NTNU

Norwegian University of
Science and Technology

Estimation with low measurement frequency

Martin Leite

Master of Science in Cybernetics and Robotics
Submission date: June 2017
Supervisor: Morten Hovd, ITK

Norwegian University of Science and Technology
Department of Engineering Cybernetics

Address: Cybernetica AS
Leirfossveien 27
N-7038 Trondheim
Norway
Phone.: +47 73 82 28 70
Fax: +47 73 82 28 71

CYBERNETICA
Model based control systems for the process industry

MASTER'S THESIS 2017

To: Martin Leite

Number of pages: 1
Date: 2016-11-25

Estimation with low measurement frequency

When online estimators are used in industrial applications for decision support systems or for control schemes based on soft sensor principals, it is common to have measurements with varying quality and measurement frequencies available. Measurements can be available both continuously and very rarely, typically through manual samples with a following lab analysis. In such cases it is not obvious which type of estimator that should be chosen, and how estimator parameters should be tuned for optimal performance. The Moving horizon Estimator (MHE) has obvious advantages when the measurements are delayed, since these measurements can be placed on the proper place in the time series, while a Kalman filter (KF) must use a delayed measurement at the time of arrival, or reject the measurement if it is too old to be relevant for the process state.

Manual measurements and chemical lab analyses can be resource demanding, and industrial companies are therefore interested in reducing the frequency of these measurements as much as possible. As a result the performance of the process may be reduced due to a larger time delay and slower control loops. If an estimator model is available, it is however possible to use measurements with different measurement frequencies or measurements with the same measurement frequency, but that are taken at different times, to provide sufficient information so that the control of the process is sufficiently fast during normal operation. It is desirable to find criteria for how to best choose measurement frequencies and when to do the measurements for a predefined model structure.

The thesis is done in collaboration with Cybernetica AS

Goals for the work

- Find criteria for how measurement and process noise should be chosen in a Kalman filter where the span of measurement frequencies is large.
- Evaluate if it is reasonable that low frequency measurements are taken at the same time, of if they should be taken at different times, in cases where it is possible to influence this.

Master's thesis

- Implement a simple model of an electrolysis cell in Cybernetica's tools. The model will be used both as a process and a model with different choices of parameters.
- Simulate the model with a Kalman filter where there is a large span of measurement frequencies. Use different choices of noise parameters and evaluate how this affects the Kalman filters performance based on e.g. prediction deviation. Come with suggestions of criteria on how to choose parameters for optimal performance of the Kalman filter.
- Simulate the model with a Kalman filter where the measurement frequency and the time the measurements are taken vary. Come up with suggestions for criteria on how to optimally choose frequencies and measurement times.

Small adjustments of the tasks may be done during the work.

Summary

The invention of the computer has triggered a technological revolution in many parts of society. For the process industry, computers have led to improvement in both production rates, safety and profits. One way this has happened is through the use of online estimators. In the process industry, such estimators are typically used for decision support systems or as soft sensors for advanced process control. The Kalman filter has been one of the more popular estimators since its introduction in 1960. Since then, many versions of the Kalman filter have been developed in order to deal with systems with different characteristics. For nonlinear systems where all measurements are not available on every time step, the Multi-rate Extended Kalman filter can typically be applied. It is not straight forward to decide how the Kalman filter should be tuned for the best possible performance when different measurements arrive at different rates.

This thesis has two main objectives. The first is to look at the possibility of determining any criteria for how the noise parameters of the Multi-rate Extended Kalman filter should be tuned to achieve the best possible estimates when the measurement frequencies vary. The second objective is to evaluate if it is reasonable to perform all the different infrequent measurements simultaneously or if they should be spread out, for cases where it is possible to decide the measurement frequencies.

To be able to test these objectives, a model of the the aluminium electrolysis process has been developed. The aluminium electrolysis process is used as an example process because it is common to use multiple infrequent process measurements like the temperature and height of the electrolysis bath, the metal height and the aluminium fluoride concentration in the monitoring of aluminium electrolysis cells. Early in the thesis, an introduction to the aluminium electrolysis process and the theory behind the Discrete, Extended and Multi-rate Extended Kalman filters is given. This lies to foundation for the model development and testing of the Kalman filter properties.

In order to investigate the objectives of the thesis, the developed model has been used both as a simulated process with one set of parameters, and as a model using a slightly different set of parameters. The difference in parameters is used to simulate modelling errors, as it will create a process-model mismatch. The Kalman filter has then been used to correct for this mismatch by estimating on the changed parameters. The obtained results seem to indicate that there is no significant difference in estimator performance between using joint measurements or shifting the time at which measurements are performed. With regard to Kalman filter tuning for varying measurement frequencies, non of the results suggest that the optimal tuning is significantly altered by a change in measurement frequency.

Sammendrag

Datamaskinens inntog har utløst en teknologisk revolusjon i store deler av samfunnet. For prosessindustrien har datamaskinene ført til store forbedringer i produksjon, inntjening og sikkerhet. En av måtene dette har hendt på er gjennom bruken av online estimatorer i drift. Disse kan typisk bli brukt av prosessindustrien til driftsstøttesystemer eller som softsensorer for avansert prosessregulering. Kalmanfilteret har vært en av de mest populære estimatorene på markedet siden det ble introdusert i 1960. Siden da har flere ulike versjoner blitt utviklet for å håndtere systemer med ulike egenskaper. For ulineære system hvor all prosessens målinger ikke er tilgjengelig på hvert tidssteg er det mulig å bruke et såkalt Multi-rate Extended Kalmanfilter. Det er ikke opplagt hvordan et slikt Kalmanfilter skal tunes når målefrekvensene i systemet varierer.

Denne avhandlingen har to hovedmål. Det første er å se om det er mulig å finne kriterier for hvordan støyparametrene i et Multi-Rate Extended Kalmanfilter bør tunes for å oppnå et best mulig estimat når målingenes frekvens endres. Det andre målet er å vurdere om det er hensiktsmessig at alle sjeldne målinger utføres på samme tidspunkt eller om de bør spres ut på ulike tidspunkt, i tilfeller hvor det er mulig å påvirke dette.

En matematisk modell av aluminiumselektrolyse-prosessen har blir utviklet for å teste ut de nevnte målene. Denne prosessen har blitt brukt som eksempel fordi det er en prosess hvor det er vanlig å utføre sjeldne målinger av både temperature og høyde i elektrolysebadet, metallhøyde og konsentrasjonen av alumimiumfluorid i overvåkningen av cellens tilstand. Teori om aluminiumselektrolyse-prosessen og Discrete, Extended og Multi-rate Extended Kalmanfilter er gitt tidlig i avhandlingen. Dette legger grunnlaget for modellutviklingen og testing av Kalmanfilterets egenskaper.

Den utviklede modellen har blitt brukt både som simulert prosess med et sett parametre og som modell med et litt annet parameter-sett for å undersøke avhandlingens mål. Forskjellen i parametre er brukt for å simulere modelleringsfeil, ettersom det vil føre til at avvik mellom oppførselen til den simularte prosessen og modellen. Kalmanfilteret har så blitt brukt til å korrigere for dette avviket ved å estimere på de endrede parametrene. De oppnådde resultatene indikerer at det ikke er noen betydelig forskjell mellom å utføre alle sjeldne målinger samtidig og å spre dem utover. Når det gjelder optimal tuning av Kalmanfilteret ved varierende målefrekvens, så ble det ikke funnet noen antydning til at denne endrer seg i noen stor grad.

Preface

This master's thesis was written by Martin Leite during the spring of 2017. It is the final part of a 5 year study program leading to a M.Sc. in Engineering Cybernetics at the Norwegian University of Science and Technology (NTNU).

The project has been done in collaboration with Cybernetica, whom suggested the topic. It is an extension of an internship at Cybernetica during the summer of 2016 and a compulsory project at NTNU during the subsequent autumn. I would like to thank everyone at Cybernetica for making me feel welcome during this past year. Most of the work this past spring has been done at Cybernetica's offices in Trondheim, and it has been inspiring to work alongside experts in process control engineering. A special thanks goes out to my co-supervisor at Cybernetica, Dr. Stein O. Wasbø, for always setting aside time when help was needed.

I would also like to thank Prof. Morten Hovd, my supervisor at NTNU, for all support along the way.

I would like to express my gratitude towards Ellen and Øyvind, my fellow students, who have made the office a fun place to be this semester.

Last but not least I would like to thank my parents, Hilde and Bård Arne Leite, for all their love and support during my years at the university.

Martin Leite,
Trondheim, June 2017

Table of Contents

Summary i

Sammendrag ii

Preface iii

Table of Contents v

List of Tables vii

List of Figures ix

List of Symbols xiii

Abbreviations xix

1 Introduction 1
- 1.1 Background and motivation . . . 1
- 1.2 Goal and Method . . . 2
- 1.3 Outline . . . 2

2 Aluminium electrolysis process 5
- 2.1 Aluminium . . . 5
- 2.2 The Hall-Héroult process . . . 5
  - 2.2.1 Anode design . . . 6
- 2.3 Cell operation . . . 7
  - 2.3.1 Cell performance . . . 7
  - 2.3.2 Heat balance . . . 7
  - 2.3.3 Measurements and control inputs . . . 8
  - 2.3.4 Normal cell operations . . . 9
  - 2.3.5 Unwanted disturbances . . . 9
  - 2.3.6 Control . . . 10

3 Kalman filter theory **11**
3.1 Discrete Kalman filter (DKF) . . . 12
3.2 Kalman filter tuning . . . 14
3.3 Extended Kalman filter (EKF) . . . 14
3.4 Multi-rate Extended Kalman filter . . . 16

4 Aluminium process modelling **19**
4.1 Physical description . . . 19
4.2 Mass balance . . . 22
4.2.1 Theory of mass balances . . . 22
4.2.2 Aluminium electrolysis cell mass balance . . . 23
4.3 Energy balance . . . 28
4.3.1 Energy balance theory . . . 28
4.3.2 Mechanisms of heat transfer . . . 31
4.3.3 Aluminium electrolysis cell energy balance . . . 33
4.4 Model inputs and outputs . . . 40
4.5 Model behaviour . . . 41
4.6 Summary of assumptions used . . . 46
4.7 Model conclusion . . . 47

5 Kalman filter testing **49**
5.1 Utilised software - Modelfit . . . 52
5.2 Shifted measurements versus joint measurements . . . 53
5.2.1 Problem description . . . 53
5.2.2 Method . . . 53
5.2.3 Results . . . 56
5.2.4 Discussion . . . 59
5.2.5 Conclusion . . . 64
5.3 Varying measurement frequencies . . . 64
5.3.1 Problem description . . . 64
5.3.2 Method . . . 64
5.3.3 Results . . . 67
5.3.4 Discussion . . . 82
5.3.5 Conclusion . . . 83

6 Overall summary, conclusion and future work **85**
6.1 Summary and conclusion . . . 85
6.2 Future work . . . 86

Bibliography **89**

Appendix **I**
C Calculation of the bath resistivity . . . I

List of Tables

5.1 Maximum absolute values for the white noise on the process measurements. The percentages are included because they describe the basis for how the absolute values are decided. A potential improvement of the measurement noise model would be to decide constant absolute max-values instead of calculating them as a percentage of the mean value of the measurement in a data set 51

5.2 Kalman filter measurement noise tuning values when using noise-free measurements for estimation 56

5.3 Kalman filter measurement noise tunings when using noisy measurements for estimation 59

5.4 Estimator performance results for days 35-105 of the long simulations shown in Figure 5.9 62

5.5 The maximum absolute value of noise on different measurements in section 5.3 66

5.6 Mean and standard deviation for the difference between the estimate of h_{top-0} and its actual process value for the same tunings as in Figure 5.14 69

5.7 Mean and standard deviation between the measurement estimates of the model and the noise-free process measurements for the tunings of Figure 5.16. The analysis is done for the last half of the data sets. Units for mean and st. dev.: Bath temp. $[^{\circ}\mathrm{C}]$, Acidity $[\mathrm{wt}\%]$, Bath height $[\mathrm{m}]$ 71

5.8 Mean and standard deviation for the difference between the estimated metal height and the actual process value. As the parameter needs some time to converge towards the correct value, the mean and standard deviation are calculated from day 35. Unit for mean and st.dev.: Metal height $[\mathrm{m}]$ 76

5.9 Mean and standard deviation of the difference between the estimate of h_{top-0} and its actual process value during the last two thirds of the data series (from approximately day 35). Tunings are the same as in Figure 5.23 78

5.10 Mean and standard deviation between the measurement estimates of the model and the noise-free process measurements. The chosen tunings are all tunings where the r_{BT} is set to 3. The analysis is done for the last two thirds (from day 35) of the data sets. Units for mean and st. dev.: Bath temp. $[^{\circ}\mathrm{C}]$, Acidity $[\mathrm{wt}\%]$, Bath height $[\mathrm{m}]$ 81

List of Figures

2.1 Prebake aluminium electrolysis cell 6

4.1 Aluminium cell model 20
4.2 Heat conduction through a wall consisting of two different materials . . . 32
4.3 Aluminium cell energy model 33
4.4 Complexity of top heat loss 38
4.5 The simulated metal production and tapping. Metal is tapped daily 42
4.6 The change in bath temperature, liquidus temperature, side ledge thickness and acidity from a step in the ACD 43
4.7 The mass of liquid cryolite increases as the side ledge melts. Bath is tapped about 17 days into the simulation. After the bath is tapped, the steady increase in cryolite mass shows that some cryolite will be produced continuously by the bath production reaction 44
4.8 Model heat distribution 45
4.9 Typical heat loss distribution for a Hall-Héroult cell. Collected from Grjotheim et al. (1993) 45

5.1 Interaction between the process, model and Kalman filter 50
5.2 Deviation between process and model states from a parameter change . . 54
5.3 Simultaneous versus shifted measurement strategies 54
5.4 Simulated process outputs 55
5.5 Estimation of h_{top-0} towards a reference value of 7 for three different tunings with noise-free measurements. The simulations have been split up to show the show the initial step of each tuning in the left plots and finer variations in the right plots 57
5.6 Deviation between noise-free measurements and estimator output for bath temperature and acidity with three different Kalman filter tunings. Estimation done with noise-free measurements 58

5.7 Estimation of h_{top-0} towards a reference value of 7 for 3 different tunings with noisy measurements. The left plots show the first 17 days of the simulation while the right plots show the remaining simulation time. Notice that the scale of the y-axis for the bottom two right plots is different from the rest 60
5.8 Deviation between noise-free measurements and estimator output for bath temperature and acidity measurements using three different Kalman filter tunings. Estimation done with noisy measurements 61
5.9 Parameter response for a process where the states vary a lot. The left plots show the initial step of the parameter while the right plots show the finer variations after the value has closed in on its correct value 62
5.10 Estimation of two parameters in the heat balance at once. The reference value is 7 for both parameters. It is clear that the system manages to find another combination of parameter values that closes the 63
5.11 Process measurements without noise used throughout section 5.3. The process output responses of the shorter simulations are given on the left while the longer ones are given on the right side 65
5.12 Response in c_{tap} from varying the $\Delta t_{m,BT}$ and the tuning of bath temperature, acidity and bath height. There are little variations as long as r_{MH} and $\Delta t_{m,MH}$ are kept constant. The legend shows the measurement frequency and the measurement noise tuning used for the simulations . . 67
5.13 Simulation results for estimating h_{top-0} with different choices of r_{BT} tunings when $\Delta t_{m,BT}$ is varied. The legend shows the measurement noise tuning used for the simulations 68
5.14 The best tuning results for each combination of $\Delta t_{m,BT}$ and the chosen r_{BT} values when estimating h_{top-0}. The legend shows the measurement noise tuning used for the simulations 70
5.15 The best obtained estimate of h_{top-0} for each $\Delta t_{m,BT}$. The legend shows the measurement noise tuning used for the simulations 71
5.16 Measurement estimates for the tunings of Figure 5.15. The legend shows the measurement noise tuning used for the simulations 72
5.17 The bath height tuning doesn't affect the estimation of c_{tap}. The legend shows the measurement noise tuning used for the simulations 73
5.18 Wide spread tuning results for estimation of c_{tap} when $\Delta t_{m,MH}$ is varied. The legend shows the measurement noise tuning used for the simulations 74
5.19 Longer simulations for estimation of c_{tap} with varying interval for the metal height measurements. The legend shows the measurement noise tuning used for the simulations 75
5.20 Comparison of estimator performance with the same tuning for different metal height measurement frequencies. The legend shows the measurement frequency and the measurement noise tuning used for the simulations 76
5.21 Comparison of estimation performance for the best tuning obtained for each estimate. The legend shows the measurement frequency and the measurement noise tuning used for the simulations 77

5.22 Wide spread tuning results for estimation of h_{top-0} when $\Delta t_{m,BT}$, $\Delta t_{m,A}$ and $\Delta t_{m,BH}$ are varied simultaneously. The legend shows the measurement noise tuning used for the simulations 79
5.23 Fine tuning results for estimation of h_{top-0} when $\Delta t_{m,BT}$, $\Delta t_{m,A}$ and $\Delta t_{m,BH}$ are varied simultaneously. The legend shows the measurement noise tuning used for the simulations 80

List of Symbols

Model

Latin letters:

A	Cross section area
A_{anodes}	Horizontal anode cross section area
A_{sl}	Side ledge cross section area
A_{top}	Horisontal area of cell inner dimensions
B	Quantity of an extrinsic variable within a control volume
B_{in}	Quantity of an extrinsic variable added to a control volume
B_{out}	Quantity of an extrinsic variable removed to a control volume
$B_{generated}$	Quantity of an extrinsic variable generated within a control volume
B_{out}	Quantity of an extrinsic variable consumed within a control volume
c_p	Specific heat capacity
$\bar{c}_{p_{Al_2O_3}}$	Average specific heat capacity of aluminium oxide
$\bar{c}_{p_{AlF_3}}$	Average specific heat capacity of aluminium fluoride
$c_{p_{an}}$	Specific heat capacity of anodes
$c_{p_{bath,l}}$	Specific heat capacity of liquid bath
$c_{p_{bath,s}}$	Specific heat capacity of solid bath
$c_{p_{cry,l}}$	Specific heat capacity of liquid cryolite
$c_{p_{cry,s}}$	Specific heat capacity of solid cryolite
$c_{p,i}$	Specific heat capacity of component i in a control volume
$c_{p_{met}}$	Specific heat capacity of aluminium
$c_{p_{wall}}$	Specific heat capacity of cell wall
c_{tap}	Parameter describing uncertainty in metal tapping weight
CE	Current efficiency
C_p	Heat capacity
E	Total energy within a control volume
E_i	Total energy of a component within a control volume
E_{in}	Energy added to a control volume through mass flows
E_{out}	Energy taken out of a control volume through mass flows
$E_{r_{bath\,prod}}$	Energy produced by bath production reaction
$E_{r_{ox}}$	Energy consumed by primary reaction
$E_{sc_{freeze}}$	Energy flow from state transformation when freezing bath
$E_{sc_{melt}}$	Energy flow from state transformation when melting side ledge
$E_{tc_{freeze}}$	Energy flow from decreasing bath to freeze temperature
$E_{tc_{melt}}$	Energy flow from increasing side ledge to melt temperature
E_u	Total energy needed to heat components added to bath
F	Faradays constant
h	Convective heat transfer coefficient
h_{acd}	Anode-cathode distance (ACD)

h_{an-0}	Convective heat transfer coefficient between anodes and ambient
$h_{bath-an}$	Convective heat transfer coefficient between bath and anodes
$h_{bath-met}$	Convective heat transfer coefficient between bath and metal pad
$h_{bath-sl}$	Convective heat transfer coefficient between bath and side ledge
h_{cat-0}	Convective heat transfer coefficient between cathode and ambient
h_{cell}	Cell height
$h_{met-cat}$	Convective heat transfer coefficient between metal pad and cathode
h_{top-0}	Proportionality constant for linear heat loss through crust
h_{wall-0}	Convective heat transfer coefficient between cell wall and ambient
H	Enthalpy
$H_{in,i}$	Total enthalpy of mass flows into a control volume
$\hat{H}_{in,i}$	Specific enthalpy of mass flows into a control volume
$\hat{H}_{out,i}$	Specific enthalpy of mass flows out a control volume
$H_{out,i}$	Total enthalpy of mass flows out a control volume
$\Delta_{fus}H_{ALF_3}$	Heat of fusion for aluminium fluoride
$\Delta_{fus}H_{AL_2O_3}$	Heat of fusion for aluminium oxide
$\Delta_{fus}H_{bath}$	Heat of fusion for the bath
$\Delta_{fus}H_i$	Heat of fusion for component i in a control volume
ΔH_{cry}	Energy required to melt one kg of side ledge
$\Delta H_{r_{bath\,prod}}$	Heat of bath production reaction
$\Delta H_{r_{ox}}$	Heat of primary reaction
I_{line}	Line current
k	Conductive heat transfer coefficient
k_{an}	Conductive heat transfer coefficient of anodes
k_{cat}	Conductive heat transfer coefficient of cathode
k_{sl}	Conductive heat transfer coefficient of side ledge
k_{wall}	Conductive heat transfer coefficient of cell wall
K_p	Proportional gain of P-controller
l_{anode}	Anode length
l_{cell}	Cell length
m_{an}	Total mass of anodes
m_{bath}	Total mass of bath
m_{cat}	Mass of cathode
m_{cry}	Mass of liquid cryolite in the bath
m_{fl}	Mass of aluminium fluoride in the bath
m_i	Mass of general substance i
m_{met}	Mass of aluminium in metal pad
m_{ox}	Mass of aluminium oxide in the bath
m_{sl}	Mass of side ledge
m_{wall}	Mass of cell wall
M_{Al}	Molar mass of aluminium
$M_{Al_2O_3}$	Molar mass of aluminium oxide
M_{AlF_3}	Molar mass of aluminium fluoride
M_{Na_2O}	Molar mass of sodium oxide
$M_{Na_3AlF_6}$	Molar mass of cryolite

N_{anodes}	Number of anodes
p	Pressure
P_{el}	Total electric power input of the cell
Q	Energy transferred as heat through the control surface of a control volume
Q_{an-0}	Heat loss from the anodes to ambient
Q_i	Energy from component i transferred as heat through the control surface of a control volume
$Q_{bath-an}$	Heat loss from the bath to the anodes
$Q_{bath-liq}$	Heat loss from the bath to surface of the side ledge
$Q_{bath-met}$	Heat loss from the bath to the metal pad
$Q_{bath-sl}$	Heat loss from the bath to the side ledge
Q_{cat-0}	Heat loss from the cathode to ambient
Q_{crust}	Heat loss through top crust
Q_{liq-sl}	Heat loss from bath to side ledge
$Q_{met-cat}$	Heat loss from metal pad to cathode
$Q_{sl-wall}$	Heat loss from side ledge to wall
Q_{wall-0}	Heat loss from cell wall to ambient
$r_{i,j}$	Reaction rate of substance i in reaction j
r_{ox}	Reaction rate of primary reaction
$r_{bath\,prod}$	Reaction rate of bath production reaction
R_i	Thermal resistance for a mass or surface i
R_{pseudo}	Pseudo-resistance, resistance between anode and cathode of an aluminium electrolysis cell
Δt	Time span
$\Delta t_{m,A}$	Measurement frequency of acidity measurement
$\Delta t_{m,BH}$	Measurement frequency of bath height measurement
$\Delta t_{m,BT}$	Measurement frequency of bath temperature measurement
$\Delta t_{m,i}$	Measurement frequency of measurement i
$\Delta t_{m,MH}$	Measurement frequency of metal height measurement
T	Temperature
T_{an}	Temperature at the center of the anodes
T_{bath}	Temperature at the center of the bath
T_{cat}	Temperature at the center of the cathode
T_i	Temperature of component i in a control volume
T_{liq}	Liquidus temperature
T_{met}	Temperature at the center of the metal pad
T_{sl}	Temperature at the center of the side ledge
T_{wall}	Temperature at the center of the cell wall
T_0	Ambient temperature
ΔT	Temperature difference
u_{an}	Anode movement
$u_{bath\,add}$	Bath addition
$u_{bath\,tap}$	Bath tapping
u_{fl}	Aluminium fluoride feed

$u_{met\,tap}$	Metal tapping
u_{ox}	Aluminium oxide feed
U_{el}	Total voltage drop in cell
U	Internal energy in a control volume
U_i	Internal energy of a component within a control volume
$\hat{U}_i$	Specific internal energy of a component within a control volume
$U_{in,i}$	Internal energy added to a control volume through mass flows
U_{loss}	Voltage loss
$U_{out,i}$	Internal energy taken out of a control volume through mass flows
V	Volume
V_{ext}	Zero current intercept of U_{el} versus I_{line} for small changes in current
w_{anode}	Anode width
w_{cell}	Cell width
w_{freeze}	Mass flow of cryolite from bath to side ledge as the bath freezes
$w_{in,i}$	Mass flow of substance i into a control volume
$w_{out,i}$	Mass flow of substance i out of a control volume
W	Energy added to a control volume from the environment as work
W_i	Energy added to component i in a control volume from the environment as work
W_{flow}	Work done on a control volume by mass flows
$W_{flow,i}$	Work done on component i by mass flows
$W_{n,i}$	Work done on component i by other things than mass flows
w_{melt}	Mass flow of cryolite from side ledge to bath as the side ledge melts
x	General wall thickness
x_{an}	Anode thickness
x_{cat}	Cathode thickness
x_{sl}	Side ledge thickness
x_{wall}	Cell wall thickness
z	Number of electrons involved in a reaction
$[Al_2O_3]$	Concentration of aluminium oxide in the bath
$[Al_2O_3]_{additives}$	Concentration of aluminium oxide in the bath additions
$[Al_2O_3]_{ref}$	Reference concentration of aluminium oxide for the P-controller
$[AlF_3]$	Concentration of aluminium fluoride in the bath
$[AlF_3]_{additives}$	Concentration of aluminium fluoride in the bath additions
$[CaF_2]$	Concentration of calcium fluoride in the bath
$[Na_2O]_{ox-feed}$	Concentration of sodium oxide in the aluminium oxide feed

Greek letters:

ρ_{bath}	Density of bath
ρ_i	Density of component i
ρ_{sl}	Density of side ledge
$\rho_{Na_3AlF_6}$	Density of cryolite

Kalman filter

Latin letters:

A	Matrix relating current and next state
B	Matrix relating current input and next state
e	A posteriori estimation error
e^-	A priori estimation error
$\hat{e}$	Estimate of the prediction error
$\tilde{e}_x$	Prediction error
$\tilde{e}_z$	Measurement residual
$f(x, u, w)$	Nonlinear function relating current input, state and process noise to the next state
$h(x, w)$	Nonlinear function relating current state and measurement noise to the measurements
H	Matrix relating current state to the measurements
I	Identity matrix
K	Kalman gain matrix
P	A posteriori error covariance matrix
P^-	A priori error covariance matrix
P_0^-	Initial a priori error covariance matrix
Q	Process white noise covariance matrix
r_A	Standard deviation of the noise on the acidity measurement
r_{BH}	Standard deviation of the noise on the bath height measurement
r_{BT}	Standard deviation of the noise on the bath temperature measurement
r_i	Standard deviation of the noise on an arbitrary measurement i
r_{MH}	Standard deviation of the noise on the metal height measurement
R	Measurement white noise covariance matrix
t	Time
u	Input vector
v	Measurement white noise vector
V	Jacobian matrix connecting measurement noise to measurements
w	Process white noise vector
W	Jacobian matrix connecting process noise to next state
x	State vector
$\tilde{x}$	Approximate state estimate
$\hat{x}$	A posteriori state estimate
$\hat{x}^-$	A priori state estimate
$\hat{x}_0^-$	Initial a priori state estimate
z	Measurement vector
$\tilde{z}$	Approximate measurement vector

Greek letters:

ε	Random variable for the error process
η	Random variable for the error process

Abbreviations

ACD	Anode-Cathode Distance
BT	bath temperature
DKF	Discrete Kalman Filter
EKF	Extended Kalman Filter
MH	Metal Height
MIMO	Multiple-Input-Multiple-Output

Chapter 1

Introduction

1.1 Background and motivation

In the past century, the world has been witness to a tremendous technological development. The process industry has taken advantage of the new possibilities that computers provide to improve production rates, safety and profits. One application for computers in industrial applications is the use of online estimators[1]. These can typically be used in decision support systems or as soft sensors for advanced process control. For the estimator to be able to help the model resemble the plant in a good way, process measurements are necessary. The quality of these measurements and the frequency at which they are obtained can vary a lot in an industrial process. There can be everything from continuous sensor signals to manual samples that need to be analysed in a lab. A common challenge with measurements obtained in the lab is that there might be a significant time delay between the moment at which the sample was taken and when the result is available. This rules in favor of choosing the *Moving Horizon Estimator*, which makes it possible to place the measurement at the correct place in the time series when the lab result arrives. When delayed measurements are not an issue, the Kalman filter can be a good choice. The so-called *Multi-rate Extended Kalman filter* is a variant of the Kalman filter which handles nonlinear processes at the same time as it allows for different measurements to arrive at different frequencies.

A common decision parameter in processes where infrequent measurements are present is the frequency at which these measurements should be taken. There is a cost tied to each performed measurement, both in the form of operators, equipment and the increased HSE risk from operators being in a dangerous environment. This indicates that companies try to keep the frequency of the measurements as low as possible. At the same time, if the measurements are too far apart, it will not be possible to use them directly as a reference in feedback control. This is where a process model and a Kalman filter estimator come in, as they provide an estimate of the process state between measurements. The more frequently different measurements arrive, the better the process estimate. A good estimate

[1] Mathematical model made to resemble the states of a plant when some states are not directly observable

can be used as a sufficiently fast reference for the process control system, and it can also cut costs by reducing consumption of energy and materials in the production. Choosing the optimal measurement strategy can then have a significant impact on the economy of a process plant.

