

UNIVERSITY OF LEEDS

Model predictive control of industrial crystallisation processes

Ryan Paul Leeming

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

School of Chemical and Process Engineering

March 2024

Declaration

I confirm that the work submitted is my own, except where work which has formed part of jointly authored publications has been included. My contribution and the other authors to this work has been explicitly indicated below. I confirm that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapters 3, 4, 5 and 7

Leeming, R., Mahmud, T., Roberts, K.J., George, N., Webb, J., Simone, E. and Brown, C.J. 2023. Development of a Digital Twin for the Prediction and Control of Supersaturation during Batch Cooling Crystallization. *Ind. Eng. Chem. Res.* **62**, pp. 11067-11081.

All experimental and modelling work described in this publication was completed by myself, in addition to the writing of the publication itself. Other authors provided supervisory support and feedback on the draft manuscript.

Chapters 3, 6 and 7

Leeming, R., Mahmud, T., Roberts, K.J., George, N., Webb, J., Simone, E. and Brown, C.J. 2023. Effect of Impurities on Supersaturation Control during Batch Cooling Crystallization of Hexamine. (Draft, working title).

All experimental and modelling work described in this publication was completed by myself, in addition to the writing of the publication itself. Other authors provided supervisory support and feedback on the draft manuscript.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Ryan Paul Leeming to be identified as Author of this work has been asserted by himself in accordance with the Copyright, Designs and Patents Act 1988.

Acknowledgements

This research was carried out at the EPSRC Centre for Doctoral Training in Complex Particulate Products and Processes (EP/L015285/1) as part of a collaborative project with Syngenta Ltd. (Jealott's Hill, Bracknell, U.K.) and the CMAC Future Manufacturing Research Hub (University of Strathclyde, U.K.), which I gratefully acknowledge. Part of this research was performed at the CMAC National Facility supported by UKRPIF (U.K. Research Partnership Fund) award from the Higher Education Funding Council for England (HEFCE) (Grant ref: HH13054).

First and foremost, I would like to thank my supervisor Dr. Tariq Mahmud for his guidance and support throughout my PhD project. He encouraged me to pursue opportunities to present my work at multiple international conferences, and his in-depth feedback was incredibly helpful in the completion of my first journal publication. He had previously been an excellent supervisor during my MEng degree, so choosing to work with him again was an easy decision to make. I am also very grateful for all of the knowledge and support provided by my other academic supervisors, Prof. Kevin Roberts and Dr. Elena Simone.

I am grateful to the supervision provided by Dr. Neil George and Jennifer Webb from Syngenta, providing me with useful insights into the industrial value of my research. I was very fortunate that Neil valued the collaboration between academia and industry so highly, often available for a relaxed discussion. I am also very grateful for the support provided by Dr. Cameron Brown, who offered the opportunity to complete part of my experimental research at the CMAC National Facility. I would like to thank Thomas McGlone for accommodating my placement and providing technical support, as well as Humera, Michael, Momina, Magdalene, Daniel, Kenny and the rest of the CMAC team for making me feel welcome (and for introducing me to a proper Scottish whiskey).

A massive thank you to Dr. Niall Mitchell for his support with gPROMS software, helping me to resolve any technical issues and reinforcing my interest in modelling techniques. Without his prompt replies over the COVID-19 lockdown, my project would have come to a complete standstill. For similar support in learning my way around PharmaMV software and helping me deal with the challenges of process control, I would like to thank John, Darren, Ravi, Ewan and Jason from Perceptive Engineering. I am also grateful for the generous discount that was provided to the CDT for access to PharmaMV for two years, giving me freedom to explore the capabilities of the software in detail.

I would like to thank Anoushka Kulikowski for her dedication to managing the CP3 CDT, keeping us all on track and organising plenty of events for us to socialise and briefly forget the stress of completing a PhD. To my colleagues and close friends Jonny and Steve; for all the pub sessions and gaming nights, the DnD campaign (which shall continue one day), the honour of being a groomsman and all the moral support you have given me over the past few years, I thank you deeply. I'm grateful for the great times spent with Liam, Brad, Charley, all of my fellow PhD students in the CP3 CDT and friends in SCAPE at the University of Leeds.

Finally, I am thankful for the praise and personal support given to me by my family, especially my partner Ryan whose love, care and patience I could not possibly have done without. Thank you for giving me countless reasons to smile throughout this whole process.

Abstract

For fine chemicals produced by batch cooling crystallisation, control of the product crystal size distribution (CSD) is necessary for good performance properties and downstream processing. To improve upon the resource-heavy process of developing an empirical model-based control system, this thesis presents a model predictive control (MPC) methodology for the control of supersaturation and product CSD using commercially available tools.

A one-dimensional (1D) population balance model (PBM) was parameterised in gPROMS FormulatedProducts software (Siemens-PSE Ltd.) for the hexamine-ethanol crystallisation system, using kinetics reported by Myerson et al. (1986). Seeded batch crystallisation experiments were run to validate the model, and scope out suitable process conditions for control performance tests. The gPROMS model was coupled with the control software PharmaMV (Perceptive Engineering Ltd.) to create a digital twin of a 500 ml crystalliser. Simulated crystallisation data was used to train statistical MPC structures for supersaturation control, in addition to a batch end-point control (EPC) strategy for the D_{50} of the product CSD.

Within a relative supersaturation set-point range of 0.012-0.036, the digital twin performed with excellent stability. When the MPC strategy was implemented on a real 500 ml crystallisation setup, supersaturation stability was slightly poorer (standard deviations of order 10^{-3} instead of 10^{-4}). A similar level of stability was maintained when perturbing the growth rate kinetics by manipulating the seed CSD and adding a growth rate inhibitor. Performance was mostly linked to whether cooling equipment could meet the demand to accurately maintain the required cooling rate.

Simulations of the crystal size EPC strategy achieved different target D_{50} values within a range of 20 µm. However, this was not reflected in physical experiments where secondary nucleation was negligible, implying that particle generation is mandatory for CSD control. Adaptations to improve the industrial applicability of this control methodology have been proposed in the conclusions.

Contents

Declaration i
Acknowledgements ii
Abstract iv
Contents v
List of Figures ix
List of Tables xiv
List of Abbreviations xvi
Nomenclature xvii
1. Introduction 1
1.1. Background 1
1.2. Research aims and objectives 3
1.3. Thesis structure 4
2. Literature Review 6
2.1. Crystallisation theory 6
2.1.1. Supersaturation 6
2.1.2. Metastable zone width (MSZW) 7
2.1.3. Primary nucleation 7
2.1.4. Secondary nucleation 11
2.1.5. Crystal growth 12
2.1.6. Agglomeration 15
2.1.7. Polymorphism 15
2.2. Process modelling of crystallisation 16
2.2.1. Population balance model (PBM) 16
2.2.2. Solution of the PBM 20
2.3. Process measurement techniques 22
2.3.1. Concentration measurement 22
2.3.2. Particle size and count measurement 23
2.4. Process control of crystallisation 24
2.4.1. Measurement-based (model-free) control 25
2.4.2. Empirical vs first-principles model-based control 26
2.4.3. Model predictive control (MPC) 28
2.5. Conclusions 30
3. Materials and Instrumentation 31
3.1. Materials 31

3.1.1. Hexamethylene tetramine (hexamine) and ethanol solvent 31
3.1.2. Impurities for growth rate inhibition 33
3.2. Experimental setup for batch crystallisation 34
3.2.1. 500 ml jacketed crystalliser for characterisation 34
3.2.2. 500 ml jacketed crystalliser for MPC validation 34
3.2.3. 1.5 ml Crystal16 system 35
3.2.4. 100 ml AutoMATE system 36
3.3. ATR-FTIR calibration 36
3.4. Seed crystal preparation 38
3.5. Product crystal analysis 39
3.6. Physical property measurement 39
3.7. Conclusions 40
4. Metastable Zone Width Determination and Seeded Batch Crystallisation of the Hexamine-Ethanol Crystallisation System 41
4.1. Introduction 41
4.2. Materials and methodology 41
4.2.1. Methodology for MSZW determination 41
4.2.2. Methodology for seeded batch crystallisation 42
4.3. Results and discussion 42
4.3.1. MSZW of the hexamine-ethanol crystallisation system 42
4.3.2. Seeded batch crystallisation analysis 44
4.4. Conclusions 49
5. Crystallisation Process Model Parameterisation with Re-estimation of Growth Rate Kinetics and Experimental Validation 50
5.1. Introduction 50
5.2. Materials and methodology 50
5.2.1. Setup of the batch crystallisation flowsheet in gPROMS FormulatedProducts 50
5.2.2. Methodology for crystallisation kinetic parameter re-estimation 54
5.3. Results and discussion 55
5.3.1. Process model validation using seeded batch crystallisations 55
5.3.2. Re-estimation of growth rate kinetic parameters 57
5.4. Conclusions 58
6. Growth Rate Inhibitor Screening and Effects on the Hexamine-Ethanol Crystallisation System 60
6.1. Introduction 60
6.2. Materials and methodology 60

6.2.1. Methodology for impurity screening 60
6.2.2. Methodology for seeded batch crystallisation with selected impurity 61
6.3. Results and discussion 61
6.3.1. Selection of impurity from screening results 61
6.3.2. Effect of citric acid on crystal growth rate during 500 ml seeded batch crystallisation 64
6.3.3. Effect of citric acid on hexamine product crystal structure 67
6.4. Conclusions 70
7. Supersaturation MPC Development and Performance with Varying Supersaturation Set-Point, Seeding Conditions and Growth Kinetics 71
7.1. Introduction 71
7.2. Materials and methodology 71
7.2.1. CrystalMV template for supersaturation MPC in PharmaMV 71
7.2.2. Training of temperature and supersaturation MPC blocks 73
7.2.3. Methodology for in-silico supersaturation MPC 74
7.2.4. Methodology for physical validation of supersaturation MPC 75
7.3. Results and discussion 77
7.3.1. Data generated for training MPC blocks 77
7.3.2. Performance of in-silico supersaturation MPC 79
7.3.3. Comparison of in-silico and physical supersaturation MPC with varying supersaturation set-point 81
7.3.4. Performance of physical supersaturation MPC with varied seeding 83
7.3.5. Performance of physical supersaturation MPC with varying impurity concentration 85
7.4. Conclusions 87
8. Crystal Size End-Point Control Development and Performance with Varying Size Set-Point 90
8.1. Introduction 90
8.2. Materials and methodology 90
8.2.1. Batch end-point control (EPC) 90
8.2.2. Training of crystal size EPC block 91
8.2.3. Methodology for in-silico crystal size EPC 92
8.2.4. Methodology for physical validation of crystal size EPC 93
8.3. Results and discussion 95
8.3.1. Data generated for training crystal size EPC block 95
8.3.2. Performance of in-silico crystal size EPC 96

8.3.3. Performance of physical crystal size EPC 98
8.4. Conclusions 102
9. Conclusions and Future Work 104
9.1. Conclusions 104
9.2. Future work 106
References 108
Appendix A – Conversion of Volume-Based to Number-Based CSD in gPROMS FormulatedProducts 117
Appendix B – List of Presentations and Publications 120
Conference Presentations 120
Journal Publications 120

List of Figures

Figure 1.1. Workflow to meet the research objectives. 4

Figure 2.1. Solubility and supersolubility curves describing supersaturation. 6

Figure 2.2. Free energy diagram for the classical theory of nucleation (Mullin, 1992). 9

Figure 2.3. Effect of temperature on the size and free energy change of formation of a critical nucleus ($T_1 < T_2$) (Mullin, 1992). 9

Figure 2.4. Rate of nucleation as a function of supersaturation (Davey and Garside, 2006). .10

Figure 2.5. Interfacial tensions at boundaries between the three phases during heterogeneous nucleation (Mullin, 1992). 11

Figure 2.6. Development of a growth spiral starting from a screw dislocation (Mullin, 1992). 13

Figure 2.7. Concentration driving forces for crystal growth via the diffusion-reaction model (Mullin, 1992). 14

Figure 2.8. Solubility profiles of (a) monotropic and (b) enantiotropic systems (Simone, 2015). 16

Figure 2.9. Schematic representation of ATR-FTIR probe. 23

Figure 2.10. (a) FBRM probe schematic and (b) chord length measurement of a crystal (Kumar et al., 2013). 24

Figure 2.11. Schematic representation of the model-based and model-free approaches for crystallisation (Nagy et al., 2008b). 25

Figure 3.1. a) Predicted morphology of hexamine from the BFDH method, b) morphology as viewed along the [110] faces. 32

Figure 3.2. Images of a) a single hexamine crystal and b) mixture of single crystals and agglomerated clusters of hexamine, both recrystallised from pure ethanol solution. 33

Figure 3.3. Structural formulae of a) hexamine, b) citric acid, c) PABA and d) A1CA. 33

Figure 3.4. Schematic of the 500 ml jacketed crystalliser setup used for characterisation of the hexamine-ethanol crystallisation system with associated in-situ measurement technology. ..34

Figure 3.5. Schematic of the 500 ml jacketed crystalliser used for physical validation of the supersaturation and size MPC strategies with associated in-situ measurement technology....35

Figure 3.6. Schematic of the AutoMATE parallel 100 ml crystalliser setup used for impurity screening. 36

Figure 3.7. Hexamine-ethanol FTIR spectra measured using the ABB MB3000 spectrometer at 36 ºC and 0.04 g/gsolvent. Highlighted peaks are those characteristic of hexamine selected for the calibration model. 37

Figure 3.8. Actual vs. predicted concentrations for the PLS calibration model calculated in HorizonMB software. 38

Figure 3.9. Measured profiles with temperature for a) hexamine-ethanol saturated solution density and b) solid hexamine specific heat capacity. 40

Figure 4.1. Metastable zone for hexamine-ethanol system obtained from Crystal16 and 500ml vessel using cooling rate of -0.3 ºC/min, solubility from Myerson et al. (1986).....43

Figure 4.2. Measured temperature profiles during the seeded linear cooling crystallisations. a) Run A, b) Run B, c) Run C, d) Run D, e) Run E and f) Run F 44

Figure 4.3. Measured concentration and supersaturation profiles during the seeded linear cooling crystallisations, starting from the point of seeding. a) Run A, b) Run B, c) Run C, d) Run D, e) Run E and f) Run F. 45

Figure 4.4. Measured square-weighted particle counts during the seeded linear cooling crystallisations, starting from the point of seeding. a) Run A, b) Run B and c) Run C.....45

Figure 4.5. a) Number-based CSD, b) volume-based CSD and c) image of the sieved hexamine seed crystals.....46

Figure 4.6. a) Number-based CSD, b) volume-based CSD and c) image of the dry milled hexamine seed crystals.....46

Figure 4.7. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run A. 46

Figure 4.8. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run B.....46

Figure 4.9. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run C.....47

Figure 4.10. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run D. 47

Figure 4.11. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run E. 47

Figure 4.12. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run F. 47

Figure 5.1. Flowsheet of a batch cooling crystallisation process in gPROMS FormulatedProducts software. 51

Figure 5.2. Comparison of CSDs, concentration and supersaturation profiles predicted by gPROMS with measured data. Simulations starting with either the measured seed CSD or a representative lognormal peak. Results correspond to seeded linear cooling a,b,c) Run D and d,e,f) Run E. 56

Figure 5.3. Predicted concentration fit to measured data from a) Run D and b) Run E for gPROMS simulations using re-estimated growth rate parameters. 57

Figure 5.4. Percentage error in predicted vs measured concentration data from a) Run D and b) Run E for gPROMS simulations using re-estimated growth rate parameters. 57

Figure 6.1. Volume-based CSD of the hexamine product crystals from impurity screening. a) Crystalliser 1 (2 wt% impurity), b) Crystalliser 2 (4 wt% impurity) and c) Crystalliser 3 (6 wt% impurity). 61

Figure 6.2. Images of hexamine product crystals from impurity screening batches containing pure hexamine in a) Crystalliser 1, b) Crystalliser 2 and c) Crystalliser 3. 62

Figure 6.3. Images of hexamine product crystals from impurity screening batches containing a) 2 wt%, b) 4 wt% and c) 6 wt% citric acid. 62

Figure 6.4. Images of hexamine product crystals from impurity screening batches containing a) 2 wt%, b) 4 wt% and c) 6 wt% PABA. 62

Figure 6.5. Images of hexamine product crystals from impurity screening batches containing a) 2 wt%, b) 4 wt% and c) 6 wt% A1CA. 63

Figure 6.6. Overview of the 500 ml seeded batch crystallisation with 2 wt% citric acid. a) Measured concentration and supersaturation profiles starting from the point of seeding, b) volume-based CSD and c) image of the hexamine product crystals. 65

Figure 6.7. SEM images of hexamine product crystals from the 500 ml seeded batch crystallisation with 2 wt% citric acid. 66

Figure 6.8. Measured a) concentration and b) supersaturation profiles during the 500 ml seeded batch crystallisations with 0.25-0.75 wt% citric acid, starting from the point of seeding. 66

Figure 6.9. Measured solubility of hexamine in ethanol with 0.000375 g/g solvent citric acid, compared to pure hexamine solubility from Myerson et al. (1986). 67

Figure 6.10. Volume-based CSDs of the hexamine product crystals from the 500 ml seeded batch crystallisations with no impurity and 0.25-0.75 wt% citric acid. 68

Figure 6.11. Images of hexamine product crystals from 500 ml seeded batch crystallisations with a) no impurity, b) 0.25 wt% citric acid, c) 0.50 wt% citric acid and d) 0.75 wt% citric acid. 68

Figure 6.12. Images of individual hexamine product crystals from 500 ml seeded batch crystallisations with a,b) no impurity, c,d) 0.50 wt% citric acid and e,f) 0.75 wt% citric acid. 69

Figure 7.1. Block diagram of the CrystalMV supersaturation control structure in PharmaMV. 73

Figure 7.2. Data flow within the gPROMS/PharmaMV digital twin during a supersaturation control simulation. 75

Figure 7.3. Data flow within the MPC strategy during a physical batch cooling crystallisation with supersaturation control. 76

Figure 7.4. PRBS test procedure and controlled variable response predicted by the gPROMS process model for training a) temperature MPC block and b) supersaturation MPC block. ... 78

Figure 7.5. PRBS test procedure and controlled variable response measured using the 500 ml jacketed crystalliser for re-training the temperature MPC block. 79

Figure 7.6. Temperature, cooling (ramp) rate, concentration and supersaturation profiles generated during in-silico supersaturation MPC, for supersaturation set-points of a,b) 0.012, c,d) 0.024 and e,f) 0.036. 80

Figure 7.7. Temperature, cooling rate, concentration and supersaturation profiles measured during physical supersaturation MPC experiments, for supersaturation set-points of (a,b) 0.012, (c,d) 0.024 and (e,f) 0.036.....81

Figure 7.8. Temperature, cooling rate, concentration and supersaturation profiles measured during physical supersaturation MPC, for batches seeded with (a,b) 3 g of 112.2 μm seeds, (c,d) 1 g of 68.9 μm seeds and (e,f) 3 g of 68.9 μm seeds. 83

Figure 7.9. Images of hexamine product crystals from physical supersaturation MPC batches seeded with a) 1 g of 112.2 μm seeds, b) 3 g of 112.2 μm seeds, c) 1 g of 68.9 μm seeds and d) 3 g of 68.9 μm seeds. 85

Figure 7.10. Temperature, cooling rate, concentration and supersaturation profiles measured during physical supersaturation MPC, for citric acid impurity concentrations of (a,b) 0.25 wt%, (c,d) 0.50 wt% and (e,f) 0.75 wt%.....86

Figure 8.1. Data flow within the gPROMS/PharmaMV digital twin during a crystal size control simulation. 93

Figure 8.2. Data flow within the MPC strategy during a physical batch cooling crystallisation with crystal size control. 94

Figure 8.3. a) Supersaturation and b) number-based D_{50} profiles generated using the digital twin for training of the crystal size EPC block. 95

Figure 8.4. a) Temperature and cooling (ramp) rate, b) supersaturation and c) number D_{50} profiles generated during in-silico crystal size EPC, for product D_{50} set-points of 140, 150 and 160 μm. 96

Figure 8.5. a) Vessel temperature and cooling (ramp) rate, b) supersaturation and c) number D_{50} profiles measured during physical crystal size EPC, for product D_{50} set-points of 140, 150 and 160 μm. 99

Figure 8.6. PVM images of hexamine crystals during physical crystal size EPC for a D_{50} set-point of 140 μm. a) 5 mins after seeding, b) 30 mins after seeding, c) 28 mins into controlled cooling, d) 40 mins into controlled cooling, e) 53 mins into controlled cooling and f) 66 mins into controlled cooling. 100

Figure 8.7. a) Vessel temperature and cooling (ramp) rate, b) supersaturation and c) number D_{50} profiles measured during physical crystal size EPC, following fixed supersaturation profiles expected to achieve D_{50} set-points of 140 μm and 160 μm. 101

List of Tables

Table 3.1. Solubility, nucleation rate and growth rate expressions for hexamine-ethanol crystallisation reported by Myerson et al. (1986). 32

Table 3.2. Solution temperatures and concentrations for which IR spectra were measured for the calibration model in HorizonMB software. 37

Table 3.3. Analysis of the PLS calibration model error. 38

Table 4.1. Process conditions applied to each seeded linear cooling crystallisation run.....42

Table 4.2. Mean sizes and spans of volume-based product CSDs and supersaturations from each seeded linear cooling crystallisation run. 48

Table 5.1. Input data for the 'Global Specifications' component of the flowsheet for hexamine-ethanol batch cooling crystallisation in gPROMS FormulatedProducts.....51

Table 5.2. Input data for the 'MSMPR Crystalliser' component of the flowsheet for hexamine-ethanol batch cooling crystallisation in gPROMS FormulatedProducts. Nucleation and growth kinetic parameters are provided from Myerson et al. (1986). 53

Table 5.3. Input data for the 'Temperature Controller' component of the flowsheet for hexamine-ethanol batch cooling crystallisation in gPROMS FormulatedProducts.....54

Table 5.4. Predicted volume mean crystal size of product from gPROMS simulations using re-estimated growth rate parameters. 58

Table 6.1. Mean sizes and spans of volume-based product CSDs from each impurity screening batch.....62

Table 6.2. Mean sizes and spans of volume-based product CSDs from the 500 ml seeded batch crystallisations with no impurity and 0.25-0.75 wt% citric acid.....68

Table 7.1. Process conditions applied to the gPROMS/PharmaMV digital twin for in-silico supersaturation MPC performance testing.....75

Table 7.2. Process conditions for the physical supersaturation MPC validation experiments.77

Table 7.3. Mean and standard deviation of supersaturation and average cooling rates during in-silico simulations and physical validation experiments for supersaturation control at varying supersaturation set-points. 81

Table 7.4. Mean and standard deviation of supersaturation and average cooling rates during physical supersaturation control experiments at varying seeding conditions. 84

Table 7.5. Mean and standard deviation of supersaturation and average cooling rates during physical supersaturation control experiments at varying citric acid impurity concentrations. 86

Table 8.1. Process conditions applied to the gPROMS/PharmaMV digital twin for crystal size EPC training and performance testing. 92

Table 8.2. Predicted vs. achieved D_{50} of the product number-based CSD resulting from in-silico crystal size EPC simulations. 97

Table 8.3. Predicted vs. achieved D_{50} of the product number-based CSD resulting from physical crystal size EPC experiments. 98

Table 8.4. Predicted vs. achieved D_{50} of the product number-based CSD resulting from physical crystal size EPC experiments, following fixed supersaturation profiles. 102

List of Abbreviations

1D	One-dimensional
A1CA	Adamatane-1-carboxylic acid
ARX	Auto regressive with exogenous input
ATR	Attenuated total reflection
BFDH	Bravais-Friedel-Donnay-Harker
CFD	Computational fluid dynamics
CLD	Chord length distribution
CSD	Crystal size distribution
DSC	Differential scanning calorimetry
FBRM	Focused-beam reflectance measurement
FTIR	Fourier-transform infrared
GMC	Generic model control
IIC	Iterative identification control
ILC	Iterative learning control
IN	Instantaneous nucleation
KBHR	Kashchiev-Borissova-Hammond-Roberts
MPC	Model predictive control
MSMPR	Mixed-suspension mixed-product removal
MSZW	Metastable zone width
PABA	Para-aminobenzoic acid
PAT	Process analytical technology
PBM	Population balance model
PI	Proportional-integral
PID	Proportional-integral-derivative
PLS	Partial least squares
PN	Progressive nucleation
PV	Process value
PVM	Particle vision and measurement
RLS	Recursive least squares
SEM	Scanning electron microscopy
SP	Set-point
SQP	Sequential quadratic programming
UTL	Upper transport layer

Nomenclature

A	Surface area	m^2
A_c	Crystal surface area	m^2
a	Activity	-
a^*	Equilibrium activity	-
B_{agg}	Birth rate (agglomeration)	$\#/s.m^3$
B_{br}	Birth rate (breakage)	$\#/s.m^3$
b	Secondary nucleation rate supersaturation order	-
C	Concentration	g/g solution
C^*	Saturation concentration	g/g solution
C_i	Interfacial concentration	g/g solution
ΔC_i	Width of size class i	m
c_p	Specific heat capacity	J/kg.K
D_{50}	Median diameter of the crystal size distribution	m
D_{AB}	Diffusivity of solute	m^2/s
D_{agg}	Death rate (agglomeration)	$\#/s.m^3$
D_{br}	Death rate (breakage)	$\#/s.m^3$
d	Diameter	m
e	Controller error for PID control	-
G	Growth rate	m/s
ΔG	Gibbs free energy change	J
$\Delta G'_{crit}$	Critical free energy for heterogeneous nucleation	J
ΔG_{crit}	Critical free energy for homogeneous nucleation	J
ΔG_s	Surface excess free energy	J
ΔG_v	Volume excess free energy	J
g	Growth rate supersaturation order	-
J	Objective function	-
J	Nucleation rate (primary)	$\#/s.m^3$
J_{sec}	Secondary nucleation rate	$\#/s.m^3$
$J_{sec,cc}$	Secondary nucleation rate for crystal-crystal collisions	$\#/s.m^3$
$J_{sec,ci}$	Secondary nucleation rate for crystal-impeller collisions	$\#/s.m^3$
j	Secondary nucleation rate impeller speed order	-
K_v	Volume shape factor of crystals	-
k	Secondary nucleation rate suspension density order	-
k	Sampling instant	-
k_B	Boltzmann constant	J/K
k_D	Derivative gain	-
k_d	Diffusive mass transfer coefficient	$mol/s.m^2$
K_G	Overall growth coefficient	$mol/s.m^2$
k_g	Growth rate constant	m/s
k_I	Integral gain	-
k_m	Mass transfer coefficient	$mol/s.m^2$
k_n	Primary nucleation rate constant	$\#/s.m^3$
$k_{n,s}$	Secondary nucleation rate constant	$\#/s.m^3$
$k_{n,cc}$	Secondary nucleation rate constant for crystal-crystal collisions	$\#/s.m^3$

$k_{n,ci}$	Secondary nucleation rate constant for crystal-impeller collisions	#/s.m^3
k_P	Proportional gain	-
k_r	Surface integration 'reaction' rate constant	mol/s.m^2
L	Characteristic length (crystal size)	m
M_T	Suspension density	kg solids/m^3 suspension
m	Mass of solute	kg
m	Control horizon	-
N	Number of measurements taken over all experiments	-
NE	Number of experiments	-
N_i	Number of particles in size class i	#/m^3
N_{imp}	Impeller speed	rpm
NM_{ij}	Number of measurements of variable j in experiment i	-
NV_i	Number of measured variables in experiment i	-
n	Population density	#/m.m^3
n	Order of primary nucleation	-
n_{cc}	Order of secondary nucleation for crystal-crystal collisions	-
n_{ci}	Order of secondary nucleation for crystal-impeller collisions	-
p	Prediction horizon	-
$p(v/\varepsilon)$	Fraction of new particles from breakage	-
q	Cooling rate	^{o}C/min
R	Universal gas constant	kJ/kmol.K
$R_{a,i}$	Net agglomeration rate in size class i	#/s.m^3
$R_{b,i}$	Net breakage rate in size class i	#/s.m^3
r	Radius of spherical nucleus	m
r_c	Critical radius of nucleus	m
R_G	Mass deposition rate	kg/m^2.s
S	Fundamental supersaturation	-
T	Temperature	K, ^{o}C
T_e	Equilibrium temperature	K, ^{o}C
ΔT	Undercooling	K, ^{o}C
t	Time	s
$u(t)$	Control signal for PID control	-
Δu	Control move	-
V_c	Crystal volume	m^3
y_r	Set-point reference	-
$y(t)$	Process variable for PID control	-
$\hat{y}$	Estimated future output	-
z_{ijk}	k^{th} predicted value of variable j in experiment i	-
$\check{z}_{ijk}$	k^{th} measured value of variable j in experiment i	-
α	Diffusivity correction factor	-
α_j	Controlled variable weighting factor	-
β	Surface shape factor	-
β_j	Manipulated variable weighting factor	-
$\beta(v,\varepsilon)$	Agglomeration kernel	-
γ	Interfacial surface tension	N/m

γ_e	Surface energy	J/m^2
$\gamma(\varepsilon)$	Number of new particles generated from breakage	
δ	Film thickness	m
ε	Energy dissipation rate	m^2/s^3
η	Dynamic viscosity	kg/m.s
θ	Contact angle	rad
ν	Kinematic viscosity	m^2/s
ρ_c	Crystal density	kg/m^3
σ	Absolute/relative supersaturation	-
σ^2_{ijk}	Variance of the k^{th} measurement of variable j in experiment i	-
λ	Molecular latent heat of crystallisation	J/mol
μ	Chemical potential	J/mol
μ_c	Critical undercooling	-
μ_j	jth moment of crystal size distribution	-
Y	Molecular volume	m^3/mol
Φ	Maximum likelihood objective function	-
ϕ	Nucleation correction factor	-

1. Introduction

1.1. Background

Crystallisation, the process of growing a solid crystal phase from solution, is employed in a wide variety of fine chemical industries including (but not limited to) pharmaceuticals and agrochemicals (Mullin, 1992). One such process is batch cooling crystallisation, which requires optimisation of the cooling profile to achieve both a high yield and purity of the crystalline product. It is considerably important to control the physical properties, such as the crystal size distribution (CSD), to enhance product performance properties such as the dissolution rate which are linked to size and shape. The product CSD also has an impact on downstream operations such as filtration and drying, which can be hindered significantly by the overabundance of fine particles. The aspect ratio of crystals can also negatively impact their filterability, therefore it may also be preferable to grow crystals of a specific shape (morphology) or polymorphic form. When the optimal CSD has been determined, industry finds the most value when a process can reach this target accurately, efficiently and in a repeatable manner.

For batch cooling crystallisation, the main objective of process control is to determine the optimal temperature profile that will result in production of the desired CSD. This is governed by the generation of supersaturation during cooling: a factor that is directly connected to the key rate mechanisms of crystal nucleation and growth which drive the crystallisation process.

The process can be tailored through the introduction of seed crystals with a known CSD as a starting point, however this can bring potential challenges to the control of supersaturation (Saleemi et al., 2012; Kardum et al., 2013; Long et al., 2016). For example, a high seed loading can reduce the frequency of nucleation events. The supersaturation trajectory can also be altered by the quality of the seed crystals; uneven surfaces formed as a result of milling can hinder the attachment of solute, leading to an unexpected reduction in the growth rate. Damaged surfaces can usually be repaired through the preparation of seed slurries for the feed, which treat the crystals through a process known as Ostwald ripening (Zhang et al., 2021).

Obtaining accurate and precise process measurements is vital for good control over crystallisation. Advancements in process analytical technology (PAT) over the past few decades has brought the potential to measure many of the key process properties directly (Doki et al., 2004; Simone et al., 2015; Zhang et al., 2017). Supersaturation can be monitored by measuring the solution concentration, through in-situ analytical techniques such as attenuated

total reflection (ATR) - Fourier transform infrared (FTIR) or ATR-UV/Vis spectroscopy. A range of well-established in-situ measurement techniques for particle imaging and size analysis also exist, including focused-beam reflectance measurement (FBRM) and particle vision and measurement (PVM).

Batch cooling crystallisation was initially controlled using simple measurement-based feedback loops, manipulating the process temperature using a heating/cooling jacket around the crystalliser (Fujiwara et al., 2002; Grön et al., 2003; Fujiwara et al., 2005; Rajagopalan et al., 2018). The main drawback this approach is the empirical nature of the controller development. Large amounts of experimental data are required to sufficiently capture the crystallisation behaviour, which is extremely demanding of both time and physical resources. The progression to model-based methods as technology advanced did not immediately resolve this, as many of the initial 'batch-to-batch' models that have been developed are also entirely empirical (Nagy, 2009; Forgione et al., 2015).

Although mechanistic or 'first-principles' models would ideally be derived only from physical concepts, these still require a small empirical contribution to function effectively. Solution of the population balance model (PBM) is widely implemented, which combines crystallisation rate kinetics with the material balance to predict the change in the CSD of a given crystal population (Randolph and Larson, 1988). This type of model is supported through experimentation to re-estimate the crystallisation kinetic parameters to improve accuracy.

Model predictive control (MPC) is a feedforward style of control that implements the PBM, generating optimal trajectories in properties such as temperature or supersaturation that will lead to the desired control objectives (Jha et al., 2017). Multiple successful applications of MPC for controlling crystal size have been demonstrated for continuous crystallisation processes (Moldoványi et al., 2005; Yang and Nagy, 2015; Tahir et al., 2017). Crystal size control has proven more challenging for batch processes given their more dynamic nature, therefore studies on this area have been more limited.

Recent digital advancements in the area of machine learning have proved applicable to crystallisation process control (Benyahia et al., 2021; Zheng et al., 2021). Some methods such as reinforced-learning have been shown to perform better than standard MPC for crystal size control, but standard MPC of supersaturation still performs to an equivalent degree of accuracy. Other techniques that employ machine learning to improve the computational efficiency of MPC appear promising, but they are still in development with no experimental validation of

their concepts. Therefore, there is still the need for a simplistic yet robust MPC strategy to be developed, which could use existing technology to carry out optimisation of batch crystallisation processes.

1.2. Research aims and objectives

The main research question upon which the work presented in this thesis was based is as follows:

"How can an MPC strategy be built using commercial engineering tools to accurately predict and control the CSD in an industrially relevant crystallisation system?"

In order to answer this research question and contribute towards the knowledge base of MPC applications for batch crystallisation, the following objectives should be met:

- To develop a population balance model that can reliably predict the CSD evolution over time during batch cooling crystallisation, without relying heavily on empirical batch data.
- To construct an MPC strategy using the validated process model that can accurately control the supersaturation, such that the crystallisation behaviour can be manipulated to produce a desired CSD.
- To test the MPC strategy's performance when applied to a realistic, impure crystallisation system relevant to industrial applications.

An opportunity to achieve these objectives arises from the interface between two pieces of existing commercial software. The process modelling software gPROMS FormulatedProducts (Siemens-PSE Ltd.) can be used to develop a first-principles, one-dimensional (1D) PBM for a selected crystallisation system. The resulting 'digital twin' of a batch crystalliser can then be coupled with the in-line monitoring and control software PharmaMV (Perceptive Engineering Ltd.) to simulate batch cooling crystallization processes where the supersaturation and product CSD are controlled. The MPC strategy developed using this software platform can be connected to a physical lab-scale crystallization setup, controlling real batch operations for validation. The purpose this methodology is to accelerate the optimisation of a given crystallisation process, reducing the total time and resources spent on developing a control system that can be applied to a new solute-solvent system without an extensive research and re-training procedure.

To study the performance of the MPC strategy on a more realistic crystallisation system, impurities were introduced that would inhibit the crystal growth rate. Experiments with impurities were used only as validation cases and not as training data for the process model. This was to see how the MPC strategy dealt with an unexpected disturbance to the growth kinetics, without the need for re-training. This was to test the robustness of the MPC strategy when there was an obvious mismatch between the training data (pure system) and physical measurements. As different impurity profiles may be present across various feedstocks used in industrial crystallisations, it was valuable to know how well the MPC strategy could adjust to handle these unknowns.

The workflow that was followed to achieve these objectives is summarised by the diagram in Figure 1.1.

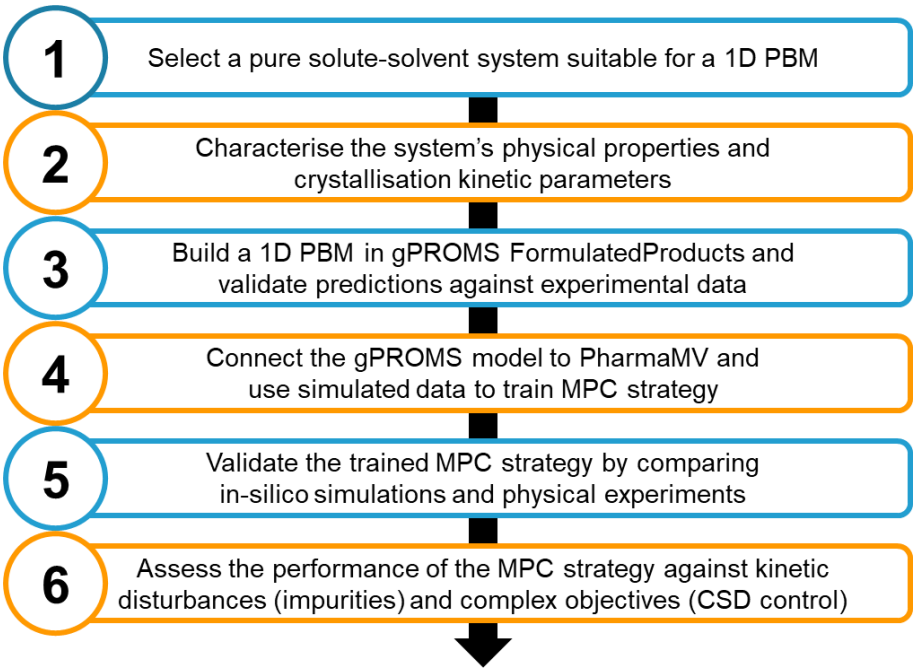


Figure 1.1. Workflow to meet the research objectives.

After evaluating the performance of an MPC strategy developed using these commercial tools, this work aims to identify any shortcomings and propose how these can be overcome through further research.

1.3. Thesis structure

Chapter 2 outlines the fundamentals of crystallisation theory, including the concepts of supersaturation, the metastable zone, nucleation and growth mechanisms and polymorphism. A literature review is then presented, starting with the modelling of crystallisation processes using population balance equations. Details are provided about the crystallisation system

chosen for study, followed by descriptions of the PAT that were implemented during the experimental work. Finally, existing measurement-based and model-based methodologies for crystallisation process control are explored to identify the current challenges.

Chapter 3 describes the solute-solvent system chosen for study, in addition to some identified growth rate inhibitors. Descriptions are also provided of the experimental instrumentation used to characterise the system and carry out batch crystallisations.

In **Chapter 4**, a series of unseeded and seeded batch cooling crystallisations are presented for characterisation of the chosen solute-solvent system. The key crystallisation mechanisms for the system are identified. By covering a wide range of operating conditions, a substantial dataset of temperature, concentration and CSDs are obtained to support the development of the crystallisation process model.

Chapter 5 outlines the parameterisation of the 1D PBM in gPROMS FormulatedProducts software, creating a 'digital twin' of the batch crystallisation process. The digital twin is utilised to simulate batch cooling crystallisations, which are validated using the experimental data obtained in Chapter 4. Re-estimation of the crystal growth rate kinetics is also described.

Chapter 6 presents the screening of various compounds to identify the most suitable crystal growth rate inhibitor. The impact of the selected impurity on the crystallisation system's key mechanisms, particularly nucleation and agglomeration, are investigated.

Chapter 7 describes the development of a supersaturation control methodology using PharmaMV software. The process of training MPC algorithms using the gPROMS digital twin is explained, followed by experimental validation and performance testing of the supersaturation control strategy.

In **Chapter 8**, the control strategy is expanded upon through the development of a crystal size control methodology. The performance of the combined MPC and batch end-point control (EPC) strategy is evaluated.

Chapter 9 concludes by highlighting the important findings of this research, expanding upon these by suggesting plans for future work.

2. Literature Review

This chapter presents an overview of the fundamental theories and understandings of crystallisation, process modelling of crystallisation utilising population balance equations, existing measurement technology and current methods that have been employed for the control of crystallisation processes. The focus of each topic will be centred around the relevance to the work outlined in the remainder of this thesis.

2.1. Crystallisation theory

2.1.1. Supersaturation

Cooling crystallisation operates based on the transition from a stable solution (in thermodynamic equilibrium with the solid phase) to a supersaturated state by reducing the temperature at a certain rate, initiating the formation of solid crystals. As outlined in the solubility graph in Figure 2.1, a stable solution transitions into the labile region where spontaneous crystallisation can occur via either cooling or evaporation, the latter of which increases the concentration of solute while maintaining a constant temperature (Mullin, 1992).

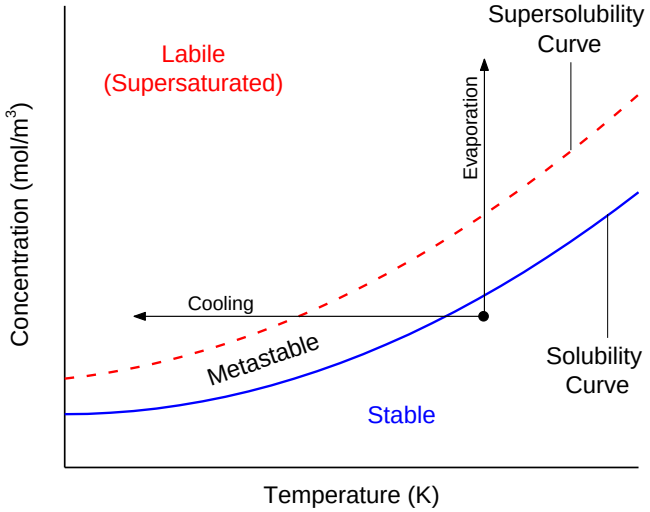


Figure 2.1. Solubility and supersolubility curves describing supersaturation.

Supersaturation can be expressed in multiple forms, including the concentration driving force ΔC, fundamental supersaturation S and absolute/relative supersaturation σ:

$$\Delta C = C - C^* \tag{2.1}$$

$$S = C/C^* \tag{2.2}$$

$$\sigma = \Delta C/C^* = S - 1 \tag{2.3}$$

where c is the bulk solution concentration and c^* is the saturation concentration at the given temperature (Mullin, 1992).

A solution is saturated when the difference between the chemical potentials of the solution μ_{sol} and solid crystal μ_s is zero, therefore the driving force for crystallisation is the difference between these and can be related to supersaturation using the activity *a*:

$$\Delta\mu = \mu_{sol} - \mu_s \tag{2.4}$$

$$\Delta\mu/RT = \ln(a/a^*) = \ln S_c \tag{2.5}$$

where *R* is the universal gas constant (8.314 kJ/kmol.K), *T* is the temperature and a^* is the equilibrium activity (Davey and Garside, 2006).

2.1.2. Metastable zone width (MSZW)

Crystallisation can also occur within the metastable zone between the solubility and supersolubility curves, but not spontaneously. The metastable zone width (MSZW) can vary widely; although the solubility curve can be determined with reasonable accuracy, the supersolubility curve is a kinetic boundary and changes due to factors such as cooling rate, agitation speed and the presence of impurities (Mullin, 1992). Through the characterisation of the MSZW of a specific solute-solvent system, the nucleation behaviour can be better understood to design batch crystallisation processes appropriately, such as deciding whether seeding is necessary and under what conditions it is feasible (Kim and Mersmann, 2001).

Most commonly the polythermal method developed by Nývlt et al. (1970) is used to identify temperatures associated with primary nucleation, however alternative model-based approaches exist which use the growth mechanics of seeded systems (Mersmann and Bartosch, 1998). The polythermal method involves the cycling of temperature of a solution of given concentration above and below its saturation temperature, leading to the identification of the temperatures of first crystal appearance and last crystal dissolution. The difference between these temperatures known as the undercooling ΔT quantifies the MSZW, allowing it to be related to supersaturation:

$$\Delta C_{max} = \frac{dC^*}{dT}\Delta T \tag{2.6}$$

where ΔC_{max} is the level of supersaturation just before crystallisation occurs (Camacho Corzo et al., 2017).

2.1.3. Primary nucleation

Nuclei are small solid bodies that can form spontaneously from solution or existing crystal surfaces, which create the starting point for crystal growth. The mechanisms of nucleation can

be divided into primary (no initial crystalline matter) and secondary (seed crystals initially present) mechanisms. Primary nucleation can further be divided into homogenous (spontaneous) and heterogenous (induced by foreign surfaces) mechanisms.

