Optimisation of flow chemistry: tools and algorithms

Carlos González Niño

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

The University of Leeds
School of Mechanical Engineering
School of Chemistry

January, 2020

The candidate confirms that the work submitted is his/her own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapter 2 of this thesis includes work relative to the first version of the milli-fluidic scale continuous stirred tank reactor design that has appeared in the following publication:

M. R. Chapman, M. H. T. Kwan, G. King, K. E. Jolley, M. Hussain, S. Hussain*, et al.*, "Simple and Versatile Laboratory Scale CSTR for Multiphasic Continuous-Flow Chemistry and Long Residence Times," *Organic Process Research & Development,* vol. 21, pp. 1294-1301, 2017.

where the author of this thesis conducted postprocessing and analysis of the RTD data for the device.

In addition, the UV setup developed by the author of this thesis for the fReactors presented in the same chapter was crucial in the research leading to the next publication:

J. A. Manson, A. D. Clayton, C. G. Niño, R. Labes, T. W. Chamberlain, A. J. Blacker*, et al.*, "A Hybridised Optimisation of an Automated Photochemical Continuous Flow Reactor," *CHIMIA International Journal for Chemistry,* vol. 73, pp. 817-822, 2019.

Chapter 4 is based on the paper:

C. González Niño, N. Kapur, M.-F. King, G. de Boer, A. J. Blacker, R. Bourne, et al., "Computational fluid dynamic enabled design optimisation of miniaturised continuous oscillatory baffled reactors in chemical processing," *International Journal of Computational Fluid Dynamics*, pp. 1-15, 2019.

All of the work carried out in relation to this publication was conducted by the author of this thesis, with the role of each of the co-authors being to provide guidance during its multiple stages, from conception, CFD modelling and response surface modelling to the optimisation of the design.

Moreover, the synchrotron XRD analysis featured in the same chapter was carried out in some of the beamtimes that led to the following publication:

M. A. Levenstein, C. Anduix-Canto, Y.-Y. Kim, M. A. Holden, C. González Niño, D. C. Green*, et al.*, "Droplet Microfluidics XRD Identifies Effective Nucleating Agents for Calcium Carbonate," *Advanced Functional Materials,* vol. 29, p. 1808172, 2019.

Works not featured in this thesis include:

K. E. Jolley, W. Nye, C. González Niño, N. Kapur, A. Rabion, K. Rossen, et al., "Highly Productive Continuous Flow Synthesis of Di- and Tripeptides in Water," *Organic Process Research & Development*, vol. 21, pp. 1557-1565, 2017.

Acknowledgements

First of all I would like to express my gratitude towards the University of Leeds, that awarded me the scholarship that made possible the four years of research that led to the culmination of this thesis.

I would like to extend my gratitude to my supervisory team, including Professors Nik Kapur, John Blacker and Harvey Thompson, and Dr. Richard Bourne. The striking number of people involved in the supervision of this thesis is explained because of the multidisciplinary nature of the project. I feel very fortunate to have counted Nik as my primary supervisor, who was a source not only of knowledge but also of personal encouragement at the times where my PhD became most problematic. A DIY legend, his insights when building electronics or mechanical prototypes were priceless. Probably, being the victim of my many and continuous assaults to his office have affected him in ways one can only speculate about. John, who is the head of the iPRD, has given valuable technical advice around the chemistry involved in the project. Harvey, who is now the Head of School of Mechanical Engineering, was able to provide pertinent advice concerning response surface modelling and optimisation, as well as motivating me with an energy that is contagious. Richard, a recognised person in the field of continuous flow automation, has provided access and understanding related to the automated rig at iPRD.

I thank Dr. Michael Chapman, Dr. Katie Jolley, Adam Clayton, Dr. Mary Bayana, Dr. Álvaro Cruz Izquierdo, Fernando Climent Barba, and Dr. Maria Kwan for their support during my time at iPRD.

I also appreciate the support of the staff at the university, especially Dave Readman, whose vast knowledge of microprocessors and electronics helped me dramatically at the time I had to build my prototypes. Tony Wiese was of great help when working at the mechanical workshop, and his ability to locate items immediately in such a large facility always amazed me. Fiona Slade and Judith Schneider were in charge of the administration, and always helped me when necessary (which was a lot of times as paperwork is not my forte).

During these years, I had the fortune to meet colleagues that eventually became friends. Dr. Maciej Napora is an encyclopaedia of automation and control, as well as a good friend of mine. Dr. Cayetano Espejo Conesa, who is one of the most brilliant people I have ever come across, is also a person of exceptional human quality that supported me both in the good and the bad

times. The same can be said of Dr. Christian Israel Aragón Briceño ‘El Cuate’. Dr. Marco-Felipe King and Dr. Gregory de Boer were also very supportive when I was writing my papers. I would also like to mention my colleagues in my first office, Will Breeze, Shunsuke Koide, Alysson Santos Bueno and specially Corentin Baudu, who I had the privilege to visit in his beautiful town, Figeac. It would be unforgivable not to mention Dr. Chung Wang, who is a very positive and supportive person.

Being an hispanophile and a coffee lover, Dr. Mark Levenstein was someone easy to make friends with. Our regular meetings over coffee to discuss Spanish history and culture, politics, and try to address the problems of the broken modern world were often the remedy for the moody behavioural consequence of cloudy days. I also have the fondest memories of our synchrotron trips to Grenoble, and the furtive excursions to Lyon and Geneva. I would also like to thank Dr. Clara Anduix Cantó for her support during the beamtimes. Along with Mark, Nik and myself, Clara formed part of the so-called “A-Team” of synchrotron experimentation.

Likewise, I would also like to thank Dr. Tim Snow and Dr. Nick Terrill of Diamond UK, and Diego Pontoni and Peter van der Linden of ESRF for their assistance during beamtimes. I also thank the staff at the Advance Research Computing facilities of the University of Leeds for making possible the use of such a great resource that helped me cut times when solving CFD problems. Let’s not forget the IT staff, specially Mark Hauptfleisch who was able to rapidly solve a problem at a very sensitive time during my thesis writing.

Finally, I would like to thank my family, and my girlfriend who supported me during my PhD. My mother, Ángela, and my sister, María believed I was capable of continue going forward from the beginning. My girlfriend, Esther, supported me every day, being my company in our cosy home in Wakefield. My father, who unfortunately passed away during the course of my studies, was the most encouraging when deciding whether to continue with a PhD after my Erasmus+ internship. I have the fondest memory of him, and his honesty and morals have always been an inspiration for me.

Abstract

The coupling of flow chemistry with automated laboratory equipment has become increasingly common and used to support the efficient manufacturing of chemicals. A variety of reactors and analytical techniques have been used in such configurations for investigating and optimising the processing conditions of different reactions. However, the integrated reactors used thus far have been constrained to single phase mixing, greatly limiting the scope of reactions for such studies. This thesis presents the development and integration of a millilitre-scale CSTR, the fReactor, that is able to process multiphase flows, thus broadening the range of reactions susceptible of being investigated in this way.

Following a thorough review of the literature covering the uses of flow chemistry and lab-scale reactor technology, insights on the design of a temperature-controlled version of the fReactor with an accuracy of ±0.3 ºC capable of cutting waiting times 44% when compared to the previous reactor are given. A demonstration of its use is provided for which the product of a multiphasic reaction is analysed automatically under different reaction conditions according to a sampling plan. Metamodeling and cross-validation techniques are applied to these results, where single and multi-objective optimisations are carried out over the response surface models of different metrics to illustrate different trade-offs between them. The use of such techniques allowed reducing the error incurred by the common least squares polynomial fitting by over 12%. Additionally, a demonstration of the fReactor as a tool for synchrotron X-Ray Diffraction is also carried out by means of successfully assessing the change in polymorph caused by solvent switching, this being the first synchrotron experiment using this sort of device.

The remainder of the thesis focuses on applying the same metamodeling and cross-validation techniques used previously, in the optimisation of the design of a miniaturised continuous oscillatory baffled reactor. However, rather than using these techniques with physical experimentation, they are used in conjunction with computational fluid dynamics. This reactor shows a better residence time distribution than its CSTR counterparts. Notably, the effect of the introduction of baffle offsetting in a plate design of the reactor is identified as a key parameter in giving a narrow residence time distribution and good mixing. Under this configuration it is possible to reduce the RTD variance by 45% and increase the mixing efficiency by 60% when compared to the best performing opposing baffles geometry.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

ACS	Amorphous Calcium Sulphate
ADC	Analog-to-Digital Converter
AMI	Absolute Mixing Index
BPR	Back Pressure Regulator
CAD	Computer-Aided Design
CCF	Face-Centered Central Composite Design
CFD	Computational Fluid Dynamics
COBR	Continuous Oscillatory Baffled Reactor
CSTR	Continuous Stirred-Tank Reactor
DC	Direct Current
DoE	Design of Experiments
EMF	Electromotive force
FIFO	First In, First Out
HPLC	High-Performance Liquid Chromatography
IC	Integrated Circuit
I^2C	Inter-Integrated Circuit
ICSP	In-Circuit Serial Programming
LCD	Liquid-Crystal Display
LOO	Leave-One-Out
LOOCV	Leave-One-Out Cross-Validation
MCCV	Monte Carlo Cross-Validation
mCOBR	Miniaturised Continuous Oscillatory Baffled Reactor
MLSM	Moving Least Squares Method
MOSFET	Metal–Oxide–Semiconductor Field-Effect Transistor
PCB	Printed Circuit Board
PDF	Probability Density Function
PI controller	Proportional-Integral controller
PID controller	Proportional-Integral-Derivative controller

PMI	Process Mass Intensity
PTFE	Polytetrafluoroethylene
PWM	Pulse-Width Modulation
RBF	Radial Basis Function
RMI	Relative Mixing Index
RMSE	Root-Mean-Square Error
RSM	Response Surface Methodology
RTD	Residence Time Distribution
SCL	Serial Clock Signal of an I^{2}C electronic bus
SDA	Serial Data Signal of an I^{2}C electronic bus
SEM	Scanning Electron Microscopy
SNR	Signal-to-Noise Ratio
STY	Space-Time Yield
TDS	Transport of Diluted Species
UART	Universal Asynchronous Receiver-Transmitter
UNF	Unified fine thread
UV	Ultraviolet
VDC	Volts Direct Current
XRD	X-Ray Diffraction

Chapter 1
Introduction and Literature Review

The inexorable demand for better reaction yields in chemical production processes, with reduced waste and costs, in the fine chemical and pharmaceutical industries have led to increasing reliance on small-scale studies for effective process optimisation.[1-4] A range of different small-scale flow chemistry platforms for reaction screening have been developed in recent years in order to determine optimal synthesis routes [5] with minimal chemical consumption [6] enabling real-time at-line [7] and in-line [8] analysis. These characteristics make possible the implementation of self-optimising routines for which the reaction outputs can be automatically explored and optimised computationally through computers connected to the laboratory instrumentation.

1.1 Flow Chemistry

In flow chemistry, a chemical reaction takes place in a flowing fluid instead of in a batch. In other words, the pumps move the fluids in conduits and when these conduits converge the fluids contact each other. If these fluids are reactants, a chemical reaction takes place. Flow chemistry is a well-established technique in the large scale manufacturing industry and is gaining momentum among the pharmaceutical and fine chemical companies.[5] Figure 1 shows a bar chart for the number of publications related to flow chemistry over the last two decades. As can be seen, interest in the topic started to timidly grow in the early 2000’s, and a steady and more pronounced growth has been taking place since 2010.

The application of the field of microfluidics to flow chemistry is a powerful tool drawing the attention for its use in the laboratory scale where microreactors are often used, in contrast with the use of the flask which has been the most important piece of equipment in chemistry labs for centuries.[9] This technology is based on the manipulation of a liquid flow in microchannels. The flow movement is generated by external pressure supplies such as mechanical external pumps, mechanically integrated micropumps, or by combinations of capillary forces and electrokinetic mechanisms.

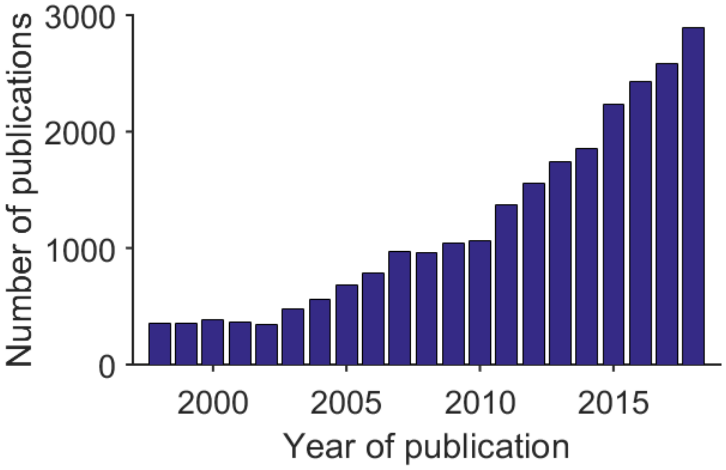

Figure 1. Number of publications containing ´flow chemistry´ in their theme field. Source: ScienceDirect.

Continuous flow operation is the most popular mode of operation with microfluidics due to its facile implementation. Flow chemistry equipment is appropriate to a lot of simple and well defined chemical and biochemical operations, and for some tasks such as chemical separation, but they are not recommended yet for tasks that require a high degree of flexibility in the manipulation of fluids. The microstructures of microfluidic channels in particular suffer from a limited ability to reconfigure the flows and a poor capacity for fault tolerance, having their tendency to clogging being identified as one of the main drawbacks of microfluidics. A big effort is being put to improve the performance of flow chemistry devices from the academic community.

1.1.1 Advantages of flow chemistry

- Reaction temperature can be higher than the atmospheric boiling point of the solvent due to the small volume of the laboratory equipment, allowing reactions to be safely run under pressure. Generally, non-compressible fluids are used without gas volume to provide a small expansion factor as a function of pressure.[10]
- The use of microfluidic flow chemistry opens the research to a wider range of temperatures, pressures and concentrations.[9, 11-14]
- Mixing can be achieved in seconds at the scales used in flow chemistry. This is highly dependent on the kind of reactors used.[9]
- Intensified heat transfer, mainly because of a high area to volume ratio. Consequently, exothermic and endothermic reactions can be performed minimising temperature changes making use of the appropriate temperature control system.

- Flow chemistry can be automated with less effort than batch chemistry. This allows for rapid operation and a reduction in the manpower required during the operations. Analysing the reactor output with different detectors makes possible to totally automate a system that could investigate a full range of different reaction parameters varying the stoichiometry, residence time and temperature.[10]
- Sequential reactions can be performed in a continuous fashion. This can be especially beneficial if the intermediate compounds are unstable, toxic or air sensitive, since they would only form momentarily and in very small quantities.
- The physical location in the flow and the temporal point of the reaction are directly related. This brings the opportunity to design a system in which more reactants are added in the flow exactly in the precise time point desired.[10]
- It is possible to design a flow system in which purification takes place at the same time that the reaction.
- Reactions in which reactants with dissolved gases are implicated can be easily managed, when in batch chemistry a high-pressure reactor would be needed.
- Liquid multiphase reactions can take place in a direct manner, with high reproducibility in a range of scales and conditions.
- Scale up can be rapidly achieved changing the reactor volume or operating multiple reactors in parallel, since the flow rates are calculated to achieve the same residence times.
- Ready accessible and scalable photochemistry. Flow chemistry provides a more efficient method to access photochemical transformations in a range of scales employing inexpensive laboratory equipment.[15-17]
- The use of microreactors for multiphase processes offers clear advantages over conventional methods due to the higher surface-to-volume ratios. Specifically, the more intimate surface contact between the different phases translates into a reduction of undesired mass-transport effects and shows improved selectivity in heterogeneous catalytic studies of multiphasic reaction systems.[18]
- Further development of methods in continuous processing was identified as the most important area of research in green chemistry and engineering for the pharmaceutical industry.[19] Flow chemistry and continuous processing contribute to the development of more efficient, environmentally friendly processes.

- Flow reactors almost always have a smaller footprint per kilogramme of product than macroscale reactors.[20]
- High selectivity. Many reactions generate more than one product from the same reactants. The use of microreactors allows choosing a reaction path over the others due to the high degree of control of the reaction conditions.

1.1.2 Disadvantages of Flow Chemistry

- Specific equipment is needed to dose in a reliable, continuous and precise manner, such as HPLC or syringe pumps and safe fluidic connections.
- Scale up some of the microeffects such as the area/volume factor is not possible and scale economy does not apply in such cases.
- It is not suitable for all kinds of reactions. Many enantioselective reactions require long reaction times to ensure high asymmetric induction. Flow chemistry is not advantageous for a wide range of time-consuming homogeneous transformations.[9]
- Solids handling, which can lead to the clogging of the reactor -especially when working with microreactors- also constitutes a limitation of this technique. Batch reactors have comparatively few issues with handling slurries and sparingly soluble reactants and products [9]. Nevertheless, it would still be possible to process solids smaller than channel dimensions.[21]
- Start up and shut down procedures have to be established.
- Safety issues for the storage of reactive material still apply.

1.2 Lab-Scale Reactor Technology

Traditionally, the optimisation of chemical reactions was a process that required expensive starting materials and was intensive in labour.[22] Furthermore, the scale-up of the results of a batch optimisation are frequently difficult as they can be constrained by mass and heat transfer limitations. Hence, microreactors have been developed over the time to overcome this limitation, providing a powerful tool for chemistry research in flow.[23] A great majority of these microreactors consist of a series of microchannels for the different reagents that eventually meet. After they meet, the reaction takes place along the microchannel –usually zigzagging in a serpentine for compactness. Microreactors have been reported to improve both mass and heat transfer rates,[24-26] apart from constituting an ideal tool to precisely control the reaction conditions.[27] The major drawback of this approach is the limited mixing at low Reynolds numbers for which diffusion is the main mixing mechanism.

1.2.1 An Overview of the Reactors Used in Automated Flow Experiments

MICROREACTORS

Microreactors have been employed in the field of chemical self-optimisation from its very early stages. Even before full experimental automation was employed, microreactors were seen as a valid tool for reaction optimisation, for which different reaction conditions were tried and the results analysed off-line.[28] In the first automated chemical self-optimisation study, Krishnadasan *et al*.[29] utilised a glass y-shaped microfluidic chip suited with a 40 cm long serpentine reactor zone. The section of the channels was 330x160 μm, and the total volume of the microreactor was ≈0.02 mL. This chip improved monodispersity of the CdSe nanoparticles generated when compared to bulk experiments. The temperature control was achieved placing the microfluidic chip over a stabilized hotplate with a high degree of thermal homogeneity.

Similarly, McMullen *et al*.[30] used a 140 μL silicon serpentine microreactor preceded by an interdigital micromixer placed immediately upstream of the microreactor to perform an automated self-optimisation. The device was built using standard photolithography and deep reactive ion etching (DRIE) and the dimensions of the channel section (400x400 μm) were larger than those of Krishnadasan. In this case, temperature control was implemented by varying the power cycle of a thermoelectric element.

Following their line of work, McMullen and Jensen[31] used in their experiments a modified version of the previous silicon reactor shown in Figure 2, for which the construction techniques and materials were kept the same. The reactor was coated with silicon nitride and capped with Pyrex to create a chemically inert environment. In this case, the microreactor included a built-in micromixer to avoid the use of an additional interdigital micromixer, reducing the number of components needed in their set up. This new microreactor was split into three different zones: a mixing zone with 200 x 400 μm channels to promote mixing; a reaction zone with 400 x 400 μm channels to act as a reaction time unit and finally, an 8 μL quench zone to terminate the reaction in the chip. Nine different reaction conditions were scaled up to a 7 mL Corning AFR (advanced flow reactor) 50-fold larger in size, the results being in good agreement with those of the microreactor under the same conditions of temperature, residence time and stoichiometry. This mesoscale commercial reactor consists of several plate reactors connected to one another, their designed based in a patented HEART cell shape that splits and recombines the fluid across its path in the reactor to maximise mixing, particularly for gas-liquid multiphasic systems,[32] although in this cases the pressure drop has been ascertained to be higher to that of a liquid-liquid system.[33] In this case, a pre-mixer consisting of a HPLC gradient mixer packed with stainless steel ball bearings was improvised and placed upstream the AFR to ensure good mixing.

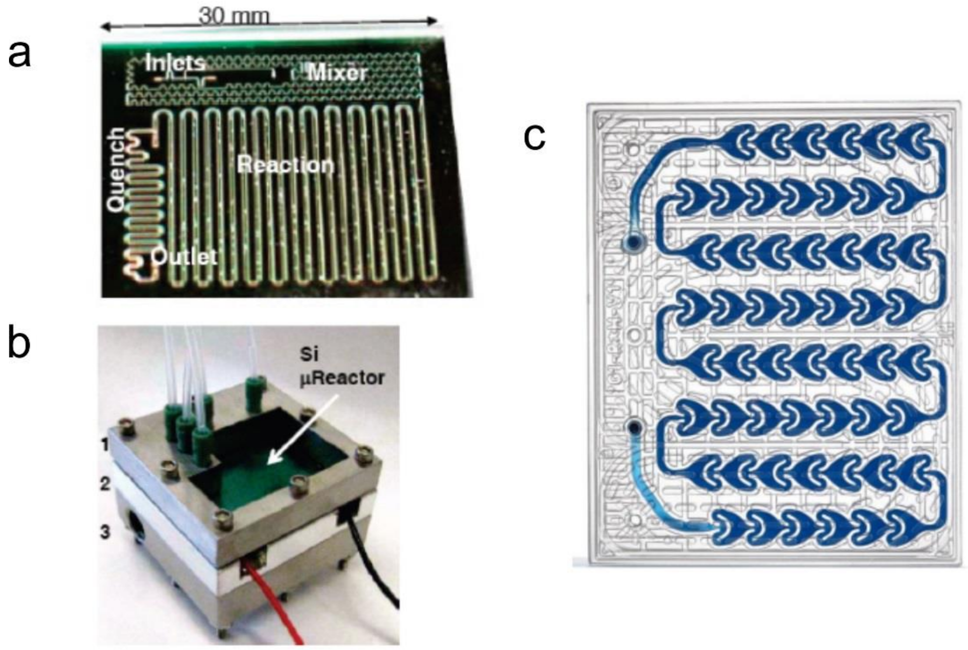

Figure 2. (a) Microreactor used by McMullen and Jensen with mixing, reaction and quench zones. Reproduced from ref.[31](b) Assembled device composed of a top plate with ports for the fluidic connections (1), housing for the microreactor and the thermoelectric device (2) and a heat exchanger to provide sufficient heat removal when necessary (3). Reproduced from ref.[31] (c) Corning AFR device with its characteristic HEART shaped cells.

In later work, McMullen and Jensen[34] carried out both a self-optimisation and a kinetics analysis in a similar microreactor. This device included an integrated mixing zone, while the reaction channel in this case had a spiral instead of a serpentine architecture. The total volume of the microreactor was 120 μL, keeping the same channel section than the previous chips used by the authors (400x400 μm). Temperature control of the device was achieved with a thermoelectric actuator for the reaction zone, and by recirculating water at 20 °C for the outlet zone, with the objective of quenching the reaction. The results were also successfully scaled up using kinetical data to a 60 mL Corning AFR reactor, 500-fold larger in size. The predicted conversion values from the kinetical data collected automatically with the microreactor agreed exceptionally well with those yielded by the operation of the mesoreactor. The same spiral microreactor was also used in an automated experimental setup with a different configuration of the laboratory equipment to demonstrate the estimation of multistep reaction kinetics.[35]

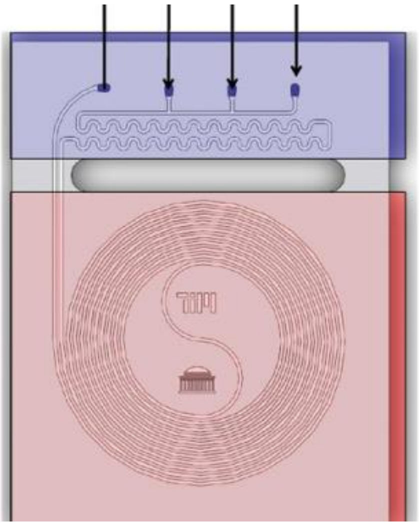

Figure 3. Microreactor utilized by McMullen and Jensen.[34] Three inlets and an outlet can be distinguished in the upper part, along with the built-in mixer (zigzagging channel). The spiral reaction zone in the bottom part of the chip is temperature-controlled, while the blue shaded area is kept at 20 °C with the aim of quenching the reaction before the fluid abandons the device. Reproduced from ref.[34]

Self-optimising approaches were rapidly adopted by various researchers and different, simpler reactors were integrated in automated flow chemistry platforms. Easy to access standard laboratory equipment such as chromatography columns were soon employed as laboratory scale reactors in automated experiments.[36-38]

PASSIVE MICROMIXERS

As commented previously for Jensen’s work, a common approach when utilising microreactors is the use of a pre-mixer immediately upstream of the reactor as a means to ensure good quality mixing of the reagents before entering the device. An interesting take on the design of micromixers can be found in the work of Hessel and co-workers.[39] The micromixer presented by the authors, namely caterpillar micromixer (CPMM) was developed by the Fraunhofer Institute for Microengineering and Microsystems (IMM) and has a ramp-like internal structure where one channel is directed up and down alternately (Figure 4). At low Reynolds numbers diffusion is the main mixing mechanism, while convection dominates at high Reynolds numbers, for which eddies are formed. However, the geometry of the channels is designed with the objective of inducing a split-and-recombine action that results in the multiplication of the number of fluid lamellae to maximise mixing even in a low Reynolds number scenario. These devices were manufactured using 3D micromilling in two different sizes for comparision with batch processing in a 1 L flask. The aqueous Kolbe-Schmitt synthesis was carried out making use of these new microreactors, and compared to the traditional flask synthesis the advantages were a reduction of reaction time by orders of magnitude and an increase in throughput by a factor of 2. The major disadvantage of this micromixer was a pronounced sensitivity to fouling.

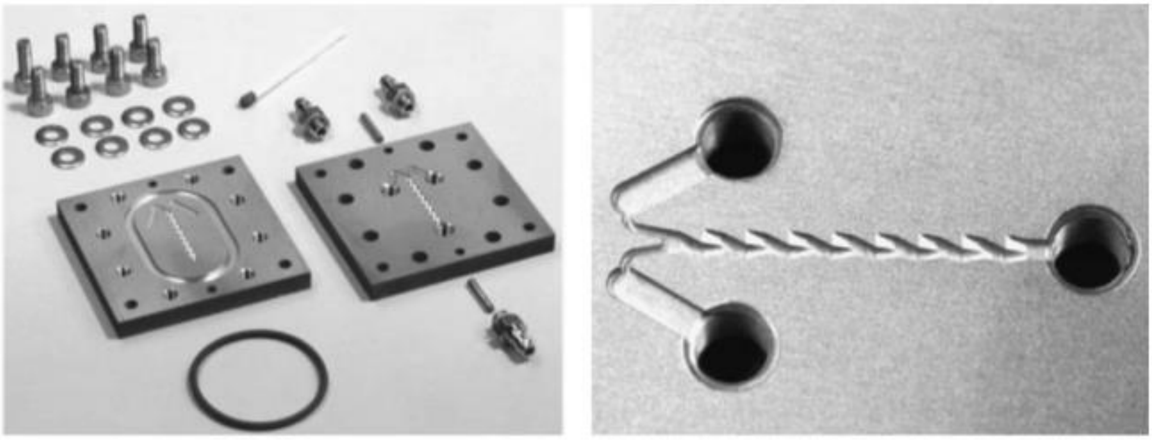

Figure 4. Disassembled CPMM micromixer used by Hessel and co-workers (left). Detail of the up-down curved ramp-type channels (right). Reproduced from ref.[39]

Similarly, Verma and co-workers[40] designed a micromixer that improved mixing by using a 3D channel system generated monolithically in poly-dimethylsiloxane (PDMS) using twisted nylon as a template (Figure 5). The steps in the fabrication of this device were presented in detail in a previous publication.[41] The nylon thread is embedded in a block of cross-linked poly (dimethylsiloxane) PDMS that is let to swell with the appropriate solvent while the thread remains unaltered. Finally the threads are removed to generate the desired microchannel. In this fashion, complex geometries such as knots or helices can be generated relatively easily using this method.

The channels were generated as helices with three or more branches, the orientation of each one changing along the length of the channel. This design led to inherent asymmetry of the channel cross-section compared with conventional circular and rectangular channel cross-sections. This was presented as a way to generate a chaotic flow even at low Reynolds numbers and compared different geometries assessing the effect of two different parameters: the channel length (L) and the helix angle (θ). It was found that the mixing efficiency increased linearly with $\theta^{2/3}$ and $L^{1/2}$. For a given geometry, and as can be seen in Figure 5 (G), as the Reynolds number grows, the mixing efficiency initially decreases (Re = 1-4) until plateauing at a constant level (Re = 5-15), before increasing at higher Reynolds numbers (Re = 15-50). This is due to a dominance of molecular diffusion over convection at low Reynolds, for which the velocity is low and consequently the residence time is longer, leaving more time to the fluids to mix. When the velocity keeps growing, the residence time keeps decreasing, but the nature of the geometry induces transverse flows. As the Reynolds number continues growing, the effect of these flows balance the decrease in residence time and the mixing efficiency remains almost constant, until a second critical value is reached for which chaotic advection is triggered and the mixing efficiency starts growing regardless of the decrease in residence time.

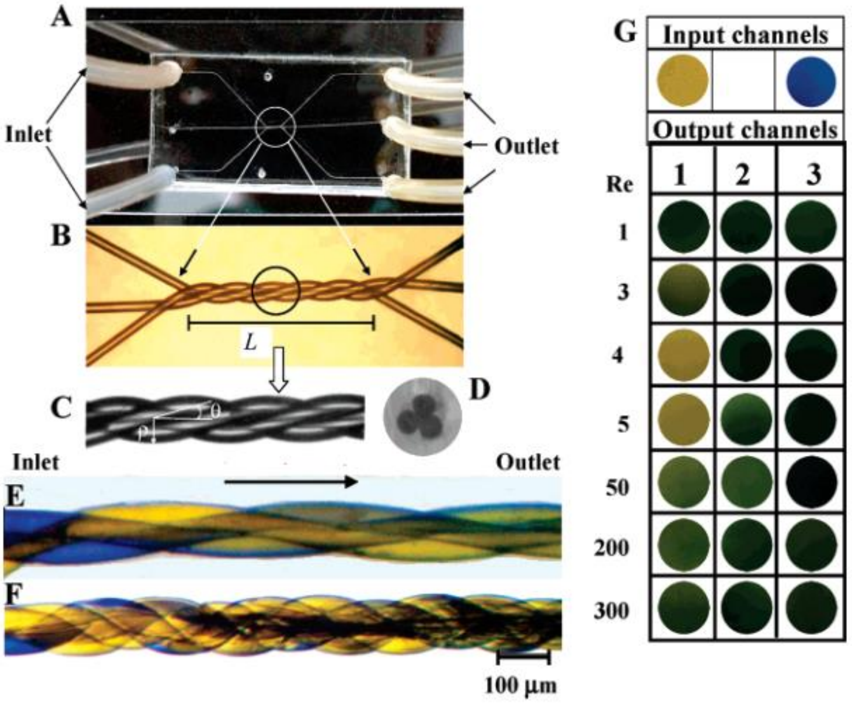

Figure 5. (A) Triple helical microchannel microreactor developed by Verma and co-workers. (B-D) Magnified top view illustrating the geometry of the channel. (E-F) Side view optical image for the progressive mixing of the liquids for different geometries. (G) Colours of the inputs and outputs under different Reynolds numbers. Reproduced from ref.[40]

Microreactors coupled with a high pressure interdigital multilamination micromixer (HPIMM) have also been used in polymerization reactions [42, 43] and their ability to produce controlled-size polymer nanoparticles ascertained. This micromixer is manufactured by IMM (Mainz, Germany) and laminates two inlet flows into 15 lamellae of 20 µm of thickness for each flow. This strategy consists in the reduction of the diffusion path to promote mixing by diffusion at low Reynolds numbers. These lamellae are arranged in a way that increases the contact area between the two fluids as compared with conventional channel micromixers. Then, the lamellae are focused into the outlet channel, which reduces the lamellae width and intensifies mixing.[44, 45]

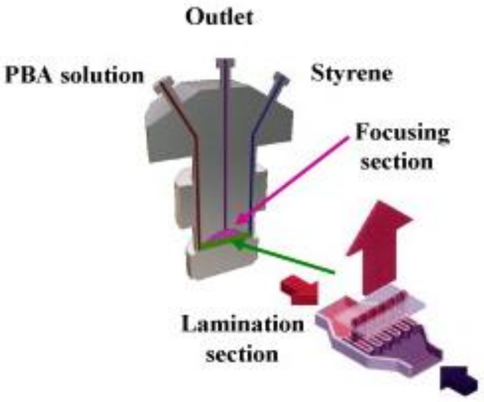

Figure 6. HPIMM multilamination micromixer. Reproduced from ref.[46]

FIXED BED TUBULAR REACTORS

Poliakoff and co-workers[14] demonstrated self-optimisation of a chemical reaction at a larger scale, since all previous work was carried out in microreactors. This study applied the technique using a fixed bed tubular reactor in supercritical carbon dioxide. Hence, this made possible to work at high pressures and temperatures, which enabled collection and analysis of volatile by-products that would otherwise have been impossible to scrutinise. This point was stressed in previous publications using the same kind of reactor and working with supercritical carbon dioxide.[47] The reactor, consisting of a 10 mL stainless steel tube (156 mm x 12 mm OD), represents a two orders of magnitude increase in volume respect to a microfluidics chip. The temperature actuators were heating cartridges within an aluminium block.

Poliakoff and co-workers' intensive work with this reactor before fully automation for self-optimisation included continuous phase separation/purification [48] and product switching,[49] among many others.[50-53] However, these early studies did not constitute examples of self-optimising rigs, as the reaction parameters were changed one at a time in a univariate approach that was not uncommon at the time.[54] To the light of these previous works, this study is particularly relevant as it constitutes the first chemical self-optimisation not carried out in a microfluidic chip as well as the first one carried out under such extreme temperature and pressure conditions, as high as 300 °C in temperature and 190 bar in pressure.[14] The authors demonstrated the versatility of self-optimisation in this reactor by optimising different parameters for a variety of reactions using different optimisation algorithms.[12, 13, 55-57]

TUBULAR REACTORS. CONTINUOUS AND SEGMENTED FLOWS.

Sans *et al*.[58] presented a simple milli-fluidic synthetic organic reactor in which they used a PTFE tubing coil as a tubular reactor, although in this case the effect of temperature was not taken into account. The tubing was cut to obtain the desired reactor volume of 3.75 mL. The simplicity of the rector employed led many other researchers to use similar approaches. A different example can be found in the work of Felpin *et al.* where the authors make use of a 5 mL reactor made of a 0.75 mm ID PEEK tubing.[59]

As Jensen's work demonstrates,[60] this approach can also be applied to the field of microfluidics by using sufficiently small tubing. A 240 µL reactor was made out of 750 µm FEP tubing and coiled inside a tailored housing the authors named 'pancake'. This housing consists of an aluminium block with a groove to fit the tubing in position. Temperature control is achieved by means of heating cartridges placed in the aluminium block. The reactor is held in position even under high pressure conditions thanks to the inclusion of an O-ring and a sheet of polycarbonate compressing the tubing against the aluminium block. Apart from the reactor design, an important novelty introduced in this work is that this reactor was operated under segmented flow conditions. In this fashion, a carrier phase separates different reacting 'slugs'. This offers comparative advantages to the continuous flow processing approach, as this method is able to screen many different reaction conditions in a small amount of time (the residence time desired) instead of waiting to achieve the steady state which is common practice in continuous flow experiments. The amount of by-products and unreacted material can be highly reduced using this approach, since the reactants are confined in the slugs, hence minimizing backmixing. In addition to that, the amount of reagents needed is much smaller than in the case of a continuous flow experiment. Nonetheless, these advantages come at the cost of a higher system complexity in which detectors are needed to ensure accurate reagent injection and sampling.

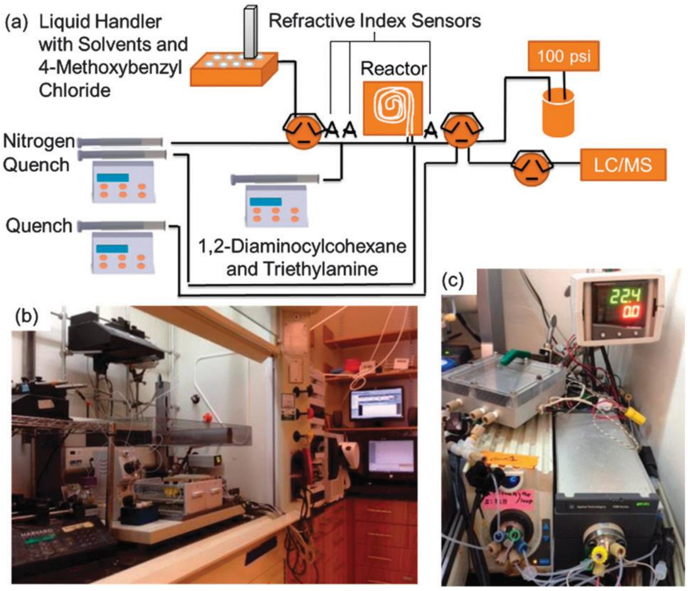

Figure 7. Slug flow self-optimisation platform used by Reizman and Jensen. (a) Diagram of the components of the system. (b) View of the system components in the laboratory, including syringe pumps, a liquid handler, LS/MS and computers. (c) Closer view of the components including the pancake reactor (top left), temperature controller (top right) and sampling valves (bottom). Reproduced from ref.[60]

Very recently, Bédard *et al*. presented a highly versatile automated system where modular interchangeable microreactor devices can be attached by the user in a matter of seconds.[61] The authors designed and demonstrated the use of a variety of reactors that fit in the bay area of their automated system, including a heated reactor (up to 120 °C), a cooled reactor (to -20 °C) a LED-based photochemical reactor and a packed bed reactor. These reactors were made using 1/16” OD PFA tubing with different inner diameters (0.02”, 0.03” and 0.04”), except the photoreactor which was constructed from PTFE (1/16” OD, 0.03” ID) The volumes of the reactors ranged between 215 µL and 860 µL. These reactors (Figure 8) are built in a similar way to the ‘pancake’ reactor discussed previously,[60] with a hosing for the tubing built by grooving a block and compressing the tubing against it by means of a glass cover.

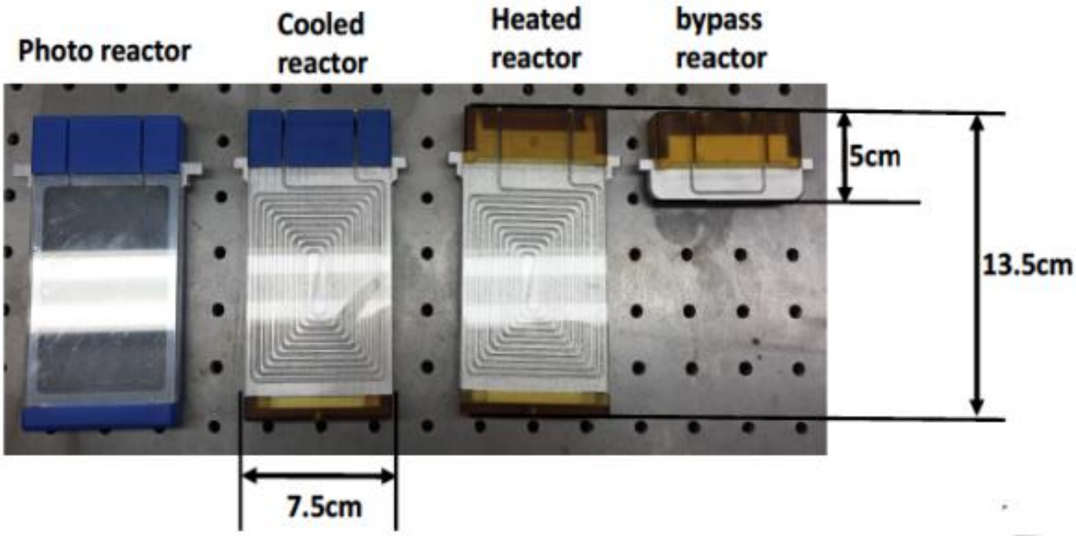

Figure 8. Variety of plug-and-play modular reactors developed by Bédard *et al*. Reproduced from ref.[61] supporting informations.

The same approach including the use and detection of slugs was used in later works,[62, 63] this time demonstrating the use of an oscillatory flow in a horseshoe shaped reactor, namely single point oscillatory flow reactor (SPOFR), made of 1/8” OD FEP clear tubing. The reacting droplets are kept oscillating in the reactor for the specified residence times until injection to analysis. In this fashion, this device is able to decouple the flow rates of the reagents and the volume of the reactor from the residence time.

From the solution of the Navier-Stokes equations for an incompressible laminar flow flowing thorough a cylindrical tube of constant cross-section, known as Hagen-Poiseuille flow,[64] it is known that the velocity profile under the described circumstances is parabolic. The maximum velocity is reached in the middle of the tube, while a no-slip condition at the walls means that the velocity is null at their location.[65] This fact is important when regarding the transport of diluted species, since the particles in a given cross-section of the tube will be flowing at different velocities. This velocity profile induces a dispersion known as Taylor dispersion.[66] This dispersion’s effect on the residence time distribution makes it diverge from plug-flow ideal behaviour, since different particles spend different times in the reactor, and can be analytically calculated from channel geometry and diffusion coefficients.[67, 68]

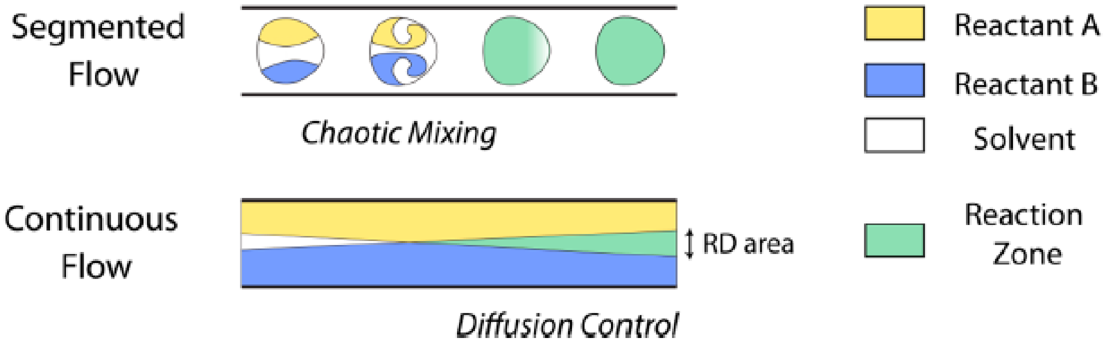

Figure 9. Segmented vs continuous flow in channels. Reproduced from ref. [69]

Making use of segmented flows, this dispersion can be effectively avoided by omitting the no-slip boundary condition at the reactor's walls,[69-72] so the particles confined in a droplet spend approximately the same time in the reactor, minimising the formation of by-products or the presence of unreacted material at the reactor's outlet.[73] While in the case of continuous flow reactors mixing is slow and only achieved by molecular diffusion, segmented flow has been proven to enhance mixing by chaotic advection in the travelling reacting slugs.[69, 74] Jensen's work demonstrates that this technique, paired with the appropriate apparatus for slug detection, makes possible to trigger the sampling loop at right time to transport a reacting slug into the online analysis.[60] This is a huge contribution in cutting times when operating automated flow chemistry systems, although it comes at the cost of a more complicated set up.

The use of 1/16” OD, 1/32” ID PTFE tubing as a reactor coiled around a Polar Bear Plus flow synthesiser has led to important advances of automated flow chemistry by Bourne and co-workers.[7, 75-77] This reactor provides active cooling in addition to heating, which greatly helps reducing optimisation times.

As Echtermeyer work demonstrates,[78] this simple kind of reactor made of commercially available PTFE tubing can also be employed as a segmented flow reactor to achieve good RTD with the appropriate apparatus for segment detection and injection. Along with Jensen's,[60] this study constitutes an example of the use of segmented flow in flow chemistry automation, although at different scales. In this case the total volume of the reactor was 10 mL and an air heater provided temperatures as high as 150 °C, with flow rates ranging from 0.001 to 10 mL/min. The slugs were detected utilising an in-line UV detector cell.

Figure 10. Polar Bear Plus Flow Synthesiser.

1.2.2 Lab-Scale Reactors For Multiphase Reactions

Processing solids in the milli and micro scale still remains a challenge. Reactors consisting of microchannels are susceptible to clogging and are not able to process solids or slurries, limiting conditions to those where reactant and product remain soluble within the flow. Flow in microchannels present low Reynolds numbers and are highly laminar: this absence of turbulence and recirculation strictly constrains the mixing mechanism to diffusion where the small dimensions of such channels make possible to quickly get rid of concentration gradients.[79, 80] Under this specific conditions, negligible effects such as particle-wall and particle-particle interactions in the centimetre or meter scale become determinant in the case of milli and microfluids.[81, 82] However, processes in the pharmaceutical industry often involve the presence of solids either as reagents, intermediates or products, and their presence is regarded as unavoidable in many important reactions,[81] ranging from the Nobel Prize winning Suzuki cross-coupling reactions [83] to the palladium-catalysed amination,[84] among many others.[85-92] In handling these solids, fixed bed reactors in small scales have been used with success, while the presence of solids in such applications as catalysts differs from that of a solid product or by-product formation resulting from a chemical reaction and having to be transported through the reactor. Under this circumstances, preventing clogging is extremely difficult in the micro and millifluidic scales or limited by the size of the solids generated, as what has been highlighted as one of the major drawbacks of the use of microreactors.[81]

Nevertheless, some techniques enable solids handling even in the case of microreactors. Segmented flows in which an organic reaction mixture is dispersed into an aqueous phase make it possible to carry out polymerization reactions avoiding clogging of the reactor.[93-95] In this fashion, the solids are confined in reacting droplets which are immiscible with the carrier phase. Other studies have applied the same technique in the production of indigo [96] and in the field of protein crystallization.[97] Xu *et al*.[98] reported the generation of monodisperse particles using this technique with a high degree of control over size and shape, achieving very narrow particle size distributions. Some innovative techniques include the use of non-contacting forces over the particles, such as acoustic,[92, 99-101] magnetic [102-104] and electrophoretic [105, 106] that have successfully been employed to effectively maintain the solids in suspension in the fluid flow.

Likewise there are a number of multiphase liquid-liquid reactions where one phase should be dispersed within a second. This has motivated recent research effort towards this goal. For immiscible aqueous-organic systems, the reactions when carried out in flow constrained in a channel generate a segmented flow consistent of droplets or slugs. The major drawback of this approach is the mass transfer resistance in the boundary between the immiscible fluids. The interfacial area plays a key role in mass transfer in such devices and a common approach has been to try to maximise this interfacial area by means of changing the characteristic length of the device.[107-109]

Furthermore, most flow reactors employed in automated studies are tubular, in common with those discussed in the previous section, and present poor mixing characteristics even for single-phase reactions. This limitation of the microfluidic devices implies that the fluids are not immediately mixed after entering a T-junction and they will only mix gradually by diffusion along the channel or pipe they are flowing through. This is the case also for coil reactors previously addressed such as the Polar Bear Plus flow synthesiser. Hence, the use of different reactors that are able to maximise mixing will also be beneficial in the single-phase reactions context.

Examples of lab-scale multiphasic reactor technologies are scarce. Nonetheless, a few examples have been identified in the academic literature and are discussed below. The relation of reactors used in automated experiments covered hitherto presented mixing characteristics directly related to the flow rates of the reagents and hence are named passive mixers. On the other hand, active mixing for which an external energy input is needed, can effectively decouple the mixing performance of a reactor from the flow rate of the reagents flowing through it. Hence, it will make possible to provide excellent mixing performance even at low flow rates. This, along with the possibility of broadening the range of reactions that could be explored and optimised to the multiphasic field, would constitute a great feat in the field of lab-scale chemical automated experiments.

An example of sequential semi-batch processes can be found in Holden's work,[110] for which a 0.5 L double-jacketed stirred reactor suited with a four-blade turbine impeller was used. Although this kind of reactor offered many advantages such as solids and slurries handling, its impractical high volume (orders of magnitude higher than any other study examined in this review) made chemical self-optimisation in the manner that it has been reviewed hitherto extremely reagent and time-consuming.

The Coflore agitated cell reactor (ACR) is based on the CSTR technology, and it includes a laterally shaking motor instead of a magnetic stirrer. The ACR consists of a plate with equal cylindrical cavities suited with a loose-fitting cylindrical stirrers inside, each one behaving like a CSTR. This commercially available reactor has been utilized by Ley and co-workers,[111] who tested its ability to handle slurries carrying out the reaction of morpholine with iodine to produce the hydroiodide salt of N-iodomorphine, since this reaction produces slurries shortly after mixing. The ACR was able to process 65.5 g of isolated product in 7.5 hours, with a 71% yield (compared to 90% in batch). In spite of that, enabling automated optimisation processes for this reactor would make possible to explore the best conditions for the production of the salt achieving even higher yields. It is important to notice that these reactors are rather big (9 mL per cell) compared with those previously reviewed in automated experiments, and will lead to unsuitable high residence times at the flow rates usually employed for chemical self-optimisation, making the process excessively long and reagent-consuming. Although an appropriate tool for multiphasic processing, this reactor is still not ideal for conducting automated experiments due to the high volume of reagents required.

Figure 11. Coflore ACR reactor block. Reproduced from ref.[111]

Baxendale *et al*.[112] made use of a Coflore agitated tube reactor (ATR) to produce triacetic acid lactone solid. This is a dynamic plug flow reactor, which presents high radial mixing. The volume is higher than that of the ACR, being each tubular reactor 100 mL. This can hardly be considered a microfluidic or millifluidic application, as the flow rates were as high as 60 mL/min during operation. However, it constitutes a good example of controlled solid processing in a tube reactor for a moderate flow rate. Similarly to the Coflore ACR, an external motor provides an agitation to the reactor in the transverse flow direction, thus inducing radial dispersion by a free moving mechanical stirrer inside the tubular reactor, as show in Figure 12. This results in high mass transfer rates in multiphasic systems, high heat transfer coefficients and good solids handling characteristics, as well as plug-flow like performance where fluid across the flow path travels through the reactor in very similar times.

Figure 12. Section of the Coflore ATR, showing the inner mechanical free moving agitator. The white arrows show the direction of agitation, transverse to that of the flow. Reproduced from the manufacturer's webpage.

Solids have also been produced in the millilitre scale using the multijet oscillating disc (MJOD) millireactor that displaces continuously a multijet disk assembly along a tube. This constitutes a rather complicated setup as can be seen in Figure 13 but was proven efficient to produce solids by Liguori and Bjørsvik.[113] The reactor is composed of four different elements: input section, reactor body, output section and oscillator. The reactor consists of two concentric stainless steel tubes (inner tube: 10 mm id, 12 mm od; outer tube: 33 mm id, 37 mm od). While the reagents are fed into the inner tube, the annular space between the two tubes is used to circulate heating or cooling liquid acting as a heat exchanger. The reactor is highly versatile, as multiple inlets and outlets can be configured for the heat exchanger, making possible the presence of different temperature zones. It is also possible to concatenate different reactor sections to modify the residence time at the same flow rate.

A shaft with 60 to 100 perpendicular discs assembled is placed inside the inner tube. The shaft is connected to the oscillator. Each disc is suited with small perforations (jets) that force the reaction mixture through them when the shaft oscillates. The outer diameter coincides with the inner diameter of the reactor tube, so the spaces between adjacent discs can be seen as different compartments connected through the jets. When the reacting mixture goes through the jets, vortices which provide extremely good mixing are generated. In this sense, the MJOD could be contemplated as a special kind of COBR, which will be discussed later in this review. The volume for a reactor consisting of 60 discs is of 38 mL.

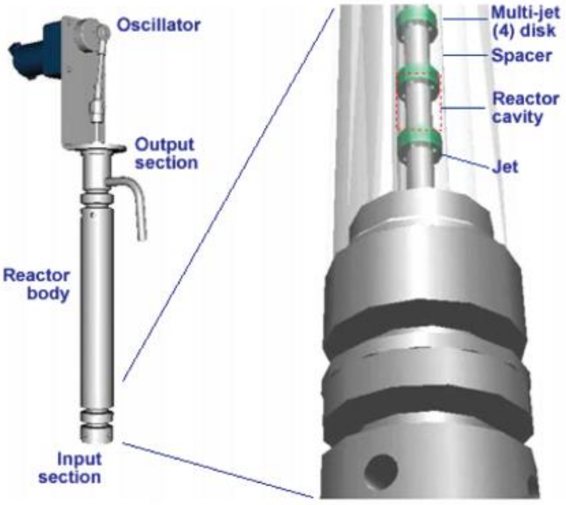

Figure 13. MJOD millireactor. Reproduced from ref.[113]

Cascades of continuous stirred tank reactors (CSTRs) are widely used in the pharmaceutical industries for continuous crystallisation in continuous mixed suspension, mixed product removal (MSMPR) crystallizers.[114] Besides, they are suited for multiphase systems [1, 115, 116] and slurries handling.[117] Mo and Jensen first evaluated the possibility of transferring this kind of reactor to the laboratory scale,[118] stressing its excellent capabilities for solid processing and high performance liquid-liquid mixing. The reactor consists of three main elements: a PTFE reactor block, a glass cover and a stainless steel cover. The reactor is a 30 x 30 mm square block, 19 mm deep. The cylindrical reservoir has a diameter of 18 mm and 10 mm depth with a volume of 2.54 mL in which a magnetic stir cross is placed. An O-ring surrounds the reservoir to provide sealing, and a stainless-steel cover along with four bolts are used to mount the device. Then, the whole setup is placed directly over a magnetic stirrer to provide active mixing. Single-phase mixing was evaluated for the device, assessing the residence time distributions for cascades of a different number of reactors, exceptionally agreeing with the theoretical ideal distributions. The authors also provided two examples of reactions involving solids. The device was tested for a solid formation reaction in which the product was insoluble in the solvent, with a high solid loading (4.4% wt) under a flow rate of 1 mL/min and with a magnetic stirring rate of 600 rpm, successfully carrying out the formation of the solid in continuous flow thanks to the action of the stir cross over the fluid keeping the particles in suspension against gravity, as well as particle-particle and particle-wall forces, thus preventing them from sticking to the reactor's chamber walls or agglomerating. The second reaction was chosen to generate a needle-shaped crystal by-product insoluble in the solvent, as crystals of this shape easily agglomerate and are avoided in industrial pharmaceutical processing when possible. At the same volumetric flow and stirring rate conditions than the previous experiment, the reactors successfully worked for 8 hours continuously in the production of the crystals before clogging, proving the small-scale version of the classic CSTR reactor a valid lab-scale tool for liquid-solid multiphasic reactions.