The aluminium electrolysis process is an example of a process where both continuous resistance measurements and infrequent measurements like the electrolysis bath temperature, metal height and aluminium fluoride concentration are used to obtain essential information about the state of the cell (Hestetun, 2009). With so many infrequent measurements, it is possible to achieve a lot of flexibility with respect to measurement frequency and pattern. This thesis looks at how the tuning and performance of a Kalman filter estimator when such factors as measurement frequency, pattern and quality changes.

1.2 Goal and Method

The goal of this thesis is to investigate how the optimal tuning and performance of a Kalman filter is affected by certain measurement qualities. To be more specific, the goals are:

- To find out whether it is possible to determine any criteria for how to tune the measurement and process noise parameters of a Kalman filter when the measurement frequency of the system varies.
- To evaluate if infrequent measurements should all be performed seldom at the same time, or if they should be spread out so that small bits of information are provided more often.

To test these properties, a simple simulation model of an aluminium electrolysis cell has been created. The model is not made to be a precise representation of an actual cell. It is instead supposed to resemble the main dynamics of it. The above goals have then been investigated by utilising this model both to create process measurements and as a process model. Some model parameters were changed in the model to simulate the fact that no model will describe the process perfectly. The idea is that the estimator will then use the process model and measurements to correct the changed parameters to compensate for the mismatch between them. Simulation results will then hopefully be able to give some indications on how measurement frequencies and strategies affect both tuning and performance of the Kalman filter.

1.3 Outline

This report is divided into several chapters. First, an introduction to the aluminium electrolysis process is given in Chapter 2. Then, Chapter 3 moves on to describe some of the fundamental theory regarding the Discrete Kalman filter, Extended Kalman Filter and the Multi-rate Extended Kalman filter. A model describing most of the main dynamics of the mass and energy balance of an aluminium electrolysis cell is then created in Chapter 4. In Chapter 5, this model is then used both as process and model to test the goals of the thesis

through simulations. All simulations have been performed using Cybernetica's Modelfit software. To test if it is possible to find any criteria for how to choose the tuning of the Kalman filter when measurement frequencies vary, simulations have been run where the frequencies of different combinations of measurements and the Kalman filter measurement noise parameters have been varied. Then the best tuning and its performance has been documented to see how they change with varying measurement frequencies. When testing if all measurements should be taken simultaneously or if they should be spread out, this has been done by looking at a wide range of different Kalman filter tunings, and how the results of different tunings are affected by the choice of measurement strategy. In the end, Chapter 6 summarises the work, draws some concluding remarks and provides propositions for future work.

Chapter 2

Aluminium electrolysis process

The purpose of the following chapter is to introduce the reader to some core concepts regarding the aluminium electrolysis process. If not otherwise stated, the information is collected from Grjotheim et al. (1993) and Hestetun (2009). Roughly the same process description is used in Leite (2016)

2.1 Aluminium

Aluminium makes up about 8 % of the Earth's crust. That means it is the third most abundant element, and the most abundant metal. Pure aluminium is rarely found in nature because of the strong binding between aluminium and oxygen. Aluminium oxide (Al_2O_3), often called alumina, is the compound used in aluminium production. A commonly used source of alumina is Bauxite. This ore typically contains 40-60 % alumina. The alumina, in the form of a white powder, is recovered from bauxite in a process called the Bayer process. Through the Hall-Héroult process the alumina is then turned into aluminium. Aluminium finds a vast number of applications in modern society. Aluminium alloys are used in vehicles, packaging like foil and cans, construction plating and household items, to mention but a few of the possible usages of this versatile material.

2.2 The Hall-Héroult process

Aluminium has been produced using the well-established Hall-Hérault process for more than a century. Working separately, both Paul Hérault and Charles Hall applied for patents on the concept in France and in the U.S. respectively in 1886. Even though a lot of incremental improvements have been made since then, to this day, the same core principle is still the only commercially viable alternative for the production of aluminium. An illustration of an aluminium reduction cell can be found in Figure 2.1. Two things about this illustration that is important to be aware of is that the dimensions are not accurate (e.g. the distance between the bottom of the anode and the top of the aluminium) at the same time

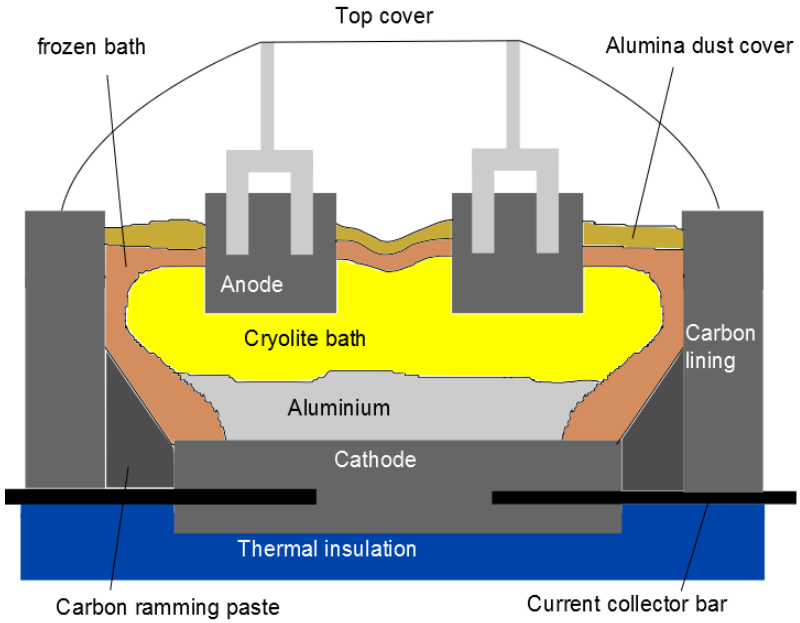

Figure 2.1: Prebake aluminium electrolysis cell

as the fact that there will usually be a gas filled space between the cryolite bath and the frozen bath above.

In the Hall-Hérault process, alumina is added to an electrolytic reduction bath. The main component of this bath is liquid cryolite (Na_3AlF_6), which dissolves the added alumina. To obtain a high production rate, a constant, large *line current*[1] is applied, usually in the range of 100-300 kA. The reduced aluminium will be drawn to the cathode, gathering as a layer between the cathode and the bath. The oxygen will move towards the anodes, reacting with the carbon in the anodes, and is released from the electrolysis cell as CO_2 or CO. The main reaction is

$$2Al_2O_3 + 3C \rightarrow 4Al + 3CO_2. \tag{2.1}$$

2.2.1 Anode design

There have historically been two cell designs used for aluminium production, Söderberg and prebake cells. The difference lies in the choice of anode structure. The Söderberg cells use one large self-baking anode, adding more anode material on top to keep production continuous. In prebake cells, the anodes are produced in a separate production process, before the anodes are installed in the electrolysis cells. This allows better control of anode baking conditions, and makes these conditions independent of the operating

[1] Multiple cells in a so-called potline are connected in series and share a common current

conditions of the aluminium cells. The result is anodes with a higher quality. The prebake cells, depicted in Figure 2.1, usually have between 18 and 32 (Hestetun, 2009) individual anodes, which has to be changed when they are worn down to $1/3$ to $1/4$ of the original size. Prebake cells offer lower emissions, a better quality product and less anode carbon consumption compared to the Söderberg design (Karuppannan and Prabhakar, 2009). All new cells today are made using prebaked anodes. The percentage of aluminium produced with Söderberg anodes has been significantly reduced in later years, from 40 % in 1993 (Grjotheim et al., 1993) to 10 % in 2010 (aluminium institute, 2011). The increased focus on cutting emissions might be an important reason for this drastic change (Karuppannan and Prabhakar, 2009).

2.3 Cell operation

2.3.1 Cell performance

Cell performance is measured using three criteria, namely current efficiency, energy efficiency and emission rate. The current efficiency is a measure of how much aluminium is produced for a certain amount of energy, compared to the theoretically optimal amount. Many factors might affect the current efficiency negatively, but the back reaction

$$2Al + 3CO_2 \rightarrow Al_2O_3 + 3CO \tag{2.2}$$

has been found to be the most significant. Normally, the current efficiency will lie in the range of 90-95 % (Grjotheim et al., 1993) for production cells. Modern cells might even be operated with a current efficiency above 95 %.

The energy efficiency is closely related to the current efficiency, but it takes the cell voltage into account as well. A lot of the energy put into the cell will be released from the cell as heat. That means that economically ideal operation requires both high current efficiency and low cell voltage (less power input as line current is constant). In later years there has also been a focus from the governmental side to reduce the emission rates of pollutants from aluminium electrolysis plants. It was found that vegetation and livestock in areas close to smelters was damaged due to harmful gas and particles from the production. In addition to this the aluminium industry produces a significant amount of greenhouse gases. As a consequence, there has been an increased focus on capturing, cleaning and reducing emissions from smelters in later years.

2.3.2 Heat balance

The aluminium reduction process usually happens at about 940 to $980\,^{\circ}\mathrm{C}$. The melting point of pure cryolite is $1009\,^{\circ}\mathrm{C}$. To lower the melting point, aluminium fluoride (AlF_3) is added to the bath. Liquid cryolite dissolves most of the substances it comes in contact with. Therefore it can't be in direct contact with the side wall of the cell for a long time. A layer of frozen cryolite separates the bath from the walls. To maintain this layer, it is necessary to keep the *superheat* relatively small and positive. The superheat is the difference between the bath temperature and the *liquidus temperature*. The liquidus temperature is the melting temperature of the cryolite bath. If the superheat turns negative for some time,

the whole bath will freeze. This must be avoided at all cost since it requires that the whole electrolysis cell is rebuilt from scratch. If the superheat gets too large, the side ledge will melt, exposing the side walls, which can result in a hole in the cell. This also leads to the need to rebuild the cell from scratch. Luckily, for a superheat in the right region, the side ledge thickness is self-regulating. If the superheat increases, the side ledge thickness will decrease, which will in turn increase the heat loss through the wall. This leads to a new balance with a thinner wall. The opposite will happen when the superheat decreases.

2.3.3 Measurements and control inputs

The line current I_{line} going through a cell is normally kept constant. The amount of power P_{el} put into the cell is then regulated by adjusting the *anode-cathode distance* (ACD) - the distance between the bottom of the anodes and the top of the layer of molten aluminium. Varying this distance will vary the total voltage drop U_{el}, which will alter the amount of power according to $P_{el} = I_{line}U_{el}$. Instead of measuring U_{el} directly, the pseudo-resistance is used instead. The pseudo-resistance R_{pseudo} is defined as

$$R_{pseudo} = \frac{U_{el} - V_{ext}}{I_{line}}, \tag{2.3}$$

where V_{ext} is the zero current intercept of U_{el} versus I_{line} for small changes in current. That means V_{ext} is a certain threshold voltage necessary before any current can pass through the bath at all. The pseudo-resistance measurement is essential in keeping the cell stable, as described in the control section below.

The rest of the common measurements are done manually at a much lower frequency. *Bath temperature* and *acidity* (AlF_3-concentration in the bath) are measured about once a day. These two measurements are important to keep the superheat and bath temperature in the right region. A high bath temperature indicates a high energy input, and should be avoided by reducing the pseudo-resistance. The alumina added to the cell contains a small amount of Na_2O. Some of the AlF_3 in the bath will react with the added Na_2O to create Na_3AlF_6 and Al_2O3 according to

$$3Na_2O + 4AlF_3 \rightarrow 2Na_3AlF_6 + Al_2O_3. \tag{2.4}$$

This will reduce the acidity of the bath. Some aluminium fluoride will also evaporate and escape the cell as carbon fluoride gases or $NaAlF_4$ gases. In total, about 30 kg/day needs to be added to keep the concentration stable. AlF_3 is the main component for adjusting the liquidus temperature, so keeping the acidity at a certain level is important to be able to keep a low bath temperature at the same time as keeping the superheat at a reasonable temperature. The *metal height* and *bath height* are also measured at similar intervals. The metal height measurement is necessary to make sure that the amount of tapped metal correspond to the amount that is produced. The bath height will vary with the side wall thickness, bath production and also with evaporation or spillage of electrolyte. Infrequently it can then be necessary to add or remove some electrolyte to adjust the electrolyte volume. All the manual measurements are taken from a hatch at the short end of the electrolysis cell. Since the cell has a large dimension and the manual measurements are taken at only one point in the cell, the measurement is not necessarily representative for the bath as a whole.

Some modern measurement concepts include the continuous cathode temperature measurement, which gives an indication of the bath temperature dynamics, and on-line bath analysis systems. One such system is called STARprobe™ (Wang et al., 2011). That specific system measures superheat, temperature, alumina concentration and bath ratio[2]. The cost and quality of these measurements decide if they will be used in future cells. All measurement intervals are evaluated based on the cost of the measurement. More frequent manual measurements require more operators, but might lead to an improved control of the cell, which in turn can reduce cost. Finding the right balance here can have a large impact on plant economics.

2.3.4 Normal cell operations

Most modern aluminium plants today have implemented so-called point-feeders that keep the feed rate of alumina steady by releasing small doses of alumina into the electrolytic bath every 1-2 minutes. Similar feeders are also used for AlF_3 in modern cells. As described above, most of the cell measurements are done manually. Other manual actions include anode change, metal tapping and bath tapping. Over time, the carbon anodes will be worn down by being consumed in the formation of CO_2. A new anode will be a lot colder than the electrolysis bath even though it is preheated. This means a layer of frozen bath will form underneath the anode, so that it will be idle for up to a day while the layer melts. If a change is made to one of the anodes closest to the hatch where the bath temperature is measured, this will affect the bath temperature measurement. Cranes are usually used for large cell alterations like anode change. As the top of the anodes are covered by a combination of frozen bath and alumina, the crust above the anode needs to be crushed before it can be removed. This will add chunks of the cover to the bath, which might alter the bath characteristics.

Both metal and bath tapping are done using siphoning[3]. The liquid is poured into a crucible that is managed by the crane or a vehicle like a fork truck with a dedicated crucible. Metal tapping happens regularly to keep the heat balance within a stable region. The tapped amount is weighed continuously to keep track of when to stop tapping. Bath tapping or addition happens more infrequently. It will be done whenever the height of liquid bath is measured to be too high or too low. The main reason for high bath levels is due to a high superheat for some time, leading to a melting of the side walls. Sometimes bath can also be tapped to help fill another cell with a too low bath level.

2.3.5 Unwanted disturbances

The anode effect is a phenomenon that may occur if the concentration of alumina in the bath gets too low. During the anode effect the formation of gases including CO, CF_4 and C_2F_6 will happen underneath the anodes. This will lead to the formation of bubbles that stick to the bottom of the anode, greatly increasing the resistance between the anode and the cathode. It can therefore be detected when the voltage increases above a certain limit. Completely avoiding the anode effect is hard to achieve in practice. The frequency has

[2]The mass ratio of NaF and AlF_3 in the bath (Grjotheim et al., 1993)

[3]Siphon: A tube that carries a liquid from a higher level up and over a barrier and then down to a lower level, with the flow maintained by gravity and atmospheric pressure as long as the tube remains filled.

been reduced a lot in later years. The introduction of point feeders keeps the feed rate of alumina a lot more stable than what was achieved with manual feeding. The reason for keeping the frequency of the anode effect low is because the increase in voltage leads to a significant increase in energy consumption when the anode effect is happening. It can also lead to an increase in bath temperature, which may disturb the heat balance of the cell. As CF_4 and C_2F_6 are very strong greenhouse gases, there is also an environmental motivation for keeping the production rate of these gases as low as possible. The anode effect is handled by increasing the alumina concentration and by getting rid of the bubbles underneath the anodes. Adjusting the alumina concentration can be done by adding a large amount of alumina in a short period of time. Modern cells remove gas bubbles by either moving the anode up and down or tilting it. The manual way of removing the bubbles is by inserting a "green pole" underneath the anodes. This will agitate the electrolyte and remove the gas layer (Hestetun, 2009).

If the alumina concentration becomes too high or too much alumina is added at once, the electrolytic bath might not be able to dissolve all the alumina. This will lead to sludge formation (Hestetun, 2009). The sludge will gather between the liquid metal and the carbon cathode, and will decrease the electrical contact between the cathode and the metal. It takes a long time to remove sludge because it requires that it is in contact with liquid bath. This only occurs along the walls in the intersection between liquid bath and metal. The steady feeding of alumina through point-feeders has reduced the amount of sludge in modern cells quite a lot.

2.3.6 Control

Aluminium electrolysis is a multi-variable process. This means each input might affect multiple measurements, and that each measurement might also be affected by multiple inputs. The fact that most measurements are obtained on a daily basis is also a major challenge when controlling the cell. The result of this is that most basic cell control schemes are mainly based on the continuous pseudo-resistance measurement. A change in the pseudo-resistance measurement can indicate both changes in anode-cathode distance, as discussed in section 2.3.3, and alumina concentration.

The main control objectives of an aluminium electrolysis cell are to keep the production rate stable at a high level as well as making sure the heat balance is in an acceptable zone. The key component in maintaining a high production rate is to keep the alumina concentration as stable as possible to avoid anode effects and sludge formation. The alumina concentration is not measured directly. Instead alumina is overfed and then underfed (with respect to the theoretical consumption rate of alumina) in cycles. The corresponding change in pseudo-resistance indicates if the current alumina concentration is too high or too low. This way the point-feeding system gets an indication of whether to increase or decrease the alumina feed rate to move towards the target concentration.

The main long-term control procedure for adjusting the heat balance is to add aluminium fluoride. Because the acidity is so slowly varying, it is possible to make adjustments to the addition of aluminium fluoride based on the infrequent measurements of bath temperature and acidity. Small short-term adjustments of the energy input can be done by adjusting the anode-cathode distance based on the pseudo-resistance measurements.

Chapter 3

Kalman filter theory

The Kalman filter is named after Rudolf E. Kalman, who published his famous paper describing it in 1960. Norbert Wiener lay the foundation for Kalman's work in the 1940's by trying to isolate the signal of a measurement containing both signal and noise (Brown and Hwang, 2012). Wiener proposed to do this by minimizing the mean-square error of the measurement in the time domain. It was common at the time to use frequency-based separation methods to accomplish this.

The recursive method presented by Kalman in 1960 was based on discrete measurements as opposed to Wiener's continuous time signals. The proposed method was also given in a state space format that worked well with so-called *Multiple-Input-Multiple-Output* (MIMO)-systems (Brown and Hwang, 2012). Rapid changes in the development of computers since the 1960's, combined with the fact that the recursive nature of the Kalman filter was well suited for implementation on a computer, made it a preferred choice for separating signals from noise. Because of this, the Kalman filter has been subject to a lot of research, which has led to a variety of different formulations.

When a mathematical model of a process is created, this will not be a perfect description of the process behaviour. It is however possible to use an estimator, e.g. a Kalman filter, to update the model in order for it to better resemble the process state. Which type of Kalman filter to use for a specific system is based on the system characteristics. There are lots of different ways in which it is possible to categorise a system. Some of these include whether the system is linear or nonlinear, deterministic or stochastic and discrete or continuous. Kalman filters are typically used for stochastic systems which can be both linear, nonlinear, continuous and discrete. The difference between stochastic and deterministic systems is that in a deterministic system, none of the signals are assumed to have a random element (Brown and Hwang, 2012). That means all measurements are perfect and without noise. In stochastic systems, noise is assumed to be present, and the Kalman filter's ability to provide good estimates in the presence of such noise is one of the key reasons why it is still used a lot for state estimation today.

The following discussion will first present the simple Discrete Kalman filter, which is a good introduction to understanding the basic concepts of the Kalman filter theory. The

reader will then be introduced to the *Extended Kalman filter*, which is a version of the Kalman filter used for nonlinear systems. Some attention will also be given to the *multi-rate extended Kalman filter*, which is the specific version of the Extended Kalman filter used in this thesis.

3.1 Discrete Kalman filter (DKF)

In Welch and Bishop (1995), the DKF is presented for the controlled, stochastic linear process. The following discussion is based both on Welch and Bishop (1995) and Brown and Hwang (2012). The process can be described in the following way in state space form:

$$x_{k+1} = A_k x_k + B_k u_k + w_k \tag{3.1}$$

and

$$z_k = H_k x_k + v_k, \tag{3.2}$$

where the subscript k indicates that a vector is given at time t_k, x is the state vector, A_k is the matrix relating the current and next state, B_k is the matrix relating the current inputs u_k to the next state, w_k is the process white noise sequence with known covariance structure, z_k is the measurement vector, H_k is the matrix connecting the measurements to the state and v_k represents the measurement noise which is assumed to be white noise with a known covariance structure. This leads to the definition of the system's covariance matrices Q and R. They are defined as

$$E\left[w_k w_i^T\right] = \begin{cases} Q_k, & i = k \\ 0, & i \neq k \end{cases} \tag{3.3}$$

and

$$E\left[v_k v_i^T\right] = \begin{cases} R_k, & i = k \\ 0, & i \neq k \end{cases}, \tag{3.4}$$

where it is assumed that there is no cross correlation between v and w.

It is important to separate the a priori and a posteriori versions of a vector on a given time step. The a priori version, denoted with a $^-$, is called the *prediction*. It is based on the a posteriori information of the previous time step. The measurements available on the current time step can then be used to improve the prediction, leading to the a posteriori *corrected* vector.

The last two components of the Kalman filter equations are the error covariance matrix P_k and the Kalman gain matrix K_k, where the error covariance matrix has both an a priori and an a posteriori version. The error covariance matrices are calculated from the estimation errors

$$e_k^- = x_k - \hat{x}_k^-,$$

and

$$e_k = x_k - \hat{x}_k$$

as

$$P_k^- = E\left[e_k^- {e_k^-}^T\right] \tag{3.6a}$$

and

$$P_k = E\left[e_k {e_k}^T\right]. \tag{3.6b}$$

The objective of the Kalman gain matrix is to distribute how much weight should be given to the a priori state and the measurement when calculating the a posteriori state and error covariance vectors. It can be seen in (3.7) below that R_k is included in the inverted part of the calculation of the Kalman gain. What this means is that if R_k is small (the measurement is really certain), K_k will be close to H_k^{-1}, meaning that the actual measurement z_k will be trusted more and more compared to the predicted measurement $H_k\hat{x}_k^-$. It must be assumed that H_k is invertible. If that is not the case, some of the states will not be possible to update by using the available measurements. It can also be seen that if P_k^- gets small, which indicates that the prediction is very close to the real state, K_k will also be small. This means the measurement z_k will be trusted less than the predicted state.

With all the important parts of the filter defined, the DKF can then be given as a recursive procedure in 4 steps. With $\hat{x}_0^-$ and P_0^- given, the procedure starts by calculating the Kalman gain

$$K_k = P_k^- H_k^T (H_k P_k^- H_k^T + R_k)^{-1}. \tag{3.7}$$

Then the state prediction is corrected with the measurement as

$$\hat{x}_k = \hat{x}_k^- + K_k(z_k - H_k\hat{x}_k^-). \tag{3.8}$$

Next the a posteriori error covariance matrix is calculated as

$$P_k = (I - K_k H_k)P_k^-, \tag{3.9}$$

before a prediction is done for the next step with the following equations as

$$\hat{x}_{k+1}^- = A_k\hat{x}_k + B_k u_k \tag{3.10}$$

and

$$P_{k+1}^- = A_k P_k A_k^T + Q_k. \tag{3.11}$$

This procedure can then be repeated at every time step. It is important to notice that many of the matrices, including the system matrices A_k and H_k, are denoted by a subscript k, indicating the sampling time. If that is the case, it is also necessary to calculate these matrices on each time step. Various alternative formulations of the Kalman filter have been developed to handle continuous systems and nonlinearity. The Extended Kalman filter is a modification of the discrete Kalman filter to account for nonlinearities in an approximate way. An explanation of this filter is given in section 3.3.

3.2 Kalman filter tuning

To obtain optimal performance with a Kalman filter, it is necessary to tune the system. The system matrices A_k, B_k and H_k should always be chosen in such a way that they resemble the system behaviour in the best possible way. That means the covariance matrices R_k and Q_k are the ones used to tune the Kalman filter.

In theory, it is possible to calculate R_k by running a sequence of test measurements and calculating the variance of the measurements based on the results. The process noise Q_k can often be considered to represent the model uncertainty. While there might sometimes be a reasonable basis for choosing both R_k and Q_k, it is often possible to obtain better performance of the filter by tuning the values (Welch and Bishop, 1995).

It is common to tune R_k and Q_k offline, and then assume that they are constant during operation. Sometimes, like for the aluminium electrolysis process, this can be quite a good approximation because the measurements are done in approximately the same place from time to time, so that the conditions usually doesn't change much. A moving system that collects information from nearby beacons is an example of a system where the measurement variance can change a lot depending on the distance between the system and the beacons (Welch and Bishop, 1995). In such cases, keeping the covariance matrix of the measurement error constant is not a good assumption.

3.3 Extended Kalman filter (EKF)

As for the DKF, the task of the EKF is to estimate the state of a discrete stochastic process. The following presentation is based on Welch and Bishop (1995). The difference between the DKF and the EKF is that the EKF is designed to handle processes that are themselves nonlinear or that have a nonlinear relationship between the measurements and states. The process that is to be estimated can be described as

$$x_{k+1} = f(x_k, u_k, w_k), \tag{3.12}$$

$$z_k = h(x_k, v_k). \tag{3.13}$$

Here, $f(\cdot)$ and $h(\cdot)$ are linear or nonlinear equations. In a multidimensional system, these expressions will be represented as a vector with multiple linear or nonlinear expressions inside. In that case, the vectors have the same dimension as the state vector and the measurement vector respectively. All of the variables are defined in the same way as in section 3.1. The EKF handles the nonlinearities by linearizing about the current mean and covariance (Welch and Bishop, 1995).

The state and measurement equations can be approximated by setting the white noise to 0. The approximate state vector $\tilde{x}_k$ can then be described as

$$\tilde{x}_{k+1} = f(\hat{x}_k, u_k, 0), \tag{3.14}$$

while the measurement is calculated as

$$\tilde{z}_k = h(\hat{x}_k, 0). \tag{3.15}$$

As in section 3.1, $\hat{x}_k$ is the *a posteriori* state vector at time step k. A linearization about these approximations then yields

$$x_{k+1} \approx \tilde{x}_{k+1} + A(x_k - \hat{x}_k) + Ww_k, \tag{3.16}$$

$$z_k \approx \tilde{z}_k + H(x_k - \tilde{x}_k) + Vv_k. \tag{3.17}$$

A, H, V and W are Jacobian matrices where the element on row i and column j is defined as

$$A_{(i,j)} = \frac{\partial f_i}{\partial x_j}(\hat{x}_k, u_k, 0) \tag{3.18}$$

$$W_{(i,j)} = \frac{\partial f_i}{\partial w_j}(\hat{x}_k, u_k, 0) \tag{3.19}$$

$$H_{(i,j)} = \frac{\partial h_i}{\partial x_j}(\tilde{x}_k, 0) \tag{3.20}$$

$$V_{(i,j)} = \frac{\partial h_i}{\partial v_j}(\tilde{x}_k, 0) \tag{3.21}$$

All of these jacobian matrices need to be recalculated for each time step, but the subscript k has been left out here for readability.

If the prediction error is now defined as

$$\tilde{e}_{x_k} \equiv x_k - \tilde{x}_k \tag{3.22}$$

and the measurement residual is defined as

$$\tilde{e}_{z_k} \equiv z_k - \tilde{z}_k, \tag{3.23}$$

it is then possible to calculate the *error process* as

$$\tilde{e}_{x_{k+1}} \approx A(x_k - \hat{x}_k) + \varepsilon_k, \tag{3.24}$$

$$\tilde{e}_{z_k} \approx H\tilde{e}_{x_k} + \eta_k. \tag{3.25}$$

The error process introduces ε_k and η_k. These are new random variables with zero mean and covariance matrices given by WQW^T and VRV^T respectively. Q and R are defined in equations (3.3) and (3.4). It is not possible to find the estimation error directly through (3.22) because x_k is not known. Instead it is possible to utilise the error process to create an estimate $\hat{e}_k$ of the prediction error $\tilde{e}_{x_k}$. The error process is linear and looks a lot like the discrete process stated in (3.1) and (3.2). That means it is possible to create a second Kalman filter, using $\tilde{e}_{z_k}$ as the measurement, to calculate $\hat{e}_k$. With $\hat{e}_k$ calculated, equation (3.22) gives an a posteriori estimate of the state as

$$\hat{x}_k = \tilde{x}_k + \hat{e}_k. \tag{3.26}$$

By assuming that $\tilde{e}_{x_k}$, η_k and ε_k have a normal distribution with a mean value of 0, and letting the predicted value of $\hat{e}_k$ be 0, the Kalman filter equation (3.8) can be used to estimate $\hat{e}_k$ as

$$\hat{e}_k = K_k\tilde{e}_{z_k}. \tag{3.27}$$

K_k is the Kalman gain stated in equation (3.7) just with the small adjustment that R_k is replaced by $V_k R_k V_k^T$. It is then possible to substitute (3.27) into (3.26) and use (3.23) to end up with the equation for calculating the a posteriori state estimate:

$$\hat{x}_k = \tilde{x}_k + K_k(z_k - \tilde{z}_k) \tag{3.28}$$

As with the DKF, it is possible to summarise the EKF as a recursive procedure in four steps. To keep the notation consistent with that of the DKF, the a priori vectors will be denoted as $\hat{x}^-$ instead of $\tilde{x}$. Assuming the initial estimates $\hat{x}_0^-$ and P_0^- have been calculated, the first thing that needs to be calculated is the Kalman gain K_k. It is given by

$$K_k = P_k^- H_k^T (H_k P_k^- H_k^T + V_k R_k V_k^T)^{-1}. \tag{3.29}$$

Then measurement is then used to correct the a priori prediction as

$$\hat{x}_k = \hat{x}_k^- + K_k(z_k - h(\hat{x}_k^-, 0)). \tag{3.30}$$

The a posteriori error covariance matrix is calculated next as

$$P_k = (I - K_k H_k) P_k^-. \tag{3.31}$$

Last, the EKF predicts the state and error covariance matrices for the next time step. These equations are given by

$$\hat{x}_{k+1}^- = f(\hat{x}_k, u_k, 0) \tag{3.32}$$

and

$$P_{k+1}^- = A_k P_k A_k^T + W_k Q_k W_k^T. \tag{3.33}$$

3.4 Multi-rate Extended Kalman filter

It is possible to achieve effective process control if all essential process variables are available frequently. When these variables are not available directly through measurement, estimators can be used as a way to calculate the process variables from available measurements. When using a classical EKF, it is assumed that all measurements are available on each time sample. For many processes, the measurements needed to calculate the essential process variables are however not available this often. One solution can then be to use a multi-rate Extended Kalman filter instead. This version of the Kalman filter looks a lot like the EKF, with the exception that the number of measurements being considered at each time sample will vary. Gudi et al. (1995) describes the concept for a process with measurements coming in at two different frequencies. Here, a time step where both frequent and infrequent measurements are available is known as a *major sampling instant* while a time step with just the frequent measurements is known as a *minor sampling instant*. For a minor sampling instant, the dimension of the measurement vector will decrease. This also affects the dimension of the linearised measurement vector H_k, the measurement noise covariance matrix R_k, the Kalman gain matrix K_k and the vector $h(\hat{x}_k^-, 0)$ connecting the measurements to the a priori version of the states on time step k.