The exact process governing spontaneous, homogeneous nucleation is still not fully understood; stable nuclei must be formed from tens to thousands of molecules depending on the substance. However, it is extremely unlikely that these could all collide at once with the correct orientation to form a fixed lattice (Mullin, 1992). The clustering of molecules in order to exceed the energy barrier for nucleation is recognised as a stochastic process, requiring statistical analysis to be fully understood (Maggioni and Mazzotti, 2017). The critical size of nuclei beyond which crystal growth occurs has been described by the classical theory of nucleation based on work by Gibbs (1948), where the free energy change required ΔG is calculated from the summation of the surface and volume excess free energies ΔG_s and ΔG_v respectively:

$$\Delta G = \Delta G_s + \Delta G_v \tag{2.7}$$

$$\Delta G_s = 4\pi r^2 \gamma \tag{2.8}$$

$$\Delta G_v = -\frac{4}{3}\pi r^3 \frac{\Delta \mu}{v} \tag{2.9}$$

where r is the radius of a spherical nucleus, γ is the interfacial surface tension and v is the molecular volume. When plotting these quantities against r as shown by the free energy diagram in Figure 2.2, a maximum where $\Delta G = 0$ can be identified as the energy required to form the critical nucleus size r_c. Not all areas of the crystalliser have the same energy level despite constant thermodynamic properties; where energy fluctuates above the mean value in a supersaturated region, nucleation will be more likely to occur (Mullin, 1992).

$$r_c = 2\gamma v / \Delta \mu \tag{2.10}$$

The primary nucleation rate J can be expressed using an Arrhenius relationship, which can further be rewritten using the Gibbs-Thomson relationship to express it in terms of the three most important factors affecting nucleation: temperature, supersaturation and interfacial tension (Mullin, 1992).

$$J = A \exp\left(-\Delta G / k_B T\right) \tag{2.11}$$

$$J = A \exp\left[-\frac{16\pi {\gamma_e}^3 v^2}{3 {k_B}^3 T^3 (\ln S_c)^2}\right] \tag{2.12}$$

where A is an empirical constant, γ_e is the surface energy and k_B is the Boltzmann constant (1.38×10^{-23} J/K).

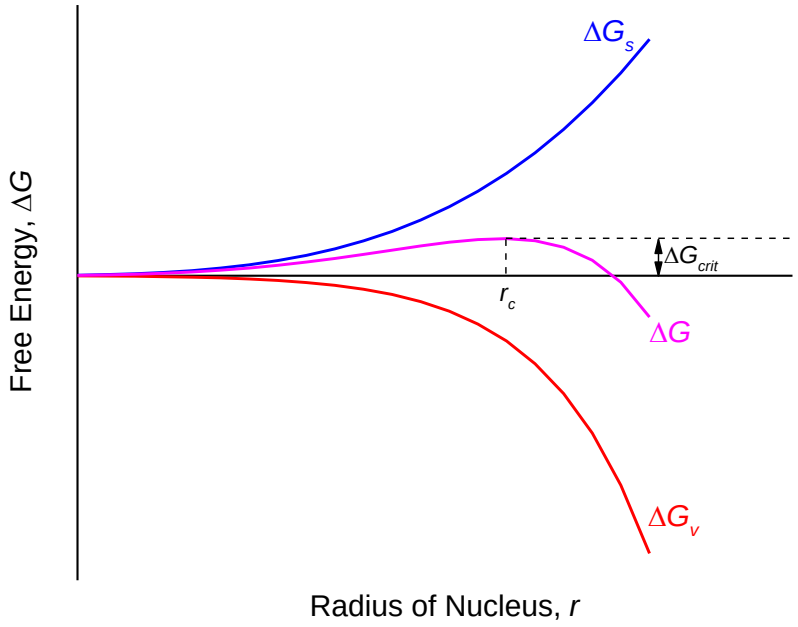


Figure 2.2. Free energy diagram for the classical theory of nucleation (Mullin, 1992).

The degree of undercooling ΔT influences nucleation due to the proportional relationships shown below, where an increase in ΔT (i.e. more cooling) reduces the critical nucleus size (Mullin, 1992). The impacts of temperature on ΔG and supersaturation on J can be seen in Figures 2.3 and 2.4 respectively.

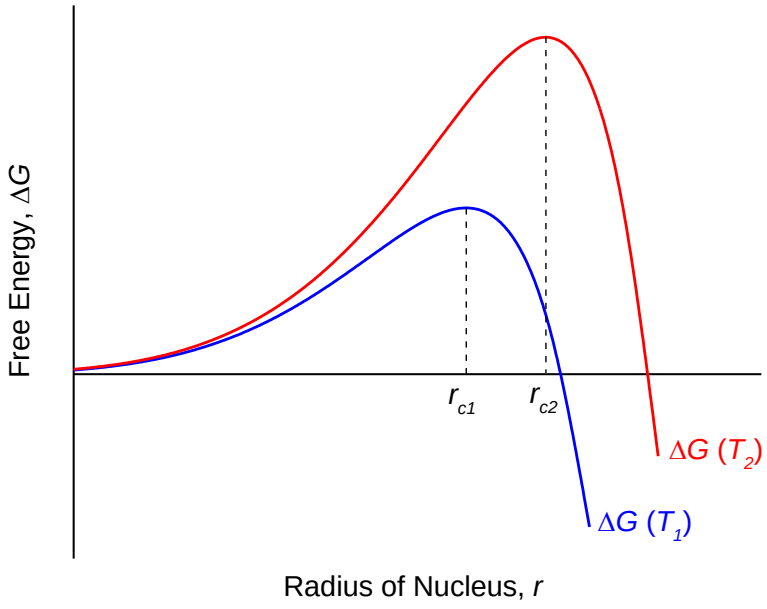


Figure 2.3. Effect of temperature on the size and free energy change of formation of a critical nucleus ($T_1 < T_2$) (Mullin, 1992).

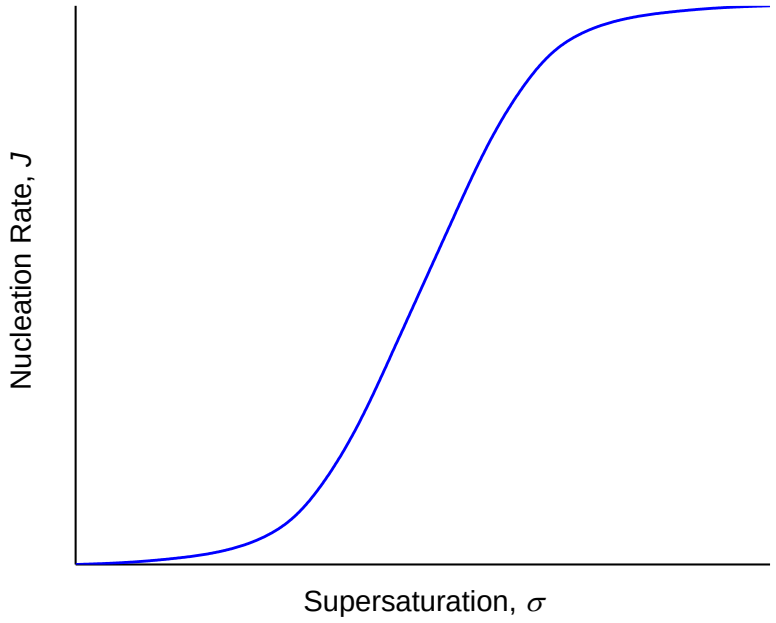


Figure 2.4. Rate of nucleation as a function of supersaturation (Davey and Garside, 2006).

$$r_c \propto (\Delta T)^{-1} \quad (2.13)$$

$$\Delta G_{crit} \propto (\Delta T)^{-2} \quad (2.14)$$

Due to certain assumptions made by the classical nucleation theory such as treating nuclei as spheres whose stability is only governed by size, research has been done into the possibility of a two-step nucleation model which better describes the process (Brown, 2012). Chen et al. (2011) used a wide range of evidence from dynamic light scattering, differential scanning calorimetry and X-ray scattering to propose a mechanism where an unordered intermediate structure is formed by solute molecules, which proceeds to reorder into a more stable nucleus. Regions of increased solution concentration within the bulk solution will typically lead to this effect occurring more readily. This theory continues to be relevant in crystallisation studies, including reactive applications such as that described by Fei et al. (2018), where tungsten carbide was found to develop into an amorphous intermediate preceding the formation of solid nuclei. Despite this, the two-step theory is not fully understood or adopted by many, instead paving the way for potential updated models in the future.

On closer inspection, many cases of homogeneous nucleation have been found to be induced in some way from the surfaces of 'heteronuclei', most commonly dust and impurities that could not be filtered out of the system (Mullin, 1992). It should be noted however that impurities can act both as nucleation inhibitors and accelerators depending on their nature. Lower degrees of undercooling are required for nucleation to occur on the surfaces of foreign bodies, adjusting

the critical free energy for homogeneous nucleation ΔG_{crit} by a certain correction factor $0 < \phi \leq 1$:

$$\Delta G'_{crit} = \phi \Delta G_{crit} \tag{2.15}$$

$$\phi = \frac{(2 + \cos\theta)(1 - \cos\theta)^2}{4} \tag{2.16}$$

$$\cos\theta = \frac{\gamma_{sl} - \gamma_{cs}}{\gamma_{cl}} \tag{2.17}$$

where $\Delta G'_{crit}$ is the critical free energy for heterogeneous nucleation, θ is the angle of contact between the crystal and foreign solid surfaces (shown in Figure 2.5) and the subscripts on the interfacial tensions correspond to the boundaries between foreign solid surface (s), crystal mass (c) and liquid (l) (Mullin, 1992).

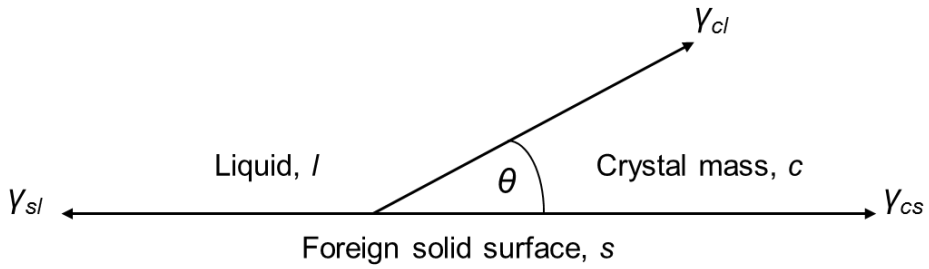


Figure 2.5. Interfacial tensions at boundaries between the three phases during heterogeneous nucleation (Mullin, 1992).

2.1.4. Secondary nucleation

The free energy reduction that comes from introducing solid surfaces for nucleation is greatest when a surface closely matches the structure of the crystalline phase (Brown, 2012). Therefore, the deliberate creation and addition of seed crystals of the same material as the solute into a crystallisation system provides the best surfaces for encouraging nucleation to occur. Similar results can also be achieved using different materials with strong crystal lattice similarity to the solute (Mullin, 1992). Industrial settings can exert precise control of the product's size distribution through effective implementation of seeding (Mullin, 1992).

Secondary nucleation is a combination of shear (activated) and contact (attrition) mechanisms, the latter occurring much more commonly. Activated nucleation occurs when the hydrodynamic shear forces generated by strong enough agitation sheds seed crystals of fine, ordered matter from their surfaces. These fine particles then become new nuclei for crystal growth (Mullin, 1992). At high supersaturations, the dendritic growth of crystals can induce secondary nuclei that are easily removed by shear forces, through a process known as 'needle

'breeding' or propagation (Davey and Garside, 1980). However, as most industrial crystallisations typically operate at low supersaturations to avoid uncontrollable primary nucleation, dendritic growth is often difficult to generate (Myerson and Ginde, 2002).

Attrition occurs when crystals collide with each other, the sides of the vessel or the agitation equipment, where the hardness of the contacting surface can affect its ability to break the crystals apart and therefore affect the nucleation rate (Mullin, 1992). This type of secondary nucleation is strongest in systems with low supersaturation where growth rate is optimum for good quality crystals, being easy to achieve based on the low energy requirement for breaking off crystal growths via collision (McCabe et al., 1993). Fragments of seed crystals become nuclei that share the original crystal's structure and crystalline form, making this a useful method of controlling polymorphism (Cui et al., 2015).

In addition to supersaturation effects, the hydrodynamics within a crystalliser have a significant effect on secondary nucleation. High agitation rates paired with larger seed crystals increase the probability and energy associated with crystal collisions, although there have been exceptions where nucleation rate stagnated after an increase in seed size (Mullin and Leci, 1972).

2.1.5. Crystal growth

Once stable crystal nuclei have formed in a supersaturated system they will begin to grow using additional solute, for which there are multiple theories on the mechanism. Although now commonly unused, surface energy theory suggested that higher index faces would grow faster due to surface energy differences such that an 'equilibrium' shape would be achieved, which would have the lowest total surface energy possible (Mullin, 1992). Theories based on adsorption originally presented by Volmer (1939) have been thoroughly investigated and adapted, but the most generally accepted theory involves an additional diffusion step prior to adsorption.

The adsorption-layer theory consists of the formation of 1-10 nm thick layers in dynamic equilibrium with the bulk solution, built up by solute adsorbing to the surface (Mullin, 1992). Some form of 'centre of crystallisation' is required to start the next layer of growth, which can be either produced via nucleation on the crystal surface or an existing dislocation in the structure can be used, resulting in different growth patterns such as in the example of a screw dislocation shown in Figure 2.6 (Mullin, 1992).

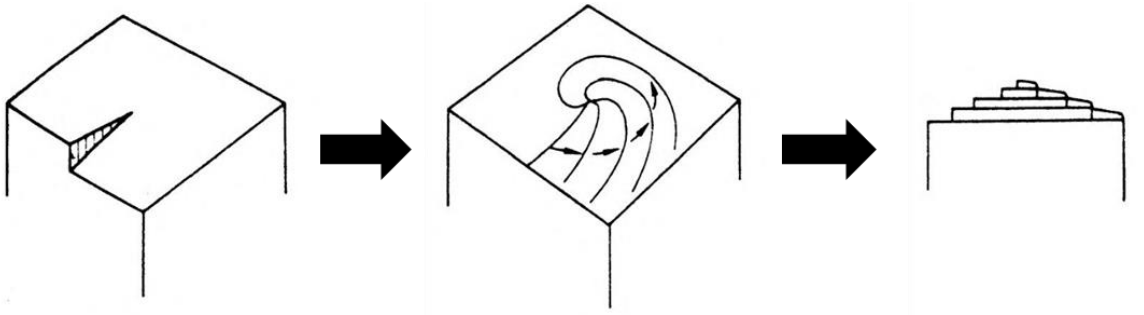


Figure 2.6. Development of a growth spiral starting from a screw dislocation (Mullin, 1992).

Originally considered by Noyes and Whitney (1897), it was proposed that the process of solute depositing onto crystal surfaces is a diffusional process and that crystallisation is essentially the reverse of dissolution, however this isn't always the case as rates can differ even at the same temperature and composition (Mullin, 1992). This provided the basis for crystal growth to be considered in terms of mass transfer and concentration differences between the surface and bulk solution:

$$\frac{dm}{dt} = k_m A(c - c^*) \tag{2.18}$$

where m is the mass of solute, t is the time of deposition, k_m is the mass transfer coefficient and A is the crystal surface area (Mullin, 1992). On the assumption that there is a stagnant film of liquid adjacent to the crystal surface that solute molecules must diffuse through, Nernst (1904) suggested an adaptation to include the thickness of the stagnant film δ, which varies depending on the degree of agitation.

$$\frac{dm}{dt} = \frac{D_{AB}}{\delta} A(c - c^*) \tag{2.19}$$

where D_{AB} is the diffusivity of solute A through medium B. Typically, $\delta < 150$ μm in slowly agitated systems, but higher agitation speeds theoretically cause the film to reduce in size down to nearly zero thickness. This would suggest an almost infinite growth rate, therefore the assumption that crystal growth is based on diffusion alone is flawed. Building upon this, a two-stage process was developed involving diffusion from the bulk solution to the interface between crystal and solution, followed by first-order 'reaction' at the surface where solute molecules arrange themselves into the appropriate lattice structure, as depicted by Figure 2.7 (Mullin, 1992).

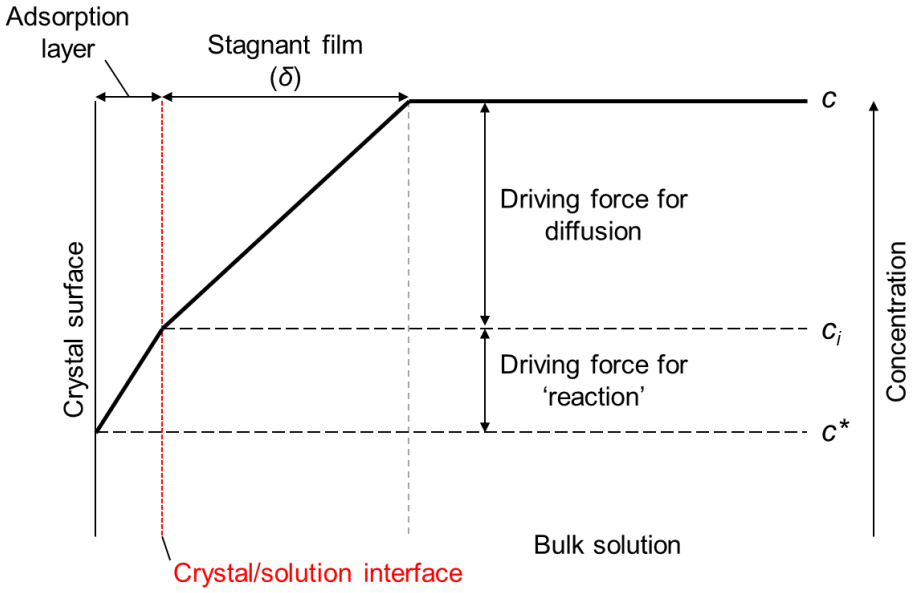


Figure 2.7. Concentration driving forces for crystal growth via the diffusion-reaction model (Mullin, 1992).

The two stages can be represented by their own mass transfer equations, however due to the use of an interfacial concentration C_i that is difficult to measure they are combined into a general crystal growth equation:

$$\frac{dm}{dt} = k_d A(C - C_i) \tag{2.20}$$

$$\frac{dm}{dt} = k_r A(C_i - C^*) \tag{2.21}$$

$$\frac{dm}{dt} = K_G A(C - C^*)^g \tag{2.22}$$

where k_d is the diffusive mass transfer coefficient, k_r is the 'reaction' rate constant, K_G is the overall crystal growth coefficient and g is the order of growth. If $g = 1$, the coefficients follow the relationship shown in Equation 2.23; this relationship suggests that if one of k_d or k_r are very large, the process of crystal growth will be controlled by either the surface integration 'reaction' or diffusion respectively (Mullin, 1992).

$$\frac{1}{K_G} = \frac{1}{k_d} + \frac{1}{k_r} \tag{2.23}$$

Equation 2.22 can be rearranged to produce a term known as the mass deposition rate R_G, which relates to the linear growth rate G of a crystal based on its volume and surface shape factors α and β respectively:

$$R_G = \frac{1}{A}\frac{dm}{dt} = K_G \Delta C^g = \frac{3\alpha}{\beta}\rho_c G \tag{2.24}$$

$$\alpha = V_c/L^3 \tag{2.25}$$

$$\beta = A_c/L^2 \tag{2.26}$$

where ρ_c is the crystal density, V_c is the crystal volume, A_c is the crystal surface area and L is the characteristic length of the crystal shape (Brown, 2012).

2.1.6. Agglomeration

A population of crystals is not always entirely made up of single crystals with the expected morphology; particles sometimes cluster together through the process of agglomeration, which can be divided into primary and secondary mechanisms.

Primary agglomerates are formed when crystals nucleate and grow from an existing crystal body, resulting in features such as twinning (two crystals joined along a particular axis) or dendrites (multiple growths from the surface of a single core crystal) (Jones, 1993). Clusters formed from dendritic growth can form organised arrangements based on a crystal's morphology, such as the six-member clusters observed in hexamethylene tetramine (Davey and Rütti, 1976).

Secondary agglomeration is the more widely understood and accepted mechanism when referring to just 'agglomeration' as a broader term. This involves the aggregation of multiple suspended particles that collide with each other during agitation. Whether particles will remain agglomerated depends on the surface properties, supersaturation and agitation speed, which all in turn affect the bond energy between particles (Brunsteiner et al., 2005). Finer particles typically appear to agglomerate more frequently than larger crystals, due to differences in the level of surface definition (Jones, 1993).

Agglomeration poses a great challenge to crystallisation process modelling, as the growth rate becomes more difficult to measure through conventional in-process techniques (Brown et al., 2017). Agglomerates are still detected as individual particles, therefore the overall particle size being measured ends up overestimated compared to the primary CSD.

2.1.7. Polymorphism

The ability of a molecule to arrange itself into multiple solid crystal forms is known as polymorphism; McCrone (1965) stated that generally all molecules have the potential to exhibit polymorphism, and the number of known forms is proportional to the time and money spent researching a given substance. There will always be a polymorph at any given temperature that

is the most thermodynamically stable, into which metastable forms will transform through various mechanisms. Solid state transformations are kinetically slow due to the high activation energy, while solution mediated polymorphic transformations are easier and more common (Sato, 1993). However, solid state transformations of excipients or active ingredients during pharmaceutical product processing and storage is a known issue, especially when dealing with metastable forms which require investigation on the molecular level (Anwar and Zahn, 2017).

Considering an example substance with two polymorphic forms, the stability of the system can either be described as monotropic or enantiotropic depending on the solubility profiles of each form, as described by Figure 2.8. In contrast to a monotropic system where one polymorph remains more soluble at any given temperature, enantiotropic systems involve a transition temperature beyond which the solubility curves of each polymorph cross, resulting in a change in which form is most soluble and therefore stable.

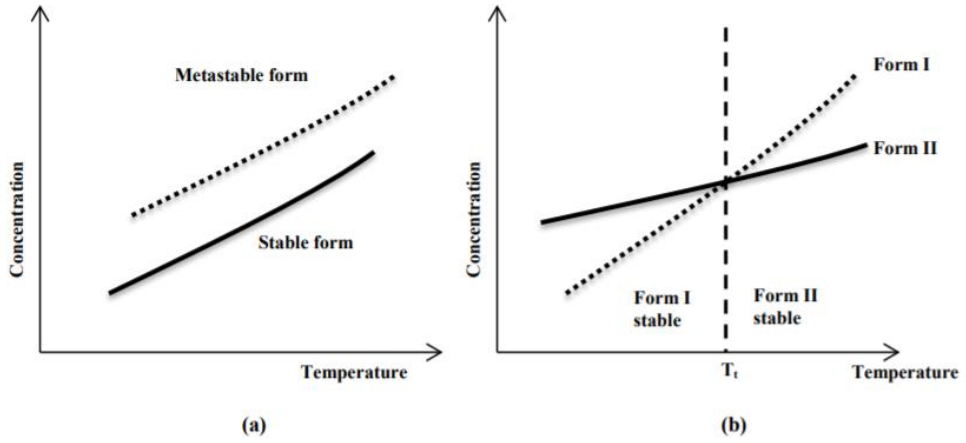


Figure 2.8. Solubility profiles of (a) monotropic and (b) enantiotropic systems (Simone, 2015).

Prediction of likely polymorphic forms has been aided recently through advances in molecular modelling, where the close contacts between molecules in a crystal lattice can be identified and related to certain structural features which can further be linked to nucleation and growth processes (Zhang et al., 2019).

2.2. Process modelling of crystallisation

2.2.1. Population balance model (PBM)

A description of the number and size of crystals developing over time in a crystallisation process can be described through a PBM, which combine the nucleation and growth rates with additional 'birth and death' terms associated with breakage/agglomeration of present crystals. The PBM varies in complexity depending on the number of dimensions considered to represent

the exact particle shape (Shaikh et al., 2015). A basic form of the 1D PBM for a single polymorph can be constructed in the following way:

$$\frac{\partial n}{\partial t} + \frac{\partial (Gn)}{\partial L} = B - D + J \tag{2.27}$$

where n is the population density, B is the birth term and D is the death term (Hounslow et al., 1988; Maggioni and Mazzotti, 2013). These different terms can be added or removed depending on the assumptions made, including whether any crystals are initially present, whether breakage or agglomeration occur and whether the system is well-mixed (Maggioni and Mazzotti, 2013). Once all appropriate kinetic parameters have been considered, the growth of crystals must be connected to the change in solute concentration of the liquid via a mass balance:

$$\frac{dc}{dt} = -3\rho_c K_v \int_0^\infty nGL^2 dL \tag{2.28}$$

where K_v is the volumetric shape factor of the crystals formed (Jha et al., 2017).

Vollmer and Raisch (2006) presented an interesting approach towards inverting the PBM to work backwards from a known CSD. Given that the particle shape is known and fixed, a single cooling profile leading to such CSD could be deduced. Although this may not be directly applicable to the process model developed in this project, their study highlights the importance of implementing optimization features into a crystal size control scheme.

In addition to the nucleation rate expression given by Mullin (1992), Nývlt's polythermal method leads to a simpler power law expression for the primary nucleation rate:

$$J = k_n \Delta C^n \tag{2.29}$$

where k_n is the primary nucleation rate constant and n is the order of primary nucleation (Camacho Corzo et al., 2017). Similar methods are available that express the supersaturation as either S (Sangwal and Wójcik, 2009) or ΔT (Kubota, 2008) rather than ΔC.

A comparison between these different approaches on the nucleation kinetics of ammonium oxalate monohydrate in water was performed by Gherras and Fevotte (2012), where it was found that despite Nývlt's and Kubota's methods leading to much lower values than Sangwal's, the nucleation orders determined by each method consistently fell within the acceptable range for organic compounds. This study also highlighted that although each approach characterises

changes with temperature well, they don't capture the stochastic, unpredictable nature of nucleation at high supersaturations (Gherras and Fevotte, 2012).

Supporting the movement towards first-principles modelling, the Kashchiev-Borissova-Hammond-Roberts (KBHR) method links Mullin's classical equations for primary nucleation to two potential mechanisms: instantaneous nucleation (IN) where all nuclei emerge simultaneously, or progressive nucleation (PN) where nuclei continue to appear throughout the crystallisation process (Kashchiev et al., 2010; Camacho Corzo et al., 2017). If the slope of the linear fit between logarithms of critical undercooling μ_c (calculated via Equation 2.30) and cooling rate q is either greater or less than 3, nucleation is considered progressive or instantaneous respectively.

$$\mu_c = \frac{\Delta T}{T_e} \tag{2.30}$$

where T_e is the equilibrium temperature.

Multiple crystallisation systems involving PN have been studied, including methyl stearate from hydrocarbon solutions and carbamazepine-saccharin co-crystals from ethanolic solutions (Camacho et al., 2014; Mohammad et al., 2017). These cases are useful examples of how the empirical factor A in Mullin's classical equation (Equation 2.12) can be determined using the relationship between q and μ_c as outlined below. However, it should be noted that accurate measurement of the number of particles at the determined nucleation point N_{det} is also required (Camacho Corzo et al., 2017).

$$\ln(q) = \ln(q_0) + 3\ln(\mu_c) - \frac{b_0}{1-\mu_c}{\mu_c}^2 \tag{2.31}$$

$$q_0 = \frac{VAT_e}{N_{det}2b_0} \tag{2.32}$$

$$b_0 = \frac{\alpha v^2 {\gamma_e}^3}{k_B T_e \lambda^2} \tag{2.33}$$

where λ is the molecular latent heat of crystallisation.

Although information on the concentration of nuclei can be obtained in cases involving IN, the acquisition of useful kinetic data is unfortunately much less documented than PN, therefore in order to quantify the primary nucleation rate an empirical power law obtained via the polythermal method remains the only solution (Toroz et al., 2015). In order to ensure a complete transition towards first-principles modelling in the future, IN needs to be better

understood from a population balance standpoint, or avoided through means such as alternative solvent selection.

Regarding equations for growth kinetics, most tend to be characterised through some form of basic power law similar to Equation 2.24. Similar to nucleation kinetics, power laws for growth can be written in terms of different types of supersaturation, most commonly ΔC. For the theory of two-step growth based on diffusive mass transfer and surface integration, subtle differences exist between the available Sherwood correlations for k_d depending on the crystal size:

< 60 µm:
$$k_d(L) = \frac{D_{AB}}{L}\left[2 + 0.8\left(\frac{\bar{\varepsilon}L^4}{\nu^3}\right)^{1/5}\left(\frac{\nu}{D_{AB}}\right)^{1/3}\right] \tag{2.34}$$

> 60 µm:
$$k_d(L) = \frac{D_{AB}}{L}\left[2 + 0.47\left(\frac{\bar{\varepsilon}^{1/3}L^{4/3}}{\nu}\right)^{0.62}\left(\frac{d_s}{d_v}\right)^{0.17}\left(\frac{\nu}{D_{AB}}\right)^{0.36}\right] \tag{2.35}$$

$$D_{AB} = \alpha\frac{kT}{6\pi\eta\frac{d_m}{2}} \tag{2.36}$$

where ε is the energy dissipation rate, ν is the solution kinematic viscosity, d_s and d_v are the stirrer and internal vessel diameters respectively, α is the diffusivity correction factor, η is the solution dynamic viscosity and d_m is the molecular diameter of the solute (Garside et al., 1990; Levins and Glastonbury, 1972).

Both the mechanisms of agglomeration and breakage are expressed by the birth and death of crystals, otherwise considered as the transfer of crystals from one size class to another (Verkoeijen et al., 2002). Integral functions used to represent each of these have been presented by Costa et al. (2007):

$$B_{agg}(v) = \frac{1}{2}\int_0^v \beta(v-\varepsilon,\varepsilon)n(v-\varepsilon,t)n(\varepsilon,t)\mathrm{d}\varepsilon \tag{2.37}$$

$$D_{agg}(v) = \int_0^\infty \beta(v,\varepsilon)n(\varepsilon,t)\mathrm{d}\varepsilon \tag{2.38}$$

$$B_{br}(v) = \int_v^\infty \gamma(\varepsilon)b(\varepsilon)p\left(\frac{v}{\varepsilon}\right)n(\varepsilon,t)\mathrm{d}\varepsilon \tag{2.39}$$

$$D_{br}(v) = -b(v)n(v,t) \tag{2.40}$$

where $\beta(v,\varepsilon)$ is the agglomeration kernel (frequency of collisions between particles of volumes v and ε to produce a particle of volume $v+\varepsilon$), $\gamma(\varepsilon)$ is the number of new particles generated from

breakage of a particle of volume ε, $b(\varepsilon)$ is the breakup rate of a particle with size ε and $p(v/\varepsilon)$ is the fraction of new particles with size between v and $v+\mathrm{d}v$ generated from breakage.

Expressing these terms with respect to particle volume rather than size was contested by Mahoney and Ramkrishna (2002), suggesting these would be inapplicable in certain systems where the mean crystal size is small enough for the volumetric growth rate to approach zero when $L = 0$. Formulating birth and death rates as a function of particle size makes agglomeration highly complex and difficult to evaluate. Therefore, it is common to only consider breakage via attrition in the form of a secondary nucleation rate J_{sec}. Evans et al. (1974) describe this rate using equations for both impeller-crystal collisions ($J_{sec,ci}$) and crystal-crystal collisions ($J_{sec,cc}$):

$$J_{sec,ci} = k_{n,ci}\sigma^{n_{ci}}\frac{N_Q}{N_P}K_v\rho_c\varepsilon\int_{L_{min,ci}}^{\infty} nL^3\mathrm{d}L \tag{2.41}$$

$$J_{sec,cc} = k_{n,cc}\sigma^{n_{cc}}\rho_c\varepsilon^{5/4}\bar{L}^4\int_{L_{min,cc}}^{\infty} nL^0\mathrm{d}L \tag{2.42}$$

where $k_{n,ci}$ and $k_{n,cc}$ are the secondary nucleation rate constants, n_{ci} and n_{cc} are the orders of supersaturation and $L_{min,ci}$ and $L_{min,cc}$ are the sizes above which crystals are prone to the given type of collision.

The overall rate of secondary nucleation is the sum of the rates for each type of collision, but it is often more simply expressed by the empirical relationship described in Equation 2.43 (Myerson and Ginde, 2002):

$$J_{sec} = k_{n,s}\sigma^b {N_{imp}}^j {M_T}^k \tag{2.43}$$

where $k_{n,s}$ is the secondary nucleation rate constant, N_{imp} is the impeller speed, M_T is the suspension density and b, j and k are empirical constants. Some power laws replace the impeller speed term with the energy dissipation rate ε, which represents the energy transfer involved in particle collisions (Xu et al., 2020).

2.2.2. Solution of the PBM

The standard method of moments is most typically applied for the solution of the PBM, which transforms the partial differential equation into a series of ordinary differential equations. Each of these equations can be solved to represent certain factors of the CSD, including crystal number, size, area and volume (Fujiwara et al., 2015). The equations below outline the process

of converting the PBM into a system of moments' equations when applied to the 1D PBM (Equation 2.27):

$$\mu_j = \int_0^\infty L^j n \, \mathrm{d}L \tag{2.44}$$

$$\frac{\mathrm{d}\mu_0}{\mathrm{d}t} = J \tag{2.45}$$

$$\frac{\mathrm{d}\mu_j}{\mathrm{d}t} = jG\mu_{j-1} \tag{2.46}$$

where μ_j is the jth moment of the CSD (Randolph and Larson, 1988).

The extended method of moments developed by Falola et al. (2013) resolved the inability of the standard method to handle size-dependent terms related to agglomeration and breakage. This method also allowed multiple moments to be tracked over time. However, the extended method still faces drawbacks given that multiple CSDs can have the same moments, meaning prior information about the distribution is required to accurately reconstruct it.

Discretisation methods alternatively calculate the size distribution directly as opposed to specific features from the moments, at the cost of being more computationally demanding (Omar and Rohani, 2017). In the context of crystallisation, these methods involve the division of crystal size into discrete bins of a given number and size, which must be optimised to save computational power. The method of classes developed by Marchal et al. (1988) is one such method applied to PBM solution, which in its simplest form in the case of batch crystallisation (no flow in/out) and negligible agglomeration/breakage effects can be described by the following equation:

$$\frac{dn_i}{dt} + \frac{G}{2\Delta C_{i+1}} n_{i-1} - \frac{G}{2\Delta C_{i-1}} n_i = 0 \tag{2.47}$$

where ΔC_i is the width of size class i. The process modelling software gPROMS FormulatedProducts utilises this method alongside high resolution finite differencing algorithms to optimise the classes and adaptive time-stepping to improve computational stability (Mitchell, 2012).

Monte Carlo methods involve the use of probability distributions to randomly decide whether a particular event occurs, such as growth, agglomeration or nucleation, to generate solutions to the PBM. These methods well represent the stochastic nature of mechanisms such as primary nucleation and agglomeration, but due to the need to track the state of every particle in a

population size typically between 10^3-10^6, they can be quite computationally demanding (Omar and Rohani, 2017).

There are four main types of Monte Carlo solution methods. Constant-number methods involve altering the process volume while keeping the original particle count (Lin et al., 2002). Constant-volume methods involve changing the number of particles within a fixed volume by considering the frequency of nucleation, agglomeration and breakage events (Maisels et al., 2004). Multi-Monte Carlo combines features from both of these methods (Zhao et al., 2005). Finally, direct simulation methods involve a periodic regulation of the number of particles across a given simulation time (Liffman, 1992). Zhao et al. (2007) compared all four methods and found that they all performed to a similar degree of accuracy, which was more linked to the total number of particles in the simulation.

2.3. Process measurement techniques

Through findings that a more accurate CSD can be achieved through on-line temperature optimisation instead of off-line, Zhang and Rohani (2003) stated that:

> "The ultimate objective in the control of crystallization processes is to develop a truly on-line optimal control algorithm with constraint-handling capabilities in the absence of sufficient process measurements".

During batch cooling crystallisation, temperature can easily be measured on-line for control purposes using thermocouples, however other properties have proven to be difficult to obtain. Thanks to advancements in process analytical technology over the past two decades or so, monitoring of solute concentration and particle size/count has become much more viable for control applications in dispersed systems. These methods can overcome common drawbacks during crystallisation control, such as moving the objective function away from maximising the weight-average crystal size to prevent the overproduction of fines (Nagy and Braatz, 2012).

2.3.1. Concentration measurement

Solute concentration measurement is vital due to its connection to supersaturation, which should usually be maintained in an optimal range that is high enough for significant crystal growth but low enough to limit nucleation (Nagy et al., 2008a). ATR-FTIR spectrometry can be applied for on-line solute concentration monitoring (Doki et al., 2004; Borissova et al., 2009). As depicted in Figure 2.9, an IR beam is passed through an ATR crystal (most commonly germanium) and reflects back to a detector from the crystal surface in contact with

the measured solution; FTIR spectroscopy is then applied to relate IR absorbance to concentration (Nagy and Braatz, 2012).

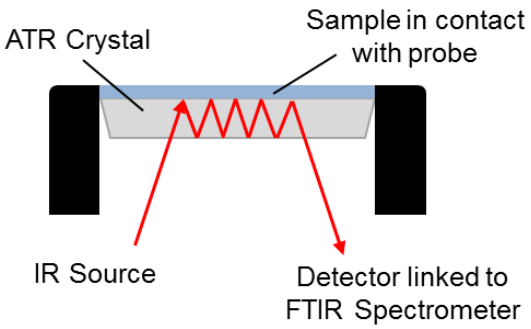


Figure 2.9. Schematic representation of ATR-FTIR probe.

Initially, ATR-FTIR was praised as a highly accurate measurement technique with precisions up to ±0.00014 g solute/g solvent and implemented into a range of crystallisation studies. Its effectiveness in process control was eventually explored by various studies (Fujiwara et al., 2002; Doki et al., 2004; Nagy et al., 2008a; Khan et al., 2011). Nagy et al.'s (2008a) study on the comparison of temperature and concentration control of paracetamol crystallisation was one example of how this developed, where the latter enabled growth of larger crystals with a more uniform size distribution and better structural quality.

Concentration measurement can also be achieved through use of ultraviolet light, known as ATR-UV/Vis spectroscopy. This involves the excitation of electrons through UV absorption followed by correlation of the absorbance spectrum to solution concentration, although applications are limited due to the transparency of most functional groups within the typical operation range of 190-800 nm (Bakeev, 2010). ATR-FTIR is usually favoured over this as it is less affected by suspended particles within the monitored solution (Zhang et al., 2017).

2.3.2. Particle size and count measurement

As crystallisation modelling requires information on the nucleation and growth kinetics, it is important to deduce an effective way of tracking the number and size of crystals developing over time (Doki et al., 2004). FBRM involves the use of a rotating laser optical probe to measure chord lengths through crystals as illustrated in Figure 2.10. The chord length distribution (CLD) can be mathematically converted into a CSD through geometric inverse modelling, though this can be relatively complex depending upon the shape (particularly aspect ratio) of the particles being measured (Nagy et al., 2008a; Kumar et al., 2013). As well as supporting crystallisation models, FBRM's potential for on-line particle sizing has been suggested as the basis of robust model-free control of crystallisers (Nagy and Braatz, 2012).

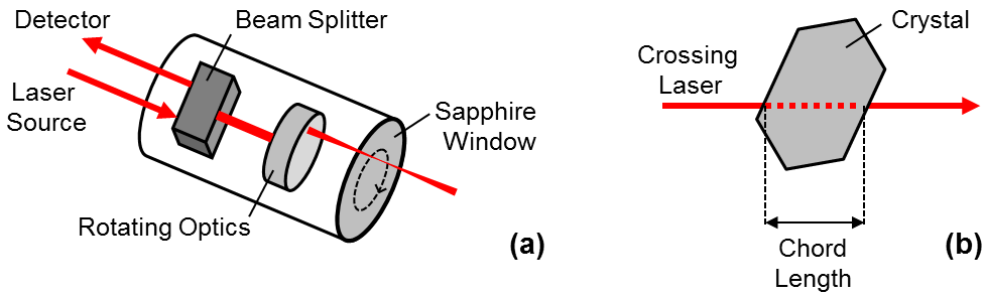


Figure 2.10. (a) FBRM probe schematic and (b) chord length measurement of a crystal (Kumar et al., 2013).

Although developments in this technique have corrected certain drawbacks such as fouling on the probe window, a major problem still exists in the method of measurement via chord lengths. The accuracy of crystal sizes calculated via this method are heavily influenced by the crystal shape; for example, a chord length measured through the width of a needle-like crystal with a large aspect ratio provides poor representation of the average size (Szilágyi and Nagy, 2018). Abu Bakar et al. (2009) noted this issue after morphology changes between seeded and product crystals, which suggested a falsely significant change in the CSD. As a result, FBRM is a more useful technique for measuring particle count over time and relating this to both the primary and secondary nucleation rates, depending on the operating conditions.

The PVM probe is an in-situ video microscopy technique that can be used to obtain images for size and shape analysis (Braatz and Hasebe, 2001). Images are captured at a rate of around 10-30 per second, then processed to provide real time particle size and count information. However, the reported size data can sometimes be unreliable depending on the image quality, as factors such as the crystal size and slurry density can hinder the detection of individual particles. Instead, PVM is more useful for monitoring crystal shape development, and observing the key crystallisation mechanisms that are occurring during the process such as nucleation or polymorphic transformations (Barret and Glennon, 2002; De Anda et al., 2005).

2.4. Process control of crystallisation

The two most important factors in the improvement of crystallisation process control are advancements in understanding of crystallisation mechanisms, and the development of more robust and accurate PAT (Nagy and Braatz, 2012). Measurement-based or ‘model-free’ control methods use PAT to operate the process temperature and solution concentration within the metastable-zone, offering a fast approach to design that can be easily automated (Nagy et al., 2008b). However, these operating curves do not always lead to the most optimal solution and are strongly affected by external disturbances. Model-based approaches implement detailed

mathematical models of the crystallisation process, which can be used to determine more optimal process conditions with a shorter experimental time requirement (Togkalidou et al., 2001). Model development can also be time and resource demanding, especially when complex practical properties such as product filterability or purity need to be considered (Nagy et al., 2008b). Figure 2.11 shows a general schematic comparing both model-free and model-based approaches (Nagy et al., 2008b).

Although industry still favours measurement-based control due to its simplicity, the application of predictive control methods such as MPC to crystallisation has proved to be an area worth further exploration (Orehek et al., 2021). If resources are available to put towards model development, MPC provides advantages that include faster reaction to changes within a process, as well as better handling of complex, multivariable systems (Jha et al., 2017).

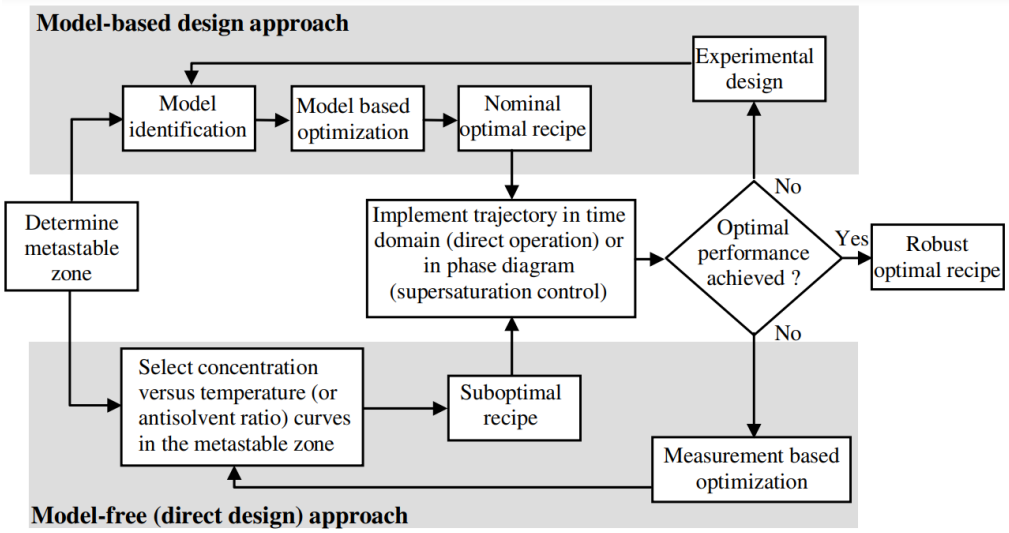


Figure 2.11. Schematic representation of the model-based and model-free approaches for crystallisation (Nagy et al., 2008b).

2.4.1. Measurement-based (model-free) control

Proportional-integral-derivative (PID) controllers are the simplest and most efficient control method available, implemented in around 90% of industrial process control applications due to their ease of use (Levine, 1996). The control signal $u(t)$ from a PID controller can be described by Equation 2.48:

$$u(t) = k_P e(t) + k_I \int_0^t e(\tau)\mathrm{d}\tau - k_D \frac{\mathrm{d}}{\mathrm{d}t} y(t) \tag{2.48}$$

where e is the control error, $y(t)$ is the process variable and k_P, k_I and k_D are the proportional, integral and derivative gains respectively (Ang et al., 2005). The last three terms cover different functions within the control scheme:

- k_P – provides an overall control signal proportional to the error.
- k_I – reduces steady-state error but extends response time.
- k_D – reduces transient response time.

