DTU Library

Performance of SARUS

A Synthetic Aperture Real-time Ultrasound System

Jensen, Jørgen Arendt; Holten-Lund, Hans; Nielson, Ronnie Thorup; Tomov, Borislav Gueorguiev; Stuart, Matthias Bo; Nikolov, Svetoslav Ivanov; Hansen, Martin; Larsen, Ulrik Darling

Published in:
Proceedings of the 2010 IEEE International Ultrasonics Symposium

Link to article, DOI:
10.1109/ULTSYM.2010.0075

Publication date:
2010

Document Version
Early version, also known as pre-print

Performance of SARUS: A Synthetic Aperture Real-time Ultrasound System

Jørgen Arendt Jensen[1], Hans Holten-Lund[2], Ronnie Thorup Nielson[2], Borislav Georgiev Tomov[1],
Matthias Bo Stuart[1], Svetoslav Ivanov Nikolov[3], Martin Hansen[2] and Ulrik Darling Larsen[2]
[1]Center for Fast Ultrasound Imaging, Dept. of Elec. Eng. Bldg. 349,
Technical University of Denmark, DK-2800 Lyngby, Denmark
[2]Prevas A/S, Frederikskaj 6, DK-2450 Copenhagen SV, Denmark
[3]BK Medical, Mileparken 34, DK-2730 Herlev, Denmark

***Abstract*— The SARUS scanner (Synthetic Aperture Real-time Ultrasound System) for research purposes is described. It can acquire individual channel data for multi-element transducers for a couple of heart beats, and is capable of transmitting any kind of excitation. It houses generous and flexible processing resources that can be reprogrammed and tailored to many kinds of algorithms. The 64 boards in the system house 16 transmit and 16 receive channels each, where data can be stored in 2 GB of RAM and processed using four Virtex 4FX100 and one FX60 FPGAs. The VHDL code can acquire data for 16 channels and perform real-time processing for four channels per board. The receive processing chain consists of three FPGAs. The beamformer FPGA houses 24 focusing units (6 x 4-way) each working in parallel at 220 MHz for parallel four-channel beamforming. The fully parametric focusing unit calculates delays and apodization values in real time in 3D space and can produce 630 million complex samples per second. The processing can, thus, beamform 192 image lines consisting of 1024 complex samples for each emission at a rate of 3200 frames a second yielding full non-recursive synthetic aperture B-mode imaging at more than 30 high resolution images a second.**

I. Introduction

Modern ultrasound scanners use multi-element transducer and advanced beamforming for obtaining a high image quality. This entails measuring and processing RF signals from 64 to 128 channels sampled at 40-70 MHz and an even larger number of channels is used for 3D ultrasound imaging. The data rates are, thus, on the order of 5 to 18 GBytes per second, which have to be processed in real time for 2D imaging. This necessitates dedicated hardware for keeping the power consumption at a reasonable level and makes it difficult, if not impossible, to acquire real-time *in-vivo* data for research purposes. Such data are needed for further developing advanced beamforming strategies like synthetic aperture (SA) imaging, adaptive beamforming, and vector flow imaging. There is, thus, a real need for devising systems, where the stringent demands for real-time processing do not preclude obtaining complete, multi-channel data sets.

A number of systems for acquiring research data have been developed. The RASMUS system [1] has 128 independent transmit channels and 64 independent receive channels, each equipped with a 2-to-1 multiplexor. The system can store more than three seconds of multi-channel data. It has limited capabilities for real-time processing, but can implement real-time conventional beamforming. Storage of real-time *in-vivo* data has been instrumental in developing vector flow imaging, SA imaging, and in conducting pre-clinical studies of this [2]. The system by Lu et al. [3] can handle 128 channels and can perform real-time imaging, but has limited capability for storing long data sequences. A number of other systems have been developed and some are described in a special UFFC issue on ultrasound systems [4].