What this means is that while the number of states stay the same on every time step, there will be less measurements available to update each state. This will in other words

affect the *observability* of the system. The observability describes if the internal states can be inferred by the external outputs or measurements. It can be seen from equation (3.30) that when there are less measurements available, the a posteriori state estimate $\hat{x}_k$ will be less affected by the measurements, and some states might even become unobservable on a time step. The unobservable states will then not be corrected by measurements on that time step, and will instead solely become $\hat{x}_k^-$, which is calculated from the process model equations. If the model does not fit perfectly with the real system, this could cause an increasing deviation over time. That is why it is also important to have the major sampling instants. The extra measurements available here will help correct for the system-model mismatch that always occurs in a complex system.

For a complex process like the aluminium electrolysis process, there might be more than two different combinations of measurements that are performed at once. Different measurements can be performed at different rates. For the measurements included in the simple model used for this thesis, the pseudo resistance measurement will be the only measurement available on every time sample. The other measurements are usually performed manually by operators. The rate at which the manual measurements should be taken is a trade-off between the enhanced precision of the states in the Kalman filter gained from having more frequent measurements, and the increased cost in both wage and HSE from employing more operators to perform the extra measurements. The economic gain from increasing the precision of the estimates can be that tighter control of the process may lead to a decreased energy consumption per ton of aluminium produced. The goal of this thesis is not to find the optimal trade-off between these two factors, but it could be an interesting discussion for another time. Instead, some of the properties of the Kalman filter have been investigated. Chapter 4 starts by developing the model needed to implement the Multi-rate Extended Kalman filter. Then, Chapter 5 investigates the goals of the thesis. These are:

- To find out whether it is possible to determine any criteria for how to tune the measurement and process noise parameters of a Kalman filter when the measurement frequency of the system varies.
- To evaluate if infrequent measurements should all be performed seldom at the same time, or if they should be spread out so that small bits of information are provided more often.

Chapter 4

Aluminium process modelling

A simple model of the aluminium electrolysis process has been created. The purpose of the model is to investigate how a Multi-rate Extended Kalman filter should be tuned for a model with characteristics similar to that of an aluminium electrolysis process. The goal is thus not to create a model that covers all aspects of a real cell and it is hence not fitted against real cell data. It is more important to keep the model simple, only including the main aspects of the mass and energy balance of the process. The model description has been separated into four parts. First, the physical aspects of the model is presented in section 4.1. This part covers the partitioning of the system into *control volumes*[1]. Then the mass balance is given in 4.2 before the energy balance is described in 4.3. In 4.4, all utilised system inputs and outputs are given. Then, section 4.5 gives some simulations to describe how the model behaves. A summary of all the assumptions used while creating the model can be found in 4.6. At the end of the chapter, 4.7 gives a conclution on what works well and what has potential for improvement.

4.1 Physical description

The design of a real aluminium electrolysis cell is given in Figure 2.1. When making a mathematical model of this process, there will be a trade-off between simplicity and realism in the design. As the core objective is for this model to be simple, its physical layout has been chosen as simple as possible. The finished design is shown in Figure 4.1. Some inspiration for the model design has been taken from Drengstig (1997), but significant alterations have been done to achieve the final model for this thesis. The model is partitioned into control volumes that resemble most of the main components of a real cell. A system can be partitioned into any number of control volumes. Additional control volumes means both increased complexity and realism. This is because the number of equations increase with each volume, which makes it possible to describe more details.

[1] A spatial representation of a system (which may be part of a larger system) that is separated from the surroundings (another system) by the systems control surface

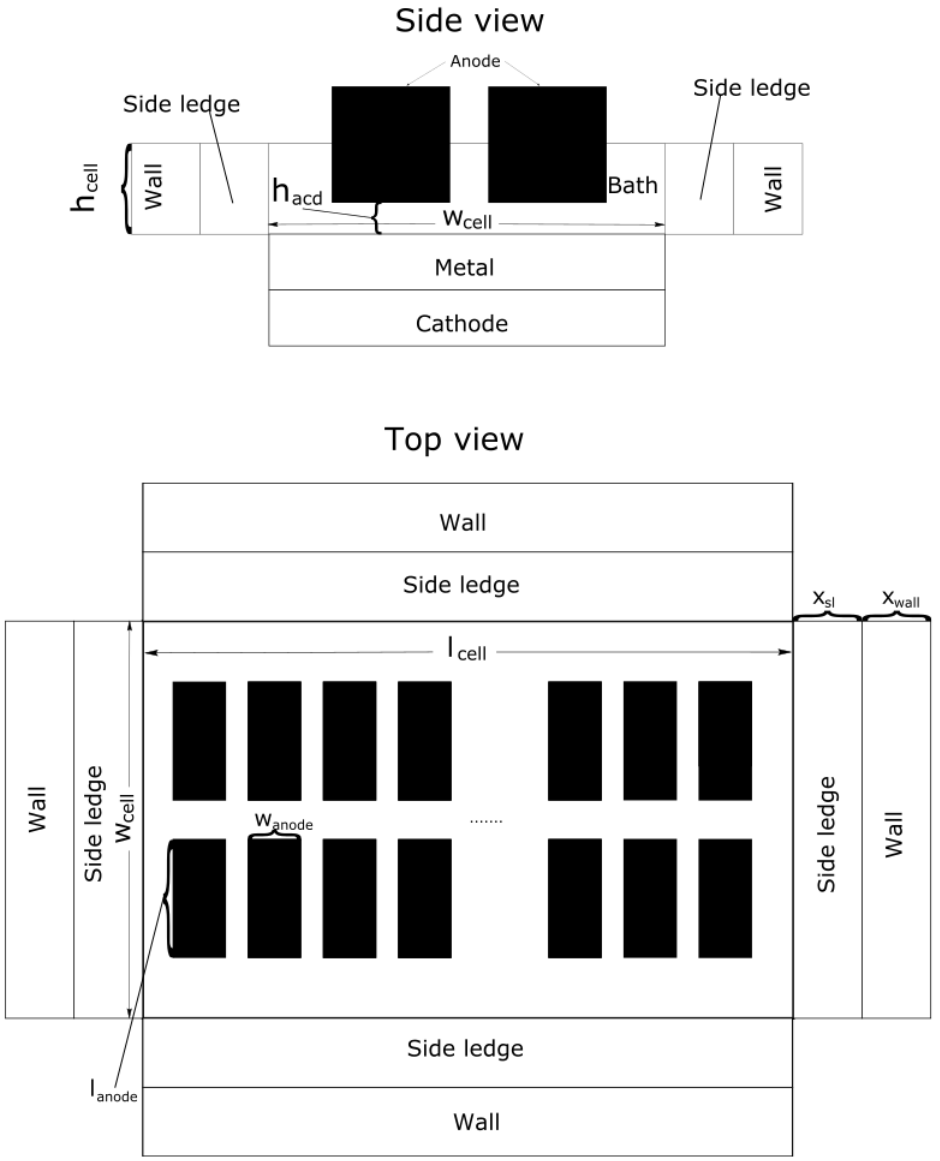

Figure 4.1: Aluminium cell model

Simplicity has been prioritised over realism in the design of this model. That means the chosen partitioning is one control volume for each significantly different part of the cell. The main control volume is the "Bath". This is where power is added to the system and also where the aluminium oxide, aluminium fluoride and bath additions are assumed to be added. Sideways from the bath, there are two control volumes. These are the "Side ledge" and "Wall" control volumes. In reality, there will be 4 physically separate side ledges and walls, one on each side of the bath. It is however assumed that all of the 4 sides behave identically. This way it is possible to describe their behaviour by only 2 control volumes. The anodes are also viewed as one control volume with a common behaviour. Beneath the bath, there will be one control volume filled with liquid aluminium, called the "Metal" control volume in addition to the "Cathode" control volume. There have been done some simplifications on the structure of the different cell components. These simplifications are described below.

The side ledge control volume is assumed to have an even thickness from top to bottom. By assuming this, the side ledge can be approximated as a cuboid[2]. This makes the calculation of the energy balance a lot simpler. The equation for the side ledge thickness is

$$x_{sl} = \frac{m_{sl}}{A_{sl}\rho_{sl}}, \tag{4.1}$$

where m_{sl} [kg] is the total mass of the side ledge, A_{sl} [m^2] is the cross section area of the side ledge and ρ_{sl} [$\mathrm{kg/m}^3$] is the side ledge density. Another simplification is that the side ledge area A_{sl} is assumed to be constant. The side ledge thickness is assumed to vary, but this variation does not affect the calculation of the side ledge area. At the same time, a fixed cell height h_{cell} [m] has been assigned for calculating the side ledge area. This yields that the side ledge area A_{sl} [m^2] can be calculated as

$$A_{sl} = h_{cell}(2l_{cell} + 2w_{cell}), \tag{4.2}$$

where l_{cell} [m] and w_{cell} [m] are the length and width of the cell, as indicated in Figure 4.1. It is also assumed that all sideways heat loss happens from the bath, so that nothing goes from the metal pad. By adding differential equations for the metal and bath height, it would be possible to avoid these simplifications. However, as the purpose of the model is to achieve a simple representation of the main attributes of the aluminium electrolysis process, and not to give a perfect replication of a real cell, it is seen as a reasonable simplification to ignore these dynamics. The "Wall" control volume represents the entire carbon and steel casing around the bath and side ledge.

Modelling the aluminium pad as a cuboid is not quite accurate because in reality the sides grow unevenly, and a closer approximation would be a trapezoid (Because of the carbon ramming paste which can be seen in Figure 2.1). That would however mean that the area between the aluminium pad and bath would change with the metal height, which would add more dynamics and additional calculations. This is not deemed necessary for a model which is only supposed to provide the main dynamics of a cell. The area of the bath (inner dimensions of the cell) seen from above is called A_{top} [m^2] and is calculated as

$$A_{top} = l_{cell}w_{cell}. \tag{4.3}$$

[2]Box with right angles, two and two parallel sides but where not all sides are the same length

This area is used to calculate how much heat is lost through the bottom of the cell in equations (4.70)-(4.73). A change in the side ledge thickness is however not assumed to change A_{top}. This is clearly wrong, but it is a simplification that makes for easier calculations.

On a real cell, the structures above the bath are the anodes and the top cover which is a crust consisting of alumina and bath. The cover is not included as a separate control volume with a separate temperature state. Instead, the heat loss through the cover is modelled as a heat loss directly from the bath. The anodes are included as one large control volume. The horizontal cross section area of the anodes is called A_{anodes} $[\mathrm{m}^2]$. It is calculated as

$$A_{anodes} = l_{anode} w_{anode} N_{anodes}, \tag{4.4}$$

where l_{anode} $[\mathrm{m}]$ and w_{anode} $[\mathrm{m}]$ are the length and width of an anode respectively, and N_{anodes} is the number of anodes in a cell. These lengths are marked in the top view of the cell in Figure 4.1. The fact that the anodes are consumed in a real cell is disregarded here, so that the total mass of anodes in the model is always the same.

To be able to change the energy input to the cell, a differential equation for the anode-cathode distance (ACD) is added to the model. The ACD is marked in Figure 4.1 as h_{acd} $[\mathrm{m}]$. It is assumed in the model that as metal is produced, the anodes will automatically move upwards to keep the ACD constant. That means it is only necessary to change the ACD when a change in the energy input of the cell is required. It can hence be described as

$$\dot{h}_{acd} = \frac{u_{an}}{\Delta t}. \tag{4.5}$$

Here, u_{an} $[\mathrm{m}]$ is the anode movement and Δt $[\mathrm{s}]$ is the time span spent on the movement.

4.2 Mass balance

This section will cover basic theory on mass balances in general and the mass balance of a simple aluminium electrolysis cell model. The theory is collected from Skogestad (2009), while the model is based loosely on Drengstig (1997) and has been developed in collaboration with Cybernetica AS.

4.2.1 Theory of mass balances

A mass balance equation is a specific example of the general concept of *balance equations*. Such equations can be used to look at how the amount of a measurable quantity within a control volume changes. The general equation is then

$$\frac{dB}{dt} = \dot{B}_{in} - \dot{B}_{out} + \dot{B}_{generated} - \dot{B}_{consumed}, \tag{4.6}$$

where B is the amount of the quantity inside the control volume, $\dot{B}_{in}$ and $\dot{B}_{out}$ are the rates at which B is added to and removed from the control volume across the control surface while $\dot{B}_{generated}$ and $\dot{B}_{consumed}$ are the rates at which B is generated or consumed within the control volume. In general, it is only possible to use the balance equation on *extensive*

variables. What characterises an extensive variable is that it is dependant on the size (e.g. mass) of the system. Some examples are mass, energy, impulse and money. One example of an intensive (non-extensive) quantity is volume.

For mass balances, it is possible to write balance equations both for the total mass and for the various components within the volume. For a chemical process control volume, the typical mass balance for a component i, given in $[\mathrm{kg/s}]$, can consist of the time derivative $\frac{dm_i}{dt}$ of the mass m_i, a mass flow rate $w_{in,i}$ in, a mass flow rate $w_{out,i}$ out and the rate $r_{i,j}$ at which the component is generated or consumed in reaction j. The resulting equation is then

$$\frac{dm_i}{dt} = \dot{m}_i = w_{in,i} - w_{out,i} + \sum_{j=1}^{n_j} r_{i,j}. \tag{4.7}$$

The total mass of a control volume has the same form as (4.7), but without the reaction rate $r_{i,j}$ because the total mass of a control volume is a conserved quantity.

4.2.2 Aluminium electrolysis cell mass balance

There are a lot of different chemical compounds and elements present in an aluminium electrolysis cell. Some of these have been found to be of more significance to the cell dynamics than others. For the given model, the dynamics of four of the substances are considered. These are:

- Aluminium (Al), main product of the process
- Aluminium oxide/alumina (Al_2O_3), raw material
- Cryolite (Na_3AlF_6), solvent
- Aluminium fluoride (AlF_3), main component in lowering the liquidus temperature

There are a lot of other elements and compounds that are usually present to some extent in the process as well. Sodium oxide accounts for a small fraction of the raw material (alumina feed) that is added to the bath. This fraction is assumed to be constant, which means that it is not necessary to keep track of the sodium oxide concentration explicitly. It can be calculated directly from the alumina feed rate. There will be some other compounds present in a cell than the ones in the list above. Their concentrations are assumed to be negligible or constant because their impact on the mass and energy balance are too small to be considered in a simple model like this. The changing of anodes has a major impact on the energy balance of a real cell. This effect is not included in the model, as such a large disturbance would require a much more advanced model. It would also complicate the tuning of the Multi-rate Extended Kalman filter, which is the actual task of this thesis. Because anode consumption is not modelled, it is also not necessary to keep track of the production rate of CO and CO_2 in the cell. The main effect these compounds have in the cell is that the CO_2 will react with Al to form Al_2O_3 and CO in the back reaction. This is the reaction which leads to the largest reduction in current efficiency. To simulate the effect this reaction has on the cell, a constant reduction in the current efficiency has been added, so that the current efficiency is not at 100 %. By doing this, there is no need to explicitly give the CO contents in the system. The CO_2 that is not part of the back

reaction will mostly leave the cell as gas. In a real cell, gas bubbles might gather under the anodes and start an anode effect if the Al_2O_3-concentration goes to low. The anode effect is neglected in this process model, so it is not necessary for the model to know how much of the produced CO_2 that is still in the system.

One major simplification that has been applied to keep the complexity of the model down is to disregard sludge formation. Sludge formation is discussed in the section about unwanted disturbances in Chapter 2. It has a significant impact on the cell dynamics as it changes the resistance between the anode and cathode. As it is both hard to model and measure, it has however been ignored here. Another simplification is that all added alumina is assumed to be dissolved in the bath at once and that all sodium oxide reacts the moment it is added to the cell. This simplifies the mass balance of the cell, as it is not necessary to keep track of the concentration of these two solids.

Three of the control volumes will have a change of mass during normal operation of the cell. The bath will be fed with alumina (u_{ox}) and aluminium fluoride (u_{fl}) frequently. The added alumina will be used to produce aluminium through the main reaction in equation (2.1). The reaction rate r_{ox} $[\mathrm{mol/s}]$ of this reaction is given by Faraday's laws of electrolysis, given on page 108 in Drengstig (1997), which states that

$$r_{ox} = \frac{CE \cdot I_{line}}{F \cdot z}. \tag{4.8}$$

Here, CE is the current efficiency, a constant between zero and one which is added to simulate the energy loss of the back reaction (2.2). I_{line} $[\mathrm{A}]$ is the line current, $F = 96486.7$ $[(\mathrm{A\,s})/\mathrm{mol}]$ is Faraday's constant and $z = 12$ is the number of electrons involved in the reaction.

Infrequently, the bath level will be too low or to high. If that is the case, it is necessary to add or remove some bath. These actions are denoted $u_{bath\,add}$ $[\mathrm{kg/s}]$ and $u_{bath\,tap}$ $[\mathrm{kg/s}]$ respectively. It is assumed that the bath that is added has a constant composition, and that it consists of cryolite, alumina aluminium fluoride and small amounts of other chemical species. The concentration of Al_2O_3 and AlF_3 in the bath additions are called $\left[Al_2O_3\right]_{additives}$ and $\left[AlF_3\right]_{additives}$. These concentrations may differ from the concentration of the same components in the bath. The *rate* at which bath addition and bath tapping is performed is defined with one input variable u_{bath} $[\mathrm{kg/s}]$. They are defined as

$$u_{bath\,add} = max(0, u_{bath}); \tag{4.9}$$

and

$$u_{bath\,tap} = max(0, -u_{bath}) \tag{4.10}$$

in the model. That way, at least one of them must be zero at any time, so that it is not possible to tap and add bath at once, which is reasonable.

As mentioned above, some Na_2O will be present in the aluminium oxide feed. This mass fraction is assumed to be small and constant, and is denoted $\left[Na_2O\right]_{ox-feed}$. It is assumed that the whole amount of sodium oxide will react with aluminium fluoride and create bath and alumina as shown in equation (2.4). The rate of this reaction in $[\mathrm{kmol/s}]$ can be described as

$$r_{bath\,prod} = \frac{\left[Na_2O\right]_{ox-feed} \cdot u_{ox}}{3M_{Na_2O}}, \tag{4.11}$$

where M_{Na_2O} [g/mol] is the molar mass of sodium oxide.

The distribution of substances in the cell is simplified compared to a real cell. The bath is assumed to have a varying amount of cryolite, alumina and aluminium fluoride. In addition, it is assumed that a constant percentage of the bath mass, about 6 % in total, is given by a combination of other chemical species. That means the total bath mass m_{bath} [kg] can be calculated as

$$m_{bath} = (m_{fl} + m_{ox} + m_{cry})/0.94, \tag{4.12}$$

while the mass fractions of aluminium oxide and aluminium fluoride are given by

$$\left[Al_2O_3\right] = \frac{m_{ox}}{m_{bath}} \tag{4.13}$$

and

$$\left[AlF_3\right] = \frac{m_{fl}}{m_{bath}}. \tag{4.14}$$

To control the alumina concentration in the bath, a simple P-controller has been added to the model. This controller is given by

$$u_{ox} = max(0, K_p(\left[Al_2O_3\right]_{ref} - \left[Al_2O_3\right])), \tag{4.15}$$

where $\left[Al_2O_3\right]_{ref}$ is the desired alumina concentration, and more alumina is added as soon as the concentration goes beneath this. K_p is the proportional gain of the controller. As it is a P-controller, there will be some stationary bias. The precise concentration of alumina is however not important when testing the Kalman filter properties, so this deviation is accepted. In a real cell, such a simple controller would not be used because the alumina concentration is not available continuously. That means over- and under-feeding periods (Hestetun, 2009) are used to determine the alumina concentration, as described in the control part of section 2.3.6. As controlling the alumina concentration of the cell is not essential in the investigation of the Kalman filter performance, and the fact that controlling the alumina concentration through resistance measurements is not straight forward, it was found reasonable to make the simplification that the alumina concentration is already stabilised, simulating this as a P-controller.

The side ledge is assumed to consist of solely cryolite in solid state, so that the cryolite in the cell moves between these two control volumes as the heat balance changes in the cell. The net cryolite mass flow between the bath and side ledge will be the difference between what freezes onto the side ledge and what melts from it. These mass flows can be calculated as

$$w_{melt} = max(0, (Q_{bath-liq} - Q_{liq-sl})/\Delta H_{cry}) \tag{4.16}$$

and

$$w_{freeze} = max(0, (Q_{liq-sl} - Q_{bath-liq})/\Delta H_{cry}). \tag{4.17}$$

Here, the max functions are added to prevent that the melting and freezing mass flow is active at once. $Q_{bath-liq}$ [W] and Q_{liq-sl} [W] are the convective heat loss from bath to the surface of the side ledge and the conductive heat loss from the surface of the side ledge to the center of it respectively. These are described in more detail in section 4.3. ΔH_{cry}

$[\mathrm{J/kg}]$ is the amount of energy required to melt one kilogram of cryolite from the side ledge.

It is assumed that all pure aluminium in the cell is present in the metal pad beneath the bath. In a real cell, there will be some aluminium in the bath that has not moved to the metal pad yet, but zero transport delay is a good enough assumption here. The "Metal" control volume is assumed to only consist of aluminium.

There will be differential equations describing the dynamics of all the important masses mentioned above. These are aluminium fluoride (m_{fl}), alumina (m_{ox}) and cryolite (m_{cry}) in the bath, cryolite (m_{sl}) in the side ledge and aluminium (m_{met}) in the metal pad.

The differential equation for the mass of aluminium oxide $\dot{m}_{ox}$ $[\mathrm{kg/s}]$ in the bath is given by

$$\begin{aligned}\dot{m}_{ox} = (1 - \left[Na_2O\right]_{ox-feed})u_{ox} - \left[Al_2O_3\right]u_{bath\,tap} \\ + \left[Al_2O_3\right]_{additives}u_{bath\,add} - \frac{2}{1000}r_{ox}M_{Al_2O_3} + r_{bath\,prod}M_{Al_2O_3},\end{aligned} \tag{4.18}$$

where $M_{Al_2O_3}$ $[\mathrm{g/mol}]$ is the molar mass of alumina and all other notation is given in the discussion above. The first part of the differential equation describes the alumina feed contribution. The second part describes the alumina that is removed when bath is tapped, while the third part is the alumina fraction of added bath. The two last parts describe the consumption and production of alumina through equations (2.1) and (2.4) respectively.

The change of aluminium fluoride mass $\dot{m}_{ox}$ $[\mathrm{kg/s}]$ in the bath is given by

$$\begin{aligned}\dot{m}_{fl} = u_{fl} - \left[AlF_3\right]u_{bath\,tap} + \left[AlF_3\right]_{additives}u_{bath\,add} \\ - 4r_{bath\,prod}M_{AlF_3},\end{aligned} \tag{4.19}$$

where M_{AlF_3} is the molar mass of aluminium fluoride. As the equation shows, additional AlF_3 is added through the fluoride feed and bath additions. On the other hand, it is removed through bath tapping and through the bath production equation (2.4).

The third component present in the bath is liquid cryolite. The differential equation describing how $\dot{m}_{cry}$ $[\mathrm{kg/s}]$ changes is

$$\begin{aligned}\dot{m}_{cry} = w_{melt} - w_{freeze} - (0.94 - \left[Al_2O_3\right] - \left[AlF_3\right])u_{bath\,tap} \\ + (0.94 - \left[Al_2O_3\right]_{additives} - \left[AlF_3\right]_{additives})u_{bath\,add} \\ + 2r_{bath\,prod}M_{Na_3AlF_6}.\end{aligned} \tag{4.20}$$

Here, $M_{Na_3AlF_6}$ is the molar mass of cryolite. The mass transfer between side ledge and bath is described by the two first terms of the differential equation¡. In addition, the bath addition and tapping contributions are included as well as the production of cryolite bath through equation (2.4). The concentration of cryolite is calculated with $0.94 - \left[Al_2O_3\right] - \left[AlF_3\right]$ because the remaining 6 % are assumed to be a constant amount as described in (4.12). The same applies for bath addition.

The differential equations for the mass of side ledge $\dot{m}_{sl}$ $[\mathrm{kg/s}]$ and metal $\dot{m}_{met}$ $[\mathrm{kg/s}]$ are really simple as they are assumed to consist of one component only. The change in side ledge mass is simply given by

$$\dot{m}_{sl} = w_{freeze} - w_{melt}, \tag{4.21}$$

while the change in metal mass is given by

$$\dot{m}_{met} = \frac{2}{1000} r_{ox} M_{Al} - c_{tap} u_{met\,tap}. \tag{4.22}$$

M_{Al} $[\mathrm{g/mol}]$ is here the molar mass of aluminium and c_{tap} is a parameter close to 1 that is used to correct for the fact that there is usually some bias in the weight used to weigh the tapped metal. Estimating this parameter will then correct for this bias. The equation for metal mass suggests a continuous metal production, while it will be tapped on regular, discrete points in time.

4.3 Energy balance

The following section describes some basic theory about energy balances, conduction and convection. In addition, it describes a simple model for the energy balance of an aluminium electrolysis cell. The theory is based on Skogestad (2009) and Jessen (2008), while the energy balance for the aluminium electrolysis cell is based on these principles and on the control volume partitioning of section 4.1.

4.3.1 Energy balance theory

The most general form of an energy balance for a system (control volume) is

$$\dot{E} = \dot{E}_{in} - \dot{E}_{out} + \dot{Q} + \dot{W}. \tag{4.23}$$

It states that the change in total energy $\dot{E}$ [W] can be given as the sum of total energy going in and out of the system through mass flows ($\dot{E}_{in}$ and $\dot{E}_{out}$), energy transferred as heat $\dot{Q}$ through the control surface of the system, and work $\dot{W}$ added to the system from the environment.

When presenting the energy balance of a system, it can often be reasonable to express it as differential equations of temperature, which is also known as temperature explicit form. The following derives the temperature explicit form of the energy balance for one component i in a control volume.