In batch cooling crystallisation, the aim of implementing PID controllers is to maintain a concentration or temperature set point. As these are both linked through the solubility of the system, temperature can be manipulated to vary the supersaturation via a feedback loop to equipment such as a heating jacket around the crystallisation vessel (Jha et al., 2017).

Model-free controllers operate much more simply than control algorithms based on first principles. For example, Fujiwara et al. (2002) calibrated a temperature controller to follow a set supersaturation profile determined empirically, which would remain close to the metastable limit and therefore ensure optimum crystal growth. This approach to control is usually referred to as direct design, which attempts to overcome the uncertainties on kinetic parameters using engineering understanding of the metastable zone (Fujiwara et al., 2005).

A major downside to this approach is the empirical nature of the controller development, requiring large amounts of experimental data to quantify mechanisms that in practice vary widely depending on the chosen operating conditions. Khan et al. (2011) studied supersaturation control during the batch cooling crystallisation of L-glutamic acid at both 20 L and 250 L scales. Each scale required a different cooling profile to maintain the same level of supersaturation, and variation in the product CSD was influenced by the more prevalent secondary nucleation occurring in the smaller scale.

2.4.2. Empirical vs first-principles model-based control

PID controllers alone cannot provide satisfactory response times and stability around the temperature set point. The progression from measurement-based to model-based control has helped to quantify the necessary understanding of a batch crystallisation process in a more generally applicable form. Due to working with a low number of dimensions, process models can be utilised to construct control schemes through the solution of a small set of ordinary differential equations (ODEs), allowing for minimal computation time and complexity (Christofides et al., 2007). This allows space for model-based algorithms to control specific crystal properties such as size and shape more accurately than measurement-based methods.

Jha et al. (2017) performed a detailed comparison between this and two types of model-based control methods: generic model control (GMC) and MPC, the latter of which is described in

more detail later in this section. The model developed by Lee and Sullivan (1988) for GMC had a simplistic structure to mirror the proportional-integral part of PID control. However, the ability to handle nonlinearity allowed the rate of temperature change to be selected as a variable. MPC acts more closely to feedforward control, where estimated parameters are used to minimise error. Overall, it was found that PID control resulted in an 18% overshoot past the temperature set point, higher than the 10% and 1% achieved by GMC and MPC respectively; similar trends were displayed for the settling time.

The enhanced performance of control systems by introducing modelling is clear, but there are still drawbacks that must be handled appropriately. Due to their initial basis on many uncertainties, the transition between different equipment (vessels, heaters etc.) can invalidate model-based control if they lack enough similarity, especially in feedforward control (Forgione et al., 2015). Models can be aided either through additional experimental analysis of crystallisation systems and physical data, or by combining different types models to cover the precision of specific areas such as computational fluid dynamics (CFD) (Bermingham et al., 2007; Forgione et al., 2015). Dealing with impurities also remains a great challenge to crystallisation modelling, as the shape variation they generate through inhibition of non-unidirectional growth and potential changes to agglomeration and breakage are difficult to characterise (Braatz, 2002).

The ideal case for a model is to be derived solely from basic physical concepts, removing the need for prior experimentation. In real cases, there is always a minimum portion of a model that is empirical; so-called ‘first-principles’ models will often involve minor experimentation to aid in model formulation and parameter estimation, as opposed to a fully empirical model that requires a large quantity of rigorous experimental characterisation to function sufficiently (Miller and Rawlings, 1994).

Empirical models used for crystallisation control operate using previous batch data, for which a large number of repeated experiments are required for the input of data to estimate model parameters. Forgione et al. (2015) investigated two forms of ‘batch-to-batch’ iterative models that use this approach, namely Iterative Learning Control (ILC) and Iterative Identification Control (IIC). Each have complementary advantages and disadvantages based on their potential performance quality. IIC has the potential to perform strongest, but widely varies due to its reliance on model assumptions for parameter estimation; performance could be improved by

developing an algorithm that selects ILC or IIC between batches depending on which could most effectively utilise the previous batch's outputs.

ILC generates a new cooling profile by comparing the supersaturation to outputs from a previous batch, therefore updating the model to match the actual supersaturation more closely. External, real-time disturbances in properties such as temperature can easily have a negative impact on the performance of ILC since it is a form of feedforward control, although this can be reduced by incorporating elements of measurement-based PID control. Griffin et al. (2016) paired ATR-FTIR with chord length and particle count measurements from FBRM to develop a dynamic model for controlling the average crystal size and yield within set batch times. This was made achievable by limiting the dimensionality of the model, compromising between computational efficiency and the level of detail and accuracy that a PBM could capture.

Aiming to reduce extensive experimental times, the first-principles approach to crystallisation modelling is focused around connecting material and energy balances to crystal population in order to optimise a given function. Additionally, parameter estimation can be shortened further using prior knowledge of the system being controlled, or by obtaining initial estimates during characterisation of the MSZW (Fujiwara et al., 2005). This approach to model design has been applied to crystals of varying morphology, such as in Miller and Rawlings' (1994) study on plate-like naphthalene crystals grown from solution in toluene. First-principles experimental design and modelling was found to be effective in forming a confident representation of the crystallisation kinetics needed to find the optimal cooling conditions.

2.4.3. Model predictive control (MPC)

A stronger form of feedforward control that allows a PBM to be implemented is MPC, which involves the generation of control moves made to the manipulated variable that most efficiently reach and maintain a controlled variable set-point (Jha et al., 2017). MPC benefits from being able to handle complex, dynamic multiple input/output systems, making it superior over other forms of control for batch cooling crystallisation due to the change in nucleation and growth rates as the supersaturation varies (Kwon et al., 2015). That being said, most successful MPC applications for controlling crystal size have been shown for continuous processes rather than batch (Moldovány et al, 2005; Yang and Nagy, 2015; Tahir et al., 2017). Lacking measurement of process parameters besides temperature has previously limited MPC applications in crystallisation, however advancements in PAT have introduced new possibilities (Nagy et al., 2008; Nagy and Braatz, 2012). Additionally, the objective function of the model, although

usually a specific CSD, could also be changed to crystal purity, morphology or bulk density, relating to performance properties such as filterability.

One example through which MPC has achieved these objective functions is described by the lysozyme protein crystallisation by Shi et al. (2005). In their study, the number of fines in the final product needed to be reduced. An increase in the volume-averaged crystal size was achieved using a hybrid MPC process; to conserve stability, a bounded nonlinear controller could be switched to conserve stability. Although, these switches were done as infrequently as possible due to the bounded controller's poorer overall performance.

An effective real-time process monitoring and control approach, built using the software PharmaMV (Perceptive Engineering Ltd.), was demonstrated by Tahir et al. (2017) to control the median crystal size within a continuous oscillatory flow crystalliser. This technique benefitted from the steady-state nature of continuous processes where a constant set-point for the median size could be applied. However, this would need to be expanded upon for a batch crystallisation, given that the crystal size is expected to change constantly as it develops towards a certain end-point throughout the process.

Most industries are undergoing a digital transformation, using technology in new ways to improve efficiency. For the fine chemical industry there exists the concept of Pharma 4.0, built on a foundation of digital maturity to develop resources, organisation and processes, information systems and the working culture (Romero-Torres, 2021). Regarding crystallisation, the use of technology has been studied thoroughly to provide more accurate process models. Szilágyi and Nagy (2016) showed how parallel GPUs running a high-resolution finite volume technique improved calculation speed significantly, enough for real-time resolution of a multi-dimensional PBM. Camacho Corzo et al. (2020) used CFD to investigate the hydrodynamics inside a batch crystalliser, suggesting that these simulations could be used to construct a compartmental model with which a morphological PBM (Ma and Roberts, 2018) could be integrated.

Another recent advancement for MPC of crystallisation processes is the use of open-loop simulations with the PBM to train machine-learning control algorithms, such as in the study by Zheng et al. (2022). They demonstrated how a recurrent neural network (RNN) model trained using a semi-empirical PBM could improve the computational efficiency of MPC to achieve an optimal target product yield and crystal size. However, advanced control methods such as this have only been tested in-silico, noting particular difficulties that may arise when applying

these MPC methods on a real crystalliser, such as the accuracy of on-line measurements. A reinforced-learning approach to improve the control of temperature, supersaturation and crystal size trajectories was taken by Benyahia et al. (2021) for the batch cooling crystallisation of paracetamol in water. This control method showed a significant improvement over MPC for controlling the crystal size, however MPC performed similarly for both temperature and supersaturation, therefore remaining a relevant and applicable approach to crystallisation process control.

2.5. Conclusions

This literature review has summarised the fundamentals of crystallisation, implementation of the PBM for predicting crystal size, and recent model-based methods of controlling these processes. From this, the need for a robust MPC strategy for achieving precise control over the product CSD from a batch crystallisation process is evident. This work aims to outline a methodology that can be adopted by industry to optimise their crystallisation processes, as well as a review of how commercial tools perform in a continuously advancing digital age.

3. Materials and Instrumentation

The solute-solvent system used in experiments and the compounds selected for screening of a suitable crystal growth rate inhibitor are described in this chapter. In addition, the experimental setups for batch crystallisation at different scales are described, along with the necessary PAT for in-situ measurements. The equipment used for seed crystal preparation and off-line analysis of the product crystals is described, along with instruments used for measurement of the solute-solvent system's key physical properties. Brief explanations of any standard experimental procedures have been provided.

3.1. Materials

An appropriate solute-solvent system was chosen that met certain criteria to reduce the complexity of the initial process model development. As well as having a good temperature dependence of solubility to ensure a high yield from large-scale crystallisation, the system was to have a crystal morphology resulting in equant growth in all directions, making it suitable for a 1D PBM.

3.1.1. Hexamethylene tetramine (hexamine) and ethanol solvent

The crystallisation system chosen for this study was hexamethylene tetramine (hexamine, purity >99%, Scientific Laboratory Supplies) crystallized from solution in pure ethanol (purity >99.8%, Sigma Aldrich). This material was used consistently throughout the experiments described in all following chapters.

Hexamine is widely used in applications such as the treatment of urinary tract infections and the manufacture of fuel tablets that burn smokelessly (Miller, 2008). Hexamine was suitable for study due to its single growth form {110}, developing a rhombic dodecahedron morphology as determined through the Bravais-Friedel-Donner-Harker (BFDH) method (Docherty et al., 1991), which has been shown in Figure 3.1. This is due to its highly symmetrical, body-centred cubic unit cell resulting in relatively equant growth in all directions, therefore being well represented by a 1D PBM (i.e., well represented by a sphere).

Hexamine's solubility in pure water increases as temperature decreases, making it an unsuitable solvent for batch cooling crystallisation. For this reason, pure ethanol was chosen instead to fit the more conventional solubility relationship, at the cost of a slightly reduced potential yield (Wang et al., 2008). Myerson et al. (1986) reported power law rate expressions for nucleation and growth of hexamine in various solvents; the data for crystallisation from ethanol is summarised in Table 3.1.

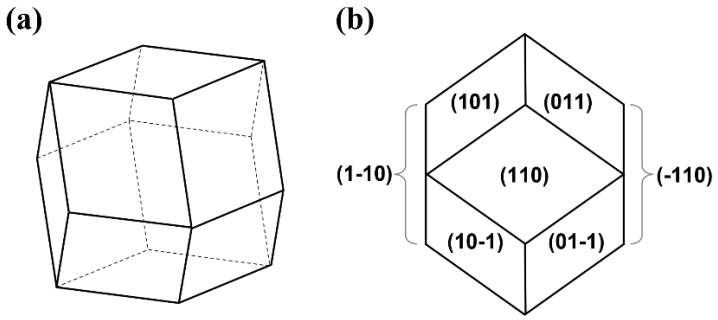


Figure 3.1. a) Predicted morphology of hexamine from the BFDH method, b) morphology as viewed along the [110] faces.

Table 3.1. Solubility, nucleation rate and growth rate expressions for hexamine-ethanol crystallisation reported by Myerson et al. (1986).

Solubility (g/g solution)	$C^* = 1.373 \times 10^{-2} + 5.729 \times 10^{-4}T + 6.707 \times 10^{-6}T^2$
Nucleation rate (#/s.m^3 solvent)	$J = (5.0 \pm 2.0) \times 10^{13}\Delta C^{2.6\pm0.9}$
Growth rate (m/s)	$G = 2.6 \times 10^{-2}\Delta C^{1.95}$

Although unspecified in their study, it was assumed that the nucleation expression would most closely represent secondary nucleation, given its determination via continuous crystallisation. Additionally, the lack of terms that account for agitation and slurry density suggested a poorly accurate description of the system's behaviour. However, this was a reasonable starting point and posed an important challenge to the MPC strategy's performance. If a good level of control could be maintained despite slight differences between the simulated and real crystallisation behaviour, it would be a useful tool for industrial applications that typically face unpredictable disturbances.

Hexamine exhibits an interesting agglomeration mechanism that falls under the category of 'primary' agglomeration described in section 2.1.6. Dendritic growth following from nucleation on the corners of hexamine's morphology leads to clusters of around 6-7 crystals, which incorporate solvent inclusions into the structure upon further crystal growth. (Denbigh and White, 1966; Davey and Rütti, 1976; Gajjar et al., 2021) Figure 3.2 shows examples of both hexamine single crystals and agglomerates. Although this mechanism has been thoroughly understood through the molecular modelling work by Nguyen et al. (2023), the agglomeration rate has not yet been quantified sufficiently for application in a PBM. Therefore, this poses an interesting challenge to the performance of a control strategy when the agglomeration rate cannot easily be modelled.

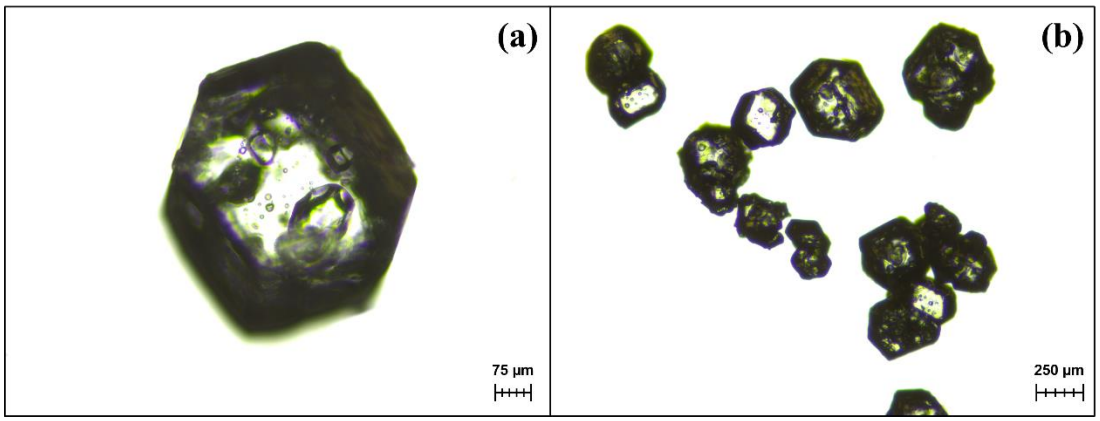


Figure 3.2. Images of a) a single hexamine crystal and b) mixture of single crystals and agglomerated clusters of hexamine, both recrystallised from pure ethanol solution.

3.1.2. Impurities for growth rate inhibition

Based on the favourable functional group interactions in an alcohol solvent environment, compounds containing a carboxylic acid group have the potential to be strongly interactive with the amine functional group in hexamine (Hunter, 2004). For this reason, various compounds were selected for screening of a potential crystal growth rate inhibitor for the hexamine-ethanol crystallisation system. Investigating how the performance of a control strategy is affected by disturbances to the growth kinetics was considered vital for understanding its applicability to industrial crystallisation processes.

Figure 3.3 shows the structural formulae of hexamine and the three growth rate inhibitors used in this work: citric acid (purity >99.5%, Sigma Aldrich), 4-aminobenzoic acid (purity >99%, Alfa Aesar) and adamantane-1-carboxylic acid (purity >99%, Alfa Aesar). Citric acid was selected due to having three carboxylic acid groups in its structure, which would enhance the strength of its interaction. 4-aminobenzoic acid, also referred to as para-aminobenzoic acid (PABA), was selected due to having good solubility in ethanol, which would ensure it remained in solution during hexamine crystallisation (Turner et al., 2016). Finally, adamantane-1-carboxylic acid (A1CA) was selected due to its structural similarity to hexamine.

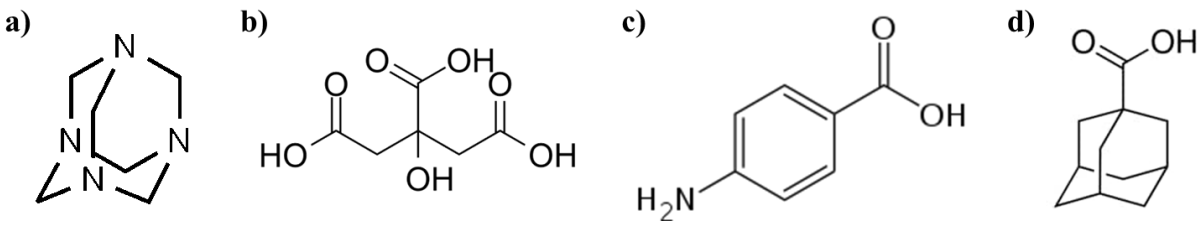


Figure 3.3. Structural formulae of a) hexamine, b) citric acid, c) PABA and d) A1CA.

3.2. Experimental setup for batch crystallisation

3.2.1. 500 ml jacketed crystalliser for characterisation

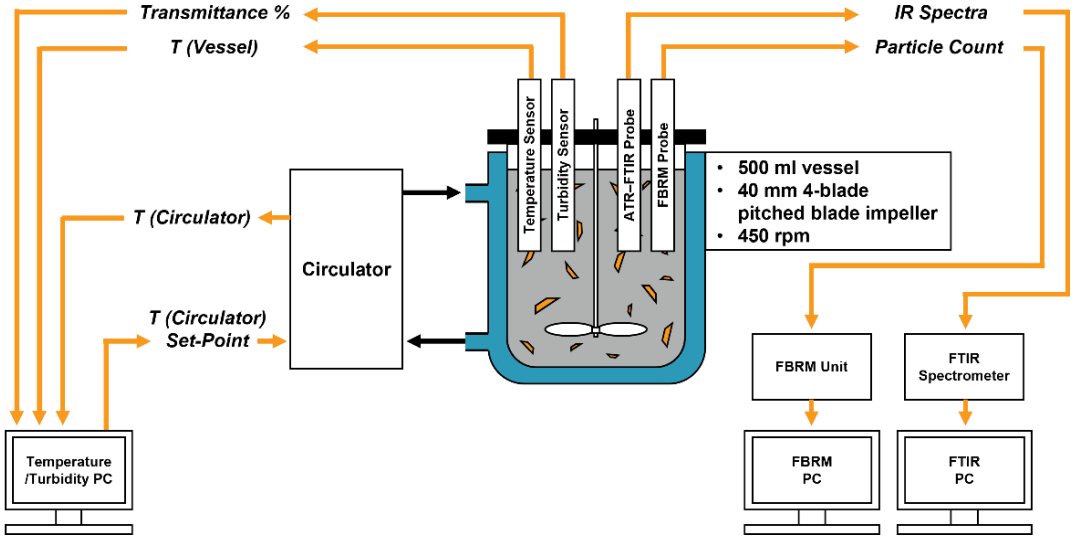


Figure 3.4. Schematic of the 500 ml jacketed crystalliser setup used for characterisation of the hexamine-ethanol crystallisation system with associated in-situ measurement technology.

Figure 3.4 shows a schematic of the experimental setup used to measure the MSZW and carry out seeded batch crystallisations in chapters 4 and 6.

A Radleys 500 ml unbaffled jacketed glass crystallisation vessel with a dish-shaped bottom was clamped to a supporting frame alongside a Radleys RS50 Control Overhead Stirrer, fitted with a 40 mm diameter 4-blade pitched blade impeller. An agitation speed of 450 rpm was used in all batch experiments. The temperature within the vessel was controlled using a Huber Ministat 230 thermostatic bath (operated via in-house LabVIEW software) to circulate water through the jacket and monitored with a PTFE Pt-100 temperature probe connected to the Huber Ministat. An in-house built turbidimetric fibre-optic probe (Grön et al., 2003) was inserted into the vessel to measure the solution turbidity and determine the clear/cloud points during heating/cooling. A Mettler-Toledo Lasentec S400A FBRM probe was used to measure particle count, while an ATR-FTIR probe connected to an ABB MB3000 FTIR spectrometer was used to measure solution concentration. The method for calibration of the IR spectra to concentration is detailed in section 3.3.

3.2.2. 500 ml jacketed crystalliser for MPC validation

Figure 3.5 shows a schematic of the experimental setup used for the physical application and validation of the MPC strategy described in chapters 7 and 8. This setup was available at the CMAC Future Manufacturing Research Hub (University of Strathclyde), where the experiments outlined in these two chapters were carried out.

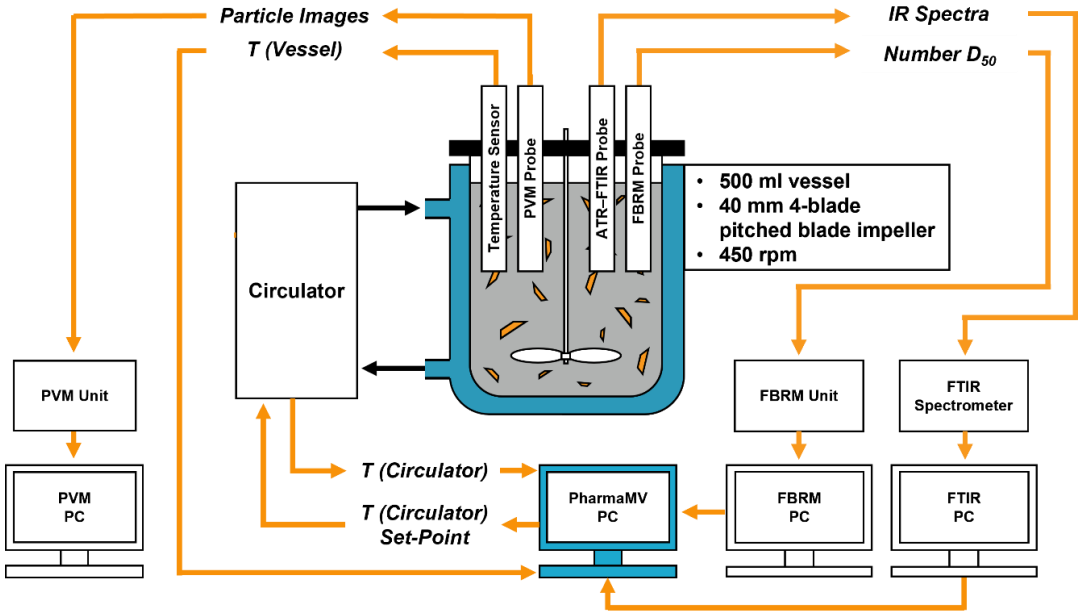


Figure 3.5. Schematic of the 500 ml jacketed crystalliser used for physical validation of the supersaturation and size MPC strategies with associated in-situ measurement technology.

In order to keep the same hydrodynamic conditions, the same models of 500 ml jacketed crystalliser and impeller were implemented as detailed in section 3.2.1, with the same agitation speed of 450 rpm. The temperature was controlled using a Lauda PROLINE RP 855 thermostatic bath, circulating water through the jacket. A Mettler-Toledo ParticleTrack G400 FBRM probe was used to measure the square-weighted primary D_{50} of the suspended crystals, while an ATR-FTIR probe connected to a Mettler-Toledo ReactIR 15 FTIR spectrometer was used to measure solution concentration. The same procedure detailed in section 3.3 was followed for calibration of the IR spectra to concentration. During one of the batch crystallisations with size control described in chapter 8, a Mettler-Toledo ParticleView V819 PVM probe was used to take images of the suspended crystals during the batch. The circulator, FTIR spectrometer and FBRM unit were connected to a PC running PharmaMV software, which calculated appropriate set-points to control the circulator temperature based on MPC strategies described in chapters 7 and 8 for supersaturation and size control respectively.

3.2.3. 1.5 ml Crystal16 system

In chapter 4, the MSZW was also measured using a Technobis Crystal16 system. This equipment holds sixteen 1.5 ml flat bottom vials, which can be programmed in groups of four to follow different temperature profiles. Agitation was provided using 7 mm diameter magnetic stirrers at 900 rpm. Temperature was controlled using a Julabo FL601 cooler to circulate water throughout the Crystal16 system. Transmittance through the solution was measured to determine the clear/cloud points during heating/cooling.

3.2.4. 100 ml AutoMATE system

Figure 3.6 shows a schematic of the HEL AutoMATE parallel reactor setup, which was used for the impurity screening described in chapter 6.

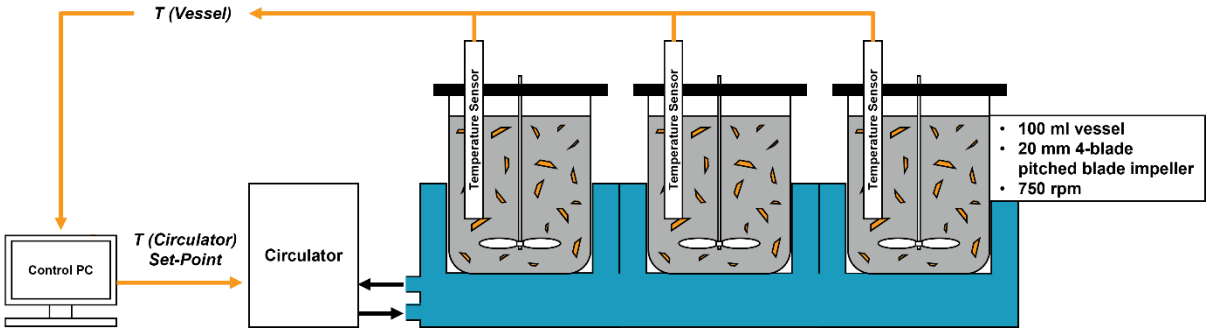


Figure 3.6. Schematic of the AutoMATE parallel 100 ml crystalliser setup used for impurity screening.

This setup consisted of three 100 ml glass-lined crystallisers with overhead stirrers, each fitted with a 20 mm diameter 4-blade pitched blade glass impeller. An agitation speed of 750 rpm was used in all screening experiments. The crystallisers were housed within a platform that allowed the temperature of each vessel to be controlled independently, using water circulated via a Julabo F32 refrigerating/heating circulator and connected PC running the associated WinISO v2.3.135 software.

3.3. ATR-FTIR calibration

In order to determine the concentration in-situ during the crystallisation experiments described in chapters 4 and 6, the IR spectra of hexamine-ethanol solutions at varying temperatures and concentrations were measured using an ABB MB3000 FTIR spectrometer with in-situ ATR probe. A partial least squares (PLS) regression model was calibrated against the known solution concentrations within the accompanying software, HorizonMB.

For each concentration selected, the appropriate amount of hexamine was dissolved in 400 g of ethanol and heated to approximately 5 ºC above the saturation temperature. Once the hexamine had been fully dissolved, spectra were measured every minute, performing 16 scans in the range of 700-1600 cm^{-1} at a resolution of 4 cm^{-1}. After holding the temperature for 10 mins, the solution was rapidly cooled in steps of 2 ºC, holding at each temperature for 10 mins. This procedure continued until primary nucleation occurred, and any spectra measured at constant temperature before the point of nucleation were used in the calibration model. Table 3.2 summarises the range of solution conditions of which IR spectra were measured.

Table 3.2. Solution temperatures and concentrations for which IR spectra were measured for the calibration model in HorizonMB software.

Temperature (ºC)	Concentration (g/g solvent)						
	0.050	0.045	0.040	0.037	0.034	0.030	0.025
44	X						
42	X						
40	X	X					
38	X	X					
36		X	X				
34		X	X				
32			X	X	X		
30			X	X	X		
28				X	X		
26				X	X	X	
24				X	X	X	
22					X	X	
20						X	X
18						X	X
16						X	X

Figure 3.7 shows a typical FTIR spectra measured for a hexamine-ethanol solution. Two distinct peaks characteristic of hexamine were selected to calibrate the spectra against concentration. These were located between wavenumbers of approximately 1227-1246 cm^{-1} and 993-1024 cm^{-1}.

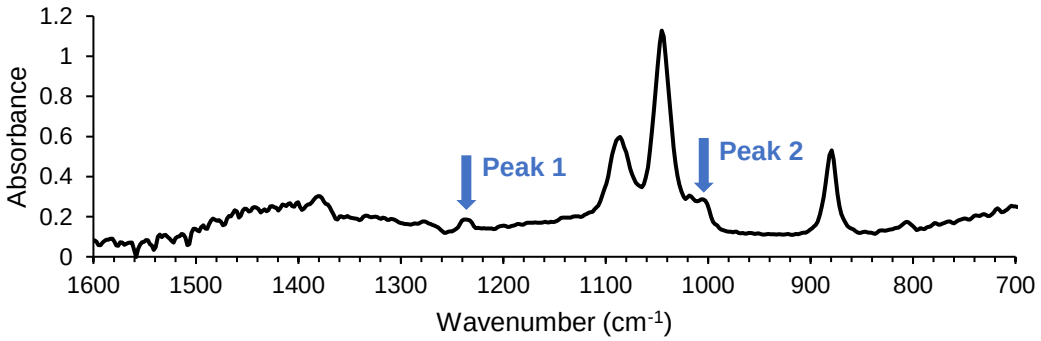


Figure 3.7. Hexamine-ethanol FTIR spectra measured using the ABB MB3000 spectrometer at 36 ºC and 0.04 g/gsolvent. Highlighted peaks are those characteristic of hexamine selected for the calibration model.

Of the total 320 measured spectra, 50 were excluded from the calibration model to be used for model validation. A second derivative transformation using a Savitzky-Golay filter with 17 points of smoothing was applied to the spectra before calibration. The PLS model calculated after pre-processing the calibration spectra was applied to the 50 validation spectra. The model

predicted concentrations and actual values are compared in Figure 3.8, for which data on the ΔC and percentage error are summarised in Table 3.3.

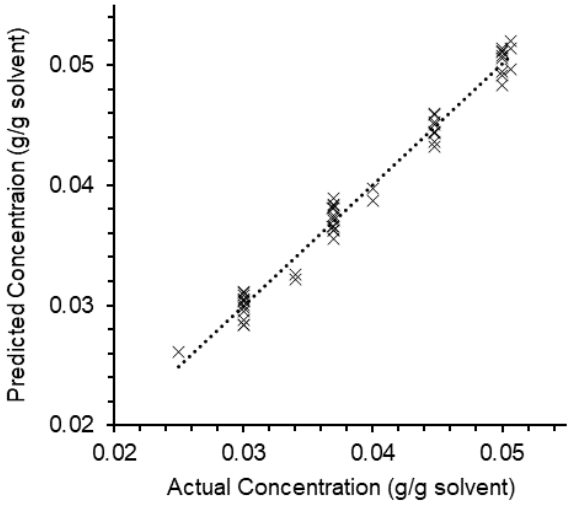


Figure 3.8. Actual vs. predicted concentrations for the PLS calibration model calculated in HorizonMB software.

Table 3.3. Analysis of the PLS calibration model error.

	Minimum	Maximum	Mean	Standard Dev.
$\lvert\Delta C\rvert$ (g/g solvent)	1.20×10^{-5}	1.88×10^{-3}	8.99×10^{-4}	5.00×10^{-4}
Error (%)	0.04	5.54	2.36	1.45

For the experiments described in chapters 7 and 8, a new concentration calibration model was developed within PharmaMV software due to the difference in instrumentation. New spectral measurements were recorded with the ReactIR system, following the same procedure and experimental conditions described above. A second derivative transformation using a Savitzky-Golay filter with 10 points of smoothing was applied to the spectra before calibration. The calibration model was incorporated into the overall strategy as its own MPC block, which communicated the measured concentration to the supersaturation MPC block.

3.4. Seed crystal preparation

Multiple sets of hexamine seed crystals were prepared through different methods. 100 g of the raw hexamine from the supplier was sieved by hand through metallic sieves between 125-250 µm, producing around 50 g of seeds with a volume mean size of 178.4 µm.

A second batch of seeds was prepared through dry milling. 50 g of raw hexamine was put through a ball mill with 1 cm^{-3} yttrium stabilised zirconia milling media for 20 hours. This produced around 40 g of seeds with a volume mean size of 69.4 µm.

A further two batches of seeds were prepared through wet milling, using an IKA magic LAB mixing unit with ULTRA-TURRAX upper transport layer (UTL) dispersing module. For each batch, around 50g raw hexamine from the supplier was suspended in 500 ml of saturated ethanol solution and pumped continuously through the mill and attached hopper. Using a pipette, 10 ml samples were removed every 5 mins to check the progress of milling using a Malvern Mastersizer 3000, stopping the milling process once the desired size had been achieved. Using a fine-toothed rotor at a speed of 20000 rpm for 30 mins, approximately 30 g of seeds with a volume mean size of 68.9 μm were produced. Using a coarse-toothed rotor at 8000 rpm for 20 mins produced approximately 50 g of seeds with a mean size of 112.2 μm.

3.5. Product crystal analysis

Hexamine product crystals from all batch crystallisations were removed and filtered under vacuum using a Buchner funnel. Images were collected using an Olympus BX51 fluorescence microscope, or through scanning electron microscopy (SEM) via a Hitachi TM3030Plus microscope. The number-based and volume-based CSDs of all seed and product crystal samples were measured using a Malvern Morphologi G3 at 5x magnification.

3.6. Physical property measurement

Various physical properties of the solid hexamine crystals and hexamine-ethanol solution were required for the parameterisation of the crystallisation process model in chapter 5.

The densities of saturated hexamine-ethanol solutions were measured using an Anton Parr DMA 4500 M density meter. Saturated solutions were chosen to acquire a density profile that varied with both temperature and concentration; this enabled the most accurate representation of the dynamic solution conditions during cooling crystallisation. Hexamine was dissolved in ethanol to produce five solutions of around 3 ml, saturated at 10 ºC intervals between 10-50 ºC (0.0238-0.0603 g/g solvent. Glass vials containing the solutions were heated to their corresponding saturation temperatures using a water bath. Using a syringe, 1 ml of solution was inserted into the U-tube of the density meter and the sample density was recorded at a constant temperature.

The specific heat capacity of hexamine was measured by differential scanning calorimetry (DSC) using a Mettler-Toledo DSC 3+ calorimeter. A method was applied within the associated software to automatically calculate the specific heat capacity based on calibrated heat flow measurements to both an empty aluminium crucible and a reference sample of sapphire (46.772 mg). A crucible containing solid hexamine (28.589 mg) was loaded into the

calorimeter alongside the two references. The temperature was held at the minimum operating value of 30 ºC for 10 min, then heated at 5 ºC/min up to 80 ºC, where the sample was then held at constant temperature for 10 min. Values for the specific heat capacity were averaged over three repetitions of this method.

The measured profiles of both the hexamine-ethanol solution density and solid hexamine specific heat capacity with temperature are displayed in Figure 3.9.

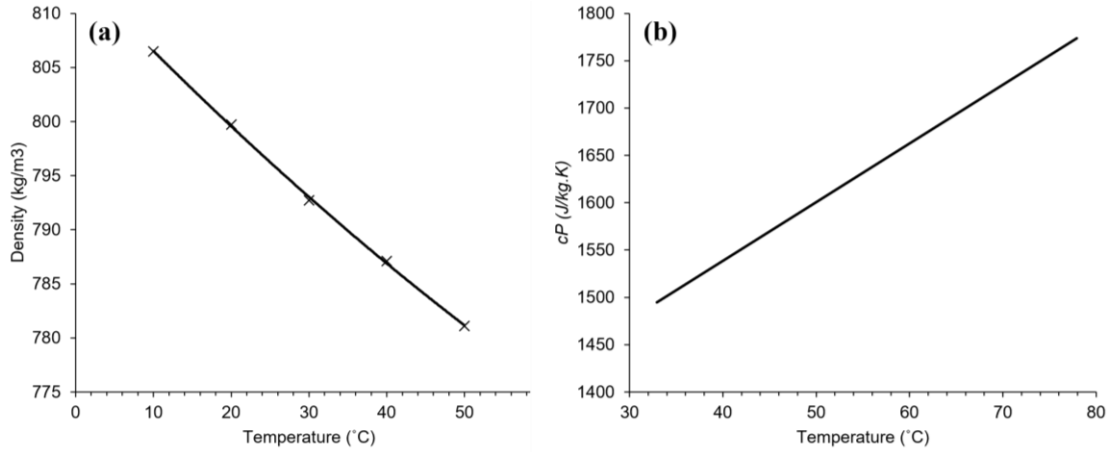


Figure 3.9. Measured profiles with temperature for a) hexamine-ethanol saturated solution density and b) solid hexamine specific heat capacity.

3.7. Conclusions

The raw materials and experimental equipment used in this study have been detailed. The solubility, nucleation rate and growth rate expressions for hexamine in pure ethanol solvent have been reported, which will be implemented into the process model in chapter 5. The measured solution density and specific heat capacity will also be implemented into the model. Methodologies for the calibration of hexamine-ethanol solution IR spectra to concentration, hexamine seed crystal preparation and physical property measurement have been outlined. Further experimental methods using the described crystallisation equipment will be outlined in the relevant chapters.

4. Metastable Zone Width Determination and Seeded Batch Crystallisation of the Hexamine-Ethanol Crystallisation System

4.1. Introduction

Prior to the development of a process model for the batch cooling crystallisation of hexamine from ethanol, a thorough understanding of the crystallisation behaviour through physical characterisation was required. Measuring the MSZW was important for the design of any seeded batch crystallisations, as the size of this region would inform of the ideal seeding temperatures and initial supersaturations. Following from this, seeded crystallisation experiments were designed and performed to cover a wide range of operating conditions, such that the process model predictions could be validated reasonably. In addition, the concentration time profiles measured during these experiments were used for the re-estimation of the crystallisation kinetic parameters, potentially improving the model for the crystal growth rate of hexamine. Analysis of the product crystals from batch crystallisations with different seeding conditions also provided further insight into the occurrence of other key mechanisms, such as secondary nucleation and agglomeration.

4.2. Materials and methodology

4.2.1. Methodology for MSZW determination

The MSZW was first measured in the 500 ml jacketed crystalliser for characterisation. Following the polythermal method, the vessel temperature was cycled three times between 5-50 ºC at a heating rate of 0.15 ºC/min and cooling rate of -0.3 ºC/min. Dissolution and crystallisation points were recorded where the turbidity probe measured 100% and < 90% transmittance respectively. The solution concentration was varied by dissolving enough hexamine into 400 g of ethanol to reach the solubility limit at 15, 25, 30 and 40 ˚C, calculated using the solubility expression in Table 3.1.

MSZW data was also obtained using the Crystal16 system; the same temperature ramp procedure as previously described was applied to vials containing around 1.5 ml (1.18 g) of ethanol, with appropriate amounts of hexamine to reach similar saturation points as in the 500 ml vessel. Tuning of the turbidity measurements (automatic adjustment of the brightness) was performed after the first heating step in all cases.

4.2.2. Methodology for seeded batch crystallisation

Six seeded crystallisation experiments were performed in total, the conditions of which are summarised in Table 4.1. These were selected such that comparisons could be made between specific conditions through different pairs of experiments, including cooling rate (Runs A and D), seed type/size (Runs A and B), seeding temperature (Runs B and E), seed mass (Runs B and C) and seeding method (Runs C and F). The solution concentration was measured in-situ using the ATR-FTIR probe. Additionally, the particle count was monitored during Runs A, B and C using the FBRM probe.

Table 4.1. Process conditions applied to each seeded linear cooling crystallisation run.

	Run A	Run B	Run C	Run D	Run E	Run F
Seeding Temperature (˚C)	40	40	40	40	38	40
Cooling Rate (˚C/min)	0.2	0.2	0.2	0.3	0.2	0.2
Seed Type	Sieved	Milled	Milled	Sieved	Milled	Milled
Seed Mean Size (µm)	178.4	69.4	69.4	178.4	69.4	69.4
Seed Mass (g)	1	1	3	1	1	3
Seeding Method	1	1	1	1	1	2

1: *Held at seeding temp.*
2: *Seeded during cooling*

At the beginning of each run, hexamine was added to 400 g of ethanol to make up an initial concentration of 0.0507 g/g solvent (saturation temperature of around 42 ˚C). The solution was heated to 50 ˚C and held for 60 mins to ensure all material had dissolved. Two seeding methods were applied; most runs followed method 1, where the solution was rapidly cooled to the selected seeding temperature and dry solid seed crystals were added. The solution was then held at constant temperature for 30 mins to allow the system to equilibrate, before cooling at the selected linear cooling rate down to 7 ˚C. Run F alternatively followed seeding method 2, where the solution was cooled linearly from 50 ˚C and solids were added at the desired seeding temperature without a holding period. The expected yield from each batch was around 57 %.

4.3. Results and discussion

4.3.1. MSZW of the hexamine-ethanol crystallisation system

At both the 1.5 ml and 500 ml scales, strong fluctuations were experienced in the transmittance (the turbidimetric probe output signal) during the initial heating step, improving significantly during dissolution of the recrystallised material in the second temperature cycle. This was a result of the high hygroscopicity of the raw hexamine feed, leading to the presence of clumps of material in the initial cycle. These caused disruptive readings when passing the turbidity

probe, as opposed to the much finer recrystallised material during the second and third cycles due to the abundance of nuclei formed. Hence, the dissolution temperatures reported were the averages calculated from only the second and third heating cycles, whereas the crystallisation temperatures were averaged between all three cooling cycles.

The MSZW curves for each scale of crystalliser are plotted in Figure 4.1. Using the Crystal16 resulted in a MSZW between 8-15 ºC, whereas in the 500 ml vessel a narrower 5-10 ºC range was observed. The relative supersaturations reached were approximately 0.30-0.61 in the Crystal16 and 0.10-0.25 in the 500 ml vessel. Both appeared to approach the solubility curve at higher temperatures.

The Crystal16 was utilised to increase the speed at which the MSZW could be measured, given the number of experiments that could be run in parallel. However, due to the differences in the MSZW across different scales, it was decided that the MSZW measured in the 500 ml vessel would be much more accurate and useful for designing future experiments at the same scale. It was noted that the narrow MSZW at -0.3 ºC/min cooling rate may make seeding within these conditions difficult if primary nucleation was to be avoided, however the results of the seeded linear cooling batches later disproved this concern.

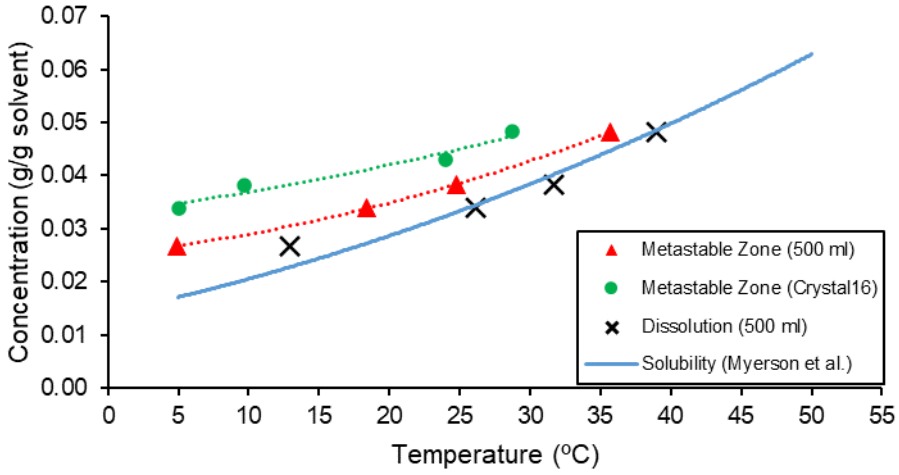


Figure 4.1. Metastable zone for hexamine-ethanol system obtained from Crystal16 and 500ml vessel using cooling rate of -0.3 ºC/min, solubility from Myerson et al. (1986).

Dissolution point measurements were generally in agreement with the solubility data from literature (see Table 3.1). The data point at 15 ºC saturation (0.0267 g/g solvent) showed the largest discrepancy, being measured at 12.9 ºC. Considering the operating range indicated by the MSZW measurements, seeding temperatures around 2-4 ˚C below saturation temperature and cooling rates close to -0.3 ˚C/min were chosen to ensure primary nucleation would be avoided during crystallisation experiments.