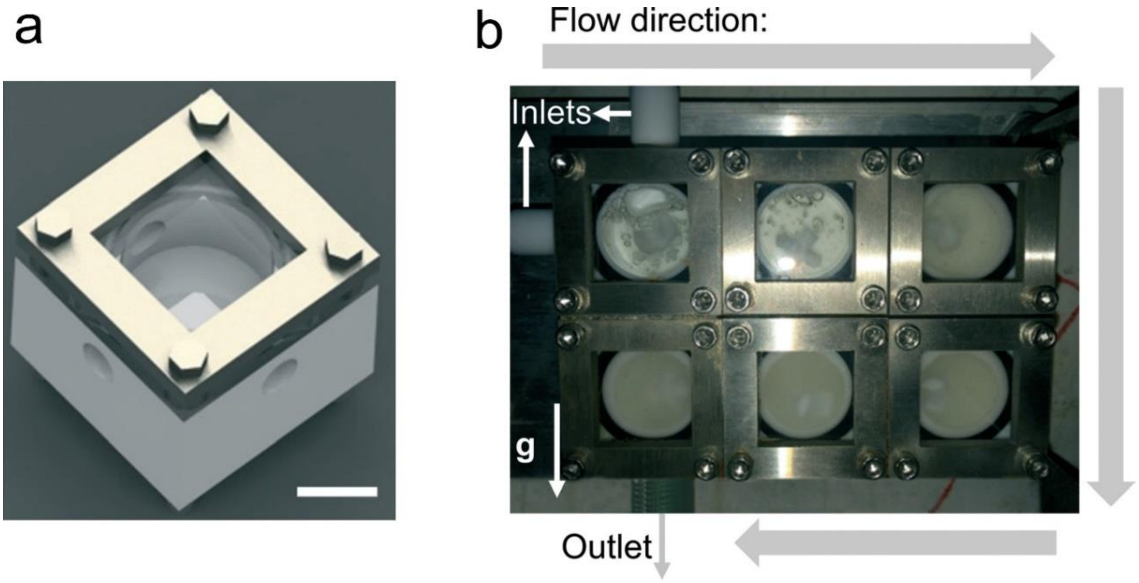

Figure 14. Miniaturised CSTR design proposed by Mo and Jensen. (a) CAD drawing of a single miniaturised CSTR. (b) CSTR cascade during operation under high solid loading. Adapted from ref. [118]

Mo, Lin and Jensen presented another similar miniaturized CSTR specially engineered for liquid/liquid biphasic reactions.[119] The main differences were the presence of baffles in the cavity or reaction chamber, the inclusion of a ceramic bearing and a thinner stir cross. The authors utilised high speed imaging to characterise the hydrodynamics of a liquid/liquid system in the reactor. Hexane and water were fed into the reactor and the hydrodynamics assessed for various rotation speeds and hexane/water ratios. It was concluded that higher rotation speeds led to smaller droplets of hexane in water (Figure 15) and thus higher interfacial areas that can give faster reaction rates. Mass transfer coefficients were estimated based on the high speed images to reveal them similar to their industrial counterparts.

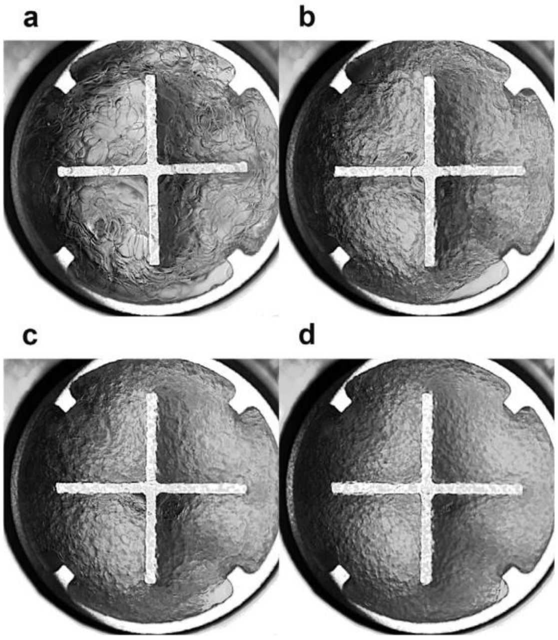

Figure 15. Hydrodynamics of a 1:1 hexane/water mixture in the miniaturized CSTR under different rotation speeds. (a) 500 rpm. (b) 800 rpm. (c) 1000 rpm. (d) 1300 rpm. Reproduced from ref.[119]

In parallel to Jensen's work, academics Nikil Kapur and John Blacker at the University of Leeds developed their own version of miniaturised CSTRs, which were named fReactors (free-to-access reactors) (Figure 16). The primary aim was to apply this technology to make lab scale flow chemistry available for researchers studying multiphase reactions, broadening the range of reactions and conditions that could be performed in flow. The plug-and-play nature of the devices rapidly attracted the attention of researchers, and the fReactor was then tested for a variety of reactions. The reactors are machined from a block of acetal plastic (which is compatible with most reagents, and also presents a high thermal resistance) and make possible to flow preheated reactants through a cascade of fReactors with a relatively low drop in temperature. Each reactor has a cylindrical 2mL reservoir in which a magnetic stir cross is placed. A PTFE gasket is placed over the body of the reactor to provide sealing compressing a lens –which makes possible to see the interior of the chamber at all times as well as perform photochemical reactions- against the block using three standard M6 bolts. In perpendicular to the chamber and from the sides of the reservoir block as many as 5 ports are drilled as inlets or outlets, to enable the use of standard HPLC fittings for the fluidic connections. In this manner, many different combinations of fReactors are possible due to their modular nature, making possible to easily configure a fReactor system for specific reactions, as a later work from Jolley, Chapman and Blacker demonstrates.[120] The residence time distribution of the device was assessed by the authors, yielding results close to the theoretical curves in the literature.[67, 121] This means that the stirrer bars are able to rapidly homogenise concentrations in the reactor's reservoir, with no evidence of bypassing or dead zones during operation.

A variety of diverse reactions were carried out in the reactors and tested against the case of conventional benchtop batch reactors. Multiphasic liquid/liquid (L/L), liquid/solid (L/S) and the gas/liquid/solid (G/L/S) systems were tested to demonstrate the ability of the fReactor to process different multiphasic system (Table 1). In 5 of the 6 reactions tested, the use of the miniaturised CSTR cascade outperformed the classic batch reactors. With residence times ranging from 2 minutes to 3 hours, the fReactor can be successfully employed in rapid reactions such as the heterocyclization as well as in slow ones like the hydrogenation, a characteristic of the fReactor the authors highlight as being unprecedented for the same geometry. Introducing active mixing in such small scales, effectively decouples mixing performance from the flow rate of the reagents, which is the reason behind the previously discussed versatility of the device under very different flow rates.

Table 1. Results for six different reactions using fReactors and their batch counterpart.

			Productivity (g L^{-1} h^{-1})	
Type of reaction	**phase (G,L,S)**	**τ_{res} (min)**	**batch**	**fReactors cascade**
IRED	L	19	6.4	17
N-chloroamine	L/L	5-50	198	826
monoacetylation	L/L	30	51	173
heterocyclization	L → L/L	2	660	1920
crystallization	L → L/S	20	8.2	31
hydrogenation	G/L/S	180	3.5	0.12

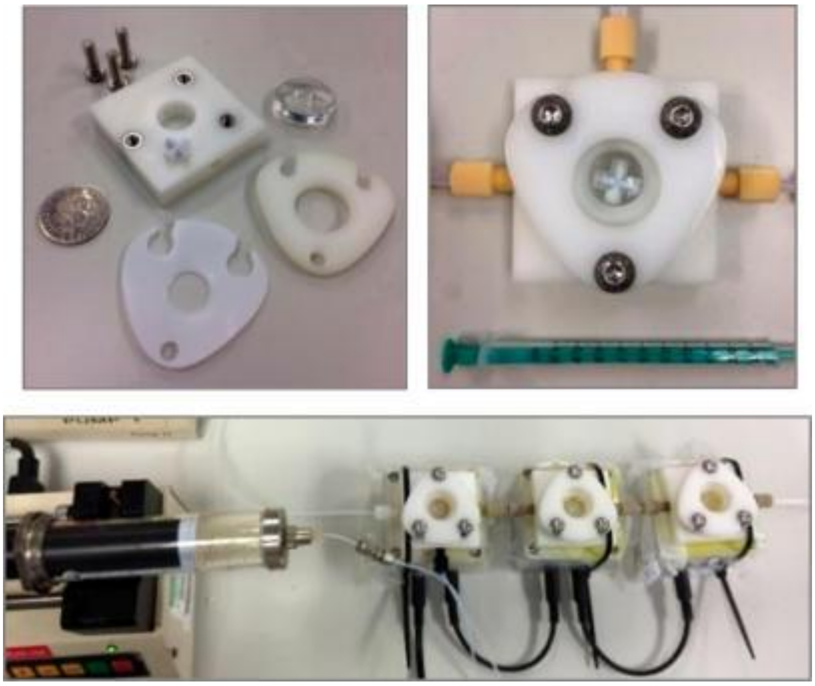

Figure 16. Top left: disassembled fReactor showing its individual components along a 10p coin for comparison. Top right: assembled fReactor. Bottom: typical cascade configuration with the fReactors mounted over their individual stirrers. Reproduced from ref.[122]

OSCILLATORY BAFFLED REACTORS

As has been previously discussed, the flow regime in micro-reactors is laminar and the mixing mechanism is dominated by diffusion due to the small scales and low flow rates employed.[123] Tubular reactors, while operating at higher flow rates and volumes at the scale of millilitres, do not provide good mixing due to low Reynolds numbers and negligible radial velocity components.[124]

Continuous oscillatory baffled reactors (COBRs) can overcome these limitations [4] to achieve near plug flow under laminar flow conditions, minimising mixing of fluid entering the reactor at different times overtaking (axial dispersion),[67] whilst providing good local mixing thereby reducing the formation of by-products. COBRs are formed from a tubular reactor with periodical internal baffles, where the mixing of the reactive species is intensified through the formation of vortices around these. An oscillatory and reversing flow is created using a pump, piston or diaphragm, and it is the combination of this oscillatory flow imposed over the constant net flow and the baffled reactor geometry that enables the generation of vortices and recirculating flow which greatly enhances mass and heat transfer.[125, 126] Figure 17 illustrates the flow behaviour at different phases within a COBR. As the flow accelerates, vortices form and expand downstream of the baffles. During flow deceleration, the vortices continue expanding until the flow reverses and they are forced into the centre of the column to finally unravel into the main stream. Then, the same process is repeated in the opposite direction.[127, 128] COBRs use several small volume cells to achieve good mixing in multiphasic reactions with near plug flow conditions at low net flow Reynolds numbers [129] and have proven effective in a range of important chemical production processes including polymerisation,[130] crystallization,[131] particle suspension,[132] multiphasic reaction [133] and biofuel production.[129]

It has been shown that near plug flow can be induced through sufficiently long cascades of continuous stirred tank reactors (CSTR) –in fact, the residence time distribution (RTD) of a COBR can be mathematically modelled like a CSTR cascade [134]- or under conditions of highly turbulent flow. A comparison of the residence time distributions for CSTRs and COBRs for a different number of reactors/cells is given in Figure 18. As the number of cells increases, the RTD curves become narrower around the mean residence time ($t/t_m=1$), indicating that different particles spend more uniform residence times in the reactor. Regarding RTD, the use of a high number of small cells COBR largely outperforms common CSTRs cascades, which even in recent miniaturised versions rarely exceed 5 reactors.[122]

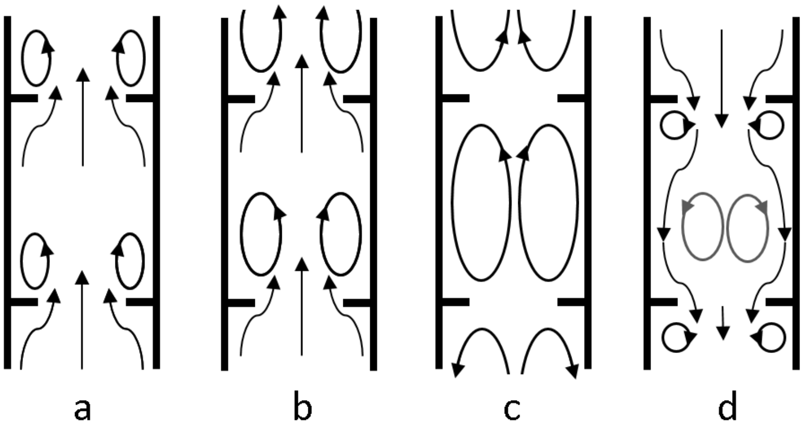

Figure 17. Diagram of the flow in a single stage of a COBR for different oscillation phases. a) Beginning of the upstroke and vortices formation. b) Flow acceleration and vortex expansion. c) Maximum vortex expansion after flow deceleration. d) Flow inversion forcing the vortices into the main flow stream as they fade. New vortices similar to those in a) but in the opposite direction starting to form. Adapted from ref.[127]

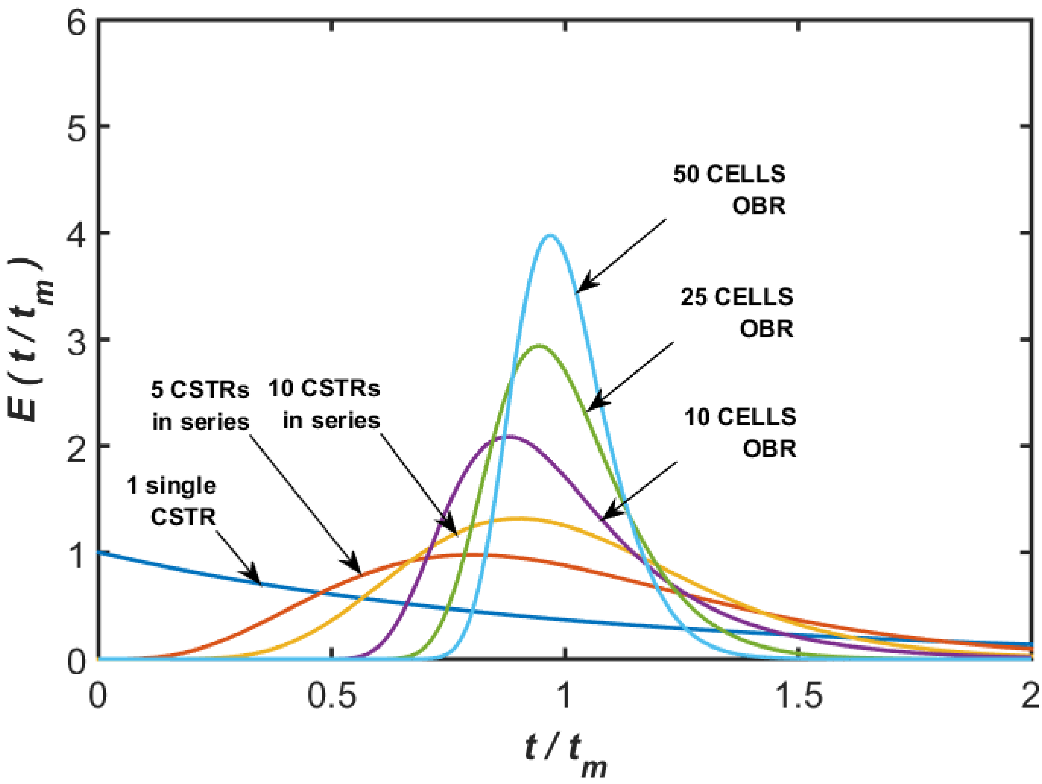

Figure 18. Residence time distributions for CSTRs and COBRs for a different number of cells. CSTR curves represent well mixed CSTRs following ref.[121] while COBR curves for 10, 25 and 50 cells were obtained from CFD simulations in this work.

Mesoscale oscillatory baffled reactors were first introduced as a laboratory scale reactor by Harvey *et al.*[135] with the objective of optimising the production of sterols in a saponification reaction, thus bringing the advantages of COBR to the area of meso-scale chemical processing. Further work concerning the characterisation of the residence times and mixing in such reactors under different fluid oscillation frequencies and amplitudes was carried out by Reis.[136] Harvey *et al.* continued the work in mesoscale COBRs by comparing the residence time distribution for different baffle designs,[137] quantifying fluid mixing properties [138] and the effect of geometrical parameters [139] by analysing the RTD and its variation throughout the flow. Their smallest mesoscale reactors had dimension of 4.4 mm outer diameter, 1.6 mm inner diameter and 7.5 mm spacing between adjacent baffles.[138] Later studies explored the feasibility of further miniaturisation of COBRs for manufacturing silver nanoparticles with a high degree of control over particle size.[140] These represented a 6-fold reduction in cell volume compared to the smallest COBR investigated previously by Harvey *et al.* Planar COBRs have previously been demonstrated to give good RTD [141] where constrictions are formed by smoothly varying channel walls.

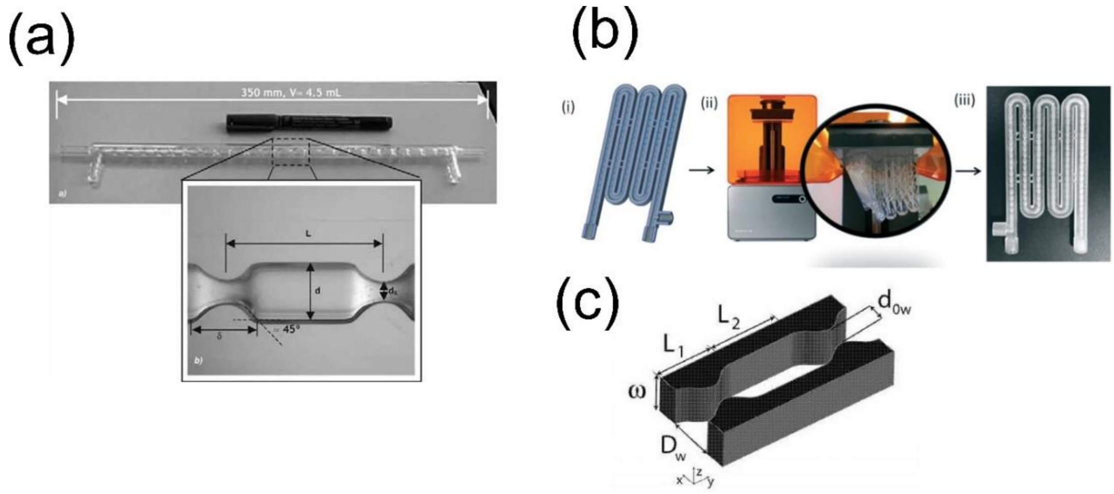

Figure 19. Different lab-scale COBR examples. (a) Reactor developed by Reis *et al.* Reproduced from ref.[142] (b) 3D printed miniaturised COBR built by Okafor *et al*. Reproduced from ref.[140] (c) Planar COBR with smooth periodic constrictions as designed by Almeida *et al*. Reproduced from ref.[141]

1.3 Analytical Techniques

1.3.1 Infra-Red Spectroscopy

An early example of the application of this technique in microfluidics is found in the work of Jensen and co-workers,[143] where a silicon microreactor was placed in a bench FTIR instrument (Nicolet FTIR 800 bench), exploiting the transparency of silicone to infrared radiation in the wavelength region of interest. The residence time and temperature of the reaction were changed in order to perform mixing studies using an acid-base reaction and compare their outcome to the results obtained from CFD simulations. Hence, the criteria for the election of this analytical technique must take into account the transparency of the micromixer to infrared radiation to be able to gather acceptable results, which reduces the application range of this technique.

The ReactIR^{TM} flow cell system developed by Mettler Toledo has been widely employed as a simple way to integrate microfluidic analysis. In 2010, Ley and co-workers used this equipment in a continuous flow processing application, monitoring in real time the reagent consumption and product formation.[37] The conclusion of the experiment was that the tool was appropriate for rapid optimisation of procedures as well as for the observation of short-lived reactive intermediates in situ –this providing further mechanistic insight into complex transformations. Further experiments were carried out by Ley and co-workers using ReactIR as the main analysis tool for multistep reactions with integrated purification employing solid-supported reagents,[144] azide synthesis and purification,[36] measure solution concentrations of carbon monoxide in a Teflon AF-2400 mediated gas-liquid contact application [145] and Grignard formation in continuous flow processing,[146] among others. A particularly interesting paper by Ley and co-workers explores the possibility of using ReactIR to precisely control the addition of reagent streams in complex reaction sequences during multi-step operations, based on the concentration of reaction intermediates.[147] This particular procedure enabled precise mixing with perfect timing, which led to a great increase in the product quality.

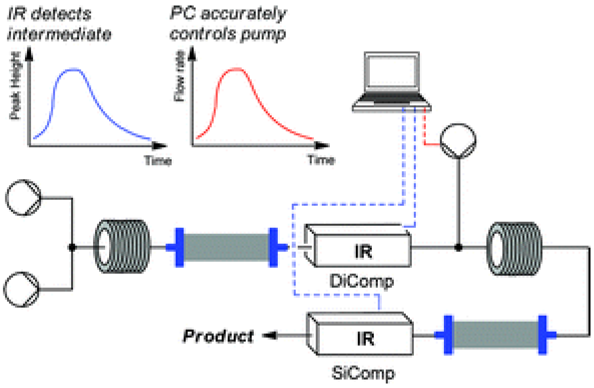

Figure 20. Setup for precise control of reagents streams in a multistep operation. Reproduced from ref.[147]

A relevant optimisation study making use of this analysis technology was carried out by Moore and Jensen[148] to maximise the production of a Paal-Knorr reaction. The advantages of FTIR over HPLC are stressed in this study, since FTIR allows to use the entire reactor effluent at the reaction concentrations rather than performing a previous dilution of a small sample, thus making possible the characterization of system fluctuations and non-destructive analysis between time unit operations. Nevertheless, the disadvantages of this analysis technique are also taken into account as any significant uncalibrated impurity could result in the alteration of the spectrum decomposition, preventing quantification. It follows that in such cases FTIR is not a valid analytical tool, with HPLC a more appropriate approach.

More recently, Moore and Jensen developed a method using conversion and residence time profiles at different temperatures to fit the reaction parameters to a kinetic model using online IR analysis.[79] Although no optimisation was performed during this study, it shows the relevance of the use of automated flow processing for sequential experiments, reporting the economy of time and starting materials and stressing its use for the generation of kinetic data.

A similar setup was used by Rueping, Bootwicha and Sugiono,[8] that used ReactIR to optimize the asymmetric organocatalytic hydrogenation of benzoxazines, quinolines, quinoxalines and 3H-indoles in continuous flow microreactors, connecting the output of their microreactor to the ReactIR flow cell, although no information about the optimization methods involved is given in the paper. The same setup was used for optimisation of photochemical reactions by Rueping *et al.* in a recent work.[149]

An excellent example of the application of IR spectroscopy for continuous flow optimisation using a ReactIR unit can be found in the work of Skilton and co-workers[55] where they used the technique coupling it with a self-optimizing feedback algorithm. In this case, online GLC data was used for calibration. The use of the FTIR as a real-time analytical technique resulted in the reduction of an order of magnitude in time and materials compared to previous studies. This feature is highly relevant since a critical limiting factor in pharmaceutical processing research is the usage of high-cost chemicals –particularly at the start of the development process when only small quantities are available.

1.3.2 Mass Spectrometry

In a pioneer publication in 1998, Adamczyk and co-workers[150] used flow injection electrospray mass spectrometry to detect the intermediates and products formed during the reaction of chemiluminescent acridinium salts under the conditions necessary for light emission. Similarly, Fabris[151] used MS for the investigation of dynamic processes in condensed phase. In both cases, the purpose of the flow chemistry system is to produce enough product for the MS analysis, which consumes all the available flow material, as mass spectroscopy is a destructive technique. There has been a considerable interest in the in-line coupling of microchip reactors to mass spectrometers for various aims such as detection of reaction product,[152] control of multicomponent reaction products and intermediates,[153, 154] reaction kinetics and mechanism studies.[155, 156]

An early example of the in-line application of this technique was published by de Mello and co-workers[157] in 2001, in which they used a silicon-machined microreactor and time-of-flight mass spectrometry (TOF-MS) (Mariner, PerseptiveBiosystems) for the generation of compound libraries based on sub-reactions of an Ugi multicomponent reaction (MCR).

Ley *et al*.[158] used an on-line miniature mass spectrometer (3500 MiD, Microsaic Systems) to monitor a continuous flow reaction. A miniature electrospray ionisation mass spectrometer was coupled to a preparative flow chemistry to monitor reactive intermediates and competing reaction paths, screen starting materials and optimise reaction conditions. This was one of the first publications to use this technique in the area of continuous flow monitoring despite being ideally suited for the application, as traditional instruments are too bulky to be coupled to flow chemistry platforms.

A recent example of the application of this technique for continuous flow optimization can be found in the work of Holmes and co-workers[77] that explores the use of online MS (Advion Expression CMS) to quantitative monitor continuous reactors, stressing its appropriateness as an analytical technique for optimising an automated flow reactor since it can determine when the steady state is reached and then calculate a product yield with minimal data manipulation. It can provide structural information and the product composition, all in real time due to its short method times. The MS signal was calibrated to HPLC using linear relative response values, and it was possible to distinguish between product adducts and isotope patterns. The method times are significantly reduced when compared to chromatography and was exploited to enable rapid optimisation.

More recently, Zare *et al.* used MS to optimise four different reactions using a Deep Reinforcement Learning approach.[159]

1.3.3 Raman Spectroscopy

Raman spectroscopy has been successfully used in the field of microfluidics flow chemistry, exploiting its capability of providing information of the quantity and type of the chemical bonds in a sample.[160] Furthermore, the Raman laser is an in-line detection technique that can be easily focused into a certain point within the microfluidic system, allowing the analysis of a very small volume of fluid.

In an early example of its application in microfluidics, Lee and co-workers[161] proved the applicability of the technique for the in situ monitoring of an imine formation reaction in a glass microfluidic chip, performing Raman measurements along the reaction microchannel as well as time-dependent measurements over the same microfluid path point.

This technique has proven really useful for the investigation of analytes in mixtures, as Raman is capable of providing accurate measurements of different analytes simultaneously. Different organic liquids can be detected with this technique, as they have a strong Raman cross-section. Fletcher *et al*.[162] used an inverted Raman microscope spectrometer to profile the spatial evolution of reactant and product concentrations for a chemical reaction within a microreactor consisting of a T-shaped channel network. Under the flow conditions used, different positions within the concentration profile correspond to different times after contact mixing of the reagents, enabling observation of the time dependence of the product formation.

Proteins such as enzymes and antibodies, as well as carriers and membrane proteins have been studied extensively in microfluidics. Microfluids have been used for the purification of proteins from other components of a cell, being this entire process continuously monitored using Raman microscopy.[163]

In the sphere of chemical synthesis optimisation, Raman spectroscopy could be used to gather real-time quantitative information on target species, allowing the determination of concentrations of chemical components in a target area. Mozharov *et al*.[164] coupled a microreactor to a Raman spectrometer to determine the order and rate of the Knoevenagel reaction. The same reaction, among other medicinally relevant ones, was monitored using Raman spectroscopy by Hamlin and Leadbeater,[165] comparing the results with the ones obtained via NMR spectroscopy. However, this publication, although relevant and extensive in its application of Raman spectroscopy for in-line continuous flow monitoring of chemical reactions, does not include an automated optimization. The authors manually tried several reaction conditions to come up with a point close to the optimum; this way of working is not only repetitive, inefficient and time consuming but gives only a rough approximation of the optimum conditions compared to the results of a fully automated optimization process.

1.3.4 Gas Chromatography

Gas chromatography was the analytical tool of choice in the research of Poliakoff's group with supercritical carbon dioxide. By adopting this technique, the collection and analysis of volatile products which under other circumstances would evaporate, becomes possible. GLC has been used by Poliakoff and co-workers for the optimization of a variety of chemical reactions in supercritical carbon dioxide where commercially available gas chromatographs Shimadzu GC17-a v3,[47, 49, 50] Shimadzu GC-14B,[12, 13] and the combined use of two instruments (Shimadzu GC-2010 and Shimadzu GC-2014) have been utilised depending on the product analysed.[57]

Felpin *et al*. made use of a GC-FID Agilent 7820A chromatograph and GC-MS analysis (TRACE GC Ultra) to optimise a Heck-Matsuda reaction, although in this case such analysis were carried offline and the optimization algorithm (a modified version of the Nelder-Mead simplex) manually fed the yield information before suggesting the next experiment's reaction conditions.[59]

Similarly, Cherkasov *et al*. utilised this analytical technique to present their open access automated flow chemistry platform OpenFlowChem. In this case, the instrument used was a Shimadzu GC-2010 with which the authors performed different hydrogenation reactions, including a self-optimisation example.[166] The authors reasoned the adequacy of the gas chromatography for liquid and gaseous products that require analysis are formed in the reactions studied.

Echtermeyer and co-workers used a specially designed sampling chamber integrated into an automated liquid sampler carousel to carry out online GC analyses for a model-based DoE utilising an Agilent 6850 device.[78]

1.3.5 HPLC

High Performance Liquid Chromatography is, with no doubt, the most extensively used technique in flow chemistry automated studies. Its availability in most chemistry laboratories, along with its versatility and low implementation time have contributed to its popularity and quick adoption by flow chemistry researchers over the last decade.[167]

McMullen *et al*.[30] utilised a commercially available HPLC (Agilent 1200 SL RRLC) in the optimisation of a Heck reaction. For every set of conditions given, steady-state data collection was ensured by waiting three residence times before the sampling loop was triggered for injection of 2 µL of the outlet stream into the analytical device by means of a six port valve. The yield of the product was calculated by comparison to an internal standard.

In a similar fashion, McMullen and Jensen[31] used an HPLC system (Waters, 1525 binary pumps, Nova-Pak C18 4 µm, 3.9 mm ×150 mm column, 2996 PDA detector, Empower software) that was integrated into the automatic optimisation rig, arguing that the incorporation of inline methods such as ATR-IR does not allow to distinguish between regio- and stereochemically different compounds. Furthermore, slight differences between the reactant and product structures can be difficult to quantify with spectroscopic measurements alone. The optimisation loop was implemented using a computer with Matlab and LabVIEW to interact with the different components of the system (pumps, heating elements) and to process the data provided by the inline HPLC in the form of chromatograms. In each step of the optimisation process, the algorithm in Matlab waited until the difference between the real temperature of the system and the set point obtained with the optimisation is less than 1.5 ˚C. To ensure steady state data collection, the system was flushed for a period of four residence times before the HPLC analysis was initialized by the computer and reaction samples injected utilizing a 6-way valve.

Bourne and co-workers have used intensively the technique in automated self-optimisation experiments with an Agilent 1100 Series HPLC instrument for a broad range of reactions including the final stage in the synthesis of an epidermal growth factor receptor (EGFR) kinase inhibitor,[7] a Claisen-Schmidt condensation [76] and a nucleophilic aromatic substitution.[75]

A system that makes use of HPLC not only to identify the optimum reaction conditions but also able to switch solvents and catalysts was reported by Reizman and Jensen.[60] In a later work, Jensen and co-workers performed a self-optimisation of a photoredox reaction in segmented flow using HPLC (Agilent 1200 series) as the analytical technique.[63]

1.3.6 NMR Spectroscopy

The application of this method represents a significant step forward, since NMR is a direct technique that requires simple methods of calibration, although at present, the low sensitivity of NMR instruments presents a limitation.[168] Some advantages of this technique that make it suitable for real-time experiments are that it is highly chemically specific, sensitive to all chemical groups, and that the signal is directly proportional to molecular concentrations, requiring only a simple calibration that is valid for every reaction and all solvents. Strikingly, high-field NMR is the only method that has not reached the compactness required to operate in a chemistry laboratory as it uses large and expensive superconducting magnets to generate the very strong and highly homogeneous magnetic fields required to maximise the sensibility and spectral resolution. Moreover, these magnets need to be constantly refrigerated as they are required to operate under their superconductivity critical temperature.

Buser and MacFarland presented an *in situ* flow method for the quantification and monitored of dissolved H_2, using a high-field NMR machine equipped with a flow probe.[169] Highly elaborated microfluidic probes for use with this kind of machines have also been reported.[170, 171]

A cost-efficient approach to this technique is the use of low-field NMR machines equipped with permanent magnets. This makes it more suitable for laboratory scale experiments, although at the cost of a loss in resolution.[172, 173]

A bench-top NMR was developed by Danieli and *et al*.[174] that based its working system in the use of permanent cryogen-free magnets (such as NdFeB or SmCo) rather than superconductors, the downside being that the field strength generated is an order of magnitude lower than the generated by the typical superconducting magnets. On the other hand, the reduction in size was of three orders of magnitude when compared to its commercial counterparts, as the arrangement of the individual magnets was based in Halbach's design.[175] However, the sensitivity achieved at this field strengths (≈ 1-2 T) was catalogued as highly valuable for the chemical analysis of simple molecules in the early days of NMR. This pioneer study allowed further research in the field and opened the possibility of self-optimisation using NMR, since the size of the device made possible its use in standard laboratory fume hoods.

The first application of this novel laboratory-scale NMR in the field of continuous flow processing was carried out by Danieli *et al*.[176], testing a transfer hydrogenation reaction by continuously pumping the reaction mixture from the reactor to the magnet and back in a closed loop. Combining flow chemistry with this technique eliminates the need for sample preparation, and owing to the progress in terms of field homogeneity and sensitivity available with the current compact NMR spectrometers, small molecules dissolved at concentrations on the order of 1 mmol L^{-1} can be characterised in single-scan measurements with 1 Hz resolution.

Sans *et al*.[58] performed the first continuous flow optimisation making use of NMR to monitor and control organic reactions in real-time utilising a commercially available benchtop NMR (Spinsolve, Magritek). Using this setup they were able to report advanced structural characterization of reaction mixtures, obtaining an unprecedented amount of chemical information in real-time. The potential of this technique was demonstrated through the optimisation of a catalytic organic reaction which employed the Nelder-Mead Simplex.

More recently, Cortés-Borda *et al.* presented an autonomous self-optimizing reactor for the synthesis of carpanone, which made use of HPLC or NMR (T-benchtop spectrometer, Spinsolve, Magritek) analysis, depending on the reaction conditions.[177] The authors reasoned that the main advantages of NMR over HPLC were the yield of richer data, including structural information that chromatography is not able to provide. The main drawback is the loss of resolution and sensitivity, and in consequence the use of NMR was restricted for concentrations of reagents higher than 0.3 M. In conclusion, they highlighted the complementarity of the two analytical techniques and the need to choose one of them depending on the characteristics of the reaction studied.

1.3.7 Other Techniques

CHARGED COUPLED DEVICE SPECTROMETRY

In the pioneering self-optimisation study, Krishnadasian *et al.* used a CCD spectrometer (S2000, Ocean Optics) to register the wavelength and intensity of the emission of the excited particles downstream of the reactor.[29] The composition of the products was not relevant in this study, only the size and dispersity of the particles' sizes.

SIZE EXCLUSION CHROMATOGRAPHY

Very recently, an autonomous flow reactor coupled to size exclusion chromatography was used to enhance process control of polymerizations by targeting molecular weights.[178] This was the first work that used a self-optimizing reactor in the field of polymer synthesis. The aim of the study was to synthesise polymers of specific molecular weights under optimal reaction conditions. The SEC, equipped with an evaporative light scattering detector (ELSD), is able to provide direct information about the molecular weight distribution and dispersity of the product.

Using this setup is possible to systematically create polymer libraries by automatically carrying out subsequent experiments targeting different molecular weights. The authors emphasised the high reproducibility that arises from the experimental results, impossible to achieve manually, and tried different polymerization reactions obtaining similar results in terms of accuracy and reproducibility. A process control mode was implemented after the optimum reaction conditions for a specified molecular weight had been achieved, to ensure steady operation and maximum product quality during production. In the case of polymerization reactions, the optimum production conditions can change over time, deviating from the optimum set of reaction conditions found during self-optimisation. The authors deduced ageing and pre-polymerization of the stocks as the main reason for deviation in this case, which the process control algorithm compensated for by reducing the residence times. In an 18 hours long experimental run, a deviation of only 1.5% in molecular weight was ascertained. Undoubtedly, this work represents a significant step forward in polymerization research and polymer libraries creation, as well as yet another different and promising application of the ever expanding field of automated self-optimisation in flow.

1.4 Optimisation Algorithms

In this section, a review of the most common optimisation algorithms used for chemical optimisation are presented. The algorithms are discussed in detail, drawing attention to their main features.

1.4.1 Gradient Descent

The gradient descent is a classical line search method, and one of the oldest minimisation methods of nonlinear functions. It starts with an initial point in the search space that the user must provide, and then the gradient is calculated. The next point will be located a given step size away from the previous one in the direction of maximum decrease of the function subject of the minimization. Formulated mathematically, if the function F(x) is defined and differentiable in a neighbourhood of a point a, then F(x) decreases fastest from point a in the opposite direction of the function's gradient vector at a. It follows that, if

$$b = a - \gamma \nabla F(a) \qquad (1)$$

for a step size (γ) small enough, then F(a) ≥ F(b), being b the next point calculated by the algorithm in its search of the minimum. Hence, the mathematical formulation of the algorithm is

$$x_{n+1} = x_n - \gamma_n \nabla F(x_n), \quad n \geq 0 \qquad (2)$$

The convergence of the method is guaranteed if F is convex and its gradient Lipschitz, and the step size is chosen via a line search that satisfies the Wolfe conditions. This implies that the selection of the step size parameter is a critical choice, as too large values will make the algorithm diverge from the minimum. One option is to choose a fixed step size that assures convergence wherever the user starts the method. A different, more refined approach is to choose a different step size at each iteration, technique known as 'adaptive step size'.

For the selection of a fixed step size, it is possible to determine the maximum step size, taking into account that any differentiable function has a maximum derivative value. If this maximum happens to be finite it is known as the Lipschitz constant and the function is said to be Lipschitz continuous.

$$\frac{\|f(x) - f(y)\|}{\|x - y\|} \leq L(f) \quad for\ any\ x, y \qquad (3)$$

Applying the same concept to the gradient of the function, if the maximum of the second derivative is finite, the function is Lipschitz continuous gradient and the maximum value of the gradient of the function will be the Lipschitz constant of the gradient.

$$\frac{\|\nabla f(x) - \nabla f(y)\|}{\|x - y\|} \leq L(\nabla f) \quad for\ any\ x, y \qquad (4)$$

It can be demonstrated then that for any $\lambda \leq 1/L(\nabla f)$, the function will be minimised with each step of the algorithm. Greater values won't guarantee convergence, being the maximum step size value that ensures convergence the inverse of the Lipschitz gradient constant.

In the case of an adaptive step size, there are methods, known as line search, that estimate the best value for the step size in each iteration after calculating the gradient. This is, the amount to move along the search direction given by the gradient. These methods choose the step size by minimising a function of the step size itself:

$$\lambda_k = h(\lambda) \tag{5}$$

Each method defines its own function, based on a different criteria. While some of them exactly minimise the function, others calculate an approximation that improves the last iteration. Three of the most relevant methods are presented below for illustration.

Cauchy[179] opted to choose the step size which minimises the objective function in each step.

$$\lambda_k = \underset{\lambda}{\operatorname{argmin}} f(x_k - \lambda \nabla f(x_k)) \tag{6}$$

This approach is called the steepest descent method and, although it intuitively seems the best possible method, it converges only linearly and is very sensitive to ill-conditioning problems. However, its importance is capital in the understanding of optimisation theory.

In 1988, Barzilai and Borwein[180] proposed an approach that minimises

$$\lambda_k = \underset{\lambda}{\operatorname{argmin}} \|\Delta x - \lambda \Delta g(x)\|^2 \tag{7}$$

$$where\ \Delta x = x_k - x_{k-1} \qquad and \qquad \Delta g(x) = \nabla f(x_k) - \nabla f(x_{k-1})$$

and in this way, it provides a two-point approximation to the secant equation that happens to be significantly faster than the classical one-point steepest descent method. In addition to that, the algorithm is also computationally cheaper and less sensitive to ill-conditioning.

A different approach to the calculation of the step size receives the name of backtracking. Backtracking is an inexact line search, since it doesn't optimise the step size, but it provides one that guarantees some descent.[181]

$$\begin{gathered} Start\ with \quad \lambda > 0 \quad and \quad \tau \in (0, 1) \\ Repeat \quad \lambda = \tau\lambda \quad until\ Armijo - Goldstein \\ or\ Wolfe\ conditions\ are\ met \end{gathered} \tag{8}$$

The conjugate gradient method calculates the new descent direction as the weighted sum of the previous descent direction and the direction of the gradient. In this manner, the common zigzagging of the descent direction which is very common in gradient descent methods with linear searches is to some extent avoided and the algorithm results less prone to become trapped during the optimum search.

It is important to notice that these methods need a differentiable function to work with. In most cases in chemical self-optimisation, derivative-free methods have been used for which models of the reaction don't have to be generated in order to calculate the objective function's gradients and the optimization procedure solely relies in the results calculated for an initial set of points and the subsequent sets of points iteratively calculated during the optimisation procedure.

Jensen's group was the first to introduce a gradient descent method in a chemical self-optimisation.[31] The authors compared three different algorithms (SNOBFIT, Simplex and Steepest Descent) in terms of performance optimising a Knoevenagel reaction, being the steepest descent the method after an initial DoE the one which optimised the yield of the reaction with a fewer number of experiments. It is stressed, though, that the algorithms were compared in order to demonstrate the variety of procedures that could be implemented in a self-optimising rig instead of assessing their convergence of efficiency.

In a later work, Jensen *et al.* aimed to optimise the production rate of a Paal-Knorr reaction making use of the steepest descent method, as well as the conjugate gradient.[148] Although different methods were used, all of them shared the same structure: first, a DoE was carried out, fitting the data to a model that allowed calculation of the gradients. After that, a succession of experiments ran following the direction of the gradient descent line until the results worsened. At this moment, another DoE is carried out around the last point that improved the objective. Iteratively, the procedure is repeated until one of the specified termination criteria is reached.

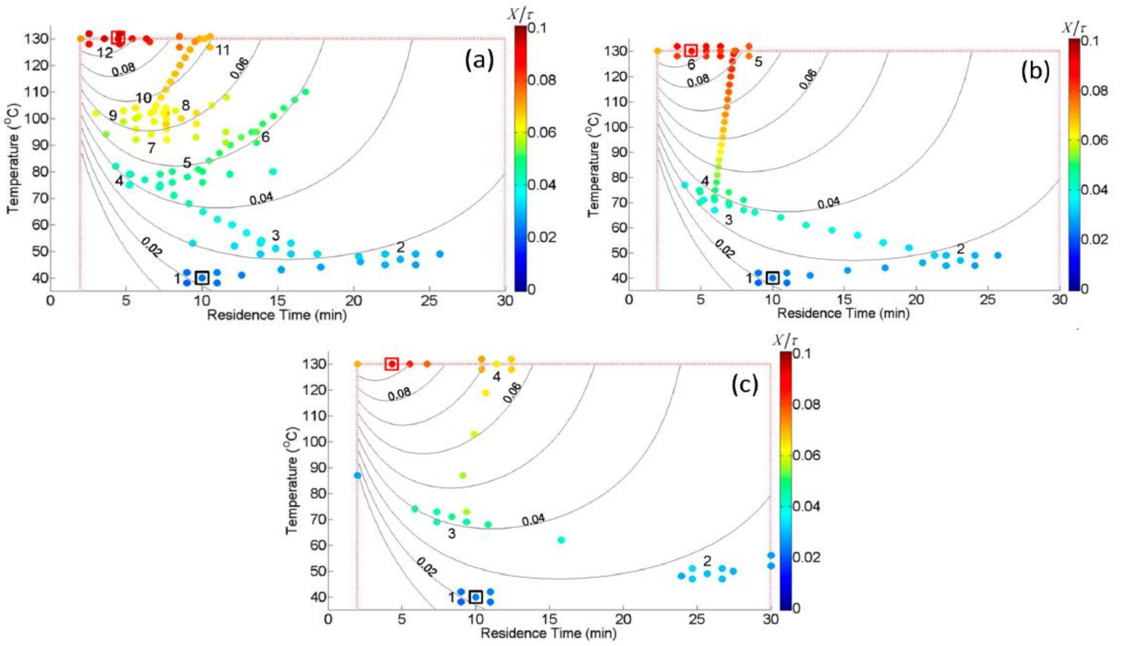

Figure 21. Routes in the maximisation of the production rate of a Knoevenagel reaction for different algorithms. Top left: steepest descent method. Top right: conjugate gradient method, which greatly reduces the number of trajectories. Bottom: conjugate gradient with an Armijo-type linear search, leading to a fewer number of experiments carried out to reach the optimum. DoE can be identified in the corners of every change of direction. Black and red boxes surround the points corresponding to the initial and final conditions, respectively. Adapted from ref.[148]

In their first experiment, the steepest descent method was implemented with a fixed step size. This was outperformed by the conjugate gradient method, which resulted in faster convergence and a lower number of trajectories towards the optimum. Nevertheless, the authors reasoned that the efficiency of the conjugate gradient method could be improved by implementing a linear search to the conjugate gradient. Hence, an Armijo-type linear search was used, effectively reducing both the number of trajectories and the number of experiments necessary to reach the optimum within the specified tolerance.

1.4.2 Simplex

This algorithm was first introduced by Spendley *et al*.[182] in 1962, which pretended to design a simple optimisation algorithm that could be applied with little effort to different fields, with a stress in the operation of chemical plants. A simplex is the k-dimensional analogue of a triangle (a geometrical figure enclosed within k+1 vertices in a k-dimensional space). The simplex is said to be regular if all its edges are equal in length. Spendley's method is based on the use of regular simplexes, being this primitive yet effective and popular algorithm often referred as Basic Simplex Method (BSM).

At the beginning, the algorithm has to be provided with an initial simplex in the design space, and then the values of the objective function are computed for its vertices. Then, the vertex in which the objective function value is the lowest is identified and reflected through the centroid of the rest of the vertices, forming a new simplex. This process is repeated until a certain vertex remains in the simplex for more than a given number of iterations M (Spendley chose M = 1.65k + 0.05k^2), when the whole simplex is contracted to continue with a finer search of the objective function's minimum.

The author stressed the simplicity of the algorithm and the possibility of its use even when precise quantitative information is not available, considering its operation mechanism drives the search taking into account solely the ranking of the observations. The simplex algorithm has been subject to modifications increasing its efficiency, namely Nelder-Mead Simplex (NMSIM) and super modified simplex (SMSIM) that have been preferred by researchers in the field of self-optimisation.

1.4.3 Nelder-Mead Simplex

Nelder and Mead proposed a modified version of the simplex method in 1965,[183] allowing the use of irregular simplexes and introducing different mechanisms for moving the simplex around in a more efficient way: reflection, contraction, expansion and shrinkage. Denoting x_{k+1} the point to be reflected and x_0 the centroid of the rest of the vertices, the algorithm works as follows:

- Reflection of the worst sample point, then the objective function is evaluated for the reflected point, this being

$$x_r = x_0 + \alpha(x_0 - x_{k+1}) \tag{9}$$

- If after the reflection the sample is still the worst, the simplex is contracted

$$x_c = x_{k+1} + \rho(x_0 - x_{k+1}) \tag{10}$$

- If after the reflection the sample is the best, then the reflected point is pushed further (expanded) in the same direction

$$x_e = x_0 + \gamma(x_0 - x_{k+1}) \tag{11}$$

- If a certain point x_1 is sufficiently old, the simplex is then shrunk to refine the search

$$x_i = x_1 + \sigma(x_i - x_1), \qquad i = 2, \dots, k+1 \tag{12}$$

Where α, ρ, γ and σ are the coefficients of reflection, expansion, contraction and shrinking. The typical values for the coefficients are shown in equation (13), being highly significant with respect to speed and accuracy of convergence. Nelder and Mead chose these values as producing the best compromise.

$$\alpha = 1, \qquad \rho = \frac{1}{2}, \qquad \gamma = 2, \qquad \sigma = \frac{1}{2} \tag{13}$$

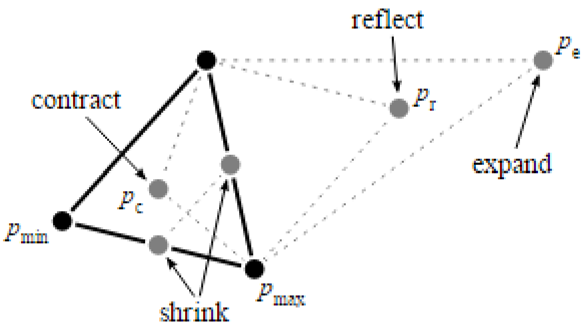

Figure 22. Nelder-Mead simplex iteration for a two-variable optimization. Reproduced from ref.[184]

This method has been one of the researcher's favourite optimization tools throughout the years, being used by different groups in a plethora of different reactions including a Knoevenagel condensation,[31] a Heck reaction,[30] imine synthesis,[58] a palladium-catalysed Heck-Matsuda reaction [59] and carpanone's synthesis [177] among others. This algorithm is generally regarded as being robust and efficient, and making use of a low number of experiments, which explains why it has been used with assiduity by different groups since the birth of self-optimising platforms.

1.4.4 Super Modified Simplex

Routh and co-workers developed a new version of the Simplex Method,[185] introducing further changes in the Nelder-Mead Simplex. The Super Modified Simplex (SMSIM) determines the location of the new vertex fitting a second-order polynomial curve to the responses at the worst response point (x_{k+1}), the centroid of the rest of the vertices (x_0) and the reflected vertex (x_r). This means that the response at the centroid must be calculated every step, in contrast with the Nelder-Mead Simplex.

Furthermore, the curve is extrapolated beyond (x_{k+1}) and (x_r) by a percentage of the x_{k+1} - x_r vector resulting in two types of curve shapes. In the concave down case, a maximum occurs within the interval. Assuming a maximisation process, the evaluation of the derivative of the curve leads to the predicted optimum whose response is subsequently evaluated, the new vertex is located at that position and the optimisation continued. If the curve results to be concave up, the response maximum doesn't occur within the interval, so the new vertex is chosen in the best boundary point according to the extrapolation of the curve. If this value results to be worse than x_r, this may remain as the new vertex and the process is continued.

The slope of the curve at the extended internal boundary may be evaluated and this value used to choose an appropriate expansion coefficient, a high slope indicating a rapid approach to the maximum (a small expansion coefficient would be enough) and low slopes indicating remoteness of the optimum and leading to the selection of high expansion coefficients.

Thus, the SMSIM allows the new vertex a vast freedom of location along the direction of the x_{k+1} - x_r vector, breaking the rigidity in the selection of the new vertex of the BS and Nelder-Mead Simplex. However, it is important to notice that if the predicted optimum response occurs very near to the centroid, the location of this vertex at this new point would reduce the dimensionality of the process and will terminate any further progress of the search in one or more dimensions. Thus, a small safety interval preventing this move is used to locate the new vertex far enough from the centroid. However, this interval can be eliminated or minimised near the optimum, as it reveals itself unnecessary when approaching it.

NMSIM is in this way improved by fitting the responses along the line through the centroid of the vertex to a second order polynomial to determine the most appropriate point to evaluate instead of relying in expansion and contraction coefficient, at the cost of evaluating one point more per simplex (its centroid) and fitting the responses.

This more refined simplex method was used intensively by Poliakoff's group in supercritical carbon dioxide. Their work features the methylation of alcohols,[12, 13] etherification reactions [14, 55] and the formation of different products from aniline, DMC and THF.[57] Poscharny *et al.* also made use of this technique in order to optimise the yield in the synthesis of oxetanes, demonstrating the possibilities of photo-flow optimisation and stressing the lack of research in this area.[149]

1.4.5 SNOBFIT

In 2008, Huyer and Neumaier[186] developed an optimisation algorithm (SNOBFIT, Stable Noisy Optimization by Branch and Fit) that addressed most of the problems that other algorithms presented, specially tailored to optimise noisy and expensive to compute functions. It combines local and global search and is a versatile optimisation tool that allows the user to specify and change the conditions of the optimisation while it is been carried out.

This algorithm internally builds local models of the function to minimise around each point, returning a set of points whose evaluations are likely to improve the models or provide better function values. One of the virtues of SNOBFIT is that it deals with a number of different issues, while the rest of the optimisation algorithms available fail to address all of them at the same time. Some of this advantages are presented below.