None of the above systems can handle 3D imaging and real-time processing of SA data or advanced vector flow imaging. The construction of a new system was therefore initiated in 2005 [5] to be able to handle advanced imaging concepts. This paper describes the resulting system and its performance. The system demands are described in Section II. The architecture and programming model are described in Sections III and IV. The real-time capabilities and performance are revealed in Sections V and VI.

II. Specifications and features

The system must be capable of acquiring RF sampled data for all different imaging modes possible to investigate with the system. This includes conventional 2D and 3D imaging as well as SA, coded, and vector flow imaging. The system should be capable of acquiring data for at least a couple of heart beats for off-line processing as well as real-time processing and image display for orientation.

Sampling: High-end transducers reach a center frequency of up to 15 MHz and 100% bandwidth and cMUT probes can go beyond this. A sampling frequency beyond $f_s = 2(f_0 + f_0/2) >$ 45 MHz should be used. The system should be able to use 2D arrays with 32×32 elements for 3D imaging. It should be possible to store data for at least a couple of heart beats for each element to later experiment with off-line beamforming and processing. The accuracy of the hardware should support generating focused fields with side-lobes less than -80 dB.

The Synthetic Aperture Real-time Ultrasound System (SARUS) can sample RF data with a sampling frequency of 70 MHz with a precision of 12 bits. Summing 64 channels this would give a possible signal-to-noise ratio of up to $72 + 6 \cdot 3 =$ 90 dB leaving some margin for an 80 dB dynamic range. The

70 MHz sampling makes it possible to experiment with 25 MHz cMUT probes.

Transmission: The system must be connected to traditional arrays, and here waveforms up to ±100 volts are used in transmission. Often coded imaging is needed for experimentation and it should be possible to transmit arbitrary waveforms on individual channels with a new waveform for each element and transmission. Side-lobes in ultrasound fields are often below -60 dB and the dynamic range of the waveforms should correspond to this. The emitted field must be focused and a delay should be applied to the individual element waveforms.

The SARUS transmission stages consist of near-linear power amplifiers working at up to ±100 volts fed by 70 MHz, 12-bit digital-to-analog converters. The DAC block is connected to a large dynamic storage from where a new waveform can be selected for each transmission and for each channel.

Processing: A vital point in obtaining high quality *in-vivo* data for later experiments with processing is to have a good orientation image during the scan. The processing should be so flexible that conventional as well as experimental frames can be interleaved and acquired at the same time.

The processing engines in SARUS are based on Xilinx Virtex-4 FPGAs. They are reprogrammable and house generous processing resources. With these it has been possible to implement the full processing chain for real-time SA imaging as described in Section V.

Programming: It is important that a research system is flexible to use, and that new and unforeseen imaging modes can be implemented quickly. It is also a requirement that as many of the details in setting up the system are hidden from the user, so they can concentrate on imaging aspects.

The solution in SARUS is to use Matlab as the front-end for the system, where a set of high-end commands are used for controlling the system. Around 50 lines of code suffice for implementing e.g. phased array imaging. This makes testing of new imaging ideas fast and fairly easy to debug.

III. System Architecture

The main blocks of a SARUS digital board are shown in Fig. 1. All functional units are based on a Virtex-4 FPGA from Xilinx that all are connected to both a large (1-2 Gbytes) dynamic RAM and a smaller DRAM of 64 Mbytes. The transmission part consists of one FPGA (2) connected to 16 digital-to-analog converters operating at 70 MHz and 12 bits. All can emit a different waveform for each emission and all waveforms are stored in the DRAM. FPGA (1) controls the sampling of data from the 16 analog-to-digital converters working at 70 MHz and 12 bits. It can selectively store data for the channels and also houses the processing of received channels. The two remaining FPGAs (3, 4) are used for the processing of data as described in Section V. All the FPGAs are connected through high-speed 3.2 Gbit/s Rocket IO links. In the processing chain four links are used both between FPGAs and between boards for a combined data rate of more than 10 Gbits/s. One Rocket IO link is also routed to all

