micromachined ultrasonic transducers (CMUTs), as well as phantom fabrication, with a current focus on stereolithographic fabrication of micro-flow- and scatter-phantoms.



Sigrid Husebø Øygard is a PhD candidate at DTU Health Tech in the Center for Fast Ultrasound Imaging. Sigrid received her BEng in Acoustical Engineering from the Institute of Sound and Vibration at the University of Southampton in 2016, where she specialized in bubble acoustics. She received her MSc in Physics with a focus on theoretical acoustics from the University of Bergen in 2018. Her research interests include 3D ultrasound imaging, super resolution imaging, and bubble acoustics.



Mikkel Schou was born in 1990. He received the M.Sc. degree in biomedical engineering from the Technical University of Denmark, Kongens Lyngby, Denmark,

and the University of Copenhagen, Copenhagen, Denmark, in 2017. He is currently pursuing the Ph.D. degree in biomedical engineering with the Center for Fast Ultrasound Imaging, Technical University of Denmark. The topic of his Ph.D. research is 3-D ultrasound Perfusion and Flow imaging using Row-Column Arrays.



Thomas Sams received his MSc in Physics at the Niels Bohr Institute, University of Copenhagen (1986), where he also did his PhD on collective vibrations in nuclei (1991). He has worked for a 10 years on patterns in natural backgrounds at the Danish Defence Research Est. Since 2005 he has been an Associate Professor in Biomedical Engineering at the Technical University of Denmark, where he heads a research group in Cellular Signaling and Bioransport and the Biomedical Engineering Section. The common theme in his research is on the understanding of collective behavior in systems composed of smaller building blocks, including bacterial cell communities, patterns in yeast, patterns in sand ripples, and neural networks. He has worked as a post-doc at Saclay, France, the Niels Bohr Institute, and as a visiting researcher in Los Alamos, USA, and at the University of Cambridge, UK.



Mathias Bo Stuart received the M.Sc. and Ph.D. degrees in Computer Engineering from the Technical University of Denmark, Lyngby, Denmark in 2006 and 2010 respectively. He is currently an Associate Professor with the Biomedical Engineering Section, Department of Health Technology at the Technical University of Denmark. His research interests include synthetic aperture methods for both anatomical and flow imaging in both 2-D and 3-D, ultrasound systems, and real-time implementations of ultrasound processing algorithms.



Christopher Beers received the M.S. degree in acoustics from Pennsylvania State University, State College, PA, USA, in 2007, where his thesis research explored end-element anomalies in medical ultrasound transducer arrays. He has been with BK Medical, State College, PA, USA, since 2007 (formerly Sound Technology), where he develops transducer technology and designs commercial medical ultrasound probes.



Erik Vilain Thomsen was born in Aarhus, Denmark, in 1964. He received the M.Sc. degree in physics from the University of Southern Denmark, Odense, Denmark, in 1992 and the Ph.D. degree in electrical engineering from the Technical University of Denmark (DTU), Kongens Lyngby, Denmark, in 1998. He is currently Professor at the Department of Health Technology, Technical University of Denmark (DTU) where he is also the Head of the MEMS Applied Sensors Group and Head of Division with responsibility for the educations in healthcare engineering. He teaches classes in solid state electronics, microtechnology, and nano and microfabrication. His current research interests include all aspects of capacitive micromachined ultrasonic transducers (CMUTs), general MEMS technology, and piezoelectric MEMS. Dr. Thomsen received the Danish National Advanced Technology Foundation Grand Solution Prize in 2017, the AEG Electron Prize in 1995, and has received several teaching awards at DTU.



Niels Bent Larsen received his MSc and PhD degrees in chemistry in 1993 and 1997, both from University of Copenhagen. He is currently Professor and Section Head at the Department of Health Technology. His current research interests are in high-resolution 3D printing of compliant hydrogel materials to recreate the vascular network

of body organs to enable advanced in vitro organ models.



Borislav Gueorguiev Tomov received a M. Sc. degree in Electronics Engineering from the Technical University of Sofia, Bulgaria, in 1996, and a Ph.D. degree in Medical Electronics from the Technical University of Denmark in 2003. He is currently Senior Researcher at the Center for Fast Ultrasound imaging, Department of Health Technology, Technical University of Denmark. His research interests include medical ultrasound signal processing, and ultrasound scanner architectures and implementations.

APPENDIX B

Papers under review

B.1 Paper I - 3D Printed Calibration Micro-Phantoms for Validation of Super-Resolution Ultrasound Imaging

3D Printed Calibration Micro-Phantoms for Validation of Super-Resolution Ultrasound Imaging

*Martin Lind Ommen*¹, *Mikkel Schou*¹, *Christopher Beers*², *Jørgen Arendt Jensen*¹, *Niels Bent Larsen*¹,
and *Erik Vilain Thomsen*¹

¹Department of Health Technology, Technical University of Denmark, Kgs. Lyngby, Denmark

²BK Medical, State College, Pennsylvania, USA

Abstract—This study evaluates the use of 3D printed phantoms for 3D super-resolution ultrasound imaging (SRI) algorithm calibration. The main benefit of the presented method is the ability to do absolute 3D micro-positioning of sub-wavelength sized ultrasound scatterers in a material having a speed of sound comparable to that of tissue. Stereolithography is used for 3D printing soft material calibration micro-phantoms containing eight randomly placed scatterers of nominal size $205\ \mu\text{m} \times 205\ \mu\text{m} \times 200\ \mu\text{m}$. The printed structures are found to expand linearly in all three dimensions by 2.6% after printing. SRI algorithm calibration is demonstrated by imaging a phantom using a $\lambda/2$ pitch 3 MHz 62+62 row-column addressed (RCA) ultrasound probe. The printed scatterers will act as point targets, as their dimensions are below the diffraction limit of the ultrasound system used. Two sets of 640 volumes containing the phantom features are imaged, with an intervolumetric axial movement of the phantom of $12.5\ \mu\text{m}$, to emulate a flow velocity of 2 mm/s at a frame rate of 160 Hz. The ultrasound signal is passed to a super-resolution pipeline to localise the positions of the scatterers and track them across the 640 volumes. After compensating for the phantom expansion, a scaling of 0.989 is found between the distance between the eight scatterers calculated from the ultrasound data and the designed distances. The standard deviation of the variation in the scatterer positions along each track is used as an estimate of the precision of the super-resolution algorithm, and is expected to be between the two limiting estimates of $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.7\ \mu\text{m}, 27.6\ \mu\text{m}, 9.5\ \mu\text{m})$ and $(\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z) = (17.3\ \mu\text{m}, 19.3\ \mu\text{m}, 8.7\ \mu\text{m})$. In conclusion, this study demonstrates the use of 3D printed phantoms for determining the accuracy and precision of volumetric super-resolution algorithms.

Index Terms—3D printing, stereolithography, phantom, hydrogel, calibration, resolution, ultrasound

I. INTRODUCTION

Super-resolution ultrasound imaging (SRI) has recently emerged as a non-invasive technique, which enables imaging of the smallest vessels of the vasculature [1], [2], [3]. Micrometer sized gas filled bubbles provide high contrast in ultrasound imaging, and their path through the vasculature can be tracked over time to reveal the fine details of the vascular network. The conventional B-mode ultrasound images are replaced by cumulated maps of the super-localised centroids of the micro-bubbles, revealing vascular features which are much smaller than the diffraction limit of conventional ultrasound. However, a fundamental problem is to validate the spatial accuracy of these new techniques. Biological structures often have extremely complex geometries, with the added compli-

cations of liquid flow and tissue motion. Phantoms with well controlled dimensions are used for validation instead. As new imaging techniques are introduced, they are typically tested initially against numerically simulated data. This data could for instance be generated in Field II [4], [5], [6]. The next step would typically be to test the imaging techniques using phantoms, which have been adapted to suit the techniques. In the case of SRI, in which the end goal is to image vasculature on the scale of only a few tens of micrometers, the precision, the accuracy, and the repeatability of the phantom fabrication method all need to be improved. At the same time, it should ideally be possible to replicate the dimensionality of the vascular networks. Viessmann et al. [7] and Christensen-Jeffries et al. [8] employed tube phantoms of 3 mm and 200 μm diameters, respectively, to validate their SRI algorithms. Both of these are significantly larger than the vessels of interest, i.e. arterioles and venules with sub-100 μm dimensions and capillaries of 5-9 μm diameters [9]. Desailly et al. presented a phantom study in which the channel dimensions were reduced to $40 \times 80\ \mu\text{m}^2$ by utilizing the high resolution of silicon micro-fabrication UV lithography in polydimethylsiloxane (PDMS) [10]. While the latter dimensions of the channel were comparable to the sizes of capillaries, the ability to expand the phantom types to three dimensions are severely limited in all cases. A completely different approach uses the vasculature of chicken embryos, which is optically visible [11]. While the features are on the correct scale, and high resolution optical images can be taken, it is impossible to obtain a three-dimensional representation of the vascular network using commonly available optical microscopes. It should be noted that this is not a limitation of the chicken embryo model itself, since this will feature complex three-dimensional structures. But the characterisation of those networks is very complex, and not possible using regular optical microscopes. Optical mapping of the structures could be performed with other more complex methods such as optical coherence tomography [12]. All of the above mentioned methods are channel based, and thereby meant to provide an outer limit for the localization of the micro-bubbles which are tracked. But that leaves the inherent problem that it is not possible to control the position of the micro-bubbles within the tubes or vessels, and therefore, the source of the signal will not be precisely known.

3D printing of phantoms is a promising new approach, which does not suffer from these limitations. It provides

complete three-dimensional flexibility in fabrication and can replicate features in the sub-100 μm range [13]. Recently, 3D printed phantoms for ultrasound were demonstrated by Jacquet et al. [14], supposedly not with SRI in mind. The phantoms contained highly scattering solid features as small as $30 \times 50 \mu\text{m}^2$ in cross section, demonstrating the exciting potential for point spread function evaluation provided by the method, as well as other possibilities for phantom features and uses.

Recently, we presented an alternative 3D printing method for phantom fabrication, namely stereolithography. The method is used to print hydrogels, a soft material with acoustic properties similar to tissue, and demonstrated that it is possible to obtain cavities of $100 \times 100 \times 100 \mu\text{m}^3$ and channels with a cross-section of $200 \times 200 \mu\text{m}^2$ in a hydrogel, which are suitable for SRI ultrasound [15]. We used that technique to print channel phantoms, which have been imaged using an row-column addressed (RCA) array to super-localize micro-bubbles in 3D [16]. The phantom consisted of a 200 μm diameter channel, and the super-resolved micro-bubble positions were determined in 3D using an RCA array. Through statistical analysis of the radial distribution of the micro-bubbles around the center line of the channel, the precision of the SRI system was determined to be less than 23 μm in all dimensions, which approaches the scale of the smallest vessels in tissue. In this case, the statistical analysis was used to mitigate the problem of not knowing the exact position of the micro-bubbles, since that of course is still a problem when using 3D printed channel based micro-phantoms.

A 3D printing solution will in principle allow for arbitrarily complex 3D structures or channel networks to be made. However, the presented work demonstrates an alternative to localisation of micro-bubbles in channel systems for SRI algorithm calibration, which allows for elimination of the issue of the uncertainty in the micro-bubble positions within micro-channels. It is based on [15] which also demonstrated that small 3D printed cavities will scatter sound. By keeping them smaller than the ultrasound wavelength, they can be used as point targets to evaluate imaging performance for regular B-mode imaging. Thereby, it becomes possible to fixate small scatterers at very precise locations within the hydrogel. These structures will be stable in time, enabling repeated imaging, in direct contrast to small channels and microbubbles.

In this paper, we demonstrate how these stable, fixated scatterers can be used as an alternative to conventional tube phantoms to determine the accuracy and precision of SRI hardware and algorithms independently in all three dimensions. The 3D printing method allows for absolute positioning accuracy and precision un-paralleled by other types of phantoms fabrication methods. It is also shown how the high degree of control of phantom features using this method can illuminate additional problems in a SRI pipeline, such as distortion. The phantoms can potentially be used to demonstrate local variations in the SRI properties based on the scatterer position within the localization field of view.

II. MATERIALS AND METHODS

This section describes the phantom fabrication method, and the details of the phantoms which have been designed for

these experiments. Additionally, the experimental procedures are described, as well as the SRI pipeline structure.

A. Fabrication of the phantoms

Calibration phantoms were fabricated by stereolithographic 3D printing of an aqueous solution of poly(ethylene glycol) diacrylate (PEGDA, M_n 700 g/mol, 455008, Sigma-Aldrich) to form a hydrogel solid. Stereolithography is based on printing a stack of individual thin layers of materials, calling for prior digital slicing of the targeted 3D design into separate layer designs matching the printing system. The method and components have previously been presented in more detail [15]. Each layer is printed by spatially confined illumination of the targeted solid areas, which leads to localized photo-chemically induced solidification of the printing solution. A custom-designed stereolithographic printer that projects a full image of the current layer was used. Each layer image is a one-to-one projection of a digital image generated on a Digital Mirror Device (DMD, DLP9500UV, Texas Instruments, TX) with a center-to-center pixel spacing of 10.8 μm in both lateral dimensions. Thus, there will inherently be a physical mapping of the targeted phantom design layers onto a square grid of 10.8 μm spacing. Phantom shapes were generated directly as a series of layer images matching the DMD pixel pitch using a MATLAB (MathWorks, MA) script and with a layer thickness of 20 μm . The aqueous printing solution contained 20% (weight by volume) PEGDA as pre-polymer, 5 mg/mL LAP (lithium phenyl-2,4,6-trimethylbenzoylphosphinate, 900889, Sigma-Aldrich) as photoinitiator, and 12 mg/mL Quinoline Yellow (309052, Sigma-Aldrich) as absorber. Each patterned layer motif was illuminated with 365 nm light at an intensity of 20 mW/cm² for 3 seconds (in the bulk of the phantom) to 23 seconds (locally on the scatterer perimeter), depending on the features being printed. The phantoms were printed on $22 \times 22 \times 0.40 \text{ mm}^3$ cover glasses (MEN-ZDA022022A4E0, Menzel Gläser, DE) pretreated with (3-glycidyloxypropyl)trimethoxysilane (440167, Sigma-Aldrich) to enhance the adhesion to the printed PEGDA. The resulting printed structures are not in equilibrium with water directly after printing, but will swell slightly when subsequently transferred to water. Previous work showed that after four hours, the printed structure reaches its equilibrium swelling.

Acoustic parameters of a hydrogel solid with a layer exposure time of 3 s, similar to the bulk of the used scatterer phantom, has been measured. The speed of sound was determined to be 1580 m/s, which correlates well to the speed of sound found in typical human tissues [17]. The attenuation was fitted as a power law model given by $\mu = a \cdot f^b$, where f is the ultrasound frequency in MHz, a is the attenuation coefficient at 1 MHz and b describes the degree of nonlinearity of the dependence on frequency [18]. The parameters were determined as $a = 0.15 \text{ dB}/[\text{MHz cm}]$ and $b = 1.5$. As the phantom is submerged in water during experimentation, the speed of sound in the entire imaged volume does not match perfectly. Beamforming compensation for different speeds of sound is a complex matter as illustrated in [19]. For this experiment, we have chosen to approximate the entire volume as having a speed of sound equal to that of water (1480 m/s).

The foundation for all experiments in this work is a phantom containing eight randomly placed scatterers. The outer dimensions of the phantom is $21.1 \times 11.9 \times 11.9 \text{ mm}^3$ with each scatterer being $205 \times 205 \times 200 \text{ }\mu\text{m}^3$. While the printing setup allows for printing significantly smaller scatterers, it was necessary with an increased size to obtain reflections with intensities larger than background scattering due to unavoidable small random print errors in the phantom. The scatterers will function as point targets in regular B-mode volumes, when the imaging wavelength is larger than the scatterer size, in this case for any frequency below 6 MHz. They were placed with a minimum separation distance of 3 mm, which will eliminate overlapping signals for any frequency above 0.5 MHz. The designed layout is shown in Fig. 1, in which the blue points represents the randomly placed scatterers. Separate droplines lead from the points out along the y -axis and along the z -axis respectively. The droplines end up 1 mm from the respective surfaces in the collapsed x - y -plane version (red) and the collapsed x - z -plane version (turquoise) of the scatterers.

Hollow structures will initially be filled with unpolymerized printing solution during printing. When submerged in water, this will over time partly be replaced by water since the hydrogel is diffusion open to water. The cavities referred to in this work are only cavities in terms of not containing solid PEGDA hydrogel.

B. Experimental setup and procedure

1) *Optical validation of phantoms:* The phantom fabrication method accuracy should be verified by another characterisation method. Although the printer specifications have been stated, they only specify the lower limit of the attainable feature sizes and accuracies. Furthermore, the phantom expansion due to post-printing swelling needs to be determined to compensate the designed feature sizes before using the phantom as a calibration tool. Optical characterisation using an optical microscope can be used to locate phantom features with high precision. Unfortunately, the printed hydrogel scatters light, rendering it impossible to use the base phantom with the eight randomly placed scatterers, since these are placed too far inside the phantom. Instead, the same coordinates were used to make two new phantoms, in which the coordinates were collapsed either into the x - y -plane and placed near the top of the phantom (red in Fig. 1), or into the x - z -plane and placed near the side of the phantom (turquoise in Fig. 1). By placing them near the surfaces, the light scattering is minimised and the scatterers become clearly visible in the optical microscope. Each scatterer was physically moved into a defined centre point in the optical viewfield using an X-Y microscope stage with integrated linear encoders for accurate readout of the actual position. This procedure circumvents possible measurement errors due to distortions in the optical components. The measurements were performed using a Zeiss LSM 700 upright microscope with a Zeiss 130x85 PIEZO stage having a positioning reproducibility of $\pm 0.6 \text{ }\mu\text{m}$.

The positioning accuracy of the procedure was assessed by repeatedly locating the same scatterer. The position was found with a standard deviation of $1.3 \text{ }\mu\text{m}$ along both the

x -axis and the y -axis ($n = 50$). The analysis procedure is sketched in Fig. 2. The distance between all scatterers can be determined from the scatterer positions. The distances can then be correlated to the corresponding design distances from the 3D model, and the correlation should be linear. The slope of the correlation is the factor by which the printed structure has expanded relative to the design. If the printed structures are a perfect replication of the design, the correlation will be a linear relationship with a slope of 1.

2) *Super-resolution ultrasound imaging calibration:* With the correlation determined, the true distances between the scatterers in the 3D version of the scatterer phantom will be known, and can be used to compare against those found by ultrasound. When aiming to measure position changes on the order of a few micrometers, vibrations of the measurement setup will be detrimental. A 3D printed holder was fitted to the phantom dimensions, enabling mounting of the phantoms on top of an absorbing polyurethane rubber sheet (Sorbothane, Inc., Kent, Ohio, USA). The holder was mounted to the bottom of a water tank which in turn was mounted on a 8MR190-2-28 rotation stage (0.01° resolution) combined with a 8MTF-75LS05 x - y translation stage ($0.31 \text{ }\mu\text{m}$ resolution) (Standa, Vilnius, Lithuania). To minimize the effect of vibrations, everything was mounted on a Newport PG Series floating optical table (Irvine, California). A sketch of the setup can be seen in Fig. 3.

The phantom was translated relative to the ultrasound probe using the translation stage along a single axis; in the first experiment along the x -axis, and in the second experiment along the y -axis. The inter-volume stage movement in both experiments was $12.5 \text{ }\mu\text{m}$, corresponding to a 2 mm/s velocity acquired at a volume rate of 160 Hz . This speed corresponds to common flow velocities in small vessels. By moving the phantom in between volume acquisitions, any differences depending on the phantom placement within the field of view of the transducer will be included in the analysis, instead of simply testing the SRI pipeline parameters locally within the transducer field of view.

The imaging probe was a prototype $62+62$ elements 3 MHz piezo-electric, row-column addressed (RCA) array [20]. The probe was connected to the experimental scanner SARUS [21], which is capable of storing channel data for offline processing. A single frame is a summation of 32 defocused emissions using a synthetic aperture (SA) imaging approach [16]. Rows were transmitting and columns were receiving, thereby resulting in 62 channels in receive per emission. The phantom was stationary while a frame was being measured to avoid intra-frame motion artefacts. In total 2×640 volumetric frames were acquired over the 2×640 positions. The volumetric frames were then passed to the SRI pipeline.

C. Super-resolution pipeline

The SRI pipeline has been described in detail in [16]. It is briefly summarised in the following. The super resolution pipeline consists of three steps. The first is SA beamforming. Each imaged volume spans a volume of $14.86 \times 14.86 \times 7.43 \text{ mm}^3$, corresponding to

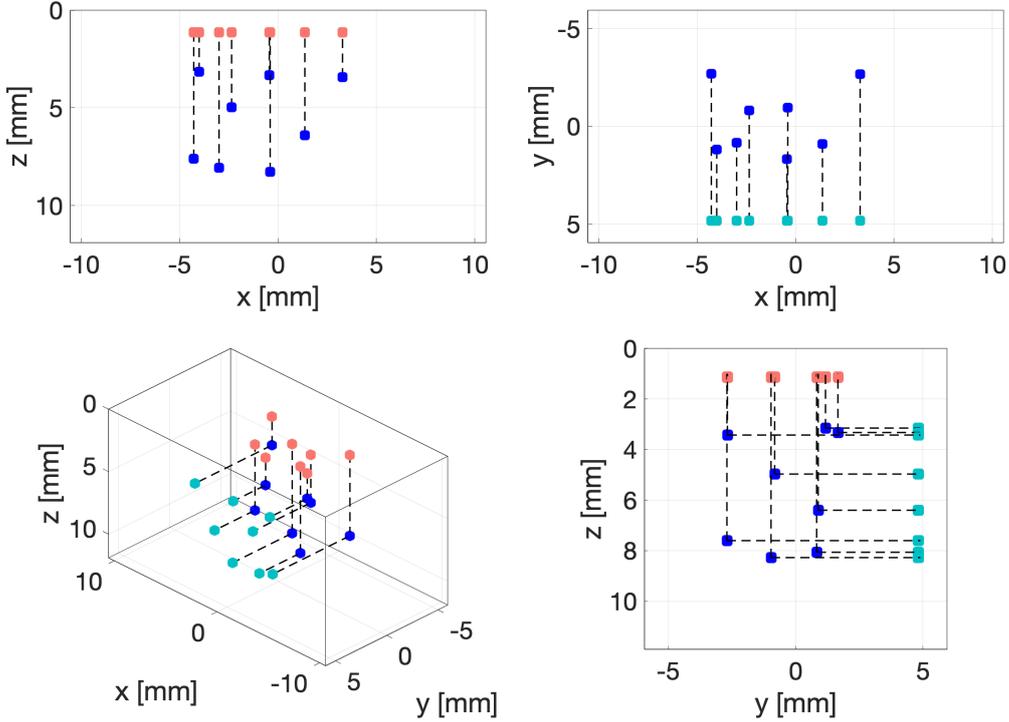


Figure 1. The designed layout of the scatterers within the $\sim 21.1 \times 11.9 \times 11.9 \text{ mm}^3$ phantom. The blue points are the randomly placed scatterers. The droplines are included to aid the 3D perception of the scatterer placement. For the optical correlation experiment, one set of phantoms had the scatterers collapsed into the x - y -plane near the top surface (red), and the other set had the scatterers collapsed into the x - z -plane near the side (turquoise).

$61 \times 61 \times 243$ voxels. Each high resolution volume was a summation of 32 volumes beamformed from 32 emissions, using a specialised beamformer [22] implemented on a GPU [23]. The volume was dynamically focused in receive (F-number of 1.5) and synthetically in transmit (F-number of 1), with an optimized sequence for SA B-mode. This was done for all 2×640 frames. In the next step, a stationary echo filter was applied to remove stationary tissue. In a micro-bubble experiment, this would remove the signal stemming from the tissue as it is stationary, leaving only the micro-bubble signal. However, since the entire phantom was translated between each frame in this experiment, the stationary echo filter would have no effect on the results. The final step is to determine the points scatterer positions based on local maxima. Sub-pixel positioning is obtained by interpolating the peak location using a second order polynomial in all three dimensions. The 3D coordinates $\{x_p, y_p, z_p\}$ of the detected points is then provided as the output from the third stage. Tracks of the individual scatterers can then be formed by collecting spatially similar coordinates across all imaged frames. The pipeline was implemented in MATLAB, and was processed offline [16].

III. RESULTS

This section presents the results of the optical validation of the printed structures, as well as the accuracy and precision of the SRI pipeline.

A. Optical validation of scatterer positions

Two replicates of each of the two projected-scatterer phantoms for optical validation were made. Each scatterer was located using the optical microscope and the translation stage coordinates of each scatterer was determined. Subsequently, the scatterer coordinates were used to determine the distance between the scatterers. The correlation between the optically measured distances and the designed distances can be seen in Fig. 4. In addition to analysing the direct correlation between measured distances and design distances, it was also investigated whether there was any difference between the two sets of cross-planes (x - y and x - z), which could potentially be explained by the anisotropic voxels. The different phantoms were also modelled as a random factor, to test and compensate for print-to-print variability. The combination of fixed and random factors makes the fitted model a linear mixed effects model. Such a model can be analysed using the lmerTest package [24] in R [25]. A summary of the data types and the factors included in the analysis can be seen in Table I. The initial mixed effects model is given as

$$Y_i = \mu + \alpha(\text{Plane}_i) + (\beta_1 + \beta_2(\text{Plane}_i)) x_{\text{design},i} + c(\text{Phantom}_i) + \epsilon_i, \quad (1)$$

where Y_i is the optically measured distances, μ is the overall intercept, $\alpha(\text{Plane}_i)$ is an intercept addition due to the Plane factor, β_1 is the average slope of the model,

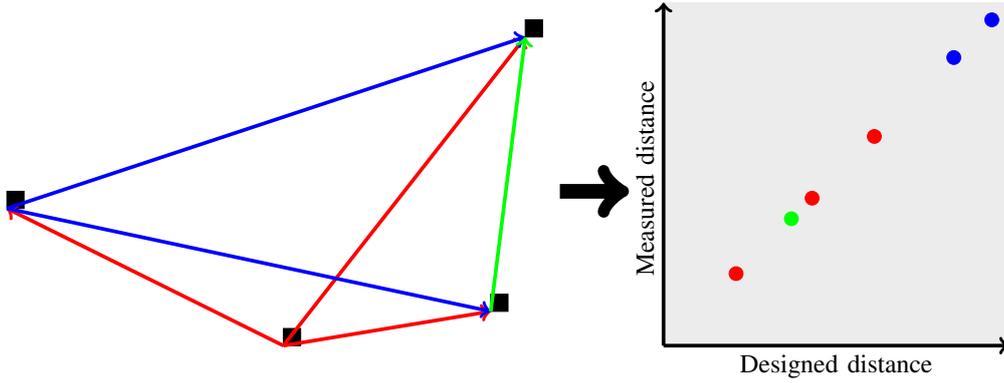


Figure 2. Sketch of the procedure to determine the accuracy of the printed phantoms. The black squares represent printed scatterers. The distances between scatterers are determined, and a correlation between the measured distances and the designed distances is made. The slope of the correlation will be the expansion factor post-printing.

Table I
SUMMARY OF THE VARIABLES AND THEIR DATA TYPES USED IN THE OPTICAL CORRELATION ANALYSIS.

	Sample values	Variable type	Description
Optical distance [mm]	4.041, 1.950, ..., 8.189	Numerical values	The response variable, measured by optical microscope
Design distance [mm]	3.973, 1.927, ..., 8.087	Numerical values	The designed distance between points
Plane	XY, XZ	Fixed factor	The cross-plane investigated
Phantom	1, 2, 3, 4	Random factor	The phantom group

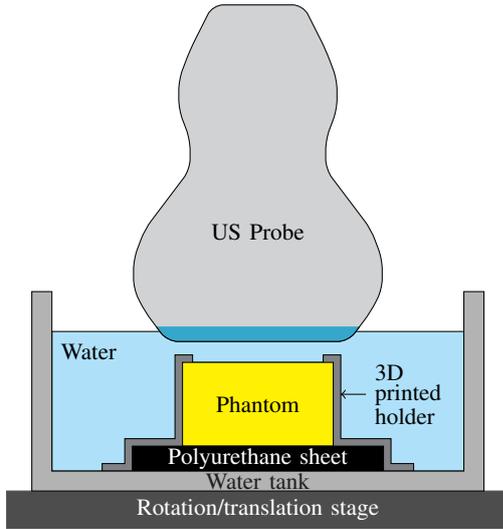


Figure 3. Sketch of the experimental ultrasound setup.

$\beta_2(\text{Plane}_i)$ is a plane dependent correction to the slope, $c(\text{Phantom}_i) \sim N(0, \sigma_{\text{Phantom}}^2)$ is a random offset from phantom to phantom, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All $c(\text{Phantom}_i)$'s and ϵ_i 's are independent.

The model reduction was conducted by removing only a single term at a time, based on a 5% level of significance. Neither the random effect of the individual phantoms ($c(\text{Phantom}_i)$), nor the Plane dependent intercept addition ($\alpha(\text{Plane}_i)$), nor the Plane dependent slope ($\beta_2(\text{Plane}_i)$) were significant at 5%.

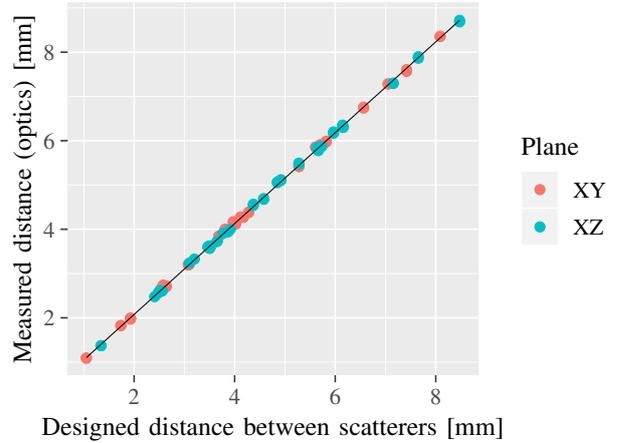


Figure 4. Correlation between the distance between the designed scatterer positions and the distances measured using an optical microscope. The black line is the final reduced model seen in Eq. (2).

Thereby the model reduction converged at the final model

$$Y_i = \mu + \beta_1 \cdot x_{\text{design},i} + \epsilon_i. \quad (2)$$

The model coefficients and confidence intervals of the reduced model can be seen in Table II. The analysis showed that the phantom swelling is isotropic, since there was no effect of the Plane factor. There was no significant difference between the four test phantoms, indicating good print repeatability. The parameter estimate of β_1 indicates that the phantom expands by approximately 2.6% along all dimensions. The residual standard error of the model is 36.6 μm . Model diagnostics showed that the residuals appeared to be normally distributed.

Table II
MODEL PARAMETER ESTIMATES OF THE FINAL REDUCED MODEL
INCLUDING CONFIDENCE INTERVALS OF CORRELATION BETWEEN
OPTICAL MEASUREMENTS AND DESIGN DISTANCES.

	Estimate	Standard Error	2.5%	97.5%
Intercept [mm]	0.023	0.009	0.005	0.042
β_1 (slope)	1.026	0.002	1.022	1.030

Thereby, the model is a good describer for the phantom expansion. The overall good correlation of all points to the straight line indicate that the expansion is uniform and isotropic in the investigated region of the print area. The analysis showed a significant intercept of 23 μm , which was unexpected. Given that the intercept lies outside of the data range of interest, it will not be analysed any further. It should be noted that the confidence interval for the intercept varies from less than a single voxel width, to four voxel widths.

B. Ultrasound super-resolution pipeline calibration

1) *Scatterer localisation*: Fig. 5 shows three selected cross planes of a B-mode volume. The coloured dots mark the localised positions of the scatterers detected in one of the 640 volumes. The example cross planes have been chosen such that they all contain the scatterer marked by a blue dot. The $x-z$ cross plane (Fig. 5c) also contains an additional scatterer (red). The selected volume contains a total of five scatterers, with the remaining scatterers not visible within the selected cross-planes. The large reflection at $x \approx 3.5$ mm and $z \approx 4$ mm does not correlate with any of the designed scatterer positions, and likely stems from a print artefact.

The localised positions of the 3D printed scatterers, accumulated over the 640 volumes, can be seen in Fig. 6. The colours group the tracked points of the individual scatterers, while the black tracks illustrate the expected tracks based on the design coordinates. The latter are included for visual confirmation that the localisations are indeed the designed scatterers. It is recommended to always include such a comparison to confirm that the localisations indeed correspond to the features of the designed phantom. Drop lines are included to aid the 3D perception. The horizontal field of view in the figures have been limited to the measured data tracks, removing parts of the black tracks. The actual cross-sectional field of view of the probe is 14.86×14.86 mm².

Although eight scatterers were printed, not all were found in the two experiments: seven scatterers were correctly localised for the movement along the x -axis (Fig. 6 a)) and five scatterers were correctly localised for the movement along the y -axis (Fig. 6 b)). In addition, the track length varies from 81 localisations to 633 localisations, across the 640 volumes. Two additional tracks, which did not align with the design coordinates, have been omitted from the images and the analyses. It is expected that these tracks stem from print artefacts, resulting in unintended cavities in the phantom, which therefore reflect the ultrasound similarly as the designed scatterers. They aligned well with the reflection seen in Fig. 5c at $x \approx 3.5$ mm and $z \approx 4$. While these print artefacts

would also be fixed in position, and be moved along the same trajectory as the designed scatterers, the print artefact geometry is not known. If a print artefact is significantly larger than the imaging wavelength, localisation of the centroid might be ambiguous, and therefore, these tracks were omitted from the analysis.

2) *Super-resolution accuracy*: The SRI pipeline accuracy was investigated in a similar manner to the optical validation, by comparing the known distances between the designed points to the measured distances between points from the ultrasound experiments. There are two main differences to the optical experiment: The scatterers are now positioned not in collapsed planes but in 3D, visualised as the blue points in Fig. 1, and the design distances are compensated for the expansions according to the results in Table II before analysing the correlation between the designed distances and those calculated from the ultrasound data. After the compensation, the correlation should be a straight line with a slope of 1, in the case of perfect correlation. Since there are two sets of experiments, one for each direction of motion of the translation stage, the variables of the analysis are the compensated design distances, the measured ultrasound distances, and a factor separating the data into the x - and y -motion, all summarised in Table III. In this experiment, the entire beamformed volume has been assumed to have a speed of sound equal to that in pure water, 1480 m/s.