- SNOBFIT is a derivative-free algorithm and, as such, does not need to calculate function gradients. However, they are estimated from surrogate functions.
- The algorithm does not make use of interpolation, but the surrogate functions are fitted to a stochastic model to take noisy function values into account. This is one of the most relevant features of the algorithm, as it allows its use for noisy functions that are usual in the case of chemical systems, as the different analysis tools present noise in the readings.
- Another problem SNOBFIT deals efficiently with is the infrequent feed of function values. In self-optimizing systems, both reaching the steady-state for a certain set of reaction conditions and obtaining quantitative information of the product via analytical techniques can take a significant amount of time. It is common then, that the computer in which the algorithm is running is used otherwise or even switched off during these waiting times. To avoid the problems related to this inconvenience, SNOBFIT stores all the intermediate results in a working file at the end of every call.
- Since SNOBFIT provides a user-specified number of new evaluation points at each call, it also allows evaluating several points simultaneously, by making several simultaneous experiments or parallel simulations. Hence, it enables parallel function evaluation, a feature which has not been exploited yet in the area of flow chemistry.

- Initial evaluation points are not needed by the algorithm, which will generate a randomised space-filling design to enhance the likelihood that the surrogate model is globally valid.
- Both local and global search are possible, and the user can control which one to emphasise, as SNOBFIT allows to classify evaluation points into five classes with a different scope.
- To deal with the possibility of having multiple local minima, the algorithm keeps a global view of the optimisation by recursive partitioning of the search box.
- SNOBFIT searches simultaneously in several promising subregions and predicts their quality by means of less costly models based on fitting the function values at the safeguarded nearest neighbours.
- It also allows the user to change the search region during the optimisation process. This could be highly promising in the case of flow chemistry, as it would allow elaborating kinematic models at the same time that the optimisation is carried out.
- It is possible to evaluate the same point various times during the same optimisation, case in which the average will be computed and included in the model. This feature has proved itself to be useful in the case of evaluating very noisy functions.

The algorithm's inputs are a set of points and their corresponding function values, the uncertainties of the function values Δf_j , a natural number n_{req} specifying the number of points to be generated, two n-vectors delimiting the space where the generated points should belong $[u', v']$, and a real number $p \in [0, 1]$.

In a call to SNOBFIT, the algorithm is fed with a set of points and their function values, and then it returns the same number of points that are supposed to improve the function evaluation. These suggested points are catalogued in five different classes, depending if they have been generated by the global or local aspect of the algorithm. Points of classes 1 to 3 are generated by the local aspect of the algorithm with the aid of local linear or quadratic models. Points of class 4 and 5 correspond to the global search of the algorithm and are generated in the large unexplored regions.

The global search region $[u, v]$ will be partitioned by the algorithm, which will generate the new set of points in a subbox $[u', v']$ after each call. As it was explained before, this subbox does not need to be contained in the original search region $[u, v]$, as the algorithm automatically expands it to include the new space.

This feature of the algorithm can be exploited to explore more thoroughly a particular subbox or expand the limits of the search region.

Huyer and Neumaier coined the term box tree for a certain kind of partitioning structure of a box. The box tree corresponds to a partition of the original search region $[u, v]$ into subboxes, each one containing a point. Hence, each node of the box tree consist of a subbox and a point in it and two children, which, again, consist of a subbox (obtained by splitting the node one along a plane perpendicular to one of the axes) and a point in it (which in the case of one of them will be the same point the father contained).

During the optimisation, if sufficient points are available, the algorithm creates local surrogate models by linear least squares fits. It is important to notice that SNOBFIT was specially tailored for noisy applications, and so the use of Δn more data points than parameters. The authors recommend the use of Δn =5 as a compromise solution. If this condition is met, the algorithm proceeds to determine the safeguarded nearest neighbours of each point $x \in X$ from the set of points for which the objective function was evaluated in the past, X. The safeguarded nearest neighbours are the closest points in X to x. If the value of the function evaluated for some nearest neighbour is not available, it is substituted by

$$f = f_{min} + 10^{-3}(f_{max} - f_{min}) \tag{14}$$

Then, a local model around each point is fitted with the aid of the nearest neighbours

$$f(x_k) - f(x) = g^T(x_k - x) + \varepsilon_k((x_k - x)^T D(x_k - x) + \Delta f(x_k) \tag{15}$$

Where x_k are the nearest neighbours of x, $D = diagonal(\Delta f(x)/\Delta x_i^2)$ and g is the gradient to be estimated. To estimate the gradient, rewriting equation (15):

$$Ag - b = \varepsilon$$

Where

$$\begin{gathered} b(x_k) = (f(x) - f(x_k))/Q_k \\ A_{ki} = (x_i - x_i^k)/Q_k \\ Q_k = (x_k - x)^T D(x_k - x) + \Delta f_k \end{gathered} \tag{16}$$

And then, making a reduced singular value decomposition $A = U\Sigma V^T$ and replacing the ith singular value Σ_{ii} by $\max\,(\Sigma_{ii}, 10^{-4}\Sigma_{11})$, to enforce a condition number less than 10^{-4}. Then the gradient is estimated as

$$g = V\Sigma^{-1}U^T b \tag{17}$$

A local quadratic fit is computed then around the best point, i.e. the point with the lowest objective function value, and a point w of class 1 is generated by minimising the local quadratic model. In order to do so, SNOBFIT makes use of another minimization algorithm, the bound constrained quadratic programming package MINQ, which is designed to find a stationary point, usually a local minimizer. The point will be part of the output of the algorithm, as a suggested evaluation point for the next iteration.

The rest of the points are generated according to a system based on the use of the boxes and trees that was described before, that allows the algorithm to keep searching in the whole design space. A point of class 2 is a guess for a putative local minimizer; by local point, the author implies all the safeguarded neighbouring points that satisfy $f < f_1 - 0.2\ (f_2 - f_1)$. These points are calculated minimizing an auxiliary objective function that includes the gradient and an estimation of the standard deviation. This ensures that the class 2 point generated is not too far from the region where higher order terms missing in the linear model are not relevant. All the points obtained in this fashion from local points are classified as class 2 points, while the points obtained in the same way from nonlocal points are taken to be in class 3.

Points of classes 2 and 3 are alternative good points for the search of the local minimizers, although they have some significance in the global search too, since they might occur near points throughout the available sample.

Class 4 points are generated in so far unexplored regions, in large subboxes of the current partition. They represent the most global aspect of the algorithm. The input parameter p denotes the desired fraction of points of class 4 among the points of class 2 to 4, letting the user define the degree with which the algorithm is focused on its global aspect of the search.

Points of class 5 are generated only if the algorithm is not capable of reaching the desired number of points of classes 1 to 4 (a case that could happen if SNOBFIT does not have at its disposal a sufficient number of points available to build the local quadratic fittings). These points are chosen from a set of random points such that their distances from the points already in the set are maximum, extending X to a space-filling design.

As said before, the branching property of the algorithm partitions the search region into subboxes, each one containing exactly one point where the function has been evaluated. After each call to the algorithm, a new set of points will be generated, some boxes will contain more than a point, and a splitting algorithm will be required.

To split the subbox $[\underline{x}, \overline{x}]$ of $[u, v]$ containing the pairwise distinct points x_k, such that each subbox contains one point. In the case of a pair of points, the *i* with a maximal $\left|x_i^1 - x_i^2\right|/(v_i - u_i)$ is chosen and the subbox is split along the ith coordinate at $y_i = \lambda x_i^1 + (1-\lambda)x_i^2$, where λ is the golden section number ρ if $f(x^1) \leq f(x^2)$ and $\lambda = 1 - \rho$ if not. This way, the subbox with the lower function value gets a larger fraction of the original box, making more likely its election for the generation of a class 4 point.

For each subbox, SNOBFIT calculates its *smallness*

$$S = round\left(log_2 \frac{\overline{x_i} - \underline{x_i}}{u_i - v_i}\right), \qquad (18)$$

This quantity gives a rough measure of how many bisections are necessary to obtain this box from the original search space. Hence, the smallness will be null for the exploration box and it will become larger for smaller boxes. For a thorough global search, boxes with low values of S should be preferably explored when selecting new points for evaluation.

This algorithm has been widely used in chemical self-optimisation by a variety of groups and reactions. First used by Krishnadasan *et al.* in the first example of a self-optimising platform,[29] the use of the algorithm was rapidly embraced by Jensen's group [30, 61] and specially by Bourne's group, which has been able to optimise a variety of reactions using SNOBFIT such as the final step in the synthesis of EFGR kinase inhibitor,[7] the amidation of methyl nicotinate [77] and more recently a Claisen-Schmidt condensation.[76] Cherkasov *et al.* report the optimisation of a custom objective function for hydrogenation reactions.[166] This approach has been usually employed to include various (and usually competing) objectives in the optimisation, as will be discussed in the next section.

1.4.6 Multi-Objective Optimisation

Process development frequently faces the challenge that the presence of different objectives poses. While the processes usually aim to maximize economic benefit, other conflicting objectives could be present such as minimizing the amount of waste generated during operation (E-factor).

In self-optimisation studies, these conflicting objectives have usually been included in the same objective function that a single-objective optimisation algorithm minimizes. This approach has been used by Sans and Cronin, aiming to maximise the yield at the highest concentration of aldehyde and at the lowest possible residence time.[58]

$$J = \frac{A_{im}}{A_{im} + A_{ald}} \, x_1 \, t_r^{-1} \qquad (19)$$

Equation (19) shows the custom objective function the authors used. A_{im} and A_{ald} are the integrated areas of the imine and aldehyde in the experimental spectra, respectively. x_1 is the volumetric fraction of the aldehyde and t_r the residence time.

Similarly, Cherkasov *et al.*[166] utilised a tailored objective function, including a squared term for the product yield and an extra term for the flow rate. This codifies mathematically the desired objectives and its relative importance, since a higher yield will always be preferred, but also a higher flow rate if this increase comes at no cost for the yield. In this way, the objectives are ranked according to their importance by means of including powers in the objective function.

However, this kind of approach leads to a single optimum rather than to a set of solutions. The optimal set of solutions of a multi-objective optimization is known as the Pareto front.[187] The main characteristic of a Pareto solution (or nondominated solution) is that it is not possible to improve one of the desired objectives without worsening another.

A series of algorithms have been developed over the last decades that are able to iteratively search for the Pareto front of solutions of a multi-optimisation problem, among which NSGA-II is generally regarded as the most popular, with many other novel algorithms being based on it.[188] NSGA-II is a genetic algorithm that aims to find the Pareto optimal set of solutions from a population of samples, creating an offspring from the best performing individuals.[189] Then they are ranked, and the process continues iteratively with the best individuals.

What differentiates NSGA-II from single-optimisation algorithms is the nature of its ranking system, which is split into two separate steps. In the first step, sets of nondominated individuals are identified comparing every solution with every other solution in the sample and counting how many individuals are dominated by them and how many dominate them. This process continues after the first non-dominated front has been identified, using the previously calculated data from the pairwise comparison and splitting the sample into different fronts F_1, F_2, … , F_i, according to their nondominance level (ie. solutions in F_1 are not dominated by any other solution, solutions in F_2 are only dominated by F_1, etc). The next generation of μ individuals is formed by the union of the best fronts in terms of nondominance. If i_μ represents the first index for which the number of solutions in the union of the first i_μ groups (ν) is higher or equal to μ, then two scenarios arise: if μ is equal to the number of elements in the first i_μ groups, then the next generation is formed by F_1 U F_2 U …F_{ik}. If, on the other hand, if $\mu < \nu$, another ranking process is needed to select the remaining individuals extracted from F_{ik+1} that will form part of the next generation, aiming to preserve diversity.

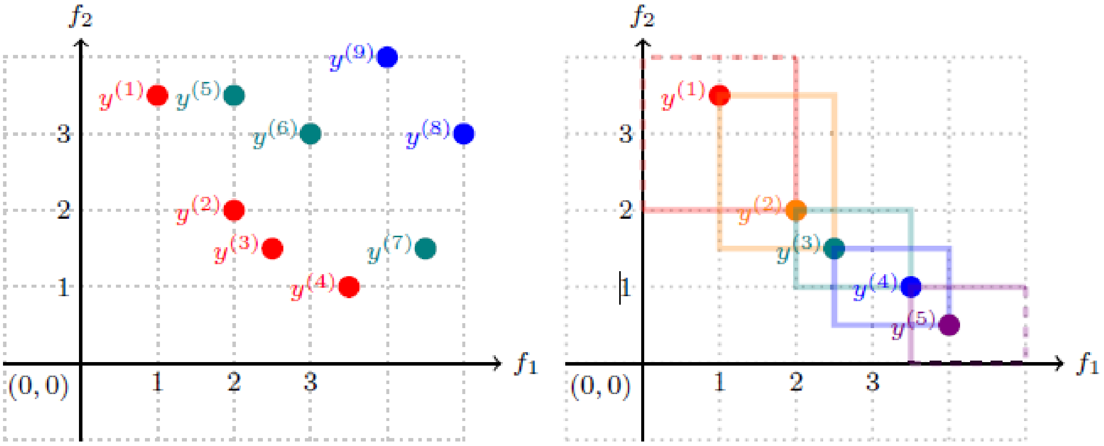

Figure 23. Left: nondominated sorting example. Red points represent individuals not dominated by any others (F1). Green points (F2) are only dominated by F1. Right: crowding ranking criteria. For a given level of dominance, points are ranked according to their crowding, which in the case of two conflicting objectives is proportional to the perimeter of the rectangles whose corners correspond to the positions of the nearest neighbours. Reproduced from ref.[188]

In this second step of the ranking process, samples in F_{ik+1} are sorted according to their crowding distance, which is a good indicator of diversity. The aim of this ranking is to maintain a good spread of solutions in the obtained set that forms the next generation. Then, the process starts over again from this set.

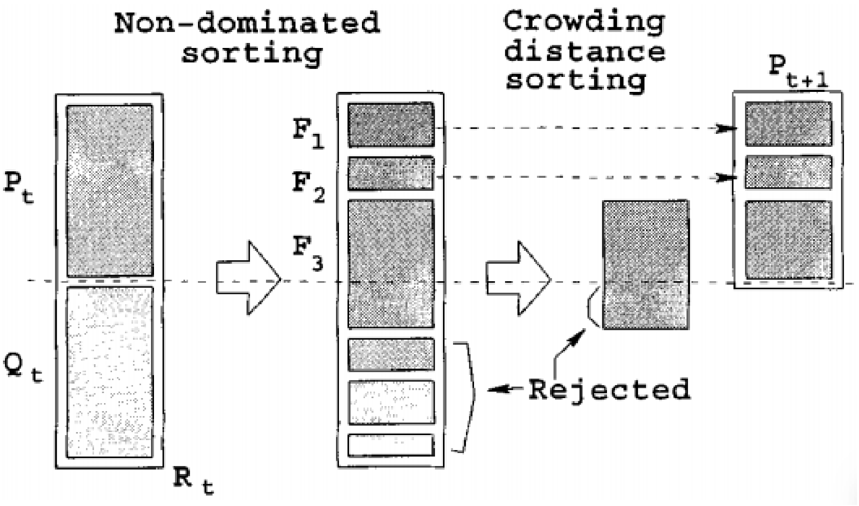

Figure 24. NSGA-II procedure. From a population P_t and using a genetic algorithm, Q_t is formed. Then, a non-dominated sorting takes place, where the best fronts are selected to form part of the next generation. After that, if necessary, a crowding distance sorting is carried out to select the rest of the individuals that form the next generation P_{t+1}. Reproduced from ref.[189]

The main drawback of the NSGA-II and generally of population-based methods is the large number of function evaluations needed, which constrains their range of application to cheap-to-evaluate functions. In the case of expensive black-box optimization, surrogate-based algorithms have been developed to minimise the number of function evaluations needed during optimization. In this kind of approach, the optimization is carried out on the surrogate models generated from an initial DoE, hence identifying promising evaluation locations using them. Then, experiments are carried out for these locations, updating the surrogate model with them and repeating this process iteratively.[188] Important examples of surrogate assisted multi-objective optimization algorithms include ParEGO [190] and the development of surrogate assisted MOAS by Emmerich *et al.,*[191] who used Gaussian processes surrogate models (Kriging) to assist an evolutionary algorithm. More examples and detailed accounts of these kind of algorithms can be found in the reviews by Emmerich[188] and Forrester and Keane.[192]

A more recent and refined multi-objective optimization algorithm, namely TSEMO, has been recently developed by Bradford *et al.* making use of spectral sampling, Gaussian processes and a genetic algorithm.[193] The authors generate Gaussian process surrogate models for every objective function at each iteration. Spectral sampling is then utilised to sample the GPs and obtain linear predictors for them, which is a condition required by Thompson Sampling.[194] NSGA-II is used on the cheap to evaluate GPs samples to come up with a set of locations that form part of the Pareto front of the objective functions previously obtained by spectral sampling. This set of locations forms the candidate set for the next evaluation, from which the location which yields the higher hypervolume increase will be selected for evaluation. After evaluation, the GPs are updated with the response and the process is repeated for a specified number of times. When compared with other algorithms, TSEMO outperformed NSGA-II, ParEGO and EHV for 4 out of 9 test problems, coming close second in 3 more, which demonstrated the potential of the algorithm. As can be noted from TSEMO, the novelty does not lie in the methodology for the search of the Pareto, which relies on a genetic algorithm, but in the combination of GPs surrogate models, hyperparameter training, spectral sampling and the hypervolume criterion.

Very recently, Schweidtmann *et al.* tested TSEMO in a self-optimization rig for two test reactions: S_NAr and N-benzylation, in which the conflicting objectives were the maximization of the space-time yield and the minimization of the E-factor in the case of the S_NAr or the percentage of impurity in the product for the N-benzylation.[75] This study is the only example yet of automated multi-objective self-optimisation, providing valuable information about the trade-off between the different objectives for both reactions, as well as the GP models that can be employed for the study of the reaction's landscape.

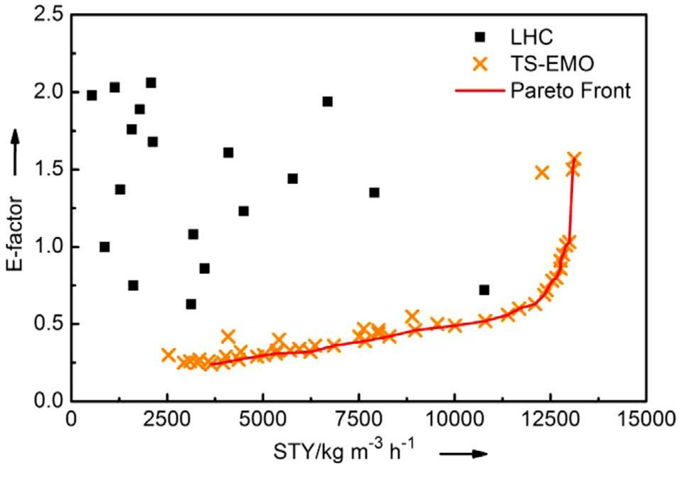

Figure 25. Experiments carried out by TSEMO for a nucleophilic substitution reaction. Black squares represent the Latin hypercube DoE from which the algorithm starts. Orange crosses correspond with the subsequent experiments carried out to determine the position of the Pareto front. Reproduced from ref.[75]

1.5 Examples of Chemistry Self-Optimisation

The earliest example of the automation of a continuous flow reactor for self-optimization of a process can be found in Krishnadasan *et al*.,[29] which used an automated microreactor to synthesise fluorescent CdSe quantum dots. The continuous synthesis of CdSe nanoparticles in a microreactor was previously demonstrated by the same group, which improved monodispersity when compared to its synthesis in bulk.[195] The system engineered by the authors used a microfluidic reactor where CdO and Se solutions were mixed and the nucleation and growth of the nanoparticles took place along the microchannel. An inline CCD spectrometer was utilised to monitor the emission spectra of the emergent particles. The acquired data was fed into the SNOBFIT control algorithm implemented in Matlab, which reduced each spectrum to a scalar 'dissatisfaction coefficient' and then updated the reaction's conditions in an effort to minimise this coefficient and drive the system's output towards the desired goal. This dissatisfaction coefficient consisted of a weighted product that aimed to simultaneously minimize the deviation from a target wavelength and maximize the intensity of the peak in the spectra, assigning different weights to these objectives that determine the relative importance of the two. The experimental equipment, consisting of two syringe pumps and a hot plate, was controlled using LabVIEW to modify the process parameters. By means of modifying the flow rates of the CdO and Se solutions, both the molar equivalents and the residence time can be controlled, while the hot plate was used to set the microreactor's temperature. Although introducing a ground-breaking experiment, the authors of this study stressed the need to develop heuristic termination criteria. This would allow the algorithm to assess if the global optimum has been reached, in which case the search would be finished and the laboratory equipment shut down.

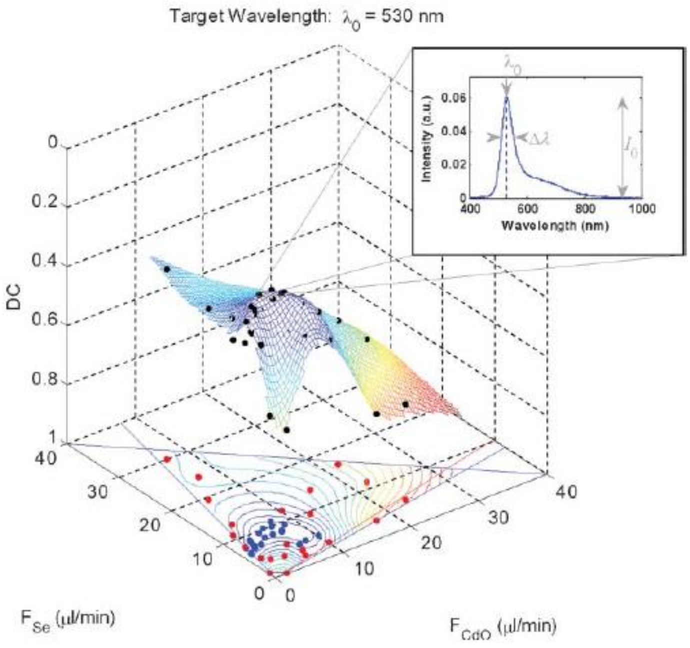

Figure 26. Graph showing the dissatisfaction coefficient for a target wavelength of 530 nm for different injection rates of CdO and Se solutions at a constant temperature of 230 °C. The spectrum for the optimum experiment that minimized the dissatisfaction coefficient is shown in the exploded graph. Reproduced from ref.[29]

After some years, the approach of Krishnadasan and co-workers started to be adapted by numerous groups, with the work in the area steadily growing ever since. McMullen *et al*.[30] developed a method for the self-optimization of a Heck reaction using the Nelder-Mead Simplex Method. This system was designed to maximise the yield of a Heck reaction by adjusting the equivalents of alkene and the residence time, but not the temperature of the microreactor, which in this case was maintained at 90 °C. The equivalents of alkene were controlled by adjusting the relative flow rate of the syringe pumps used, meanwhile the residence time was set by adjusting the total flow rate delivered by the three syringes used. The analyses were carried out using an inline HPLC. A schematic of the system can be found in Figure 27. Once the optimisation was carried out in a microreactor, it was scaled-up 50-fold and performed at the optimal conditions found in a mesoscale flow reactor (Corning Advanced-Flow Reactor), the results being in good agreement with the ones obtained in the microreactor. This is a good example of how multivariable optimisation in microfluidic flow chemistry can be used to find the optimal conditions for the production of a product and then easily scaled-up to industrial production, highlighting the advantages of both flow chemistry and self-optimization platforms and their suitability for process development in the fine chemicals and pharmaceutical industries.

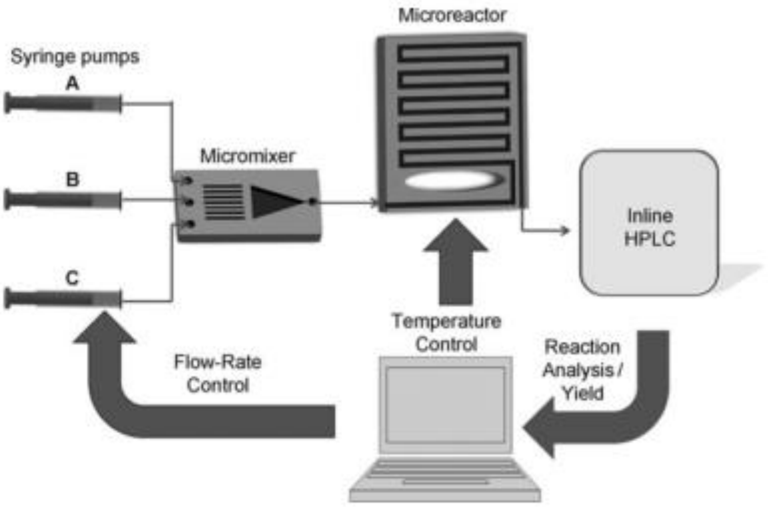

Figure 27. Automated optimisation system used by MacMullen and Jensen. Inline HPLC provides a spectrum of the product, from which the yield of the reaction is calculated. Then, NMSIM provides the conditions of the next experiment, which the computer sets by changing the flow rate of the syringe pumps. Reproduced from ref.[30]

Following their research in this area, McMullen and Jensen[31] performed optimisations for the Knoevenagel condensation and the oxidation of benzyl alcohol to benzaldehyde by chromium trioxide. For the condensation reaction, they used three different optimisation algorithms, comparing the results obtained by them. In each optimisation step, the temperature and residence time were varied to maximise a weighted objective function that took into account the experimental and theoretical maximum yields and flow rates for the reaction shown in eq. (20). The aim of selecting this objective function was to look at shorter residence times without sacrificing yield. This is a good example of how the objective function can be other than the reaction yield to include more parameters. In addition, the objective function could be changed to include coefficients taking into account the cost of the reactants, weighing it to experimentally find the optimum conditions that minimise costs and residence times without sacrificing yield.

$$f_1 = \max_{T,\tau} \frac{Q_{exp} Y_{exp}}{Q_{max} Y_{\max\,(theory)}} Y_{exp} \qquad (\ 20\ )$$

Three different optimisation algorithms were used to find the optimum reaction conditions of the Knoevenagel condensation: Simplex, SNOBFIT and the Steepest Descent Method. Both Simplex and SNOBFIT are non-derivative methods that don’t require a mathematical expression of the function to be optimised. To apply the Steepest Descent Method, which is a derivative method, a response surface model is needed. This method converges faster, the main drawback being that the function being optimised is an approximation of the real system based on the experimental values of the points taken into account during the modelling of the response surface. In this case, the program performed a two-level factorial with three repeats at the centre and fitted a quadratic model to the responses. A summary of this optimisation results shown in Figure 28 can be found in Table 2.

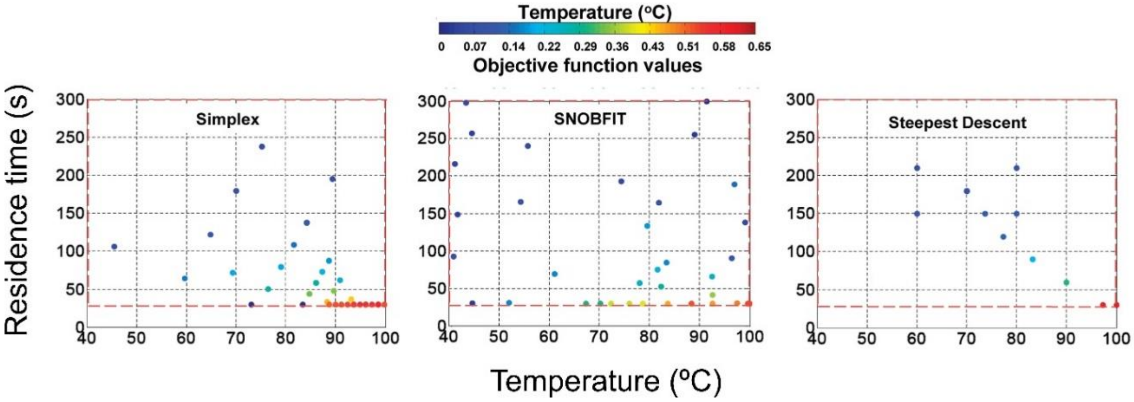

Figure 28. Optimization results from the Knoevenagel condensation. In this case, the steepest descent reached the optimum within a lower number of experiments. Adapted from ref.[31]

Table 2. Optimisation results for Knoevenagel reaction. Data from ref.[31]

Optimal point

Algorithm	Total experiments	T (°C)	tau (s)	Objective function value	Yield at optimum	Total time (h)
Steepest Descent	13	100	30	0.60	77%	4.5
Simplex	30	99	30	0.58	76%	8
SNOBFIT	36	99	30	0.54	74%	11

Poliakoff *et al*.[14] explored the possibility of application of this technique for continuous reactions in supercritical carbon dioxide, which is the main area of expertise of the group. In particular, this study demonstrated the methodology valid for optimisation of the yield of the target product, as well as optimising for multiple products from the same reaction mixture and operate on a larger scale (mL/min) than previous experiments reported in the literature.

The setup consisted of HPLC pumps to pressurize the reactants and CO_2 which pass through the reactor and a back pressure regulator (BPR) to achieve pressure control. The temperature was controlled making use of K-type thermocouples and a PID heating control, although no further information was given about the heating element used in the setup. The pressure was measured by pressure transducers and the product composition by on-line GLC. A schematic of this setup can be found in Figure 29 (left).

All these data were used as inputs to the control algorithm, which modified the optimisation parameters (temperature, pressure and flow rate of CO_2) in order to maximise the objective function (which in this case was the yield of the reaction product). The control algorithm used was based on the super modified simplex method (SMSIM). The flow rate of the organic pump was kept constant at 0.2 mL/min and was not subjected to change, as it was not included as an optimisation parameter in the study, although it would be in future publications of the group.[12]

Various reactions were optimised using the setup, with a focus on the dehydration of ethanol over (γ)-alumina because it is well understood in the literature and also formed the basis of the proof-of-concept experiments of the original automated reactor of the group.[47] It is also a good candidate for testing the optimisation procedure because it produces multiple products, including the desired product diethyl ether, ethane, and acetaldehyde. This provides the opportunity of optimising the yield of different products at the same time, and this is the reason why the optimisation stops the process at only 87% conversion of (γ)-alumina, because increasing the conversion also would reduce the selectivity to diethyl ether and increase the formation of by-products. The optimisation path is illustrated in Figure 29 (right).

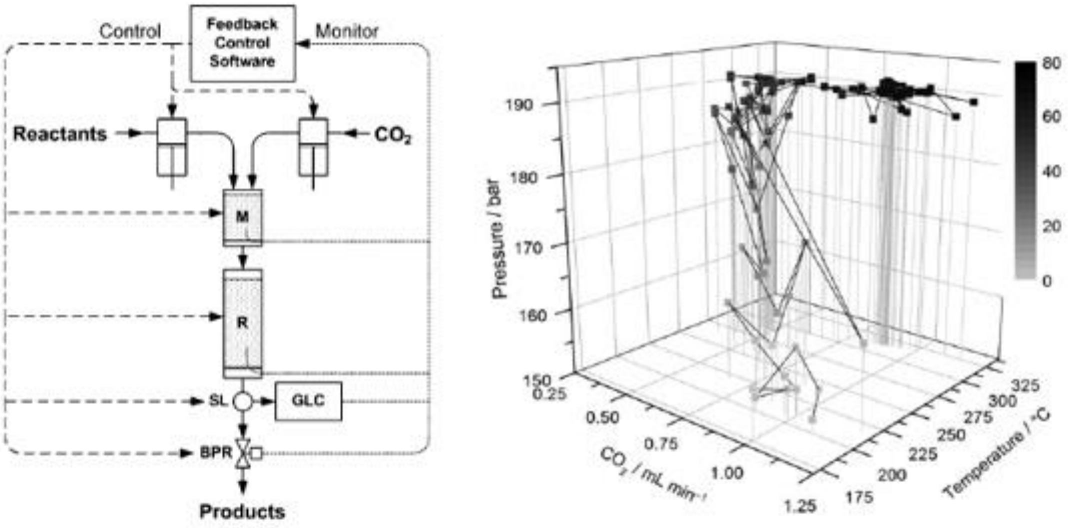

Figure 29. Setup used by Parrot and co-workers (left) and optimisation path of the experiment (right), where the shading of the points indicates the yield of the desired product diethyl ether. Reproduced from ref.[14]

The authors reasoned that the optimisation time is a function of the thermal characteristics of the reactor and the time required for the analysis, so the procedure could be accelerated by the development and implementation of new heating techniques and faster analytical methods. This would lead to better optimisation times, saving time and costs in reactants to obtain the optimal conditions values for the production. However, developing efficient and rapid heating tools would not lead to better optimisation times by itself unless the heating stage is the bottleneck of the process. Otherwise, the amount of time taken by the optimisation would be mostly determined by the rate of the product analysis regardless of how quick is the heating process.

In a later publication, Poliakoff and colleagues[12] performed an optimization for the continuous methylation of alcohols in supercritical carbon dioxide using again the Super-Modified Simplex algorithm, after optimizing manually for optimum yield reactions in continuous reactors using supercritical carbon dioxide, which is a laborious method and usually doesn’t allow to optimize all the reaction variables due to time constraints. This system used online gas chromatographic analysis (GLC) and the optimisation was performed through modifying four parameters: (i) temperature, (ii) pressure, (iii) carbon dioxide flow rate and (iv) ratio of the methylating agent. This publication is particularly complete, as it provides accurate and detailed information about the automation system used and the way it operates, as well as data relative to the communication between the software and the different setup elements, including an explanatory exposition of the optimisation algorithm employed.

The setup was used for optimisation of acid-catalysed reactions in CO_2, specifically the dehydration of ethanol over (γ)-alumina and the production of pentyl-methyl ether from 1-pentanol and dimethyl carbonate (DMC).

In the first case, the etherification was optimised with the following four parameters: (i) temperature of the preheater and the (γ)-alumina catalyst bed, (ii) pressure, (iii) flow rate of CO2 and (iv) the molar ratio of the 1-pentanol to the methylating agent (MA, either DMC or MeOH). This ratio was varied controlling the different flow rates using HPLC pumps; one pump delivered a constant flow of 1:1 molar ratio of 1-pentanol and MA, while the other pumped a flow of MA that was the parameter controlled by the optimisation algorithm from the PC. Temperature and pressure were monitored using a PicoLog data logger coupled with K-type thermocouples and RDP pressure transducers, respectively. The communication between the computer and all the rest of the elements was implemented using serial communication as all of them used RS-232 modules.

Once the different values for the optimisation parameters and the yield had been received, they were processed using the Super Modified Simplex in Matlab to determine the reaction conditions for the next experiment in which the yield (the objective function for this experiment) would be evaluated using GLC again. This iterative method led the following experiments to a greater yield. The results obtained were excellent, as a 98% yield was obtained. The study also concluded that the results were better when using DMC as the alkylating agent.

A similar self-optimising experiment was carried out by Bourne *et al.* in 2013,[55] where the optimisation process was increased in speed by changing the characterization technique to FT-IR, since the previous process took more than 2 days for a four parameters optimisation, mainly due to the time required for the GLC analysis. However, this change of the analytical technique has some limitations and the system could not be pressurised because of the more restrictive pressure limitations of the IR probe. The system was completely solvent-free and required further calibration, unlike previous IR systems reviewed before.[37, 148] It combined the online FT-IR with GLC to calibrate for the varying concentrations and independent validation of the FT-IR analysis. The calibration was carried out with an additional pump providing the reference compound, ethyl acetate, and adjusting its flow to the same flow rate of the 1-pentanol. This paper constitutes an excellent demonstration of continuous calibration when combining various analytical techniques. A 3^3 grid of experiments for different 1-pentanol flow rates, DMC flow rates and reactor temperatures was conducted, duplicating the measurements using GLC analysis and ATR using the Mettler-Toledo ReactIR simultaneously. This made possible to identify peaks and quantify the relative amount of pentyl methyl ether and ethyl acetate, and in consequence calculate the yield. The spectra of the FT-IR was correlated with the measured yield by GLC, generating a calibration curve that was used to calculate the yield in real time from FT-IR spectra.

Using the SMSIM algorithm once again, the optimisation was complete within 150 minutes, with an average of 3.2 minutes for each data point compared to 35 minutes on their previous work using GLC.[12] The optimisation was also carried out using the SNOBFIT algorithm as shown in Figure 30, which can be slower than SMSIM but presents a greater confidence that the global optimum has been located. However, the faster sampling rate allows a greater number of experiments to be performed.

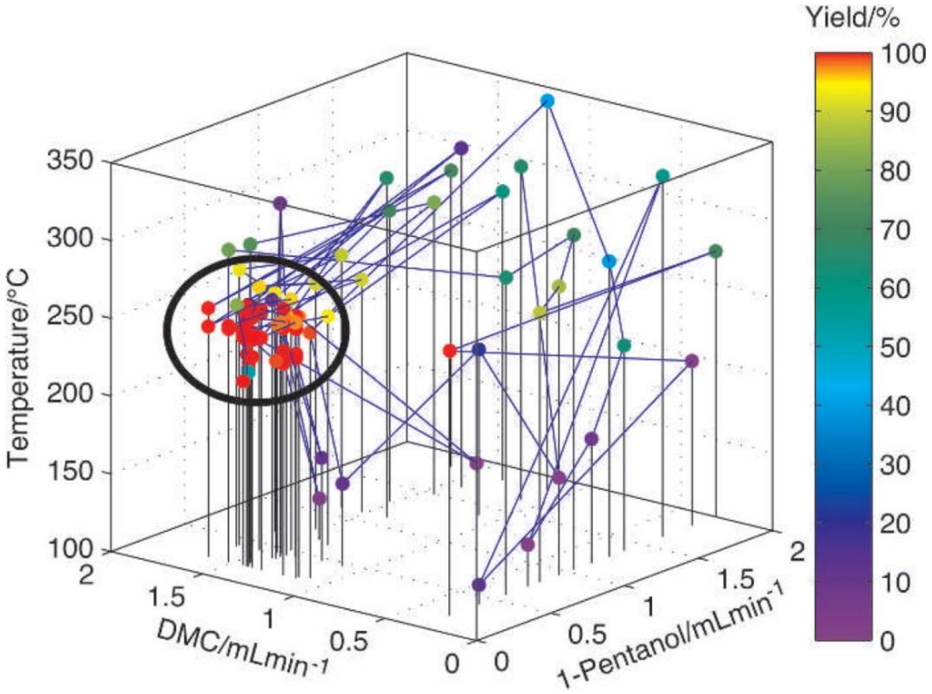

Figure 30. SNOBFIT optimisation path for the continuous methylation of 1-pentanol with DMC in supercritical CO_2. Reproduced from ref.[55]

The outcome of the experiment reflected that both algorithms were able to locate the conditions for the maximum yield which, in both approaches, presented a conversion of 99%. Interestingly, the optimal conditions located for these yields were different, suggesting that the optimum yield for this reaction is not located in a point but rather in a region of the search space. Further research in green chemistry could address the problem of selecting the point in this optimal region that minimises the energy input of the system or the production of toxic or contaminant by-products.

The speed of the FTIR analysis encouraged the authors to carry out a map of the whole region, taking advantage of both the speed of the analytical technique and the automation of the system to measure the yield under 252 different conditions spaced at regular intervals (Figure 31). This demonstrates the enormous potential the automation of flow chemistry, enabling visualisation of the impact of the different reaction conditions and their interdependence, and thus making possible to reach a new level of understanding in chemical processing.

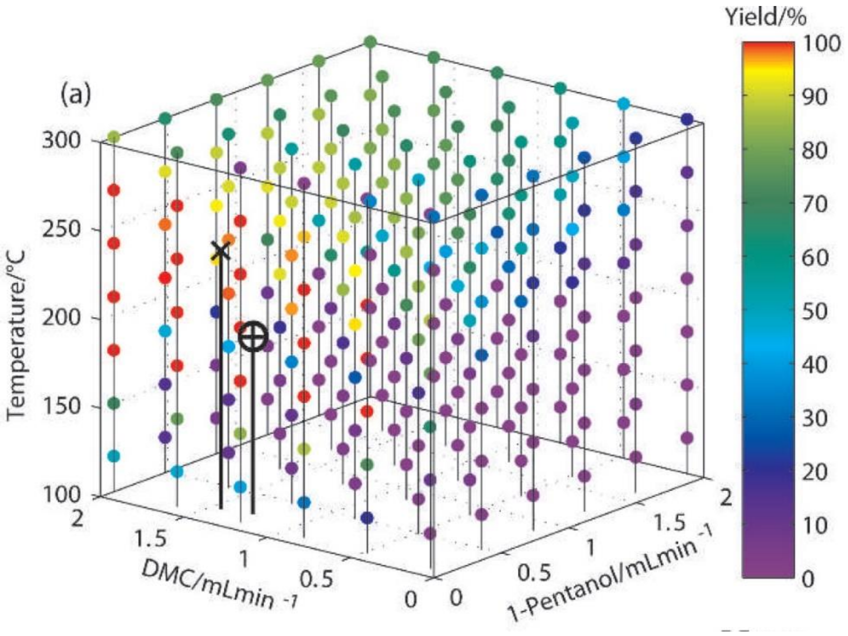

Figure 31. Mapping of the whole search region with the optimal conditions reached by SMSIM (+) and SNOBFIT (x) algorithms indicated. Reproduced from ref.[55]

Holmes and co-workers[77] performed an optimisation using SNOBFIT and DoE. For both approaches, the system recognised a steady state when the deviation of the three last samples was less than 0.75%. SNOBFIT optimisation took 12 hours and 21 experiments to reach the optimum, leading to a maximum yield of 93%. The DoE-based model was developed from a CCF design, and the response surface was obtained statistically. To improve the efficiency, the reaction conditions were performed in order of ascending temperature and then randomised within these blocks, as waiting for the heating and cooling of the reactor (Polar Bear Plus Flow Synthesizer) was the biggest contributor to the total optimisation time. Further improvement of the system could integrate quicker, more efficient heating techniques and configurations. The model for the composition of the product stream was generated by creating a saturated polynomial including all square and interaction terms and then manually removing any nonsignificant terms using MODDE v 10.1.1 (Umetrics). The model was generated based on the value of the yield under 14 different reaction conditions (Figure 32) and the data gathering with this configuration took 5.5 hours, leading to a maximum yield of 96%.

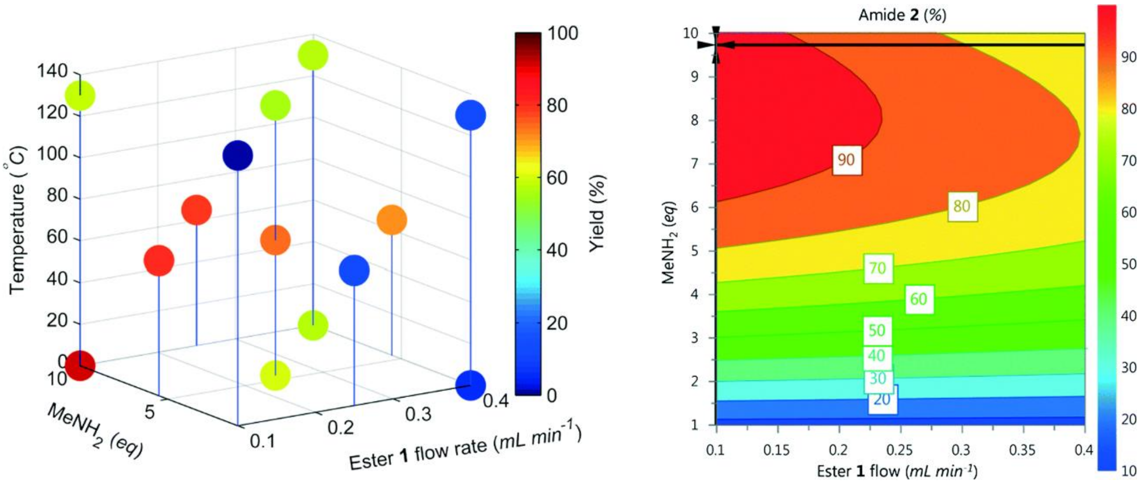

Figure 32. DoE for the RSM (left) and Contour Plot for the optimum conditions derived from the CCF model. Reproduced from ref.[77]

Whilst it is true that SNOBFIT took longer, it is important to stress that it is an algorithm that ensures that it has reached the global maximum, meanwhile the optimization based on the RSM of a CFF DoE gives only an approximation. However, in the case of this publication, the two showed good agreement, which reveals DoE and surrogate modelling as a powerful tool in the study of reaction´s landscapes.

The optimization of the final bond-forming stage in the synthesis of EGFR kinase inhibitor AZD9291 was carried out by Holmes *et al.,* demonstrating the applicability of this approach to complex systems, including telescoped reactions for the synthesis of pharmaceuticals.[7] The setup includes three HPLC pumps for the reagents, which meet in two mixing tees before entering the tubular reactor. The reaction mixture then leaves the reactor and enters the sample loop, which delivers an aliquot of reaction mixture to the mobile phase of the HPLC without the need of a previous dilution. The optimisation algorithm used for this experiment was SNOBFIT, being the flow rate of aniline, the molar equivalents of acid chloride and trimethylamine, and the reactor temperature the variables.

The bottleneck of the process is the time in which the reactor (Polar Bear Plus) reaches the temperature set point, being the biggest contribution to the overall optimisation duration, this being the reason why the SNOBFIT algorithm was chosen. Furthermore, SNOBFIT allows generating sets of experiments that can be performed in order of ascending temperature. This algorithm presents another feature that is exploited in the experiment, this is; since SNOBFIT provides a set of points to analyse, the next reaction can be started as soon as the previous one has provided a sample to the analysis. In this case, HPLC analysis was chosen as it is ideally suited for quantification of the complex range of species generated. In order to minimise the quantity of reagents used during the optimisation process, the flows were reduced to 0.02 mL/min until the reactor reached the temperature set point required.

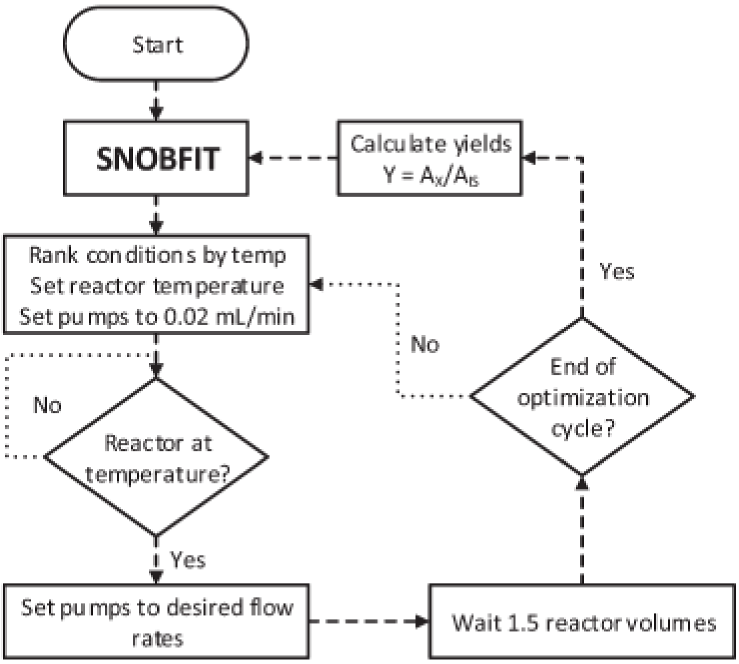

Figure 33. Algorithm used to minimise reagent consumption for an SNOBFIT-based optimisation developed by Holmes *et al*. Reproduced from ref.[7]

1.6 Research Gap

From a detailed review of the literature, the impact of the proposed research will extend the reach of the optimisation technology currently constrained to single phase flow to multiphase flow in chemical automation. The design of new reactors will also allow the processing of solids in flow, opening the possibility of particle-size optimisation or even of automated crystallization studies in flow. This is particularly pertinent as it will extend conditions of flow based laboratory studies to those found in the majority of industrial chemical manufacture.

As has been highlighted numerous times in the literature, slow temperature controls are often found to be the bottleneck of self-optimizing platforms,[7, 14] especially of those operating at the millilitre scale. Hence, laboratory based optimisation processes would benefit from the development and integration of quick and precise temperature controls, which would lead to reducing development times and minimizing the amount of the often costly reagents employed during optimization.

Equally, very little attention has been paid to response surface modelling in the context of chemical self-optimization, which in most occasions is limited to the use of polynomial cubic models fitted by least squares [7, 31, 148] or to the fitting of *a priori* known models.[78] In consequence, the inclusion of different surrogate modelling techniques, coupled with various cross-validation procedures could help in the automatic discovery and study of different reactions' outputs, as well as yield more robust models.

In summary, the work proposed in this thesis explores the possibility of the design and integration of multiphasic millireactors in a self-optimization platform. These reactors will be suited with a tailored quick and efficient temperature control to help reduce the major bottleneck identified in the literature, as well as being capable of carrying out sequential automated multiphasic experiments in flow for the first time. An emphasis will be made in employing different metamodeling techniques rather than standard polynomial cubic fittings, as well as the use of cross-validation procedures to provide more robust models.

1.7 Aim and Objectives

From the literature review and the research gap identified, the aim and objectives of the present thesis were defined. Mainly, it will introduce multiphasic processing capabilities, broadening the variety of reactions susceptible of being explored in an automated fashion, by means of the design and integration of new reactors. Secondly, surrogate modelling and cross-validation techniques largely ignored hitherto in the context of reaction optimisation will be implemented by making use of a Machine Learning approach, in order to outperform the least squares fitting of polynomials that was used in this context until now. In addition, these same techniques will be applied to the CFD-enabled design of a COBR reactor to achieve optimal behaviour in terms of residence time distribution and mixing performance.

1.7.1 Research Aim

The aim of this project is to develop flow reactors for automated exploration of reaction space. The work will focus on two platforms: (i) the fReactor CSTR and (ii) a continuous oscillatory baffled reactor. A key aim is to underpin both with optimisation techniques. In the case of the fReactor, such techniques will focus around the single and multi-objective optimisation of the output of multiphasic reactions; in the case of the COBR, they will be used to optimise the design of the device paying attention to different flow features.

1.7.2 Research objectives

From the research gap identified in the previous section, it was decided that there are five main objectives to this project:

a) The design and development of a temperature-controlled version of the fReactor.
b) Integration of said equipment into an automated rig existent at iPRD in order to enable automatic exploration of the reaction space for multiphasic reactions.
c) Analysis of the data obtained by the use of such equipment employing a variety of different metamodeling and cross-validation techniques, aiming to outperform the widely used least squares polynomial fitting.
d) Investigation and quantification of the performance of a plate miniaturised continuous oscillatory baffled reactor via computational fluid dynamics.
e) Computational fluid dynamics – enabled design optimisation of the plate mCOBR via surrogate modelling.

1.8 Structure of the thesis

This thesis has a length of five chapters. The current chapter introduced the project, including a comprehensive literature review and identification of the research gap.

Chapter 2 focuses on the design and integration of a temperature-controlled fReactor cascade. This includes insights around the mechanical design of the fReactors and also around the electronic design of the controller. Special attention is paid to microprocessor programming in order to provide a quick and stable temperature control with appropriate routines to effectively communicate with laboratory instrumentation. It also features additional hardware developed for the fReactors; these consist of a tailored stirrer plate and a photochemical setup that enabled UV photochemistry.

Chapter 3 focuses on the analysis of the data gathered from automated flow chemistry experiments in the fReactors. A Claisen-Schmidt condensation is assessed, building metamodels for different metrics by comparison of different cross-validation techniques over which different optimisations are exemplified. How different metamodeling techniques can help obtaining more precise response surfaces is demonstrated by means of the in-silico metamodeling of a nucleophilic substitution reaction. Additionally, the use of the fReactor as a tool for *in situ* X-Ray Diffraction (XRD) is demonstrated.

Chapter 4 explores the CFD-enabled optimisation of a miniaturized plate COBR, assessing the effect of baffle offsetting and including a multi-objective optimisation between competing objectives of maximising transversal mixing and minimising the variance of the RTD.

Finally, Chapter 5 outlines the conclusions of the project in relation with the aim and objectives introduced in the previous sections of this chapter, highlighting its main contributions where relevant. Furthermore, it also recommends future research direction based on the findings of the thesis.

Chapter 2

Design and Development of a heated fReactor version

2.1 Introduction

The implementation of a heating system for this miniature CSTR cascade of fReactors presents various challenges. The design, as well as the selection and development of the different elements of the controller for this application was an important part of this research, that included several multi-disciplinary aspects. Accordingly, the chapter is structured to cover all the different parts that, in conjunction, lead to a fully integrated temperature-controlled version of the fReactor in the automated rig at iPRD, expanding the automation of a range of flow chemistry to multiphasic studies across temperature. The chapter is structured as follows:

- Section 2.2 presents the mechanical design of a temperature controlled version of the fReactor.
- Section 2.3 gives insights on the electronic design of the controller, including the electrical design of the device that was transferred to a PCB and finally enclosed in a box for safety.
- Section 2.4 focuses on the microprocessor's programming, software architecture, data acquisition, serial communication protocols and PID tuning implementation. The final program and implementation notes are given in Appendix C.
- Section 2.5 introduces additional hardware of a stirring and photochemical light source specially tailored for the application.

2.2 Mechanical Design

2.2.1 Background

The fReactors, introduced in the literature review of this thesis, are a miniaturised version of a continuous stirred tank reactor, first developed by Professors Kapur and Blacker at the University of Leeds Institute of Process Research & Development (iPRD). A fReactor consists of a base with a reservoir and conduits for the fluids allowing interconnection. The reservoir contains a magnetic stir bar to provide enhanced and uniform mixing. A gasket and a lens are used to seal the reservoir, which are clamped onto the base of the fReactor by means of a PTFE lid that can be tightened using M5 bolts for which threaded holes are provided in the body of the device. The use of the transparent window enables visual monitoring of the reaction in the reservoir, as well as the implementation of a photochemical setup for the fReactors that will be introduced later in section 2.5 of this chapter.

