

Stochastic Control Theory Optimization of Energy Systems

Brok, Niclas Brabrand

Publication date: 2021

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

Link back to DTU Orbit

Citation (APA): Brok, N. B. (2021). *Stochastic Control Theory Optimization of Energy Systems*. Technical University of Denmark.

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.



Stochastic Control Theory Optimization of Energy Systems

PhD Thesis Niclas Brabrand Brok





Stochastic Control Theory

Optimization of Energy Systems

PhD Thesis 14. of December, 2020

By:	Niclas Brabrand Brok(DTU Compute).			
Main-supervisor:	Henrik Madsen (DTU Compute).			
Co-supervisors:	John Bagterp Jørgensen (DTU Compute), Uffe Høgsbro Thygesen (DTU Compute), Niels Kjølstad Poulsen (DTU Compute.			
Copyright:	Reproduction of this publication in whole or in part must include the customary bibliographic citation, including author attribution, report title, etc.			
Cover photo:	Fie Brabrand Brok, 2020.			
Published by:	DTU, Department of Applied Mathematics and Com- puter Science, Building 303b, 2800 Kgs. Lyn- gby Denmark www.compute.dtu.dk			
ISSN:	[0909-3192] (electronic and printed versions)			

Preface

This PhD thesis was prepared at the Department of Applied Mathematics and Computer Science at the Technical University of Denmark in fulfilment of the requirements for acquiring a PhD degree in Applied Mathematics and Computer Science.



Signature

14th of December 2020

.....

Date

Abstract (in English)

Future low-carbon societies will be driven by renewable energy sources (e.g. wind and solar power). This will flip the characteristics of our power system from a production-tracking-consumption paradigm, to a consumption-tracking-production paradigm. This will increase the need of complex coordination of our power consumption as power grids require a strict balancing between power production and consumption. This dissertation investigates the potential of applying nonlinear model predictive control algorithms to solve complex power market coordination problems, where flexible consumers leverage the more volatile balancing power prices and thereby in-directly help neutralizing production and consumption imbalances.

This dissertation only considers continuous-discrete stochastic models. Paper A provides a tutorial on how to formulate the entire algorithm-stack of nonlinear model predictive control algorithms where system dynamics are governed by stochastic differential equations with discrete-time observations. The techniques introduced in Paper A are applied in four case-studies relating to energy systems. Paper F introduces a new filtering technique that generalizes the observational model to contain general likelihood models.

Paper B and Paper C propose an optimal control problem to operate the aeration equipment at wastewater treatment plants with the criteria to minimize the operational costs and the accumulated nutrient concentrations of the discharged effluent.

Paper D considers the operation of a non-invasive ice-tank module added to a small retail refrigeration system located at Danfoss' test-facility in Nordborg, Denmark. It is shown that the icetank is an efficient method for curtailing the power consumption of the refrigeration system, without re-arranging and modifying the entire piping and general hardware infrastructure.

Paper E presents an optimization technique to optimally leverage the volatile power prices observed in the Northern European regulating power market using a Vanadium redox-flow battery. This paper shows that the payback time of investing in grid-scale flowbatteries are approximately seven years, when implementing this optimization-based trading strategy.

Abstract (in Danish)

Fremtidens samfund uden fossile brændstoffer vil primært være drevet af bæredygtige energikilder som eksempelvis vind og solenergi. Dette vil betyde, at vores elsystem vil gennemgå et paradigmeskift fra et system, hvor elproduktionen følger elforbruget, til et system, hvor elforbruget skal følge den varierende elproduktion. Dette vil øge behovet for en kompleks koordinering af vores elforbrug, i og med at vores elsystem kræver en tæt balancering mellem elproduktion og forbrug. Denne afhandling undersøger potentialet i at bruge ikke-lineær model prædiktiv regulering til at løse kompleks koordinering blandt elforbrugere i vores fælles elmarked. Fleksible forbrugere kan udnytte høj variation i volatile regulerkraftspriser og derigennem indirekte hjælpe elsystemet med at neutralisere ubalancer i vores elsystem.

Denne afhandling er afgrænset til udelukkende at håndtere kontinuertdiskrete stokastiske modeller. Artikel A giver en introduktion til, hvordan man kan formulere hele algoritme-behovet for ikkelineære model prædiktiv reguleringsmetoder, der er drevet af stokastiske differentialligninger med diskrete observationer. Artikel A danner grundlag for de fire anvendelser, der er arbejdet med i denne afhandling med relation til energisystemer. Artikel F introducerer en ny filterteknik, der generaliserer observationsligningen til at kunne håndtere generelle *likelihood* observationsmodeller.

Artikel B og Artikel C foreslår en optimeringsbaseret metode til at regulere beluftningsudstyret hos spildevandsanlæg. Kriteriet for denne metode er at reducere de operationelle omkostninger samtidigt med at reducere den akkumulerede koncentration af næringsstoffer i det udledte spildevand.

Artikel D undersøger, hvordan man kan koordinere det simultane system bestående af en istank og et kølesystem. Denne testopstilling er sat op hos Danfoss i Nordborg, Sønderjylland. Formålet i Artikel D er at bekræfte hypotesen, at det er muligt at aflaste strømforbruget af et kølesystem ved at bruge en ikke-invasiv istankstilkobling og dermed undgå at skulle omlægge rør og ændre den generelle hardware infrastruktur i kølesystemet.

Artikel E præsenterer en optimeringsbaseret teknik, hvormed man kan udnytte de volatile elpriser, der optræder på det nordiske regulerkraftsmarked. Dette studie viser, hvordan man kan bruge et vanadium redox-flow batteri til at skabe et sikkert afkast, hvor man udnytter tidsvarierende elpriser på tværs er forskellige elmarkeder. Artikel E påviser, at tilbagebetalingstiden på batteriet, ved at anvende en optimeringsbaseret metode, er ca. 7 år.

Acknowledgements

First and foremost, I would like to extend my gratitude to Henrik Madsen - my main-supervisor - for giving me the scientific liberty to explore - at times - bold research ideas while also giving me consistent and critical feedback. I would also like to thank John Bagterp Jørgensen for teaching me how nonlinear model predictive control should be done(!) and I would like to thank Uffe Høgsbro Thygesen for giving me valuable insights and intuitions into the complex world of stochastic differential equations. I also want to thank Niels Kjølstad Poulsen for many interesting discussions and meetings with our master thesis students.

I have been extremely fond of the work-place environment at the section for Dynamical Systems at DTU Compute and I would like to thank all my colleagues for making the last 3+ years very enjoyable. I would especially like to say thank you to all my "office friends": Mathias Blicher Bjerregård; Maksim Mazuryn; Frederik Banis and Sebastian Wolf. I am also grateful for the valuable feedback on this dissertation from Rune Grønborg Junker and all the ping-pong I have played in the basement with Phillip Brinck Vetter and Peter Stentoft.

I also want to extend my gratitude to some industrial partners I have met during my studies: Henning Parbo (Energinet), for your incredible knowledge and ability to explain the regulating power market; Jan Hvidberg (Ørsted), for initially turning my attention to this market; Torben Green (Danfoss) and Christian Heerup (Technological Institute), for our joint work on Paper D; and Peter Stentoft (Krüger), for our joint work on Paper B and Paper C.

Finally, I would like to thank Shmuel Oren (from UC Berkeley) for hosting me for 6 months during my PhD studies (just in time before the covid-19 pandemic!).

List of publications

Peer-reviewed and acceptet

- A Niclas Brok, Henrik Madsen and John Bagterp Jørgensen. "Nonlinear Model Predictive Control for Stochastic Differential Equation Systems". 6th IFAC Conference on Nonlinear Model Predictive Control, August 19-22, 2018.
- B Niclas Brok, Thomas Munk-Nielsen, Henrik Madsen and Peter Stentoft. "Flexible Control of Wastewater Aeration for Cost-Efficient and Sustainable Treatment". 2019 IFAC Workshop on Control of Smart Grid and Renewable Energy Systems, June 10-12, 2019.
- C Niclas Brok, Thomas Munk-Nielsen, Henrik Madsen and Peter Stentoft. "Unlocking Energy Flexibility of Municipal Wastewater Aeration using Predictive Control to Exploit Price Differences in Power Markets". Applied Energy Volume 280, December 15, 2020.

Submitted

- D Niclas Brok, Torben Green, Christian Heerup, Shmuel Oren and Henrik Madsen. "Optimal Operation of an Ice-Tank for a Supermarket Refrigeration System". Control Engineering Practice, August 7, 2020.
- E Niclas Brok, Henrik Madsen, Anders Bentien and Rune Grønborg Junker. "CO2 Neutral Energy Arbitraging using Vanadium Redox-Flow Batteries". Applied Energy, December 13, 2020.

In preparations

F Niclas Brok, Uffe Høgsbro Thygensen, Jan Kloppenborg Møller, and Henrik Madsen. "Optimization-Based State Estimation for Stochastic Differential Equation Systems".

Contents

	Preface	ii
	Abstract (English)	iii
	Abstract (Danish)	iv
	Acknowledgements	V
	List of publications	vi
Pa	art I Summary Report	1
1	Introduction	2
	1.1 Stochastic modelling in continuous-time	3
	1.2 Stochastic control theory	3
	1.3 Nonlinear model predictive control	4
	1.4 Applications in energy systems	4
	1.5 Outline	5
2	Stochastic modelling in continuous time	7
	2.1 Stochastic differential equations	7
	2.2 State estimation	9
	2.3 Parameter estimation	18
3	Nonlinear model predictive control	23
	3.1 The model predictive control principle	24
	3.2 Pontryagin's minimum principle	24
	3.3 Direct methods	25
	3.4 Examples	29
4	The Northern European power market	33
	4.1 The day-head market	35
	4.2 The intra-day market	35
	4.3 The regulating power market	36
5	Applications in energy systems	39
	5.1 Optimal operation of wastewater treatment plants	39
	5.2 Optimal operation of a joint ice-tank and refrigeration	
	system	41
	5.3 Optimal energy arbitraging using a Vanadium redox-flow	
	battery	42
6	Summary and future perspectives	45
	6.1 Advances in state and parameter estimation	45
	6.2 Forecasting the regulating power market	46
	6.3 Portfolio optimization of energy units	46
Bi	ibliography	48

Part II Publications

A	Nonlinear Model Predictive Control for Stochastic Differential Equation Systems	55
в	Flexible Control of Wastewater Aeration for Cost-Efficient and Sustainable Treatment	62
с	Unlocking Energy Flexibility of Municipal Wastewater Aeration using Predictive Control to Exploit Price Differences in Power Markets	69
D	Optimal Operation of an Ice-Tank for a Supermarket Refriger- ation System	79
E	CO2 Neutral Energy Arbitraging using Vanadium Redox-Flow Bat- teries	93
F	Optimization-Based State Estimation for Stochastic Differen- tial Equation Systems	.05

Part I

Summary Report

1 Introduction

In 2019, Denmark passed their first ever Climate Bill, committing to reduce greenhouse gas emissions by more than 70 % by 2030 compared to the emission levels of 1990 [1]. This will increase the demand of renewable energy sources (e.g. wind and solar power) in the Danish power system to satisfy the existing power consumption of the power system, but also to saturate the future power demand from e.g. large power-to-x plants.

One of the most important features of our power system is grid balancing. Traditionally, it has been large centralized power plants which have had the role of maintaining coherency between power production and consumption. However, with an increasing share of renewable energy sources - and thereby indirectly a decreasing share of traditional power production from e.g. power plants - the grid balancing responsibility is slowly transitioning from being production-based to being consumption-based. This will increase the importance of coordination between transmission system operators, distribution system operators and consumers. Consumers who are able to shift their power demand in time, are often framed as flexible consumers in the literature [2]; but to unlock this flexibility, consumers have to be incentivised [3, 4]. In Scandinavia (Denmark, Finland, Norway and Sweden), the regulating power market provides a market-based mechanism for consumers to price their flexibility [5] and through this mechanism help the transmission system operator to balance the power system.

This dissertation applies stochastic modelling in continuous-time and nonlinear model predictive control to optimize a range of industrial processes to achieve cost-efficient operation. This is done by leveraging time-varying power prices, from different power markets available in Northern Europe, to monetize the potential flexibility embedded in the industrial processes. The optimization framework is applied to three industrial processes:

- Wastewater treatment: bio-chemical processes, that remove nutrients from the incoming wastewater, can be a very efficient tool for providing short-term flexibility. Paper B and Paper C implement and backtest methods for cost-efficient operation of aeration equipment at wastewater treatment plants.
- Supermarket refrigeration: using an ice-tank it is possible to curtail the power load of large refrigeration systems. Paper D tests the cost-savings potential of adding an ice-tank to a small refrigeration system located at the Danfoss test-facility in Nordborg, Denmark.
- Vanadium redox-flow battery: flow batteries have a long durability and can be used to arbitrage time-varying power prices across multiple power markets and horizons. Paper E presents

an optimization method based on an optimal control problem, which can arbitrage varying power prices across multiple markets.

The publications described above, all use an optimal control formulation and apply the nonlinear model predictive control principle. In Paper A, a tutorial, on how to formulate the entire algorithmstack for nonlinear model predictive control methods, is given. The entire formulation is based upon continuous-discrete systems with a continuous-time stochastic model given by a stochastic differential equation and discrete-time observational model.

The following sections briefly introduce key concepts contained in this dissertation and puts the scientific work in context to current state-of-the-art.

1.1 Stochastic modelling in continuous-time

The starting-point for any dynamical optimization problem is a dynamical model. This dissertation only focuses on stochastic modelling in continuous-time with discrete-time observations. The class of models considered are stochastic differential equations on the form

$$x(t) = x_0 + \int_0^t f(x(\tau), u(\tau); \theta) \, \mathrm{d}\tau + \int_0^t g(x(\tau), u(\tau); \theta) \, \mathrm{d}\omega(\tau), \tag{1.1}$$

where x is the state vector, u is the input vector and θ is the parameter vector. A practical introduction to stochastic differential equations is given in the next chapter and a more theoretical introduction to this topic is given in [6, 7].

1.2 Stochastic control theory

Stochastic control theory provides the ultimate formalism to formulate optimization problems of dynamical system with stochastic disturbances. [8] and [9] define a general stochastic control problem according to

$$\min_{x,u} \quad \mathbb{E}\left[\int_0^T q(x(t), u(t)) \, \mathrm{d}t + q_f(x(T))\right],\tag{1.2a}$$

s.t.

$$x(t) = x_0 + \int_0^t f(x(\tau), u(\tau); \theta) \,\mathrm{d}\tau + \int_0^t g(x(\tau), u(\tau); \theta) \,\mathrm{d}\omega(\tau), \qquad (1.2b)$$

$$u(t) \in \mathcal{U}(t)$$
 for $t \in [0,T]$. (1.2c)

In (1.2a), the objective function is defined using the expectation operator, $E[\cdot]$. The applications considered in this dissertation, typically define the objective function according to time-varying power prices. Thus, the states or controllable inputs will typically relate to the power consumption of the system and (1.2a) will relate this consumption to an economic cost. (1.2b) is the dynamical model (a stochastic differential equation) and (1.2c) is the physical constraints of the controllable variable, u.

In general, it is numerically (and analytically) intractable to solve a general stochastic control problem of the form (1.2) using existing methods. However, in some special cases it is possible to solve (1.2); e.g. if the stochastic model in (1.2b) is linear (or Gaussian) and the objective function (1.2a) is quadratic (or linear). A Gaussian model with a quadratic objective function defines the so-called linear-quadratic regulator [8, 9]. This regulator is an optimal linear feedback strategy (i.e. u(t) = Kx(t) for a suitable, possibly time-varying, matrix K). However, for general nonlinear and non-Gaussian models, there exists few (if any) computationally tractable solution methods.

One numerically tractable framework for solving (1.2) is nonlinear model predictive control. This framework uses the perpetual data-flow (or system feedback) to compute the optimal control strategy, u, recursively [10-12].

1.3 Nonlinear model predictive control

Model predictive control is an advanced method for optimizing the operation of dynamical systems; it is referred to as nonlinear model predictive control, when the dynamics are nonlinear. The general optimal control problem - solved recursively in nonlinear model predictive control algorithms - have the form

$$\min_{x_k,u} \left\{ \int_{t_k}^{t_k+T} q(x_k(t), u(t)) \, \mathrm{d}t + q_f(x_k(T)) \right\}, \tag{1.3a}$$

s.t.

$$x_k(t) = x_0(t_k) + \int_{t_k}^t f(x_k(\tau), u(\tau); \theta) \, \mathrm{d}\tau, \tag{1.3b}$$

$$u(t) \in \mathcal{U}(t)$$
 for $t \in [t_k, t_k + T]$. (1.3c)

Comparing to (1.2), x_k denotes a single trajectory, representing the time-varying probability distribution defined in (1.2b) and $x_0(t_k)$ denotes the initial value; this parameter will be updated using a filtering technique, when new data, at time t_k , is available. The problem (1.3) is solved recursively, when new data is available; this recursive approach will be denoted the nonlinear model predictive control principle.

1.4 Applications in energy systems

This dissertation contains three applications relating to costefficient operation of energy systems.

The first application is the operation of the aeration equipment at wastewater treatment plants. During the treatment process of wastewater, oxygen is blown into treatment tanks to stimulate certain bio-chemical processes under aerobic conditions. This aeration process consumes a large fraction of the total power consumption at wastewater treatment plants [13]. Periods with aerobic conditions have to be succeeded by periods under anaerobic conditions to stimulate different bio-chemical processes. This phase-switching strategy can be very flexible under some circumstances and can be used to leverage time-varying power prices to ensure cost-efficient operation [14-16]. The results obtained, through this dissertation, show that significant cost-savings are available when intelligently coordinating the power prices across multiple power markets with the dynamics governing nutrient concentrations at wastewater treatment plants.

The second application is the operation of a joint refrigeration system consisting of a retail super market and an ice-tank module. The ice-tank can switch between three modes: idle, charging (accumulating ice) and discharging (curtailing the refrigeration system by melting the stored ice). The ice-tank is added to the refrigeration system using a sub-cooler attached to the returnpipe from the refrigerated display-cases. In practice, this could be the return-pipe from any refrigeration system (e.g. an HVAC system). The results from this project show that an ice-tank is an efficient, non-invasive method, for intelligent curtailment of refrigeration systems.

The third application is a simulation study of optimal energy arbitraging using a Vanadium redox-flow battery. A fixed-strategy for the Northern European day-ahead market is derived and short-term optimal control problems are used to optimally leverage the more volatile regulating power prices. The impact of special regulation due to a bilateral agreement between Denmark and Germany is investigated, as this market feature demonstrates very attractable power prices. The main conclusion is that the pay-back horizon of flow batteries can be heavily reduced by operating a battery across multiple power markets.

1.5 Outline

This dissertation is structured as follows:

- The first two chapters describe the theoretical scope. The first chapter introduces stochastic modelling of continuous-discrete systems. The second chapter describes numerical methods for solving (deterministic) optimal control problems of the form (1.3).
- The third chapter introduces the Northern European power market and discusses how market participants can leverage time-varying power prices to optimize and coordinate industrial processes.
- The fourth chapter summarizes the peer-reviewed scientific output of this dissertation.
- The last (fifth) chapter discusses future work and the perspectives of the methods and results contained in this dissertation.

2 Stochastic modelling in continuous time

This chapter introduces stochastic modelling in continuous-time using stochastic differential equations (SDEs). Two important topics within the field of modelling using SDEs are introduced; namely state and parameter estimation. The modelling methods applied in this dissertation are often referred to as continuous-discrete methods, since the dynamical model of the states is a continuous-time model but the system is only observed in discrete-time. There will be little theoretical attention given to a formal introduction of SDEs. Instead, this chapter will focus on the more applied details of using SDEs for optimal control purposes.

2.1 Stochastic differential equations

The general form of an SDE is

$$x(t) = x_0 + \int_0^t f(x(\tau), u(\tau); \theta) \,\mathrm{d}\tau + \int_0^t g(x(\tau), u(\tau); \theta) \,\mathrm{d}\omega(\tau), \tag{2.1}$$

or in short

$$dx = f(x, u; \theta) dt + g(x, u; \theta) d\omega, \qquad (2.2)$$

where $x: \mathbb{R} \to \mathbb{R}^{n_x}$ denotes the states, x_0 is the initial value of the states, $u: \mathbb{R} \to \mathbb{R}^{n_u}$ is the known input variables, $\theta \in \mathbb{R}^{n_\theta}$ is the model parameters and $\omega: \mathbb{R} \to \mathbb{R}^{n_\omega}$ denotes a standard (possibly multivariate) Brownian motion. Brownian motion is defined by its independent increments which satisfy that for each $s, t \in \mathbb{R}$, $\omega(t) - \omega(s)$ is normally distributed with zero mean and covariance $I_\omega | t - s |$ (where $I_\omega \in \mathbb{R}^{n_\omega \times n_\omega}$ is the identity matrix). Another important property of a Brownian motion, is that increments of ω on non-overlapping intervals are independent random variables. $f: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x} \times \mathbb{R}^{n_\omega}$ is called the diffusion function. For further reading, the author refers to [6, 7], which are more theoretical introductions to SDEs.

The systems considered in this dissertation are continuous-discrete systems with continuous-time dynamics governed by (2.2) and discrete-time observations. The most common observation equation seen in the literature is given by

$$y_k = h(x(t_k); \theta) + v_k. \tag{2.3}$$

Given $N_{\text{obs}} + 1$ observations $\{y_k\}_{k=0}^{N_{\text{obs}}}$, the random variables $\{v_k\}_{k=0}^{N_{\text{obs}}}$ are assumed to be independent and identically distributed; often it is assumed that they follow a normal distribution. $h: \mathbb{R}^{n_x} \to \mathbb{R}^{n_y}$ is the observation function that relates the hidden state, x, with an observation at time t_k , y_k .

This chapter will use an example model to demonstrate and evaluate the performance of the methods introduced later. This example model is given by

$$dx = \begin{bmatrix} x_2 \\ \lambda(1-x_1)x_2 - x_1 \end{bmatrix} dt + \begin{bmatrix} 0 & 0 \\ 0 & \sigma \end{bmatrix} d\omega, \qquad (2.4)$$

where $x = (x_1, x_2)'$. The parameter $\sigma > 0$ is the diffusivity and $\lambda > 0$ defines the stiffness [17] of the system and will be subject to parameter estimation in a subsequent section. This SDE (2.4) is often referred to as a stochastic extension to the Van der Pol oscillator model [18].

It is very challenging to find analytical solutions to most SDEs and hence numerical schemes have to be applied. The simplest method for simulating a single realization of an SDE is called the Euler-Maruyama method [6, 7]. This method is given by the recursion

$$x(t_{i+1}) - x(t_i) = f(x(t_i), u(t_i); \theta)(t_{i+1} - t_i) + g(x(t_i), u(t_i); \theta) \,\mathrm{d}\omega_i, \tag{2.5}$$

where $d\omega_i$ is normally distributed with zero mean and covariance $I_{\omega} |t_{i+1} - t_i|$. This numerical scheme is also used in the theoretical construction of Itô integrals [6, 7]. The integrals in (2.1) can be constructed theoretically according to many different limiting schemes [6, 7]; however, in this dissertation, these integrals will be interpreted as Itô integrals.

In Figure 2.1, three realizations of the Van der Pol model (2.4) are shown for parameter values $\sigma = 1/2$ and $\lambda \in (1,2,5)$; increasing stiffness of the system (2.4) is clearly observed for increasing values of λ .



Figure 2.1: Three SDE realizations of the Van der Pol model (2.4) using the Euler-Maruyama method (2.5).

2.2 State estimation

State estimation is a critical element in the algorithm-stack that constitutes nonlinear model predictive control algorithms. The main goal in state estimation is to reconstruct the system state x at time t_k given the previous observations $\{y_j\}_{j=0}^k$. The literature contains many state estimation methods for continuous-discrete models: exact filtering [6]; the Kalman filter [19]; the extended Kalman filter [6]; the unscented Kalman filter [20]; the unscented particle filter [21] and the efficient iterated filter [22]. The state estimation method is potentially a time-critical component of the decision-making algorithm and hence the filtering method used has to be computationally tractable. In general, exact filtering methods are computationally intractable and hence filter approximations have to be applied [6].

This section focuses on two different filtering techniques for continuous-discrete models: the extended Kalman filter and a class of Bayesian filtering techniques. But before these methods are introduced and discussed, the filtering principle and the exact filtering method are introduced.

2.2.1 The filtering principle

The ultimate goal for any state estimation method is to aggregate the information given by the process model and the information coming from the observations. In the continuous-discrete case, where the process model is given by an SDE (2.1) and the observation model is given by a likelihood function, $d_{k,{\rm lik}}$, the filtering principle defines two different distributions:

- the predicted (or prior) distribution $d_{k,\text{prior}}$: this is the conditional density function of $x(t_k)$ given all observations prior to t_k ,
- the estimated (or posterior) distribution $d_{k,\text{post}}$: this is the conditional density function of $x(t_k)$ given all observations up until (and including) t_k .

The main goal of the filtering principle is to associate $d_{k,\text{prior}}$ and $d_{k,\text{post}}$ in a recursive scheme using the observation model $d_{k,\text{lik}}$ at every new observation. In the interval (t_k,t_{k+1}) , x evolves according to the SDE (2.1). Let $d_k \colon \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}$ denote the conditional density function of x given the information accumulated until t_k . For $t \geq t_k$, d_k is governed by the forward Kolmogorov equation [6]

$$d_k(x,t) = d_{k,\text{post}}(x) - \int_{t_k}^t \nabla \cdot \left[\nu(x,\tau;\theta) d_k(x,\tau) - D(x,\tau;\theta) \nabla d_k(x,\tau) \right] d\tau \qquad (2.6)$$

where the advective flow field, $\nu \colon \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}^{n_x}$, and the diffusivity tensor, $D \colon \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}^{n_x \times n_x}$, are given by

$$D = \frac{1}{2}gg', \qquad (2.7a)$$

$$\nu = f - \operatorname{div} D. \tag{2.7b}$$

The new prior distribution at t_{k+1} is defined according to

$$d_{k+1,\text{prior}}(x) = d_k(x, t_{k+1}). \tag{2.8}$$

Bayes rule associates this prior distribution with the posterior distribution using the new information observed at t_{k+1} such that

$$d_{k+1,\text{post}} \propto d_{k+1,\text{prior}} d_{k+1,\text{lik}}.$$
(2.9)

For the observational model (2.3) with normally distributed noise variables with covariance R_k , the likelihood is given by

$$d_{k+1,\text{lik}}(x) = (2\pi)^{-n_x/2} \det(R_k)^{-1} \exp\left(-\frac{1}{2}(y_k - h(x;\theta))'R_k^{-1}(y_k - h(x;\theta))\right). \quad (2.10)$$

To summarize; the filtering principle is defined by a time and data update according to the following recursive scheme:

- 0. Set initial posterior distribution $d_{0,\text{post}}$ and set k=0.
- 1. Time update: solve (2.6) to obtain $d_{k+1,\text{prior}}$ from (2.8).
- 2. Data update: compute $d_{k+1,\text{post}}$ according to Bayes rule (2.9).
- 3. Repeat from step [1.].

In many practical use-cases, the time-update is numerically intractable as it involves solving a potentially high-dimensional advection-diffusion problem as given by the forward Kolmogorov equation (2.6). Hence, approximations have to be applied. The following two sections introduce two methods to approximate this exact filtering method.

2.2.2 Continuous-discrete extended Kalman filtering

The continuous-discrete extended Kalman Filter (CDEKF) is a computationally tractable alternative to the filtering problem stemming from the continuous-discrete system in (2.2)-(2.3). The CDEKF consists of two schemes: a prediction scheme (time update) and an updating scheme (data update) [6, 18]. These schemes will be defined in the following where $\hat{x}_{k|k} \in \mathbb{R}^{n_x}$ and $\hat{P}_{k|k} \in \mathbb{R}^{n_x imes n_x}$ denote the filtered state and covariance estimates at time t_k , and, $\hat{x}_k \colon \mathbb{R} \to \mathbb{R}^{n_x}$ and $\hat{P}_k \colon \mathbb{R} o \mathbb{R}^{n_x imes n_x}$ denote the predicted state and covariance values on the interval (t_k, t_{k+1}) . The CDEKF assumes independent and normally distributed random variables $\{v_k\}_{k=0}^{N_{\rm obs}}$. In this section it is assumed that v_k follows a normal distribution with zero mean and covariance R_k for any $k=0,...,N_{obs}$.

The prediction scheme (time update) Given the initial conditions

$$\hat{x}_k(t_k) = \hat{x}_{k|k},$$
 (2.11a)

$$\dot{P}_k(t_k) = \dot{P}_{k|k},\tag{2.11b}$$

the state- and covariance are predicted by solving the system of ordinary differential equations (ODEs) given by

$$\frac{\partial \hat{x}_k}{\partial t} = f(\hat{x}_k, u; \theta), \tag{2.12a}$$

$$\frac{\partial \hat{P}_k}{\partial t} = \frac{\partial f}{\partial x}(\hat{x}_k, u; \theta)\hat{P}_k + \hat{P}_k\frac{\partial f}{\partial x}(\hat{x}_k, u; \theta)' + g(\hat{x}_k, u; \theta)g(\hat{x}_k, u; \theta)'.$$
(2.12b)

The one-step predictions of the mean and covariance are obtained as the solution of (2.11)-(2.12) at the new sample point t_{k+1} . Consequently, the predictions of the state and covariance are

$$\hat{x}_{k+1|k} = \hat{x}_k(t_{k+1}), \tag{2.13a}$$

$$P_{k+1|k} = P_k(t_{k+1}). \tag{2.13b}$$

The updating scheme (data update)

The literature contains many methods for the updating scheme of extended Kalman filter algorithms. They all compute the innovation (or the one-step prediction error) by

$$e_{k+1} = y_{k+1} - h(\hat{x}_{k+1|k}; \theta), \qquad (2.14)$$

the Kalman filter gain by

$$K_{k+1} = \hat{P}_{k+1|k} \frac{\partial h}{\partial x} (\hat{x}_{k+1|k}; \theta)' \left[\frac{\partial h}{\partial x} (\hat{x}_{k+1|k}; \theta) \hat{P}_{k+1|k} \frac{\partial h}{\partial x} (\hat{x}_{k+1|k}; \theta)' + R_{k+1} \right]^{-1}$$
(2.15)

and the filtered state estimate by

$$\hat{x}_{k+1|k+1} = \hat{x}_{k+1|k} + K_{k+1}e_{k+1}.$$
(2.16)

The key difference appears in the formulation of the calculation of the filtered covariance. Two standard updating schemes for the covariance are

$$\hat{P}_{k+1|k+1} = \left[I - K_{k+1}\frac{\partial h}{\partial x}(\hat{x}_{k+1|k};\theta)\right]\hat{P}_{k+1|k}$$
(2.17a)

$$= \hat{P}_{k+1|k} - K_{k+1} \left[\frac{\partial h}{\partial x} (\hat{x}_{k+1|k}; \theta) \hat{P}_{k+1|k} \frac{\partial h}{\partial x} (\hat{x}_{k+1|k}; \theta)' + R_{k+1} \right] K_{k+1}'. \quad (2.17b)$$

Numerical implementations based on either (2.17a) or (2.17b) may give rise to inferior performance (even divergence), as the numerically computed values are not guaranteed to be both positive (semi-) definite and symmetric. One alternative to (2.17a) and (2.17b) is the Joseph stabilization form [18]. This formulation (which is mathematically equivalent to (2.17a) and (2.17b)) of the filtered covariance estimate is numerically guaranteed to be symmetric positive (semi-) definite [23]. Other numerically stable alternatives are the so-called array and square root algorithms [24]. However, these methods are omitted here.

Comparison to the exact filtering principle

The probability density function, $p: \mathbb{R}^n \times \mathbb{R}^{n \times n} \times \mathbb{R}^n \to \mathbb{R}$, of an *n*-dimensional normally distributed random variable is given by

$$p(\xi; \Sigma, \mu) = (2\pi)^{-n/2} \det(\Sigma)^{-1/2} \exp\left(-\frac{1}{2}(\mu - \xi)'\Sigma^{-1}(\mu - \xi)\right), \qquad (2.18)$$

where $\mu\in\mathbb{R}^n$ and $\Sigma\in\mathbb{R}^{n\times n}$ are the mean and covariance, respectively. The CDEKF implements two critical assumptions in the filtering scheme:

Stochastic Control Theory

- 1. the prior distribution, $d_{k+1,\text{prior}}$, is approximated by the normal density function $p\left(\cdot; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}\right)$,
- 2. and the posterior distribution, $d_{k+1,\text{prior}}$, is approximated by the normal density function $p\left(\cdot; \hat{P}_{k+1|k+1}, \hat{x}_{k+1|k+1}\right)$.

Thus, the CDEKF relies heavily on the assumption that these Gaussian approximations are adequate. Furthermore, the CDEKF does not (directly) accommodate observation models with different structures than defined in (2.3); the CDEKF linearizes the possibly nonlinear observation function, h, to derive a closed-form solution to the filtered state and covariance estimates. However, this linearization is not strictly necessary. This is explained in more details in the next section.

Example: the Van der Pol model

This section presents an example of using the CDEKF which is shown in Fig. 2.2. It is assumed that the random variables from the observation equation (2.3) are independent and follow a zero mean normal distribution with covariance

$$R_k = \frac{1}{100} \begin{bmatrix} 1 & 0\\ 0 & 1 \end{bmatrix}, \quad \forall k \in \{0, ..., N_{\text{obs}}\}.$$
 (2.19)

It is also assumed that h is the identity mapping; i.e. that h(x) = x. The sampling time is assumed to be constant; i.e. $t_k - t_{k-1} = 2.5$ for any $k \in \{1, ..., N_{obs}\}$. The nonlinearity of the system is clearly observed from the expanding and tightening prediction intervals within one prediction horizon (t_k, t_{k+1}) . In particular, this is observed when x_2 attains values greater than 2 or less than -2.



Figure 2.2: State estimation using the CDEKF with $\lambda = 1$ and $\sigma = 1/2$. The sampling time is assumed to be $t_k - t_{k-1} = 2.5$ for any $k \in \{1, ..., N_{obs}\}$.



Figure 2.3: Two steps of the CDEKF filtering process with $\lambda = 1$ and $\sigma = 1/2$. The sampling time is assumed to be $t_k - t_{k-1} = 2.5$ for any $k \in \{1, ..., N_{obs}\}$.

In Fig. 2.3 two steps of the CDEKF state estimation process are shown; the shaded grey area denotes the prior distribution generated by the CDEKF prediction equations (2.12), the shaded green area denotes the sampling distribution and the shaded red area denotes the posterior distribution of the filtered state estimates. As expected, it is observed that the posterior distribution has a a tighter confidence region compared to the prior distribution.