As was mentioned in Section III-B1 and shown in Fig. 6, an unequal number of scatterers were localised by the SRI pipeline in the two experiments, and the tracks were of unequal length. This means there will be more data for the x -direction of motion, resulting in an unbalanced dataset from a statistical point of view. In addition, our analysis of the variation in the data showed that the data was heteroscedastic. Modelling the correlation of the raw distances between points might be heavily biased toward certain parts of the data simply due to the large number of samples. Instead, a weighted least squares analysis of the distance distributions was conducted. This was performed by modelling the mean distance between each point across all measurements, with each mean value being weighted by the variance of the measurements contributing to that mean. The correlation between the compensated design distances and the mean of the distances calculated by the SRI pipeline is shown in Fig. 7.

The initial linear model is given as

$$Y_i = \mu + \alpha(\text{Motion}_i) + (\beta_1 + \beta_2(\text{Motion}_i)) x_{\text{design},i} + \epsilon_i, \quad (3)$$

where Y_i is the mean of the distance between points calculated from the SRI pipeline output, μ is the overall intercept, $\alpha(\text{Motion}_i)$ is an intercept addition due to the Motion factor, β_1 is the average slope of the model, $\beta_2(\text{Motion}_i)$ is a Motion dependent correction to the slope, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All ϵ_i 's are independent.

The model reduction was conducted by removing only a single term at a time, based on a 5% level of significance. Neither

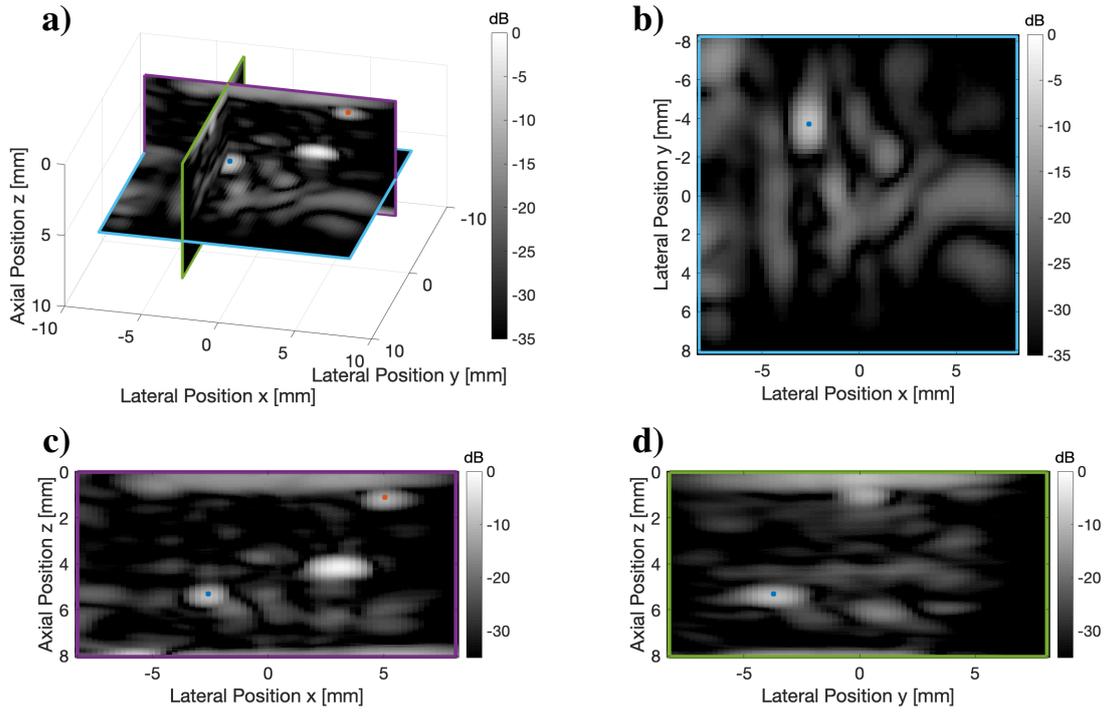


Figure 5. a) B-mode volume containing scatterers. Three cross planes of the B-mode volume are shown, b) x - y , c) x - z , and d) y - z . The super-localised positions of the scatterers are marked by coloured dots.

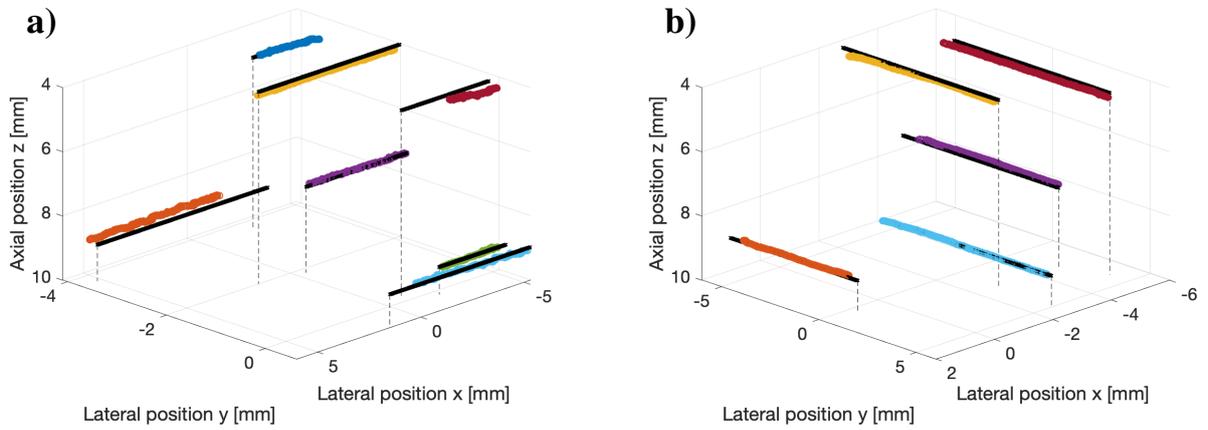


Figure 6. Cumulated localized scatterers acquired over 640 volumes. The phantom was translated in two separate experiments, along the transducer x -axis (a), and along the transducer y -axis (b). The black tracks illustrate the expected tracks based on the design coordinates. Droplines end on the $z=10$ mm plane, and are included to aid the 3D perception.

Table III
SUMMARY OF THE VARIABLES AND THEIR DATA TYPES USED IN THE ULTRASOUND CORRELATION ANALYSIS.

	Sample values	Variable type	Description
Ultrasound distance [mm]	8.717, 3.730, ..., 6.279	Numerical values	The distance between points calculated through the super-resolution pipeline
Compensated design distance [mm]	8.719, 3.811, ..., 6.384	Numerical values	The compensated designed distance between points
Motion	X, Y	Fixed factor	The axis of translation

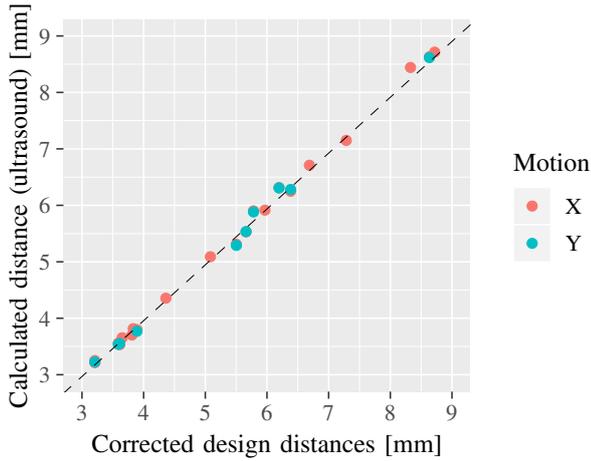


Figure 7. Correlation between the compensated design distances and the mean of the distances calculated by the SRI pipeline. The line represent the final reduced model seen in Eq. (4).

Table IV
MODEL PARAMETER ESTIMATES OF THE FINAL REDUCED MODEL INCLUDING CONFIDENCE INTERVALS OF CORRELATION BETWEEN ULTRASOUND DISTANCES AND COMPENSATED DESIGN DISTANCES.

	Estimate	Standard Error	2.5%	97.5%
β_1 (slope)	0.989	0.003	0.982	0.996

the overall intercept (μ), nor the direction of motion dependent addition to the intercept ($\alpha(\text{Motion}_i)$), nor the direction of motion dependent correction to the slope ($\beta_2(\text{Motion}_i)$) were significant at 5%, and were therefore removed. Thereby the model reduction converged at the final model

$$Y_i = \beta_1 \cdot x_{\text{design},i} + \epsilon_i. \quad (4)$$

The model coefficient and confidence interval of the reduced model are presented in Table IV. The analysis showed no dependence of the direction of motion, nor any intercept of the correlation. The modelled average behaviour of the fitted line has a slope of 0.989, close, yet not equal, to a perfect correlation with a slope of 1. Based on the heteroscedastic assumption of the data, a direct estimate of the residual standard error is not meaningful.

3) *Super-resolution precision*: The same ultrasound data was used to estimate the SRI pipeline precision. The precision was estimated by investigating the variation of the individual localisations relative to the trajectories of the translated scatterers. The tracks with motion along the x -direction were used to estimate the precision in y . The tracks with motion along the y -direction were used to estimate the precision in x . Both datasets were used to estimate the precision in z . To visualise the variation, the mean x -, y - and z -coordinate were subtracted from each individual track, to centre the tracks around the transducer coordinate-system origin. This is illustrated in Fig. 8, where two cross-planes (x - y and x - z) are shown for the tracks with motion along the x -axis, which corresponds to the tracks in Fig. 6 a). The colour of the points represent the

Table V
ESTIMATED PRECISION FOR THE SUPER-RESOLUTION ALGORITHM.

	Average trajectory	Individual trajectories
$\tilde{\sigma}_x$ [μm]	17.7	17.3
$\tilde{\sigma}_y$ [μm]	27.6	19.3
$\tilde{\sigma}_z$ [μm]	9.5	8.7

tracks of the different design points, and are matched to those of the tracks in Fig. 6 a). The movement was uni-axial along the translation stage x -axis. However, slight misalignment between the ultrasound transducer and the translation stage have resulted in the localisation tracks not being perfectly aligned to the transducer axes. This can be observed in Fig. 8, in which the black line is the average trajectory of all tracks in the dataset. It should be noted however, that the axes are not equally scaled in the main plots, but only in the small inserts. The misalignment angle is 0.49° in the x - y plane, and 0.79° in the x - z plane. This misalignment should be compensated for when determining the variation of the tracks. The scatterers are fixed in the phantom and have been moved collectively by the translation stage. Then all tracks should have moved in the same direction, and the average trajectory of the tracks would be a good estimate of that. An estimate of the precision could be determined as the variation relative to the average trajectory. The precision along all three dimensions based on the variability relative to the average trajectory is displayed in Table V (“Average trajectory”). However, the coloured lines indicate that the tracks are in fact not parallel, but at small angles to each other. It is fairly small angles relative to the average trajectory, with the largest angle in any plane being 3.1° . This indicates that there is an error somewhere in the SRI pipeline, and that determining the precision relative to the average trajectory might be misleading. As an alternative, the estimate of the precision could be determined relative to the individual trajectories of the tracks. The precision along all three dimensions based on the variability relative to the individual trajectories is displayed in Table V (“Individual trajectories”). However, given that the tracks should have been parallel, this latter estimate of the precision might also be misleading. It is expected that the two presented estimates of the precision are limiting cases, and that the true precision of the SRI pipeline will lie somewhere in between.

IV. DISCUSSION

The 3D printed phantoms have successfully been used for SRI pipeline characterisation. The presented results illustrate that it is possible to obtain estimates for precision and accuracy, using these specialised phantoms. The obtained precision is an improvement of at least a factor of 18 compared to the ultrasound wavelength. It is particularly worth noting that even the worst obtained estimates for precision are comparable to the size of the smallest vessels in tissue. Thereby it is clear that the used method is suitable for resolving features at the size of the smallest vessels in tissue in three dimensions, and the stability of the phantom features allows for documentation of this.

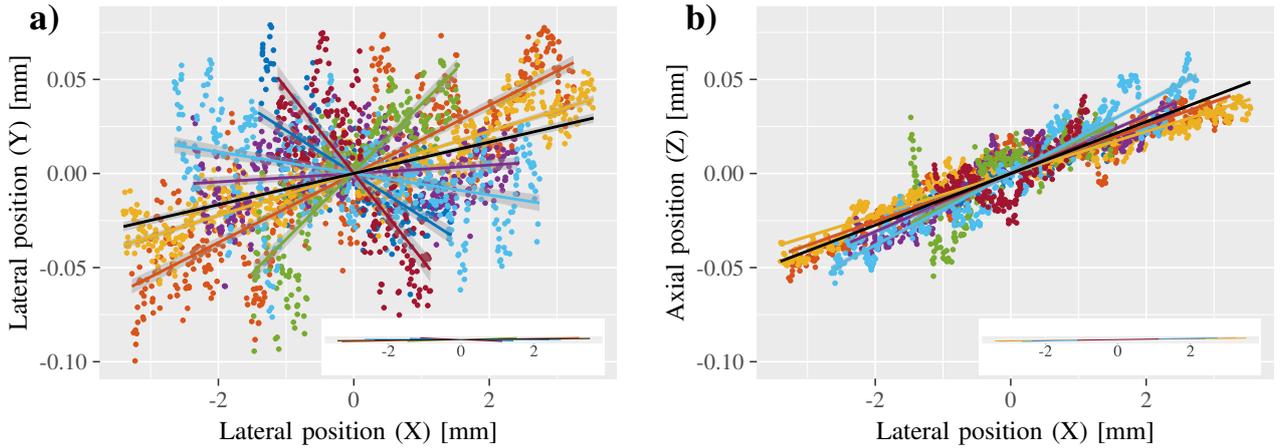


Figure 8. Crossplanes of the tracks with motion along the x -axis, offset to be centred around the coordinate system origin. The black lines show the average trajectory, while the coloured lines are linear fits to the individual trajectories of the different scatterers. The main plot axes do not have equal scaling. The inserts shows the same linear fits, with equally scaled axes.

It should be noted that there is good agreement between the SRI pipeline property estimates determined using these calibration phantoms, and the estimates for precision presented in [16], which were obtained using a 3D printed channel phantom and micro-bubbles, using the same SRI pipeline. In that paper, the precision was found to be less than $23\ \mu\text{m}$ in all directions. Note that the precision was determined through analysis of the radial distribution of micro-bubbles within the channel. By doing that, all precision estimates are a mixture of the axial precision and lateral precision, which are not expected to be equal. By instead measuring the distribution of fixated points as presented here, the analysis is not limited to a radial distribution, but independent estimates of the precision along x , y , and z have been obtained.

For the initial optical characterisation of the phantoms, a $36.6\ \mu\text{m}$ residual error was found for the correlation between the designed distances and those measured using an optical microscope. This is significantly larger than the position repeatability claimed by the microscope stage manufacturer and the experimentally validated position repeatability which was tested. A possible explanation might be that the experiment to determine the position repeatability was made by locating the same scatterer multiple times. On the other hand, the correlation in Fig. 4 was made localising many different scatterers. Local distortion of the printed structures might make the scatterer shapes slightly unequal, resulting in localisation of comparative features (for instance a specific corner) more difficult between scatterers, than when locating the same feature on the same scatterer. It should be noted that the model diagnostics showed that the residuals appeared to be normally distributed, indicating that the model is a good descriptor for the phantom expansion.

The high positioning control has allowed for the detection of distortion in the SRI pipeline, through the non-parallel tracks, which would not have been possible using conventional phantoms. The tracks should have been parallel given that the scatterers are fixated in the phantom, and that they have only been moved collectively using the translation stage.

The distortion is the reason for the discrepancy between the precision estimates. However, it was quite small with an angular distortion of at most 3.1° . A possible explanation could be that the experiment has been conducted assuming a speed of sound of $1480\ \text{m/s}$ in the entire beamformed volume. This was chosen, since the phantom was submerged in water, and the phantom itself consists of $\approx 75\%$ water. However, the speed of sound of the phantom has been measured to be $\approx 1580\ \text{m/s}$, which will lead to distortion.

An alternative or additional explanation could be that the ultrasound system has both a spatially dependent sensitivity and a spatially dependent point spread function, which changes in shape and intensity. This would not only explain the non-parallel tracks, but could also explain the difference in the number of tracks detected in the two ultrasound experiments, and that the eighth scatterer was not localised in either experiment. A consequence of a spatially dependent point spread function could be that full calibration of a SRI pipeline should perhaps be performed with local parameter estimates throughout the field of view of the probe instead of globally, as presented here. Thus the properties of a SRI pipeline would then be given by accuracy and precision estimates, both as functions of the x , y , and z coordinates. This might even be necessary, illustrated by the results in this paper, as proper thresholding can become difficult to implement globally in the field of view.

The presented phantom illustrates an alternative solution for SRI pipeline calibration to regular tube phantoms. However, this does not mean that it is irrelevant to create phantoms, which allow flow of micro-bubbles to be tracked. Given that it is a 3D printing method, these could easily be made, as was demonstrated in previous work [15], [16]. The 3D printing method allows for creating any arbitrary complex structure, making it possible to mimic complex vascular systems. The complexity is in principle only limited by the printer specifications and the printing material governing the achievable minimum voxel size.

The presented phantom concept could be expanded to

investigate other aspects of super-resolution algorithms and systems, such as resolvability and separability. The separation of 3 mm was chosen to ensure no overlap between the reflected signals from the individual scatterers, thereby mimicking how many SRI pipelines work today. Thereby, the resolution that can be expected from a SRI pipeline will be given by the variability of the positions, presented here as the σ values in Table V. Phantoms could be developed with scatterers placed much closer, to tune algorithms to be able to separate signals from partially overlapping reflections. This would be highly relevant for instance for some of the new types of SRI schemes which seek to be able to separate reflections much closer than the wavelength [26].

The 205 by 205 by 200 μm^3 scatterer size limits the phantom to be used with imaging frequencies equal to or less than 6 MHz. However, imaging probes capable of using larger frequencies are widely used. The phantom feature sizes are by no means the limit of the printing system. In [26] we presented another 3D printed phantom containing 45 by 45 by 1000 μm^3 scatterers for 2D imaging. Integration of the signal across the elevation plane allows for an increased intensity even though the scatterers were significantly smaller in cross-section than those used in this work. That is not possible for 3D imaging. However, no optimization of the scatterer size has been done for this work. Additional optimization through local dose changes resulting in local acoustic parameter changes might allow for obtaining even larger intensities from the same sized scatterer. A different approach would be to not only consider increasing the intensity from the scatterers, but also decreasing the background noise from the bulk of the phantom. Some of the unintended structures observed in ultrasound might originate from issues in the printer system, which could potentially be optimised.

The precision, accuracy and repeatability of the 3D printed phantoms would be incredibly difficult to achieve, if not impossible, using the traditional types of tube phantoms or chicken embryos. Yet, it still provides the opportunity of creating complex three-dimensional phantom features, providing the opportunity for full volumetric characterisation of an ultrasound system, which is not offered by any other phantom fabrication method available today.

V. CONCLUSION

We have presented 3D printed micro-phantoms containing absolute 3D micro-positioned sub-wavelength sized ultrasound scatterers for calibration of super-resolution ultrasound imaging (SRI) pipelines. The presented phantoms contains fixated scatterers and can therefore be used for long acquisitions, producing repeatable results, unlike traditional tube phantoms. The resulting printed structures have been characterised using an optical microscope, and it has been shown that the printed structures systematically expand isotropically by 2.6% relative to the design. The phantom was used to calibrate a super-resolution ultrasound imaging (SRI) pipeline by correlating the distances between the cavities calculated through the SRI pipeline to the phantom design distances. The analysis showed a correlation slope of 0.989, close

to a perfect correlation of 1. The variability of the super-localised positions of the individual scatterers across 640 volumes were used as an estimate of the precision of the SRI pipeline. Based on the analysis, it is expected that the precision of the SRI pipeline lies between the two limiting estimates of $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.7 \mu\text{m}, 27.6 \mu\text{m}, 9.5 \mu\text{m})$, when estimated relative to the average trajectory off all tracks, and $(\tilde{\sigma}_x, \tilde{\sigma}_y, \tilde{\sigma}_z) = (17.3 \mu\text{m}, 19.3 \mu\text{m}, 8.7 \mu\text{m})$, when estimated relative to the individual track trajectories. Both of these precision estimates are on the same scale as the features intended to be resolved *in vivo*, namely vessels of only a few tens of micrometers.

The presented phantom has proven to be a useful tool in validating accuracy and precision for a SRI pipeline, as well as unveiling distortion in the SRI pipeline, the latter of which would have been impossible using a traditional tube phantom. The study demonstrates the use of 3D printed phantoms for determining the accuracy and precision of volumetric super-resolution algorithms.

ACKNOWLEDGMENT

This work was financially supported by grant 7050-00004B from Innovation Fund Denmark, and from BK Medical, Herlev, Denmark.

REFERENCES

- [1] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433–440, February 2015.
- [2] M. A. O'Reilly and K. Hynynen, "A super-resolution ultrasound method for brain vascular mapping," *Med. Phys.*, vol. 40, no. 11, pp. 110 701–7, 2013.
- [3] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture, and M. Tanter, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, pp. 499–502, November 2015.
- [4] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 39, no. 2, pp. 262–267, 1992.
- [5] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comp.*, vol. 10th Nordic-Baltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1, pp. 351–353, 1996.
- [6] —, "A multi-threaded version of Field II," in *Proc. IEEE Ultrason. Symp.* IEEE, 2014, pp. 2229–2232.
- [7] O. M. Viessmann, R. J. Eckersley, K. Christensen-Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, pp. 6447–6458, 2013.
- [8] K. Christensen-Jeffries, J. Brown, P. Aljabar, M. Tang, C. Dunsby, and R. J. Eckersley, "3-D in vitro acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 64, no. 10, pp. 1478–1486, 2017.
- [9] H. C. H. Ko, B. K. Milthorpe, and C. D. McFarland, "Engineering thick tissues - the vascularisation problem," *Eur. Cells Mater.*, vol. 14, pp. 1–19, 2007.
- [10] Y. Desailly, O. Couture, M. Fink, and M. Tanter, "Sono-activated ultrasound localization microscopy," *Appl. Phys. Lett.*, vol. 103, no. 17, p. 174107, 2013.
- [11] C. Huang, M. R. Lowerison, F. Lucien, P. Gong, D. Wang, P. Song, and S. Chen, "Noninvasive contrast-free 3D evaluation of tumor angiogenesis with ultrasensitive ultrasound microvessel imaging," *Scientific Reports*, vol. 9, no. 1, p. 4907, 2019.
- [12] G. Liu, W. Qi, L. Yu, and Z. Chen, "Optimized doppler optical coherence tomography for choroidal capillary vasculature imaging," in *Proc. SPIE*, vol. 7889, 2011, pp. 1–10.

- [13] H. Gong, B. P. Bickham, A. T. Woolley, and G. P. Nordin, "Custom 3D printer and resin for $18\ \mu\text{m} \times 20\ \mu\text{m}$ microfluidic flow channels," *Lab Chip*, vol. 17, no. 17, 2017.
- [14] J. R. Jacquet, F. Ossant, F. Levassort, and J. M. Grégoire, "3-D-printed phantom fabricated by photopolymer jetting technology for high-frequency ultrasound imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 65, no. 6, pp. 1048–1055, 2018.
- [15] M. L. Ommen, M. Schou, R. Zhang, C. A. V. Hoyos, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, "3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [16] J. A. Jensen, M. L. Ommen, S. H. Øygaard, M. Schou, T. Sams, M. B. Stuart, C. Beers, E. V. Thomsen, N. B. Larsen, and B. G. Tomov, "Three-dimensional super resolution imaging using a row-column array," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, pp. 1–9, 2019.
- [17] J. A. Jensen, *Estimation of Blood Velocities Using Ultrasound: A Signal Processing Approach*, 3rd ed. Department of Electrical Engineering, 2013.
- [18] R. A. Crescenti, J. C. Bamber, M. Partridge, N. L. Bush, and S. Webb, "Characterization of the ultrasonic attenuation coefficient and its frequency dependence in a polymer gel dosimeter," *Phys. Med. Biol.*, vol. 52, pp. 6747–6759, 2007.
- [19] M. Jakovljevic, S. Hsieh, R. Ali, G. C. L. Kung, D. Hyun, and J. J. Dahl, "Local speed of sound estimation in tissue using pulse-echo ultrasound: Model-based approach," *J. Acoust. Soc. Am.*, vol. 144, no. 1, pp. 254–266, 2018.
- [20] H. Bouzari, M. Engholm, S. I. Nikolov, M. B. Stuart, E. V. Thomsen, and J. A. Jensen, "Imaging performance for two row-column arrays," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 7, pp. 1209–1221, 2019.
- [21] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, "SARUS: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.
- [22] M. F. Rasmussen, T. L. Christiansen, E. V. Thomsen, and J. A. Jensen, "3-D imaging using row-column-addressed arrays with integrated apodization — Part I: Apodization design and line element beamforming," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 5, pp. 947–958, 2015.
- [23] M. B. Stuart, M. Schou, and J. A. Jensen, "Row-column beamforming with dynamic apodizations on a GPU," in *Proc. SPIE Med. Imag.*, 2019, pp. 1–7, paper number 10955-20.
- [24] A. Kuznetsova, P. B. Brockhoff, and R. H. B. Christensen, "lmerTest package: Tests in linear mixed effects models," *Journal of Statistical Software*, vol. 82, no. 13, pp. 1–26, 2017.
- [25] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2018. [Online]. Available: <https://www.R-project.org/>
- [26] J. Youn, M. L. Ommen, M. B. Stuart, E. V. Thomsen, N. B. Larsen, and J. A. Jensen, "Ultrasound multiple point target detection and localization using deep learning," in *Proc. IEEE Ultrason. Symp.*, 2019.

B.2. PAPER J - DETECTION AND LOCALIZATION OF ULTRASOUND SCATTERERS USING CONVOLUTIONAL NEURAL NETWORKS

B.2 Paper J - Detection and Localization of Ultrasound Scatterers Using Convolutional Neural Networks

Detection and Localization of Ultrasound Scatterers Using Convolutional Neural Networks

Jihwan Youn, Martin Lind Ommen, Matthias Bo Stuart, Erik Vilain Thomsen,
Niels Bent Larsen, Jørgen Arendt Jensen, *Fellow, IEEE*

Abstract—Delay-and-sum (DAS) beamforming is unable to identify individual scatterers when their density is so high that point spread functions overlap each other. This paper proposes a convolutional neural network (CNN)-based method to detect and localize high-density scatterers, some of which are closer than the resolution limit of DAS beamforming. A CNN was designed, which takes radio frequency (RF) channel data as input and returns non-overlapping Gaussian confidence maps. The scatterer positions were estimated from the confidence maps by identifying local maxima. The RF channel data for training, validation, and evaluation were simulated in Field II pro by placing scatterers randomly in the region of interest, and transmitting three steered plane waves. Evaluation was performed on the simulated test sets at the scatterer densities from 0.49 mm^{-2} to 4.88 mm^{-2} . A precision of 0.999 and a recall of 0.911 were achieved, and localization uncertainties after excluding outliers were $\pm 46 \mu\text{m}$ (0.16λ) (outlier ratio: 0.04) in the lateral direction and $\pm 26 \mu\text{m}$ (0.09λ) (outlier ratio: 0.01) in the axial direction. Also, a mean resolved rate of 0.67 was achieved for the two scatterers laying closer than the theoretical resolution limit of DAS beamforming. Two PEGDA 700 g/mol hydrogel phantoms containing cavities were 3-D printed, and were imaged using a 5.2 MHz linear array transducer to test the method on measured data. The RF channel data were acquired by the synthetic aperture real-time ultrasound system (SARUS). A new CNN was trained for the phantom study using the modified training sets according to the physical properties of the phantom. On the grid scatterer phantom, a precision of 0.98 and a recall of 1.00 were achieved and localization uncertainties after excluding outliers were $\pm 101 \mu\text{m}$ (0.33λ) (outlier ratio: 0.01) in the lateral direction and $\pm 37 \mu\text{m}$ (0.12λ) (outlier ratio: 0.01) in the axial direction. On the random scatterer phantom, a precision of 0.59 and a recall of 0.63 were achieved, and localization uncertainties after excluding outliers were $\pm 132 \mu\text{m}$ (0.43λ) (outlier ratio: 0) in the lateral direction and $\pm 44 \mu\text{m}$ ($0.70 \pm 0.15 \lambda$) (outlier ratio: 0) with a bias of $22 \mu\text{m}$ in the axial direction. This method can potentially be extended to detect highly concentrated microbubbles in order to shorten data acquisition times of super-resolution ultrasound imaging.

Index Terms—high-density scatterers, convolutional neural network, super-resolution ultrasound imaging, ultrasound localization microscopy

I. INTRODUCTION

DELAY-AND-SUM (DAS) beamforming [1] is simple and effective for B-mode image generation, but the spatial resolution is limited by wave diffraction. The resolution of conventional ultrasound imaging depends on wavelength, f-number, and excitation pulse bandwidth. Recently, ultra-

sound localization microscopy (ULM) and the resulting super-resolution ultrasound imaging (SRI) was devised to overcome the diffraction limit [2]–[6]. The microvasculature, composed of vessels that are separated by less than a half-wavelength, was mapped by deploying microbubbles (MBs) as contrast agents. SRI can be achieved by detecting and tracking the centroids of individual MBs over time.

ULM-based SRI, however, requires long data acquisition times since the MB detection still relies on conventional ultrasound images. The ultrasound images are generally DAS beamformed and diffraction-limited as a consequence. Therefore, the MB concentration should be low to avoid the overlaps of point spread functions (PSFs) for accurate and reliable MB detection. This constrains the number of detectable MBs in a frame, and it leads to long data acquisition times for mapping the entire target structure.

A novel method is proposed in this paper to detect and localize high-density scatterers by using convolutional neural networks (CNNs). Deep learning has had a profound impact on processing complex data and making associated decisions. By training deep neural networks with a large number of examples, impressive improvements were achieved in various challenging problems such as image classification [7]–[10], object detection [11], [12], semantic segmentation [13]–[15], and single-image super-resolution [16], [17]. It would be nearly impossible to attain those improvements using traditional logic programming or model-based approaches. The same principles can be applicable to ultrasound signals. It is hypothesized that the CNN-based method can identify scatterers laying closer than the resolution limit of DAS beamforming directly from radio frequency (RF) channel data.

In optics, where localization microscopy was firstly proposed [18]–[20], several studies were conducted to incorporate deep learning in super-resolution localization microscopy [21]–[23]. Those studies used CNNs to localize fluorescent molecules and showed that deep learning-based methods can drastically reduce data acquisition times and data processing times while achieving state-of-the-art performance.

Similar attempts also exist in ultrasound SRI. Van Sloun *et al* [24] proposed Deep-ULM that outputs high-resolution images where the pixel values correspond to scattering intensities, given image patches of contrast-enhanced ultrasound (CEUS) acquisitions. This is similar to our approach in the sense that it handles high-density scatterer detection using CNNs but Deep-ULM takes beamformed signals as input, whereas the proposed method only uses RF channel data without beamforming. Allman *et al* [25] tried to locate and

This work was supported in part by the Fondation Idella.

The authors are with the Department of Health Technology, Technical University of Denmark, 2800 Kogens, Lyngby, Denmark (email: jihyoun@dtu.dk).

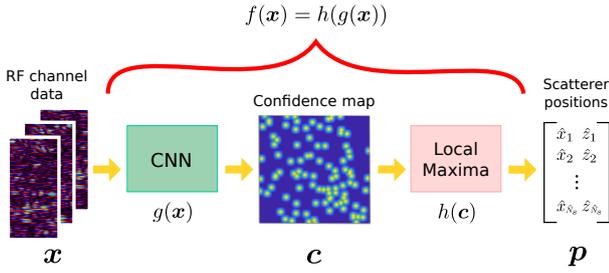


Fig. 1. Overview of the CNN-based scatterer detection and localization.

classify sources and artifacts from pre-beamformed photoacoustic channel data using Faster R-CNN [26] with VGG16 [27]. However, only up to 10 sources were considered, and classification for artifact removal is not necessary for scatterer detection.

Deep learning techniques can be used to achieve better ultrasound image quality. A fully connected neural network beamformer improved image contrast by suppressing off-axis scattering [28]. Hyun *et al* [29] proposed a CNN beamformer that reduces speckle and eventually enhances contrast while preserving resolution. Generative Adversarial Network (GAN) [30], an architecture that generates output following the same distribution as the training dataset, were applied to improve image quality without sacrificing frame rate. Multi-focus line-by-line images were synthesized from single-focus line-by-line images [31] and image quality comparable to using thirty one plane waves was achieved using three plane waves [32].

In this work, CNNs were trained to learn a mapping from RF channel data to confidence maps, and scatterer positions were then estimated from the confidence maps by identifying local maxima. The RF channel data were directly fed to the CNNs without beamforming to avoid the information loss caused by overlapping PSFs. We have shown the potential of the CNN-based method in [33]. However, the training was performed at a fixed scatterer density and its performance was not fully investigated. In this paper, two CNNs were trained and evaluated using simulated RF channel data with one plane wave or three plane waves. The training sets were generated at four different scatterer densities, and the test sets were generated at ten different scatterer densities. Two phantoms with water-filled cavities were 3-D printed and scanned to examine the feasibility of the CNN method on measured data.