Although the fReactor was proven to be a very easy lab-scale tool that enabled multiphasic reactions in flow, most of its applications have been constrained to room-temperature reactions. The demonstration of temperature-controlled reactions was limited to situations in which the fReactors were placed over a commercially available multi-position stirrer heating plate and the user waited until a steady fluid temperature was reached.[122] With this arrangement, the temperature of the hot plate had to be manually increased at various times until the steady-state temperature of the fluid at the outlet was the one required for the experiment. This time-consuming approach to temperature-controlled reactions also needs to be carried out again if the flow rate is changed between experiments, since this would change the dynamics of the process. Besides this, the thermal characteristics of Delrin such as its low thermal conductivity (0.31 W/K·m) and high specific heat (1465.39 J/kg·K) do not make it appropriate for heat transfer applications, as even a thin layer of such material can lead to high temperature gradients across both sides. Furthermore, using a hot plate does not allow to independently control the temperature of various fReactors in a cascade, a feature that would enable applications such as cooling crystallisation.

In order to overcome these limitations, a series of modifications were undertaken in the design of the fReactor with the aim of enabling temperature control of the fluids in the reservoir. A controller makes use of two basic elements to control a process: a measurement device and an actuator. In order to provide quick and precise temperature control, each fReactor in the cascade was fitted with a thermocouple and a cartridge heater. Among the many options available in the market, mineral insulated K-type thermocouples were chosen as they include a long 1.5 mm diameter probe that can make use one of the ports on the base of the fReactor to be physically inserted in its reservoir and sealed with generic flangeless fittings. Different fittings were tested for the application, with the optimal being the use of VICI collapsible ferrules. When the nut compresses the ferrule, which is positioned around the thermocouple's probe, it is deformed and tightly wraps around it, effectively sealing the fReactor even under pressures of 250 p.s.i. In this regard, the body of the reactor does not need to be modified, since it uses an existing fluidic connection to integrate the temperature measurement device, as shown in Figure 34. The heaters chosen for the application, in view of their low resistance and rapid response, were 40 W ceramic cartridge heaters. These heaters are commonly used as a hot end in 3D printing applications due to their rapid heat transfer characteristics. Accordingly, the new version of the fReactor designed using SolidWorks (Dassault Systèmes, SOLIDWORKS Corporation) included a cylindrical aperture that made possible to fit the heaters in. The diameter of the hole was that of the heaters, 6 mm, and the cartridges were fit into these apertures with the use of Arctic Silver 5 thermal compound (k = 8.9 W/m·K), providing good contact between both surfaces. Considering the length of the heater, 20mm, the holes were 35 mm long to position the heater in the centre of the device. A smaller hole was included in the design communicating the bottom of the heater's hole with the opposite side of the reactor to facilitate removal of the cartridge with the aid of a thin rod if in need of replacement. However, the high temperatures reached by the cartridge, which greatly surpasses Delrin's melting point (175 °C), do not allow use of the heater in conjunction with a Delrin reactor.

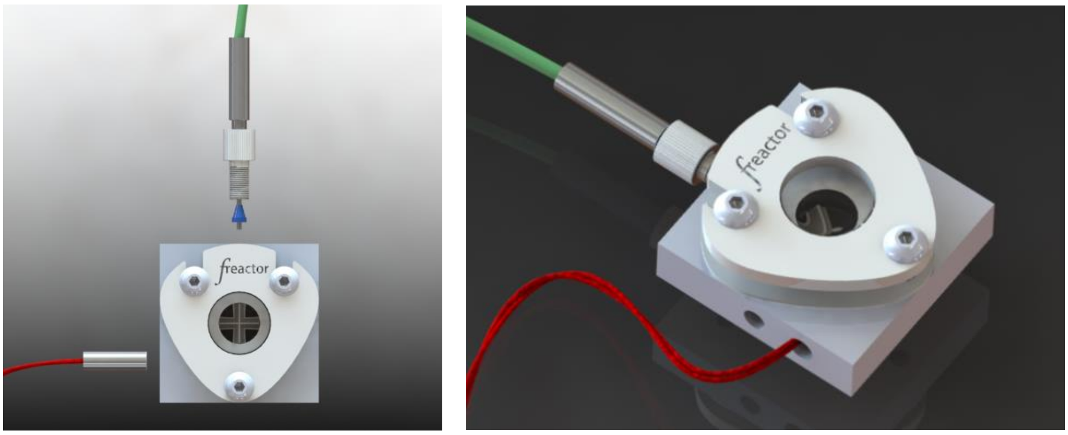

Figure 34. Left: top view of the proposed modified fReactor with the thermocouple and cartridge heater out. Right: 3D view of a fully assembled fReactor.

In consequence, a different material was chosen that enabled the use of the heaters in the reactor's body without safety concerns. It was decided to proceed with the design of 316-grade stainless steel reactors. This is the second most common austenitic steel, also referred to as A4 stainless or marine grade stainless. This kind of stainless steel presents a thermal conductivity of 14.6 W/K·m, 47 times higher than Delrin's, and a specific heat of 450 J/Kg·K (almost three times lower than Delrin's) and melting point from 1390 °C that ensures the safety of the application. Another benefit derived from the use of this material is its high resistance to corrosion, including the following features [196]:

- Highly resistant to sulfuric acid for concentrations below 10%.
- Resistant to the attack of phosphoric acid.
- Ideal for handling acetic acid, especially when combined with formic acid.
- Not damaged by aldehydes or amines.
- Fatty acids only affect it for extreme temperature conditions (above 260 °C).

The body of this version of the reactor was manufactured from 50×50×12 mm marine grade stainless blocks. Two computer-aided design (CAD) models for different versions of the fReactor were produced using SolidWorks. The first one, shown in Figure 35 (left), is the standard version of the device that provides an inlet and an outlet, apart from a port for the thermocouple. The only difference in the second version is that it is equipped with two inlets, allowing mixing of two incoming reagent streams directly in the fReactor’s reservoir instead of in a tee fluidic connection before it. Detailed technical drawings for both versions can be found in Appendices A and B, respectively.

In both cases, the reaction chamber was drilled perpendicularly at the centre of the top 50×50 mm face, consisting in a cylindrical space that is 15 mm in diameter and 10 mm deep. A cross-section running across the reaction chamber is shown in the in Figure 35 (left, section A). Along with it, three standard M5 threaded through-holes are machined that enable clamping the lid. Conduits for the entrance and exit of the fluids are drilled perpendicular to the side faces of the fReactor. All these connections are equipped with standard 6 mm long UNF ¼-28 threaded ports for fluidic connections that are tapped from 8 mm long #3 drill bits (5.41 mm diameter). These are compatible with most flangeless fittings in the market. A PTFE gasket is used along with a convex lens to seal the reaction chamber, clamping it under a lid that allows three M5 bolts to tighten it. The lid is laser cut from Delrin with a distinctive shape that characterises the fReactors. In this way, the fReactors form an easy assembly plug-and-play module (Figure 35 (right)) that enables not only multiphasic but enhances single phase mixing.

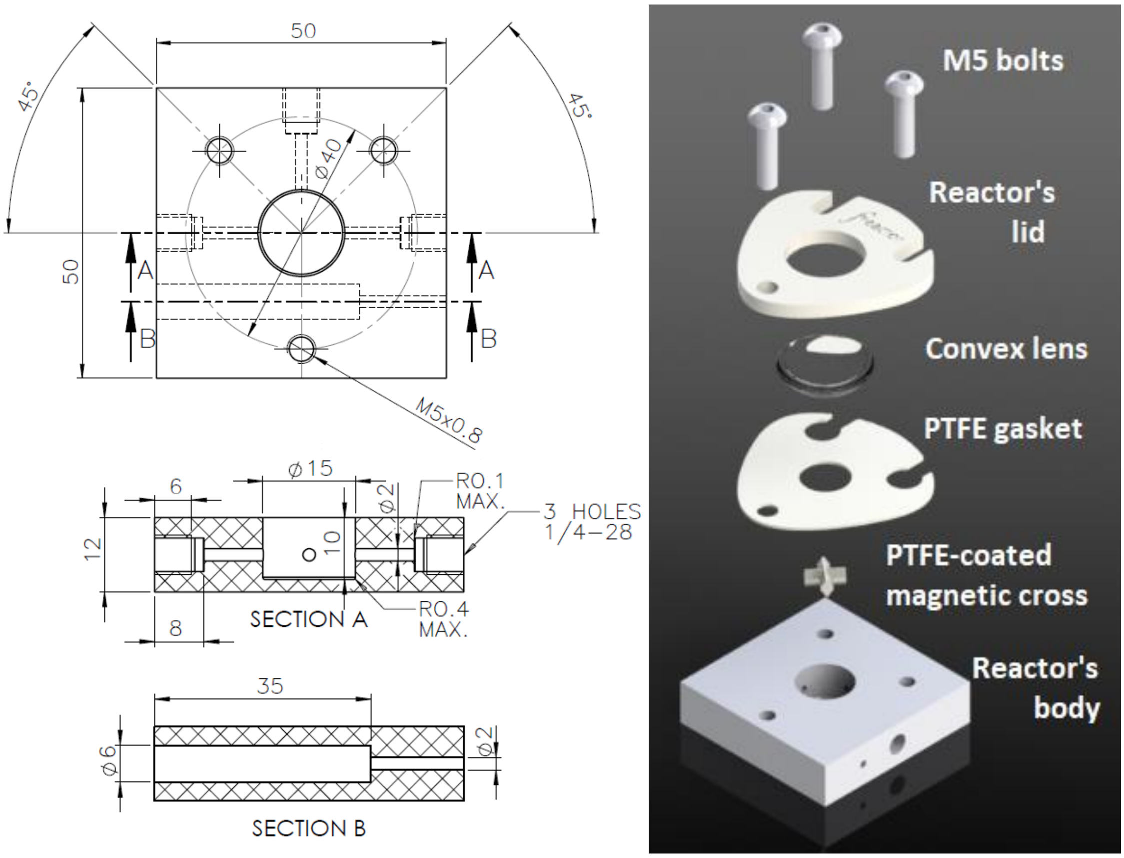

Figure 35. Left: Technical drawing for the fReactor's stainless steel body. Right: Exploded view of the assembly, showing the elements that form a fReactor.

Device fabrication was undertaken at the EPSRC National Facility for Innovative Robotic Systems at the University of Leeds School of Mechanical Engineering using state of the art machine tools. All threads in the base of the fReactor were manually retouched using the appropriate tap sets. The PTFE gasket and lids were laser cut from 0.5 mm and 5 mm thick sheets from CorelDraw X7 (Corel Corporation, Ottawa, Canada) designs. The magnetic crosses and lenses were purchased from external providers. The assemble can be easily fitted making use of button head socket screws and an Allen key.

2.2.2 Characterisation

It is obvious that during the operation of the device, different elements of fluid will spend different times in the reactor, due to them following different pathways between the entrance and exit. For instance, some particles could flow quickly through the reactor, while others could spend longer times circulating in the reactor to the action of the magnetic cross. The distribution of these different times for a stream of flow leaving the reactor is known as the Residence Time Distribution (RTD) of the reactor. This distribution conveys important information about the performance of the reactor, and it is possible to quantify the fractions of liquid that spend between different times within the reactor.

A common non-chemical method of experimentally determining the RTD of a reactor is the use of a known mass of a dye as a tracer, injecting it in the reactor and analysing its concentration or a related quantity at the stream of flow leaving the reactor. This is known as the tracer or pulse experiment. Registering these data, is possible to obtain a graph of the concentration over time. This graph can be modified to obtain the RTD, simply dividing the concentration values by the area under the curve. For an ideal CSTR, the tank-in-series model states [121]:

$$E(t) = \frac{t^{n-1}}{(n-1)!\, t_m^n}\; e^{-t/t_m} \qquad (21)$$

where t is the time, t_m the mean residence time per tank, and n the number of reactors in the cascade.

In order to compare the experimental values with the theoretical, a pulse experiment was performed by iPRD fellow Maria Kwan, using a 10 second pulse of 2.5% (v/v) red food dye tracer which was injected in the reservoir of the first reactor in the cascade. The UV-vis absorbance (516 nm), that is proportional to the concentration of the dye in the sample was measured by means of a standard laboratory spectrophotometer at regular intervals of 30 s. These data were imported using Matlab and plotted against the ideal curves obtained from equation (21), as shown in Figure 36.

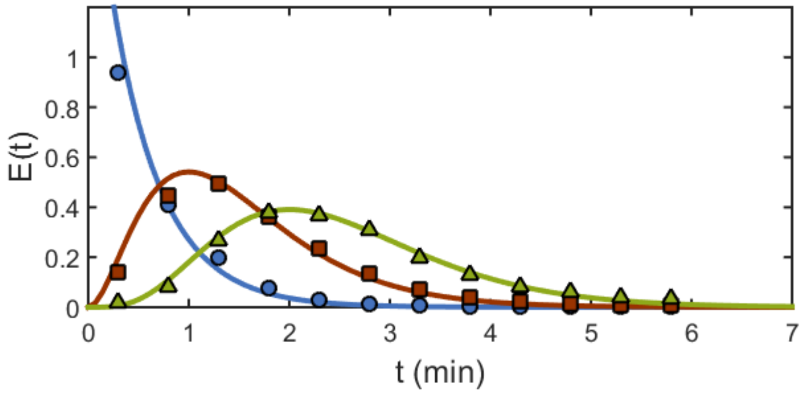

Figure 36. RTD for a single fReactor (blue) and cascades of three (red) and five (green) fReactors. Lines: theoretical curves. Markers: experimental values.

The ideal curves are based on the premise that the tracer is instantaneously and perfectly mixed once injected into the reactor. Naturally, this behaviour is difficult to achieve experimentally due to the non-ideal nature of mixing and flow patterns, and the length of the tubing used to feed the cascades, that will create an additional delay between the tracer reservoir and the reactor.

Whilst it is possible to create a superposition of the theoretical RTD for the feed pipe and the reactor, here the mean delay of 1.7 minutes was subtracted from the experimental data in order to compare it with the theoretical curves for $E(t)$. The experimental data shows good agreement with theoretical for all three examples (1, 3 and 5 CSTRs). From these plots, it can be inferred that the reactor does not comprise dead-zones (i.e. no long theoretical tails are present), nor bypass effects (i.e. no overestimated value of $E(t)$ is observed).

2.3 Electronic Design

Controlling the temperature of a system requires the use of various electronic elements. For a control loop, it is needed to i) measure the process variable using a sensor, ii) feed this information to a microcontroller which will process it and produce an output signal that iii) needs to be amplified to control the actuator element in an effort to reach and maintain the desired process variable setpoint (SP).

A schematic of the electronic prototype designed and built for the application can be found in Figure 37, where the different components and the connections between them can be appreciated.

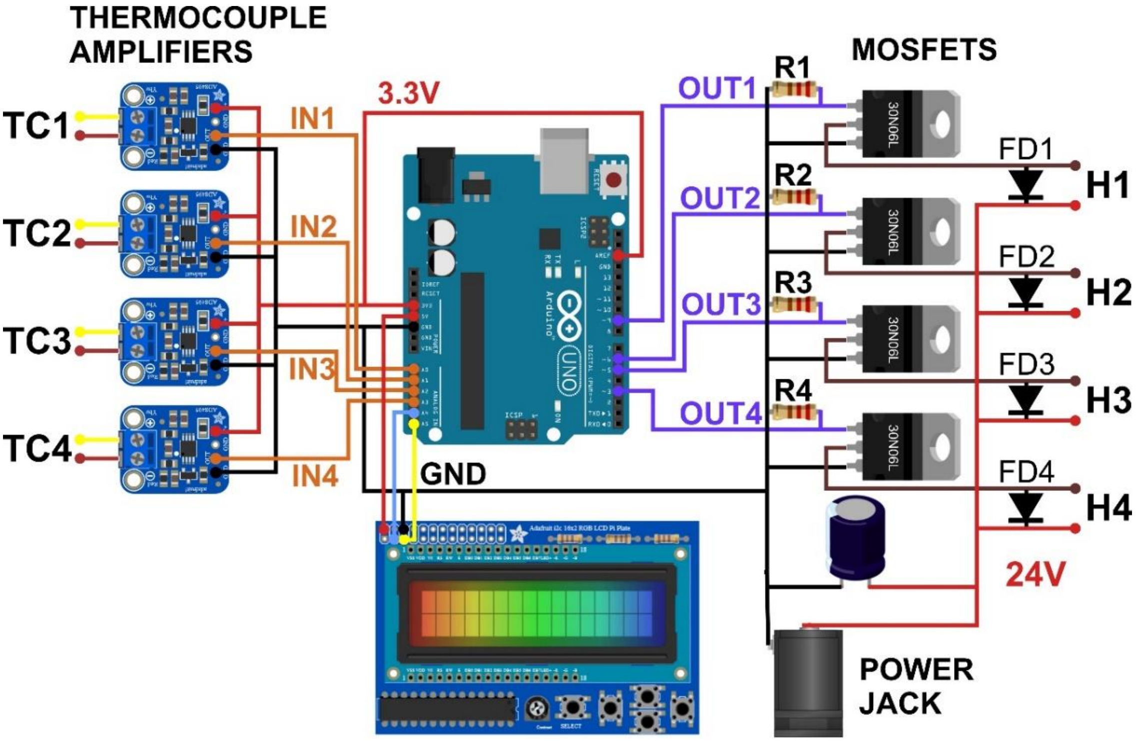

Figure 37. Schematic of the electronic prototype.

This electronic setup allows for the control of up to four different CSTRs, but the nature of the software developed for the microprocessor makes possible to easily arrange the system for a different number of reactors. The connections for the temperature sensors, named TC1 to TC4, can be found in the left part of the figure. These are connected to four thermocouple amplifiers whose role is to provide a readable electronic signal proportional to the temperature (IN1 to IN4, in orange) to the Arduino UNO board, pictured in the centre of the figure. These signals are read and processed by the microcontroller, which produces the output control signals (OUT1 to OUT4, in violet) that are amplified by means of the four MOSFETs and a 24 V external power source used to power the fReactor heaters. The different elements highlighted here are further discussed in the following sections.

2.3.1 Microcontroller

The microcontroller is the core of this automated application, where the processing takes place. Summarising, an Arduino UNO consists of a microcontroller ATmega328 with 14 input/output pins (of which 6 can be used as a Pulse Width Modulated outputs), 6 analogic inputs, a 16 MHz ceramic resonator, a USB type-B connection, a power jack connector, an in-circuit serial programming (ICSP) connector consisting of 6 pins and a reset button.[197] This board was selected as it is the most robust and documented board of the Arduino family, providing all the necessary features for the use of a microcontroller, being a widely used low-cost tool for prototyping of automated applications.

There are a number of different ways in which the Arduino board can be powered.[197] In the matter at hand, and since a serial communication between the master computer and the microcontroller is needed and can be achieved through the USB type-B port, the same connection is used to power the board. In this way, this connection is the mean for both the powering and the communication of the board.

2.3.2 Data Acquisition

A variety of temperature sensors are available in the market including thermocouples, resistor temperature detectors (RTDs), thermistors, infrared sensors and semiconductor-based ICs, among others. Of these, thermocouples are widely used as they present a series of advantages such as simplicity, broad temperature range, robustness, rapid response and low cost.[198] In the current case, their reduced size represents also an advantage, for the reason that it will make possible to fit the sensors in the reactors using one of the built-in ports reserved for this purpose as explained before in the mechanical design section of this chapter (section 2.2). In this way, it is possible to place the tip of the sensor in contact with the fluid in the chamber of the fReactors to provide an accurate measurement of the temperature of the fluid, as opposed to previous approaches that relied on measuring the temperature of the body of the reactor instead, which assume that the temperature of the fluid is the same as the reactor body.

A thermocouple consists of two wires of different metals joined at one end, known as either “hot” or “measurement” junction. At the other end, each of these wires is connected to a conductor (usually the copper tracks of a PCB), forming a “cold” or “reference” junction. Under this configuration, an electric voltage is created as a result of a temperature difference between the hot and cold junctions and the use of dissimilar metals in what is known as the Seebeck effect.[199] For a range of temperatures known as the linear region of the thermocouple, this voltage is proportional to the temperature difference between the junctions and in consequence can be used to determine the temperature of the measurement junction, provided that the temperature of the reference junction is known. Therefore, a second temperature sensor must be used in conjunction with the thermocouple to measure the temperature of the reference junction if an absolute value of temperature for the measurement junction is wanted. This procedure is termed “cold junction compensation”.[200] It is important, thus, to place the temperature sensor as close as possible to the reference junction of the thermocouple to minimise measurement errors.

Mineral insulated K-type (chromel/alumel) thermocouples with pot seal were selected for their use in this application and purchased from Thermosense. This kind of thermocouple is inexpensive and presents a wide linear region, what makes it optimal for the application. Furthermore, the model selected includes a sheath that provides housing for the measurement junction of the thermocouple, protecting it from chemicals while at the same time giving it a rigidity that makes easier to work with it. It is built from 310 stainless steel, with a diameter of 1.5 mm that makes it ideal for its use with the ports in the reactor, enabling positioning of the tip inside the fReactors, where the sheath will be in direct contact with the fluids in the interior. The use of standard 1/16’’ flangeless fittings make possible this positioning and seals the reactor, avoiding leaking through the ports reserved for the thermocouples. This kind of thermocouple presents a sensitivity of 41 µV/°C (Seebeck coefficient at 25 °C), and therefore the output voltage that the difference of temperatures creates in the thermocouple is minimal. For instance, for a temperature difference of 200°C, a K-type thermocouple in its linear range would yield 8.2 mV (200°C × 41µV/°C), which is a very low voltage for either the 5 or 3.3 V logic level of the 10-bit analogic input pins of the board to read. In consequence, it is necessary to make use of an amplifier in order to provide a higher voltage temperature signal subject to be read using Arduino's analogic input pins.

The K-type precision thermocouple amplifier AD8495 manufactured by Analog Devices provides a solution to both problems pointed out in the previous paragraphs, as it includes a precalibrated amplifier and an internal cold junction compensation.[201] The amplifier produces a high-level output of 5 mV/°C from the small thermocouple voltages by applying a gain of 122.4. In this case, a temperature difference of 200 °C would produce a 1 V output signal. A built-in temperature sensor allows for cold junction compensation of any temperature changes in the reference junction. The chip also includes a reference pin for offset adjustment, allowing reading of negative temperatures. The expression for the output voltage of the chip is:

$$V_{out} = \left(T_{mj} \times 5 \left[{}^{mV}/_{\circ C}\right]\right) + V_{ref} \qquad (22)$$

where T_{mj} is the temperature of the measurement junction and V_{ref} is the reference voltage externally supplied to the chip. It is easy to see how V_{ref} makes reading negative temperature values possible by offsetting V_{out}. Isolating T_{mj}, the temperature can be easily calculated in the microprocessor from the amplifiers output voltage according to equation (23).

$$T_{mj} = \frac{V_{out} - V_{ref}}{5 \left[{}^{mV}/_{\circ C}\right]} \qquad (23)$$

The AD8495 can be purchased soldered in a breakout board along with the ancillary components to the chip suggested within technical documentation provided by the manufacturer. These include several features including a low-pass filter, a ground connection, a 1.25 V precision voltage reference and a decoupling capacitor, as well as a ferrite bead for high-frequency noise suppression. The low-pass filter is included with the aim to remove both common-mode (16 kHz) and differential (1.3 kHz) noise in the thermocouple lines. A ground connection through a 1 MΩ resistor on the negative input of the amplifier enables open thermocouple detection and, in addition, allows potential static charges to discharge to ground. A 0.1 µF decoupling capacitor is used on the positive power pin to provide clean power supply voltages, with the ferrite bead blocking line noise. The precision voltage reference TLVH431 [202] is used to provide an accurate V_{ref} of 1.25 V in order to enable negative temperature readings. Figure 38 shows the chip layout, while the pin connections of the breakout board are presented in Table 3.

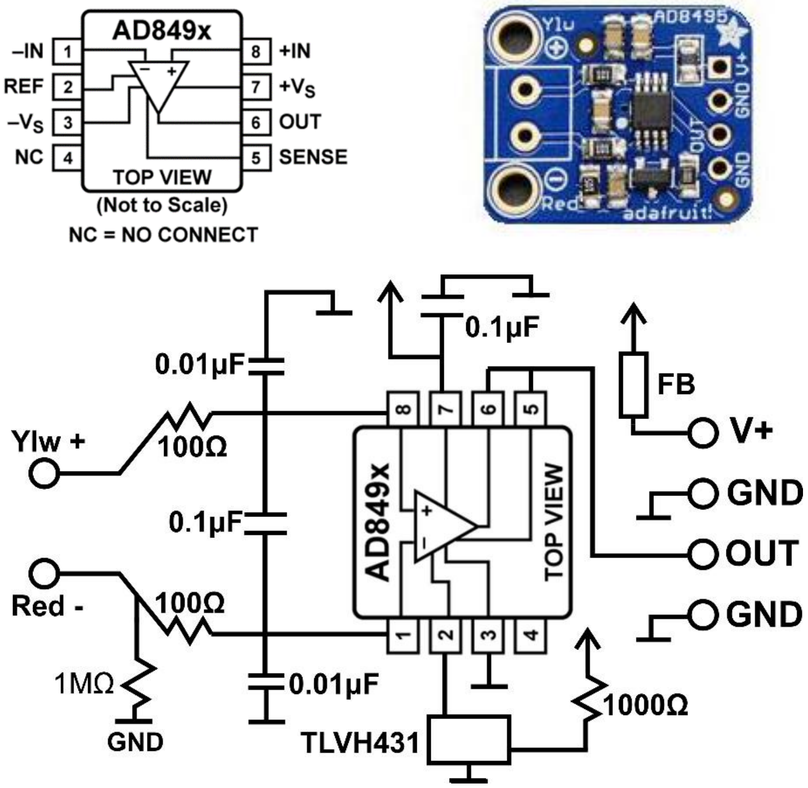

Figure 38. Top left: Analog Devices AD849X amplifiers family layout. Top right: Adafruit’s AD8495 breakout board, in which the amplifier with eight pins can be seen in the centre. Bottom: electric schematic of the breakout board.

Table 3. AD8495 breakout board connections.

PIN	DESCRIPTION
Ylw (+)	Thermocouple positive terminal connection.
Red (-)	Thermocouple negative terminal connection.
GND	Ground. The chip includes two ground pins to facilitate measurement of the output voltage if needed.
V+	Connection for the external power supply.
OUT	Analogic output of the chip.

This chip amplifies the output of any K-type thermocouple, what makes the thermocouples easily interchangeable and replaced if needed, without affecting the readings. It accepts supply voltages between 2.7 and 18 V, with both the 5 and 3.3 V logic levels that the Arduino UNO board is able to provide falling within this range. Thus, the thermocouple amplifiers can be driven from these pins (5 or 3.3 V) on the Arduino board, which in turn is powered from a computer via the USB type-B connector.

The temperature values will be computed in the microcontroller according to equation (23), using a tailored function developed in the Arduino IDE and called *getTemp,* that will be introduced in the software section (section 2.4.4) of this chapter.

2.3.3 Output Amplification

From the temperature values calculated using the thermocouples and their respective amplifiers, the PID control algorithm will produce a numeric result that can be used to provide a proportional “analogic” PWM electrical output. This signal, which is digital in nature, aims to mimic an analogic signal by means of switching between its low (0 V) and high (5 V) states for a portion of time of a complete cycle.[203, 204] The amount of time the signal stays in its high level compared to the duration of a complete cycle is referred to as the duty cycle. The evolution of the PWM signal’s voltage over time is illustrated in Figure 39, which shows the kind of PWM outputs produced by Arduino’s analogWrite function (whose argument ranges from 0 to 255). What stands out in this figure is the continuous switching of the voltage to imitate an analogic signal. For example, in the case of the 25% duty cycle, the signal stays high for 25% of each cycle, yielding a virtual voltage of 1.25 V (5 V × 0.25).

However, the microcontroller board powered from a computer is only capable of providing a logic level low-power signal that is unable to drive a heating element by itself, as the signal switches between 0 and 5 V, with a maximum current of only 40 mA (W-6). In consequence, the signal produced by the board needs to be amplified in order to drive the heating elements, which are designed to work at much higher voltages (24 V) and currents (up to 1.33 A) that the ones the PWM output signal from Arduino can provide.

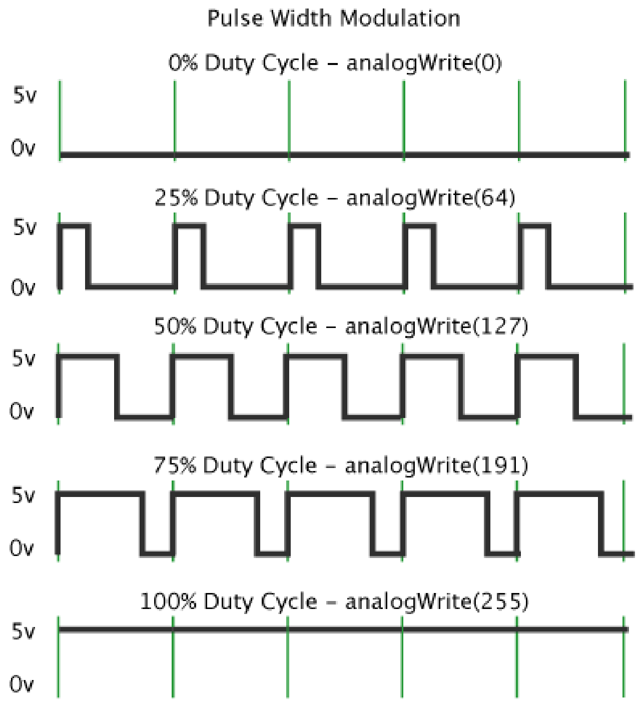

Figure 39. Arduino PWM output. Reproduced from ref.[205]

An effective low-cost option to amplify electronic signals is the use of an external high-power source along with a metal-oxide-semiconductor field-effect transistor, commonly referred to as its acronym, MOSFET.[203] This element presents three terminals called Source, Drain and Gate. The Gate makes possible to control the flow of high currents between the Source and Drain terminals with a logic level signal. Consequently, a PWM signal can be used to switch the flow of current from the high-power source proportionally to its duty cycle. As depicted in Figure 37 (right part), the source of each transistor is connected to the ground of the external power supply, which in turn is connected to the ground in the Arduino board. The Drain is connected to a lead of the heating element, with the other lead connected to the external power supply positive terminal. When the PWM output of the board drives the gate of the transistor it switches at high frequency, allowing the high-power current to flow when the PWM is at its high level, while restraining it from flowing when it is low, essentially allowing the high-power supply current to flow proportionally to the duty cycle of the signal to drive the heater, simulating a high-power analogic signal. However, the main drawback of this approach is the generation of electrical noise due to the high frequency switching of the power supply. This difficulty was addressed by the introduction of a series of elements that will be presented later in section 2.3.5 and that aim for the reduction of the noise in the circuitry.

However, while a wide range of MOSFETs are available in the market, not all of them are suitable for this specific application due to limitations in the maximum ratings for current and voltage. In a MOSFET the drain current is limited depending on the voltage applied to the gate, and some of them cannot be activated working with a logic level of 5V. This is an important selection and thus, the spreadsheets and ratings of different transistors were evaluated to come up with a suitable option. The Fairchild FQP30N06L N-channel MOSFET [206] was selected because of its high maximum current (32A) and voltage (60V) ratings and fast switching transfer characteristics, that make possible to drive the gate at the 5 V logic level of the PWM output Arduino provides. It can be appreciated from Figure 40 that for a Gate-Source Voltage of 5 V, the maximum drain current that the transistor would allow to flow through it is ~50A, which is higher than the maximum current required by the application. This ensures that the transistors could be used without degrading rapidly or causing excessive overheating, as both the voltage and current values present during operation are well under the maximum limits of operation of the transistor.

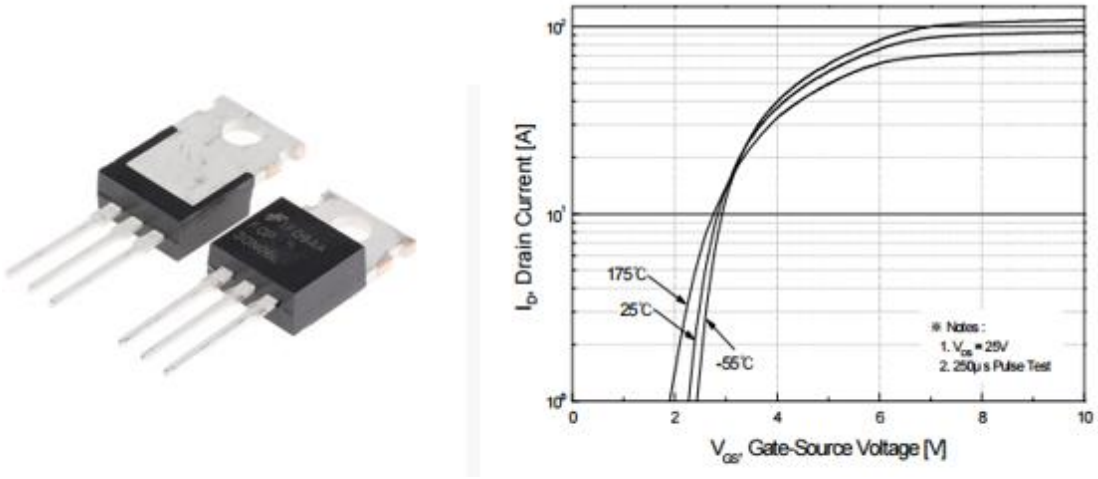

Figure 40. Fairchild FQP30N06L MOSFET view (left) and transfer characteristics (right).

2.3.4 LCD Screen

The electronic setup includes a 2×16 LCD screen to continuously show the temperature values of the different reactors to the user. Although these values will also be accessible from the computer screen via Matlab, it was considered convenient to provide a visual reference to the users that will work with the automated system in the laboratory. The screen is powered from the 5 V source in the Arduino, eliminating the need of an extra power supply. With this 5 V level, there are a total of three different voltage levels in the circuit: the thermocouple amplifiers, which work at 3.3 V; the LCD screen, that is fed using Arduino’s built-in 5 V supply; and the heaters, that use 24 V. However, all of these parts share a common ground, a good practice to provide the same voltage reference to the whole circuit.[207]

The required number of pins needed to run the LCD screen would normally be of 6 for screen control and an extra 3 for RGB backlight control, amounting to a total of 9 pins.[208] However, as the number of pins on the Arduino UNO board is quite constrained, an alternative approach was found that permitted the use of the screen while controlling it with a minimal number of pins involved. The Adafruit RGB LCD shield kit includes a port expander that makes communication possible by using solely 2 pins and a communication protocol. This protocol, known as I^2C (Inter-Integrated Circuit), enables communication over two lines, namely SDA and SCL, which are the serial data and serial clock, respectively. The data is transferred over the serial data line, while the clock line is used to synchronize the data transfer by sharing the same clock signal between the master and the slave devices.[209] Arduino provides SDA and SCL pins, corresponding to the analogic A4 and A5 pins, that can be used to implement I^2C communication with the port expander integrated in the LCD breakout board, the 16-bit I^2C I/O MCP2307, that in turn drives the LCD screen using as many pins as the LCD requires. The integration of the LCD in the electrical design is shown in Figure 37, being the light blue and yellow connections the SDA and SCL lines, respectively.

2.3.5 Other Technical Considerations

PULL-DOWN RESISTORS

Although the basic features of the electrical device have been already explained, some other considerations must be taken into account when designing high-power applications, where oftentimes the electrical noise caused by switching of the power electronics elements can create interferences that affect the state of a MOSFET gate and make it “float” in a state that is neither recognised as high or low. In this regard, pull-down resistors (named from R1 to R4 in Figure 37) are used to hold the gate to the logical low level (0 V) whenever the microcontroller is not sending a high signal, by means of connecting the gate of the transistor to the ground through a resistor.[204] This resistor lowers the voltage of the line to ground if nothing else is driving it up. A high value of resistance of 9.1 kΩ was chosen for these elements, as it prevents excessive currents from flowing through them. In summary, the addition of pull-down resistors prevents transient voltages in the lines from accidentally making them appear high due to noise, thus increasing the robustness of the design.

DECOUPLING

Similarly, it is important to decouple the electronics from the power supply. As it has been previously explained, the controller works by switching the heaters on and off at high frequencies. Every time there is a sudden current demand, if drawn directly from a power supply, transient voltages are created, generating electrical noise that can potentially prevent the circuit from working. Decoupling capacitors are used as a mean to filter out the noise.[207] These basic elements store charge, opposing changes in voltage. If the voltage suddenly drops, the capacitor discharges by supplying a current in an effort to maintain the previous voltage. On the other hand, if the voltage is quickly raised, the capacitor stores charge from the current with the same aim of preventing changes in the voltage.[210] By placing these elements in parallel to the rest of the circuit, the transient currents can be supplied from the capacitors instead of directly from the power supply, avoiding problems such as ringing in the power supply. In this case, two capacitors of 820 and 770 µF were placed across the power supply line, effectively removing the voltage spikes generated by the constant high frequency switching. These voltage spikes prevented the LCD from working correctly, either showing random characters or turning it off during the development stage; however, the problem disappeared as soon as the capacitors were included in the circuitry.

FREEWHEELING DIODES

Sometimes, inductance is present in the actuators of the system. Whenever inductance is present in a switching application, a high voltage transient is generated because of Faraday's law of induction [211] every time the current is switched off. In fact, this transient can be large enough to destroy the transistors.[212] According to Lenz's law,[213] the voltage will be produced in the direction that opposes the change in current. Taking this into account, the element that presents inductance will reverse its polarity to keep the current flowing, using the energy stored in its magnetic field to this end.[212] This can be envisioned as an analogous case to the capacitor: while the capacitor opposes changes in voltage, the inductor opposes changes in current. Having a look at equation (24) , the sudden change in current produced when switching off will create a very large back electromotive force (EMF), since the drop in intensity is abrupt and, in consequence, ${dI}/{dt}$ will be very high.

$$\varepsilon = -L\,\frac{dI}{dt} \qquad (24)$$

The addition of a freewheeling diode in antiparallel (in parallel, with its polarity inverted) to the inductive device provides a path for the current generated as a consequence of Faraday’s law of induction when the circuit is switched off, creating a loop between the inductor and the diode.[214] This prevents the current through the inductor from dropping abruptly, as the voltage the inductor creates forward biases the diode, allowing a path for the current to recirculate and avoiding the build-up of the voltage spike. When the transistor is switched on, the diode is reverse biased and has no effect over the circuit. These elements are named FD1 to FD4 in Figure 37. The choice of diodes were the popular 1N4007, rated for a maximum DC voltage of 1000 V, which can sustain a current of up to 1 A continuously and a peak current of 30 A during 8.3 ms.[215] With the addition of these elements to the design, it is ensured that the rest of the circuit is protected from the back EMF that could have been generated by any potential inductance of the actuators of the system, which in this case are heating cartridges.

2.3.6 PCB Design

A printed circuit board integrating soldering pads for all the electronic elements described hitherto was designed in order to increase the robustness and durability of the device. The PCB design software EAGLE (CadSoft Computer GmbH, Germany) was used to model the board, which was shaped in a form that resembles that of Arduino, and designed as a shield for the latter, including soldering pads for all the pins in the board. Straight male headers were soldered into these pads, making the shield stackable into the Arduino, forming a compact hardware unit (Figure 42(left)) that could be placed inside a protection box. In this way, all Arduino pins are accessible from the PCB, simplifying both the connection of all the electronic components and the hardware’s debugging.

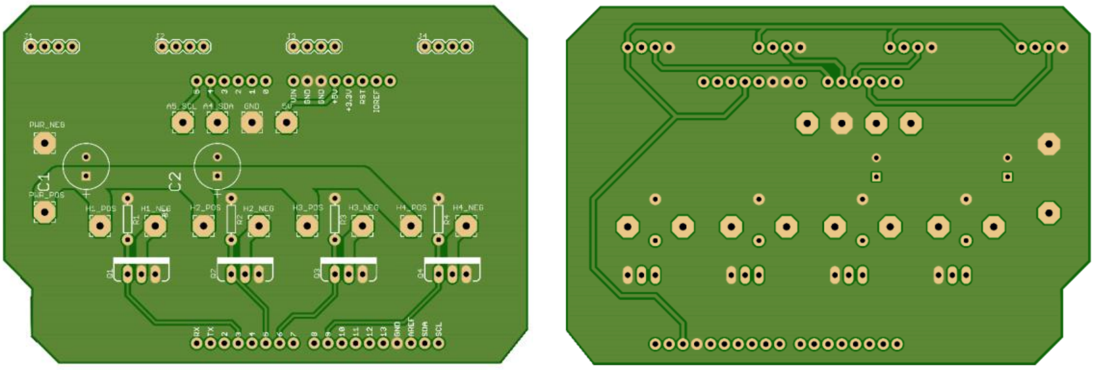

Figure 41. Top (left) and bottom (right) of the PCB designed especially for the temperature controller.

Attending at Figure 41 (left), all power connections, as well as the LCD connections are placed in the top part of the PCB. A wide trace can be recognised that connects the power supply (two big pads on the left side of the board) to the respective positive terminal connections of the four heaters. The width of this trace was calculated following the generic standard on PCB design IPC-2221 [216] as to allow the flow of the high currents needed in order to drive the heaters without producing significant overheating in the conductor. The minimum conductor's width and thickness can be obtained by using equation (25):

$$I = k\,\Delta T^{0.44}\,A^{0.725} \qquad (25)$$

where I is the current, k is a constant that equals 0.048 for outer layers like the ones used in the design and A is the cross section of the track in squared mils. A mil, or a thousandth of an inch, is a unit commonly used in the manufacture of PCBs. Usually, the conductive layer used in PCB manufacturing is 1 oz / ft^2 copper,[217] equivalent to 1.37 mils thickness. In this occasion, and in order to reduce the width of the traces needed in the board to accommodate the peak current demand, a copper layer thickness of 2 oz / ft^2 (2.74 mils) was preferred. Isolating A in equation (25) and using a trace width of 3.81 mm (150 mils) the equation yields a temperature rise of ~1.14 ºC for the maximum intensity case (4 A). It is worth noticing that such temperature rise will be produced only in the first, smaller stretch of the trace before part of the current is derived to the first heater, being lower for the rest of the length of the trace. In the same way, the traces connecting each MOSFET's drain with the pad for the negative terminal of the heater is wider than the standard signal tracks that can be noticed for the rest of the connections in the board. This is a 1.778 mm (70 mils) wide track able to handle the maximum individual heater intensity of 1.33 A while producing only a slight temperature increase of 0.33 °C. The rest of the connections are routed with a 0.4064 mm (16 mils) trace width, leading to negligible temperature rises due to the low maximum currents at the logic level (40 mA). In this way, excellent thermal stability is granted for the board, widely accommodating the current requirements of the application.

Specialized soldering pads for the MOSFETs can be recognised in the lower part of the upper side of the board (Figure 41 (left)). The gate of each transistor is connected to the corresponding PWM pin of the board that governs the action of its heater. Each gate is also connected to the ground through its pull-down resistor, for which the PCB provides soldering pads. A space reserved for the decoupling capacitors can also be identified (C1 and C2 white circles) designed to connect the wide power trace to the circuit's ground, decoupling the circuit from the power supply.

The connections regarding the thermocouple amplifiers are located on the bottom part of the board (Figure 41 (right)), including 4 pins for each one. It is also on this side of the board that the Arduino 3.3 V generator is connected to its external reference pin AREF, being all the thermocouple amplifiers fed from said generator in parallel. Each amplifier's 'OUT' pin is connected to its respective analogic pin on the Arduino board. The rest of the bottom part of the board forms a vast ground plane, to which the ground pins of all elements are connected, creating a return path for the current of all components in the board. This is a well-established way of avoiding ground loops, providing the same ground voltage reference to all devices on the board.[218] In addition, and considering that the wide power tracks of the upper side and the ground plane are parallel plates separated by a dielectric, this configuration (which is the same as the one found in a capacitor) also adds extra capacitance to the entire circuit board. In summary, including a ground plane is a simple yet effective way of achieving better signal integrity and more resistance to interference in the board, simplifying the circuit's layout and avoiding the use of additional traces to connect all grounds together.

2.3.7 Enclosure

The PCB with the electronics soldered on it was protected by using an ideally suited plastic box. The enclosure contains all the relevant electronics that form the core of the controller, presenting connections for the external elements needed; power supply, USB communication, thermocouples and heaters. All these elements were equipped with the appropriate fittings, carefully chosen to ensure the operating conditions in the periods of peak current demand stay under their maximum current and voltage ratings. The enclosure with all its connections is shown in Figure 42 (right).

In the top part of the box, four rectangular cut-outs allowed the placement of the panel mount miniature socket thermocouple connectors.[219] This is achieved easily by using the fascia included with the socket without the need of either gluing or screwing. The connectors have unequally sized pins that prevent accidental reverse polarity connections, precluding human error. The thermocouples were fitted with the appropriate male plugs [220] supplied by the same manufacturer and both the plugs and the sockets were marked with their corresponding tags from T1 to T4 to prevent mistakes, since every thermocouple is associated with a specific heater in an individual temperature control for each reactor.

The side of the enclosure placed closer to both the laboratory computer and the wall sockets in the fume hood comprises the connections for the USB communication and power supply. The former will be achieved simply by cutting out a square destined to fit the Arduino USB port, allowing it to slightly stick out of the enclosure when the board is positioned in its fixed place inside of the box. In order to hold the board in position, four pairs of M3 nylon hexagonal female standoffs and bolts are used along the built-in Arduino mounting holes. The standoffs are attached to the bottom of the box using epoxy adhesive, providing a strong bond that results in a stable configuration.

The fitting for the main power connection is meant to safely endure the highest current drawn during the controller's operation, produced when the device is working at full capacity. This peak demand is of 4 A, while the voltage stays at 24 V. Due to these high-power requirements, standard DC jack fittings were discarded for the application, with specialised DC high-power fittings being selected for this connection. The Lumberg snap-in DC power socket [221] and plug [222] were selected as they present current and voltage ratings of 10 A and 24 V, respectively, able to provide a safe connection for the controller with minimal contact resistance (~30 mΩ). The same fittings were used for the four heater connections on the opposite side of the box, for which the flyback diodes were mounted inside the plastic cover of the plugs. Again, both the sockets and the plugs were tagged with their corresponding identification codes (from H1 to H4) to avoid incorrect connections by the user.

A space fitting the LCD screen was cut out in the box cover, connecting it with soldered wires to the corresponding pins of the PCB. The screen shield is attached to the box cover by means of screws and nuts through its mounting holes. An especially tailored laser cut part of PTFE was placed directly around the screen, in a close fit to prevent dust or other particles to enter the box as far as possible.

In this fashion, the box constitutes a piece of equipment that works with standard parts and can be easily connected, disconnected and transported.

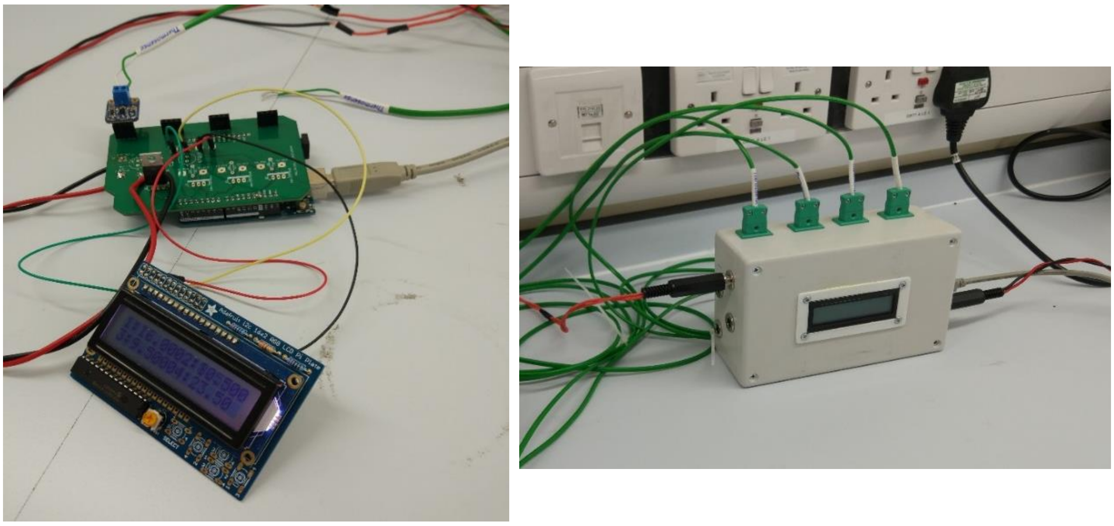

Figure 42. Left: PCB mounted over Arduino UNO during the development stage, with the LCD shield connected to it. The USB connection can be identified to the right of the set. Right: enclosure of the temperature controller.

2.4 Microprocessor Programming

2.4.1 Specifications

The microprocessor programming is designed in order to meet the following requirements:

- Effectively control the temperature of the system, providing a stable and accurate temperature for each reactor. This will be achieved regarding different aspects:
  - Data acquisition and smoothing.
  - Dynamic PID control.
- Successfully communicate with the master computer, which will control the process using algorithms developed in Matlab. This will include various subroutines with different objectives:
  - Set the temperature setpoint required by the master computer.
  - Send the current temperature or setpoint back when the appropriate command is received.
  - Turn the heating off when demanded.

The iterative process of optimisation of a reaction will generally run as follows: the master computer will set the operation conditions for an experiment, for which the chemical yield will be evaluated. Amongst these variables is the temperature of the system. A function in Matlab will communicate to the microprocessor the required setpoint over the serial port, and the microprocessor then will set the requested temperature. It will also provide temperature data when demanded. Finally, it will change the temperature of the reactors again if needed or will turn the heaters off if the appropriate command from Matlab is received.

This approach will require the development of an effective communication strategy that could be used by different optimisation algorithms and for different reactions.

2.4.2 Programming Language

Arduino boards use a combination of C and C++ as a programming language.[197] As an open-source project, Arduino also provides its own development environment, Arduino IDE, which will be used to write and upload programs to the microcontroller and also to monitor the serial communication between Arduino and a PC over the USB cable when needed. The latter is possible because of the built-in serial monitor that the Arduino IDE includes, which enables sending and receiving information from the board, and was extensively used during the development of the controller, being the main tool during debugging of the system and serial communication programming.

2.4.3 Program Structure

The program was developed iteratively and adapted through a series of evaluations of approaches. Once the program was proven successful, small modifications were included over time to improve its performance and avoid bugs. The final software structure is shown in Figure 43. A full version of the final code, along with detailed, specific notes on programming elements are given in Appendix C.

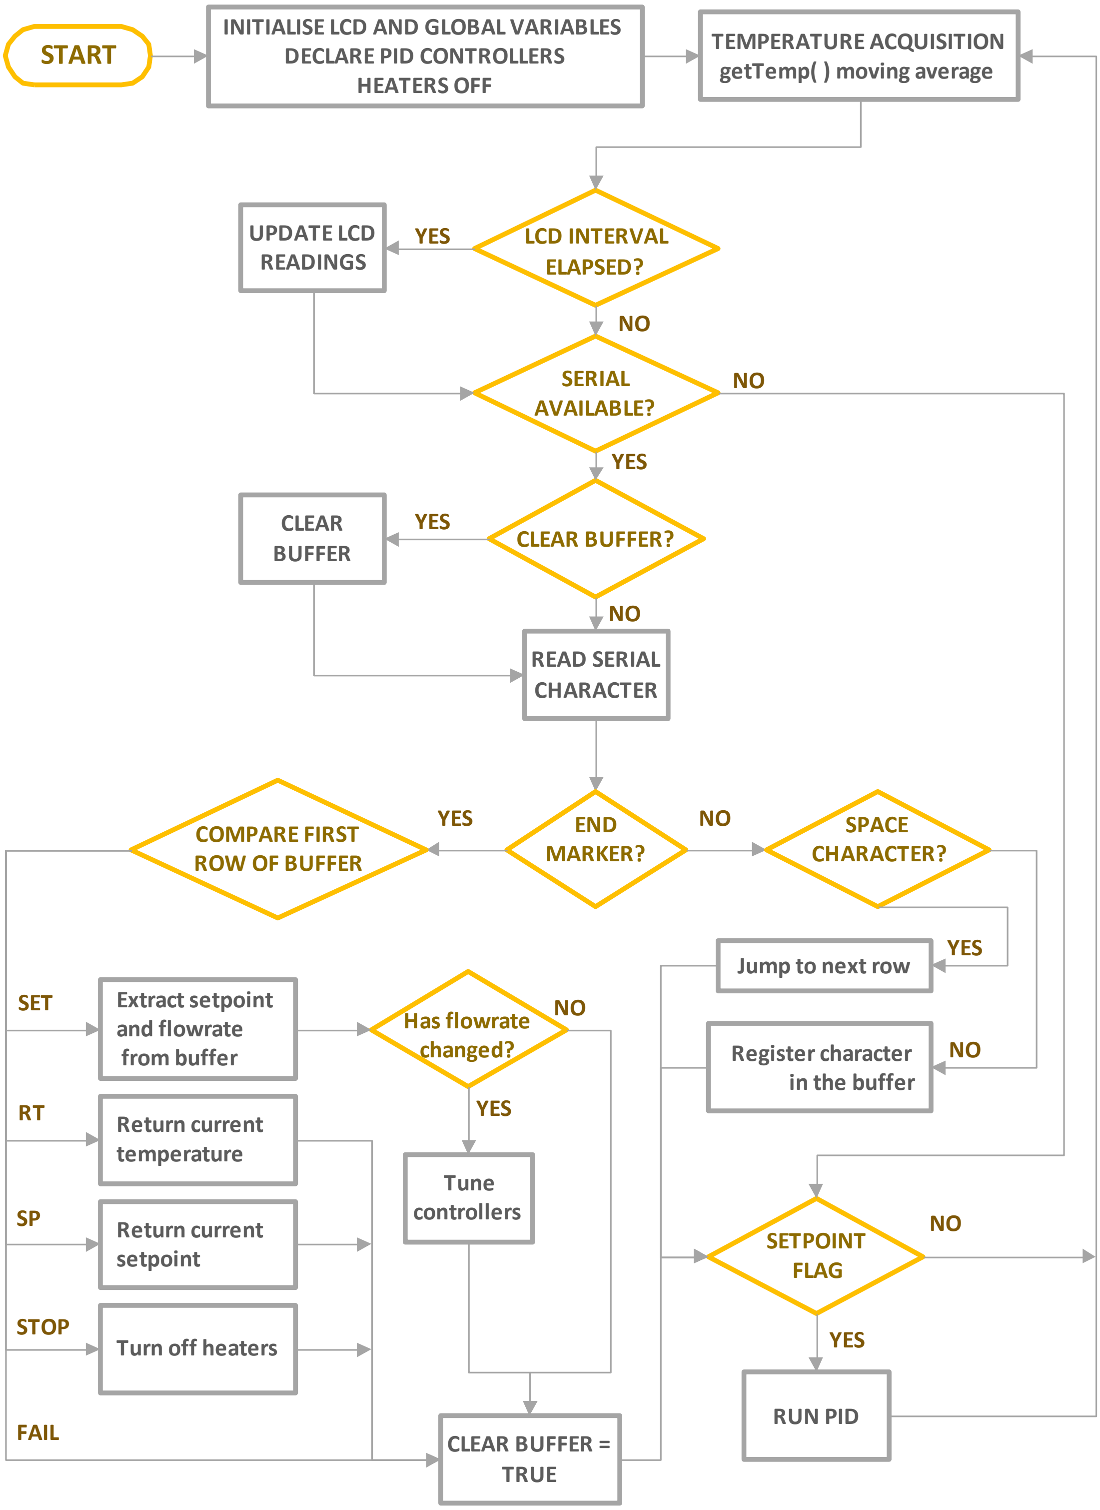

Figure 43. Flow chart of the software developed for the controller.