2.2.3 Bayesian filtering

In this section, an alternative solution to the filtering problem is proposed. Using the prior information from the SDE model (2.1) and the sampling information from the observations (2.3), it is possible to define an optimization problem that maximizes the conditional posterior distribution (conditioned on the newest information) of the unobserved state $x(t_k)$. This technique is called maximum a posteriori (MAP) estimation. The prediction scheme from the extended Kalman filter is used to approximate the prior distribution. Using this approximation of the prior distribution, the posterior distribution can be factorized from Bayes rule according to (identical to the data update defined in (2.9))

$$d_{k+1,\text{post}}(x) \propto p\left(x; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}\right) d_{k+1,\text{lik}}(x)$$
 (2.20)

The optimal state reconstruction can then be defined as the $x \in \mathbb{R}^{n_x}$ that maximizes the (non-normalized) posterior distribution (2.20)

$$\hat{x}_{k+1|k+1} = \operatorname*{arg\,max}_{x \in \mathcal{X}} \left\{ p\left(x; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}\right) d_{k+1, \text{lik}}\left(x\right) \right\},$$
(2.21)

where \mathcal{X} is some feasibility set (this might be \mathbb{R}^{n_x}). Due to potential numerical issues, it is often beneficial to maximize the natural logarithm of the posterior distribution. This transformation preserves the maximum of the optimization problem [25]. The logarithm of the prior density function is given by

$$\log p\left(x; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}\right) = -\frac{1}{2} \left[n_x \log 2\pi + \log \det \left(\hat{P}_{k+1|k} \right) + (\hat{x}_{k+1|k} - x)' \hat{P}_{k+1|k}^{-1} (\hat{x}_{k+1|k} - x) \right].$$
(2.22)

Assuming that the observation model is of the form (2.3) with independent and normally distributed random variables $\{v_k\}_{k=0}^{N_{\rm obs}}$ (with zero mean), the log transformation of the sampling distribution is given by

$$\log d_{k+1,\text{lik}}(x) = \log p(h(x;\theta); R_{k+1}, y_{k+1}) = -\frac{1}{2} \left[n_y \log 2\pi + \log \det (R_{k+1}) + (y_{k+1} - h(x;\theta))' R_{k+1}^{-1}(y_{k+1} - h(x;\theta)) \right].$$
(2.23)

The two first terms within the brackets of (2.22) and (2.23) are independent of x. Hence, the MAP problem (2.21) will be invariant to these terms. Define $l: \mathbb{R}^n \times \mathbb{R}^{n \times n} \times \mathbb{R}^n \to \mathbb{R}$ according to

$$l(x; \Sigma, \mu) = \frac{1}{2}(\mu - x)'\Sigma^{-1}(\mu - x), \qquad (2.24)$$

then the MAP problem has the equivalent log-transformed formulation

$$\hat{x}_{k+1|k+1} = \underset{x \in \mathcal{X}}{\arg\min} \left\{ l\left(x; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}\right) + l\left(h(x; \theta); R_{k+1}, y_{k+1}\right) \right\}.$$
(2.25)

Note that (2.25) is defined as a minimization program to eliminate the minus sign. Assuming that the solution to the optimization problem (2.25) is in the interior of \mathcal{X} , then the solution, $x^* \in \mathbb{R}^{n_x}$, satisfies the first-order optimality condition

$$\hat{P}_{k+1|k}^{-1}(\hat{x}_{k+1|k} - x^*) + \frac{\partial h}{\partial x}(x^*;\theta)' R_{k+1}^{-1}(y_{k+1} - h(x^*;\theta)) = 0.$$
(2.26)

If h is a linear function

$$h(x) = Cx, \tag{2.27}$$

for some $C \in \mathbb{R}^{n_y imes n_x}$, then (2.26) can be refactored into

$$\hat{P}_{k+1|k}^{-1}\hat{x}_{k+1|k} + C'R_{k+1}^{-1}y_{k+1} = \left[\hat{P}_{k+1|k}^{-1} + C'R_{k+1}^{-1}C\right]x^*.$$
(2.28)

Adding and subtracting $C\hat{x}_{k+1|k}$ to y_{k+1} in (2.28) yield that

$$\hat{P}_{k+1|k}^{-1} \hat{x}_{k+1|k} + C' R_{k+1}^{-1} y_{k+1}
= \hat{P}_{k+1|k}^{-1} \hat{x}_{k+1|k} + C' R_{k+1}^{-1} (y_{k+1} - C \hat{x}_{k+1|k} + C \hat{x}_{k+1|k})
= \left[\hat{P}_{k+1|k}^{-1} + C' R_{k+1}^{-1} C \right] \hat{x}_{k+1|k} + C' R_{k+1}^{-1} (y_{k+1} - C \hat{x}_{k+1|k}),$$
(2.29)

using this and solving for x^* result in

$$x^* = \hat{x}_{k+1|k} + \left[\hat{P}_{k+1|k}^{-1} + C'R_{k+1}^{-1}C\right]^{-1}C'R_{k+1}^{-1}(y_{k+1} - C\hat{x}_{k+1|k}).$$
(2.30)

The Woodbury matrix identity (or the matrix inversion lemma) [26, 27] defines the matrix identity

$$\hat{P}_{k+1|k}C'\left[C\hat{P}_{k+1|k}C'+R_{k+1}\right]^{-1} = \left[\hat{P}_{k+1|k}^{-1}+C'R_{k+1}^{-1}C\right]^{-1}C'R_{k+1}^{-1}.$$
(2.31)

Using this relation together with (2.30) yields the same filtered state estimate as the estimate defined by the CDEKF (2.16)

$$\hat{x}_{k+1|k+1} = x^* = \hat{x}_{k+1|k} + \hat{P}_{k+1|k}C' \left[C\hat{P}_{k+1|k}C' + R_{k+1}\right]^{-1} (y_{k+1} - h(\hat{x}_{k+1|k})). \quad (2.32)$$

The covariance of $\hat{x}_{k+1|k+1}$ can be computed using the observed Fischer Information matrix [28]

$$\hat{P}_{k+1|k+1} = \left[\frac{\partial^2}{\partial x^2} \left(l\left(\hat{x}_{k+1|k+1}; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}\right) + l\left(h(\hat{x}_{k+1|k+1}; \theta); R_{k+1}, y_{k+1}\right) \right) \right]^{-1}, \qquad (2.33)$$

where

$$\frac{\partial^2 l}{\partial x^2} \left(x; \hat{P}_{k+1|k}, \hat{x}_{k+1|k} \right) = \hat{P}_{k+1|k}^{-1}, \qquad (2.34)$$

and

$$\frac{\partial^2 l}{\partial x^2} \left(h(x;\theta); R_{k+1}, y_{k+1} \right) = \frac{\partial h}{\partial x} (x;\theta)' R_{k+1}^{-1} \frac{\partial h}{\partial x} (x;\theta) + \sum_{i=1}^{n_y} \left[R_{k+1}^{-1} (y_k - h(\hat{x}_{k+1|k+1};\theta)) \right]_i \frac{\partial^2 h_i}{\partial x^2} (x;\theta).$$
(2.35)

Thus

$$\hat{P}_{k+1|k+1} = \left[\hat{P}_{k+1|k}^{-1} + \frac{\partial h}{\partial x}(\hat{x}_{k+1|k+1};\theta)'R_{k+1}^{-1}\frac{\partial h}{\partial x}(\hat{x}_{k+1|k+1};\theta) + \sum_{i=1}^{n_y} \left[R_{k+1}^{-1}(y_k - h(\hat{x}_{k+1|k+1};\theta))\right]_i \frac{\partial^2 h_i}{\partial x^2}(\hat{x}_{k+1|k+1};\theta)\right]^{-1}.$$
(2.36)

If h is a linear function, as in (2.27), then (2.36) simplifies into

$$\hat{P}_{k+1|k+1} = \left[\hat{P}_{k+1|k}^{-1} + C'R_{k+1}^{-1}C\right]^{-1}.$$
(2.37)

Using the Woodbury matrix identity, (2.37) can be expanded into

$$\hat{P}_{k+1|k+1} = \hat{P}_{k+1|k} - \hat{P}_{k+1|k}C' \left[C\hat{P}_{k+1|k}C' + R_{k+1}\right]^{-1}C\hat{P}_{k+1|k}.$$
(2.38)

This is exactly the covariance updating formula (2.17) defined by the CDEKF. Thus, the CDEKF and the MAP methods are equivalent if the observation function h is linear.

Generally, when h is nonlinear, there are two main differences between the CDEKF and the MAP methods defined above:

- the MAP method computes the filtered state estimate as the root of a nonlinear system of equations (2.26),
- 2. and the MAP method includes a higher order term of h in the filtered covariance estimate (2.36).



Figure 2.4: The first two derivatives of the nonlinear observation function h. (2.41).

Comparison to the exact filtering principle

One advantage of the MAP method is the ability to have more general sampling distributions than the common normal distribution assumption defined in (2.3). In the general case, the filtered state estimate is given by

$$\hat{x}_{k+1|k+1} = \operatorname*{argmin}_{x \in \mathcal{X}} \left\{ l(x; \hat{P}_{k+1|k}, \hat{x}_{k+1|k}) - \log d_{k+1, \text{lik}}(x) \right\}.$$
(2.39)

Assuming that $d_{k+1,{\rm lik}}$ is twice-differentiable, the filtered covariance MAP estimate is given by

$$\hat{P}_{k+1|k+1} = \left[\hat{P}_{k+1|k}^{-1} - \frac{\partial^2}{\partial x^2} \log d_{k+1,\text{lik}}(x)\right]^{-1}.$$
(2.40)

Thus, the MAP method implements the same two critical Gaussian assumptions concerning the prior and posterior distributions as the CDEKF, but contains no assumptions concerning the likelihood distribution (except of it being twice-differentiable).

Example: the Van der Pol model with nonlinear observations

This section presents an example of using the MAP method for recursive state estimation of the van der Pol model with nonlinear observations. It is assumed that the random variables from the observation equation (2.3) are independent and follow a zero mean normal distribution with covariance given as in (2.19). The observation function is defined as

$$h(x) = \begin{bmatrix} \log(1 + \exp(x_1)) \\ \log(1 + \exp(x_2)) \end{bmatrix}.$$
 (2.41)

In Fig. 2.4, h_1 and h_2 and their first two derivatives are shown; the Hessian of these functions clearly has significant values in

the range (-3,3), which is the primary range of the hidden state x observed in Fig. 2.1. Using the same constant sampling time and model parameters as in Fig. 2.2, Fig. 2.5 shows the state estimation results by recursively applying the MAP method. Fig. 2.6 shows two



Figure 2.5: State estimation using the MAP method with $\lambda = 1$ and $\sigma = 1/2$. The sampling time is assumed to be $t_k - t_{k-1} = 2.5$ for any $k \in \{1, ..., N_{obs}\}$.



Figure 2.6: Two steps of the MAP filtering process with $\lambda = 1$ and $\sigma = 1/2$. The sampling time is assumed to be $t_k - t_{k-1} = 2.5$ for any $k \in \{1, ..., N_{obs}\}$.

steps of the MAP filtering process. Again, it is observed that the posterior distribution has a tighter confidence region compared to the prior distribution.

2.2.4 Numerical comparison of the filtering methods

In Tab. 2.1 the two filtering methods presented in the previous sections have been compared numerically. 100 data-sets have been used to quantify the performance of the two filtering methods. The first column is the log-likelihood of the observations given the true parameters, but using the two different filtering techniques. The Gaussian approximation of the likelihood function (2.47) (which is defined in the next section) has been used to approximate the likelihood. The second column shows the log-density value of the true state measured in the posterior distribution and the third column shows the log-density value of the true state measured in the prior distribution. The last row in each of the table-blocks shows the percentage of which either method outperforms the other in the 100 data-sets. From this table, it is clearly observed (especially from the second column) that the MAP method outperforms the CDEKF in reconstructing the true state value. This is observed consistently across all three sampling times tested. From the first and third columns it is observed that improved filtered estimates also seem to imply statistical improvements in the one-step predictions.

2.3 Parameter estimation

This section presents a maximum likelihood (ML) method for parameter estimation of continuous-discrete systems with continuous-time dynamics governed by (2.1) and discrete-time observations given by (2.3). It is assumed that $\{v_k\}_{k=0}^{N_{\rm obs}}$ are independent and normally distributed with zero mean and covariance R_k . Define the information (or observation set) as

$$\mathcal{Y}_{N_{\rm obs}} = \{y_k\}_{k=0}^{N_{\rm obs}}.$$
(2.42)

The likelihood function, $L: \mathbb{R}^{n_{\theta}} \to \mathbb{R}$, is then defined as the joint probability density of the available information $\mathcal{Y}_{N_{\text{obs}}}$

$$L(\theta; \mathcal{Y}_{N_{\text{obs}}}) = d_{\theta}(\mathcal{Y}_{N_{\text{obs}}}).$$
(2.43)

where $d_{\theta} \colon \mathbb{R}^{n_{\theta}} \to \mathbb{R}$ is the conditional probability density function of the observations given the parameters θ . Using the product rule of conditional densities, (2.43) can be factorized according to

$$L(\theta; \mathcal{Y}_{N_{\text{obs}}}) = d_{\theta}(y_0) \prod_{k=1}^{N_{\text{obs}}} d_{\theta}(y_k | \mathcal{Y}_{k-1}).$$
(2.44)

The ML principle defines the ML parameter estimates, $\theta_{\rm ML},$ as the parameters that maximize the likelihood function

$$\theta_{\rm ML} = \underset{\theta \in \Theta}{\arg\max} \Big\{ L(\theta; \mathcal{Y}_{N_{\rm obs}}) \Big\}, \qquad (2.45)$$

Table 2.1: This table shows a statistical comparison between the CDEKF and the MAP method. 100 data-sets have been used to quantify the performance of the two filtering methods. The first column shows the likelihood of the observations, the second column shows the density function value of the true state measured in the posterior distribution and the third column shows the density function value of the true state measured in the prior distribution. It is the log-scores that have been used as performance metric in this table and it is the log-likelihood that is used as metric in the first column. The last row in each of the table-blocks shows the percentage of which either methods outperforms the other in each of the 100 data-sets. The 10 %, 50 % and 90 % percentiles (across data-sets) are shown for each metric.

	Log score $h(\hat{x}_{k+1 k})$	Log score $\hat{x}_{k+1 k+1}$	Log score $\hat{x}_{k+1 k}$
1475	Sampling rat	e: $t_k - t_{k-1} = 2.5$	
	0 100	0 700	1 1 (1
10 %	-0.128	-0.726	1.161
50 %	0.780	-0.01/	2.197
90 %	2.749	1.341	4.863
out of 100	52 %	83 %	58 %
CDEKF			
10 %	-0.130	-0.441	1.154
50 %	-0.745	0.507	2.252
90 %	2.935	2.447	5.249
out of 100	48 %	17 %	42 %
	Sampling rat	e: $t_k - t_{k-1} = 1.0$	
MAP			
10 %	-0.764	-1.034	0.063
50 %	-0.580	-0.839	0.334
90 %	-0.311	-0.566	0.762
out of 100	70 %	82 %	65 %
CDEKF			
10 %	-0.762	-0.995	0.072
50 %	-0.571	-0.769	0.331
90 %	-0.322	-0.499	0.744
out of 100	30 %	18 %	35 %
	Sampling rat	e: $t_k - t_{k-1} = 0.5$	
MAP			
10 %	-1.048	-1.241	-0.417
50 %	-0.924	-1.069	-0.264
90 응	-0.785	-0.828	0.022
out of 100	59 %	76 %	61 %
CDEKF			
10 %	-1.048	-1.215	-0.417
50 %	-0.925	-1.053	-0.272
90 %	-0.778	-0.811	0.046
out of 100	41 %	24 %	39 %

where Θ is some feasibility set (this might be $\mathbb{R}^{n_{\theta}}$). Due to numerical limitations, it is often preferable to consider the loglikelihood as the objective function. The log-transformation is monotonically increasing and hence preserves the maximum of the optimization problem. Thus, the parameter estimation problem is re-defined according to

$$\theta_{\mathrm{ML}} = \operatorname*{arg\,max}_{\theta \in \Theta} \bigg\{ \log d_{\theta}(y_0) + \sum_{k=1}^{N_{\mathrm{obs}}} \log d_{\theta}(y_k | \mathcal{Y}_{k-1}) \bigg\}.$$
(2.46)

In [29] it is discussed that under mild regularity conditions, it is reasonable to assume that the conditional densities in (2.46) can be well approximated as Gaussian densities. In [29] this Gaussian approximation is defined such that

$$d_{\theta}(y_k|\mathcal{Y}_{k-1}) \approx p\left(y_k; \frac{\partial h}{\partial x}(\hat{x}_{k|k-1}; \theta) \hat{P}_{k|k-1} \frac{\partial h}{\partial x}(\hat{x}_{k|k-1}; \theta)' + R_k, h(\hat{x}_{k|k-1}; \theta)\right). \quad (2.47)$$

In (2.47) it is important to emphasize that $\hat{x}_{k|k-1}$ and $\hat{P}_{k|k-1}$ also depend on θ . The covariance of the parameters can be computed using the observed Fischer information matrix; i.e. as the negative inverse of the Hessian of the objective function in (2.46).

2.3.1 Estimating the van der Pol model

This section presents a case-study applying the parameter estimation method defined above. The van der Pol model is used as a test model with the true parameters being $(\lambda, \sigma) = (1, 1/2)$. It is assumed that the random variables from the observation equation (2.3) are independent and follow a zero mean normal distribution with covariance given as in (2.19). It is also assumed that σ is known, but $\lambda > 0$ is unknown and has to be estimated from data. The observations are sampled according to the sampling time $t_k - t_{k-1} = 1/2$ for all observations and h is the nonlinear function defined in (2.41). 100 observations are simulated and used to estimate λ . In Fig. 2.7 the log-likelihood is shown as a function of the unknown parameter λ using the two filtering methods presented previously in this chapter. Using a simple golden-section search algorithm [30], the estimated values of λ are $\lambda_{\text{ML,MAP}} = 0.929$ and $\lambda_{\text{ML,CDEKF}} = 0.914$.



Figure 2.7: Negative log-likelihood as a function of λ .

3 Nonlinear model predictive control

Model predictive control is an advanced method for optimizing the operation of dynamical systems. This chapter describes the nonlinear model predictive control (NMPC) principle for continuousdiscrete systems with continuous-time dynamics governed by an SDE (2.1) and discrete-time observations given by a general likelihood function $d_{k,\text{lik}}$. This likelihood function could e.g. be given by the observation model given in (2.3). The previous chapter introduced



Figure 3.1: This figure shows an overview of the sub-components that constitutes NMPC software. This figure is a modified version of a figure used in [18].

the concepts of state and parameter estimation. These methods are crucial elements in NMPC software: the parameter estimation process calibrates the model to the most recent data (this is often done offline) and the state estimation algorithm recursively readjusts the initial value used in the optimal control algorithm according to new data (this is an online process). Fig. 3.1 shows an overview of the sub-components that constitute NMPC software. The final component of the NMPC stack shown in Fig. 3.1 is the dynamical optimization component. Several methods can be used to solve dynamic optimization problems: the Hamilton-Jacobi-Bellmann equations (or dynamical programming) [31]; indirect methods (based on Pontryagin's minimum principle and calculus of variations) [32, 33] or direct methods (e.g. shooting and collocation methods) [34-36]. This chapter will only focus on the indirect and direct solution methods to solve the dynamical optimization problem. Methods based on dynamical programming can only be used on lowdimensional systems ($n_x \leq 3$) as the underlying principle is cursedby-dimensionality [31]. However, this principle may be used for very specific dynamical systems (e.g. linear systems).

3.1 The model predictive control principle

The general form of the optimal control problem, considered in this dissertation, has the structure

$$\min_{x_k,u} \left\{ \int_{t_k}^{t_k+T} q(x_k(t), u(t)) \, \mathrm{d}t + q_f(x_k(t_k+T)) \right\}, \tag{3.1a}$$

s.t.
$$x_k(t_k) = \hat{x}_{k|k},$$
 (3.1b)

$$\frac{\partial x_k}{\partial t}(t) = f(x_k(t), u(t); \theta) \quad \text{for} \quad t \in [t_k, t_k + T], \quad (3.1c)$$

$$u(t) \in \mathcal{U}(t) \qquad \qquad \text{for} \qquad t \in [t_k, t_k + T], \qquad (3.1d)$$

where $q: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ is the Lagrangian and $q_f: \mathbb{R}^{n_x} \to \mathbb{R}$ is the endpoint cost (or penalty). $\mathcal{U}(t)$ is the feasibility set defining the set of admissible input signals; typical constraints are box-type or rate-of-change-type constraints. T > 0 is the optimization horizon. The formulation (3.1) shows the clear relation to the filtering method used; $\hat{x}_{k|k}$ is the posterior estimate of the hidden state $x(t_k)$ at time t_k given the most recent information set \mathcal{Y}_k . The dynamical model (3.1b)-(3.1c) is identical to the prediction model from the prediction scheme from (2.12).

The NMPC principle is defined according to the following recursive scheme:

- 0. Set initial posterior estimate $\hat{x}_{0|0}$ and set k=0.
- 1. Optimize: solve (3.1) and implement the optimal input signal, u^* . u^* is implemented until a new observation at time t_k is observed.
- 2. Data update: given new observation, $y_{k+1}, \mbox{ compute } \hat{x}_{k+1|k+1}$ using a suitable filter.
- 3. Repeat from step [1.].

3.2 Pontryagin's minimum principle

Let $\lambda_k \colon \mathbb{R} \to \mathbb{R}^{n_x}$ denote the adjoint variables (or co-states) of (3.1), then λ_k satisfies the (backward) ODE given by

$$-\frac{\partial\lambda_k}{\partial t}' = \frac{\partial q}{\partial x}(x_k, u) + \lambda_k' \frac{\partial f}{\partial x}(x_k, u; \theta), \qquad (3.2)$$

with boundary condition

$$\lambda_k(t_k+T)' = \frac{\partial q_f}{\partial x}(x_k(t_k+T)). \tag{3.3}$$

This formulation is derived in e.g. [8] using calculus of variations. The combined boundary value problem of the state and adjoint variables are given by

$$\frac{\partial x_k}{\partial t} = f(x_k, u^*; \theta), \qquad \qquad x_k(t_k) = \hat{x}_{k|k}, \qquad (3.4a)$$

$$-\frac{\partial \lambda_k}{\partial t}' = \frac{\partial q}{\partial x}(x_k, u^*) + \lambda_k' \frac{\partial f}{\partial x}(x_k, u^*; \theta), \qquad \lambda_k(t_k + T)' = \frac{\partial q_f}{\partial x}(x_k(t_k + T)), \quad (3.4b)$$

0

where the optimal controllable input, u^* , is defined point-wise as the solution to the optimization problem [8]

$$u^*(t) = \underset{u \in \mathcal{U}(t)}{\operatorname{arg\,min}} \left\{ q(x_k(t), u) + \lambda_k(t)' f(x_k(t), u; \theta) \right\} \quad \text{for} \quad t \in [t_k, t_k + T]. \quad (3.5)$$

The combined conditions (3.4)-(3.5) yield one method for solving the optimal control problem (3.1). This method is called an indirect solution method, as it defines the optimality conditions of the optimal control problem (3.1) in continuous-time. However, numerical methods still have to be used in order to solve the combined conditions (3.4)-(3.5). This might be challenging if the feasibility set, $\mathcal{U}(t)$, imposes non-trivial constraints.

3.3 Direct methods

The previous section describes a class of indirect methods that can be defined based on Pontryagin's minimum principle. This principle defines the optimality conditions of (3.1) in continuous-time before discretization. This approach is often referred to as *first optimize*, then discretize. A different class of methods called direct methods, adapts the reverse approach to *first discretize*, then optimize. In this solution method paradigm, there exists two main branches of numerical methods: shooting and collocation methods. These methods are introduced in the next two sections

Common for any direct method, is that a discretization of u has to be given. For simplicity, it is in this dissertation assumed that u is discretized according to the piece-wise definition

$$u(t) = \sum_{i=1}^{N_u} v_i \mathbb{1}_{\tau_{i-1}^u, \tau_i^u}(t) \quad \text{and} \quad v = (v_1 \dots v_{N_u}). \tag{3.6}$$

The indicator function, $1\!\!1$, is defined as

$$\mathbb{1}_{a,b}(t) = \begin{cases} 1, & a \le t < b, \\ 0, & \text{otherwise,} \end{cases}$$
(3.7)

and $\mathcal{T}_u = \{ au_i^u \}_{i=0}^{N_u}$ satisfies the ordering

$$t_k \le \tau_0^u < \dots < \tau_{N_u}^u = t_k + T.$$
(3.8)

3.3.1 The single shooting method

One of the main motivations of shooting methods is that given an initial value, $\hat{x}_{k|k}$, and input trajectory, u, the state trajectory, x_k , is uniquely determined. Hence, there seems to be little need to consider x_k as an unknown trajectory as implicitly stated in (3.1). Using the input discretization (3.6), x_k is uniquely determined by $\hat{x}_{k|k}$ and v. This means that the optimal control problem (3.1) can be condensed into

$$\min_{v} \left\{ \mathcal{J}(v) = \sum_{i=1}^{N_u} \int_{\tau_{i-1}^u}^{\tau_i^u} q(x_k(t; v, \hat{x}_{k|k}), u_i) \, \mathrm{d}t + q_f(x_k(\tau_{N_u}; v, \hat{x}_{k|k})) \right\}, \quad (3.9a)$$

$$v_i \in \mathcal{U}(\tau_i^u)$$
 for any $i = 1...N_u$. (3.9b)
In this formulation, $x_k(t; v, \hat{x}_{k|k})$ stresses the dependence on $\hat{x}_{k|k}$ and v. The formulation (3.9) is called the single (or sequential) shooting method. The derivatives of the objective function in (3.9) can be computed using the adjoint variables defined in (3.4b) [37, 38]

$$\frac{\partial \mathcal{J}(v)}{\partial v_i} = \int_{\tau_{i-1}^u}^{\tau_i^u} \left[\frac{\partial q}{\partial u} (x_k(t), v_i) + \lambda_k(t)' \frac{\partial f}{\partial u} (x_k(t), v_i; \theta) \right] \mathrm{d}t. \tag{3.10}$$

Using the single shooting formulation (3.9) and the derivative calculation (3.10), it is straight-forward to implement this method numerically using e.g. **IPOPT** [39], which is a framework for solving large-scale nonlinear optimization problems. There exists many methods for discretizing the ODE (3.1c), objective function (3.1a) and derivative calculation (3.10). However, these methods are not elaborated any further; one should simply use a suitable and well-established ODE solver (e.g. the C++ libraries odeint [40] or SUNDIALS [41]).

One well-known limitation with the single shooting method is the lack of robustness concerning dynamical systems with unstable regions [35]. The single shooting method will fail if an input sequence v drives the dynamical system (3.1c) to an unstable region, causing the state trajectory, x_k , to diverge (numerically explode in values).

3.3.2 The collocation method

The class of direct methods called collocation methods take a different approach to the discretization of (3.1) compared to the single shooting method. Instead of sequentially computing the input sequence, v, and state trajectory, x_k , the collocation methods computes these variables simultaneously. Let $\mathcal{T}_x = \{\tau_l^x\}_{l=0}^{N_x}$ be defined such that

$$t_k \le \tau_0^x < \dots < \tau_{N_r}^x = t_k + T. \tag{3.11}$$

The values of x_k evaluated in \mathcal{T}_x is defined according to

$$s_l = x_k(\tau_l^x), \qquad l = 0...N_x \qquad \text{and} \qquad s = (s_0...s_{N_x}).$$
 (3.12)

To illustrate the structure of collocation methods, the forward Euler method is used to discretize the state trajectory and for simplicity it is assumed that $\mathcal{T}_x = \mathcal{T}_u$. Using this convention, the optimal control problem (3.1) is discretized into the nonlinear program

$$\min_{s,v} \left\{ \sum_{i=1}^{N_x} \frac{(\tau_i^x - \tau_{i-1}^x)}{2} \left(q(s_i, v_i) + q(s_{i-1}, v_i) \right) + q_f(s_{N_x}) \right\},$$
(3.13a)

s.t. $s_0 = \hat{x}_{k|k},$ (3.13b)

$$\begin{aligned} s_{i+1} - s_i &= (\tau_{i+1}^x - \tau_i^x) f(s_i, v_i; \theta) & \text{for any} \quad i = 1 \dots N_x, \quad (3.13c) \\ v_i &\in \mathcal{U}(\tau_i^u) & \text{for any} \quad i = 1 \dots N_u. \quad (3.13d) \end{aligned}$$

The mid-point rule has been used to discretize the integral in the objective function (3.13a). Using a collocation method to solve the dynamical optimization problem (3.1), it is very simple to compute both the gradient and the hessian of the objective function and the Jacobians of the constraints (3.13b)-(3.13d). The hessian of the objective function and the Jacobian of the constraints will also have a systematic sparse structure, making optimization problems originating from collocation schemes very suitable for implementation using **IPOPT** [18, 35, 42]. One immediate limitation of any collocation methods is the large amount of decision variables. In [43] it is described how numerical methods with spectral convergence can be used to discretize problems of the form (3.1) and thereby lower the number of decision variables needed to attain a sufficient accuracy of the solution.

3.3.3 The multiple shooting method

The multiple shooting method is a hybrid method of the single shooting and collocation methods [34, 35]. The main idea behind the multiple shooting method is to divide the prediction horizon $(t_k, t_k + T)$ into sub-intervals (much like the collocation method), but apply the single shooting framework on each of these elements. Let \mathcal{T}_x denote discretization that defines the sub-intervals and for simplicity assume that $\mathcal{T}_x = \mathcal{T}_u$. Let x_k^i denote the *i*-th trajectory defined by the ODE

$$\frac{\partial x_k^i}{\partial x} = f\left(x_k^i, v_i; \theta\right) \quad \text{on} \quad t \in [\tau_{i-1}^x, \tau_i^x] \quad \text{for any} \quad i = 1...N_x. \quad (3.14)$$

To make these partial trajectories consistent, a continuity condition must be imposed. This condition can be implemented in multiple ways. In this dissertation, it is implemented from the following conditions

$$x_k^i(\tau_i^x) = s_i$$
 for any $i = 1...N_x$, (3.15a)

$$x_k^i(au_{i-1}^x) = s_{i-1}$$
 for any $i = 1...N_x,$ (3.15b)

where $s = (s_0...s_{N_x})$ is the initial values of the N_x state trajectories from (3.14). The optimal control problem (3.1) can then be discretized according to

$$\min_{s,v} \quad \left\{ \mathcal{J}(s,v) = \sum_{i=1}^{N_x} \int_{\tau_{i-1}^x}^{\tau_i^x} q\left(x_k^i(t;s_{i-1},v_i),v_i\right) \mathrm{d}t + q_f(s_{N_x}) \right\}, \tag{3.16a}$$

s.t.
$$s_0 = \hat{x}_{k|k},$$
 (3.16b)

$$s_{i-1} = x_k^i(\tau_{i-1}^x; s_{i-1}, v_i)$$
 for any $i = 1...N_x$, (3.16c)

$$s_i = x_k^i(\tau_i^x; s_{i-1}, v_i)$$
 for any $i = 1...N_x$, (3.16d)

 $v_i \in \mathcal{U}(\tau_i^u)$ for any $i = 1...N_u$. (3.16e)

The derivatives of this optimization problem can be computed very similarly to the derivatives computed for the single shooting problem (3.10). However, the adjoint trajectory, λ_k^i , will have to

be defined on the individual sub-intervals [34]. Define the adjoint trajectory on the $i\mathchar`-th$ interval according to

$$-\lambda_k^{i\,\prime} = \frac{\partial q}{\partial x} \left(x_k^i, v_i \right) + \lambda_k^{i\,\prime} \frac{\partial f}{\partial x} \left(x_k^i, v_i \right) \quad \text{on} \quad t \in [t_{i-1}^x, t_i^x] \quad \text{for} \quad i = 1...N_x, \quad (3.17)$$

with boundary condition

$$\lambda_k^i(t_i^x) = 0. \tag{3.18}$$

The derivatives of the objective function with respect to the control variables, v, are then defined according to [37, 38]

$$\frac{\partial \mathcal{J}(s,v)}{\partial v_i} = \int_{\tau_{i-1}^x}^{\tau_i^x} \left[\frac{\partial q}{\partial u} \left(x_k^i(t;s_{i-1},v_i), v_i \right) + \lambda_k^i(t)' \frac{\partial f}{\partial u} \left(x_k^i(t;s_{i-1},v_i), v_i \right) \right] \mathrm{d}t, \quad (3.19)$$

for any $i=1\ldots N_u$. The derivatives with respect to the initial conditions are given by

$$\frac{\partial \mathcal{J}(s,v)}{\partial s_i} = \lambda_k^i(\tau_{i-1}^x), \qquad (3.20)$$

for any $i=0...N_x-1$ where the last nodal point, $s_{N_x},\ {\rm defines}$ the derivative

$$\frac{\partial \mathcal{J}(s,v)}{\partial s_{N_x}} = \frac{\partial q_f}{\partial x}(s_{N_x}). \tag{3.21}$$

The Jacobian of the constraint (3.16d) also requires the derivatives

$$\frac{\partial x_k^i(\tau_i^x; s_{i-1}, v_i)}{\partial v_i} \quad \text{and} \quad \frac{\partial x_k^i(\tau_i^x; s_{i-1}, v_i)}{\partial s_{i-1}}. \quad (3.22)$$

However, these are simply computed from the sensitivity systems [38]

$$\frac{\partial}{\partial t}\frac{\partial x_k^i(t;s_{i-1},v_i)}{\partial s_{i-1}} = \frac{\partial f}{\partial x}(x_k^i(t;s_{i-1},v_i),v_i)\frac{\partial x_k^i(t;s_{i-1},v_i)}{\partial s_{i-1}}, \qquad t \in [\tau_{i-1}^x,\tau_i^x], \quad (3.23a)$$

$$\frac{\partial}{\partial t}\frac{\partial x_k^i(t;s_{i-1},v_i)}{\partial v_i} = \frac{\partial f}{\partial x}(x_k^i(t;s_{i-1},v_i),v_i)\frac{\partial x_k^i(t;s_{i-1},v_i)}{\partial v_i} + I_u, \quad t \in [\tau_{i-1}^x,\tau_i^x], \quad (3.23b)$$

where $I_u \in \mathbb{R}^{n_u \times n_u}$ is an identity matrix. The initial values of the sensitivity system is given by

$$\frac{\partial x_k^i(\tau_{i-1}^x; s_{i-1}, v_i)}{\partial s_{i-1}} = I_x,$$
(3.24a)

$$\frac{\partial x_k^i(\tau_{i-1}^x; s_{i-1}, v_i)}{\partial s_{i-1}} = 0_u, \tag{3.24b}$$

where $I_x \in \mathbb{R}^{n_x \times n_x}$ is an identity matrix and $0_u \in \mathbb{R}^{n_u \times n_u}$ is a zero matrix.

3.4 Examples

This section presents the NMPC framework applied on two test cases:

- A case-study where the observation model is nonlinear but with additive Gaussian noise.
- A case-study where the observation model is nonlinear and with a Poisson observation model.

The case-studies will use **CasADi** [44] to interface the discretized optimal control problem to **IPOPT**. The MAP estimation method is used to conduct the state estimation in order to recursively solve the optimal control problem (3.1) according to the NMPC principle. The van der Pol oscillator model is used as test model with parameters $(\lambda, \sigma) = (1, 1/10)$. The goal in both case-studies is for x_1 to track a set-point trajectory, \overline{x}_1 . This set-point trajectory is shown in Fig. 3.2. The optimal control problem, which repeatedly will be solved according to the NMPC principle, is defined according to

$$\min_{x_k,u} \left\{ \int_{t_k}^{t_k+T} \left[\left(x_{k,1}(t) - \overline{x}_1(t) \right)^2 + \varepsilon u(t)^2 \right] \mathrm{d}t \right\},\tag{3.25a}$$

s.t.
$$x_k(t_k) = \hat{x}_{k|k},$$
 (3.25b)
 $\frac{\partial x_k}{\partial t}(t) = \begin{bmatrix} x_{k,2}(t) \\ (1 - x_{k,1}(t))x_{k,2}(t) - x_{k,1}(t) + u(t) \end{bmatrix}$ for $t \in [t_k, t_k + T],$ (3.25c)

$$\overline{\partial t}^{(t)} = \left[(1 - x_{k,1}(t)) x_{k,2}(t) - x_{k,1}(t) + u(t) \right] \quad \text{for } t \in [t_k, t_k + T], \quad (3.25c)$$
$$u(t) \in [-1, 1] \quad \text{for } t \in [t_k, t_k + T]. \quad (3.25d)$$

As seen from from (3.25c) it is assumed that it is only possible to control the system via the dynamics governing x_2 . Via (3.25d) it is imposed that the control signal of u is constrained to be in the (box-type) interval [-1,1]. $\varepsilon > 0$ is a small regularization coefficient, which penalizes non-zero values of u.



Figure 3.2: Set-point trajectory for the case-studies.

3.4.1 Nonlinear observations with additive Gaussian noise

In the first case-study, it is assumed that the observational model is nonlinear, but with additive Gaussian noise. Thus, the observation equation is of the type defined in (2.3); i.e. that

$$y_k = \begin{bmatrix} \log(1 + \exp(x_1(t_k))) \\ \log(1 + \exp(x_2(t_k))) \end{bmatrix} + v_k.$$
(3.26)

It is assumed that the random variables in (3.26) are independent and follow a zero mean normal distribution with covariance given as in (2.19). It is assumed that the sampling time is $t_k - t_{k-1} = 1/5$ between any observations. The regularization coefficient is chosen such that $\varepsilon = 1/1000$.

In Fig. 3.3, a simulation using the NMPC principle, to track the set-point trajectory shown in Fig. 3.2, is shown. Using the NMPC principle it is clearly possible to track \overline{x}_1 with the true state, x_1 , oscillating around the set-point trajectory. It is important to note that the sampling time is very important; especially for systems with strong diffusion. If the sampling time is too long, the feedback is not sufficient to recursively re-adjust the optimal input, u. However, this is expected as the optimal control problem only optimizes against a simplified version of the true dynamical system.



Figure 3.3: NMPC case-study results with nonlinear observations with additive Gaussian noise.