II. METHODS

Consider RF channel data $\mathbf{x} \in \mathbb{R}^{N_a \times N_l \times N_t}$ induced by scatterers $\mathbf{p} \in \mathbb{R}^{N_s \times 2}$ where N_a is the number of samples along the axial direction, N_l is the number of active elements of a transducer in reception, N_t is the number of transmissions, N_s is the number of scatterers, and 2 is the number of the spatial dimensions (in the lateral and axial positions). The nonlinear mapping $f: \mathbb{R}^{N_a \times N_l \times N_t} \rightarrow \mathbb{R}^{N_s \times 2}$ needs to be found to estimate scatterer positions from the RF channel data, which satisfies

$$\mathbf{p} = f(\mathbf{x}). \quad (1)$$

TABLE I
RF CHANNEL DATA SIMULATION PARAMETERS

Category	Parameter	Value
Transducer	Center frequency	5.2 MHz
	Pitch	0.20 mm
	Element width	0.18 mm
	Element height	6 mm
	Number of elements	192
	Imaging	Number of TX elements
Number of RX elements (N_l)		64
Steered angles		$-15^\circ, 0^\circ, 15^\circ$
Environment	Speed of sound (c)	1480 m/s
	Field II sampling frequency	120 MHz
	RF data sampling frequency	29.6 MHz
Scatterer	Number of scatterers (N_s)	$20 \cdot i, \forall i \in \{1, 2, \dots, 10\}$
	Lateral position range	$(-3.2, 3.2)$ mm
	Axial position range	$(14.8, 21.2)$ mm

Here, N_s varies depending on the given RF channel data \mathbf{x} so the mapping f needs to adjust N_s adaptively, but this is not straightforward. Therefore, the mapping f is decomposed into two functions g and h to handle the varying N_s . The mapping $g: \mathbb{R}^{N_a \times N_l \times N_t} \rightarrow \mathbb{R}^{N_h \times N_w}$ forms a confidence map $\mathbf{c} \in \mathbb{R}^{N_h \times N_w}$ where N_h and N_w are the number of samples in the axial and lateral directions, respectively. The confidence map \mathbf{c} represents a region of interest (ROI) where the pixel values indicate confidences of scatterer presence in each pixel. The mapping $h: \mathbb{R}^{N_h \times N_w} \rightarrow \mathbb{R}^{N_s \times 2}$ detects and locates scatterers from the confidence map. The mapping in (1) can be rewritten using g and h as follows:

$$\begin{aligned} \mathbf{p} &= f(\mathbf{x}) \\ &= h(g(\mathbf{x})) = h(\mathbf{c}), \end{aligned} \quad (2)$$

where

$$\mathbf{c} = g(\mathbf{x}). \quad (3)$$

The overview of the proposed method is illustrated in Fig. 1. The mapping g was modeled by a fully CNN and the mapping h corresponded to local maxima identification with thresholding. The RF channel data simulation and confidence map generation are explained in Section II-A and II-B, respectively. The architecture of the proposed CNN is introduced in Section II-C. Scatterer detection from the confidence maps is explained in II-D and the phantom fabrication is described in Section II-E. A baseline method for comparison is introduced in Section II-F.

A. RF Channel Data Simulation

Field II pro [34]–[36] was used to simulate RF channel data to generate data sets for training, validation and evaluation. The parameters for the simulation are listed in Table I. The transducer was modeled after a commercial 5.2 MHz 192-element linear array transducer, and a measured impulse response [37] was applied to make the simulated RF channel data as close to measured data as possible [38].

For each frame, a certain number of point scatterers were placed randomly within a region of 6.4 mm \times 6.4 mm where

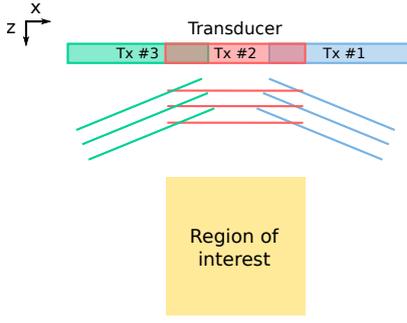


Fig. 2. An illustration of the imaging scheme. Scatterers were placed in the region of interest, and three steered plane waves were transmitted for each frame. The aperture was shifted to insonify only the region of interest.

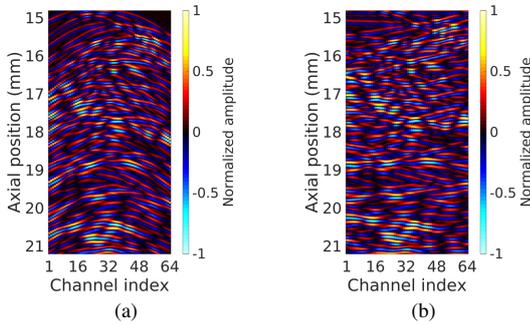


Fig. 3. An example of simulated RF channel data with one plane wave without steering. (a) is simulated raw RF channel data and (b) is delayed RF channel data. Note that the delay here is different from the delay in DAS beamforming.

the center of the region was 18 mm away from the transducer, and three steered plane waves were transmitted using 32 elements. Individual scatterers had the same scattering intensity. Motion and flow were not considered, therefore, scatterers were static among three plane wave transmissions in a frame and the scatterer positions were independent between frames. The aperture was shifted for each steered angle to insonify only the ROI, as shown in Fig. 2. The elements used in transmission were the 105th to the 136th (-15°), the 81st to the 112nd (0°), and the 57th to the 88th (15°) elements. Backscattered waves were received with 64 elements in the center of the transducer.

The simulated RF channel data were not beamformed but delayed based on the time-of-flight calculated by

$$\tau_i(x, z) = \left(\sqrt{(x - x_i)^2 + z^2} + z \right) / c. \quad (4)$$

Here, τ_i is the time-of-flight of the i -th transmission, (x, z) is the data point, x_i is the center of the i -th transmission aperture, and c is the speed of sound. This preprocessing helped the CNN solve the problem by making wavefronts more like straight lines, instead of parabolas, as shown in Fig. 3.

It is required for the input and output of the proposed CNN to have the same number of samples along the axial direction. The delayed RF channel data were accordingly resampled to match the same number of samples as confidence maps along the axial direction ($N_a = N_h$). Essentially, the sampling

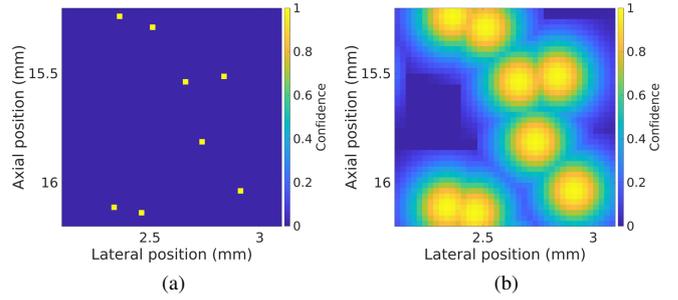


Fig. 4. An example of cropped confidence maps. (a) is a binary confidence map and (b) is a non-overlapping Gaussian confidence map created from (a).

frequency of the RF channel data was determined by the pixel size of the confidence maps, and N_a was determined by the sampling frequency and the ROI. After preprocessing, the size of RF channel data \mathbf{x} for one frame was $256 \times 64 \times 3$ before being fed to a CNN.

B. Non-overlapping Gaussian Confidence Map

Binary confidence maps were firstly created, whose pixel values indicate presence (1) or absence (0) of a scatterer in the corresponding location, as shown in Fig. 4a. CNNs, however, could not be trained using such confidence maps because most of their pixel values were zero. The sparse confidence maps provided small gradients during optimization and made CNNs easily converge to wrong optimal solutions, returning only zero confidence maps regardless of input.

A non-overlapping Gaussian confidence map (Fig. 4b) was proposed to solve the imbalance problem of binary confidence maps. Applying 2-D Gaussian filtering to sparse labels can improve training stability and guide CNNs to correct solutions [21], [24], [39]. But simply applying 2-D Gaussian filtering is problematic since peaks do not correspond to scatterer positions when scatterers are closer than a certain distance. A 1-D example is shown in Fig. 5a. To keep peaks at scatterer positions in the confidence maps, the Gaussian filter was applied one by one at each scatterer position in the binary confidence maps. Notably, when the Gaussian filter values induced by different scatterers were overlapped, the maximum values were taken. By doing so, clearly separated peaks can be obtained at the true scatterer positions, as shown in Fig. 5b.

The parameters for non-overlapping Gaussian confidence maps are listed in Table II. The 2-D Gaussian filter is defined by

$$G(u, v; \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}} \quad (5)$$

where u and v are the pixel distances from the scatterer position in the lateral and axial directions, respectively, and σ is the standard deviation. The filter size was fixed to $4\sigma+1$ and the standard deviation was chosen by cross-validation among 3, 5, and 7 pixels. Scatterer positions were quantized according to pixel size since the confidence maps are on discrete grids. Here, the pixel size was set to $25 \mu\text{m}$ ($\approx \lambda/10$); the lateral and axial localization uncertainties are $\pm 12.5 \mu\text{m}$ in ideal situation. The confidence map size was 256×256 ($N_w = N_h = 256$) given the pixel size and the area of the ROI.

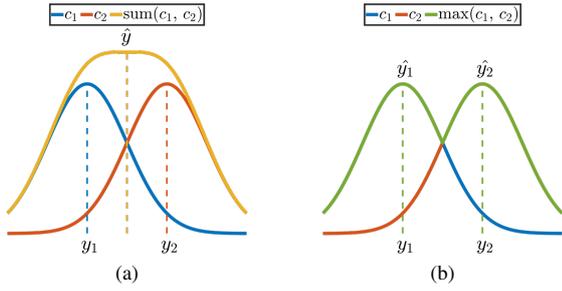


Fig. 5. A comparison of 1-D Gaussian confidence maps created by (a) summation and (b) maximum operation. There are two scatterers y_1 and y_2 , and c_1 and c_2 are their confidence maps, respectively. The yellow line in (a) is the sum of c_1 and c_2 . The green line in (b) is the maximum of c_1 and c_2 . In (a), one scatterer \hat{y} is found at a wrong position whereas in (b), two scatterers \hat{y}_1 and \hat{y}_2 can be recovered at correct positions in the confidence map.

TABLE II
CONFIDENCE MAP PARAMETERS

Parameter	Value
Pixel size	25 μm
Confidence map size ($N_h \times N_w$)	256 \times 256
Gaussian filter size	21 pixels
Gaussian filter standard deviation	5 pixels

C. Convolutional Neural Network Architecture

The proposed CNN has an encoder-decoder structure with pooling and unpooling. The encoder-decoder structure was adopted to transform the input in the channel data domain to the confidence maps in the ultrasound image domain. In the encoding path, information is extracted from the RF channel data, and in the decoding path, the confidence map is reconstructed based on the extracted information.

The overview of the CNN architecture and its components are shown in Fig. 6. It mainly consists of four *down-blocks*, one *conv-block*, and four *up-blocks*. In the *down-blocks*, the feature map size is decreased by strided convolution to reduce the amount of parameters, and in the *up-blocks*, the feature map size is increased to the confidence map size by pixel shuffle [40]. An 11×1 convolution layer prior to the encoding path extracts per-channel features, and two convolution layers after the decoding path refine the feature maps and return the confidence maps.

The pre-activation residual units [9] (Fig. 6a) were used instead of common convolution and rectified linear unit (ReLU) layers to improve the network performance. Batch normalization (BN) in the residual units helped ease the optimization, limited covariate shift, and had the effect of regularization [41]. Dropout [42] was additionally attached after the shortcut for further regularization. Leaky ReLU [43] and Sigmoid were chosen as non-linear activation. CoordConv [44] was added to transfer spatial information over convolution layers.

D. Scatterer Detection from Confidence Maps

The scatterer positions can be found by locating the pixels whose confidences are one in the true confidence map c .

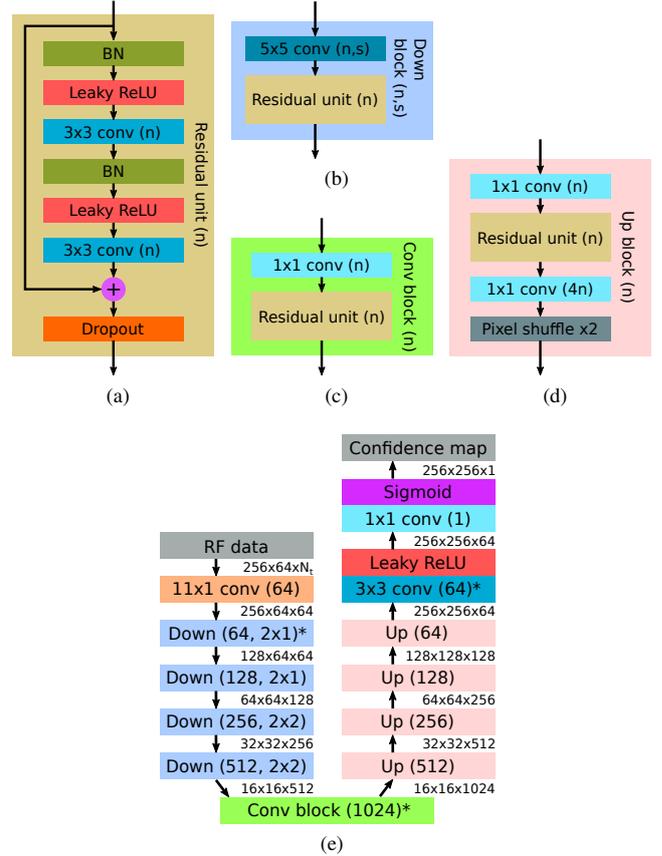


Fig. 6. The proposed network architecture and its components: (a) residual unit, (b) down-block, (c) conv-block, (d) up-block, and (e) the network overview. The n and s in the parenthesis are the number of kernels and stride. In (e), the sets of three numbers are the feature map size between two blocks, and the asterisk indicates that CoordConv was applied at the first convolution in the block.

However, the estimated confidence map $\hat{c} = g(x)$ acquired from a trained CNN is an approximation of c . It is not guaranteed that the confidences are one where scatterers are located in \hat{c} . Therefore, the algorithm relies on that fact that pixels containing scatterers are local peaks. The scatterer positions were recovered by finding the local maxima whose confidence is larger than a certain decision value. The chosen decision value was 0.9 in this work.

E. Phantom Fabrication

Two PEGDA 700 g/mol hydrogel phantoms were 3-D printed [45], [46] to assess the CNN method on measured data. The phantom contained water-filled cavities which acted as scatterers. The volume of each cavity was $45 \mu\text{m} \times 1000 \mu\text{m} \times 45 \mu\text{m}$. The cavities were designed to be elongated in the elevation direction to increase the intensity of received signals.

For the first phantom, 100 cavities were placed on a 10×10 grid with a spacing of $518 \mu\text{m}$ in the lateral direction and $342 \mu\text{m}$ in the axial direction, as illustrated in Fig. 7. This grid scatterer phantom had the spacing larger than the resolution limit of DAS to show that the CNN method works on measured data. The second phantom, on the other hand, had 100 cavities randomly distributed with a minimum spacing of $190 \mu\text{m}$

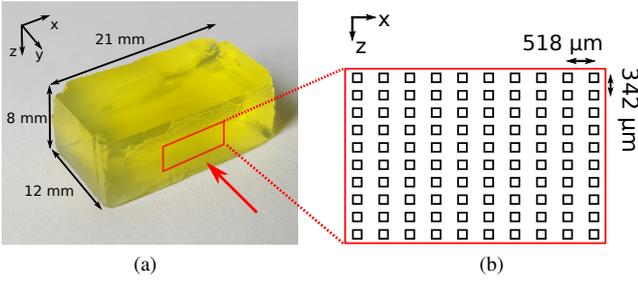


Fig. 7. Fabricated 3-D phantom with cavities: (a) photograph of the phantom and (b) 100 cavities placed on a 10×10 grid.

to present that the CNN method can resolve targets closer than the conventional resolution limit. The minimum spacing between cavities were constrained due to the cavity size and the 3-D printer voxel size.

F. Baseline Method

Local peak detection on the beamformed images was chosen as a baseline method for comparison. RF channel data were DAS beamformed in the region of interest with the same pixel size as the confidence map, and, for three plane wave transmissions, individual beamformed images were coherently compounded [47]. The baseline method detected and located scatterers in the envelop detected and log-compressed B-mode images with a dynamic range of 40 dB. The B-mode images were smoothed to avoid the situation where more than one pixel correspond to a peak, and scatterer positions were estimated by finding local maxima.

Deconvolution using an estimated PSF is one of the commonly used techniques for microbubble localization [5]. However, it was not considered in this work since its performance was sensitive to the parameters when the PSFs were overlapped, and the spatially varying PSF of ultrasound imaging led to imprecise scatterer position estimation.

III. EXPERIMENTS

A. Training Details

CNNs, the mapping g in (2), were trained to return the corresponding confidence map c_i given RF channel data x_i by minimizing the mean squared error (MSE), given by

$$\mathcal{L}_{\text{MSE}}(x_i, c_i; g) = \frac{1}{N} \sum_{i=1}^N \|c_i - g(x_i)\|_F^2, \quad (6)$$

where N is the number of samples and $\|\cdot\|_F$ is the Frobenius norm.

One data set consisted of frames simulated at the same scatterer density, and four training sets and four validation sets were generated at the scatterer densities of 0.49 mm^{-2} , 0.98 mm^{-2} , 2.44 mm^{-2} , and 4.88 mm^{-2} , i.e., the numbers of scatterers are 20, 40, 100, and 200 in one frame, respectively. Each training set and validation set had 10240 and 1280 frames, respectively.

The kernel weights were initialized by orthogonal initialization [48] and optimized with ADAM [49] by setting

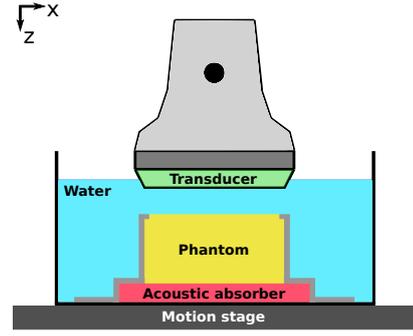


Fig. 8. Illustration of the experimental setup for phantom measurement.

$\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-7}$. Firstly, the training was performed only using the training set at the scatterer density of 2.44 mm^{-2} . The initial learning rate was 10^{-4} and it was halved every 100 epochs. After 600 epochs, the learning rate was set to 10^{-5} and the training continued using all the training sets while the learning rate was halved every 50 epochs. The mini-batch size was 32, and each batch was composed of frames from all four training sets after 600 epochs. The CNN was implemented in Python using Tensorflow [50], and were trained on a server equipped with a NVIDIA TESLA V100 16 GB PCIe graphics card. The total number of training epochs was 800, and the training took approximately 40 hours.

During training, the RF channel data and confidence maps were flipped along the lateral direction at random with a probability of 0.5 to augment the training sets. White Gaussian noise was added to the RF channel data for generalization along with BN and dropout. The signal-to-noise ratio after noise addition was 6 dB, and the dropout rate was 0.3. The RF channel data and confidence maps were then normalized to be in the range $[-1, 1]$ and $[0, 1]$. Validation was performed every epoch to monitor the training, and also after training for cross-validation to choose hyper-parameters.

For both simulation and phantom experiment, two CNNs were trained and compared: one CNN acting on the data from one plane wave (0°) and the other CNN acting on the data from three plane waves ($-15^\circ, 0^\circ, 15^\circ$).

B. Simulation Experiment

The CNNs were evaluated on simulated test sets firstly. One test set consisted of 3840 frames simulated at the same scatterer density, and 10 test sets were created by varying the number of scatterers from 20 to 200 with intervals of 20. The parameters in Table I were used again, apart from the number of scatterers. The evaluation was performed in ten test sets to evaluate how the performance changes over different scatterer densities and how well the CNNs were generalized in terms of scatterer density.

C. 3-D Printed Phantom Experiment

1) *RF Channel Data Acquisition*: The 3-D printed phantoms were scanned using the 5.2 MHz 192-element linear array transducer which has the same parameters as in Table

I. The raw RF channel data were acquired by the synthetic aperture real-time ultrasound system (SARUS) experimental ultrasound scanner [51]. The same imaging scheme and processing as in the simulation were applied.

The experimental setup is shown in Fig. 8. The transducer was fixed, and a water tank containing the phantom was placed on a motion stage. The phantom was aligned with the transducer by the motion stage, capable of moving in the x - and y -axis, and rotating around the z -axis. During measurement, the motion stage was moved along the x -axis in steps of $50\ \mu\text{m}$ between frames, and 33 frames were acquired for each phantom experiment.

2) *Training Set Modification*: The training sets were modified and new CNNs were trained from scratch for the phantom experiment. In the simulation, it was assumed that scatterers were infinitesimally small points. However, the cavities in the phantom were squares, as shown in Fig. 7b, if the elevation direction is ignored. Scattering, therefore, happens twice at each cavity: once when a wave goes into a cavity and the other when the wave comes out of the cavity. Additionally, the first scattering experiences a phase reversal because the acoustic impedance of the phantom is higher than that of water.

RF channel data for training were, therefore, re-simulated by modeling each scatterer using two points separated by the cavity size axially and with a phase reversal. To remain consistent, the same scatterer positions of the original training set were used.

3) *Depth Correction*: The speed of sound in the phantom is higher than in water. The axial positions of the estimated scatterers were corrected to compensate for the different speed of sound in the phantom by

$$\hat{z}^* = (\hat{z} - d_{\text{pht}}) \cdot \frac{c_{\text{water}}}{c_{\text{pht}}} + d_{\text{pht}} \quad (7)$$

where \hat{z} and \hat{z}^* are the axial position before and after correction, c_{water} and c_{pht} are the speed of sound in water and in the phantom, respectively, and d_{pht} is the distance from the transducer to the phantom surface.

D. Evaluation Metrics

Three evaluation criteria were considered to assess the CNNs: detection, localization, and resolution. The positive and negative detections were firstly determined by pairing estimated scatterers with true scatterers based on their pair-wise distances, as stated in Algorithm 1. To be a positive detection, an estimated scatterer should be exclusively matched with a true scatterer within a certain localization precision. This localization precision can be translated to a target resolution for ULM without tracking and it was set to be the half of full width at half maximum (FWHM) in this work. More An ellipse whose major axis and minor axis are half of FWHM_x and half of FWHM_z , respectively, was used as the desired localization precision, where FWHM_x is the lateral FWHM and FWHM_z is the axial FWHM. This bi-directional matching process is extended from the left-right consistency check [52], [53] for stereo matching in computer vision. It conforms to the uniqueness constraint; one true scatterer can be paired with at most one estimated scatterer.

Algorithm 1 Algorithm for determining positive or negative detections

Input: $\mathbf{p} \in \mathbb{R}^{N_s \times 2}$ and $\hat{\mathbf{p}} \in \mathbb{R}^{\hat{N}_s \times 2}$, where \mathbf{p} is true scatterer positions and $\hat{\mathbf{p}}$ is estimated scatterer positions

Output: Positive or negative detection $\mathbf{a} \in \mathbb{R}^{\hat{N}_s \times 1}$

```

1:  $\mathbf{a} \leftarrow \mathbf{0} \in \mathbb{R}^{\hat{N}_s \times 1}$ 
2:  $D \leftarrow \left\{ (d_{ij}) \in \mathbb{R}^{N_s \times \hat{N}_s} \mid d_{ij} = \|\mathbf{p}_i - \hat{\mathbf{p}}_j\|_2 \right\}$ 
3: for  $j = 1$  to  $\hat{N}_s$  do
4:    $\hat{i} \leftarrow \arg \min D_{*,j}$ 
5:   if  $j = \arg \min D_{\hat{i},*}$  and  $\frac{(p_{i1} - \hat{p}_{j1})^2}{(\text{FWHM}_x/2)^2} + \frac{(p_{i2} - \hat{p}_{j2})^2}{(\text{FWHM}_z/2)^2} < 1$ 
6:     then
7:        $a_j \leftarrow 1$ 
8:     else
9:        $a_j \leftarrow 0$ 
10:   end if
end for

```

The detection capability was assessed by quantifying wrong detections and missed detections using precision, recall, and F_1 score which are defined as follows:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (8)$$

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (9)$$

and

$$F_1 \text{ score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}, \quad (10)$$

where TP is the number of true positives (correct detections), FP is the number of false positives (wrong detections), and FN is the number of false negatives (missed detections).

Localization uncertainties were measured by calculating the lateral and axial position errors. Only positive detections were considered for the localization assessment.

Spatial resolution, the ability to separate two points that are close together, was investigated statistically. For two isolated true scatterers, it was checked whether they were detected. A pair of scatterers was set to *resolved* if both scatterers were detected. It was set to *non-resolved* if only one of them was detected. And it was not considered if none of them were detected, as this would be a detection problem. The resolved rates were calculated in $20\ \mu\text{m} \times 20\ \mu\text{m}$ bins by

$$\text{Resolved rate} = \frac{N_{\text{res}}}{N_{\text{res}} + N_{\text{non-res}}} \quad (11)$$

where N_{res} is the number of resolved pairs and $N_{\text{non-res}}$ is the number of non-resolved pairs in a bin.

IV. RESULTS

The CNN method results on the simulated data and the 3-D printed phantom measured data are presented in this Section. Quantitative evaluation comparing one plane wave and three plane waves was performed as specified in Section III-D. The results of the baseline method on the same test data are also presented.

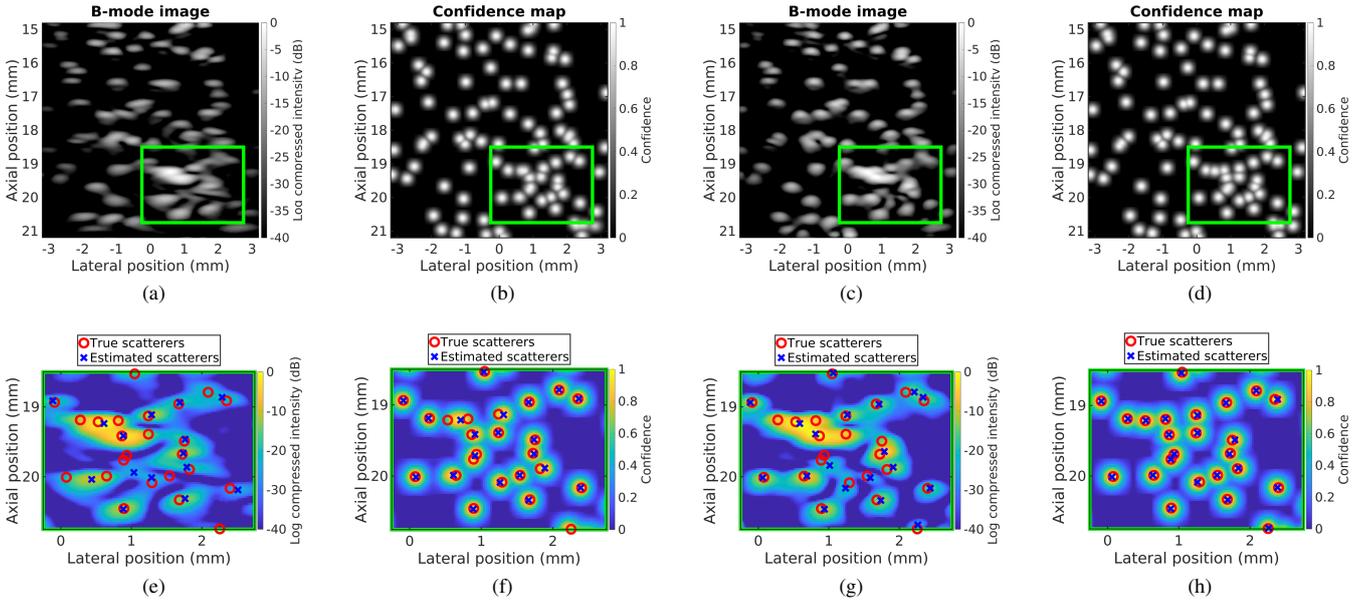


Fig. 9. Comparison of scatterer detection between baseline method and CNN method in a simulated test frame. (a) and (c) are DAS beamformed B-mode images with one and three plane waves, respectively. (b) and (d) are estimated confidence maps by CNNs with one and three plane waves, respectively. (e) - (h) show true scatterers and estimated scatterers from their corresponding results above in the same column in the green box region.

TABLE III
PRECISION, RECALL, AND F_1 SCORE COMPARISON IN THE SIMULATED TEST DATASETS

Method	One plane wave			Three plane waves		
	Precision	Recall	F_1	Precision	Recall	F_1
Peak	0.83	0.51	0.63	0.93	0.62	0.75
CNN	0.99	0.83	0.90	1.00	0.91	0.96

A. Simulation Experiment

The qualitative comparison between the peak and CNN methods is shown in Fig 9. The proposed CNN method successfully detected and localized high-density scatterers when the peak method failed due to overlaps of PSFs. This can also be confirmed quantitatively.

The detection results in the simulated test sets are shown in Table III. The CNN method achieved the better precision, recall, and F_1 score for both one and three plane transmissions. Also, when more number of transmissions were involved, the detection performance was improved for both methods. The detection capabilities over different scatterer densities were also investigated, as shown in Fig. 10. The recalls dropped as the scatterer density increased while the precisions were relatively kept high. Additionally, the recalls of the baseline method decreased more drastically as the scatterer density increased, which led to the lower F_1 scores.

For localization, box-and-whisker plots along with violin plots were used to display the results, as presented in Fig 11 and 12. The bottom and top edges of the blue boxes indicate the 25th (q_1) and 75th percentiles (q_3), and the center red lines indicate the medians. The whiskers, vertically extended lines from the boxes, indicate the range of values except outliers.

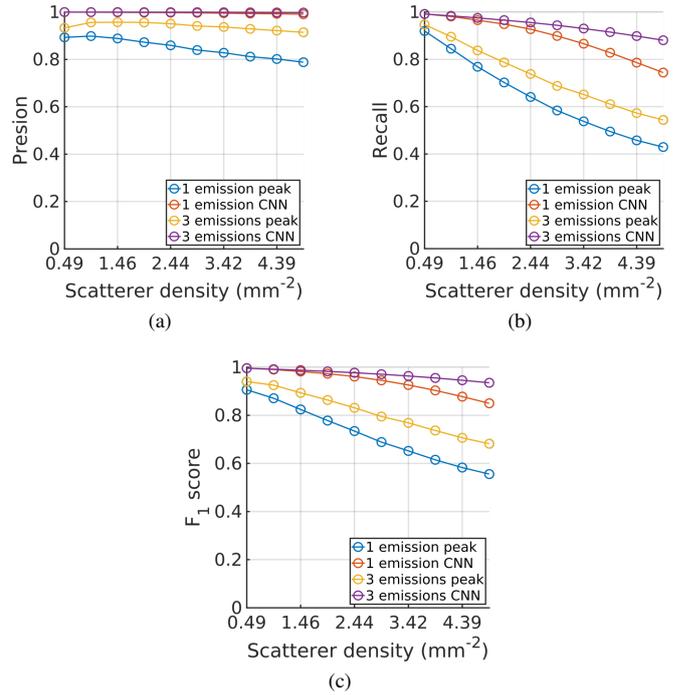


Fig. 10. Detection capabilities of the baseline and CNN methods over different scatterer densities with one plane wave and three plane waves: (a) precision, (b) recall, and (c) F_1 score.

The outliers are greater than $q_3 + 1.5 \times (q_3 - q_1)$ or less than $q_1 - 1.5 \times (q_3 - q_1)$. To show the error distribution directly, violin plot, the shaded area, was also demonstrated.

The lateral position error was higher than the axial position error for both methods, and the CNN method achieved more than two times better localization uncertainties than the

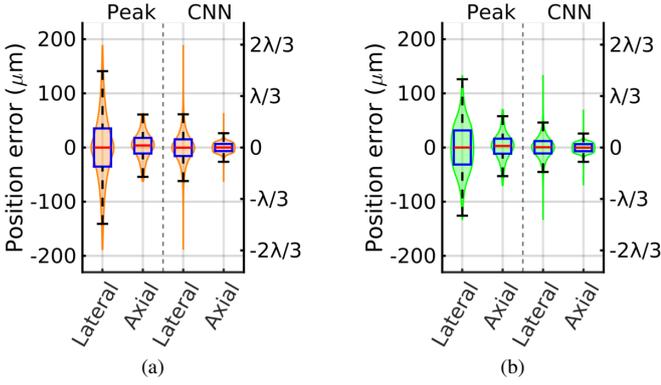


Fig. 11. Localization uncertainties of baseline and CNN methods on the simulated test sets. (a) and (b) are the results with one plane wave and three plane waves, respectively.

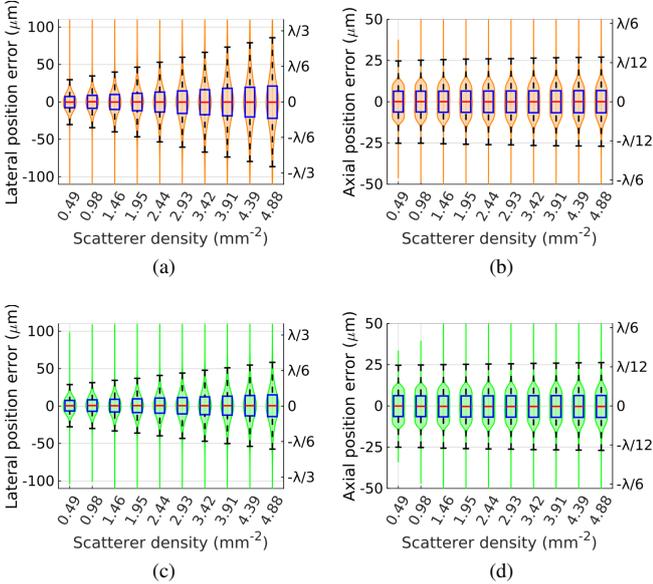


Fig. 12. Localization uncertainties of the CNN method on simulated test sets at different scatterer densities: the lateral position errors with (a) one plane wave and (b) three plane waves, and the axial position errors with (c) one plane wave and (d) three plane waves.

baseline method, as shown in Fig. 11. The medians were mostly very close to zero, indicating that the scatterer position estimation was unbiased in both directions. The localization was also improved when more number of plane waves were transmitted. The CNN method localization uncertainties at different scatterer densities are shown in Fig. 12. Neither scatterer density nor number of emissions had much impact on the axial position errors. The lateral position errors, on the other hand, gradually increased as the scatterer density increased.