2.4.4 Data Acquisition and Signal Filtering

The Arduino analogic channels used to read analogic inputs include a 10-bit analogic converter (ADC). This means that input voltages are mapped into 2^{10} (1024) integer values (between 0 and 1023), known as quantization levels or ADC codes. The voltage reference of Arduino, which is 5 V by default, can also be set to 3.3 V, enabling the device to resolve 4.9 mV or 3.2 mV, respectively. These are the smallest changes that can be detected using the built-in ADC. However, this resolution does not allow the voltage values generated by the K-type thermocouple to be read directly by the board, forasmuch as they generate extremely low voltages, requiring a good amplifier with a cold-compensation reference in order to produce a readable signal in the range of the microcontroller logic level.

In the case of the amplifier selected for the application (AD8495), the device can be powered from 3 to 18 VDC and its output converted to temperature following the equation provided by the manufacturer [201]:

$$Temperature = \frac{V_{OUT} - 1.25}{0.005\ V} \qquad (26)$$

Using 5 V logic, and according to equation (26) the minimum and maximum temperatures susceptible of being read range from -250 to 750 °C, respectively; an interval of 1000 °C for which the Arduino ADC assigns 1024 codes. This means that the temperature resolution in this particular case is 1000 / 1024 = 0.9766 °C, a poor value for the application that would lead to a noisy signal. Nevertheless, Arduino has an internal 3.3 V generator that can be used in order to improve the resolution by setting the analogic reference to external in a single line of code and connecting said generator to the voltage reference pin of the board (AREF). Referring again to equation (26), it can be concluded that under this configuration the system will be able to read temperatures between -250 and 410 °C, improving the resolution when compared to using the default 5 V reference level, as the range has been reduced from 1000 to 660 °C but the number of ADC codes remains the same. In this case, the resolution achieved is of 0.6445 °C (660 °C / 1024). Although improving the resolution value achieved by default operating at a 5 V reference level, this remains a poor value. Fortunately, some techniques exist that can be used to exploit the signal’s noise and increase the resolution.

IMPROVING RESOLUTION

As discussed before, the combination of the thermocouple with the temperature amplifier is able to measure over a wide range of temperatures. However, this comes at the cost of poor resolution, as the number of ADC codes to map the entire temperature range remains 2^{10}, due to the resolution of Arduino’s built-in 10-bits ADC. Unfortunately, such configuration makes difficult to measure small changes in temperature (with the resolution constrained to 0.6445 °C for the 3.3 V reference, as demonstrated before). There are a few alternatives to address this problem, of which some are exposed here:

- The use of the 2^{10} ADC codes to map temperature values only in the range of temperatures of interest. As the values of interest in the current case may vary between room temperature and 120 °C, it is a waste to use a 3.3 V reference (which, as discussed before, enables reading temperature values up to 410 °C). Using an external voltage regulator, the range can be reduced to provide better resolution. Due to the nature of the temperature amplifier, the minimum temperature the device is able to read remains -250 °C independently of the reference voltage; but using a 1V external regulator and connecting it to AREF would lead to a range between said value and 150 °C, yielding a resolution of 400/1024 = 0.39 °C. This is a better resolution than the one achieved with 3.3 V logic, but nevertheless not considered precise enough for the application.
- Use an external 16-bits ADC. These chips would greatly improve resolution, being able to resolve approximately a hundredth of a degree with a voltage reference of 3.3 V (660 °C / 2^{16}) but since it uses the I^2C communication protocol it will come at the cost of communication-induced delays. The drawbacks of this approach are its expensiveness and the additional complexity it brings to the electrical design.
- Oversampling and averaging.[223] This is a common software-based low-pass filter used to improve resolution, at the cost of a more computational intensive algorithm and reduced throughput. The technique consists in collecting various temperature readings and averaging them, yielding improved resolution values while increasing the signal to noise ratio (SNR) in the process.

Amongst the three approaches introduced here, oversampling is the only technique that avoids the inclusion of new hardware elements, while still being able to enhance the resolution of the device. As such, it will be implemented in the microcontroller relatively easily through developing the pertinent software. It is convenient here to briefly describe how the technique produces improved resolution values from the 10-bits ADC.

OVERSAMPLING AND AVERAGING

For this technique to work, two criteria must be met: that noise shall exist in the sensor's signal and that such noise must fit a Gaussian distribution.[224] The rationale for this follows from the fact that, if noise is not present in the signal, all the steady voltage values would fall into the same ADC code. For example, in the case examined, working at a 3.3 V voltage reference provides a resolution of 3.2 mV/°C. If two adjacent quantization levels were to correspond to voltages of 3.2 and 6.4 mV, the input voltage will be rounded to its nearest level. If the input is 4 mV in absence of noise, the ADC would round it to the 3.2 mV level every single time the input signal is sampled, and oversampling and averaging would have no effect, producing a value of 3.2 mV for the response and incurring in a large error. On the other hand, if Gaussian noise is present in the signal, such that the signal is able to vary between at least the two ADC codes, and it is sampled for a large enough number of times, the ADC will round the value to 3.2 mV three times more than to 6.4 mV. Thus, a valid reading of 4 mV would be calculated by averaging the values produced by the ADC. Undoubtedly, the amplitude of the noise should be large enough to change randomly between two adjacent ADC codes (0.5 LSB) for the method to have any effect.[225] This is the principle that the technique exploits in order to provide enhanced resolution from ADC readings.[226]

In order to add b extra bits of resolution, the number of samples (s) to average must be:

$$s = 4^b \qquad (27)$$

Meaning that for each additional bit of resolution, the number of samples needs to be increased by a factor of four (for a mathematical proof of this see reference[224]). More specifically, the resolution achieved in the current case using 3.3 V logic, for which the temperature range is 660 °C, can be expressed as a function of the additional bits of resolution:

$$resolution = \frac{temperature\ range}{2^{10+b}} = \frac{660}{2^{10}\ 2^b} \qquad (28)$$

Taking into account equation (27) and isolating b:

$$s = 4^b \Rightarrow s = 2^{2b} \Rightarrow \log_2 s = 2b \Rightarrow b = \frac{\log_2 s}{2} = \log_2 \sqrt{s} \qquad (29)$$

And substituting now in equation (28), an expression for the resolution is found:

$$resolution = \frac{660}{2^{10}\sqrt{s}} = \frac{original\, resolution}{\sqrt{s}} \qquad (30)$$

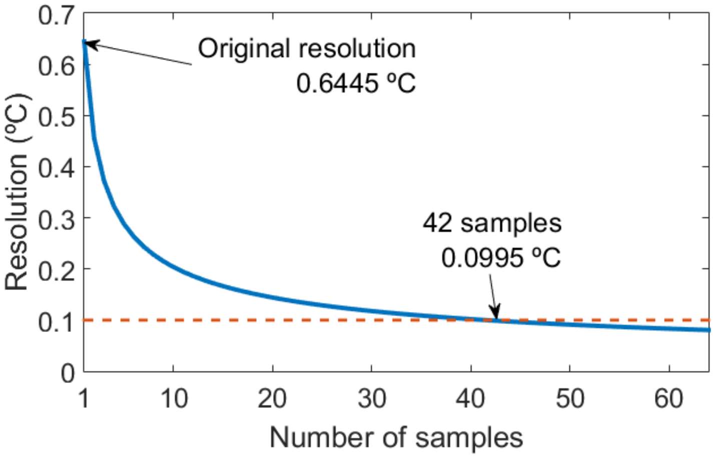

Figure 44. Resolution vs Number of samples used by the oversampling and averaging technique for the hardware employed in the application.

Figure 44 shows the resolution as a function of the number of samples used for averaging. Attending at the shape of the curve, a great improvement over the original resolution can be made with a relatively low number of samples. According to equation (30), 16 samples allow cutting the original resolution by a factor of 4, yielding a resolution of 0.16 °C. However, improving the resolution becomes more costly as the number of samples involved in the process increases. In order to achieve an effective resolution equal or lower than 0.1 °C, and in virtue of equation (30), a minimum of 42 samples are required, yielding a resolution of 0.0995 °C. This case is highlighted in Figure 44, and it will be the number of samples used to implement the filter in Arduino.

When working with normal averaging, this increase in resolution comes at the cost of throughput, as the s samples must be taken before the mean can be computed, filling a buffer and reducing the frequency at which the calculated temperature values are received. Nonetheless, a rolling average (also referred to as ‘moving’ or ‘running’ average in the literature), equivalent to a low-pass filter, can be implemented in Arduino IDE to avoid this reduction in throughput. This enables providing temperature values at roughly the same rate the sensor does. Instead of waiting for the buffer to be filled with the s values read by the ADC before calculating their average, a FIFO (first in, first out) approach is taken. In this fashion, once the buffer is filled for the first time, the next incoming value replaces the oldest one in the buffer. However, the advantages of averaging come at the cost of a slight time delay, as the calculated temperature values are the mean of the previous s values sampled. Let f_s be the sampling frequency, then the delay Δt can be expressed as:

$$\Delta t = {}^{s}\!/_{2 * f_s} \qquad (31)$$

By now, it should be apparent that there is a trade-off between resolution and time delay, as a consequence of the increasing number of samples needed to improve the former. As the frequency at which Arduino was capable of collecting data from the sensors was of ~25 Hz, the time delay is kept under a second for the desired resolution of 0.1 °C.

Oversampling and averaging was implemented relatively easily within the Arduino. The developed function, called *getTemp*, accepts a thermocouple number (ranging from 0 to 3, one for each thermocouple connected) as an argument and returns the oversampled and averaged value of the temperature in °C for the corresponding thermocouple. The function samples the mapped ADC values and averages them, before implementing equation (26) and returning the temperature value in double format. It makes use of global variables such as A, which is a 2D array with *RASIZE* rows and 4 columns. The former is the number of samples used in the oversampling, s, while the latter is the number of thermocouples used in the setup that the controller has access to.

First, the function checks if the number of elements used in the oversampling has been reached or not, in order to provide the appropriate mean of the readings. For each thermocouple, the integer variable *‘stindex’* tracks the element of the readings array being filled, while the type double variable ‘*total*’ is the sum of all the values already stored in the readings buffer. This is done in order for the function to always return the average of the values it has received even before the array is completely full (*stindex<RASIZE*) by dividing the updated values of *total* by *stindex* in every call. When the first condition in the function is met and the array is full, the state of the boolean variable ‘*start*’ is changed to false to bypass this part of the code once the readings array has been completely filled for the first time.

```
double getTemp(int tcnumber)
{
  if(start[tcnumber] == true && stindex[tcnumber]<RASIZE)
  {
      A[stindex[tcnumber]][tcnumber] = analogRead(tcnumber);
      total[tcnumber] += A[stindex[tcnumber]][tcnumber];
      stindex[tcnumber]++;
        if(stindex[tcnumber] == RASIZE)
        {
          start[tcnumber] = false;
        }
  }
```

From this moment on, the program will replace the oldest value in the array with the new reading from the thermocouple amplifier. The oldest value stored is subtracted from the auxiliary variable ‘*total*’ before adding the new value to it. The integer variable ‘*index*’ tracks the value being replaced in the array, and once it reaches the end of it (*RASIZE*), it is set to 0 again (since in C++ arrays are zero indexed). In this fashion, the oldest value is replaced in the array every time the function is called with a minimum use of memory, as opposed to the commonly found technique in which all the values are displaced in the array every time a new one is included.

```
 else
  {
      total[tcnumber] = total[tcnumber] -
A[index[tcnumber]][tcnumber];
      A[index[tcnumber]][tcnumber] = analogRead(tcnumber);
      total[tcnumber] += A[index[tcnumber]][tcnumber];
      index[tcnumber]++;
         if(index[tcnumber] >= RASIZE)
           index[tcnumber] = 0;
  }
 double avg = total[tcnumber]/stindex[tcnumber];
 double Vin=(3.3*avg);
 float vinl=Vin/1024;
 double Vout=vinl-1.25;
 double T;
     T = Vout/0.005;
     return T;
}
```

The implementation of this function in the main code of the microcontroller allows to gather temperature values and analyse them. In order to assess the real effect of this approach, the room temperature was tracked using both 5 V and 3.3 V, with and without oversampling and averaging in both cases. The results are shown in Figure 45.

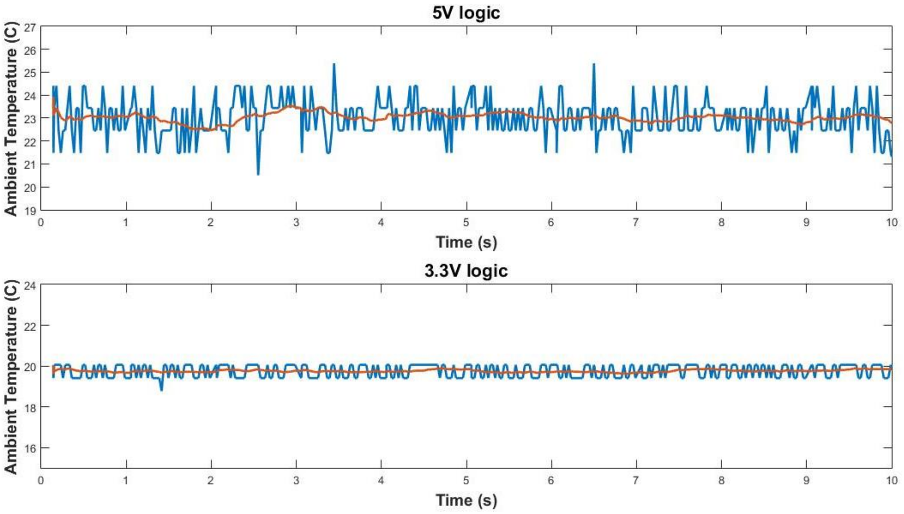

Figure 45. Raw (blue) and filtered (red) temperature readings for 5V (top) and 3.3V (bottom) logic levels.

Both working with a reference of 3.3 V and using oversampling and averaging (running average filter) provide a smooth signal to the rest of the program, this being especially significant in the case of a control algorithm in which any peaks from the input signal could potentially lead to instability and decreased performance.

Table 4. Statistic values for the analysed signals.

Logic	Signal	Max Temp °C	Min Temp °C	Temp Median	Temp Mean	Temp Mode	σ	σ^2
5V	**Raw**	25.39	20.51	23.44	22.99	22.46	0.8722	0.7607
	Filtered	23.93	22.46	23.01	23.00	23.01	0.2128	0.0453
3.3V	**Raw**	20.06	18.77	20.06	19.76	20.06	0.3275	0.1072
	Filtered	20.06	19.62	19.75	19.76	19.75	0.0612	0.0038

Table 4 contains the values of the maximum, minimum, median, mean, mode, standard deviation and variance for the raw and filtered signals for each logic level. The values of the standard deviation are better for the filtered signals in each case, improving further when 3.3V logic is used. The overall best values of both standard deviation and variance correspond to the filtered signal with 3.3V logic, which reveals itself as the most stable and reliable signal analysed in this study and consequently this filter was chosen to be used in the temperature control.

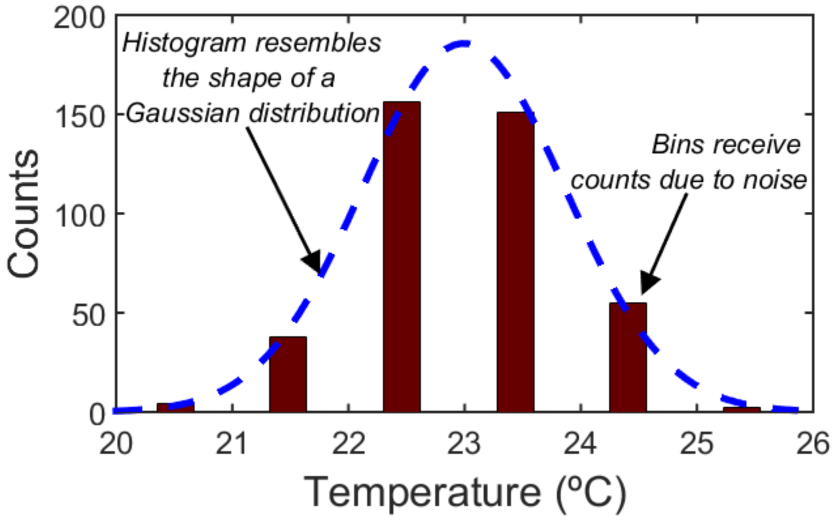

Figure 46. Histogram of unfiltered temperature samples at room temperature using a 5 V reference.

Figure 46 shows a histogram for the temperature samples at room temperature for a 5 V reference voltage. As explained before, and as every temperature value comes from the conversion of an ADC level, they fall into discrete temperature levels. When a Gaussian PDF is fitted to the data, it becomes apparent that the nature of the noise affecting the readings is that of white noise, and accordingly the application fits the criteria for which oversampling and averaging is effective.[224] When a 3.3 V reference level is used, the values fall within only two to three levels, meaning that not only the resolution is higher, but that the noise is greatly reduced in amplitude. However, both levels share the Gaussian nature of the noise.

CALIBRATION

Attending to Figure 45, a clear offset exists between the 5 and 3.3 V signals. This is a well-known issue when working with this kind of chips with different voltage references. To make sure the temperature is accurate, it was calibrated using as a reference the thermocouple submerged in water flask on a hot plate, systematically registering the temperature read by Arduino for different hot plate temperatures of 30, 40, 50 and 60 °C. Results are compiled in Table 5.

Table 5. Calibration results.

T Hot Plate (°C)	30	40	50	60
T Arduino 5V	30	40	50	60
T Arduino 3.3V	27.2	37.2	47.2	57.2

The values obtained by Arduino with a 5V logic are in excellent agreement with the values given by the hot plate thermocouple, while an offset of 2.8 °C is detected when using 3.3 V logic for the thermocouple amplifiers. In addition, each thermocouple amplifier produces a different offset, and accordingly specific calibrations for each of the four amplifiers were included in the code to make sure the temperature values agree with those read by the reference. This is enabled by including a switch statement in the getTemp code before calculating the temperature, shown here for convenience.

```
switch (tcnumber)
 {
     case 0:
       T = Vout/0.005 + 2.8;
       break;
     case 1:
       T = Vout/0.005 + 1.385;
       break;
      case 2:
       T = Vout/0.005 + 4.22;
       break;
       case 3:
       T = Vout/0.005 + 3.45;
       break;
     default:
       break;
 }
```

2.4.5 Serial Communication

The Arduino UNO board employed in this project has a number of different ways to communicate with a computer or another device. The chip ATmega328 provides UART TTL (5 V) serial communication, which is available on digital pins 0 (RX) and 1 (TX). However, an ATmega16U2 chip acts like a bridge, channelling this serial communication over the USB port, appearing as a virtual COM port on the computer. This USB connection will be used to establish a serial communication with Matlab, enabling logging and monitoring temperature values over the Arduino IDE serial port.

Serial communication allows the transference of characters between Arduino and a program running in Matlab as shown in Figure 47. It can be noticed that contrarily to Matlab’s function *fscanf*, which automatically reads the character string sent by Arduino and returns a number in float format, a function in C/C++ has to be developed for the Arduino in order to read the characters one by one until a certain terminator (which here was chosen to be \n, the newline character) is found. Then, the *atof* function can be used to transform the character string into a float number.

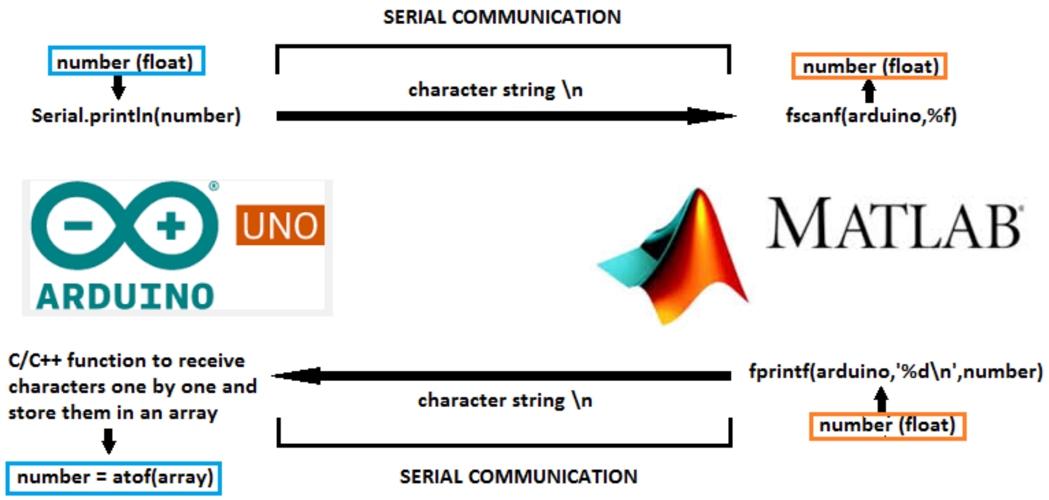

Figure 47. Data flow for the serial communication between Arduino and Matlab.

A string of characters with a predefined end marker will be sent over the serial communication line from the computer to Arduino to i) select a new setpoint, ii) request the current temperature or setpoint of the system or iii) stop the controller’s operation, turning off all heaters.

ARDUINO SIDE

The controller will behave as a slave device to the laboratory’s computer, which will send different commands over the serial port. This command is received as stream of characters that needs to be analysed in order to come up with the action of the controller required by the master. The commands follow a predetermined form for the incoming character string where the end marker is a newline character (\n):

ACTION VALUE1 VALUE2\n

Action specifies the type of action required from the controller, while value1 and value2 (optional) provide extra parameters. The spaces act like separators between the different elements of the command, and the newline character specifies the end of the command. The programming strategy implemented to read this type of command was to create a matrix in which different elements of the command are stored in different rows. In this way, ACTION is stored in the first row, and if a space is received then the program will jump to the second row and store the characters for VALUE1 in the second row. Again, if a space is received, the program jumps to the next row to store VALUE2. Whenever the terminator is received the operation stops and the algorithm proceeds to check the type of operation required from the first row received. This is a simple State Machine. The commands implemented for this purpose can be found in Table 6.

Table 6. List of commands sent from Matlab.

COMMAND	ACTION
SET	Fix the temperature setpoint of the system to the number specified in the second row.
RT	Return temperature of the system to Matlab.
SP	Return current setpoint.
STOP	Interrupt the algorithm and turn off the heaters.

With an algorithm of this kind, it is easy to include more data to select the most appropriate tuning parameters for the temperature controller from a predefined list elaborated from previous experimental data from step responses. For example, the master computer could feed Arduino with data relative to the set point and flow rate of the next experiment:

SET 45.3 1.25\n

And the algorithm can proceed to change the setpoint to 45.3 °C and select the most appropriate PID tuning for it taking into account the flow rate through the CSTR cascade in said experiment. This approach is known as ‘gain scheduling’ and will be explained in detail in section 2.4.6.

MATLAB SIDE

The GUI that has been developed over the years for the automation of HPLC analyses in flow chemistry includes a set of functions to interact with different temperature controllers, allowing the user to include new controllers, provided that they develop the required set of functions for them. A structure in Matlab must be created for every controller, where every functionality has a name that points to a function specially tailored for it. The structure is composed of 33 fields, among which there are fields for temperature, pressure and flow rate control. Among these, the most important functions developed for the application can be found conveniently reproduced and commented in detail in Appendix D.

Next, the structure for the Arduino Temperature Controller along with the set of functions developed for it are included into the accessible Matlab's path during the GUI operation. Once this is done, a pull-down menu shows *Arduino Temperature Controller* as an option that can be selected, for which a COM port has to be chosen. The device is connected to the USB port of the computer by means of a USB cable type A/B. This structure has been proven to work successfully, without producing any incident in the many hours of automated operation of the system since it was implemented.

2.4.6 PID Control

Proportional, Integral and Derivative or PID control is a popular approach to process control that provides stability to the process variable.[227] A PID library developed by Brett Beauregard for Arduino exists and was used for the implementation of temperature control.[228] However, the stability of the control is highly dependent on an effective tuning of the controller. In essence, a PID is used to control a process variable ($PV(t)$) by continuously monitoring it and calculating the error $e(t)$ between it and the desired setpoint ($SP(t)$) using feedback. The algorithm then provides a control output, $CO(t)$, that is a function of this error. As its own name states, a PID control consists in the combination of three different terms:

$$CO(t) = K_p e(t) + K_i \int e(t)dt + K_d \frac{de(t)}{dt} \qquad (32)$$

$$where \quad e(t) = SP(t) - PV(t)$$

- The first term, known as the proportional term, provides a response proportional to the error between the process variable and the setpoint. By itself, it constitutes one of the most basic kinds of control, with the main drawback being the induction of a steady-state or stationary error.
- The second term integrates the error over time, with the aim of eliminating the stationary error that a proportional term by itself is likely to introduce. If the process variable is close to the setpoint, such that the proportional term does not produce an output high enough to reduce the stationary error, the integral term remains summing the errors continuously until the stationary error is removed.
- The third term in the equation, or derivative term, is used to speed up the control, increasing the overall output when the rate of error change is high,

and constraining it when it is low, finally reaching a value of zero when the process variable remains constant over time.

In consequence, the controller's output can be envisioned as a weighted sum of the error, $e(t)$, the integral of the error over time, $\int e(t)dt$ and the time derivative of the error $de(t)/dt$ that uses the sum weights K_p, K_i and K_d, respectively. The process by which these values, known as the PID gains, are calculated in order to produce a reliable, stable and accurate control is known as PID tuning.

PID TUNING

Once a serial communication structure has been implemented for the controller, effective tuning can be carried out, allowing the user full control of the device through the Arduino IDE serial port. The tuning of a PID controller is the appropriate selection of the proportional, integral and derivative gains. Most of them are based on the analysis of the response of the system to some sort of input. Different tuning methods have been developed over the years such as the popular Ziegler-Nichols,[229] Cohen-Coon [230] and Lambda [231] tuning methods. In this case, Lambda tuning was chosen for a number of reasons, including:

- Ziegler-Nichols is based in the increase of the proportional gain until the output of the controller produces stable and consistent oscillations, which in the present case could lead to excessively long waiting times.
- Lambda tuning rules are less sensitive to potential errors in the calculation of the dead time of the process for a step response. Both Ziegler-Nichols and Cohen-Coon can lead to bad results if the dead time is measured incorrectly.
- The tuning obtained is robust and the controller will remain stable even when the conditions of the process differ significantly from those of the tuning process.
- The system will absorb perturbations, which will only minimally affect the process. This is the main reason for its popularity within industrial applications, especially in the pulp industry.
- The user is able to specify the desired response time (time constant) of the closed control loop, within the specified limits.
- Overshoot is highly minimised or completely removed from the closed-loop response. This is an important feature in this application in order to produce a quick temperature controller in the absence of active cooling.

DETERMINATION OF THE PROCESS CHARACTERISTICS

Lambda tuning consists in analysing the behaviour of the system for a step input. The controller's output is set to a constant value and the evolution of the temperature for the open-loop response is monitored. For a first order dynamic, the response of the system will be similar to the one showed in Figure 48. Lambda tuning allows the calculation of the tuning parameters by the analysis of this response, being this a straightforward tuning method. However, this kind of tuning provides only the proportional and integral gains of the controller; in other words, it produces a PI control. Nevertheless, the derivative term is quite often removed in industrial applications since it may introduce instabilities caused by the noise in the signals. Alternatively, additional filtering is used in the error calculations for the derivative part of the controller's output.

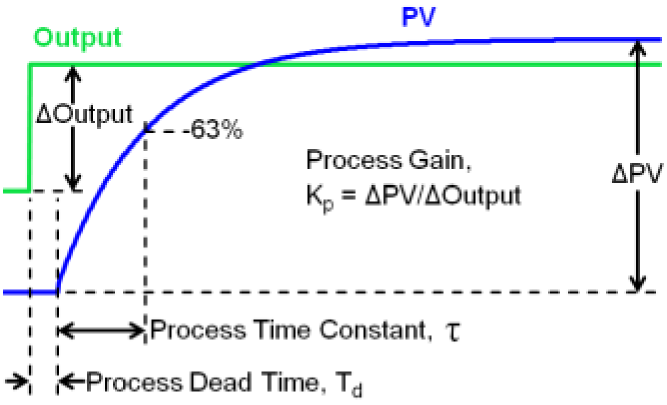

Figure 48. Process parameters from the step response of the system. Reproduced from ref.[232]

Attending at Figure 48, the process gain can be expressed as:

$$K_P = \frac{\Delta PV}{\Delta Output} \tag{33}$$

To calculate the process dead time (T_d) is necessary to measure the time lapse between the step in the output signal until the process variable starts to raise. The process time constant (τ) is defined as the elapsed time between the time at which the process variable starts to rise until it reaches 63% of its final steady-state value. The temperature response to an input signal was recorded using Arduino IDE's serial monitor and transferred to Matlab in order to accurately calculate the dead time and the process time constant, and ultimately tune the controller.

LAMBDA PARAMETER SELECTION

A characteristic feature of this tuning method is that it allows the user to define the time constant for the control loop, lambda (λ), within certain limits. It is worth considering that a large value of this parameter produces a stable but slow control, while low values lead to faster and more aggressive controls. In general, λ values are chosen between one and three times the value of the dead time of the process.

$$T_d < \lambda < 3T_d \tag{34}$$

In this case, and in order to produce a quick temperature control, lambda was chosen equal to the dead time of the process to calculate the parameters of the controller. This was done with the aim of achieving a quick temperature control. However, if the control results unstable, this initial result can be refined manually to obtain a slower but more stable one. There is a trade-off between speed and stability, and although the rule in equation (34) is designed to provide stable controls there is a certain degree of flexibility for the user to choose lambda.

$$\lambda = T_d \tag{35}$$

GAINS SELECTION

Let equation (36) be the PI controller's output equation written in the standard form:

$$CO(t) = K_c\left[e(t) + \frac{1}{T_i}\int e(t)\right] \tag{36}$$

Lambda tuning rules allow easy calculation of the proportional gain using the values yielded from the system's open-loop step response analysis by using:

$$K_c = \frac{\tau}{K_{proc}(\lambda + T_d)} \tag{37}$$

The rule for the integral time states that:

$$T_i = \tau \tag{38}$$

Given that the proportional and integral gains used by the PID library refer to the parallel form for the PI expression:

$$CO(t) = K_p e(t) + K_i \int e(t) \qquad (39)$$

A simple parameter conversion allows to use the lambda rules for the controller. Writing K_p and K_i as functions of the parameters obtained from the open-loop step response:

$$K_p = K_c = \frac{\tau}{K_{proc}(\lambda + T_d)} \qquad (40)$$

$$K_i = \frac{K_c}{T_i} = {}^{1}\!/_{K_{proc}(\lambda + T_d)} \qquad (41)$$

DYNAMIC TUNING

The controller produced by the newly introduced equations is supposed to be used under the same process dynamics for which the step response experiment was carried out. However, in practice it will be used over a broad range of flow rate conditions, from extremely low flow rates in the μL/min range to a few mL/min. It is clear that the dynamics of the process vary depending on the flow rate, so that if the same tuning parameters were kept for all different flow rates, the stability of the control would be compromised or it would become more sluggish.[233] For example, at low flow rates in the range of μL/min, the process will be much more responsive to changes in the controller's output than when running at a few mL/min. In this regard, when operating at low flow rates, the controller would only need to apply a small change in the output to overcome any error in the temperature; while if working at high flow rates the same output would produce only a minimal change in temperature. Accordingly, gain scheduling was implemented, so that the controller can make use of different proportional and integral gains depending on the flow rate that is flowing through the system (the gain scheduling variable).

In order to implement this strategy, the system was tuned for different flow rates across the operative range of the system (from 0 to 4 mL/s), recording the step response for each one of them. The dead time (T_d) and process time constant (τ) parameters were extracted from the responses, calculating a first approximation of the tuning parameters for each flow rate assessed by means of equations (40) and (41). Software was developed that enabled the user to connect the controller to a laptop, which accepted '*START'* as a command once the heaters were required to heat, being the response of the system to the step input recorded from this moment on. This enabled precise measurement of both the dead and process times. As the range of usual flow rates used in optimisation experiments range from low flow rates such as 0.25 mL/min to a few mL/min, step (open-loop) responses were recorded for flow rate values of 0, 0.25, 0.5, 1, 2 and 4 mL/min.

A cascade was formed with the reactors using standard 1/8' PTFE tubing and flangeless fittings with ¼-28 threads compatible with the ports in the fReactors. Two large 60 mL syringes were loaded with water and toluene, respectively, and fed into the reactors with equal flow rates using a Harvard Apparatus 33 syringe pump, creating a mixture of water and toluene similar to those that are encountered during operation in the laboratory. Finally, the PI controllers were tested for each flow rate to manually introduce the last changes in order to produce a quick, stable and precise temperature control, avoiding overshooting. In the vast majority of the cases, these were limited to modifying the integral gain to remove any remaining stationary errors incurred by the controller.[227]

Table 7. Kp and Ki values for optimal closed-loop performance.

Flow Rate (mL/min)	**Kp**	**Ki**
0	65	0.02
0.25	70	0.02
0.5	75	0.03
1	75.5	0.03
2	85.5	0.04
4	115	0.05

PID PERFORMANCE

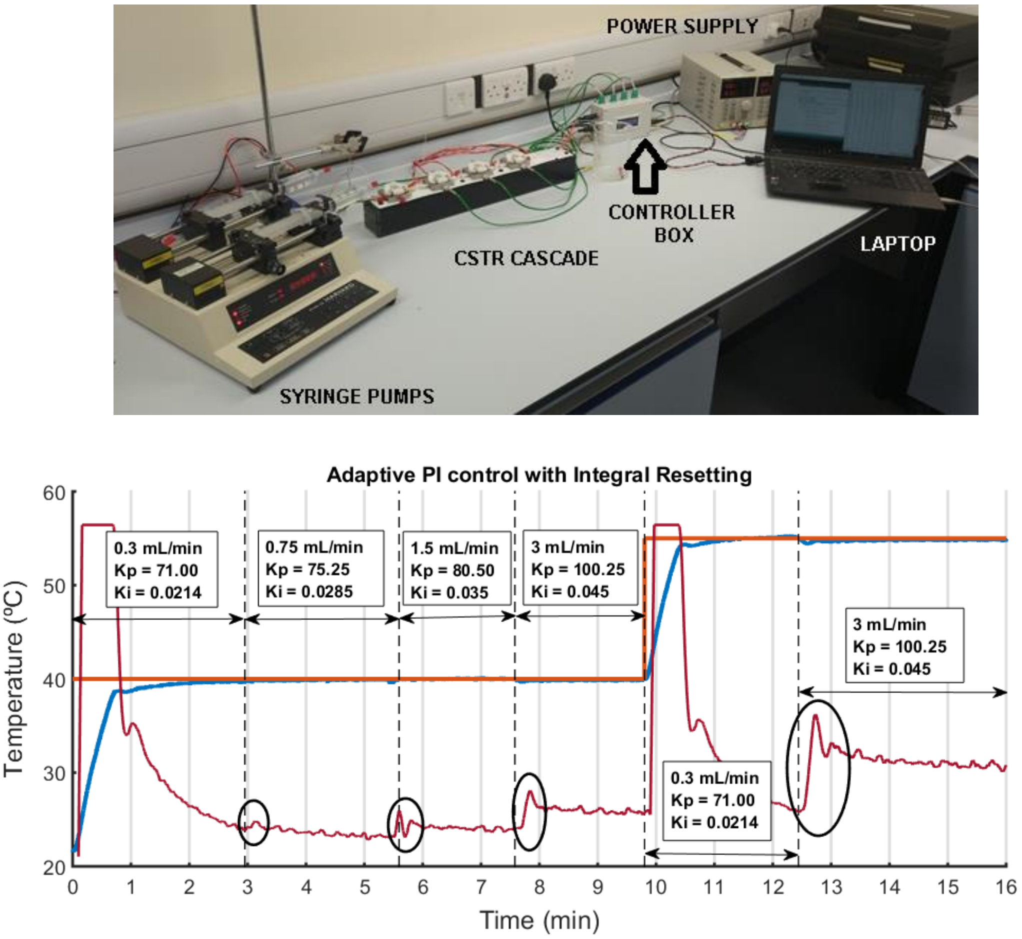

Figure 49. System development workbench (top). Temperature control achieved for the PI gain-scheduled controller under different conditions (bottom). Blue line: temperature reading. Red line: setpoint required. Garnet line: controller output just for illustrative purposes (not to scale).

After calculating the optimal values of the gains of the controller for closed-loop performance, the gain scheduler was implemented. A piecewise linear interpolation of the gains obtained (Table 7) as a function of the flow rate was used in order to provide reliable controllers in the regions between the sets of data for which the gains were calculated. Figure 49 shows the behaviour of the controller under different conditions and changes in the setpoint and the process and gain scheduling variables. The flow rate conditions were selected to be the most adverse possible, in the mid-points between the adjacent sets of data for which the controller was tuned. It can be appreciated that at the moment when the setpoint remains unchanged but the flow rate is increased, the temperature minimally drops, remaining close to the setpoint, as the change in the controller output rapidly prevents a bigger drop in temperature from happening, spiking as soon as the temperature begins dropping (shown in the black circles in Figure 49 (bottom)). The proportional and integral gains are modified every time the flow rate changes, producing larger outputs for the same temperature error as the flow rate becomes larger. The interesting features in the figure are the quick temperature control and the robustness achieved through its dynamic tuning. An accuracy of ±0.3 °C is achieved, which represents a great improvement when compared with previous setups used in iPRD that presented ±1 °C. Looking at Figure 50, the fReactors are also a quicker alternative to the Polar Bear Plus that has been used at iPRD. In the case illustrated, the fReactors are able to reach the setpoint accurately within three minutes, while the Polar Bear is only able to reach it slightly under six minutes; for these specific conditions the fReactor is almost twice as fast as the Polar Bear. This is important when considering not only time savings, but also the amount of sometimes expensive reagents that are wasted during the transition between experiments carried out at different temperatures, which under these circumstances could be halved by using the fReactor. Furthermore, the temperature readings of the fReactor refer to the fluid in the reactors, while in the case of the Polar Bear is the temperature of the reactor itself, so longer waiting times are needed to ensure the fluid reaches the same temperature as the reactor before being able to carry out the experiment.

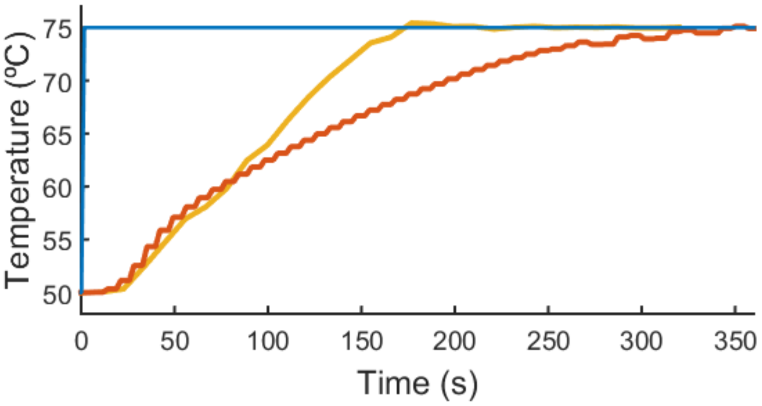

Figure 50. fReactor (orange line) and Polar Bear Plus (red line) closed-loop performance comparison. Data for both devices was recorded during two different optimisations once integrated in the system, for which the temperature change between experiments was the same (from 50 to 75 °C).

2.5 Additional Hardware

2.5.1 Stirrer Plate

As commented before, the fReactors could be used over a commercially available multi-position hot plate to induce the stirring. However, it was noticed that poor magnetic contact prevented the magnetically coupled cross to spin smoothly, provoking continuous bouncing and misalignment. This may bring benefit in terms of mixing, but the desire here was to create a well-controlled flow field. In order to overcome this effect and provide a good and robust magnetic coupling, a specially tailored multi-position stirrer was designed and built for the fReactors.

The unit is comprised of five stirring modules. Each of these is built from 12 V computer fans onto which a metal disc (Ø ~30 mm) with 2 high-performance (3.2 kg pull) neodymium magnets glued on is mounted using epoxy. Placing these at the right distance to the bottom of the fReactors provided excellent magnetic coupling, allowing the magnetic cross to perfectly spin in the centre of the reservoir without producing any bouncing. A 500×56 mm strip of Delrin was designed that permitted holding five of these modules by means of through holes aligned with the four mounting holes in the corners of the fans, using nylon bolts and nuts, as shown in Figure 51 (top). Circular holes (Ø 30 mm) were cut in the holder that provided space to the spinning magnets. All the fans were connected in parallel with the aid of connection terminal blocks glued on the bottom part of the holder. The strip with all the modules fitted on it was then slid in a square-section PVC tube, and then hung from its top part by means of M6 bolts going through manually cut holes in the top of the case that matched the threads created with a tap set in the Delrin strip. This allowed placement of the holder in the interior of the tube, allowing fine height positioning and providing excellent and robust stirring never achieved before on this type of reactor with commercially available hot plates at iPRD. The extremes of the tube were closed using the matching lids purchased from the same provider, where one of them was fitted with a standard jack connection that allowed powering the stirrers. A PTFE insulation layer was designed to be placed on top of the stirrer, considering that the high temperatures the fReactors may reach could damage the PVC case. A second layer with hollowed out spaces for the fReactors and the heads of the go-through M6 bolts allows precise positioning of the fReactors on top of the device, as shown in Figure 51 (bottom). All Delrin and PTFE designs were laser cut, with only the M6 threads in the holder needing manual completion.

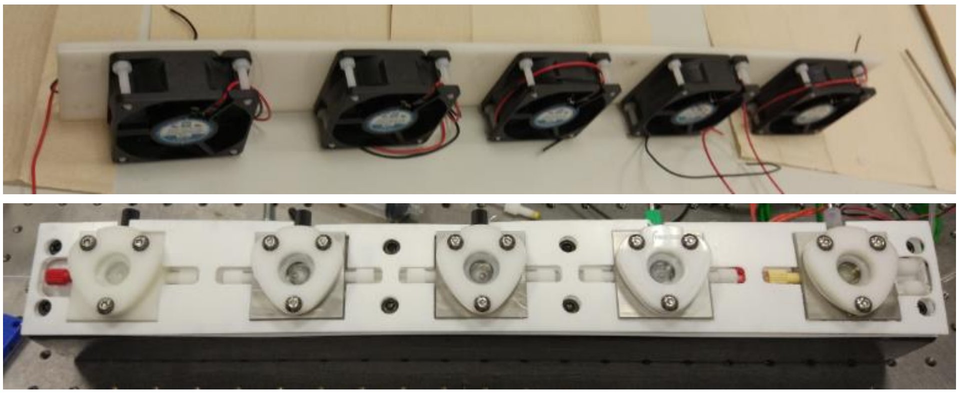

Figure 51. Top: bottom view of the inner stripe of the stirrer with the fans fitted on it. Bottom: finished assembly with the reactors placed in position and showing excellent stirring.

The resulting unit provides excellent stirring at higher rates than those obtainable with the previously used hot plate, establishing a handy plug-and-play device that only requires the user to plug in the electrical jack connection and position the fReactor cascade on top of the device using the PTFE layer designed to this end, which is aligned by using the heads of the M6 bolts as a reference.

2.5.2 Photochemical Setup

Interest in photochemical fluid flow applications has grown over the last decade,[15, 16, 63, 234] with some of its recent applications being the photoredox reactions via segmented flow [62, 63, 235] and automated quenching screening and Stern-Volmer analysis.[236] Within lab-scale flow photochemistry, segmented flow has been the preferred technique to avoid clogging when handling solids, but these reactions were highly constrained due to the small cross-section of the reactors used previously.[237] In view of the current limitations, it was decided that an add-in for the fReactors would help expanding their use to the UV photochemical field.

A high-power UV LED (LedEngin LZ4, 365 nm emitter) was selected for the application. The LED was purchased assembled in a star-shaped MCPCB that allows fitting it to a heat sink, as these devices generate heat when producing light, what could lead to overheating and ultimately LED failure. To dissipate the excess heat, a heat sink was used that provided special thermal interface material and matching threaded holes to fit the LED to it, forming the basic UV unit. These heat sinks consist of aluminium alloy cylinders with thin longitudinal fins in the radial direction. Standard optical components purchased from ThorLabs were used to allow manual positioning of the UV units over the fReactor. Optical posts (TR200/M) were mounted over a 12 mm thick 500×200 mm Perspex sheet onto which M6 taps were created to be used in conjunction with the matching setscrews provided for the posts. Right-angle clamps (RA90/M) including hand screws were attached to the posts allowing holding 75mm long optical posts (TR75/M) horizontally. These posts include M4 setscrews at one end that were used to hold 0.5” thick threaded cages (CP02T/M).

The UV units were placed over the upper part of these cages by means of long M3 threaded rods, nuts and washers. The rods were mounted through two holes in diagonally opposite corners of the cage and fitted between adjacent fins of the heat sink. The assembly was tightened using nuts and washers both on the bottom side of the threaded cage and on the top part of the heat sinks. This provided secure location of the UV units. One of the rods was also used for every unit to hold 5 V computer fans to allow for heat dissipation via forced convection, as the initial tests showed that the heat sinks did not dissipate enough heat with natural convection alone to avoid failure of the LEDs. In this regard, the flexible nature of the assembly allows easy positioning of the UV units by manually adjusting the posts using the thumbscrews in the right-angle clamp and could be used with both the tailored stirring block and commercially-available multi-position hotplates.

However, using the system under this configuration, and due to the wide viewing angle of the LEDs (110 degrees), would waste much of the energy, as only a fraction of the light would fall on the window over the fReactor reservoir. In order to improve the performance of the kit, the addition of a lens was used to focus the UV light on the fReactors’ window. From basic optics calculations, an aspheric condenser lens (ACL2520U-A , ThorLabs, Ø 25 mm) produced the best results used in conjunction with a lens tube (SM1L10) that was screwed in the central hole on the bottom side of the threaded cage. This configuration allows precise adjustment of the lens height inside the tube by using 2 retaining rings in its threaded interior surface, clamping it. This was manually refined to produce the optimal results when placing the lens tubes in direct contact with the lid of the fReactor over its window, as Figure 52 shows.

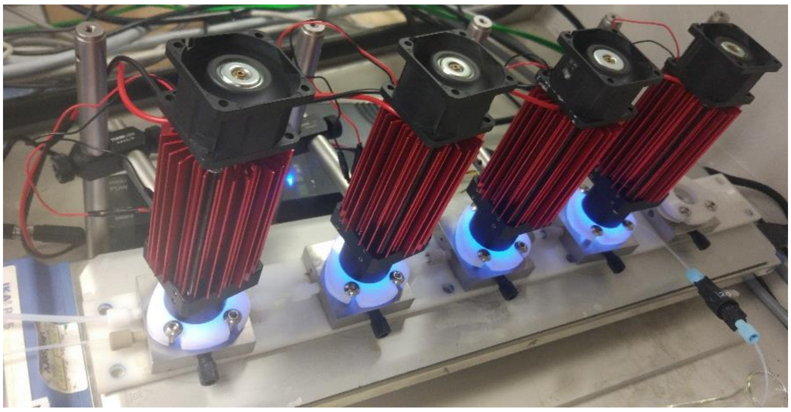

Figure 52. fReactors' UV kit during operation using a commercially available multi-position stirrer.

A high-power driver for the LEDs (eldoLED 90D) was used and programmed to feed all four LEDs from a 20 V AC/DC adapter. A jack connector was fitted to the high-power driver, connecting in parallel the driver and a DC/DC converter that produced 5 V in order to feed the cooling fans. The kit was proven to be a reliable and robust tool intensively employed at iPRD. Its use led to ground-breaking research featuring the automated experimentation of multiphasic flow photochemistry in this kind of reactors for the first time,[238] leading the way for other pioneering groups such as Jensen's at MIT.[239]

2.6 Summary

This chapter focused on the design and development of a heated fReactor cascade, with the aim of using it in an automated flow chemistry rig. It presented insights on the mechanical design and the RTD characterisation of the device. The introduction of a quick response heating cartridge and a thermocouple that can be used along with standard flangeless fittings using one of the fluidic ports in the body of the reactor makes possible a quick -cutting heating times by almost 50% when compared to the PolarBear reactor- and precise (±0.3 ºC compared to ±1ºC) temperature control.

These improvements in performance were achieved by means of a carefully engineered electronic system, for which a decimating and averaging technique was applied to filter the raw temperature readings. In addition, the approach taken frees the laboratory PC from additional workload, since all calculations relative to the PI temperature control are carried out in the Arduino. Such control was tuned under different flow rate conditions to finally come up with a dynamic controller that is able to deal efficiently with abrupt changes in flow rate.

A set of communication routines were developed that enabled secure communication with the laboratory PC. The conjunction of all these features led to a reliable piece of equipment that worked without problems during the many hours of its experimental use in the laboratory.

In addition, other hardware consisting in a stirrer plate and a UV photochemical setup were developed that enabled automated studies of multiphasic photochemistry for the first time.

Overall, a new piece of equipment was designed and built that enabled automated multiphasic flow chemistry, as well as outperforming the previous reactor’s temperature control both in rapidness and precision. In the next chapter, the use of this equipment is demonstrated for the automated experimentation of a multiphasic reaction.

Chapter 3
fReactor derived data and optimisation

3.1 Introduction

This chapter focuses on the analysis of the data obtained by the use of the fReactors using the equipment developed in Chapter 2. Two different studies were carried out: (i) a Claisen-Schmidt reaction where a sequence of experiments under different reaction conditions was carried out automatically using the integrated temperature-controlled fReactors in the automated setup at iPRD. The experimental data was supplemented with published data from a nucleophilic substitution, allowing metamodeling to be used to exemplify the different algorithms and cross-validation procedures for this type of study. (ii) a shorter study is described where the fReactors were used for a synchrotron-based X-ray diffraction study showing the crystallisation of material within flowing CSTRs.