3.4.2 Nonlinear and non-Gaussian likelihood function

In the second case-study, it is assumed that the observational model is nonlinear and with non-additive Gaussian noise. The likelihood of an observation y_k at time t_k is given by

$$d_{k,\text{lik}}(x) = \left(\frac{\gamma(x_1(t_k))^{y_{k,1}}}{y_{k,1}!} \exp(-\gamma(x_1(t_k)))\right) \left(\frac{\gamma(x_2(t_k))^{y_{k,1}}}{y_{k,2}!} \exp(-\gamma(x_2(t_k)))\right). \quad (3.27)$$

This likelihood function is the conditional density of y_k given the hidden state $x(t_k)$. Note that y_k attains only discrete values. The γ function in (3.27) is defined according to

$$\gamma(x) = 10 \log(1 + \exp(x)). \tag{3.28}$$

It is assumed that the sampling time is $t_k - t_{k-1} = 1/20$ between any observations. The regularization coefficient is chosen such that $\varepsilon = 1/1000$.

In Fig. 3.4, a simulation using the NMPC principle, to track the set-point trajectory shown in Fig. 3.2, is shown. Again, it is clearly possible to track \overline{x}_1 with the true state state, x_1 . It is important to note that this is only possible using the MAP method defined in (2.39) for state estimation; the CDEKF does not accommodate observational models of the form (3.27).



Figure 3.4: NMPC case-study results with nonlinear observations and non-Gaussian noise.

4 The Northern European power market

The main field of application in this dissertation has been energy systems. This chapter will introduce the Northern European power market and discuss which opportunities that exist for market participants which operate dynamical systems under (potentially) time-varying power prices. Hence, while the previous two chapters were dedicated to the theoretical and computational aspect of nonlinear model predictive control algorithms, this chapter will bridge nonlinear model predictive control with applications relating to energy systems. The power prices (and the power markets under consideration) will define the structure of the objective function (3.1a) in the optimization problems used to operate the dynamical energy systems.

The Northern European power market consists of many markets which



Figure 4.1: Nord Pool price zone topology.

all yield different opportunities for market participants. However, this chapter will focus on three of these markets:

- the day-ahead market,
- the intra-day market,
- and the regulating power market.

Thus, it is only the markets trading with the shortest horizons that are considered, and the financial contracts traded on e.g. the NASDAQ commodities exchange (with horizons of weeks, months and years) are omitted from discussion in this dissertation. The day-ahead and intra-day markets are operated and cleared by Nord Pool [45]. Nord Pool's market clearing algorithm for the day-ahead market is responsible for setting the day-ahead (or spot) price in 15 inter-connected price zones as shown in the topology-map in Fig. 4.1. The Nordic Transmission System Operators (TSOs) in Denmark (Energinet), Norway (Statnett), Sweden (Svenska Kraftnätt) and Finland (Fingrid) are responsible for managing and clearing the regulating power market. In Fig. 4.2, the different horizons and key settlement time-points are shown; this figure should be used as a reference in the following sections. The following sections, describing the three power markets listed above, are inspired by and follow the power market descriptions given in Paper C with some visual modifications. Details concerning the day-ahead and intra-day markets are given in [45] and an elaborate introduction to the regulating power market is given in [5].



Figure 4.2: Different horizons and market closures in the Northern European power market.

4.1 The day-head market

The day-ahead market has the largest trading volume and electricity can be purchased and sold for the up-coming day. Bidding to the day-ahead market closes at 12pm on the day prior to the day of operation (as shown in Fig. 4.2), and hence the planning horizon is 12-36 hours. Participants can choose between a wide range of bidding-types; e.g. block-bids, where bids for multiple hours can be aggregated into one combined bid, or traditional single-hourbids, where bids only are valid for the given hour. The spot price is settled as the intersection between the ordered bids of power purchases and sales, constrained by the physical limitations of the power grid and the possible constraints of the variety of bids available to the participants. It is voluntary for market participants to participate in the day-ahead market.

4.2 The intra-day market

The intra-day market is open for trading when the day-ahead market closes. This market has traditionally been used to eliminate imbalances between the expected power production/consumption and the power production/consumption plans traded in the day-ahead market. The trading volume in the intra-day market is significantly smaller than the trading volume in the day-ahead market. The intra-day market closes 45 minutes prior to the hour of operation but is open for trading up until then. Fig. 4.3 compares the trading volume in DK1 of the day-ahead, intra-day and regulating power markets during the first 11 months of 2020; as stated, the trading volume of the intra-day market is significantly smaller than the day-ahead market volume.

The power prices settled in the day-ahead and intra-day markets have the least volatile prices of the three markets considered in this chapter. However, these markets play a key role in optimal



power market speculation [46] as all three markets typically clear very different power prices; thereby create interesting arbitraging opportunities.

4.3 The regulating power market

Simultaneously with the closure of the intra-day market, bidding to the regulating power market also closes. The regulating power market is a market operated and settled by the TSOs in the Nordic countries and all regulating power bids are aggregated within the Nordic Operational Information System (NOIS). The regulating power market is one (of many) mechanism the TSOs use to maintain grid stability. The regulating power market is relevant if the grid imbalance (a discrepancy between power production, consumption and net import) is expected to be of significant volume and duration. Given that the system imbalance is expected to be significant, the sign of the imbalance define:

- ↑ If the imbalance is negative, there is a deficit of electricity - an increase of the production or a decrease in the consumption is needed. This is called up-regulation.
- ↓ If the imbalance is positive, there is a surplus of electricity
 a decrease of the production or an increase in the consumption is needed. This is called down-regulation.

The TSOs work as intermediaries in the regulating power market. The TSOs sell electricity when there is up-regulation and buy electricity when there is down-regulation. The counter parties responsible for these imbalances are settled in the balancing power market. In this market the TSOs also function as intermediaries.



Figure 4.4: Regulating power prices in DK1 during the first 11 months of 2020.

The balancing power market is settled after the operating hour when the actual meter measurements are available. Thus, the regulating power market works as a neutralization mechanism to alleviate imbalances and the balancing power market is a physical market which holds the responsible parties accountable for their grid imbalances.

The regulating power market functions, similarly to the day-ahead market, across 12 of the 15 interconnected price zones shown in Fig. 4.1 (Estonia, Lithuania and Latvia are not included in the Nordic regulating power market). The regulating power price might differ between price zones but is always settled as the intersection between the ordered bids of down- and up-regulation and the activated regulating power volume. Hence, all activated bids are activated at the same price. The day-ahead price is used as the basis for the regulating- and balancing power pricing procedures:

- The down-regulation price is less than or equal to the dayahead price.
- The up-regulation price is greater than or equal to the dayahead price.
- Consumption imbalances with an opposite sign of the system imbalance is settled at the regulation power price in the balancing power market. Production imbalances with an opposite sign of the system imbalance is settled at the day-ahead price.
- Imbalances with the same sign as the system imbalance is settled at the regulating power price in the balancing power market.

This structure implies that consumption imbalances are settled at the regulating power prices in the balancing power market, while production imbalances are not. This settlement structure for production imbalances implies a net profit for the TSOs. This profit is used to partially finance the reserve capacity and thereby implicitly lowering the system tariff [5]. This pricing procedure yields hourly regulating power prices.

In Fig. 4.4, a distribution of the observed regulating power prices are shown. The regulating power prices are measured relative to the day-ahead prices, since the day-ahead prices are used as a reference-point in the regulating power market. This means that negative prices in Fig. 4.4 imply down-regulation and positive prices imply up-regulation. The prices shown in Fig. 4.4 are the prices observed in DK1 (the Western Danish price zone) during the first 11 months of 2020. The right-axis shows the level of congestion of the Skagerak transmission line, connecting DK1 to NO2 (a Southern Norwegian price zone). From Fig. 4.4 it is seen, that when the transmission line between DK1 and NO2 has spare capacity, the regulating power prices become very competitive; they only deviate slightly from the spot prices. This is because NO2 (often) has a large capacity of hydro-power [47] which has very low marginal power costs. On the other hand, when this transmission line is congested, the regulating power prices show larger deviations from the spot prices. This is due to less liquidity of the regulating power market in DK1, when NO2 is unable to provide regulation to DK1.

4.3.1 Special regulation

In recent years there has been a rapid development of the activated volumes of a market feature called special regulation in DK1. This is also observed from Fig. 4.3; the demand for special down-regulation has been the largest source of regulation during the first 11 months of 2020 in DK1. Special regulation can occur as both down- and up-regulation. Special regulation is in effect when the TSOs skip the usual pricing procedure and in the event of special regulation, the activated bids are paid-as-bid.

During the last three years (2017-2019), one of the main sources of special down-regulation in DK1 has been excessive wind power production in Northern Germany. In Tab. 4.1 the total special downregulation volumes activated in DK1, due to special regulation from the German TSO, TenneT, is shown. This table shows that in 2019 a total of 1,312 GWh was activated for special down-regulation in DK1 due to large (positive) imbalances in Northern Germany. This table also shows that in 2019, 32 % of the special regulation was handled by curtailment of wind turbines in Denmark. Energinet is only allowed to publish yearly average prices; in 2019 the average special down-regulation price was -12.3 EUR/MWh.

In general, Tab. 4.1 shows that the need for special down-regulation in DK1 is increasing with corresponding power prices being vary attractive for consumers. Thus, there is a great potential for flexible consumers to participate in this market and ideally eliminate the need of curtailing Danish wind turbines.

Table 4.1: GWh of Special down-regulation received from the German TSO, TenneT, for 2017-2019. The activated volumes are activated volumes for Danish participants and the average price is in EUR/MWh.

	2019	2018	2017
Special regulation in GWh	1,312	1,114	781
Thermal power plants	46 %	53 %	64 %
Electric boilers	22 %	21 %	22 %
Wind turbines	32 %	26 %	14 %
Average price in EUR/MWh	-12.3	-9.3	-7.7

5 Applications in energy systems

This dissertation presents 6 papers (3 acceptet, 2 submitted and 1 in preparation); 4 of these are applications in energy systems and 2 are methods relevant for general nonlinear model predictive control. Paper A provides a tutorial introduction for formulating the entire algorithm-stack for continuous-discrete systems governed by SDEs as presented in Section 3.4.2 and Section 3.4.1. Paper F presents the theoretical contribution shown in Section 2.2.3 where an optimization-based filtering method is introduces and compared to the continuous-discrete extended Kalman filter. This chapter is dedicated to summarizing the key findings and contributions of the 4 other papers, dedicated to application in energy systems.

5.1 Optimal operation of wastewater treatment plants

Paper B and Paper C both consider the operational challenge of cost-efficient operation of a wastewater treatment plant. Paper B states that approximately 1 % of a country's total electricity consumption is due to wastewater treatment and Paper C explains that the aeration process is responsible for approximately 40 % to 75 % of the total power consumption at a treatment plant. Fig. 5.1 shows the schematics of the control strategies implemented in Paper B and Paper C; the main goal is to reduce the nutrient concentrations of e.g. ammonium (NH₄) and nitrate (NO₃) while also minimizing the power consumption due to aeration. The aeration process is a switching process, meaning that it can be switched on or off. Paper B and Paper C uses the same two-dimensional SDE to model the response of the nutrient concentrations (the states, x) to the aeration switching process (the input, u).



Figure 5.1: This figure is taken from Paper C. This figure shows a schematic overview of the components in the optimal control problem and how the energy flexibility is unlocked by a coupling to power markets. NOIS represents participation in the regulation power market.

Paper B introduces a concept framed as a flexibility-diagram that describes how the power consumption responds to changing power prices; a similar concept is introduced in e.g. [2]. Paper C conducts a backtest which investigates the performance using dayahead and regulating power prices compared to using a constant power price. This indicates how much flexibility that is available at wastewater treatment plants. Fig. 5.2 demonstrates 24 hours of



Figure 5.2: This figure is taken from Paper C. Example day from the backtest: 20th March 2019. The regulating power price model with $p_s=0$ EUR/MWh (price assumption when there is special regulation) is used to design the optimal aeration cycles. The top plot shows the price input to the price models: the left y-axis measures the demand for special down regulation and the right y-axis measures the day-ahead and regulating power prices. The middle plot shows the state realization using the aeration sequence shown in the same plot. This is illustrated for the regulating power and special regulating price model and the baseline with constant electricity price (dashed lines). The bottom plot shows the savings relative to the constant price model.

simulated operation using the control strategy defined in Paper C. The top plot shows the power market data, the middle plot shows the simulated state values and optimal input (output from an optimal control problem) and the bottom plot shows the accumulated savings compared to operating against a constant power price. The main conclusion from Paper B and Paper C is that high variations of the the power prices used in design of the optimal input, u, leads to larger economical savings, compared to optimization against a constant power price. Cost-savings of up to 40 % are observed in the backtest simulations.

5.2 Optimal operation of a joint ice-tank and refrigeration system

Paper D presents a method for operating the joint refrigeration system consisting of a small retail refrigeration system and an ice-tank. The test-system is physically located at Danfoss' testcenter in Nordborg, Denmark. Fig. 5.3 shows a schematic overview of the test-system; the ice-tank is added in the return loop from the refrigerated racks via a sub-cooler. The user can choose between three states of the ice-tank: idle, charging and discharging. The idle-mode is defined as a no-action mode, where the ice-tank is staying idle. In the charging-mode, the ice-tank starts accumulating ice. In the discharging-mode, the ice-tank starts curtailing the refrigeration system by melting ice and lowering the temperature of the return refrigerant. For many refrigeration systems, the outdoor temperature is very significant for the performance; to eliminate this feature and only consider the effect of operating an ice-tank, the outdoor temperature is simulated using a water heater with a fixed temperature.

A first-order and one-dimensional SDE is proposed as model of the response of the compressor capacity, from the refrigeration system, to switching modes of the ice-tank. The model is calibrated using 13 hours of training-data by applying the parameter estimation method described in Section 2.3. Fig. 5.4 shows the cost-savings



Figure 5.3: This figure is taken from Paper D. This figure shows an overview of the total refrigeration system used to test the performance of optimally controlling an ice-tank.



Figure 5.4: This figure is taken from Paper D. This figure visualizes the effect of having the ice-tank installed. The top plot shows the accumulated cost of actively operating the ice-tank vs. not having the ice-tank. The bottom plot shows the relative savings of operating the ice-tank. The bottom plot shows how the predicted savings compare to the realized savings. It is assumed that the ice-tank has been charged with zero cost.

of optimally curtailing the refrigeration system using the icetank compared to not having an ice-tank available. During the 24 hours of online tests, the control strategy shows savings of up-to 20 %. The main conclusion from Paper D is that an ice-tank is a non-invasive method for curtailing the refrigeration system during periods with high power prices and thereby reducing the operational costs of operating the joint ice-tank and refrigeration system.

5.3 Optimal energy arbitraging using a Vanadium redox-flow battery

Paper E presents a method for optimal energy arbitraging across multiple power markets using a Vanadium redox-flow battery. This paper tests different configurations of the nominal power to physical storage ratio and investigates the importance of high-quality forecasts of the regulating power market. The paper also investigates the influence of CO2-taxes in energy arbitraging.

Fig. 5.5 and Fig. 5.6 present two interesting dependencies. Fig. 5.5 shows the relation of the quality of forecasts and the possible revenues generated by multi-market energy arbitraging. Testing different assumptions on the accuracy of forecasts on the regulating power market; it is shown that small forecasting deficiencies aren't too critical. This deficiency parameter models the hourly decay rate of the forecasted regulating power prices to the day-ahead prices. Paper E uses these simulated forecasts to evaluate the importance of the forecasting capabilities on the regulating power market. Fig. 5.6 shows a relation between increasing CO2-taxes and revenues generated by arbitraging power markets. This figure shows that it is possible to heavily increase CO2 reduc-

tions without affecting revenues too much. These figures show a Pareto-like structure [48] as e.g. used in portfolio theory in finance where risk-adjusted metrics are used to quantify portfolio performance (e.g. the efficient frontier generated by the Sharpe ratio [49, 50]).



Figure 5.5: This figure is taken from Paper E. This figure shows the relation between the quality of the forecasts (shown along the first-axis) and the revenues generated by arbitraging multiple power markets.



Figure 5.6: This figure is taken from Paper E. This figure shows the dependency between CO2-taxes (shown along the first-axis) and the revenues generated by arbitraging multiple power markets. Each point in the figure represents a unique CO2-tax parameter (with units EUR/t).

6 Summary and future perspectives

The main focus of this dissertation has been centered around formulating nonlinear model predictive control problems for costefficient operation of energy systems. The main conclusion is that nonlinear model predictive control algorithms present a huge opportunity to optimize and coordinate the operation of energy systems.

In Paper B and Paper C, it is demonstrated how the operation of aeration equipment at wastewater treatment plants can be optimized and coordinated according to time-varying power prices, thereby reducing the operational costs of removing nutrient from the continuousflow of incoming wastewater.

Paper D demonstrates how a joint refrigeration system consisting of an ice-tank and a small retail refrigeration system can be coordinated and optimized according to time-varying power prices to reduce the aggregated electricity costs of this combined refrigeration system.

Paper E presents a method for optimal arbitraging in multiple power markets and across different operating hours. It is shown how a multi-objective criteria can be used to consider the simultaneous optimization of maximizing revenues and minimizing CO2 emissions; this approach will become increasingly relevant as governments implement taxation of greenhouse gas emissions.

Paper A shows a complete tutorial on how to formulate the entire algorithm-stack from nonlinear model predictive control methods based on continuous discrete models driven by stochastic differential equations (2.1) with discrete observations (2.3). Paper F condenses the work shown in Section 2.2.3 and describes a new filtering technique for state estimation for stochastic differential equation models with discrete-time observational models.

The following sections discuss future perspectives of the methods and applications considered in this dissertation.

6.1 Advances in state and parameter estimation

Section 2.2.3 introduces an optimization-based formulation of the continuous-discrete extended Kalman filter based on maximum a posteriori estimation. Using this alternative formulation, it is possible to implement state estimation on more general observation models given by conditional density (or likelihood) functions; an example of this is shown in Section 3.4.2. In Section 2.2.4, a comparison between the optimization-based filter method and the continuous-discrete extended Kalman filter method is shown. The main conclusion is that the optimization-based filter method significantly outperforms the continuous-discrete extended Kalman filter extended Kalman filter in the ability to efficiently reconstruct the hidden state modelled by the stochastic differential equation. This work is currently being condensed in Paper F and will be submitted in the beginning of 2021.

[29] describes a method for parameter estimation of parameters embedded in stochastic differential equation models of the form (2.1) with (possibly nonlinear) additive Gaussian observational models as given according to (2.3). This method is also briefly described in Section 2.3 with an example given in Section 2.3.1. Given the simulation study from Section 2.2.4, it would be very interesting to develop a parameter estimation framework using the optimization-based filtering technique from Section 2.2.3 and consider state and parameter estimation simultaneously in a joint optimization problem. [51] defines a similar framework, but for more general statistical models (nonlinear random effect models), using the Laplace approximation.

6.2 Forecasting the regulating power market

In Paper B, Paper C, Paper D and Paper E, one of the key assumptions is very good forecasting capabilities of the regulating power market. This market is described in Section 4.3. However, looking at the current state-of-the-art in the scientific literature, little work exists concerning accurate forecasting of the regulating power market in Northern Europe. The work that has been conducted is given in [52], [53] and [54], where black-box models (e.g. autoregressive and moving-average models with exogenous regressors) are used to forecast the regulating power prices. However, one critical draw-back of these methods is the inability to accurately model the effect of congestion and speculation without introducing too many model parameters. Thus, to efficiently model the prices cleared in the regulating power market, a more physical representation of the market is needed.

6.3 Portfolio optimization of energy units

The applications considered in this dissertation have been focused on optimizing the operation of individual energy units (e.g. one wastewater treatment plant, or one refrigeration system). However, in an efficient implementation of energy systems in the future low-carbon society, coordination of multiple energy consuming (or producing) units is necessary to fully unlock the flexibility of the joint portfolio of these units. The concept of portfolio optimization of energy units have been investigated in the literature. [55] considers the problem of coordinating the power consumption of multiple data-centers according to time-varying power prices. [56] describes how a distributed collection of model predictive control algorithms can be used to optimize the joint power consumption of a portfolio of flexible consumers. [57] defines an efficient numerical algorithm for solving optimization problems arising from the joint coordination of model predictive control algorithms. The proposed framework in [58] and [59] demonstrates control-based methods that adopt price signals as the economic driver to stimulate consumers' response and to unlock the flexibility needed to balance largescale integration of fluctuating renewable power production. In this framework, various system operators broadcast price signals independently to fulfil their requirements for ancillary services.

However, little work exists within the intersection of distributed systems governed by stochastic differential equation systems and optimal decision-making hereof. This dissertation has only focused on the application of nonlinear model predictive control for solving the more general stochastic control problem defined initially in (1.2). A different approach would be to consider a scenariobased method and apply concepts from stochastic programming [60]. Stochastic differential equations provide a simple method for generating a batch of scenarios; these scenarios can then be used to implement principles from stochastic programming. This would also involve formulating the objective function using statistical criteria; e.g. the conditional-value-at-risk [61] measure or the Sharpe ratio [62] - to optimize a risk-adjusted criteria. This opens the discussion of formulating efficient algorithms to iteratively find the optimal input to the stochastic system. One interesting method to consider is the application of stochastic gradient methods [63] where the scenario-batch is altered in very iteration of the optimization-process. Stochastic optimization algorithms are implemented in e.g. PyTorch [64], which is framework for machine-learning (e.g. deep-learning, natural language processing and reinforcement learning) in python.

Bibliography

- [1] "Denmark Strikes Deal to Slash CO2 Emissions by 70% in a Decade". In: Bloomberg (2019).
- [2] Rune Grønborg Junker, Armin Ghasem Azar, Rui Amaral Lopes, Karen Byskov Lindberg, Glenn Reynders, Rishi Relan, and Henrik Madsen. "Characterizing the energy flexibility of buildings and districts". In: Applied Energy 225 (2018), pp. 175-182.
- [3] Clay Campaigne and Shmuel S Oren. "Firming renewable power with demand response: an end-to-end aggregator business model". In: Journal of Regulatory Economics 50.1 (2016), pp. 1-37.
- [4] Henrik Madsen, Jacopo Parvizi, Rasmus Fogtmann Halvgaard, Leo Emil Sokoler, John Bagterp Jørgensen, Lars Henrik Hansen, and Klaus Baggesen Hilger. "Control of Electricity Loads in Future Electric Energy Systems". In: Handbook of Clean Energy Systems (2015).
- [5] Energinet Website (market regulations). https://energinet.dk/El/Elmarkedet /Regler-for-elmarkedet/Markedsforskrifter. Accessed: 2020-12-03.
- [6] Andrew H Jazwinski. Stochastic Processes and Filtering Theory. San Diego, CA, USA: Academic Press, 1970.
- [7] Bernt Øksendal. Stochastic differential equations: an introduction with applications. Springer Science & Business Media, 2013.
- [8] Robert F Stengel. "Stochastic optimal control: theory and application." In: New York (1986).
- [9] Karl J Åström. Introduction to stochastic control theory. Courier Corporation, 2012.
- [10] Eduardo F Camacho and Carlos Bordons Alba. Model predictive control. Springer Science & Business Media, 2013.
- [11] James Blake Rawlings and David Q Mayne. *Model predictive* control: Theory and design. Nob Hill Pub., 2009.
- [12] Frank Allgöwer and Alex Zheng. Nonlinear model predictive control. Vol. 26. Birkhäuser, 2012.
- [13] PA Stentoft, L Vezzaro, PS Mikkelsen, M Grum, T Munk-Nielsen, P Tychsen, H Madsen, and R Halvgaard. "Integrated model predictive control of water resource recovery facilities and sewer systems in a smart grid: example of full-scale implementation in Kolding". In: Water Science and Technology (2020).
- [14] Peter A Stentoft, Daniela Guericke, Thomas Munk-Nielsen, Peter S Mikkelsen, Henrik Madsen, Luca Vezzaro, and Jan K Møller. "Model predictive control of stochastic wastewater treatment process for smart power, cost-effective aeration". In: Ifac-papersonline 52.1 (2019), pp. 622-627.

- [15] Niclas Brabrand Brok, Thomas Munk-Nielsen, Henrik Madsen, and Peter A Stentoft. "Flexible Control of Wastewater Aeration for Cost-Efficient, Sustainable Treatment". In: IFAC-PapersOnLine 52.4 (2019), pp. 494-499.
- [16] Niclas Brabrand Brok, Thomas Munk-Nielsen, Henrik Madsen, and Peter A Stentoft. "Unlocking energy flexibility of municipal wastewater aeration using predictive control to exploit price differences in power markets". In: Applied Energy 280 (2020), p. 115965.
- [17] Desmond J Higham and Lloyd N Trefethen. "Stiffness of odes". In: BIT Numerical Mathematics 33.2 (1993), pp. 285-303.
- [18] Niclas Laursen Brok, Henrik Madsen, and John Bagterp Jørgensen. "Nonlinear model predictive control for stochastic differential equation systems". In: *IFAC-PapersOnLine* 51.20 (2018), pp. 430-435.
- [19] Rudolph Emil Kalman. "A new approach to linear filtering and prediction problems". In: (1960).
- [20] Eric A Wan and Rudolph Van Der Merwe. "The unscented Kalman filter for nonlinear estimation". In: Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373). IEEE. 2000, pp. 153-158.
- [21] Rudolph Van Der Merwe, Arnaud Doucet, Nando De Freitas, and Eric A Wan. "The unscented particle filter". In: Advances in neural information processing systems. 2001, pp. 584-590.
- [22] Erik Lindström, Edward Ionides, Jan Frydendall, and Henrik Madsen. "Efficient iterated filtering". In: IFAC Proceedings Volumes 45.16 (2012), pp. 1785-1790.
- [23] René Schneider and Christos Georgakis. "How to NOT make the extended Kalman filter fail". In: Industrial & Engineering Chemistry Research 52.9 (2013), pp. 3354-3362.
- [24] Dimitri Boiroux, Rune Juhl, Henrik Madsen, and John Bagterp Jørgensen. "An efficient UD-based algorithm for the computation of maximum likelihood sensitivity of continuous-discrete systems". In: 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE. 2016, pp. 3048-3053.
- [25] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex* optimization. Cambridge university press, 2004.
- [26] Max A Woodbury. "Inverting modified matrices". In: Memorandum report 42.106 (1950), p. 336.
- [27] Max Welling. "The kalman filter". In: Lecture Note (2010), pp. 92-117.
- [28] Bradley Efron and David V Hinkley. "Assessing the accuracy of the maximum likelihood estimator: Observed versus expected Fisher information". In: *Biometrika* 65.3 (1978), pp. 457-483.
- [29] Niels Rode Kristensen, Henrik Madsen, and Sten Bay Jørgensen. "Parameter estimation in stochastic grey-box models". In: Automatica 40.2 (2004), pp. 225-237.

- [30] William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. Numerical recipes 3rd edition: The art of scientific computing. Cambridge university press, 2007.
- [31] Richard Bellman. "Dynamic programming". In: Science 153.3731 (1966), pp. 34-37.
- [32] Lev Semenovich Pontryagin. Mathematical theory of optimal processes. Routledge, 2018.
- [33] Donald E Kirk. Optimal control theory: an introduction. Courier Corporation, 2004.
- [34] Hans Georg Bock and Karl-Josef Plitt. "A multiple shooting algorithm for direct solution of optimal control problems". In: IFAC Proceedings Volumes 17.2 (1984), pp. 1603-1608.
- [35] Moritz Diehl, Hans Georg Bock, Holger Diedam, and P-B Wieber. "Fast direct multiple shooting algorithms for optimal robot control". In: Fast motions in biomechanics and robotics. Springer, 2006, pp. 65-93.
- [36] Oskar Von Stryk. "Numerical solution of optimal control problems by direct collocation". In: Optimal control. Springer, 1993, pp. 129-143.
- [37] John Bagterp Jørgensen. "Adjoint sensitivity results for predictive control, state and parameter estimation with nonlinear models". In: 2007 European Control Conference (ECC). IEEE. 2007, pp. 3649-3656.
- [38] W Grimm and A Markl. "Adjoint estimation from a direct multiple shooting method". In: *Journal of optimization theory and applications* 92.2 (1997), pp. 263-283.
- [39] Andreas Wächter and Lorenz T Biegler. "On the implementation of an interior-point filter line-search algorithm for largescale nonlinear programming". In: *Mathematical programming* 106.1 (2006), pp. 25-57.
- [40] Karsten Ahnert and Mario Mulansky. "Odeint-solving ordinary differential equations in C++". In: AIP Conference Proceedings. Vol. 1389. 1. American Institute of Physics. 2011, pp. 1586-1589.
- [41] Alan C Hindmarsh, Peter N Brown, Keith E Grant, Steven L Lee, Radu Serban, Dan E Shumaker, and Carol S Woodward. "SUN-DIALS: Suite of nonlinear and differential/algebraic equation solvers". In: ACM Transactions on Mathematical Software (TOMS) 31.3 (2005), pp. 363-396.
- [42] Lorenz T Biegler and Victor M Zavala. "Large-scale nonlinear programming using IPOPT: An integrating framework for enterprise-wide dynamic optimization". In: Computers & Chemical Engineering 33.3 (2009), pp. 575-582.
- [43] David Benson. "A Gauss pseudospectral transcription for optimal control". PhD thesis. Massachusetts Institute of Technology, 2005.
- [44] Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. "CasADi - A software framework for nonlinear optimization and optimal control". In: *Mathematical Programming Computation* 11.1 (2019), pp. 1-36.

- [45] Nord Pool AS Website (trading introduction). https://www.nordpoolgroup.com/trading. Accessed: 2020-12-03.
- [46] Alexander W Dowling, Ranjeet Kumar, and Victor M Zavala. "A multi-scale optimization framework for electricity market participation". In: Applied energy 190 (2017), pp. 147-164.
- [47] Frode Kjaerland. "A real option analysis of investments in hydropower-The case of Norway". In: Energy Policy 35.11 (2007), pp. 5901-5908.
- [48] Yair Censor. "Pareto optimality in multiobjective problems". In: Applied Mathematics and Optimization 4.1 (1977), pp. 41-59.
- [49] Harry M Markovitz. Portfolio selection: Efficient diversification of investments. John Wiley, 1959.
- [50] William F Sharpe. "Capital asset prices: A theory of market equilibrium under conditions of risk". In: The journal of finance 19.3 (1964), pp. 425-442.
- [51] Kasper Kristensen, Anders Nielsen, Casper W Berg, Hans Skaug, and Brad Bell. "TMB: automatic differentiation and Laplace approximation". In: arXiv preprint arXiv:1509.00660 (2015).
- [52] Tryggvi Jónsson. "Forecasting of electricity prices accounting for wind power predictions". PhD thesis. Technical University of Denmark, 2008.
- [53] Magnus Olsson and Lennart Soder. "Modeling real-time balancing power market prices using combined SARIMA and Markov processes". In: IEEE Transactions on Power Systems 23.2 (2008), pp. 443-450.
- [54] Stefan Jaehnert, Hossein Farahmand, and Gerard L Doorman. "Modelling of prices using the volume in the Norwegian regulating power market". In: 2009 IEEE Bucharest PowerTech. IEEE. 2009, pp. 1-7.
- [55] Mahdi Ghamkhari, Adam Wierman, and Hamed Mohsenian-Rad. "Energy portfolio optimization of data centers". In: IEEE Transactions on Smart Grid 8.4 (2016), pp. 1898-1910.
- [56] Rasmus Halvgaard, Lieven Vandenberghe, Niels Kjølstad Poulsen, Henrik Madsen, and John Bagterp Jørgensen. "Distributed model predictive control for smart energy systems". In: IEEE Transactions on Smart Grid 7.3 (2016), pp. 1675-1682.
- [57] Kristian Edlund, Jan Dimon Bendtsen, and John Bagterp Jørgensen. "Hierarchical model-based predictive control of a power plant portfolio". In: *Control Engineering Practice* 19.10 (2011), pp. 1126-1136.
- [58] Giulia De Zotti, S. Ali Pourmousavi, Henrik Madsen, and Niels Kjølstad Poulsen. "Ancillary Services 4.0: A Top-To-Bottom Control-Based Approach for Solving Ancillary Services Problems in Smart Grids". eng. In: *Ieee Access* 6 (2018), pp. 11694-11706.
- [59] Giulia De Zotti, Seyyed Ali Pourmousavi Kani, Juan M. Morales, Henrik Madsen, and Niels Kjølstad Poulsen. "A Control-based Method to Meet TSO and DSO Ancillary Services Needs by Flex-

ible End-Users". eng. In: Ieee Transactions on Power Systems 35.3 (2020), pp. 1868-1880.

- [60] Peter Kall, Stein W Wallace, and Peter Kall. *Stochastic pro*gramming. Springer, 1994.
- [61] R Tyrrell Rockafellar, Stanislav Uryasev, et al. "Optimization of conditional value-at-risk". In: Journal of risk 2 (2000), pp. 21-42.
- [62] William F Sharpe. "The sharpe ratio". In: Journal of portfolio management 21.1 (1994), pp. 49-58.
- [63] Léon Bottou. "Large-scale machine learning with stochastic gradient descent". In: Proceedings of COMPSTAT'2010. Springer, 2010, pp. 177-186.
- [64] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. "Automatic differentiation in PyTorch". In: (2017).

Part II

Publications



Available online at www.sciencedirect.com





IFAC PapersOnLine 51-20 (2018) 430-435

Nonlinear Model Predictive Control for Stochastic Differential Equation Systems

Niclas Laursen Brok, Henrik Madsen, John Bagterp Jørgensen

Department of Applied Mathematics and Computer Science, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

Abstract: Using the Van der Pol oscillator model as an example, we provide a tutorial introduction to nonlinear model predictive control (NMPC) for systems governed by stochastic differential equations (SDEs) that are observed at discrete times. Such systems are called continuous-discrete systems and provides a natural representation of systems evolving in continuous-time. Furthermore, this representation directly facilities construction of the state estimator in the NMPC. We provide numerical details related to systematic model identification, state estimation, and optimization of dynamical systems that are relevant to the NMPC.

© 2018, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.

Keywords: Nonlinear Model Predictive Control, Continuous-Discrete Extended Kalman Filter, Maximum Likelihood Estimation, Stochastic Differential Equations, Van der Pol Oscillator

1. INTRODUCTION

This paper provides a tutorial of how to use the nonlinear model predictive control (NMPC) principle to regulate a stochastic system governed by stochastic differential equations (SDEs). The systems considered in this paper are continuous-discrete systems of the form (Jazwinski, 1970)

$$d\boldsymbol{x}(t) = f(\boldsymbol{x}(t), u(t); p) dt + g(\boldsymbol{x}(t), u(t); p) d\boldsymbol{\omega}(t), \quad (1a)$$

$$\boldsymbol{y}(t_k) = h(\boldsymbol{x}(t_k)) + \boldsymbol{v}_k, \quad (1b)$$

where \boldsymbol{x} , \boldsymbol{u} and \boldsymbol{p} are the states, inputs and time-invariant parameters. $\boldsymbol{v}_k \sim N_{\rm iid}(0, R_k)$ is the measurement noise and $\boldsymbol{\omega}(t)$ is a standard Brownian motion. Brownian motion is defined by its independent increments which satis first that for each $s, t \in \mathbb{R}$, $\omega(t) - \omega(s)$ is normally distributed with zero mean and covariance I(t-s); i.e. $d\boldsymbol{\omega}(t) \sim N_{\rm iid}(0, I \, dt)$. The SDE model representation (1a) provides a natural way to represent physical systems as they evolve in continuous-time. In contrast to discretetime models, a priori knowledge about the system can be included and the estimated parameters do not depend on the sampling time. The representation of noise in continuous time also allow for a parsimonious representation that is independent of the sampling time. While these advantages of the continuous-discrete model (1) are wellknown in the systems identification community (Garnier and Young, 2012; Kristensen et al., 2004; Rao and Unbehauen, 2006), most NMPC methods rely on either 1) a deterministic discrete-time model, 2) a stochastic discretetime model, or 3) a deterministic continuous-time model for which the noise terms in the estimators are added in an ad hoc manner. The key insight is that the continuousdiscrete model (1) provides a systematic way to obtain an estimation (filtering and prediction) algorithm that is used in the offline system identification, the online stateand parameter-estimation, and the prediction of the dynamic optimization. The diffusion model, $q(\boldsymbol{x}(t), u(t); p)$, represents a convenient and powerful way of representing complex stochastic processes and model-plant mismatch as needed for the filtering and prediction algorithm in NMPC. Boiroux et al. (2016a,b,c, 2010) and Mahmoudi et al. (2016a, 2017, 2016b) demonstrate systematic use of the continuous-discrete model (1) for system identification, nonlinear filtering and prediction, fault detection, and NMPC in an artificial pancreas for people with type 1 diabetes. Mahmood and Mhaskar (2012) uses a continuousdiscrete model (1) and a Lyapunov-based NMPC to stabilize a reaction model around an unstable equilibrium. Buehler et al. (2016) uses a similar setup where a bioreactor model is controlled according to a desired set-point probability density. Jørgensen and Jørgensen (2007a,b) apply the continuous-discrete stochastic model (1) for linear MPC using transfer function representations of the model.