The 2-D histograms in Fig. 13 show the resolved rates of two isolated scatterers measured in $20\ \mu\text{m} \times 20\ \mu\text{m}$ bins. The green lines represent the theoretical resolution limit of DAS beamformed images, assuming that the 6 dB contour of a PSF is an ellipse. The FWHM was measured on a simulated PSF in the center of the ROI. For one plane wave, the FWHM

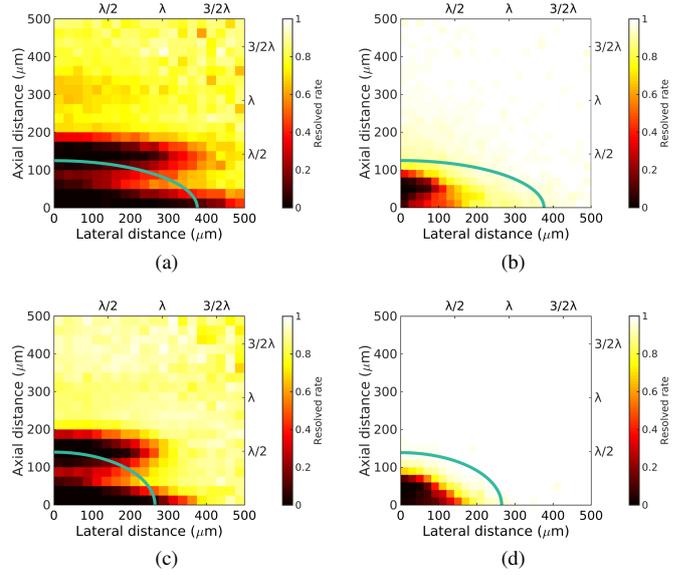


Fig. 13. Resolved rate of (a), (c) baseline methods and (b), (d) CNN methods in the simulated test sets where (a) and (b) are with one plane wave and (c) and (d) are with three plane waves. The green lines represent the theoretical resolution limit of DAS beamforming.

was $376\ \mu\text{m}$ ($1.32\ \lambda$) laterally and $125\ \mu\text{m}$ ($0.44\ \lambda$) axially. For three plane waves, the FWHM was $265\ \mu\text{m}$ ($0.93\ \lambda$) laterally and $140\ \mu\text{m}$ ($0.49\ \lambda$) axially. The resolution results clearly show that the CNN method can resolve scatterers closer than the limit. The mean resolved rates in the area under the green line for the peak and CNN methods were 0.16 and 0.68 with one plane wave, and 0.17 and 0.67 with three plane waves, respectively.

B. 3-D Printed Phantom Experiment

The qualitative results of the baseline and CNN methods on the grid and random scatterer phantoms are presented in Fig. 14, and their quantitative comparison is shown in Table IV and Fig. 15. With one plane wave, the side lobe level was so high that the DAS beamforming was unable to identify individual scatterers of the grid phantom properly, as shown in Fig. 14a. The low precision was achieved as peaks were shifted toward wrong directions due to the overlaps of PSFs. The CNN method also achieved poor results with one plane wave on the grid phantom data, as shown in Fig. 14b since the CNN was not generalized enough to handle regularly placed scatterers as the training frames were generated by placing scatterers randomly. Most of the scatterers in the first and the last columns were correctly detected while the other scatterers were missed so the precision was higher than the baseline but the recall was poorer. On the contrary, with three plane waves, the baseline method found all the individual scatterers without any false detection. The CNN method also achieved comparable detection and localization results with three plane waves, showing that more number of transmissions helped generalization of the CNN. On the random scatterer phantom, the CNN method achieved better detection and localization for both one and three plane waves.

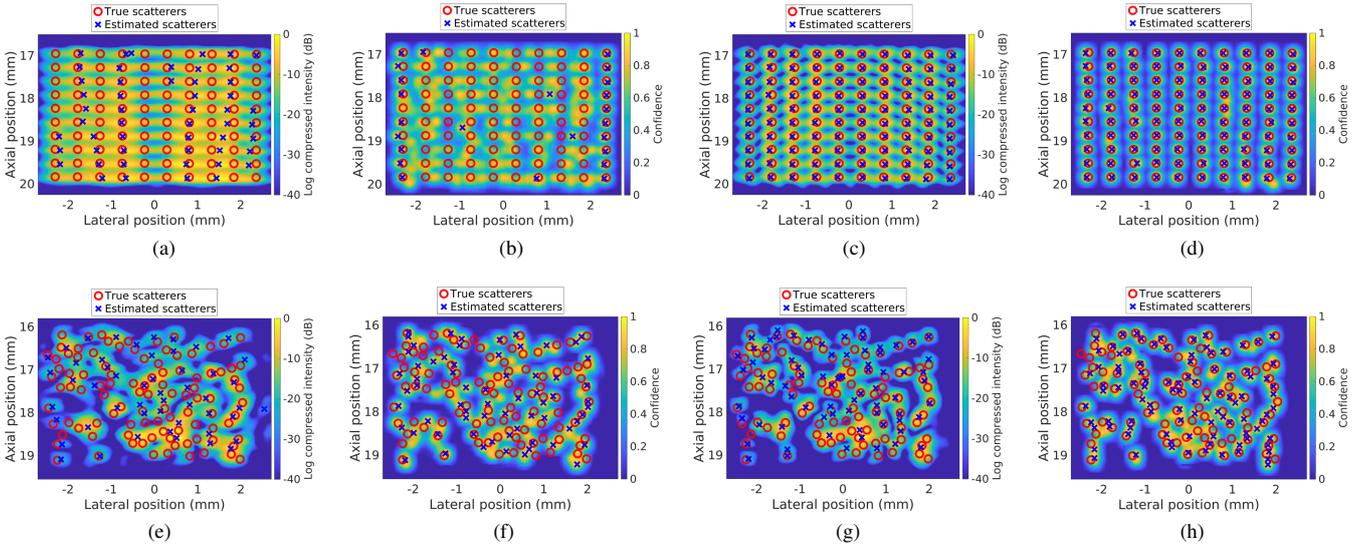


Fig. 14. Comparison of scatterer detection between baseline method and CNN method on phantom measured frames. (a) - (d) are results of the grid phantom and (e) - (h) are results of the random phantom. B-mode images with (a), (e) one plane wave and (c), (g) three plane waves and confidence maps with (b), (f) one plane wave and (d), (h) three plane waves are shown with true scatterers and estimated scatterers.

TABLE IV
PRECISION, RECALL, AND F_1 SCORE COMPARISON IN THE PHANTOM TEST DATASETS

Phantom	Method	One plane wave			Three plane waves		
		Precision	Recall	F_1	Precision	Recall	F_1
Grid	Peak	0.82	0.41	0.54	1.00	1.00	1.00
	CNN	0.89	0.22	0.35	0.98	1.00	0.98
Random	Peak	0.47	0.23	0.31	0.49	0.32	0.39
	CNN	0.53	0.37	0.44	0.59	0.63	0.61

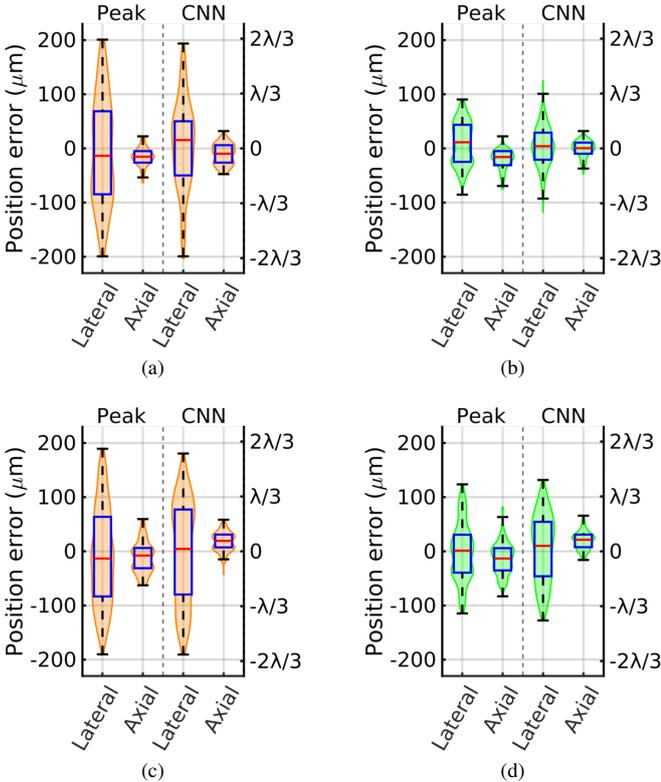


Fig. 15. Localization uncertainties of baseline and CNN methods on phantom measured data: (a) and (b) are results on the grid scatterer phantom with one and three plane waves, respectively. (c) and (d) are results on the random scatterer phantom with one and three plane waves, respectively.

V. DISCUSSION

A CNN-based scatterer detection and localization method is presented. Instead of end-to-end training, the CNNs were trained to learn the mapping from RF channel data to non-overlapping Gaussian confidence maps, and scatterers were detected and localized from the confidence maps by looking for local maxima. This two-step framework made it possible to handle varying numbers of scatterers (N_s). By obtaining non-overlapping Gaussian confidence maps from RF channel data without beamforming, it was able to identify high concentrations of scatterers which cannot be separated by conventional ultrasound imaging due to the overlaps of PSFs, as seen in Fig. 9. This method also has an advantage of fast processing time by exploiting GPU computation. The proposed CNN implicitly included beamforming since it is a mapping from the channel domain to the ultrasound image domain, which is a bottleneck of current ultrasound imaging. For the CNNs, processing time for a frame was 16 ms on average in a PC equipped with a NVIDIA Titan V GPU.

It was essential to use non-overlapping Gaussian confidence maps to make training work. The binary confidence maps were initially used to train CNNs with advanced loss functions such as weighted cross entropy [13], jaccard loss [54], or focal loss

[55], as well as simple loss functions such as MSE or mean absolute error, but all of them failed. The binary confidence maps were too sparse to be handled by simply manipulating the loss function. Non-overlapping Gaussian confidence maps, however, relaxed the sparsity of the binary confidence maps, while being able to recover scatterer positions by taking the maximum of overlapping Gaussians. Therefore, the greater gradients were provided during training and the CNNs were able to be guided to the correct solutions stably.

The delayed RF signal induced by a scatterer lies across all the channels and at several depths depending on the lateral location of the scatterer. Hence, it was necessary for CNNs to have large receptive fields so four *down* and four *up blocks* were used. It was tried to incorporate skip connections into the proposed CNN by, if necessary, applying upsampling to the feature maps in the contracting path to match the size of their corresponding feature maps in the expanding path. For image segmentation, the skip connections play an important role to recover the lost spatial information. The resulting reconstructed images have more fine details and, as a result, provide better localized semantic segmentation [13], [56]. However, the skip connections hindered training and the CNNs learned zero confidence maps. We presume that the feature maps extracted from RF channel data in the contracting path are not directly related to the reconstruction of confidence maps, unlike image segmentation. CoordConv [44] was applied instead to cope with the spatial information loss. The CNNs with the CoordConv reconstructed non-overlapping Gaussians more precisely and achieved the better recall and localization precision on the validation sets.

The training was firstly performed in the training set at the scatterer density of 2.44 mm^{-2} , and further fine-tuned on the whole training sets. Interestingly, the CNNs trained at the scatterer density of 2.44 mm^{-2} were already well generalized at the higher scatterer densities than 2.44 mm^{-2} . On the other hand, the CNNs achieved poor precision at the lower scatterer density and localization as two peaks appeared laterally near a true scatterer position in the confidence maps. The training sets, therefore, had more frames at the lower scatterer densities. It was also investigated to train CNNs using the whole training sets from the beginning of the training but the proposed way was more efficient; CNNs converged to the solutions with less iterations.

The comparison in the simulated test sets clearly shows that the proposed method outperforms the baseline method. The performance drop was much more severe for the baseline method at the higher scatterer densities. Deep-ULM is another CNN-based method which localizes high-density targets from B-mode images that contain overlapping PSFs. To compare the proposed method with Deep-ULM, the recall and localization errors were re-calculated following the way as the supplementary Fig. 1 in [24] was generated. The threshold values for determining positive detection was $\lambda/7$ and Euclidean distances between the true and estimated targets were calculated. The results are presented in Fig. 16. Both methods showed good performance at high target densities but the proposed method achieved slightly better recall and localization precision. Deep-ULM recovered roughly 1.80 mm^{-2} while the

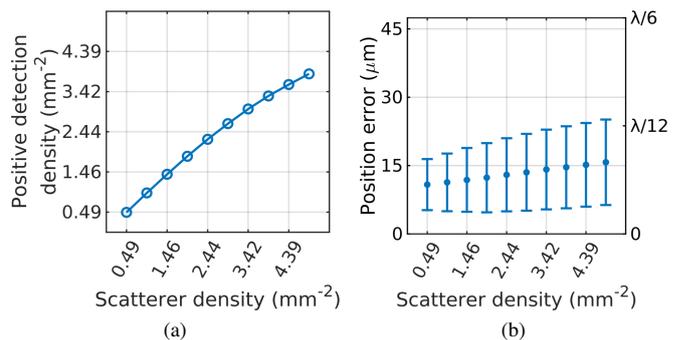


Fig. 16. Recall and localization precision re-calculated to compare CNN method to Deep-ULM: (a) Positive detection density and (b) median of Euclidean position errors with one standard deviation bars at different scatterer densities.

proposed method recovered 2.26 mm^{-2} at the target density of 2.44 mm^{-2} , and Deep-ULM recovered roughly 2.10 mm^{-2} at the target density of 3.53 mm^{-2} when the proposed method recovered 3.00 mm^{-2} and at the target density of 3.42 mm^{-2} . The median of Euclidean errors of Deep-ULM was approximately $\lambda/12$ but the proposed method achieved smaller errors than that. It is difficult to conclude that the proposed method is better than Deep-ULM since the evaluation was not performed in the same test data. This, however, shows the potential of the methods directly employing RF channel data.

To validate the proposed method for real world applications, two 3-D printed phantom were measured. The benefit of using the 3-D printed phantoms is that true scatterer positions and the dimensions of the phantom and scatterers (cavities) are known. It was important to modify the scatterers in the training sets to match the cavity dimensions. The CNNs trained for the simulation experiment failed showing too many wrong detections when they were tested on the measured data. The CNNs trained with the modified training set, however, achieved results comparable with the DAS method. It is notable that this was achieved only with the simulated training set, since it is extremely difficult to obtain sufficient training data with ground truth for these kinds of experiments.

The phantom experiment shows that the CNN method is transferable to measured data by modelling scatterers properly in the training data simulation. The baseline method performed slightly better for the trivial case: the grid scatterer phantom with three plane waves but the CNN method performed better for the other cases. The CNN method on the random scatterer phantom yet presented relatively large number of false positives compared to the simulation results. This could be because of factors not considered in the simulation such as attenuation in the phantom medium, different scattering intensity of the cavities, etc. The more accurate simulation and better generalized CNN model would be able to increase the CNN method performance on the measure data.

The proposed method gives 2-D images using a 1-D transducer. This limits the view of the target structure along the elevation direction. The 3-D printed phantoms were essentially 2-D phantoms which have elongated cavities and the dimension along the elevation direction was not captured in the

results. This limitation can be solved by using 2-D transducers such as fully addressed transducers or row-column addressed transducers.

There are several expected problems in order to apply the CNN method to MB detection for SRI. MBs are not static but move with different velocities depending on the vessel size. This should be considered during training data generation. The signals from tissue also need to be rejected. Clutter filtering is normally used to remove these, but this will change the signals from the scatterers too and the CNNs method might fail. Therefore, a way of rejecting the tissue signals without hurting the performance of CNNs needs to be investigated. Also, it is important to model MBs properly in simulations since their sizes are different. It was necessary to remodel scatterers following the real physical structure for phantom experiment. This is expected to be an important factor for applying the CNN method on the measured MB signals.

Lastly, further research on the optimal imaging scheme and scalability of CNN is required. Plane waves were used to support the hypothesis in a small region. In practice, however, a larger field of view is needed. Also, the more correlated data are available, the better estimation can be achieved. The CNNs with three plane waves achieved better performance than the CNN with one plane wave in all evaluation criteria, but this increases the required GPU memory. In addition, the imaging scheme would affect the capability of the CNN method and plane waves might not be the optimal choice. It is necessary to examine how other imaging schemes, such as focused or defocused waves affect the CNN method, or a new imaging scheme could be developed.

VI. CONCLUSION

The CNN-based scatterer detection and localization method is presented. CNNs were trained to return non-overlapping Gaussian confidence maps from simulated RF channel data, and the scatterer positions were estimated from the confidence maps. The simulation results show that the proposed method can identify high-density scatterers successfully even when some of them are closer than the resolution limit of conventional ultrasound imaging. It is also shown that the CNN method can be transferred to real measured data by modeling scatterers following the true scatterer structure. The CNN method can potentially be extended to replace DAS beamforming for high concentration MB detection and thus reduce the long data acquisition times of SRI using ULM.

ACKNOWLEDGMENT

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan V GPU used for this research.

REFERENCES

- [1] F. L. Thurstone and O. T. von Ramm, "A new ultrasound imaging technique employing two-dimensional electronic beam steering," in *Acoustical Holography*, P. S. Green, Ed., vol. 5. New York: Plenum Press, 1974, pp. 249–259.
- [2] O. Couture, B. Besson, G. Montaldo, M. Fink, and M. Tanter, "Microbubble ultrasound super-localization imaging (MUSLI)," in *Proc. IEEE Ultrason. Symp.*, 2011, pp. 1285–1287.
- [3] O. M. Viessmann, R. J. Eckersley, K. Christensen-Jeffries, M. X. Tang, and C. Dunsby, "Acoustic super-resolution with ultrasound and microbubbles," *Phys. Med. Biol.*, vol. 58, pp. 6447–6458, 2013.
- [4] M. A. O'Reilly and K. Hynynen, "A super-resolution ultrasound method for brain vascular mapping," *Med. Phys.*, vol. 40, no. 11, pp. 110701–7, 2013.
- [5] C. Errico, J. Pierre, S. Pezet, Y. Desailly, Z. Lenkei, O. Couture *et al.*, "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, pp. 499–502, November 2015.
- [6] K. Christensen-Jeffries, R. J. Browning, M. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433–440, February 2015.
- [7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [9] —, "Identity mappings in deep residual networks," in *Eur. Conf. Computer Vision*, 2016, pp. 630–645.
- [10] G. Huang, Z. Liu, L. v. d. Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2017, pp. 2261–2269.
- [11] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *IEEE Int. Conf. Computer Vision*, 2017, pp. 2980–2988.
- [12] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," [arXiv:1804.02767v1 \[cs.CV\]](https://arxiv.org/abs/1804.02767v1), 2018.
- [13] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [14] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2017, pp. 2881–2890.
- [15] L. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Eur. Conf. Computer Vision*, 2018, pp. 801–818.
- [16] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2017, pp. 105–114.
- [17] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2017, pp. 1132–1140.
- [18] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino *et al.*, "Imaging intracellular fluorescent proteins at nanometer resolution," *Science*, vol. 313, no. 5793, pp. 1642–1645, 2006.
- [19] S. T. Hess, T. P. K. Girirajan, and M. D. Mason, "Ultra-high resolution imaging by fluorescence photoactivation localization microscopy," *Biophysical Journal*, vol. 91, no. 11, pp. 4258–4272, 2006.
- [20] M. J. Rust, M. Bates, and X. Zhuang, "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)," *Nature methods*, vol. 3, no. 10, pp. 793–795, 2006.
- [21] E. Nehme, L. E. Weiss, T. Michaeli, and Y. Shechtman, "Deep-STORM: super-resolution single-molecule microscopy by deep learning," *Optica*, vol. 5, no. 4, pp. 458–464, Apr 2018.
- [22] W. Ouyang, A. Aristov, M. Lelek, X. Hao, and C. Zimmer, "Deep learning massively accelerates super-resolution localization microscopy," *Nature biotechnology*, 2018.
- [23] N. Boyd, E. Jonas, H. Babcock, and B. Recht, "Deeploco: Fast 3d localization microscopy using neural networks," [bioRxiv 267096](https://arxiv.org/abs/267096), 2018.
- [24] R. J. G. van Sloun, O. Solomon, M. Bruce, Z. Z. Khaing, H. Wijkstra, Y. C. Eldar *et al.*, "Super-resolution ultrasound localization microscopy through deep learning," [arXiv:1804.07661v2 \[eess.SP\]](https://arxiv.org/abs/1804.07661v2), 2018.
- [25] D. Allman, A. Reiter, and M. A. L. Bell, "Photoacoustic source detection and reflection artifact removal enabled by deep learning," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1464–1477, 2018.
- [26] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [27] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Int. Conf. Learning Representations*, 2015.

- [28] A. C. Luchies and B. C. Byram, "Deep neural networks for ultrasound beamforming," *IEEE Trans. Med. Imag.*, vol. 37, no. 9, pp. 2010–2021, 2018.
- [29] D. Hyun, L. L. Brickson, K. T. Looby, and J. J. Dahl, "Beamforming and speckle reduction using neural networks," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 66, no. 3, pp. 898–910, 2019.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair *et al.*, "Generative adversarial nets," in *Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [31] S. Goudarzi, A. Asif, and H. Rivaz, "Multi-focus ultrasound imaging using generative adversarial networks," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2019, pp. 1118–1121.
- [32] X. Zhang, J. Li, Q. He, H. Zhang, and J. Luo, "High-quality reconstruction of plane-wave imaging using generative adversarial network," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [33] J. Youn, M. L. Ommen, M. B. Stuart, E. V. Thomsen, N. B. Larsen, and J. A. Jensen, "Ultrasound multiple point target detection and localization using deep learning," in *Proc. IEEE Ultrason. Symp.*, 2019.
- [34] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 39, no. 2, pp. 262–267, 1992.
- [35] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comp.*, vol. 10th Nordic-Baltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1, pp. 351–353, 1996.
- [36] —, "A multi-threaded version of Field II," in *Proc. IEEE Ultrason. Symp.* IEEE, 2014, pp. 2229–2232.
- [37] B. G. Tomov, S. E. Diederichsen, E. V. Thomsen, and J. A. Jensen, "Characterization of medical ultrasound transducers," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [38] J. A. Jensen, "Safety assessment of advanced imaging sequences, II: Simulations," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 63, no. 1, pp. 120–127, 2016.
- [39] A. Gomariz, W. Li, E. Ozkan, C. Tanner, and O. Goksel, "Siamese networks with location prior for landmark tracking in liver ultrasound sequences," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2019, pp. 1757–1760.
- [40] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop *et al.*, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *IEEE Conf. Computer Vision and Pattern Recognition*, 2016, pp. 1874–1883.
- [41] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Int. Conf. Machine Learning*, 2015, pp. 448–456.
- [42] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, pp. 1929–1958, 2014.
- [43] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *ICML Workshop on Deep Learning for Audio, Speech, and Language Processing*, 2013.
- [44] R. Liu, J. Lehman, P. Molino, F. P. Such, E. Frank, A. Sergeev *et al.*, "An intriguing failing of convolutional neural networks and the coordconv solution," in *Neural Information Processing Systems*, 2018, pp. 9605–9616.
- [45] M. L. Ommen, M. Schou, R. Zhang, C. A. V. Hoyos, J. A. Jensen, N. B. Larsen *et al.*, "3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging," in *Proc. IEEE Ultrason. Symp.*, 2018, pp. 1–4.
- [46] M. L. Ommen, M. Schou, C. Beers, J. A. Jensen, N. B. Larsen, and E. V. Thomsen, "3d printed calibration micro-phantoms for validation of super-resolution ultrasound imaging," in *Proc. IEEE Ultrason. Symp.*, 2019, pp. 1–4.
- [47] G. Montaldo, M. Tanter, J. Bercoff, N. Benech, and M. Fink, "Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 56, no. 3, pp. 489–506, March 2009.
- [48] A. M. Saxe, J. L. McClelland, and S. Ganguli, "Exact solutions to the nonlinear dynamics of learning in deep linear neural networks," [arXiv:1312.6120v3 \[cs.NE\]](https://arxiv.org/abs/1312.6120v3), 2013.
- [49] D. Kingma and L. Ba, "Adam: A method for stochastic optimization," [arXiv:1412.6980 \[cs.LG\]](https://arxiv.org/abs/1412.6980), 2015.
- [50] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2011, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [51] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten *et al.*, "SARUS: A synthetic aperture real-time ultrasound system," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 9, pp. 1838–1852, 2013.
- [52] C. Chang, S. Chatterjee, and P. R. Kube, "On an analysis of static occlusion in stereo vision," in *IEEE Conf. Computer Vision and Pattern Recognition*, 1991, pp. 722–723.
- [53] P. Fua, "A parallel stereo algorithm that produces dense depth maps and preserves image features," *Mach. Vis. Appl.*, vol. 6, no. 1, pp. 35–49, 1993.
- [54] P. Jaccard, "The distribution of the flora in the alpine zone," *New phytologist*, vol. 11, no. 2, pp. 37–50, 1912.
- [55] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *IEEE Int. Conf. Computer Vision*, 2017, pp. 2999–3007.
- [56] M. Drozdal, E. Vorontsov, G. Chartrand, S. Kadoury, and C. Pal, "The importance of skip connections in biomedical image segmentation," [arXiv:1608.04117v2 \[cs.CV\]](https://arxiv.org/abs/1608.04117v2), 2016.

APPENDIX C

Papers in preparation

C.1 Paper K - Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary?

Reduced Cavity Pressure in Fusion Bonded Devices: Is a Wafer Bonder Necessary?

Martin Lind Ommen and Erik Vilain Thomsen

Department of Micro and Nanotechnology, Technical University of Denmark, Kgs. Lyngby, Denmark

Abstract—Fusion bonding is a common wafer bonding technique, which can be used to form sealed cavities. The bonding process is often conducted in a wafer bonder, which can provide alignment bonding of two substrates as well as control of the bonding conditions in terms of force on the wafer stack, wafer temperature, and the atmospheric composition. A reduced cavity pressure is often a desired or integral property in applications such as absolute pressure sensors and in capacitive micro-machined ultrasonic transducers, and is commonly obtained by bonding inside a vacuum in a wafer bonder. However, as demonstrated in this paper, a wafer bonder is not necessary to obtain a reduced cavity pressure. Fusion bonding of silicon and silicon nitride plates to substrates containing cavities formed in silicon dioxide in four different atmospheres, all result in similar cavity pressures. Hence, the final cavity pressure is determined during the subsequent bond anneal, and not during the pre-bond, making it possible to obtain a reduced cavity pressure without using a wafer bonder. On the other hand, bonding in a vacuum does not ensure a vacuum cavity in the finished device.

Index Terms—Bonding, direct bonding, fusion bonding, wafer bonder

I. INTRODUCTION

Wafer bonding is a common processing technique for combining multiple wafers into a single structure. It can be used to stack structures which would otherwise not be possible to combine by epitaxy or film deposition. Examples of wafer bonding applications are in the production of silicon on insulator (SOI) wafers and in the formation of sealed cavities, such as in pressure sensors [1] or in capacitive micro-machined ultrasonic transducers (CMUTs) [2], [3].

The original fusion bonding method was first described in the literature in 1986 [4]. In this work, two hydrophilic, mirror-polished silicon wafers were fusion bonded by being brought in contact in a clean environment at room temperature. The wafer stacks were then heated to 1000 °C, which completed the bonding process. Since the initial introduction, the overall processing has not changed; it still consists of a cleaning and/or surface activation process; a pre-bond in which the wafers to be bonded are placed in contact, typically with an applied pressure on the wafer stack; and a subsequent high-temperature bond-anneal. However, sophisticated wafer bonders have been made, which are typically used during the pre-bond. Such pieces of equipment allow for more advanced bonding methods, including alignment bonding when structures on the top and bottom wafers should be aligned. They can also provide control of the bonding conditions, in terms of the bonding temperature, the pressure applied on the wafer stack during the pre-bond, and the pre-bond atmosphere. Additionally, it has been shown that it is possible to fusion bond other materials together, such as silicon dioxide (SiO₂) to silicon [5], silicon nitride (Si₃N₄) to Si₃N₄ [6] or in fact any combination of the three materials. The potential to obtain

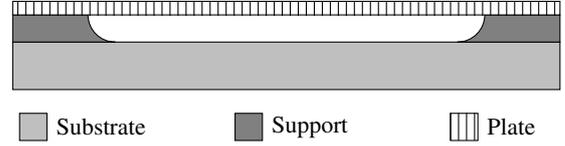


Figure 1. Cross section of a cavity device consisting of a substrate wafer, a support structure defining the cavities, and a plate which is bonded on top.

a reduced cavity pressure, by bonding in a vacuum, is in some cases used as an integral part of the functional device [1], [2], [3].

In this paper, it is shown that it is possible to obtain a reduced cavity pressure, without bonding in a vacuum. The results not only show that the bond anneal is the essential process in determining the final cavity pressure, they also indicate that even though the bonding is conducted in a vacuum chamber, the resulting cavity pressure is not a vacuum. Essentially, this means that unless alignment bonding or elevated pre-bond temperatures are required, a wafer bonder will not be necessary for obtaining reduced pressures in fusion bonded cavities.

II. MATERIALS AND METHODS

A. Experimental Design

Test structures were made to determine any differences in the cavity pressure after the bonding process. Simple wafer bonded cavity test structures enable indirect determination of the cavity pressure, by measurement of the deflection of a plate suspended over the cavities in an ambient environment. A cross section of such a device can be seen in Fig. 1. The center deflection of an isotropic circular plate, w_0 , can be expressed according to [7]

$$w_0 = \frac{3}{16} \frac{(1 - \nu^2)^2 a^4}{E h^3} \Delta p, \quad (1)$$

where ν is Poisson's ratio, Δp is the pressure difference across the plate, a is the radius of the plate, E is Young's modulus, and h is the plate thickness. Hence, any difference in plate deflection between the devices is proportional to the difference in cross-plate differential pressure, and thus, to the cavity pressure as well.

To test the effect of the pre-bond environment on the resulting cavity pressure, four bonding conditions were compared: three formed inside a wafer bonder and one formed directly by hand, referred to as hand-bonded. In the wafer bonder, the atmospheric environment was changed between 2×10^{-4} mbar (Vacuum), air, and argon. Assuming a perfect seal of the cavities, the three different atmospheres should result in different cavity pressures. For the devices bonded in a vacuum, the cavity pressure should be 0 bar,

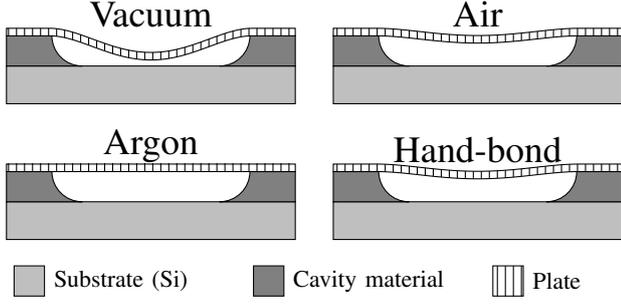


Figure 2. Schematic of the expected deflections for the four different bonding conditions.

and $\Delta p = 1$ bar when the ambient pressure is 1 bar. For the devices bonded in 1 bar of argon $\Delta p = 0$ bar, as the argon atmosphere is inert and should remain intact. For the devices bonded in air $\Delta p \approx 0.2$, since air is composed of 78% nitrogen, 21% oxygen and 1% argon, of which the 21% oxygen will be consumed in oxidation of any silicon surfaces of the cavities during the high-temperature bond-anneal. The oxygen consumption has previously been described in [8], [9]. Finally, for the devices bonded directly in hand $\Delta p \approx 0.2$ as the atmospheric environment is the same as that of the air devices. As the plate deflection is linear in pressure, these differences in Δp should correspond directly to the relative differences in plate deflections. The maximum deflection is expected for the Vacuum devices, whereas the Air and Hand-bond devices would only deflect one fifth of the vacuum devices, and the Argon devices should not deflect at all. These expectations are illustrated in Fig. 2.

B. Material Choices - Silicon Nitride Plates

There are a number of ways to fabricate fusion bonded cavities, as fusion bonding can be made with the combination of any two substrates with a surface of either silicon, SiO_2 , or Si_3N_4 . For ease of fabrication, the cavities are etched in a SiO_2 layer, which is grown on a silicon wafer. This provides control of the cavity depth when using a selective wet etchant, due to essentially an etch stop once the etchant reaches the silicon below the SiO_2 . The plate can then be fabricated using either silicon, SiO_2 , or Si_3N_4 . To obtain a device layer of a few μm as would be required for the chosen design when using silicon as the plate material, commercial SOI wafers are typically thinned by chemical mechanical polishing. This processing unfortunately results in a thickness variation between 300 nm and 500 nm, providing large variations in deflection as (1) scales with h^3 . Therefore, two alternative types of plates were considered. Firstly, an SiO_2 plate grown on a silicon wafer. SiO_2 would provide control of the layer thickness and uniformity, but also introduce built-in compressive stress when grown on a silicon substrate. This could result in buckling of the plate with a direct influence on the deflection measurements. Secondly, an Si_3N_4 film deposited on a silicon wafer. Low pressure chemical vapor deposition (LPCVD) of Si_3N_4 layers can, similarly to growth of SiO_2 films, provide good control of thickness and uniformity compared to the SOI wafers. There will be a built-in stress, but as it is tensile, it

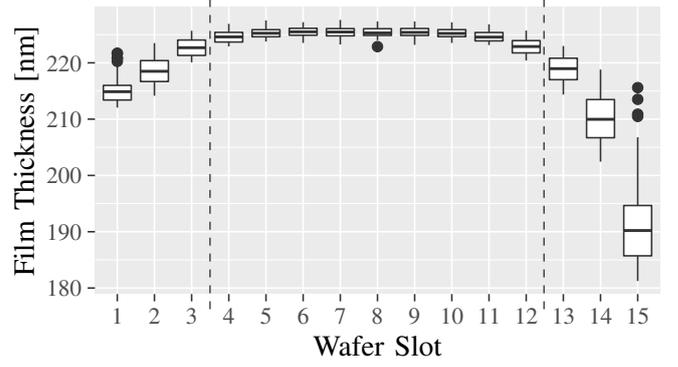


Figure 3. Box-plot of the film thickness distribution across a furnace boat of Si_3N_4 films. The film thickness is $225.0 \text{ nm} \pm 0.8 \text{ nm}$ between the dashed lines, and therefore suitable for comparative experimentation.

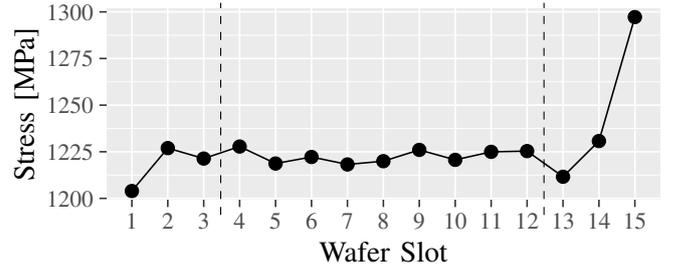


Figure 4. Stress distribution across a furnace boat of Si_3N_4 films. The dashed lines indicate the same region selected in Fig. 3. The stress within that region is $1223 \text{ MPa} \pm 4 \text{ MPa}$, and therefore suitable for comparative experimentation.

will not result in buckling, and is therefore acceptable.