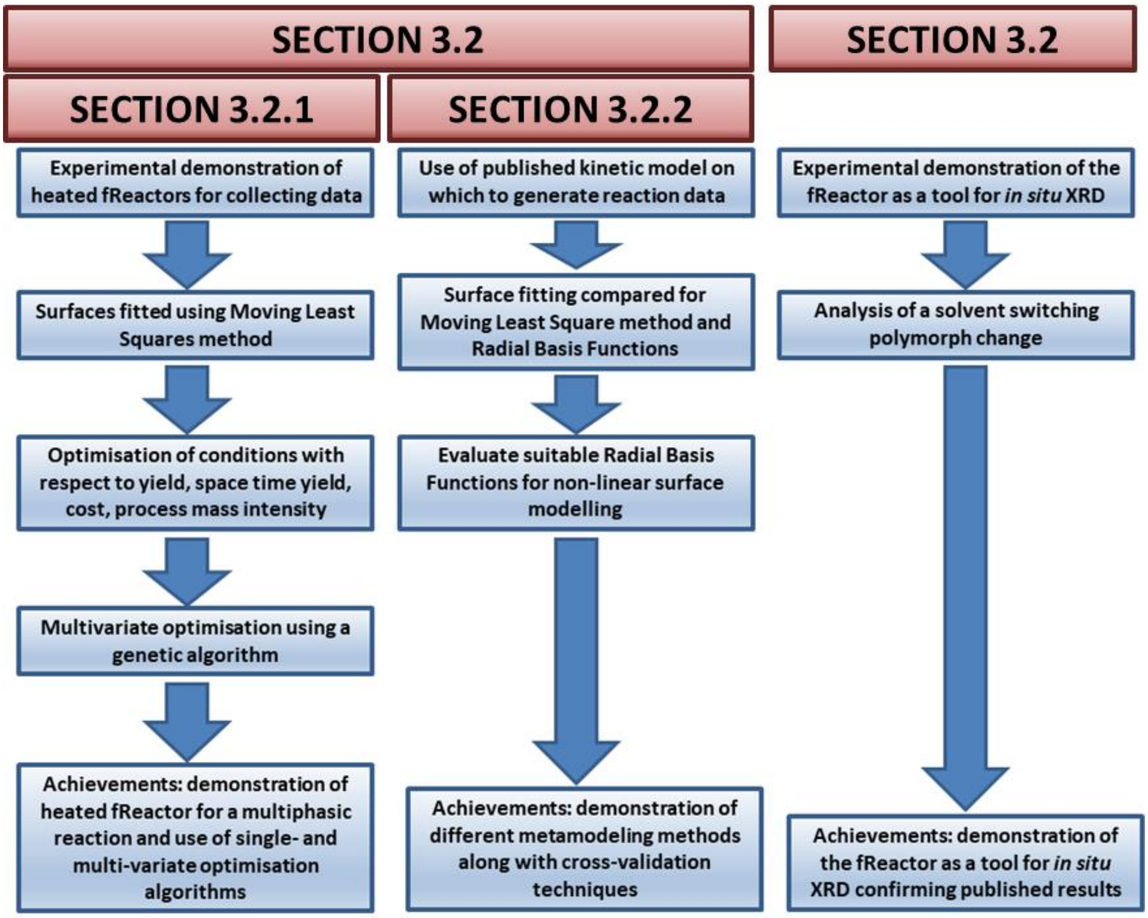

Figure 53. Overview of Chapter 3.

3.2 Reaction Modelling

3.2.1 Claisen-Schmidt Condensation

BACKGROUND

The purpose of this study was to demonstrate the use of the temperature-controlled fReactors integrated into the automated system developed within the iPRD for optimising reaction conditions. A Claisen-Schmidt condensation, a type of ‘aldol’ (aldehyde + alcohol) condensation,[240] was chosen to illustrate how different modelling strategies apart from those commonly used in this field (which, as discussed in the review of the literature of this thesis mainly rely on the least-squares fitting of polynomial models) could be implemented for a multiphasic flow reaction. This specific reaction was chosen as it was previously carried out -although in a different configuration- at iPRD,[76] and because of its well-known reaction mechanisms that make it commonly used as a teaching model reaction in organic chemistry.[241]

The reaction, discovered independently by Claisen[242] and Schmidt[243] in the early 1880’s, consists of the reaction a ketone with an aldehyde in the presence of a base. In the reaction chosen, whose intermediates are shown in Figure 54 for convenience, acetone (**1**) reacts with benzaldehyde (**2**) in the presence of NaOH to produce the target product benzalacetone (**3**) and the undesired by-product dibenzalacetone (**4**). Acetone has two α-hydrogens (hydrogens united to the first carbon apart from the functional group, which in this case is the carbonyl group) that can be deprotonated in the presence of a base, generating a molecule of water from the acetone’s proton and an hydroxide ion from the base. After losing the proton (hydrogen cation), the acetone becomes a nucleophilic enolate anion as shown in Figure 54 after the first step.

The carbonyl group in the benzaldehyde molecule is more electrophilic than that of the acetone, and in consequence reacts with the enolate in the second step of the reaction forming the intermediate alkoxide, which in turn is protonated in the third step to produce the β-hydroxyketone that contains both functional groups hydroxide and carbonyl. A base-catalysed dehydration follows in the next step to give the corresponding hydroxyenolate, which loses the hydroxide group forming a pi bond to finally generate the desired product benzalacetone (**3**).

If the product undergoes the described sequence in full again, dibenzalacetone (**4**) is formed, which under the scope of this study will be considered a by-product or impurity.

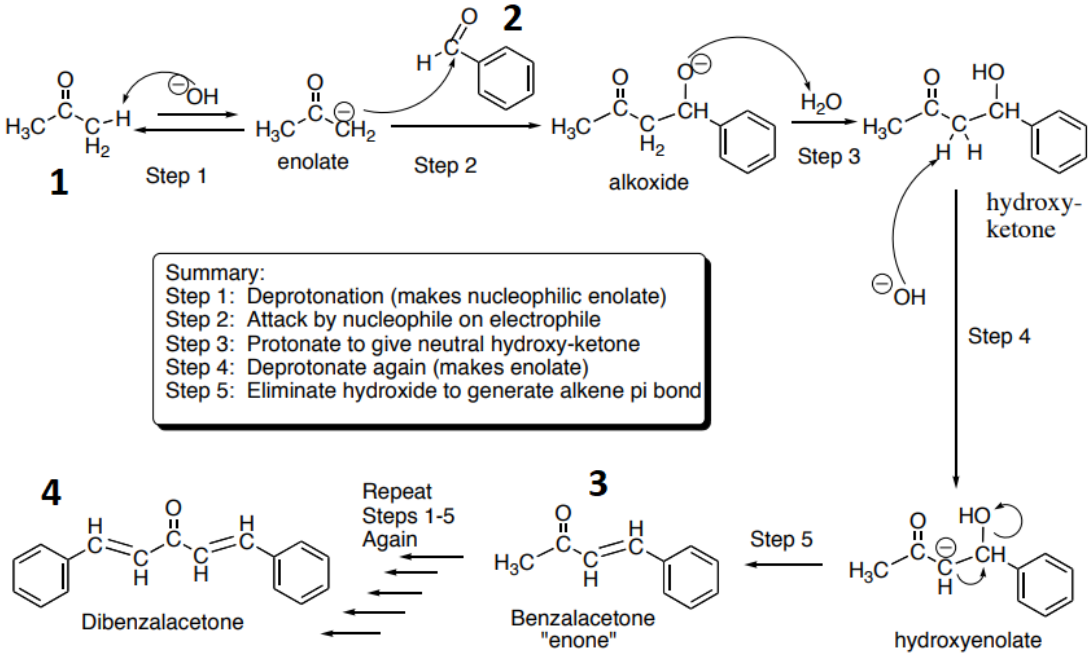

Summary:
Step 1: Deprotonation (makes nucleophilic enolate)
Step 2: Attack by nucleophile on electrophile
Step 3: Protonate to give neutral hydroxy-ketone
Step 4: Deprotonate again (makes enolate)
Step 5: Eliminate hydroxide to generate alkene pi bond

Figure 54. Claisen-Schmidt condensation mechanisms. Adapted from ref.[244]

In parallel with these transformations, a side reaction may occur that solely involves the acetone to form mesityl oxide and its corresponding polymer. In the previous experiments carried out at iPRD, Jeraal *et al.*[76] constrained their study of the reaction to temperatures under 80 °C in order to prevent the formation of these undesired products that eventually led to clogging in the Polar Bear Synthesiser tubular reactor. In addition, the formation of dibenzalacetone also clogged the reactor, a fact that clearly highlights the need for a different kind of reactor to carry out the reaction under a broader range of conditions. In comparison, the integration of the temperature controlled fReactors in the automated rig at iPRD enabled both:

- Assessment of the reaction at temperatures higher than 80 °C, avoiding clogging.
- Processing a multiphasic flow, including an aqueous NaOH solution and an organic phase.

EXPERIMENTAL SETUP

Figure 55 shows the elements of the automated system at iPRD. Its main components are a rack of HPLC pumps, a cascade of fReactors, a sample loop that injects a small volume of the product stream into the analytical device, and an HPLC instrument. The latter is connected to a desktop computer that stores the chromatograms and is able to modify the flow rates of the pumps, the temperature of the fReactors and to activate the sampling loop. Three Jasco PU-

980 HPLC pumps are used to dose the reagents into the fReactor cascade. A temperature-controlled cascade of three fReactors (total volume 6 mL) was used for the experiment, whose construction and temperature control implementation were covered in detail in the previous chapter of this thesis. After leaving the fReactor cascade, the fluid goes through a 75 psi back pressure regulator (BPR) into a laboratory scale Zaiput SEP-10 liquid-liquid separator,[245] with the aqueous phase being transferred to waste. The organic phase is deviated to a Vici Valco sample loop that enables injection into HPLC analysis (Agilent 1100 series) at the times specified by the computer, which registers the spectra of the product under different reaction conditions. All the connections present in the fReactor, the BPR and the Zaiput separator are standard ¼-28 UNF connections, making the components easily interchangeable and the setup easy to arrange in a plug-and-play fashion. 1/16'' tubing was used throughout the system, except between fReactors, for which 1/8'' tubing was preferred.

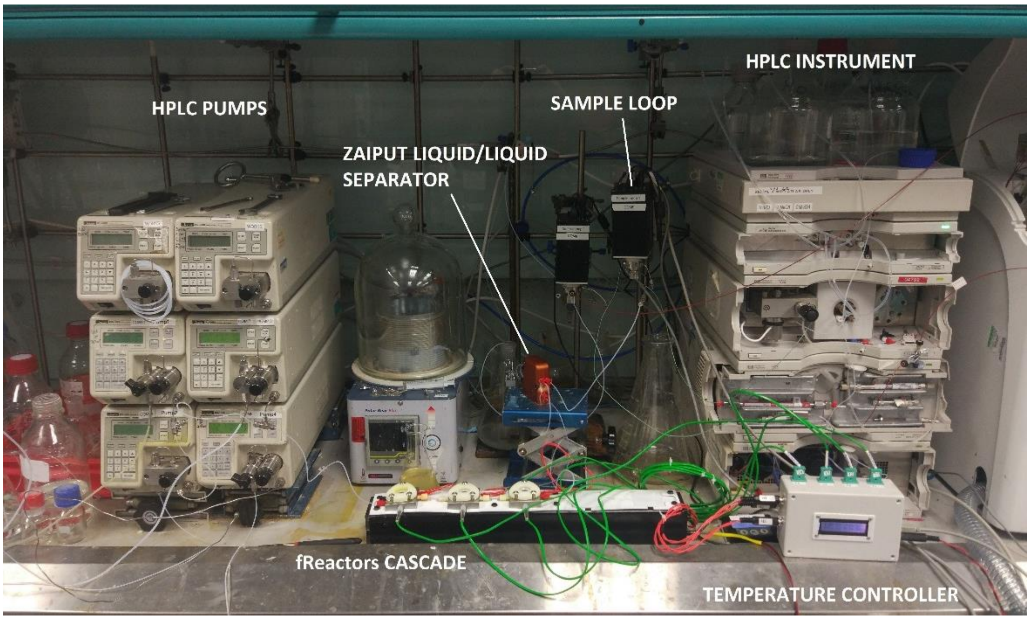

Figure 55. Picture of the automated rig with its elements marked.

The reaction study was proposed contemplating sodium hydroxide equivalents (eq_{NaOH}), mean residence time (τ) and temperature (T) as the variables of the system. As commented previously, a liquid-liquid multiphasic system was tried, including an organic and an aqueous phase. Two pumps were used to dose previously made solutions of the organic phase: the first one contained a solution of benzaldehyde (>98.0%, Fluorochem) in toluene (HPLC grade, VWR), as well as the internal standard anisole (>99.8%, Sigma-Aldrich) in order to provide accurate concentration calculation from the peak areas of the HPLC chromatogram; the second one consisting solely of acetone (laboratory grade, VWR). These two streams met in a Swagelok tee before reaching the first fReactor of the cascade. The third pump provided an aqueous solution of NaOH that was supplied directly in the fReactor, where it meets the organic phase and the reaction takes place.

However, the space design for the proposed experiment is subjected to the concentrations of the bulk solutions and the limitations of the system, specifically to the minimum flow rate that the Jasco pumps are available to provide with confidence (0.05 mL/min). In order to precisely calculate a feasible range of conditions under these circumstances, a system of equations relating the variables explored and the flow rates of the system was developed.

Given the mean residence time, it is straightforward to calculate the total flow rate for the experiment using equation (42). In this experiment, the organic and aqueous flow rates were kept the same at all times, allowing easy calculation of the aqueous flow rate of the sodium hydroxide solution according to equation (43). Using this value, the bulk concentrations of the sodium hydroxide (C_{NaOH}) and benzaldehyde (C_{BA}) solutions and the molar equivalents of sodium hydroxide relative to benzaldehyde (eq_{NaOH}) it is possible to calculate the flow rate of the latter's solution, as shown in equation (44). Calculation of the flow rate of acetone ($\dot{V}_{AC}$) is trivial when considering the equality of flow rates for the organic and aqueous phase condition imposed upon the experiment (equation (45)).

$$\dot{V}_{TOT} = {}^{V_R}\!/_{\tau} \tag{42}$$

$$\dot{V}_{NaOH} = \dot{V}_{TOT}/2 \tag{43}$$

$$\dot{V}_{BA} = \frac{\dot{V}_{TOT}}{2}\,\frac{C_{NaOH}}{C_{BA}\;eq_{NaOH}} \tag{44}$$

$$\dot{V}_{AC} = \frac{\dot{V}_{TOT}}{2} - \dot{V}_{BA} \tag{45}$$

Imposing the conditions that all the flow rates provided by the pumps ($\dot{V}_{NaOH}$, $\dot{V}_{BA}$ and $\dot{V}_{AC}$), must be equal or greater than 0.05 mL/min, the following system of inequalities is obtained:

$$\tau < {}^{V_R}\!/_{2 * 0.05} \tag{46}$$

$$\tau < \frac{V_R}{2 * 0.05}\,\frac{C_{NaOH}}{C_{BA}\;eq_{NaOH}} \tag{47}$$

$$\tau < \frac{V_R}{2 * 0.05}\left(1 - \frac{C_{NaOH}}{C_{BA}\ eq_{NaOH}}\right) \qquad (\ 48\ )$$

where C_{NaOH}/C_{BA} is the concentration ratio between the aqueous NaOH and the organic benzaldehyde bulk solutions.

Figure 56 shows the restriction on reaction conditions caused by imposing the minimum flow rate condition of the pumps for different relations of concentrations between the sodium hydroxide and benzaldehyde bulk solutions. In every case the intersection of the areas under the curves represents the set of conditions achievable by the system, meaning that all possible experimental conditions are contained in this area. This feasible operating space is shaded for one of the cases shown in Figure 56, under the blue lines for a bulk concentration ratio of 0.1. Ascending lines represent the condition in acetone flow rate (equation (48)) and descending lines on benzaldehyde flow rate (equation (47)).

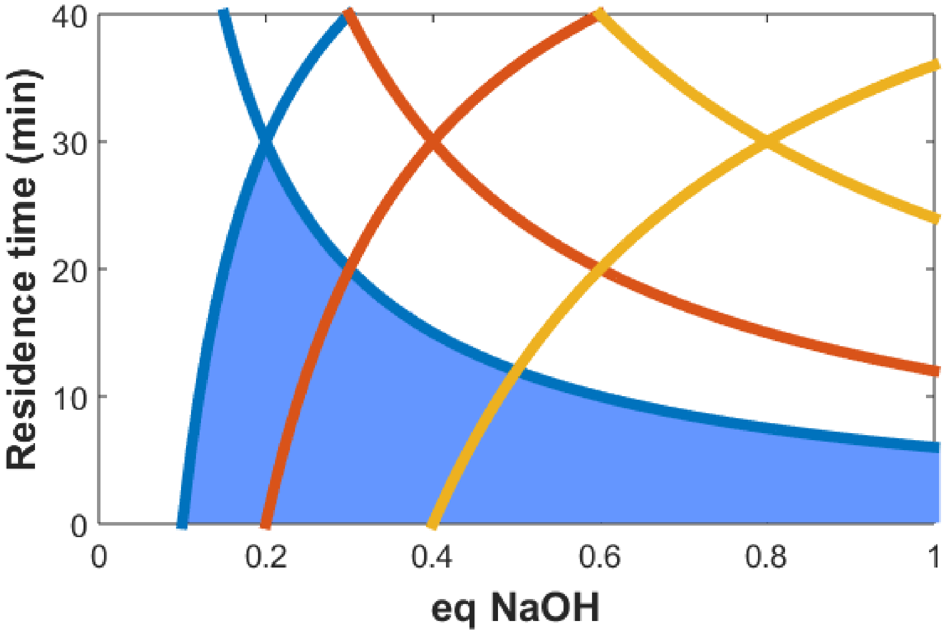

Figure 56. Plot of the different feasible design spaces for according to the restriction of a minimum flow rate of 0.05 mL/min. Blue, red and yellow lines represent the restriction for bulk concentration ratios of 0.1, 0.2 and 0.4, respectively. Residence time in minutes.

From this operability plot, a bulk ratio of 0.2 was selected, as this allowed a greater range of sodium hydroxide equivalents for residence times up to 10 minutes. Accordingly, 0.5 and 0.1 M solutions of benzaldehyde and sodium hydroxide, respectively, were made. The former was prepared by mixing 51.02 mL of benzaldehyde in 948.98 mL of toluene, including 7.61 mL of anisole as the internal standard, yielding 1 L of solution. The latter was produced by dissolving 4 g of NaOH in 1L of deionised water under stirring at ambient conditions.

HPLC mobile phases were A H_2O (18.2 MΩ), B MeCN, both buffered with 0.1% TFA. The method used was 10% to 90% B 3.5 min, 90% to 10% B 0.5 min, 10% B 1 min, flow rate 1.75 mL/min, column temperature 20 °C. Detector wavelength = 210 nm. Retention times (min): benzaldehyde = 1.91; benzylideneacetone = 2.33; dibenzylideneacetone = 3.36.

RESPONSE SURFACE MODELLING OF EXPERIMENTAL DATA

A response surface modelling methodology was selected, for which the objective is to build a model of the objective function using scattered data under different reaction conditions, defining a design of experiments or DoE. In consequence, a sampling plan had to be decided first. In this case, and due to the time-consuming nature of the experiment, a face-centered central composite design was selected (Appendix E). This is a sampling plan that only makes use of 15 data points in a three-dimensional problem: 5 points for each one of the three levels. The bottom and upper levels include points in the vertices of the cube, while the middle level has them in the middle of the edges. Every level includes a central point. In addition to these, 8 additional points forming a nested DoE were included which were located in the mid-point of the lines between the central point and the corners of the space design.

The conditions of the experiments were then automatically calculated and sorted in order of ascending temperature to reduce waiting times between them. Implementation of equations (43) to (45) in Matlab allowed calculation of the flow rates of the different pumps for every set of conditions in the DoE, which were fed to the software available at iPRD, using a GUI developed alongside the equipment.

Four different metrics were calculated from the spectra of the organic product. The first one, the yield of benzalacetone, is the percentage of said substance in the organic stream of the product, calculated from the concentrations of the different species in the organic product stream (dibenzalacetone, benzalacetone and the starting material benzaldehyde). This is a metric widely explored in the literature for different reactions in single-objective optimisations, in order to maximise the yield of a reaction, the purity of the product obtained.[7, 77] In the current case and in addition to the yield of the target product, different metrics were modelled so the trade-offs between them could be explored.

The cost per kilogram of product as defined in equation (49) was included in the study, as obviously this economical variable is of great relevance in industrial processes. Isolation of this product might need further steps that are not accounted for in the metric, which considers exclusively the cost of the reagents and not of the rest of the costs involved in the process, such as manpower and energy consumption, among many others. The numerator includes the summed cost of the stream of each pump, while the denominator accounts for the mass flow rate of the product, with its explicit formula in equation (50).

$$Cost = \frac{\sum_1^3 Cost_i \cdot \dot{V}_i}{\dot{m}_{product}} \left[{}^{£}/_{kg}\right] \tag{49}$$

$$\dot{m}_{product} = \dot{V}_{org} \cdot C_{prod} \cdot M_{prod} \tag{50}$$

The few multi-objective optimisations performed in flow chemistry to date mainly address the use of two objectives; for this study two further metrics were selected to exemplify the different trade-offs achievable with this approach. Hence, the space-time yield (STY), a measure of the productivity, was also calculated. This metric has been previously used in single-objective optimisations, although included in a custom equation in an attempt to achieve a balance between higher yields and productivity,[13] a common technique used to combine different objectives in single-objective optimisation. The STY, as defined by Haber and Greenwood,[246] is the quantity of product per unit volume per unit time; consequently, maximising the STY, aims to find production conditions that maximise the quantity of material per time but also concentration in the stream.

Finally, an objective to minimise the environmental impact of the process was included in order to balance it with the other factors. The process mass intensity (PMI) is the ratio between the mass flow rate of the process and the mass flow rate of the product obtained. This mass flow rate takes into account all regents and solvents, including both the aqueous and organic phases. The PMI is relatively simple to calculate for most processes and has been widely used in the pharmaceutical industry over the last 10 years.[247] In fact, it has been suggested by the sustainability accounting standards board (SASB) as the preferred metric for the process material efficiency that should be disclosed to investors.[248] The reader is referred to Jiménez González *et al.*[249] for further reasons why this metric has been selected by the pharmaceutical industry as the indicator of the overall greenness of a process.

$$PMI = \frac{\dot{m}_{process}}{\dot{m}_{product}}$$

$$\dot{m}_{process} = \sum_{1}^{3} \rho_i \cdot \dot{V}_i$$

To show the results of the experiments, the Moving Least Squares Method (MLSM) was used to create response surfaces which fit the data. The MLSM is a generalisation of the least squares method where weights are functions of the Euclidean distance between the DoE points and the point at which the surrogate model is being created.[250] A Gaussian weight decay function is used to determine the weighting of points in the regression analysis at each design point $\underline{x}$ [251]:

$$w_i(\underline{x}) = \exp(-\theta \left\|\underline{x} - \underline{x}_i\right\|^2) \qquad (51)$$

where $\left\|\underline{x} - \underline{x}_i\right\|^2$ is the Euclidean distance between the design point $\underline{x}$ at which the surrogate is being calculated, the ith DoE point $\underline{x}_i$ and θ is a *closeness of fit* parameter which provides a means of tuning the surrogate model. In the current study, the design variables are normalised into the unit cube to avoid scaling issues, and a second-order polynomial base with all terms is used for MLSM. The value of θ is optimised using the Leave One Out cross-validation (LOOCV) [251] and Monte Carlo cross-validation (MCCV) [252] techniques. LOOCV removes one point from the DoE data set and uses the rest of the DoE points to construct the surrogate model and calculate the error at the removed point. This procedure is repeated for every point in the DoE, and the root mean squared error (RMSE) computed. In the present study, the value of θ is determined by using a golden search algorithm to minimise the RMSE. MCCV [253] is similar to LOOCV but instead removes a random group of *k* points iteratively to be used as a validation data set while the remaining DoE points are used to create the surrogate model. In this study, the number of iterations N^2 suggested by Zhang[254] is used, where N is the number of data points. This produces results close to carrying out cross-validation over all unique possible training sets. After the closeness of fit has been determined in this fashion, this value is used to provide the final metamodel by using all datapoints in the DoE.

The cross-validation procedures were carried out for all four metrics studied, including LOO and MCCV with between 2 and 7 points (around 30% of the data) being left out. Among the seven models produced for each metric, the ones with an appropriate trade-off between minimisation of the cross-validation error and also of the ratio between the errors are presented in the following table for all metrics. The table includes the number of points left out during cross-validation, the closeness of fit θ, the cross-validation RMSE that was minimised during the process, the training error -defined here as the RMSE of the resulting model over the whole original dataset- and the ratio between both errors.

Table 8. Models selected by cross-validation.

	$k_{\text{left out}}$	**θ**	**CV RMSE**	**Training RMSE**	**CV/TRAIN**
YIELD	4	0.26109	29.4688	15.350	1.9198
STY	5	0.98137	3.62724	1.4255	2.5445
1/COST	4	0.62081	0.00941	0.0044	2.1517
1/PMI	4	0.65762	0.02366	0.0110	2.1600

Due to the definition of the different objective functions, and to the presence in the responses of points with null conversion, the cost and PMI computations would lead to a value of infinity in the some locations of the DoE, corrupting both test RSME and training error calculations. In these cases, models for the inverse values of the metrics were fitted, avoiding this problem. Isosurface representations are presented in Figure 57, aiding visualisation of different features of the responses that can be commented upon.

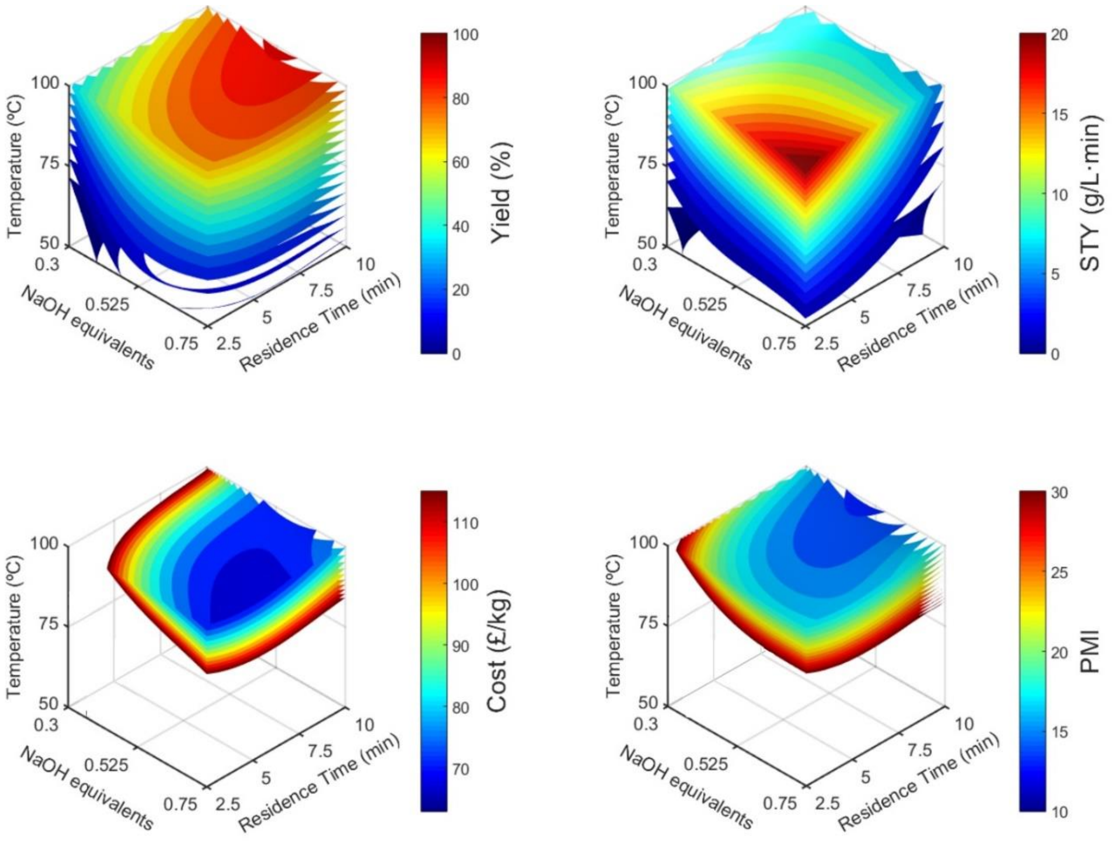

Figure 57. Representation of the objective functions using isosurfaces.

The yield of benzalacetone is higher at high temperatures, also increasing with residence time, while for every residence time an optimal NaOH equivalents exists that leads to maximum conversion. From this point on, increasing the equivalents of NaOH decreases the conversion. This may be related to the formation of the undesired product dibenzalacetone from benzalacetone in presence of high NaOH equivalents.

In the case of the space-time yield of the target product, a totally different trend can be identified. This metric is greatly improved at low residence times and high NaOH equivalents, being only similar to the yield in terms of temperature aiding maximisation. High flow rates (low residence times) increase the throughput, indicating that this increment in throughput outweighs the loss in concentration.

In the case of the cost, an interesting feature is that its minimisation is aided by high values of NaOH equivalents; this helps minimisation of the consumption of benzaldehyde, which is the most expensive reagent employed in the reaction. The process mass intensity presents similarities with the conversion, as maximising the product presence leads to low mass intensities.

OPTIMISATION

The results of single-objective optimisations aiming to maximise yield and STY and minimise the cost and PMI can be found in Table 9. Using particle swarm optimisation [255] on each response surface with a swarm size of 20 and setting the maximum number of iterations to 1000 and the error goal to 10^{-10}, the optimal set of variables was calculated.

Table 9. Single objective optimisation results for the four process metrics evaluated.

	NaOH equivalents	**Residence Time (min)**	**Temperature (ºC)**	**value**
YIELD	0.54	10	100	91.13 %
STY	0.75	2.5	100	21.74 g/L·min
COST	0.75	4.75	100	63.50 £/kg
PMI	0.49	10	100	12.67

Two two-objective optimisations were carried out to demonstrate the trade-offs between different aims. An implementation of the NSGA-II algorithm (a commonly used multi-objective genetic algorithm) [189] included in Matlab's optimisation toolbox (gamultiobj) was used on the response surfaces generated previously in this section with the appropriate settings that ensured convergence. A population of 250 individuals was used with a tolerance of 10^{-10} for a maximum limit of 300 generations. The selection of the individuals was carried out in a tournament of size 2, for which the individuals are compared pairwise with the best ones selected as parents. A crossover function was chosen that creates children from a weighted average of the parents. The share of population created by the crossover function in the next generation was set to a fraction of 0.8. When migration occurs the worst individuals of a subpopulation are replaced by the best of another. This fraction was set to 0.2 to happen at 20 generations intervals. Finally, the Pareto front population fraction was set to 0.35.

First, simultaneous maximisation of the yield and minimisation of the cost was demonstrated, generating the Pareto front shown in Figure 58. It is easy to see that the Pareto front obtained is almost linear except for the region in which the values of the yield are over 89%. In the linear region, spanning from 79.5 to 89%, increasing the yield of benzalacetone by one percent raises the cost 0.49 pounds per kilogram of product. For higher yields between 89 and 90.8%, yield can be enhanced at very little increase of the cost. Additionally, between 90.8 and 91.13% (corresponding to its single-objective optimum as can be seen in Table 9) the cost slope increases again to produce a higher purity product. This discontinuity in the slope of the front is due to the algorithm reaching the constraint of the problem for the residence time, being forced to simply reduce the NaOH equivalents in order to increase the yield.

In the case of the trade-off when trying to minimise the cost and maximise the STY at the same time, a smooth curve is obtained as the Pareto front. Paying attention to its curvature, the cost to maximise the STY (slope of the curve) raises as the STY is increased. Between 17 and 18 the increase is very cheap in terms of the cost of the product. The increase of STY between 17 and 18 raises the cost per kilogram only 20 pence, while the increase between 20 and 21 supposes a raise of the cost of 1.57 pounds per kilogram of product, almost eight times more. An STY around 20 g/L·min seems appropriate in terms of trade-off between these variables, since to obtain better STY values the cost increases very rapidly from this compromise solution.

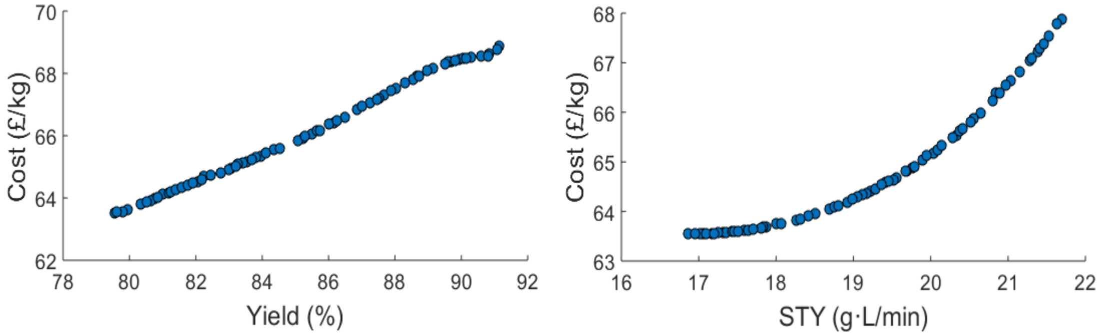

Figure 58. Pareto front for the optimisation of yield and cost (left) and of space-time yield and cost (right).

Next, a three objective optimisation was performed, aiming maximisation of the yield and minimisation of both the cost and the PMI. In this case, the options for the *gamultiobj* algorithm were kept the same as those used in the two-objective optimisations except for the population size, that was doubled from 250 to 500 individuals to provide extra resolution in the case of a three dimensional solution. The result produced for the Pareto front is shown in Figure 59, where it forms a 3D curve. More precisely, it is a narrow 3D surface, which indicates that the objectives for yield and PMI conflict only minimally, so the trade-offs achievable between them are remarkably small. The rounded feature in Figure 59 where the front becomes truly a curve and its shape changes forming a 'tail' is due to reaching the limits of the design space. This illustrates that at higher yields, increasing the cost leads to lower yields but decreases PMI, which in this region is very low. For any level of one of the three variables, the other two can be unequivocally determined in this region, what is produced by two of objectives (maximisation of the yield and minimisation of the cost) not conflicting in this region.

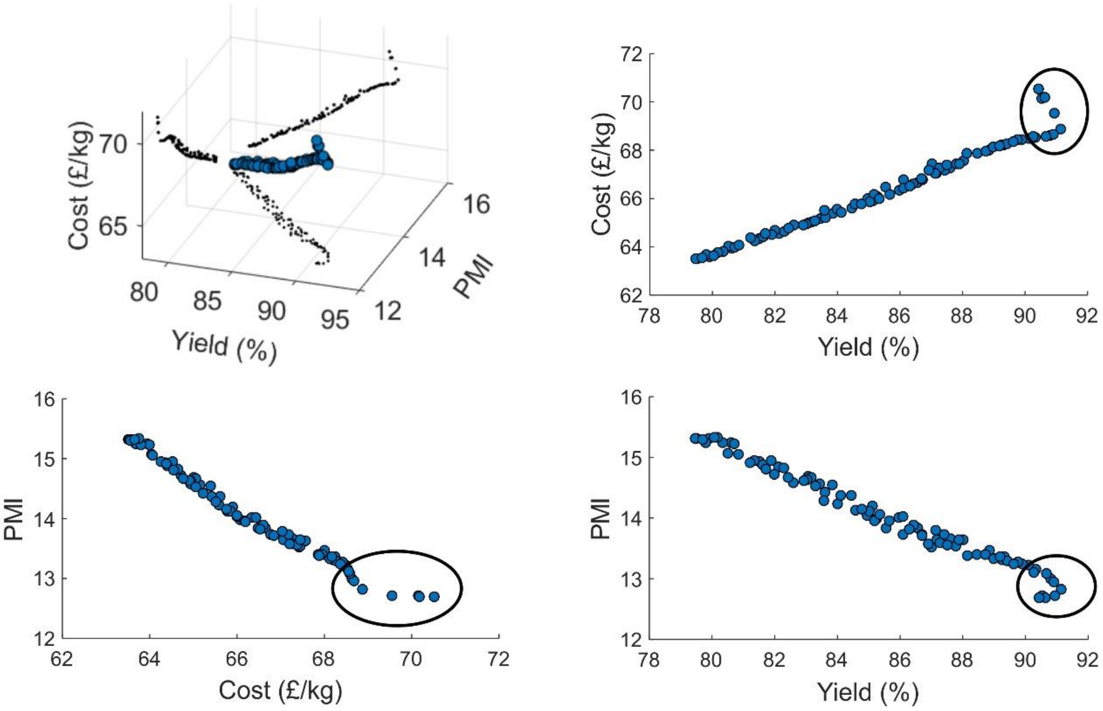

Figure 59. Pareto front obtained for the simultaneous optimisation of yield, process mass intensity and cost.

An extra three-objective optimisation was carried out in order to maximise the STY and minimise both cost and PMI. Figure 60 shows a representation of the Pareto front, where the surface has been represented via interpolation of the solution's scattered data, which is also projected in the three directions of the axes to aid visualisation. In this case, the limit of the feasible region is convex, meaning that it is not possible to move out of this surface without worsen the objective of at least one of the metrics. For example, considering different cost levels, and attending to Figure 60, it is easy to see that for each level, a 2D trade-off curve could be generated by the intersection of the cost level with the 3D Pareto front, this curve having the same meaning as the same 2D Pareto fronts in the previous two-objective studies, so for each level of cost the optimal set of solutions for the PMI and STY can be accessed, from which the user can balance them out to select the conditions of production of benzalacetone. The selection of a particular solution among the optimal set has been explored in the literature,[256] although a degree of subjectivity is present in the decision making that leads to a compromise solution.

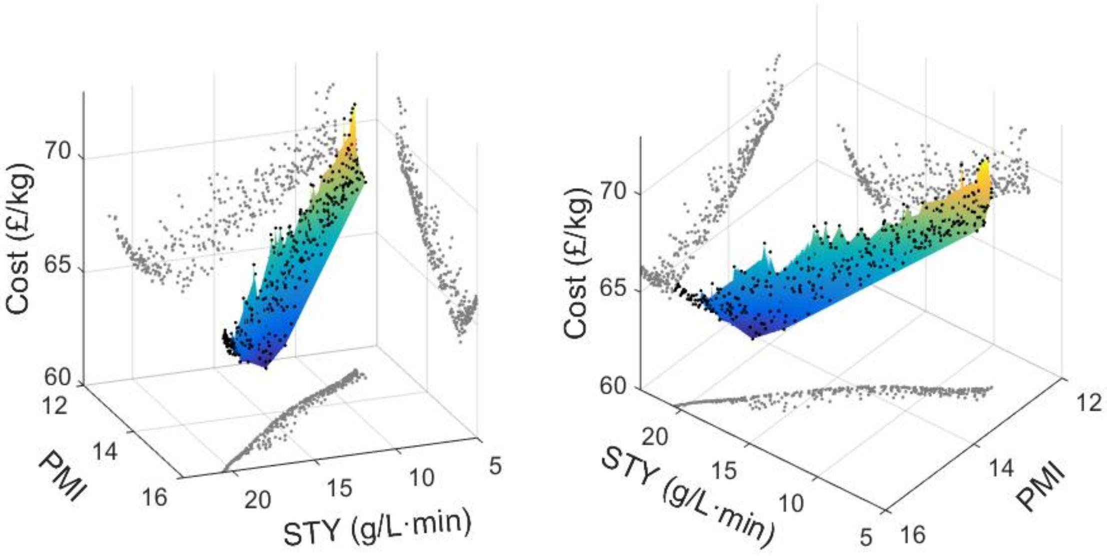

Figure 60. Three dimensional Pareto front solution for simultaneous Cost and PMI minimisation and STY maximisation.

3.2.2 In-silico metamodeling of an S_NAr reaction

To further demonstrate the potential uses of optimisation algorithms, an *in-silico* method has been used to simulate a flow reaction, giving the composition of products as a function of the operating conditions of the reaction. This allows a more detailed investigation into the performance of the algorithms to be carried out and provide recommendations for processing of the types of data shown in section 3.2.1.

BACKGROUND

Jensen et al.[35] demonstrated the automated continuous-flow determination of the reaction kinetics of a multistep nucleophilic substitution reaction using a microreactor. In their setup, 2,4-dichloropytimidine (1) reacted with morpholine (2) to yield the target 2-substituted aminopyrimidine (4) along with the 4-substituted (3) and 2,4-substituted (5) by-products. Figure 61 shows a scheme of the reactions involved, including the rate constants of each one. The kinetics were determined under the following assumptions: that all the reactions involved followed second-order, bimolecular reaction kinetics, and that their microreactor could be modelled as a plug-flow reactor. From this premise, the authors developed a system of 5 differential equations that could be used to calculate the yield of (4) with the residence time, initial concentrations of (1), molar equivalents of (2) and the temperature as inputs. The objective here is to sample this in-silico reaction from the equations provided in Jensen's work with a space-filling DoE, in order to use different metamodeling techniques and assess their performance, attending at just three of the variables: mean residence time, equivalents of morpholine and temperature, while the initial concentration of (1) will be fixed at 0.150 M throughout the study.

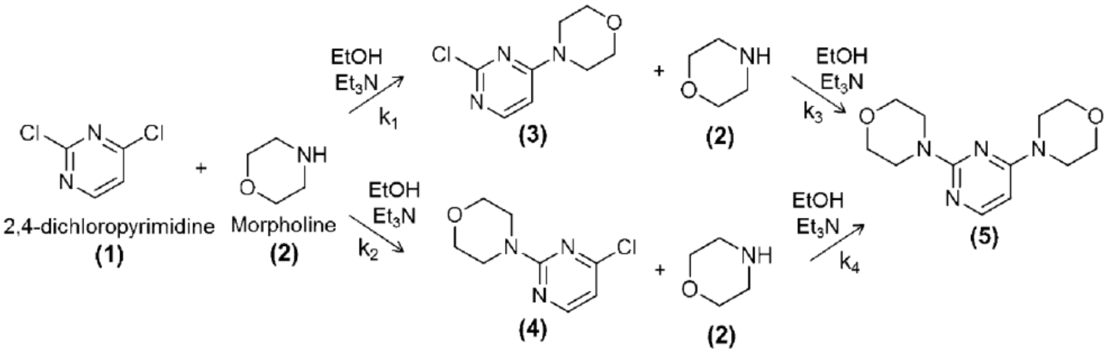

Figure 61. Scheme of the S_NAr reaction investigated by Reizman et al and reproduced from their work in ref.[35]

In the present case, and due to the rapid acquisition of data possible, two design-of- experiments were used. A Morris-Mitchells optimised Latin hypercube design [257] consisting of 36 samples was used (Appendix F), that would provide both training and test sets for model tuning by iterative splits as explained before for LOO and MCCV. Noise was added to these samples by selecting a random number from a standard normal distribution in order to replicate the behaviour of a real analysis instrument. This is achieved by using Matlab's built-in function randn with no arguments, which results in a distribution with a mean of zero and standard deviation of one. The DoE, kept as a holdout dataset, was a 5^3 full factorial design including 125 samples obtained from the differential equations without adding noise, so the predictions of the different metamodels computed from the first DoE could be directly compared to the exact yield results in these locations to give a measure of its accuracy on completely unseen data evenly scattered in the design space. The differential equations were solved using software developed in Matlab by Dr. Nicholas Holmes, that was based in the ode23tb Matlab ODE (ordinary differential equations) solver. This solver uses the trapezoidal rule in combination with a backward differentiation formula to integrate a system of differential equations between an initial and final time, being an implementation of TR-BDF2, an implicit Runge-Kutta formula.[258]

MLSM METAMODELLING

As in the Claisen-Schmidt condensation (section 3.2.1), the MLSM was used to generate metamodels for the reaction's yield of the target product. A second-degree polynomial base including all 10 terms was used for the fitting, combined with the LOO and MCCV cross-validation procedures for groups of different sizes of points left out. Zhang's suggested iteration number for MCCV ($N^2 = 1296$) was also used in this case.[254] The optimal θ was found by means of a golden search algorithm for every cross-validation procedure undertaken. The results are presented in Table 10, where CV RMSE is the RMSE over the points left out over all iterations of the cross-validation algorithm, Training is the RMSE of the resulting model over all 36 points and Test RMSE is the same over the 5^3 full factorial holdout dataset. It is the latter metric the one that offers a quantitative measure of the true performance of the algorithm, although it is important to take into account that under normal circumstances this could not be assessed due to the high number of experiments involved (that would make the process time-consuming), and in consequence the golden search is carried out towards the minimisation of the Test RMSE.

Table 10. MLSM cross-validation results.

	θ	**CV RMSE**	**Training RMSE**	**CV/Train**	**CV σ^2**	**TEST RMSE**
0 (LSM)	0		2.9144			7.9480
1 (LOO)	1.38561	4.14357	2.4718	1.6764	17.66	7.5945
3	1.11610	4.25198	2.5572	1.6628	6.65	7.6475
5	1.04339	4.28241	2.5802	1.6597	4.30	7.6631
7	0.92822	4.40861	2.6168	1.6848	3.39	7.6890
9	0.94915	4.48147	2.6101	1.7169	2.966	7.6842
11	0.86158	4.65410	2.6379	1.7643	3.083	7.7046
13	0.93556	4.78854	2.6144	1.8316	2.852	7.6873

In this instance, that could be considered of data scarcity, LOO produced the best results. It can be seen that, as the number of points being left out grows, the cross-validation error grows: this is a natural consequence of an algorithm that makes use of a fewer number of samples predicting responses over a bigger test set. Both the cross-validation error and the trade-off between it and the training error would generally lead the user to the selection of a particular model.[259] In this case, the minimum ratio between the errors that presents a reasonable test error is achieved for a number of samples left out of 5, which corresponds to a quantity data left out of almost 15%. In this regard, this is in well agreement with the previous study, in which similar quantities seemed appropriate for MCCV. However, attending to the holdout error, it can be appreciated that it grows as the quantity of data left out grows, being the LOO procedure the one yielding the overall best performing metamodel. This is indicative that under a situation of data scarcity, including as many points as possible in the training set during cross-validation dataset is highly beneficial; however any cross-validation procedure of those assessed here still outperforms the standard least-squares polynomial fitting. In the light of these results, it can be concluded that the integration of both MLSM and simple LOO cross-validation constitutes a promising technique in the field of chemical modelling, from which optimisation algorithms based in response surfaces could benefit.

RADIAL BASIS FUNCTIONS

The radial basis functions approach has been traditionally used for a large number of data samples (DoE points). However, over the last decade research has been conducted on the use of RBFs with limited number of samples.[260] Basically, an RBF model generates a prediction (approximation to a function) from a set of observations according to equation (52):

$$pred(\boldsymbol{x}) = \sum_{i=1}^{n} w_i \, \phi(\|\boldsymbol{x} - \boldsymbol{x_i}\|) \qquad (52)$$

where $\boldsymbol{x}$ is the position of the enquiry point, $\boldsymbol{x_i}$ an array of the sampling locations (also called RBFs centres) in the model, $r = \|\boldsymbol{x} - \boldsymbol{x_i}\|$ is the Euclidean norm, ϕ the basis function and w_i the weights associated with each specific centre. In essence, this is a linear combination of the basis functions, yet it has been demonstrated to be able to predict nonlinear responses with greater accuracy than other methods.[261] There is choice for the type of the basis function employed, and although in most cases this is an *a priori* choice, here four standard types were assessed (Gaussian, multi-quadric, inverse multi-quadric and polyharmonic spline), along with a compactly supported positive definite RBF first introduced by Wu[262] and suggested as one of the globally best-performing RBFs in the literature.[263] These are presented in equations (53) to (57), respectively.

$$\phi(r) = e^{-\beta^2 r^2} \tag{53}$$

$$\phi(r) = \sqrt{\beta^2 + r^2} \tag{54}$$

$$\phi(r) = {1}\Big/{\sqrt{\beta^2 + r^2}} \tag{55}$$

$$\phi(r) = r^{\beta} \ln \beta \tag{56}$$

$$\begin{gathered}\phi(r) = (1-t)^7(5 + 35t + 101t^2 + 147t^3 + 101t^4 + 35t^5 + 5t^6) \\ where\ t = r/\beta\end{gathered} \tag{57}$$

The weights are easy to calculate by applying the interpolating condition at every RBF centre, as presented in equation (58) and in its matrix form, where $\boldsymbol{\Phi}$ is the Gram matrix ($\boldsymbol{\Phi_{i,j}} = \phi\left(\left\|\boldsymbol{x}^{(j)} - \boldsymbol{x_i}\right\|\right)$) in equation (59):

$$\sum_{i=1}^{n} w_i\, \phi\left(\left\|\boldsymbol{x}^{(j)} - \boldsymbol{x_i}\right\|\right) = y^{(j)} \tag{58}$$

$$\boldsymbol{\Phi w} = \boldsymbol{y} \tag{59}$$

RBFs AND PARAMETER SELECTION

Most often, the user selects the RBF type before carrying out cross-validation in order to select the RBF's shape parameter β. In the study here, software was developed that enabled automatic selection of both the type of RBF and its β that minimised test error over cross-validation. The software proceeds with a golden search of β that produces the best (lowest) cross-validation RMSE using LOO cross-validation. This procedure takes place iteratively for all the five types or RBFs included in the study and returns the best performing metamodel among all of them. The results from this strategy can be accessed in Table 11, since the results for each RBF optimal parameters and errors are stored in a variable for further assessment.

Table 11. LOO optimal β selection.

RBF	Optimal β	CV RMSE	Train RMSE	Test RMSE
Gaussian	1.3740	5.5151	0	8.0792
Multi-quadric	0.4472	4.4235	0	8.0945
Inverse multi-quadric	0.6290	4.3613	0	7.8592
Polyharmonic spline	2.5352	4.6179	0	8.7240
Wu's	46.18034	4.87194	0	7.4744

It can be noticed by comparing the errors incurred over the holdout set (Test RMSE) that most of these metamodels are outperformed by MLSM, which suggests that under a data scarcity situation, RBFs may be less appropriate choice than MLSM. However, using the developed selection routine in Matlab yields a result that betters those obtained when using the MLSM. This corresponds to the use of the optimised compactly supported function and highlights the importance in the selection of an appropriate RBF, being in agreement with the findings by Fang and Horstemeyer,[263] that recommended the use of such functions over all the commonly used RBFs.

Carrying out a MCCV with groups of 3 samples left out yields the results that appear in Table 12. It was noticed during this process that the use of MCCV for radial basis functions is very computationally expensive, leading to much longer computing times than MLSM. The same search algorithm used for LOO was modified to allow the use of MCCV in the search of the best performing metamodel, minimising the test error.

Table 12. MCCV with 3 samples left out.

RBF	Optimal β	CV RMSE	Train RMSE	Test RMSE
Gaussian	1.3525	5.4956	0	8.1387
Multi-quadric	0.4892	4.5237	0	7.9875
Inverse multi-quadric	0.6600	4.4607	0	7.8551
Polyharmonic spline	2.5425	5.0456	0	8.5336
Wu's	46.18034	5.04437	0	7.4752

These do not prevent overfitting, since the condition imposed to find the weights of each centre forces this algorithm to interpolate. To avoid overfitting, a regularisation parameter λ is added as specified in equation (60) that accounts for deviation and allows a degree of disagreement in the centre of the RBFs. Developing the appropriate software to search for β and λ by means of a genetic algorithm, and minimising the test error, the results indicated that this is a highly sensitive and problematic search in the case of RBFs, since minimising the cross-validation error is possible but this does not guarantee a good performing model, as this could lead to higher values of the error over the holdout dataset. In light of this behaviour, colourmaps of the holdout error as a function of β and λ were developed from which decisions could be made in order to reduce the holdout error. The results presented in Table 13 are obtained in this fashion, where the CV RMSE is calculated by calling to the cross-validation function for the specific values picked from the graphs. The cross validation method here is constrained to LOO, since as discussed before MCCV when using RBFs is computationally intensive and time consuming.

$$\boldsymbol{w} = (\boldsymbol{\Phi} + \lambda\,\mathbf{I})^{-1}\,\boldsymbol{y} \qquad (60)$$

Table 13. Examples with an appropriate regularisation parameter.

RBF	Optimal β	Optimal λ	CV RMSE	Training RMSE	Test RMSE
Gaussian	1.3740	0.00025	4.6583	0.8469	7.3684
Multi-quadric	0.85	0	4.7028	0	7.6635
Inverse multi-quadric	0.925	0.0002	4.2735	1.0537	7.6414
Polyharmonic spline	2.5352	0	4.6179	0	8.7240
Wu’s	46.18034	0	4.87194	0	7.4744

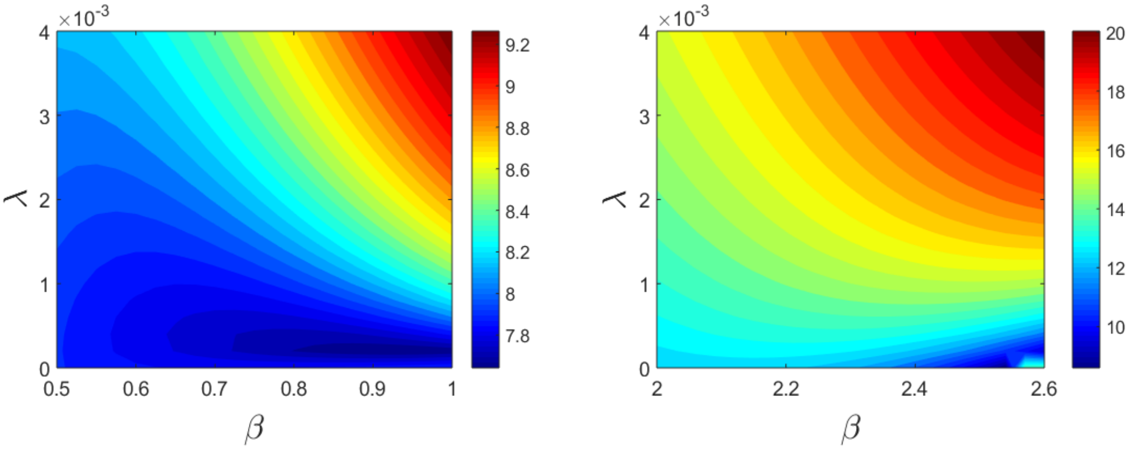

Figure 62. Colormaps for the holdout RMSE.

Obviously, the resulting metamodels benefit from the inclusion of the additional regularisation parameter and get generally better than their non-regularised counterparts, depending on the RBF type. However, both the polyharmonic spline and Wu's function appear to reach their optimal performance for null values of the regularisation parameter and do not benefit from the more complex search in this occasion. Additionally, these are results that would not be accessible during cross-validation in an expensive data environment and were used solely to assess the effect of the introduction of a regularisation parameter. More sophisticated searches, including multi-objective optimisation of the test error and the holdout error (or the ratio between them) could be developed and integrated in the algorithm, but this is out of the scope of this study.