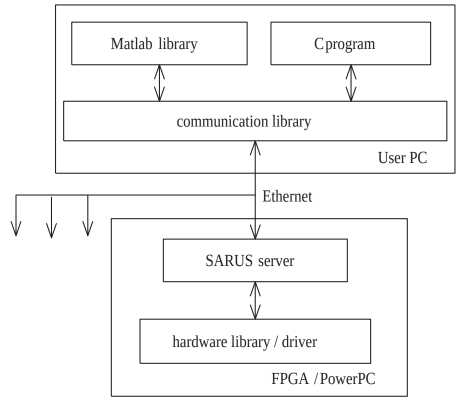

Fig. 2. Structure of the SARUS software.

FPGAs from the central controlling FPGA (5) for sending and receiving parameters for the set-up of the system and reading of RAM content.

The transmission sequence is controlled from the PowerPC FPGA (5). The sequence can consist of several different image types with different lengths, active channels, and waveforms, and the time between pulse emissions T_{prf} can vary from emission to emission. It is also possible to have rapid switching between imaging modes to have sequences for orientation and ones where data are stored only. The sequencer also controls the sampling of data and here the channels to sample and depth range can be controlled individually for each channel and emission. The FPGA (5) is also connected to a 1 Gbit Ethernet network and houses a PowerPC running Linux. This is used for controlling the system and setting parameters. The network is also used for reading data from the RAM of the other FPGAs to store e.g. RF data on a remote Linux storage cluster.

IV. Programming Model

The prime reason for a research system is flexibility, and this is achieved by providing the system with a programming language through Matlab. The structure of the SARUS software is shown in Fig. 2.

A communication layer is being executed on the user Linux PC and can be called from either Matlab or a C executable (single executable scanner). That layer is configured by a text file describing the hardware to be used: number and IP address of the SARUS digital boards, timing source board, and channel connectivity to the transducer connectors. The commands to SARUS are sent through the TCP/IP protocol to SARUS servers. These run within Linux on the embedded PowerPC processors in FPGA (5) on each digital board. The result from the execution - an error code and in some cases parameters or data - are returned to the user PC and passed back to Matlab or the executable.

The user accessible logic resources on a SARUS digital board - register sets and RAM blocks - are mapped into resources using an XML file. The resources are joined into