To be able to compare the deflections of the different bonded structures illustrated in Fig. 2, the variability of thickness and the stress in the Si_3N_4 films is a critical parameter. Engholm et al. [10] showed that both thickness and stress will influence how much a plate deflects, describing the centre deflection, w_0 , of a plate with a built in tensile stress as

$$w_0 = \frac{\Delta p a}{2} \sqrt{\frac{C D}{N_t^3}} \left(\frac{1 - I_0\left(\sqrt{\frac{N_t}{C D}} a\right)}{I_1\left(\sqrt{\frac{N_t}{C D}} a\right)} \right) + \frac{\Delta p a^2}{4 N_t}, \quad (2)$$

where C is a constant based on the solution to the plate equation, D is the flexural rigidity of the plate, I_n is the modified Bessel function of first kind, and $N_t = \sigma h$ is the stress resultant, where σ is the planar biaxial stress in the plate. Since it is not possible to change the bonding atmosphere or method of pressure application locally on a single wafer, the comparison of the four bonding conditions will need to be between devices fabricated on separate wafers. Therefore, it is essential that the inter-wafer variability in thickness and stress of the plate is not so large, that it makes distinguishing between the expected differences in deflection across the different wafers impossible. Fig. 3 shows how the LPCVD Si_3N_4 film thickness varies across a full quartz boat of silicon wafers after a single batch process. The thickness has been measured using an ellipsometer. Each box in the box-plot represents the individual wafers in the quartz boat, and consists of 49 thickness measurements distributed across each wafer, thereby showing the intra-wafer variability. The wafers in the central region of the quartz boat, marked by the dashed lines, are uniform both in terms of the inter-wafer

and intra-wafer variability. The average thickness between the dashed lines is $225.0 \text{ nm} \pm 0.8 \text{ nm}$. Fig. 4 shows the stress variation across the same batch of wafers. The stress has been calculated from the wafer curvature. Within the central region, the average stress is $1223 \text{ MPa} \pm 4 \text{ MPa}$. In both cases, the uncertainty is the standard deviation.

Using (2), it is possible to determine the expected deflections of the four different types of devices, and more importantly, the smallest expected difference between the devices. The expected differential pressure across the plates for the four different bonding conditions were 1 bar (Vacuum), 0.2 bar (Air), 0 bar (Argon) and 0.2 bar (Hand-bond), which means that the smallest difference in differential pressure between any two devices would be 0.2 bar. Thus, for the devices with a cavity radius of $32 \mu\text{m}$, the smallest difference in deflection for the average values of plate thickness and plate stress can be calculated using (2) to be

$$\Delta w_{0,\text{min}} = 18.0 \text{ nm}.$$

The inter-wafer variability of the data in Fig. 3 and Fig. 4 can be used to estimate how large the expected variations in deflection are. Assuming the sources of variation are independent and random, the propagation of error in the centre plate deflections, δw_0 , can be estimated using (2) to calculate

$$\delta w_0 = \sqrt{\left(\frac{\partial w_0}{\partial h} \delta h\right)^2 + \left(\frac{\partial w_0}{\partial \sigma} \delta \sigma\right)^2}, \quad (3)$$

where δh is the uncertainty in the plate thickness, and $\delta \sigma$ is the uncertainty in the stress, and it is assumed that h and σ are the only varying parameters [11]. The uncertainty in the plate thickness and stress are estimated as the standard deviations of the central regions of Fig. 3 and Fig. 4. The calculated uncertainty in plate deflection will increase with the applied cross plate differential pressure. For the largest pressure difference ($\Delta p = 1 \text{ bar}$), the uncertainty in plate deflection will be

$$\delta w_0 = 0.4 \text{ nm}. \quad (4)$$

This difference from the processing uncertainties is much smaller than the expected difference between devices. Consequently, the expected deflection differences due to different cavity pressures should be distinguishable when choosing Si_3N_4 as the plate material.

C. Fabrication of Test Devices

An illustration of the process flow can be seen in Fig. 5. A $405 \text{ nm} \pm 0.5 \text{ nm}$ SiO_2 layer was grown on a batch of both single side polished four inch silicon (100) wafers and double side polished four inch silicon (100) wafers in a dry thermal oxidation process at $1100 \text{ }^\circ\text{C}$ (a). The single side polished wafers were used as substrates, in which the cavities were to be etched. The double side polished wafers were used as support substrates for the plate layers. The plate wafers were transferred directly to an LPCVD furnace for deposition of a $226 \text{ nm} \pm 0.8 \text{ nm}$ Si_3N_4 layer (b). After the deposition, the plate wafers were transferred to an oxidation furnace for oxidation of the Si_3N_4 layer, which has been shown to improve the bonding strength between SiO_2 and Si_3N_4 [6]. The plate wafers were left inside the furnace until

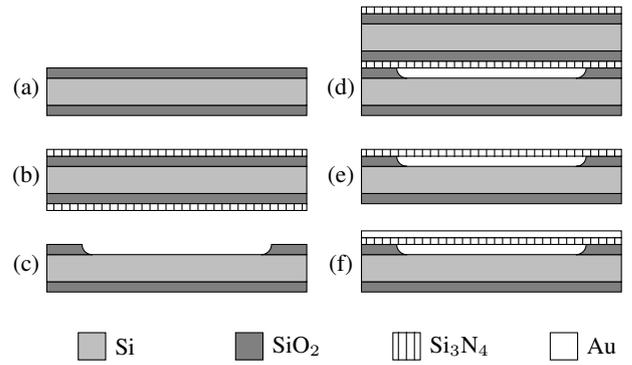


Figure 5. Process flow for the fabricated test devices.

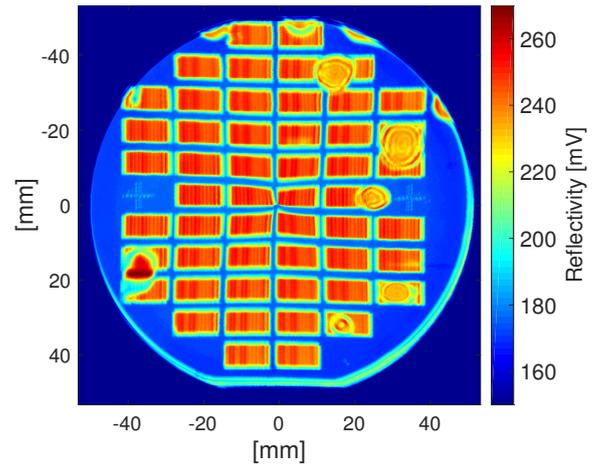


Figure 6. Infrared reflectance map of one of the device wafers bonded directly in hand. Most of the interface is void-free.

needed for bonding to minimize particle contamination. The cavities on the substrate wafers were defined in a lithographic process with a radius of $a = 32 \mu\text{m}$. They were then etched in a wet BHF etch to define the 405 nm deep cavities (c). The substrate wafers were RCA cleaned to remove particles, directly after which the substrate and plate wafers were fusion bonded together, under the four different bonding conditions (d). A Süss SB6 wafer bonder (Garching, Germany) was used to bond the non-hand-bond devices. The Vacuum devices were bonded at a pressure of $2 \times 10^{-4} \text{ mbar}$, the Air devices were bonded without pumping down the chamber, and the Argon devices were bonded in an argon atmosphere at a pressure of 1 bar. All of the devices bonded in the wafer bonder had a 600 mbar pressure applied on the wafer stack during the pre-bond. All bonds were made at room temperature. After the pre-bond, all bonded structures were annealed at $1100 \text{ }^\circ\text{C}$ in 1 bar of N_2 for 3 hours. The bonding interfaces were then characterized by infrared reflectance measurements using the infrared photoluminescence system Accent RPM2000 Compound Semiconductor Photoluminescence System to check for voids. An infrared reflectance map for one of the hand-bonded wafers can be seen in Fig. 6. A few voids are visible, but most of the interface has been properly bonded. The handle layer of the top wafer was etched away using a sequential combination of dry etching and wet KOH etching to release the nitride plates (e). Finally, a layer of gold was sputtered on top of the wafer to increase the reflectivity of the surface for the subsequent analysis (f).

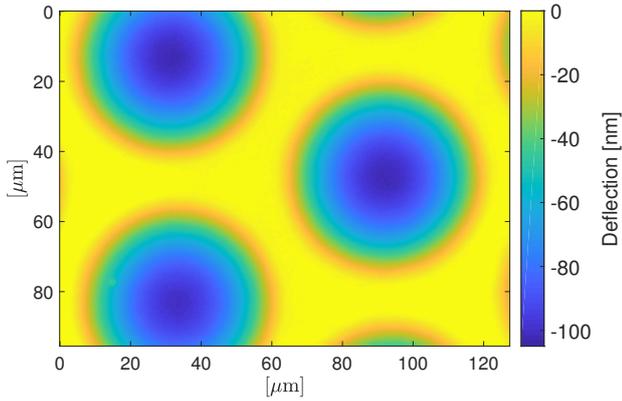


Figure 7. Optical profile of one of the Vacuum devices. The blue regions are deflecting plates.

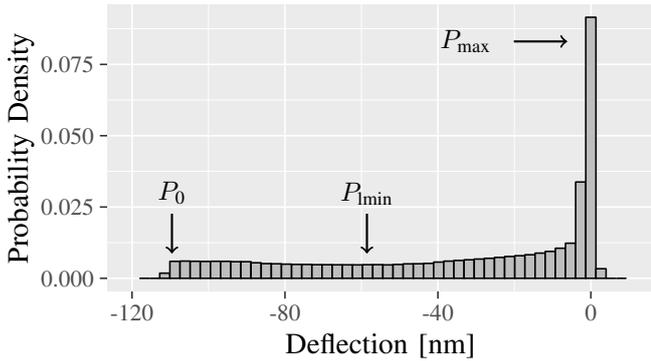


Figure 8. Histogram of the data shown in Fig. 7. The distinct behaviour, with points of interest marked, allows for automatic detection of the maximum deflection.

III. RESULTS AND DISCUSSION

A. Deflection Measurements

Eight different plates on each wafer were measured using the Sensorfar PLu Neox Optical Profiler (Sensorfar, Terrassa, Barcelona) to determine their deflections. An example of a such an optical height distribution is shown in Fig. 7 for a Vacuum devices. The blue regions are the deflecting plates. The corresponding histogram can be seen in Fig. 8. These histograms are very distinct, and can be used to determine a systematic estimate of the deflection of each measurement. The highest probability density, P_{\max} , correlates with the area between the cavities, the yellow region in Fig. 7. This peak in the histogram can be used to offset the data to align all measurements to the same reference point. For decreasing values (larger deflections) the density initially decreases rapidly, reaching a local minimum, P_{\min} , between -40 nm and -80 nm in the case of Fig. 8 before increasing again slightly and finally dropping to zero. This non-monotonic behaviour means it is not possible to set a lower density threshold and use that to find the maximum deflection value, as it could result in the deflection value corresponding to P_{\min} . Also, choosing the lowest probability value increases the susceptibility to data outliers. By locating the first bin in the histogram with a value larger than the tenth quantile of the density data, and choosing the deflection corresponding to this as the deflection, it is possible to systematically determine the deflection of the plates near P_0 , while avoiding the risks listed previously. Fig. 9 shows a comparison of the height distributions of the four different types of bond conditions. All plates of

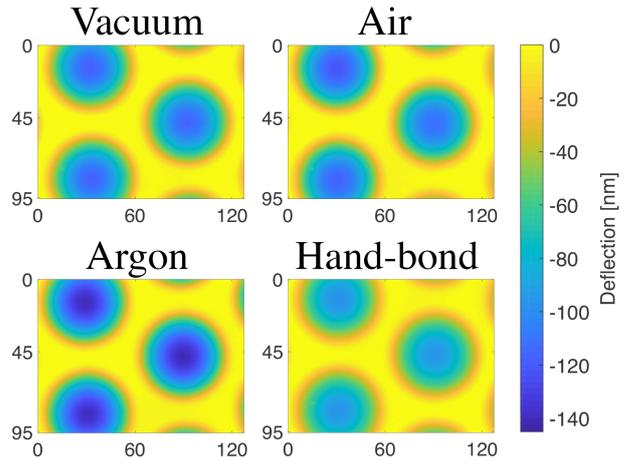


Figure 9. Comparisons of typical optical profiles of the four different bond conditions. The lateral dimensions are all in μm .

the test structures deflect significantly and almost the same amount. The deflection data can be seen in Fig. 10. It should be noted that the magnitude of deflections of all devices is large, regardless of the bonding conditions. This is a remarkable result, as the Argon devices were not expected to deflect at all. However, it seems that there are two groups, namely the devices bonded in the wafer bonder which all deflect ≈ 110 nm, and the hand-bonded devices which all deflect ≈ 60 nm. The Air devices and Hand-bond devices are directly comparable in terms of the pre-bond atmosphere, but the Air devices which were bonded in a wafer bonder deflect significantly more. These measurements indicate that whether the devices were pre-bonded in a vacuum, in air, or in argon, is not critical for the final cavity pressure, and that the cavity pressures end up being similar regardless, but whether the pre-bond is done in a wafer bonder or in hand will have an effect. Although the deflection of the Hand-bond devices is lower than that of the other devices, they still deflect more than the expected fifth of the Vacuum devices. Finally, the magnitude of the intra-wafer variation on most of the devices as well as the inter-wafer variation, does not correlate with the uncertainty estimate presented in Sec. II-B, which must mean that there is a source of variability not accounted for. However, it should not have an effect on the Argon devices, since there should be no pressure difference across the plate according to the hypothesis.

B. Hypothesis for Bond Interface Diffusion

The expectations for the experiments presented in Sec. II-A were based on the cavities being sealed during the pre-bond. However, if the bond-interfaces are not leak tight after the pre-bond, a gas exchange between the cavities and the external environment can occur. During the bond-anneal, the temperature is increased to 1100 $^{\circ}\text{C}$ in 1 bar of nitrogen. According to the ideal gas law, the high temperature will increase the pressure inside of the cavities by a factor of about 4.5. As illustrated in Fig. 11, any pressure gradient will be able to drive gas diffusion between the cavities and the external environment, potentially equilibrating the pressures. For the Vacuum devices, the pressure inside the cavities is initially 0 bar, while the external pressure in the furnace is 1 bar. Therefore, gas will diffuse into the cavities during annealing.

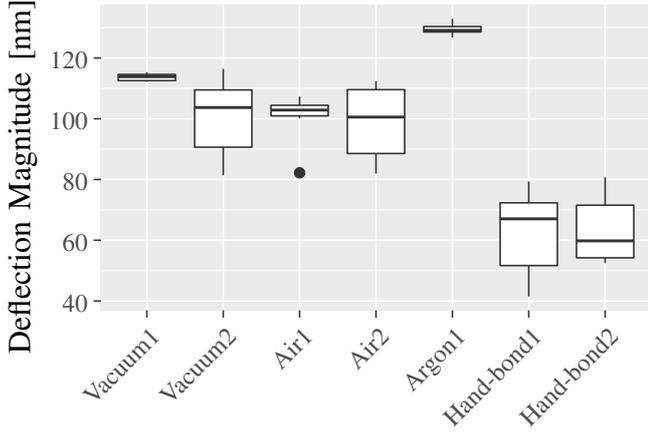


Figure 10. Box-plot showing the deflection measurements of each of the devices. Each box consists of measurements of eight different cavities distributed across a wafer.

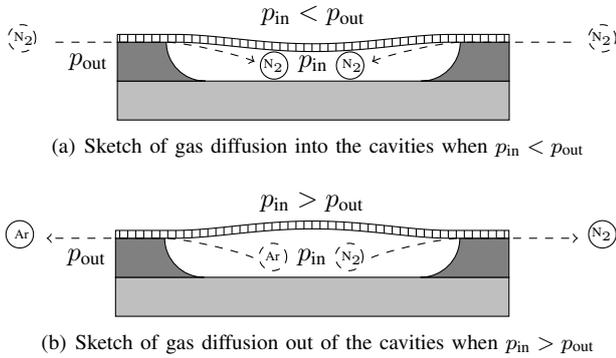


Figure 11. Sketches of gas diffusion during the annealing process, determined by the direction of the pressure gradient between the cavities and the ambient environment. (a) represents the Vacuum devices, where the pressure initially is larger outside of the cavities. (b) represents the Air, Argon and Hand-bond devices, where the pressure initially is larger inside the cavities. Note that the deflecting plates are used to illustrate the pressure variation. During the annealing process, the handle layer of the top wafer is still attached, meaning practically no deflection will occur.

For the Air, Argon and Hand-bond devices, the pressure inside the cavities will initially be around 4 bar, and gas will diffuse out of the cavities during the anneal. Once the bond-anneal is finished, and the cavities are sealed, the cavity pressure will be reduced by the same factor of about 4.5 when the temperature is returned to room temperature. This explains the large deflection of the Argon devices, when the hypothesis of a sealed cavity predicts no deflection at all. It also explains why most wafers deflect the same amount, due to the pressure inside the cavities having been equilibrated to the same value. Both the intra-wafer and inter-wafer deflection variations which are seen could potentially be explained by the sealing of the cavities being obtained at different points in time for the different cavities. If the gas diffusion has not managed to equilibrate the pressure, a sealing of the cavity would result in an off-equilibrium static pressure in the cavities. This could suggest that the pre-bond of the Hand-bond devices was stronger than that of the other devices.

C. Deflection Test With a Silicon Plate

In order to test whether the out-diffusion of gas was a unique effect of using Si_3N_4 as the plate material, another set of wafers were fabricated, using SOI wafers as the top plate.

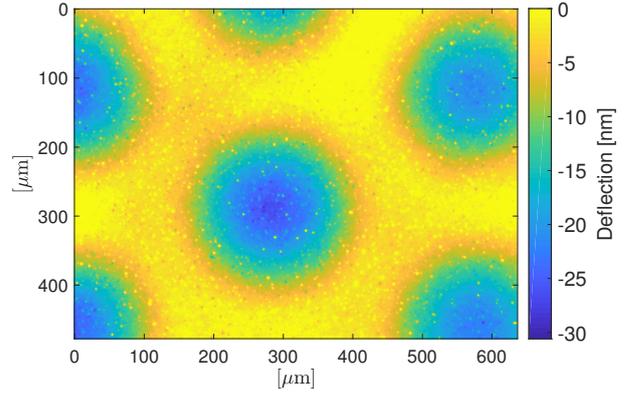


Figure 12. Deflection measurement of a hand-bonded cavity device fabricated with a silicon top plate.

Table I

DEFLECTION MEASUREMENTS BEFORE AND AFTER 7 HOURS AT 2 BAR HELIUM. THE BEFORE VALUE HAS BEEN MEASURED FIVE MONTHS AFTER THE BONDING PROCESS. THE VACUUM2 WAFER WAS BROKEN IN HALF.

	Deflection after five months [nm]	Deflection after helium exposure [nm]	Deflection difference [nm]
Vacuum1	-113.2	-115.3	-0.8
Vacuum2	-134.4	+10.1	144.5
Air1	-105.4	-107.3	-1.9
Air2	-102.4	-103.5	-1.1
Argon1	-129.4	-128.4	1.0
Hand-bond1	-42.4	-43.2	-1.9
Hand-bond2	-45.6	-48.6	-3.0

Apart from the change of plate material and the metal being used as an etch mask, the process flow was the same as described in Sec. II-C. The thickness of the SOI device layer was $3 \mu\text{m} \pm 0.5 \mu\text{m}$, as specified by the manufacturer. A height distribution of the surface after fabrication of a silicon plate device which was hand-bonded can be seen in Fig. 12. The grainy appearance is not noise in the data, but roughness of the surface due to the etching of the metal.

By using (1), the centre deflection is expected to be $w_0 \approx 6 \text{ nm}$ when $\Delta p = 0.2 \text{ bar}$. To obtain the deflection of $w_0 \approx 30 \text{ nm}$ seen in the figure, it would be required that $\Delta p \approx 1 \text{ bar}$. This is similar to the Si_3N_4 devices, and shows that it is plausible that the same out-diffusion is obtained when bonding silicon to SiO_2 , as when bonding Si_3N_4 to SiO_2 .

D. Bond Interface Leak Rate Test

Deflection measurements do not reveal differences in leak rates of the bonding interfaces for the four different types of devices. In particular, such measurements do not indicate whether there is a difference in leak rate between the devices bonded in a wafer bonder and those bonded directly in hand. They only document a static situation in a very narrow period of time. A simple method for investigating the leak rates is to measure the deflection after long periods of time. Table I shows a collection of measurements conducted at different points in time. The second column consists of deflection measurements conducted five months after the device fabrication. Considering the large variation in deflections observed for each device type in Fig. 10, it is imperative that the exact same cavities are compared. Unfortunately, the deflection measurements conducted directly after fabrication were not

logged with exact cell placement, leaving direct comparison to the deflections immediately after processing impossible. However, the values after 5 months align well with those presented in Fig. 10, suggesting that the leak rate must be small. Another method of assessing the leak rate would be by helium leak testing. Helium is commonly used for testing of leak rates due to its small atomic size enabling faster propagation through porous structures, than would be the case for e.g. the molecules of atmospheric air.

Each device was exposed to 2 bar of helium for 7 hours, and the deflections were measured within 30 minutes. Table I also lists the comparative measurements after the helium test, as well as the difference to the measurement before. Positive differences in deflection means the cavity pressure has decreased, and that some amount of helium has reached the device cavities, as would be the expectation. A number of interesting points can be made from these measurements. Firstly, that most differences are small, suggesting a low permeability of the bonding interface. Secondly, that the Vacuum2 device behaves significantly different than the others. The measurements are highlighted in red font in the table. This particular wafer had broken in half during the latter half of processing, certainly setting it apart from the others, and likely influencing the bonding interface. However, the wafer had a large deflection after five months, indicating very little leak rate at that point. Only the exposure to helium changed the deflection significantly, demonstrating that a difference can actually be detected simply by changing the gas to helium. In addition, cavities next to the one investigated deflected similarly to the measurement conducted five months after bonding, showing that this is a very local effect, simply due to the breakage. Thirdly, most values are negative, corresponding to fewer gas particles being in the cavities after helium bombardment, which does not make physical sense. More likely, small changes in the exact region which is investigated before and after helium testing, influence the deflection estimate more than the leakage flow of helium. Thus, the listed differences in pressure are indicators of the uncertainty of each individual measurement. As a consequence, it does not make sense to calculate individual leak rates for the devices. However, by estimating that the true change in deflection is no larger than the measurement uncertainty, estimated to be the largest deflection difference in Table I namely 3 nm, it is possible to calculate an upper limit of the leak rate. The simplest estimate of the leak rate, will be calculated as

$$L = \frac{dp}{dt}V, \quad (5)$$

where V is the cavity volume. Using (2), a 3 nm deflection difference is found to correspond to a 34 mbar pressure difference. Combining this with the geometrical parameters of the cavity, namely a radius of $a = 32 \mu\text{m}$ and a height of $h = 405 \text{ nm}$, the upper estimate for the leak rate is becomes

$$L \leq 1.76 \times 10^{-15} \frac{\text{mbar L}}{\text{s}}. \quad (6)$$

This order of magnitude agrees with the literature [12], albeit having been determined through a different method.

IV. CONCLUSION

A study of the resulting cavity pressure after fusion bonding has been presented. Fabrication of cavity test structures

consisting of etched SiO_2 cavities, with Si_3N_4 plates fusion bonded on top, enabled measurements of the deflection of the nitride plates as an indirect measure of the cavity pressure. Four sets of bonding conditions were used, three in a wafer bonder in atmospheres of vacuum, air and argon, and the last set was bonded directly in hand in atmospheric conditions. Qualitative arguments and observations of the plate deflections over time and after helium testing revealed a maximum leak rate of the fusion bonded structures of $1.76 \times 10^{-15} \text{ mbar L s}^{-1}$. Comparison of the test devices revealed similar deflections for all devices. The same phenomenon was observed for devices fabricated with silicon plates, showing that this is not an isolated feature of using Si_3N_4 plates. This has led to the conclusion, that the initial pre-bond of the wafers did not provide a leak-tight bond-interface. Instead, gas is able to diffuse from and to the cavities during the subsequent bond-anneal, until reaching an equilibrium pressure between the cavities and the surrounding atmosphere, prior to the cavities being sealed. The cavities reaching an equilibrium pressure explains why the different bonding conditions resulted in similar plate deflections, and reveals that the bonding conditions do not influence the final cavity pressure, and even bonding in vacuum does not ensure a vacuum cavity. Thus, whether the pre-bond is made in a wafer bonder or directly in hand does not matter, and therefore, it is not necessary to use a wafer bonder in order to obtain a reduced cavity pressure.

REFERENCES

- [1] G. Fragiaco, T. Pedersen, O. Hansen, and E. V. Thomsen. Intrinsic low hysteresis touch mode capacitive pressure sensor. *IEEE Sensors J.*, pages 2279–2282, 2011.
- [2] Y. L. Huang, A. S. Ergun, E. Haeggstrom, M. H. Badi, and B. T. Khuri-Yakub. Fabricating capacitive micromachined ultrasonic transducers with wafer-bonding technology. *J. Microelectromech. Syst.*, 12(2):128–137, 2003.
- [3] S. E. Diederichsen, J. M. F. Hansen, M. Engholm, J. A. Jensen, and E. V. Thomsen. Output pressure and pulse-echo characteristics of cmuts as function of plate dimensions. *2017 IEEE International Ultrasonics Symposium (IUS)*, page 8092352, 2017.
- [4] M. Shimbo, K. Furukawa, K. Fukuda, and K. Tanzawa. Silicon-to-silicon direct bonding method. *J. Appl. Phys.*, 60(8):2987–2989, 1986.
- [5] A. Berthold, B. Jakoby, and M. J. Vellekoop. Wafer-to-wafer fusion bonding of oxidized silicon to silicon at low temperatures. *Sens. Actuators, A: Physical*, 68(1-3):410–413, 1998.
- [6] K. Reck, C. Østergaard, E. V. Thomsen, and O. Hansen. Fusion bonding of silicon nitride surfaces. *J. Micromechanics Microengineering*, 21(12):125015, 2011.
- [7] S. Timoshenko. *Theory of Plates and Shells*. McGraw-Hill Book Company, New York City, USA, 1940.
- [8] J. B. Lasky. Wafer bonding for silicon-on-insulator technologies. *Applied Physics Letters*, 48(1):78–80, 1986.
- [9] C. Harendt, B. Hofflinger, H. G. Graf, and E. Penteker. Silicon direct bonding for sensor applications - characterization of the bond quality. *Sens. Actuators, A: Physical*, 25(1-3):87–92, 1991.
- [10] M. Engholm, T. Pedersen, and E. V. Thomsen. Modeling of plates with multiple anisotropic layers and residual stress. *Sens. Actuators, A: Physical*, 240:70–79, 2016.
- [11] J. R. Taylor. *An introduction to error analysis*. University Science Books., Sausalito, CA 94965, USA, 1997.
- [12] K. Schjolberg-Henriksen, N. Malik, A. Sandvand, G. Kittilsland, and S. T. Moe. Leak rates and residual gas pressure in cavities sealed by metal thermo-compression bonding and silicon direct bonding. *ECS Trans.*, 64(5):305–314, 2014.

APPENDIX D

Presented posters

D.1 Poster 1 - 3D printed flow phantoms with fiducial markers for super-resolution ultrasound imaging



3D PRINTED FLOW PHANTOMS WITH FIDUCIAL MARKERS FOR SUPER-RESOLUTION ULTRASOUND IMAGING



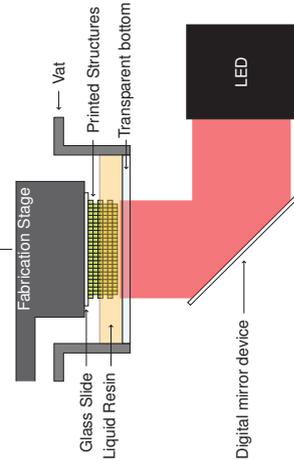
Martin Lind Ommen¹, Rujing Zhang¹, Mikkel Schou², Carlos Armando Villagómez Hoyos², Jørgen Arendt Jensen², Niels Bent Larsen¹ and Erik Vilain Thomsen¹

¹Department of Micro- and Nanotechnology, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark
²Center for Fast Ultrasound Imaging, Technical University of Denmark, Kgs. Lyngby, Denmark

Objective

- Make geometrically stable ultrasound SRI tissue-like phantoms by 3D printing using stereolithography
- Manufacture phantoms within 5% tolerance from the design
- Print perfusable flow channels on relevant size scale
- Print fiducial markers for improved alignment

Stereo lithography

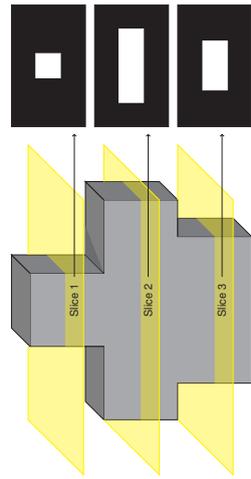


Sketch of the stereolithography setup. Light from an LED illuminates a digital mirror device, which reflects the light in the desired pattern through the transparent primer vat bottom. The illuminated resin in the vat will start crosslinking. The initial layer of the printed structures is crosslinked to a glass slide mounted on the movable fabrication stage.

Voxel size - $10 \times 10 \times 20 \mu\text{m}^3$ - > 2.000.000 voxels per layer
The specific components are listed in the manuscript

3D Model Preparation

3D model is converted to slices and build list by the open source software slic3r (www.slic3r.org)



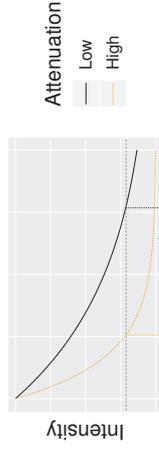
A 3D model, on the left, is decomposed into slices, which are exported as .png files. The images on the right are examples of slices corresponding to the yellow planes.

Increased precision can be obtained by manual .png construction

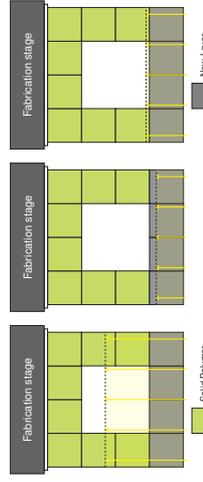
Liquid Resin

Three components:

- Liquid pre-polymer (PEGDA 700 g/mol)
- Photoinitiator (LAP 5 mg/mL)
- Photoabsorber (QY 12 mg/mL)



Light intensity against depth for two (arbitrary) different attenuation coefficients. d_1 and d_2 mark the depths, at which the light dose (intensity times exposure duration) has decreased to the threshold of resin solidification, marked by the grey horizontal dotted line.



(a) Too little attenuation (b) Too much attenuation (c) Ideal attenuation
 Sketches of the crosslinking depth for different levels of attenuation. The dashed lines mark the depth of the threshold dose. (a) has too little attenuation, and previous layers are re-exposed. (b) has too much attenuation, and the newly exposed layer is unable to crosslink with the structures in the previous layer. (c) has sufficient attenuation, with only a small overlap of the exposed region to the structures in the previous layer.

Final hydrogel contains $\approx 75 \text{ wt\%}$ water - similar to tissue

Manufacturing Accuracy

Quantified from optical images and ultrasound images

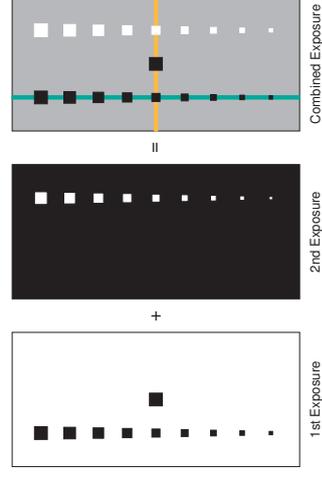


	Measured Pitch [µm]	Relative mean deviation
Designed Optical	2074	-
Optical	2115 ± 9	2 %
Ultrasound	2078 ± 71	0.2 %

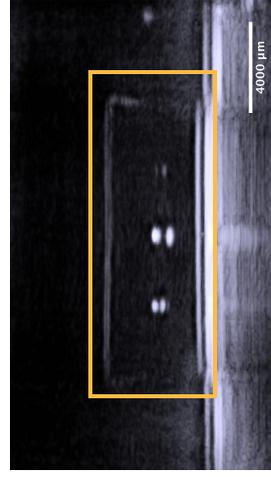
Fiducial Marker Tests

Marker Design Schematic

Relative sizes not to scale

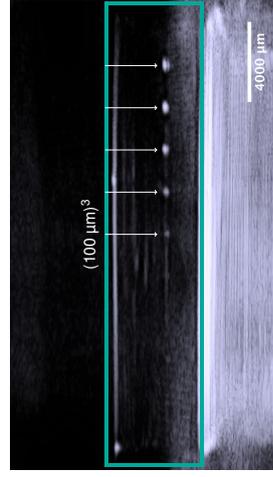


Marker Type



Ultrasound image showing the signal of two types of fiducial markers: Hollow (left and centre) and solid (right). The intensity from the hollow marker is 28 dB larger than the solid marker.

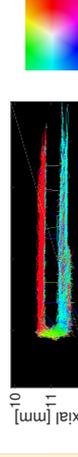
Marker Size



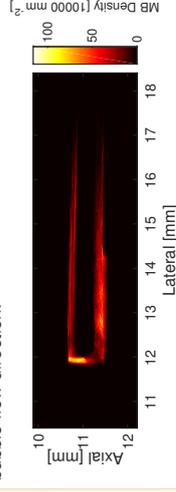
Ultrasound image showing the signal from hollow markers of different design sizes. The markers decrease from (140 µm)³ on the far right until (60 µm)³ on the far left in steps of 10 µm. The clearly visible markers are indicated by arrows.

3D Printed Flow Phantom

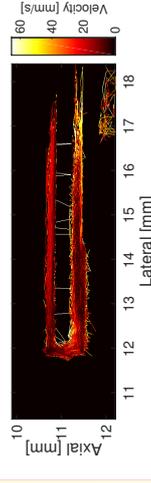
Measurements using a 3D printed phantom with (200 µm)² cross sectional channels. Liquid enters the lower channel to the right, and leaves the upper channel to the right.