4.3.2. Seeded batch crystallisation analysis

The temperature profiles of the crystallisation vessel and circulator during each of the seeded linear cooling crystallisations are shown in Figure 4.2. The concentration profiles measured by the ATR-FTIR probe are presented in Figure 4.3, alongside the relative supersaturation that was calculated using the saturation concentration. To improve the stability of the concentration measurements, the profiles shown in Figure 4.3 have been averaged over 5 data points (collected over 5 mins). For the three batches that included the FBRM probe (Runs A, B and C), the square-weighted particle counts of fine (< 100 µm) and coarse (> 300 µm) material are shown in Figure 4.4. The number-based and volume-based CSDs of the seed crystals and product crystals from each batch are shown alongside microscope images of each sample in Figures 4.5-4.12. Further numerical details of the volume-based product CSDs have been summarised in Table 4.2 alongside the mean supersaturation values, for comparison between batches.

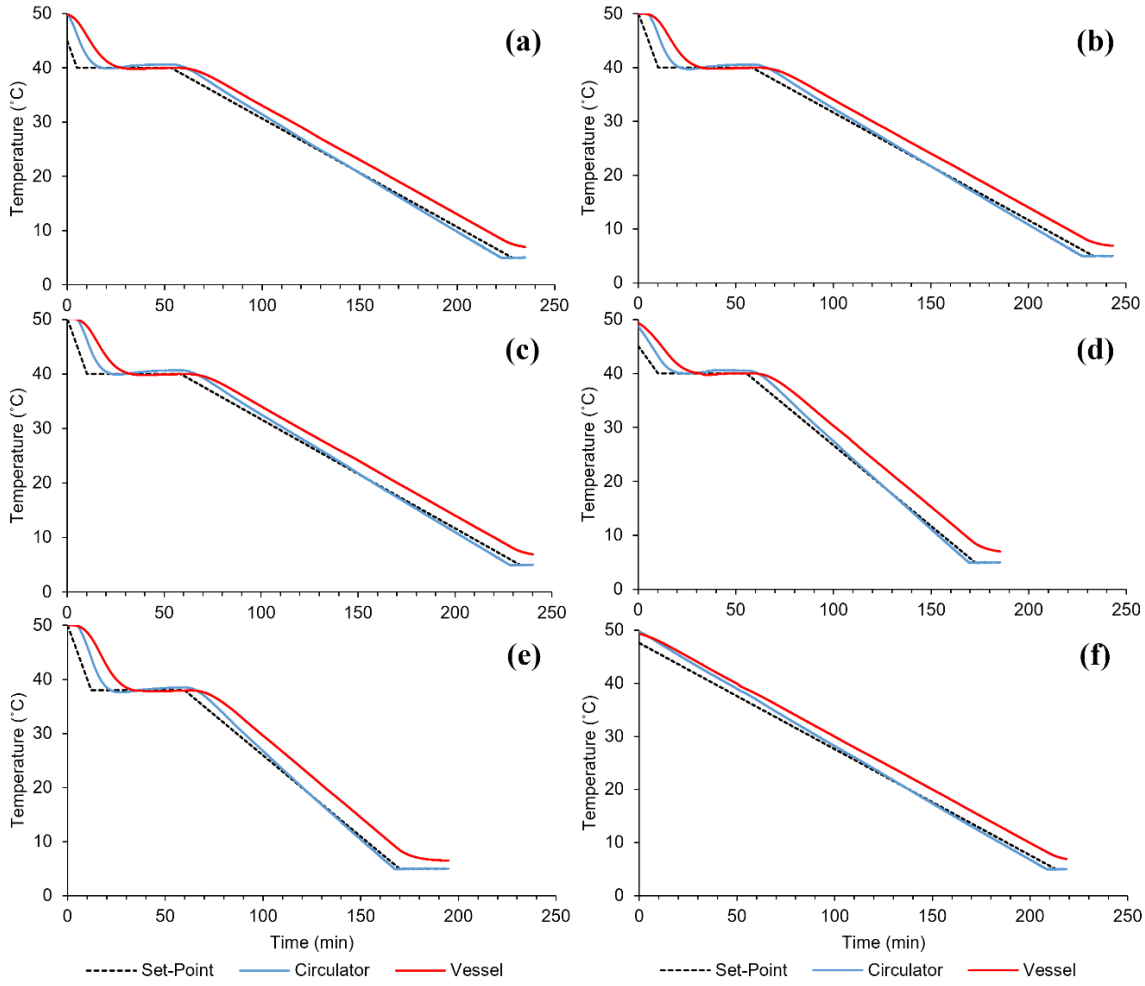


Figure 4.2. Measured temperature profiles during the seeded linear cooling crystallisations. a) Run A, b) Run B, c) Run C, d) Run D, e) Run E and f) Run F

Although there was a time lag between the vessel set-point and measured process temperatures during each batch, the correct cooling rates were still achieved during the main linear cooling segments. Each batch was held at the final temperature for slightly different times to allow supersaturation to fall to zero before removal of the product crystals. This remained consistently at 7 ˚C, meaning that the amount of hexamine that left solution in each batch remained the same.

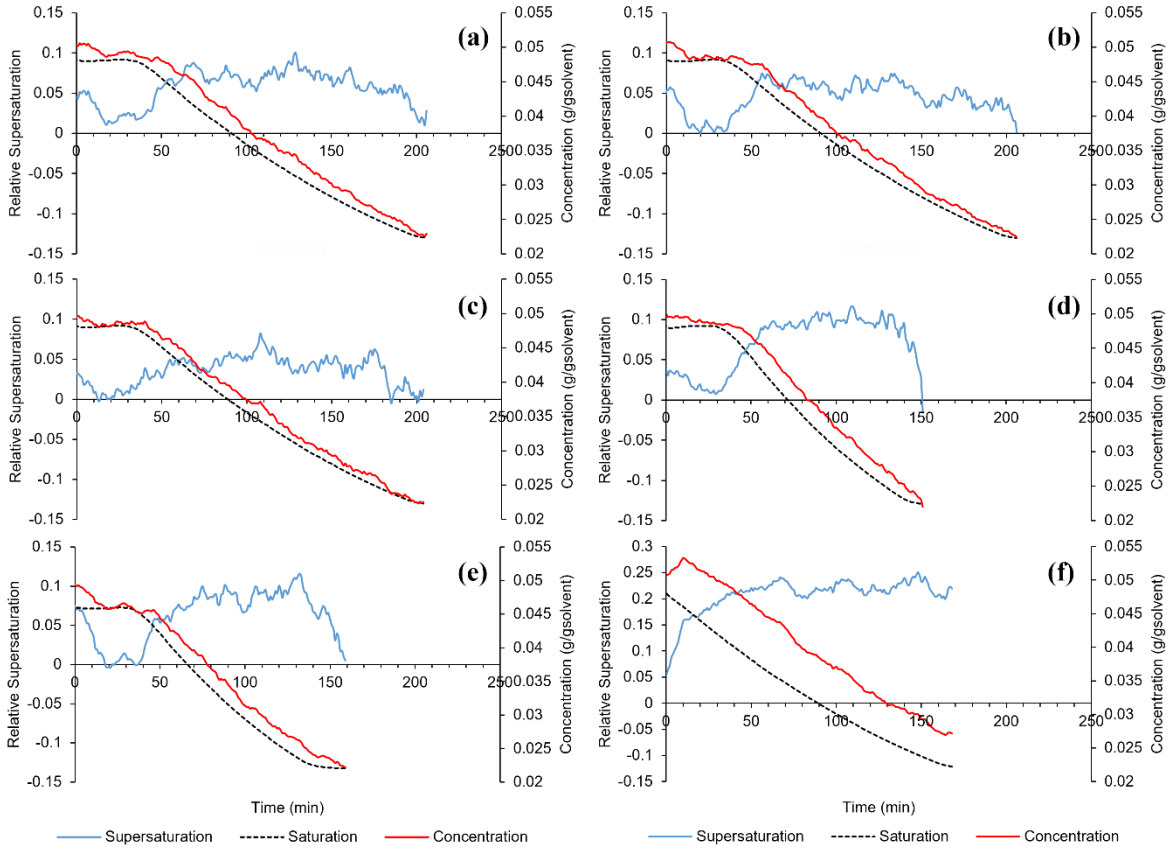


Figure 4.3. Measured concentration and supersaturation profiles during the seeded linear cooling crystallisations, starting from the point of seeding. a) Run A, b) Run B, c) Run C, d) Run D, e) Run E and f) Run F.

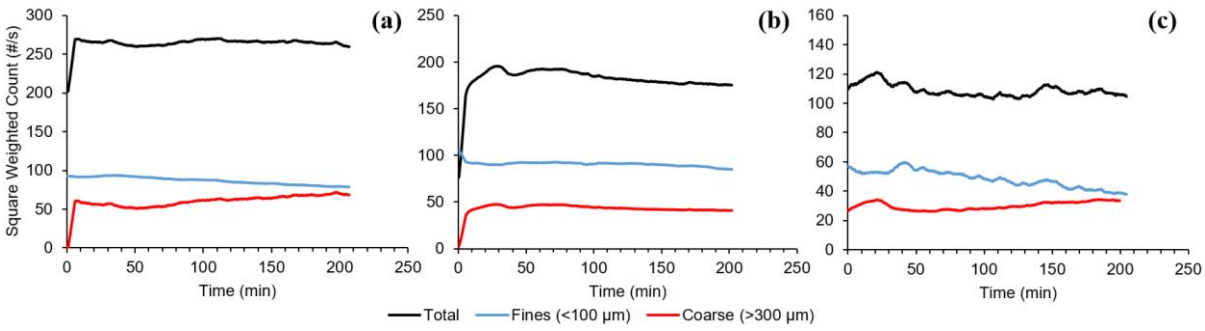


Figure 4.4. Measured square-weighted particle counts during the seeded linear cooling crystallisations, starting from the point of seeding. a) Run A, b) Run B and c) Run C.

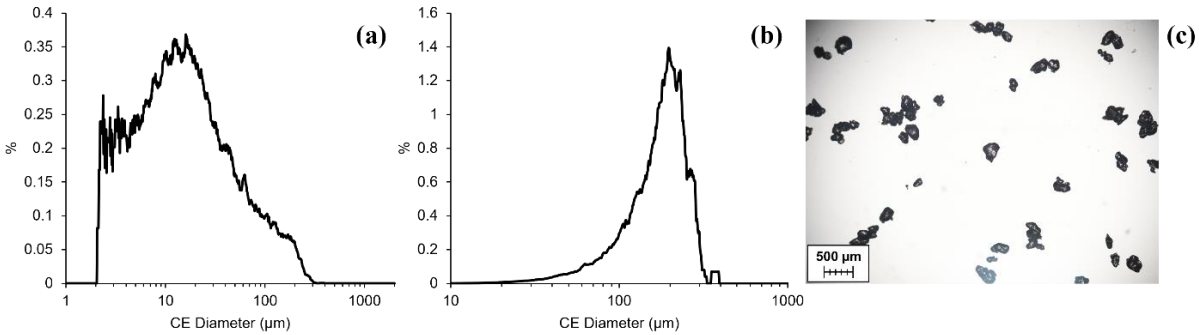


Figure 4.5. a) Number-based CSD, b) volume-based CSD and c) image of the sieved hexamine seed crystals.

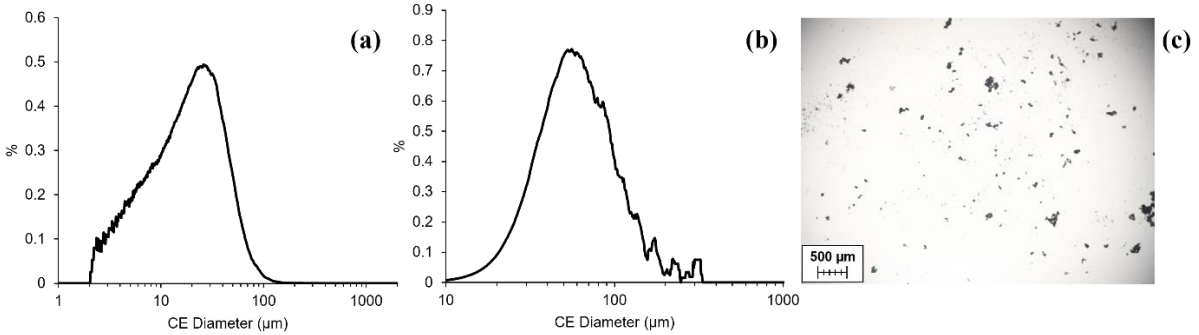


Figure 4.6. a) Number-based CSD, b) volume-based CSD and c) image of the dry milled hexamine seed crystals.

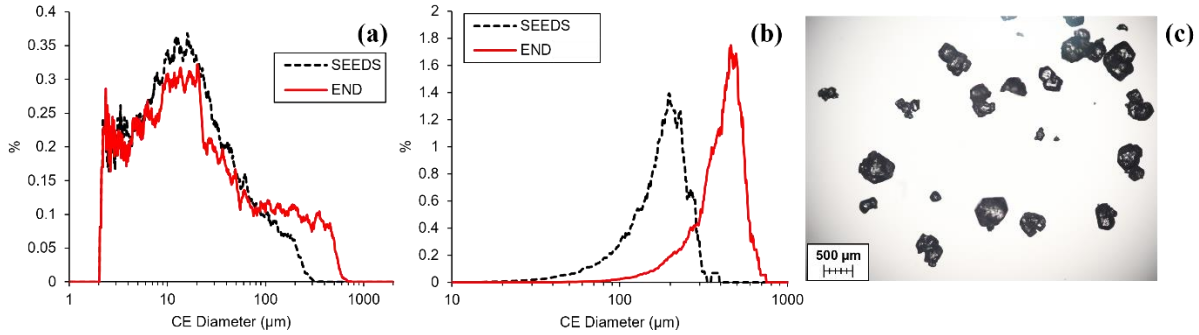


Figure 4.7. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run A.

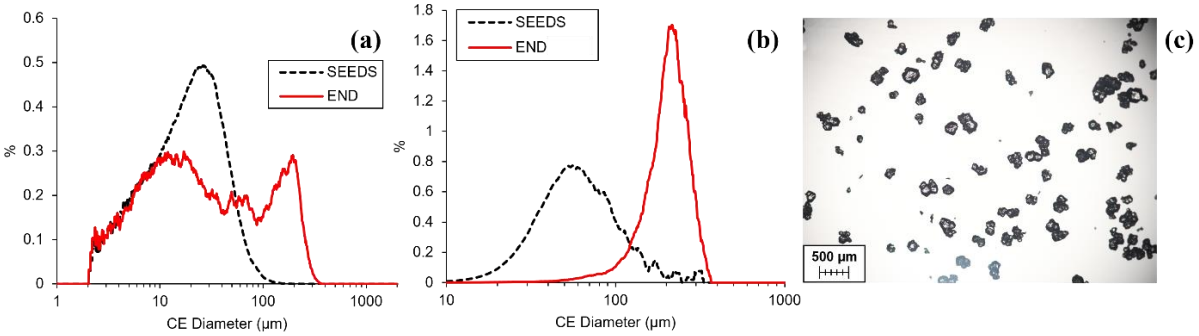


Figure 4.8. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run B.

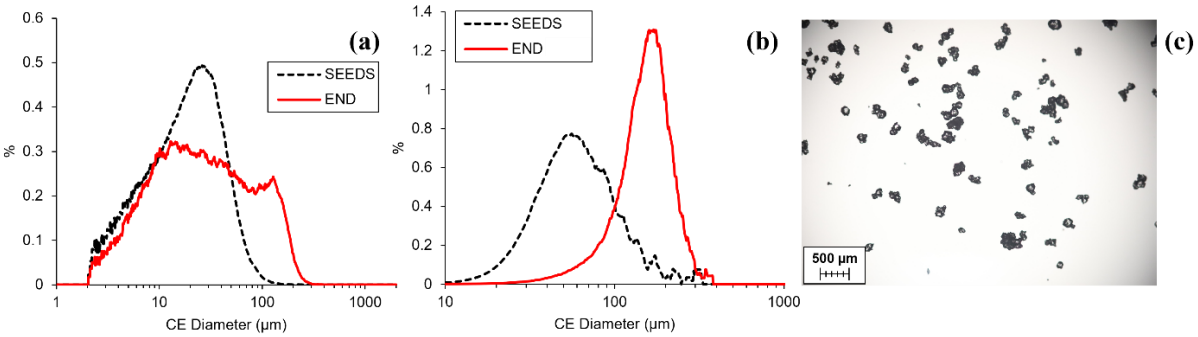


Figure 4.9. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run C.

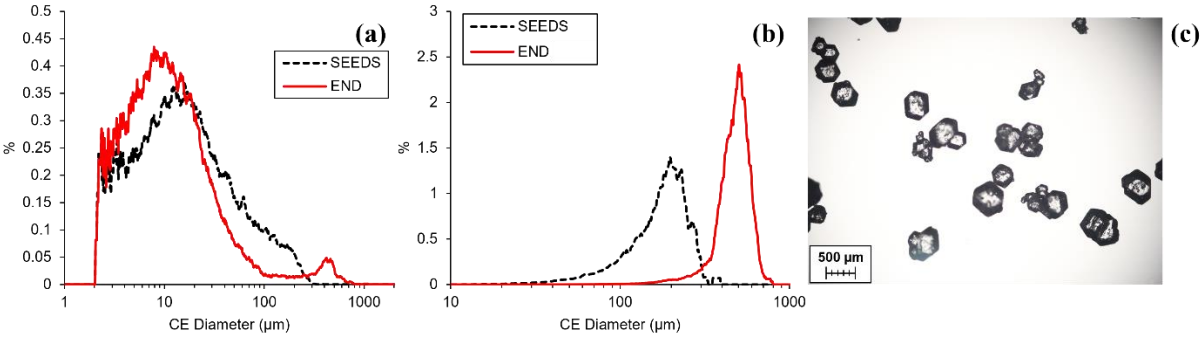


Figure 4.10. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run D.

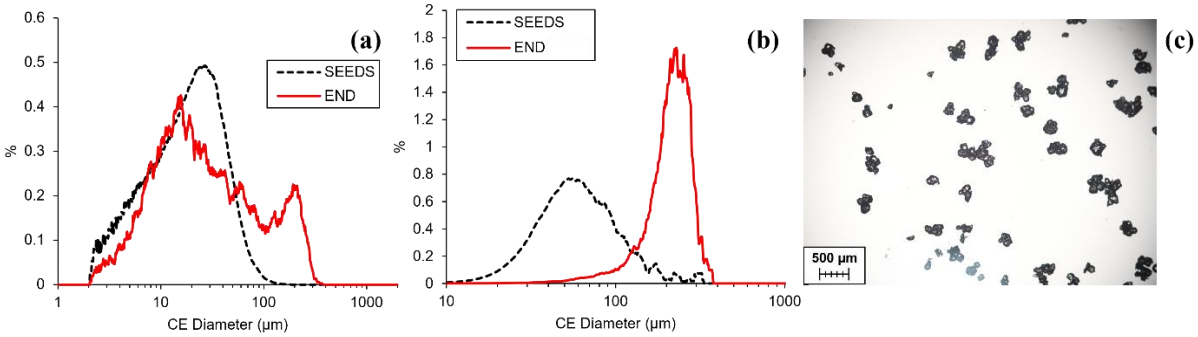


Figure 4.11. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run E.

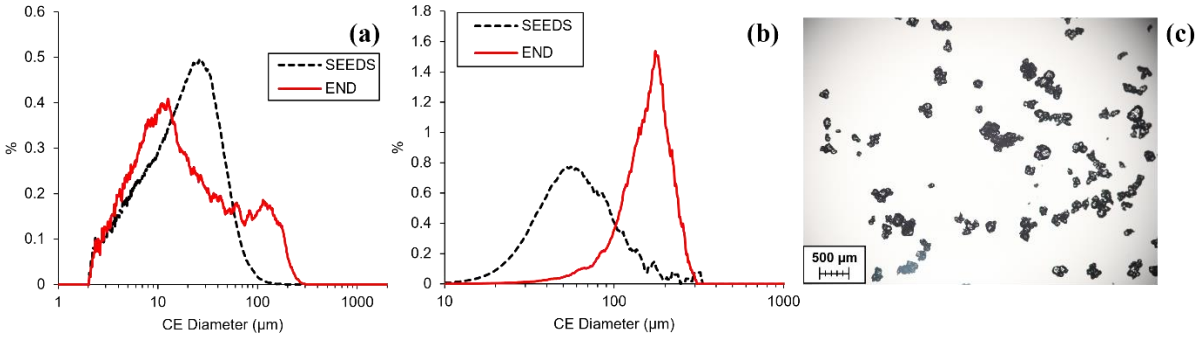


Figure 4.12. a) Number-based CSD, b) volume-based CSD and c) image of the hexamine product crystals from seeded linear cooling crystallisation Run F.

Table 4.2. Mean sizes and spans of volume-based product CSDs and supersaturations from each seeded linear cooling crystallisation run.

	Run A	Run B	Run C	Run D	Run E	Run F
Volume Mean Size (μm)	407.7	207.4	157.8	478.5	214.4	160.3
Span of Volume-Based CSD	0.74	0.69	0.86	0.48	0.68	0.80
Mean Supersaturation	0.056	0.044	0.035	0.070	0.062	0.207

In most runs there was no noticeable drop in concentration during the initial holding period after adding the seed crystals. This event was more prevalent in Run E, where a lower seeding temperature was chosen. The resulting mean crystal size in Run E was higher than that from Run B, for which all other conditions were identical. This suggested a predominance of crystal growth over secondary nucleation during the holding period for most experiments. As the growth rate may have been too slow to offset the amount of nucleation occurring at the beginning of Run B, the greater crystal population developed into a product with a smaller final mean size.

The saturation point was reached most rapidly in Run C, where a larger mass of seed crystals was added than in Run B. Given the larger surface area available for crystal growth, the supersaturation remained low throughout the entire batch. At a higher seed loading, there was a reduced final mean crystal size, resulting from the same mass of solute being distributed around a larger number of particles. This is an important consideration for supersaturation control, as a larger seed mass would likely make it more difficult to maintain the supersaturation at a constant level.

Another key factor that would affect process control is the cooling rate. Comparing Runs A and D, which only differed in cooling rate, it was observed that a lower maximum supersaturation was generated when cooling more slowly. This in turn implies that faster cooling rates will be required to maintain higher levels of constant supersaturation, which may be limited by the cooling power achievable through control of the jacket temperature. Given that only two cooling rates were tested, there was insufficient evidence to suggest a maximum feasible value for the supersaturation set-point.

From the particle count measured during Runs A, B and C, it was found that each batch followed the same general trends. The total particle counts remained fairly constant over time, whereas there was a more obvious increase in the coarse counts and decrease in the fine material. This suggested that the crystallisation under these conditions was dominated by crystal growth, without any significant secondary nucleation occurring.

In Run F, a substantially larger supersaturation was generated in the absence of a holding period after seed addition (compared to Run C with equal cooling rate and seed mass). The decrease in measured concentration was accompanied by a similar rate of decrease in saturation over most of the duration. Even in this case, the minimal difference between the final CSD between Runs C and F suggested no significant increase in nucleation, despite the mean supersaturation of 0.207 falling within the upper bounds of the range predicted to lead to primary nucleation based on the narrow MSZW.

Comparison between the images of the seed and product crystals in Figures 4.5-4.12 further supported the 'growth only' hypothesis. There was no obvious increase in the number of fine particles at the end of batches that used either the larger sieved or smaller milled seed crystals.

A more significant difference observed was the increase in agglomeration when starting with the smaller seeds. This meant that agglomeration could only be neglected in the PBM for experiments involving larger seed crystals for the crystallisation behaviour accurately predicted. Hence, larger seed sizes were favourably used during the process control experiments described in later chapters, to ensure the MPC strategy performed strongly during the validation phase.

4.4. Conclusions

The MSZW of the hexamine crystallisation at a cooling rate of -0.3 ˚C/min has been reported for the 1.5 ml and 500 ml scale. The narrow operating range of 5-10 ˚C in the 500 ml crystalliser was considered when selecting the operating conditions for seeded batch cooling crystallisations. The temperatures, concentrations and particle counts during seeded batch crystallisation experiments with a wide range of conditions were analysed alongside the measured CSDs and images of product crystals. Through determination of the effects of cooling rate, seed size/mass and seeding temperature/method on the key crystallisation mechanisms of nucleation and growth, a good understanding of the hexamine-ethanol system has been established. The seeded crystallisation experiments performed will be used for the validation of the process model parameterised in the following chapter, and the identification of experiments resulting in 'growth-only' conditions will be implemented to allow isolation and re-estimation of the growth rate kinetics.

5. Crystallisation Process Model Parameterisation with Re-estimation of Growth Rate Kinetics and Experimental Validation

5.1. Introduction

The aims of this chapter are: (1) to provide an outline of the procedure that was followed to build an appropriate flowsheet for batch cooling crystallisation in gPROMS FormulatedProducts software, (2) to summarise the necessary input data for the chosen crystallisation system and (3) to simulate a series of seeded batch crystallisations for validation of the model. The formulation of such a digital twin representing the hexamine-ethanol crystallisation was simplified here through the parameterisation of a 1D PBM in the gPROMS software platform. The accuracy of the model's predictions of concentration and crystal size over time were validated against the measurements recorded for physical seeded crystallisations from chapter 4. The impact of approximating the initial seed CSD as a single lognormal peak instead of the directly measured CSD data was assessed, as this would potentially allow a reduction in computational complexity of the MPC strategy.

The potential inaccuracy of the literature growth rate expression describing the hexamine-ethanol crystallisation system was expected to result in some inaccuracy in the model's predictions. To improve this, the parameter estimation feature in gPROMS was used to re-estimate selected crystallisation kinetic parameters. It was demonstrated in chapter 4 that the majority of seeded batch crystallisations resulted in negligible secondary nucleation; therefore, only modifications to the growth rate expression have been considered. An assessment of the re-estimated parameters was made before deciding on the final state of the model to be carried forward for controller training.

5.2. Materials and methodology

5.2.1. Setup of the batch crystallisation flowsheet in gPROMS FormulatedProducts

The mixed-suspension, mixed-product removal (MSMPR) crystalliser module in gPROMS FormulatedProducts v2.2.0 was used to create the digital twin of the hexamine-ethanol crystallisation system. Figure 5.1 shows the flowsheet for a batch cooling crystallisation process, built within the software to provide the structure for resolving a 1D PBM of the hexamine-ethanol system in a simulated 500 ml jacketed crystalliser. The software applies discretisation via the method of classes outlined in section 2.2.2 to solve the PBM and calculate the predicted volume CSD. Size distribution and composition (solute concentration) sensors were used to closely monitor the CSD and supersaturation respectively during the simulation.

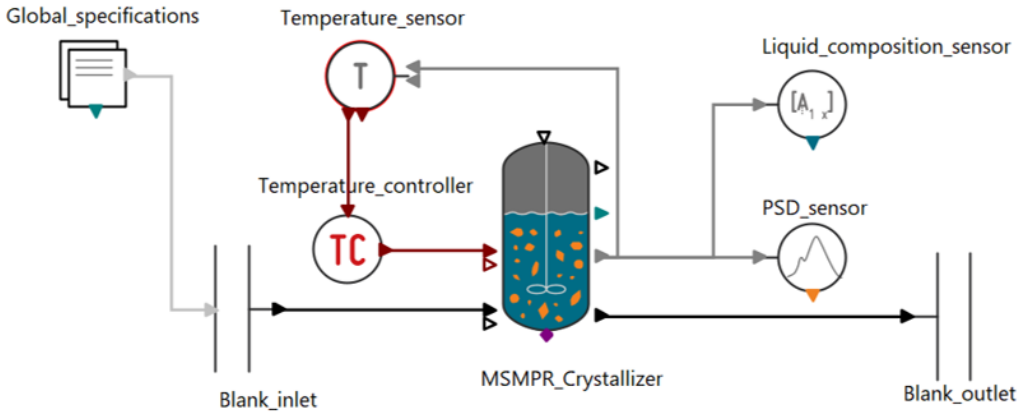


Figure 5.1. Flowsheet of a batch cooling crystallisation process in gPROMS FormulatedProducts software.

Table 5.1. Input data for the 'Global Specifications' component of the flowsheet for hexamine-ethanol batch cooling crystallisation in gPROMS FormulatedProducts.

Section/Property	Input	Source
System information		
Components	Hexamine, Ethanol	-
Liquid properties		
Molecular weights (kg/mol)	Hexamine: 0.140186 Ethanol: 0.04607	- -
Mass density coefficients a, b, c	1172.3, -1.8681, 0.002	Measured
Mass specific heat coefficients Hexamine: a, b Ethanol: a, b, c, d	 -402.78, 6.1976 6271.4, 93.036, -0.3528, 0.0005	 Measured (Miyazawa et al., 2012)
Dynamic viscosity coefficients a, b, c, d	 0.0733, -0.0006, $2x10^{-6}$, $-2x10^{-9}$	 (Kaye and Laby, 1995)
Crystal properties		
Volumetric shape factor	Equivalent to a sphere	-
Mass density (kg/m^3)	1330	(Rohani and Bourne, 1990)
Solubility		
Solubility calculated using	Generic polynomial	-
Solubility coefficients a, b, c	 $5.354x10^{-1}$, $-4.13x10^{-3}$, $8.2x10^{-6}$	 (Myerson et al., 1986)
Grid parameters		
Number of grid points	200	-
Grid type	Logarithmic	-
Minimum particle size (μm)	0.1	-
Maximum particle size (μm)	1500	-

Table 5.1 summarises the most relevant input data for the 'Global Specifications' component, which contained information about the selected crystallisation system including the components' molecular weights, physical properties, crystal shape factor and the solubility curve. It was assumed that the presence of hexamine would have a minimal impact on solution viscosity due to the low solute loading, therefore the viscosity of pure ethanol was applied to the model (Kaye and Laby, 1995). A logarithmic grid was chosen to ensure a fine resolution around the smaller particle sizes, which is a critical region of the CSD when considering nucleation.

Table 5.2 summarises the input data for the 'MSMPR Crystalliser' component, which contained specifications of the geometry and operating conditions for the simulated 500 ml crystalliser, as well as the secondary nucleation and growth rate kinetic parameters. The kinetic parameters listed in Table 3.1 were incorporated into the model, using the nucleation rate expression as attrition-based secondary nucleation and selecting values for the rate constant/supersaturation order in the middle of the reported variance. The initial seed crystal mass and solution composition were also specified here; the values used were the same as in Runs A-E of the seeded crystallisation experiments described in chapter 4. The initial seed CSDs were input through two methods: the measured volume CSD from the Morphologi G3 (see Figures 4.5b and 4.6b), or a lognormal peak that approximated the measured CSD. The lognormal peaks were calculated using the location parameters (D_{50}) and standard deviations reported in Table 5.2.

To simulate these experiments using the gPROMS process model, the measured temperature profiles were reconstructed in the 'Temperature Controller' component of the flowsheet. The piecewise linear profiles that were constructed to match each of the measured temperature profiles are summarised in Table 5.3.

Simulations of each of the five seeded batches were run with a reporting time interval of 5 seconds. Runs D and E were selected as cases for model validation, as the differences between their operating conditions (seed size, seeding temperature and cooling rate) provided the widest and most detailed overview of the model's accuracy. Simulations of Runs D and E were also performed with secondary nucleation switched off, which provided evidence to support the decision not to re-estimate the kinetic parameters for this mechanism in the following section.

Table 5.2. Input data for the 'MSMPR Crystalliser' component of the flowsheet for hexamine-ethanol batch cooling crystallisation in gPROMS FormulatedProducts. Nucleation and growth kinetic parameters are provided from Myerson et al. (1986).

Section/Property	Input				
Equipment and operation					
Equipment volume (L)	0.5				
Liquid and solid outflow	Batch operation				
Specific power output	Calculated				
Impeller diameter (m)	0.04				
Impeller frequency (rpm)	450				
Impeller power number	1.3				
Impeller pumping number	0.8				
Energy balance					
Heat transfer	Not considered (Adiabatic)				
Secondary nucleation (attrition)					
Secondary nucleation (attrition)	Power law kinetics				
Driving force	Absolute supersaturation				
Rate constant (ln(#/m^3.s))	31.543				
Supersaturation order	2.6				
Activation energy (J/mol)	0				
Specific power input order	0				
Minimum particle size (μm)	0				
Growth and dissolution					
Growth and dissolution	Power law kinetics				
Driving force	Absolute supersaturation				
Size dependent solubility	Inactive				
Growth rate constant (m/s)	2.6×10^{-2}				
Supersaturation order	1.95				
Activation energy (J/mol)	0				
Initial conditions	**Run A**	**Run B**	**Run C**	**Run D**	**Run E**
Mass (kg)					
Solid crystal phase	0.001	0.001	0.003	0.001	0.001
Liquid phase	0.42028	0.42028	0.42028	0.42028	0.42028
Temperature (˚C)	40	40	40	40	38
Initial conditions: Liquid					
Mass fraction					
Ethanol	0.95175	0.95175	0.95175	0.95175	0.95175
Hexamine	0.04825	0.04825	0.04825	0.04825	0.04825
Initial conditions: Solid					
Specify CSD	As a density distribution function				
Location parameter (μm)	199	58	58	199	58
Standard deviation (μm)	40	30	30	40	30

Table 5.3. Input data for the 'Temperature Controller' component of the flowsheet for hexamine-ethanol batch cooling crystallisation in gPROMS FormulatedProducts.

	Interval 1	Interval 2	Interval 3
Run A			
Duration (min)	32	165	2
Final Temperature (˚C)	40	7	7
Run B			
Duration (min)	32	165	2
Final Temperature (˚C)	40	7.1	7
Run C			
Duration (min)	32	165	2
Final Temperature (˚C)	40	7.1	7
Run D			
Duration (min)	33	107.3	8
Final Temperature (˚C)	40	7.8	7
Run E			
Duration (min)	35	102.7	7
Final Temperature (˚C)	38	7.2	7

5.2.2. Methodology for crystallisation kinetic parameter re-estimation

After evaluating the performance of the literature model kinetics, the parameter estimation feature in gPROMS software was used to investigate whether a better fit to experimental data could be achieved. In addition to re-estimating the parameters for the power-law growth expression, implementation of the two-step growth expression was also tested. In total, three different methods of parameter estimation were employed:

- Re-estimation of the rate constant k_g for the power law growth expression, keeping the supersaturation order fixed
- Re-estimation of the supersaturation order g for the power law growth expression, keeping the rate constant fixed
- Estimation of the diffusivity correction factor α for the two-step growth expression (Equation 2.36), keeping the rate constant and supersaturation order fixed

The parameter estimation method built into gPROMS follows the optimisation of a maximum likelihood function, solved using a non-linear sequential quadratic programming (SQP) algorithm (Gill et al., 2001). The objective function can be described as follows (Siemens-PSE Ltd., 2021):

$$\Phi = \frac{N}{2}\ln(2\pi) + \frac{1}{2}\min_{\theta}\left\{\sum_{i=1}^{NE}\sum_{j=1}^{NV_i}\sum_{k=1}^{NM_{ij}}\left[\ln\left(\sigma_{ijk}^2\right) + \frac{\left(\tilde{z}_{ijk} - z_{ijk}\right)^2}{\sigma_{ijk}^2}\right]\right\} \quad (5.1)$$

This overall function consists of a constant term, variance term and residual term. For a fixed number of experiments and variance model, the constant and variance terms remain fixed during the parameter estimation. Therefore, only the contribution from the residual term is minimised for each of the measured variables included in the optimisation problem:

$$\Phi = \min_{\theta}\left\{\sum_{i=1}^{NE}\sum_{j=1}^{NV_i}\sum_{k=1}^{NM_{ij}}\left[\frac{\left(\tilde{z}_{ijk} - z_{ijk}\right)^2}{\sigma_{ijk}^2}\right]\right\} \quad (5.2)$$

Each parameter estimation strategy was applied using the experimental data from seeded Runs A, B and C. After inputting the temperature profiles for each performed experiment, the software aimed to modify the selected parameters until a good fit was made between the measured concentration profiles and the predicted profiles constructed by the process model. The variance model for concentration in each case was a constant variance of 0.0005 g/g solution, which was chosen based on the precision error of the ATR-FTIR calibration model. As only the initial and final CSDs had been measured, there was insufficient data to also include a size-related property as an additional optimisation objective. Once acceptable parameters had been determined by the estimation algorithm, the concentration profiles and final CSDs from the unused Runs D and E were used as validation data.

5.3. Results and discussion

5.3.1. Process model validation using seeded batch crystallisations

Figure 5.2 shows a comparison between the predicted and measured volume-based CSDs, concentration profiles and supersaturation profiles for seeded Runs D and E. When the seed CSD was defined using the measured distribution, the gPROMS process model underpredicted the final crystal size in both cases. Improvements were observed when specifying the seeds as a lognormal peak, which had a more significant effect on Run D due to the elimination of the fine material tail in the distribution. The predicted volume mean crystal size increased from 216 µm to 393 µm for Run D, and 99 µm to 113 µm for Run E.

An improved agreement with the measured concentration profile during the batch was also achieved using the lognormal CSD, most significantly for Run D. Once again considering the elimination of the fines, this would reduce the available surface area of the seed crystals and therefore slow the growth rate, hence resulting in higher supersaturations throughout the batch.

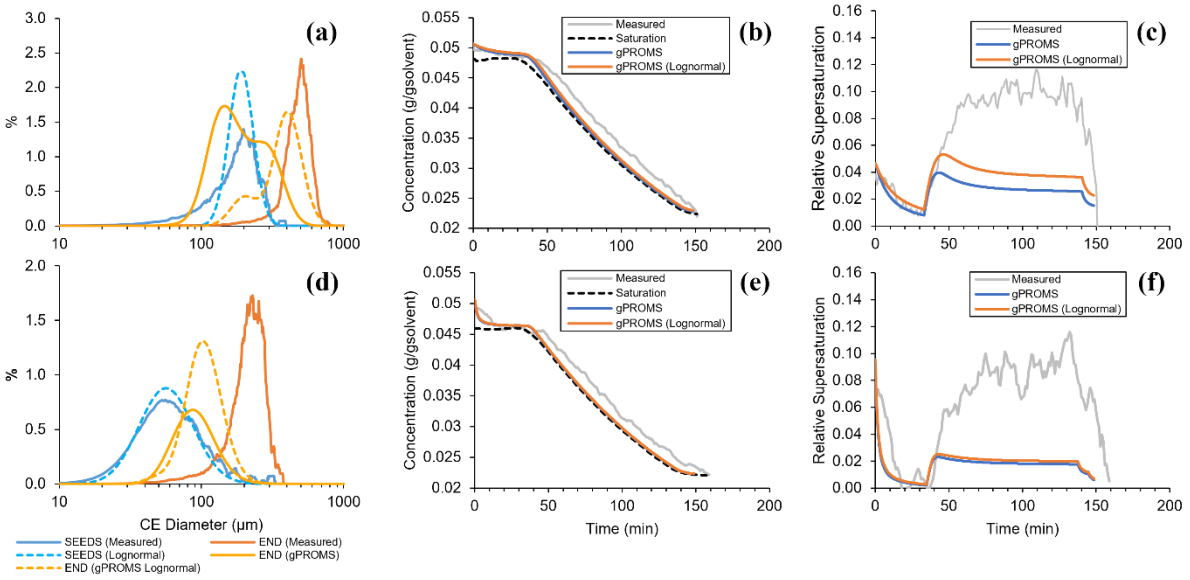


Figure 5.2. Comparison of CSDs, concentration and supersaturation profiles predicted by gPROMS with measured data. Simulations starting with either the measured seed CSD or a representative lognormal peak. Results correspond to seeded linear cooling a,b,c) Run D and d,e,f) Run E.

Predictions of the rate at which supersaturation declined during the isothermal holding period were more accurate for Run D than Run E; the measured supersaturation in the latter case was consumed more slowly than the model prediction. This could have been due to the use of milled seed crystals with a more uneven shape than the assumed spherical particles in the gPROMS model. The growth rate in reality may have been slowed due to poor attachment of solute molecules to the uneven surfaces, leading to a slower rate of decrease in supersaturation than expected.

Given that the absence of an agglomeration mechanism in the PBM could be the larger cause of poor crystal size predictions, this alone was not enough to suggest the crystallisation kinetic parameters of Myerson et al. (1986) applied in the gPROMS model were not entirely accurate. However, both cases showed a significant underprediction of the level of supersaturation experienced during the linear cooling section. A potential reason for this could be the inaccuracy of the secondary nucleation kinetics, which would result in the supersaturation being consumed faster than in reality if the predicted rates were too high.

After performing simulations of Runs D and E with secondary nucleation switched off, it was found that this mechanism actually made little contribution to the predicted CSDs in both cases. The final volume mean size predicted was 400 µm for Run D (previously 393 µm), and 119 µm for Run E (previously 113 µm). For both cases, this increase was still not enough to bring the predicted sizes high enough to match the measured values. Therefore, it was decided that

there would be benefit only in re-estimating the growth rate kinetic parameter prior to training predictive controllers using the digital twin.

5.3.2. Re-estimation of growth rate kinetic parameters

Figure 5.3 shows a comparison between the concentration predictions for Runs D and E obtained through re-estimation of various growth rate parameters. An evaluation of the percentage error for each parameter estimation method is shown in Figure 5.4 for both runs. The predicted volume mean sizes of the product crystals for each case are listed in Table 5.4.

For both of the selected validation cases, transitioning from a power law to two-step growth expression proved to weaken the model's prediction of both concentration (3-4 % increase in mean error) and crystal size.

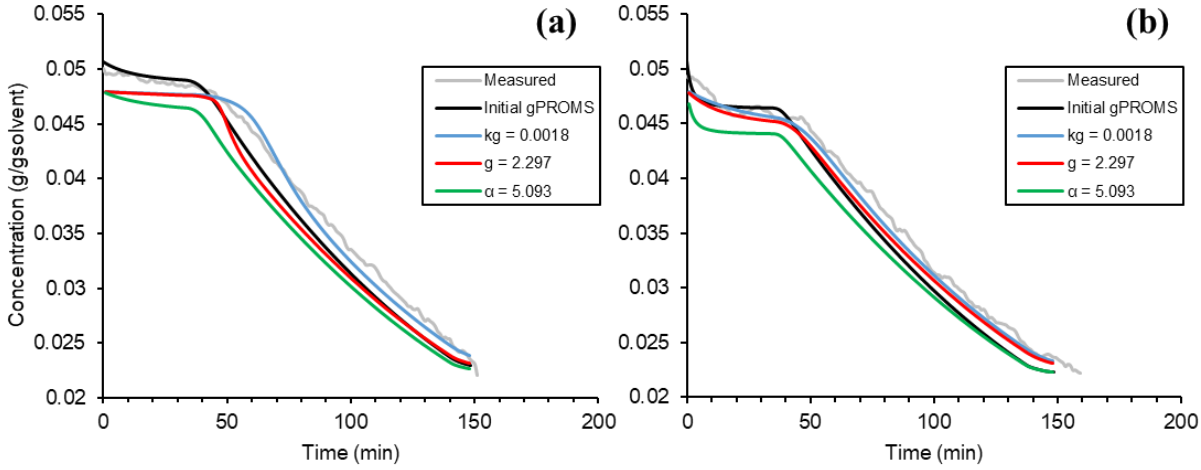


Figure 5.3. Predicted concentration fit to measured data from a) Run D and b) Run E for gPROMS simulations using re-estimated growth rate parameters.

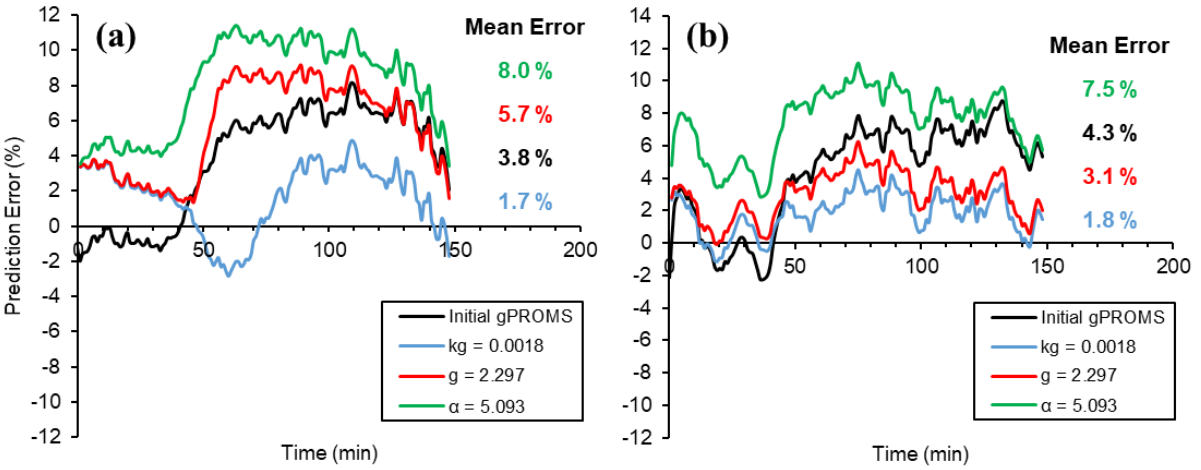


Figure 5.4. Percentage error in predicted vs measured concentration data from a) Run D and b) Run E for gPROMS simulations using re-estimated growth rate parameters.