First, it is possible to approximate $E_i = U_i$, where U_i is the internal energy of the component and E_i is the total energy of it. By doing this approximation, the main contributions that are disregarded are potential and kinetic energy. For most processes, these are negligible compared to the internal energy, so that it is reasonable to simplify the energy balance (4.23) to

$$\dot{U}_i = \dot{U}_{in,\,i} - \dot{U}_{out,\,i} + \dot{Q}_i + \dot{W}_i. \tag{4.24}$$

It is further possible to separate work into

$$W_i = W_{flow,\,i} + W_{n,\,i}, \tag{4.25}$$

where $W_{flow,\,i}$ is the work done on the component by the mass flows, while $W_{n\,i}$ is all other kinds of work. The enthalpy of a flow is defined as

$$H = U + W_{flow} = U + pV. \tag{4.26}$$

Here, p is the pressure of the flow and V is the volume of it. That means it is possible to express (4.24) as

$$\begin{aligned} \dot{U}_i &= \dot{H}_{in,\,i} - \dot{H}_{out,\,i} + \dot{Q}_i + \dot{W}_{n,\,i} \\ &= w_{in,\,i}\hat{H}_{in,\,i} - w_{out,\,i}\hat{H}_{out,\,i} + \dot{Q}_i + \dot{W}_{n,\,i}, \end{aligned} \tag{4.27}$$

where $w_{in,\,i}$ [kg/s] and $w_{out,\,i}$ [kg/s] are the mass flows of the component in and out of the system and $\hat{H}_{in,\,i}$ [J/kg] and $\hat{H}_{out,\,i}$ [J/kg] are the specific enthalpies of these flows. In the same way, U_i can be defined as

$$U_i = \hat{U}_i m_i, \tag{4.28}$$

with $\hat{U}_i$ [J/kg] as the specific internal energy and m_i [kg] as the mass of the component. That means

$$\frac{d(\hat{U}_i m_i)}{dt} = m_i \frac{d\hat{U}_i}{dt} + \hat{U}_i \frac{dm_i}{dt} = w_{in,\,i}\hat{H}_{in,\,i} - w_{out,\,i}\hat{H}_{out,\,i} + \dot{Q}_i + \dot{W}_{n,\,i}. \tag{4.29}$$

If working with solids and liquids, as is the case for the aluminium electrolysis model derived here, it is now possible to do the approximation

$$\hat{H}_i \approx \hat{U}_i \tag{4.30}$$

for the left hand side of (4.27). To justify this,

$$\hat{H}_i = \hat{U}_i + \frac{p}{\rho_i} \tag{4.31}$$

from Sælid (1984) is used. Here, ρ_i [kg/m^3] is the density of component i while p [N/m^2] is the pressure in the system. This equation is the equivalent of (4.26) per unit mass, so that the equation is given in [J/kg]. By comparing the contribution in energy increase from the pressure and enthalpy, it can be established that the pressure does contribute next to nothing in energy increase for the system in question. This comparison is done for cryolite, which is the main component in the electrolysis bath when producing aluminium. The heat capacity is obtained from Anovitz et al. (1987). It is set to $C_p = 394.7$ [J/(°C mol)] for liquid cryolite. The molar mass of cryolite is $M_{Na_3AlF_6} = 209.9$ [g/mol] (bibliography: molar mass cryolite, 2017), while the density of cryolite in [g/cm^3] at $T = 1000\,°\text{C}$ has been found from Thonstad et al. (2001) to be approximately

$$\rho_{Na_3AlF_6} = 3.032 - 0.937 \cdot 10^{-3}T = 2.095\,[\text{g/cm}^3] = 2095\,[\text{kg/m}^3]. \tag{4.32}$$

The process is not closed, so the pressure is always 1 [atm] = 101325 [Pa]. That means pressure contributes with

$$\frac{p}{\rho_{Na_3AlF_6}} = \frac{101325\,[\text{N/m}^2]}{2095\,[\text{kg/m}^3]} = 48\,[\text{J/kg}]. \tag{4.33}$$

The enthalpy increase related to a temperature increase ΔT of one degree for one kg of cryolite under constant pressure can be found by

$$C_p = \left(\frac{\partial H}{\partial T}\right)_p. \tag{4.34}$$

If the heat capacity C_p [J/°C] is assumed constant here, this gives

$$\Delta H = \frac{C_p \Delta T}{M_{Na_3AlF_6}} = \frac{394.7}{0.2099} = 1880.4\,[\text{J/kg}]. \tag{4.35}$$

That means that the pressure accounts for less than 3 % of the change in internal energy for a temperature increase of 1 °C. The contribution will be even smaller for larger temperature changes. Because of this, (4.30) is a good approximation.

Using this approximation, it can be found that

$$\frac{d\hat{U}}{dt} \approx \frac{d\hat{H}}{dt} = \frac{\partial \hat{H}}{\partial T}\frac{dT}{dt} = c_p \frac{dT}{dt}, \tag{4.36}$$

where the specific form of equation (4.34) has been used. c_p $[\mathrm{J}/(\mathrm{kg}\,^\circ\mathrm{C})]$ is hence the specific heat capacity of the system. Putting (4.36) into (4.29) and moving $\hat{U}_i \frac{dm_i}{dt}$ to the right side of the equation then leads to

$$\frac{dT_i}{dt} = \dot{T}_i = \frac{1}{m_i \cdot c_{p,\,i}}\left(w_{in,\,i}\hat{H}_{in,\,i} - w_{out,\,i}\hat{H}_{out,\,i} + \dot{Q}_i + \dot{W}_{n,\,i} - \hat{U}_i\frac{dm_i}{dt}\right). \tag{4.37}$$

Now, the mass differential equation (4.7) is inserted into (4.37). That leads to

$$\begin{aligned}\dot{T}_i = &\frac{w_{in,\,i}\hat{H}_{in,\,i} - w_{out,\,i}\hat{H}_{out,\,i} + \dot{Q}_i}{m_i \cdot c_{p,\,i}} \\ &+ \frac{\dot{W}_{n,\,i} - \hat{U}_i(w_{in,\,i} - w_{out,\,i} + \sum_{j=1}^{n_j} r_{j,\,i})}{m_i \cdot c_{p,\,i}}.\end{aligned} \tag{4.38}$$

If the contents of the control volume is assumed to be perfectly mixed, and (4.30) is applied, then the specific enthalpy of the out flow from the control volume has the same composition as the contents of the control volume, and hence

$$\hat{U}_i \approx \hat{H}_{out,\,i} \tag{4.39}$$

leads to

$$\begin{aligned}\dot{T}_i = &\frac{w_{in,\,i}(\hat{H}_{in,\,i} - \hat{H}_{out,\,i}) + \dot{Q}_i}{m_i \cdot c_{p,\,i}} \\ &+ \frac{\dot{W}_{n,\,i} - \hat{H}_{out,\,i}\sum_{j=1}^{n_j} r_{i,j}}{m_i \cdot c_{p,\,i}}.\end{aligned} \tag{4.40}$$

This is the temperature specific energy equation for one component i in a control volume. The reaction term says that the reaction contribution will depend on the state of the contents in the control volume. $\hat{H}_{out,\,i}$ is assumed to include both a temperature change contribution of heating a component and a reaction enthalpy or state change enthalpy that will be valid for a energy change due to a reaction. For a control volume of multiple components, things are a bit more complex if no further simplifications are done. It is reasonable to say that there will be a common temperature for all the components. At the same time, all components can have a contribution to the heat loss and work terms. Then it is reasonable to sum the contribution from each component and look at e.g. a common heat loss from one control volume to the next. The reaction terms and mass flows might get more complex when trying to join terms from different components. Because of this it might be reasonable to only simplify these as seem fit in each separate case. The total enthalpy of a mixture of components will be different than for the same amounts of separate components. This can e.g. affect the heat capacity of the mixture. This mixture enthalpy has been left out of this simple model because it adds a level of complexity that is not necessary when the goal is to provide a model with characteristics that are only similar to that of a real cell, instead of providing a precise representation of it. That means the total heat capacity of a mix of components is given as the sum of the products between the mass fraction and the constant specific heat capacity of each component.

4.3.2 Mechanisms of heat transfer

Heat can be transferred from one physical system to another in many different ways. In the aluminium electrolysis process, conduction and convection are the main principles of heat transfer. A short introduction will therefore be given of these here, based on Jessen (2008).

Conduction is the transfer of heat through a material. The mathematical expression for conduction is given by Fourier's law, which states that the heat flow Q [W] is given by

$$Q = -kA\frac{\partial T}{\partial x}. \tag{4.41}$$

This equation says that the conductive heat flow through a body is proportional to the temperature gradient $\frac{\partial T}{\partial x}$ [°C/m] in the direction of heat flow and to the cross section area A [m^2] perpendicular to the direction of heat flow. The proportionality constant k [W/(m °C)] describes the conductive properties of the material.

For the simple case of one directional heat flow through a plane wall with a fixed cross section area, the steady state heat flow through a wall with thickness x [m] is described by

$$Q = kA\frac{T_1 - T_2}{x}, \tag{4.42}$$

where $T_1 > T_2$.

Convection is the transfer of heat through the motion of a fluid. An example can be that heat will be transferred when a fluid flows over a surface, given that their temperature differ. The equation for convection is similar to that of conduction. It states that the heat transfer Q [W] is given by

$$Q = hA(T_s - T_f). \tag{4.43}$$

Here, A [m^2] is the contact surface between the surface and liquid, h [W/(m^2 °C)] is the proportionality constant which is specific for the combination of materials, T_s °C is the surface temperature and T_f °C is the fluid temperature.

For a system with multiple layers, where each layer have different heat conductive properties and steady state is assumed, it is possible to sum the heat loss of multiple layers to look at the total heat loss through those layers expressed by only the temperatures on each end. This can also be combined with convective heat transfer. The following example with two walls, which is depicted in Figure 4.2, illustrates the concept. In this system, the temperatures T_1, T_2 and T_3 are the surface temperatures of the left side, connection between the different sides and right side respectively. The heat transfer through the left section of the wall is

$$Q_l = k_1 A\frac{T_1 - T_2}{x_1} \tag{4.44}$$

while the heat transfer in the right section is

$$Q_r = k_2 A\frac{T_2 - T_3}{x_2}. \tag{4.45}$$

Because the system is assumed to be in steady state, Q is the same through the whole wall, so that $Q_l = Q_r$. That means

$$k_1 A\frac{T_1 - T_2}{x_1} = k_2 A\frac{T_2 - T_3}{x_2}, \tag{4.46}$$

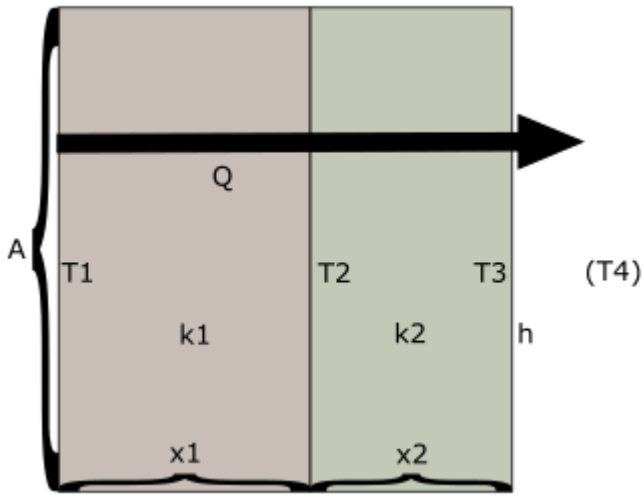

Figure 4.2: Heat conduction through a wall consisting of two different materials

which can be solved for T_2 to give

$$T_2 = \frac{k_1 x_2 T_1 + k_2 x_1 T_3}{k_2 x_1 + k_1 x_2}. \tag{4.47}$$

By putting this expression for T_2 into (4.44) and simplifying, the heat conduction from T_1 to T_3 is given by

$$Q = \frac{T_1 - T_3}{\frac{x_1}{k_1 A} + \frac{x_2}{k_2 A}} \tag{4.48}$$

For N transitions and $N + 1$ temperatures, the resulting heat transfer becomes

$$Q = \frac{T_1 - T_N + 1}{\sum_{i=1}^{N} R_i}, \tag{4.49}$$

where

$$R_i = \begin{cases} \frac{x_i}{k_i A}, & \text{conduction} \\ \frac{1}{h_i A}, & \text{convection.} \end{cases} \tag{4.50}$$

If convection should be added to the above example, it could be in the form of convection between the wall, with a temperature T_3, and the surrounding air, with a temperature T_4. The resulting heat transfer from T_1 to T_4 would then be

$$Q = \frac{T_1 - T_4}{\frac{x_1}{k_1 A} + \frac{x_2}{k_2 A} + \frac{1}{hA}} \tag{4.51}$$

This expression is obtained by assuming that both the convective and conductive heat flow is equal for all parts of the equation.

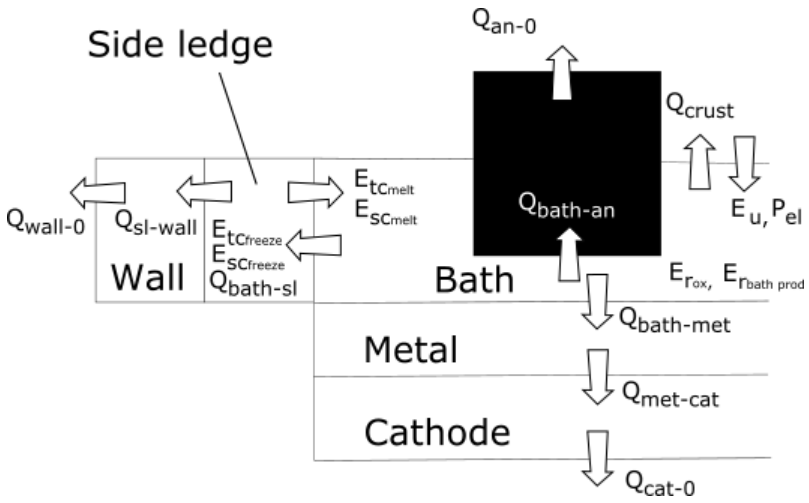

Figure 4.3: Aluminium cell energy model

4.3.3 Aluminium electrolysis cell energy balance

The energy model of the aluminium electrolysis cell uses the same control volume partitioning as the mass balance. It has been assumed that the substances within a control volume is perfectly mixed, so that the properties of the material are the same throughout it. Another assumption is that there is a uniform temperature in the electrolysis bath. The sideways heat loss is further assumed to be piecewise stationary and one directional. Piecewise stationary means that the heat loss is assumed to be the same from the middle of one control volume to the middle of the next one at any point in time. It is however not necessarily the same through all the control volumes at once. This means that there will be separate energy balances for the side ledge and wall.

Figure 4.3 summarises the equations that will be described below. Instead of having separate equations for each of the 4 cell walls, the sideways area has instead been summed up to one area called A_{sl}. As with the mass, it is assumed that the bath is the only place where energy is put into the cell. The sideways and downwards heat losses are calculated through conduction and convection. The heat loss upwards is a combination of conduction and convection through the anodes and conduction, convection and radiation through the crust. The heat loss through the crust is modelled as a linearised function. The reason for this is given in the discussion of equation (4.73) below.

First, the energy exchange in the bath is presented, before moving on to the sideways heat loss. Then, the downwards and upwards heat loss are presented before everything is summarised in differential equations at the end.

To generate the heat needed to keep the bath liquefied, a large current is applied to the cell. This current will be kept constant, so that the power put into the cell can be regulated by rising or lowering the anodes. This will adjust the resistance in the gap between the anodes and the cathode. This so-called pseudo-resistance R_{pseudo} [Ω], which is described in equation 2.3, can be calculated from the resistivity equation (Sears et al., 1987), which

gives

$$R_{pseudo} = \frac{h_{acd} \cdot \rho_{bath}}{A_{anodes}}. \tag{4.52}$$

A_{anodes} and h_{acd} are calculated from (4.4) and (4.5) respectively. ρ_{bath} $[\Omega/\mathrm{m}]$ is the specific resistance of the bath. It is assumed to be constant, and the calculation can be found in Appendix C.

The total voltage drop U_{el} [V] of the cell can then be calculated from equation (2.3). This equation has been modified with a voltage loss U_{loss} to simulate the fact that some of the voltage will be lost other places than in the bath. The modified version of the equation is

$$U_{el} = R_{pseudo} I_{line} + V_{ext} + U_{loss}. \tag{4.53}$$

P_{el} [W] describes the total energy input of the cell. It is given by

$$P_{el} = U_{el} I_{line}. \tag{4.54}$$

In addition to the electrical energy put into the cell, there are a few other energy contributions happening directly in the bath. The primary reaction is endothermic. The heat $E_{r_{ox}}$ [W] consumed by this reaction is

$$E_{r_{ox}} = r_{ox} \Delta H_{r_{ox}}, \tag{4.55}$$

where r_{ox} is given by (4.8) and ΔH_r $[\mathrm{J/mol}]$ is the heat of the primary reaction.

The bath production reaction is exothermic. This reaction produces cryolite from the Na_2O in the alumina feed according to the reaction described in equation (2.4). This reaction will add the heat $E_{r_{bath\,prod}}$ [W] to the bath. It is described by

$$E_{r_{bath\,prod}} = 1000 r_{bath\,prod} \Delta H_{r_{bath\,prod}}, \tag{4.56}$$

where $\Delta H_{r_{bath\,prod}}$ $[\mathrm{J/mol}]$ is a negative constant describing the heat of the bath production reaction and $r_{bath\,prod}$ $[\mathrm{kmol/s}]$ is given by equation (4.11).

When alumina, aluminium fluoride and frozen bath are added to the liquid bath, these will have a temperature that is lower than the bath. Because of this, energy is needed both to increase the temperature of the substances and to melt them. In general, the energy needed to heat a solid substance of mass m_i from T_1 to a liquid with temperature T_2 can be described by

$$Q = m_i(c_{p,\,i}(T_2 - T_1) + \Delta_{fus} H_i), \tag{4.57}$$

where $c_{p,\,i}$ $[\mathrm{J/(^\circ C\,kg)}]$ is the specific heat capacity of the compound and $\Delta_{fus} H_i$ $[\mathrm{J/kg}]$ is the *heat of fusion*[3]. The heat capacity is here assumed to be constant, and its value is calculated as the average heat capacity of the solid substance from a typical T_1 to a typical T_2. The energy needed to heat and melt all these substances has been gathered into E_u [W] which is written as

$$\begin{aligned} E_u = {} & (\bar{c}_{p_{Al_2O_3}} u_{ox} + \bar{c}_{p_{AlF_3}} u_{fl} + c_{p_{bath,s}} u_{bath\,add})(T_{bath} - T_{in}) \\ & + \Delta_{fus} H_{Al_2O_3} u_{ox} + \Delta_{fus} H_{AlF_3} u_{fl} + \Delta_{fus} H_{bath} u_{bath\,add}. \end{aligned} \tag{4.58}$$

[3]The amount of energy required to melt one kilogram of a substance

The Na_2O-fraction of the alumina feed will be ignored here because it is so small. For the bath additions, the heat capacity and heat of fusion will be assumed constant, and they will be approximations because it it is hard to obtain precise values for these. $c_{p_{bath,s}}$ is the specific heat capacity of solid bath. The temperature T_{in} [°C] is assumed to be constant and equal for all the additions. Calculating the average heat capacity of alumina and aluminium fluoride has been done with

$$\bar{c}_{p_{Al_2O_3}} = \frac{\int_{T_{in}}^{960} c_{p_{Al_2O_3}}(\tau)d\tau}{960 - T_{in}} \tag{4.59}$$

and

$$\bar{c}_{p_{AlF_3}} = \frac{\int_{T_{in}}^{960} c_{p_{AlF_3}}(\tau)d\tau}{960 - T_{in}}. \tag{4.60}$$

Here, 960 °C was considered a normal bath temperature while $c_{p_{Al_2O_3}}(\tau)$ and $c_{p_{AlF_3}}(\tau)$ were found in bibliography: heat capacity alumina (2017) and bibliography: heat capacity aluminium fluoride (2017) respectively.

The **sideways energy balance** can be partitioned into two parts. One is the conductive and convective heat loss from the cell while the other is the energy being transferred from the bath to the side ledge and the other way when cryolite melts from the side ledge or freezes onto it. The heat transfer from the bath to the center of the side ledge is modelled in two steps - the convective heat loss $Q_{bath-liq}$ [W] from the bath to the surface of the side ledge and the conductive, one-dimensional heat loss Q_{liq-sl} [W] from this surface to the center of the solid side ledge. $Q_{bath-liq}$ and Q_{liq-sl} are described by

$$Q_{bath-liq} = h_{bath-sl}A_{sl}(T_{bath} - T_{liq}) \tag{4.61}$$

and

$$Q_{liq-sl} = \frac{2k_{sl}A_{sl}(T_{liq} - T_{sl})}{x_{sl}} \tag{4.62}$$

where T_{sl} [°C] is the center temperature of the side ledge, T_{liq} [°C] is the liquidus temperature as described in equation (4.67), T_{bath} [°C] is the uniform bath temperature, $h_{bath-sl}$ [W/(m^2 °C)] is the heat transfer coefficient in the transition between bath and side ledge, k_{sl} [W/(m °C)] is the thermal conductivity of the side ledge and x_{sl} [m] is the thickness of the side ledge as shown in Figure 4.1.

From the middle of the side ledge to the middle of the wall outside it, the conductive heat transfer $Q_{sl-wall}$ [W] is described by

$$Q_{sl-wall} = \frac{A_{sl}(T_{sl} - T_{wall})}{x_{wall}/(2k_{wall}) + x_{sl}/(2k_{sl})}. \tag{4.63}$$

In (4.63), T_{wall} [°C] is the temperature in the middle of the wall, k_{wall} [W/(m °C)] is the heat conductivity of it while x_{wall} [m] describes its constant thickness.

From the middle of the wall, the heat is transported to the outside of the cell by conduction through the rest of the wall and then by convection against the air outside the cell. The temperature in the outer part of the wall will be so low that radiating heat from the

outside of the wall can be neglected. The air outside of the cell is assumed to have a constant ambient temperature T_0 [°C]. Summarised, the heat Q_{wall-0} [W] transferred from the middle of the wall to the air outside the cell can be described by

$$Q_{wall-0} = \frac{A_{sl}(T_{wall} - T_0)}{1/h_{wall-0} + x_{wall}/(2k_{wall})}. \tag{4.64}$$

As mentioned above, there will also be a transfer of energy between the side ledge and the bath through the freezing and melting of side ledge. There will be two important contributions altering the energy balance of the cell which are related to the melting and freezing of side ledge. These contributions are the energy needed to change the temperature of the mass flowing from one control volume to the other and the energy transfer due to a phase transition of the mass. The temperature change (tc) can be described by the equations

$$E_{tc,\,freeze} = w_{freeze} c_{p_{cry,l}} (T_{bath} - T_{liq}) \tag{4.65}$$

and

$$E_{tc,\,melt} = w_{melt} c_{p_{cry,s}} (T_{liq} - T_{sl}), \tag{4.66}$$

where only one of these can be active at once. Here, w_{melt} and w_{freeze} are the masses transferred to and from the bath. These are described by (4.16) and (4.17). $c_{p_{cry,l}}$ and $c_{p_{cry,s}}$ are the specific heat capacities of liquid and solid cryolite respectively. The temperature on the surface between the molten bath and the side ledge is called the liquidus temperature. It is the temperature at which the cryolite bath melts. There are a lot of equations available for calculating the liquidus temperature based on the concentrations of different compounds in the bath. The one used in this thesis is a simplified version of the formulation proposed by Kolås (2008). It is given by

$$\begin{aligned} T_{liq} =& 1011 + 0.63\big[AlF_3\big] - 0.135\big[AlF_3\big]^{2.2} \\ & - \frac{(8.1 + 0.000022\big[AlF_3\big]^{3.5}\big[Al_2O_3\big])\big[Al_2O_3\big]}{1 + 0.119\big[Al_2O_3\big] - 0.012\big[Al_2O_3\big]^{1.5}} \\ & + 0.0178\big[AlF_3\big]^2(1 - e^{-(2-0.08[AlF_3])[CaF_2]}) \\ & - 0.0012(\big[CaF_2\big]\big[AlF_3\big])^2 - C_0, \end{aligned} \tag{4.67}$$

where C_0 replaces the parts of the equation that are assumed to be constant. $\big[CaF_2\big]$ is also assumed to be constant. One important simplification that has been done in (4.66) is that the temperature difference is calculated from T_{liq} at the surface of the side ledge and all the way in to the middle of it. This is a simplification to keep the number of equations down. The alternative would be to discretise the side ledge into multiple control volumes with additional temperature equation, which is unnecessary for the precision wanted for this model. The second contribution to energy change due to melting and freezing is the state change (sc). This describes the energy taken from the bath to melt the side ledge with $E_{sc,\,melt}$ [W] and the energy added to the bath when some liquid cryolite freezes on to the side ledge with $E_{sc,\,freeze}$ [W]. These are calculated as

$$E_{sc,\,melt} = w_{melt} \Delta H_{cry} \tag{4.68}$$

and

$$E_{sc,\,freeze} = w_{freeze}\Delta H_{cry}, \tag{4.69}$$

where ΔH_{cry} [J/kg] is the specific heat needed to melt or freeze cryolite.

The downwards heat transfer is described in a similar way as the sideways heat transfer with conduction and convection. From the bath, the first part of the heat transfer goes to the middle of the metal pad. The pure aluminium is assumed to conduct heat so good that it is not necessary to include the conductive heat loss here. A consequence of this is that the metal temperature T_{met} [°C] is assumed to be the same in the whole control volume. That means the total heat transfer can be expressed as

$$Q_{bath-met} = A_{top}h_{bath-met}(T_{bath} - T_{met}), \tag{4.70}$$

where $h_{bath-met}$ [W/(m^2 °C)] is the heat transfer coefficient between the bath and the metal. From there, the heat goes on through the cathode. The equation for the heat transfer $Q_{met-cat}$ [W] between the metal pad and the middle of the cathode is

$$Q_{met-cat} = \frac{A_{top}(T_{met} - T_{cat})}{1/h_{met-cat} + x_{cat}/(2k_{cat})}. \tag{4.71}$$

Here, T_{cat} [°C] is the temperature in the middle of the cathode, $h_{met-cat}$ [W/(m^2 °C)] is the heat transfer coefficient between the metal pad, k_{cat} [W/(m °C)] is the thermal conductivity of the cathode while x_{cat} [m] is the thickness of the cathode. The heat transfer from the cathode to the outside of the cell is called Q_{cat-0} [W] and has the same form as (4.64). It is given by

$$Q_{cat-0} = \frac{A_{top}(T_{cat} - T_0)}{1/h_{cat-0} + x_{cat}/(2k_{cat})}, \tag{4.72}$$

where h_{cat-0} [W/(m^2 °C)] is the heat transfer coefficient between the cathode and the air gap underneath the cell.

It is a little more complex to model **the upwards heat loss** than the sideways and downwards heat losses. As shown in Figure 4.4, it is necessary to break the top crust to add Al_2O_3 and AlF_3 to the bath. It is also necessary to change anodes when they are worn down. Both these things makes it challenging to get a model that is correct under all considerations. This combined with the fact that there is both radiation from the bath to the crust, convection on both sides of the crust, conduction through it and that there are usually no measurements of the thickness of the crust makes it hard to create a good heat loss model for it. Instead, the crust heat loss Q_{crust} [W] is modelled as a linear function which is quite precise in the normal area of operation for the cell. The model used here is thus

$$Q_{crust} = (T_{bath} - T_0)(A_{top} - A_{anodes})h_{top-0}, \tag{4.73}$$

which is based on the equation for a convective heat loss, but where the heat transfer coefficient h_{top-0} [W/(m^2 °C)] is used as a tuning parameter to make this one equation represent the whole heat loss upwards from the bath to the outside of the cell.

In addition to the heat loss through the crust, there will also be a significant heat loss through the anodes. This is modelled in two stages, one from the bath to the middle of the

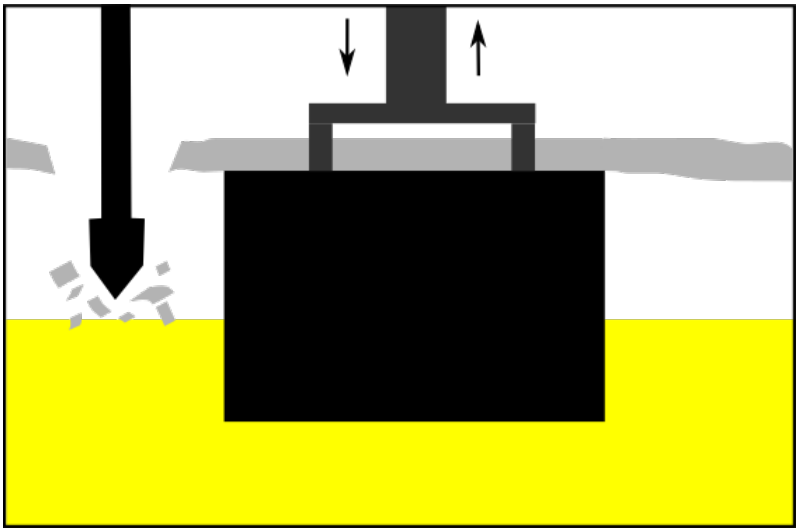

Figure 4.4: Complexity of top heat loss

anodes and one from the middle and all the way to the outside of the cell. The heat loss $Q_{bath-an}$ [W] from the bath to the center of the anodes can be modelled as

$$Q_{bath-an} = \frac{A_{anodes}(T_{bath} - T_{an})}{1/h_{bath-an} + x_{an}/(2k_{an})}, \tag{4.74}$$

where T_{an} [°C] is the center temperature in the anodes, $h_{bath-an}$ [W/(m^2 °C)] is the heat transfer coefficient between the cryolite bath and the anodes, k_{an} [W/(m °C)] is the heat conductivity of the carbon in the anodes and A_{anodes} [m^2] is the total anode area as seen from above. Next, the heat Q_{an-0} [W] travels from the anodes to the surroundings. This can be expressed as

$$Q_{an-0} = \frac{A_{anodes}(T_{an} - T_0)}{1/h_{an-0} + x_{an}/(2k_{an})}, \tag{4.75}$$

where h_{an-0} [W/(m^2 °C)] is the heat transfer coefficient between the anodes and the surroundings.

Now, the above discussion can be summarised in the differential equations describing the time derivative of the center temperature of each control volume in the cell. The expressions are based on equation 4.40, where no significant work $\dot{W}_n$ is assumed to be present in any of the control volumes. First off, the bath temperature equation is presented. The heat $\dot{Q}$ transferred across the control surface is here expressed as the the difference between the added electrical power P_{el} (4.54) and the heat losses Q_{crust} (4.73), $Q_{bath-an}$ (4.74), $Q_{bath-met}$ (4.70) and $Q_{bath-liq}$ (4.61). Further, the energy change due to mass transfer is described by E_u (4.58), which is the energy needed to heat all bath additions, and the heating of melted cryolite $E_{tc,\,melt}$ (4.66). It is assumed that all the energy needed for melting and gained from freezing is taken from the bath. It is described by $E_{sc,\,melt}$ (4.68) and $E_{sc,\,freeze}$ (4.69), which described the energy change related to the melting and

freezing of side ledge. The reaction energy is given by $E_{r_{ox}}$ (4.55) and $E_{r_{bath\,prod}}$ (4.56). Combining all of these, the equation for $\dot{T}_{bath}$ [°C/s] is

$$\begin{aligned}\dot{T}_{bath} = &\frac{P_{el} - E_u - E_{r_{ox}} - E_{r_{bath\,prod}} - Q_{crust} - Q_{bath-an}}{m_{bath} c_{p_{bath,l}}} \\ &+ \frac{E_{tc,\,melt} - E_{sc,\,melt} + E_{sc,\,freeze}}{m_{bath} c_{p_{bath,l}}} \\ &- \frac{Q_{bath-met} + Q_{bath-sl}}{m_{bath} c_{p_{bath,l}}},\end{aligned} \tag{4.76}$$

where m_{bath} [kg] is the mass of bath calculated with (4.12) and $c_{p_{bath,l}}$ [J/(°C kg)] is the specific heat capacity of the liquid cryolite bath. This specific heat capacity is assumed to be constant.