Super-resolution image of micro-bubble trajectories in a 3D printed channel. The colour wheel indicates the micro-bubble flow direction.



Super-resolution density map of micro-bubbles in a 3D printed channel. The density is not uniform throughout the channel system.



Super-resolution velocity map of micro-bubbles in a 3D printed channel. The velocity profile is not uniform throughout the channel system.

Summary and Outlook

Phantoms have been printed by stereolithography.

- Similar acoustical properties to tissue due to 75 wt% water
- Hollow fiducial markers provide 28 dB more signal than solid fiducial markers
- Fiducial markers of only (100 µm)³ provide well defined point spread functions
- 3D printed flow phantom with (200 µm)² cross sectional channels has been presented
- Structures are printed with a 2% increase in dimensions

In house testing has demonstrated (50 µm)² perfusable channels. The limit has yet to be reached.

D.2. POSTER 2 - REDUCED CAVITY PRESSURE IN FUSION BONDED DEVICES: IS A WAFER BONDER N

D.2 Poster 2 - Reduced cavity pressure in fusion bonded devices: is a wafer bonder necessary? (no)



REDUCED CAVITY PRESSURE IN FUSION BONDED DEVICES: IS A WAFER BONDER NECESSARY? (No)



Martin Lind Ommen¹ and Erik Vilain Thomsen¹

¹Department of Micro- and Nanotechnology, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

Fusion Bonding in Short

A three step process:

- Cleaning of the wafers
- Formation of pre-bond
- Anneal of the bond

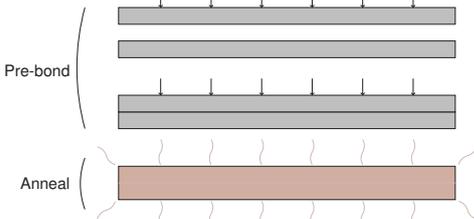


Plate Deflection

The center deflection, w_0 , of a plate at a given pressure depends on the plate material.

For unstressed materials, such as silicon

$$w_0 = \frac{3(1-\nu^2)^2 a^4}{16 E h^3} \Delta p \quad (1)$$

where ν is Poisson's ratio, a is the radius of the plate, Δp is the pressure difference across the plate, E is Young's modulus, and h is the plate thickness.

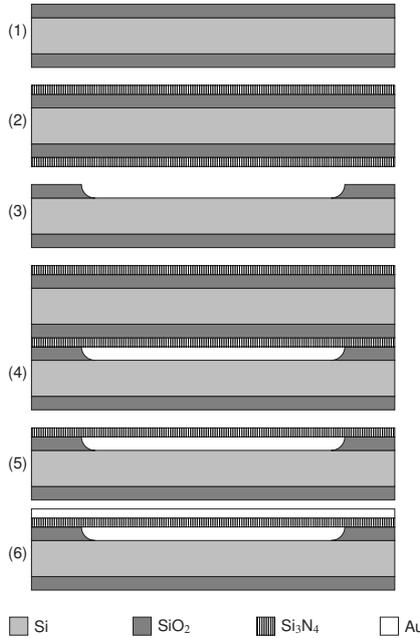
For plate materials with tensile stress, such as Si_3N_4

$$w_0 = \left(\frac{a}{2} \sqrt{\frac{CD}{N_t^3}} \left(\frac{1 - I_0\left(\sqrt{\frac{N_t}{CD}} a\right)}{I_1\left(\sqrt{\frac{N_t}{CD}} a\right)} \right) + \frac{a^2}{4N_t} \right) \Delta p$$

where a is the radius, C is a constant based on the solution to the plate equation, D is the flexural rigidity of the plate, I_n is the modified Bessel function of first kind, N_t is the stress resultant, and Δp is the pressure difference across the plate. In both cases

$$w_0 \propto \Delta p$$

Fabrication and Design

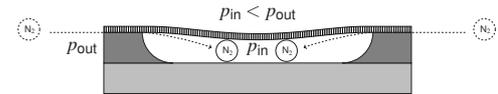


1. Top and bottom wafer oxidation
2. Top wafer Si_3N_4 deposition
3. Bottom wafer SiO_2 etch
4. Fusion bond and anneal top and bottom wafer
5. Handle wafer etch
6. Gold deposition for increased reflectivity

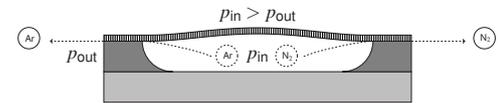
Wafer bonder parameters: Pre-bond at RT, 600 mbar stack pressure, 5 minutes
Anneal parameters: 1100 °C, 1 bar N_2 , 3 hours

Modified Hypothesis: Gas Diffuses In and Out of Cavities During Anneal

The wafers are annealed at 1100°C in 1 bar of N_2 . As the temperature increases, the initial pressure inside the cavities will change according to the ideal gas law, increasing approximately 4-fold for the air, handbond and argon devices, while the vacuum devices remain unchanged.



When the cavity pressure on the other hand is larger than the pressure outside of the cavity, such as in the air, handbond and argon devices, the pressure gradient will enable gas to diffuse out of the cavity during the high temperature anneal.

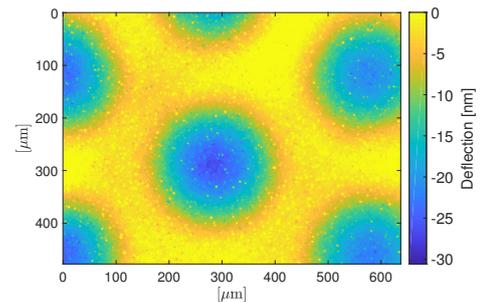


When the cavity pressure on the other hand is larger than the pressure outside of the cavity, such as in the air, handbond and argon devices, the pressure gradient will enable gas to diffuse out of the cavity during the high temperature anneal.

All devices will thus acquire a similar cavity pressure, corresponding to 1 bar at 1100°C, or ~0.2 bar at room temperature assuming gas diffusion continues until equilibrium. Once the interface has been annealed sufficiently, the cavities are sealed. This also means that interfaces may seal themselves before the equilibrium has been reached, which would result in different plate deflections.

Comparative Silicon Plate Study

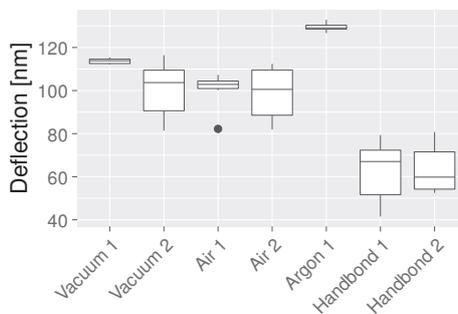
To verify whether the gas diffusion is only happening with Si_3N_4 plates, a set of handbonded wafers with SOI device layers as plates were made.



Applying the unstressed plate Equation (1), -30 nm should correspond to $\Delta p \sim 1$ bar, significantly more than the expected ~0.2 bar.

Deflection Measurements

The cavity pressure was indirectly measured by measuring the plate deflections using a Sensofar PLU Neox Optical Profiler. This enables measuring the deflection without applying a force on the plate, as would be the case for any stylus measurement system.

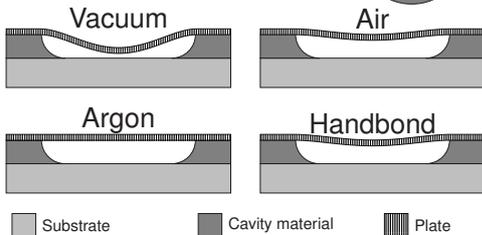
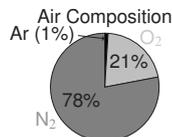


All devices deflect significantly. The argon device is particularly noticeable, since according to the hypothesis this should not deflect at all. This indicates that the initial hypothesis can not be correct, and there must be a gas exchange between the cavities and the external environment. The handbonded devices are not deflecting as much. However, the deflection is still larger than the expected 5th of the maximally deflecting device.

Working Hypothesis: The Bonding Atmosphere is the Cavity Atmosphere

Four pre-bond conditions:

- Bonding in vacuum (WB)
- Bonding in air (WB)
- Bonding in argon (WB)
- Bonding directly in hand



The three first pre-bond conditions are all acquired in a wafer bonder (WB). Considering only the gases present during the pre-bond, the vacuum bonded devices would be expected to deflect the most, the air and hand bonded devices should deflect about a 5th of the vacuum devices, due to the oxygen of the atmosphere being used to oxidize the silicon surfaces during the high temperature anneal (corresponding to a pressure decrease of 21%), and the argon in the argon bonded devices should not react at all during the anneal, and the plates should therefore not deflect.

Conclusion

- All devices deflect similarly
- The bonding interface is not leak tight directly after the pre-bond
- Gas diffusion during the high temperature anneal equilibrates the cavity pressure to the furnace pressure
- Lowering of the temperature after the anneal results in reduced cavity pressure regardless of the bonding atmosphere
- Bonding in a vacuum does not result in a post-anneal cavity vacuum
- The effect is seen for both silicon- and Si_3N_4 plates

No, a wafer bonder is not required to obtain a reduced cavity pressure.

E.1 Optical characterisation of scatterer sizes

This section provides a larger overview of the statistical model of the printed scatterer size analysis presented in Section 7.2.3, along with model diagnostics and model reduction. The combination of fixed and random factors makes the fitted model a linear mixed effects model. Such a model can be analysed using the `lmerTest` package [144] in R [145].

The data was analysed using a linear model of the form

$$\begin{aligned}
 Y_i = & \mu + \alpha_1(\text{Shape}_i) + \alpha_2(\text{DoseScheme}_i) \\
 & + \alpha_3(\text{Shape:DoseScheme}_i) \\
 & + [\beta_1 + \beta_2(\text{Shape}_i) + \beta_3(\text{DoseScheme}_i) \\
 & + \beta_4(\text{Shape:DoseScheme}_i)] x_{\text{design},i} \\
 & + \gamma_1(\text{Row}_i) + \gamma_2(\text{PrintColumn}_i) \\
 & + \gamma_3(\text{Row}_i^2) + \gamma_4(\text{PrintColumn}_i^2) \\
 & + \gamma_5(\text{Row:PrintColumn}_i) \\
 & + d(\text{Phantom}_i) + \epsilon_i
 \end{aligned} \tag{E.1}$$

where Y_i is the measured printed side length or diameter, μ is the overall intercept, $\alpha_1(\text{Shape}_i)$ is an intercept addition due to the Shape factor, $\alpha_2(\text{DoseScheme}_i)$ is an intercept addition due to the DoseScheme factor, $\alpha_3(\text{Shape:DoseScheme}_i)$ is an intercept addition due to the interaction between the Shape and the DoseScheme factor, β_1 is the overall slope of the model for correlation with the design size, $\beta_2(\text{Shape}_i)$ is a correction to the slope depending on the Shape factor, $\beta_3(\text{DoseScheme}_i)$ is a correction to the slope depending on the DoseScheme factor, $\beta_4(\text{Shape:DoseScheme}_i)$ is a correction to the slope depending on the interaction between the Shape and the DoseScheme factors, $\gamma_1(\text{Row}_i)$ is a slope addition due to the Row factor, $\gamma_2(\text{PrintColumn}_i)$ is a slope addition due to the PrintColumn factor, $\gamma_3(\text{Row}_i^2)$ is a quadratic addition due to the Row factor, $\gamma_4(\text{PrintColumn}_i^2)$ is a quadratic addition due to the PrintColumn factor, $\gamma_5(\text{Row:PrintColumn}_i)$ is an addition due to the interaction between the Row and the PrintColumn factor, $d(\text{Phantom}_i) \sim N(0, \sigma_{\text{phantom}}^2)$ is a random offset from phantom to phantom, and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All $d(\text{Phantom}_i)$'s and ϵ_i 's are independent.

Model diagnostics of the initial model shows no issues with the residuals. They appear to be normal distributed, and show no obvious tendencies with the tested factors. This can be seen in

Table E.1: BIC values of the compared models for the scatterer size analysis. Each model has only had a single term removed from the previous model, which is noted under column “Removed term”. *Model 5 is the final model as Model 6 increases BIC.

	Removed term	Degrees of freedom	BIC
Model 1	\sim	19	4325.9645
Model 2	$\gamma_3(\text{Row}^2)$	18	4315.2182
Model 3	$\beta_4(\text{Shape:DoseScheme})$	16	4309.0732
Model 4	$\beta_2(\text{Shape})$	15	4301.4394
Model 5	$\gamma_5(\text{Row:PrintColumn})$	14	4291.3226
Model 6*	$\alpha_3(\text{Shape:DoseScheme})$	12	4295.1404

Figure E.1. Normalised residuals are plotted against theoretical $N(0,1)$ quantiles in a quantile-quantile (Q-Q)-plot (Figure E.1(a)), against model predicted values (Figure E.1(b)), and against the main factors investigated (Figure E.1(c), E.1(d), and E.1(e)). The dashed lines mark the range within which 95% of the residuals is expected to lie. Note that with a total of 720 measurements, 36 are expected to lie outside of the region.

The model was reduced based on the bayesian information criterion (BIC). The BIC values of the compared models can be seen in Table E.1. Model 1 represents the full model in Equation (E.1). The other models have had a single term removed from the previous model, based on the highest p – value, while making sure not to remove terms included in higher order effects. The removed factor in each model is noted under column “Removed term”.

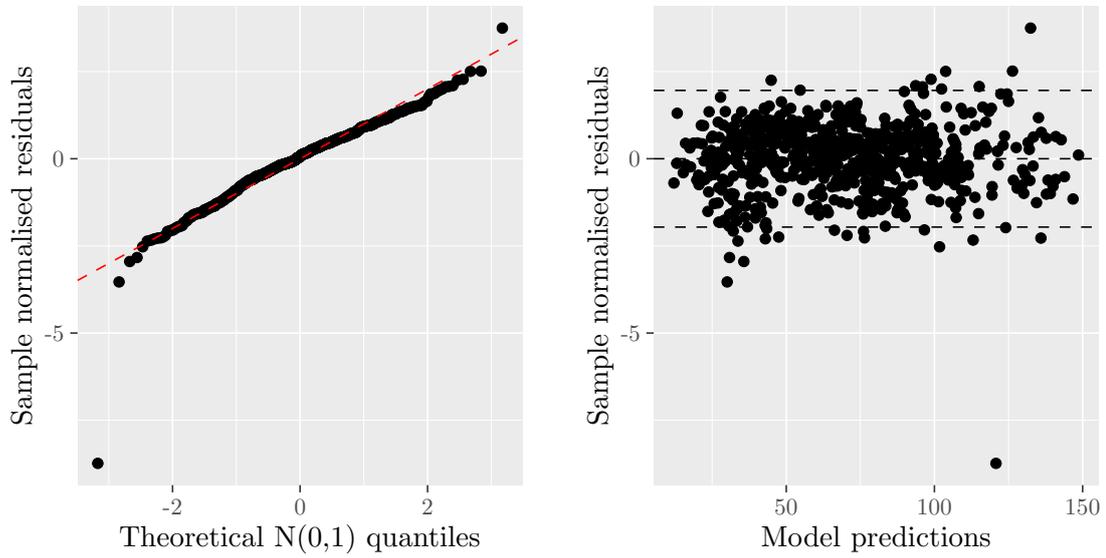
The model reduction converges at Model 5, as the BIC increase in Model 6. This model was presented in the main text as

$$\begin{aligned}
 Y_i = & \mu + \alpha_1(\text{Shape}_i) + \alpha_2(\text{DoseScheme}_i) \\
 & + \alpha_3(\text{Shape:DoseScheme}_i) \\
 & + (\beta_1 + \beta_3(\text{DoseScheme}_i)) x_{\text{design},i} \\
 & + \gamma_1(\text{Row}_i) + \gamma_2(\text{PrintColumn}_i) \\
 & + \gamma_4(\text{PrintColumn}_i^2) \\
 & + d(\text{Phantom}_i) + \epsilon_i
 \end{aligned} \tag{E.2}$$

Model diagnostics of the final model once again showed no issues with the residuals. These can be seen in Figure E.2. The dashed lines mark the range within which 95% of the residuals is expected to lie. Note that with a total of 1440 measurements, 72 are expected to lie outside of the region.

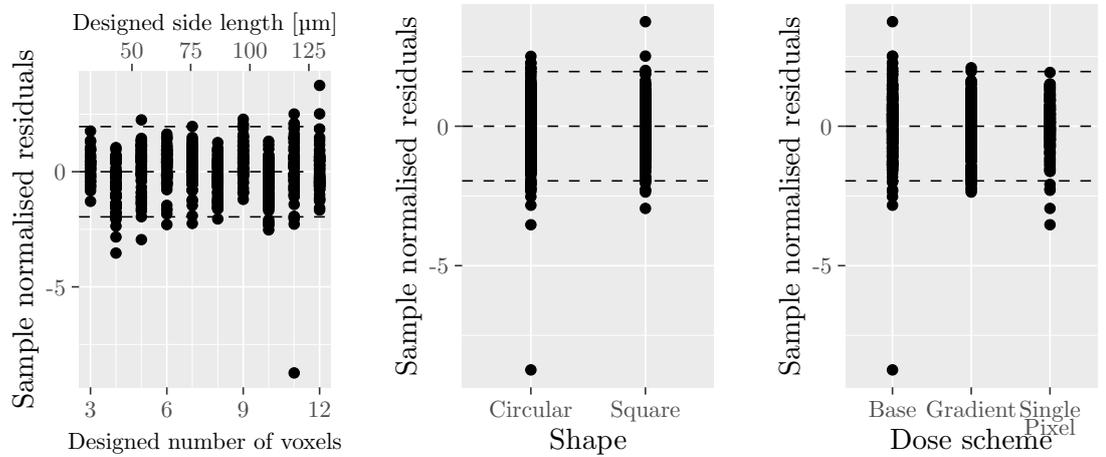
E.2 Scattering strength

This section provides a larger overview of the statistical model of the scatterer intensity analysis presented in Section 7.2.4, along with model diagnostics and model reduction. The combination of fixed and random factors makes the fitted model a linear mixed effects model. Such a model can be analysed using the `lmerTest` package [144] in R [145].



(a) Q-Q-plot of normalised residuals

(b) Normalised residuals against model prediction

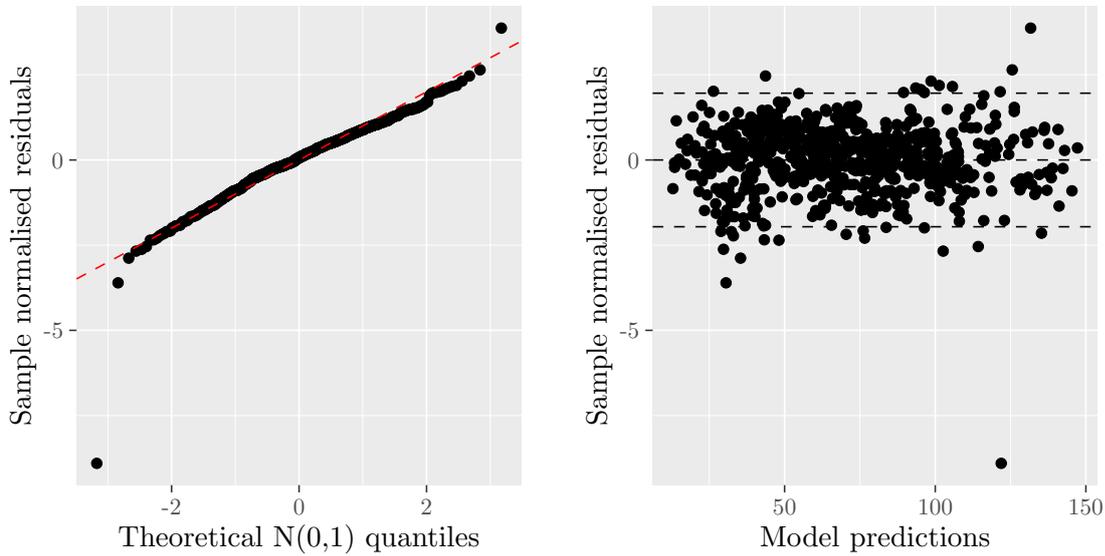


(c) Normalised residuals against design size

(d) Normalised residuals against shape

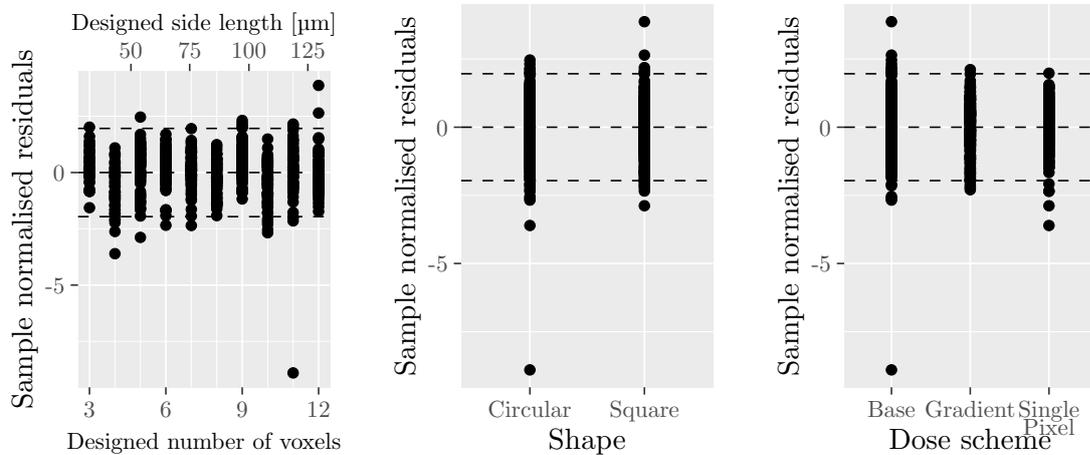
(e) Normalised residuals against dose scheme

Figure E.1: Residual analysis of initial model of the scatterer size. (a) shows a Q-Q-plot of the residuals. Apart from a single extreme residual, all residuals fall on the expected line, indicating normal distributed data. (b) shows the normalised residuals plotted against model prediction value. (c), (d), and (e) shows the normalised residuals against the main factors. The dashed lines mark the range within which 95% of the residuals is expected to lie. No abnormal tendencies are seen in the data.



(a) Q-Q-plot of normalised residuals

(b) Normalised residuals against model prediction



(c) Normalised residuals against design size

(d) Normalised residuals against shape

(e) Normalised residuals against dose scheme

Figure E.2: Residual analysis of final model of the scatterer size. (a) shows a Q-Q-plot of the residuals. Apart from a single extreme residual, all residuals fall on the expected line, indicating normal distributed data. (b) shows the normalised residuals plotted against model prediction value. (c), (d), and (e) shows the normalised residuals against the main factors. The dashed lines mark the range within which 95% of the residuals is expected to lie. No abnormal tendencies are seen in the data.

The data was analysed using a linear model of the form

$$\begin{aligned}
Y_i = & \mu + \alpha_1(\text{Shape}_i) + \alpha_2(\text{DoseScheme}_i) \\
& + \alpha_3(\text{Shape:DoseScheme}_i) \\
& + [\beta_1 + \beta_2(\text{Shape}_i) + \beta_3(\text{DoseScheme}_i) \\
& + \beta_4(\text{Shape:DoseScheme}_i)] x_{\text{design},i} \\
& + \gamma_1(\text{Row}_i) + \gamma_2(\text{Row}_i^2) \\
& + \gamma_3(\text{ImageColumn}_i) \\
& + \gamma_4(\text{PrintColumn}_i) + \gamma_5(\text{Column}_i^2) \\
& + \gamma_6(\text{Row:PrintColumn}_i) \\
& + d(\text{Phantom}_i) + e(\text{Phantom:Flip}) + \epsilon_i,
\end{aligned} \tag{E.3}$$

where Y_i is the measured scattering intensity in dB, μ is the overall intercept, $\alpha_1(\text{Shape}_i)$ is an intercept addition due to the Shape factor, $\alpha_2(\text{DoseScheme}_i)$ is an intercept addition due to the DoseScheme factor, $\alpha_3(\text{Shape:DoseScheme}_i)$ is an intercept addition due to the interaction between the Shape and the DoseScheme factor, β_1 is the overall slope of the model for correlation of intensity with the design size in voxels, $\beta_2(\text{Shape}_i)$ is a correction to the slope depending on the Shape factor, $\beta_3(\text{DoseScheme}_i)$ is a correction to the slope depending on the DoseScheme factor, $\beta_4(\text{Shape:DoseScheme}_i)$ is a correction to the slope depending on the interaction between the Shape and the DoseScheme factors, $\gamma_1(\text{Row}_i)$ is a slope addition due to the Row factor, $\gamma_2(\text{Row}_i^2)$ is a quadratic addition due to the Row factor, $\gamma_3(\text{ImageColumn}_i)$ is a slope addition due to the ImageColumn factor, $\gamma_4(\text{PrintColumn}_i)$ is a slope addition due to the PrintColumn factor, $\gamma_5(\text{Column}_i^2)$ is a quadratic addition due to the PrintColumn or ImageColumn factor (indistinguishable), $\gamma_6(\text{Row:PrintColumn}_i)$ is an addition due to the interaction between the Row and the PrintColumn factor, $d(\text{Phantom}_i) \sim N(0, \sigma_{\text{Phantom}}^2)$ is a random offset from phantom to phantom, $e(\text{Phantom:Flip}_i) \sim N(0, \sigma_{\text{Phantom:Flip}}^2)$ is a random offset further separated into and $\epsilon_i \sim N(0, \sigma^2)$ is the residual error, with $N(\mu, \sigma^2)$ being a normal distribution with mean μ and standard deviation σ , all for the i th response. All $d(\text{Phantom}_i)$'s and ϵ_i 's are independent.

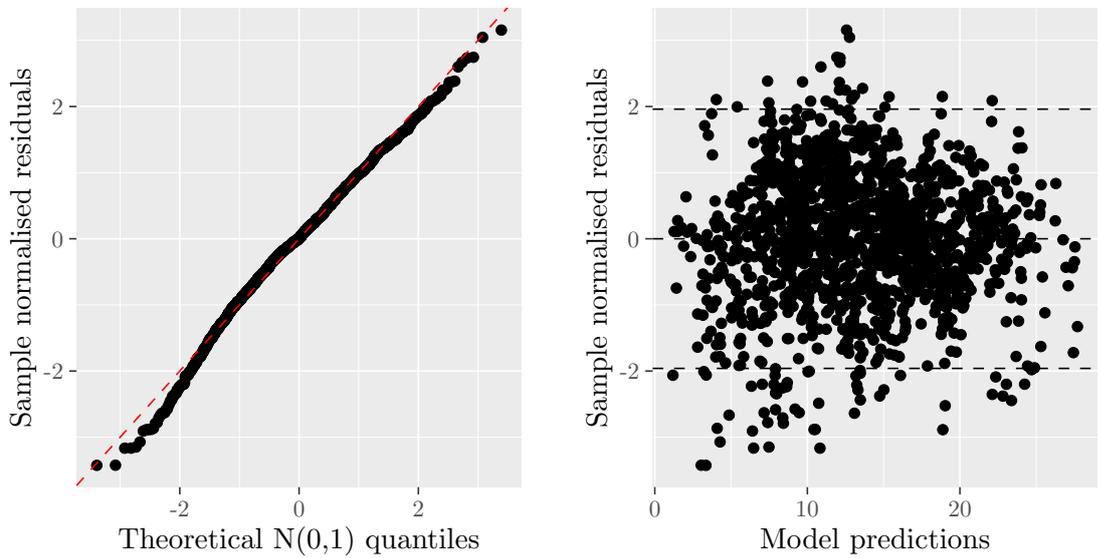
Model diagnostics of the initial model shows no issues with the residuals. They appear to be normal distributed, and show no obvious tendencies with the tested factors. This can be seen in Figure E.3. Normalised residuals are plotted against theoretical $N(0,1)$ quantiles in a Q-Q-plot (Figure E.1(a)), against model predicted values (Figure E.1(b)), and against the main factors investigated (Figure E.1(c), E.1(d), and E.1(e)). The dashed lines mark the range within which 95% of the residuals is expected to lie. Note that with a total of 1440 measurements, 72 are expected to lie outside of the region.

The model was reduced based on the BIC. The BIC values of the compared models can be seen in Table E.2. Model 1 represents the full model in Equation (E.3). The other models have had a single term removed from the previous model, based on the highest p -value, while making sure not to remove terms included in higher order effects. The removed factor in each model is noted under column "Removed term".

The model reduction converges at Model 7, as the BIC increase in Model 8. This model was presented in the main text as

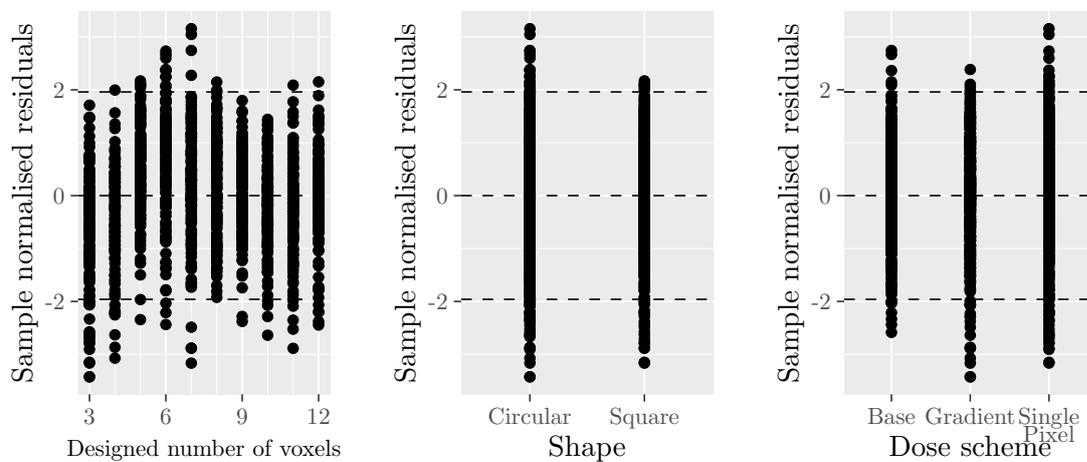
$$\begin{aligned}
Y_i = & \mu + \alpha_2(\text{DoseScheme}_i) \\
& + [\beta_1 + \beta_2(\text{Shape}_i) \\
& + \beta_3(\text{DoseScheme}_i)] x_{\text{design},i} \\
& + \gamma_1(\text{Row}_i) + \gamma_2(\text{PrintColumn}_i) \\
& + \gamma_3(\text{Row}_i^2) + \gamma_4(\text{Column}_i^2) \\
& + d(\text{Phantom}_i) + \epsilon_i,
\end{aligned} \tag{E.4}$$

Model diagnostics of the final model once again showed no issues with the residuals. These can be seen in Figure E.4. The dashed lines mark the range within which 95% of the residuals is



(a) Q-Q-plot of normalised residuals

(b) Normalised residuals against model prediction

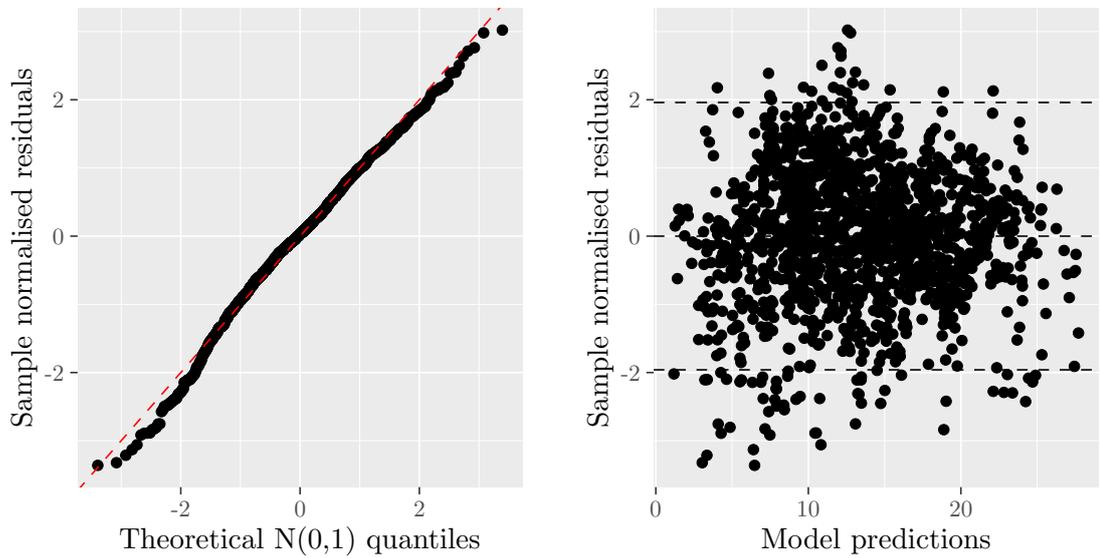


(c) Normalised residuals against design size

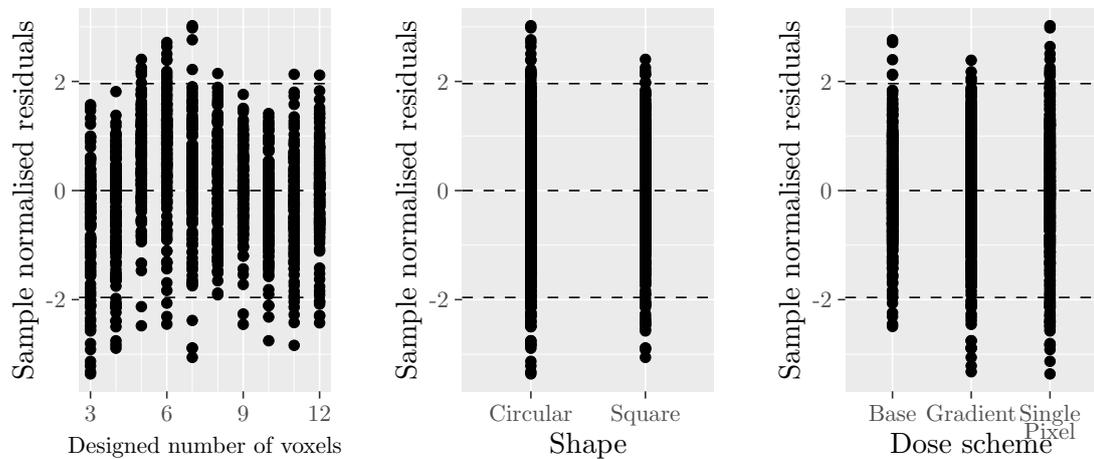
(d) Normalised residuals against shape

(e) Normalised residuals against dose scheme

Figure E.3: Residual analysis of initial model of the scatterer intensity. (a) shows a Q-Q-plot of the residuals. Apart from a single extreme residual, all residuals fall on the expected line, indicating normal distributed data. (b) shows the normalised residuals plotted against model prediction value. (c), (d), and (e) shows the normalised residuals against the main factors. The dashed lines mark the range within which 95% of the residuals is expected to lie. No abnormal tendencies are seen in the data.