AUGMENTED RADIAL BASIS FUNCTIONS

Even when RBFs have been proven to be a good choice for nonlinear responses, their downside is their unsuitability to predict linear responses. A commonly technique employed in order to overcome this disadvantage is the inclusion of a polynomial in addition to the RBFs, in what is called an augmented RBF:

$$pred(\boldsymbol{x}) = \sum_{i=1}^{n} w_i\, \phi(\|\boldsymbol{x} - \boldsymbol{x_i}\|) + \sum_{j=1}^{p} c_j g_j\,(\boldsymbol{x}) \tag{61}$$

where $g_j(\boldsymbol{x})$ is a linear polynomial function, c_j its coefficients and p is the number of terms of the polynomial. In this case a second-order polynomial with all 10 terms was used to augment the RBFs. The coefficients of the polynomial and weights of the RBFs can be simultaneously calculated imposing both the interpolation condition on the weights and coefficients, and the orthogonality condition solely on the coefficients, as shown in equations (62) for the latter and (63) for both conditions implemented in matrix form.

$$\sum_{i=1}^{n} w_i \, g_j(\boldsymbol{x}_i) = 0 \qquad for\; j = 1, 2, \dots, p \tag{62}$$

$$\begin{pmatrix} \boldsymbol{\Phi} & \boldsymbol{G} \\ \boldsymbol{G}^T & \boldsymbol{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{c} \end{pmatrix} = \begin{pmatrix} \boldsymbol{f} \\ \boldsymbol{0} \end{pmatrix}$$
$$where\; G_{i,j} = g_j\,(\boldsymbol{x_i}) \tag{63}$$

This new scheme for RBFs was included in the previously detailed search method to find the shape parameter that minimises the test error for every augmented RBF in the study. Results are shown in Table 14, where it can be seen that the model achieving a lower test error also results the best one in global terms, minimising the holdout error. This may point that this search method is appropriate even when using a low number of samples. In this case only the augmented polyharmonic spline achieves a holdout error under 7. Strikingly, this was the worst performing RBF when using not-augmented RBFs (see Table 11), which highlights the tremendous impact that the augmented RBFs can have. In all cases, the global performance of the algorithm is enhanced, with the holdout errors reduced to variable extents from 0.74% for Wu's function to 20% for the polyharmonic spline.

Table 14. Results for augmented RBFs.

RBF	Optimal β	RP	Test RMSE	Train RMSE	HOLDOUT RMSE
Gaussian	3.7635	0	4.0190	0	7.3555
Multi-quadric	0.1877	0	3.9660	0	7.0899
Inverse multi-quadric	0.3097	0	3.9933	0	7.1786
Polyharmonic spline	2.2565	0	3.9211	0	6.9780
Wu's	40.10557	0	4.86335	0	7.4192

It is also a possibility to let the regularisation parameter vary on these augmented RBFs as well, by modifying the Gram matrix in the same way as in the case of the non-augmented RBFs.

$$\begin{pmatrix} \boldsymbol{\Phi} + \lambda\, \mathbf{I} & \boldsymbol{G} \\ \boldsymbol{G}^T & \boldsymbol{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{w} \\ \boldsymbol{c} \end{pmatrix} = \begin{pmatrix} \boldsymbol{f} \\ \boldsymbol{0} \end{pmatrix} \tag{64}$$

Using this method, and minimising the holdout dataset error, it is found that the polyharmonic spline is again the best performing RBF, but also that its error is only reduced in a 0.14%. In the light of these results, it is clear that in this case, the model almost imperceptibly benefitted from the inclusion of the regularisation parameter, which complicates excessively the search of a well-performing model. In consequence, the use of non-regularised augmented RBFs is advocated in this case.

Table 15. Results for regularised augmented RBFs.

RBF	Optimal b	RP	Test RMSE	Train RMSE	HOLDOUT RMSE
Gaussian	1.8	0.0035	3.9218	1.1074	7.136
Multi-quadric	0.36	0	4.0374	0	7.0439
Inverse multi-quadric	0.95	0.000175	3.8463	0.9157	7.0242
Polyharmonic spline	2.4	0.0001	3.8946	0.2324	6.9675
Wu's	40.10557	0	4.86335	0	7.4192

Figure 63 in the next page shows representations of (a) the real yield obtained, (b) along a least squares polynomial fit, (c) the best performing MLSM and (d) the augmented polyharmonic spline predictions of it.

Both polynomial models fail to capture the shape of the highly nonlinear response, while the metamodel that used augmented radial basis functions, which is able to capture this behaviour, especially in the design space where the yield is lower than 75% (orange isosurface in the graph). However, it becomes somewhat biased as can be seen from the 80% yield level where a bump can be spotted that does not correspond to the straight shape that appears in (a). Taking into account that the model was computed just from 36 samples, this result can be considered good, which highlights the benefits that state of the art metamodeling techniques in conjunction with appropriate cross-validation may bring to automated flow chemistry.

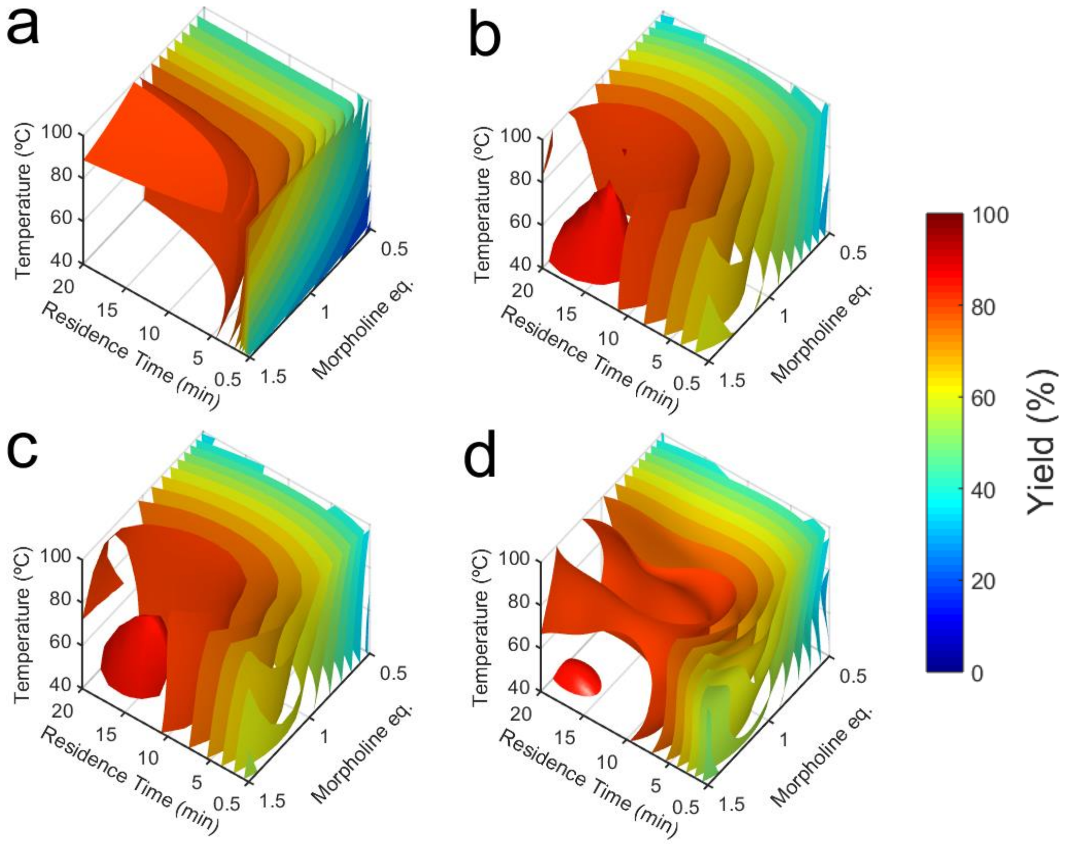

Figure 63. Isosurfaces representation of the yield. (a) Real reaction landscape. (b) Least-squares polynomial fitting. (c) MLSM with optimal β found by LOO. (d) Augmented polyharmonic spline RBF.

The benefits of the use of RBFs are twofold in this field: on the one hand, and since chemical reactions can exhibit highly nonlinear behaviour as it was the case, they are an appropriate choice to develop metamodels. In addition to this, increasing the performance of a metamodeling technique could be of relevance for the self-optimisation algorithms that generally use polynomial least squares fitting: this could be translated in cutting operational times and reagent consumption. Since typically the kind of response that is going to be obtained is not known *a priori*, the use of the proposed search algorithm across different augmented RBFs results of great assistance, as it is an unassuming method that is able to provide good performance for both linear and nonlinear responses, displacing MLSM in this regard.

The use of kinetic models for evaluating experimental design is a promising approach that gives a better understanding of the trade-off between number of experiments and precision of model. This study was in part inspired by the experimental work described in section 3.2.1, and a further extension described in section 5 (conclusions of thesis) would be to repeat these experiments with RBF to develop the metamodels.

3.3 *In situ* XRD analysis of calcium sulphite formation using fReactors

During the course of the research, the opportunity to evaluate the fReactors at a synchrotron beamline (Diamond Light Source, UK) became available. The sole aim of this study was to demonstrate the suitability of the fReactors as a tool for fluid flow synchrotron experimentation, where its capabilities were exploited to carry out experiments under conditions that would otherwise lead to clogging when using the most common sample environment for this approach, the microreactor. In order to meet this aim, a case study regarding the crystallisation of calcium sulphite in different polymorphs was identified and previous experimental results confirmed by the use of the fReactors for in situ synchrotron XRD.

3.3.1 Background

Calcium sulphite ($CaSO_4$) crystallises in different forms according to its degree of hydration. It is found in nature mainly as the dihydrate form ($CaSO_4 \cdot 2H_2O$), commonly referred as gypsum, and also as an anhydrate (anhydrous $CaSO_4$), with both minerals being highly relevant in the mass production of plaster, fertiliser and cement.[264] However, and although a much less naturally available mineral (mainly found in dry lake beds and interlayed with gypsum in caves) the hemi-hydrate ($CaSO_4 \cdot 0.5H_2O$), known as bassanite,[265] is one of the most industrially produced inorganic materials worldwide. It has a key role in construction -where it is known as ‘plaster of Paris’- due to its suitability for use in plaster products.[266] The most wide-spread method to obtain bassanite is the calcination of dry gypsum at temperatures between 130 and 170 °C, causing the release of part of its hydration water.[267] Temperatures above 170 °C lead to the formation of the undesired anhydrite form. Production of bassanite at lower temperatures (75 °C) has been proven possible by Guan *et al.*[268] by using methanol-water solutions. When water is added to this product, a strong exothermic hydration reaction takes place that causes its return to the preferred dihydrate form (gypsum), a process that explains the hardening of plaster and cement.[269]

A recent study by Tritschler *et al.* showed that pure bassanite can be obtained at ambient conditions by quenching aqueous $CaSO_4$ solutions with ethanol.[270] Reducing water content to values under 33 wt% resulted in bassanite becoming the main precipitating polymorph. The authors reasoned that this striking effect on polymorph selection was related to the quick change in solvent polarity. This caused the gypsum present in aqueous solutions to form bassanite nanorods, constituting a green alternative to the production of the mineral when compared to the highly energy-consuming process of calcination of gypsum. In addition, a more uniform distribution of crystal sizes for the bassanite nanorods is achieved when the method is employed. Just as a demonstration of the potential of the fReactors for *in situ* X-ray analysis, a confirmation of these results was undertaken. To the author's knowledge this represents both the first demonstration of lab-scale CSTRs employed in synchrotron X-ray analysis and also the first *in situ* confirmation of Tritschler's results.

3.3.2 Experimental Procedure

Experimental time was granted at the UK's national synchrotron facility Diamond Light Source following an application of fellow PhD student, Mark Levenstein, that included other experiments connected with the study of crystallisation within droplets.[271] Beamline I22 (SAXS/WAXS) was specifically requested, as it allows simultaneous small and wide angle X-ray scattering. While WAXS helps study the crystalline structure of a material, using SAXS enables extracting information about bigger features such as the size and shape of macromolecules, although this capability was not exploited in this particular experiment. The standard beam size of the beamline is of 320 horizontal × 80 vertical μm. After the beam interacts with the sample, the diffraction pattern is registered by a Dectris Pilatus 2M detector.[272] A second detector is located meters away of the sample stage for collection of SAXS patterns, for which a vacuum tube avoids scattering of environmental air in the laboratory.

The fReactors were slightly modified for the experiment from the previously described version, substituting the plastic tubing between adjacent reactors by Kapton tubing (3.188 mm ID, 0.095 mm wall thickness) in order to enable XRD. Kapton was chosen as it presents good chemical resistance characteristics, in addition to the near-transparency to X-rays that makes it popular in sample preparation for X-ray analysis.[273] The cascade was fitted on top of a multi-position stirrer (IKA 3690500), which in turn was attached to the I22's hutch translational goniometer using standard optomechanical components purchased from ThorLabs, as shown in Figure 64. The use of such setup allows time-resolved analysis, as the effluent of each reactor can be analysed individually once a steady-state is reached. Precise positioning of the cascade was possible by using an automated process that is briefly described here. First, X-ray exposures in the X and Y directions are carried out, from which the computer generates transmission plots as a function of position. The minimum is identified and the setup moves to this new position. This process can be repeated iteratively to find the position of minimal transmission (centre of the filled tube) with a high degree of accuracy.

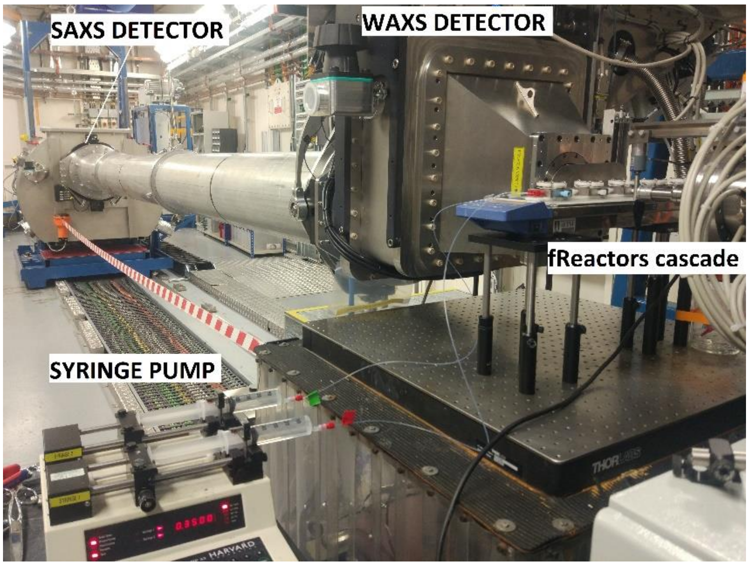

Figure 64. I22 hutch experimental setup. Both WAXS and SAXS detectors can be appreciated, as well as the dual syringe pump and the fReactor cascade.

150 mM solutions of Na_2SO_4 (anhydrous, >99%, Fischer) in deionised water and $CaCl_2$ (dihydrate, >99%, Sigma-Aldrich) in ethanol (>99.5%) were employed in the synthesis of $CaSO_4$, being fed into the fReactors by a Harvard Apparatus 33 Dual syringe pump at varying ratios. The study contemplated two scenarios. Initially, both solutions were mixed at a 1:1 ratio, which led to the rapid formation of gypsum, visible through the window of the fReactor as a cloudy white solution. After an X-ray scan was carried out, the ratio was changed to 2:1, varying the flow rates by at the same time rising the flow rate of ethanol to 2 mL/min and decreasing the aqueous solution flow rate to 1 mL/min. A second scan was taken to assess if bassanite was formed under these circumstances. Both scans included a total of 20 frames with an exposure time of 100 ms.

3.3.3 Postprocessing

Raw data were processed using a Matlab script developed by PhD fellow Mark Levenstein that performed an averaging of the frames for each scan, followed by a background subtraction. Once the process is finished, the intensity in the 2D plots was azimuthally integrated using DAWN, the XRD analysis software developed by Diamond.[274, 275] This takes into account all the beamline's available information like precise detector position which enables to account for its tilting and possible defects during integration.

First, averaging of the 20 frames for each scan was computed, followed by a background subtraction to distinguish between the diffraction pattern of the material of interest and the scattering of the rest of materials (the solvents and air in the path of the X-ray beam) using Matlab. Next, a mask, shown as a green area in Figure 65 (a) and (c), is applied that ignores problematic sections of the detectors encompassing broken zones (dead pixels) and the space between different panels of the detector. This makes possible to ignore intensity values at these positions, avoiding distortion of the azimuthal integration results. From these results, 1D plots were computed using Matlab in order to identify changes in the crystal morphology of the product at the outlet of the first reactor.

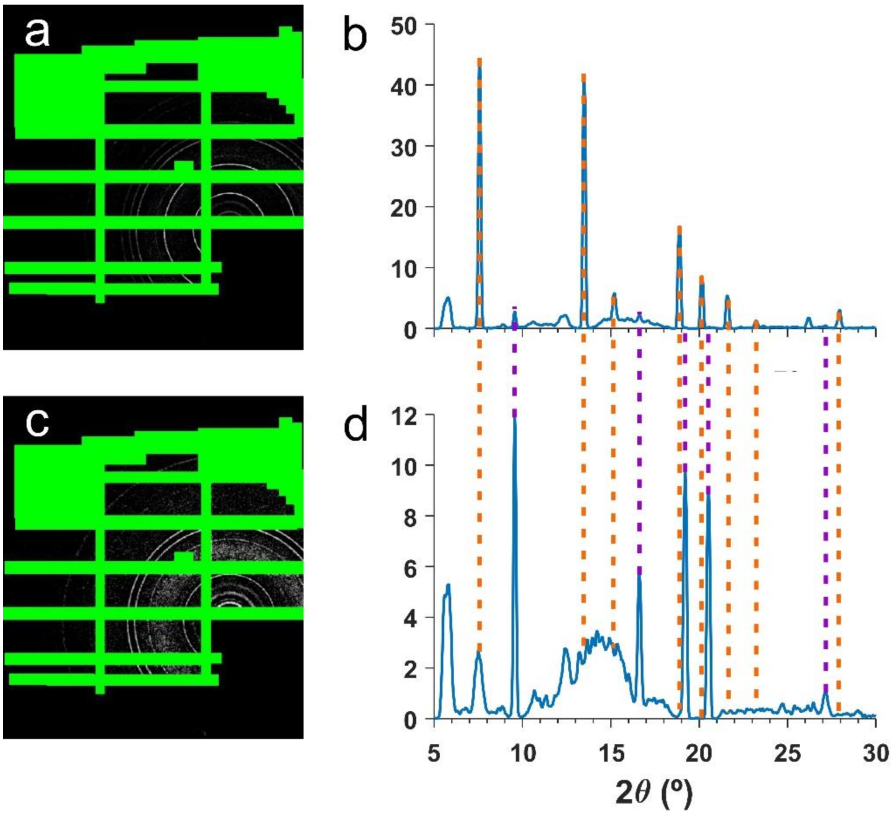

Figure 65. (a) 2D representation of the intensity in the detector. (b) Plot of the azimuthal integration.

Figure 65 (b) and (d) show the 1D plots for the obtained azimuthal integrations of the first and second experiments carried out, respectively. Analysis of the peaks was based in data accessed from the American Mineralogist Crystal Structure Database.[276] However, the data corresponded to X-rays analyses carried out with a wavelength of 1.54 Å, while the synchrotron's wavelength used in the experiment was of 1 Å. Use of Bragg's Law,[277] shown in equation (65), allows easy transformation of the peaks reported in the database into their corresponding angles at 1 Å. It is immediate to find the expression of the sine of the glancing angle for the same material (same order of reflection n and interplanar distance d) at two different wavelengths λ_i (equation (66)). From these, a relationship between the sines is found that can be used to finally calculate the glancing angle when a wavelength λ_2 is used, as shown in equation (68).

$$n\lambda = 2dsin(\theta) \tag{65}$$

$$sin(\theta_1) = \frac{n\lambda_1}{2d} \quad ; \quad sin(\theta_2) = \frac{n\lambda_2}{2d} \tag{66}$$

$$sin(\theta_2) = \frac{\lambda_2}{\lambda_1}\sin(\theta_1) \tag{67}$$

$$\theta_2 = arcsin\left(\frac{\lambda_2}{\lambda_1}\sin(\theta_1)\right) \tag{68}$$

After this transformation, X-ray data reported by Schofield *et al.* for gypsum [278] were compared with the experimentally obtained 1D plot (Figure 65 (b)). The latter exceptionally agreed with the former, presenting high intensity peaks for 2θ values of 7.59 and 13.46°, followed by a less intense one located at 15.22°. It also has the characteristic presence of three successive peaks descending in intensity at 18.88, 20.12 and 21.60° that are present in the results. Two smaller peaks at 23.36 and 27.81° are also identifiable in the pattern obtained for the experiment. According to these results, the presence of gypsum is accredited, as was expected for an experiment at a solvent ratio of 1:1, which equals a water weight percentage of 55.9 wt%.

Applying the same transformation for the second case, peaks are identified in the intensity graph that correspond to those reported by Ballirano *et al.* for bassanite.[279] Specifically, peaks at 9.55, 16.59, 19.17, 20.54 and 27.39° are easily identifiable and confirm the conversion of gypsum to bassanite. However, taking a closer look at Figure 65, small peaks in the characteristic positions of the bassanite can be identified in the first plot, as well as small peaks corresponding to gypsum in the second. This demonstrates the presence of the other polymorph in both solutions, although in a low extent, as was also reported by Tritschler. In summary, going from 55.9 to 38 wt% in water induced a rapid conversion to bassanite. The great deal of scattering present in the plots between 10 and 20° suggests that an amorphous phase had been also formed, most likely amorphous calcium sulphate (ACS), that seems to be present in higher quantities in the second plot.

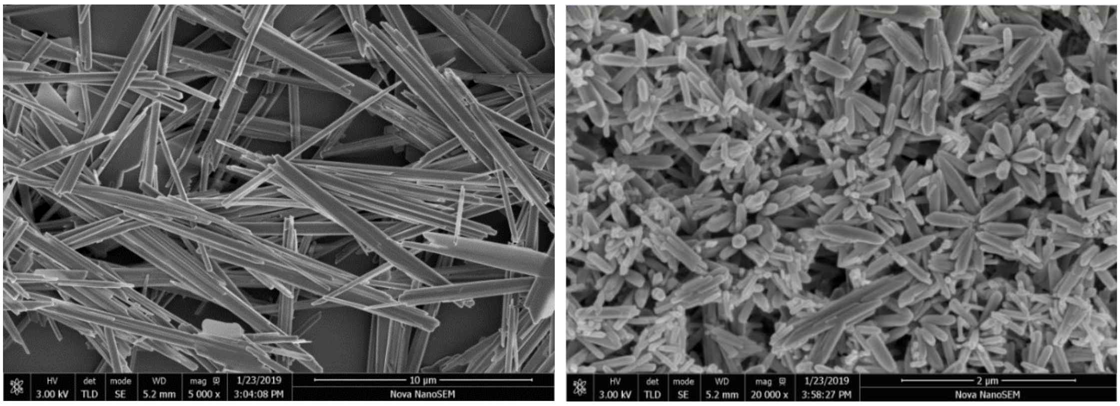

Figure 66. SEM analysis of samples collected before (left) and after (right) solvent ratio modification.

Further confirmation of these results was achieved by replicating the methodology described here and carrying an *ex situ* SEM analysis (Figure 66), which revealed the presence of mainly gypsum crystals for the first experiment that almost completely transformed into bassanite.

Additionally, it is important to highlight the suitability of the setup employed in this study, as this allows rapid analysis at varying water/ethanol ratios in the mixture, which can be achieved effortlessly by changing the flow rates of the pumps. This could be done remotely from the user office, avoiding security risks and human errors associated with the mandatory protocols that allow access to the hutch every time conditions have to be changed.

In summary, an *in situ* confirmation of Tritschler's results was successfully undertaken that demonstrated the appropriateness of the fReactors as a novel sample environment in synchrotron XRD analysis. To the author's knowledge, this study represents both the first time such a miniaturised CSTR device has been used in the context of *in situ* XRD analysis and also the first *in situ* confirmation of Tritschler's findings. In the work carried out here, the coupling of the capabilities of the fReactor -allowing processing of a high solid mass content- and of flow chemistry -allowing rapid solvent switching- with synchrotron XRD analysis yielded results that closely approached those of Tritschler. Concluding, the use of the fReactor as a novel XRD sample environment was demonstrated as a new capability of the device.

3.4 Summary

This chapter demonstrated the use of the equipment developed in Chapter 2. The cascade of fReactors was integrated into a previously existing rig in order to automatically perform a series of experiments. For these, the mean residence time and the molar equivalents between the reactants were changed by means of varying the flow rate of the pumps, while the temperature was changed by making use of the temperature control explained in the previous chapter.

From the data gathered from these experiments, various modelling and cross-validation techniques were used that outperformed the standard approach in this field, the least squares fitting of polynomial models. The MLSM was used along with different cross-validation routines in order to develop surrogate models that were used in single (particle swarm) and multi-objective (NSGA II) optimisation routines. Furthermore, scattered data extracted from an in-silico kinetic model was fitted using both MLSM and RBFs, leading again to improved results that resulted in increased performance: using LOOCV the error over the held out dataset was reduced by almost 5% when using the MLSM, and it could be further reduced between 6 and 15% when using RBFs.

In addition, a demonstration of the fReactors as a novel sample environment for synchrotron XRD was carried out, where a polymorph change due to solvent switching was ascertained.

Chapter 4
CFD-enabled COBR design optimisation

4.1 Introduction

As it has been demonstrated through the previous work in this thesis, the combined application of surrogate modelling and optimisation algorithms reveals itself a powerful tool in the field of chemical optimisation and discovery, which has a meaningful impact in the real world in the form of cost reduction in the pharmaceutical and fine chemical industries. This approach can also be utilised to minimise waste generation -which in some cases includes costly processes for the treatment of by-products- in an effort to make a process ‘greener’ and more respectful with the environment. In this regard, the same chemical reaction can be subject to a variety of multi-objective optimisation problems, trying to minimise more than one variable to come up with the appropriate production conditions. For example, a common approach will be to minimise the production costs while at the same time minimising waste generation. Likewise, the same techniques can be applied in the design of chemical reactors, aiming to maximise different performance metrics and exploring the trade-offs between them.

Hitherto, the reactor of choice was the fReactor, a miniaturised version of a CSTR. As has been demonstrated, the fReactor within this thesis constitutes a powerful tool for process research and development with many applications ranging from multiphasic reaction optimisation and modelling to *in situ* XRD of suspended solids flowing through it. This novel reactor helped to bridge the gap between flow chemistry automated experimentation and multiphasic reactions, overcoming the limitations of most of the reactors that operate at the same scale. However, a downside of this approach, as a countenance to its excellent mixing and particle suspension capabilities, is the broad residence time distributions generated as a result of the stirring, which increases fluid back-mixing dramatically. Therefore, the excellent mixing performance comes at the cost of different “particles” or elements of fluid spending more widely differing times in the reactor compared to its micro and milli-scale counterparts, which can lead to the formation of overreacted by-products or, on the other hand, to the channelling of particles of unreacted materials thus decreasing the efficiency of the device.

Whilst it is possible to achieve narrower RTDs by the in series combination of a larger number of CSTRs, this would create a larger volume reactor (unless further minimisation of the design is undertaken, which at the moment is considered impractical). Amongst the variety of reactors examined in the review of the literature, continuous oscillatory baffled reactors (COBRs) pose an alternative to the fReactor that can provide relatively good multiphasic behaviour and a much narrower RTD. It is possible then to apply the surrogate modelling and cross-validation techniques employed for the modelling of chemical reactions in the previous chapter to optimise the performance of a COBR reactor that could be used in chemical synthesis with excellent RTD.

Series of rules of thumb have been developed in the field of COBR design, motivating the extensive research works of Professors Ni [127, 280, 281] and Harvey,[4, 135, 137-139, 142, 282-284] although CFD as a tool for study is being used increasingly to understand flow fields within COBRs and to determine axial dispersion and shear strain rates. Mannien *et al.*[285] used CFD to evaluate axial dispersion and mixing performance of a COBR, introducing the concept of the velocity ratio R_v for a unit cell. Mazubert *et al.,*[286, 287] explored the effect of baffle geometry on the velocity ratio, pressure drop, and energy dissipation for five different COBR designs and demonstrated that COBR geometry had a huge impact on the flow behaviour.[286] The second part of their study [287] introduced novel characterisation techniques, using particle tracking to determine radial and axial fluid stretching, as well as the shear strain rate history. Reis *et al.*[142] proposed using CFD to understand the flow in COBRs through the radial and axial velocities and the vorticity, analysing the number, position and area of the vortex rings, which they assumed to be elliptical. However, methods for formal optimisation of such a problem have not yet been established, providing a strong motivation for this work.

The present chapter proposes the first CFD-enabled design optimisation methodology for COBRs. It develops innovative post-processing techniques to determine the mixing efficiency index and residence time distributions for a range of different COBR geometries and uses the data to carry out the first multi-objective optimisation of COBRs. The planar design of the COBR brings with it the potential for ease of manufacture, with flexibility in material choices, a small operating footprint and scalability. The chapter is organised as follows. The problem specification is given in section 4.2 and the CFD and post-processing methods are described in section 4.3. The effect of baffle offsetting is then discussed in section 4.4. The multi-objective optimisation methodology is given in section 4.5 with a series of illustrative results. Conclusions are drawn in section 4.6.

4.2 Problem Specification

4.2.1 Geometry Specification

The proposed plate mCOBR to be optimised in this study is shown in Figure 67. It is a simple and versatile device, consisting of an insert in which the geometry of the mCOBR has been hollowed out. This insert is held between two plates of plastic through standard M5 screws, with the upper plate providing ¼”-28 ports for the different inlets and the outlet. Such plate designs have been proven successful for heat exchangers,[288] electrochemical cells [289] and reactors where a complex geometry enhances mixing.[290] Within each channel in the insert there is a series of baffles whose geometry is specified in terms of the parameters shown in Figure 68.

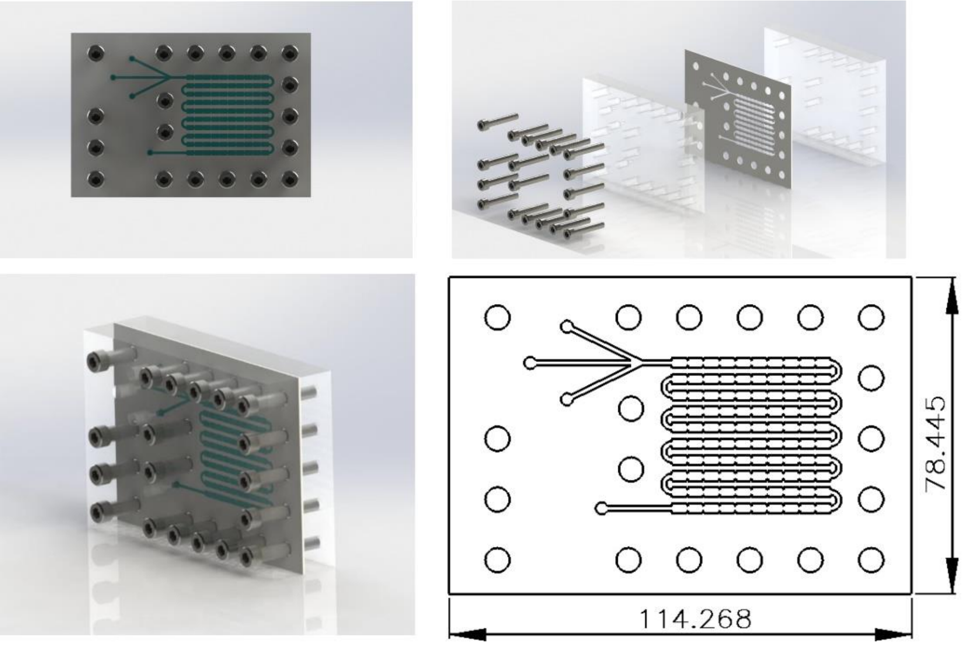

Figure 67. Proposed design of a plate mCOBR for process research and development. Top left: view of the front of the device, with three inlets and an outlet. Top right: exploded view showing the components. Bottom left: 3D view of the device. Bottom right: drawing of a possible insert for the device, with units in mm.

The relationship between different geometry parameters and its effect on mixing and residence time distribution (RTD) have been explored experimentally for tubular COBRs.[284, 291] Traditionally, the geometrical parameters studied that govern the performance of the reactor are the baffle spacing L (often as a ratio of the diameter to the reactor (L/D)), the free baffle area (e/D in the case of the planar design presented here) and the baffle thickness (b). Whilst there has been some work on smoothly varying periodic baffles in two-dimensions,[141] the design space is constrained in this study to discontinuous changes in channel width to minimise the degrees of freedom within the optimisation framework, as the main reason for the design of smooth periodic constrictions [292] was to make COBR technology available for shear sensitive bioprocesses.

The performance of the COBR is strongly linked to the geometric spacing of the baffles, which determines the hydrodynamics of the flow. When the baffle spacing is small, the vortices formed in each cell reach the next one before fully developing, thus reducing radial dispersion and deviating from plug-flow behaviour. For large baffle spacing, stagnant zones will form in each cell as a result of the vortex only extending partway through the cell. In light of this behaviour, the optimal baffle spacing will enable the full development of the vortices in the cell, generating minimal or no stagnant pockets. The baffle open area is also a very influential parameter. If it is too big, then the small baffles generate narrow and weak vortices, a situation in which the axial movement of the fluid prevails, leading to poor mixing. Such configurations also lead to a high degree of channelling (also referred to as bypassing in the literature), deviating from plug flow.[281] Contrarily, when the baffle open area is too small, the vortices generated neither expand through the entirety of the cell's cross-section or length, resulting in stagnant zones that have an adverse effect on both mixing and residence time distribution.[281] For most practical applications in tubular COBRs, the ratio baffle open area lies between 0.2 and 0.3.[284] The results of studies on the effect of baffle thickness indicate that thinner baffles promote the generation of vortices.[291]

The COBR geometries considered here are formed by combining identical two-dimensional cells with the parameters shown in Figure 68. A 2D modelling approach was undertaken, as the outcome of 2D CFD simulations for channels of varying geometry are generally regarded as a good representation of 3D models unless there are complex 3D flow structures driven by the 3D geometry.[293] Three design variables are considered: the length to width ratio (L/D), the ratio of the open space between baffles and the cell width (e/D) and the ratio of the offset distance between the baffles, a, to the total cell length, L, a/L. To the author's knowledge, the latter parameter (this is, the inclusion of an offset between opposing baffles) has not been proposed before and constitutes a novelty in the design of planar COBRs. Hence, its effect on the flow features will be a matter of research in this work. For simplicity and minimisation of computational costs in all cases considered throughout the study, the baffle thickness (b) will be set to 0.25 mm and the radius of the cell corners (R) to 0.3 mm, small values which are however above the minimum feature size our laser cutter can provide. The effect of these parameters is not considered, and the design is limited to keeping *b* reasonably low (which as has been discussed promotes vortex formation), and including rounded corners of radius *R*, aiming to reduce the stagnant pockets (also known as dead-zones) in the reactor's cells.[294] The unit cell area (A_{uc}) will be kept to 8.75 mm^2 in all cases throughout this study, regardless of the design variables. This unit cell area was chosen to give a reactor volume of 1.75 mL (similar to other mesoscale reactors) based on a manufacturable design with 100 cells cut into a plate of thickness 2 mm.

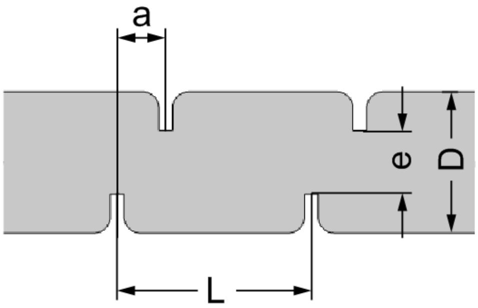

Figure 68. COBR Cell Geometry Parameters. L is the distance between two adjacent baffles, D is the width of the cell, e is the open space between the baffles and a is the offset between baffles.

4.3 CFD Modelling

Two-dimensional, laminar flow within a COBR is governed by the incompressible Navier-Stokes and continuity equations:

$$\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\nabla \cdot \left(\frac{p}{\rho}\right) + \nu \nabla^2 \boldsymbol{u} \qquad (69)$$

$$\nabla \cdot \boldsymbol{u} = 0 \qquad \nu = \frac{\mu}{\rho} \qquad (70)$$

At the inlet, the following time-dependent velocity profile for the average inlet velocity as a sinusoidal function of time is specified with a frequency of 5 Hz and a period of 0.2 seconds, defined so that the residence time in a hundred cells long reactor is approximately one minute, a typical time-scale of chemical reactions:

$$u_{av}(t) = 61.7 \cdot (0.25 + \sin(10\pi t)) \,/\, e \;\; [mm/s] \qquad (71)$$

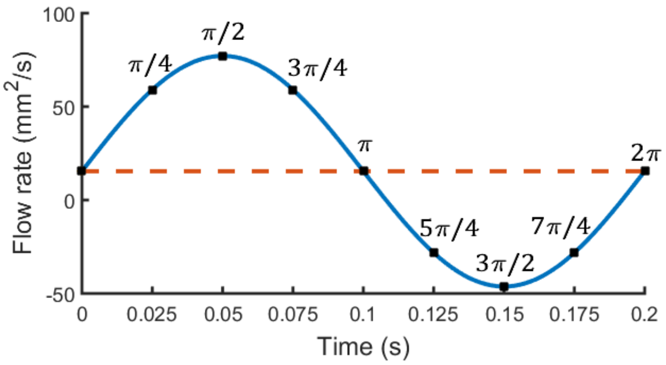

Figure 69. Inlet flow rate (Q_A) for the 2D mCOBR CFD simulation for the first oscillation (blue line). The net flow is also plotted for comparison (dashed red line). Black squares with the phase of the oscillation indicated represent the points in the cycle for which solutions were computed.

The inlet velocity has both positive and negative values and is defined to ensure that the flow rate and thus the mean residence time for all the geometries is the same regardless of their inlet width, *e*, as the flow rate is the same for all cases evaluated, so equal areas are swept in equal times by the fluid ($Q_A = u_{av} \cdot e$). A representation of the inlet flow rate can be found in Figure 69. The phase of the oscillation $\phi(t) = 10\pi t$ will be taken into account to compare the fluid flow at different moments during an oscillation.

A no-slip boundary condition is imposed for the stationary walls of the reactor, and at the outlet boundary the pressure is set to *0 Pa*. All the properties of the fluid are set to be those of water at *293.15 K,* with $\mu = 0.0010$ *Pa·s* and $\rho = 1000$ *kg/m³.*

After the flow field is determined, the transport of a dilute chemical species through advection (as a consequence of the flow) and diffusion (as a consequence of chemical gradients) is used to characterise the mixing and residence time. This transport of diluted species study (TDS) is governed by the following advection-diffusion equation:

$$\frac{\partial c}{\partial t} = \nabla \cdot (D\nabla c) - \nabla \cdot (\boldsymbol{u}c) \qquad (72)$$

where the diffusion coefficient *D* was set to 10^{-9} m^2/s, representative of a typical small molecule in an aqueous solvent.[295]

At the walls of the reactor, a condition of no flux of diluted species is imposed,

$$n \cdot (-D\nabla c) = 0 \qquad (73)$$

and at the outflow boundary the diluted species is transported by advection only.

GEOMETRY BUILDING

For every geometry governed by the set of parameters *L/D*, *e/D* and *a/L*, and knowing the area of the unit cell, it is possible to obtain the necessary values of *L*, *D*, *e* and *a* by solving a system of equations.

Writing the expression for the unit cell's area (A_{uc}) and being b the baffle thickness and R the radius of the cell's corner:

$$A_{uc} = (L - b)\, D - 4R^2\left(1 - {}^{\pi}\!/_{4}\right) + be \tag{74}$$

Where the first term accounts for the area in the space between the baffles, supposing the corners of the cell are squared. The next term discounts the area left out when the curvature is applied to the corners. Finally, the last term introduces the area of the reactor in the space left between the opposing baffles. Including the known parameters in eq. (74) and solving for *D*:

$$A_{uc} = \left({}^{L}\!/_{D}\ D - b\right) D - 4R^2\left(1 - {}^{\pi}\!/_{4}\right) + b\ {}^{e}\!/_{D}\ D \tag{75}$$

$${}^{L}\!/_{D}\ D^2 - b\left(1 - {}^{e}\!/_{D}\right) D - 4R^2\left(1 - {}^{\pi}\!/_{4}\right) - A_{uc} = 0 \tag{76}$$

$$D = \frac{b\left(1 - {}^{e}\!/_{D}\right) \pm \sqrt{b^2\left(1 - {}^{e}\!/_{D}\right)^2 + 4\ {}^{L}\!/_{D}\ 4R^2\left(1 - {}^{\pi}\!/_{4}\right)}}{2\ {}^{L}\!/_{D}} \tag{77}$$

$$L = {}^{L}\!/_{D}\ D, \qquad e = {}^{e}\!/_{D}\ D, \qquad a = {}^{a}\!/_{D}\ D \tag{78}$$

Equations (77) and (78) were implemented in Comsol Multiphysics (Comsol Inc., USA), so that for every set of design parameters inserted by the user in the parameters table in the model builder tree, the desired geometry is automatically generated following a predefined sequence of geometrical operations using Comsol's built-in CAD. In the first operation, the bottom half of a single a cell is built by means of combining additions and subtractions of rectangular entities, along with a fillet to round its corners with the specified radius, *R* (Figure 70 (a)). Once this half-cell is completed, the array feature is used to concatenate this entity as many times as specified in the parameters table in the model (Figure 70 (b)). This way, the bottom half of the reactor is built. Then, a mirrored copy of this bottom half is displaced '*a*' towards the x-direction to generate the baffle offsetting (Figure 70 (c)). Finally, two rectangles are used to complete the reactor without abrupt changes in section, forming the reactor's inlet and outlet (Figure 70 (d)). The domain formed in this way is then split into cells in order to proceed with the meshing (Figure 70 (e)). All dimensions and distances are specified in the program in terms of *L*, *D*, *e*, *b* and *R*, to make possible the automatic generation of any given geometry with minimal user intervention, which will be constrained to provide the values of *L/D*, *e/D* and *a/L* in the parameters table under ‘Global Definitions’ in the model.

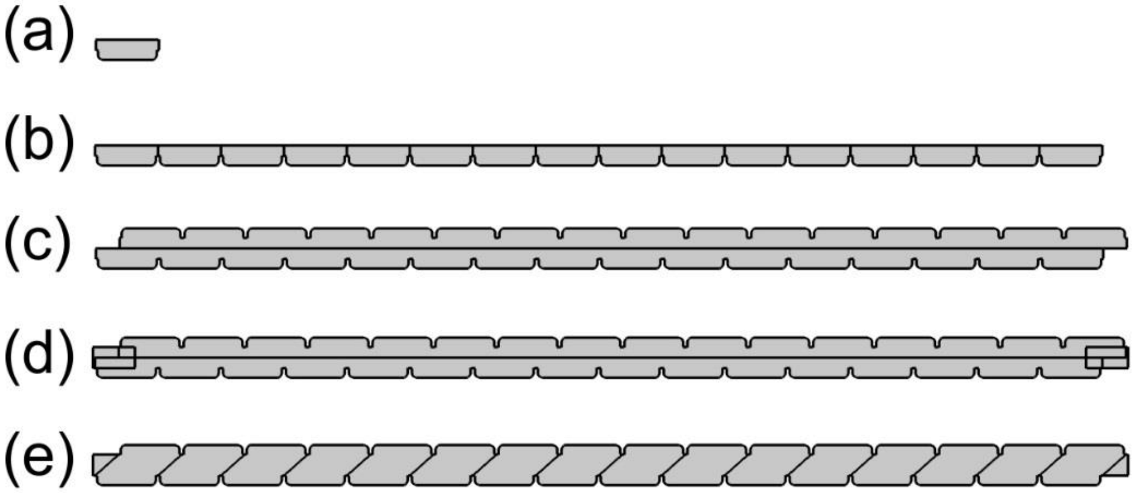

Figure 70. Phases in the automated construction of the geometry, which is entirely carried out within Comsol’s CAD. (a) The bottom half of a single cell is created using a combination of addition and subtraction operations over rectangles. (b) The array feature is used to concatenate the cell as many times as necessary. (c) The previous geometry is copied, mirrored and translated to generate the baffle offset specified by the user. (d) Two rectangles are used to provide appropriate inlet and outlet zones. (e) The cells are separated in different domains to simplify both meshing and post-processing. The geometry represented to illustrate the process corresponds to L/D=1.5625, e/D=0.525, a/L=0.375.

4.3.1 Spatial and Temporal Periodicity

The COBR studied here consists of 16 cells, as shown in Figure 70 (e). This is chosen as a geometry under which a periodic flow across a number of cells is established; optimising this shorter geometry will provide the design of a cell which can be repeated within a longer geometry to give a reactor of a pre-determined volume.

In the case of the mixing efficiency index used for this study, this metric will yield the same value regardless of the length of the reactor, since it is the same for any cell of the domain once spatial and temporal periodicity has been assured. On the other hand, in the case of the residence time distribution, the length of the reactor is a crucial parameter. The approach followed in this study allows the best geometry to be obtained in terms of RTD for a longer reactor to be extrapolated from a shorter, 16-cells reactor on which computations can efficiently be carried out. This is achieved by using a 1D axial dispersion model,[67] which, as will be ascertained later on in this sub-chapter, ensures that the ranking of the different geometries in terms of the variance tested in the study will remain the same for a longer reactor. A pivotal issue to analysing the performance of a COBR design is to be able to determine when spatial and temporal periodicity is achieved within the COBR. This will be considered next.

STEADY FLOW

As a first approximation a steady flow with a constant inlet velocity, set to the maximum positive inlet velocity, is analysed in order to indicate where inlet boundary effects can be neglected. A grid independence test was carried out evaluating the velocity magnitude in the central point of the cell. The study was computed for five different meshes using the finite element software Comsol Multiphysics (Comsol Inc., USA), ranging between 1264 and 11484 mesh elements per cell. It was ascertained that a mesh of 5276 elements produced a mesh-independent solution, as presented in Figure 71 (a). Consequently, meshes of this size were employed across all geometries examined in the present study. A histogram for vertex skewness is presented in Figure 71 (b). The mean skewness is 0.1853. It can be seen that the vast majority of elements have a skewness less than 0.4, indicative of a high-quality mesh. A typical mesh of 5000 elements used in this study is shown in Figure 71 (c).

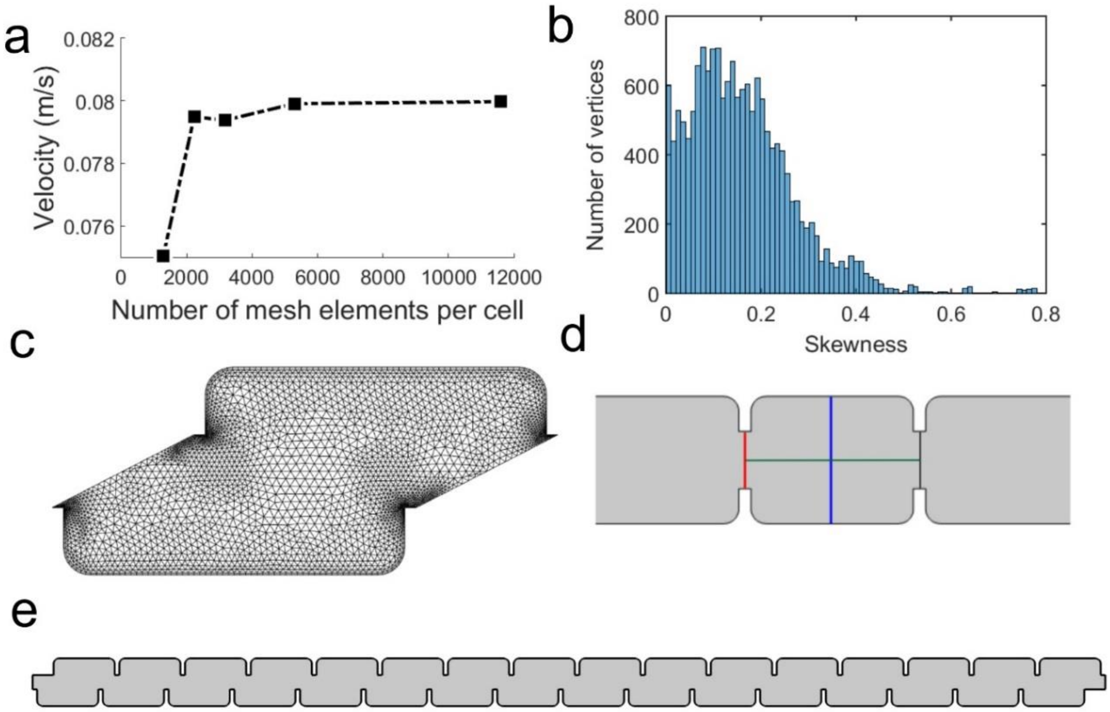

Figure 71. a) Velocity at the middle point of the geometry for a variety of different mesh sizes. b) Histogram showing element skewness. c) Mesh with 5816 elements per cell as used in this study. d) Cut lines on which solutions are obtained. e) Example of a full 16 cells long computational domain.

Fluid velocities on the green longitudinal cut lines shown in Figure 71 (c), for cells one to four (counted from the inlet) are given in Figure 72. The relative error between the velocities in cells *i* and *j* for every mesh node along the line is calculated and expressed as a percentage according to

$$error\ i,j = 100\left|\frac{v_i - v_j}{v_j}\right|, \quad i < j \tag{79}$$

where v_j is in the denominator because the velocities in the cells as we move apart from the inlet are supposed to be free from any inlet effects and, in consequence, expected to be closer to the expected periodic solution.

Figure 72 shows that the maximum percentage error in velocity between cells 3 and 4 is less than 0.125%. On this basis, it is taken that the inlet boundary effects are negligible from cell three onwards and hence, spatial periodicity is reached from this point.

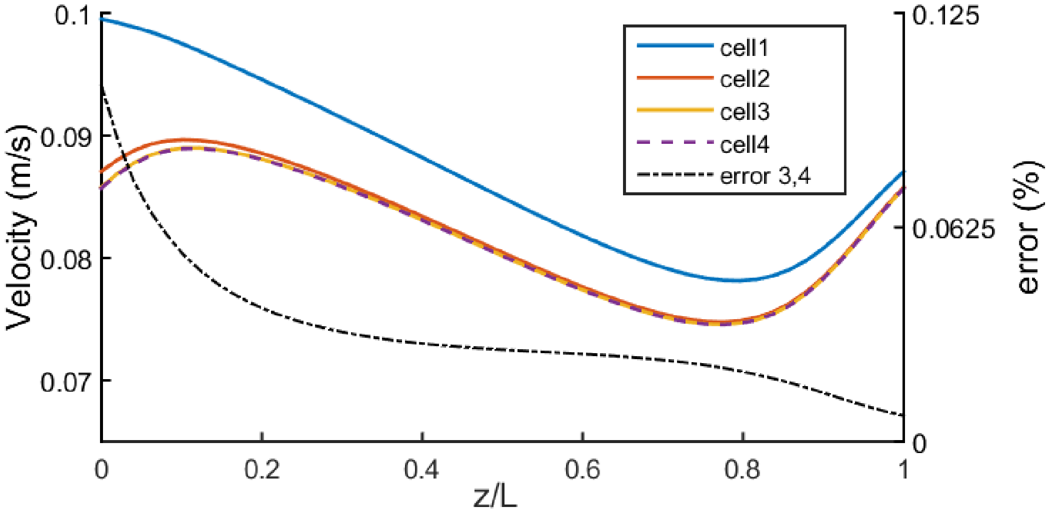

Figure 72. Longitudinal velocities on the green cut line for the first four COBR cells.

In the same fashion, the periodicity for the stationary solution is also evaluated for the other end of the reactor, in the interest of determining from which cell apart from the outlet, its effects can be considered negligible. Unlike eq. (79), this time the denominator of the error expression will be the velocity of the cell further from the outlet, for the same reasons considered before (which means in this case equation (79) holds but with *i>j*). As depicted in Figure 73, there is a significant discrepancy between cells 16 and 15, the error nearing 4% at the outlet of the cells, that drops significantly for the next pair of cells (14 and 15) for which the error is negligible and clearly decreases as we move away from the outlet towards the inlet. Comparing with cells 3 and 4, cells 14 and 15 present lower errors in terms of spatial periodicity. Hence, we can consider that the solution has reached spatial periodicity from cell 3 up to cell 15. Nevertheless, further testing is needed to assess if the solution for the time-dependent oscillatory flow problem can also be considered periodical over this range of cells.