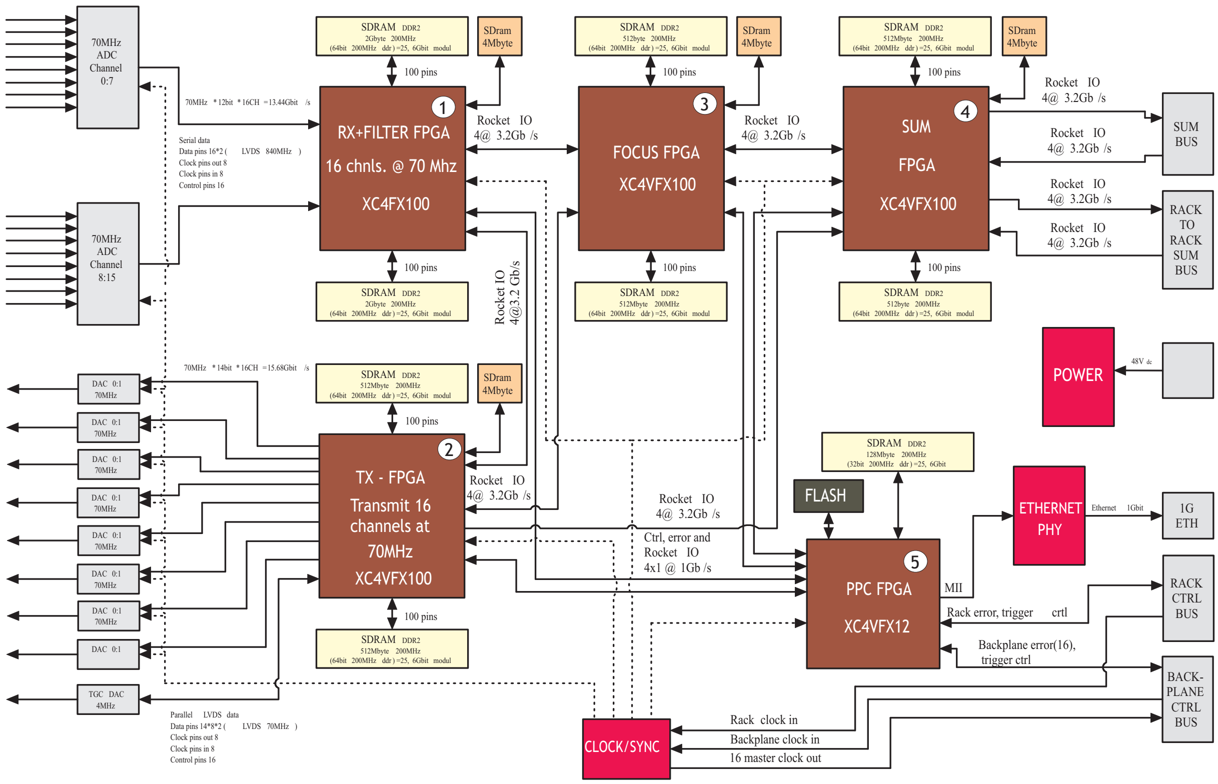

Fig. 1. Block diagram of the DAUP board.

functional units. An FPGA contains several of these units. The passing of parameters and data is done using a serial bus with a tree topology, as shown in Fig. 3. The originating command/data blocks propagate on the bus through switches, in which the first byte of the command block is stripped and used as address (multiplexer position), thus determining which branch of a switch will be used for the further propagation of the command and the returning data.

The user programming is done entirely in Matlab and executed on the PC connected to the net. The Matlab commands operate at a high level, and the user is only concerned about the imaging situation such as selecting the virtual focusing points, which emission codes to use and how to process the data. The C part of the code calculates the relevant parameters and sets them in the hardware. After an image sequence has been set-up, it can be executed either in an infinite loop with a real-time display of the image, or it can be stopped and the data stored on disk. This happens in parallel for all 64 boards in the system through the 1 Gbit/s Ethernet ports, which are connected through fast switches to a Linux storage cluster connected to fast RAID storage. The storage server can later be used for processing the data. Another interesting feature of the system is that data can be kept in the system and then further processed by it. Through this it is easy to change the parametric beamformer or the matched filter and experiment with alternative processing schemes. This is attractive since the processing power of the system roughly corresponds to

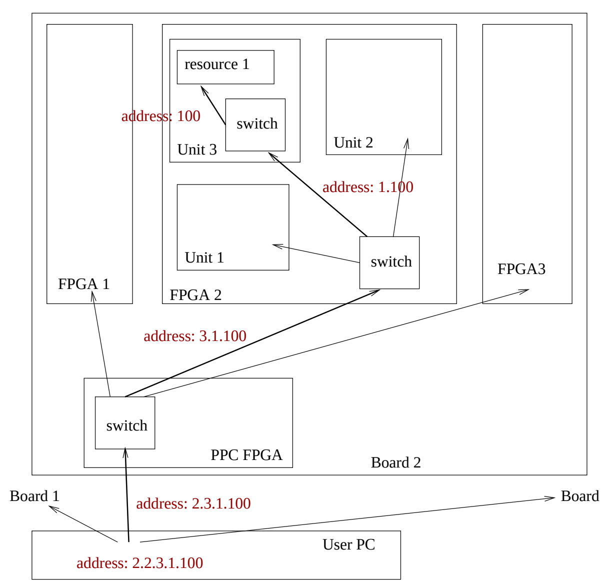

Fig. 3. Communication of a parameters and data in SARUS.