The side ledge temperature change $\dot{T}_{sl}$ [°C] is given by

$$\dot{T}_{sl} = \frac{Q_{liq-sl} - Q_{sl-wall} + E_{tc,\,freeze}}{m_{sl} c_{p_{cry,s}}}, \tag{4.77}$$

where m_{sl} [kg] is the mass of the side ledge and $c_{p_{cry,s}}$ [J/(°C kg)] is the specific heat capacity of frozen cryolite. m_{sl} can be calculated with (4.21). The heat transport to the outside of the cell will change the side ledge temperature according to the difference between Q_{liq-sl} (4.62) and $Q_{sl-wall}$ (4.63). Additional change in the side ledge temperature is caused by the melting and freezing of side ledge, which can be described by $E_{tc,\,freeze}$ (4.65) and $E_{tc,\,melt}$ (4.66). $E_{sc,\,freeze}$ and $E_{sc,\,melt}$ are not added because the state change energy is assumed to only be taken from and added to the bath.

The change of temperature in the middle of the side walls, T_{wall} [°C], can be expressed by the differential equation

$$\dot{T}_{wall} = \frac{Q_{sl-wall} - Q_{wall-0}}{m_{wall} c_{p_{wall}}}, \tag{4.78}$$

where m_{wall} [kg] is the constant mass of the walls surrounding the bath and side ledge and $c_{p_{wall}}$ [J/(°C kg)] is the average specific heat capacity of the materials in the wall. The temperature will increase or decrease based on whether the heat added to the wall by $Q_{sl-wall}$ (4.63) or the heat removed from it with Q_{wall-0} (4.64) is the largest of the two.

For the time derivative of the metal temperature T_{met} [°C/s], it is possible to give the following equation:

$$\dot{T}_{met} = \frac{Q_{bath-met} - Q_{met-cat} + 0.002 M_{Al} r_{ox} c_{p_{met}} (T_{bath} - T_{met})}{m_{met} c_{p_{met}}} \tag{4.79}$$

$Q_{bath-met}$ (4.70) and $Q_{met-cat}$ (4.71) describe the heat transfer through the metal cap while the last term in the numerator describes the energy increase gained from cooling down the aluminium which is produced by the primary reaction. m_{met} [kg] and $c_{p_{met}}$ [J/(°C kg)] are the mass and specific heat capacity of the metal pad respectively. The metal mass is calculated with equation (4.22).

The cathode temperature time derivative $\dot{T}_{cat}$ [°C] is calculated as

$$\dot{T}_{cat} = \frac{Q_{met-cat} - Q_{cat-0}}{m_{cat} c_{p_{cat}}}. \tag{4.80}$$

m_{cat} [kg] is the constant mass of the cathode while its specific heat capacity is given by $c_{p_{cat}}$ [J/(°C kg)]. As with the wall, the temperature of the cathode will increase or decrease based on whether more heat is added through $Q_{met-cat}$ (4.71) or removed through Q_{cat-0} (4.72)

The last temperature it is necessary to keep track of is the anode temperature T_{an}. This is a common temperature for the center of each of the anodes as the heat is assumed to be distributed equally in the whole bath and out through the anodes. The differential equation of this temperature is given by

$$\dot{T}_{an} = \frac{Q_{bath-an} - Q_{an-0}}{m_{an} c_{p_{an}}}, \tag{4.81}$$

where m_{an} [kg] is the sum of the masses of all the anodes and $c_{p_{an}}$ [J/(°C kg)] is their specific heat capacity. The temperature will vary with $Q_{bath-an}$ (4.74) and Q_{an-0} (4.75) in the same way as for the wall and cathode.

4.4 Model inputs and outputs

The above discussion has presented simple mass and energy balances for an aluminium smelter. These equations show the impact different inputs have on the process. During operation of a plant, it is essential to keep track of key process quantities. This is usually done with measurements. The operators can use information from the obtained measurements to correct the inputs so that a stable cell operation is obtained. This section finishes the model presentation by giving an overview of the inputs and the measurements included in the model. In Chapter 2 above, more specifically section 2.3.3, the usual measurements and inputs are presented. It is however reasonable to give a short presentation of which of these are included in the model, and how they affect it.

There are in total 5 different ways in which to alter the system. In a long-term perspective, the current is the input that has the largest influence on the economics of a cell. The current decides the cell production rate. To keep the production rate up, the current will usually be chosen to be constant and as high as the *potline*[4] can handle. The maximum line current is decided by the cell design. Reducing it could be useful to simulate e.g. a problem with delivering enough current to the cell.

In a shorter time span, it is possible to control the power throughput of the electrolysis cell by moving the anodes. As the molten aluminium leads electricity really well, it is considered to be a part of the cathode electrically. That means that as the amount of metal increases, the anodes are slowly moved upwards to keep the gap between the anodes and the metal pad constant. This will make sure the pseudo-resistance, and hence the power input to the cell, is kept as stable as possible. If it is desirable to increase or decrease the

[4]Multiple cells placed next to each other share a common line current

bath temperature, this can be done by altering the anodes in such a way that the anode-cathode distance is changed. It will then lead to a change in the power input to the cell. In this model, the anodes are assumed to automatically follow the change in metal height. Because of this, it is only necessary to use this input when a different anode-cathode distance is wanted.

It is assumed that both alumina and aluminium fluoride are both fed to the cell with automatic feeding systems to keep the control of the cell as good as possible. That means it is reasonable to simulate these with small amounts being fed quite often. In the model, this is done by adding a small amount of fluoride every 10 samples. As described in the model, the alumina feed is already handled as a P-controller. It is hence not considered a model input. The model has however support for adding additional aluminium fluoride to the cell. This will increase the acidity and lower the liquidus temperature of the cell.

The two last ways of affecting the cell from outside are by tapping metal and bath. In the model, both of these are assumed to happen within one sample. Even though tapping will usually take more time than this, doing such a simplification will not have a large impact on the quality of the model.

The measurements included in the model are some of those that are most common to find in all aluminium smelters. To obtain measurement series, the model is run with a certain configuration, and then the measurement values are collected directly from the model in this configuration. Noise is added to the measurement series by adding a random signal to the nominal value using MATLAB. The pseudo-resistance is the only continuous measurement available. This is collected directly from the model equations. The same is the case for the bath temperature measurement and the acidity measurement. Both the bath temperature and the acidity are essential process values to control the heat balance of the cell. Both the bath and metal heights are used as measurements in the model as well. These are calculated based on the bath and metal masses in addition to the densities of these control volumes. One simplification that is done here is that when the bath height measurement is calculated, the fact that some of the bath lies above the bottom of the anodes is ignored, so that it can be calculated directly from the mass and density of the bath.

4.5 Model behaviour

The simulation model manages to recreate the main aspects of the mass and energy balances quite well. As it has not been fitted against data from a real cell, it is hard to say if the dynamics are quantitatively precise. The model does however seem to behave qualitatively reasonable. The following will highlight some of the limitations and strengths of the model. The most uncertain parameters in the model are the heat transfer coefficients for convective heat transfer. There are several conditions which decide what the values should be, so that the best way of deciding them are often empirically (Jessen, 2008). Because of this, these are typical values chosen for the Kalman filter to adjust in order to fit the model to measured data from a real cell. In addition, there is also included some uncertainty in the metal tapping. This is added to simulate the fact that the weight used to evaluate how much metal has been tapped does not give perfect results. These are then the parameters that will be used for estimation when running different tests on the Kalman

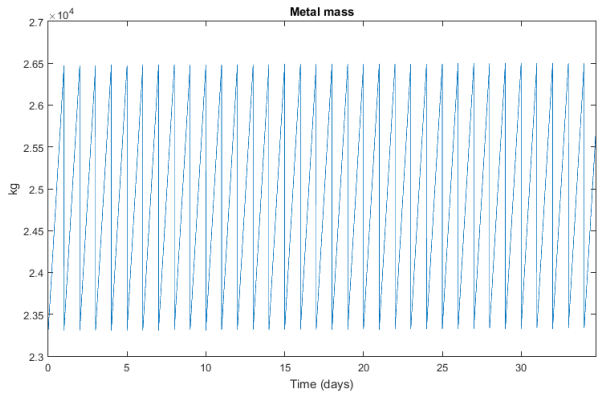

Figure 4.5: The simulated metal production and tapping. Metal is tapped daily

filter performance.

The mass balance seems to behave as expected. It can be seen in Figure 4.5 that it has been decided to tap metal once a day. The tapped amount is approximately 3200 $\mathrm{kg/day}$. In Grjotheim et al. (1993), it is stated that on cells of approximately 300 kA produce 2300 $\mathrm{kg/day}$. In this model, the current is set to 420 kA, which is a reasonable amount for a modern cell (Kvande and Drabløs, 2014). The ratio between current and metal production is approximately the same for both cells, so this shows that the chosen production rate is reasonable. The difference from a real cell is that in reality, the production rate will not be constant because the current efficiency is not constant. Some uncertainty in the metal mass can be introduced by setting the metal tapping parameter to a value different from 1. A weakness with the way the model is designed is that the ACD used in the electrical model to adjust the systems resistance is decoupled from the metal height. This is done to simulate that the anodes will automatically follow the change in metal height. By doing it this way, the ACD and hence resistance is kept constant if a change in the power input to the cell is not desired. The problem about this decoupling is that when uncertainty is introduced in the metal tapping, this will lead to a deviation between what the metal height actually is and what the model calculates it to be. To calculate the effect of this, lets say the weight used to weigh the tapped metal consequently shows e.g. 5 % more than what is actually tapped. The model will accept the value given by the weight, leading to it believing it has 5 % more metal than what is actually present in the process. This will however not impact the electrical model, since R_{pseudo} is only based on the value of the ACD. In our simulated process, the metal height is found to vary with about 2.3 $\mathrm{cm/day}$. That means the error in metal height between the simulated process and the model would increase with 1.15 $\mathrm{mm/day}$. Such an error would lead to an increase in resistance deviation of 0.17 μΩ which is about 2.5 % of the nominal cell resistance. This is a problem if the deviation keeps increasing for several days. In this thesis, the error is assumed to be compensated for by the Kalman filter through the metal height measurements. The weakness of this approach is that there will be some deviation in the resistance model before the Kalman filter manages to compensate for this modelling

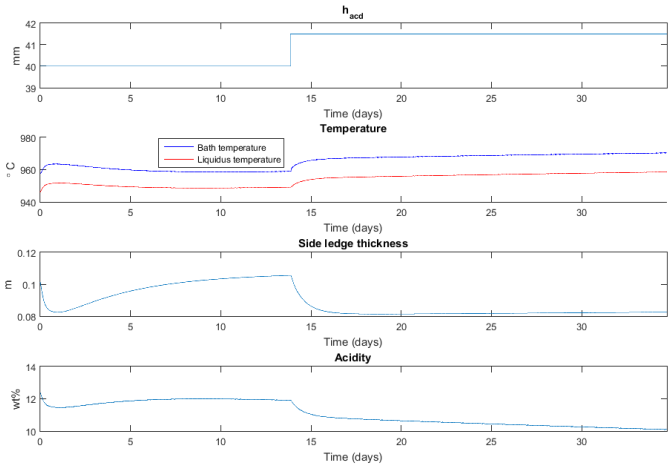

Figure 4.6: The change in bath temperature, liquidus temperature, side ledge thickness and acidity from a step in the ACD

error. The positive part of it is that it makes for really simple control of the anodes. An alternative here would be to control it using a PI-controller. This would however increase the complexity of the model. As one of the most important aspects in this model is that it is supposed to be simple, this given model structure is then kept.

The balance between the heat balance and the side ledge thickness shows reasonable behaviour in a qualitative manner. As the model has not been checked against real cell data, it is hard to say if the correct amount of side ledge is added or removed from a certain increase in bath temperature. Figure 4.6 shows that the model manages to obtain stable operation at a thinner side ledge when the temperature is increased. The same figure also displays that when the side ledge, consisting solely of cryolite, melts, the concentration of aluminium fluoride goes down, which again increases the liquidus temperature, restoring a stable heat balance for the cell. It can be seen in Figure 4.7 that when the side ledge melts, the mass of cryolite in the bath increases. This is as expected. As the uncertain parameters in the model have been chosen to be some of the heat transfer coefficients, it is possible that this makes the model calculate more or less melting/freezing of side ledge than what should be expected. This will again change the bath height. That means that the bath height measurement can be weakly connected to the estimation of the uncertain heat transfer coefficients. The resulting difference in bath height from the uncertainty in melting/freezing of side ledge will however be very small, so this is not a precise way of correcting the heat transfer coefficients. The tapping and addition of bath is assumed to be perfectly executed in the model, so that there is no uncertainty connected to this. In the model, bath was tapped half way into the simulation. This can be seen in Figure 4.7. The amount of bath will vary with the state of the cell, so it should only be tapped when it is necessary. This will then not happen regularly. Tapping bath does not affect the heat balance of the cell to a large extent. It makes sense because the mass that is left in the bath

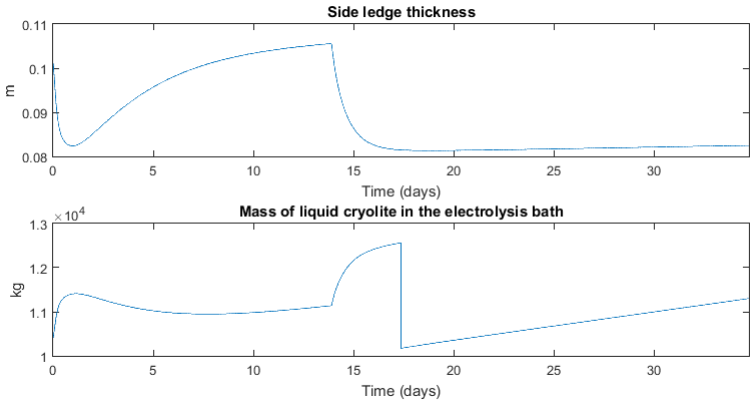

Figure 4.7: The mass of liquid cryolite increases as the side ledge melts. Bath is tapped about 17 days into the simulation. After the bath is tapped, the steady increase in cryolite mass shows that some cryolite will be produced continuously by the bath production reaction

has the same composition and temperature as what is taken out.

The heat lost from the cell is distributed in all directions as it is supposed to. As can be seen in Figure 4.8, about 60 % of the energy is released through the top of the cell, 25 % is released through the side walls while the remaining 15 % goes downwards. If this is compared to the typical heat loss distribution of a cell in Figure 4.9, it can be seen that the heat loss from the top of the cell is a little high. At the same time, the heat loss going sideways is too small. In the discussion around Figure 4.9 in Grjotheim et al. (1993), it is made clear that a heat loss over the top of between 40-60 % is reasonable. That means the obtained value is within reasonable limits. The sideways heat loss is probably still a little low. When adjusting the model to obtain a stable operating state, increasing the sideways heat loss and lowering the top heat loss led to a severe melting of the side ledge. The final choice of cell parameters achieves the goal of keeping the cell at a stable side ledge thickness at the same time as the heat distribution of the cell is quite good. A potential improvement to the heat balance would be to split each control volume into multiple smaller control volumes. That would make a more precise description of the temperature distribution of the cell. The cost of doing this would however be to increase the model complexity. This goes against the goal of creating a simple simulator model with a behaviour close to the main features of an aluminium cell.

One major weakness of the model is the P-controller used to handle the alumina concentration. This is a large part of the challenge of controlling real aluminium electrolysis cells. The reason why a P-controller is used in this model is because it would require a more complicated resistance model connecting it to the alumina concentration in the bath to be able to control it in the way it is usually done (Aalbu et al., 1988). The operators are not usually able to keep the alumina concentration completely stable by using this method. A real cell is quite sensible to changes in the alumina concentration compared to e.g. changes in the acidity. That means trying to stabilise the alumina concentration manually for each simulation would be quite time consuming. It was then found to be a

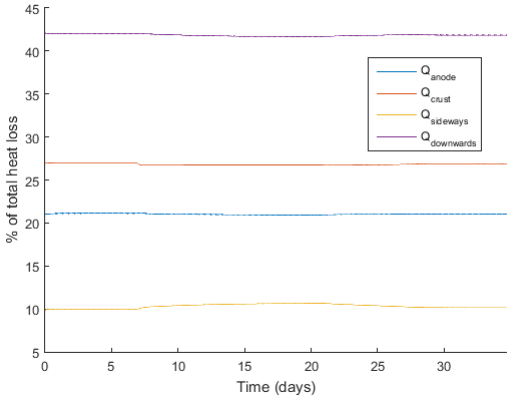

Figure 4.8: Model heat distribution

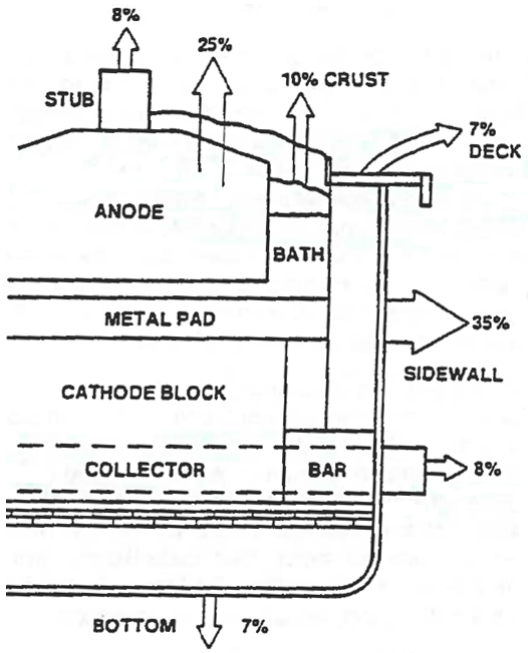

Figure 4.9: Typical heat loss distribution for a Hall-Héroult cell. Collected from Grjotheim et al. (1993)

acceptable simplification to assume that the alumina concentration was already handled. This removes some of the cell dynamics but allows for more effective simulation of other cell features.

4.6 Summary of assumptions used

This section is meant as a summary of the model assumptions which are relevant when testing the performance of the Kalman filter on the system, so that the reader does not have to look through the whole model specification to find the essential assumptions. Perhaps the most important thing to establish is that this is a simple simulator model that does not strive to cover all aspects of the cell dynamics. The key goal of the model is to create a system whose dynamics have similar behaviour to a simple version of an aluminium electrolysis cell. To keep the model simple, all unwanted disturbances like sludge formation, anode effect and power outages are disregarded in the model. All of these effects happen regularly, but they all tend to cause large deviations from normal operation. As the main objective of this thesis is to look at how the Kalman filter reacts to changes in the measurement methods, introducing such deviations will only make the results less certain. Another periodical operation that is not included in the model is anode change. When an anode is changed, there will be a period of time where the new anode does not function properly 2.3.4. This is an unpredictable effect that is hard to model. It is therefore excluded. When the changing of anodes is not a part of the model, it is also reasonable to assume that they are not consumed, and that all anodes can instead be modelled to be of a common average anode height. Then it is not necessary to model how the oxygen of the alumina binds to the carbon on the anodes to create carbon dioxide.

The physical dimensions and properties of the cell is assumed to be constant and known during all simulations. This includes the anode, cathode and side ledge masses, lengths, widths and height of the cell and the conductivity constants of the different materials. This is a reasonable assumption in the time frame that the model is going to be used for. Over a longer period of time, the cathode could experience some swelling because sodium and bath penetrates it, but this is a rather slow process, and it is hence regarded as negligible (Grjotheim et al., 1993).

For the mass balances, the assumptions made are regarding which substances should be included in each control volume and whether they are constant or dynamic. The bath is assumed to have a dynamic amount of Al_2O_3, AlF_3 and Na_3AlF_6. In addition, a constant fraction of the bath (6 %) is assumed to be made up of a combination of other minor components such as LiF, CaF_2, MgF_2 and KF. Some cryolite will be interchanged between the side ledge, which is assumed to consist of pure cryolite, and the bath. Further is the produced metal assumed to go instantly to the metal pad beneath the bath when it is produced.

With regard to the measurements, it is assumed that the time delay of obtaining the acidity in the lab is ignorable compared to the slow dynamics of the cell. The effect of this is that the measurement can be placed in the time series on the point where the results are ready instead of when the sample was obtained. Another important aspect regarding the measurements is that the alumina concentration is not assumed to influence the resistance measurement. That means that it is not possible to use the resistance measurement to cal-

culated the alumina concentration. It is therefore assumed that the alumina concentration is already controlled.

The ACD is assumed to always be the same unless it is changed manually. What this means is that the anodes increases the same amount as the metal height. In a real cell, it would happen automatically to keep the power input constant. This is because the resistance would change when the ACD changes. The problem is here that the resistance is independent of the metal height. If the actual metal height is then not equal to the modelled value, this change is not picked up by the resistance measurement. This deviation is assumed to be compensated for by the Kalman filter estimator.

4.7 Model conclusion

It seems the model achieves what it was designed to do. The main dynamics of the cell, how everything is connected, seems to be working as expected. When the power input is increased, this leads to an increased bath temperature, which again melts away some of the side ledge. By melting the side ledge, the heat loss of the cell is increased. At the same time, the additional melted cryolite in the bath will lead to a decrease in in its aluminium fluoride concentration, and hence an increase in the liquidus temperature. Both these factors will then lead to a new stable cell state. Overall, the model gives a good representation of the connections between acidity, bath temperature, side ledge thickness and the heat balance of the cell. At the same time, it seems that metal production and tapping has a qualitatively reasonable behaviour. The whole heat balance could of course be significantly improved by using data from an actual cell to adjust the heat distribution throughout the cell.

The largest weakness of the model is its inability to manipulate the alumina feed. A varying alumina feed could make it possible to get better control of the cell resistance. With an improved cell resistance model, it could be possible to simulate disturbances like anode changing, anode effect and sludge formation. All of these are regularly present in a real aluminium cell and will affect the cell resistance. Adding a more advanced feeding system for alumina would however have a negative effect on the simplicity of running the simulations. Because of this, choosing to assume that it is already stabilised is a good assumption for a simulator, even though it would not work to control a real cell.

Chapter 5

Kalman filter testing

In an industrial process, measurements are normally used to provide essential state information for controlling the system. The rate at which these measurements are taken is usually a trade-off between two things. One is the cost of each measurement, both in personnel, HSE risk and equipment. The other is how often the measurements are needed to obtain satisfactory control/estimator performance, and how much it is possible to save by achieving better control of the cell. Aluminium production plants will typically have a continuous resistance measurement which is obtained automatically. In addition, multiple manual measurements are taken infrequently. These are typically performed by operators that are available in the plant regardless of how often the measurements are taken. It is then reasonable to think that the cost of a measurement is more dependent on how many measurements are taken (HSE risk: More measurements means more exposure to hot surfaces, high currents, liquid metal and dangerous gas. Equipment cost: rods used for height measurements and thermocouples use for temperature measurements are worn down in the corrosive electrolysis bath) than on the measurement strategy (whether each unique measurement is taken at a separate time or if they are all performed at once). In a full scale plant, which can have hundreds of cells, the cost of doubling the measurement frequency on every cell can be substantial. It is then interesting to look at how varying the measurement intervals will affect the control/estimator performance.

The model developed in Chapter 4 is used in this chapter to investigate how manipulating the measurement strategy (shifted versus joint measurements) will affect the performance of the Multi-rate Extended Kalman filter estimator. It is also investigated how the optimal tuning of the Kalman filter is affected when the measurement intervals are varied. These cases have been studied through simulations in *Modelfit*, a tool developed by Cybernetica AS for offline estimation of states and parameters. Measurement series are created by simulating the process with a fixed set of parameters. Then, one or more parameters are changed to simulate the fact that it is hard to design a perfect model. The task is then for the Kalman filter to use the infrequent measurements to bring all changed parameters back to the value they had when the measurement series were created. When tuning the Kalman filter, the process noise parameters are chosen to have the same values throughout

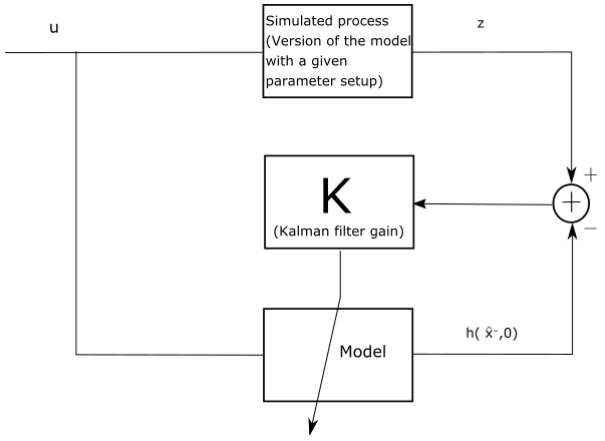

Figure 5.1: Interaction between the process, model and Kalman filter

all simulations while the measurement noise parameters are changed in order to alter the filter's characteristics. In Figure 5.1, the interaction between the simulated process, the process model and the Kalman filter is described. This is equivalent to what happens in equation (3.30). The difference in output between the process and the model is multiplied by the Kalman gain in order to correct the model in such a way that the difference in output is minimised.

To be able to separate the different measurement noise parameters used throughout this chapter and define how they are tuned, it is necessary to establish some notation. The measurement covariance matrix R can for the Kalman filter used here be defined as

$$R = \begin{bmatrix} r_R^2 & 0 & 0 & 0 & 0 \\ 0 & r_{BT}^2 & 0 & 0 & 0 \\ 0 & 0 & r_A^2 & 0 & 0 \\ 0 & 0 & 0 & r_{MH}^2 & 0 \\ 0 & 0 & 0 & 0 & r_{BH}^2 \end{bmatrix}. \tag{5.1}$$

Here, r_R [µΩ], r_{BT} [°C], r_A [wt%], r_{MH} [m] and r_{BH} [m] are the standard deviations of the measurement noise v on each of the measurements. These are the values that are changed in order to tune the Kalman filter in the discussions below. The subscript abbreviations are used to separate the measurement noise parameters of the different measurements. They are defined as:

- R = Resistance
- BT = Bath Temperature
- A = Acidity
- MH = Metal Height

- BH = Bath Height

When r_i is used, it refers to the standard deviation of any or all of the measurements in general. In addition to the symbols used for the standard deviation of measurement noise, it is convenient to define notation for the measurement frequencies of different measurements. These are given by $\Delta t_{m,i}$, where i can be any one of the abbreviations in the list above.

To add white noise on the measurements, MATLABs "rand"-function is used to create a vector of random numbers in a specific range around the noise-free measurement. This range is decided individually for each measurement, and is calculated as a certain percentage of the mean of all noise-free measured values of a measurement in a given data set. Table 5.1 summarises what percentage is chosen for each measurement, and the resulting size of the white noise. At about 1000 °C, the accuracy of different types of thermocouples is usually in the area of 1-4 °C (International Electrotechnical Commission (IEC), 1982). From this, an accuracy of 0.1 %, which gives approximately 1 °C is chosen. It is hard to find any specific information about the accuracy of the resistance, acidity and height measurements of a cell. The chosen percentages are based on what was found reasonable during a conversation with a supervisor at Cybernetica AS. When thinking about potential ways of improving the simulations after they were done, it is clear that a better way of choosing the standard deviation of the white noise would be to give it a constant value instead of calculating it as a percentage of the mean value of a measurement in the data set. This mean will change slightly between simulations, giving a slightly different maximum value for the measurement noise. It is more reasonable that the maximum value of the noise is rather defined by the precision of the measurement equipment than the state of the process. Another thing to notice is that both height measurements are usually measured using the same measurement technique. 5 % of the measurements mean value for these measurements was considered as a reasonable uncertainty in the measurement, as this gives an uncertainty of about 1 cm for the metal measurement. This was found to be reasonable for a measurement that is performed by sticking a pole into the bath and metal pad. It is however clear that as the mean height of metal is more than 50 % larger than the bath height, this indicates that the bath height measurements are quite a lot more precise than what is accomplished for the metal height measurements. As the two measurements do not provide information about the same states of the system, and they are in the same range, this configuration is still considered to give reasonable results.

Table 5.1: Maximum absolute values for the white noise on the process measurements. The percentages are included because they describe the basis for how the absolute values are decided. A potential improvement of the measurement noise model would be to decide constant absolute max-values instead of calculating them as a percentage of the mean value of the measurement in a data set

	Resistance [μΩ]	Bath temperature [°C]	Acidity [wt%]	Metal height [m]	Bath height [m]
%	1	0.1	1	5	5
Absolute values	0.066	0.967	0.101	0.009	0.0057

5.1 Utilised software - Modelfit

Modelfit is a software tool developed by Cybernetica AS. It is used for offline estimation of states and parameters for mathematical models. The Modelfit tool is designed to work with models created in a Cybernetica framework. These models are usually written in the C programming language. The purpose of the program is to use simulations to get a best possible match between a theoretical model and real process measurements. When developing a first principles model of a process, it will be based on theoretical equations and constants. It is well established that while these equations and constants may give a good description of the system in an ideal, noise-free setting, they seldom manage to recreate the behaviour of a real system without any alterations. This is where Modelfit comes in. By using measurement data from a real cell, it is possible to do adjustments to the model so that the deviation between the model output and real measurement data is minimised. Such offline model fitting should be done before the model is installed in an online application.