(a) Q-Q-plot of normalised residuals (b) Normalised residuals against model prediction



(c) Normalised residuals against design size (d) Normalised residuals against shape (e) Normalised residuals against dose scheme

Figure E.4: Residual analysis of final model of the scatterer intensity. (a) shows a Q-Q-plot of the residuals. Apart from a single extreme residual, all residuals fall on the expected line, indicating normal distributed data. (b) shows the normalised residuals plotted against model prediction value. (c), (d), and (e) shows the normalised residuals against the main factors. The dashed lines mark the range within which 95% of the residuals is expected to lie. No abnormal tendencies are seen in the data.

Table E.2: BIC values of the compared models for the scatterer intensity analysis. Each model has only had a single term removed from the previous model, which is noted under column “Removed term”. *Model 7 is the final model as Model 8 increases BIC.

	Removed term	Degrees of freedom	BIC
Model 1	\sim	21	7246.7330
Model 2	$e(\text{Phantom:Flip})$	20	7239.4606
Model 3	$\gamma_6(\text{Row:PrintColumn})$	19	7225.2174
Model 4	$\gamma_3(\text{ImageColumn})$	18	7213.0861
Model 5	$\beta_4(\text{Shape:DoseScheme})$	16	7203.2678
Model 6	$\alpha_3(\text{Shape:DoseScheme})$	14	7188.9880
Model 7	$\alpha_1(\text{Shape})$	13	7181.8277
Model 8*	$\gamma_5(\text{Column}_i^2)$	11	7282.0486

expected to lie. Note that with a total of 1440 measurements, 72 are expected to lie outside of the region.

Process optimisation and analysis scripts

F.1 Thin film thermal processing time

A description of the furnace processing time script is given in the following. The script is fairly long and most parts are only of local relevance due to the way processes are logged. Therefore it is not included directly, but conceptually described. The script takes the film thickness as an input, and provides the processing time required to obtain that thickness, along with the residual error of the statistical model based on the data log from the furnace used. Instead of creating a purely theoretical model to predict idealised processing times, the actual process logs from the used furnaces are used as input, providing very good results. The data is based on measurements of a new wafer centrally placed in the quartz boat during every process. Thus the residual error marks the variability of the film thicknesses from deposition to deposition. The issue with using processing logs is that the furnace might change behaviour over time, meaning not all data is necessarily equally representative of what to expect from the furnace when it is actually going to be used. Therefore, it is possible to input exactly how many of the log entries to include, resulting in only that amount of the newest entries being analysed. Another issue is the human factor, with users occasionally inputting bad data in the process log. This is tested by iteratively analysing the residuals of the model, and discarding outlier values which do not fit the distribution of outliers sufficiently well, before remodelling the remaining data.

Figure F.1 shows two the output plots of the furnace data and the fitted model. Figure F.1(a) shows the output from an oxidation furnace, and Figure F.1(b) shows the output from an low pressure chemical vapour deposition (LPCVD) Si_3N_4 furnace. The processes are different with the thermal oxidation being described by the Deal-Grove model, and LPCVD deposition being a linear deposition. The models are shown in the figures. The script recognises the process based on the input furnace log file name.

A few input errors can be seen for short processing time in Figure F.1(a), likely due to a user entering for instance 1 hour as 1 minute, or a 0 thickness value near 40 minutes in Figure F.1(b), likely simply because the user forgot to enter a value. The erroneous values should be removed automatically.

Some furnaces provide multiple different process recipes, in which case the data log first needs to be filtered based on this. Once the correct data has been selected, the data is fitted to either the Deal-Grove model or a linear model. The fitting function in MATLAB has trouble fitting a square root expression for the film thickness, and would often provide ambiguous results. For oxidations, the relationship is flipped, fitting time as a function of thickness instead, making the fitting model depend on x^2 , which works without problems. After a model has been fitted to the data, an

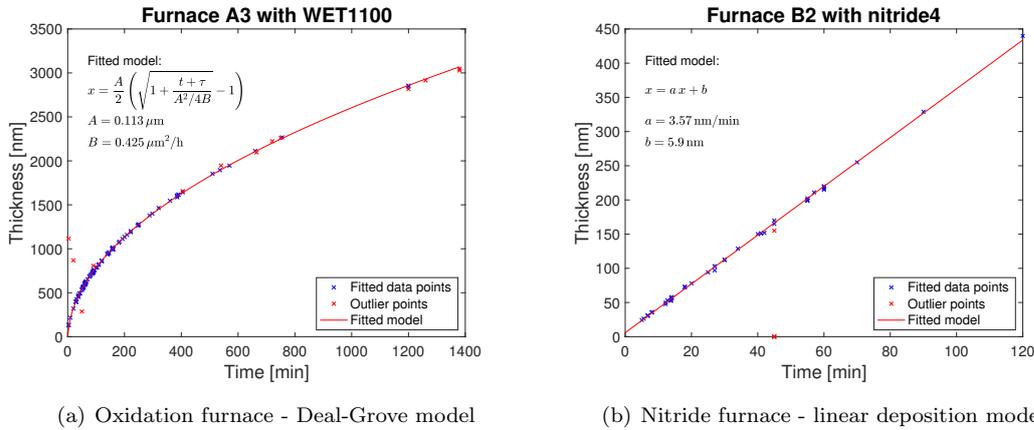


Figure F.1: Output plots from the furnace script. Figure F.1(a) shows the output from an oxidation furnace, and Figure F.1(b) shows the output from an LPCVD Si_3N_4 furnace. The fits are based on different models for the two types of deposition.

outlier detection loop starts. The following procedure is the same regardless of the fitted model. The residuals of the data log points to the fitted model are calculated. These are normalised to the standard deviation of the residuals, and any normalised residual larger than 3 are marked as outliers, seen as red crosses in the figures. The model is then fitted again without the outlier values, resulting in slight changes to the model. This procedure of refitting and outlier detection can be run any number of times as input in the furnace script. For stable processes, which the oxidation furnaces typically run, very small deviations from the average model will be marked as outliers. LPCVD processes are typically less stable, and the furnaces even change behaviour over time. Therefore, it is also possible to enter the number of data entries to use for the fit, with only the newest entries corresponding to the input value being considered. Finally, the last remaining input value is the desired thickness. The last fitted model is used to determine what processing time will result in the desired thickness. The residuals of the fitted model can be used to determine the magnitude of the variation in the furnace operation.

F.2 Film thickness map

The film thickness map script is included below. The script allows for creating a wafer map of thicknesses, or other parameters, based on a set of measurements at different coordinates. The script works with a variety of number of samples, regardless of distribution of the points, as the coordinates are also used as input.

The script allows for using a companion file containing the values of the mean square error (MSE) of the thickness measurements. The distribution of MSE values is considered, and any thickness values which are deemed outliers are discarded, and the wafer contour map will be created based on the remaining valid points. The distribution of MSE values are typically one-sided, with a majority of values around a fairly low value, with outliers always being high. It is not uncommon with extreme MSE outliers. A single extreme MSE value might create a large offset to the mean value of the population. Instead, the median MSE value was used. The difference from each point to the median value was calculated for all points. Outliers were flagged as values differing by more than 1.5 standard deviations away from the median value. The exact distance varied with film type.

Note that for better performance, the differences should have been normalised to the variation of the difference to the median value instead of the pure standard deviation.

```
1 # -*- coding: utf-8 -*-
```

```

2  """
3  Created on Mon Mar 21 11:17:55 2016
4
5  @author: s113044
6  """
7
8  import numpy as np
9  #from matplotlib.mlab import griddata # uses x y and z columns
10 from scipy.interpolate import griddata # uses matrix of x y and z data
11 import matplotlib.pyplot as plt
12 #from mpl_toolkits.axes_grid1 import make_axes_locatable
13 import os
14
15 os.chdir("/Volumes/NTCH$/Group_members/MartinLindOmmen/Projects/
    ↪ FurnaceCharacterisation/A3") # change directory to the directory of your
    ↪ datafiles
16 # Load Files
17
18 plt.rcParams.update({'font.size': 13})
19 dataT = np.loadtxt(open("a3Slot29Thick.txt","rb"), skiprows=2) # change filename
    ↪ to the thickness .txt file
20 # load MSE data to be able to replace badly fitted values. If MSE data is not
    ↪ available, code should be modified manually below
21 dataMSE = np.loadtxt(open("a3Slot29MSE.txt","rb"), skiprows=2) # change filename
    ↪ to the MSE .txt file
22
23 # separate the x, y and z data - mse comes later
24 xT = dataT[:,0]
25 yT = dataT[:,1]
26 zT = dataT[:,2]
27 # Optional - uncomment line below to subtract mean of Z to observe variation
    ↪ around mean (and get nice variation in the contours - negative values give
    ↪ dashed lines)
28 #zT = zT - np.mean(zT)
29
30 # define the edges of the grid from the min/max values
31 min_xT = np.min(xT)
32 min_yT = np.min(yT)
33 max_xT = np.max(xT)
34 max_yT = np.max(yT)
35 dim_xT = np.size(xT)
36 dim_yT = np.size(yT)
37
38
39 # Create interpolated plot of full dataset, without replacing measured values
40 #create interpolation grid
41 interpFactor = 3 # set the factor for how much denser you would like your
    ↪ gradients
42 x = np.linspace(min_xT, max_xT, interpFactor*dim_xT)
43 y = np.linspace(min_yT, max_yT, interpFactor*dim_yT)
44 # convert to matrix
45 X,Y = np.meshgrid(x, y)
46 #interpolate Z data in new grid

```

```

47 Z = griddata((xT,yT), zT, (X,Y), method='cubic') # nearest, linear or cubic -
    ↪ fill value writes NaN outside of the measured region as a default
48
49 # draw wafer outline
50 WaferR = 5 #wafer radius
51 flatX = [-1.58,1.58] # flat x values
52 y1 = 4.7+0.16 # lower bound for plotting area (cuts of part of circle)
53 flatY = [y1*1.007,y1*1.007]
54
55 #begin figure
56 fig = plt.figure(1)
57 ax = fig.add_subplot(111) # add to control plot behaviour
58 ax.set_aspect('equal') # add to set aspect ratio
59 # contour the gridded data, plotting dots at the nonuniform data points.
60 v = np.linspace(1200, 2400, 100, endpoint=True) # optional line - fix color bar
    ↪ range
61 CS = plt.contour(X, Y, Z, 12, linewidths=0.5, colors='k') # draw contour lines.
    ↪ 12 is the default number chosen
62 CS = plt.contourf(X, Y, Z, 100, cmap='gist_rainbow') # syntax (xdata,ydata,zdata,
    ↪ numberOfColors,linewidth...), for fixed color range, change numberOfColors
    ↪ to v
63 cbar = plt.colorbar(label = 'Thickness [nm]') # draw colorbar
64 tick_locator = plt.MaxNLocator(nbins=8) # define the ticks for the colorbar
65 cbar.locator = tick_locator
66 cbar.update_ticks()
67 # # plot data points.
68 plt.scatter(xT, yT, marker='o', c='b', s=5) # draw the originally measured points
    ↪ as blue dots
69 circle2=plt.Circle((0,0.15),5,color='k', fill=False, linewidth=1) #Perimeter of
    ↪ wafer
70 fig.gca().add_artist(circle2)
71 plt.plot(flatX,flatY, linewidth=1, color = 'k') # Wafer flat
72 plt.xlim(-WaferR*1.02, WaferR*1.02)
73 plt.ylim(-WaferR*1.02, y1*1.01)
74 plt.axis('off')
75 plt.gca().invert_xaxis() # invert x (and y) axis, to orient wafer properly
76 plt.gca().invert_yaxis() # invert y (and x) axis, to orient wafer properly
77 #plt.show()
78 # fig.savefig('afterDevelopmentProximityAbsolute.png', format='png', dpi=500,
    ↪ bbox_inches='tight')
79
80
81
82 #####
83
84
85 # now modify data to consider MSE data, to automatically remove badly fitted
    ↪ values
86 zMSE = dataMSE[:,2]
87 medianMSE = np.median(zMSE) # median is chosen due to the distribution not being
    ↪ even, and the results are therefore better (prettier). This can be
    ↪ optimised
88 #meanMSE = np.mean(zMSE)

```

```

89 stdMSE = np.std(zMSE)
90 mseResiduals = zMSE - medianMSE
91
92 stResiduals = mseResiduals/stdMSE
93 outlierIND = np.nonzero(stResiduals>1.5) # outliers defined as being more than 3
    ↪ standard deviations from median. the amount may vary depending on film
    ↪ type
94
95 xTrep = np.delete(xT,outlierIND)
96 yTrep = np.delete(yT,outlierIND)
97 zTrep = np.delete(zT,outlierIND)
98
99 xTOutlier = xT[outlierIND]
100 yTOutlier = yT[outlierIND]
101
102 # define the grid
103 min_xT = np.min(xTrep)
104 min_yT = np.min(yTrep)
105 max_xT = np.max(xTrep)
106 max_yT = np.max(yTrep)
107 dim_xT = np.size(xTrep)
108 dim_yT = np.size(yTrep)
109
110 # start by creating interpolated plot of full dataset
111
112 #create interpolation grid
113 x = np.linspace(min_xT, max_xT, interpFactor*dim_xT)
114 y = np.linspace(min_yT, max_yT, interpFactor*dim_yT)
115 # convert to matrix
116 X,Y = np.meshgrid(x, y)
117 #interpolate Z data in new grid
118 Z = griddata((xTrep,yTrep), zTrep, (X,Y), method='cubic') # nearest, linear or
    ↪ cubic - fill value writes NaN outside of the measured region as a default
    ↪ ,fill_value = 'nan'
119
120 #begin figure
121 fig = plt.figure(2)
122 ax = fig.add_subplot(111) # add to control plot behaviour
123 ax.set_aspect('equal') # add to set aspect ratio
124 # contour the gridded data, plotting dots at the nonuniform data points.
125 v = np.linspace(1200, 2400, 100, endpoint=True) # optional line - fix color bar
    ↪ range
126 CS = plt.contour(X, Y, Z, 12, linewidths=0.5, colors='k') # draw contour lines.
    ↪ 12 is the default number chosen
127 CS = plt.contourf(X, Y, Z, 100, cmap='gist_rainbow') # syntax (xdata,ydata,zdata,
    ↪ numberOfColors,linewidth....), for fixed color range, change numberOfColors
    ↪ to v
128 cbar = plt.colorbar(label = 'Thickness [nm]') # draw colorbar
129 tick_locator = plt.MaxNLocator(nbins=8) # define the ticks for the colorbar
130 cbar.locator = tick_locator
131 cbar.update_ticks()
132 # # plot data points.
133 plt.scatter(xT, yT, marker='o', c='b', s=5) # draw the originally measured points

```

```
    ↪ as blue dots
134 plt.scatter(xTOutlier, yTOutlier, marker='o', c='r', s=5) # draw data points
    ↪ marked as outliers as red dots
135 circle2=plt.Circle((0,0.15),5,color='k', fill=False, linewidth=1) #Perimeter of
    ↪ wafer
136 fig.gca().add_artist(circle2)
137 plt.plot(flatX,flatY, linewidth=1, color = 'k') # Wafer flat
138 plt.xlim(-WaferR*1.02, WaferR*1.02)
139 plt.ylim(-WaferR*1.02, y1*1.01)
140 plt.axis('off')
141 plt.gca().invert_xaxis() # invert x (and y) axis, to orient wafer properly
142 plt.gca().invert_yaxis() # invert y (and x) axis, to orient wafer properly
143 plt.show()
```

F.3 Characterisations of furnaces

The distribution of film thicknesses on a wafer boat for a nitride furnace was presented in the main text. Similar characterisations were made for other furnaces in the cleanroom. These are included here.

Figure F.2 shows the thickness uniformity across a furnace boat in the oxidation furnace A3. The two subfigures compares a full boat of wafers with that of half a boat of wafers. It is not uncommon to use a set of dummy wafers at each end of the boat to increase the temperature uniformity of the remaining wafers. This is likely the effect seen in the figures as the thickness decreases towards either end of the wafer boat. Note that depending on the uniformity requirements, it might be necessary with more than two wafers at each end. With only half a boat, the same decreasing tendencies are observed, only moved in towards the centre. This shows that it is in fact the dummy wafers at the ends which stabilises the process.

Another oxidation furnace is shown in Figure F.3.

Finally, the B2 nitride furnace is repeated here in Figure F.4.

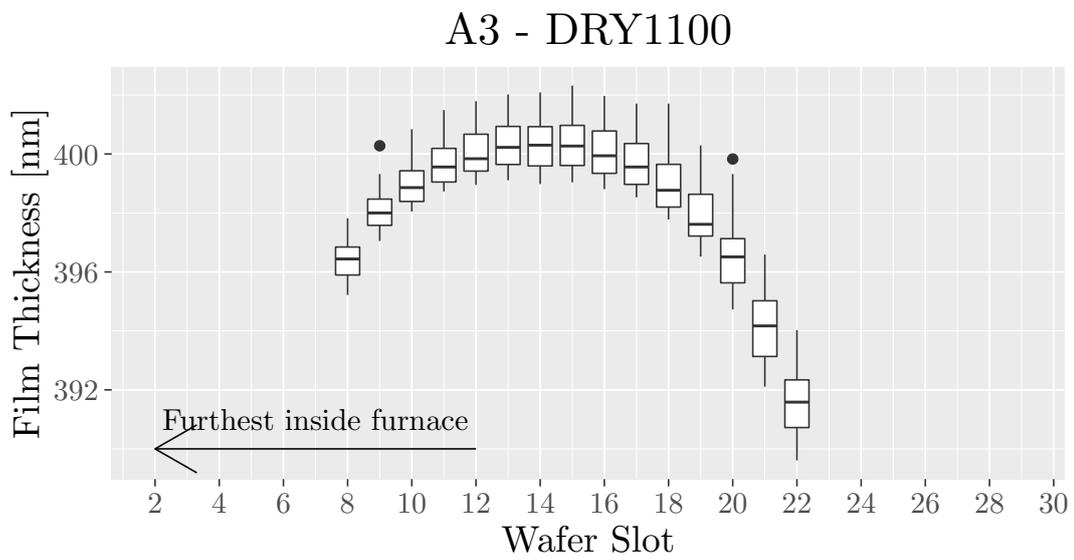
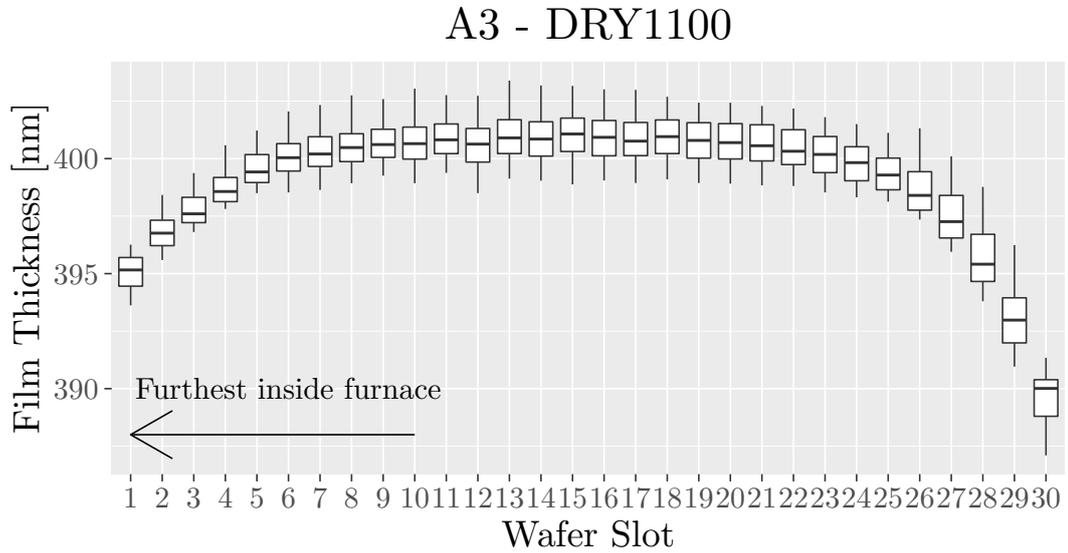


Figure F.2: Film thicknesses across a wafer boat of wafers after thermal oxidation using DRY1100 in the A3 oxidation furnace. (a) shows the thickness distributions across a full wafer boat. (b) shows the thickness distributions when run with only half a boat of wafers.

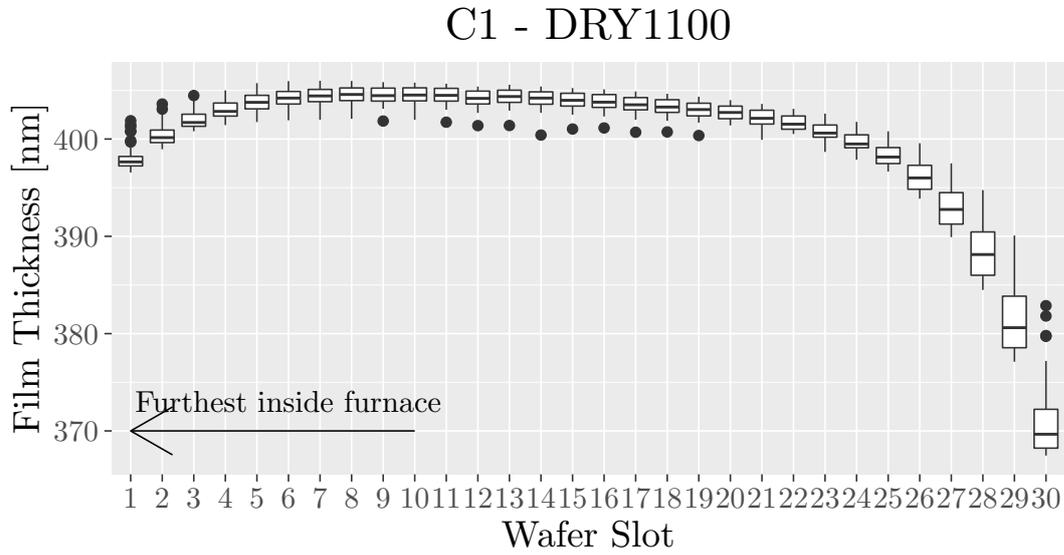


Figure F.3: Film thicknesses across a wafer boat of wafers after thermal oxidation using DRY1100 in the C1 oxidation furnace.

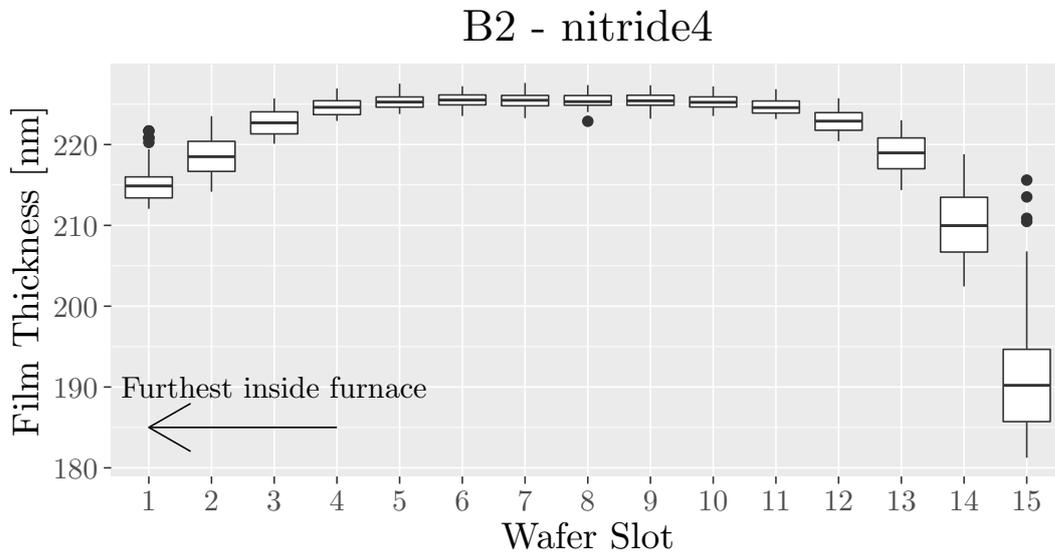


Figure F.4: Film thicknesses across a wafer boat of wafers after nitride deposition using nitride4 in the B2 nitride deposition furnace.

Phantom generation scripts

G.1 Flow channel scripts

The flow phantom script was conceptually described in the main text, Section 8.3. The channel generation scripts are included here.

G.1.1 Cylindrical channels

```
1 function [phantom, path] = axilinearParametricCylinderChannel(phantom, diameter,
   ↪ channelCurvature, startCoordinate, endCoordinate)
2
3 % determine size of the phantom matrix in order to ensure nothing is done
4 % for negative matrix indices or indices beyond the original phantom size
5 [zDim, xDim, ~, yDim] = size(phantom);
6
7 coordDim = size(startCoordinate);
8 if numel(diameter) == coordDim(1)-1
9     % do nothing - everything is fine
10 elseif numel(diameter)==1
11     % same diameter for all segments - just repmat
12     %diameter = repmat(diameter,1,coordDim(1)-1);
13     % edit to not break
14 else
15     fprintf('Error: The number of diameters passed does not match with the number
   ↪ of channel segments')
16 end
17
18 % first, convert input diameter to pixel values
19 pixelSize = 10.8;
20 pixDiameter = diameter/pixelSize;
21 pixCurvature = channelCurvature/pixelSize;
22
23 layerSize = 20;
24
```

```

25 % correction factor due to the anisotropy of the different axes - might
26 % need modification depending on the design orientation - right now, the
27 % prints are defined on the side, essentially switching z and y. -
28 % therefore the correction relates to y - but if the designs are not
29 % rotated, it will relate to z.
30 anisotropyCorrection = pixelSize/layerSize; %
31
32
33
34 %rescale the coordinates to the voxel grid of 10.8x10.8x20 micrometer3
35 % note rescaling is made according to the rotated coordinate system, with y
36 % being the layer height - each y is 20 micrometer
37 startCoordVox = round([rescale(startCoordinate(1),1,xDim,'InputMin',0,'InputMax',
    ↪ xDim*10.8), ...
38     rescale(startCoordinate(2),1,yDim,'InputMin',0,'InputMax',yDim*20), ...
39     rescale(startCoordinate(3),1,zDim,'InputMin',0,'InputMax',zDim*10.8)]);
40 endCoordVox = round([rescale(endCoordinate(1),1,xDim,'InputMin',0,'InputMax',xDim
    ↪ *10.8), ...
41     rescale(endCoordinate(2),1,yDim,'InputMin',0,'InputMax',yDim*20), ...
42     rescale(endCoordinate(3),1,zDim,'InputMin',0,'InputMax',zDim*10.8)]);
43
44 %% there are a ton of special cases
45 % first outer loop case will be that if multiple coordinates are passed,
46 % the very first segment starts out straight, but any following segments
47 % will be rounded in the connections between each bend - until the very
48 % last one which will also be straight.
49
50 % In addition, if one or more of the
51 % coordinates dont change (i.e. start x is the same as end-x) - this will both
52 % affect the linear segments, as well as the curve segments
53
54 % it will be necessary to potentially adjust the length of the segments
55 % with the radius of curvature of the curved segments - if they are present
56
57 % make start-coordinate/end-coordinate a parametric curve
58 % first section along the x-axis
59
60
61
62
63 if startCoordVox(1) == endCoordVox(1) % if start and endcoordinate of x is the
    ↪ same,
64     % the 'x' part of the path will not exist
65     pathAX = [];
66     pathAY = [];
67     pathAZ = [];
68 else
69     if startCoordVox(1) < endCoordVox(1)
70         fac = 1;
71     else
72         fac = -1;
73     end
74     if startCoordVox(2) ~= endCoordVox(2) || startCoordVox(3) ~= endCoordVox(3) %

```

```

75     % if both y and z start and end in the same coordinate (no
76     % segments) no correction will need to be made to the length of the
77     % x segment. If even just one of them changes (as above), the length
78     % reduced by the radius of curvature
79     correction = round(pixCurvature/2);
80     else
81         correction = 0;
82     end
83     pathAX = startCoordVox(1):fac:(endCoordVox(1)-fac*correction);
84     pathAY = repmat(startCoordVox(2),1,length(pathAX));
85     pathAZ = repmat(startCoordVox(3),1,length(pathAX));
86 end
87
88 if startCoordVox(2) == endCoordVox(2) % if start and endcoordinate of y is the
89     ↪ same,
90     % the 'y' part of the path will not exist
91     pathBX = [];
92     pathBY = [];
93     pathBZ = [];
94 else
95     if startCoordVox(2) < endCoordVox(2)
96         fac = 1;
97     else
98         fac = -1;
99     end
100    if startCoordVox(1) == endCoordVox(1) && startCoordVox(3) == endCoordVox(3) %
101        % four cases - x start = x end, z start = z end, xstart = xstart and
102        % ↪ zstart = zstart and both not equal
103        correction1 = 0;
104        correction2 = 0;
105    elseif startCoordVox(1) == endCoordVox(1) && startCoordVox(3) ~= endCoordVox
106        ↪ (3)
107        correction1 = 0;
108        correction2 = round(pixCurvature/2);
109    elseif startCoordVox(1) ~= endCoordVox(1) && startCoordVox(3) == endCoordVox
110        ↪ (3)
111        correction1 = round(pixCurvature/2);
112        correction2 = 0;
113    elseif startCoordVox(1) ~= endCoordVox(1) && startCoordVox(3) ~= endCoordVox
114        ↪ (3)
115        correction1 = round(pixCurvature/2);
116        correction2 = round(pixCurvature/2);
117    end
118    pathBY = (startCoordVox(2)+round(fac*correction1*anisotropyCorrection)):fac:(
119        ↪ endCoordVox(2)-round(fac*correction2*anisotropyCorrection));
120    pathBX = repmat(endCoordVox(1),1,length(pathBY));
121    pathBZ = repmat(startCoordVox(3),1,length(pathBY));
122 end
123
124 if startCoordVox(3) == endCoordVox(3) % if start and endcoordinate of y is the
125     ↪ same,

```

```

120 % the 'z' part of the path will not exist
121 pathCX = [];
122 pathCY = [];
123 pathCZ = [];
124 else
125     if startCoordVox(3) < endCoordVox(3)
126         fac = 1;
127     else
128         fac = -1;
129     end
130     if startCoordVox(1) ~= endCoordVox(1) || startCoordVox(2) ~= endCoordVox(2)
131         correction1 = round(pixCurvature/2);
132     else
133         correction1 = 0;
134     end
135     pathCZ = (startCoordVox(3)+fac*correction1):fac:(endCoordVox(3));
136     pathCX = repmat(endCoordVox(1),1,length(pathCZ));
137     pathCY = repmat(endCoordVox(2),1,length(pathCZ));
138 end
139
140
141 % Add corner rounding
142 % at each bend, the channel should be curved, connecting the two segments
143 % at first, it is made very detailed, and rescaled to the coordinate grid.
144 % Then only unique combinations are kept
145 % depending on the direction of segment 1 and 2, the rotation will be
146 % different
147
148 % note that since the data has already been converted to the voxelgrid
149 % which is anisotropic, there needs to be a correction factor for anything
150 % involving y. - the factor is created above
151
152 t = 1:100;
153
154 if startCoordVox(1) ~= endCoordVox(1) && startCoordVox(2) ~= endCoordVox(2) % if
    ↪ there is both an x and y segment, they need to be joined by a curved
    ↪ segment
155     if startCoordVox(1) < endCoordVox(1) && startCoordVox(2) < endCoordVox(2)
156         curveAX = round(pathAX(end)+round(pixCurvature/2)*cos(t/100*pi/2-pi/2));
157         curveAY = round(pathAY(end)+round(pixCurvature/2*anisotropyCorrection)*(1+
    ↪ sin(t/100*pi/2-pi/2)));
158         curveAZ = repmat(pathAZ(end),1,length(t));
159     elseif startCoordVox(1) < endCoordVox(1) && startCoordVox(2) > endCoordVox(2)
160         curveAX = round(pathAX(end)+round(pixCurvature/2)*cos(t/100*pi/2-pi/2));
161         curveAY = round(pathAY(end)-round(pixCurvature/2*anisotropyCorrection)*(1+
    ↪ sin(t/100*pi/2-pi/2)));
162         curveAZ = repmat(pathAZ(end),1,length(t));
163     elseif startCoordVox(1) > endCoordVox(1) && startCoordVox(2) < endCoordVox(2)
164         curveAX = round(pathAX(end)-round(pixCurvature/2)*cos(t/100*pi/2-pi/2));
165         curveAY = round(pathAY(end)+round(pixCurvature/2*anisotropyCorrection)*(1+
    ↪ sin(t/100*pi/2-pi/2)));
166         curveAZ = repmat(pathAZ(end),1,length(t));
167     else