Table 5.4. Predicted volume mean crystal size of product from gPROMS simulations using re-estimated growth rate parameters.

	Volume Mean Size (μm)	
	Run D	**Run E**
Measured	478.5	214.4
Initial gPROMS prediction (power-law growth rate, $k_g = 2.60$, $g = 1.95$)	393.0	113.0
Power-law growth rate, estimated $k_g = 0.0018$	181.4	121.7
Power-law growth rate, estimated $g = 2.297$	75.6	121.8
Two-step growth rate, estimated $\alpha = 5.093$	152.7	83.4

The re-estimated supersaturation order g led to a closer fit to the measured data in both concentration (1.2 % decrease in mean error) and size for Run E. However, this improvement was not reflected for Run D; there was a 1.9 % increase in concentration error and significant underprediction of the final mean crystal size, which was around five times smaller than the already underpredicted original value. Despite the re-estimated rate constant k_g leading to an improved concentration fit for both validation cases (1.3-2.1 % decrease in mean error), a similarly poor prediction of the crystal size was observed for Run D.

Given that the crystal size was only measured at the start and end of the batches, the parameter estimation could not fit to this parameter due to the lack of data points. Therefore, a heavy weight was placed on fitting the concentration to the measured profiles that were provided, neglecting the effect on the final crystal size reached. As these results suggested unreliability as to whether the new growth rate parameters would improve the model predictions depending on the operating conditions, the re-estimated parameters were rejected and the original model parameters were retained.

5.4. Conclusions

A digital twin for the batch cooling crystallisation of the hexamine-ethanol system was created in gPROMS FormulatedProducts software, using measured physical property data alongside solubility and nucleation/growth rate kinetic parameters from literature. Two of the seeded batch crystallisations presented in the previous chapter were simulated to compare the process model's predictions of concentration and crystal size to experimentally measured values.

It was found that the model parameterised using the nucleation and growth kinetics from literature underpredicted the maximum level of supersaturation reached during both batches.

Additionally, the mean crystal size of the product was also often underpredicted. This was likely due to the lack of an expression for agglomeration included in the process model, which would be difficult to introduce given the complexity of the nucleation-related agglomeration mechanism that hexamine exhibits (described in section 3.1.1).

The implementation of lognormal peaks for the initial seed crystal CSDs showed improvement in both concentration and size predictions, particularly in the case where the original measured seed CSD contained excessive fine material. Re-estimation of the growth rate kinetics using the concentration profiles measured during three other batch experiments was able to decrease the mean supersaturation error by 1.3-2.1 %. However, this came with the drawback of worsening the predictions of the final crystal size.

Training the MPC strategy using this digital twin may provide certain limitations based on the slight mismatch between predictions and experimental results. For example, given the underprediction of the supersaturation, attempting to control this property at levels beyond the maximum predicted by the model may have a negative impact on controller performance. This was taken into consideration when selecting the conditions for the assessment of the MPC strategy for both supersaturation and crystal size.

6. Growth Rate Inhibitor Screening and Effects on the Hexamine-Ethanol Crystallisation System

6.1. Introduction

If the performance of the MPC strategy was tested solely on a pure solute-solvent system, there would not be major relevance towards industrial crystallisation systems that typically face impurities in their raw material feeds. These are typically biproducts from synthesis that interact with the crystallising material, thereby altering the growth rate kinetics even when present at very low levels. Industrial crystallisation control must therefore be robust and ideally adaptable to the effects of various impurities.

In this chapter, three potential impurities have been screened via seeded crystallisation at the 100 ml scale, in order to select an additive that significantly disturbed the crystallisation kinetics of the hexamine-ethanol system. The selected impurity was further tested during seeded crystallisation experiments in the 500 ml jacketed crystalliser, as it was important to assess the effects on the change in solution concentration and crystal size at this scale, where the MPC strategy would later be validated. Impurities can also alter the thermodynamics and therefore equilibrium solubility of a crystallisation system. Therefore, the dissolution temperatures of varying hexamine concentrations were measured with the selected impurity present to confirm whether the solubility expression used in the gPROMS model would remain accurate.

The results of this study would influence the selection of impurity concentrations to be implemented during validation of the supersaturation MPC strategy, assessing its ability to adapt to the disturbances introduced by the impurities.

6.2. Materials and methodology

6.2.1. Methodology for impurity screening

The effects of the three potential growth inhibitors (citric acid, PABA, A1CA) were tested using the 100 ml AutoMATE system. The rationale for selecting these impurities is outlined in section 3.1.2. In each crystalliser, hexamine was added to 80 g of ethanol to make up an initial concentration of 0.0507 g/g solvent (saturation temperature of around 42 ˚C). A control experiment without any impurity was first run in each of the three crystallisers, to account for potential differences in heat transfer or agitation between each vessel. Further experiments were performed where a given impurity was dissolved in each crystalliser to reach a certain

weight percentage with respect to the amount of hexamine dissolved. In each case, impurity concentrations of 2, 4 and 6 wt% were implemented in crystallisers 1, 2 and 3 respectively.

Each solution was heated to 50 ˚C and held for 60 mins to dissolve all solid material before rapidly cooling to a seeding temperature of 40 ˚C. After adding 0.2 g of the sieved hexamine seed crystals (mean size of 178.4 μm), the solution was held at constant temperature for 30 mins before cooling at -0.2 ˚C/min down to 7 ˚C. The product crystals were filtered and dried before imaging and measurement of the CSD, to assess the effects of the impurity on the crystallisation kinetics.

6.2.2. Methodology for seeded batch crystallisation with selected impurity

The impurity selected after screening was introduced into seeded batch crystallisations using the 500 ml jacketed crystalliser setup for characterisation. The same initial hexamine concentration and temperature profile from the screening experiments were implemented, maintaining the same seed loading by using 1 g of the sieved seeds. Batch crystallisations were carried out with impurity concentrations of 0.25, 0.50, 0.75 and 2 wt%. The selection of these concentrations was based on the results presented in section 6.3.2. SEM images were taken of the product crystals using a Hitachi TM3030Plus microscope.

6.3. Results and discussion

6.3.1. Selection of impurity from screening results

The volume-based CSDs of the seed and product crystals from each impurity screening experiment are shown in Figure 6.1, where the results have been grouped depending on which of the three crystallisers was utilised. Numerical details of the volume-based product CSDs are summarised in Table 6.1 for further comparison between the effects of each impurity. Images of the product crystals from each screening experiment are presented in Figures 6.2-6.5.

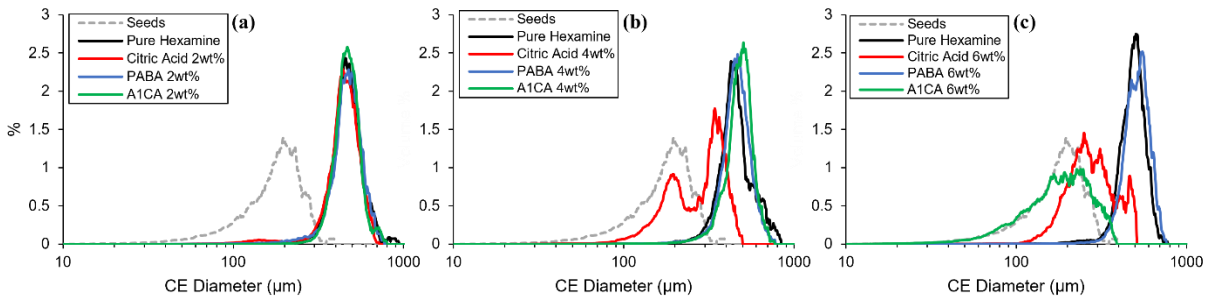


Figure 6.1. Volume-based CSD of the hexamine product crystals from impurity screening. a) Crystalliser 1 (2 wt% impurity), b) Crystalliser 2 (4 wt% impurity) and c) Crystalliser 3 (6 wt% impurity).

Table 6.1. Mean sizes and spans of volume-based product CSDs from each impurity screening batch.

	Pure Hexamine	Citric Acid	PABA	A1CA
Volume Mean Size (µm)				
Crystalliser 1 (2 wt% impurity)	474.9	448.1	472.3	472.1
Crystalliser 2 (4 wt% impurity)	464.3	282.5	464.6	482.1
Crystalliser 3 (6 wt% impurity)	488.4	288.1	516.3	188.0
Span of Volume-Based CSD				
Crystalliser 1 (2 wt% impurity)	0.477	0.491	0.475	0.402
Crystalliser 2 (4 wt% impurity)	0.607	0.811	0.454	0.439
Crystalliser 3 (6 wt% impurity)	0.404	0.935	0.412	1.132

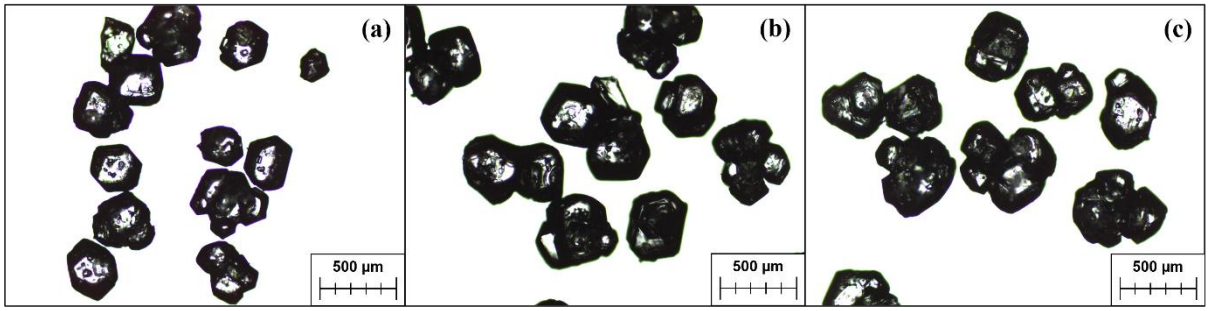


Figure 6.2. Images of hexamine product crystals from impurity screening batches containing pure hexamine in a) Crystalliser 1, b) Crystalliser 2 and c) Crystalliser 3.

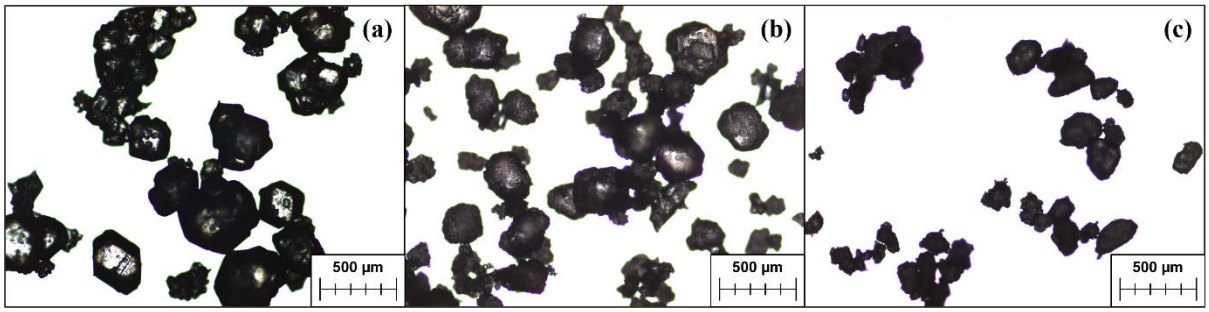


Figure 6.3. Images of hexamine product crystals from impurity screening batches containing a) 2 wt%, b) 4 wt% and c) 6 wt% citric acid.

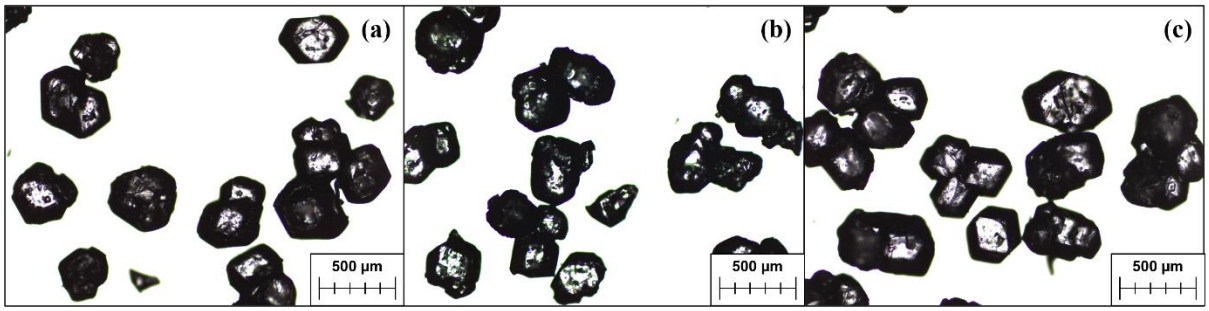


Figure 6.4. Images of hexamine product crystals from impurity screening batches containing a) 2 wt%, b) 4 wt% and c) 6 wt% PABA.

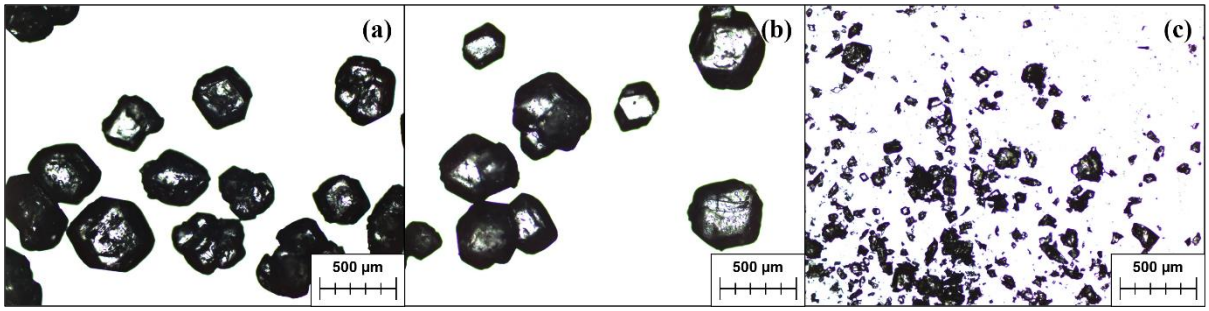


Figure 6.5. Images of hexamine product crystals from impurity screening batches containing a) 2 wt%, b) 4 wt% and c) 6 wt% A1CA.

It was obvious that PABA had little to no effect on the crystallisation of hexamine from ethanol. Images showed no clear difference between the size and shape of the mostly single prismatic crystals, even at the highest impurity concentration. Additionally, the measured mean sizes from batches containing PABA were almost identical to those with pure hexamine, following the same trend of having the smallest mean size in crystalliser 2 and the largest in crystalliser 3. This did at least confirm that there were slight physical differences between each crystalliser, which was taken into consideration while comparing other experiments.

Addition of A1CA also resulted in little change in the product CSD at impurity concentrations of 2 and 4 wt%. However, the latter produced crystals with a mean size around 20 µm larger than the pure hexamine batch in crystalliser 2, where the size was expected to be slightly lower than the other crystallisers. There is the possibility that due to the chemical structural similarity of A1CA to hexamine, some of the impurity was incorporated into the crystal structure and added to the overall growth of the crystals.

At a concentration of 6 wt%, A1CA had a much more drastic effect on the crystallisation behaviour, leading to a final CSD that was similar to that of the original seed material. From the large population of small, fragmented crystals observed in Figure 6.5c, it could be inferred that this was a result of a heightened secondary nucleation rate, likely a result of the increased supersaturation generated by the inhibition of crystal growth. Alternatively, given the close similarity to the original seed CSD, the growth rate may have been completely inhibited and the solution prevented from desaturating entirely. However, this is unlikely since the relative supersaturation that would be generated (around 1.32) would be far beyond the metastable limit measured in chapter 4 (maximum of 0.61 at the 1.5 ml scale). This would result in large amounts of primary nucleation that would form a noticeable peak in the smaller size ranges of the product CSD, which was not observed in the measured data. In a typical industrial process, impurities are able to disturb the crystallisation kinetics at lower concentrations than these.

Therefore, A1CA was rejected as it was deemed unrepresentative of an industrial crystallisation system.

The most effective growth inhibitor of the three impurities was citric acid, which generally showed a consistent trend of decreasing mean size and broadening span of the product CSD with increasing concentration. The mechanism of primary nucleation likely came into play, where the supersaturation increase from the inhibition of growth was large enough to exceed the metastable limit. Images provided clear evidence to support this, as they showed an increasing population of small, agglomerated crystals as the amount of citric acid increased. At a concentration of 6 wt%, what appeared to be larger crystals were theorised to also be agglomerates, supported by the rounded shape and darkening of these particles compared to the normally translucent, prismatic crystals produced without impurities.

Even at a concentration as low as 2 wt% citric acid, the CSD in Figure 6.1a showed the formation of a secondary peak of finer crystals at around 150 μm, which became much more prevalent at 4 wt% and eventually merged with the main peak at the highest concentration. This suggested that there could be a region somewhere below 2 wt% where citric acid would inhibit the growth rate of hexamine without causing uncontrollable nucleation, which made it a suitable choice for investigation at a larger scale. Alternatively, the similarity of the shape and size range of this smaller peak to the original seed CSD could imply that it was simply an indicator of seeds that had been completely restricted from growing due to interaction with citric acid.

6.3.2. Effect of citric acid on crystal growth rate during 500 ml seeded batch crystallisation

Figure 6.6 summarises the results of the seeded batch crystallisation performed in the 500 ml jacketed crystalliser at an impurity concentration of 2 wt% citric acid, including measurements of the concentration during the batch via the attached ATR-FTIR probe.

Concentration measurements revealed that the generation of smaller particles with citric acid present in the system was due to primary nucleation. In this case, nucleation began to occur after around 55 mins into the linear cooling period, where the measured concentration started to rapidly increase. Rather than decreasing as expected, significant fouling of hexamine around the ATR probe resulted in a falsely higher concentration, as the build-up of solids would have increased the intensity of the IR peaks characteristic to hexamine.

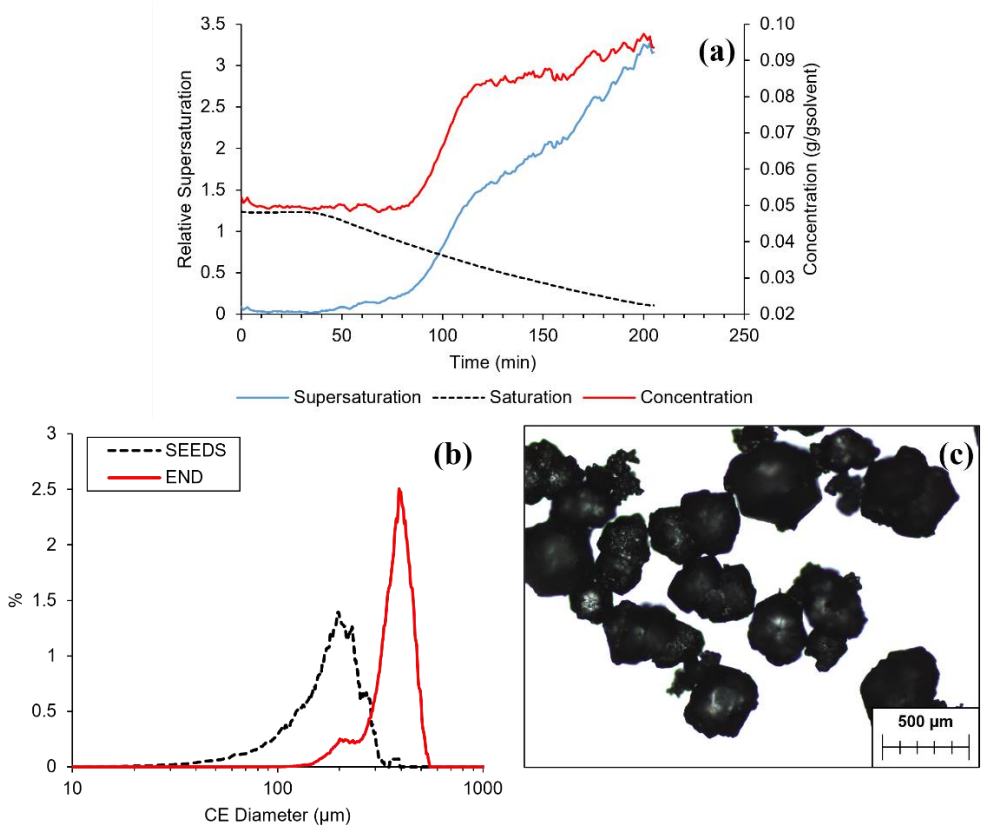


Figure 6.6. Overview of the 500 ml seeded batch crystallisation with 2 wt% citric acid. a) Measured concentration and supersaturation profiles starting from the point of seeding, b) volume-based CSD and c) image of the hexamine product crystals.

Other differences resulting from scale-up of the process included a more pronounced secondary peak in the product CSD at around 200 µm. Images of the product crystals appeared more similar to those from screening experiments at greater impurity concentrations. The SEM images of the product crystals shown in Figure 6.7 provide further evidence that smaller nucleated crystals were present in the system, which agglomerated into clusters around the size of 150-200 µm (i.e. the location of the secondary peak in the CSD).

As this suggested that the effect of citric acid may be intensified by the increased scale, concentrations between 0.25-0.75 wt% were investigated to attempt avoiding the issues caused by primary nucleation. Additionally, these lower concentrations were also considered to be more representative of typical raw material feeds, making the model system more relevant to industrial applications. The concentration profiles measured during seeded batch crystallisations with 0.25, 0.50 and 0.75 wt% citric acid and calculated relative supersaturation profiles are shown in Figure 6.8, compared with those measured during a batch without impurities.

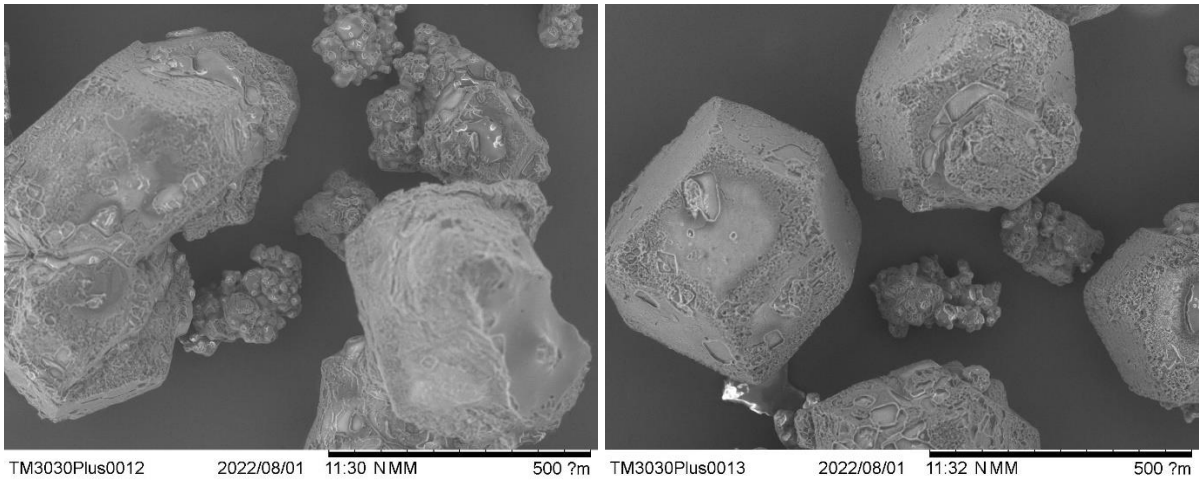


Figure 6.7. SEM images of hexamine product crystals from the 500 ml seeded batch crystallisation with 2 wt% citric acid.

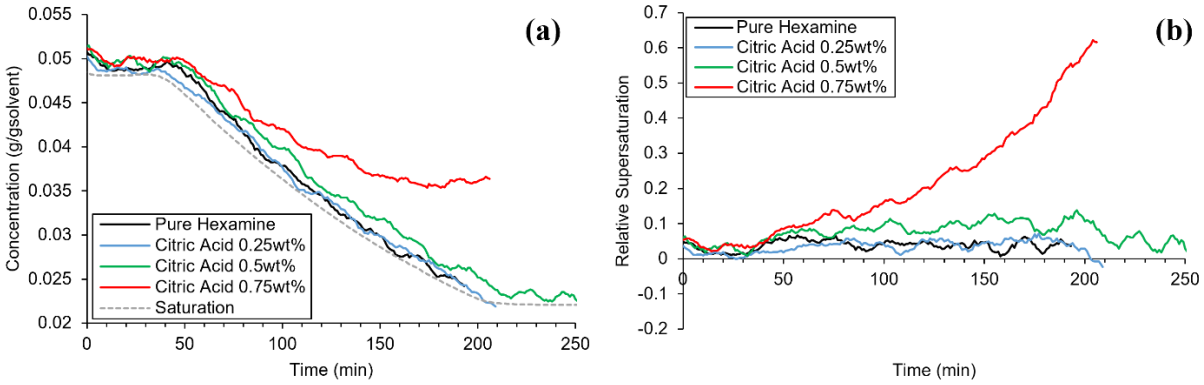


Figure 6.8. Measured a) concentration and b) supersaturation profiles during the 500 ml seeded batch crystallisations with 0.25-0.75 wt% citric acid, starting from the point of seeding.

Within this range of impurity concentrations, a more subtle effect on the crystallisation of hexamine was observed. No significant change was observed at 0.25 wt% citric acid, but at 0.50 wt% there was a noticeable increase in the maximum supersaturation achieved throughout the batch. It was inferred that this was due to inhibition of the crystal growth rate, as slower consumption of the dissolved hexamine lead to overall higher concentrations while the solution was cooled.

The continued increase in supersaturation during the batch with 0.75 wt% citric acid initially implied that some degree of fouling from nucleation on the probe may have occurred once again. However, the measured concentration did not behave in the same way as the batch with 2 wt% impurity (Figure 6.6a). Instead, it decreased at a reduced rate before appearing to level off at a high solution concentration of around 0.035 g/g solvent. To alleviate concerns that the addition of citric acid may have altered the solubility of hexamine in ethanol, dissolution

temperatures were measured for varying concentrations of hexamine with the impurity present. The results of this investigation are presented in Figure 6.9, which shows the measured solubility in a solution of 0.000375 g/g solvent of citric acid in ethanol (the same concentration as the experiments at 0.75 wt% citric acid).

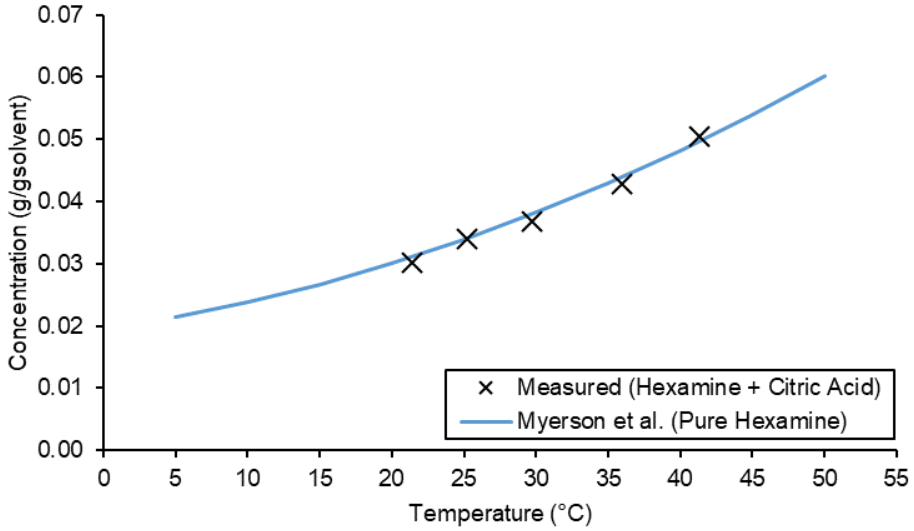


Figure 6.9. Measured solubility of hexamine in ethanol with 0.000375 g/g solvent citric acid, compared to pure hexamine solubility from Myerson et al. (1986).

There was no significant deviation from the solubility of hexamine in pure ethanol solution used throughout this study, therefore the hypothesis that this concentration of citric acid had changed the solubility limit was rejected. It was instead theorised that the rate of crystal growth had been slowed so significantly that, given enough time at the final temperature, the supersaturation would still have eventually been consumed.

6.3.3. Effect of citric acid on hexamine product crystal structure

The volume-based CSDs of the seed and product crystals from each batch are shown in Figure 6.10, with further numerical details of the CSDs summarised in Table 6.2 for comparison between batches. Images of the product crystals from each batch are presented in Figure 6.11.

There was no particular trend in the final crystal size with increasing concentration of citric acid, however each of the batches with the impurity present produced crystals with an overall larger mean size than the pure hexamine batch. Although the similarity of the CSDs provided little evidence to support a theory for this, images of the crystals showed a more noticeable variation between each sample. With increasing concentration of citric acid, the shape of the crystals appeared more rounded than prismatic, which matched observations during the screening experiments.

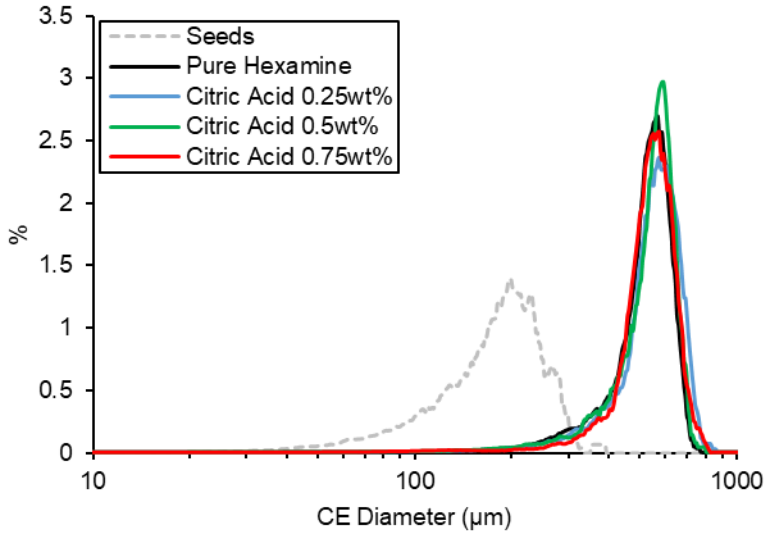


Figure 6.10. Volume-based CSDs of the hexamine product crystals from the 500 ml seeded batch crystallisations with no impurity and 0.25-0.75 wt% citric acid.

Table 6.2. Mean sizes and spans of volume-based product CSDs from the 500 ml seeded batch crystallisations with no impurity and 0.25-0.75 wt% citric acid.

	Pure Hexamine	0.25 wt% Citric Acid	0.50 wt% Citric Acid	0.75 wt% Citric Acid
Volume Mean Size (µm)	519.9	548.3	533.0	538.9
Span of Volume-Based CSD	0.492	0.515	0.482	0.424

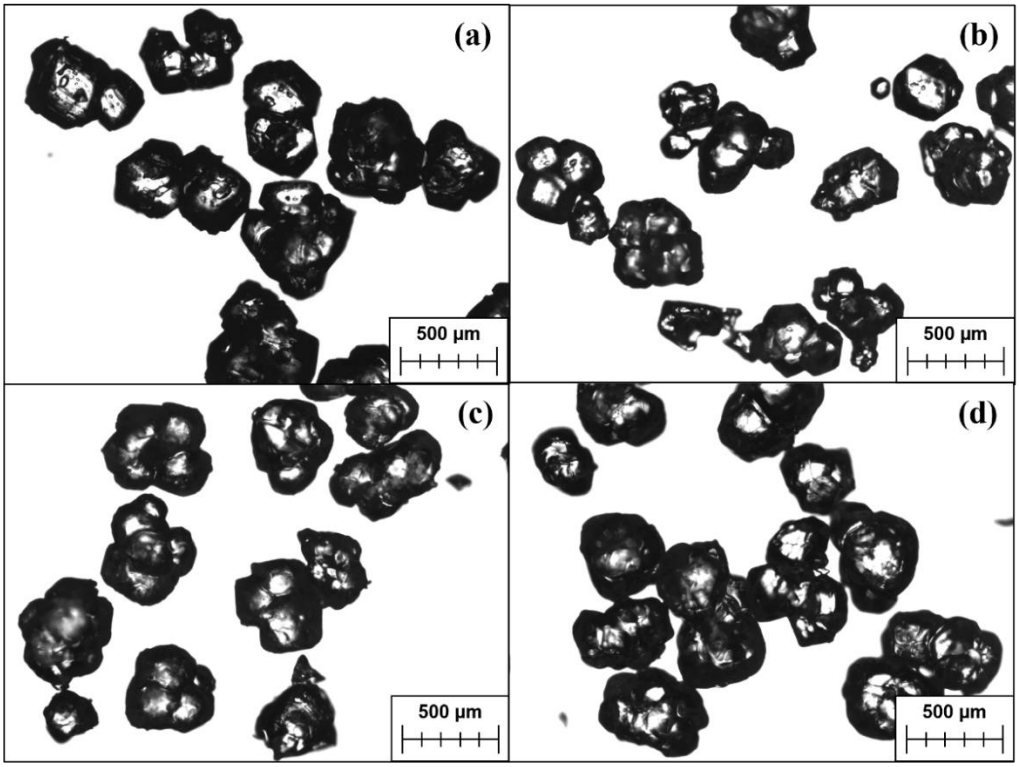


Figure 6.11. Images of hexamine product crystals from 500 ml seeded batch crystallisations with a) no impurity, b) 0.25 wt% citric acid, c) 0.50 wt% citric acid and d) 0.75 wt% citric acid.

The reason for the increase in mean crystal size could have been the more even distribution of hexamine around each particle, caused by the reduced growth rate. Considering the agglomeration mechanism exhibited by hexamine, the spaces formed between the central crystal and the nuclei that form surrounding it would have been filled more gradually by additional material, hence the alteration of the crystal shape. Further evidence of this can be seen in the close-up images of individual product crystals shown in Figure 6.12. A reduced number of solvent inclusions was apparent in the structure of hexamine crystals produced with increasing impurity concentrations. This suggests a disruption to the agglomeration mechanism that would normally lead to the trapping of solvent in the gaps between crystals, which has been influenced by the change in growth rate kinetics.

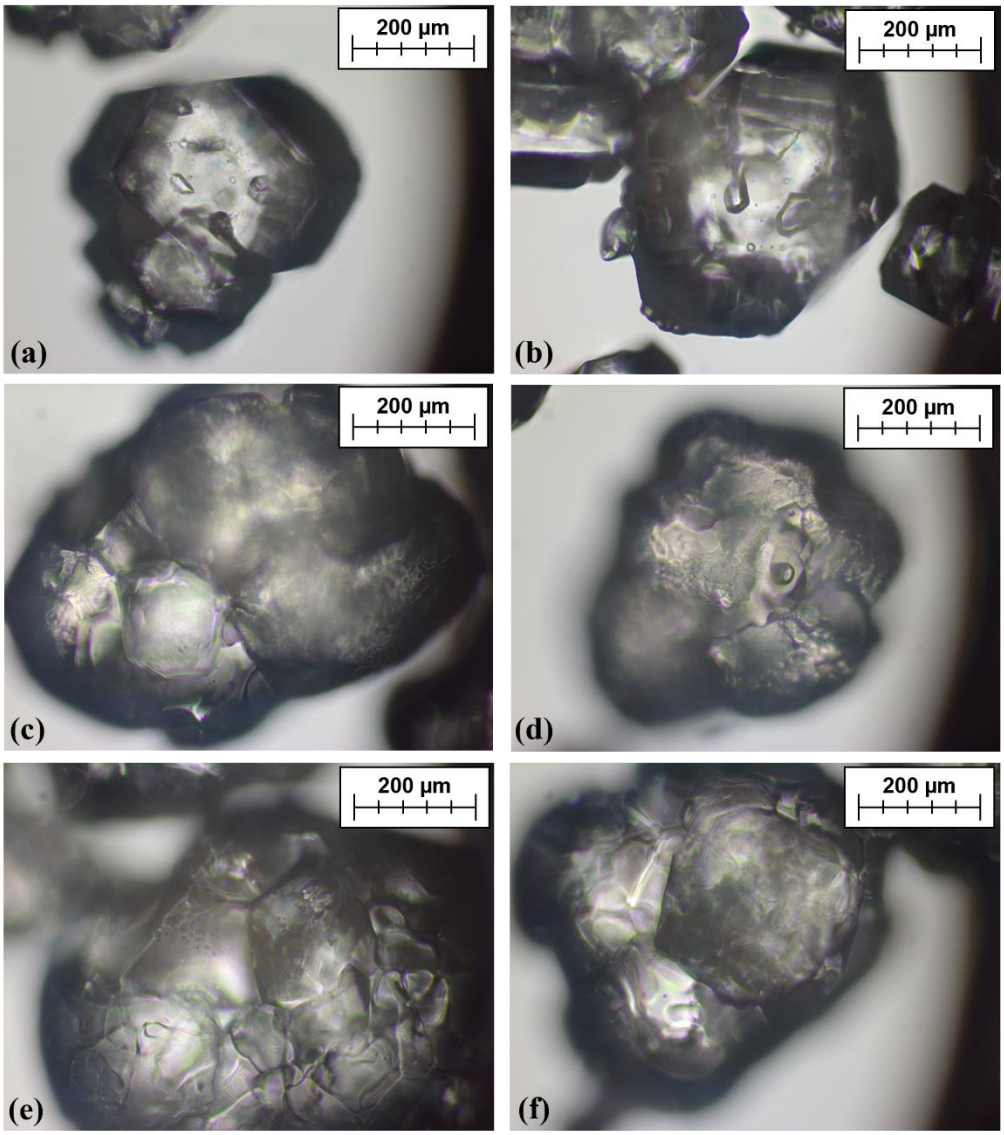


Figure 6.12. Images of individual hexamine product crystals from 500 ml seeded batch crystallisations with a,b) no impurity, c,d) 0.50 wt% citric acid and e,f) 0.75 wt% citric acid.

6.4. Conclusions

Seeded batch crystallisations were performed at the 100 ml scale to investigate the effects of citric acid, PABA and A1CA on the crystallisation kinetics of the hexamine-ethanol system, starting with impurity concentrations between 2-6 wt%. PABA was observed to have little to no effect on hexamine's crystallisation behaviour. Addition of A1CA at lower concentrations also had little effect, however at 6 wt% there was a significant increase in nucleation leading to a product CSD close to that of the original seed crystals. Citric acid was selected as the most suitable growth inhibitor due to the clear trends shown between impurity concentration and product CSD. Given that even 2 wt% citric acid was enough to reduce the growth rate such that supersaturation exceeded the boundary for primary nucleation, the impurity concentration could also be reduced to more industrially relevant levels.

At the 500 ml scale, the primary nucleation induced by addition of 2 wt% citric acid resulted in fouling of hexamine around the ATR probe used to measure solution concentration. To avoid this, lower impurity concentrations between 0.25-0.75 w% were employed, where a more subtle increase in supersaturation was observed due to the reduced crystal growth rate. It was confirmed through dissolution experiments that the solubility of hexamine in ethanol was unaffected by the presence of citric acid up to 0.75 wt%.

A potential influence over the agglomeration mechanism exhibited by hexamine was observed with increasing concentration of citric acid. The slower growth rate was hypothesised to fill spaces between agglomerated crystals more evenly, leading to rounded particles with a reduced amount of solvent inclusions. Overall, the effects of citric acid on both growth rate and agglomeration pose an interesting challenge to the MPC strategy. Specifically, the performance of the controllers that have been trained using a process model that does not directly incorporate these impurity effects needs to be assessed, which will speak for the adaptability of the strategy to disturbances typical of industrial crystallisation processes.

7. Supersaturation MPC Development and Performance with Varying Supersaturation Set-Point, Seeding Conditions and Growth Kinetics

7.1. Introduction

The validated process model representing the hexamine-ethanol crystallisation system was coupled with PharmaMV software, to develop a suitable MPC strategy for supersaturation control during a batch cooling crystallisation. Additionally, selection of citric acid as an effective growth inhibitor allowed the crystallisation kinetics to be modified during physical crystallisation experiments. This was important for testing the performance of the MPC strategy not only to represent an industrially relevant system, but also to investigate how accurate the model's kinetic parameters needed to be.

This chapter outlines the training of a series of model predictive controllers in PharmaMV using the gPROMS crystallisation process model that was parameterised in chapter 5. The developed MPC strategy maintains the requested supersaturation set-point through manipulation of the vessel temperature profile over time. The supersaturation control strategy was tested in-silico using the digital twin of a 500 ml jacketed crystalliser to simulate measurements of temperature and solution concentration. After carrying out these tests for a range of supersaturation set-points, each of these simulations were validated by using the MPC strategy to control the circulator in a physical 500 ml crystallisation setup. For further assessment of the performance of the physical supersaturation MPC strategy, additional batches were performed with varying seed crystal masses/sizes and amounts of citric acid impurity. Through analysis of the mean value and standard deviation of the supersaturation during these experiments, the accuracy and precision of the MPC strategy over a wide range of operating conditions was established.

7.2. Materials and methodology

7.2.1. CrystalMV template for supersaturation MPC in PharmaMV

The process control software PharmaMV v8.1.1445 was used as the platform for development of the MPC strategy. The software had the capability to perform real-time control of specific parameters during a physical crystallisation by taking measurements with appropriate in situ PAT. Prior to this, in-silico simulation using the digital twin in gPROMS was utilised to test the MPC strategy's performance, as well as to investigate the range of operating conditions that would be feasible for real processes.

Most information in the PharmaMV real-time control software is treated as either an actuation signal (sent from PharmaMV to external software/equipment) or a measured signal (calculated by PharmaMV or measured from external source). For example, the temperature set-point value communicated to a heater/chiller would be an actuation signal, whereas a measurement of the current bath temperature would be recorded as a measured signal. A typical MPC 'block' within PharmaMV involves a statistical relationship being calculated between an actuation signal and one or more measured signals.

The resulting changes in the controlled variables are calibrated to an auto-regressive with exogenous input (ARX) model structure using the recursive least squares (RLS) method (Shamekh et al., 2008). The objective function to be optimised for MPC is shown in Equation 7.1 (Clarke et al., 1987):

$$J = \sum_{j=1}^{p} \alpha_j \left(y_r \left(k + \frac{j}{k} \right) - \hat{y}(k + j/k) \right)^2 + \sum_{j=1}^{m} \beta_j \Delta u(k + j - 1/k)^2 \qquad (7.1)$$

where J is the objective function to be minimised, k is the sampling instant, p is the prediction horizon, m is the control horizon, y_r is the set-point reference, $\hat{y}$ is the estimated future output, α_j and β_j are weighting parameters for the controlled and manipulated variables respectively, and Δu is the control move (i.e. the incremental change in the manipulated variable to be calculated by the MPC algorithm).

The objective function is designed to track the set-point reference trajectory over a specified length of time – the prediction horizon p – and determine the optimal value for the control move Δu that will minimise the set-point error. Within the MPC blocks in PharmaMV, empirical tuning parameters can be manipulated either before or during a control run, to refine the stability of these control moves. For example, if changes in the cooling rate were too large and resulted in repeated over/undershooting of the supersaturation, the 'move weight' could be increased to limit changes to the cooling profile and keep the supersaturation level steady.

The MPC strategy described in this chapter was built upon an existing template provided by Perceptive Engineering Ltd., labelled 'CrystalMV'. This template consisted of two MPC blocks, connected as shown in Figure 7.1 such that the specified supersaturation set-point would lead to manipulation of the circulator temperature. The supersaturation MPC block compared the current supersaturation to the specified set-point, providing an output of the cooling rate required to approach the set-point value. The optimal cooling rate was calculated

based on a PLS model built into the MPC block; the procedure for training these models is described in section 7.2.2.