Modelfit offers two approaches to accomplish a best possible fit between the model and measurements. One is to estimate states or parameters using a Kalman filter or Moving Horizon Estimator. This way it is possible to establish the noise model of the process and measurements from the real system. It also gives the opportunity to easily determine which states and parameters should be estimated upon. The other key method featured in Modelfit is parameter fitting. Here, an optimisation problem is solved in which the goal is to minimise the difference between the model output and the process measurements. This is done by adjusting a weighted set of chosen parameters to achieve the best possible match between model and process.

When using the tool for the purpose of this thesis, the Kalman filter is used. Modelfit is a tool that is well suited for running the required tests. As mentioned above, the task is to see how different measurement frequencies affect the estimator performance. The framework allows the user to easily choose the measurement frequency without generating a completely new measurement series. At the same time, it is possible to choose which parameters to adjust in order for the Kalman filters measurement predictions to fit the process measurements without changing anything else in the model setup. Other important features include easy export of simulated data to MATLAB and flexible display options for states, measurements, inputs and so on. In addition, the Kalman filter can be easily tuned by adjusting the noise parameters of both the process and the measurements.

5.2 Shifted measurements versus joint measurements

5.2.1 Problem description

As mentioned, it is reasonable to assume that the measurement cost of infrequent measurements in an aluminium electrolysis cell varies more with the amount of measurements than the measuring strategy. This gives the flexibility of deciding when each separate measurement should be taken. It is reasonable to assume that there might be a difference between getting all the available measurement information at once, and by splitting up the measurements so that a smaller part of information comes more often. The following discussion compares these two approaches both with and without noise on the measurements.

5.2.2 Method

To test how the Kalman filter responds to splitting up the measurements compared to having them joint in one place, all the infrequent measurements were assumed to be taken daily. The tests were run with one uncertain parameter. The chosen parameter was h_{top-0} from equation (4.73). This equation is a simplified linearisation of a complicated part of the system, and hence considered to be so uncertain that it is well suited for fitting the model to the measurements. To see how well the Kalman filter performs, there has to be some deviation between the estimator outputs and the actual process measurements for it to correct. This is done by creating a measurement series for the real cell where all parameters are set to their correct value. A deviation is then created by shifting the parameter value so that the model state is no longer exactly the same as the state of the actual cell at the start of the simulation. The effect of changing the parameter is best seen in the acidity and bath temperature measurements. This is illustrated in Figure 5.2. It is worth noticing that the model will not deviate more and more from the process, but instead stabilise at a different system state.

It is reasonable that the bath temperature and acidity measurements are the ones that are most strongly connected to the heat balance, and hence the ones that will correct the parameter most. There will be two simulations for each tuning. The first will have all measurements arriving at the same time. On the second, the four infrequent measurements are spread out evenly throughout the day. To make sure the important information is spread out as much as possible, the bath temperature and acidity measurements are placed half a day apart, with the height measurements in between. This is illustrated in Figure 5.3.

Interesting results to look for is then whether spreading the measurements out will help in closing in on the correct parameter value faster than when they are joint together. It is also of interest to look at how much the parameter varies because of noise after its initial step. If the values of each measurement noise parameter r_i in the Kalman filter is set too low, the process will follow the noise more than the actual process value. If the opposite is the case, the measurements will not be trusted by the Kalman filter at all, and this will lead to the estimator not being able to replicate the system state.

When testing different tunings, the simulator ran for about 35 days. This was enough time for the Kalman filter to adjust the parameter properly, and to see how large the variations were after the parameter had closed in on the correct value. When a good tuning was found, that tuning was tested for a simulation of about 104 days, and with a lot of

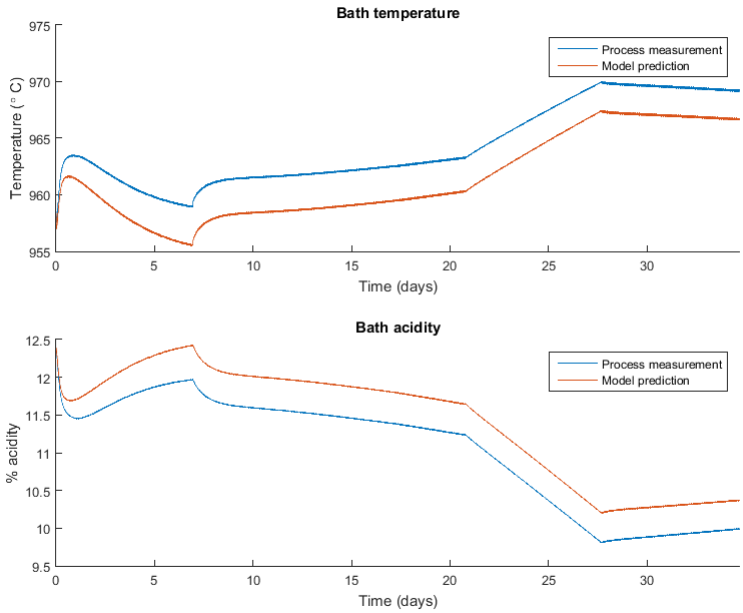

Figure 5.2: Deviation between process and model states from a parameter change

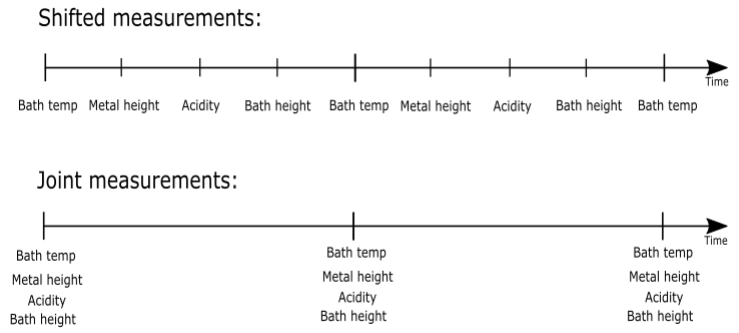

Figure 5.3: Simultaneous versus shifted measurement strategies

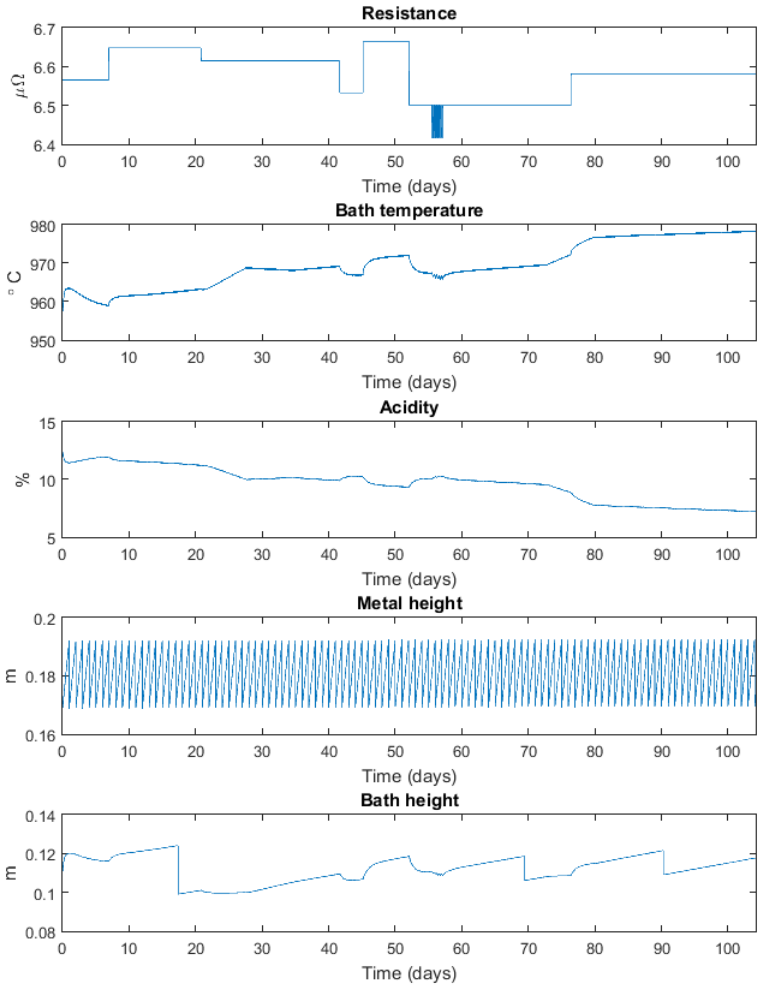

Figure 5.4: Simulated process outputs

Table 5.2: Kalman filter measurement noise tuning values when using noise-free measurements for estimation

	r_R [μΩ]	r_{BT} [°C]	r_A [wt%]	r_{MH} [m]	r_{BH} [m]
Aggressive tuning	0.0001	0.00001	0.000001	0.00000001	0.00000001
Slack tuning	0.1	5	0.5	0.1	0.1
Medium tuning	0.1	1	0.1	0.05	0.05

changes in the system inputs to see how the estimator would manage to follow the variations. The process response of the long simulation, which can be found in Figure 5.4, shows the systems response to the varying inputs. This response is caused by varying both the amount of aluminium fluoride fed to the cell and the ACD, while the bath was tapped when necessary and the metal was tapped evenly throughout the entire period.

5.2.3 Results

When no noise is added to the measurements, the Kalman filter seems to be able to estimate the right parameter value for a wide range of measurement noise tunings. In Figure 5.5, three vastly differently tuned Kalman filters have estimated this parameter for both joint and shifted measurements. Each simulation has been split into two parts. This is done to be able to show both the initial step in the parameter from 7.5 to 7, in addition to getting the smaller variations when the parameter value has stabilised. It is clear that when less measurement noise is assumed in the Kalman filter, the Kalman gain increases. Then, a more aggressive estimation of the parameter is achieved. This will lead to a quicker parameter estimation. The negative effect of this aggressive tuning is however that there will be some overshoot and then a small deviation after the parameter has stabilised close to the correct value. This effect is smaller for the less aggressive tunings. The effect this has on the system is however very small, which can be seen in Figure 5.6. It leads to a deviation of about 0.1 °C for the bath temperature and an acidity deviation of 0.02 % for the aggressive tuning. The deviation seems to be smaller for the less aggressive ones. The tunings used can be found in Table 5.2. The reason for the large gap between the aggressive and medium tuning is to show that the estimator is stable in a very large area for the given setup. In a real application, choosing a tuning with as little measurement noise as what is chosen for the aggressive tuning here is completely unreasonable. Because of this, more aggressive tunings than this were not tested. The performed tunings in the area between the medium and aggressive ones confirm the trend that as less noise is assumed on the measurements, the faster the estimation, and the larger the deviation.

When noise is added to the measurements, the impact the tuning has on the estimator performance changes a lot compared to the noise-free situation. The maximum value of the white noise added to the simulated process measurements is approximately the same as what is given in Table 5.1. Simulation results from using three different tunings can be found in Figure 5.7. Table 5.3 displays the tuning used in each case. The resulting Kalman

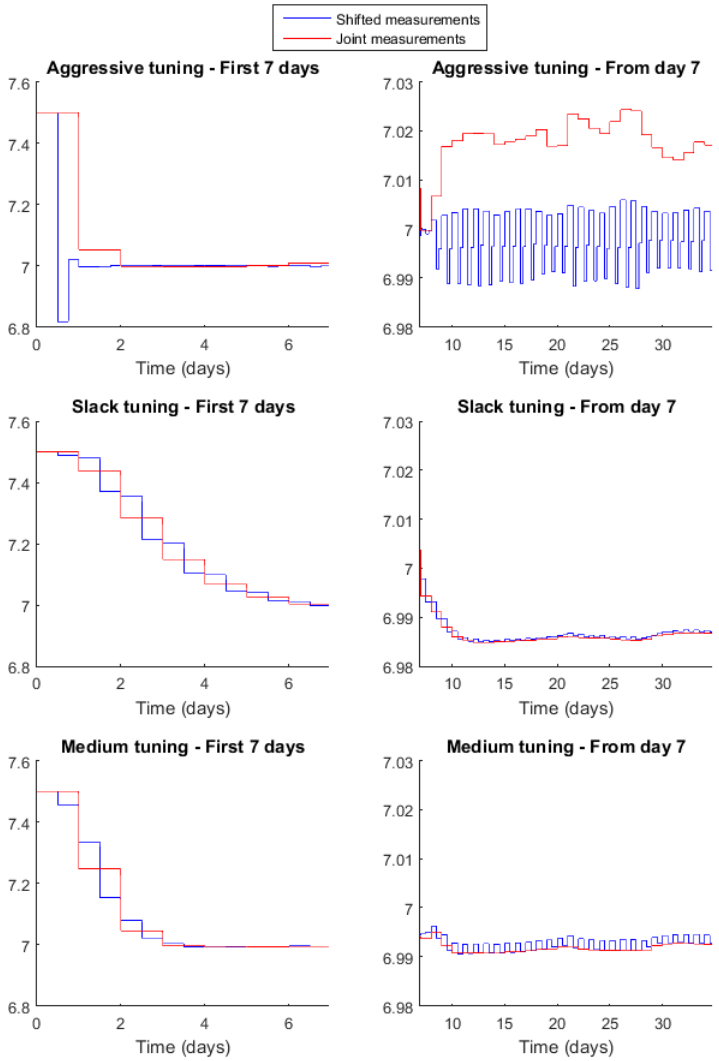

Figure 5.5: Estimation of h_{top-0} towards a reference value of 7 for three different tunings with noise-free measurements. The simulations have been split up to show the show the initial step of each tuning in the left plots and finer variations in the right plots

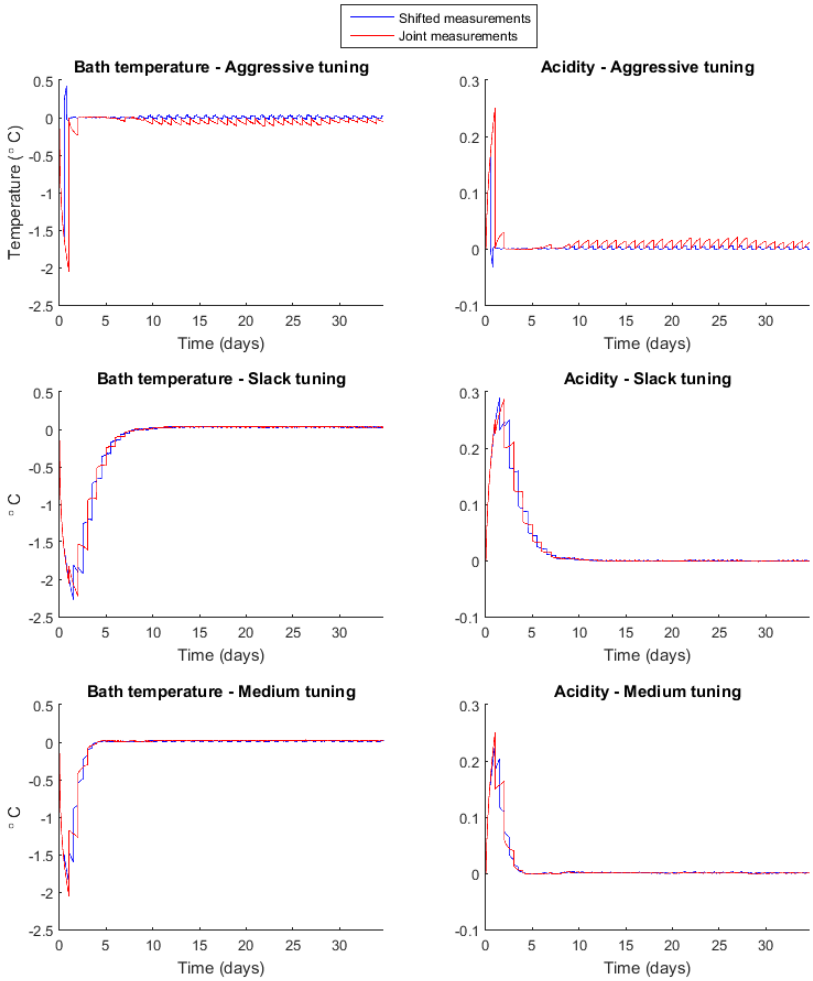

Figure 5.6: Deviation between noise-free measurements and estimator output for bath temperature and acidity with three different Kalman filter tunings. Estimation done with noise-free measurements

Table 5.3: Kalman filter measurement noise tunings when using noisy measurements for estimation

	r_R [μΩ]	r_{BT} [°C]	r_A [wt%]	r_{MH} [m]	r_{BH} [m]
Aggressive tuning	0.1	0.01	0.005	0.0001	0.0001
Slack tuning	0.1	15	1	0.1	0.1
Medium tuning	0.1	4	0.6	0.1	0.1

filter prediction deviation for bath temperature and acidity can be seen in Figure 5.8.

The ”Slack” and ”Medium” tunings both seem to achieve close to the same precision once the parameter has closed in on the correct value. The medium tuning moves towards the correct value faster. There seems to be little difference between splitting up the measurements compared to joining all of them at one time. To quantify the difference in performance between the shifted and joint measurements for each of these tunings, a longer simulation of a little more than a hundred days was performed. The real process outputs, which are the noise-free process measurements, are given in Figure 5.4. These outputs are the ones that the estimator should ideally follow, and will therefore be used as the reference when checking the Kalman filters performance.

Figure 5.9 shows how well the estimator manages to estimate the parameter in these cases. There is some overshoot for the medium tuning, so they settle completely at about the same time, even though the medium tuning comes to the right area at about half the time for both the shifted and the joint measurement cases. The overshoot seems to be a bit smaller for the joint measurements than for the shifted measurements. To test how well the estimation worked once the estimate had gotten close to the correct value, the mean and standard deviation of the difference between the estimator output and the process outputs shown in Figure 5.4 was checked. This was done for the last part of the measurement series, from day 35 and out. The results for each of the four setups can be found in Table 5.4. The mean and standard deviation in metal height was exactly 0 for all of the setups, which is reasonable as the metal production and tapping is assumed to not be directly connected to the heat balance of the cell. As discussed, the resistance is not affected by changes in the heat balance, so that the variations in resistance is also negligible. The magnitude of the mean and standard deviation for the resistance estimate was in the area of less than 3e-16 μΩ. These values are hence not included in Table 5.4 as they are so small they are not considered of interest in the discussion.

5.2.4 Discussion

Table 5.4 gives a lot of interesting information. As for the main task of looking at how shifting the measurements to separate instants affects the estimator performance, the results suggest that the effect of doing this is insignificant. With the same tuning, all simulation results suggest that the performance is very close to the same for shifted and joint measurements both in parameter convergence rate and estimate precision once it has gotten close to the correct value. For the medium tuning, the table shows that the joint measure-

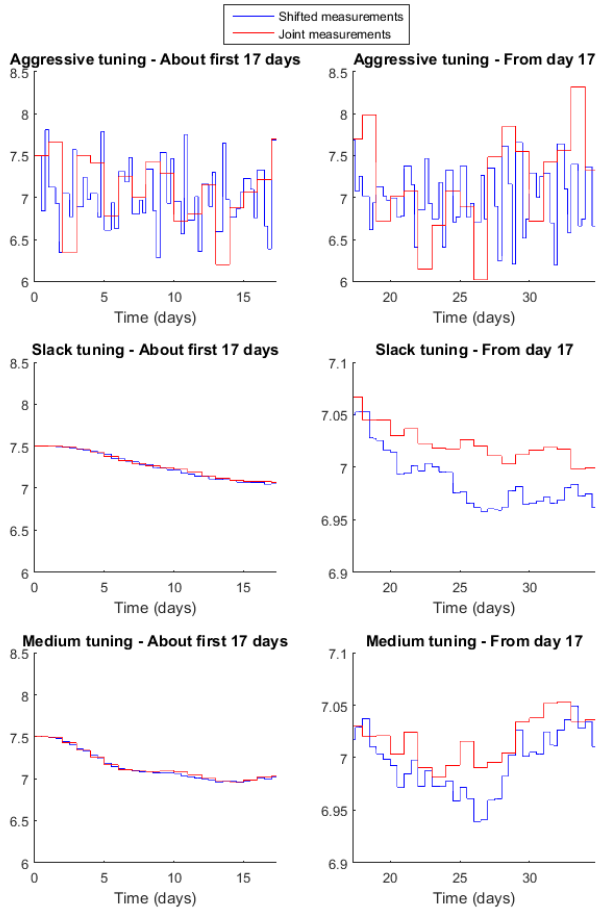

Figure 5.7: Estimation of h_{top-0} towards a reference value of 7 for 3 different tunings with noisy measurements. The left plots show the first 17 days of the simulation while the right plots show the remaining simulation time. Notice that the scale of the y-axis for the bottom two right plots is different from the rest

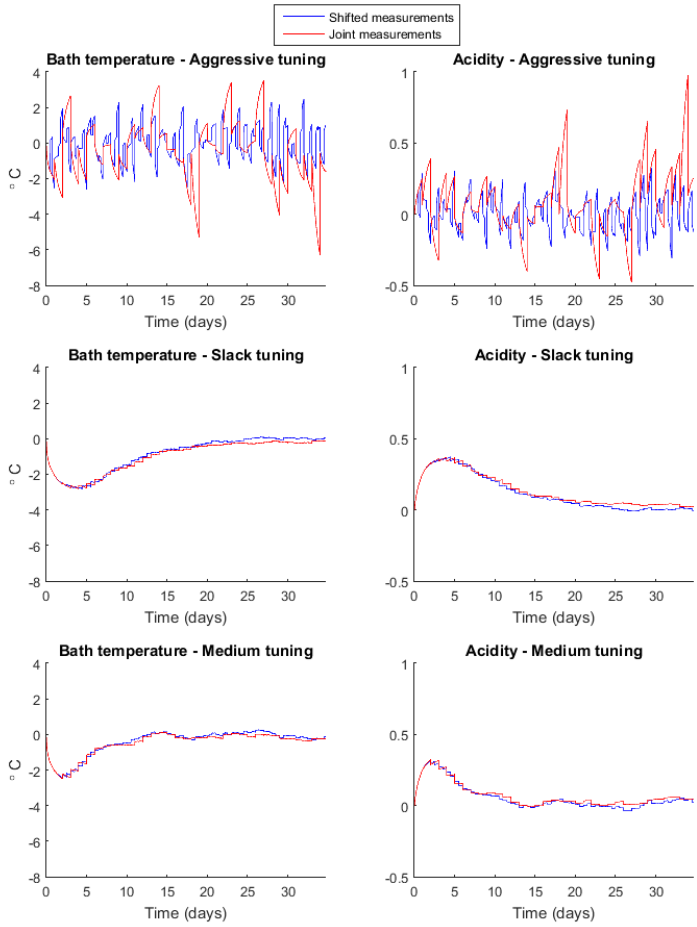

Figure 5.8: Deviation between noise-free measurements and estimator output for bath temperature and acidity measurements using three different Kalman filter tunings. Estimation done with noisy measurements

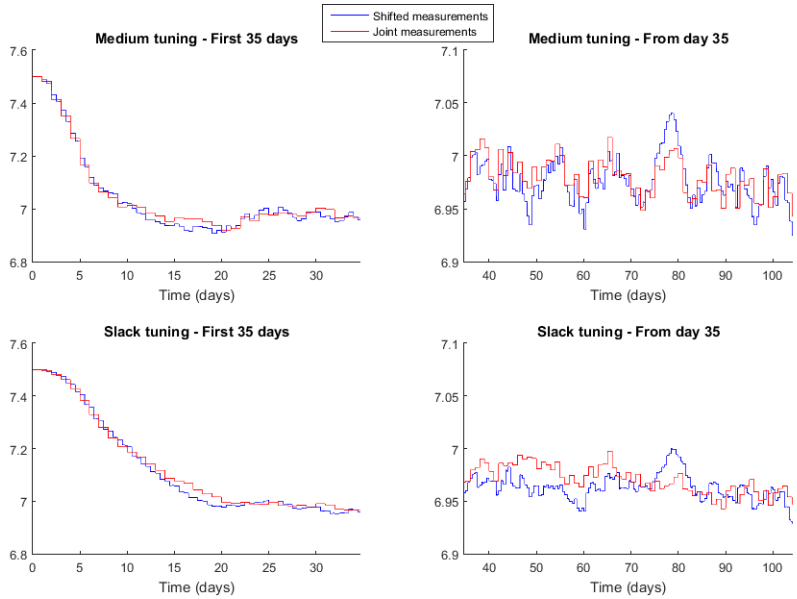

Figure 5.9: Parameter response for a process where the states vary a lot. The left plots show the initial step of the parameter while the right plots show the finer variations after the value has closed in on its correct value

Table 5.4: Estimator performance results for days 35-105 of the long simulations shown in Figure 5.9

Tuning	Bath temperature [°C]		Acidity [wt%]		Bath height [m]	
	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.
Medium tuning, joint meas.	0.005	0.084	0.007	0.014	6.34e-05	1.77e-04
Medium tuning, shifted meas.	-0.081	0.476	0.017	0.061	-3.64e-05	6.38e-04
Slack tuning, joint meas.	0.010	0.058	0.006	0.009	1.16e-05	1.18e-04
Slack tuning, Shifted meas.	0.045	0.055	0.001	0.010	9.02e-05	1.35e-04

ments give both a mean and a standard deviation that is many times better for both bath temperature and acidity. The numerical values are however very small for both tunings, and well within the anticipated magnitude of the noise. This means the variations can just as well be caused by noise as an actual performance difference. This is also backed up by the results from the slack tuned system, where the variance is close to identical and there is a very small mean value for both systems.

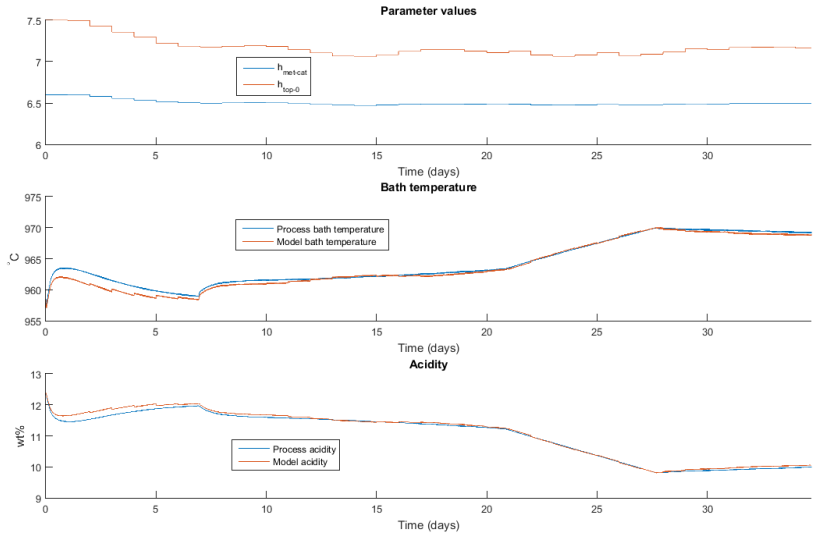

Figure 5.10: Estimation of two parameters in the heat balance at once. The reference value is 7 for both parameters. It is clear that the system manages to find another combination of parameter values that closes the

The investigation has in this case been limited to being used on one parameter. An interesting way of extending the testing would be to look if estimation of more parameters and with additional relevant measurements would give a different result. Initial tests on estimating two different heat transfer coefficients with the current measurements indicated that there might be a problem with the parameter identifiability for some combinations of more than one variable. Figure 5.10 clearly shows that the estimator manages to find a different combination of the two parameters that also gives close to 0 deviation between the process and model responses. Because of these results, estimating on two parameters in the heat balance at once has not been investigated further. Such an identifiability problem could possibly be fixed by adding e.g. additional temperature measurements like the cathode temperature measurement. This would give the estimator an indication of how much of the heat that actually disappears in a specific direction from the bath.

5.2.5 Conclusion

The obtained simulation results all point towards the same conclusion. If the number of measurements is the same, and the time between each instant of the same measurement is constant, it does not seem to make a difference if the measurements are taken all at the same time or if they are spread out. The difference in performance seems to be minimal. The variations that are observed is more likely caused by the difference in noise on the measurements when each measurement is taken than by an actual performance increase caused by the choice of measurement strategy.

5.3 Varying measurement frequencies

5.3.1 Problem description

The aluminium electrolysis process has a lot of different measurements. These offer varying precision, which can to some extent be affected by both the quality of equipment and the skill of the operators. At the same time, it is possible to affect how often each measurement is taken. The following will test how varying the measurement intervals changes the optimal tuning of the Kalman filter in addition to how much the estimator performance is affected by it.

Simulations are run to decide the optimal tuning of the Kalman filter. The different simulated process outputs which are used as measurements (with additional noise added) throughout section 5.3 are given in Figure 5.11. As can be seen by the figure, there is one short and one longer simulation. The long one is used when more data is required to get a good result. The optimal performance is decided by looking at the mean and standard deviation of the difference between the process outputs and the estimator output estimates once the process parameters have closed in on the correct process value. The convergence rate of the parameters can also be used as an evaluation criterion if the performance does not differ much for the two other factors.

The investigation was limited to looking for trends in how performance changes with measurement frequency. What this means is that there has not been made any attempt at finding a perfect tuning for a realistic setup where the measurement noise, process noise and measurement intervals have been adjusted according to a real cell. This is because such a tuning would be very system specific, which means it would probably need to be changed anyway if it was to be used on a real cell with other properties than the model. To include some additional uncertainty compared to the tests in Section 5.2, estimation has here also been added on the metal tapping parameter c_{tap} of equation 4.22 in addition to that of h_{top-0}. The offset in this parameter is supposed to simulate the fact that there might be a bias in the weight used to weigh the tapped metal.