```

```

168     curveAX = round(pathAX(end)-round(pixCurvature/2)*cos(t/100*pi/2-pi/2));
169     curveAY = round(pathAY(end)-round(pixCurvature/2*anisotropyCorrection)*(1+
    ↪ sin(t/100*pi/2-pi/2)));
170     curveAZ = repmat(pathAZ(end),1,length(t));
171 end
172 else
173     curveAX = [];
174     curveAY = [];
175     curveAZ = [];
176 end
177
178 if startCoordVox(2) ~= endCoordVox(2) && startCoordVox(3) ~= endCoordVox(3)
179     if startCoordVox(2) < endCoordVox(2) && startCoordVox(3) < endCoordVox(3)
180         curveBX = repmat(pathBX(end),1,length(t));
181         curveBY = round(pathBY(end)+round(pixCurvature/2*anisotropyCorrection)*(
    ↪ cos(t/100*pi/2-pi/2)));
182         curveBZ = round(pathBZ(end)+round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
183     elseif startCoordVox(2) < endCoordVox(2) && startCoordVox(3) > endCoordVox(3)
184         curveBX = repmat(pathBX(end),1,length(t));
185         curveBY = round(pathBY(end)+round(pixCurvature/2*anisotropyCorrection)*(
    ↪ cos(t/100*pi/2-pi/2)));
186         curveBZ = round(pathBZ(end)-round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
187     elseif startCoordVox(2) > endCoordVox(2) && startCoordVox(3) < endCoordVox(3)
188         curveBX = repmat(pathBX(end),1,length(t));
189         curveBY = round(pathBY(end)-round(pixCurvature/2*anisotropyCorrection)*(
    ↪ cos(t/100*pi/2-pi/2)));
190         curveBZ = round(pathBZ(end)+round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
191     else
192         curveBX = repmat(pathBX(end),1,length(t));
193         curveBY = round(pathBY(end)-round(pixCurvature/2*anisotropyCorrection)*(
    ↪ cos(t/100*pi/2-pi/2)));
194         curveBZ = round(pathBZ(end)-round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
195     end
196 else
197     curveBX = [];
198     curveBY = [];
199     curveBZ = [];
200 end
201
202 if startCoordVox(1) ~= endCoordVox(1) && startCoordVox(3) ~= endCoordVox(3) &&
    ↪ startCoordVox(2) == endCoordVox(2)
203     if startCoordVox(1) < endCoordVox(1) && startCoordVox(3) < endCoordVox(3)
204         curveCX = round(pathAX(end)+round(pixCurvature/2)*(cos(t/100*pi/2-pi/2)));
205         curveCY = repmat(pathAY(end),1,length(t));
206         curveCZ = round(pathAZ(end)+round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
207     elseif startCoordVox(1) < endCoordVox(1) && startCoordVox(3) > endCoordVox(3)
208         curveCX = round(pathAX(end)+round(pixCurvature/2)*(cos(t/100*pi/2-pi/2)));
209         curveCY = repmat(pathAY(end),1,length(t));

```

```

210     curveCZ = round(pathAZ(end)-round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
211 elseif startCoordVox(1) > endCoordVox(1) && startCoordVox(3) < endCoordVox(3)
212     curveCX = round(pathAX(end)-round(pixCurvature/2)*(cos(t/100*pi/2-pi/2)));
213     curveCY = repmat(pathAY(end),1,length(t));
214     curveCZ = round(pathAZ(end)+round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
215 else
216     curveCX = round(pathAX(end)-round(pixCurvature/2)*(cos(t/100*pi/2-pi/2)));
217     curveCY = repmat(pathAY(end),1,length(t));
218     curveCZ = round(pathAZ(end)-round(pixCurvature/2)*(1+sin(t/100*pi/2-pi/2))
    ↪ );
219 end
220 else
221     curveCX = [];
222     curveCY = [];
223     curveCZ = [];
224 end
225
226
227 combine = [pathAX curveAX pathBX curveBX curveCX pathCX ; ...
228           pathAY curveAY pathBY curveBY curveCY pathCY; pathAZ curveAZ pathBZ curveBZ
    ↪ curveCZ pathCZ];
229
230 path = unique(combine','rows','stable');
231
232
233
234 for l = 1:length(path)
235     for i = (path(1,1)-round(pixDiameter/2)):(path(1,1)+round(pixDiameter/2))
236         for j = (path(1,2)-round(pixDiameter/2)):(path(1,2)+round(pixDiameter/2))
237             for k = (path(1,3)-round(pixDiameter/2)):(path(1,3)+round(pixDiameter
    ↪ /2))
238                 if i <= 0 || j <= 0 || k <= 0 || i > xDim || j > yDim || k > zDim
239                     %
240                 else
241                     if sqrt((i-path(1,1))^2+((j-path(1,2))/anisotropyCorrection)
    ↪ ^2+(k-path(1,3))^2) < pixDiameter/2
242                         % note anisotropy correction again
243                         phantom(k,i,1,j) = 0;
244                     end
245                 end
246             end
247         end
248     end
249 end
250
251
252 end

```

G.1.2 Square channels

```

1 function phantom = axilinearCylinderAndSphereChannel(phantom, diameter,
    ↪ startCoordinate, endCoordinate)
2
3
4
5
6 % first, convert input diameter to pixel values
7 pixelSize = 10.8;
8 pixDiameter = diameter/pixelSize;
9
10 % determine size of the phantom matrix in order to ensure nothing is done
11 % for negative matrix indices or indices beyond the original phantom size
12 [zDim, xDim, ~, yDim] = size(phantom);
13
14
15 %first channel along x
16 for j = (startCoordinate(2)-round(pixDiameter/2)):(startCoordinate(2)+round(
    ↪ pixDiameter/2))
17     for k = (startCoordinate(3)-round(pixDiameter/2)):(startCoordinate(3)+round(
        ↪ pixDiameter/2))
18         if j <= 0 || k <= 0 || j > yDim || k > zDim
19             %
20         else
21             if sqrt((j-startCoordinate(2))^2+(k-startCoordinate(3))^2) <
                ↪ pixDiameter/2
22                 phantom(k,startCoordinate(1):endCoordinate(1),1,j) = 0;
23             end
24         end
25     end
26 end
27
28 % then channel along y
29 for i = (endCoordinate(1)-round(pixDiameter/2)):(endCoordinate(1)+round(
    ↪ pixDiameter/2))
30     for k = (startCoordinate(3)-round(pixDiameter/2)):(startCoordinate(3)+round(
        ↪ pixDiameter/2))
31         if i <= 0 || k <= 0 || i > xDim || k > zDim
32             %
33         else
34             if sqrt((i-endCoordinate(1))^2+(k-startCoordinate(3))^2) < pixDiameter
                ↪ /2
35                 phantom(k,i,1,startCoordinate(2):endCoordinate(2)) = 0;
36             end
37         end
38     end
39 end
40
41 % then channel along z
42 for i = (endCoordinate(1)-round(pixDiameter/2)):(endCoordinate(1)+round(
    ↪ pixDiameter/2))

```

```

43     for j = (endCoordinate(2)-round(pixDiameter/2):(endCoordinate(2)+round(
44         ↪ pixDiameter/2))
45         if i <= 0 || j <= 0 || i > xDim || j > yDim
46             %
47         else
48             if sqrt((i-endCoordinate(1))^2+(j-endCoordinate(2))^2) < pixDiameter/2
49                 phantom(startCoordinate(3):endCoordinate(3),i,1,j) = 0;
50             end
51         end
52     end
53 end
54
55 % each joint between sections should also be rounded
56
57
58 % round corner between entrance to first section section by distance in 3D
59 for i = (startCoordinate(1)-round(pixDiameter/2):(startCoordinate(1)+round(
60     ↪ pixDiameter/2))
61     for j = (startCoordinate(2)-round(pixDiameter/2):(startCoordinate(2)+round(
62         ↪ pixDiameter/2))
63         for k = (startCoordinate(3)-round(pixDiameter/2):(startCoordinate(3)+
64             ↪ round(pixDiameter/2))
65             if i <= 0 || j <= 0 || k <= 0 || i > xDim || j > yDim || k > zDim
66                 %
67             else
68                 if sqrt((i-startCoordinate(1))^2+(j-startCoordinate(2))^2+(k-
69                     ↪ startCoordinate(3))^2) < pixDiameter/2
70                     phantom(k,i,1,j) = 0;
71                 end
72             end
73         end
74     end
75 end
76
77 % round corner between x and y section by distance in 3D
78 for i = (endCoordinate(1)-round(pixDiameter/2):(endCoordinate(1)+round(
79     ↪ pixDiameter/2))
80     for j = (startCoordinate(2)-round(pixDiameter/2):(startCoordinate(2)+round(
81         ↪ pixDiameter/2))
82         for k = (startCoordinate(3)-round(pixDiameter/2):(startCoordinate(3)+
83             ↪ round(pixDiameter/2))
84             if i <= 0 || j <= 0 || k <= 0 || i > xDim || j > yDim || k > zDim
85                 %
86             else
87                 if sqrt((i-endCoordinate(1))^2+(j-startCoordinate(2))^2+(k-
88                     ↪ startCoordinate(3))^2) < pixDiameter/2
89                     phantom(k,i,1,j) = 0;
90                 end
91             end
92         end
93     end
94 end

```

```

87 end
88
89 % round corner between y and z section by distance in 3D
90 for i = (endCoordinate(1)-round(pixDiameter/2)):(endCoordinate(1)+round(
    ↪ pixDiameter/2))
91     for j = (endCoordinate(2)-round(pixDiameter/2)):(endCoordinate(2)+round(
    ↪ pixDiameter/2))
92         for k = (startCoordinate(3)-round(pixDiameter/2)):(startCoordinate(3)+
    ↪ round(pixDiameter/2))
93             if i <= 0 || j <= 0 || k <= 0 || i > xDim || j > yDim || k > zDim
94                 %
95             else
96                 if sqrt((i-endCoordinate(1))^2+(j-endCoordinate(2))^2+(k-
    ↪ startCoordinate(3))^2) < pixDiameter/2
97                     phantom(k,i,1,j) = 0;
98                 end
99             end
100         end
101     end
102 end
103
104 % round corner at end of z section by distance in 3D
105 for i = (endCoordinate(1)-round(pixDiameter/2)):(endCoordinate(1)+round(
    ↪ pixDiameter/2))
106     for j = (endCoordinate(2)-round(pixDiameter/2)):(endCoordinate(2)+round(
    ↪ pixDiameter/2))
107         for k = (endCoordinate(3)-round(pixDiameter/2)):(endCoordinate(3)+round(
    ↪ pixDiameter/2))
108             if i <= 0 || j <= 0 || k <= 0 || i > xDim || j > yDim || k > zDim
109                 %
110             else
111                 if sqrt((i-endCoordinate(1))^2+(j-endCoordinate(2))^2+(k-
    ↪ endCoordinate(3))^2) < pixDiameter/2
112                     phantom(k,i,1,j) = 0;
113                 end
114             end
115         end
116     end
117 end
118
119
120 end

```


H.1 Hand-bonding recess

The handbonding recess model was created in Autodesk Inventor and can be seen in Figure H.1. It was made to aid in supplying an initial pre-bond pressure in the middle of the wafer stack. The centre of the recess is slightly elevated compared to the boundary, only allowing for pressure to be applied in the middle.

H.2 Pressure chamber

The pressure chamber model was created in Autodesk Inventor and can be seen in Figure H.2. A recess for a 4" wafer is milled in the centre. The pressure chamber lid is fastened using bolts and nuts, and a small rim fitted for an o-ring is made to seal the chamber. Two tube-connectors are placed at the top of the lid, to allow for controlled flushing of the chamber.

H.3 3D printed holder for optical measurements of swelling

A small sample holder was designed for mounting of the hydrogel samples for inspection of sample dimensions in optical microscope. The samples should be inspected while submerged in water, since exposure to air will dehydrate the samples which will affect the size of the sample. The hydrogels cannot themselves be glued to the bottom of a petri dish, but it would be possible to glue a polylactic acid (PLA) 3D printed sample holder.

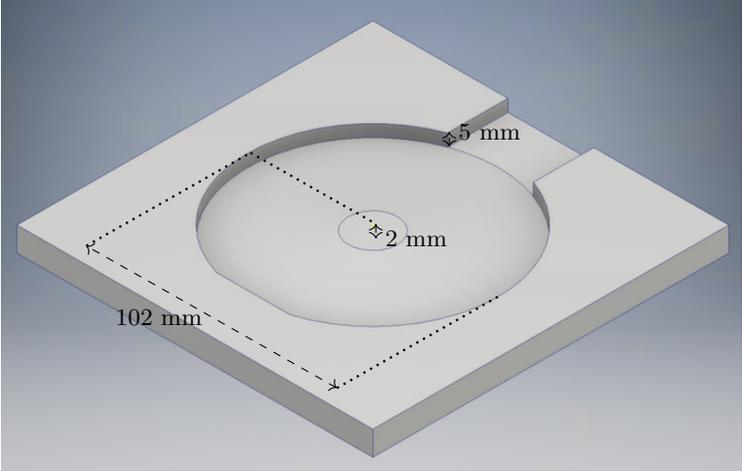


Figure H.1: Isometric view of the handbonding recess 3D model made to aid in supplying an initial pre-bond pressure in the middle of the wafer stack. The centre of the recess is slightly elevated compared to the boundary, only allowing for pressure to be applied in the middle.

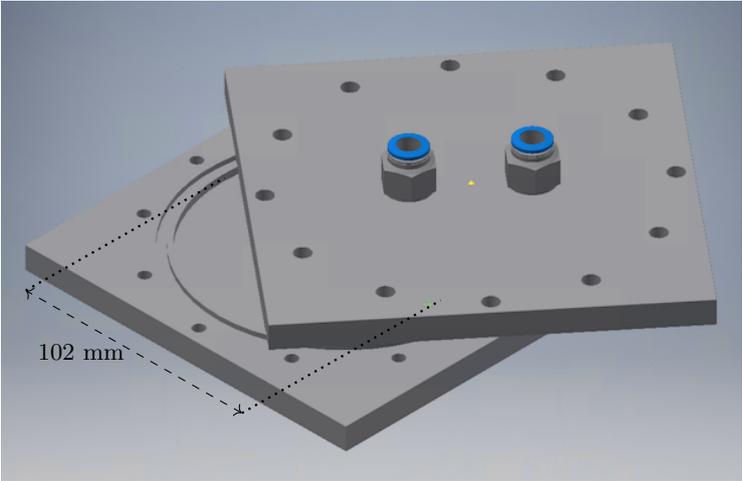


Figure H.2: Isometric view of the pressure chamber 3D model. A recess for a 4" wafer is milled in the centre. The pressure chamber lid is fastened using bolts and nuts, and a small rim fitted for an o-ring is made to seal the chamber. Two tube-connectors are placed at the top of the lid, to allow for controlled flushing of the chamber.

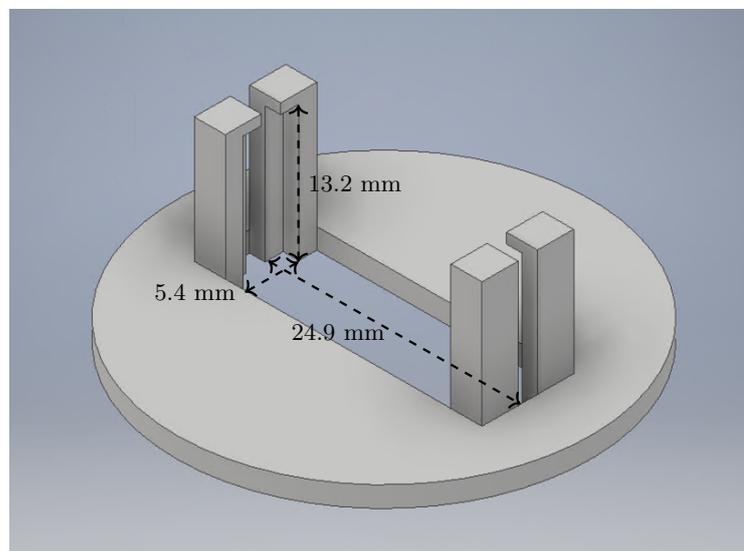
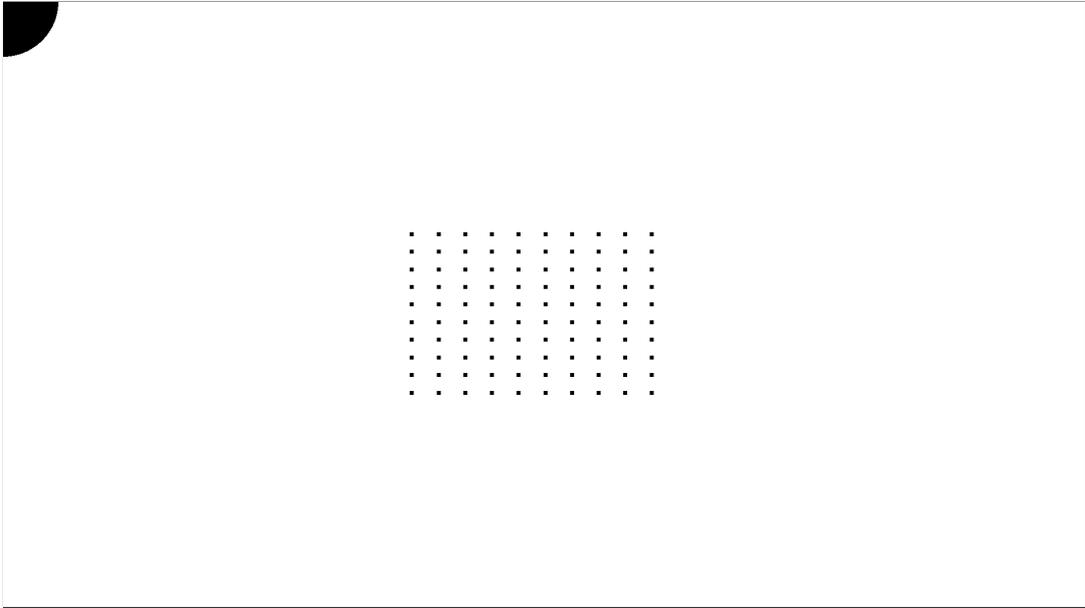


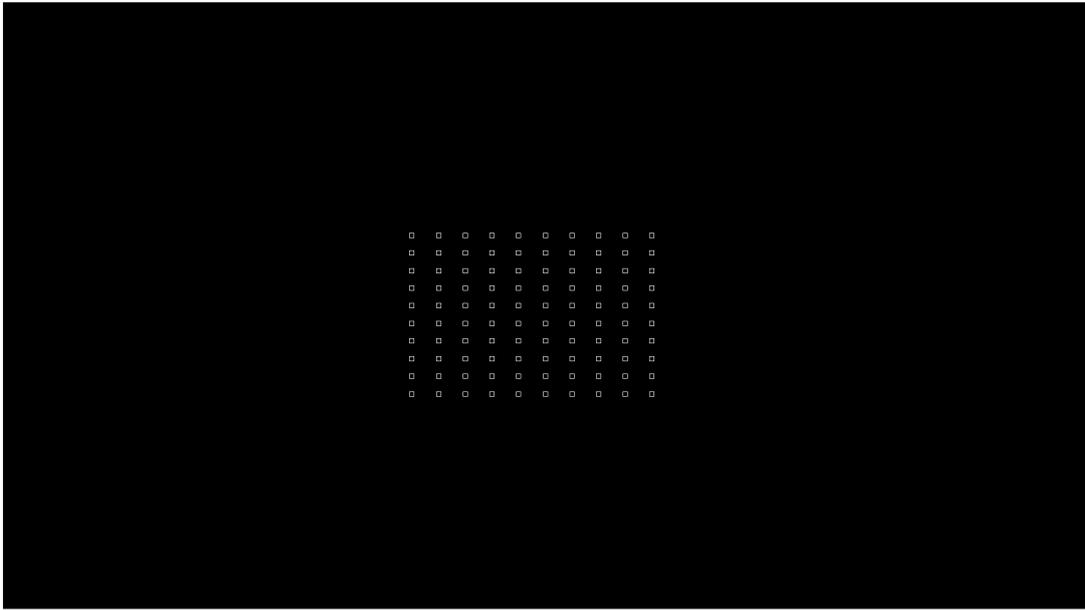
Figure H.3: Hydrogel sample holder model, created to fixate samples while submerged in water for inspection of sample dimensions under water.

I.1 Scatterer phantom for neural network testing

The scatterer phantom designed for testing of a neural network can be seen in Figure I.1, in which Figure I.1(a) is the base exposure, and Figure I.1(b) is the single pixel overexposure around each scatterer. The figures show the 10 by 10 grid of scatterers designed for neural network scatterer detection. Each scatterer is designed to be 7 by 7 voxels, which results in actual printed dimensions of 55.3 μm by 55.3 μm according to the scatterer size statistical model.



(a) Base exposure



(b) Single pixel exposure

Figure I.1: 10 by 10 grid of scatterers designed for neural network scatterer detection. Each scatterer is designed to be 7 by 7 voxels, which results in actual printed dimensions of 55.3 μm by 55.3 μm according to the scatterer size statistical model.

3D printing of hydrogels

J.1 True PEGDA concentration

The actual concentration of poly(ethylene glycol) diacrylate (PEGDA) might be slightly higher than 20%. Figure J.1 shows a comparison of two resin constituent solutions, with PEGDA on the left, and lithium phenyl-2,4,6-trimethylbenzoylphosphinate (LAP) on the right. Both are supposed to be 5 ml in total volume. However, the PEGDA volume is smaller. The solution is made from 4 g PEGDA and 1.4 ml of water. It seems that water and PEGDA are able to pack slightly closer than the constituents. This difference means that the actual PEGDA concentration will be slightly higher in the mixed solution. Based on measurements directly on the shown image, the final concentration of PEGDA ends up being 21% instead of 20%. If one wants a true 20% solution, it will be necessary to determine how much water needs to be mixed into the PEGDA solution, before the combined solution reaches 5 ml.

J.2 Optical microscope images of scatterers

Additional optical microscope images of each scatterer size and type is shown in Figure J.2. The images have been cropped to enlarge the scatterers. However, the dose gradient of the dose gradient scatterers is no longer visible.

J.3 Additional process optimisation and printing issues

J.3.1 Stress-induced bending of hydrogel samples

The phantoms swell after being printed. However, the bottom part of the phantoms which are printed on the cover glass cannot swell freely since they are fixated laterally on the cover glass. The initial thought was that the stress induced bending would be a temporary state, with the stress being released as soon as the phantom was removed from the cover glass. However, if enough time passes, it appears that the stress induced bending might be frozen into the printed structure. This is seen in Figure J.3. Figure J.3(a) is a phantom which has been removed from the cover glass shortly after the print was finished. Figure J.3(b) is a phantom which was kept on the cover glass for more than a month. The second image was taken a week after the phantom had been removed from the cover glass. It did not recover. The phantom is permanently deformed.

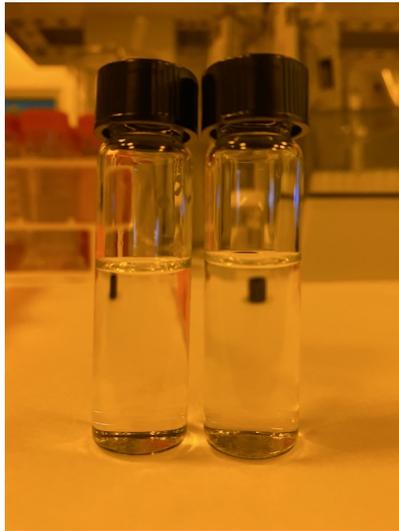


Figure J.1: Two resin constituent solutions with PEGDA on the left, and LAP on the right. Although both are supposed to result in a volume of 5 ml, the PEGDA appears to be higher, resulting in a lower total volume.

J.3.2 Ghost image

An example of the ghost image was illustrated in Section 6.3 for the 100% PEGDA sample. In the main text, it was only observed to a lesser extent, with the intended separation between the two parts of the phantom being partly polymerised. A more severe case was also seen. This is shown in Figure J.4.

J.3.3 Minimum scatterer separation

The minimum separation study in Section 7.2.5 showed that it is possible to place scatterers only 1 voxel apart, seemingly without issue. However, the scatterer shapes seemed to change depending on where in the field of view of the printer it was. Two different groups of scatterers separated by 1 voxel can be seen in Figure J.5. The printed shapes of the scatterers are different even though they are designed to be the same. This is not simply a matter of the placement in the printer field of view, since the shapes even differ within each group of scatterers.

J.3.4 Scratches in the film

A transparent film is placed on the transparent bottom of the vat, with a non-stick surface to minimise the hydrogel adhesion to the vat bottom. Over time, the film will inevitably become scratched, and the scratch pattern will influence the printed structures. These structures are visible to the naked eye all the way through the print, and is the likely cause for the noisy static background seen in ultrasound images of hydrogel samples. The scratches are directly visible on the surface of the phantoms, and can be seen in the line patterns in previously shown hydrogel microscope images. Figure J.6 shows two different scatterers from two different phantom, printed in the same exact position of the phantom. The same scratch pattern is visible in the top right of both images, since they are printed on top of the same part of the film. The slight difference in position of the scratch pattern relative to the scatterers show the precision of vat placement from print to print.

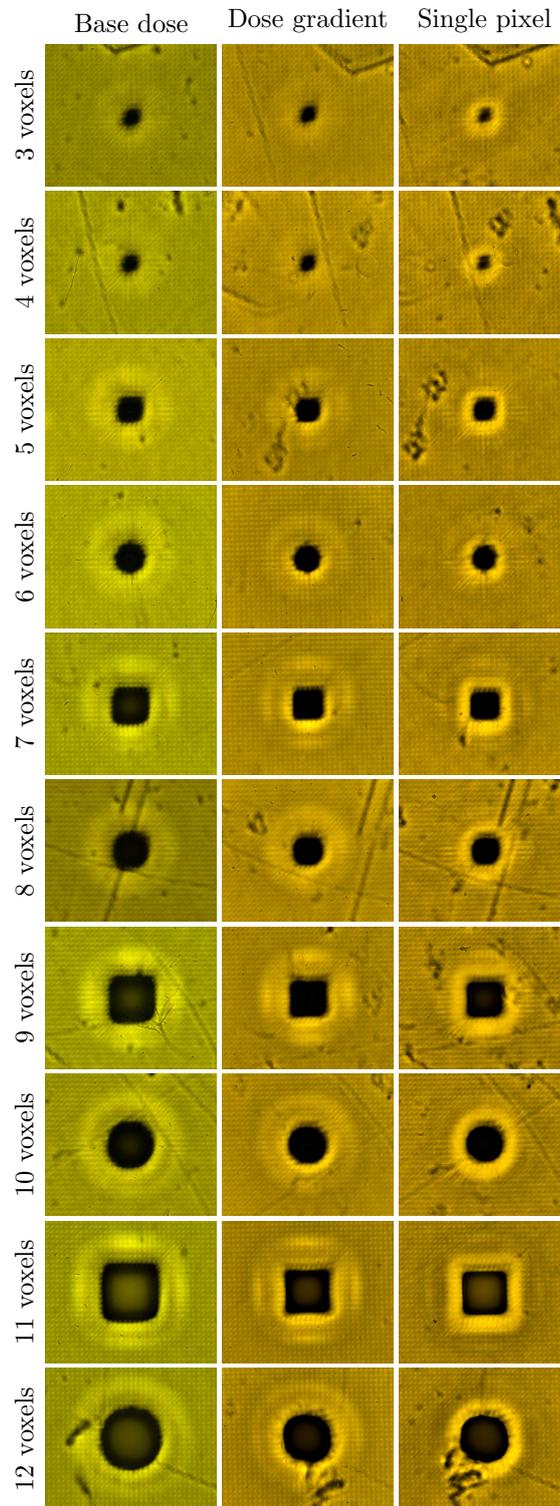


Figure J.2: Examples of the optical microscope images with different sizes and different dosing schemes. Note that the images have been cropped, resulting in the dose gradient of the dose gradient scatterers not being visible as it was in the images in the main text.

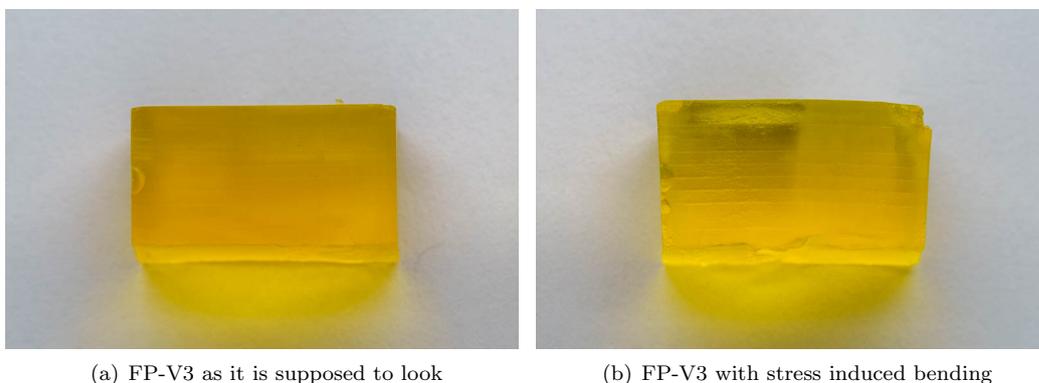


Figure J.3: Phantoms which are not removed from the cover glass might suffer permanent deformation due to the fixation at the cover glass. (a) is a phantom which has been removed from the cover glass shortly after the print was finished. (b) is a phantom which was kept on the cover glass for more than a month. Image was taken a week after it was removed from the cover glass. The phantom is permanently deformed.

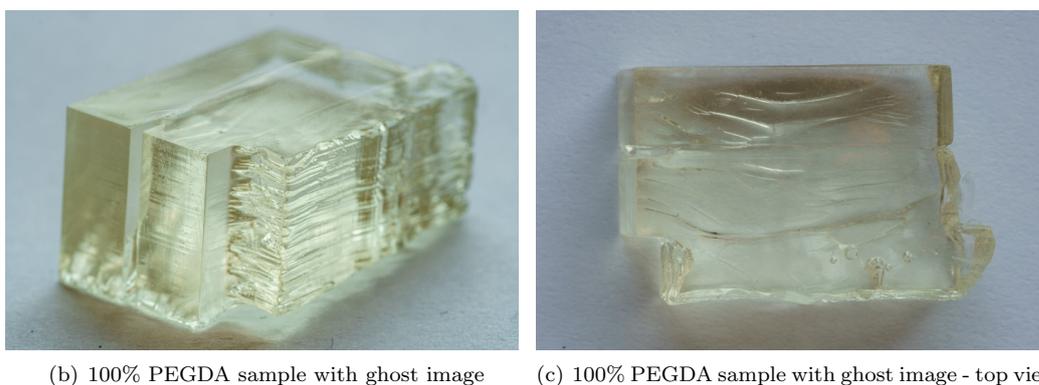
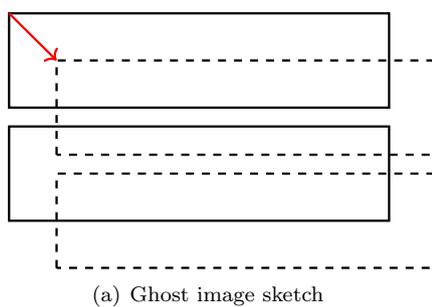


Figure J.4: Ghost image. (a) illustrates the concept, and how the solid line pattern is offset to the dashed line regions. (b) and (c) are images of the most severe sample showing the ghost image.

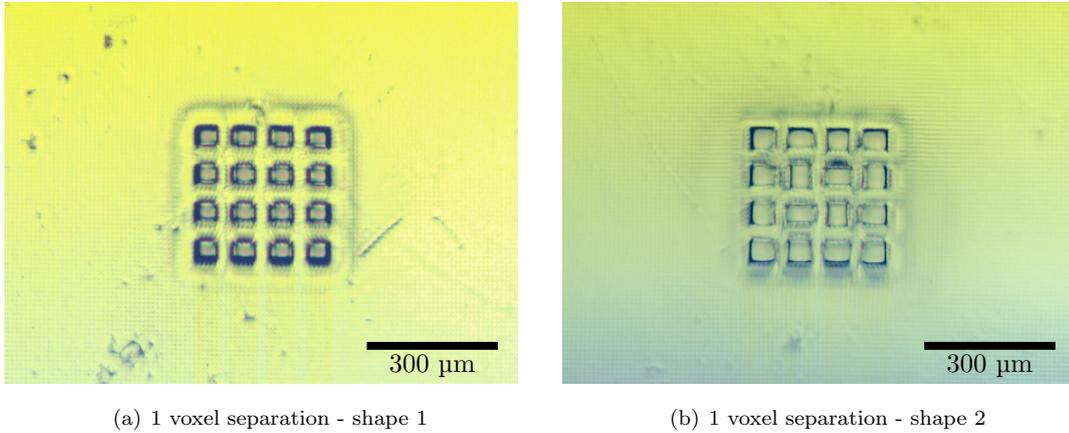


Figure J.5: Images of scatterer separation phantom groups. The voxels in both groups are separated by only a single voxel. The shapes are different, even though they are designed to be similar.

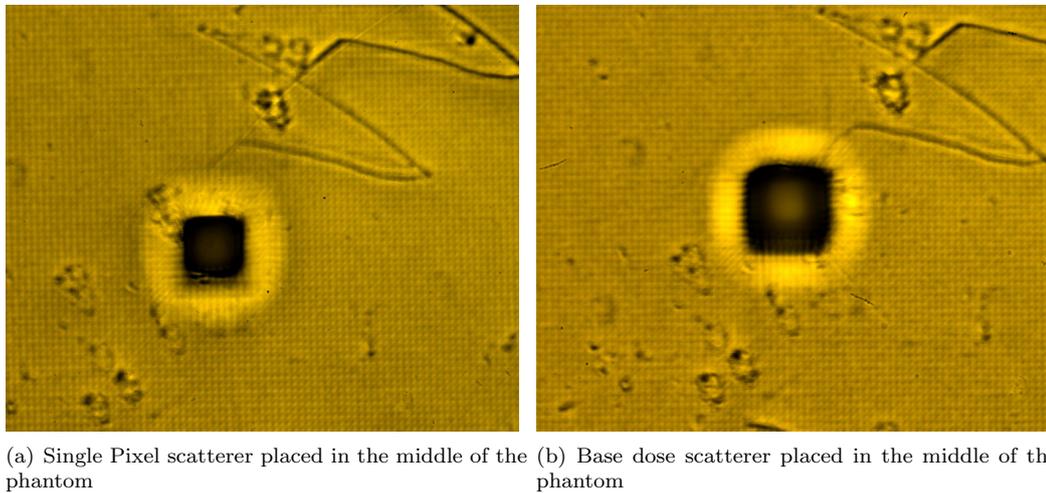


Figure J.6: Two different scatterers from two different phantom, printed in the same exact position of the phantom. The same scratch pattern is visible in the top right of both images, since they are printed on top of the same part of the film. The slight difference in position of the scratch pattern relative to the scatterers show the precision of vat placement from print to print.