1.1 Components and software of the NMPC

Allgöwer et al. (1999), Johansen (2011), Grüne and Pannek (2011), and Rawlings et al. (2017) describe state-ofthe-art NMPC technology. A system for NMPC consists of an offline method for identification of the model as well as an online part - where the online part consists of a state- and parameter estimation algorithm and an algorithm for dynamic optimization. Fig. 1 schematically illustrates this structure. This tutorial fills a missing gap in existing NMPC literature, by systematically formulating all components in the NMPC software system based on the continuous-discrete model (1). All components in the NMPC, presented in this paper, use the same continuousdiscrete extended Kalman filter (CDEKF) for (1) in the one-step prediction of the offline system identification, in the online state- and parameter-estimation algorithm, and in the prediction of the dynamic optimization algorithm.

^{*} Innovation Fund Denmark is acknowledged for partly funding the work as a part of Center for IT-Intelligent Energy Systems (CITIES - IFD 1035-00027B).

^{2405-8963 © 2018,} IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved. Peer review under responsibility of International Federation of Automatic Control. 10.1016/j.ifacol.2018.11.071



Fig. 1. Overview of the closed-loop control structure.

The rigorous solution to the filtering and prediction problem is obtained by solving the Fokker-Planck equations (Jazwinski, 1970). However, for systems with more than a couple of states, this solution is computationally intractable. The CDEKF is a computationally tractable alternative for filtering and prediction in (1). While other filters such as the unscented Kalman filter (UKF), the ensemble Kalman filter (EnKF), the particle filter (PF), and the moving horizon estimator (MHE) can also be used instead of the (CDEKF), the CDEKF represents the best balance between performance and computational tractability for many processes (Simon, 2006). This is particularly true, when the maximum-likelihood method is used for estimation of the parameters in the filter and predictor. For good performance of the CDEKF, it is also important to notice that we implement it using a differential equation solver with adaptive time step and using the Joseph stabilization scheme (Schneider and Georgakis, 2013). By including a disturbance model as part of the model, the CDEKF is used for online estimation of the states as well as selected rapidly varying parameters.

The offline system identification is based on a maximum likelihood (ML) formulation, where the conditional densities of the state equations are approximated by Gaussian densities. Using this assumption, it is possible to derive an optimization problem which uses the CDEKF to compute the likelihood of the parameters, p, given a set of observations (Kristensen et al., 2004).

The dynamic optimization component of the NMPC consists of the solution of a deterministic open-loop optimal control problem. The optimal control problem considered is a Bolza problem with input constraints, i.e.

$$\min_{x,u} \quad \int_{t_k}^{t_k+T} l(x(t), u(t)) \,\mathrm{d}t \ + l_f(x(t_k+T)), \qquad (2a)$$

s.t.
$$x(t_k) = \hat{x}_{k|k},$$
 (2b)

$$\dot{x}(t) = f(x(t), u(t); p), t \in [t_k, t_k + T], \qquad (2c) u(t) \in \mathcal{U}(t), \qquad t \in [t_k, t_k + T], \qquad (2d)$$

where $\hat{x}_{k|k}$ denotes the filtered state estimates (from the filter) and p is the parameter estimates (from the online- or offline estimation method). $T = T_c = T_p$ is the control and prediction horizon. Several indirect and direct methods exists for the numerical solution of this optimal control problem (Binder et al., 2001). In this paper we use a direct local collocation method (von Stryk, 1993).

We use a stochastic extension of the van der Pol oscillator model to illustrate the components of the NMPC system. The numerical methods for the simulation study are implemented in python and the source code is available via \texttt{GitHub}^1 .

1.2 Paper organization

The paper is organised as follows. Section 2 presents the CDEKF, while Section 3 derives the use of the CDEKF for online and offline parameter estimation. Section 4 presents the local collocation method for numerical optimal control. Section 5 illustrates these components of the NMPC using the stochastic van der Pol oscillator model. Finally, Section 6 contains a short summary.

2. THE EXTENDED KALMAN FILTER

We present the CDEKF used in the NMPC as well as for the offline system identification. $\hat{x}_{k|k}$ and $\hat{P}_{k|k}$ denote the filtered state- and covariance estimates. $\hat{x}_{k|k-1}$ and $\hat{P}_{k|k-1}$ denote the predicted (one-step predictions) stateand covariance values.

2.1 The prediction scheme

Given the initial conditions

 $\hat{x}_{k-1}(t_{k-1}) = \hat{x}_{k-1|k-1}, \quad \hat{P}_{k-1}(t_{k-1}) = \hat{P}_{k-1|k-1},$ (3) the state- and covariance are predicted by solving the system of ordinary differential equations (ODEs) given by

$$\dot{\hat{x}}_{k-1}(t) = f(\hat{x}_{k-1}(t), u(t); p),$$
(4a)

$$\hat{P}_{k-1}(t) = A(t)\hat{P}_{k-1}(t) + \hat{P}_{k-1}(t)A(t)' + G(t)G(t)', \quad (4b)$$

where

$$A(t) = \frac{\partial f}{\partial x}(\hat{x}_{k-1}(t), u(t); p), \quad G(t) = g(\hat{x}_{k-1}(t), u(t); p).$$

The one-step predictions of the mean and covariance of the states are obtained as the solution of (3)-(4) at the new sample point, t_k . Consequently, the predictions of the state- and covariance are

$$\hat{x}_{k|k-1} = \hat{x}_{k-1}(t_k), \quad \hat{P}_{k|k-1} = \hat{P}_{k-1}(t_k).$$
 (5)

2.2 The updating scheme

The literature contains many methods for the updating scheme of extended Kalman filter algorithms. They all compute the innovation by

$$e_k = y_k - h(\hat{x}_{k|k-1}),$$
 (6)

the Kalman filter gain, K_k , by

$$C_k = \frac{\partial n}{\partial x}(\hat{x}_{k|k-1}),\tag{7a}$$

$$R_{k|k-1} = C_k \dot{P}_{k|k-1} C'_k + R_k, \tag{7b}$$

$$K_k = P_{k|k-1} C'_k R^{-1}_{k|k-1}, (7c)$$

and the filtered state estimate, $\hat{x}_{k|k}$, by

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k e_k.$$
(8)

The key difference is how they compute the filtered co-variance, $P_{k|k}$. Two standard updating schemes for the covariance are

$$\hat{P}_{k|k} = (I - K_k C_k) \,\hat{P}_{k|k-1}$$
 (9a)

$$= \hat{P}_{k|k-1} - K_k R_{k|k-1} K'_k.$$
(9b)

¹ https://github.com/niclasbrok/nmpc_vdp.git

Numerical implementations based on either (9a) or (9b) may give rise to bad performance and even divergence, as the *numerically* computed values are not guaranteed to be both positive (semi-)definite and symmetric. The Joseph stabilization form

$$\hat{P}_{k|k} = (I - K_k C_k) \hat{P}_{k|k-1} (I - K_k C_k)' + K_k R_k K'_k.$$
 (10)
for updating the filtered covariance estimate guarantees
that the numerical value of $P_{k|k}$ is symmetric posi-
tive (semi-)definite. The CDEKF is implemented using
(10) rather than (9) for reasons of numerical stability
and robustness (Schneider and Georgakis, 2013). Numer-
ically stable alternatives based on array- and squareroot-
algorithms also exist (Boiroux et al., 2016c), but are less
straightforward to implemente compared to (10).

3. PARAMETER ESTIMATION

In this section, we outline the application of the CDEKF for online and offline parameter estimation.

3.1 Online identification using the CDEKF

Using a disturbance model, the CDEKF can be used for parameter estimation in addition to state estimation. One way to do this is by augmenting the SDE by as many states as parameters undergoing the estimation

$$\mathrm{d}\boldsymbol{p} = \Sigma \,\mathrm{d}\boldsymbol{\omega}_p,\tag{11}$$

where $\boldsymbol{p} = (\boldsymbol{p}_1, ..., \boldsymbol{p}_{N_p})^{\top}$ are the parameters to be estimated and $\Sigma = \text{diag}(\sigma_1, ..., \sigma_{N_p})$. Success of this approach depends on detectability of the augmented system. Defining $\boldsymbol{z} = (\boldsymbol{x}, \boldsymbol{p})$, the augmented SDE has the form

$$d\boldsymbol{z}(t) = f_z(\boldsymbol{z}(t), \boldsymbol{u}(t); \tilde{p}) dt + g_z(\boldsymbol{x}(t), \boldsymbol{u}(t); \tilde{p}) d\boldsymbol{\omega}, \quad (12)$$

where

$$f_z(\boldsymbol{x}, u; \tilde{p}) = \begin{pmatrix} f(\boldsymbol{x}, u; p, \tilde{p}) \\ 0 \end{pmatrix}, g_z(\boldsymbol{x}, u; \tilde{p}) = \begin{pmatrix} g(\boldsymbol{x}, u; \tilde{p}) & 0 \\ 0 & \Sigma \end{pmatrix}$$

 \tilde{p} denotes the remaining parameters that are not estimated. Using this approach, the parameters are represented as disturbance states of the system since the observation equation is still given by

$$\boldsymbol{y}(t_k) = h_z(\boldsymbol{z}(t_k)) + \boldsymbol{v}_k = h(\boldsymbol{x}(t_k)) + \boldsymbol{v}_k.$$
(14)

The online parameter estimates are given by the filtered values from the augmented state vector, i.e. $\hat{z}_{k|k} = (\hat{x}_{k|k}; \hat{p}_{k|k})$. This augmentation method may be used for parameter estimation as well as disturbance estimation in offset free control (Morari and Maeder, 2012).

3.2 Offline identification using an ML formulation

Another use of the CDEKF is to estimate the parameters for a batch of data (Kristensen et al., 2004) in an offline optimization. The parameter estimates are the parameter set that maximizes the likelihood of the one-step prediction errors.

Let $\{y_j\}_{j=1}^N$ denote N observations relating to the sample points $\{t_j\}_{j=1}^N$ in (1b). Define the information accumulated up until the k-th sample point as $\mathcal{Y}_k = \{y_j\}_{j=1}^k$. Then the likelihood function, \mathcal{L} , can be defined as

$$\mathcal{L}\left(p \mid \mathcal{Y}_{N}\right) \propto \phi\left(\mathcal{Y}_{N} \mid p\right),\tag{15}$$

where ϕ is the joint density function of the observations, \mathcal{Y}_N . Using the definition of conditional probabilities, the right hand side can be rewritten into

M

$$\phi\left(\mathcal{Y}_{N}\mid p\right) = \prod_{k=1}^{n} \phi\left(y_{k}\mid \mathcal{Y}_{k-1}, p\right), \qquad (16)$$

such that the log-likelihood function can be expressed by $_{N}$

$$\log\left(\mathcal{L}\left(p \mid \mathcal{Y}_{N}\right)\right) = \sum_{k=1} \log\left(\phi\left(y_{k} \mid \mathcal{Y}_{k-1}, p\right)\right).$$
(17)

Consequently, the ML parameter estimates, $p_{\rm ML}$, is given by

$$p_{\mathrm{ML}} = \operatorname*{arg\,max}_{p \in \mathbb{R}^{N_p}} \log \left(\mathcal{L}(p \mid \mathcal{Y}_N) \right)$$
(18a)

$$= \underset{p \in \mathbb{R}^{N_p}}{\operatorname{arg\,max}} \sum_{k=1}^{N} \log(\phi(y_k \mid \mathcal{Y}_{k-1}, p)).$$
(18b)

Since the SDE in (1a) is driven by a Brownian motion and since the increments of a Brownian motion are Gaussian it is reasonable to assume, under some regularity conditions, that the conditional densities in (16) can be well approximated by Gaussian densities

$$\phi(y_k \mid \mathcal{Y}_{k-1}, p) = \frac{\exp\left(-\frac{1}{2}e'_k R_{k|k-1}^{-1} e_k\right)}{\sqrt{\det(R_{k|k-1})(2\pi)^{n_y}}},$$
 (19)

where n_y is the number of output variables.

4. NUMERICAL OPTIMAL CONTROL

In this section, we briefly present the algorithm for optimization of (2) used by the NMPC considered in this tutorial paper. The algorithm is based on a direct local collocation method presented by von Stryk (1993).

For simplicity, we let the Mayer term of (2) be zero, i.e. $l_f(x(t_k + T)) = 0$, and only consider optimal control problems with a Lagrange term. We consider only bounds on the inputs. This implies that $\mathcal{U}(t)$ denotes these bound constraints, i.e.

$$\mathcal{U}(t) = [u_{\min}, u_{\max}]. \tag{20}$$

4.1 A local collocation method

When formulating a direct solution method, the first step is to introduce a parametrization of the manipulated variable, u(t). The simplest parametrization is to approximate u(t) as a piecewise constant function. For reasons of simlicity, we adopt this parametrization method in the following. Hence, u(t) is parametrized via the values $\{q_k\}_{k=1}^N$ and time points $\{\tau_k\}_{k=0}^N$ such that

$$u = \sum_{k=1}^{N} q_k \chi_{[\tau_{k-1}, \tau_k[},$$
(21)

where χ_I denotes the characteristic function associated with the set *I*. The time points that define the subintervals of u(t) also constitute the global collocation points (GCPs) of the collocation method. Fig. 2 provides a schematic overview of the GCPs in relation to the local collocation points (LCPs), $\{\gamma_j\}_{j=0}^M$. The collocation points satisfy the relations

$$t_k = \tau_0 < \tau_1 < \dots < \tau_N = t_k + T,$$
 (22a)

$$0 = \gamma_0 < \gamma_1 < \dots < \gamma_M = 1. \tag{22b}$$



Fig. 2. Schematic overview of the collocation points.

A local collocation method approximates the ODEs on the smaller subintervals and apply a quadrature rule to impose a finite dimensional approximation. The ODEs on the interval $[\tau_{k-1}, \tau_k]$ can be formulated as an integral equation of the form

$$\boldsymbol{x}(\tau_k) - \boldsymbol{x}(\tau_{k-1}) = \int_{\tau_{k-1}}^{\tau_k} f(\boldsymbol{x}(t), q_k; p) \,\mathrm{d}t.$$
(23)

Using the forward Euler scheme, the right hand side can be approximated via the LCPs by

$$\int_{\tau_{k-1}}^{\tau_k} f(\boldsymbol{x}(t), q_k; p) \,\mathrm{d}t \tag{24a}$$

$$= \sum_{j=1}^{M} \int_{\tau_{k-1}+\gamma_j \Delta \tau_k}^{\tau_{k-1}+\gamma_j \Delta \tau_k} f(\boldsymbol{x}(t), q_k; p) \,\mathrm{d}t$$
(24b)

$$\approx \Delta \tau_k \sum_{j=1}^{M} \Delta \gamma_j f(\boldsymbol{x}(\tau_{k-1} + \gamma_{j-1} \Delta \tau_k), q_k; p), \quad (24c)$$

where $\Delta \tau_k = \tau_k - \tau_{k-1}$ and $\Delta \gamma_j = \gamma_j - \gamma_{j-1}$. Next, introduce the discrete state vector as

$$s_{k,j} = \boldsymbol{x}(\tau_{k-1} + \gamma_j \Delta \tau_k). \tag{25}$$

Using this notation, the local collocation scheme can be formulated as

$$s_{k-1,M} - s_{k-1,0} = \sum_{j=1}^{M} \Delta t_{k,j} f(s_{k-1,j-1}, q_k; p), \quad (26)$$

where $\Delta t_{k,j} = \Delta \tau_k \Delta \gamma_j$. (26) solves the ODEs on the local intervals - to obtain a meaningful ODE solution, a continuity condition must be imposed, together with an initial value constraint

$$s_{0,0} = x_0$$
 and $s_{k-1,M} = s_{k,0}$. (27)

Finally, the objective function is also approximated using the forward Euler method with the collocation points defined in (22)

$$\int_{t_k}^{t_k+T} l(\boldsymbol{x}(t), u(t)) \, \mathrm{d}t \approx \sum_{k=1}^N \sum_{j=1}^M \Delta t_{k,j} l(s_{k-1,j-1}, q_k).$$
(28)

Thus, the finite dimensional NLP for numerical solution of the optimal control problem (2) can be defined as

$$\min_{s,q} \sum_{k=1}^{N} \sum_{j=1}^{M} \Delta t_{k,j} l(s_{k-1,j-1}, q_k)$$
(29a)

s.t.
$$s_{0,0} = x_0$$
 (29b)

$$s_{k-1,M} = s_{k,0} \tag{29c}$$

$$s_{k-1,M} - s_{k-1,0} = \sum_{j=1} \Delta t_{k,j} f(s_{k-1,j-1}, q_k)$$
 (29d)

$$q_k \in \mathcal{U}(\tau_k). \tag{29e}$$

We solve (29) using ipopt and python. The numerical implementation in python uses the pyipopt package to

interface to ipopt (Wächter and Biegler, 2006). Using ipopt it is possible to exploit the sparse structure that appears in the Jacobian of the constraint function of (29).

5. NUMERICAL CASE STUDY

To illustrate the methodology presented in this paper, we use the stochastic van der Pol oscillator model that is defined as

 $d\boldsymbol{x}(t) = f(\boldsymbol{x}(t), u(t); \lambda) dt + g(\boldsymbol{x}(t), u(t); \sigma) d\boldsymbol{\omega}(t), \quad (30a)$ where

$$f(\boldsymbol{x}, u; \lambda) = \begin{pmatrix} x_2 \\ -x_1 + \lambda(1 - x_1^2)x_2 + u \end{pmatrix}, \quad (30b)$$

$$g(\boldsymbol{x}, u; \sigma) = \begin{pmatrix} 0 & 0 \\ 0 & \sigma \end{pmatrix}.$$
 (30c)

 $\lambda > 0$ is a parameter governing the stiffness of the system and $\sigma > 0$ is a parameter related to model deficiency. The SDE model (30) is used as a model for simulating the plant as well as in the CDEKF of the NMPC. Fig. 3 provides a comparison between three realizations of the SDE (30) and the corresponding ODE. The effect of the added noise to x_2 is clearly visible.

5.1 Set-point tracking using online parameter estimation

The set-point trajectory, $\overline{x}_1(t)$, for $x_1(t)$ is defined as the step function

$$\overline{x}_1(t) = \begin{cases} 0, & t \le 15\\ 1, & 15 \le t \le 30\\ 0, & 30 \le t \end{cases}$$
(31)

and the corresponding integrand of the control objective, l, is defined as

 $l(x(t), u(t)) = (1 - \alpha)(x_1(t) - \overline{x}_1(t))^2 + \alpha u(t)^2, \quad (32)$ where $\alpha = 1/1000$ is regularization parameter. The noise parameters are defined as

$$R_k = \sigma_{\varepsilon}^2 I, \quad \sigma_{\varepsilon} = 1/100, \quad \sigma = 1/5.$$
 (33)

It is assumed that both states are directly observable, i.e. h(x(t)) = x(t). The NLP has been constructed with an equidistant mesh such that

$$\Delta \tau_k = 1/(N-1), \quad \Delta \gamma_k = 1/(M-1), \quad (34)$$

where N = 51 and M = 21. The control and prediction horizon of the optimal control problem (2) is $T = T_c = T_p = 20$. It is assumed that observations occur equidistantly with $T_s = t_k - t_{k-1} = 0.4$. The control signals are constrained by the sets, $\mathcal{U}(\tau_k) = [-1, 1]$.

Fig. 4 shows a closed-loop simulation. For this simulation, the controller has to estimate λ . Fig. 5 shows the true value of λ as well as the online parameter estimate, $\hat{\lambda}_{k|k}$. Fig. 5 also shows how the online estimation method performs when an unmodelled disturbance is introduced. The unmodelled disturbance is introduced at t = 45 where the true value of λ is changed from $\lambda = 1$ to $\lambda = 3$.

Fig. 6 shows the offline parameter estimates and the corresponding log-likelihood functions based on 100 observations of the plant. The offline estimation is tested in two cases; a case where $\lambda = 1$ and a case where $\lambda = 10$. The parameters are estimated to be $\hat{\lambda} = 1.017$ and $\hat{\lambda} = 9.953$, respectively.



Stochastic Realizations

Fig. 3. The difference between an ODE realization (deterministic) and three different realizations of the SDE (stohcastic).



Set-Point Tracking

Fig. 4. A closed-loop simulation where x_1 has to track the set-point trajectory given in (31). The top-left plot shows the true state values of x_1 together with the set-point trajectory. The bottom-left plot shows the manipulated variables, u(t), computed by the NMPC. The phase plot to the right illustrates the path the NMPC chooses when a set-point change is imposed.

The results presented in Fig. 4 and Fig. 5 show that the NMPC is able to simultaneously control the system to the desired set-point trajectory (31) and accurately estimate the unknown parameter, λ . However, from Fig. 5 it is seen that the parameter estimate is sensitive to the set-point change. For the estimates around $t \in \{15, 30\}$ in Fig. 5, the EKF estimates significant parameter changes despite of the fact that the true value is kept constant at $\lambda = 1$. The results from Fig. 4 also show that the resulting control signal, u(t), is active around (and on) the upper bound when $\overline{x}_1(t) = 1$. This is a result of the fact that (1,0) is not an equilibrium for (30). Hence, the controller has to actively change x_2 to keep x_1 close to \overline{x}_1 .

6. SUMMARY

We provide a tutorial overview of how to construct an NMPC to regulate a stochastic system governed by SDEs. Based on the CDEKF, an online and an offline method for parameter estimation are presented. The dynamic optimization module is based on a local collocation scheme, where the forward Euler method has been used as dis-



Fig. 5. The online parameter estimates of λ , where at t = 45 an un-modelled disturbance is introduced (λ shifts from 1 to 3).

cretization method for the dynamical equations. The performance of the closed-loop controller is investigated for a stochastic extension to the van der Pol oscillator model. The source code for the tutorial is available via GitHub. The key contribution and insight is to use the continuousdiscrete model (1) and the same associated CDEKF in all components of the NMPC.



Fig. 6. Examples of offline parameter estimation of λ . Two cases are shown; the top plot shows an offline estimation problem where the true value is $\lambda = 1$ and the bottom plot shows an offline estimation problem where the true value is $\lambda = 10$. Both problems are based on 100 observations with the noise parameters given in (33).

REFERENCES

- Allgöwer, F., Badgwell, T.A., Qin, J., Rawlings, J., and Wright, S.J. (1999). Nonlinear predictive control and moving horizon estimation - an introductory overview. In Advances in Control. Highlights of ECC'99. Springer, London, GB.
- Binder, T., Blank, L., Bock, H.G., Bulirsch, R., Dahmen, W., Diehl, M., Kronseder, T., Marquardt, W., Schlöder, J.P., and von Stryk, O. (2001). Introduction to model based optimization of chemical processes on moving horizons. In M. Grötschel, S.O. Krumke, and J. Rambau (eds.), Online Optimization of Large Scale Systems. Springer.
- Boiroux, D., Hagdrup, M., Mahmoudi, Z., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016a). An ensemble nonlinear model predictive control algorithm in an artificial pancreas for people with type 1 diabetes. In 2016 European Control Conference (ECC), 2115–2120.
- Boiroux, D., Hagdrup, M., Mahmoudi, Z., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016b). Model identification using continuous glucose monitoring data for type 1 diabetes. *IFAC-PapersOnline*, 49(7), 759–764.
- Boiroux, D., Juhl, R., Madsen, H., and Jørgensen, J.B. (2016c). An efficient square root algorithm for the computation of maximum likelihood sensitivity of continuous-discrete systems. In 55th IEEE Conference on Decision and Control (CDC), 3048–3053. IEEE.
- Boiroux, D., Finan, D.A., Jørgensen, J.B., Poulsen, N.K., and Madsen, H. (2010). Meal estimation in nonlinear model predictive control for type 1 diabetes. *IFAC Proceedings Volumes*, 43(14), 1052–1057.
- Buehler, E.A., Paulson, J.A., and Mesbah, A. (2016). Lyapunov-based stochastic nonlinear model predictive control: Shaping the state probability distribution functions. In *American Control Conference (ACC)*, 2016, 5389–5394. IEEE.
- Garnier, H. and Young, P.C. (2012). What does continuous-time model identification have to offer? *IFAC Proceedings Volumes*, 45(16), 810–815.
- Grüne, L. and Pannek, J. (2011). Nonlinear Model Predictive Control. Theory and Algorithms. Springer, London, GB.
- Jazwinski, A.H. (1970). Stochastic Processes and Filtering Theory. Academic Press, San Diego, CA, USA.

- Johansen, T.A. (2011). Introduction to nonlinear model predictive control and moving horizon estimation. In M. Huba, S. Skogestad, M. Fikar, M. Hovd, T.A. Johansen, and B. Rohal-Ilkiv (eds.), Selected Topics on Constrained and Nonlinear Control, chapter 5, 187– 240. STU Bratislava - NTNU Trondheim, Bratislava, Slovakia.
- Jørgensen, J.B. and Jørgensen, S.B. (2007a). Comparison of prediction-error-modelling criteria. Proceedings of the 2007 American Control Conference, 140–146.
- Jørgensen, J.B. and Jørgensen, S.B. (2007b). MPCrelevant prediction-error identification. Proceedings of the 2007 American Control Conference, 128–133.
- Kristensen, N.R., Madsen, H., and Jørgensen, S.B. (2004). Parameter estimation in stochastic grey-box models. *Automatica*, 40(2), 225–237.
- Mahmood, M. and Mhaskar, P. (2012). Lyapunov-based model predictive control of stochastic nonlinear systems. *Automatica*, 48(9), 2271–2276.
- Mahmoudi, Z., Boiroux, D., Hagdrup, M., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016a). Application of the continuous-discrete extended Kalman filter for fault detection in continuous glucose monitors for type 1 diabetes. In 2016 European Control Conference (ECC), 714–719.
- Mahmoudi, Z., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2017). Fault and meal detection by redundant continuous glucose monitors and the unscented kalman filter. *Biomedical Signal Processing and Control*, 38, 86–99.
- Mahmoudi, Z., Wendt, S.L., Boiroux, D., Hagdrup, M., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016b). Comparison of three nonlinear filters for fault detection in continuous glucose monitors. In 2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 3507–3510.
- Morari, M. and Maeder, U. (2012). Nonlinear offset-free model predictive control. Automatica, 48(9), 2059–2067.
- Rao, G.P. and Unbehauen, H. (2006). Identification of continuous-time systems. *IEE Proceedings - Control Theory and Applications*, 153(2), 185–220.
- Rawlings, J.B., Mayne, D.Q., and Diehl, M.M. (2017). Model Predictive Control: Theory, Computation, and Design. Nob Hill Publishing, Madison, WI, USA.
- Schneider, R. and Georgakis, C. (2013). How to NOT make the extended Kalman filter fail. *Industrial & Engineering Chemistry Research*, 52(9), 3354–3362.
- Simon, D. (2006). Optimal state estimation: Kalman, H_{∞} , and nonlinear approaches. John Wiley & Sons.
- von Stryk, O. (1993). Numerical solution of optimal control problems by direct collocation. In R. Bulirsch, J. Stoer, and K. Well (eds.), Optimal Control: Calculus of Variations, Optimal Control Theory and Numerical Methods, 129–143. Birkhäuser Basel, Basel.
- Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical* programming, 106(1), 25–57.


Available online at www.sciencedirect.com

ScienceDirect



IFAC PapersOnLine 52-4 (2019) 494-499

Flexible Control of Wastewater Aeration for Cost-Efficient, Sustainable Treatment

Niclas Brabrand Brok * Thomas Munk-Nielsen ** Henrik Madsen * Peter A. Stentoft *,**

* Department of Applied Mathematics and Computer Science, Technical University of Denmark (DTU), Lyngby, Denmark (e-mail: nlbr@dtu.dk) ** Krüger A/S, Veolia Water Technologies, Søborg, Denmark (e-mail: pas@kruqer.dk)

Abstract: With increasing fluctuations in electricity production due to prioritization of renewable energy sources, new applications that can adjust quickly to changes in demand/supply will be needed. Wastewater treatment use a significant amount of electricity to reduce nutrients in wastewater before discharge. The treatment process demands electricity in some selected periods which can be controlled, and hence the time of consumption is changeable. Here we suggest a novel predictive control strategy which enhances the flexibility in electricity demand by accounting for electricity price and probability of up or down regulation. The strategy is demonstrated in simulation experiments, where the concept is illustrated and the potential savings are estimated. Furthermore flexibility is investigated as a function of regulating prices, and it is shown that when difference between electricity price and regulating price increases, so does flexibility of the system.

© 2019, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.

Keywords: Wastewater Treatment, Stochastic Model Predictive Control, Smart Energy Application, Balancing Market, Day-ahead market, Optimization, Price-based control

INTRODUCTION

Advancements in wind turbines and solar panels coupled with societal requests, declines in cost and favourable regulation have all contributed to a rapid increase in renewables in electricity grids (Ueckerdt et al., 2015). While this is an obligatory development to reach the desired fossil-free energy system (Jacobson and Delucchi, 2011), fluctuating renewable energy sources impose a challenge in securing sufficient supply to cover demand at all times. According to the EU, adaptation to the increasing amount of renewables requires massive development in smart energy systems. Hence, building heat, vehicle charging and industrial cooling (e.g. (Zemtsov et al., 2018)) have received attention in their ability to smartly prioritize electricity consumption in selected periods without significant loss of utility for the users. However, as more renewables penetrate the electricity grids, more smart applications will be needed to maintain a stable system (Morales et al., 2014).

Wastewater treatment plants (WWTP) use approximately 1% of a country's total electricity consumption (Shi, 2011) meaning that e.g. Germany and USA spend 4.4 and 30.2 TWh/yr on wastewater treatment respectively (Haberkern et al., 2008; Pabi et al., 2013). From this follows (i) that electricity is a major economic issue for plant operation corresponding to 25-50% of operational costs (e.g. (Huang et al., 2013)) and (ii) that the greenhouse gas emissions (GHG) related to electricity consumption of wastewater treatment are noteworthy (Mizuta and Shimada, 2010).

The most electricity consuming process on a WWTP is aeration which accounts for 40-75% of total electricity demand of a WWTP (Rosso et al., 2008). Aeration is typically carried out in large, engineered tanks where specialized bacteria need aerobic conditions (oxygen present) to convert ammonium from e.g. urine to nitrate. Then other bacteria convert nitrate to nitrogen gas under anoxic conditions (oxygen not present) and hence nitrogen is removed from the water. This implies, that the ideal process requires both aerobic and anoxic periods or areas in the tank. Hence advanced control of wastewater treatment aims at turning aeration on and off in feedback loops to secure good treatment (Zhao et al., 2004).

In this paper we suggest a novel optimization strategy which controls aeration with respect to both the nonlinear biochemical processes and the electricity market. We use Nonlinear Model Predictive Control (NMPC) methods to solve the control problem. In other words we satisfy the wastewater treatment requirements regarding treatment and equipment constraints while we control electricity demand in a flexible way that allows for trading electricity. Last, we show through an example study that the control leaves satisfactory effluent concentrations of the investigated nutrients, and that flexibility in power usage can be envoked. Finally we show how different regulating prices influence the flexibility and costs.

THEORY AND METHODS

We briefly describe how wastewater treatment can be modelled using Stochastic Differential Equations (SDE).

2405-8963 © 2019, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved. Peer review under responsibility of International Federation of Automatic Control. 10.1016/j.ifacol.2019.08.259

Furthermore we resume the Nordpool electricity market as an example case. Then, with respect to the market design, the optimal control problem is defined, and finally the numerical implementation is briefly mentioned.

Wasterwater Treatment: Applied SDE Modelling

SDEs are used in a wide range of applications. A general form of an SDE is

$$x(t) = x_0 + \int_0^t f(x(\tau), u(\tau)) \,\mathrm{d}\tau + \int_0^t g(x(\tau), u(\tau)) \,\mathrm{d}\omega(\tau),$$

or in short

or in short

$$dx(t) = f(x(t), u(t)) dt + g(x(t), u(t)) d\omega(t), \qquad (1)$$

where $x \colon \mathbb{R} \to \mathbb{R}^{n_x}$ denotes the states, x_0 is the initial distribution of the states, $u \colon \mathbb{R} \to \mathbb{R}^{n_u}$ is the input variables and $\omega \colon \mathbb{R} \to \mathbb{R}^{n_\omega}$ denotes a standard (possibly multivariate) Brownian motion. Brownian motion is defined by its independent increments which satisfy that for each $s, t \in \mathbb{R}$, $\omega(t) - \omega(s)$ is normally distributed with zero mean and covariance I(t-s). $f \colon \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x} \times \mathbb{R}^{n_\omega}$ is called the diffusion function.

SDEs provide a powerful stochastic continuous-time modelling framework which can be used for both parameter and state estimation (Kristensen et al., 2004). This framework has been applied to wastewater treatment in Halvgaard et al. (2017); Stentoft et al. (2018) where the activated sludge model in Henze et al. (2000) is reduced to a lower-order SDE model. Here, we use the model

$$dx_{1}(t) = a_{1}(a_{2} - x_{1}(t)) dt$$

$$- u(t)a_{3}\frac{x_{1}(t)}{a_{4} + x_{1}(t)} dt + \sigma_{1} d\omega_{1}(t),$$

$$dx_{2}(t) = a_{1}(a_{5} - x_{1}(t)) dt + a_{3}u(t)\frac{x_{1}(t)}{a_{4} + x_{1}(t)} dt$$

$$- (1 - u(t))a_{6}\frac{x_{2}(t)}{a_{7} + x_{2}(t)} dt + \sigma_{2} d\omega_{2}(t),$$

(2)

Table 1. Example parameters of (2).

Shorthand	Description	Value
$x_1(0)$	Initial ammoinum conc.	1.12
$x_2(0)$	Initial nitrate conc.	0.87
a_1	Incoming wastewater rate	0.00067
a_2	Mean incoming ammonium	36.9
a_3	Nitrification rate	0.073
a_4	Monod kinetic constant	0.1
a_5	Mean incoming nitrate	2.00
a_6	Denitrification rate	0.300
a_7	Monod Kinetic constant	7.84
σ_1	Ammonium noise parameter	0.0085
σ_2	Nitrate noise parameter	0.026

where the parameters are estimated from online measurements from a WWTP using the estimation method given in Kristensen et al. (2004) and applied to wastewater treatment in Stentoft et al. (2018). The parameter estimates and a short description are presented in Table 1.

The input function, u, in (2) is a binary-valued function, which is 1 if the aeration system is activated and 0 otherwise. This means that the system may be characterized as a switched dynamical system where the optimal control is related to control of switching times between the two systems (Egerstedt et al., 2003).

The initial values, $x(0) = (x_1(0), x_2(0))^{\top}$, will in an online implementation be estimated when new measurements become available. However, in the simulations carried out in the later case-study, these values will constitute the initial values in the open-loop control problem.

The Nordic Electricity Market

In Northern Europe, electricity is traded in a common market called Nord Pool which consists of 15 interconnected price areas. The market which trades with the largest volume is called the day-ahead market. Here, electricity is purchased and sold for the upcoming day. When the dayahead market closes, the intra-day market opens. In this market, electricity can be traded until 45 minutes prior to the operating hour.

One of the primary challenges when operating transmission systems is to guarantee grid stability. The Nordic Transmission System Operators (TSOs) have many methods for dealing with this challenge - one of them being a common balancing market (EnergiNet, 2016). In the balancing market, market participants have the option to make a bid that defines how much a participant is willing to change their production- or consumption schedule in a given operating hour. The balancing market also closes 45 minutes prior to the operating hour. Hence, when approaching the operating hour, the TSO has the possibility to activate balancing bids ahead of time and avoid possible imbalances. Three scenarios can take place in the balancing market:

- (\uparrow) If the imbalance is negative, there is a deficit of electricity in the price area and hence an increase in the production or a decrease in the consumption is needed. This is called up regulation.
- (\downarrow) If the imbalance is positive, there is a surplus of electricity in the price and hence a decrease in the production or an in the consumption is needed. This is called down regulation.
- (-) If the imbalance is too small or the duration is too short, the imbalance is not offered in the balancing market.

In a situation with up regulation, electricity is sold while in the situation with down regulation, electricity is purchased. The structure of the balancing market requires that the up regulation price is greater than the day-ahead price while the down regulation price is lower than the dayahead price. Let p(t) denote the day-ahead price, $p^{(\uparrow)}(t)$ the up regulation price and $p^{(\downarrow)}(t)$ the down regulation price. We then have that

$$p^{(\downarrow)}(t) \le p(t) \le p^{(\uparrow)}(t). \tag{3}$$

In Fig. 1 a Nord Pool price example of these prices is shown. In the price area DK1, a large share of the total production comes from wind turbines - this is a source of energy which is very difficult to predict and hence one of the primary reasons to imbalances - in Parbo (2014) it is suggested that approximately 65% of the total imbalances are due to forecast errors of the wind power production.