5.3.2 Method

The tests are done by varying the measurement frequencies of some different combinations of measurements. Simulations are then used to find the choice of measurement noise parameters (r_i) that gives the best performance for each setup. First, the performance is

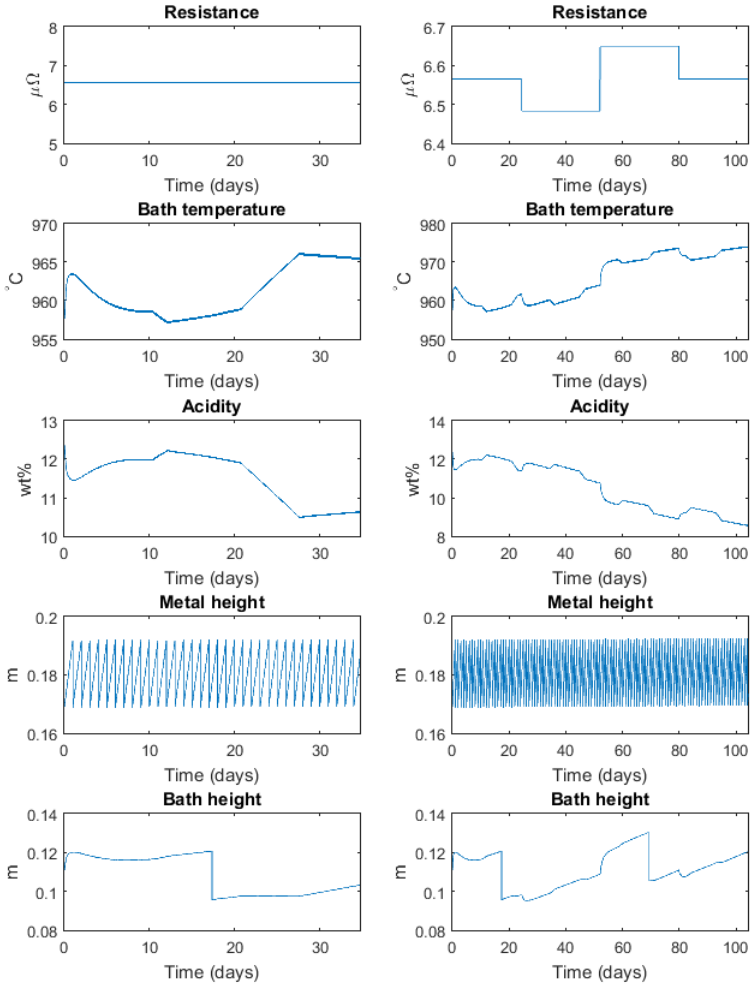

Figure 5.11: Process measurements without noise used throughout section 5.3. The process output responses of the shorter simulations are given on the left while the longer ones are given on the right side

tested when the measurement interval of the bath temperature is varied. The periods between each measurement is chosen to be half a day, 1 day, 2 days and 4 days. The same periods are then tested for the metal height separately, and in the end for a combination of bath temperature, bath height and acidity. The reason for using these two separate measurements in addition to the given combination is that the bath temperature alone shows how a change in the measurement frequency of one important measurement affects the performance when the rest of the measurements that affect that parameter are still kept constant. At the same time, the metal height is the only measurement with a significant impact on the model state representing the metal mass, and hence the metal height output of the estimator. That means it is possible to see how the change in measurement interval affects an estimate that is solely affected by that one measurement. The last combination of measurements is chosen because it includes all the states that usually impact the estimation of h_{top-0}. The bath height measurement does however not have a great impact on the estimation of this heat transfer coefficient, so changing r_{BH} has not been used while trying to find the best possible tuning for the filter.

Three criteria are chosen to evaluate which simulation shows the best performance. These are the convergence rate, mean and standard deviation. The last two are only evaluated after the initial step in parameter value. Since the aluminium electrolysis process has rather slow dynamics, a slow Kalman filter will be able to follow the changes of the process quite well without lagging behind. That means the time it takes for the parameters to get close to the correct value is not as critical as giving a good replication of the state once the parameter value has closed in on it. Evaluating the estimation performance on the relevant parameter deviation instead of on the measurement deviation is found to give the same results since the results will usually only vary on one parameter at a time.

The simulation testing is performed in stages to decide what is a good tuning in each case. First, a wide spectra of r_i-values are tested to get an implication of what area the best tuning should be in. Then, a second round of simulations is performed to find the best performing tuning in this smaller area. The performance criteria above are then used to look at the performance difference among the best candidates of the above simulations. In some of the simulations, it was necessary to extend the measurement time. That gave other input values and also a slightly different maximum value for the noise on each measurement. The maximum absolute values for the noise on each measurement for both of the measurement series is given in Table 5.5. The resistance is not included in the table because the noise value is not varied. Its noise is fixed to 0.0657 μΩ because it does not help in estimating the heat and metal mass parameter to the same extent as the other measurements.

Table 5.5: The maximum absolute value of noise on different measurements in section 5.3

Simulation type	Bath temperature [°C]	Acidity [wt%]	Metal height [m]	Bath height [m]
Short	0.96	0.12	0.09	0.006
Long	0.96	0.10	0.09	0.006

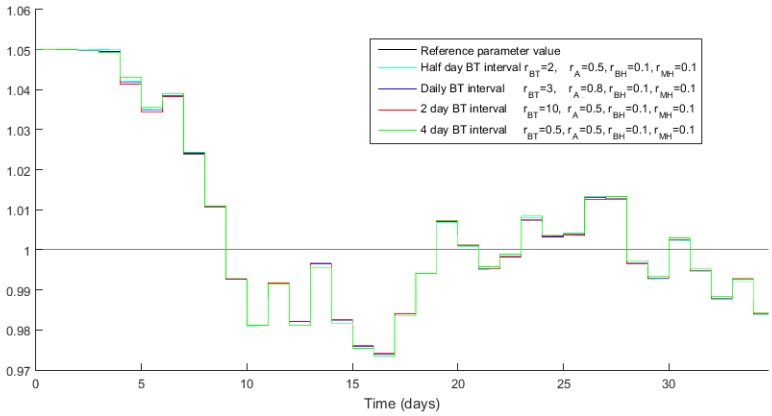

Figure 5.12: Response in c_{tap} from varying the $\Delta t_{m,BT}$ and the tuning of bath temperature, acidity and bath height. There are little variations as long as r_{MH} and $\Delta t_{m,MH}$ are kept constant. The legend shows the measurement frequency and the measurement noise tuning used for the simulations

5.3.3 Results

Varying bath temperature measurement frequency

The first step in finding the ideal tuning for a given setup is to run simulations with a lot of different tunings to narrow down the area in which the optimal tuning might lie. Figure 5.13 shows some different results from this first round of simulations when the bath temperature measurement frequency ($\Delta t_{m,BT}$) is varied. Varying the $\Delta t_{m,BT}$ will not have much of an effect on the estimation of c_{tap}. This can be seen in Figure 5.12. As long as r_{MH} and $\Delta t_{m,MH}$ are kept constant, the estimation of c_{tap} will be close to unaltered from changing the tuning of r_{BT}, r_A and r_{BH} in addition to their measurement intervals. Therefore, performance is here evaluated on the estimation of h_{top-0} alone. In the legend of Figure 5.13, r_{BT}, r_A, r_{BH} and r_{MH} describe the used value for the standard deviation of the covariance for the four infrequent measurements in the Kalman filter. This is better explained in the introduction to this chapter. The resistance measurement was not found to have an impact on the simulation results, so the value of r_R is kept constant at 0.1 for all simulations. The results obtained in Figure 5.13 indicates that a r_{BT} of 0.5 is too small for all choices of $\Delta t_{m,BT}$. With a r_{BT} of 10, the BT does not seem to have much impact on the estimation. All the different measurement frequencies show approximately the same parameter behaviour, which indicates that the BT measurement has become so uncertain that the other measurements are prioritised for estimating h_{top-0}. For the half day interval, there is some small changes in the parameter value because of the BT measurement, but it does not look like it changes the quality of the estimate to a large degree. With this tuning, the parameter always seem to have a bias compared to the reference. This bias does not seem to appear to the same degree for the other two tunings. There seems to be some more variation in the estimates when using a r_{BT} of 2 than when it is 10. Because of this, tunings in the area between a r_{BT} of 2 and 10 were tested for fine tuning. More specifically, all

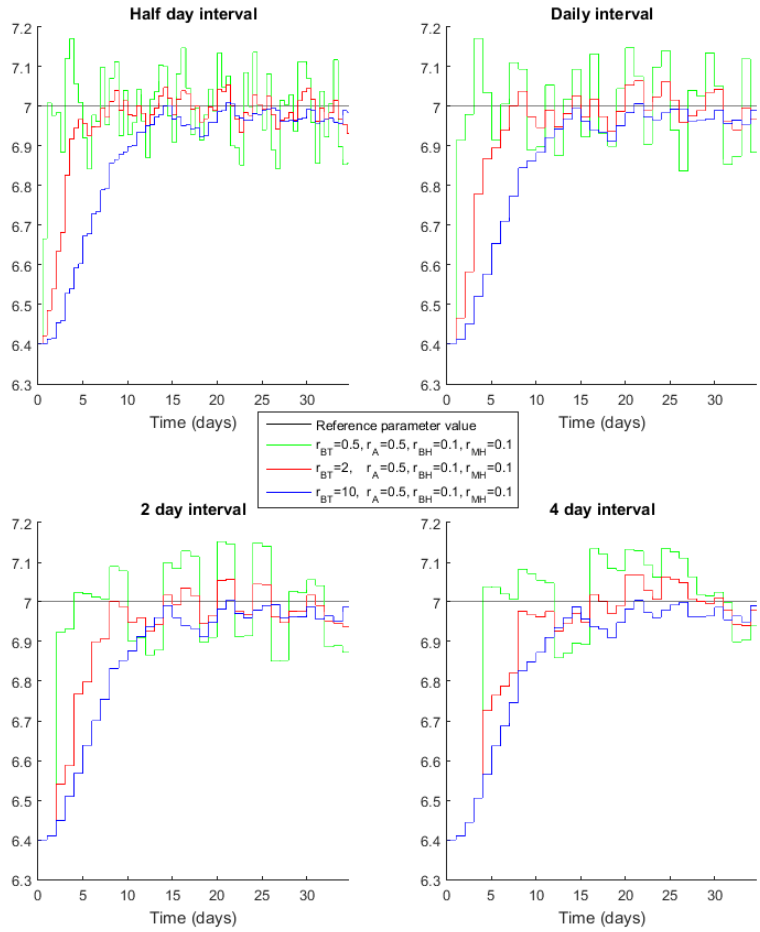

Figure 5.13: Simulation results for estimating h_{top-0} with different choices of r_{BT} tunings when $\Delta t_{m,BT}$ is varied. The legend shows the measurement noise tuning used for the simulations

Table 5.6: Mean and standard deviation for the difference between the estimate of h_{top-0} and its actual process value for the same tunings as in Figure 5.14

Interval	Tuning	Mean	St. dev.
Half day	$r_{BT} = 3, r_A = 0.8, BH = 0.1, r_{MH} = 0.1$	-0.0067	0.0218
	$r_{BT} = 5, r_A = 0.8, BH = 0.1, r_{MH} = 0.1$	-0.0152	0.0142
	$r_{BT} = 7, r_A = 0.8, BH = 0.1, r_{MH} = 0.1$	-0.0230	0.0124
Daily	$r_{BT} = 3, r_A = 0.8, BH = 0.1, r_{MH} = 0.1$	-0.0083	0.0297
	$r_{BT} = 5, r_A = 0.8, BH = 0.1, r_{MH} = 0.1$	-0.0159	0.0209
	$r_{BT} = 7, r_A = 0.6, BH = 0.1, r_{MH} = 0.1$	-0.0260	0.0204
2 day	$r_{BT} = 3, r_A = 0.4, BH = 0.1, r_{MH} = 0.1$	-0.0198	0.0292
	$r_{BT} = 5, r_A = 0.5, BH = 0.1, r_{MH} = 0.1$	-0.0265	0.0223
	$r_{BT} = 7, r_A = 0.4, BH = 0.1, r_{MH} = 0.1$	-0.0280	0.0256
4 day	$r_{BT} = 3, r_A = 0.4, BH = 0.1, r_{MH} = 0.1$	-0.0175	0.0291
	$r_{BT} = 5, r_A = 0.5, BH = 0.1, r_{MH} = 0.1$	0.0246	0.0246
	$r_{BT} = 7, r_A = 0.4, BH = 0.1, r_{MH} = 0.1$	-0.0267	0.0266

combinations of a r_{BT} of 3, 5 or 7 and a value of r_A (The acidity measurement noise parameter) of 0.3, 0.4, 0.5, 0.6, 0.7 or 0.8 were tested for each $\Delta t_{m,BT}$. In Figure 5.14, the r_A that gives the best performance for each combination of r_{BT} and $\Delta t_{m,BT}$ can be seen. It is clear that there is not a lot of difference between the performance of these tunings. Numerical values for the mean and standard deviation between all the parameter estimated plotted in Figure 5.14 and the reference value during the second half of the data set can be found in Table 5.6. For the half day interval, choosing whether $r_{BT} = 3$ or $r_{BT} = 3$ is best is quite hard. While $r_{BT} = 3$ gives a smaller mean value for the last half of the data set, the parameter value also varies quite a bit more than the other two in this period. It is seen as more important to have a small mean than standard deviation. Because of this, and the fact that it provided faster estimation, a r_{BT} of 3 combined with a r_A of 0.8 was found to be the best obtained tuning for the half day interval. For the daily interval, it can be seen that the results are quite similar to the ones obtained for the half day tuning. The mean values obtained are quite alike, while the standard deviation is better for the half day interval. Again, $r_{BT} = 3$ and $r_A = 0.8$ looks to give the best result, because it has a very small mean value compared to the alternatives. $r_{BT} = 5$ and $r_A = 0.8$ is however also an alternative for the best tuning because of a rather large improvement in standard deviation compared to a r_{BT} of 3. Both for a bath temperature measurement frequency of 2 and 4 days, the tuning with a r_{BT} of 3 and r_A of 0.4 looks to give the best result. The obtained mean and standard deviation is also very similar for the two setups. It is interesting that the mean is bigger with measurements coming more often for the parameter when comparing the results for these two measuring frequencies. This might be caused by large noise on some of the measurement instants that are not included in the 4 day bath measurement, that the estimate naturally can be more off for the start of the period that is checked or it is a possibility that the bath temperature measurement actually makes the estimate worse when it come so infrequently.

In Figure 5.15, the best estimate of h_{top-0} for each measurement frequency can be

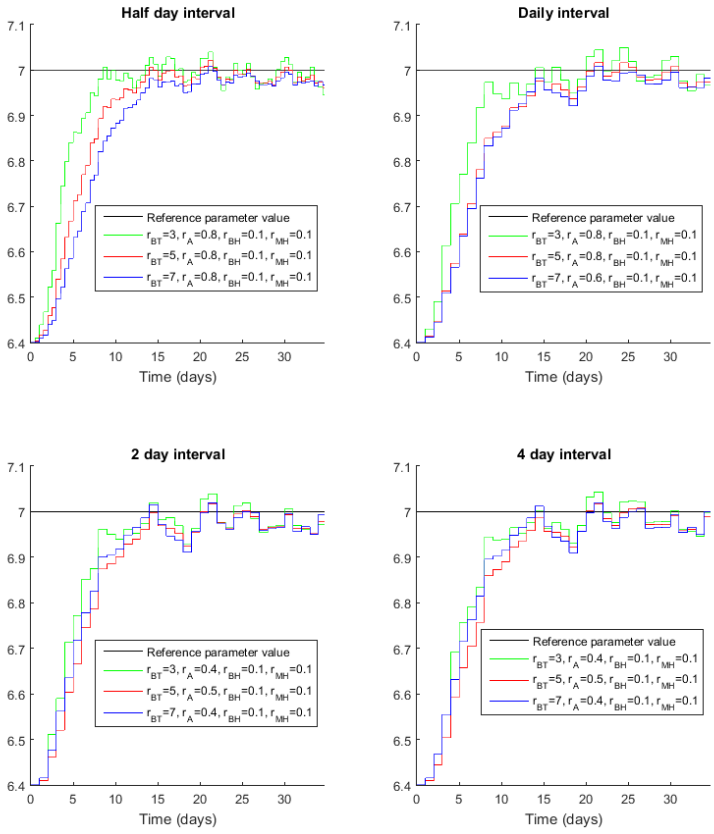

Figure 5.14: The best tuning results for each combination of $\Delta t_{m,BT}$ and the chosen r_{BT} values when estimating h_{top-0}. The legend shows the measurement noise tuning used for the simulations

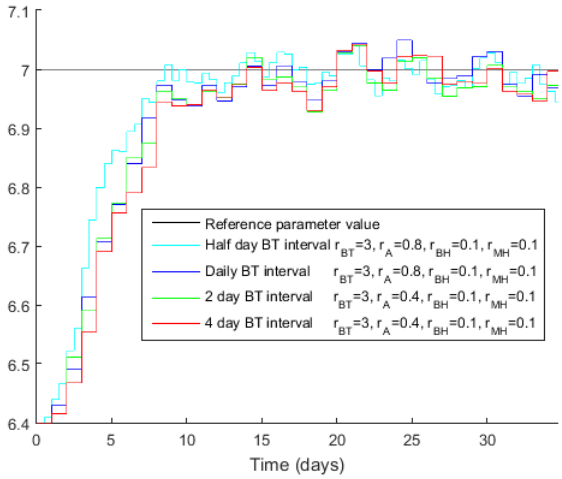

Figure 5.15: The best obtained estimate of h_{top-0} for each $\Delta t_{m,BT}$. The legend shows the measurement noise tuning used for the simulations

Table 5.7: Mean and standard deviation between the measurement estimates of the model and the noise-free process measurements for the tunings of Figure 5.16. The analysis is done for the last half of the data sets. Units for mean and st. dev.: Bath temp. $[^{\circ}\mathrm{C}]$, Acidity $[\mathrm{wt}\%]$, Bath height $[\mathrm{m}]$

Measurement	Interval	Mean	St. dev.
Bath temperature	Half day	0.0284	0.1326
	Daily	-0.0477	0.1884
	2 day	0.0721	0.1854
	4 day	0.0191	0.1995
Acidity	Half day	-0.0027	0.0196
	Daily	0.0097	0.0274
	2 day	-0.0080	0.0266
	4 day	-9.4459e-05	0.0287
Bath height	Half day	-2.6695e-05	1.7642e-04
	Daily	-1.0737e-04	2.4669e-04
	2 day	7.6547e-05	2.3813e-04
	4 day	1.1532e-08	2.6012e-04

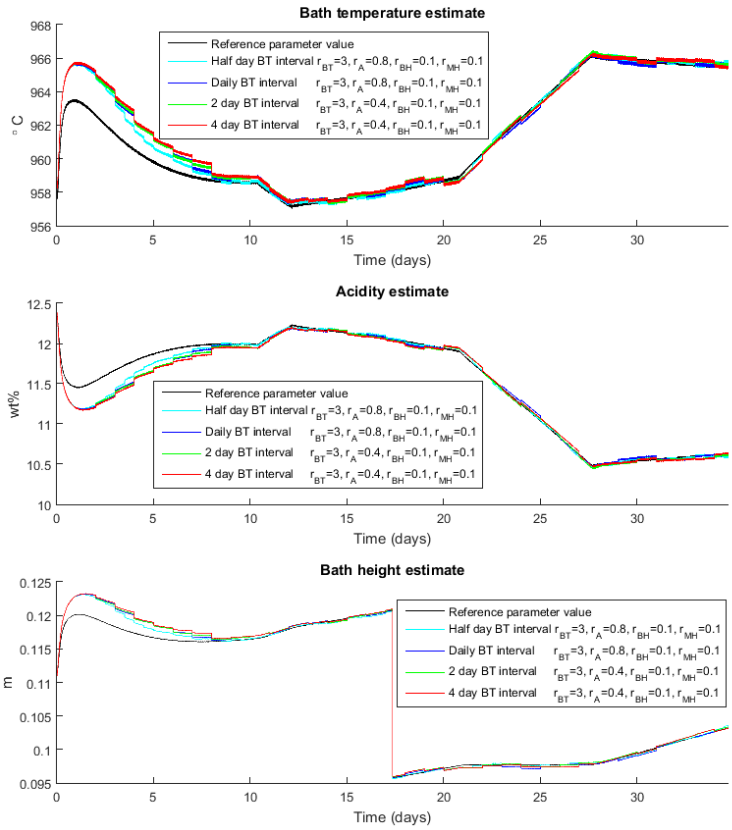

Figure 5.16: Measurement estimates for the tunings of Figure 5.15. The legend shows the measurement noise tuning used for the simulations

found. For each of the setups, the bath temperature, bath height and acidity estimates are displayed in Figure 5.16, while their mean and standard deviation are given in Table 5.7. The results show that the effect of varying $\Delta t_{m,BT}$ is very small, both in mean value and standard deviation. For the half day measurement interval, the deviations are smaller and the estimation is faster, but overall, there is very little deviation in all results. This might indicate that with proper tuning, the other measurements manages to compensate for the reduced frequency of BT measurements.

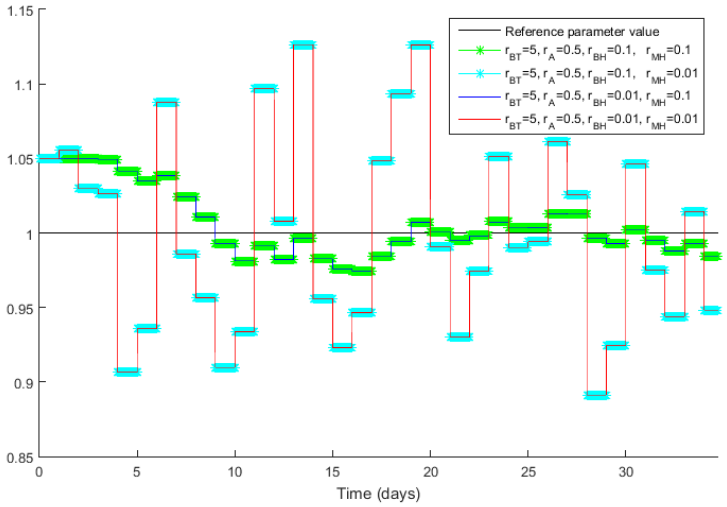

Figure 5.17: The bath height tuning doesn't affect the estimation of c_{tap}. The legend shows the measurement noise tuning used for the simulations

Varying metal height measurement frequency

The metal height measurement is the one primarily used by the Kalman filter to estimate c_{tap}. That means that the relevant parameter to estimate while altering the metal height measurement frequency ($\Delta t_{m,MH}$) is c_{tap}. As both metal height and bath height are measured using the same measuring technique, it is reasonable to also choose the same tuning for both of them. That means that even though the interesting relation to look at here is how the optimal tuning and estimation performance changes with the frequency of the metal height measurements, it is reasonable to also change the tuning of the bath height measurement noise parameter (r_{BH}) along with it. This is of course not a necessity, and it will alter the estimate of h_{top-0}, but the small change this has on the heat balance will not be relevant for this discussion. Figure 5.17 shows that altering the bath height tuning has an negligible impact on the estimation of c_{tap}. The plot legends throughout the discussion are described in the introduction of the chapter.

To establish approximately what region the optimal metal height tuning is located in, some spread out initial tests were performed. Some of the results from these tests are shown in Figure 5.18. The initial tests seem to indicate that setting the metal height measurement noise parameter (r_{MH}) to 0.01 is too low for all simulations, as it means the MH measurement is trusted so much that the estimate is very much affected by measurement noise. For the 4 day interval, it can look like a r_{MH} of 0.1 might be too high as it shows a significant bias. The problem can however be that the estimator is not given enough time to correct for the bias since there are only 8 available metal height measurements during the simulation time with this $\Delta t_{m,MH}$. For the rest of the measurement frequencies, there seems to be little to no bias, and that the amount of variance gets smaller as r_{MH} gets

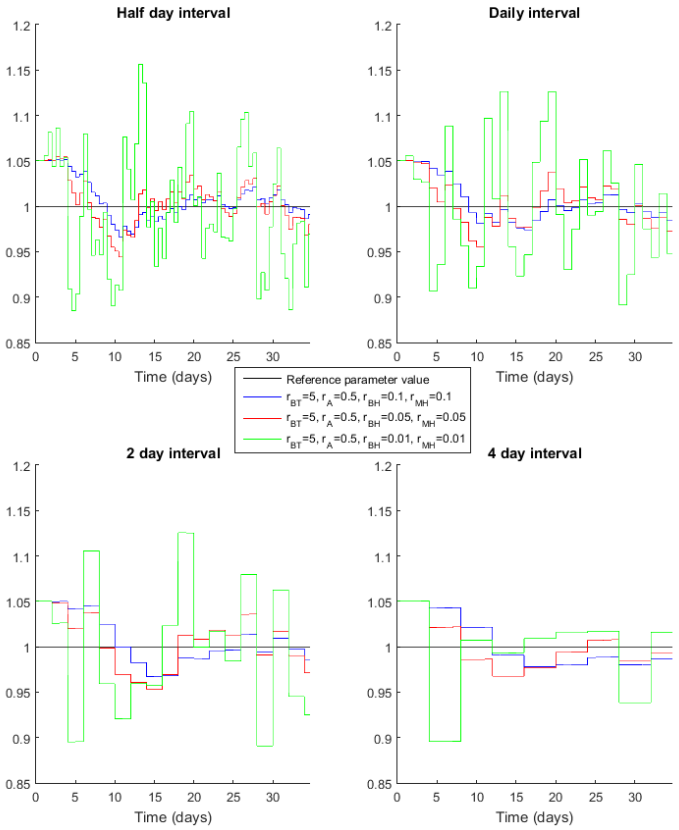

Figure 5.18: Wide spread tuning results for estimation of c_{tap} when $\Delta t_{m,MH}$ is varied. The legend shows the measurement noise tuning used for the simulations

bigger.

Some additional simulations were performed for all measurement intervals around a r_{MH} of 0.1 as the parameter seemed to close in on the correct value with that tuning and as it showed the least variance. The results from these simulations provided two important insights. All simulations seem to crash when r_{MH} gets too large. It is believed that this happens because the covariance keeps increasing if the gap between the measurements and their predictions is not closed. When the measurements are trusted to little (r_i is too high) or the measurements are too far apart, the covariance will increase to a level where it causes numerical problems for the solver. This happens for a r_{MH} of about 0.4 for the half day interval, 0.3 for the 1 day interval, 0.2 for the 2 day interval and 0.1 for the 4 day interval. In these simulations, Euler's method was used. This is the simplest solver available in Modelfit. It was used because it was seen as sufficiently accurate and as it is faster than

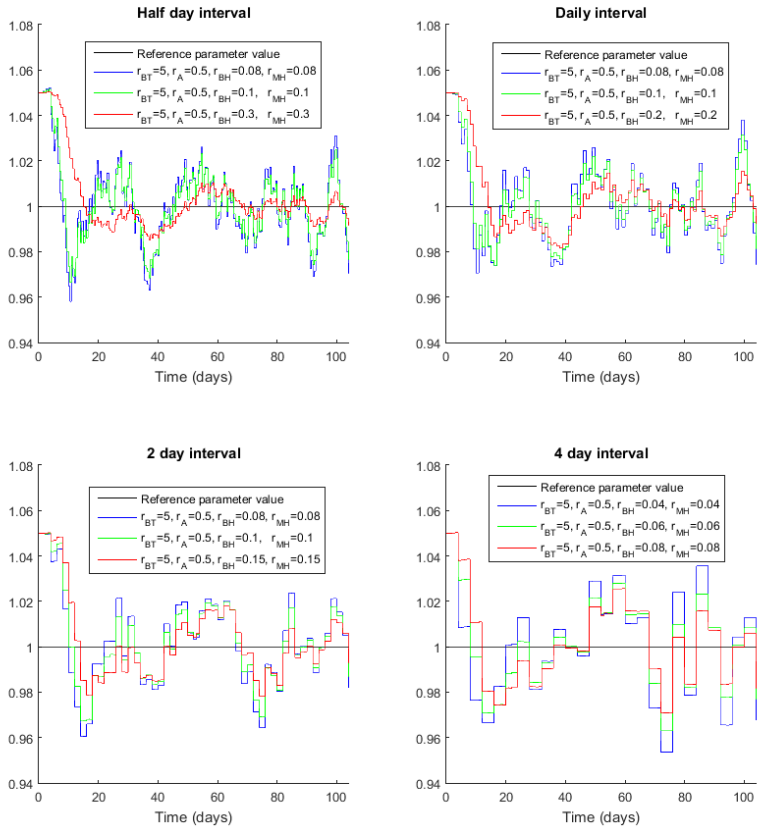

Figure 5.19: Longer simulations for estimation of c_{tap} with varying interval for the metal height measurements. The legend shows the measurement noise tuning used for the simulations

more intricate solvers. It is possible that the numerical issues could be avoided by using a different solver. The second insight is that it is hard to conclude which simulations are best in this time frame. To get a better look at the parameter estimate once it had closed in on the correct value, some longer simulations were performed.

Some of the results from these simulations can be found in Figure 5.19. These graphs show that the estimate will oscillate around the correct value for all the chosen tunings, but this effect seems to be smaller for larger values of r_{MH}. What this means is that the best values are the largest ones that does not crash the simulation. It would then be interesting to see if another solver would make it possible to run simulations with larger values for r_{MH}. The plots clearly show that the variations are smaller for the half day interval because it is possible to use a larger r_{MH} here. In Figure 5.20, the parameter estimate for c_{tap} is shown with identical tuning for the different measurement frequencies.