5000 PCs.

V. Real-time VHDL software

The processing chain is shown in Fig. 4. Each of the 64 boards house 16 transmit and 16 receive channels, where data can be stored in 2 GB of RAM and processed using four

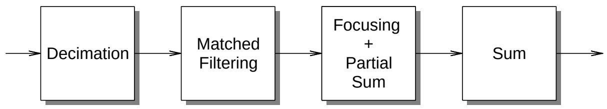

Fig. 4. Steps in the real-time processing in SARUS.

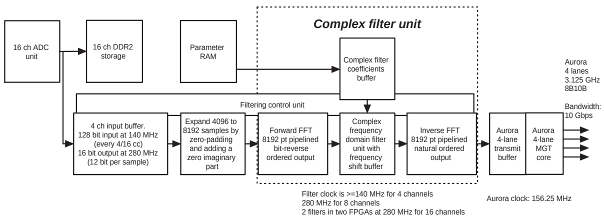

Fig. 5. Overview of the Fourier based matched filter.

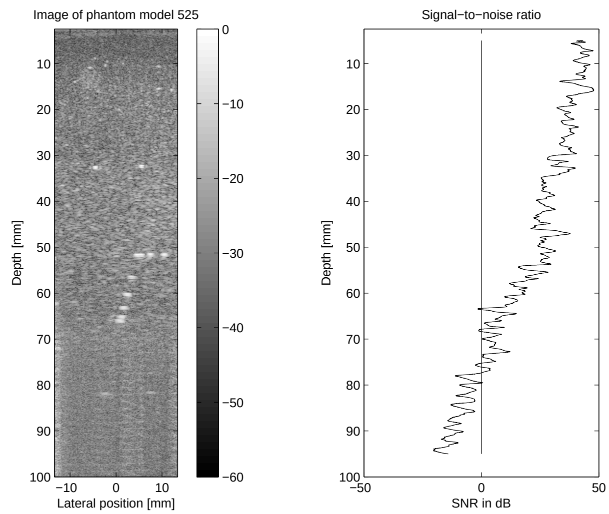

Fig. 6. Linear array image image created by the SARUS scanner using a BK Medical 8804 7 MHz transducer (left). Corresponding signal-to-noise ratio as a function of depth in a phantom with an attenuation of 0.5 dB/[MHz cm] (right).

Virtex 4FX100 FPGAs. One of the FPGAs is used for transmit focusing and sending out arbitrary coded waveforms on the individual channels. The current incarnation of the VHDL code acquires data for 16 channels and performs real-time processing for four channels per board.

The processing chain consists of three FPGAs. The first acquires, decimates, and stores data in RAM. It houses 2 x 8192-point FFT-cores (forward and inverse) for arbitrary matched filtration on four channels at 140 MHz as shown in Fig. 5. 24 DSP slices are used for each FFT (each FPGA houses 160 DSP slices). The time to process four channels of 8192 samples each is 32,768 clock cycles allowing a repetition frequency of 4272 Hz. The unit can handle complex filters and shifts in the frequency domain.

The next FPGA contains a fully parametric focusing unit that calculates delays and apodization values in real time in 3D space [6]. The FPGA houses 24 focusing units (6 x 4-way) each working in parallel at 220 MHz for parallel four-channel beamforming. For each channel it can produce 630 million complex samples per second for either traditional, SA, or 3D beamforming in real time. The output samples are subsequently summed in the last FPGA and transferred over the backplane to be summed with data from the other cards. The processor can beamform 192 image lines consisting of 1024 complex samples for each emission at a rate of 3200 frames a second, yielding full synthetic aperture B-mode imaging at more than 30 high resolution images a second. The focusing units use 150 DSP slices and 36,000 FPGA slices out of the 160 DSP and 42,176 FPGA slices available.