Fig. 1. Nord Pool price example for DK1 for the first five days of April 2018.

The Optimal Control Problem

 τ_{c}

In this section we define an optimization problem, which can compute the optimal input signal, u. We will assume that u can be parameterized by

$$\tau = \{(\tau_{\text{off},k}, \tau_{\text{on},k})\}_{k=0}^{N},$$
(4)

such that

$$u(t;\tau) = \begin{cases} 1, \ t \in [\tau_{\text{on},k-1}, \tau_{\text{off},k}], & k = 1, ..., N\\ 0, \ t \in [\tau_{\text{off},k}, \tau_{\text{on},k+1}], & k = 0, ..., N-1 \end{cases}$$
(5)

where we impose the cycle structure that we start with the aeration being deactivated (i.e. that $\tau_{off,0} = 0$). The cycle parameters are ordered such that

$$\tau_{\text{on},k} \le \tau_{\text{off},k}, \quad k = 0, \dots, N, \tag{6a}$$

$$\tau_{n,k+1} \ge \tau_{\text{off},k}, \quad k = 0, ..., N - 1.$$
 (6b)

The main objective of a WWTP is to remove nutrients from incoming wastewater before discharging it back to environment at a sustainable level at minimum operational cost. A sustainable level is arguably depending on the vulnerability of the environment where the wastewater is discharged, and hence it is typical for a WWTP to implement hard constraints on ammonium- and total-N concentrations (here x_1 and $x_1 + x_2$) in the effluent. In this study hard constraints on concentrations are not imposed, but rather, we define a cost on the discharge of total-nitrogen to the environment which means that large concentrations are preferably avoided. This is similar to the current Danish legislation where discharge of total-N is taxed by 30 DKK/kg-N. In addition to the cost related to total-N, the cost of electricity related to the aeration process is considered in the minimization of total cost where we assume that the WWTP has the possibility to bid (and be activated) in the balancing market.

To contain all of this into a univariate function, we propose a scenario-based structure, consisting of:

- an up regulation scenario,
- a down regulation scenario,
- a neutral scenario, where no regulation is activated,

where each of the scenarios contain a set of cycle parameters: $\tau^{(\uparrow)}$, $\tau^{(\downarrow)}$ and τ , respictively. In Fig. 2 we have made a schematic overview of these cycle parameters.

Invariant of which scenario is active, the WWTP has to trade electricity such that a nominal operation strategy can be deployed. We define this nominal strategy as the decision relating to the neutral scenario. Hence, the electricity cost for each of the scenarios may be defined as



Fig. 2. Schematic overview of the cycle parameters, τ , $\tau^{(\uparrow)}$ and $\tau^{(\downarrow)}$.

$$c_p^{(0)}(\tau) = \sum_{k=1}^N \int_{\tau_{\text{on},k-1}}^{\tau_{\text{off},k}} p(t) \,\mathrm{d}t,\tag{7a}$$

$$c_p^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right) = \sum_{k=0}^{N} \left[\int_{\tau_{\text{on},k}}^{\tau_{\text{on},k}^{(\uparrow)}} p^{(\uparrow)}(t) \,\mathrm{d}t + \int_{\tau_{\text{off},k}}^{\tau_{\text{off},k}} p^{(\uparrow)}(t) \,\mathrm{d}t \right], \quad (7b)$$

$$c_{p}^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right) = \sum_{k=0}^{N} \left[\int_{\tau_{\text{on},k}^{(\downarrow)}}^{\tau_{\text{on},k}} p^{(\downarrow)}(t) \,\mathrm{d}t + \int_{\tau_{\text{off},k}}^{\tau_{\text{off},k}^{(\downarrow)}} p^{(\downarrow)}(t) \,\mathrm{d}t \right], \quad (7c)$$

where $c_p^{(0)}(\tau)$ denotes the electricity cost for the nominal strategy, $c_p^{(\uparrow)}(\tau, \tau^{(\uparrow)})$ denotes the electricity cost in the up regulation scenario and $c_p^{(\downarrow)}(\tau, \tau^{(\downarrow)})$ denotes the electricity cost in the down regulation scenario. Hence, the expected cost, c_p , when participating in the balancing market is

$$c_p\left(\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}\right) = c_p^{(0)}(\tau) - \theta^{(\uparrow)}c_p^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right) + \theta^{(\downarrow)}c_p^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right),$$
(8)

where $\theta^{(\uparrow)}$ and $\theta^{(\downarrow)}$ denote the probability of the upand down regulation scenarios, respectively. Note that we have, with probability one, to pay the cost $c_p^{(0)}(\tau)$ for the nominal strategy, since this defines the consumption schedule from which we can participate in the balancing market. Similarly, we define the taxation cost (cost of effluent load) for the scenarios as

$$c_t^{(0)}(\tau) = \int_{\tau_{\text{off},1}}^{\tau_{\text{on},N}} \left(x_1^{(0)}(t) + x_2^{(0)}(t) \right) \mathrm{d}t \qquad (9a)$$

$$c_t^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right) = \int_{\tau_{\text{out},1}^{(\uparrow)}}^{\tau_{\text{out},1}^{(\bullet)}} \left(x_1^{(\uparrow)}(t) + x_2^{(\uparrow)}(t)\right) \mathrm{d}t \qquad (9b)$$

$$c_t^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right) = \int_{\tau_{\text{off},1}^{(\downarrow)}}^{\tau_{\text{off},N}^{(\downarrow)}} \left(x_1^{(\downarrow)}(t) + x_2^{(\downarrow)}(t)\right) \mathrm{d}t, \qquad (9c)$$

where $c_t^{(0)}(\tau)$ denotes the taxation cost for the nominal strategy, $c_p^{(\uparrow)}(\tau, \tau^{(\uparrow)})$ denotes the taxation cost in the up regulation scenario and $c_p^{(\downarrow)}(\tau, \tau^{(\downarrow)})$ denotes the taxation cost in the down regulation scenario. Note that this formulation of the taxation amount assumes that the flow of the incoming/outgoing water is constant over time. The expected taxation cost is then given by

$$c_t\left(\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}\right) = \left(1-\theta^{(\uparrow)}-\theta^{(\downarrow)}\right)c_t^{(0)}(\tau) +\theta^{(\uparrow)}c_t^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right) +\theta^{(\downarrow)}c_t^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right).$$
(10)

Note that the taxation cost of the nominal strategy is multiplied by $1-\theta^{(\uparrow)}-\theta^{(\downarrow)}$, since we might be activated for

up or down regulation. Thus, the total cost of operating the WWTP can be modelled as

$$c\left(\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}\right) = c_t\left(\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}\right) + c_p\left(\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}\right).$$
(11)

However, from an optimization perspective, c, is not a suitable choice for objective function; c will prefer cycle parameters which are as small as possible to minimize the optimization horizon. To eliminate this preference, we consider the time-averaged analogues of (7a)-(7c) and (9a)-(9c). We define these time-averaged versions according to

$$\hat{c}_p^{(0)}(\tau) = \frac{c_t^{(0)}(\tau)}{\tau_{\text{on},N}},$$
 (12a)

$$\hat{c}_{p}^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right) = \frac{c_{p}^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right)}{\tau_{0NN}^{(\uparrow)}},$$
(12b)

$$\hat{c}_{p}^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right) = \frac{c_{p}^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right)}{\tau_{\mathrm{on},N}^{(\downarrow)}},\qquad(12\mathrm{c})$$

and

$$\hat{c}_t^{(0)}(\tau) = \frac{c_t^{(0)}(\tau)}{\tau_{\text{on},N}},$$
(13a)

$$\hat{c}_t^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right) = \frac{c_t^{(\uparrow)}\left(\tau,\tau^{(\uparrow)}\right)}{\tau_{\text{on }N}^{(\uparrow)}},\tag{13b}$$

$$\hat{c}_t^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right) = \frac{c_t^{(\downarrow)}\left(\tau,\tau^{(\downarrow)}\right)}{\tau_{\text{on }N}^{(\downarrow)}}.$$
(13c)

Thus, we can define the optimal control problem as

$$\min_{\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}} \hat{c}\left(\tau,\tau^{(\uparrow)},\tau^{(\downarrow)}\right),\tag{14a}$$

subject to

$$\left(\tau, \tau^{(\uparrow)}, \tau^{(\downarrow)}\right) \in \mathcal{T},$$
 (14b)

$$\dot{x}^{(0)}(t) = f\left(x^{(0)}(t), u(t;\tau)\right), \quad t \in [0, \tau_{\text{on},N}], \quad (14c)$$

$$\dot{x}^{(\uparrow)}(t) = f\left(x^{(0)}(t), u\left(t; \tau^{(\uparrow)}\right)\right), \ t \in \left[0, \tau_{\text{on}, N}^{(\uparrow)}\right], \ (14\text{d})$$

$$\dot{x}^{(\downarrow)}(t) = f\left(x^{(0)}(t), u\left(t; \tau^{(\downarrow)}\right)\right), \ t \in \left[0, \tau_{\text{on},N}^{(\downarrow)}\right], \ (14e)$$

where \mathcal{T} defines the set of permissible cycle structures, which might be constraints such as minimum- and maximum cycle length and minimum- and maximum levels of flexibility bid to the balancing market via $\tau^{(\uparrow)}$ and $\tau^{(\downarrow)}$. f is the system model which is given from (2). In an actual application of the framework presented above we would need to apply the principle of Nonlinear Model Predictive Control as described in Brok et al. (2018) to obtain a closed-loop control system, where we would solve an optimization problem of the form (14a)-(14e) each time we receive new measurements from the WWTP.

Numerical Implementation

The optimization problem (14a)-(14e) has been solved in julia using ipopt (Bezanson et al., 2017; Wächter and Biegler, 2006). ipopt is a gradient-based optimization method - hence, the derivatives of the optimization problem (14a)-(14e) have to be provided. The dynamical systems (14c)-(14e) are switched dynamical systems which implies that the gradient of the objective function can be computed using the adjoint equations of the optimal control problem (14a)-(14e) (Egerstedt et al., 2003). The numerical method we have implemented is in the literature often referred to as a single shooting method (Bock and Plitt, 1984).

EXAMPLES AND DISCUSSION

The control strategy is tested with respect to the model presented in (2) and parameters in Table 1. We assume that \mathcal{T} in (14b) is defined as the set containing the cycle parameters satisfying the constraints defined in Table 2. Note that these constraints are also implemented across scenarios. We also assume that the following data is given:

- (1) the day-ahead electricity price, p(t),
- (2) up regulation price and probability, $p^{(\uparrow)}$ and $\theta^{(\uparrow)}$,
- (3) down regulation price and probability, $p^{(\downarrow)}$ and $\theta^{(\downarrow)}$.

Table 2. Cycle parameter constraints.

Constraint Description	Value
Minimum aeration length	6.0
Minimum no-aeration length	20.0
Maximum aeration length	60.0
Maximum no-aeration length	120.0

We will in the following simulation study assume that p is a constant function of time with value 10 price units. Hence we optimize costs with known prices and probabilities of up- or down regulation. We also assume that it is only the first cycle shift from on to off and the second cycle shift from off to on that can be bid to the balancing market. This means that \mathcal{T} also imposes the constraints $\tau_{\text{on},k} = \tau_{\text{on},k}^{(\uparrow)} = \tau_{\text{on},k}^{(\downarrow)}$ for $k \in \{0, 1, 3, ..., N\}$ and $\tau_{\text{off},k} = \tau_{\text{off},k}^{(\uparrow)} = \tau_{\text{off},k}^{(\downarrow)}$ for $k \in \{0, 2, 3, ..., N\}$.

Nominal Control Strategy

In this, first, simulation, we assume that $\theta^{(\uparrow)} = \theta^{(\downarrow)} = 0$. The resulting simulation is shown in Fig. 3. We observe that when the aeration is on, the ammonium concentration, x_1 , decreases while the nitrate concentration, x_2 , increases and vice versa when the aeration is off. From 3 we also observe that the accumulated cost is increasing steadily over time with approximately 70% of the total cost associated with the price of electricity. Finally, we note that after the first cycle, the aeration cycles are repeated with almost the same period. This is expected since the price of electricity and the flow-rate are constant.



Fig. 3. Baseline simulation showing the optimal control given constant electricity price and no regulation. Note that regime refers to the aeration state (on/off).



Fig. 4. Scenarios where regulation is expected and activated vs the nominal strategy as shown before.

Up and Down Regulation Scenarios

In the following simulations, we have two different scenarios:

- (†) A scenario where the probability of being activated for up regulation is $\theta^{(\uparrow)} = 1/10$ at an up regulation price of $p^{(\uparrow)}(t) = 20$.
- (\downarrow) A scenario where the probability of being activated for down regulation is $\theta^{(\downarrow)} = 1/10$ at an down regulation price of $p^{(\downarrow)}(t) = 0$.

Note that for the (\downarrow) -scenario, there is a small probability that the WWTP can consume electricity for free.

From Fig. 4 we see that for both scenarios it is a long term beneficial strategy to participate in the balancing market. We observe that the total cost, relative to the nominal strategy shown in Fig. 3, is ≈ 50 price units less than the nominal cost for the nominal strategy without probability of activation. From Fig. 4 we can also note that after the initial activation, the strategies re-prioritize reducing the total concentration of the nutrients and that the concentrations converge towards the result of the nominal strategy. In Fig. 5 we see that the nominal strategies with balancing market participation have a total cost which is very close to the cost of the nominal strategy with no participation. This is a desired property, since we don't want to impose a risk of huge total cost of the nominal strategies which primarily will define the state of operation.

Flexibility Diagram

To illustrate how flexibility can be invoked in the control, the effect of different up- and down regulation prices is investigated. In Fig. 6 we have shown a flexibility diagram. This diagram shows how much flexibility the method builds into the system for different up- and down regulation prices and how these additional strategies affect the nominal cost, expected cost and regulating cost. The total flexibility is defined as the amount the system is willing to change the nominal schedule given the up or down regulation price. Hence, it defines the available flexibility in the system. The flexibility diagram is generated based on:

- (\uparrow) In the up regulation region, the probability of being activated for up regulation is $\theta^{(\uparrow)} = 1/10$.
- (\downarrow) In the down regulation region, the probability of being activated for down regulation is $\theta^{(\downarrow)} = 1/10$.



Fig. 5. Scenarios where regulation is expected but not activated vs the nominal strategy as shown before.

We observe in fig. 6 that larger savings are expected as the regulation prices becomes more favorable (relative to the fixed day-ahead price) and that the nominal cost is flat with increased sensitivity towards down regulation.

The flexibility diagram depends on the parameters used in the prediction model in (2) and the probabilities and prices used. Thus in an online application, parameters are expected to be frequently updated and hence the flexibility diagram will also change. Also, the framework presented in this paper assumes that the scenario probabilities are invariant over time, this is however by no means a necessary assumption. Hence, we imagine that these could be generated from forecasts of the balancing market (even though this is a very difficult market to predict). For closed-loop strategies, where we rely on feedback to get efficient operation, we imagine that flexibility diagrams might be key tool in aggregation of multiple WWTPs where the total consumption and flexibility is traded in the day-ahead, intra-day and balancing markets. In applications of the optimization method, we would also need to consider how to realistically mimic the market structure of the Nordic electricity markets.

Lastly, we remark that the presented optimization method only considers biological nitrogen removal. An extension to include biological phosphorous removal would be useful for many WWTPs. Furthermore, legislation in many countries requires that ammonium and total-nitrogen are kept below a certain limit (i.e. hard constraints on x_1 and $x_1 + x_2$). For improved applicability this should also be included in future implementations.



Fig. 6. Flexibility diagram showing the effect of changing the up- and down regulation price. Price level 10.0 corresponds to no regulation (or regulation price equal zero).

CONCLUSION

In this paper we presented a novel stochastic model predictive control strategy of wastewater aeration for cost-efficient and sustainable treatment. The optimization model trade-off taxation cost of nutrients discharged into the environment with the cost of operating the WWTP.

The output from the optimization model is a nominal strategy and strategies for up- and down regulation, respectively. We show via a simulation study, that by allowing the optimization model to utilize the balancing market, the operational costs can be decreased compared to only considering a nominal strategy with no participation in the balancing market. The simulation study also shows that after the activation period, the regulation-based strategies converges (over time) to the nominal strategy. We furthermore suggest a flexibility diagram which shows how different up and down regulation price levels influence the operational costs and flexibility. Furthermore it is suggested that this can be a tool for aggregation of multiple WWTPs.

Finally we consider this paper as a step towards integrating wastewater treatment in the balancing market(s). Ultimately, our approach can help wastewater treatment operation adjust to rapid changes in electricity supply/prices and thereby make them more resilient to increasing amounts of renewables in the transmission grid.

ACKNOWLEDGEMENTS

This work is partly funded by the Innovation Fund Denmark (IFD) under File No. 7038-00097B Peter A. Stentofts industrial PhD study; *Stochastic Predictive Control of Wastewater Treatment Processes and File No. 7038-*00097B, and the IFD project CITIES (1305-0027B).

REFERENCES

- Bezanson, J., Edelman, A., Karpinski, S., and Shah, V.B. (2017). Julia: A fresh approach to numerical computing. *SIAM review*, 59(1), 65–98.
- Bock, H.G. and Plitt, K.J. (1984). A multiple shooting algorithm for direct solution of optimal control problems. *IFAC Proceedings Volumes*, 17(2), 1603–1608.
- Brok, N.L., Madsen, H., and Jørgensen, J.B. (2018). Nonlinear model predictive control for stochastic differential equation systems. *IFAC-PapersOnLine*, 51(20), 430– 435.
- Egerstedt, M., Wardi, Y., and Delmotte, F. (2003). Optimal control of switching times in switched dynamical systems. 42nd Ieee International Conference on Decision and Control (ieee Cat. No.03ch37475), 3, 2138,2139,2140,2141,2142,2143, 2138–2143 Vol.3. doi: 10.1109/CDC.2003.1272934.
- EnergiNet (2016). Regulation C2: The balancing market and balance settlement. *EnergiNet*.
- Haberkern, B., Maier, W., and Schneider, U. (2008). Steigerung der Energieeffizienz auf kommunalen Klranlagen (English: Enhanced energy efficiency in waste water treatment plants). iat-Ingenieurberatung fr Abwassertechnik, Darmstadt und Stuttgart.
- Halvgaard, R.F., Vezzaro, L., Munk-Nielsen, T., Tychsen, P., and Madsen, H. (2017). Stochastic Greybox Modeling

for Control of an Alternating Activated Sludge Process. DTU Compute. (DTU Compute-Technical Report-2017, Vol. 08), Lyngby, Denmark.

- Henze, M., Gujer, W., Mino, T., and van Loosdrecht, M.C.M. (2000). Activated Sludge Models: ASM1, ASM2, ASM2d and ASM3. Scientific and Technical Report no. 9. IWA Publishing, London, UK.
- Huang, X., Han, H., and Qiao, J. (2013). Energy consumption model for wastewater treatment process control. *Water Science and Technology*, 67(3), 667–674. doi: 10.2166/wst.2012.563.
- Jacobson, M.Z. and Delucchi, M.A. (2011). Providing all global energy with wind, water, and solar power, part i: Technologies, energy resources, quantities and areas of infrastructure, and materials. *Energy Policy*, 39(3), 1154–1169. doi:10.1016/j.enpol.2010.11.040.
- Kristensen, N.R., Madsen, H., and Jørgensen, S.B. (2004). Parameter estimation in stochastic grey-box models. *Automatica*, 40(2), 225–237.
- Mizuta, K. and Shimada, M. (2010). Benchmarking energy consumption in municipal wastewater treatment plants in japan. Water Science and Technology, 62(10), 2256– 2262. doi:10.2166/wst.2010.510.
- Morales, J.M., Conejo, A.J., Madsen, H., Pinson, P., and Zugno, M. (2014). Integrating Renewables in Electricity Markets. Springer. doi:10.1007/978-1-4614-9411-9.
- Pabi, S., Armanath, A., Goldstein, R., and Reekie, L. (2013). Electricity Use and Management in the Municipal Water Supply and Wastewater Industries. Electric Power Research Institute (EPRI).
- Parbo, H. (2014). Balancing management with large shares of renewables. In 2nd Annual European Forum on Grid Integration and Electricity Anciliar Services.
- Rosso, D., Larson, L.E., and Stenstrom, M.K. (2008). Aeration of large-scale municipal wastewater treatment plants: state of the art. *Water Science and Technology*, 57(7), 973–978. doi:10.2166/wst.2008.218.
- Shi, C.Y. (2011). Mass flow and energy efficiency of municipal wastewater treatment plants. Water Intelligence Online, 10. doi:10.2166/9781843393825.
- Stentoft, P.A., Munk-Nielsen, T., Vezzaro, L., Madsen, H., Mikkelsen, P.S., and Møller, J.K. (2018). Towards model predictive control: Online predictions of ammonium and nitrate removal by using a stochastic asm. WRRmod 2018 proceedings.
- Ueckerdt, F., Brecha, R., and Luderer, G. (2015). Analyzing major challenges of wind and solar variability in power systems. *Renewable Energy*, 81, 1–10. doi: 10.1016/j.renene.2015.03.002.
- Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical* programming, 106(1), 25–57.
- Zemtsov, N., Hlava, J., Frantsuzova, G., Madsen, H., and Jørgensen, J.B. (2018). Economic nonlinear mpc for a population of thermostatically controlled loads. *Computer Science - Research and Development*, 33(1-2), 157–167. doi:10.1007/s00450-017-0354-7.
- Zhao, H., Freed, A.J., Dimassimo, R.W., Hong, S.N., Bundgaard, E., and Thomsen, H.A. (2004). Demonstration of Phase Length Control of BioDenipho Process Using On-line Ammonia and Nitrate Analyzers at Three Full-Scale Wastewater Treatment Plants. WEFTEC.

Applied Energy 280 (2020) 115965

Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apenergy



Unlocking energy flexibility of municipal wastewater aeration using predictive control to exploit price differences in power markets



Niclas Brabrand Brok^{a,*}, Thomas Munk-Nielsen^b, Henrik Madsen^a, Peter A. Stentoft^{a,b}

^a Department of Applied Mathematics and Computer Science, Technical University of Denmark (DTU), Lyngby, Denmark ^b Krüger A/S, Veolia Water Technologies, Søborg, Denmark

ARTICLE INFO

Keywords: Smart energy application Energy flexibility Optimal control Wastewater treatment Optimization Balancing market Price-based control

ABSTRACT

To accommodate the increasing amount of renewable energy sources in electricity grids it is crucial to utilize the flexibility of all electricity consumers. Municipal wastewater treatment plants consume approximately 1 % of the electricity consumption of a country's total electricity consumption, to reduce nutrient concentrations from the incoming wastewater before discharging the water back into the environment. In this paper, a novel economic optimal control strategy is proposed for unlocking the available energy flexibility in wastewater treatment. The strategy suggests that the power consumption in wastewater treatment can be flexible since the water is treated in large tanks with long retention times where specialized aeration equipment is repeatedly switched on and off. By controlling these switching times with respect to nutrient concentrations, electricity consumption can be predicted and shifted in time and hence provide short-term demand side flexibility. The proposed principle is used to reduce the operating costs of a wastewater treatment plant by enabling the flexibility to distribute the aeration load to periods with less expensive power prices. The performance of the proposed method is demonstrated for the operation of a single wastewater treatment plant and the strategy is backtested on Nord Pool market data from 2019. This shows modest savings of 1.15% if only the dayahead market is considered. However, if the regulating and special regulating power prices are included in the optimization the realized savings are in the magnitude of 7.23% and 27.32%, respectively. Thus this study is considered as a step towards exploiting flexibility for the benefit of reducing the combined taxation and aeration cost of wastewarer treatment.

1. Introduction

Advancements in wind turbines and solar panels, coupled with societal demand for low carbon energy, declines in cost and favorable regulation, have all contributed to an increase in power production based on fluctuating renewable energy resources in power grids [1]. While this is a natural development to reach the desired fossil-free energy system, sustainable integration of these power sources requires a paradigm shift in electricity demand [2]. This means a transition from electricity demand independent of production towards a demand side which adapts to the weather driven power production, without significant loss of utility for the users. Such adaptation is suggested and quantified for several applications, including buildings that can maintain a constrained temperature by optimally exploiting electricity prices [3] or by exploiting local photovoltaics [4], and electric vehicles that optimize their charging schedules using photovoltaic power [5]. However, as more fluctuating and intermittent energy sources are introduced, more flexible applications will be needed to efficiently utilize these renewable energy sources [6].

Municipal Wastewater Treatment Plants (WWTPs) use electricity in various processes including pumping, aeration, mixing, and desanding. In total, WWTPs are responsible for approximately 1% of a country's total electricity consumption [7], meaning that Germany and the USA spend 4.4 and 30.2 TWh each year on wastewater treatment, respectively [8,9]. Furthermore it has been found that electricity is costly for plant operation, with cases of electricity costs corresponding to 25-50% of operating costs [10]. Different studies have investigated the electricity consumption of WWTPs and concluded that the process with the highest power demand is the aeration process, constituting 40-75% of the total power consumption [11,12]. Aeration of the wastewater is carried out in large, engineered tanks in which specialized bacteria need aerobic conditions (oxygen present) to convert ammonium from e.g. urine to nitrate. Other bacteria then convert nitrate to nitrogen gas under anoxic conditions (oxygen not present) and hence nitrogen is removed from the water. These processes can be modeled using activated sludge models [13] and predicted using real-time data and

https://doi.org/10.1016/j.apenergy.2020.115965

Received 5 June 2020; Received in revised form 10 September 2020; Accepted 29 September 2020

Available online 16 October 2020 0306-2619/© 2020 Elsevier Ltd. All rights reserved.

^{*} Corresponding author.

E-mail addresses: nlbr@dtu.dk (N.B. Brok), pas@kruger.dk (P.A. Stentoft).

process knowledge [14]. The biochemical process requires both aerobic and anoxic periods or areas in the tank in order to reduce nutrient concentrations. To secure this, some advanced aeration controls aim at turning aeration on and off in feedback loops to achieve the desired effluent concentrations [15].

Model Predictive Control (MPC) is an advanced control method that aims to find the optimal control strategy by predicting the effect of different control sequences and then choosing the sequence that minimizes some objective/cost function. Using the MPC principle, it is possible to incorporate time-varying signals (e.g. time-varying power prices) into the optimal control problem and use these signals directly in decision making. Hence, this can be seen as an alternative to using a battery and then have "constant", averaged prices. To the authors' knowledge, the idea of including electricity prices in control and operation of WWTPs was first mentioned three decades ago in [16]. Since then, several studies have concluded that wastewater transport and treatment can provide flexibility in their power demand [17–19]. [20] minimized aeration costs with respect to varying energy tariffs using a benchmark model of a WWTP [21]. [22] used an MPC approach to minimize total costs in terms of tax payment on effluent nitrogen and time-varying power prices for a small Danish WWTP. In [23] the pumps that direct wastewater to the WWTP from upstream sewage storage were controlled using MPC with respect to time-varying power prices.

1.1. Paper contribution

The existing literature contains (to the knowledge of the authors) very limited work on predictive optimization of the power consumption at municipal WWTPs with simultaneous consideration of a spot market and balancing power market. This paper presents an optimal control strategy for optimal control of the aeration equipment at WWTPs. In [24], this control strategy is defined in more details, however, in this paper the performance of the strategy is tested extensive on a full year of market data.

Due to time-varying power prices, WWTPs can shift the aeration process schedule such that it favors less expensive power prices and thereby unlocking the energy flexibility available within the biochemical process in the wastewater aeration tanks. Such strategies yield potential within energy balancing [25] as well as solving ancillary service problems in future smart grids [26]. The optimal control strategy proposed in this paper optimally balances costs associated with the nonlinear biochemical process (tax on the effluent total nitrogen concentration) and the power consumption. The cost associated with the power consumption will be subject to four different pricing assumptions. The main contribution of this paper is to demonstrate how WWTPs can optimize their operational costs by shifting their power consumption to operating hours with favorable power prices and discuss the potential of using this flexibility in power consumption in multiple electricity markets. The performance of the optimal control strategy is backtested on Nord Pool market data from 2019 and the benefit of utilizing time-varying power prices is shown. To make the backtest computational tractable the MPC principle has been applied such that the optimal control strategy is solved repeatedly two days ahead with only the first 24 h being implemented.

2. The Northern European electricity market

In Northern Europe, electricity is traded in a common market which consist of 15 interconnected price areas where electricity can be purchased and sold in a range of markets. On the NASDAQ Commodities Exchange financial derivatives are traded and market participants can hedge against e.g. grid congestion and general price volatility. These financial derivatives are in general traded with a long horizon — prices can be hedged for the coming week, month or year. The Nord Pool Exchange is responsible for operating two power markets in Northern

Europe: the day-ahead market and the intra-day market. The dayahead market has the largest trading volume and electricity can be purchased and sold for the up-coming day. Bidding to the day-ahead market closes at 12 pm on the day prior to the day of operation, and hence the planning horizon is 12-36 h. The day-ahead price is settled as the intersection between the ordered bids of power purchases and sales constrained by the physical limitations of the power grid. The intra-day market is open for trading when the day-ahead market closes. This market has traditionally been used to eliminate imbalances between the expected power production/consumption and the power production/consumption plans traded in the day-ahead market. The trading volume in the intra-day market is significantly smaller than the trading volume in the day-ahead market. The intra-day market closes 45 min prior to the hour of operation but is open for trading up until then. Simultaneously with the closure of the intra-day market, bidding to the regulating power market also closes. The regulating power market is a market operated and settled by the Transmission System Operators (TSOs) in the Nordic countries and all regulating power bids are aggregated within the Nordic Operational Information System (NOIS). The regulating power market is one (of many) mechanism the TSOs use to maintain grid stability. The regulating power market comes into play if the grid imbalance is expected to be of significant volume and duration. Given that the system imbalance is expected to be significant, the sign of the imbalance define:

- (1) If the imbalance is negative, there is a deficit of electricity an increase of the production or a decrease in the consumption is needed. This is called up regulation.
- (1) If the imbalance is positive, there is a surplus of electricity a decrease of the production or an increase in the consumption is needed. This is called down regulation.

The TSO works as an intermediary in the regulating power market. The TSO sell electricity when there is up regulation and buy electricity when there is down regulation. The counter parties responsible for the these imbalances are settled in the balancing power market. In this market the TSO also functions as an intermediary. The balancing power market is settled after the operating hour when the actual meter measurements have been made. Thus, the regulating power market works as a neutralization mechanism to alleviate imbalances and the balancing power market is a physical market which holds the responsible parties accountable for their grid imbalances.

The regulating power market functions similarly to the day-ahead market across the 15 interconnected price areas. The regulating power price might differ between price areas but is always settled as the intersection between the ordered bids of down- and up regulation and the activated regulating power volume. Hence, all activated bids are activated at the same price. The day-ahead price is used as the basis for the regulating- and balancing power pricing procedures:

- The down regulation price is less than or equal to the day-ahead price.
- The up regulation price is greater than or equal to the day-ahead price.
- Consumption imbalances with an opposite sign of the system imbalance is settled at the regulation power price in the balancing power market. Production imbalances with an opposite sign of the system imbalance is settled at the day-ahead price.
- Imbalances with the same sign as the system imbalance is settled at the regulating power price in the balancing power market.

This structure implies that consumption imbalances are settled at the regulating power prices in the balancing power market while production imbalances always are settled at the highest price of the regulating power price and the day-ahead price. This settlement structure for production imbalances imply a net profit for the TSO, which is used to partially finance the reserve capacity and thereby implicitly lowering the system tariff.

Table 1

GWh of Special down regulation received from the German TSO, TenneT, for 2017– 2019. The activated volumes are activated volumes for Danish participants and the average price is in EUR/MWh.

	2019	2018	2017
Special down regulation	1,312	1,114	781
Thermal power plants	46 %	53 %	64 %
Electric boilers	22 %	21 %	22 %
Wind turbines	32 %	26 %	14 %
Average price (EUR/MWh)	-12.3	-9.3	-7.7

In recent years there has been a rapid development of the activated volumes of a market feature called special regulation in the Western Danish price area, DK1. Special regulation can occur as both downand up regulation. Special regulation is in effect when the TSO skips the normal pricing procedure and in the event of special regulation, the activated bids are paid as bid. During the last three years, one of the main sources of special down regulation in DK1 has been excessive wind power production in Northern Germany. In Table 1 the total special down regulation volumes activated in Denmark due to special regulation from the German TSO, TenneT, is shown. This table shows that in 2019 a total of 1312 GWh was activated for special down regulation in Denmark due to unforeseen large (positive) imbalances in Northern Germany. This table also shows that in 2019, 32% of the special regulation was handled by curtailment of wind turbines in Denmark. Energinet (Danish TSO) is only allowed to publish yearly average prices - in 2019 the average special down regulation price was -12.3 EUR/MWh. In general Table 1 shows that the need for special down regulation in DK1 is increasing with corresponding power prices being vary attractive for consumers. Thus, there is a great potential for flexible consumers to participate in this market and ideally eliminate the need of curtailing Danish wind turbines.

In Fig. 1 a comparison between the market volumes in the dayahead market, intra-day market and the regulating power market are shown. The left y-axis shows the weekly volumes in the day-ahead market and intra-day market and the right y-axis shows the weekly volumes in the regulating power market. This figure clearly shows that the day-ahead market has the largest trading volume and that demand for special down regulation is significantly larger than the demand for regular down regulation. In Fig. 2 a monthly comparison between the day-ahead prices and the regulating power prices are shown (measured on the left y-axis). The solid black line (measured on the right yaxis) represents the monthly average price difference (the monthly average day-ahead price minus the monthly average regulating power price). This figure shows a tendency towards cheaper average regulating power prices. This figure also shows that the power prices were the most expensive in the first quarter of 2019. Similarly to Fig. 2, Fig. 3 shows the hourly price differences between the day-ahead prices and the regulating power prices. As expected, during the morning peak hours (07-10) and the evening peak hours (17-21) the power prices tends to be more expensive compared to the nightly power prices. This figure also shows that the regulating power prices on average are 1-3 EUR/MWh less expensive than the day-ahead prices leading up to these peak hours. In general the power prices (both the day-ahead prices and the regulating power prices) exhibit large, systematic variations implying a great potential for optimizing according to time-varying power prices.

The aeration equipment at municipal WWTPs use a significant amount of electricity when it is switched on. However, operators of such plants are able to schedule this aeration process such that it favors less expensive power prices. This paper will focus on four different power price assumptions where the WWTP is price-taker of:

- · a constant power price,
- a time-varying power price equal to the day-ahead price,
- · a time-varying power price equal to the regulating power price,

 and a time-varying power price equal to the regulating power price, but where the WWTP also can bid for the attractive special down regulation prices.

It will be assumed that these prices are fully known to the WWTP when scheduling the aeration equipment. This also means that the WWTP is assumed to be a price-taker of all the prices available. This is usually what happens in real-life application; consumers deal with an intermediary (e.g. an aggregator) which manages the power trading. In this paper it is assumed that this intermediary has perfect market information and makes the actual prices available to the WWTP. In reality the prices will be subject to uncertainty and the intermediary will typically charge a fee and take a margin to manage the power trading.

3. Wastewater treatment modeling and optimization

For real-time applications, models that are suitable for learning from all available/relevant data and sensors are needed. This section briefly describes how concentrations of selected nutrients in wastewater can be modeled using Stochastic Differential Equations (SDEs). This section also describes an optimal control strategy for optimally balancing the tax of nutrient concentrations in the effluent wastewater and the power cost associated with the wastewater aeration process. A similar model and optimization strategy is introduced and implemented in [24].