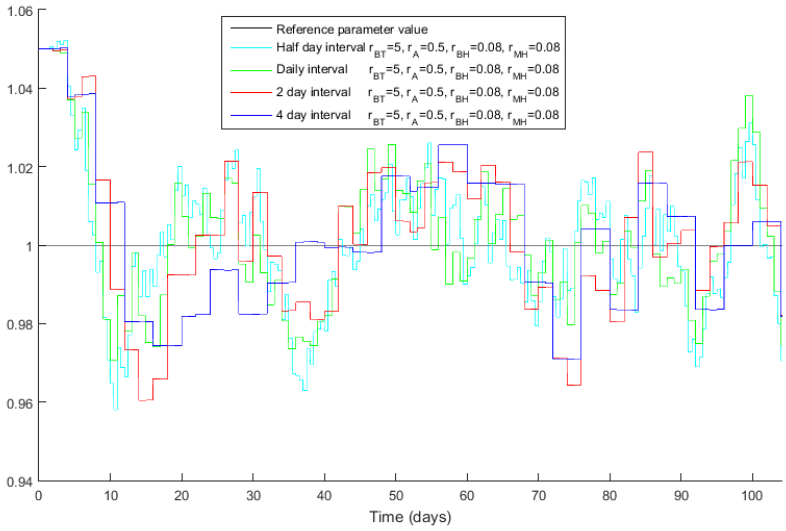

Figure 5.20: Comparison of estimator performance with the same tuning for different metal height measurement frequencies. The legend shows the measurement frequency and the measurement noise tuning used for the simulations

The frequency does not seem to have a large affect on the parameter estimate in this case. The parameter estimates seem to vary about the same amount for all frequencies. The variations for the longer intervals do however look like they are slower, so there are less peaks for these estimates. This could indicate that this estimate will be slower to adapt to a change in the process conditions. The overshoot and extra time needed for the parameter estimate to close in on the reference parameter value in Figure 5.20 is also an indication of this.

The deviation between the Kalman filter's best estimate of the metal height for each $\Delta t_{m,MH}$ and the reference metal height state of the process (without noise) is shown in Figure 5.21. Further, Table 5.8 displays the standard deviation and mean values of these estimates. The values in the table are calculated from day 35 and out because this gives

Table 5.8: Mean and standard deviation for the difference between the estimated metal height and the actual process value. As the parameter needs some time to converge towards the correct value, the mean and standard deviation are calculated from day 35. Unit for mean and st.dev.: Metal height [m]

Interval	Tuning	Mean	St. dev.
Half day	$r_{BT} = 5, r_A = 0.5, r_{BH} = 0.3, \quad r_{MH} = 0.3$	-2.7681e-04	0.0016
Daily	$r_{BT} = 5, r_A = 0.5, r_{BH} = 0.2, \quad r_{MH} = 0.2$	5.8526e-04	0.0021
2 day	$r_{BT} = 5, r_A = 0.5, r_{BH} = 0.15, r_{MH} = 0.15$	5.2624e-04	0.0030
4 day	$r_{BT} = 5, r_A = 0.5, r_{BH} = 0.08, r_{MH} = 0.08$	0.0010	0.0036

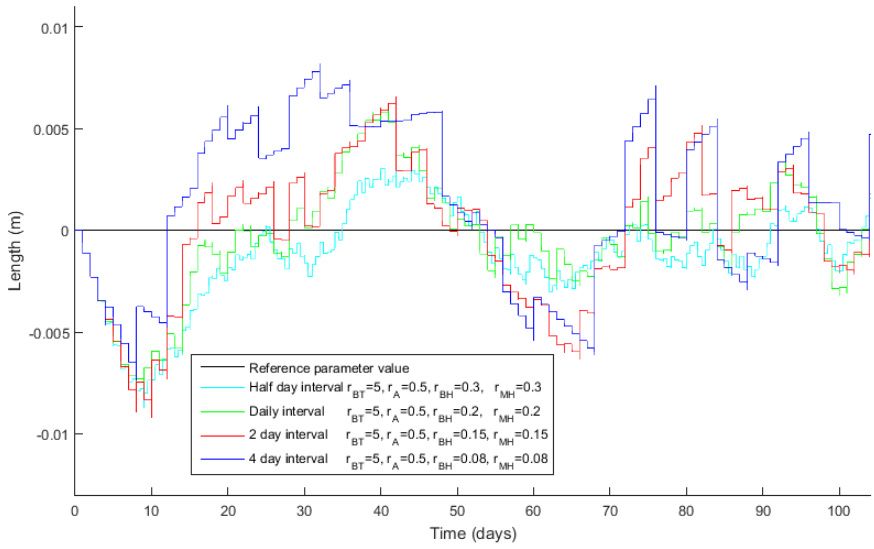

Figure 5.21: Comparison of estimation performance for the best tuning obtained for each estimate. The legend shows the measurement frequency and the measurement noise tuning used for the simulations

the parameter some time to close in on the correct parameter value. It is clear from these results that the mean values are very good for all estimates, where the largest deviation in mean value is 1 mm. The largest deviations are shown in Figure 5.21 to be less than 1 cm, which should be considered good given that the uncertainty of the measurements used to estimate the value are of the same magnitude. The standard deviation is more than twice as large for the 4 day interval compared to the half day interval. Other results could however be possible if it is possible to prevent the simulations from crashing, so it is hard to draw any certain conclusions regarding the performance of the Kalman filter based on the obtained results, other than that with the given solver, the measurement frequency affects what tuning it is possible to choose for the filter. This again seems to affect on the performance of the estimator.

Table 5.9: Mean and standard deviation of the difference between the estimate of h_{top-0} and its actual process value during the last two thirds of the data series (from approximately day 35). Tunings are the same as in Figure 5.23

Interval	Tuning	Mean	St. dev.
Half day	$r_{BT} = 2, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0102	0.0317
	$r_{BT} = 3, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0141	0.0228
	$r_{BT} = 4, r_A = 0.8, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0158	0.0207
Daily	$r_{BT} = 2, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0016	0.0352
	$r_{BT} = 3, r_A = 0.8, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0035	0.0286
	$r_{BT} = 4, r_A = 0.8, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0096	0.0218
2 day	$r_{BT} = 2, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0037	0.0381
	$r_{BT} = 3, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0095	0.0257
	$r_{BT} = 4, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0133	0.0249
4 day	$r_{BT} = 2, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0056	0.0534
	$r_{BT} = 3, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0135	0.0432
	$r_{BT} = 4, r_A = 0.4, r_{BH} = 0.08, r_{MH} = 0.08$	-0.0189	0.0404

Varying bath temperature, acidity and bath height measurement frequencies

In the following, the measurement frequency of all three measurements that have an effect in estimating h_{top-0} will be varied at the same time. That means both the bath temperature measurement frequency ($\Delta t_{m,BT}$), acidity measurement frequency ($\Delta t_{m,A}$) and bath height measurement frequency ($\Delta t_{m,BH}$) will be varied at once. The measurement frequency will stay the same for all of them, and the same frequencies as in the previous discussion will be investigated. Firstly, a wide set of tunings are checked. Some results from these initial tests can be found in Figure 5.22. For a explanation of the plot legend, see the introduction to the chapter. These results indicate that if the r_{BT} is used as the primary tuning variable and the r_A as the fine tuning parameter, the optimal tuning is achieved for a r_{BT} larger than 1, as this value causes quite a lot more variance than the other presented values. For the half day and daily intervals, it does also look like a r_{BT} of 5 causes the estimate to have some bias. It is hard to evaluate if this same bias is present for the 2 day and 4 day intervals because they should have some more time to settle. An important observation here is that getting measurements more frequently clearly improves the time the parameter uses on the initial step in parameter value.

To find an optimal tuning for each of the measurement frequencies, additional simulations were performed where r_{BT} was given a value of 2, 3 and 4, as 3 looked to be a good compromise between mean value and standard deviation for the initial simulations. For each of these values for r_{BT}, 0.4, 0.8 and 1.2 were tested as the r_A. The best resulting combination of measurement frequency and r_{BT} is displayed in Figure 5.23. In order to better see how the parameters performed after the initial step in parameter value, it was found reasonable to increase the simulation time. To evaluate the performance of these plots, the mean and standard deviations for the last two thirds (from approximately day 35) of each data set has been calculated. The resulting values are given in Table 5.9. It is a clear trend in the data that the standard deviation for all choices of r_{BT} increases

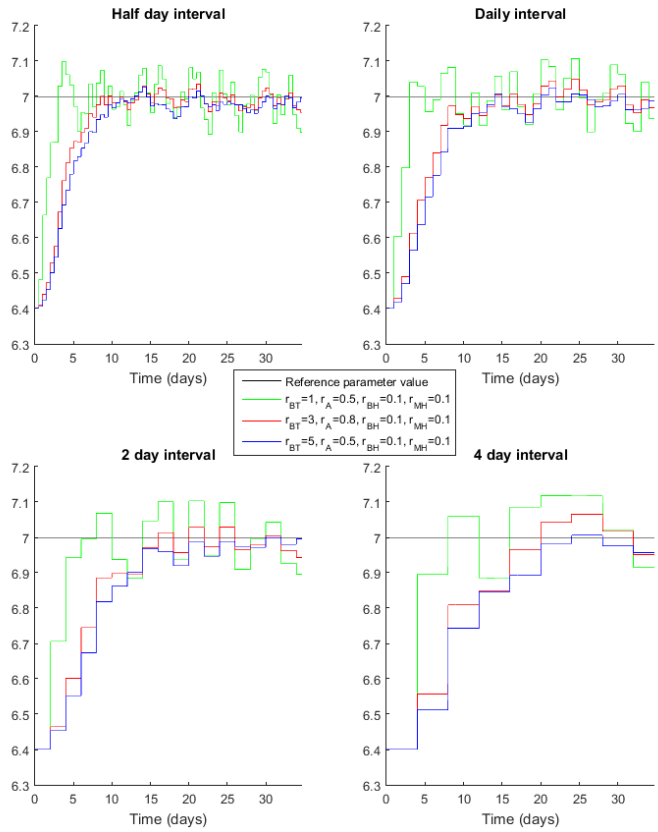

Figure 5.22: Wide spread tuning results for estimation of h_{top-0} when $\Delta t_{m,BT}$, $\Delta t_{m,A}$ and $\Delta t_{m,BH}$ are varied simultaneously. The legend shows the measurement noise tuning used for the simulations

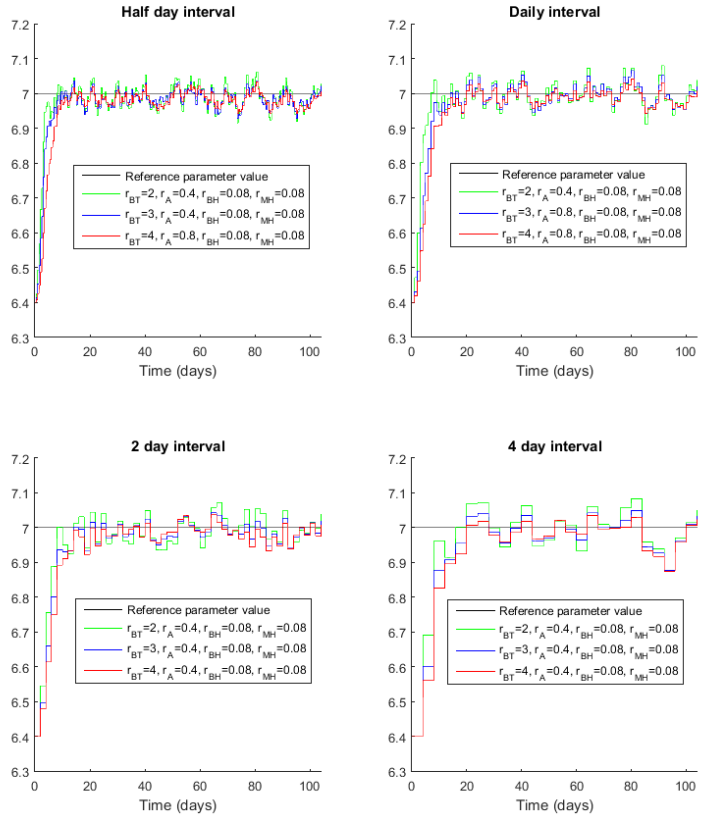

Figure 5.23: Fine tuning results for estimation of h_{top-0} when $\Delta t_{m,BT}$, $\Delta t_{m,A}$ and $\Delta t_{m,BH}$ are varied simultaneously. The legend shows the measurement noise tuning used for the simulations

along with the measurement intervals. The only exception is that it decreases between the daily and 2 day intervals when the r_{BT} is set to 3. Such an exception can probably be explained by excessive noise on some of the daily measurements that are avoided when taking measurements every second day. There does not seem to be a pattern giving how much the standard deviation changes with a change in measurement intervals. Because a half day interval has more measurements, the possibility of encountering a measurement with a lot of noise on it is higher. This can have a significant impact on the estimate for some time, which means that it is reasonable that the mean might be more off even though the measurements comes more frequently. There are no clear patterns to how the mean value changes with tuning and measurement frequency. It is however of interest to notice that the mean value is always smaller than the standard deviation. This is good since it means that the estimate will have values on both sides of the real parameter value.

Table 5.10: Mean and standard deviation between the measurement estimates of the model and the noise-free process measurements. The chosen tunings are all tunings where the r_{BT} is set to 3. The analysis is done for the last two thirds (from day 35) of the data sets. Units for mean and st. dev.: Bath temp. $[^{\circ}\mathrm{C}]$, Acidity $[\mathrm{wt}\%]$, Bath height $[\mathrm{m}]$

Measurement	Interval	Mean	St. dev.
Bath temperature	Half day	0.0297	0.1102
	Daily	-0.0413	0.1578
	2 day	-0.0211	0.1389
	4 day	-0.0250	0.2289
Acidity	Half day	3.1629e-04	0.0177
	Daily	0.0114	0.0252
	2 day	0.0093	0.0222
	4 day	0.0091	0.0380
Bath height	Half day	5.8180e-05	2.0711e-04
	Daily	-1.0294e-04	2.8365e-04
	2 day	-4.6624e-05	2.5614e-04
	4 day	-6.4010e-05	4.4760e-04

The obtained results seem to indicate that the tuning is not noticeably affected by a change in measurement frequency. There are no clear indication from the data that there is a shift in where the estimator shows the best performance. In general, it seems the estimate behaves quite similar to the same tuning, regardless of what the measurement frequency is. An example is how the peaks in the estimate has a large increase of about the same magnitude for all measurement frequencies when a r_{BT} of 1 is chosen, as shown in Figure 5.22.

It is hard to choose one tuning from Table 5.9 for each measurement frequency that is clearly better than the other tunings. Sometimes at least the mean value calculated can be off because of a large deviation for a short amount of time during the simulation. At the same time, there is a difficult trade-off between mean and standard deviation that means it is not straight forward to decide which result is best. To look at how much the measurement estimates deviate from the simulated process measurements, the mean and standard deviation of the bath temperature, acidity and bath height measurements have been calculated. It was chosen to look at all tunings from Figure 5.23 with a r_{BT} of 3. This was done because it is hard to choose one tuning that is without doubt better than the others for each measurement interval. It is for all of them a trade-off between minimising the mean and standard deviation. The results are shown in Table 5.10. All of the estimates seem to have both a mean value and standard deviation with an absolute value that is clearly smaller than the maximum value of the noise applied on the measurements. This indicates that the estimates are good because it means that the estimator will in general achieve about the same precision as the measurement equipment. The estimator's performance would probably not be as good if the reference process would be a real process with behaviour that is not covered in the model. By comparing the results of Table 5.10 to their equivalent values from varying only the frequency of the bath temperature measurement, it can be seen that the differences are small. The best obtained tuning does also seem to be quite similar for

the two cases, which might be another indication that the best tuning is not significantly affected by the measurement frequencies.

5.3.4 Discussion

All simulations have confirmed that the larger the measurement noise parameter is put on a measurement, the less that measurement is used to update the parameter value. On one hand, this tends to cause less variance on the parameter estimate, which is a good thing. There will not be large changes to the parameter estimate due to one or a few measurement instant(s) with a lot of noise on it. The negative effect of a slow measurement update is that estimation gets slower. It can mean that the estimator will use more time to come back if the parameter should drift for some reason. Another negative effect can be that the estimate can get a bias.

When changing the measurement frequency of the bath temperature measurement alone, there was not found to be a massive improvement in the estimator performance. The best obtained tuning was found to be the same for the half day and 1 day intervals and for the 2 and 4 day intervals. It can be seen in Figure 5.15 that the r_A is trusted more when the measurement frequency is decreased. The reason why there is not a large change in performance can then be that when there are fewer BT measurements available, the estimator can compensate by trusting the other measurements, which are still obtained daily, more.

The results from varying $\Delta t_{m,MH}$ show that measuring the metal height more often will make the simulation more robust, allowing for higher values of r_{MH} without crashing the simulation. With identical tuning, there was not found to be a big difference between the different measurement intervals. It would therefore be interesting to see if another solver would allow for larger values for r_{MH} in each case, which could lead to a different result.

Varying the measurement frequencies of both bath temperature, acidity and bath height simultaneously have removed the possibility that the measurements that are still kept at the same frequency might compensate for the lack of one measurement when estimating h_{top-0}. The obtained results show that there will be some improvement in performance from having more frequent measurements. For the given setup, the difference is however not found to be very large. The mean and standard deviations of the difference between the estimated and process measurements are for all measurement frequencies found to be smaller than the maximum absolute value of the noise applied to the measurements. The results are quite similar compared to the ones obtained when only the bath temperature is varied. Another observation is that regardless of the measurement frequency, the same tuning seem to have a similar effect on the estimates in changing the mean and standard deviations. This means that there was not found any indication of a shift in optimal tuning as a result of varying the measurement frequencies.

All optimal tunings give results that show less standard deviation than the maximum values of noise added to the measurements. This can be because the model is used both as the simulated process (to obtain measurements) and as the process model. That means the model is ”perfect” once the estimated parameter reaches the correct value. Then, measurement noise will be the only source of deviation between measurements and the models measurement estimates. It seems the optimal tuning is obtained by putting the

measurement noise parameters quite large. This is to minimise the impact of measurement noise on the estimates. If there is a larger deviation between the model and the process measurements, it is possible that it would be more optimal to use smaller measurement noise parameter values for the times where the model does suddenly deviate a lot from the real behaviour of the process. This would both give faster adjustments for deviations between process and model, and higher sensitivity to measurement noise. One simple way of adding such a deviation in the model would be to change the value of a parameter that is not estimated upon. This means that for the Kalman filter to make the process measurements fit the model's measurement estimates, the estimated parameter of the model would also need to get another value than what was used to create the process measurements (since the model is used both as the simulated process and as the model).

Figure 5.22 shows that the initial step is handled significantly faster when the measurement frequency is increased. This shows that it might be an idea to increase the measurement frequency during periods where an estimator experiences more deviation than usual (for an imperfect model as discussed above) and for fast initiation of the estimator. It is however reasonable to think that the improvement will be limited to how fast the dynamics of a process is. For a slow process, with a time constant of several days, the states will barely change at all during a 1 minute interval, so there is no reason for measuring this often.

5.3.5 Conclusion

The above discussion indicates that all estimates fit their measurements well. There are of course some performance improvements from measuring once every half day compared to taking measurements 4 days apart. These differences are however quite small here once the parameter values have closed in on their correct values. All results indicate that the model estimates will mostly deviate less from the process value than the maximum absolute value of the measurement noise. It is reasonable to think that this happens because a "perfect" model has been used. When testing how the tuning should be changed for estimating c_{tap}, the best results were achieved by setting r_{MH} as high as possible without the simulations crashing. A large r_i indicates that the Kalman filter will prioritise the measurements less compared to the model value when updating the estimate. If the measurement noise is the primary source of deviations, it does then make sense to trust the model more than the noisy measurements when updating the Kalman filter. None of the trends in the data indicate that there will be a significant change in how the Kalman filter should be tuned because of how often the measurement is performed.

Even though the tuning does not seem to change much with the measurement frequency, it looks like increasing the measurement frequency has a positive effect on how fast the parameter closes in on its correct value. This could be used to speed up the initialisation of the Kalman filter for an actual process, or to correct deviations faster when the model does not fit well to the actual state of the process.

Chapter 6

Overall summary, conclusion and future work

6.1 Summary and conclusion

This thesis has looked at how infrequent measurements should be used in a Multi-rate Extended Kalman filter. The reason for doing this is that there is a cost related to each process measurement. If it is possible to achieve the same estimator performance with fewer measurements, it could be possible to save money on measuring equipment. Another positive effect of performing fewer measurements is that the operators then spend less time close to hot surfaces, liquid metal, high amperage and dangerous gases. This is an advantage from a HSE-perspective.

To test the different properties, a suitable process model was developed. The aluminium electrolysis process is an industrial process where the frequency of the measurements range from continuously to once in several days, so it was found to be well suited for the task. Hence a simple model of an aluminium electrolysis cell was created. This model described many of the main aspects of the mass and energy balance of the cell, though with some limitations. The finished model was found to give a good qualitative description of the relations between the acidity, bath temperature, side ledge thickness and heat balance of the cell. The metal production and tapping did also show reasonable behaviour. One potential model improvement would be to get a good model for the relation between the resistance measurement, the ACD and the alumina concentration of the cell. In addition, it would increase the realism of the model if it would be possible to describe the effect of disturbances like the anode effect, sludge formation and anode changing.

With the model created, it was used to create measurement series with one set of parameters. This was thought to represent the process behaviour. White noise was in addition added to the measurements. The same model was then used as a process model, just with a change in one or two of the parameters. By then estimating on the changed parameters, the task was for the Kalman filter to correct the parameter values so that the model fitted

the measurements. Two different cases were tested to see how the infrequent measurements altered the performance of the Kalman filter. The first case was to look at whether it would give a difference in performance to perform all the different measurements simultaneously or if it would be more reasonable to split them up so that a small piece of information would arrive more frequently. The obtained results suggested that the choice of measurement strategy did not have a significant impact on estimation results.

The other case was to see how the optimal tuning and performance of the Kalman filter estimate would change from varying the measurement frequencies. It was found was that the standard deviation of the estimate was a little better with more frequent measurements. The difference was however smaller than expected, as both the mean and standard deviation of the difference between the measurement estimates and the noise-less process measurements were smaller than the maximum absolute value of the measurement noise. A possible explanation is that since the estimation was done with a ”perfect” model, where the only model deviation was in the parameters that were estimated on, the only source of estimate deviations was then measurement noise when the Kalman filter had managed to make the parameter value close in on its correct value. The performed simulations did not give any indications of a change in how the Kalman filter should be tuned for a different measurement frequency in one or more of the measurements.

6.2 Future work

The model worked well for testing the cases discussed above. As a general, simple model of an aluminium electrolysis cell, it is however possible to do some improvements in order for the model to describe even more of the important system dynamics like the connection between resistance, ACD and alumina concentration and the most important system disturbances.

It would be interesting to test the model with measurements from a real cell. As the measurement strategy and frequencies would then be given, it would not be possible to use such measurements to test the cases of this thesis. This could however give an indication of how well the model represents the behaviour of a real cell.

For the testing of the Kalman filter, there are some additional tests that would be interesting to perform. First of all, it would be useful to verify what has been found by testing it on a cell with different inputs, parameters and with a different amount of noise on the measurements.

A clear thing to test next for both of the cases would be to use an imperfect model in which some of the parameters are wrong, but not changed by the estimator. This could cause larger deviations between the process measurements and the model. That could potentially cause the estimates to differ more in performance, making it clearer which tuning gives the best estimate.

With an imperfect model, there will be some periods where it will experience significant deviations from the actual process state. An interesting think to investigate further would then be if it would be possible to increase the measurement frequency or make the measurement noise parameters more aggressive during periods where the model experiences large deviations. This could potentially give tight estimation during periods with much deviation between process behavior and the model, in addition to avoiding that the

estimates are prone to noise when the model fits the process measurements well.

Another potential extension would be to look at whether it is possible to create a more dynamic model by estimating on multiple parameter in the heat balance at once. As mentioned in section 5.2.4, initial tests on this indicated that the system is not always identifiable when estimating two heat balance variables at once with the available measurements. It would be interesting to add e.g. a continuous cathode temperature measurement to see if that would help the estimator in obtaining the correct value for both parameters.

Bibliography

Aalbu, J., Moen, T., Aalbu, M. S., Borg, P., Aug. 23 1988. Method of controlling the alumina feed into reduction cells for producing aluminum. US Patent 4,766,552.

aluminium institute, I., 2011. Results of the 2010 Anode Effect Survey Report on the Aluminium Industry's Global Perfluorocarbon Gases Emissions Reduction Programme.

Anovitz, L. M., Hemingway, B. S., Westrum, E. F., Metz, G. W., Essene, E. J., 1987. Heat capacity measurements for cryolite (na3alf6) and reactions in the system na fe al si o f. Geochimica et Cosmochimica Acta 51 (12), 3087–3103.

bibliography: heat capacity alumina, 2017. Accessed 2017-03-14.
URL http://www-ferp.ucsd.edu/LIB/PROPS/PANOS/al2o3.html

bibliography: heat capacity aluminium fluoride, 2017. Accessed 2017-03-14.
URL http://webbook.nist.gov/cgi/cbook.cgi?ID=C7784181&Mask=1A8F&Type=JANAFS&Table=on#JANAFS

bibliography: molar mass cryolite, 2017. Accessed 2017-03-08.
URL https://www.chemicalaid.com/tools/molarmass.php?formula=Na3AlF6

Brown, R. G., Hwang, P. Y., 2012. Introduction to Random Signals and Applied Kalman Filtering with MATLAB Exercises, 4th Edition. Wiley.

Drengstig, T., 1997. On process model representation and alf_3 dynamics of aluminium electrolysis cells. Ph.D. thesis.

Grjotheim, K., Kvande, H., Foosnæs, T., Huglen, R., Lillebuen, B., Mellerud, T., Naterstad, T., 1993. Introduction to aluminium electrolysis, 2nd Edition. Aluminium-Verlag.

Gudi, R. D., Shah, S. L., Gray, M. R., 1995. Adaptive multirate state and parameter estimation strategies with application to a bioreactor. AIChE Journal 41 (11), 2451–2464.

Hestetun, K., 2009. Use of data from anode current distribution for state and parameter estimation and fault detection in an aluminium prebake electrolysis cell. Ph.D. thesis, NTNU.

International Electrotechnical Commission (IEC), 1982. Thermocouples, part 2, tolerances. Tech. Rep. IEC 584-2.

Jessen, S. W., 2008. Mathematical modeling of a hall héroult aluminum reduction cell. Ph.D. thesis, Technical University of Denmark.

Karuppannan, C., Prabhakar, R., 2009. Analyzing of soderberg cell technology performance and possibilities analýza možností a výkonu soderbergovej bunkovej technológie.

Kolås, S., 2008. Estimation in nonlinear constrained systems with severe disturbances. Ph.D. thesis.

Kvande, H., Drabløs, P. A., 2014. The aluminum smelting process and innovative alternative technologies. Journal of Occupational and Environmental Medicine 56 (5 Suppl), S23.

Leite, M., 2016. Decision support for aluminium electrolysis, technical report, Specialization project, Departement of Engineering Cybernetics, NTNU.

Sælid, S., 1984. Forelesningsnotater i modellering av dynamiske prosesser. Tech. rep., Technical Report 84-2-X, Division of Engineering Cybernetics, NTH, Norway.(in Norwegian).

Sears, F. W., Zemansky, M. W., Young, H. D., 1987. University Physics 7th edition, 7th Edition. Addison-Wesley publishing company.

Skogestad, S., 2009. Prosessteknikk, 3rd Edition. Tapir akademisk forlag.

Thonstad, J., Fellner, P., Haarberg, G., Hives, J., Kvande, H., Sterten, A., 2001. Aluminium Electrolysis: Fundamentals of the Hall-Héroult Process. Aluminium-Verlag.

Wang, X., Tarcy, G., Batista, E., Wood, G., 2011. Active Pot Control using Alcoa STARprobe^{TM}. Springer International Publishing, pp. 491–496.

Welch, G., Bishop, G., 1995. An introduction to the kalman filter.

Appendix

C Calculation of the bath resistivity

The calculation of the bath resistivity is based on the equation for electrical conductivity χ[1] $[\mathrm{S/cm}]$ in Grjotheim et al. (1993). This equation states that

$$\begin{aligned} ln\chi = 2.0156 - \frac{2068.4}{t+273} + 0.4349(\text{bath ratio}) - 0.0207(\text{mass\% } \mathrm{Al_2O_3}) \\ - 0.005(\text{mass\% } \mathrm{CaF_2}) - 0.0166(\text{mass\% } \mathrm{MgF_2}) + 0.0178(\text{mass\% LiF}) \\ + 0.0077(\text{mass\% } \mathrm{Li_3AlF_6}). \end{aligned} \tag{6.1}$$

t is the temperature given in °C. The concentration of different substances in the bath is inspired by Hydro's cells, and is therefore confidential. The bath ratio is defined as the mass ratio between NaF and AlF_3 in the bath. When calculating these masses, the fractions are based both on how much of the substance is present in the bath as the actual compound and on how much is part of Na_3AlF_6, where each molecule can be viewed as a combination of three NaF molecules and one AlF_3 molecule. χ is calculated for a typical set of compound concentrations and for a standard temperature, and is then assumed to be a constant. Bath resistivity ρ_{bath} $[\Omega/\mathrm{m}]$ is defined as the reciprocal of bath conductivity, and can then be obtained through

$$\rho_{bath} = \frac{1}{100 \cdot \chi}. \tag{6.2}$$

[1] Equation (37) in chapter 2, page 52