VI. Performance

The performance of the system has been investigated using a 192 element linear array transducer (BK 8804, BK Medical, Herlev, Denmark) operating at 7 MHz. A linear array scan was performed with this array using 64 active elements, a transmit focus at 6 cm, and dynamic receive focusing. No transmit apodization was used and a constant F# number equal to 2 in receive with a von Hann apodization. The resulting image is shown on the left in Fig. 6. The image quality is comparable to the previous research scanner RASMUS [1] and commercial scanners. A penetration depth of approximately 6.5 cm (0 dB SNR) is obtained. The signal-to-noise as a function of depth is shown on the right in Fig. 6. This was calculated by measuring in the same direction 100 times and find the mean value of the response. Subtracting the mean from the responses yields the noise.

The time stability of the system was tested by comparing repeated measurements. Here 100 emissions were made in the same direction and the time delay between emissions was determined through a cross-correlation estimator. The results are shown in Fig. 7. The mean time delay between emissions is -1.41 ps, and the standard deviation of the delays is 48.3 ps. For a 7 MHz velocity estimation system, this deviation will introduce an error of roughly 0.03%, which is well below the accuracy of all current velocity estimation systems. The variation is also random and can therefore easily be averaged out. Alternatively it makes it possible for the system to find shifts in positions between two emissions down to 37 nm. When correlating data from the first emission with subsequent emissions a slight drift of -1.22 ps/emission is seen. This could also be attributed to a small drift in speed of sound due to a change in temperature of the phantom used.

VII. Conclusion

A photo of the SARUS system is shown in Fig. 8. A fully equipped 256 channel system is currently running, and boards for a 1024 system have been fabricated. The system now functions on the same level as a commercial scanner in terms on penetration depth and time accuracy. Its beamformer is 100-200 times faster than traditional commercial scanners. It can also store data at different stages in the processing chain from single channel RF data, focused data to the final beamformed

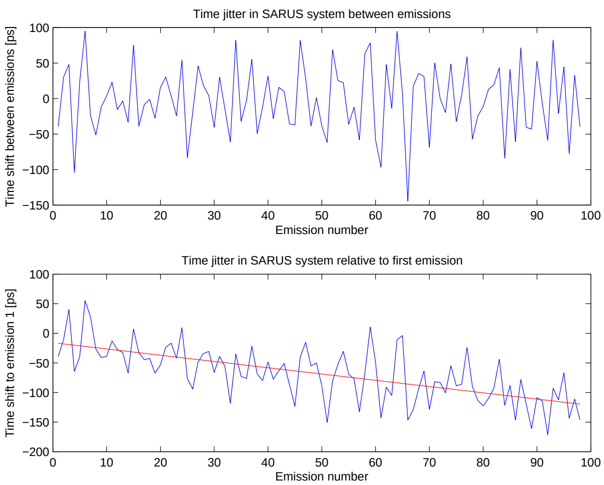

Fig. 7. Time jitter between emissions performed using the same emission. the top graph shown jitter between to adjacent channels and the lower graph shows the delay from the first emission to subsequent emissions.

and summed response, making it ideally suited for research purposes.

ACKNOWLEDGEMENT

This work was supported by grant 9700883, 9700563 and 26-04-0024 from the Danish Science Foundation and by B-K Medical Aps. This project is also supported by grant 024-2008-3 from the Danish Advanced Technology Foundation

REFERENCES

[1] 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.*, 52 (5):881–891, May 2005.

[2] J. A. Jensen, S. Nikolov, K. Løkke Gammelmark, and M. Høgholm Pedersen. Synthetic aperture ultrasound imaging. *Ultrasonics*, 44:e5–e15, 2006.

[3] J. Lu, J. Cheng, and J. Wang. High frame rate imaging system for limited diffraction array beam imaging with square-wave aperture weightings. *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, 53(10):1796–1812, 2006.