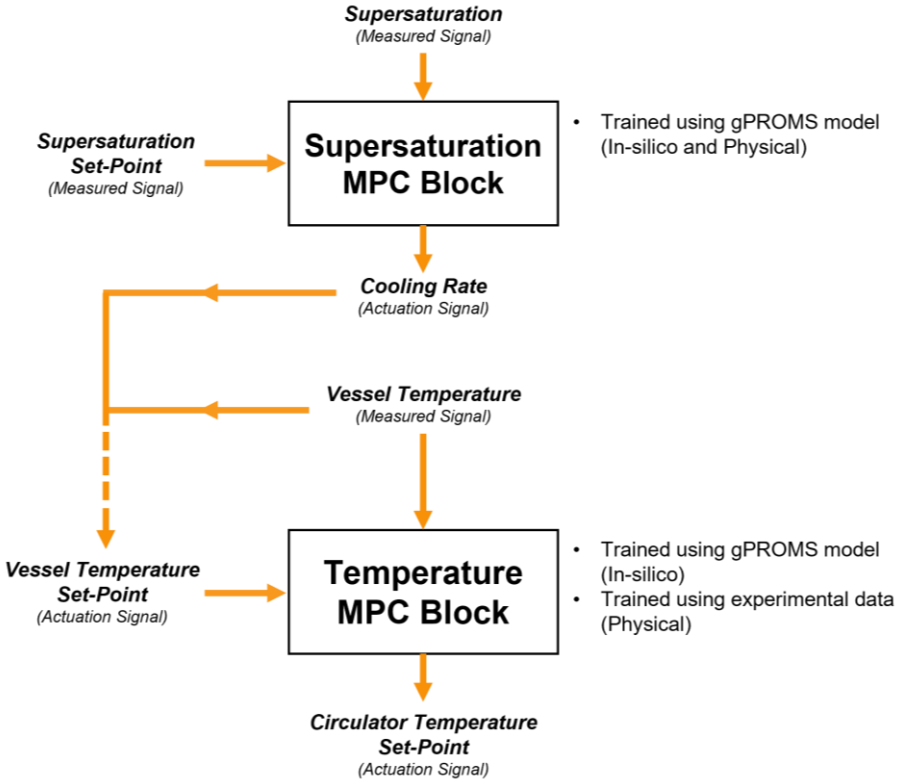


Figure 7.1. Block diagram of the CrystalMV supersaturation control structure in PharmaMV.

The cooling rate was used alongside the measured vessel temperature to construct a temperature set-point profile (calculated via a Python script) that was the input for the temperature MPC block. This second block represented the heat transfer between the crystallisation vessel and jacket, giving an output of the necessary circulator temperature needed to achieve the correct cooling profile within the vessel.

The temperature and supersaturation MPC blocks made decisions to update their associated manipulated variables every 15 seconds and 60 seconds respectively. The reason for choosing a longer time interval for the supersaturation controller was that this block benefited from having more data available to determine each control move, helping to improve the stability.

7.2.2. Training of temperature and supersaturation MPC blocks

Each MPC block was trained through pseudo-random binary sequence (PRBS) testing, where the corresponding manipulated variable for each controller was varied between set limits to measure the response of the controlled variables. Initially, both MPC blocks were trained using the gPROMS process model to simulate the measurements of temperature and concentration

during a hexamine-ethanol batch crystallisation. The initial conditions for the simulation were similar to those of seeded Run A described in chapter 4 (500 ml vessel, 40 ˚C seeding temperature, 1 g of seeds with mean size of 178.4 μm). Only the growth rate model was activated; nucleation was deactivated due to experimental observations of little to no nucleation. The seed CSD was specified using the lognormal peak parameters (199 μm location parameter, 40 μm standard deviation).

For the temperature MPC block, the circulator temperature was stepped between 5-50 ºC, which covered the full range of temperatures expected during the batch crystallisation experiments. For the supersaturation MPC block, a cooling ramp was run using the trained temperature controller, starting from 40 ºC. This generated a change in supersaturation that was measured and calibrated to by stepping the cooling rate between -0.1 and -1.5 ºC/min. Although this maximum cooling rate was higher than what would be feasible in physical experiments, generating a training data set with larger changes in supersaturation ensured the MPC block could cover a broad range of operating conditions.

To ensure good controller performance during physical control experiments, the temperature MPC block had to be re-trained with measurements from the 500 ml jacketed crystalliser for MPC validation (see section 3.2.2) to capture more realistic heat transfer dynamics. The crystalliser was filled with around 500 ml of water initially held at 20 ˚C and agitated at 450 rpm, stepping the circulator temperature between 15-25 ˚C.

7.2.3. Methodology for in-silico supersaturation MPC

For the in-silico testing of the supersaturation MPC strategy, the gPROMS process model was connected with PharmaMV to provide the necessary temperature, concentration and other process measurement data, as described by the diagram in Figure 7.2. A batch crystallisation simulation in gPROMS and data sampling in PharmaMV were run simultaneously, where the latter received signals for the current temperature and supersaturation (calculated from solution concentration and solubility data) within the vessel. The two MPC blocks functioned as described in section 7.2.1, providing an output of the circulator temperature set-point that was fed back to the temperature controller within the gPROMS model flowsheet. Data was transferred through this loop every 5 seconds, stopping the simulation automatically when a set end temperature was reached.

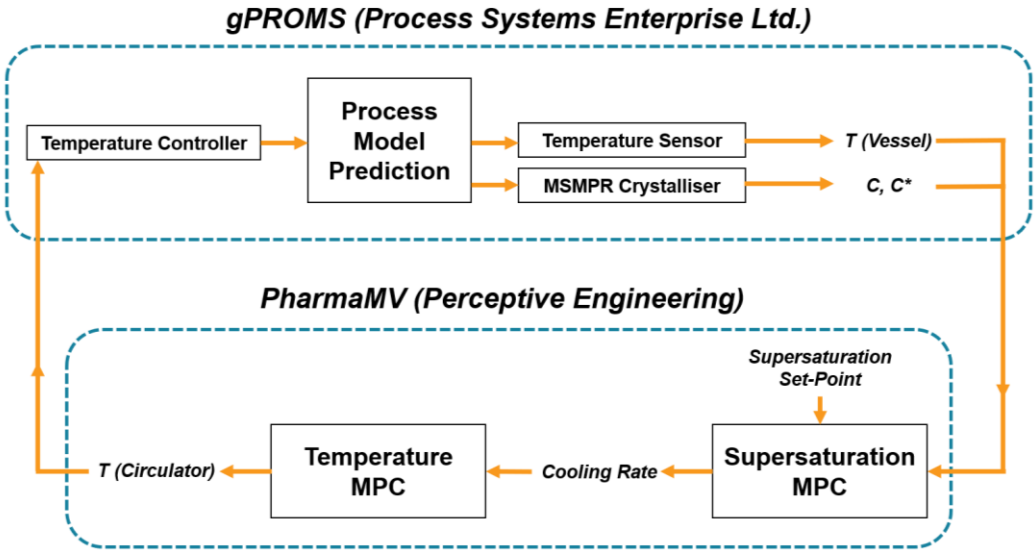


Figure 7.2. Data flow within the gPROMS/PharmaMV digital twin during a supersaturation control simulation.

To test the fully trained digital twin, examples of batch crystallisations with supersaturation control were simulated. Table 7.1 summarises the process conditions applied in these simulations, where the only variation between them was the chosen supersaturation set-point between a range of 0.012-0.036. The parameters describing the seed CSD were selected based on those produced by wet milling for physical validation (mean size of 112.2 µm). During each simulation, the temperature was held at the initial value for 30 mins before switching on both of the MPC blocks, which turned off automatically when the vessel temperature reached 5 ˚C. Only the growth rate model was activated in gPROMS for these simulations.

Table 7.1. Process conditions applied to the gPROMS/PharmaMV digital twin for in-silico supersaturation MPC performance testing.

Crystalliser volume (ml)	500
D_{50} of seed CSD (µm)	110
Standard deviation of seed CSD (µm)	25
Initial seed mass (g)	1
Initial temperature (ºC)	40
End temperature (ºC)	5
Supersaturation set-point	0.012, 0.024, 0.036

7.2.4. Methodology for physical validation of supersaturation MPC

The supersaturation control strategy developed in PharmaMV was implemented at the CMAC Future Manufacturing Research Hub using the 500 ml jacketed crystalliser for MPC validation, described in section 3.2.2. Compared to the crystallisation setup at the University of Leeds

utilised in chapters 4 and 6, this setup had been adjusted to allow for automated operation of the circulator through communication with the PharmaMV software platform. Figure 7.3 illustrates how the MPC data flow schematic in Figure 7.2 was adapted for the physical crystallisation setup. It should be noted that FBRM and PVM data were not collected for the experiments described in this chapter, hence only the ATR-FTIR probe has been shown.

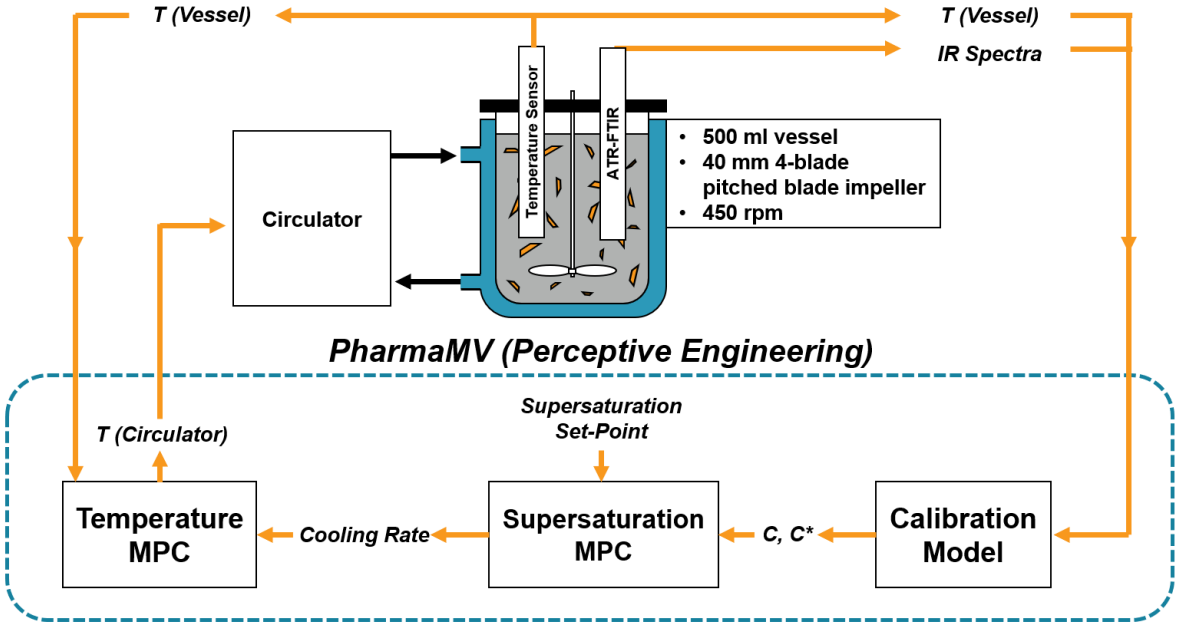


Figure 7.3. Data flow within the MPC strategy during a physical batch cooling crystallisation with supersaturation control.

A concentration calibration model was developed within PharmaMV software using new spectral measurements with the ReactIR system, following the same procedure and experimental conditions described in section 3.3. A second derivative transformation using a Savitzky-Golay filter with 10 points of smoothing was applied to the spectra before calibration. The calibration model was incorporated into the overall strategy as its own block, which communicated the measured concentration to the supersaturation MPC block.

Table 7.2 summarises the differences between the process conditions of each physical supersaturation control experiment. These experiments were grouped to investigate the performance of the MPC strategy with varying supersaturation set-point (Batches A-C), seeding conditions (Batches B and D-F) and impurity concentration (Batches B and G-I).

At the beginning of each physical control experiment, hexamine was added to 400 g of ethanol to make up an initial concentration of 0.0507 g/g solvent (saturated at 42 ˚C). For experiments with the impurity present, citric acid was also added at this stage to make up the chosen concentration in wt%. The solution was heated to 50 ˚C and held for 60 mins to ensure all

material had dissolved. The solution was then cooled at the maximum rate achievable by the circulator (-0.6 ˚C/min) down to 40 ˚C, where the chosen amount of dry seed crystals were added. Following this, the crystalliser was held at constant temperature for 30 mins to allow the system to equilibrate, after which the temperature and supersaturation MPC blocks in PharmaMV were switched on. Once the final temperature of 5 ˚C was reached, the product crystals were removed and filtered under vacuum for analysis.

Table 7.2. Process conditions for the physical supersaturation MPC validation experiments.

Batch	Supersaturation Set-Point	Seed Mass (g)	Seed Mean Size (μm)	Impurity Concentration (wt%)
A	0.012	1	112.2	0
B	0.024	1	112.2	0
C	0.036	1	112.2	0
D	0.024	3	112.2	0
E	0.024	1	68.9	0
F	0.024	3	68.9	0
G	0.024	1	112.2	0.25
H	0.024	1	112.2	0.50
I	0.024	1	112.2	0.75

7.3. Results and discussion

7.3.1. Data generated for training MPC blocks

Figure 7.4 shows the results of the PRBS tests for both MPC blocks. Throughout this chapter, temperatures are specified either as process values (PV) for measurements of the current circulator or vessel temperatures (in this case, simulated by gPROMS) or set-points (SP) for the values determined by the controllers in PharmaMV.

The minimum cooling rate of -0.1 ˚C/min was low enough to induce a drop in supersaturation; the model predicted that the consumption of supersaturation due to crystal growth was dominant over its generation through lowering the solubility limit. This resulted in a clear trend of increasing/decreasing supersaturation at high/low cooling rates that could be used to reasonably train the MPC block. It is important to note that a given change in cooling rate may not always induce the same response in supersaturation. This would likely be due to variation in the total crystal surface area available for growth resulting from changes in the number of particles, such as by secondary nucleation or agglomeration. Hence, the cooling rate was cycled

between the upper and lower limits multiple times in order for the MPC block to capture an average and therefore more generalised response for the crystallisation system.

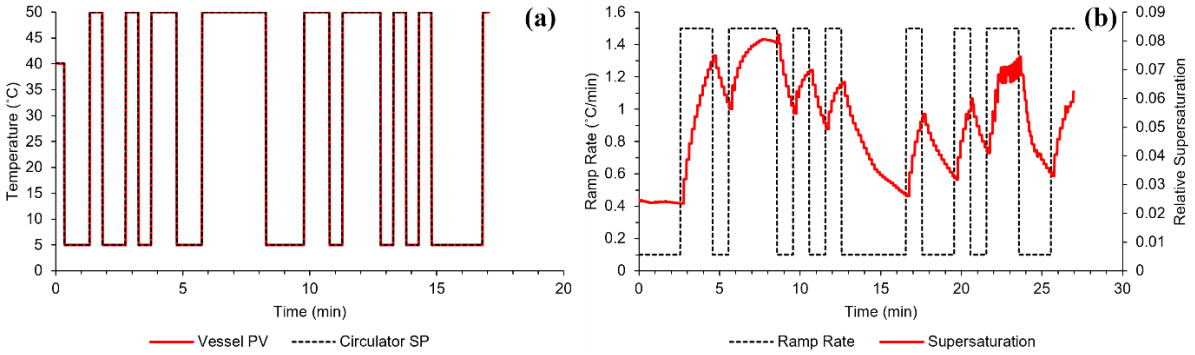


Figure 7.4. PRBS test procedure and controlled variable response predicted by the gPROMS process model for training a) temperature MPC block and b) supersaturation MPC block.

At the faster cooling rate, it was often the length of the step that determined the highest supersaturation that could be achieved before the supersaturation trajectory levelled off. However, comparing the relatively constant supersaturations at around 8 and 23 minutes in Figure 7.4b, the maximum supersaturation noticeably dropped from 0.08 to 0.07. This was likely due to the growth and subsequent increase in surface area available of the suspended crystals, resulting in a faster consumption of hexamine from solution and therefore limiting the maximum supersaturation the longer the batch continued. It was hypothesised that this would be reflected during supersaturation control experiments, where a faster cooling rate would be requested by the MPC block as the process continued, in order to keep the supersaturation at the desired set-point.

As heat transfer between the jacket fluid and the vessel contents was not considered in the gPROMS model, there was an instantaneous response between the circulator SP and vessel PV temperatures. It is important to note that this consequently affected the response of the supersaturation to changes in the cooling rate. Very fast alterations were made to the supersaturation trajectory given the immediate changes in cooling rate, which would not be feasible in a real crystallisation process.

The results of the PRBS test for re-training the temperature MPC block for physical experiments are shown in Figure 7.5. This new training procedure was able to capture not only the real heat transfer between the circulator and vessel, but also the internal temperature control within the circulator. The circulator bath temperature was found to overshoot the set-point when large, rapid changes in this value were made, which took around 4-6 minutes to correct. Fortunately, this was only the worst-case scenario when operating at cooling rates of around -

2.5 ˚C/min, which was much greater and unfeasible cooling rate for real crystallisation experiments. In fact, given that the maximum cooling rate achieved in the vessel appeared to be around -0.6 ˚C/min, the controller was limited to cooling rates below this value to prevent the circulator from struggling to keep up with its set-point profile.

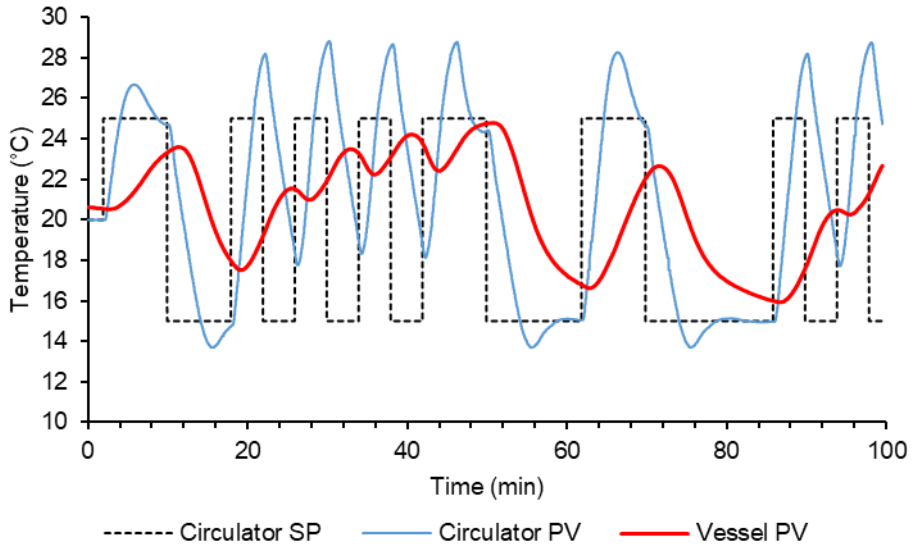


Figure 7.5. PRBS test procedure and controlled variable response measured using the 500 ml jacketed crystalliser for re-training the temperature MPC block.

7.3.2. Performance of in-silico supersaturation MPC

The cooling and supersaturation profiles resulting from the in-silico supersaturation control simulations at three different supersaturation set-points are presented in Figure 7.6. In each case, the supersaturation decreased as expected during the initial 30 min holding period, where the seeds crystals were left to grow. The increase in supersaturation at the beginning of cooling would typically result in an offshoot above the set-point, which the MPC strategy was able to correct by manipulating the cooling rate. The circulator set-point temperature and the measured circulator and vessel temperatures always overlapped, therefore only a single temperature profile has been plotted in Figure 7.6. This is due to the gPROMS process model assuming almost instantaneous heat transfer dynamics (i.e. no time delay between change in circulator and response in vessel temperatures), which were captured during training of the MPC block.

For the lowest set-point of 0.012, the simulation was ended early; although the level of control was very good after the initial offshoot, the low cooling rates chosen by the controller were anticipated to lead to a long batch time. This would likely be considered ineffective from an industrial productivity viewpoint. Experiments were completed in a more reasonable time for the remaining set-points, due to the overall higher cooling rates required to maintain high supersaturations. There was a noticeable increase in the instability of control as the set-point

increased, caused by the more drastic changes made to the cooling rate by the supersaturation MPC block.

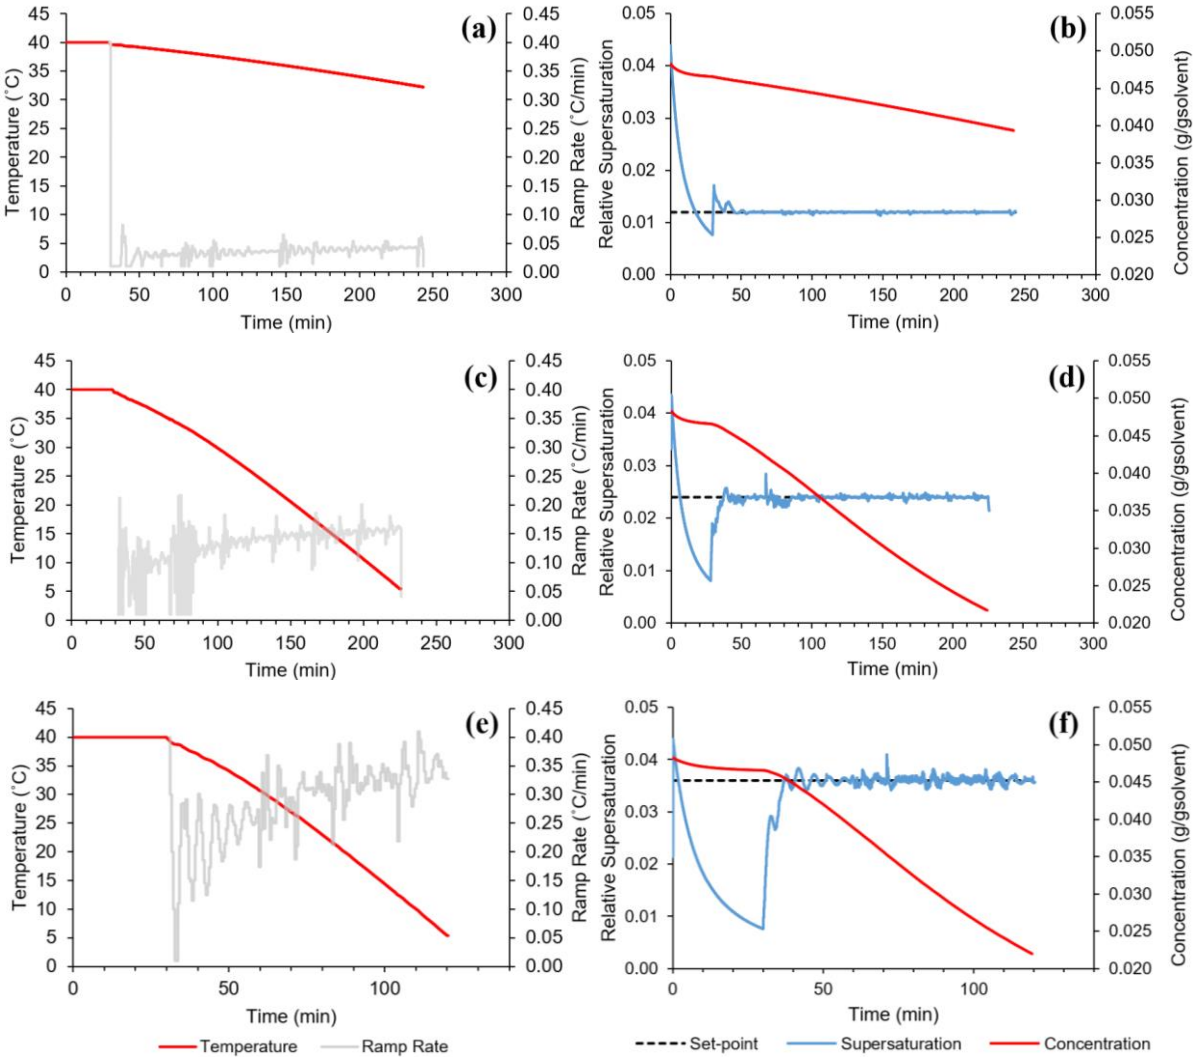


Figure 7.6. Temperature, cooling (ramp) rate, concentration and supersaturation profiles generated during in-silico supersaturation MPC, for supersaturation set-points of a,b) 0.012, c,d) 0.024 and e,f) 0.036.

The tuning weights within the MPC block could be adjusted to reduce these unstable fluctuations in supersaturation by preventing large changes in the cooling rate set-point. However, a variable tuning would have to be applied to dampen these responses only when the measured supersaturation is close to the set-point. Otherwise, the MPC strategy would take too long to bring the supersaturation up to the set-point in the early stages of the crystallisation process.

7.3.3. Comparison of in-silico and physical supersaturation MPC with varying supersaturation set-point

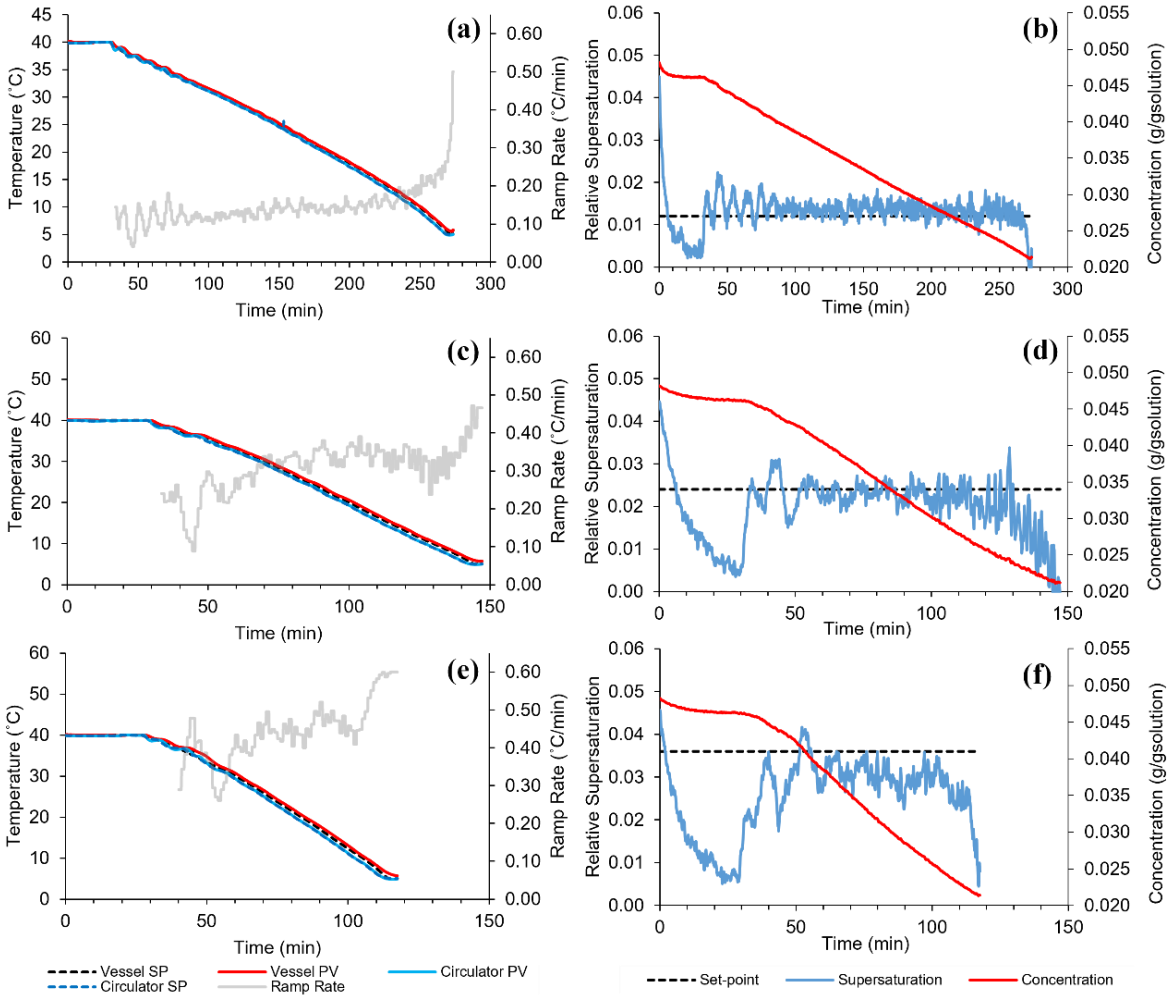


Figure 7.7. Temperature, cooling rate, concentration and supersaturation profiles measured during physical supersaturation MPC experiments, for supersaturation set-points of (a,b) 0.012, (c,d) 0.024 and (e,f) 0.036.

Table 7.3. Mean and standard deviation of supersaturation and average cooling rates during in-silico simulations and physical validation experiments for supersaturation control at varying supersaturation set-points.

	Supersaturation Mean		Supersaturation Std. Dev.		Average Cooling Rate (°C/min)	
Supersaturation Set-Point	In-Silico	Experimental	In-Silico	Experimental	In-Silico	Experimental
0.012	0.0121	0.0137	4.79×10^{-4}	2.03×10^{-3}	0.04	0.16
0.024	0.0239	0.0227	4.76×10^{-4}	3.15×10^{-3}	0.13	0.31
0.036	0.0359	0.0300	8.05×10^{-4}	4.00×10^{-3}	0.29	0.44

The cooling and supersaturation profiles resulting from physical supersaturation control experiments performed at different supersaturation set-points are presented in Figure 7.7. To assess the performance of the MPC strategy, the means and standard deviations of the measured

supersaturation profiles have been calculated and listed in Table 7.3, alongside those calculated from the in-silico simulated profiles for the same conditions. The average cooling rates throughout each batch are also included in this table.

Generally, it was observed that higher cooling rates were required to reach and maintain a given supersaturation set-point when compared with simulations; for example, at a supersaturation set-point of 0.012, the cooling rate was approximately four times higher. This resulted in the final temperature of 5 ºC being reached in a much shorter time than in the simulation. Higher cooling rates often led to more unstable control, given the work required by the circulator to respond to its frequently changing temperature set-point. This may have been a contributing factor towards the noticeable increase in instability of the supersaturation from simulations to physical experiments, as reflected by the standard deviations being an order of magnitude higher. The trend of increasing instability at higher supersaturation set-points matched that observed in simulations.

Considering potential factors related to the software, the heat transfer dynamics captured by the training data provided to each MPC block may have been the cause of instability. The temperature MPC block could be adjusted to match the dynamics achievable by the circulator and jacketed vessel quite easily. However, the gPROMS model-trained supersaturation MPC block was still operating under the assumption of perfect temperature control. This meant that the MPC strategy would expect an immediate change in supersaturation upon choosing a new cooling rate, which was not feasible to achieve in the physical experimental setup. Tuning of the control parameters of the supersaturation MPC block could not fully eliminate this issue, which could instead be resolved by re-training the controllers using a more accurate heat transfer model.

The mean supersaturation values were much further from the set-points for the physical experiments, as expected when considering the disturbances caused by imperfect temperature control and concentration measurement accuracy and precision. With a supersaturation set-point of 0.012, this was the only case where the mean value exceeded the set-point. With the concentration declining very gradually due to slower growth rates at low supersaturations, it is likely that the system was more sensitive to increases in supersaturation when the cooling rate was set too high.

The opposite effect could explain why the mean supersaturation fell short of the higher set-points of 0.024 and 0.036; in these cases, the system appeared to be more sensitive to a drop in

supersaturation should the cooling rate be insufficient. For the highest set-point, the MPC strategy struggled to bring the supersaturation up to this value for most of the batch. This suggested that a supersaturation of 0.036 may have been above the limit achievable by the cooling rates allowed by the MPC strategy. These cooling limits could not be increased without creating a large temperature differential between the vessel and circulator temperatures, which would have made precise control over the cooling rate difficult to maintain.

7.3.4. Performance of physical supersaturation MPC with varied seeding

The cooling and supersaturation profiles resulting from physical supersaturation control experiments performed using different seed masses and sizes are presented in Figure 7.8. Table 7.4 summarises the average cooling rates for each batch experiment, alongside the means and standard deviations of the measured supersaturation profiles. The results from Batch B (1 g of 112.2 µm seeds) are included for comparison.

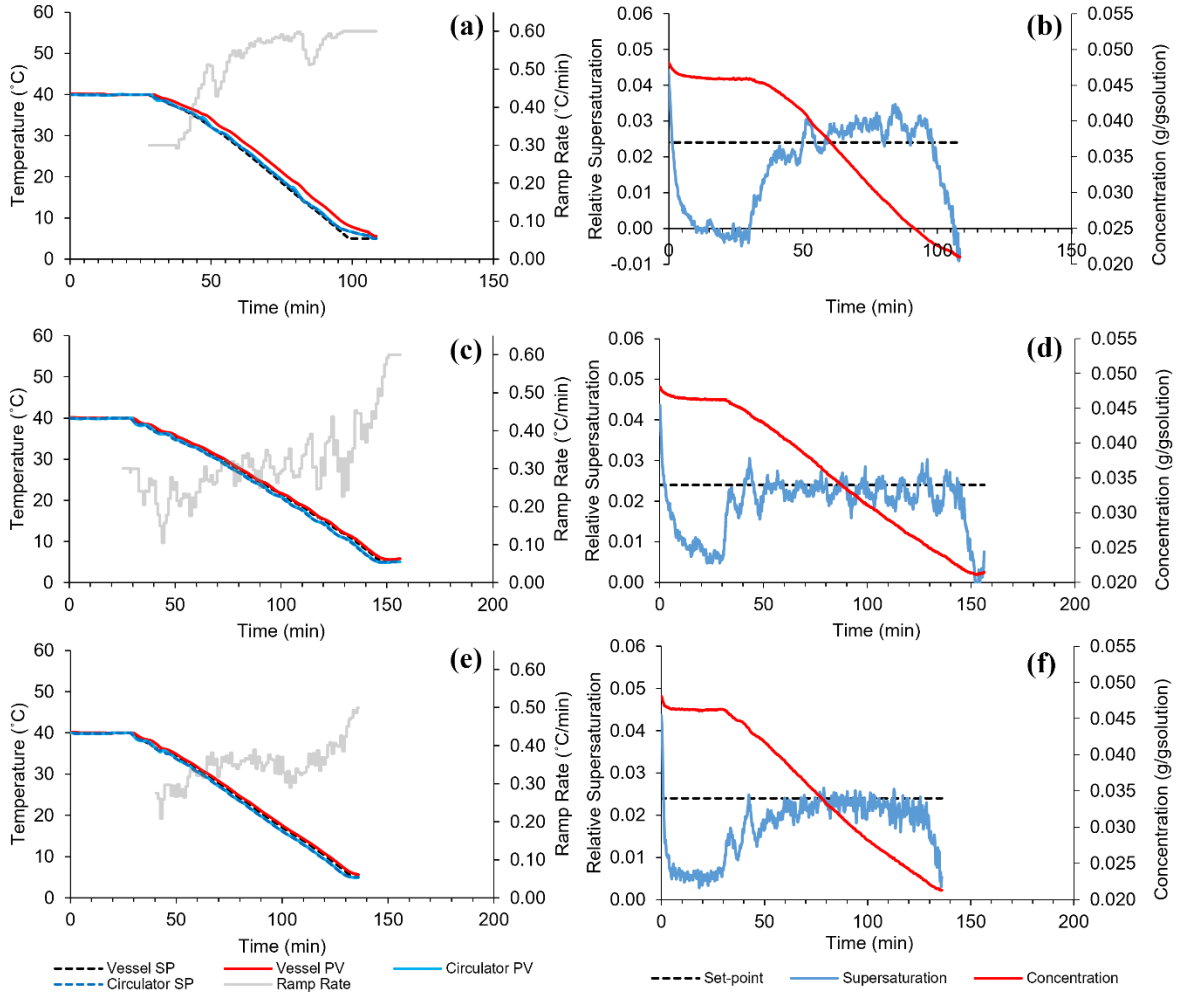


Figure 7.8. Temperature, cooling rate, concentration and supersaturation profiles measured during physical supersaturation MPC, for batches seeded with (a,b) 3 g of 112.2 µm seeds, (c,d) 1 g of 68.9 µm seeds and (e,f) 3 g of 68.9 µm seeds.

Table 7.4. Mean and standard deviation of supersaturation and average cooling rates during physical supersaturation control experiments at varying seeding conditions.

Seeding Conditions	Supersaturation Mean	Supersaturation Std. Dev.	Average Cooling Rate (˚C/min)
1 g of 112.2 μm seeds	0.0227	3.15×10^{-3}	0.31
3 g of 112.2 μm seeds	0.0266	3.75×10^{-3}	0.42
1 g of 68.9 μm seeds	0.0224	2.59×10^{-3}	0.32
3 g of 68.9 μm seeds	0.0212	2.44×10^{-3}	0.35

When introducing a larger mass of the 112.2 μm seeds, the MPC strategy seemed to overcompensate for the increased rate of desupersaturation, resulting from the increased surface area of the seeds. The average cooling rate selected by the supersaturation MPC block was the second fastest between all nine of the batch experiments presented in this chapter. In doing so, the mean supersaturation was raised higher than that observed in the batch with a smaller mass of seeds, as the controller failed to make the decision to lower the cooling rate enough for supersaturation to be consumed by the system.

Interestingly, this excessive cooling rate was not observed when using smaller seeds instead of increasing the seed mass, despite this also leading to an increase in available surface area. There was overall no difference between the experiments starting with 1 g of each seed size, with the mean and standard deviation of the measured supersaturation only slightly lowered when decreasing the seed size 112.2 μm to 68.9 μm. A reason for this could be the increased amount of agglomeration when starting with smaller seeds, which was also observed in the seeded experiments in chapter 4. The uneven surfaces of agglomerates may have altered the true crystal growth rate for single crystals, keeping it at a similar level to the batch with a larger seed size despite the difference in surface area should the suspended crystals had a more even, prismatic shape. However, images of the product crystals from this set of validation experiments, shown in Figure 7.9, suggested that the degree of agglomeration between Batches D and E was similar.

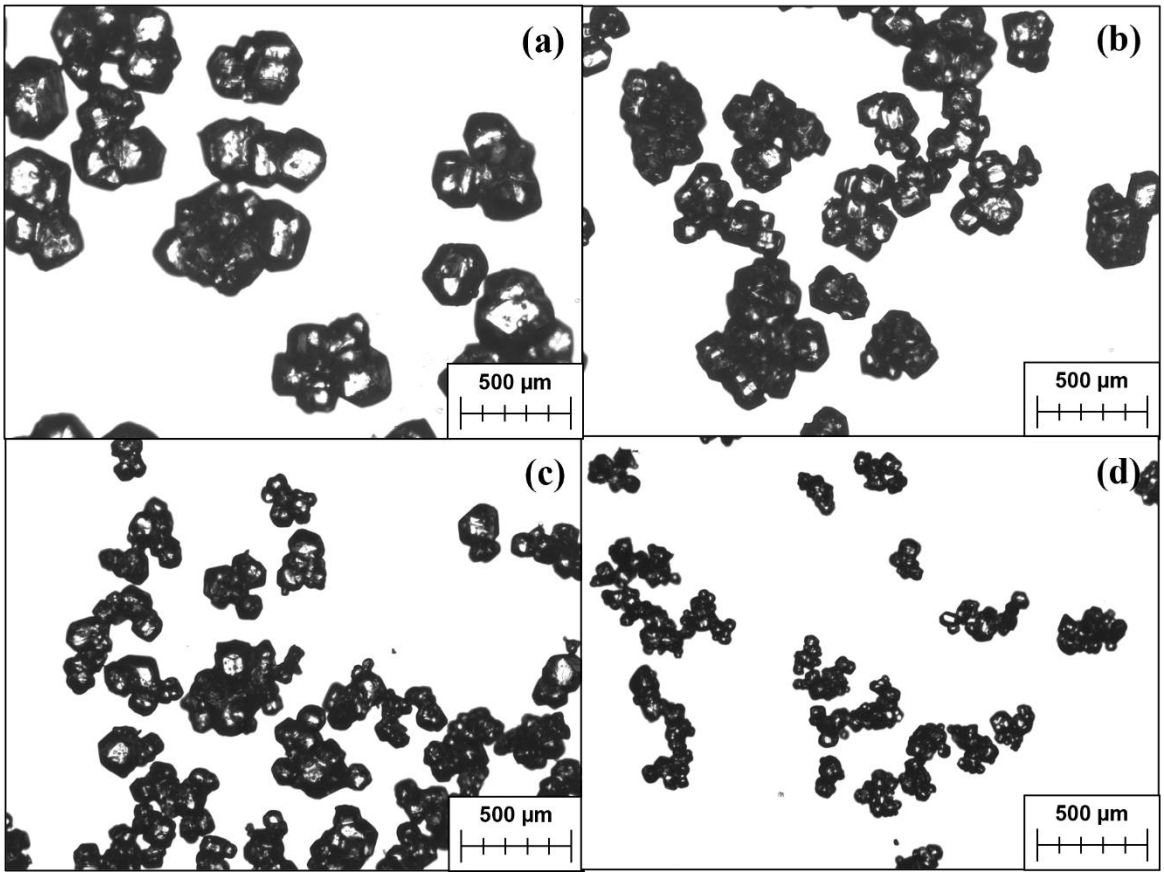


Figure 7.9. Images of hexamine product crystals from physical supersaturation MPC batches seeded with a) 1 g of 112.2 µm seeds, b) 3 g of 112.2 µm seeds, c) 1 g of 68.9 µm seeds and d) 3 g of 68.9 µm seeds.

It was only through the combination of both a larger seed mass and smaller seed size that the resulting supersaturation profile behaved in a more expected way based on the change in crystal surface area. In this case, the mean supersaturation was the lowest between all batches with a set-point of 0.024, which was likely due to the MPC strategy struggling to compensate for the increased consumption of solute compared to the simulations it was trained on. Despite this, the cooling rate was steadier than previous batches, leading to an improvement in the overall precision of control.

Overall, the stability of supersaturation control during each of these batches was better than the worst case observed at the highest supersaturation set-point (Batch C). Therefore, it can be said that the MPC strategy was able to handle variations in the seeding conditions well, even when operating outside of the training range of crystallisation behaviour.

7.3.5. Performance of physical supersaturation MPC with varying impurity concentration

The cooling and supersaturation profiles resulting from physical supersaturation control experiments performed with different concentrations of citric acid impurity are presented in

Figure 7.10, with Table 7.5 summarising the means and standard deviations of the measured supersaturation profiles. The results from Batch B (no impurity) are included in the table for comparison. Note that for Batches H and I (Figure 7.10c-f), profiles have been shown only for the period that the MPC blocks were switched on.

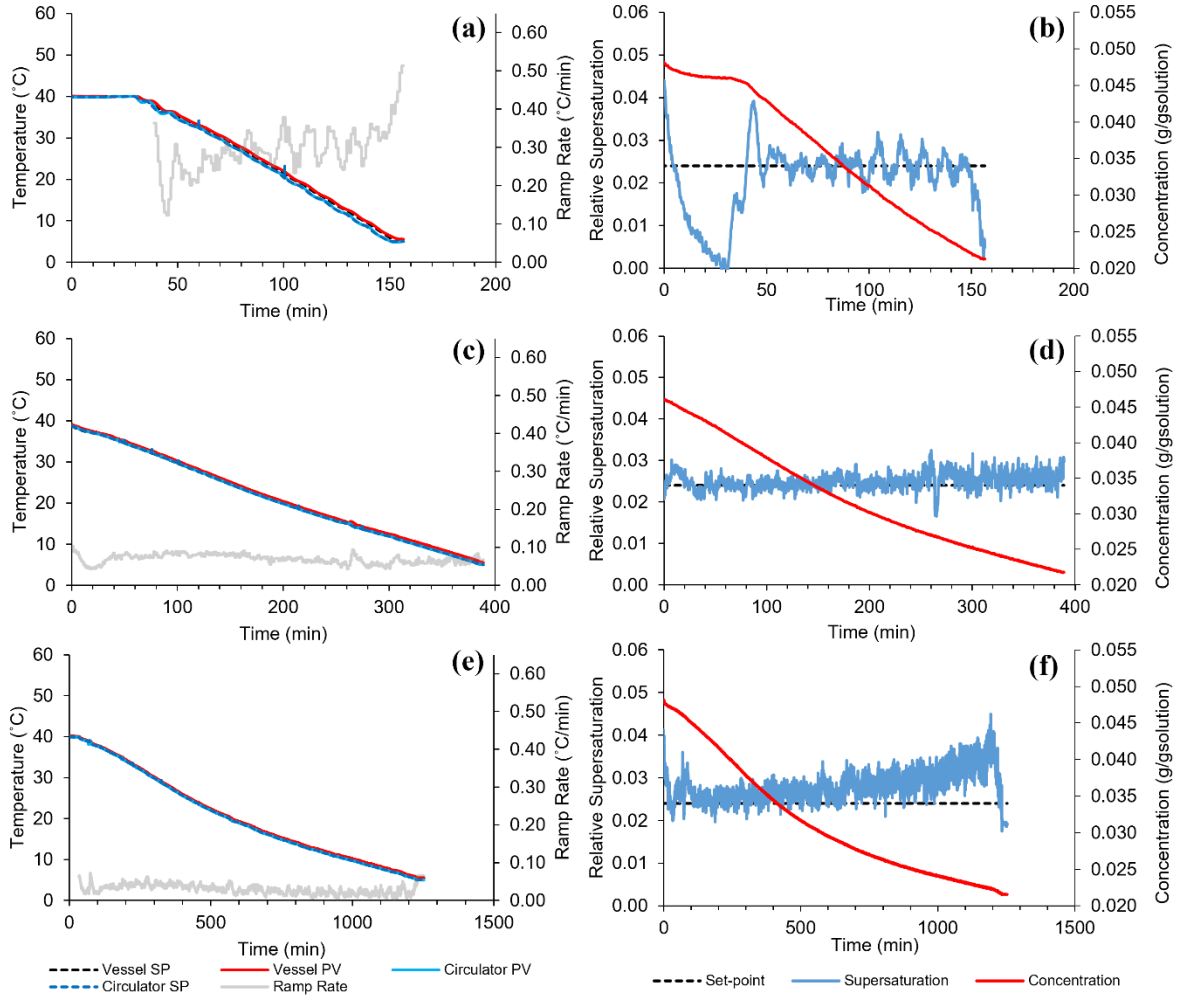


Figure 7.10. Temperature, cooling rate, concentration and supersaturation profiles measured during physical supersaturation MPC, for citric acid impurity concentrations of (a,b) 0.25 wt%, (c,d) 0.50 wt% and (e,f) 0.75 wt%.