3.1. Wastewater treatment: Applied SDE modeling

Physical systems can be represented using SDEs which provide a natural method to represent how mean and variance of a phenomenon evolves in continuous time. As opposed to discrete time models, prior knowledge about the differential equations governing the system dynamics can be included, and the estimated parameters do not depend on the sampling time. Using the same notation as [24], a general form of an SDE written in integral form is

$$x(t) = x_0 + \int_0^t f(x(\tau), u(\tau)) d\tau + \int_0^t g(x(\tau), u(\tau)) d\omega(\tau),$$

or in short

$$dx(t) = f(x(t), u(t)) dt + g(x(t), u(t)) d\omega(t),$$
(1)

where $x: \mathbb{R} \to \mathbb{R}^{n_x}$ denotes the states, x_0 is the initial value of the states, $u: \mathbb{R} \to \mathbb{R}^{n_u}$ is the input variables and $\omega: \mathbb{R} \to \mathbb{R}^{n_\omega}$ denotes a standard (possibly multivariate) Brownian motion. Brownian motion is defined by its independent increments which satisfy that for each $s, t \in \mathbb{R}$, $\omega(t) - \omega(s)$ is normally distributed with zero mean and covariance I|t - s|. $f: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ is often referred to as the drift function, while $g: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x} \times \mathbb{R}^{n_\omega}$ is called the diffusion function. For further reading the authors refer to [27,28] which represent the foundation for this use. Thereby, SDEs can provide a continuous time stochastic modeling framework which can be suitable for both parameter and state estimation [28,29]. This framework has previously been applied to wastewater treatment in [14,24,30], where the activated sludge model in [31] is reduced to a low-order (number of states are reduced from 13 to 2) SDE model. This paper applies the reduced order SDE model described in [24] and repeated here

$$dx_{1}(t) = a_{1}(a_{2} - x_{1}(t)) dt$$

$$- u(t)a_{3} \frac{x_{1}(t)}{a_{4} + x_{1}(t)} dt + \sigma_{1} d\omega_{1}(t),$$

$$dx_{2}(t) = a_{1}(a_{5} - x_{2}(t)) dt + a_{3}u(t) \frac{x_{1}(t)}{a_{4} + x_{1}(t)} dt$$

$$- (1 - u(t))a_{6} \frac{x_{2}(t)}{a_{7} + x_{2}(t)} dt + \sigma_{2} d\omega_{2}(t),$$

(2)

in the optimal control problem (described later). In the model (2), x_1 represents the ammonium concentration and x_2 represents the nitrate concentration of the wastewater. The parameters are estimated

1 (0)

/

(1) 1/



Fig. 1. Weekly day-ahead, intra-day and regulating power volumes in 2019.



Fig. 2. Monthly prices differences between the day-ahead market and the regulating power market in 2019. The filled regions define the histograms and the dashed lines represent the 25%, 50% and 75% percentiles, respectively.



Fig. 3. Hourly prices differences between the day-ahead market and the regulating power market in 2019. The filled regions define the histograms and the dashed lines represent the 25%, 50% and 75% percentiles, respectively.

from online measurements from a WWTP using the estimation method in [29] and applied in [14]. This methodology has previously been used to model predict plant performance up to 24 h ahead with good performance [14]. The same model has been used as basis for comparing different control strategies in [22]. The parameter estimates used in this paper are given in Table 2. The input function, *u*, in (2) is a binary-valued function, which is 1 if the aeration system is activated and 0 otherwise. This means that the system may be characterized as a switched dynamical system, where the optimal control is related to control of switching times between the two regimes [32]. In an online implementation, the initial state values, $x(0) = (x_1(0), x_2(0))'$, will be estimated when new measurements become available. In the backtest conducted later, these values will be assumed to be known and it will be assumed that the model used in optimal control problem is exact. Hence, no filtering techniques or re-estimation of parameters are needed.

3.2. The optimal control problem

This section defines the optimal control problem that is used to compute the optimal aeration cycles for the WWTP. The notation used and the structure in the optimal control problem is inspired by [24]. The aeration function, u, will be parameterized by

$$\tau = \left\{ \left(\tau_{\text{off},k}, \tau_{\text{on},k} \right) \right\}_{k=0}^{N},\tag{3}$$

such that

$$u(t;\tau) = \begin{cases} 1, & t \in [\tau_{\text{on},k-1}, \tau_{\text{off},k}[, k=1,\dots,N] \\ 0, & t \in [\tau_{\text{off},k}, \tau_{\text{on},k+1}[, k=0,\dots,N-1]. \end{cases}$$
(4)



Fig. 4. This figure shows a schematic overview of the components in the optimal control problem and how the energy flexibility is unlocked by a coupling to electricity markets. NOIS represents participation in the regulation power market.

Table 2

Values and description of the parameters for the model given in (2). The units are in milligrams of nitrogen per liter (except for the rates which are unit-less). The parameters are estimated from real data in [24].

-		
Shorthand	Description	Value
x ₁ (0)	Initial ammonium conc.	1.12
$x_2(0)$	Initial nitrate conc.	0.87
<i>a</i> ₁	Incoming wastewater rate	0.00067
a ₂	Mean incoming ammonium	36.9
a3	Nitrification rate	0.073
a_4	Monod kinetic constant	0.1
a5	Mean incoming nitrate	2.00
a ₆	Denitrification rate	0.300
a7	Monod kinetic constant	7.84
σ_1	Ammonium noise parameter	0.0085
σ_2	Nitrate noise parameter	0.026

In (4), N defines the number of aeration cycles, and this is an integer value that defines the number of decision variables of the optimal control problem. Hence, different values of N will lead to different optimal aeration cycles. A real-life WWTP often has policies restricting the minimum and maximum length of the aeration cycles and the time between these cycles. This is to reduce strain of the aeration equipment and to avoid sludge settling. To accommodate cycle length constraints, the aeration parameters in (3) are ordered according to

$$\tau_{\text{on},k} \le \tau_{\text{off},k}, \quad k = 1, \dots, N, \tag{5a}$$

$$\tau_{\text{on},k+1} \ge \tau_{\text{off},k}, \quad k = 0, \dots, N-1,$$
 (5b)

This will enable the definition of cycle policies by the WWTP. There is no start-up cost for initiating/ending an aeration cycle. There is a natural dynamical delay until the aeration equipment has distributed the air in the entire tank, however, this is delay is captured by the SDE model (2) [14].

The main objective of a WWTP is to reduce nutrients from the incoming wastewater to a sustainable level before discharging the water back into the environment. A sustainable level arguably depends on the vulnerability of the local environment into which the wastewater is discharged, and hence it is typical for a WWTP to implement constraints on ammonium and the total nitrogen concentrations in the effluent. These concentrations can be computed according to x_1 and $x_1 + x_2$, respectively. In this study, hard constraints on concentrations are not imposed, but rather a cost is imposed for the total discharged nitrogen concentration. This means that large concentrations are penalized in the objective function. This is similar to the current Danish legislation, where the total discharged nitrogen concentration is taxed. In addition to the cost related to the total discharged nitrogen concentration, the cost of electricity related to the aeration process is considered. Fig. 4 shows a schematic overview of the components of the wastewater treatment considered in the optimal control problem. It will be assumed

that the WWTP receives a single price signal, p (from e.g. a retailer). This input function will be subject to various price model assumptions (as described earlier). The special regulating price will be subject to a qualitative assumptions relating to the average prices listed in Table 1.

The objective function used in the optimal control problem will consider a monetary cost of the total discharged nitrogen and a monetary cost of the power consumption related to the aeration cycles. The monetary cost related to power consumption will be based on a single price signal, p, which the WWTP receives from e.g. a retailer. In the backtest conducted later, this price signal will be subject to different pricing assumptions. The optimal control problem can be defined as

$$\min_{\tau} k_p \sum_{k=0}^{N-1} \int_{\tau_{\text{on},k}}^{\tau_{\text{off},k+1}} p(t) \, \mathrm{d}t - k_t \int_0^T \left(x_1(t) + x_2(t) \right) \, \mathrm{d}t.$$

subject to

$$t \in I$$
, (6a)

$$\dot{x}(t) = f(x(t), u(t; \tau)), \quad t \in [0, T],$$
(6b)

$$\mathbf{x}(0) = \mathbf{x}_0. \tag{6c}$$

In (6a), \mathcal{T} defines the set of permissible cycle structures. This might include constraints such as minimum and maximum cycle length. In (6b) f is the system model which is given from (2) with the parameters given in Table 2. T > 0 denotes the optimization horizon and bounds all the decision variables, τ . $k_p > 0$ is the power used when the aeration equipment is in operation and hence the first term of the objective function defines the integration cost of the power consumption. $k_t > 0$ is the taxation factor which models the monetary cost of the accumulated effluent flow. This construction implies that as the integrated flow through the plant increases, the total tax paid increases similarly. This works as a penalty that favors low nutrient concentrations. x_0 is the initial condition. The initial condition to initiate the backtest is given in Table 2. Note that as ammonium concentrations (x_1) and nitrate concentration (x_2) are typically observed as milligram nitrogen (N) per liter (including in this study), no stoichiometric weights are required to estimate the total N concentration.

3.3. Gradient calculations and numerical implementation

To enable efficient numerical solution of the optimal control problem (6), gradient information has to be made available to the optimization algorithm. The numerical method used in this paper is a single-shooting method [33,34], where the optimization algorithm solves the dynamical Eqs. (6b) internally. Hence, the only constraints implemented in the optimization are the aeration cycle constraints (6a). The constraints defined in (6a) are linear and hence the derivatives are trivial. However, the derivatives of the objective function in (6) are non-trivial to compute. As stated earlier in this paper, the system considered is a switched dynamical system. In [32] it is derived how



Fig. 5. Relative savings for the 2019 backtest. The constant pricing model is used as the benchmark for the relative savings.

the derivatives of the objective function with respect to the switching times can be computed using the adjoint (or co-state) equations of the optimal control problem (6). The optimal control problem (6) has been solved in C++ using odeint [35] to integrate the dynamical equations in (6b), CppAD [36] to compute the derivatives of the control problem, and ipopt [37] is used as the nonlinear optimizer.

4. Results: 2019 backtest

This section presents the results from backtesting the optimal control strategy in (6) on Nord Pool market data from 2019. The method is tested on the four price models described earlier. The constant price model will assume that *p* is equal to the average power price in 2019. The average power price in DK1 in 2019 was approximately 38.5 EUR/MWh. Most municipal WWTPs in Denmark have a constant price agreement. The solution to the optimal control problem with a constant power price will therefore be used as benchmark for evaluating the value of the WWTP flexibility. $k_p = 1$ and $k_t = 1/2$ are split such that the electricity cost corresponds to approximately 70% of the total cost and the taxation to approximately 30% of the total cost for the constant price model. In order to make the backtest computationally tractable the MPC principle has been used:

- · the optimal control problem is solved for 48 h ahead, and
- the first 24 h are implemented.

The structure is repeated for the entire 2019. It is assumed that the dynamical model is exact and hence the dynamical model in (6b) is also used to simulate the biochemical process 24 h ahead.

The constant price model is used as the benchmark for evaluating the performance and value of the WWTP flexibility for the remaining three price models. The three time-varying price models follow the assumptions:

- For the day-ahead price model, it is assumed that the WWTP is price-taker of the day-ahead price.
- *p* is equal to the day-ahead price 0–48 h ahead.
- For the regulating power price model, it is assumed that the WWTP is price-taker of the regulating power price. *p* is equal to the regulating power price 0–48 h ahead.
- For the regulating power price model with special regulating, it is assumed that the WWTP is price-taker of the regulating power price but is activated for special down regulation. *p* is equal to $p_s \in \mathbb{R}$ if there is demand for special regulation and everywhere else equal to regulating power price.

For real-life applications, these assumptions are not valid and these prices will to be subject to forecasting. In this paper, a 48h optimization

horizon (or prediction horizon) is used. However, this horizon parameter could probably be chosen to be shorter. But, further investigations of this is outside the scope of this paper.

Fig. 5 shows the relative savings of the three price models. It is assumed that $p_s = 0$ EUR/MWh for the regulating power price model with activation for special regulation shown in Fig. 5. In Table 3 the performance metrics for the entire 2019 is shown. The costs are normalized by the total cost in 2019 for the constant price model. From Fig. 5 and Table 3 the following is observed:

- The constant price model yields the lowest taxation cost, but the highest electricity cost.
- The time-varying price models with no activation for special down regulation underperform the constant price model in the first quarter of 2019.
- The price model with activation for special down regulation consistently outperforms the price models with no activation for special down regulation. The most conservative of these price models ($p_s = 0$ EUR/MWh) outperforms the constant price model and reduces the total costs by 27.32% in 2019. However, this model increases the taxation cost by 9.79%.
- The regulating power price models exhibit significantly larger savings than day-ahead price model with a limited added taxation cost of the effluent.

In Fig. 6 the operation on the 20th of March 2019 is shown as an example. The optimal aeration cycles have been computed using the regulation power price model with activation for special regulation with $p_s = 0$ EUR/MWh. It is clearly observed that the aeration cycles condense around the demand for special regulation (when the WWTP can use electricity for free) and the aeration equipment is in general switched on when the regulating power price is low. The effluent total nitrogen concentrations exhibit are clearly burst when the aeration cycles condense - this happens at around 08:00 AM. Similarly, the ammonium effluent concentrations peak when aeration is turned off for longer periods of time. This is the case at around 7:00 AM. The figure also shows the baseline scenario where the on/off switching is controlled with respect to constant electricity prices. In this case it is noted that the control signal follows the same cycle-pattern during the 24 h period. The average effluent concentrations during this period are 1.19 and 1.37 mgN/L for ammonium and nitrate respectively in the constant electricity price scenario. This is 7.6% and 2.1% lower than the effluent concentrations in the other scenario. However, this increase in effluent results in a decrease in total costs (savings) of 22.5% while helping the grid with regulating power. Whether this is a feasible trade-off will require some holistic considerations of the receiving environment, and possibly constraints on the effluent. This is further discussed in the following section.

Table 3

Relative savings for the 2019	backtest. The constant price	e model is used as the benchman	k for the relative savings.
-------------------------------	------------------------------	---------------------------------	-----------------------------

Strategy	Taxation cost	Electricity cost	Total cost
Constant price model	0.31	0.69	1.00
Day-ahead price model	0.33 (4.67%)	0.66 (-3.79%)	0.99 (-1.15%)
Regulating power price model	0.34 (8.34%)	0.59 (-14.32%)	0.93 (-7.23%)
Regulating power price model w. $p_s = 0$ EUR/MWh	0.34 (9.79%)	0.38 (-44.22%)	0.73 (-27.32%)
Regulating power price model w. $p_s = -5$ EUR/MWh	0.35 (11.29%)	0.34 (-49.92%)	0.69 (-30.77%)
Regulating power price model w. $p_s = -10$ EUR/MWh	0.36 (13.70%)	0.30 (-56.93%)	0.65 (-34.83%)
Regulating power price model w. $p_s = -15$ EUR/MWh	0.36 (16.24%)	0.24 (-64.76%)	0.61 (-39.41%)



Fig. 6. Example day from the backtest: 20th March 2019. The regulating power price model with $p_s = 0$ EUR/MWh is used to design the optimal aeration cycles. The top plot shows the price input to the price models: the left *y*-axis measures the demand for special down regulation and the right *y*-axis measures the day-ahead and regulating power prices. The middle plot shows the state realization using the aeration sequence shown in the same plot. This is illustrated for the regulating power and special regulating price model and the baseline with constant electricity price (dashed lines). The bottom plot shows the savings relative to the constant price model.

5. Discussion and future work

The presented control strategy demonstrates that flexibility in power consumption can be enabled for the aeration process at a WWTP. This section discusses some of the important factors that influence the operational costs of applying the control strategy, and ideas for enabling more flexibility are suggested.

5.1. Flexibility in aeration equipment control

The results presented in the previous section are subject to multiple key assumptions that may be too simplistic. The parameters used in the model are in this paper kept constant. In an online application, parameters are expected to be frequently updated as suggested in [14]. This means that the plant can be more flexible in some periods (e.g. periods with low incoming nutrient loads which cause a decreased demand for aeration) and less flexible in other periods (e.g. during rain events that increase the demand for aeration by sending more wastewater to the plant). This could be accommodated by forecasting incoming ammonium loads and/or the weather.

The aeration control sequence is limited by hard constraints concerning the time the equipment can be switched on and off. This implies that the constraints govern how often changes can be made in the aeration control and thereby in the electricity consumption. These constraints will vary among different WWTPs, and hence some plants may have different potential in unlocking their flexibility.

The control strategy only considers biological nitrogen removal, meaning that the biological phosphorous removal is neglected. However, this would be useful to include for many WWTPs, as phosphate is a nutrient often managed in the biological treatment (i.e. by controlling the aeration equipment). This could be added by inclusion of

N.B. Brok et al.

phosphate in the prediction model and in the optimal control problem, as suggested in [38]. This might reduce the flexibility potential, but further investigations should be made to show this effect. It is noted that phosphate can also be managed by using chemical precipitation if aeration is insufficient.

The legislation in many places requires that ammonium and nitrogen concentrations are kept below a certain limit (i.e. hard state constraints). For improved applicability, this should also be investigated in a future implementation. This would likely add an upper bound to the flexibility, as priority would be given to satisfy legislation rather than costs of electricity consumption. This could be implemented by penalizing high nutrient concentrations in the objective function, as suggested in [22]. For increased flexibility in the urban water system additional processes of the WWTP and drainage system could be included and aggregated. This might increase the amount of flexibility of the combined drainage system can deliver, but also enable the combined drainage system to shift more power demand in time. E.g. by integrating pumping stations and basins in the sewer system as demonstrated in [39]. Alternatively, this could be obtained by using biogas for producing electricity in some scenarios, thereby exploiting that larger WWTPs are prosumers.

5.2. Market price forecasting

The control strategy defined in this study utilizes the multi-market structure of the Nordic electricity market and assumes perfect information about the market prices 48 h ahead. This means that the results presented in this paper will work as an upper bound for the cost reductions feasible in real-life applications. There is primarily two limiting assumptions for the feasibility of the results presented in the previous section:

- the perfect forecast assumption of the regulating power prices 48 h ahead, and
- the assumption about recurrent activation for special down regulation at the predefined price, p_s .

The regulating power price is (almost by definition) very difficult to forecast multiple hours ahead. In the Western Danish price area, DK1, grid imbalances are often related prediction errors in power production originating from wind turbines. One extension that would consider this is to consider a stochastic optimization problem where the price input consists of a set of realistic price scenarios of the regulating power price. For the regulating power price models with activation for special regulation, it is assumed that the WWTP can choose when it is favorable to take the price p_s and define how much and which part of the corresponding hour that will be subject to aeration and hence subject to a power consumption. This might be difficult to achieve for a single WWTP as it is not given when demand for special regulation occurs, and the TSO might be very precise concerning when the activation must happen. However, an aggregator with multiple WWTPs might be more adequate for participation with special regulation as it will be less constraining to distribute activation for special regulation across multiple plants. This aggregator approach might be solved using a similar optimal control method as described in this study by augmenting the dynamical equations with the individual plants in the aggregator portfolio. A second approach could be to utilize the method described in [40] where a penalty signal (e.g. a price input) is used to model the dynamical response to the power load.

In this study it is assumed that the WWTP is price-taker of different price models. An extension to this would be to also allow the WWTP to bid in the day-ahead and the intra-day markets while also being a pricetaker of a given price model (e.g. the regulating power price model). The day-ahead market could then be used to leverage the volume that would be subject to activation in the regulating power market. This would increase the cost reduction potential but also increasing the risk associated with participating in the electricity market. This approach will naturally be bound by the forecasting potential of e.g. the regulating power market.

One practical issue with the implementation of efficient power consumption strategies (e.g. in Denmark), is that the electricity price only constitutes a small part (around 15% in Denmark) of the total electricity bill — whereas time-invariant and constant tariffs and taxes constitute the rest. However, the control strategy proposed in this paper can easily be extended to also consider varying tariffs and taxes (e.g. a time-varying price relating to the CO2-emissions of the power consumption). Thereby the outlined method might lead to even larger potential savings.

6. Conclusion

This paper presents an optimal control problem for optimal distribution of aeration cycles across a day of operation. The potential of the proposed method is tested under different price model assumptions with an objective function that balance the taxation cost of nutrients discharged into the environment with the cost associated with operating the aeration equipment of a Municipal Wastewater Treatment Plant (WWTP). The summarized findings are:

- A constant price model is not a suitable price model for Municipal WWTPs as these plants have the opportunity of shifting their consumption in time to match less expensive power prices, and hence unlocking their energy flexibility. Thus, the inherent dynamics of WWTPs favors time-varying power prices.
- The operational costs related to wastewater treatment aeration can be reduced significantly by participating in the regulating power market. These costs can be reduced even further by also bidding for special down regulating.
- In the backtest for 2019 it is observed that the regulating power market exhibits greater potential than the day-ahead market for reducing operational costs at Municipal WWTPs.
- Further development should be made to fully exploit the flexibility in control of aeration equipment and to make the strategy widely applicable at different wastewater treatment plants (and potentially aggregate multiple WWTPs).

CRediT authorship contribution statement

Niclas Brabrand Brok: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Visualization, Project administration. Thomas Munk-Nielsen: Resources, Data curation, Supervision, Project administration, Funding acquisition. Henrik Madsen: Resources, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition. Peter A. Stentoft: Conceptualization, Methodology, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Visualization, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank Henning Parbo (Chief Economist at Energinet) for many fruitful discussions about the Danish power markets. This work is partly funded by the Innovation Fund Denmark (IFD) under File No. 7038-00097B — Peter A. Stentofts industrial PhD study; *Stochastic Predictive Control of Wastewater Treatment Processes and File No. 7038-00097B, and the IFD project CITIES (1305-0027B).*

N.B. Brok et al.

References

- Ueckerdt F, Brecha R, Luderer G. Analyzing major challenges of wind and solar variability in power systems. Renew Energy 2015;81:1–10.
- [2] Denholm P, Hand M. Grid flexibility and storage required to achieve very high penetration of variable renewable electricity. Energy Policy 2011;39(3):1817–30.
- [3] Halvgaard R, Poulsen NK, Madsen H, Jørgensen JB. Economic model predictive control for building climate control in a smart grid. In: 2012 IEEE pes innovative smart grid technologies (isgt). IEEE; 2012, 6175631.
- [4] Salpakari J, Lund P. Optimal and rule-based control strategies for energy flexibility in buildings with PV. Appl Energy 2016;161:425–36.
- [5] van der Kam M, van Sark W. Smart charging of electric vehicles with photovoltaic power and vehicle-to-grid technology in a microgrid; a case study. Appl Energy 2015;152:20–30.
- [6] Morales JM, Conejo AJ, Madsen H, Pinson P, Zugno M. Integrating renewables in electricity markets: operational problems, Vol. 205. Springer Science & Business Media; 2013.
- [7] Shi CY. Mass flow and energy efficiency of municipal wastewater treatment plants. IWA Publishing; 2011.
- [8] Haberkern B, Maier W, Schneider U. Enhanced energy efficiency in waste water treatment plants; steigerung der energieeffizienz auf kommunalen klaeranlagen. 2008.
- [9] Pabi S, Amarnath A, Goldstein R, Reekie L. Electricity use and management in the municipal water supply and wastewater industries, Vol. 194. Electric Power Research Institute, Palo Alto; 2013.
- [10] Huang X, Han H, Qiao J. Energy consumption model for wastewater treatment process control. Water Sci Technol 2013;67(3):667–74.
- [11] Longo S, d'Antoni BM, Bongards M, Chaparro A, Cronrath A, Fatone F, et al. Monitoring and diagnosis of energy consumption in wastewater treatment plants. a state of the art and proposals for improvement. Appl Energy 2016;179:1251–68.
- [12] Gu Y, Li Y, Li X, Luo P, Wang H, Robinson ZP, et al. The feasibility and challenges of energy self-sufficient wastewater treatment plants. Appl Energy 2017;204:1463–75.
- [13] Henze M, Gujer W, Mino T, van Loosdrecht MCM. Activated sludge models: ASM1, ASM2, ASM2d and ASM3. Scientific and Technical Report no. 9, London, UK: IWA Publishing; 2000.
- [14] Stentoft PA, Munk-Nielsen T, Vezzaro L, Madsen H, Mikkelsen PS, Møller JK. Towards model predictive control: online predictions of ammonium and nitrate removal by using a stochastic ASM. Water Sci Technol 2019;79(1):51–62.
- [15] Zhao HW, Freed AJ, DiMassimo RW, Hong SN, Bundgaad E, Thomsen HA. Demonstration of phase length control of biodenipho process using on-line ammonia and nitrate analyzers at three full-scale wastewater treatment plants. Proc Water Environ Fed 2004;2004(15):215–25.
- [16] Tassou SA. Energy conservation and resource utilisation in waste-water treatment plants. Appl Energy 1988;30(2):113–29.
- [17] Lisk B, Long H. Demand management strategies resulting in "zero and low cost" energy saving opportunities for water and wastewater treatment facilities. In: 86th annual water environment federation technical exhibition and conference, Weftee 2013, Vol. 3. Water Environment Federation; 2013, p. 1705–14.
- [18] Kirchem D, Lynch M, Bertsch V, Casey E. Market effects of industrial demand response and flexibility potential from wastewater treatment facilities. In: International conference on the european energy market, Eem. IEEE Computer Society; 2018, 8469974.
- [19] Reinhofer-Gubisch M, Pucker J. "Loadshift" Lastverschiebung in kommunalen Kläranlagen: Ein relevantes Potenzial für Smart Grids?, "LoadShift": Load shifting in municipal treatment plants: Does it offer real potential for smart grids? In: Osterreichische Wasser- Und Abfallwirtschaft, Vol. 66. 2014, p. 424–32.
- [20] Aymerich I, Rieger L, Sobhani R, Rosso D, Corominas L. The difference between energy consumption and energy cost: Modelling energy tariff structures for water resource recovery facilities. Water Res 2015;81:113–23.

- [21] Gernaey KV. Benchmarking of control strategies for wastewater treatment plants. IWA Publishing; 2014.
- [22] Stentoft PA, Guericke D, Munk-Nielsen T, Mikkelsen PS, Madsen H, Vezzaro L, et al. Model predictive control of stochastic wastewater treatment process for smart power, cost-effective aeration. In: Proceedings of Dycops 2019. 2019.
- [23] Evald Bjerg J, Grum M, Courdent VAT, Halvgaard RF, Vezzaro L, Mikkelsen PS. Coupling of weather forecasts and smart grid-control of wastewater inlet to Kolding WWTP (Denmark). In: Proceedings of the 10th international urban drainage modelling conference. International Association for Hydro-Environment Engineering and Research (IAHR); 2015, p. 47–59.
- [24] Brok NB, Munk-Nielsen T, Madsen H, Stentoft PA. Flexible control of wastewater aeration for cost-efficient, sustainable treatment. IFAC-PapersOnLine 2019;52(4):494–9.
- [25] Madsen H, Parvizi J, Halvgaard RF, Sokoler LE, Jørgensen JB, Hansen LH, et al. Control of electricity loads in future electric energy systems. In: Handbook of clean energy systems. Wiley; 2015.
- [26] Dominković DF, Junker RG, Lindberg KB, Madsen H. Implementing flexibility into energy planning models: Soft-linking of a high-level energy planning model and a short-term operational model. Appl Energy 2020;260:114292.
- [27] Øksendal B. Stochastic differential equations: an introduction with applications. Springer Science & Business Media; 2013.
- [28] Jazwinski AH. Stochastic processes and filtering theory. Courier Corporation; 2007.
- [29] Kristensen NR, Madsen H, Jørgensen SB. Parameter estimation in stochastic grey-box models. Automatica 2004;40(2):225–37.
- [30] Halvgaard RF, Vezzaro L, Grum M, Munk-Nielsen T, Tychsen P, Madsen H. Stochastic greybox modeling for control of an alternating activated sludge process. DTU Compute; 2017.
- [31] Henze M, Jr CPLG, Gujer W, Marais GvR, Matsuo T. Activated sludge model no. 1. Technical Report, IAWPRC; 1987.
- [32] Egerstedt M, Wardi Y, Delmotte F. Optimal control of switching times in switched dynamical systems. In: 42nd IEEE international conference on decision and control (IEEE Cat. No. 03CH37475), Vol. 3. IEEE; 2003, p. 2138–43.
- [33] Diehl M, Bock HG, Diedam H, Wieber P-B. Fast direct multiple shooting algorithms for optimal robot control. In: Fast motions in biomechanics and robotics. Springer; 2006, p. 65–93.
- [34] Bock HG, Plitt K-J. A multiple shooting algorithm for direct solution of optimal control problems. IFAC Proc Vol 1984;17(2):1603–8.
- [35] Ahnert K, Mulansky M. Odeint–solving ordinary differential equations in c++. In: AIP conference proceedings, Vol. 1389. American Institute of Physics; 2011, p. 1586–9.
- [36] Bell BM. Cppad: a package for c++ algorithmic differentiation. Comput Infrastruct Oper Res 2012;57(10).
- [37] Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math Program 2006;106(1):25–57.
- [38] Lindstrøm Sørensen M, Dahl P, Stentoft PA, Munk-Nielsen T, Kloppenborg Møller J. Stochastic model predictive control of phosphorus concentration for smart power, cost-effective municipal wastewater treatment. In: Watermatex 2019, 10th IWA symposium on modelling and integrated assessment. 2019.
- [39] Stentoft PA, Vezzaro L, Mikkelsen PS, Grum M, Munk-Nielsen T, Tychsen P, et al. Integrated model predictive control of water resource recovery facilities and sewer systems in a smart grid: example of full-scale implementation in kolding. Water Sci Technol: J Int Assoc Water Pollut Res 2020;81(8):1766–77.
- [40] Junker RG, Azar AG, Lopes RA, Lindberg KB, Reynders G, Relan R, et al. Characterizing the energy flexibility of buildings and districts. Appl Energy 2018;225:175–82.

Optimal Operation of an Ice-Tank for a Supermarket Refrigeration System *

Niclas Brok^a, Torben Green^b, Christian Heerup^c, Shmuel S. Oren^d, Henrik Madsen^a

^aDepartment of Applied Mathematics and Computer Science at the Technical University of Denmark.

^bDanfoss A/S.

^cDanish Technological Inistitute.

^dDepartment of Industrial Engineering and Operations Research at the University of California at Berkeley.

Abstract

The increasing proportion of renewable energy sources in power grids leads to challenges concerning balancing production and consumption. One solution to this grid challenge is to utilize demand-side flexibility. To use the full potential of demand-side flexibility, dynamical models and optimal control methods must be used. This paper demonstrates how demand-side flexibility can be enabled for a refrigeration system using an add-on ice-tank module to actively curtail the refrigeration system and thereby leveraging time-varying power prices. The operation of the ice-tank is the solution to an optimal control problem that minimize the integrated electricity costs. This optimal control problem is solved numerically and the performance of the strategy is successfully tested in a real experiment where cost-savings of approximately 20 % are observed (compared to not having an ice-tank available). The dynamical relation between the operation of the ice-tank and the power consumption (the compressor capacity) is modeled using stochastic differential equations. This differential equation model is calibrated on 13 hours of training data using the continuous-discrete Kalman filter and the maximum likelihood framework.

Key words: Optimal Control, Refrigeration System, Ice Storage, Electricity Markets, Optimal Switching Times, Optimization, Stochastic Differential Equations

Introduction

Recently, Denmark agreed on its first ever Climate Bill, committing to reduce greenhouse gas emissions by more than 70 % by 2030 compared to the Danish emission levels in 1990 [1]. This will increase the need for renewable energy sources and efficient integration will increase the need for demand-side flexibility [25]. This paper demonstrates a real example of short-term demand-side flexibility by curtailment of the power consumption of a refrigeration system using an add-on ice-tank module. The

ice-tank operation (i.e when the ice-tank curtails the refrigeration system) is the solution to an optimal control problem which minimize the integrated power costs. The refrigeration system is located at the Danfoss test facility in Nordborg, Denmark. This refrigeration system resembles a small retail or supermarket refrigeration system.

Denmark has approximately 4,500 supermarkets distributed across the country. These supermarkets consume more than 550,000 MWh per year, which constitutes about 2 % of the annual Danish power consumption [14]. The power consumption of supermarkets comes from e.g. electric heating, lighting, and cooling. The installed cooling capacity varies a lot between supermarkets but is typically in the range from 10 to 200 kW, depending on the size of the refrigeration system. The refrigerated goods in the refrigeration units can have a large thermal capacity, thereby enabling flexibility in the refrigeration system.

^{*} This work is partially funded by the CITIES project (Danish Innovation found. Grant DSF 1305-00027B) and the ELFORSK project ELIS (project number 345-028). The work has partially been conducted during an external research stay with CITRIS at the University of California at Berkeley where Niclas Brok visited Shmuel S. Oren for 5 months. This external research stay has been partially funded by the Danish Minstry of Higher Science and Education.



Fig. 1. This figure shows an overview of the total refrigeration system used to test the performance of optimally controlling an ice-tank.

The literature contains a lot of work concerning online control and forecasting of supermarket refrigeration systems and their power consumption. In [15], a continuous time model based on ordinary differential equations is introduced. This model is used in an economic model predictive control algorithm, where the total cost of the electricity consumption associated with the refrigeration system is minimized. This economic model predictive control algorithm is extended in [26] to also consider a balancing market to further reduce the operational cost of the refrigeration system. Similarly, [20] applies a model predictive control scheme to optimize the daily operation of a refrigeration system to reduce the power consumption. [11] uses a control strategy to use the demand-side flexibility in a refrigeration system to reduce the peak power demand and thereby reduce the electricity costs. [24] introduces a discrete time model for load forecasting a supermarket refrigeration system. This model distinguishes between supermarket opening hours and closing hours to improve longer horizon forecasts. [2, 12] discusses the potential of demand-side flexibility for refrigerated warehouses. In [13], a rule-based control method is used (charging the ice-storage during the night, and melting during daytime) to investigate whether ice-storage is financially feasible for retail CO2-based refrigeration systems. [4] presents a simulation study based on electricity tariffs from New Zealand, where a price-based control strategy is used to optimize a thermal storage for a freezer. [21] compares different control systems to optimize ice-storage for food refrigeration. In [9], a financial analysis of actively operating a battery under Californian market conditions is investigated. This paper demonstrates that, under perfect information conditions, optimally managing the battery and bidding in multiple electricity markets can create much larger revenues compared to participation only in a traditional spot market.

Key contributions and paper organization

The existing literature presents few real experiments which tests the proposed strategies. In the context of the reviewed literature, this work answers the following research questions:

- How can continuous-discrete stochastic systems be used to model the joint refrigeration and ice-tank system?
- What are the cost-savings potential of optimally operating the joint refrigeration and ice-tank system?

To answer these research questions, an optimal control problem is formulated; this optimization problem is based on a stochastic differential equation model which is calibrated to a training data-set with discrete observations. Using this differential equation model, the optimal control problem yields a numerical algorithm that optimally curtails a small retail refrigeration system using the add-on ice-tank module. This model formulation has two main advantages: (1) the estimation process is independent of the sampling rate and can easily manage irregular sampled observations, and (2) the estimation process is based on a maximum likelihood framework. The performance of the optimal control algorithm is tested in a real physical experiment where cost-savings of approximately 20 % are observed (compared to not having an ice-tank available).

This paper is structured as follows: first, the joint refrigeration and ice-tank is introduced. This section also contains an example of operation of this joint system. The second section describes how a continuous-discrete stochastic model (formulated using a stochastic differential equation with discrete observations) can be calibrated to a data-set. This section also presents a onestate model of the compressor capacity of the joint refrigeration and ice-tank system. The third section introduces the Nordic power market. The fourth section



Fig. 2. This figure shows the results of a full-horizon simulation using the first-order model. The dashed grey line is the model predictions, the grey dots are the sub-sampled observations, the solid grey line is the high-frequent observations and the solid black line is the ice-tank input.

defines the optimal control algorithm. The fifth section presents the results of operating the ice-tank in a real experiment using the optimal control algorithm. The paper concludes with a discussion of the results, a future outlook for the proposed optimal control algorithm and a brief summary.

The Refrigeration System

The test refrigeration system is located at a Danfoss test facility in Nordborg, Denmark. The system compressors have a rated power consumption of approximately 12.4 kW. The power consumption of the compressors is the only source of power consumption considered in this paper. At the Danfoss test facility it is possible to simulate an outdoor temperature. In this study, this temperature is fixed at an average of approximately 35 °C. This corresponds to a very warm Danish Summer day. The ice-tank used is an Ice Bear 40, manufactured by Ice Energy¹. It is possible to alternate the mode of this ice-tank using an API service. The ice-tank can be in the following three modes:

- CHARGE, in this mode the ice-tank builds up the icestorage. The rated power consumption in this mode is approximately 3 kW.
- IDLE, in this mode the ice-tank does nothing. This mode has a small power consumption (approximately 9 W), and due to imperfect insulation there is also a small thermal loss (i.e. a small reduction of the ice-storage).
- MELT, in this mode the ice-tank melts the ice-storage and starts curtailing the power consumption of the

refrigeration system. In this mode the ice-tank has a small power consumption (approximately 250 W).

The latency of the API service is in the order of magnitude of 1-2 minutes due to communication via a third party server from which the ice-tank pulls data and commands.