[4] P. Tortoli and J. A. Jensen. Introduction to the special issue on novel equipment for ultrasound research. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 53(10):1705–1706, oct 2006.

[5] J. A. Jensen, M. Hansen, B. G. Tomov, S. I. Nikolov, and H. Holten-Lund. System architecture of an experimental synthetic aperture real time ultrasound system. In *Proc. IEEE Ultrason. Symp.*, pages 636–640, Oct. 2007.

[6] S. Nikolov, J. A. Jensen, and B. G. Tomov. Fast parametric beamformer for synthetic aperture imaging. *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, 55(8):1755–1767, 2008.

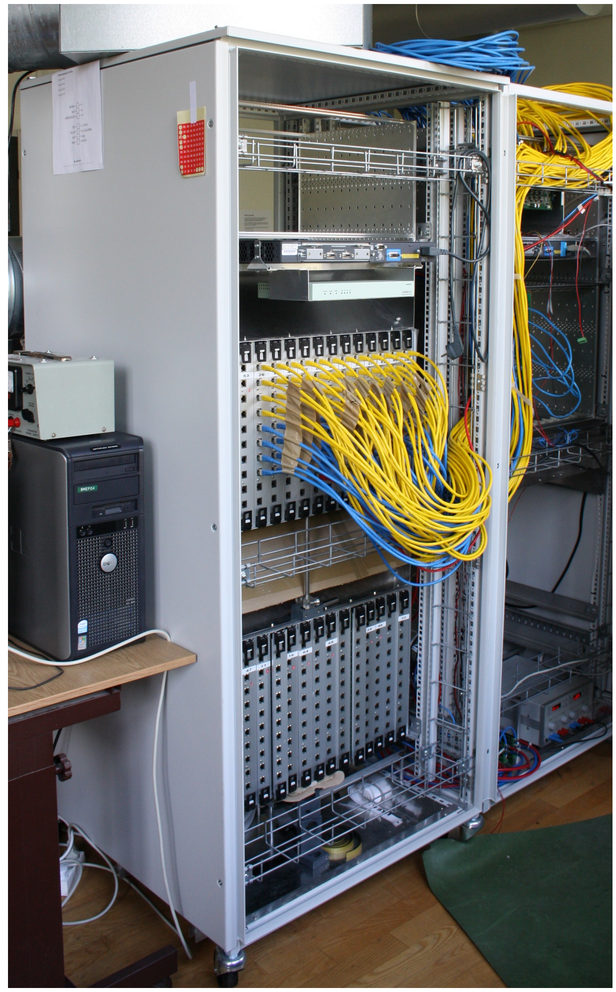

Fig. 8. Photo of SARUS connected for 192 channel scanning. Boards are inserted for a 512 channel system.

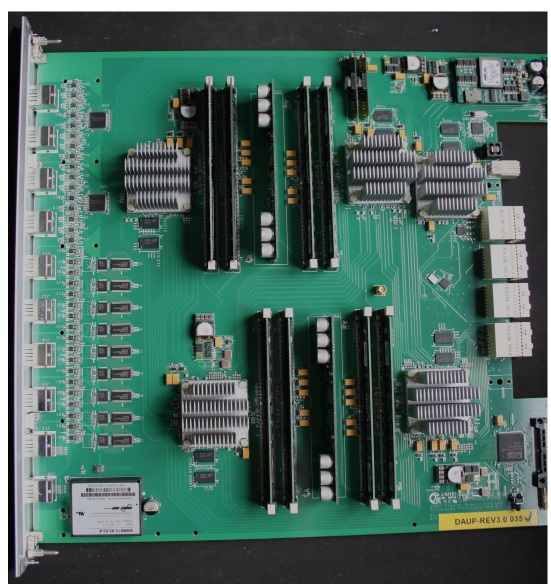

Fig. 9. Photo of daup board. The size of the board is 37 $\times$ 41 cm and it has 20 layers.