Table 7.5. Mean and standard deviation of supersaturation and average cooling rates during physical supersaturation control experiments at varying citric acid impurity concentrations.

Impurity Concentration (wt%)	Supersaturation Mean	Supersaturation Std. Dev.	Average Cooling Rate (˚C/min)
0	0.0227	3.15×10^{-3}	0.31
0.25	0.0246	3.53×10^{-3}	0.29
0.50	0.0250	1.91×10^{-3}	0.07
0.75	0.0280	3.23×10^{-3}	0.03

The most noticeable change caused by the introduction of impurities was the extension of the batch time with increasing concentrations of citric acid. As much slower cooling rates were required to maintain the same level of supersaturation, the controlled cooling that typically took around 2 hours for the pure system extended beyond 6 hours with 0.50 wt% citric acid, and 21 hours with 0.75 wt%. In the latter case, even a cooling rate close to the minimum allowed by the supersaturation MPC block (-0.01 ˚C/min) was insufficient to prevent the supersaturation from climbing towards the end of the batch.

There was a general trend of increasing mean supersaturation with increasing impurity concentration. This was due to the heightened sensitivity of supersaturation generation when cooling too quickly, which declined more slowly when the cooling rate was lowered due to the inhibition of the crystal growth rate. Interestingly, the combination of this sensitivity and the low cooling rates below -0.1 ˚C/min gave very stable control during the batch with 0.50 wt% citric acid. With the supersaturation MPC block making very small adjustments to the cooling rate at each decision point, the standard deviation of the measured supersaturation was found to be the lowest out of all nine batches.

There was a minimal difference of 0.2 ˚C/min in the average cooling rate between the batches with pure hexamine and 0.25 wt% citric acid. However, there was greater instability in control of the supersaturation due to large moves being made in the cooling rate for the batch with impurities, especially when the MPC blocks had just been activated.

The particularly large overshoot of supersaturation shown at around 45 mins in Figure 7.10b was an issue resulting from problems with the internal control of the circulator. During the initial cooling to approach the supersaturation set-point from zero, the circulator bath was cooled much lower than the set-point requested by the temperature MPC block, creating a large disturbance in the temperature control that required heating of the bath to correct. This in turn impacted the vessel temperature, preventing it from accurately following the set-point profile calculated by the MPC strategy and resulting in a persistently poor stability of control throughout the rest of the batch. This highlighted the importance of ensuring the heat transfer dynamics of the crystallisation equipment remain similar to when it is used to train the temperature MPC block.

7.4. Conclusions

An MPC strategy for controlling the supersaturation during a batch cooling crystallisation was developed using the CrystalMV control structure in PharmaMV software. Using the gPROMS

process model of the hexamine-ethanol crystallisation system, simulated temperature and concentration data were used to train the MPC blocks within the structure to capture the crystallisation kinetics. The temperature MPC block was retrained using measurements from a physical 500 ml crystalliser and associated PAT tools in order to capture the real heat transfer dynamics between the circulator and vessel contents.

The digital twin of a 500 ml jacketed crystalliser was used to perform in-silico simulations of batch crystallisations with controlled supersaturation, at levels ranging between 0.012-0.036 relative supersaturation. After a small overshoot during the initial approach from zero up to the set-point, the MPC strategy performed excellently at controlling the supersaturation, although this was expected given the minimal disturbances in the simulation. Control stability was found to decrease at higher supersaturation set-points, resulting from the larger moves made to the cooling rate when faster cooling was required to prevent supersaturation from declining.

Performance of the MPC strategy during physical validation on a real 500 ml jacketed crystalliser was not as strong as in in-silico simulations, with respect to the mean supersaturations achieved. However, stability of the controlled supersaturation remained good with standard deviations between 2-4 $x10^{-3}$ compared to the in-silico range of 4-8 $x10^{-4}$. A clearer trend of decreasing stability with increasing supersaturation set-point was observed. The MPC strategy struggled to achieve the highest set-point of 0.036, likely due to the faster growth rates at higher supersaturations leading to a high cooling demand to compete with the consumption of hexamine from the solution.

The increase in surface area resulting from an increased mass of seeds was found to have a stronger impact on supersaturation control than that resulting from a decrease in the mean crystal size. In fact, although the mean supersaturation achieved using smaller seed crystals fell further below the set-point, the overall stability of control was much improved due to the smaller changes in cooling rate made by the supersaturation MPC block. It was hypothesised that given the greater amount of agglomeration typically occurring with finer seed material, the uneven surfaces of the crystals slowed down the growth rate compared to what one would expect. However, image analysis of the product crystals suggested that there was a similar degree of agglomeration between both the cases with an increased mass and decreased mean size. In the case where the surface area had been altered the most (both larger seed mass and smaller seed size), the mean supersaturation was noticeably the lowest between all batches with

a set-point of 0.024. This is because the MPC strategy struggled to cool fast enough to keep up with the heightened consumption of solute.

Increasing concentrations of citric acid impurity, known to inhibit the growth of hexamine, resulted in an increase in the mean supersaturation and extension of the overall batch time. This was due to the significantly lower cooling rates required to maintain the desired supersaturation set-point, as well as the heightened sensitivity of the system towards raising the supersaturation when the cooling rate was too high. During a few experiments but most noticeably the batch with 0.25 wt% citric acid present, the stability of control in the early stages of cooling was found to be negatively impacted by the internal control of the circulator.

Considering the overall performance of the supersaturation MPC strategy, it appeared able to respond well to disturbances caused by varying both the seeding conditions and growth rate kinetics. This supports the potential implementation of the MPC strategy at pilot scale, in order to more closely relate to an industrial crystallisation system. However, it is important to highlight a few additional adjustments or steps that would improve its performance. Regarding the equipment, it is vital that the heat transfer achievable by the circulator is adequate to reach the desired cooling rates, and that the internal PI control of the circulator during operation matches the behaviour during controller training. Regarding the control software, further tuning of the MPC blocks would likely improve the stability of control by restricting unnecessarily large changes in the cooling rate set-point, which should vary depending on the vicinity of the current supersaturation to the set-point value. Should these suggestions be taken into consideration, this MPC strategy would be able to control more complex, industrially relevant crystallisation systems with a good degree of accuracy.

8. Crystal Size End-Point Control Development and Performance with Varying Size Set-Point

8.1. Introduction

The MPC strategy for constant supersaturation control outlined in chapter 7 can be applied to achieve more complex control objectives functions once appropriate adjustments are made. Primarily, the product crystal size distribution is a desirable objective function due to its relationship with efficiency of downstream operations (e.g. filtration) and the performance properties of many fine chemicals (e.g. dissolution rate). The supersaturation control methodology involved selection of a single set-point value for the full duration of the batch cooling crystallisation process. However, this approach would not be viable for control of the crystal size given its dynamic nature, constantly changing throughout the batch with the potential to follow different routes towards a given product CSD.

In this chapter, a second type of controller known as a batch EPC block was added to the MPC structure described in chapter 7. The EPC block was trained to capture the effect of supersaturation on the D_{50} of the number-based CSD, using simulated data from the digital twin. The controller could then generate a supersaturation profile that would result in the desired D_{50} end-point being reached at a specified final temperature. Similarly to the procedure for the supersaturation MPC strategy, the crystal size EPC strategy was first tested in-silico using the digital twin of a 500 ml jacketed crystalliser to simulate the required measured properties. Simulated control batches for a range of crystal size end-points were validated using a physical 500 ml crystallisation setup at CMAC before an assessment was made of the EPC strategy's performance, based on the accuracy of the product crystal size.

8.2. Materials and methodology

8.2.1. Batch end-point control (EPC)

The EPC strategy was developed using features available in PharmaMV v8.1.1445 control software. In contrast to the MPC block, measurements of the controlled variables associated with an EPC block are not required during a controlled batch process. Rather than operating based on the difference between measured and set-point values, it is instead trained to operate within a fixed window of conditions. Starting from an average trajectory of the manipulated variable, the EPC block adjusts this profile depending on where the desired end-point falls within the range of achievable end-points it has been provided with. In addition to outputting a new set-point for the manipulated variable at each decision point, the block also updates its

prediction of the expected end-point, which approaches the set-point value as the batch progresses.

In the case of the crystal size EPC block, the manipulated variable used as the controller output was the supersaturation set-point, which would be used as the input for the supersaturation MPC block. As precise control over the full product CSD would be a difficult challenge, the control objective was simplified to be a target D_{50} of the number-based CSD. Should strong performance of the EPC strategy be observed, other controlled variables could be introduced in the objective function alongside this (e.g. D_{10}, D_{90} and span of the CSD) to increase the precision of the final CSD.

The EPC block required a fixed number of decision points to be specified for its operation. Given that control of the CSD would implicitly require different cooling rates between batches to achieve the desired nucleation and growth rates, it was highly unlikely each batch would consist of the same number of fixed time intervals. To resolve this, each batch crystallisation was run between the same temperature range (40 to 5 ˚C) such that a fixed number of temperature intervals of 0.5 ˚C could be used as decision points for the controller to update the supersaturation set-point. The EPC block would only move to the next decision point if the associated measurement was an increasing value. Therefore, the inverse of the measured vessel temperature was used as a measured signal for this purpose.

Given that the EPC block could only operate within a set window of experimental conditions, it was important to consider the limitations of the controller. The seed CSD would need to match closely to the one used during controller training. Otherwise, the variation in surface area would impact the growth of the overall mass of suspended solids, and therefore alter the relationship between supersaturation and the evolution of the D_{50} over time. For this reason, only one particular seed CSD was implemented throughout the simulations and validation experiments presented in this chapter.

8.2.2. Training of crystal size EPC block

The crystal size EPC block was trained using the gPROMS process model to simulate a series of seeded hexamine-ethanol batch crystallisations. The initial conditions for the training simulations are summarised in Table 8.1, which were identical to those applied during the in-silico testing of the supersaturation MPC strategy in Table 7.1. As before, the parameters describing the seed CSD were selected based on those produced by wet milling (mean size of

112.2 μm). Both the secondary nucleation and growth models were activated in gPROMS for all simulations outlined in this chapter.

Table 8.1. Process conditions applied to the gPROMS/PharmaMV digital twin for crystal size EPC training and performance testing.

Crystalliser volume (ml)	500
Volume D_{50} of seed CSD (μm)	110
Standard deviation of seed CSD (μm)	25
Initial seed mass (g)	1
Initial temperature (ºC)	40
End temperature (ºC)	5

Measurements of the number-based D_{50} were required from the process model to generate the training data. However, the solution of the PBM in gPROMS involves calculations related to the volume-based CSD. Therefore, a conversion between the volume-based and number-based distributions was required, which was applied using custom calculations within the batch crystallisation flowsheet in gPROMS software. These calculations were based on the assumption that particles within a particular size class were spherical, with a diameter equal to the midpoint of that particular size class. The full calculations for this conversion have been outlined in Appendix A.

Six seeded linear cooling crystallisations were simulated using the digital twin. At the start of each simulation, the solution was held at 40 ˚C for 30 mins before cooling at the chosen linear cooling rate down to 5 ˚C. Cooling rates between 0.2-0.7 ˚C were applied in 0.1 ˚C intervals; this range was selected as not to exceed the maximum achievable cooling power of the circulator available in the laboratory for the physical validation experiments.

8.2.3. Methodology for in-silico crystal size EPC

After incorporating the trained EPC block into the existing MPC strategy, the digital twin was used to perform in-silico testing of the crystal size controller. The gPROMS process model provided simulated measurements of temperature, concentration and number-based D_{50}, as described by the diagram in Figure 8.1 that shows the adaptations made to the data flow structure from Figure 7.2.

During a controlled batch crystallisation simulation, both MPC blocks and the crystal size EPC block were activated simultaneously as soon as cooling began. The EPC block initially followed the average supersaturation trajectory calculated from the mean of the six training

profiles. In order to give the controller time to develop a projected D_{50} end-point and prevent communication issues between gPROMS and PharmaMV, this trajectory was followed for two decision points (i.e. 1 ˚C temperature drop) before the EPC block was allowed to make adjustments to reach the requested size set-point. As in supersaturation control experiments, the temperature and supersaturation MPC blocks continued to update their output signals at time intervals of 15 and 60 seconds respectively. Data was transferred through this loop every 5 seconds, stopping the simulation automatically at the set end temperature of 5 ˚C.

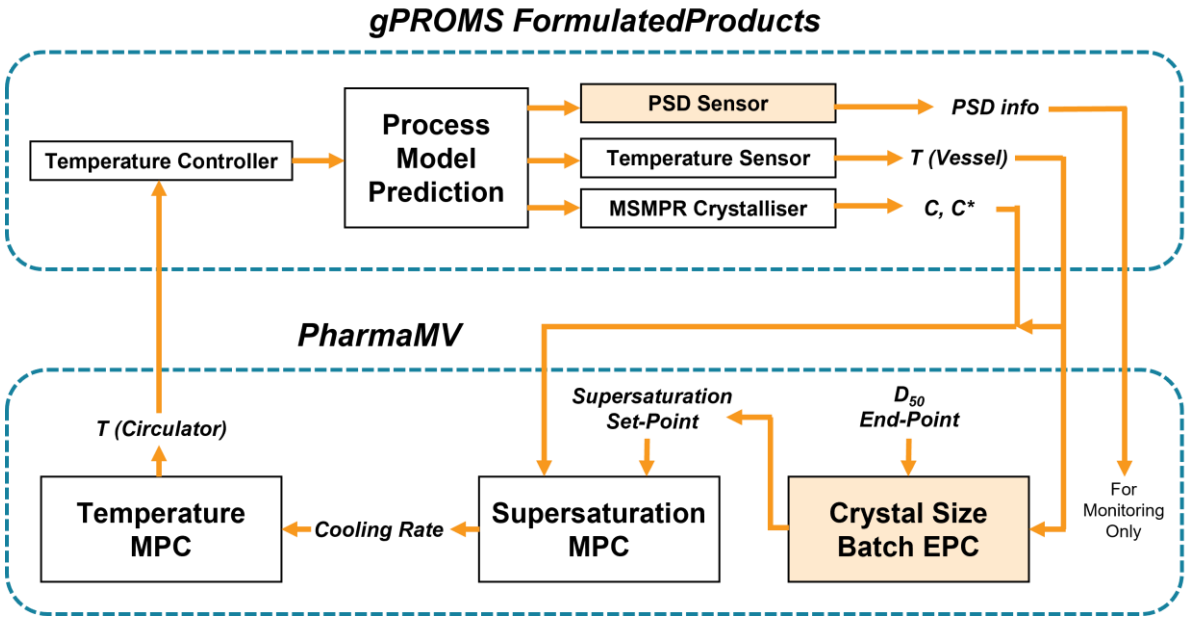


Figure 8.1. Data flow within the gPROMS/PharmaMV digital twin during a crystal size control simulation.

To test the adapted control strategy, three batch crystallisations were simulated with control of the final crystal size. The initial conditions listed in Table 8.1 for the training simulations were also applied to these batches. In each case, the temperature was held at the initial value for 30 mins before switching on all control blocks, which turned off automatically when the vessel temperature reached 5 ˚C. Set-points of 140, 150 and 160 µm were chosen for the D_{50} end-point, which were based on the smallest and largest sizes resulting from the training batches (see section 8.3.1).

8.2.4. Methodology for physical validation of crystal size EPC

The crystal size control strategy developed in PharmaMV was implemented into the 500 ml jacketed crystalliser for validation at CMAC, which was also used previously for validation of the supersaturation MPC strategy. Figure 8.2 illustrates how the data flow schematic in Figure 8.1 was adapted for the physical crystallisation setup.

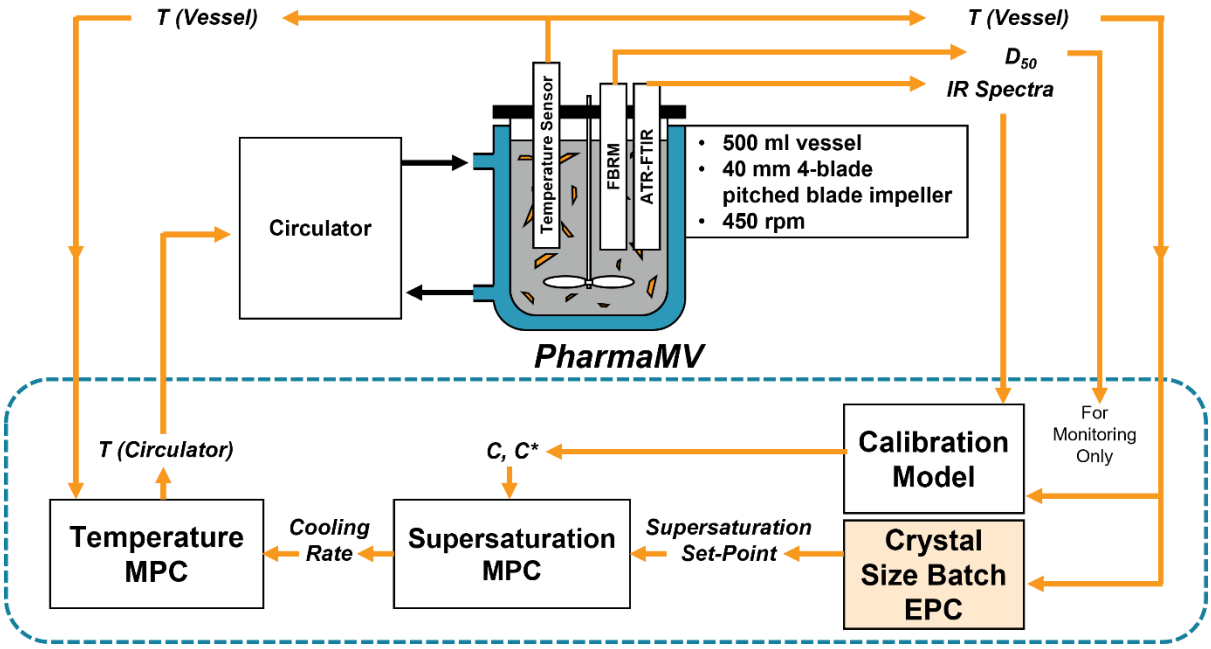


Figure 8.2. Data flow within the MPC strategy during a physical batch cooling crystallisation with crystal size control.

The FBRM probe was included in these experiments to monitor the number D_{50}. Measurements were taken under the assumption that, for the relatively evenly shaped hexamine crystals, the CLD was approximately equivalent to the CSD without additional processing of the measured data. The PVM probe was also used to capture images of the suspended crystals during the batch experiments, however data from this probe was not transferred to PharmaMV and used in the control strategy.

At the beginning of each physical control run, hexamine was added to 400 g of ethanol to make up an initial concentration of 0.0507 g/g solvent. The solution was heated to 50 ˚C and held for 60 mins to ensure all material had dissolved. The solution was then cooled at the maximum rate achievable by the circulator (-0.6 ˚C/min) down to 40 ˚C, where 1 g of the dry seed crystals were added. Following this, the crystalliser was held at constant temperature for 30 mins to allow the system to equilibrate, after which the MPC and EPC blocks in PharmaMV were switched on. Once the final temperature of 5 ˚C had been reached, the product crystals were removed and filtered under vacuum for analysis.

Each of the three in-silico simulations were replicated using the physical crystallisation setup, using set-points of 140, 150 and 160 µm for the number D_{50} end-point. Two additional experiments were carried out to assess the control strategy's ability to follow a prescribed supersaturation profile, which would allow good repeatability between batches. The average supersaturation trajectory was replaced with the training profiles for the 0.7 and 0.2 ˚C/min

linear cooling batches, which were expected to result in D_{50} end-points around 140 and 160 µm respectively. To prevent the EPC block from making adjustments to the supersaturation trajectory, the 'maximum move' value (the maximum change allowed from the programmed supersaturation profile at each decision point) was set to 0.

8.3. Results and discussion

8.3.1. Data generated for training crystal size EPC block

Figure 8.3 shows the supersaturation and number-based D_{50} profiles resulting from gPROMS simulations used to train the crystal size EPC block. Faster cooling rates were consistently found to lead to higher supersaturations for the whole duration of the batch, as expected. The more important relationship observed was that higher supersaturations were generating a smaller crystal size. Overall, the final D_{50} values fell between a range of around 140-162 µm.

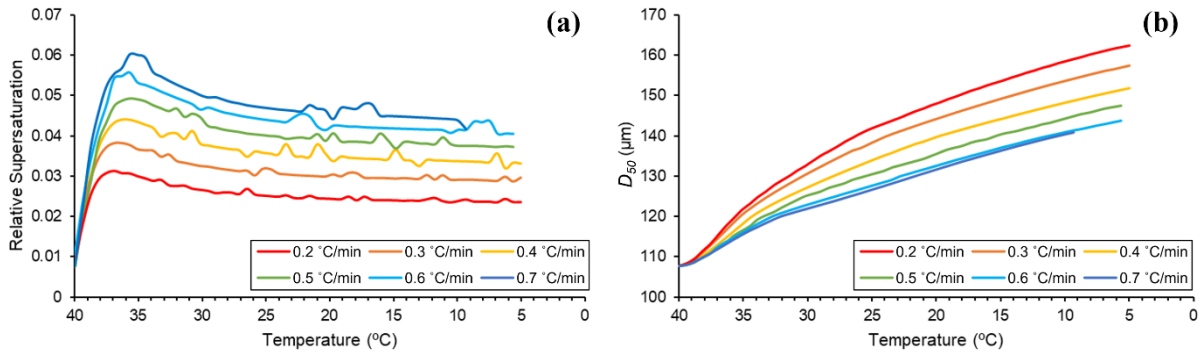


Figure 8.3. a) Supersaturation and b) number-based D_{50} profiles generated using the digital twin for training of the crystal size EPC block.

It was clear from this relationship that nucleation was the key mechanism in driving a change in the final crystal size. Given that a secondary nucleation expression was active in the gPROMS process model, it was predicted that higher supersaturations would result in greater nucleation rates, hence a smaller overall crystal size. As nucleation has been observed to be practically negligible during many of the experiments outlined in previous chapters, it was hypothesised that the EPC strategy may struggle to reach sizes towards the lower end of the predicted range when implemented into the physical crystallisation setup.

It was predicted by the process model that the peak supersaturation levels at the start of each batch would range between approximately 0.03-0.06. Although the cooling rates chosen should fall within the feasible operating range of the physical circulator, the system had previously struggled to reach a supersaturation of 0.036 (see Figure 7.7f). This suggested that attempting to follow a supersaturation profile to achieve the smaller size set-points could pose a difficult challenge for the EPC strategy.

8.3.2. Performance of in-silico crystal size EPC

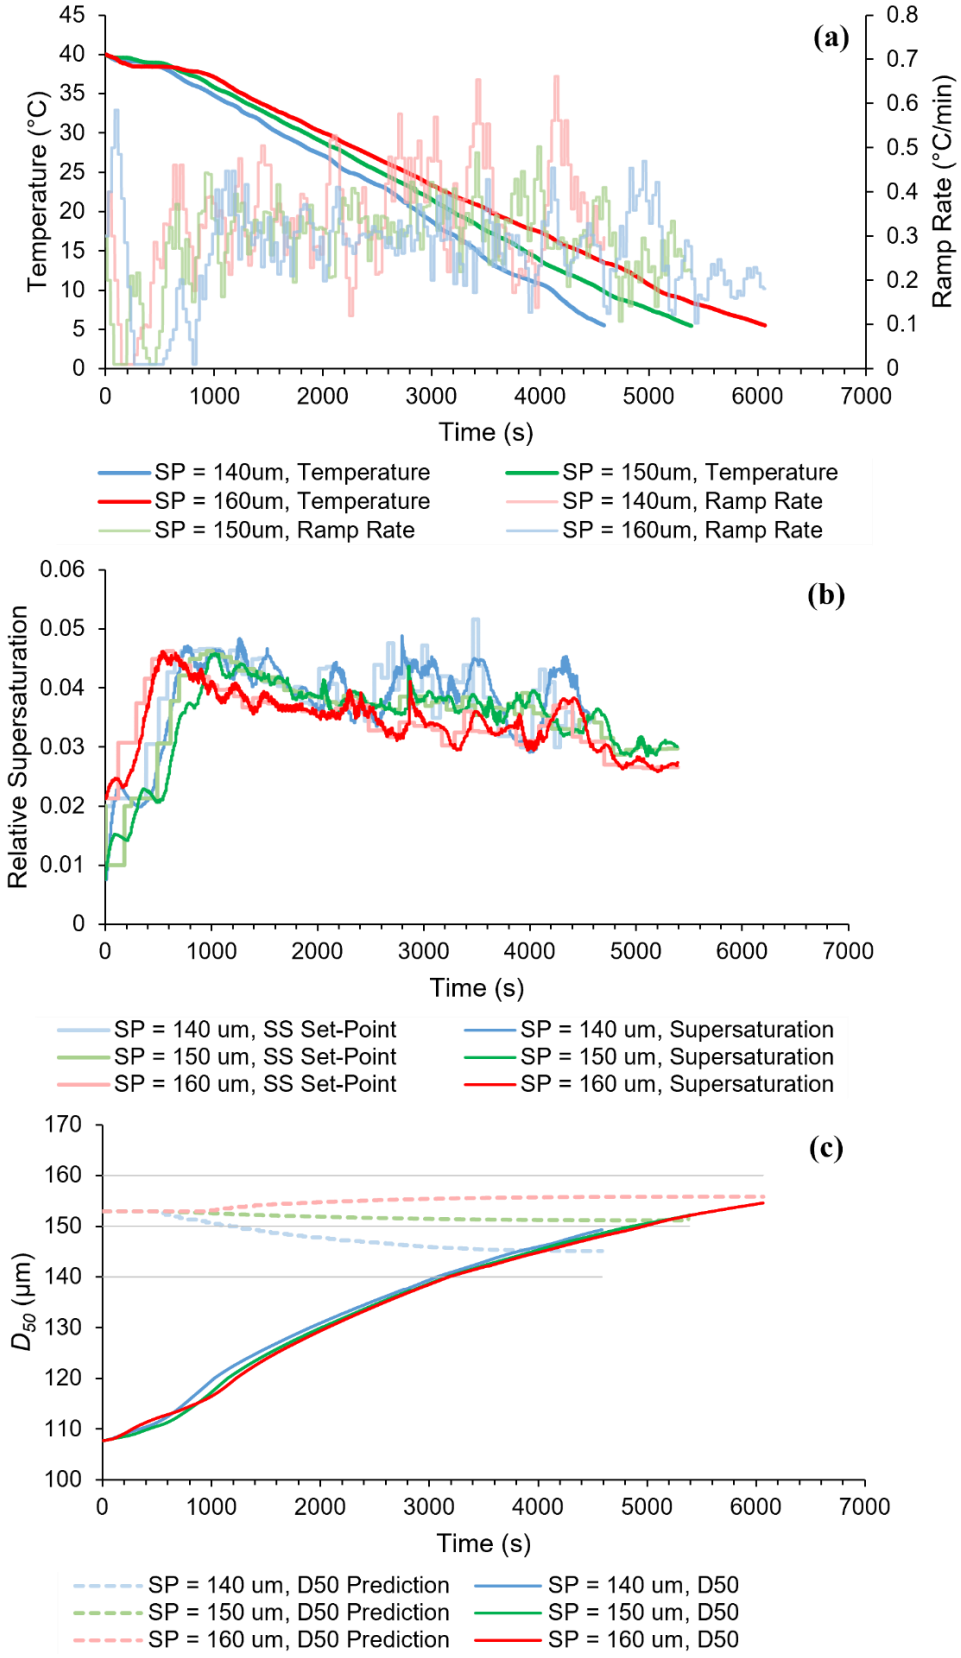


Figure 8.4. a) Temperature and cooling (ramp) rate, b) supersaturation and c) number D_{50} profiles generated during in-silico crystal size EPC, for product D_{50} set-points of 140, 150 and 160 µm.

Table 8.2. Predicted vs. achieved D_{50} of the product number-based CSD resulting from in-silico crystal size EPC simulations.

D_{50} Set-Point (μm)	Predicted End D_{50} (μm)	Achieved End D_{50} (μm)
140	145.1	149.3
150	151.2	152.1
160	155.8	154.6

The cooling, supersaturation and number-based D_{50} profiles resulting from the in-silico crystal size control simulations at three different size set-points are presented in Figure 8.4. The final D_{50} values predicted by the EPC block and achieved at the end of the simulated batches have been summarised in Table 8.2 for comparison.

The predicted end-points presented as dashed lines in Figure 8.4c show how the crystal size EPC block typically operates. At the start of each batch, the expected final D_{50} was always a value of 152.9 μm, based on the average supersaturation trajectory programmed into the control block. If this value was higher than the set-point, the controller made adjustments to its programmed supersaturation trajectory to have the predicted end-point approach this value, i.e. the supersaturation would be increased. Conversely, if the prediction was lower than the set-point value, supersaturation would be decreased. It should be noted that these predictions were ultimately based on the supersaturation profile constructed by the EPC block, ignoring the actual progression of the D_{50} during the batch.

The range of predicted sizes was considerably narrower than that suggested by the training data, with the EPC block expecting to only achieve sizes between 145-156 μm when following its adjusted supersaturation profiles. The final crystal sizes simulated by the process model had an even narrower range between 149.3-154.6 μm, with the worst case being around 9 μm away from the desired set-point. Comparison between the supersaturation profiles in Figure 8.4b suggested that there was little deviation from the initial average trajectory, significant changes only occurring after around 2500 s: 26-33 out of a total 70 decision points into the batch.

Given that the initial peak in supersaturation remained at a similar level for all three batches, it is likely a key step in the crystallisation process for governing the final crystal size that was missed. For example, since the relative supersaturation had not been raised to around 0.06 as was experienced during the 0.7 ˚C/min linear cooling, there was a reduced nucleation rate and therefore smaller population of crystals that continued to grow throughout that training batch, resulting in a larger final size than expected. It was therefore implied that the EPC block needed

to be less restrictive of its changes made to the original supersaturation trajectory at each decision point.

The EPC strategy performed well at controlling to the set-point of 150 μm, which was already close to the predicted value from the average supersaturation trajectory. When only having to make minimal changes to this profile, the final crystal size achieved was just 2 μm from the intended value. This result showed that following the optimal supersaturation profile precisely would lead to accurate control of the crystal size.

8.3.3. Performance of physical crystal size EPC

The cooling, supersaturation and number-based D_{50} profiles resulting from physical crystal size control experiments performed at different size set-points are presented in Figure 8.5. The final D_{50} values predicted by the EPC block and measured at the end of the batches have been summarised in Table 8.3 for comparison.

The cooling profiles measured during the physical experiments showed even less variation than those generated during the in-silico simulations, once again a result of the similarity between each calculated supersaturation profile. This was particularly evident in the case of the smallest D_{50} set-point of 140 μm; the supersaturation followed an almost identical profile to that of the batch with a set-point of 150 μm, despite the supersaturation set-point requested to the supersaturation MPC block being evidently higher after 2200 s. As was hypothesised based on the results of supersaturation control experiments in section 7.3.3, the control strategy failed to increase the supersaturation above a limit of approximately 0.036. Especially since this value was requested later into the batch, where consumption of hexamine from solution occurred at a faster rate due to the increase in available surface area from the initial seed material. As a result, the final crystal size was over 20 μm above the desired end-point, despite the crystal size EPC block predicting a size within 5 μm of the set-point.

Table 8.3. Predicted vs. achieved D_{50} of the product number-based CSD resulting from physical crystal size EPC experiments.

D_{50} Set-Point (μm)	Predicted End D_{50} (μm)	Achieved End D_{50} (μm)
140	144.5	162.7
150	151.0	165.0
160	157.0	159.4

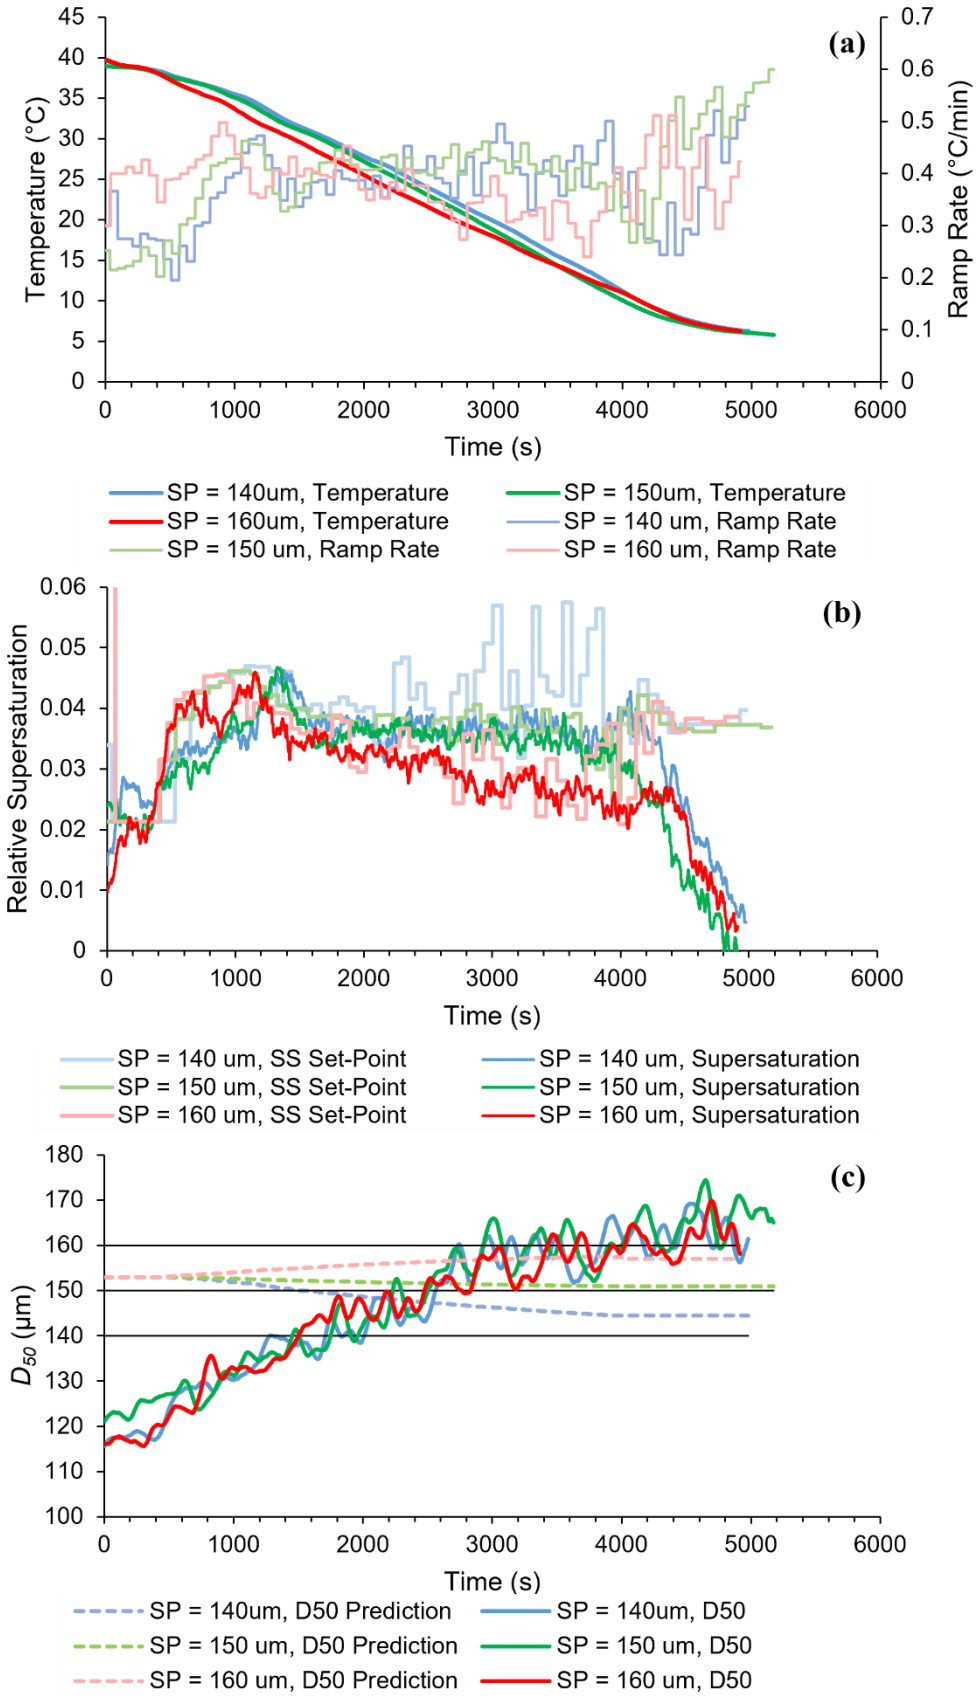


Figure 8.5. a) Vessel temperature and cooling (ramp) rate, b) supersaturation and c) number D_{50} profiles measured during physical crystal size EPC, for product D_{50} set-points of 140, 150 and 160 µm.

Although the EPC strategy had achieved good control around a set-point of 150 μm in the in-silico simulations (section 8.3.2), this was instead exceeded for the corresponding physical experiment by 15 μm. Overall, the final D_{50} values measured at the end of every batch were very similar at around 162±3 μm. As nucleation was identified as the main driving mechanism for crystal size control, this suggested that the lack of secondary nucleation occurring in the physical hexamine-ethanol crystallisation system may have prevented any precise control over the final crystal size. If crystal growth was the only mechanism occurring during these experiments, as with many of the seeded batch crystallisations described in chapter 4, then the same final crystal size would inevitably be reached given that each experiment started with the same solution concentration, seed mass and seed CSD.

Figure 8.6, which shows PVM images taken at various stages throughout the batch with a set-point of 140 μm, provides further evidence to support this. During the holding period after addition of the seeds, the relatively small agglomerated particles remained fairly unchanged for the full 30 mins (Figures 8.6a and b). As the batch progressed into the controlled cooling period, the particles most consistently observed were either large single crystals with hexamine's expected cubic morphology (Figures 8.6e and f), or agglomerates made up of a few of these larger crystals (Figure 8.6d). No evidence of any significant secondary nucleation was observed, despite the in-silico simulations predicting that the nucleation rate should have been the highest out of all three batches.

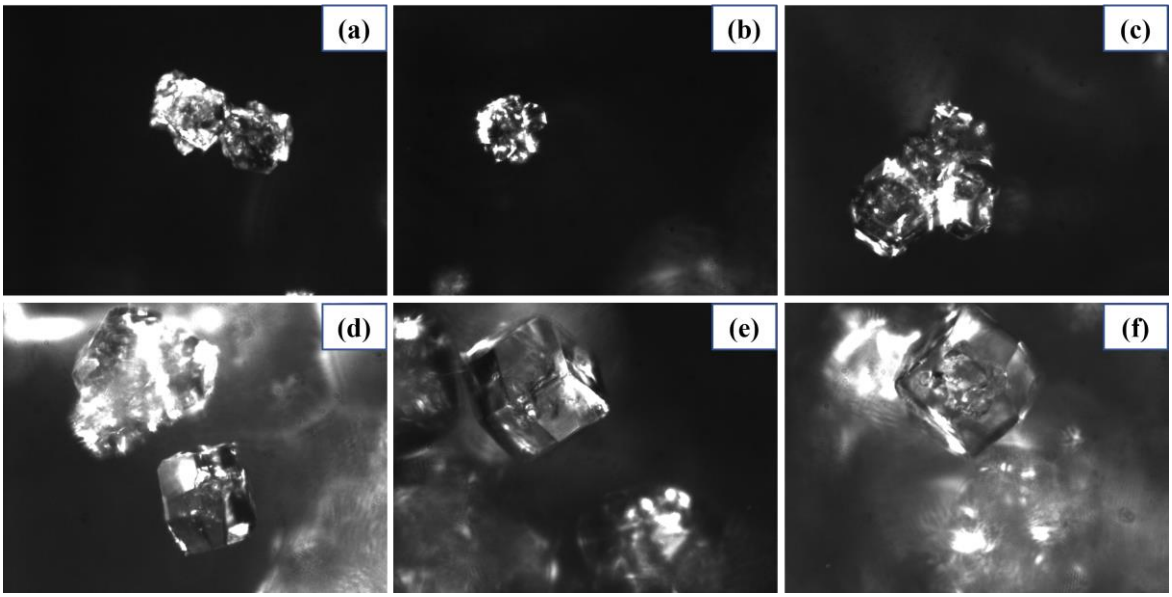


Figure 8.6. PVM images of hexamine crystals during physical crystal size EPC for a D_{50} set-point of 140 μm. a) 5 mins after seeding, b) 30 mins after seeding, c) 28 mins into controlled cooling, d) 40 mins into controlled cooling, e) 53 mins into controlled cooling and f) 66 mins into controlled cooling.

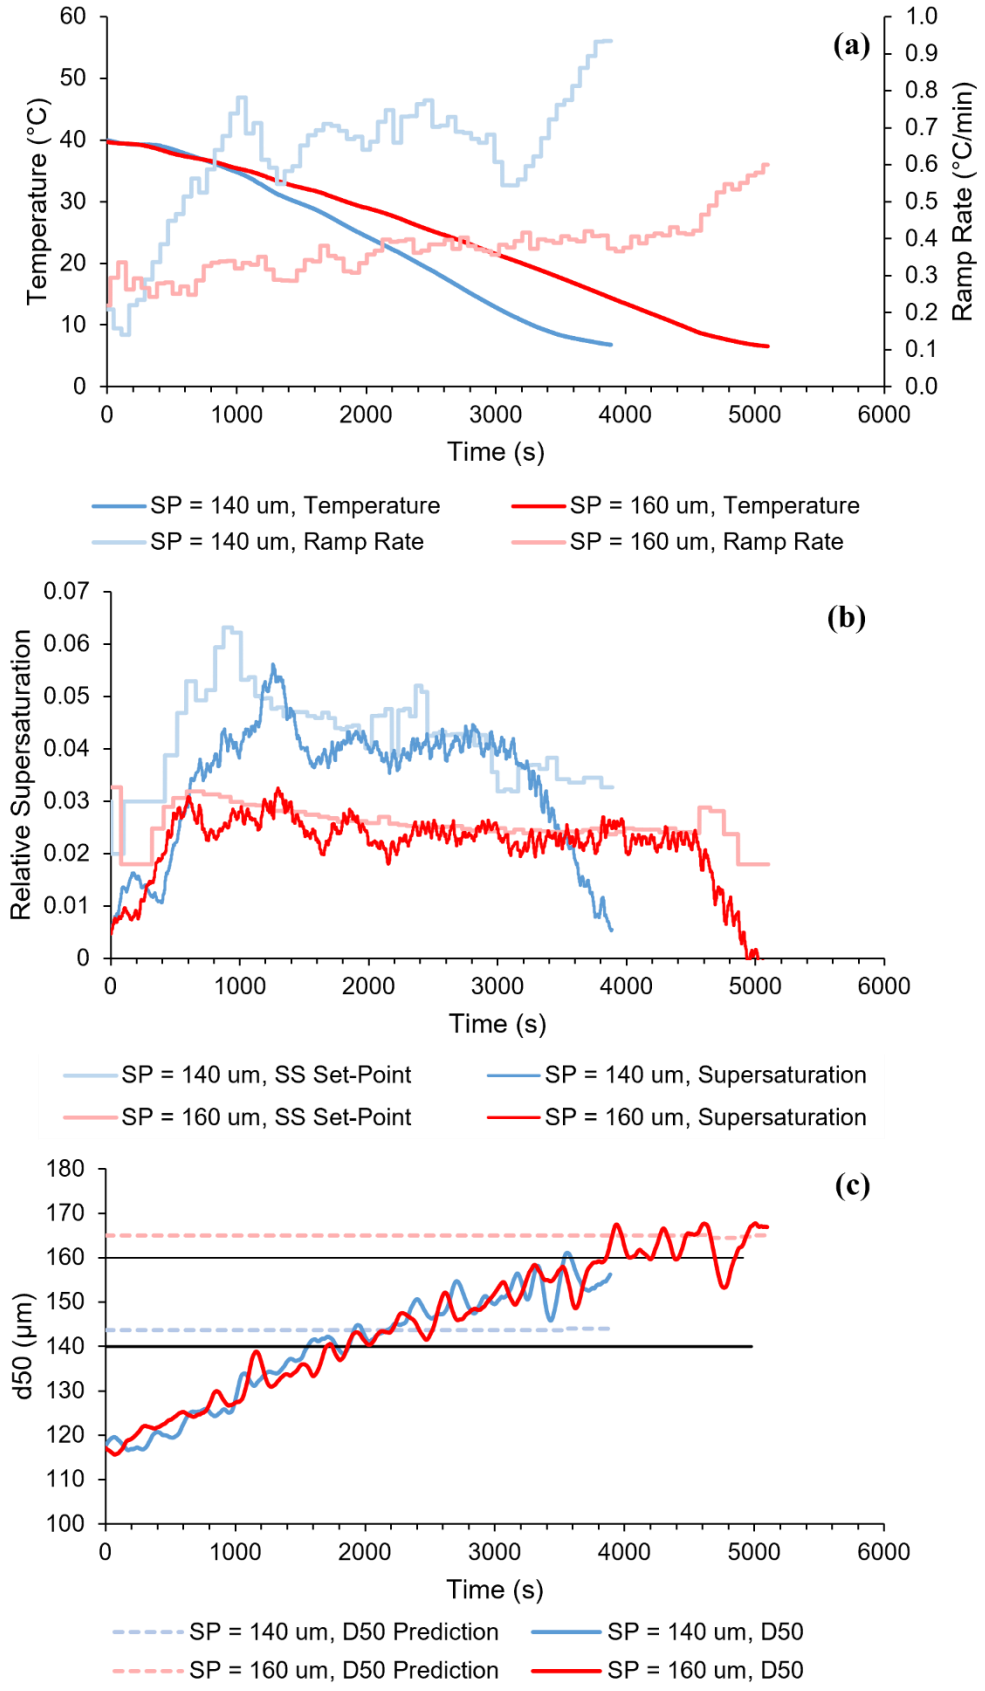


Figure 8.7. a) Vessel temperature and cooling (ramp) rate, b) supersaturation and c) number D_{50} profiles measured during physical crystal size EPC, following fixed supersaturation profiles expected to achieve D_{50} set-points of 140 µm and 160 µm.