Fig. 1 shows a schematic overview of the test setup. The left part of this illustration is the ice-tank, which is the controllable component of the total refrigeration system. The middle part is the existing refrigeration system, which will be considered as a non-controllable system. The ice-tank is connected to the refrigeration system via a sub-cooler. To artificially emulate an outdoor temperature, a water heater is connected to the refrigeration system's evaporator via a water pump. This water heater can be used to emulate a fixed outdoor air temperature and this has been done in this study.

Note: Tests with lower outdoor temperatures have also been done. However, these results will not be presented, although they will briefly be discussed later.

In the next section, a dynamical model of the dynamical interplay between the refrigeration system and the ice-tank is defined. To calibrate this model to the actual system under consideration, a test data-set has to be generated. To generate this data-set, the ice-tank is alternated through four melt-cycles:

- A cycle with 30 minutes of MELT followed by 30 minutes of IDLE (from 13:00 until 14:00).
- A cycle with 45 minutes of MELT followed by 45 minutes of IDLE (from 14:00 until 15:30).
- A cycle with 60 minutes of MELT followed by 60 minutes of IDLE (from 15:30 until 17:30).

¹ Company website: http://ice-energy.com/



Fig. 3. This figure shows the results of the initial ice-tank test. This data is used to calibrate the dynamical model which will be used in the optimal control problem defined later. The top plot shows the compressor capacity (measured in %) and the bottom plot shows the ice-tank regime and the outdoor temperature (measured in $^{\circ}$ C). The refrigeration system is sampled every six seconds.

• The final melt cycle is 90 minutes long and begins from 17:30. When this cycle ends, the ice-tank remains in the IDLE mode.

These cycles - together with the compressor capacity and the outdoor temperature - can be seen in Fig. 3. The effect of the ice-tank switching to and from the MELT regime is clearly observed. From Fig. 3 it is also seen that the simulated outdoor temperature also exhibits larger oscillations when a mode-change of the ice-tank is implemented. This is due to the change in the work-load of the refrigeration compressors.

In the next section, a method for calibrating a dynamical model (in terms of an SDE) to the test data-set shown in Fig. 3 will be presented.

Modelling using SDEs

Modelling physical systems using SDEs provides a natural method to represent the phenomenon as it evolves in continuous time. In contrast to discrete time models, a priori knowledge about the system can be included, and the estimated parameters do not depend on the sampling time. The representation of noise in continuous time also allows for a parsimonious representation that is independent of the sampling time.

This section describes how to use SDEs to model the dynamical interaction between the refrigeration system and the ice-tank. This joint system can be considered as a continuous-discrete stochastic system [16]. The model used in this paper is an SDE with discretely sampled observations defined according to

$$dx(t) = f(x(t), u(t); p) dt + g(x(t), u(t); p) d\omega(t),$$
(1a)

$$y_k = h(x(t_k); p) + v_k,$$
(1b)

where x, u, y and p are the states, inputs, observations and time-invariant parameters. $v_k \sim N_{\rm iid}(0, R_k)$ is the measurement noise and ω is a standard Brownian motion. Brownian motion is defined by its independent increments which satisfy that for each $s, t \in \mathbb{R}$, $\omega(t) - \omega(s)$ is normally distributed with zero mean and covariance I(t - s); i.e. $d\omega(t) \sim N_{\rm iid}(0, I dt)$. $f: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ is often referred to as the drift function, while $g: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x} \times \mathbb{R}^{n_\omega}$ is called the diffusion function. In this notation, n_x, n_u and n_ω denote the number of states, inputs and Brownian motions, respectively. The stochastic model defined in (1) is also referred to as a continuous-discrete stochastic statespace model where the dynamics are defined according to an SDE with discrete observations.

For the joint system consisting of the ice-tank and the refrigeration system, the state of the system, x, will be the compressor capacity and the input, u, will be the ice-tank mode (i.e. a binary switch, modeling if the ice-tank is in IDLE or MELT mode). The parameters, p, will parameterize the dynamical relations between the ice-tank mode and the compressor capacity. The definitions of the functionals, f, g and h, and the unknown parameters, p, will be defined after the introduction of the continuous-discrete extended Kalman filter (CDEKF) and the associated maximum likelihood (ML) method. For further reading concerning the CDEKF and ML methods for SDEs, the authors refer to [16, 22].

Estimating parameters embedded in SDEs

The method presented next uses the CDEKF to evaluate the likelihood for a batch of data [19]. The CDEKF is based on two schemes: a prediction scheme and an updating scheme. These schemes are briefly introduced in the following. \hat{x}_{k-1} and \hat{P}_{k-1} will denote the mean and covariance predictions of the state, x, while $\hat{x}_{k-1|k-1}$ and $\hat{P}_{k-1|k-1}$ will denote the mean and covariance filtered estimates of x (or reconstructions).

The prediction scheme

Given the initial conditions

$$\hat{x}_{k-1}(t_{k-1}) = \hat{x}_{k-1|k-1}, \quad \hat{P}_{k-1}(t_{k-1}) = \hat{P}_{k-1|k-1}, \quad (2)$$

the state and covariance are predicted by solving the system of ordinary differential equations (ODEs) given by

$$\hat{x}_{k-1}(t) = f(\hat{x}_{k-1}(t), u(t); p),$$
(3a)
$$\dot{\hat{P}}_{k-1}(t) = A(t)\hat{P}_{k-1}(t) + \hat{P}_{k-1}(t)A(t)' + G(t)G(t)',$$
(3b)

where

$$A(t) = \frac{\partial f}{\partial x}(\hat{x}_{k-1}(t), u(t); p), \quad G(t) = g(\hat{x}_{k-1}(t), u(t); p).$$

The one-step predictions of the mean and covariance of the states are obtained as the solution of (2)-(3) at the new sample point, t_k . Consequently, the predictions of the mean and covariance are

$$\hat{x}_{k|k-1} = \hat{x}_{k-1}(t_k), \quad \hat{P}_{k|k-1} = \hat{P}_{k-1}(t_k).$$
 (4)

The updating scheme

The literature contains many methods for the updating scheme in extended Kalman filter algorithms. They all compute the innovation by

$$e_k = y_k - h(\hat{x}_{k|k-1}), \tag{5}$$

the Kalman filter gain, K_k , by

$$C_k = \frac{\partial h}{\partial x}(\hat{x}_{k|k-1}),\tag{6a}$$

$$R_{k|k-1} = C_k \hat{P}_{k|k-1} C'_k + R_k, \tag{6b}$$

$$K_k = \hat{P}_{k|k-1} C'_k R^{-1}_{k|k-1}, \tag{6c}$$

and the filtered state estimate, $\hat{x}_{k|k}$, by

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k e_k.$$
(7)

The key difference is how they compute the filtered covariance, $P_{k|k}$. Two standard updating schemes for the covariance are

$$\hat{P}_{k|k} = (I - K_k C_k) \,\hat{P}_{k|k-1}$$
 (8a)

$$= \hat{P}_{k|k-1} - K_k R_{k|k-1} K'_k.$$
 (8b)

Numerical implementations based on either (8a) or (8b) may give rise to bad performance and even divergence, as the *numerically* computed values are not guaranteed to be both positive (semi-)definite and symmetric. The Joseph stabilization form

$$\hat{P}_{k|k} = (I - K_k C_k) \, \hat{P}_{k|k-1} \, (I - K_k C_k)' + K_k R_k K'_k.$$
(9)

for updating the filtered covariance estimate guarantees that the numerical value of $P_{k|k}$ is symmetric positive (semi-)definite.

Maximum Likelihood Estimation

Using the one-step prediction errors, from the prediction and updating schemes, the likelihood of the model parameters given the discretely sampled observations can be computed; the parameters that maximize this likelihood computation will be used as the model parameters for the model used in the optimal control problem. In this section, this likelihood calculation is introduced.

Let $\{y_j\}_{j=1}^M$ denote M observations relating to the sample points $\{t_j\}_{j=1}^M$ in (1b). Define the information accumulated up until the k-th sample point as $\mathcal{Y}_k = \{y_j\}_{j=1}^k$. Then the likelihood function, \mathcal{L} , can be defined as

$$\mathcal{L}\left(p \mid \mathcal{Y}_{M}\right) \propto \phi\left(\mathcal{Y}_{M} \mid p\right),\tag{10}$$

where ϕ is the joint density function of the observations, \mathcal{Y}_M . Using the definition of conditional probabilities, the right hand side can be decomposed into

$$\phi\left(\mathcal{Y}_{M}\mid p\right) = \prod_{k=1}^{M} \phi\left(y_{k}\mid \mathcal{Y}_{k-1}, p\right), \qquad (11)$$

such that the log-likelihood function can be expressed by

$$\log \left(\mathcal{L} \left(p \mid \mathcal{Y}_M \right) \right) = \sum_{k=1}^M \log \left(\phi \left(y_k \mid \mathcal{Y}_{k-1}, p \right) \right).$$
(12)

Consequently, the ML parameter estimates, $p_{\rm ML}$, are

given by

$$p_{\mathrm{ML}} \in \underset{p \in \mathbb{R}^{n_p}}{\operatorname{arg\,max}} \log \left(\mathcal{L}(p \mid \mathcal{Y}_M) \right)$$
$$= \underset{p \in \mathbb{R}^{n_p}}{\operatorname{arg\,max}} \sum_{k=1}^{M} \log(\phi(y_k \mid \mathcal{Y}_{k-1}, p)),$$
(13)

where n_p denotes the number of parameters. The SDE in (1a) is driven by a Brownian motion, and since the increments of a Brownian motion are Gaussian, it is reasonable to assume that, under some regularity conditions, the conditional densities in (11) can be well approximated by Gaussian densities

$$\phi(y_k \mid \mathcal{Y}_{k-1}, p) = \frac{\exp\left(-\frac{1}{2}e'_k R^{-1}_{k|k-1}e_k\right)}{\sqrt{\det(R_{k|k-1})(2\pi)^{n_y}}}, \qquad (14)$$

where n_y is the number of output variables.

Estimating the refrigeration system model

The drift function, f, will be parameterized according to

$$f(x(t), u(t); p) = \begin{cases} p_1(p_2 - x(t)), & t \in \text{MELT} \\ p_3(p_4 - x(t)), & t \in \text{IDLE.} \end{cases}$$
(15)

By defining the input function, u, as

$$u(t) = \left(\mathbb{1}(t \in \text{MELT}), \mathbb{1}(t \in \text{IDLE})\right)'.$$
(16)

where

$$\mathbb{1}(t \in I) = \begin{cases} 1, & t \in I \\ 0, & t \notin I \end{cases},$$
(17)

the drift function can be defined according to

$$f(x(t), u(t); p) = u_1(t)p_1(p_2 - x(t)) + u_2(t)p_3(p_4 - x(t)).$$
(18)

It is assumed that the incremental covariances are constant, and hence the diffusion function, g, is defined as a positive parameter, p_5 . g is parameterized as

$$g(x(t), u(t); p) = p_5.$$
 (19)

The compressor capacity is observed directly. The function h is therefore defined as

$$h(x(t_k); p) = x(t_k).$$
 (20)

The variance of the uncertainty of the observations is also defined as a positive parameter

$$R_k = p_6, \tag{21}$$

which means that the variance is assumed to be timeinvariant. The parameters are estimated using CTSM-R [17] and the parameter estimates are listed in Table 1. Table 1

Estimated parameters of the SDE.

Parameter	Description	Value	Unit
p_1	MELT rate	0.00183	%/s
p_2	MELT assymptotic level	66.92400	%
p_3	IDLE rate	0.00085	%/s
p_4	IDLE assymptotic level	94.89100	%
p_5	Diffusion coefficient	0.28130	-
p_6	Observation variance	1.96580	-

The data sampled from the Danfoss test system is sampled at a very high frequency (around every six seconds) compared to the time-constants of the dynamical system. The observations have been sub-sampled such that the observations of the data-set used within CTSM-R are sampled every 15 minutes (one observation every 15 minutes). In Fig. 2, a full-horizon simulation using the estimated first-order, two-regime model introduced above, is shown. The solid black line represents the schedule of the switching function, u, the filled grey dots are the subsampled observations, the solid grey line represents the observations of the full data-set and the dashed grey line is the compressor capacity forecast given by the model introduced above. The forecast is not updated for each observation and is therefore a full horizon forecast (or simulation); the dashed grey line is the solution to the prediction scheme of the CDEKF.

The Nordic Electricity Market

In Northern Europe, electricity is traded in a common market called Nord Pool, which consists of 15 interconnected price areas. The market trading with the largest volume is called the day-ahead market. Here, electricity is bought and sold for the upcoming day and this market sets the spot price. When the day-ahead market closes, the intra-day market opens. In this market, electricity can be traded until 45 minutes prior to the operating hour.

One of the primary challenges when operating transmission systems is to guarantee grid stability. The Nordic Transmission System Operators (TSOs) have many methods for dealing with this challenge; one of them being a common balancing market [10]. In the balancing market, market participants have the option to make a bid that defines how much a participant is willing to change their production or consumption schedule in a given operating hour. The balancing market also closes 45 minutes prior to the operating hour. Hence, when approaching the operating hour, the TSO has the possibility to activate balancing bids ahead of time and thereby reduce the risk of imbalances. Three scenarios can take place in the balancing market:

- (\uparrow) If the imbalance is negative, there is a deficit of electricity in the price area, and hence an increase in power production or a decrease in power consumption is needed. This is called *up* regulation.
- (\downarrow) If the imbalance is positive, there is a surplus of electricity in the price, and hence a decrease in power production or an increase power consumption is needed. This is called *down* regulation.
- (-) If the imbalance is too small or the duration is too short, the imbalance is not offered in the balancing market.

In a situation with up regulation, electricity is sold, while in the situation with down regulation, electricity is bought. The structure of the balancing market requires that the up regulation price is greater than the day-ahead price, while the down regulation price is lower than the day-ahead price. In the price area DK1, a large share of the total power production is generated by wind turbines. This is a source of energy that is very difficult to predict, and hence it is one of the primary reasons to imbalances. In fact, in [23] it is suggested that approximately 65 % of the total imbalances are due to forecast errors of wind power production.

In the optimal control problem defined later, it is assumed that the combined refrigeration and ice-tank system is a price-taker of the spot price. The potential of also participating in the balancing market is not considered. However, it should be emphasized that the savings shown later in this paper should be regarded as a lower bound of the savings one can expect from e.g. also participating in a balancing market.

The Optimization Problem

In this section, the optimization problem for the optimal switching times is defined. The system dynamics embedded into the optimization problem formulation are ordinary differential equations (ODEs) defined by the drift function in (15) with the parameters, p, listed in Tab. 1.

Optimal control by optimal switching times

The drift function, f, in (15) is a regime-type function defined by a set of switching times. Define the set of switching times $\tau = \{\tau_{i,\text{MELT}}, \tau_{i,\text{IDLE}}\}_{i=0}^{N}$ for which the

 $\operatorname{structure}$

$$\tau_{i,\text{MELT}} \leq \tau_{i+1,\text{IDLE}}, \qquad (22a)$$

$$\tau_{i,\text{IDLE}} \leq \tau_{i,\text{MELT}}, \qquad (22b)$$

is imposed. The variables $\{\tau_{i,\text{MELT}}\}_{i=0}^{N}$ denote the temporal switches for when the ice-tank curtails the refrigeration system, and hence de-loads the system, and the variables $\{\tau_{i,\text{IDLE}}\}_{i=0}^{N}$ denote the temporal switches for when the ice-tank stops curtailing the refrigeration system and the system returns to normal operation. The number of switching times, N, is a design parameter of the optimal control problem; different values of N will lead to different control strategies. In this paper N = 5 is chosen.

Using these temporal switches, f can be defined according to the temporal decomposition given by

$$f(x(t), t; p) = \sum_{i=0}^{N-1} \mathbb{1} (t \in I_{i,\text{MELT}}) f_{\text{MELT}}(x(t); p) + \sum_{i=0}^{N} \mathbb{1} (t \in I_{i,\text{IDLE}}) f_{\text{IDLE}}(x(t); p),$$
(23)

where the intervals $I_{i,\text{MELT}}$ and $I_{i,\text{IDLE}}$ are defined according to

$$I_{i,\text{MELT}} = [\tau_{i,\text{MELT}}, \tau_{i+1,\text{IDLE}}]$$
(24a)
$$I_{i,\text{IDLE}} = [\tau_{i,\text{IDLE}}, \tau_{i,\text{MELT}}],$$
(24b)

and the functionals f_{MELT} and f_{IDLE} according to

$$f_{\text{MELT}}(x;p) = p_1(p_2 - x)$$
 (25a)

$$f_{\text{IDLE}}(x;p) = p_3(p_4 - x).$$
 (25b)

The convention $\tau_{0,\text{IDLE}} = 0$ and $\tau_{N,\text{MELT}} = T$ will be used, for which the parameter T > 0 will denote the simulation horizon for the optimal control problem defined later in this section.

Let x represent the only state of the system. This state models the compressor capacity. Besides the electricity cost of running the compressors, there is also a cost associated with operating the ice-tank in the two regimes. Let c(x(t), t) denote the total cost of operating the combined system consisting of the refrigeration system and the ice-tank at time t with a compressor capacity of the refrigeration system of x(t). The total cost rate c is then



Fig. 4. This figure shows the online results of actually implementing the optimal switching times. The top plot shows the predicted compressor capacity (left axis) together with the emulated outdoor temperature (right axis). The bottom plot shows the accumulated cost (left axis) and the used spot price (right axis).

given by

$$c(x(t), t; \tau) = q(t)$$

$$\times \left(kx(t) + k_{\text{MELT}} \sum_{i=0}^{N-1} \mathbb{1} \left(t \in I_{i,\text{MELT}} \right) + k_{\text{IDLE}} \sum_{i=0}^{N} \mathbb{1} \left(t \in I_{i,\text{IDLE}} \right) \right), \qquad (26)$$

where q(t) represents the electricity cost at time t, k is the rated capacity of the compressor, k_{MELT} is the electricity consumption of the ice-tank cost when it is curtailing the refrigeration system, and k_{IDLE} is the electricity consumption of the ice-tank when it is idle. Integrating this instantaneous total cost rate, c(x(t), t), yields the accumulated total cost

$$J(x,\tau) = \int_0^T c(x(t),t;\tau) \,\mathrm{d}t. \tag{27}$$

This functional will be the objective function of the optimal control problem defined next.

Using the definitions and variables defined above, the optimal control problem is defined according to

$$\min_{x,\tau} \left\{ J(x,\tau) = \int_0^T c(x(t),t;\tau) \,\mathrm{d}t \right\},\qquad(28a)$$

s.t.

$$\tau \in \mathcal{T},\tag{28b}$$

$$\dot{x} = f(x; p, \tau), \quad \text{in } [0, T],$$
 (28c)

$$x(0) = x_0, \tag{28d}$$

where the set \mathcal{T} defines the temporal structure given in (22) together with the temporal budget constraint given

by

$$\sum_{i=0}^{N} (\tau_{i+1,\text{IDLE}} - \tau_{i,\text{MELT}}) \le \overline{\tau}.$$
(29)

This budget constraint is due to the fact that the icetank only has a finite amount of ice available and the constraint will also ensure that the trivial case, where the ice-tank is only curtailing the refrigeration system, becomes infeasible (assuming that $\overline{\tau}$ is chosen properly). The constraint in (29) models the maximum allowed time of curtailment. In this paper, this budget parameter is defined as $\overline{\tau} = 8$ hours. In a

Note: An extra argument to the functions c and f in (28a) and (28c) has been included to indicate the dependence of the switching times, τ .

Analytical gradient expression

To enable efficient numerical solution of the optimal control problem (28), gradient information has to be made available to the optimization algorithm. The Jacobians of the temporal constraints in (28b) are trivial, since these constraints are all linear. Hence, the only derivatives that are non-trivial, are the derivatives that relate to the objective function (28a). The gradient of the objective function contains the elements

$$\frac{\partial J}{\partial \tau_{j,\text{MELT}}}(x,\tau) \quad \text{and} \quad \frac{\partial J}{\partial \tau_{j,\text{IDLE}}}(x,\tau).$$
 (30)

Analytical expressions of these will be derived in the following. These expressions are based on derivations and results from [3].

First, derivatives with respect to $\tau_{k,\text{MELT}}$ are considered. Inserting the definition of J yields

$$\frac{\partial J}{\partial \tau_{j,\text{MELT}}}(x,\tau) = \frac{\partial}{\partial \tau_{j,\text{MELT}}} \int_0^T c(x(t),t;\tau) \,\mathrm{d}t, \qquad (31)$$

Using the definition of the indicator function of the temporal decomposition, the integral in (31) can be defined as

$$\int_{0}^{T} c(x(t), t; \tau) dt = \int_{0}^{T} kq(t)x(t) dt$$
$$+ k_{\text{MELT}} \sum_{i=0}^{N-1} \int_{\tau_{i,\text{MELT}}}^{\tau_{i+1,\text{IDLE}}} q(t) dt \qquad (32)$$
$$+ k_{\text{IDLE}} \sum_{i=0}^{N} \int_{\tau_{i,\text{IDLE}}}^{\tau_{i,\text{MELT}}} q(t) dt.$$

Using (32), the derivatives with respect to $\tau_{k,\text{MELT}}$ simplify into

$$\frac{\partial J}{\partial \tau_{j,\text{MELT}}}(x,\tau) = \frac{\partial}{\partial \tau_{j,\text{MELT}}} \int_{0}^{T} kq(t)x(t) \,\mathrm{d}t + (k_{\text{IDLE}} - k_{\text{MELT}}) q(\tau_{j,\text{MELT}}),$$
(33)

where the first term can be computed using the adjoint states (or co-states) for the optimal control problem [3]. The co-states (to be denoted by λ) satisfy the dynamical equations

$$\dot{\lambda} = -\frac{\partial f}{\partial x}(x; p, \tau)'\lambda - \frac{\partial c}{\partial x}(x, \cdot; \tau)', \quad \text{in } [0, T] \quad (34)$$
$$\lambda(T) = 0.$$

Using the results from [3], the derivatives with respect to $\tau_{k,\text{MELT}}$ can be computed from

$$\frac{\partial J}{\partial \tau_{j,\text{MELT}}}(x,\tau) = \lambda(\tau_{j,\text{MELT}})'(f_{\text{IDLE}}(x(\tau_{j,\text{MELT}});p) - f_{\text{MELT}}(x(\tau_{j,\text{MELT}});p)) + (k_{\text{IDLE}} - k_{\text{MELT}})q(\tau_{j,\text{MELT}}).$$
(35)

The result in (35) is not valid for $\tau_{N,\text{MELT}}$. However, this variable is by convention fixed ($\tau_{N,\text{MELT}} = T$) and will not be subject to the optimization.

Following the same steps as above, a similar result can be derived for the derivatives of the objective function (28a) with respect to $\tau_{j,\text{IDLE}}$. The derivatives with respect to these variables can be computed from

$$\frac{\partial J}{\partial \tau_{j,\text{IDLE}}}(x,\tau) = \lambda(\tau_{j,\text{IDLE}})'(f_{\text{MELT}}(x(\tau_{j,\text{IDLE}});p) - f_{\text{IDLE}}(x(\tau_{j,\text{IDLE}});p)) + (k_{\text{MELT}} - k_{\text{IDLE}})q(\tau_{j,\text{IDLE}}).$$
(36)

The result in (36) is not valid for $\tau_{0,\text{IDLE}}$. However, this variable is by convention fixed ($\tau_{0,\text{IDLE}} = 0$) and will not be subject to the optimization.

Note: The implementation of the numerical solution of the optimal control problem (28) is done using a singleshooting formulation [5, 7]. This means that the dynamical equations (28c)-(28d) are solved internally of the implemented objective function. Thus, the only constraints needing to be implemented are the linear set of constraints defined by the temporal set of constraints in (28b).

Numerical implementation

The optimal control problem has been solved numerically in python. A single-shooting approach has been applied where the solve_ivp function from the scipy package has been used to solve the dynamical equations (the state and co-state equations). The Runge-Kutta 5(4) method with adaptive step size has been chosen as the numerical method [8, 27]. The optimization algorithm used is a constrained trust-region method, which is available via the minimize function from the scipy package [6].

Results

This section presents the results of implementing the optimal control algorithm defined in (28) for the total refrigeration system illustrated in Fig. 1. A single openloop iteration of the optimal control problem is used to compute the optimal switching times. These switching times are used to operate the ice-tank for a period of 24 hours (or one day). This implies that no feedback from the refrigeration system nor the ice-tank is used. The power prices are assumed to be known for the entire 24 hours; this is a reasonable assumption as many retailers offer power prices which follows the spot price. The spot price is always known 12 hours to 36 hours ahead (depending on the time of day); hence, optimal control problem is run when there is at least 24 hourly power prices available (this is e.g. the case everyday at 4pm).

The spot prices from a random day at Nord Pool have been used to generate a load input in terms of a price signal. The outdoor temperature is emulated using the water heater shown in Fig. 1, and this temperature is set



Fig. 5. This figure visualizes the effect of having the ice-tank installed. The top plot shows the accumulated cost of actively operating the ice-tank vs. not having the ice-tank. The bottom plot shows the relative savings of operating the ice-tank. The bottom plot shows how the predicted savings compare to the realized savings. It is assumed that the ice-tank has been charged with zero cost.

to average around 35 °C such that the compressors of the existing refrigeration system would operate close to 100 % when the ice-tank is in the IDLE mode. The temporal budget constraint parameter, $\overline{\tau}$, in (29) is defined such that the ice-tank can only be in the MELT mode for 8 h. It is assumed (in the financial analysis in Fig. 5) that the ice-tank has been charged with a zero cost. The values of the power consumption parameters k, k_{MELT} and k_{IDLE} used in the objective function of the optimal control problem are given in Table 2.

Table 2

Ċ.

Parameter	Description	Value	Unit
k	Rated compressor capacity	12.4	kW
$k_{ m MELT}$	MELT power consumption	250.0	W
$k_{\mathtt{IDLE}}$	IDLE power consumption	9.0	W

The optimal control algorithm optimally distributes five MELT periods throughout the 24 h of operation (starting at 4pm on the first day of testing). These cycles are distributed according to (rounded to nearest second):

- 1. MELT starts at 16:00:08 until 16:40:06.
- 2. MELT starts at 19:57:34 until 21:29:45.
- 3. MELT starts at 23:32:58 until 01:44:14.
- 4. MELT starts at 02:45:01 until 05:58:25.
- 5. MELT starts at 14:28:35 until 14:51:43.

In Fig. 4, the results of operating the ice-tank using these optimal switching times are shown. The top plot shows the difference between the predicted compressor capacity (predicted by the SDE model) and the observed values. The bottom plot shows the spot prices used in the

objective function together with predicted accumulated cost. In Fig. 5, the performance of operating the ice-tank is presented. The top plot shows the difference between actively operating the ice-tank and not having an icetank at all. The bottom plot shows the predicted savings and the realized savings. The realized savings are computed based on the observed compressor capacity values shown in the top plot of Fig. 4. From Fig. 4 and Fig. 5 the following observations are made:

- Fig. 4 shows that when the ice-tank is in MELT, the compressor capacity drops to approximately 67 %. This corresponds to the parameter value of p_2 given in Table 1. From Fig. 4 it is also seen that the SDE model predicts too low compressor capacities when the ice-tank switches from MELT to IDLE. This is observed after the first three MELT cycles.
- From Fig. 4 it is observed that the emulated outdoor temperature exhibits a higher degree of variability when the ice-tank implements a mode change.
- Fig. 5 shows that there is some discrepancy between the predicted savings of operating the ice-tank and the realized savings. However, the predicted and realized savings are both in the order of approximately 20 %. This discrepancy is expected, as the predicted savings are based on forecasts with a horizon of up to 24 hours.

Discussion & Future Outlook

This section discusses some of the observed deficiencies observed for the SDE model used in the optimal control problem. A possible method for economical and efficient charging of the ice-tank is also discussed. In connection with this, it is also discussed how the optimal control problem can be extended to include other load inputs than the spot price (e.g. CO2 emissions). Finally, the applicability of the ice-tank to other regions (than Scandinavia) and hence other electricity markets is discussed.

Extensions to the dynamical model

The dynamical model used in the optimal control problem in (28) has difficulties in explaining the transient dynamics observed in Fig. 3-4 when the ice-tank implements a mode change. One extension to the regime-based first-order model used to model the drift function (15) could be to include the outdoor temperature as a state in the SDE model such that the higher variability of the outdoor temperature might be explained by the model. A second extension could be to consider higher order models that better describe the dynamics observed after a mode change. The data presented in Fig. 3-4 suggest that a suitable model might be nonlinear, since the compressor capacity tends to drop both when the ice-tank switches from IDLE to MELT and from MELT to IDLE.

The SDE model used in this paper is estimated from data where the refrigeration system is measured under similar conditions. Hence, this model is indifferent to changing conditions such as e.g. opening and closing hours. In [24] it is shown that there is a clear difference in a supermarket's power consumption during opening and closing hours. Thus, in a real-life application the model has to be extended to accommodate such conditions. One solution to this might be to consider two different models for opening and closing hours respectively. Neither is the SDE model suitable under conditions where the outdoor temperature exhibits time-varying dynamics. Under such conditions, the model parameters, p, might all depend on the value of the outdoor temperature. A suitable extension to accommodate this might be (again) to include the outdoor temperature as a state in the SDE model.

Both of these extensions also address the importance of robustness of the controller. The method proposed in this paper has no state or parameter update step; however, to accommodate changing external conditions such as opening hours and outdoor temperatures, recalibration of the dynamical model becomes very important. Using continuous-discrete stochastic systems, as the system introduced in (1), the re-calibration process can easily be included.

Optimal and sustainable charging of the ice-tank

In the calculations of the realized and predicted savings presented in Fig. 5, it is assumed that the ice-tank has been charged with zero cost. Naturally, under varying market conditions this will not always be feasible. However, one of the main advantages of using the ice-tank as the flexible component in the refrigeration system is

that the only two constraints are the minimum and maximum capacities of the amount of ice that can be stored. If the refrigeration system itself had been subject to a control strategy as done in e.g. [11, 15, 20, 26], then the system would also be subject to hard constraints on e.g. the temperature of the display units. Hence, the ice-tank can be used to provide flexibility in a balancing market and thereby reduce the price of generating the stored ice. Furthermore, the optimal control problem can even be extended such that the ice-tank can be used to optimally distribute in which electricity market the power consumption is traded (day-ahead market, intraday market, balancing market, etc.). It is expected that such formulations would heavily increase the observed savings, as the price differences become much larger. However, to take full advantage of this, a better model of the short-term dynamics is needed, together with a model that generalizes to varying exogenous conditions (e.g. the outdoor air temperature).

In the objective function used in the formulation of the optimal control problem in (28), only the cost associated with power consumption is considered. However, this objective function can easily be extended to also accommodate a cost associated with e.g. the CO2 emissions of power consumption. This type of combination of objectives for an electricity consumer is investigated in e.g. [18].

Other electricity markets

The results presented in this paper rely heavily on the assumption of high outdoor temperatures such that the compressors yield a high power consumption. This is a constraining assumption for the applicability in Northern Europe, as the savings presented in Fig. 5 will only be feasible during the summer months. Experiments with lower temperatures have also been conducted. For these experiments, savings in the range from 5 % to 10 % are observed. Thus, lower outdoor temperatures will extend the payback period of the investments associated with the installation and maintenance of the ice-tank.

One region with higher average temperatures than Denmark and Scandinavia in general, is California. The Californian electricity market offers a variety of markets which might be economically beneficial to consumers who can deliver demand-side flexibility. In [9] it is shown that participation in the Californian energy markets (day-ahead market, 15-minute market and real-time market) and bidding for ancillary services (non-spinning reserves, spinning reserves and regulation) hugely increases the revenue potential of actively operating a battery. The paper shows that under perfect information conditions, the revenue generated from participation in the full-stack of electricity markets might be up to 600 % larger than the revenue generated by only bidding in the day-ahead market. The ice-tank presented in this paper has many similarities to a battery. Thus, it is expected that larger savings (than shown in Fig. 5) are obtainable by actively bidding the flexibility generated by the ice-tank in multiple markets on different time-scales. However, this will require more accurate dynamical models of the interaction between the ice-tank and the refrigeration system.

Conclusion

The goal of this paper was to investigate the economic potential of actively operating an ice-tank connected to a refrigeration system using an optimal control algorithm. Based on the experiments presented in this paper, it is demonstrated that it is possible to lower the total electricity costs by approximately 20 % with a temporal budget constraint for the ice-tank of 8 h. It is expected that larger savings are obtainable by actively bidding the flexibility generated by the ice-tank in multiple electricity markets. However, this will require more accurate dynamical models of the interaction between the ice-tank and the refrigeration system.

Conflict of Interest

None.

References

- [1] Denmark strikes deal to slash CO2 emissions by 70% in a decade. *Bloomberg*, 2019.
- [2] Joy E Altwies and Douglas T Reindl. Passive thermal energy storage in refrigerated warehouses. International Journal of refrigeration, 25(1):149–157, 2002.
- [3] H Axelsson, M Egerstedt, Y Wardi, and G Vachtsevanos. Algorithm for switching-time optimization in hybrid dynamical systems. In Proceedings of the 2005 IEEE International Symposium on, Mediterrean Conference on Control and Automation Intelligent Control, 2005., pages 256–261. IEEE, 2005.
- [4] Reza Barzin, John JJ Chen, Brent R Young, and Mohammed M Farid. Peak load shifting with energy storage and price-based control system. *En*ergy, 92:505–514, 2015.
- [5] Hans Georg Bock and Karl-Josef Plitt. A multiple shooting algorithm for direct solution of optimal control problems. *IFAC Proceedings Volumes*, 17(2):1603–1608, 1984.
- [6] Andrew R Conn, Nicholas IM Gould, and Ph L Toint. Trust region methods, volume 1. Siam, 2000.
- [7] Moritz Diehl, Hans Georg Bock, Holger Diedam, and P-B Wieber. Fast direct multiple shooting algorithms for optimal robot control. In *Fast motions in biomechanics and robotics*, pages 65–93. Springer, 2006.

- [8] John R Dormand and Peter J Prince. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.
- [9] Alexander W Dowling, Ranjeet Kumar, and Victor M Zavala. A multi-scale optimization framework for electricity market participation. *Applied Energy*, 190:147–164, 2017.
- [10] DK Energinet. Regulation c2–the balancing market and balance settlement. 2017.
- [11] Miha Glavan, Dejan Gradišar, Iztok Humar, and Damir Vrančić. Refrigeration control algorithm for managing supermarket's overall peak power demand. *IEEE Transactions on Control Systems Technology*, 2018.
- [12] Sasank Goli, Aimee McKane, and Daniel Olsen. Demand response opportunities in industrial refrigerated warehouses in california. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2011.
- [13] C Heerup and T Green. Load shifting by ice storage in retail CO2 systems. In 11th IIR Gustav Lorentzen Conference on Natural Refrigerants, Hangzhou, China, 2014.
- [14] Tobias Gybel Hovgaard, Lars FS Larsen, and John Bagterp Jørgensen. Flexible and cost efficient power consumption using economic mpc a supermarket refrigeration benchmark. In 2011 50th IEEE Conference on Decision and Control and European Control Conference, pages 848–854. IEEE, 2011.
- [15] Tobias Gybel Hovgaard, Lars FS Larsen, Morten J Skovrup, and John Bagterp Jørgensen. Power consumption in refrigeration systems-modeling for optimization. In 2011 International Symposium on Advanced Control of Industrial Processes (AD-CONIP), pages 234–239. IEEE, 2011.
- [16] Andrew H Jazwinski. Stochastic Processes and Filtering Theory. Academic Press, San Diego, CA, USA, 1970.
- [17] Rune Juhl, Jan Kloppenborg Møller, and Henrik Madsen. Parameter estimation in stochastic greybox models. arXiv, 2016.
- [18] Rune Grønborg Junker, Armin Ghasem Azar, Rui Amaral Lopes, Karen Byskov Lindberg, Glenn Reynders, Rishi Relan, and Henrik Madsen. Characterizing the energy flexibility of buildings and districts. *Applied energy*, 225:175–182, 2018.
- [19] Niels Rode Kristensen, Henrik Madsen, and Sten Bay Jørgensen. Parameter estimation in stochastic grey-box models. *Automatica*, 40(2):225–237, 2004.
- [20] Lars FS Larsen, Claus Thybo, and Henrik Rasmussen. Potential energy savings optimizing the daily operation of refrigeration systems. In 2007 European Control Conference (ECC), pages 4759– 4764. IEEE, 2007.
- [21] MD Murphy, MJ O'Mahony, and J Upton. Comparison of control systems for the optimisation of ice storage in a dynamic real time electricity pricing environment. *Applied Energy*, 149:392–403, 2015.