Table 8.4. Predicted vs. achieved D_{50} of the product number-based CSD resulting from physical crystal size EPC experiments, following fixed supersaturation profiles.

D_{50} Set-Point (μm)	Predicted End D_{50} (μm)	Achieved End D_{50} (μm)
140 (Fixed Supersaturation)	144.0	160.2
160 (Fixed Supersaturation)	165.0	166.5

The cooling, supersaturation and number-based D_{50} profiles resulting from physical crystal size control experiments following prescribed supersaturation trajectories are presented in Figure 8.7. The final D_{50} values predicted by the EPC block and measured at the end of the batches have been summarised in Table 8.4 for comparison.

Given that fixed supersaturation profiles were being followed, the predicted D_{50} remained constant throughout both batches, as shown by the dashed lines in Figure 8.7c. In contrast to the previous set of experiments, two distinctly different cooling and supersaturation profiles were generated via the approach of using fixed set-point trajectories within the crystal size EPC block. Some difficulty was still observed where higher supersaturations were required, but there was noticeable improvement over the previously poor manipulation of the supersaturation set-point by the EPC block.

Unfortunately, despite the final D_{50} achieved for the higher set-point falling close to the EPC block's prediction, the one measured at the end of the batch with a smaller D_{50} set-point was once again exceeded by over 15 μm. The implication of this when combined with the results of the previous three batches is that the hexamine-ethanol crystallisation system was unsuitable for assessing the performance of the crystal size EPC strategy. The major reason for this was the mismatch between the crystallisation behaviour predicted by the literature nucleation kinetics and the behaviour observed in this study.

Considering the trend in the D_{50} end-points achieved during in-silico simulations, although the accuracy stood to be improved, the EPC block was at least able to increase/decrease the crystal size in relation to its initial prediction. It is therefore believed that if an accurate gPROMS process model can be developed for a crystallisation system with more significant secondary nucleation present, that the crystal size EPC strategy may exhibit stronger performance.

8.4. Conclusions

A model-trained system for controlling the crystal size at the end-point of a batch cooling crystallisation was implemented into the existing MPC strategy. Using the gPROMS model of the hexamine-ethanol crystallisation system, the supersaturation and number-based D_{50} of the

CSD were simulated for a series of linear cooling crystallisations and applied to train an EPC block within PharmaMV software.

The digital twin of a 500 ml jacketed crystalliser was used to perform in-silico simulations of batch crystallisations, controlling the crystal size to achieve D_{50} end-points between 140-160 μm. The crystal size EPC strategy only achieved sizes of 149.3-154.6 μm when asked to control to these extremes. The controller displayed the ability to manipulate the supersaturation profile appropriately in the second half of the batch. However, its tendency to follow the average supersaturation trajectory closely in the early stages meant that a key step in the crystallisation process was missed, being the degree of secondary nucleation required to generate the necessary population of crystals.

Performance of the EPC strategy was poorer during physical validation on a real 500 ml jacketed crystalliser, where the D_{50} in each batch tended towards similar values of around 162±3 μm. This was an especially large error for the lowest set-point of 140 μm, where the supersaturation MPC portion of the control strategy failed to increase the supersaturation beyond 0.036 for the majority of the batch time. PVM images taken during this experiment confirmed that there was negligible secondary nucleation occurring in the system, which may have been due to the low supersaturation. However, when the crystal size EPC block was forced to follow a fixed supersaturation profile that was expected to reach a D_{50} end-point of 144 μm, this size was still exceeded by over 15 μm.

It was determined that the mismatch between the amount of secondary nucleation between the gPROMS process model and real crystallisation behaviour made the hexamine-ethanol crystallisation system unsuitable for assessing the performance of the crystal size EPC strategy. Without generation of new particles during the batch, a 'growth-only' crystallisation process would always result in the same product crystal size. A different crystallisation system with more prevalent secondary nucleation and available, accurate rate kinetics describing this could be investigated to assess the capability of this control methodology further. Alternatively, a way of inducing nucleation of hexamine, such as increased agitation for secondary nucleation or operating beyond the metastable zone for primary nucleation, could allow continued investigation of this crystallisation system.

9. Conclusions and Future Work

9.1. Conclusions

This thesis outlines the methodology for development of MPC strategies for supersaturation and crystal size control during batch cooling crystallisation, with its performance assessed and validated by both in-silico simulation and physical experimentation.

The hexamine-ethanol solute-solvent system was characterised through a series of cooling crystallisations in a 500 ml crystalliser with various process conditions. Despite having a narrow MSZW, little to no nucleation occurred throughout most of the seeded crystallisation experiments carried out in this study. Only the mechanisms of crystal growth and agglomeration were initially deemed important for the purposes of modelling the crystallisation process, the significance of the latter lessened when using a sufficiently large seed size. Experiments were designed such that the cooling rate, seed mass/size and seeding temperature could be varied across a small number of experiments. This minimised the time required to gather sufficient solution concentration and CSD measurements to support the development of a crystallisation process model that could be applied across a wide range of operating conditions.

Using measured physical property data for the hexamine-ethanol system alongside solubility and crystal growth/nucleation rate kinetic parameters from literature, a 1D PBM was parameterised in gPROMS FormulatedProducts software that acted as a digital twin of the batch cooling crystallisation process. Comparing the results of simulated seeded crystallisations to those from identical physical experiments, the model was found to slightly underpredict the evolution of supersaturation over time. This error was reduced by re-estimating the growth rate kinetics using experimental data from ‘growth-only’ crystallisations. However, this came at the cost of worsening predictions of the product crystal size, leading to the re-estimated kinetics being rejected. This was attributed to the large difference in number of data points gathered between solution concentration and CSD over time, which forced the fitting of kinetic parameters to focus too strongly on the former.

Out of the three compounds tested for their suitability as a crystal growth rate inhibitor, only citric acid with its many interactive functional groups (carboxylic acid) showed significant changes to the growth rate kinetics of the hexamine-ethanol system. Concentrations above 2 wt% inhibited crystal growth to the extent that high enough supersaturations were reached to exceed the MSZW. This resulted in excessive primary nucleation that not only led to

unpredictable agglomeration of the fines that were generated, but also fouling around the ATR-FTIR probe that prevented in-situ concentration measurement. Lowering the impurity concentration resolved these problems, introducing a new effect that was theorised to reduce the presence of solvent inclusions as agglomerates grew more slowly towards the expected single crystal morphology. It was noted that impurity effects such as these should ideally be quantified for process modelling of more complex industrial crystallisation systems, which can often include a variety of impurities.

Simulated temperature and concentration data from the gPROMS process model was used to train the MPC structure in PharmaMV. Although the supersaturation controller was unchanged between in-silico simulations and physical experimentation, the temperature MPC block needed re-training with measurements from a real 500 ml crystallisation setup in order to accurately capture the heat transfer dynamics. The power of the available heater/chiller used to drive the cooling rate was a source of various challenges; the maximum achievable cooling rate (~0.6 ˚C/min) put limitations on the supersaturations that could be reached and controlled steadily (σ ~ 0.036). Changes to the system that resulted in supersaturation being consumed faster, such as increasing the seed surface area through higher loadings/smaller sizes, increased the demand on the heater/chiller and led to poorer control of the supersaturation level. Conversely, factors that slowed the consumption of supersaturation, such as impurities reducing the crystal growth rate, led to a more stable supersaturation due to the lower cooling rates required to maintain a given supersaturation level. Overall, the supersaturation MPC strategy responded well to these disturbances to the crystallisation behaviour, maintaining a similar level of stability regardless of the conditions (standard deviations of supersaturation between 2-4 $x10^{-3}$). It was noted that some minor tuning of the MPC blocks within PharamMV software could improve stability, by restricting the size of changes to the cooling rate set-point depending on the current supersaturation level.

To train an EPC block that could manipulate the supersaturation profile and achieve a target D_{50} of the product CSD, further in-silico simulations were performed using the gPROMS process model to generate the required supersaturation and number-based D_{50} data. It was deduced that the control strategy must vary the amount of nucleation occurring during the cooling crystallisation process to achieve a different product CSD. Therefore, the literature secondary nucleation kinetics for the hexamine-ethanol system were included in the gPROMS process model when training the EPC block.

Simulations using the digital twin predicted that, within a feasible cooling rate range of 0.2-0.7 ˚C/min, the D_{50} of the product CSD could be manipulated within a range of 144-165 μm. In-silico tests of the crystal size EPC strategy struggled to achieve target D_{50} values at the extremes of this range, on account of its tendency to follow the starting supersaturation trajectory too closely. However, it was demonstrated that the EPC strategy still had the capability to vary the final D_{50} within the range of 149-155 μm.

There was an even smaller variation in the final D_{50} achieved at the end of physical batch cooling crystallisation experiments (162±3 μm), even when the EPC block was implemented to follow two distinct supersaturation profiles where the digital twin predicted a much larger difference in size (144 and 165 μm). This was a consequence of the secondary nucleation being negligible in all controlled crystallisation experiments, which did not match the nucleation rate kinetics implemented in the gPROMS process model. This suggested that, without altering the experimental conditions to induce secondary nucleation in some way, the hexamine-ethanol system was not suitable for adequately testing the performance of the crystal size control strategy.

Ultimately, the answer to the research question posed in section 1.2 is as such: for an MPC strategy built using commercial tools to accurately predict and control the CSD, there is an essential requirement to train the predictive controllers using a model that captures the real nucleation behaviour of the solute-solvent system as accurately as possible. One of the greatest challenges to batch crystallisation MPC identified in this work is the demand on the circulator. This would be particularly challenging at larger scales, where heat transfer is slower and there would be a much greater power requirement to minimise time lag within the control system.

9.2. Future work

The MPC strategy shows promise for a single-component crystallisation system; the complexity of real industrial crystallisation systems, however, needs to be accounted for before implementing the same strategy at a larger scale. The simple 1D PBM that incorporated only crystal growth and nucleation could be adjusted/improved in the following ways:

- For modelling of accurate crystal size and shape evolution over time for crystals with non-equant growth rates (e.g. needles), a 2D or morphological PBM would be preferred over the 1D PBM.

- Heat transfer and fluid dynamics that are more representative of the conditions experiences in a real crystalliser could be introduced, through means such as CFD modelling.

For the performance of the crystal size control strategy to be fully assessed using the hexamine-ethanol crystallisation system as the focus of study, a means of altering the particle count during the batch crystallisation process must be determined. Working within the metastable zone where crystal growth is the only mechanism would require a variation in the seed CSD to change that of the final product. As the EPC strategy offers no means of predicting the required seed CSD for a particular end-point target, this would not be a reliable method of controlling crystal size.

Inducing nucleation in the system would be another possible route. Increasing the agitation rate or altering the agitator geometry could increase the generation of nuclei by introducing more energy into the system. Alternatively, if the cooling power of the chiller could be increased to raise the maximum achievable cooling rate, the supersaturation could be pushed beyond the MSZW to induce primary nucleation, though this would pose a challenge to control reliably.

If another solute-solvent system is to be studied instead, it should ideally meet the following criteria:

- The crystallisation system should have known secondary nucleation rate kinetics, which have been quantified in a way that is applicable to process modelling.
- The secondary nucleation kinetics should have a significant and easily measurable variation with supersaturation, at levels of supersaturation that require realistically achievable cooling rates during batch cooling crystallisation experiments.

Machine learning is a growing area of technology that could support the advancement of the crystallisation MPC strategy. One limitation imposed by PharmaMV software is the inability of its controllers to adapt to disturbances to the nucleation and growth kinetics during control operations, instead requiring the controllers to be re-trained while offline. A suggested adaptation to the strategy would be to prepare multiple predictive controllers, each trained using data from simulated crystallisations with a range of kinetic parameters. Measurements of the rate of change of supersaturation could then be used by a machine learning algorithm to determine the best model parameters to use at any given time, and switch to an alternate controller accordingly to maintain better stability of the supersaturation.

References

Abu Bakar, M.R., Nagy, Z.K. and Rielly, C.D. 2009. Seeded Batch Cooling Crystallization with Temperature Cycling for the Control of Size Uniformity and Polymorphic Purity of Sulfathiazole Crystals. *Organic Process Research & Development*. **13**(6), pp. 1343-1356.

Ang, K.H., Chong, G. and Li, Y. 2005. PID Control System Analysis, Design and Technology. *IEEE Transactions on Control Systems Technology.* **13**(4), pp. 559-576.

Anwar, J. and Zahn, D. 2017. Polymorphic phase transitions: Macroscopic theory and molecular simulation. *Advanced Drug Delivery Reviews.* **117**, pp. 47-70.

Bakeev, K.A. 2010. *Process Analytical Technology.* 2nd Edition. Chichester: John Wiley & Sons.

Barrett, P. and Glennon, B. 2002. Characterizing the metastable zone width and solubility curve using lasentec FBRM and PVM. *Chemical Engineering Research & Design*. **80**, pp. 799–805.

Benyahia, B., Anandan, P.D. and Rielly, C. 2021. Control of batch and continuous crystallization processes using reinforcement learning. *Computer Aided Chemical Engineering.* **50**, pp. 1371−1376.

Bermingham, S.K., Trivellato, S., Vernier, M. and Bezzo, F. 2007. Hybrid CFD-Multizonal Modelling of Polymorphs and Agglomeration Phenomena in Crystallisation Process. *Chemical Engineering Transactions.* **11**, pp. 149-154.

Borissova, A., Khan, S., Mahmud, T., Roberts, K.J., Andrews, J., Dallin, P., Chen, Z-P. and Morris, J. 2009. In Situ Measurement of Solution Concentration during the Batch Cooling Crystallization of L-Glutamic Acid using ATR-FTIR Spectroscopy Coupled with Chemometrics. *Crystal Growth & Design.* **9**(2), pp. 692-706.

Braatz, R.D. 2002. Advanced control of crystallization processes. *Annual Reviews in Control*. **26**, pp. 87-99.

Braatz, R.D. and Hasebe, S. 2001. Particle Size and Shape Control in Crystallization Processes. *AIChE Symposium Series*. **98**, pp. 307-327.

Brown, C., McGlone, T. and Florence, A. 2017. Continuous crystallisation. In: Kleinebudde, P., Khinast, J. and Rantanen, J. (eds.), *Continuous Manufacturing of Pharmaceuticals*. John Wiley & Sons Inc, pp. 169-226.

Brown, C.J. 2012. *Characterization of Crystallization Processes with Video Imaging*. Ph.D. thesis, Heriot-Watt University.

Brunsteiner, M., Jones, A.G., Pratola, F., Price, S.L. and Simons, S.J.R. 2005. Toward a Molecular Understanding of Crystal Agglomeration. *Crystal Growth & Design.* **5**(1), pp. 3-16.

Camacho Corzo, D.M., Ma, C.Y., Ramachandran, V., Mahmud, T. and Roberts, K.J. 2017. Crystallisation Route Map. In: Roberts, K.J., Docherty, R. and Tamura, R. (eds). *Engineering Crystallography: From Molecule to Crystal to Functional Form. NATO Science for Peace and Security Series A: Chemistry and Biology*. Dordrecht: Springer, pp. 179-213.

Camacho Corzo, D.M., Ma, C.Y., Mahmud, T. and Roberts, K.J. 2020. Digital Design of Batch Cooling Crystallization Processes: Computational Fluid Dynamics Methodology for Modeling Free-Surface Hydrodynamics in Agitated Crystallizers. *Organic Process Research & Development*. **24**, pp. 2565-2582.

Camacho, D.M., Borissova, A., Hammond, R.B., Kashchiev, D., Roberts, K.J., Lewtas, K. and More, I. 2014. Nucleation mechanism and kinetics from the analysis of polythermal crystallisation data: methyl stearate from kerosene solutions. *CrystEngComm.* **16**, pp. 974-991.

Chen, J., Sarma, B., Evans, J.M.B. and Myerson, A.S. 2011. Pharmaceutical Crystallization. *Crystal Growth & Design*. **11**, pp. 887-895.

Christofides, P.D., El-Farra, N., Li, M., and Mhaskar, P. 2007. Model-based control of particulate processes. *Chemical Engineering Science*. **63**, pp. 1156-1172.

Clarke, D.W., Mohtadi, C. and Tuffs, P.S. 1987. Generalized predictive control—Part I. The basic algorithm. *Automatica.* **23**(2), pp. 137-148.

Costa, C.B.B., Maciel, M.R.W. and Filho, R.M. 2007. Considerations on the crystallization modeling: Population balance solution. *Computers & Chemical Engineering*. **31**(3), pp. 206-218.

Cui, Y., Jaramillo, J.J., Stelzer, T. and Myerson, A.S. 2015. Statistical Design of Experiment on Contact Secondary Nucleation as a Means of Creating Seed Crystals for Continuous Tubular Crystallizers. *Organic Process Research & Development.* **19**(9), pp. 1101−1108.

Davey, R. and Garside, J. 1980. Invited Review Secondary Contact Nucleation: Kinetics, Growth and Scale-Up. *Chemical Engineering Communications.* **4**, pp. 393-424.

Davey, R. and Garside, J. 2006. *From Molecules to Crystallizers – An Introduction to Crystallization.* Oxford: Oxford University Press.

Davey, R. and Rütti, A. 1976. Agglomeration in the crystallisation of hexamethylene tetramine from aqueous solution. *Journal of Crystal Growth.* **32**(2), pp. 221-226.

De Anda, J.C., Wang, X.Z., Lai, X. and Roberts, K.J. 2005. Classifying organic crystals via in process image analysis and the use of monitoring charts to follow polymorphic and morphological changes. *Journal of Process Control*. **15**, 785–797.

Denbigh, K.G. and White, E.T. 1966. Studies on liquid inclusions in crystals. *Chemical Engineering Science*. **21**, pp. 739–753.

Docherty, R., Clydesdale, G., Roberts, K.J. and Bennema, P. 1991. Application of Bravais-Friedel-Donnay-Harker, attachment energy and Ising models to predicting and understanding the morphology of molecular crystals. *Journal of Physics D: Applied Physics*. **24**, pp. 89-99.

Doki, N., Seki, H., Takano, K., Asatani, H., Yokota, M. and Kubota, N. 2004. Process Control of Seeded Batch Cooling Crystallization of the Metastable α-Form Glycine Using an In-Situ ATR-FTIR Spectrometer and an In-Situ FBRM Particle Counter. Crystal Growth and Design. **4**(5), pp. 949-953.

Evans, T.W., Sarofim, A.F. and Margolis, G. 1974. Models of Secondary Nucleation attributable to Crystal-Crystallizer and Crystal-Crystal Collisions. *AIChE Journal.* **20**(5), p. 959.

Falola, A., Borissova, A. and Wang, X.Z. 2013. Extended method of moment for general population balance models including size dependent growth rate, aggregation and breakage kernels. *Computers and Chemical Engineering.* **56**, pp. 1-11.

Fei, L., Gan, X., Ng, S.M., Wang, H., Xu, M., Lu, W., Zhou, Y., Leung, C.W., Mak, C-L. and Wang, Y. 2018. Observable Two-Step Nucleation Mechanism in Solid-State Formation of Tungsten Carbide. *ACS Nano.* **13**, pp. 681-688.

Forgione, M., Birpoutsoukis, G., Bombois, X., Mesbah, A., Daudey, P.J. and Van de Hof, P.M.J. 2015. Batch-to-batch model improvement for cooling crystallisation. *Control Engineering Practice*. **41**, pp. 72-82.

Fujiwara, M., Chow, P.S., Ma, D.L. and Braatz, R.D. 2002. Paracetamol crystallization using laser backscattering and ATR-FTIR spectroscopy: metastability, agglomeration and control. *Crystal Growth & Design*. **2**, pp. 363–70.

Fujiwara, M., Nagy, Z.K., Chew, J.W. and Braatz, R.D. 2005. First-principles and direct design approaches for the control of pharmaceutical crystallization. *Journal of Process Control*. **15**, pp. 493-504.

Gajjar, P., Nguyen, T.T.H., Sun, J., Styliari, I.D., Bale, H., McDonald, S.A., Burnett, T.L., Tordoff, B., Lauridsen, E., Hammond, R.B., Murnane, D., Withers, P.J., and Roberts, K.J. 2021. Crystallographic tomography and molecular modelling of structured organic polycrystalline powders. *CrystEngComm*. **23**, 2520-2531.

Garside, J., Mersmann, A. and Nývlt, J. 1990. *Measurement of crystal growth rates*. Munich: European Federation of Chemical Engineering, Working Party on Crystallization.

Gherras, N. and Fevotte, G. 2012. Comparison between approaches for the experimental determination of metastable zone width: a case study of the batch cooling crystallization of ammonium oxalate in water. *Journal of Crystal Growth*. **342**, pp. 88-98.

Gibbs, J.W. 1948. *Collected Works, Vol. I, Thermodynamics.* New Haven: Yale University Press.

Gill, P.E., Murray, W., Saunders, M.A. and Wright, M.H. 2001. *User's Guide for NPSOL 5.0: A Fortran Package for Nonlinear Programming*. USA.

Griffin, D.J., Grover, M.A., Kawajiri, Y. and Rousseau, R.W. 2016. Data-Driven Modeling and Dynamic Programming Applied to Batch Cooling Crystallization. *Industrial & Engineering Chemistry Research.* **55**, pp. 1361-1372.

Grön, H., Borissova, A. and Roberts, K. J. 2003. In-Process ATR-FTIR Spectroscopy for Closed-Loop Supersaturation Control of a Batch Crystallizer Producing Monosodium Glutamate Crystals of Defined Size *Industrial & Engineering Chemistry Research*. **42**(1), pp.198–206.

Hounslow, M.J., Ryall, R.A. and Marshall, V.R. 1988. A Discretized Population Balance for Nucleation, Growth and Aggregation. *AIChE Journal*. **34**(11), pp. 1821-1832.

Hunter, C.A. 2004. Quantifying Intermolecular Interactions: Guidelines for the Molecular Recognition Toolbox. *Angewandte Chemie International Edition*. **43**, pp. 5310 – 5324.

Jha, S.K., Karthika, S. and Radhakrishnan, T.K. 2017. Modelling and control of crystallization processes. *Resource-Efficient Technologies*. **3**(1), pp. 94-100.

Jones, A.G. 1993. Particle Breakage, Abnormal Growth and Agglomeration During Industrial Crystallization. *Analytical Proceedings.* **30**, pp. 456-457.

Kardum, J. P., Hrkovac, M. and Leskovac, M. 2013. Adjustment of process conditions in seeded batch cooling crystallization. *Chemical Engineering & Technology*. **36**(8), pp. 1347-1354.

Kashchiev, D., Borissova, A., Hammond, R.B. and Roberts, K.J. 2010. Dependence of the critical undercooling for crystallization on the cooling rate. *Journal of Physical Chemistry B.* **114**(16), pp. 5441-5446.

Kaye, G.W.C. and Laby, T.H. 1995. *Tables of Physical and Chemical Constants*. 16th Ed. Harlow; Longman

Khan, S., Ma, C.Y., Mahmud, T., Penchev, R.Y., Roberts, K.J., Morris, J., Özkan, L., White, G., Grieve, B., Hall, A., Buser, P., Gibson, N., Keller, P., Shuttleworth, P. and Price, C.J. 2011. In-Process Monitoring and Control of Supersaturation in Seeded Batch Cooling Crystallisation of l-Glutamic Acid: From Laboratory to Industrial Pilot Plant. *Organic Process Research & Development.* **15**(3), pp. 540-555.

Kim, K-J. and Mersmann, A. 2001. Estimation of metastable zone width in different nucleation processes. *Chemical Engineering Science.* **56**, pp. 2315-2324.

Kubota, N. 2008. A new interpretation of metastable zone width measured for unseeded solutions. *Journal of Crystal Growth*. **310**(3), pp. 629-634.

Kumar, V., Taylor, M.K., Mehrotra, A. and Stagner, W.C. 2013. Real-Time Particle Size Analysis Using Focused Beam Reflectance Measurement as a Process Analytical Technology Tool for a Continuous Granulation–Drying–Milling Process. *AAPS PharmSciTech.* **14**(2), pp. 523-530.

Kwon, J.S-I., Nayhouse, M., Orkoulas, G., Ni, D. and Christofides, P.D. 2015. Run-to-Run-Based Model Predictive Control of Protein Crystal Shape in Batch Crystallization. *Industrial & Engineering Chemistry Research*. **54**, pp. 4293-4302.

Lee, P.I. and Sullivan, G.R. 1988. Generic model control. *Computers & Chemical Engineering.* **12**(6), pp. 573-580.

Levine, W.S. 1996. *The Control Handbook*. Piscataway, NJ: IEEE Press.

Liffman, K. 1992. A direct simulation Monte-Carlo method for cluster coagulation. *Journal of Computational Physics*. **100**, pp. 116−127.

Lin, Y., Lee, K. and Matsoukas, T. 2002. Solution of the populationbalance equation using constant-number Monte Carlo. *Chemical Engineering Science*. **57**, pp. 2241−2252.

Long, B. W., Yang, H. T. and Ding, Y. G. 2016. Impact of seed loading ratio on the growth kinetics of mono-ammonium phosphate under isothermal batch crystallization. *Korean Journal of Chemical Engineering*. **33**(2), pp. 623-628.

Maggioni, G.M. and Mazzotti, M. 2013. A Stochastic Population Balance Equation Model for Nucleation and Growth of Crystals with Multiple Polymorphs. *Crystal Growth & Design.* **19**, pp. 4698-4709.

Maggioni, G.M. and Mazzotti, M. 2017. Stochasticity in Primary Nucleation: Measuring and Modeling Detection Times. *Crystal Growth & Design.* **17**, pp. 3625-3635.

Mahoney, A.W. and Ramkrishna, D. 2002. Efficient solution of population balance equations with discontinuities by finite elements. *Chemical Engineering Science.* **57**, pp. 1107-1119.

Maisels, A., Einar Kruis, F. and Fissan, H. 2004. Direct simulation Monte Carlo for simultaneous nucleation, coagulation and surface growth indispersed systems. *Cemical Engineering Science.* **59**, pp. 2231−2239.

Marchal, P., David, R., Klein, J.P. and Villermaux, J. 1988. Crystallization and precipitation engineering—I. An efficient method for solving population balance in crystallization with agglomeration. *Chemical Engineering Science*. **43**(1), pp. 59-67.

McCabe, W.L., Smith, J.C. and Harriott, P. 1993. *Unit Operations of Chemical Engineering.* 5th Ed. Singapore: McGraw-Hill.

McCrone, W.C. 1965. Polymorphism. In: Fox, D., Labes, M.M. and Weissberger, A. (eds.), *Physics and Chemistry of the Organic Solid State*. New York: Interscience.

Mersmann, A. and Bartosch, K. 1998. How to predict the metastable zone width. *Journal of Crystal Growth*. **183**(1-2), pp. 240-250.

Miller, A.M. 2008. Characterization of Hexamine Squarate. *The Microscope*. **56**(4), pp. 147-154.

Miller, S.M. and Rawlings, J.B. 1994. Model Identification and Control Strategies for Batch Cooling Crystallizers. *AIChE Journal*. **40**(8), pp. 1312-1327.

Mitchell, N. 2012. *Numerical Modelling of Cooling Crystallisation: Process Kinetics to Optimisation*. Ph.D. thesis, University of Limerick.

Miyazawa, T., Kondo, S., Takuya, S. and Sato, H. 2012. Specific Heat Capacity at Constant Pressure of Ethanol by Flow Calorimetry. *Journal of Chemical & Engineering Data*. **57**, pp. 1700-1707.

Mohammad, K.A., Rahim, S.A. and Abu Bakar, M.R. 2017. Kinetics and nucleation mechanism of carbamazepine–saccharin co-crystals in ethanol solution. *Journal of Thermal Analysis and Calorimetry.* **130**, pp. 1663-1669.

Moldovány, N., Lakatos, B.G. and Szeifert, F. 2005. Model predictive control of MSMPR crystallizers. *Journal of Crystal Growth*. **275**(1-2), pp. 1349-1354.

Mullin, J.W. 1992. *Crystallization.* 3rd Edition. Oxford: Butterworth-Heinemann.

Mullin, J.W. and Leci, C.L. 1972. Desupersaturation of seeded citric acid solutions in a stirred vessel. *AIChE Symposium Series, No. 121*. **68**, pp. 8-20.

Mumtaz, H.S., Hounslow, M.J., Seaton, N.A. and Paterson, W.R. 1997. Orthokinetic aggregation during precipitation: A computational model for calcium oxalate monohydrate *Chemical Engineering Research and Design.* **75**(2), pp. 152-159.

Myerson, A.S. and Ginde, R. 2002. Crystals, Crystal Growth, and Nucleation. In: Myerson, A.S. (ed.), *Handbook of Industrial Crystallization*. Woburn: Butterworth-Heinemann, pp. 33-65.

Myerson, A.S., Decker, S.E. and Weiping, F. 1986. Solvent Selection and Batch Crystallisation. *Industrial & Engineering Chemistry Process Design & Development*. **25**, pp. 925-929.

Nagy, Z.K. 2009. Model based robust batch-to-batch control of particle size and shape in pharmaceutical crystallisation. *IFAC Proceedings Volumes.* **42**(11), pp. 195-200.

Nagy, Z.K. and Braatz, R.D. 2012. Advances and New Directions in Crystallization Control. *Annual Review of Chemical and Biomolecular Engineering*. **3**, pp. 55-75.

Nagy, Z.K., Chew, J.W., Fujiwara, M. and Braatz, R.D. 2008a. Comparative performance of concentration and temperature controlled batch crystallizations. *Journal of Process Control*. **18**(3-4), pp. 399-407.

Nagy, Z.K., Fujiwara, M. and Braatz, R.D. 2008b. Modelling and control of combined cooling and antisolvent crystallization processes. *Journal of Process Control.* **18**, pp. 856-864.

Nernst, W. 1904. Theorie der Reaktionsgeschwindigkeit in heterogenen Systemen. *Zeitschrift fur Physikalische Chemie*. **47**, pp. 52-55.

Nguyen, T.T.H., Gajjar, P., Sun, J., Hammond, R.B., Murnane, D., Withers, P.J. and Roberts, K.J. 2023. Understanding the agglomerate crystallisation of hexamine through X-ray microscopy and crystallographic modelling. *Journal of Crystal Growth*. **603**, article no: 126986.

Noyes, A.A. and Whitney, W.R. 1897. Rate of solution of solid substances in their own solution. *Journal of the American Chemical Society*. **19**, pp. 930-934.

Nývlt, J., Rychlý, R., Gottfried, J. and Wurzelová, J. 1970. Metastable Zone-Width of Some Aqueous Solutions. Journal of Crystal Growth. **6**, pp. 151-162.

Omar, H.M. and Rohani, S. 2017. Crystal Population Balance Formulation and Solution Methods: A Review. *Crystal Growth & Design.* **17**, pp. 4028-4041.

Orehek, J., Teslić, D. and Likozar, B. 2021. Continuous Crystallization Processes in Pharmaceutical Manufacturing: A Review. *Organic Process Research & Development.* **25**, pp. 16-42.

Rajagopalan, A.K., Bötschi, S., Morari, M. and Mazzotti, M. 2018. Feedback control for the size and shape evolution of needle-like crystals in suspension. II. Cooling crystallisation experiments. *Crystal Growth & Design.* **18**, pp. 6185-6196.

Randolph, A.D. and Larson, M.A. 1988. *Population Balances: Theory of Particulate Processes*. 2nd Edition. San Diego: Academic Press.

Rohani, S. and Bourne, J.R. 1990. A Simplified Approach to the Operation of a Batch Crystallizer. *The Canadian Journal of Chemical Engineering*. **68**, pp. 799-806.

Romero-Torres, S. 2021. The role of digital transformation in achieving Pharma 4.0. *European Pharmaceutical Review.* **26**(3), article no: 157165.

Saleemi, A.N., Rielly, C.D. and Nagy, Z.K. 2012. Comparative investigation of supersaturation and automated direct nucleation control of crystal size distributions using ATR-UV/Vis spectroscopy and FBRM. *Crystal Growth & Design*. **12**, pp. 1792-1807.

Sangwal, K. and Wójcik, K. 2009. Investigation of metastable zone width of ammonium oxalate aqueous solutions. *Crystal Research & Technology*. **44**(4), pp. 363-372.

Sato, K. 1993. Polymorphic transformations in crystal growth. *Journal of Physics and Applied Physics*. **26**, pp. B77-B84.

Shaikh, L., Pandit, A. and Ranade, V. 2013. Crystallisation of ferrous sulphate heptahydrate: Experiments and modelling. *The Canadian Journal of Chemical Engineering.* **91**, pp. 47-53.

Shaikh, L.J., Bari, A.H., Ranade, V.V. and Pandit, A.B. 2015. Generic Framework for Crystallization Processes Using the Population Balance Model and Its Applicability. *Industrial & Engineering Chemistry Research.* **54**(42), pp. 10539-10548.

Shamekh, A., Lennox, B., Sandoz, D. and Marjanovic, O. 2008. PLS and its Application within Model Predictive Controllers. *IFAC Proceedings Volumes*. **41**(2), pp. 12389-12394.

Shi, D., Mhaskar, P., El-Farra, N.H. and Christofides, P.D. 2005. Predictive control of crystal size distribution in protein crystallization. *Nanotechnology.* **16**(7), pp. 562-574.

Siemens-PSE Ltd. 2021. *gPROMS FormulatedProducts v2.2.0 Documentation.* [Accessed 07 March 2023].

Simone, E. 2015. *Application of process analytical technology (PAT) tools for the better understanding and control of the crystallization of polymorphic and impure systems*. Ph.D. thesis, Loughborough University.

Simone, E., Saleemi, A.N. and Nagy,Z. K. 2015. In situ monitoring of polymorphic transformations using a composite sensor array of Raman, NIR, and ATR-UV/vis spectroscopy, FBRM, and PVM for an intelligent decision support system. *Organic Process Research & Development.* **19**,167−177.

Szilágyi, B. and Nagy, Z.K. 2016. Graphical processing unit (GPU) acceleration for numerical solution of population balance models using high resolution finite volume algorithm. *Computers and Chemical Engineering*. **91**, pp. 167-181.

Szilágyi, B. and Nagy, Z.K. 2018. Aspect ratio distribution and chord length distribution driven modeling of crystallization of two-dimensional crystals for real-time model-based applications. *Crystal Growth & Design*. **18**(9), pp.5311-5321.

Tahir, F., Krzemieniewska-Nandwani, K., Mack, J., Lovett, D., Siddique, H., Mabbott, F., Raval, V., Houson, I. and Florence, A. 2017. Advanced control of a continuous oscillatory flow crystalliser. *Control Engineering Practice.* **67**, pp. 64-75.

Togkalidou, T., Fujiwara, M., Patel, S. and Braatz, R.D. 2001. Solute concentration prediction using chemometrics and ATR-FTIR spectroscopy. *Journal of Crystal Growth*. **231**, pp. 534–543.

Toroz, D., Rosbottom, I., Turner, T., Camacho, D.M., Hammond, R.B. and Roberts K.J. 2015. Towards an understanding of the nucleation of alpha-para amino benzoic acid from ethanolic solutions: a multi-scale approach. *Faraday Discussions*. **179**, pp. 79-114.

Turner, T.D., Camacho Corzo, D.M., Toroz, D., Curtis, A., Dos Santos, M.M., Hammond, R.B., Lai, X., and Roberts, K.J. 2016. The Influence of Solution Environment on the Nucleation Kinetics and Crystallisability of Para-Aminobenzoic Acid. *Physical Chemistry Chemical Physics*. **18**, pp. 27507−27520.

Verkoeijen, D., Pouw, G.A., Meesters, G.M.H. and Scarlet, B. 2002. Population balances for particulate processes – A volume approach. *Chemical Engineering Science*. **57**, pp. 2287-2303.

Vollmer, U. and Raisch, J. 2006. Control of batch crystallization – A system inversion approach. Chemical Engineering and Processing. **45**, pp. 874-885.

Volmer, M. 1939. *Kinetic der Phasenbildung*. Leipzig: Steinkopff.

Wang, L., Dai, L-Y., Lei, M. and Chen, Y. 2008. Solubility of Hexamethylenetetramine in a Pure Water, Methanol, Acetic Acid, and Ethanol + Water Mixture from (299.38 to 340.35) K. *Journal of Chemical Engineering Data*. **53**(12), pp. 2907-2909.

Xu, S., Hou, Z., Chuai, X. and Wang, Y. 2020. Overview of Secondary Nucleation: From Fundamentals to Application. *Industrial & Engineering Chemistry Research*. **59**, pp. 18335−18356.

Yang,Y. and Nagy, Z.K. 2015. Application of Nonlinear Model Predictive Control in Continuous Crystallization Systems. In: *2015 American Control Conference (ACC)*. Chicago: IEEE, pp. 4282−4287.

Zhang, B., Hou, X., Dang, L. and Wai, H. 2019. Selective polymorphic crystal growth on self-assembled monolayer using molecular modelling as an assistant method. *Journal of Crystal Growth*. **518**, pp. 81-88.

Zhang, F., Liu, T., Wang, X.Z., Liu, J. and Jiang, X. 2017. Comparative study on ATR-FTIR calibration models for monitoring solution concentration in cooling crystallization. *Journal of Crystal Growth.* **459**, pp. 50-55

Zhang, F., Shan, B., Wang, Y., Zhu, Z., Yu, Z.O. and Ma, C.Y. 2021. Progress and Opportunities for Utilizing Seeding Techniques in Crystallization Processes. *Organic Process Research & Development*. **25**, pp. 1496-1511.

Zhang, G.P. and Rohani, S. 2003. On-line optimal control of a seeded batch cooling crystallizer. *Chemical Engineering Science.* **58**(9), pp. 1887-1896.

Zhao, H., Maisels, A., Matsoukas, T. and Zheng, C. 2007. Analysis of four Monte Carlo methods for the solution of population balances in dispersed systems. *Powder Technology*. **173**, pp. 38−50.

Zhao, H., Zheng, C. and Xu, M. 2005. Multi-Monte Carlo method for coagulation and condensation/evaporation in dispersed systems. *Journal of Colloid and Interface Science.* **286**, pp. 195−208.

Zheng, Y., Wang, X. and Wu, Z. 2022. Machine learning modelling and predictive control of the batch crystallization process. *Industrial & Engineering Chemistry Research*. **61**, pp. 5578−5592.

Appendix A – Conversion of Volume-Based to Number-Based CSD in gPROMS FormulatedProducts

The following custom variables were created in the 'gPROMS language' tab of the batch crystallisation flowsheet built in gPROMS FormulatedProducts software, which relate to properties of the number-based size distribution:

```
[87] particles_per_bin   AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[88] mass_per_bin   AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[89] vol_per_bin    AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[90] vol_per_particle    AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[91] area_per_bin   AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[92] area_per_particle   AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[93] total_surface_area  AS no_type

[94] length_per_bin AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[95] total_number_particles AS no_type
```

```
[96] cumulative_number_dist AS
ARRAY(PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts) OF
no_type

[116] FOR i:=1 TO
PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts DO

[117]     mass_per_bin(i) =
PSD_sensor.PSD_sensor_block.PSD_analysis.mass_dist(i) *
PSD_sensor.PSD_sensor_block.PSD_analysis.total_mass ; #[kg] =
[kg/kg] * [kg]

[120]     vol_per_bin(i) = mass_per_bin(i) /
PSD_sensor.PSD_sensor_block.PSD_analysis.density_solid(1,i); #
[m^3] = [kg] / [kg/m^3]

[123]     vol_per_particle(i) =
(PSD_sensor.PSD_sensor_block.particle_bin_mid_point(i)*(1E-
6))^3 * PSD_sensor.PSD_sensor_block.PSD_analysis.shape_fac_vol
; # [m^3] = [m^3] * [-]

[126]     particles_per_bin(i) = vol_per_bin(i) /
vol_per_particle(i) ; # [#] = [m^3] / [m^3]

[129]     area_per_particle(i) =
(PSD_sensor.PSD_sensor_block.particle_bin_mid_point(i)*(1E-
6))^2 *
PSD_sensor.PSD_sensor_block.sensor_information.shape_fac_area(
'Hexamine crystal') ; # [m^2] = [m^2] * [-]

[132]     area_per_bin(i) = particles_per_bin(i) *
area_per_particle(i) ; # [m^2] = [#] * [m]^2]

[135]     length_per_bin(i) =
PSD_sensor.PSD_sensor_block.particle_bin_mid_point(i) ; # [µm]
= [µm]

[138] END
```

```
[140] total_surface_area = SIGMA(area_per_bin()) ; # [m^2] =
[m^2]

[143] total_number_particles = SIGMA(particles_per_bin()) ; #
[#] = [#]

[146] FOR i:=1 TO
PSD_sensor.PSD_sensor_block.PSD_analysis.no_grid_pts DO

[147]     cumulative_number_dist(i) =
SIGMA(particles_per_bin(1:i)) / total_number_particles ;

[149] END
```

Appendix B – List of Presentations and Publications

Conference Presentations

- Leeming, R., Mahmud, T., Roberts, K.J., Simone, E., George, N. and Brown, C.J. *Model predictive control of industrial crystallisation systems: Experimental characterisation and population balance modelling of batch cooling crystallisation of hexamine in ethanol.* BCA/BACG Online Joint Spring Meeting, April 2021 (Poster)
- Leeming, R., Mahmud, T., Roberts, K.J., Simone, E., George, N. and Brown, C.J. *Model predictive control of industrial crystallisation systems: Experimental characterisation and population balance modelling of batch cooling crystallisation of hexamine in ethanol.* ISIC 21 - 21st International Symposium on Industrial Crystallization, August 2021 (Poster and Oral)
- Leeming, R., Mahmud, T., Roberts, K.J., Simone, E., George, N., Webb, J and Brown, C.J. *Model predictive control of supersaturation during batch cooling crystallisation of hexamine from ethanol solvent.* CGOM14 - 14th International Conference on Crystal Growth of Organic Materials, September 2022 (Poster)
- Leeming, R., Mahmud, T., Roberts, K.J., Simone, E., George, N., Webb, J and Brown, C.J. *Effect of impurities on supersaturation control during batch cooling crystallisation of an organic compound.* ICCGE-20 - International Conference on Crystal Growth and Epitaxy, July 2023 (Oral)
- Leeming, R., Mahmud, T., Roberts, K.J., Simone, E., George, N., Webb, J and Brown, C.J. *Model predictive control of supersaturation and crystal size during batch cooling crystallisation of hexamine from ethanol solution.* ISIC22/BACG52, September 2023 (Oral)

Journal Publications

- Leeming, R., Mahmud, T., Roberts, K.J., George, N., Webb, J., Simone, E. and Brown, C.J. 2023. Development of a Digital Twin for the Prediction and Control of Supersaturation during Batch Cooling Crystallization. *Ind. Eng. Chem. Res.* **62**, pp. 11067-11081.
- Leeming, R., Mahmud, T., Roberts, K.J., George, N., Webb, J., Simone, E. and Brown, C.J. 2023. Effect of Impurities on Supersaturation Control during Batch Cooling Crystallization of Hexamine. (Draft, working title).