- [22] Bernt Oksendal. Stochastic differential equations: an introduction with applications. Springer Science & Business Media, 2013.
- [23] Henning Parbo. Balancing management with large shares of renewables. 2nd Annual European Forum on Grid Integration and Electricity Anciliar Services, 2014.
- [24] Lisa Buth Rasmussen, Peder Bacher, Henrik Madsen, Henrik Aalborg Nielsen, Christian Heerup, and Torben Green. Load forecasting of supermarket refrigeration. *Applied energy*, 163:32–40, 2016.
- [25] Falko Ueckerdt, Robert Brecha, and Gunnar Luderer. Analyzing major challenges of wind and solar variability in power systems. *Renewable energy*, 81:1–10, 2015.
- [26] Kasper Vinther, Torben Green, E Shafiei, Luminita C Totu, Roozbeh Izadi-Zamanabadi, and Tobias G Hovgaard. Control strategies and challenges for utilizing supermarket refrigeration systems in a smart energy context. In 2016 IEEE Conference on Control Applications (CCA), pages 593– 598. IEEE, 2016.
- [27] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0– fundamental algorithms for scientific computing in python. arXiv preprint arXiv:1907.10121, 2019.

Multi-Market Energy Arbitraging Using a Vanadium Redox-Flow Battery

Niclas Brok^a, Henrik Madsen^a, Anders Bentien^b, Rune Grønborg Junker^a

^aDepartment of Applied Mathematics and Computer Science at the Technical University of Denmark. ^bDepartment of Engineering at University of Aarhus.

Abstract

Power grids coordinate a diverse set of markets and energy systems to ensure that supply and demand are matched on different time-scales. This coordination is made increasingly difficult by a large and expanding share of renewable energy sources in power systems. To accommodate this balancing of supply and demand, grid operators facilitate balancing markets where market participants can bid for attractive power prices and help balance the power grid. This paper presents an optimal control algorithm for operating a vanadium redox-flow battery that leverage the multi market structure that constitute a key element in the power grid coordination. This work considers two markets within the Northern European power system: the day-ahead and the regulating power market. These markets can be leveraged by battery operators to create different sources of energy arbitraging: an instantaneous arbitrage where electricity in the same operating hour is different in the different markets considered and a time-shifted arbitrage where the battery utilize that the power prices are time varying. Doing so requires price forecasts, and so the effect of varying the quality of these forecasts has been taken into account. Furthermore, given an increasing focus on CO2 emissions, the effect of CO2-taxes is also considered. This energy arbitraging is formulated as an optimal control problem, where the state equation models the state of charge of the battery. The optimal control method is backtested on market data from the Northern European power system for a full year, consisting of 1 month of 2019 and 11 months of 2020. The redox-flow battery is found to be economically viable, both in the power systems of today, and for the expected future power grids, irrespective of the size of CO2-taxes. Forecasting precision is found to dominate the pay-back period, making it range from 7 to 9 years (using current battery prices) and ranging from 4-6 years (using expected future battery prices).

Keywords: Optimal Control, Vanadium Redox-Flow Battery, Electricity Markets, Optimization, Arbitraging

1. Introduction

During the final months of 2019, Denmark agreed on its first ever Climate Bill, committing to reduce greenhouse gas emissions by more than 70 % by 2030 compared to the Danish emission levels in 1990 [1]. This will increase the need for introducing renewable energy sources and efficient integrations of these energy sources. The Nordic transmission system operators (TSOs) coordinate a diverse set of energy systems to ensure that supply and demand are matched on multiple time-scales.

In this paper, the primary focus will be given to the day-ahead and the regulating power markets. These markets will define an optimal control problem for operating a battery in multiple markets simultaneously to create energy arbitrage and thereby creating an attractive revenue profile for battery operators. The battery subject to the optimization problem will be a simulation model of a vanadium redox-flow battery, with parameters and cost structure provided by [2].

The literature contains a broad range of applications that study methods which can help penetrate more renewable energy sources. A big area of research is the applicability of demand-side flexibility to help balance the seemingly stochastic nature of e.g. wind power production. [3]

Preprint submitted to Elsevier

presents a method to offer flexibility of HVAC systems in a contractual framework which minimizes the cost of operating the HVAC system. [4] discusses the flexibility potential of office buildings and the potential impact on building performance; e.g. the indoor climate. [5, 6] describe how wastewater treatment plants can utilize the flexibility of the biochemical process in the wastewater tanks by optimal scheduling of the aeration equipment. These methods all impose technical challenges of the systems involved, as these systems often have physical constraints that need to be satisfied; wastewater treatment plants in Denmark e.g. has to ensure that the daily average nutrient concentrations of the effluent is kept below a politically defined threshold.

Batteries, on the other hand, have few physical constraints and can implicitly be used to penetrate more renewable energy sources by participating in e.g. balancing markets. [7, 8] present methods to optimally operate batteries in multiple Californian energy markets. [7] shows that it is possible to create very attractive revenues by also operating batteries in balancing markets. [9] shows similar results, but under Danish market conditions; energy arbitraging is more lucrative when also participating in balancing markets. With regulations and rules, relating to climate-change, evolving rapidly, the business-case of investing in batteries is tested against (potential) future CO2-taxes. This CO2-tax parameter will define how the optimal control problem weighs time-varying power prices against timevarying CO2 emissions.

1.1. Key assumptions

This paper implements two critical assumptions:

- 1. The strategy used to operate the battery is pricetaker of the regulating power prices.
- 2. This strategy has the ability to perfectly bid for special down-regulation using a given price parameter.

The first assumption is reasonable under the condition that the volume subject to activation is in-significant compared to the general volume activated in the regulating power market. A simple bidding strategy to always be activated for regulating power is to bid the day-ahead price (this is used as reference price in the regulating power market). This bidding strategy is similar to what is known as front-running in the high-frequency trading literature [10]. The second assumption is justified based on the fact that the average special down-regulation volume (during hours with activation) was approximately 800 MWh during the first 11 months of 2020. This is a very significant volume, with the point being that this volume will not go unnoticed in the power market. The special down-regulation price parameter can be interpreted as an average activation price of the proposed optimization-based strategy.

1.2. Key contributions and paper organization

This paper investigates the economical feasibility of operating and investing in grid scale Vanadium redox-flow batteries by arbitraging two Northern European power markets: the day-ahead market and the regulating power market. The operation of the battery is scheduled as the output from an optimal control problem, where the model predictive control principle is applied to recursively find optimal schedules of battery operation. The impact of activation for special down-regulation is investigated under various price assumptions. This market is treated as a special-case of the regulating power market. The impact of future CO2-taxes is also considered.

The paper is structured as follows: The next section describes the Northern European power market, the section thereafter defines the optimal control problem which is used to operate the battery, then follows a backtest of the proposed method (tested on market data from the December 2019 until December 2020). The paper finishes with a discussion about extensions to the proposed method and a section with concluding remarks.

2. The Northern European Power Markets

In Northern Europe, electricity can be traded on various markets with different time horizons and market designs. The Nordic and Baltic countries constitute a common market of 15 interconnected price zones (2 Danish, 5 Norwegian, 4 Swedish, 1 Finish and 3 Baltic).

2.1. The day-ahead market

The Nord Pool Power Exchange is responsible for the day-ahead and intra-day markets. The day-ahead market contains the largest trading volume and market participants can purchase and sell electricity for the individual hours for the up-coming day. Bidding in the day-ahead market closes at 12pm on the day prior to the day of operation. This means that the planning horizon is 12-36 hours ahead. The day-ahead price (or spot price) is settled as the intersection between the ordered bids of power purchases and sales with constraints enforcing the physical flow limitations of the power grid.

2.2. The intra-day market

The intra-day market is open for continuous trading and is settled using traditional order matching. This market has traditionally been used to eliminate imbalances between the day-ahead planned production/consumption and the current expected power production/consumption plan. Trading on the intra-day market stops 45 minutes prior to the given hour of operation with the trading volume on the intra-day market in general being far less than the trading volume observed for the day-ahead market. After end-of trading on the intra-day market the market participants cannot modify their production/consumption schedule and potential system imbalances are cleared on the regulating power market.

2.3. The regulating power market

The regulating power market is operated and cleared by the Nordic TSOs. Before the system imbalance is cleared in this balancing market, the Nordic TSOs coordinate if there is opposite imbalances in the neighbouring price zones and adjust the flow of the inter-connectors to eliminate (potentially only some of) these opposing imbalances. This balancing operation is called netting. The current standard in Northern Europe for this netting process between inter-connected Nordic price zones is not using an efficient market mechanism, but is done by ad-hoc coordination between the Nordic TSOs. Given that the imbalance persist through this netting process it is cleared on the regulating power market. The sign of the imbalance defines the type of regulation activated; if the imbalance is negative, up-regulation bids are activated, and if the imbalance is positive, down-regulation bids are activated. Imbalances are only offered to the regulating power market if it is of significant magnitude and duration. If the system imbalance is insignificant all participants with imbalances are



Figure 1: This figure shows the price distribution of the regulating power prices relative to the day-ahead prices. The price distribution is generated based on the observed prices during the first 11 months of 2020 in DK1. The left second-axis shows the frequency of the observed power price difference and the right second-axis shows the cable congestion between DK1 and NO2.

settled at the day-ahead price. All activated regulating power bids are activated at the most favorable price, this means that for up-regulation the bid with the highest price defines the up-regulation price and for down-regulation the bid with the lowest price defines the down-regulation price.

Fig. 1 shows the price distribution of the regulating power prices relative to the spot prices. The left secondaxis shows the frequency of the given power price difference and the right second-axis shows the cable congestion from DK1 to NO2. The down-regulation occurs more frequently than up-regulation in DK1, but the up-regulation prices more frequently attain extreme prices. The congestion towards NO2 is a very significant explanatory variable to explain extreme power price difference observed in the regulating power market in DK1. NO2 (and Norway in general) contains a large capacity of hydro-power. This means that if NO2 can provide down or up-regulation to e.g. DK1, the regulating power prices become very competitive. However, if this connection is congested, the regulating power prices are more volatile and more frequently attain more extreme values.

Remark: Other balancing mechanisms come into play if the expected imbalance is of too small magnitude or too short duration, however the potential of further participation in these markets are not considered in this paper. Hence, these balancing mechanisms are not elaborated further in this paper.

2.4. Special down-regulation

In recent years there has been a rapid development of a market feature called special regulation in DK1. Special regulation can occur as both down and up-regulation and comes in effect when a Nordic TSO skips the usual pricing procedure. In the event of special regulation, the activated bids are paid-as-bid. In Table 1 aggregated values for special down-regulation in DK1 is shown for 2017-2019. During these three years, one of the main sources of special down-regulation in DK1 has been excessive wind power production in Northern Germany and during these years operators of wind turbines in DK1 have shown an increasing willingness to curtail their wind turbines in the event of special down-regulation. A market report [11] by the Danish TSO, Energinet, and the Northern German TSO, TenneT, describes the countertrade agreement between the two TSOs explaining the economical incentives behind why e.g. Danish wind turbines are curtailed to accommodate excess German power production. In the market report it is stated that the agreement is economically very beneficial and that in many operating hours, excess German power production is used as an up-regulation mechanism to eliminate negative imbalances in the Danish power system; this happens to approximately 30% of the countertrading volume originating from Northern Germany. This is claimed to be economically efficient.

Table 1 shows that in 2019 a total of 1,312 GWh was



Figure 2: Monthly day-ahead, intra-day and regulating power volumes from the first 11 months of 2020.

activated for special down-regulation in Denmark due to the balancing agreement with the German TSO, TenneT. This table also shows that in 2019, 32 % of this special down-regulation was handled by curtailment of wind turbines in Denmark. Energinet (the Danish TSO) is only allowed to publish yearly average prices; in 2019 the average special down-regulation price was -12.3 EUR/MWh. In general Table 1 shows that the demand for special downregulation in DK1 is increasing with corresponding power prices becoming very attractive. Thus, there is a great potential for flexible producers and consumers to participate in this market and ideally eliminate the need for curtailing Danish wind turbines.

2.5. Trading volumes

In Fig. 2, the trading volumes on the day-ahead, intraday and regulating power markets are shown. As stated initially, the day-ahead market has (by far) the largest trading volume. However, during the first 11 months of 2020 there has also been a massive demand for special down-regulation in the Danish price zone, DK1. This is especially observed during the Winter months, which typically contain a large amount of wind power production.

3. Battery Modelling and Energy Arbitraging

This section defines a continuous-time optimal control problem based on a simple integrator model of a battery. This dynamical model of the battery is inspired by the storage model discussed in e.g. [12], but includes a power dependent efficiency coefficient for both charging and discharging of the battery.

3.1. Dynamical battery model

The dynamical model of the vanadium redox-flow battery considered in this paper is given by

$$q\dot{x} = \eta_l(u) \max(u, 0) + \eta_g(u) \min(u, 0).$$
(1)

In this model, q denotes the storage capacity (in MWh) of the battery. x denotes the state of charge and this variable will be constrained such that $x(t) \in [\underline{x}, \overline{x}]$ for any t; these bounds $0 \le x < \overline{x} \le 1$ denote the physical bounds of the battery's depth of charge and discharge. u denotes the power at which the battery is operated at and this variable is constrained such that $u(t) \in [u, \overline{u}]$; these bounds $u < 0, \overline{u} > 0$ denote the maximum power for which the battery can be charged and discharged. The functionals η_l and η_q denote the charging and discharging efficiencies, respectively. These functionals will be assumed to be linear and is defined such that when the battery is idle (i.e. u(t) = 0), the efficiency is 100 % and when the battery is operating near the power limits (i.e. $u(t) = \underline{u}$ or $u(t) = \overline{u}$), the efficiency is 85 %. Hence, these functionals are defined according to

$$\eta_l(u) = 1 - 0.85 \frac{u}{\overline{u}}, \text{ and } \eta_g(u) = 1 - 0.85 \frac{u}{\underline{u}}.$$
 (2)

These definitions of the efficiency functionals imply that the round-trip efficiency varies from approximately 72 % to 100 % depending on the implemented levels of charging and discharging.

3.2. The optimal control problem

In the previous section, the multi-market structure of the Northern European electricity market was described. From this section it is clear that the implemented power,



Figure 3: This figure shows an example of a solution to the optimal control problem defined in (5). The top plot shows the input spot price function (in blue), p_{da} , the regulating power prices (in dark red), the regulating power price function (in transparent dark red), p_{reg} , and the co2-emissions. The bottom plot shows the state of charge and the implemented power, u.

u, can be purchased and sold on multiple markets. In order to accommodate this in the optimal control problem, the power variable, u, is decomposed in terms of the volume purchased or sold on the day-ahead market and the volume subject to down or up-regulation in the regulating power market

$$u = u_{\rm da} + u_{\rm reg}.\tag{3}$$

In (3), u_{da} denotes the day-ahead position acquired and u_{reg} denote the regulating power position (or activation).

In the optimal control problem defined next, it is assumed that two price functions, p_{da} and p_{reg} , of the future spot and regulating power prices are given. The backtest conducted in the next section, will consider different assumptions concerning the accuracy of the regulating power price function, p_{reg} . Due to the structure of the day-ahead market, the day-ahead prices are always known at least 12 hours ahead. It is also assumed that the CO2-emission function, c, of the future CO2-emissions are given. These CO2-emissions will be related to an economic penalty via a taxation parameter, k_{tax} . The objective subject to minimization is the integrated cost of operating the battery in the day-ahead and regulating power markets together with the CO2 emission costs during the optimization interval [0,T]

$$\mathcal{J}(u_{\rm da}, u_{\rm reg}) = \int_0^T \left[p_{\rm da}(t) u_{\rm da}(t) + p_{\rm reg}(t) u_{\rm reg}(t) + k_{\rm tax} c(t) u(t) \right] \mathrm{d}t.$$
(4)

Using the definitions and bounds defined above the optimal control problem used to operate the battery is defined according to

$$\min_{x, u_{\rm da}, u_{\rm reg}} \mathcal{J}\left(u_{\rm da}, u_{\rm reg}\right),\tag{5a}$$

s.t.

$$u = u_{\rm da} + u_{\rm reg},\tag{5b}$$

$$q\dot{x} = \eta_l(u)\max(u,0) + \eta_g(u)\min(u,0), \qquad (5c)$$

$$x_0 = x(0) \tag{5d}$$

$$\underline{x} \le x, \qquad \overline{x} \ge x, \tag{5e}$$

 $\underline{u} \le u, \qquad \overline{u} \ge u, \tag{5f}$

$$u_{\rm da} = u_{\rm spec}.$$
 (5g)

where x_0 denotes the initial state of charge of the battery. u_{spec} is an input function which the acquired dayahead position. The optimal control problem (5) will be
used to backtest the revenue potential of operating a Vanadium redox-flow battery by optimally leveraging price differences between the day-ahead and regulating power markets.

3.3. Example solution to the optimization problem

In Fig. 3, an example of the output from the optimization problem (5) is shown. The optimization horizon spans 48 hours starting at 6pm on the 19th of January 2020. The top plot shows on the left second-axis the power prices used as input to the optimization problem. The forecasting ability will be subject to varying efficiencies; in Fig. 3, the efficiency of these forecasts decay by 25 % per hour from the true regulating power prices to the spot prices and p_{reg} attains a constant value after the initial 12 hours. Similarly, the CO2-emissions are assumed to be known for the first 12 hours, but c attains a constant value after the initial 12 hours. These 12 hours are critical since the spot prices are always well-defined at least 12 hours ahead. The battery used in Fig. 3 has a nominal charging and discharging power of 1/3 MW with a storage capacity q = 1 MWh and $u_{\rm spec} = -1$ MW. Thus, the solution shown in Fig. 3 always sell 1 MWh in the day-ahead market. The optimal control problem is solved using CasADi [13] with an hourly zero-order-hold discretization of the optimization variables.

3.4. Vanadium Redox-Flow Price Structure

Large batteries for stationary storage for arbitraging on multiple energy markets is a relatively large investment. In this paper a flow-battery cost-model based on the nominal power, the capacity of the battery and a fixed initial start-up cost is used. The cost model captures the fundamental features of a flow battery where the charge/discharge power is determined by the flow battery stack, while the capacity is determined by the amount of redox-active Vanadium stored in the tanks. Additional cost related to the balance-of-plant (pumps, tanks, sensors, etc.) scales either with the power or capacity.

In this paper, the used values are initial start-up costs of 2,545 EUR, cost of 1,523 EUR per kW and a cost of 318 EUR per kWh. These values are obtained from a

Table 1: Activated special down-regulation (in GWh) in DK1 received from the German TSO, TenneT, in 2017-2019. The activated volumes are activated volumes for Danish participants and the average price is in EUR/MWh.

	2019	2018	2017
up-regulation (MWh)	602	484	429
% of total	31 $\%$	30 %	35 $\%$
Special down-regulation	1,312	1,114	781
Thermal power plants	46 $\%$	53 $\%$	64 $\%$
Electric boilers	22 $\%$	21 $\%$	$22\ \%$
Wind turbines	32 %	26 $\%$	14 $\%$
Avg. price (EUR/MWh)	-12.3	-9.3	-7.7

two dimensional linear parameterization of present vanadium flow battery sales costs provided by the battery company VisBlue [2]. Commercial Vanadium flow-batteries are a relatively new technology and have a large costreduction potential. Vanadium raw materials, which are relatively easily processed into a ready-to-use solution, are only about 70 EUR per kWh [14], while material costs related to power-hardware stacks are many times lower than actual costs per kW because of small production numbers [2]. For this reason significant cost-reductions in the coming years are expected. In this work, it is assumed that the future (in five years) cost-model with start-up cost, cost per kW and cost per kWh will drop to 1,697 EUR, 761 EUR per kW and 177 EUR per kWh, respectively. Both the current and five year costs are fully in line with the numbers published by the International Renewable Energy Agency [15].

Costs include power electronics, but not installation and building costs. The dimensioning of the storage will be fixed at 1 MWh such that the nominal power to physical storage ratio is the only design parameter of the battery. Based on the backtest performed in the next section it is possible to discuss the optimal configuration of the battery.

Remark: Given a nominal power of e.g. 1 kW, the battery has the physical capacity of actually charging and discharging at +1.5 kW and -1.5 kW, respectively. Thus, the physical bounds on the power limits of the battery are 1.5 times higher than the nominal power which is used to price the battery.

3.5. Backtest

The regulating power price function, $p_{\rm reg}$, is assumed to be time-varying for the initial 12 hours, but constant (the average value of the day-ahead prices in the data-set used in the backtest) in the following 36 hours. $p_{\rm reg}$ is assumed to be a weighted average of the spot prices and actual observed regulating power prices. This weighting will decay with varying deficiency parameter to quantify the importance of regulating power forecasts. Similarly, the CO2-emissions, c, is assumed to follow the exact CO2emissions during the initial 12 hours, but equals a constant (the average value of the CO2-emissions in the data-set used in the backtest) during the remaining 36 hours. This is inline with the example shown in Fig. 3.

The parameter q equals 1 MWh. Hence, the battery has a storage capacity of 1 MWh. The depth of charge and discharge parameters, \underline{x} and \overline{x} , equal 5 % and 95 %, respectively. The power bounds, \underline{u} and \overline{u} , are subject to variations as different battery configurations are investigated in the backtest; however, these will be 1.5 times larger in magnitude compared to the the nominal power, which is the primary battery-design parameter investigated in this study. The optimization horizon equals 48 hours. This paper will consider a fixed day-ahead strategy which always sells 2 times the rated discharge power, \underline{u} , of the battery. The model predictive control principle will be used to implement the optimization strategy: the optimal

Table 2: This table shows the aggregated results of the backtests conducted in th	his paper. The first four columns shows: the special							
down-regulation price assumption, the power configuration of the battery, the foreca	cast deficiency decay-rate and the CO2-tax parameter.							
The remaining six columns shows the performance metrics used in this paper: the total revenue from arbitraging the power markets, the								
accumulated CO2 reductions, the costs of investing in the given battery, the payback horizon of this investment, the costs of investing in the								
given battery using the future battery prices and the payback horizon of this investment.								
Price D. L. COOL D. D. L. C								

Price assumption (FUD (MWb)	Power (MW)	Loss (%)	CO2 tax (EUR/t)	Revenue (tEUR)	Reduc. CO2 (t)	Cur. inv. (tEUR)	Payback (years)	Fut. inv. (tEUR)	Payback (years)
(EUK/MWII)							· · ·		
0	1/5	0	0	52	1/	625	12.01	331	6 36
-10	1/0	0	0	75	14	025	8 35	- 551	0.50
-10			0	08	11		6.40	_	3 30
-20	1/3		0	77	11	828	10.40	433	5 58
-10	1/0	_	0	114	12	- 020	7.24		3 78
-20	_	_	0	151	12 12	_	5.47	_	2.85
20	1/2	_	0	101	19	1.082	10.42	559	5.38
-10	- 1/2	_	0	158	12	1,002	6.84	-	3.50
-20	-	_	0	213	13	_	5.09	-	2.63
				210	10		0.00		2.00
-10	1/3	10	0	114	15	828	7.26	432	3.79
-	-	25	-	113	15	-	7.33	-	3.82
-	-	50	-	109	16	-	7.57	-	3.95
-	-	10	25	114	19	-	7.26	-	3.79
-	-	25	-	113	18	-	7.33	-	3.83
-	-	50	-	109	20	-	7.56	-	3.95
-	-	10	50	114	23		7.27	-	3.79
-	-	25	-	113	21		7.34	-	3.83
-	-	50	_	109	22		7.58	-	3.96
-	-	10	75	114	26		7.27	-	3.80
-	-	25	-	113	24	-	7.35	-	3.84
-	-	50	-	109	25	-	7.59	-	3.96
-	-	10	100	114	29	-	7.28	-	3.80
-	-	25	-	113	26	-	7.36	-	3.84
-	-	50	-	109	27		7.61	-	3.97
-	-	25	250	111	37		7.45	-	3.89
-	-	-	500	109	46		7.61	-	3.97
-	-	-	1000	105	52		7.86	-	4.10
	-	-	5000	97	55	-	8.54	-	4.46
-10	1/3	5	0	114	14	828	7.25	432	3.78
-	-	20	-	113	14		7.30	-	3.81
-	-	33	-	112	15		7.39	-	3.86
-	-	75	-	104	18	-	7.95	-	4.15
-	-	100	-	90	24	-	9.15	-	4.77

control problem is solved after each operating hour to design the optimal charge and discharging signal, but it is only the first hour of the optimal solution which will be implemented.

The economical impact of activation for special downregulation is also investigated in the backtest. Activation for special down-regulation is defined when:

- 1. there has been demand for special down-regulation,
- 2. and the battery is not activated for regulating power; i.e. not activated at a more attractive power price in the regulating power market.

Different activation prices (inspired by Table 1) will be subject to investigation.

4. Results

This section presents the results of backtesting three different nominal power configurations of a Vanadium redoxflow battery (1/5 MW, 1/3 MW, 1/2 MW) together with three different special down-regulation activation prices (0 EUR/MWh, -10 EUR/MWh, -20 EUR/MWh). The forecast deficiency parameter (the decay from the true regulating power prices to the spot prices) are tested in the range from 0 % decay (i.e. perfect forecasts) to 100 % decay (i.e. that $p_{\rm reg}$ equals the spot prices during the initial 12 hours). Multiple values of $k_{\rm tax}$ is considered in the range from 0 EUR/t (no CO2-tax) to 5000 EUR/t (an extremely high CO2-tax).

Tab. 2 shows the aggregated results from the backtests. From this table the following general comments are relevant:

- 1. The revenue increases with more extreme special downregulation activation prices.
- 2. Large CO2 reductions are only observed for very high CO2-tax parameters, k_{tax} .
- 3. The revenues are highest for batteries with a high nominal power to physical storage ratio.

Comparing the revenue under different forecasting deficiencies (the bottom block in Tab. 2), the proposed strategy shows some dependency of the forecasting deficiency parameter (denoted loss in Tab. 2); when the quality of the forecasts decays, the revenue decreases as well. This relation is also shown in Fig. 4. Similarly, when the CO2tax increases - and the optimization problem weighs CO2 reductions higher - the power price arbitrage revenues decrease as well. This relation is shown in Fig. 5.

Fig. 6 shows one week of operation using the optimization strategy defined previously. From the top plot, it is clearly seen that this is a period with large amount of demand for special down-regulation. This has the implication that the day-ahead strategy yields high-rewards as the battery can be used to provide down-regulation at very attractive power prices (-10 EUR/MWh in this example). It is also observed that the battery discharges during periods with the highest power prices. During this week of operation, the strategy reduces the CO2 emissions by approximately 350 kg and creates a revenue of approximately 4,000 EUR.

5. Discussion

The previous section presented the results from backtesting the optimal control algorithm (5) during one year of operation (from December 2019 to December 2020). Two key assumptions in this study are: (1) the strategy is pricetaker of the regulating power prices, and (2) the strategy perfectly bids on activation for special down-regulation. Another implicit assumption is perfect state information about the state of charge of the battery.

5.1. Imperfect market information

The backtests conducted in this paper assumes a perfect ability to bid in the day-ahead and regulating power markets. However - as stated initially - these assumptions can be justified as one viable strategy to achieve this is to front-run both of these markets. The day-ahead biddingstrategy would be to bid very low prices to the day-ahead market (thereby ensuring the acquisition of the necessary day-ahead position). The regulating power bidding-strategy would be to bid the day-ahead price on the regulating power market (thereby ensuring the activation of the necessary regulating power volume). However, these strategies are only viable for insignificant volumes; applying this strategy to a high-volume battery portfolio will imply inferior pricing in both markets. Thus, this bidding strategy has to be extended in order to accommodate energy portfolio with higher volumes.

The regulating power price function, u_{reg} , assumes correct bias of the future regulating power price, as this function is a weighted average of the spot prices and actual observed regulating power prices. From Tab. 2 it is observed that - for a 1/3 MW nominal power battery with storage capacity 1 MWh and with a special down-regulation price assumption on -10 EUR/MWh - the difference between perfect forecasts (payback time of 3.8 years using future battery prices) and 100 % imperfect forecasts (payback time of 4.8 years using future battery prices) is an extension to the payback horizon of approximately 1.0 year (or 26 %). Thus, high-quality forecasts will dramatically improve the performance of the strategy. However, the regulating power market contains very large volumes of down-regulation (in particular special down-regulation) and is the primary explanation to the very attractive payback horizons shown in Tab. 2.

5.2. Supervised day-ahead trading

The day-ahead strategy proposed in this paper is very simple: sell three times the nominal discharging power



Figure 4: This figure shows the empirical dependency between the realized revenues from power price arbitraging as a function of the forecast deficiency of the regulating power prices.



Figure 5: This figure shows the empirical dependency between the realized revenues from power price arbitraging as a function of the realized CO2 reductions.



Figure 6: This figure shows the 1 week of operation in January 2020. The top plot shows the actual spot prices, regulating power prices and the CO2 emissions. The middle plot shows the state of charge and the implemented power. The bottom plot shows the accumulated revenue and the accumulated CO2 reductions realized by implementing the optimization strategy. The forecast deficiency is 25 %, the CO2 tax is 100 EUR/t and the special down-regulation price is assumed to be -10 EUR/MWh.

of the battery. An interesting case-study would be to consider extensions of this; relevant power system data might be the expected wind power production or the expected congestion on the Skagerak-cable to the Norwegian price zones, which contain high capacities of hydro-power. Wind power production volumes are relevant as the special down-regulation demand in DK1 often is due to excessive wind power production in Northern Germany [16]. The importance of congestion on the Skagerak-cable is perfectly illustrated in Fig. 1; the regulating power prices attain more extreme values under conditions with congestion on the Skagerag-cable. Combining these market features, it should be possible to formulate a supervised day-ahead (and possible also intra-day) strategy that trades relevant positions in the day-ahead and intra-day markets to accommodate attractive activation prices in the regulating power market.

5.3. Imperfect battery model

The optimal control problem (5) contains a one dimensional state equation, (1), that models the dynamical response of the state of charge, x, of the battery to the charged and discharged power, u. The model contains power dependent efficiencies of the charge and discharge

efficiencies, η_l and η_g , respectively. The parameters used in this paper are based on charge and discharge efficiencies provided by the Danish company VisBlue [2]. In actual implementations of the framework presented in this paper, the model would be calibrated (and re-calibrated repeatedly over time) to actual data of the battery subject to the optimal control algorithm. This can be done using methods presented [17] and [18]. Doing this, it might also be interesting to consider more elaborate and detailed dynamical battery models; e.g. equivalent circuit models [19] or more complicated electrochemical models [20?] of flow batteries.

6. Conclusion

This paper investigates the economic potential of operating Vanadium redox-flow batteries on two power markets in Northern Europe: the day-ahead and regulating power markets. From a large batch of backtests under multiple battery configurations and markets assumptions, payback horizons (using current battery prices) of 5.1 to 12.0 years are observed; the longest payback horizon of 12.0 years is observed using conservative price estimates of special down-regulation activation prices and the shortest payback horizon of 5.1 years is observed using more optimistic price estimates of these prices. These payback horizons reduce to 2.6 and 6.4 years, respectively, using expected battery prices in five years.

The importance of high-quality forecasts of the regulating power market is also investigated together with the potential impact of future CO2-taxes. The conclusion is: perfect forecasts improve the revenues by 26 % compared to having no forecasts available and - similarly - introducing high CO2-taxes reduces revenues by approximately 14 % but, increases the CO2 reductions by 340 %.

References

- Denmark strikes deal to slash CO2 emissions by 70% in a decade, Bloomberg (2019).
- [2] VisBlue A/S, Bautavej 1A, 8210 Aarhus V, Denmark, https://visblue.com/, 2020. [Online; accessed 30-June-2020].
- [3] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, D. S. Callaway, Model predictive control approach to online computation of demand-side flexibility of commercial buildings hvac systems for supply following, in: 2014 American Control Conference, IEEE, 2014, pp. 1082–1089.
- [4] K. Aduda, T. Labeodan, W. Zeiler, G. Boxem, Y. Zhao, Demand side flexibility: Potentials and building performance implications, Sustainable cities and society 22 (2016) 146–163.
- [5] N. B. Brok, T. Munk-Nielsen, H. Madsen, P. A. Stentoft, Flexible control of wastewater aeration for cost-efficient, sustainable treatment, IFAC-PapersOnLine 52 (2019) 494–499.
- [6] P. A. Stentoft, D. Guericke, T. Munk-Nielsen, P. S. Mikkelsen, H. Madsen, L. Vezzaro, J. K. Møller, Model predictive control of stochastic wastewater treatment process for smart power, costeffective aeration, Ifac-papersonline 52 (2019) 622–627.
- [7] A. W. Dowling, R. Kumar, V. M. Zavala, A multi-scale optimization framework for electricity market participation, Applied energy 190 (2017) 147–164.
- [8] H. Mohsenian-Rad, Optimal bidding, scheduling, and deployment of battery systems in california day-ahead energy market, IEEE Transactions on Power Systems 31 (2015) 442–453.
- [9] C. K. Ekman, S. H. Jensen, Prospects for large scale electricity storage in denmark, Energy Conversion and Management 51 (2010) 1140–1147.
- [10] R. S. Miller, G. Shorter, High frequency trading: Overview of recent developments, volume 4, Congressional Research Service Report, Washington, DC, 2016.
- [11] TenneT, DK1-DE COUNTERTRADE FOLLOWING JOINT DECLARATION 2018, Technical Report, TenneT, 2019.
- [12] K. Heussen, S. Koch, A. Ulbig, G. Andersson, Energy storage in power system operation: The power nodes modeling framework, in: 2010 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe), IEEE, 2010, pp. 1–8.
- [13] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, M. Diehl, CasADi – A software framework for nonlinear optimization and optimal control, Mathematical Programming Computation 11 (2019) 1–36.
- [14] VanadiumPrice.com, https://vanadiumprice.com/, 2020. [Online; accessed 30-June-2020].
- [15] P. Ralon, M. Taylor, A. Ilas, H. Diaz-Bone, K. Kairies, Electricity storage and renewables: Costs and markets to 2030, International Renewable Energy Agency: Abu Dhabi, UAE (2017).
- [16] N. B. Brok, T. Munk-Nielsen, H. Madsen, P. A. Stentoft, Unlocking energy flexibility of municipal wastewater aeration using predictive control to exploit price differences in power markets, Applied Energy 280 (2020) 115965.
- [17] N. L. Brok, H. Madsen, J. B. Jørgensen, Nonlinear model predictive control for stochastic differential equation systems, IFAC-PapersOnLine 51 (2018) 430–435.

- [18] N. R. Kristensen, H. Madsen, S. B. Jørgensen, Parameter estimation in stochastic grey-box models, Automatica 40 (2004) 225–237.
- [19] Y. Zhang, J. Zhao, P. Wang, M. Skyllas-Kazacos, B. Xiong, R. Badrinarayanan, A comprehensive equivalent circuit model of all-vanadium redox flow battery for power system analysis, Journal of Power Sources 290 (2015) 14–24.
- [20] C. Blanc, A. Rufer, Multiphysics and energetic modeling of a vanadium redox flow battery, in: 2008 IEEE International Conference on Sustainable Energy Technologies, IEEE, 2008, pp. 696–701.

Optimization-based state estimation for stochastic differential equation systems

Niclas Brok, Uffe Høgsbro Thygesen, Jan Kloppenborg Møller, Henrik Madsen

Department of Applied Mathematics and Computer Science at the Technical University of Denmark.

Abstract

Stochastic differential equation models provide a natural way to represent physical systems as they evolve in continuous-time. In contrast to discrete-time models, a priori information about the system can be built into the model and the estimated model parameters will be independent of the sampling time. This paper presents an optimization-based filtering method for state estimation of general nonlinear stochastic differential equations with general discrete-time likelihood models. The performance of the optimization-based filtering method is compared to the performance of the continuous-discrete extended Kalman filter and the methods are tested in a simulation study using a stochastic extension to the van der Pol oscillator model. The optimization-based filtering method significantly outperforms the continuous-discrete extended Kalman filter in the ability to efficiently reconstruct the hidden state modelled by the stochastic differential equation.

Key words: Stochastic Differential Equations, Optimization, Maximum a Posteriori, Filtering, State Estimation

"Future low-carbon societies will be driven by renewable energy sources (e.g. wind and solar power). This will flip the characteristics of our power system from a production-tracking-consumption paradigm, to a consumption-tracking-production paradigm. This will increase the need of complex coordination of our power consumption as power grids require a strict balancing between power production and consumption. This dissertation investigates the potential of applying nonlinear model predictive control algorithms to solve complex power market coordination problems, where flexible consumers leverage the more volatile balancing power prices and thereby in-directly help neutralizing production and consumption imbalances."

Technical University of Denmark Department of Applied Mathematics and Computer Science

Richard Petersens Plads, Building 324 2800 Kgs. Lyngby Tel.: +45 4525 3031

www.compute.dtu.dk