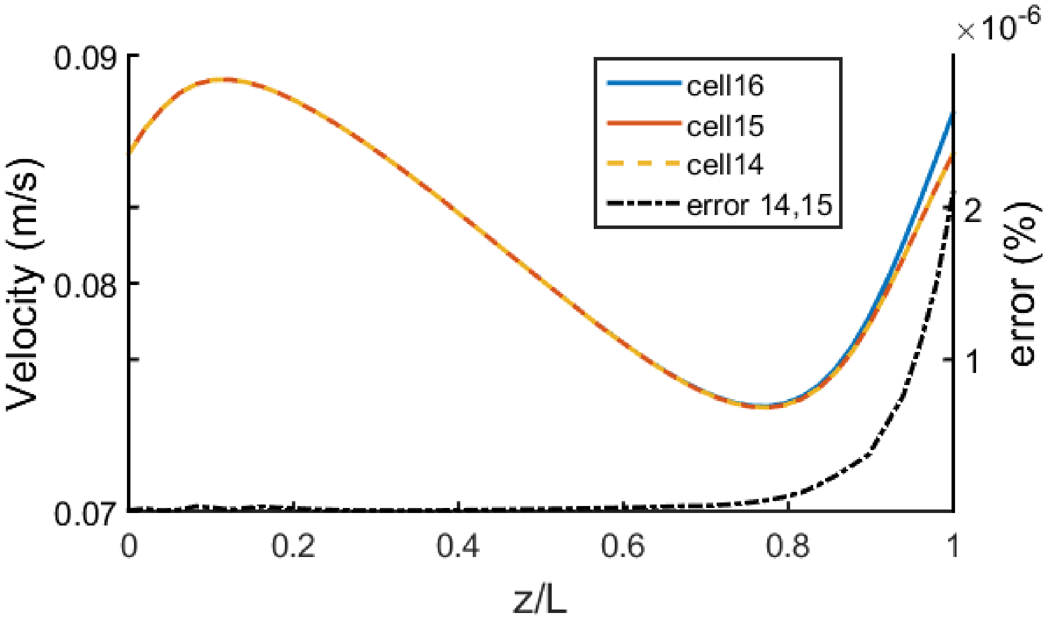

Figure 73. Longitudinal velocities on the green cut line for the last three COBR cells.

TEMPORAL PERIODICITY

Temporal periodicity is considered next using the mesh optimised for the steady-state simulation. The time-stepping for the time-dependant solver is set to 'strict', to ensure solutions are calculated at the specified times in the study settings section. In this fashion, the solution is calculated at 0.025 s (or π/4 in the case of the phase) intervals, amounting eight times per cycle, although this configuration does not constrain the overall number of steps taken by the solver (meaning that it will take as many steps as necessary to safeguard convergence, but the computation of solutions at the designated times is ensured). This precaution is taken to avoid working with interpolated data instead of actual solutions computed at the specified times within the cycles. The transverse velocity profiles along the centre line (blue cut line in Figure 71 (c)) of the third cell were determined for a time-dependent simulation over one second, corresponding to five complete periods. Figure 74 shows the differences between the velocity profiles of cycles 4 and 5 for the third cell in the reactor. It can be observed that the discrepancy amongst the 4th and 5th time period was typically less than 0.3%. Likewise, Figure 75 depicts the same metrics for cell 15, for which the conclusions are analogous and the solution can continue to be considered periodic in this case.

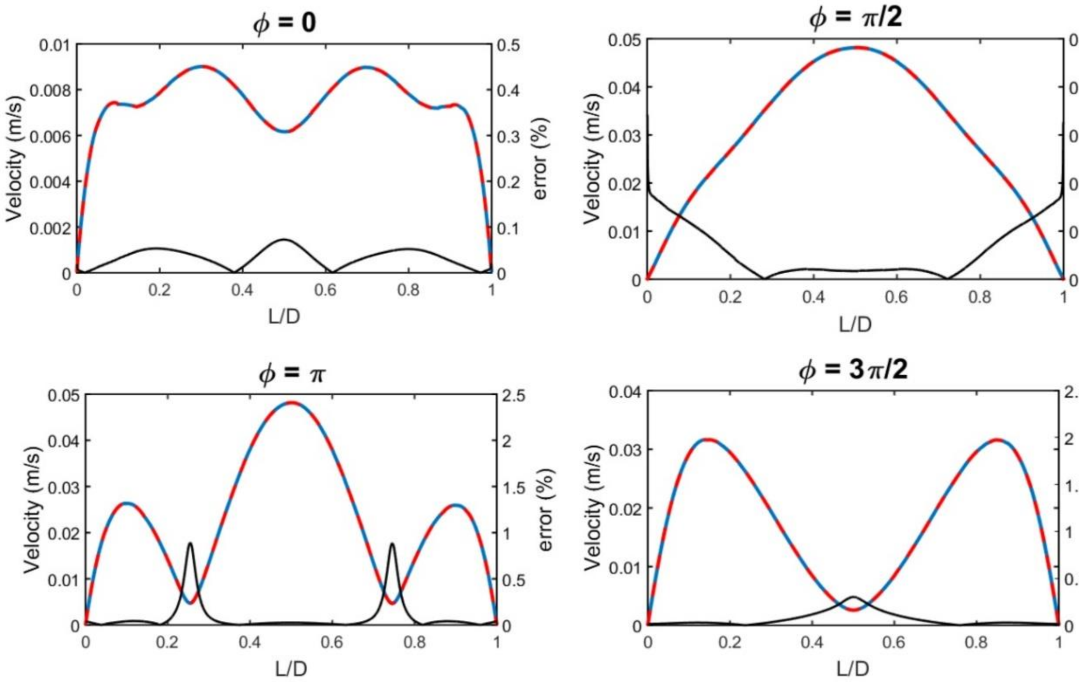

Figure 74. Transverse velocity profiles and the discrepancy between cycles 4 and 5 for cell 3 for different oscillation phases. Blue lines: velocity for the 4th cycle. Dashed red line: velocity for the 5th cycle. Black line: error.

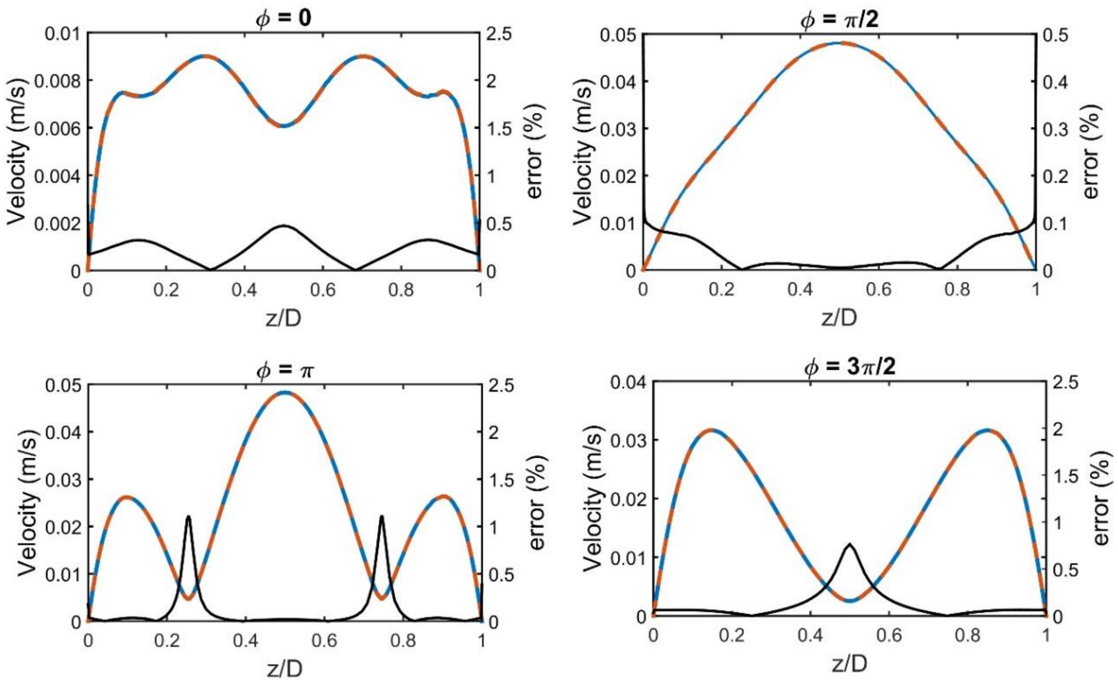

Figure 75. Transverse velocity profiles and the discrepancy between cycles 4 and 5 for cell 15 for different oscillation phases. Blue lines: velocity for the 4th cycle. Dashed red line: velocity for the 5th cycle. Black line: error.

Spatial periodicity was again considered for the time-dependent scenario to verify the conclusions drawn from the preliminary steady study. A comparison of the different velocity profiles along the three cut lines defined in Figure 71 (d) for cells 3 and 4 and the 5th cycle of the solution calculated previously revealed that the discrepancy was significantly below 2% at any time of the fifth cycle, indicating spatial periodicity of the time-dependent solution from cell 3 onwards. Table 16 presents the values of the maximum relative error between the velocity profiles and for different oscillation phases of the fifth cycle. Similarly, Table 17 lists the maximum relative errors between velocities for the 14th and 15th cells, for which the errors are significantly lower and hence, completely acceptable for further evaluation. Based on these results, the periodicity of the solution is guaranteed from cells 3 to 15, this being the domain free from inlet or outlet effects.

Table 16. Maximum relative error between cells 3 and 4 at different oscillation phases of the 5th cycle and cutlines (%).

ϕ	0	$\pi/4$	$\pi/2$	$3\pi/4$	π	$5\pi/2$	$3\pi/2$	$7\pi/2$
Blue	0.65	0.42	0.25	0.64	0.70	0.60	0.57	0.59
Red	0.49	0.35	0.55	0.40	0.31	0.45	0.37	0.60
Green	0.61	0.28	0.27	0.48	0.65	1.20	0.38	0.57

Table 17. Maximum relative error between cells 14 and 15 at different oscillation phases of the 5th cycle and cutlines (%).

ϕ	0	$\pi/4$	$\pi/2$	$3\pi/4$	π	$5\pi/2$	$3\pi/2$	$7\pi/2$
Blue	3.17×10^{-3}	2.83×10^{-4}	1.74×10^{-4}	5.79×10^{-5}	9.23×10^{-5}	1.66×10^{-4}	5.426×10^{-4}	1.398×10^{-3}
Red	9.95×10^{-3}	7.35×10^{-3}	9.86×10^{-4}	9.28×10^{-4}	2.98×10^{-3}	2.34×10^{-3}	9.61×10^{-3}	4.10×10^{-3}
Green	1.26×10^{-2}	2.21×10^{-3}	4.23×10^{-4}	1.32×10^{-4}	1.31×10^{-4}	8.61×10^{-4}	1.53×10^{-3}	1.96×10^{-3}

Based on these findings, CFD velocity solutions for every geometry assessed were computed for five complete cycles. As has been proven before, the periodic time dependant flow from the fifth cycle (lasting 0.2s, ranging from 0.8 to 1s) of this solution constitutes the representative flow field of a complete oscillation.

A transport of diluted species study follows the laminar flow study. Attending to eq. (72), velocity values ($\boldsymbol{u}$) are needed to solve the concentration field. By default, the software would proceed by calculating the flow field by solving eq. (69) before eq. (72) for every step taken by the solver, incurring high computational times. Instead, a strategy was implemented to exploit the periodicity of the representative cycle calculated in the first laminar flow study, eliminating the need to solve eq. (69) again. Regardless of the instant being evaluated during the time-dependant study of the transport of diluted species, the velocity field for it can be accessed from the corresponding phase of the representative cycle, instead of being re-calculated. This ultimately means that the transport of diluted species study can be carried out for a long time interval, circumventing the need to solve the flow field over this timeframe, thus improving computational efficiency.

The first step towards this goal is the definition of a function such that for any time of the TDS study t, the time t' of the laminar flow study that corresponds to the same oscillation phase is returned. The sawtooth function an1(t) pictured in Figure 76 (a) accomplishes this – noting that the flow field is simulated over 5 periods to achieve convergence. Consequently, the converged solution is contained between 0.8 and 1 where it is representative flow field cycle. The use of this function within the built-in function ‘withsol’, which is a pointer to a previously computed solution, allows the transport of diluted species study to access the stored velocity solutions of the previous laminar flow study, avoiding calculation of the velocity field for every step taken by the solver. In the problem tree of the TDS study, the transport properties section can be modified to implement this strategy (Figure 76 (b)). Under ‘convection’ the software allows the velocity field to be specified by the user, with separate fields for the x and y directions. The commands withsol(‘sol1’,u,setval(t,an1(t))) and withsol(‘sol1’,v,setval(t,an1(t))) point towards the corresponding solutions of the velocity field in the x and y directions, respectively.

In summary, the software solves eq. (72) to find the concentration field extracting $\boldsymbol{u}$ from the velocity solutions of the representative flow field calculated in the previous laminar flow study (which is only computed for five cycles until periodicity is ensured). This approach dramatically reduces computational times when compared to solving both eq. (69) and eq. (72) for the full length of the study. This is especially relevant in the case of the 100 cycles long studies carried out in this chapter, for which under the default conditions of the CFD software the velocity solution would have been calculated for 100 instead of 5 cycles.

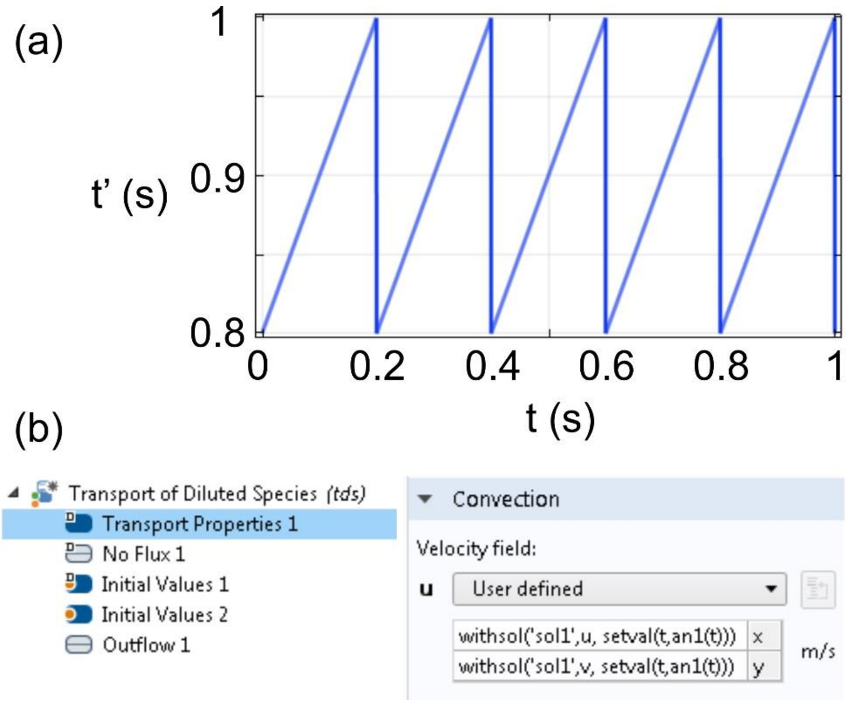

Velocity field:

u User defined

withsol('sol1',u, setval(t,an1(t))) x

withsol('sol1',v, setval(t,an1(t))) y

m/s

Figure 76. (a) Sawtooth function relating the time of the TDS study (t) to the corresponding time of the fifth cycle for the fluid flow solution (t'). (b) Comsol's model builder tree, showing the convection section for which the velocity field can be defined by the user.

VALIDATION

To ensure confidence in the numerical scheme, validation was carried out by comparing a CFD model to the experimental data taken from Armaly[296] for a backward-facing step under laminar flow conditions. This was chosen as the flow features within this problem are similar to those found within the COBR. The number of mesh elements for the validation was chosen to give a similar density between the two simulations. Consequently, 1237 mesh elements were used for the region of the backward-facing step domain which is geometrically similar to a quarter of a COBR cell, as can be appreciated in Figure 77.

In order to contrast the meshes employed by the solver, a quarter of a COBR cell is compared with the equivalent geometry in the backward-facing step. These geometries are different in size, but the relationships of the region of interest's features with the step size (baffle height in the case of the COBR) are kept the same as shown in Figure 77.

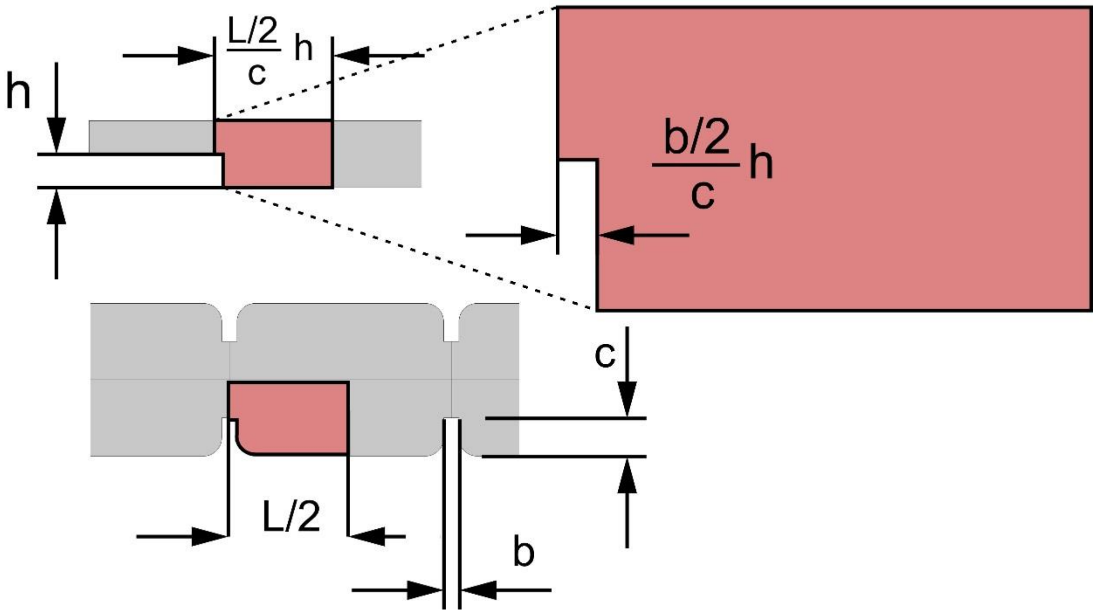

Figure 77. Region of interest selected for comparison of mesh densities for the backward-facing step and COBR geometries.

The number of mesh elements for the highlighted area in the backward-facing step experiment is 1312 compared with 1237 in the case of the quarter of COBR cell. The quadrangular elements of the top boundary layer for the backward-facing step were discounted, as this boundary does not exist in the case of the COBR. The similarity of mesh densities is high, with less than 6% difference in the number of elements generated by Comsol's predefined physics-controlled 'fine' mesh.

A graphical representation of the model for the backward-facing step problem can be found in Figure 78. The step size, h, is selected to be 5 mm, following Armaly's experiment. The computational domain is chosen to be long enough to neglect exit effects (20·h). Shown in white, the line that meets $u = 0$ (null axial velocity). The distance between the step and the intersection of this line with the bottom wall is known as the reattachment length (X_r), commonly defined as the length from which the fluid resumes flowing with a positive velocity in the flow direction in all its cross-section. This metric was calculated under different Reynolds numbers and the results compared with the experimental values obtained by Armaly. For every Reynolds number, the mean inlet velocity for the CFD model is calculated using eq. (80) and isolating u_{avg}.

$$Re = \frac{\rho u_{avg} D_h}{\mu} \tag{80}$$

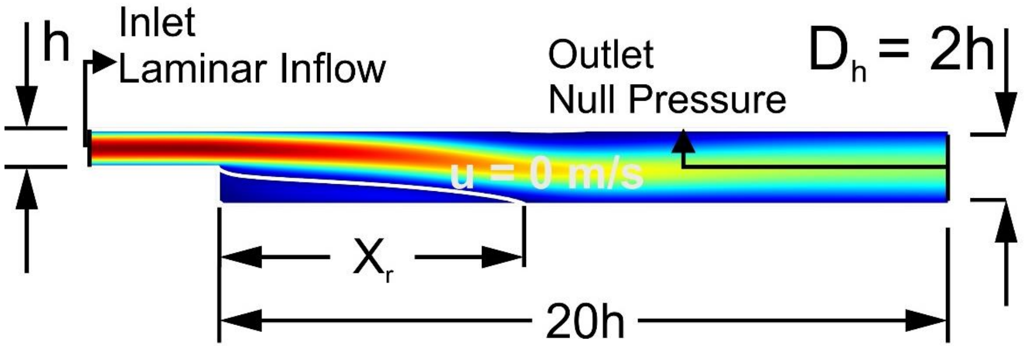

Figure 78. CFD geometry parameters for the backward-facing step study. Colour: velocity magnitude for the case Re=400.

The CFD model was shown to be converged and that there was good agreement between these predictions and the experimental data in terms of the reattachment length for Reynolds numbers up to 500 (Figure 79) and velocity profiles at different lengths from the step and under different Reynolds numbers (Figure 80 for Re = 100 and Figure 81 for Re = 389).

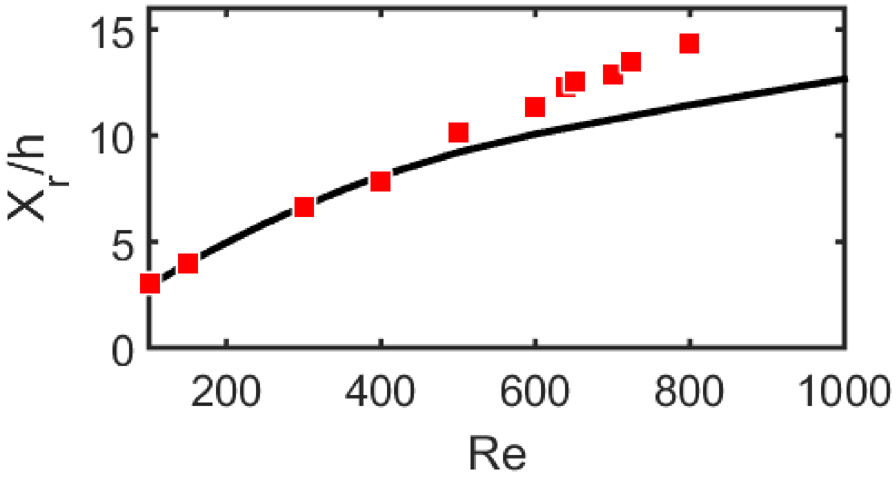

Figure 79. Re-attachment length normalised by the step size under different Reynolds numbers. Black line: CFD simulation. Red squares: Armaly's experimental data.[296]

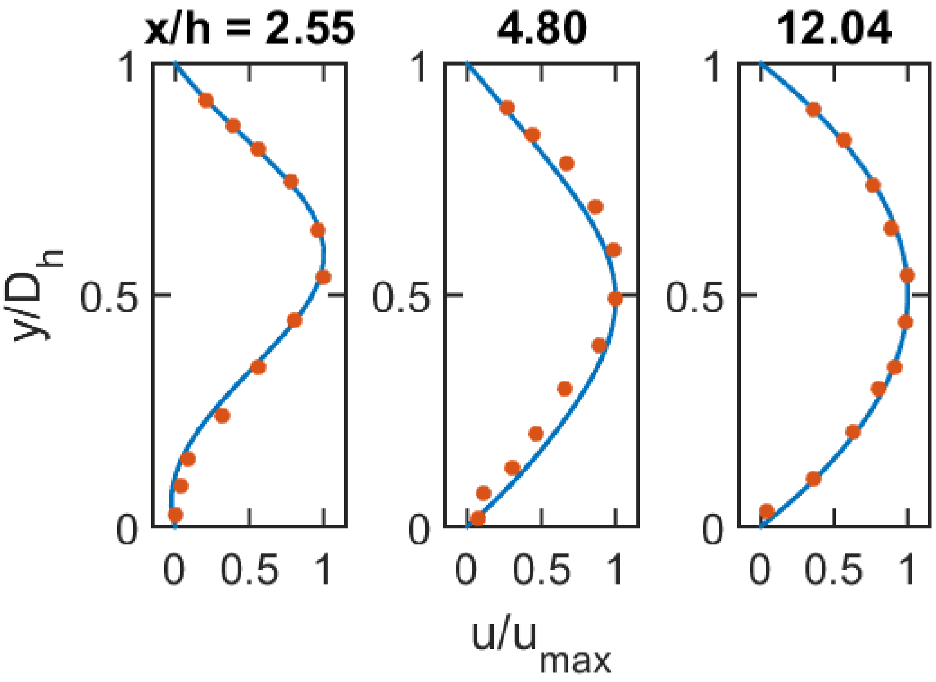

Figure 80. Axial velocity profiles for three different sections downstream of the step for Re = 100. Height is normalised by the height of the channel and velocities by the maximum velocity for better comparison. Blue line: CFD results. Red dots: Armaly's experimental data.[296]

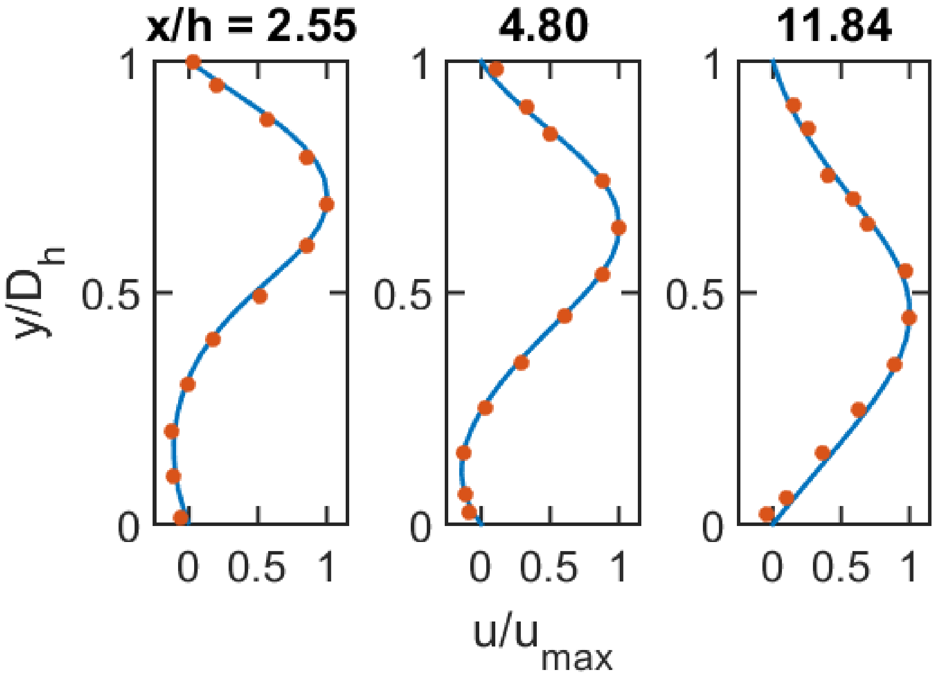

Figure 81. Axial velocity profiles for three different sections downstream of the step for Re = 389. Height is normalised by the height of the channel and velocities by the maximum velocity for better comparison. Red dots: Armaly's experimental data.[296]

In this context, the Reynolds number for the COBR can be defined as

$$Re = \frac{\rho u_{avg} \, {}^{D}/_{2}}{\mu} \tag{81}$$

to allow direct comparison with the Reynolds number defined for the backward-facing step. As previously presented in this section, *D* is the cell width and *D/2* is the length related to D_h in the case of the geometry of a cell. Taking this into account, the highest Reynolds number reached for an mCOBR simulation across all cases evaluated is *Re=129*. Looking at Figure 79, this value falls well below 500 and in the flow regime for which CFD simulations are considered to fit exceptionally well with the experimental data. For instance, for the case *Re=150*, the relative error between the proposed CFD model and the experimental data is less than 0.5%. Due to the similarity of the flow features and operating conditions, this gives confidence in the ability to capture the behaviour of the COBR using CFD.

4.3.2 Post-processing of CFD Results

MIXING EFFICIENCY INDEX

The CFD solutions are used to provide information on how effectively a reactor mixes the chemical species and to determine the residence time distribution in the flow. A range of methods have been used to assess the mixing capabilities of different reactor geometries with, perhaps, the simplest being to calculate the standard deviation of the concentration (or some associated property, such as the pixel intensity when using a camera and dye marker in experimental studies) in a cross-section of the reactor.[297-299] The major drawback of this metric is that it does not allow comparison between different studies since it is not dimensionless. The corresponding dimensionless quantity, where the standard deviation is divided by the mean concentration, is referred to as the absolute mixing index (AMI).[300] This index, although dimensionless, still depends on variables such as lighting conditions and the properties of the type of ink used in experimental studies. A related and more refined approach is the use of a relative mixing index (RMI) which compares the standard deviation of the concentration in a cross-section with the standard deviation of the unmixed state as depicted in Figure 82 (top) at time 0.[301, 302] The RMI is generally considered more versatile for comparing different studies. Other approaches include the use of particle tracking to determine mixing efficiency by calculating radial and axial fluid stretching,[287] albeit at a higher computational cost. The relative mixing index is used in this study and defined as

$$\eta = \left[1 - \left(\sqrt{\frac{1}{N}\sum_1^N (C_i - C_{p,i})^2} \Big/ \sqrt{\frac{1}{N}\sum_1^N (C_{p,i} - C_{0,i})^2}\right)\right] \qquad (82)$$

where C_i represents the concentration of the species at the location *i* for the case being evaluated, and $C_{p,i}$ and $C_{0,i}$ are the concentrations of the perfectly mixed case and the totally unmixed state, respectively, at the same location *i*. It can be easily verified then, that η will yield values of 0 and 1 for the unmixed case and the perfectly mixed case, respectively.

With an oscillating flow and complex shape of the cells, the relative mixing index η varies with both time and position of the cross-section evaluated in the cell. Besides, the inclusion of the baffle offsetting leads to the cell outlets of different geometries varying in length and orientation depending on *a/L*, complicating direct comparison of this metric, which is usually applied at the outlet of the geometry, between different geometries. A new strategy was therefore developed here to assess the mixing efficiency in a time-dependant flow field. At initialisation, the reactor is split in half as shown in Figure 82 with the bottom half of the reactor set to a concentration of *100 mol/m*^{3} (effectively a concentration of 100%) while the upper half is set to a concentration of *0 mol/m*^{3}. This is the most adverse situation, corresponding to a completely unmixed state for which $\eta = 0$. Then the species is allowed to transport over time according to equation (72) using the flow field obtained from a single representative period as described previously. As in the case of the laminar flow study, the time stepping in the solver is set to ‘strict’ to provide accurate solutions at the eight specified oscillation phases for each cycle, based on the real solutions provided by the software and not merely interpolations.

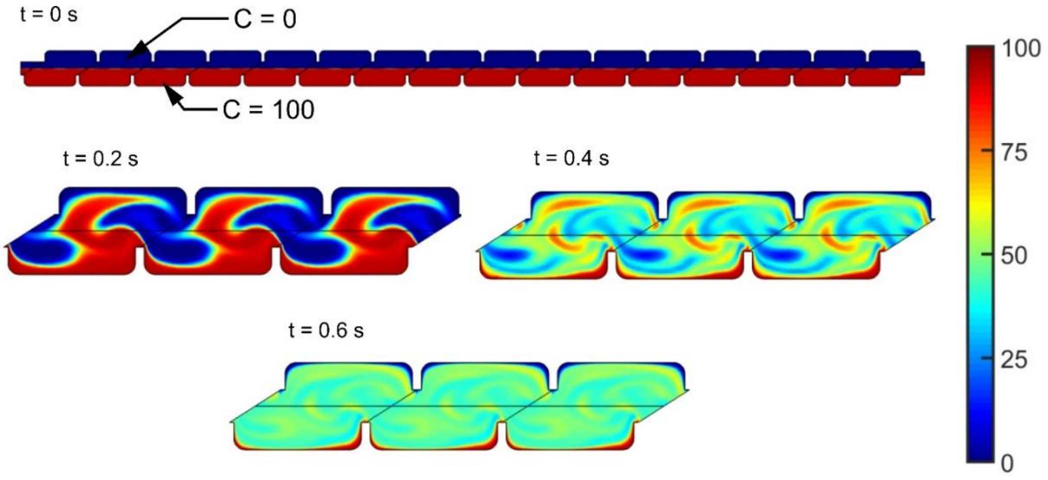

Figure 82. Initial state of the reactor (top figure) and state of three adjacent cells for three different times. The geometry parameters in this case are: L/D=1.5, e/D=0.5 and a/L=0.25. Colour: concentration [mol/m^3].

The degree of mixing is then calculated within the 13th cell from the inlet by evaluating η on a grid of 50 by 25 evenly-spaced points, covering the whole area of the cell, at 8 separate time points within a cycle. In this fashion, instead of calculating the mixing efficiency just at the outlet or a particular cross-section of the reactor, the mixing state of the whole cell over time is evaluated. This cell was selected to avoid inlet and outlet effects, falling in the range of spatial and temporal periodicity ascertained in section 4.3.1 of this chapter. The assumption is that, for a model of a sufficiently long reactor, two consecutive cells can be found for which the concentration field is not affected by the entrance or exit effects. Consequently, with identical fluid velocity field evolution over time and also equal initial concentrations, the concentration field will also be the same in the two adjacent cells for different oscillation phases, thus allowing to determine the mixing efficiency index without being affected by spurious or transitory effects.

This will be achieved by means of comparing the concentration fields of two adjacent cells (13 and 14) that are included in the domain for which the flow field was ascertained to be free of inlet or outlet effects (termed "the regular domain"). Figure 83 (a) depicts the concentrations in cell 13 for an instant early in the time series, where a high concentration gradient still exists. To verify the assumption that the concentration fields of different cells in the regular domain can be considered identical, the relative error between the concentrations of cells 13 and 14 (eq. (84)) is plotted in Figure 83 (b). It can be noted that all the high relative error values are comprehended in an area corresponding with the low concentration region of cell 13 (blue area in Figure 83 (a)). This is indicative of a small value of the denominator in eq. (84). In order to ascertain that this is the case, the absolute error was also calculated and represented in Figure 83 (c), from which it is immediate to notice that for the locations where the relative error is high (yellow areas in Figure 83 (b)), the absolute error falls well under *1.25* mol/m^3 on an overall scale of 100 mol/m^3 – ie <1.25% error This means that the initial assumption is sensible and that the concentration fields are sufficiently similar between the cells of the regular domain to yield almost identical mixing efficiency values. In fact, the mixing efficiency indexes yielded for cells 13 and 14 in the state depicted in the figure were of 0.1409 and 0.1416, respectively, amounting to a relative error of 0.5% between them. Repeating the study for 18 complete cycles and assessing η for 8 phases per cycle, the maximum relative error between the cells was of 1.33%, with over 94.4% (136 out of 144) of the cases assessed yielding a maximum relative error lower than 0.5%. This proves the validity of the approach: regardless of the cell in the regular domain for which η is assessed, the obtained values are so similar that any of them can be used to characterise the mixing performance of the reactor as a whole.

$$error\ 13{,}14 = 100\left|\frac{c_{13} - c_{14}}{c_{13}}\right| \qquad (83)$$

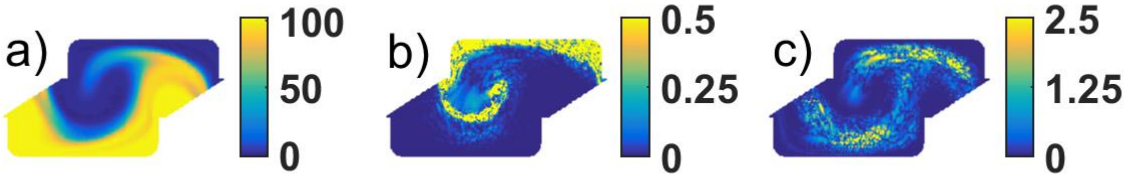

Figure 83. a) Colour map for the concentrations (mol/m^3). b) Relative error between cells (dimensionless). c) Absolute error between cells [mol/m^3].

In order to calculate η from the CFD solutions, the concentration data gathered were exported as a .txt file. A data set with the selection of the 13th cell was defined under 'Results>Data Sets' to export the data evaluated in the grid of points specified under 'Export>Data>Output'. The text file generated can be easily imported using Matlab, for which a script was written to process the raw concentration data contained in the file. A three-dimensional matrix is built in which the raw data is split and can be accessed in terms of cycle number and time stamp number within each cycle. In this way, it is easy to implement a loop to calculate η according to equation (82) for every cycle and timestamp given. Finally, the mean of η across all eight phases of each cycle is computed, giving an overall value of η for each cycle.

Figure 84 shows the calculated η values for 4 different reactor designs. For every geometry and cycle, the mean η of the cycle is plotted. It is interesting to see that, using this method, the ranking of the geometry is the same regardless of the cycle in which the mixing is being assessed, although the differences in η vary. The mixing performance will be captured at the 6th time period since this shows good sensitivity between different designs and also minimises the consumption of computational resources. In summary, the metric that is going to be used as an objective function during the optimisation procedure is calculated as the mean of η for the eight oscillation phases for which CFD solutions were collected for the 6th cycle.

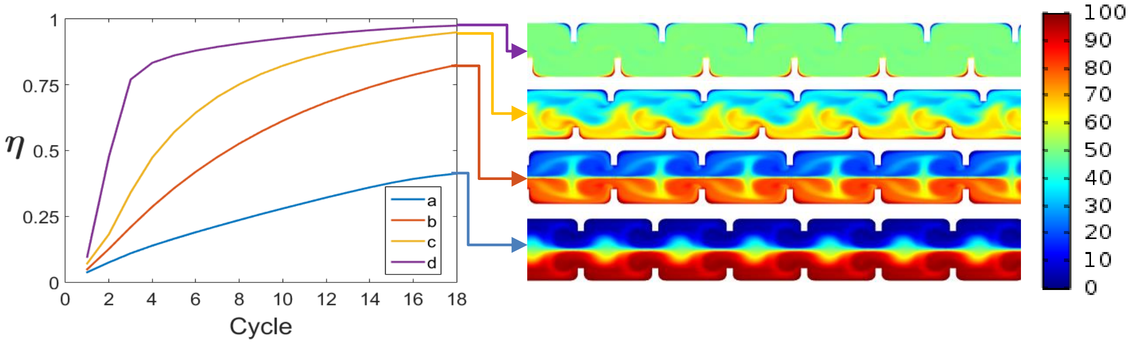

Figure 84. Temporal evolution of the mixing efficiency index for different geometries (left) and representations of the concentrations in the reactors for the last timestamp of the 6th cycle (right). The geometries, from top to bottom, correspond to values of *L/D*, *e/D* and *a/L* of (1.375, 0.3, 0.5), (1.75, 0.6, 0.25), (1.5, 0.45, 0) and (1, 0.6, 0), respectively.

RESIDENCE TIME DISTRIBUTION (RTD)

Following previous CFD studies for COBRs, [285] the RTD was determined using a tracer study approach.[67, 121] Initially, the concentration was set to *100 mol/m*^{3} at the cut line perpendicular to the longitudinal axis of the reactor in the 4th cell shown in Figure 85 (t=0s, marked “tracer”), with the concentration in the rest of the reactor set to zero. The average concentration in the cross-section 10 cells downstream of this point was then plotted as a function of time. In this way, the area of the reactor in which the species evolution is analysed was fully encompassed between cells 3 and 15, and hence in the zone of the reactor for which spatial periodicity is ensured, meaning that velocity fields for every cell in this region can be considered the same (with a minimal discrepancy as demonstrated in section 4.3.1 of this chapter). Therefore, spurious inlet and outlet effects are avoided in obtaining the RTD. Figure 85 illustrates the transport of the tracer material over time, where dimensions for the starting position of the tracer and the cut line used as a dataset in the 14th cell (for which the average concentration data will be gathered over time) are provided. Both the tracer and the cut line occupy the whole cross-section of the reactor (*D*), avoiding the locations of the baffles.

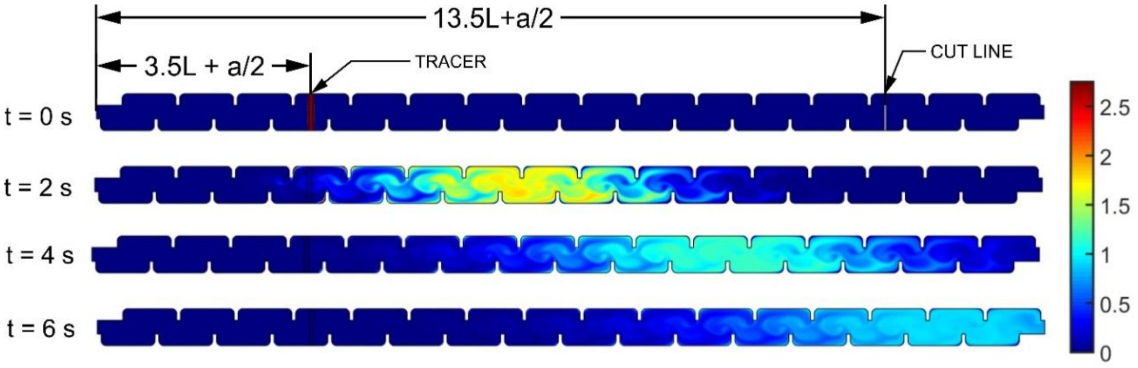

Figure 85. Concentrations at different times (0, 2, 3 and 6 seconds) for the tracer study. The geometry corresponds to values of L/D=1.5, e/D=0.3 and a/L=0.25.

An example of a concentration curve, obtained by calculating the mean of the concentration at the cut line shown in Figure 85, is shown in Figure 86. Due to the oscillatory nature of the flow, the curve has a series of peaks and valleys superimposed onto some overall response. Data were exported from Comsol into a .txt file for the phases for which the solutions are calculated, which again is at *π/4* phase intervals or 0.025 s. Splitting the data into different oscillation's phases results in smooth curves without this perturbation. In this way, 8 different RTD curves were calculated from the concentration data, each one corresponding to a phase of oscillation; concentration curves for cycle phases of *3π/4* and *7π/4* are shown in Figure 86. All concentration values lie between these phases for the case illustrated and the curve has a bell shape similar to that reported for tanks-in-series,[67, 121] being lightly skewed towards the left.

The concentration curves for the 8 different phases (for which, as explained before, the software was forced to provide solutions) within a time period of 20s were calculated using Matlab and the RTD determined from a velocity weighted average of this data, since as it was specified previously, the average velocity of the flow changes over time. With the aim of calculating this velocity-averaged-RTD, the split data generated for each oscillation phase, which provides concentration data at 0.2 s (the oscillation period) intervals had to be interpolated to provide concentration values every 0.025 s for each phase. Matlab's built-in function '*interp1*' was used to provide a piecewise cubic Hermite interpolating polynomial for every phase of the oscillation representing concentration over time. After they have been interpolated, the velocity averaged RTD can be easily calculated. The RTD function, $E(t)$, is then easily obtained by dividing the resulting concentration curve by the area under it, as indicated in equation (84).

$$E(t) = \frac{C(t)}{\int_0^\infty C(t)dt} \tag{84}$$

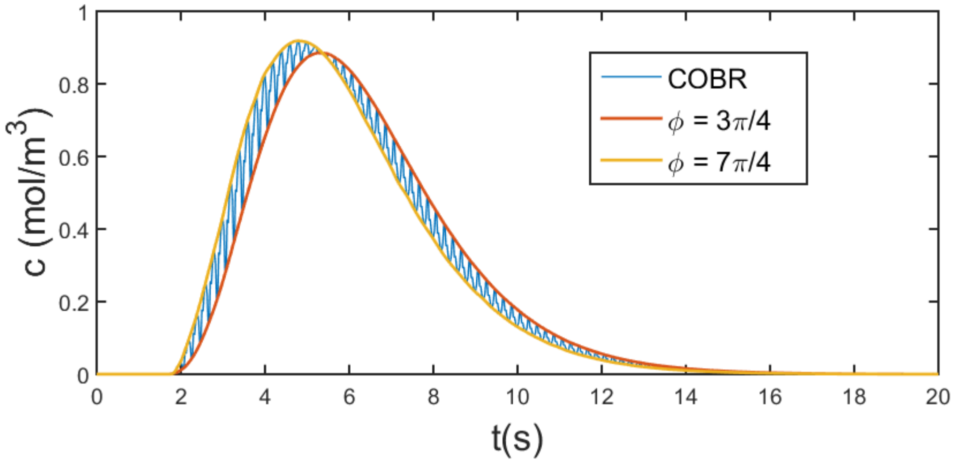

Figure 86. Concentration curve of a COBR obtained from a pulse input simulation for the geometry *L/D*=1.375, *e/D*=0.45 and *a/L*=0.

The dispersion of an RTD can be quantified by its second moment, its variance,[121] that gives a measure of the 'spread' of the distribution, with small variances indicating a narrow RTD. This is defined by

$$\sigma_t^2 = \int_0^\infty (t - t_m)^2 \, E(t)dt \tag{85}$$

This value is calculated by the algorithm, which also provides values for the third (skewness) and fourth (kurtosis) moments of the $E(t)$.

The RTD's variance yielded for these 10 cells is considered here to be representative of the performance of a longer reactor, being the central assumption that different geometries will rank the same regardless of the number of cells assessed. For convenience and minimisation of computational resources, the hypothesis that the RTD yielded over this 10 cells section of a 16 cells long reactor can be used to calculate the variances of longer reactors by using a 1D axial dispersion model was proposed. In order to confirm this assertion for the current study, a longer reactor consistent of 56 cells is considered, for which RTD curves will be obtained at cut lines separated 10, 20, 30, 40 and 50 cells apart from the initial position of the tracer in the reactor. The RTD variance data will then be fit to a 1D axial dispersion model presented in Levenspiel's book [67] that has already been used by other authors to successfully model the RTDs of COBRs,[282] for which a large deviation from plug flow ($D/uL > 0.01$) is assumed. This model, developed by Levenspiel and Smith in the late '50s,[303] introduces asymmetries in the RTD to account for the deviation from plug flow, skewing the distribution towards the left, a feature that is clearly recognizable in Figure 86. As the flow is undisturbed both as it flows through the cross-section containing the tracer at *t=0* and through the different cut lines defined for concentration measurement in the simulation, it is reasonable to use open-open boundary conditions, under which the equation of this model of the variance takes the form

$$\sigma_\theta^2 = 2\frac{D_a}{ul_m} + 8\left(\frac{D_a}{ul_m}\right)^2 \tag{86}$$

where l_m is the length across the reactor, u is the axial velocity of the fluid and D_a is the axial dispersion coefficient. The latter characterizes the degree of back-mixing during flow, and in the present case will depend on the characteristics of the oscillatory flow such as amplitude and frequency of the oscillation, and more importantly here, on the geometry of the cell. Since the flow features are kept the same throughout all simulations, D_a can be considered here a function of geometrical parameters. It is important to notice that the equation this model provides is for σ_θ^2, which is the variance of the normalised RTD function $E(\theta)$, where θ is the normalised time t/t_m, and not for the dimensional σ_t^2.

The dimensionless term

$$\frac{D_a}{ul_m} = \frac{movement\ by\ axial\ dispersion}{movement\ by\ bulk\ flow} \tag{87}$$

named the vessel dispersion number, provides a quantification of the extent of axial dispersion. Ideal plug flow would have a null vessel dispersion number (negligible axial dispersion compared to bulk flow), while reactors that deviate from plug flow will present higher vessel dispersion numbers as they diverge more from plug flow (as the axial dispersion and back-mixing become more important in comparison to the bulk flow). The inverse of this parameter is often wrongly referred to as the Péclet number in the literature, as the axial dispersion coefficient used in the vessel dispersion number is different from the molecular diffusion coefficient used in the calculus of the Péclet number.

Knowing that

$$\sigma_t^2 = t_m^2\sigma_\theta^2 \tag{88}$$

and also that, since this is a 1D model, the mean residence time can be expressed as

$$t_m = \frac{l_m}{u} \tag{89}$$

And finally, substituting equations (86) and (89) into equation (88) yields an expression for the variance of $E(t)$:

$$\sigma_t^2 = 2\frac{D_a l_m}{u^3} + 8\frac{D_a^2}{u^4} \tag{90}$$

Two issues arise from the use of the previous expressions in the context of this 2D study. As the main aim is to compare the RTD's variance for different geometries, according to the definitions of the variables for the 1D model, u will vary depending on the geometry assessed (since the volumetric flow rate stays the same for all cases, but the widths of the cells and open spaces between the baffles are different). Furthermore, comparing the geometries in terms of length apart from the initial position of the tracer in the longitudinal direction of the reactor (l_m) is not helpful in this case since the number of cells covered by l_m will be different depending on the input geometry parameters. It will yield information about which configuration of cells is better for a given length of the reactor (regardless of the volume), while the objective of this study is to find the best distribution of the volume in a cell, and it is required then that the same volumes are compared. In other words, the aim here is to obtain an expression of $\sigma_t(A)$, where $A = NA_{uc}$ so different geometries can be compared for the same number of cells.

Paying attention to Figure 68 again, it is immediately possible to calculate the unit cell area A_{uc}:

$$A_{uc} = (L-b)D - 4R^2\left(1 - {}^{\pi}/_{4}\right) + be \tag{91}$$

In order to find the mean axial velocity across the device for a given geometry, the area of the unit cell is written in the form of a rectangle without baffles that encompasses the same area than a reactor's cell A_{uc}, for which the cross-section can be easily calculated, and this value used to find the mean u with which the flow advances through the cells.

$$A_{uc} = LD^* \tag{92}$$

$$u = \frac{Q_A}{D^*} = \frac{Q_A L}{A_{uc}} \tag{93}$$

As can be appreciated in eq. (93), u depends on the geometry of the cell, more specifically on the cell length, L, as the volumetric flow rate Q_A and the unit area A_{uc} remain the same for all simulations. The next step consists in finding an expression for l_m in terms of the geometry parameters of the cell

$$A(l_m) = \frac{l_m}{L} A_{uc} \Rightarrow l_m = \frac{LA}{A_{uc}} \tag{94}$$

where $A(l_m)$ is the area of the reactor encompassed in a longitudinal distance l_m. This expression holds when $N = l_m/L$ is a natural number ($N \in \mathbb{N}$), which is the only case of interest in this study.

Substituting equations (93) and (94) in eq. (86):

$$\sigma_\theta^2 = 2\frac{D_a A_{uc}^2}{Q_A L^2 A} + 8\left(\frac{D_a A_{uc}^2}{Q_A L^2 A}\right)^2 \tag{95}$$

Moreover, knowing that, in this case, the equivalent of the 1D mean time in 2D can be expressed as

$$t_m = \frac{l_m}{u} = \frac{A}{Q_A} \tag{96}$$

Then, following eq. (90)

$$\sigma_t^2 = 2\frac{D_a A_{uc}^2}{Q_A^3 L^2} A + 8\frac{D_a^2 A_{uc}^4}{Q_A^4 L^4} \tag{97}$$

the equation for σ_t^2as a function of *A* is obtained. Finally, for simplicity

$$F = \frac{D_a A_{uc}^2}{L^2}; \qquad \sigma_t^2 = 2\frac{F}{Q_A^3} A + 8\frac{F^2}{Q_A^4} \tag{98}$$

As has been previously discussed, and since D_a only depends on geometrical parameters, the same is true for *F*. Taking into consideration that the volumetric flow rate Q_A is kept the same throughout all simulations, it is simple to see that the geometries that yield higher *F* values will count with both a higher slope and y-intercept term for the σ_t^2 equation. This means that the geometry that yields a lower σ_t^2 for a concatenation of 10 cells will also outperform the rest when a different number of cells is considered. This ensures that, if the mCOBR fits the model, the ranking of the variances of the RTD will stay the same regardless of the number of cells taken into account to obtain the RTD, as both the y-intercept term and the slope of the line grow with *F*. Evaluating σ_t^2 for the tracer study at the different lengths and taking into account that

$$N = \frac{A}{A_{uc}} \tag{99}$$

it is possible to find a linear relationship between *N* and σ_t^2. Processing the RTD curves for one of the geometries that will be assessed for the current study (*L/D=1.375*, *e/D=0.3* and *a/L=0.25*) in the manner described herein, and obtaining their variances for different numbers of cells, a fit is found for the model with *R^2=0.9992*. This was achieved using Matlab curve-fitting tool and including the expressions of the models as a custom equation, as Figure 87 shows. This demonstrates that our family of reactors fits the 1D model proposed and that the best performing geometry for a given length of the reactor will also be the best when compared with the other geometries for a different number of cells. In virtue of eq. (88), the latter is also valid for σ_t^2 too, which will be the parameter employed in the optimisation of the reactor, and from now on typed solely σ^2 to avoid confusions with the closeness of fit parameter of the modelling algorithm.

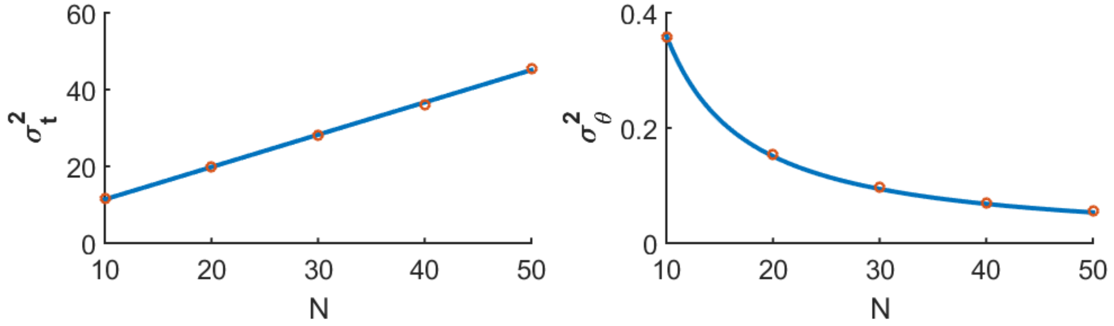

Figure 87. Regressions for the 1D dispersion model showing excellent agreement. Red unfilled dots: results for the experiment carried out for a different number of cells (10, 20, 30, 40 and 50). Blue lines: model regression.

This means that not only the ranking of different geometries has been ensured regardless of the number of cells considered, but also, taking into account the excellent agreement with the model, that the variance for reactors with a higher number of cells can be calculated with a high degree of accuracy just carrying out two CFD simulations to obtain the RTD for different lengths of the reactor, exploiting the linear relationship between *N* and σ_t^2.

For the case assessed in Figure 87, even in the most adverse situation for which D_a/uL is higher (*N=10*) and matching the variance to the theory, $D_a/ul = 0.1214$, which meets the criterion of applicability of the model being well within the limits suggested by Levenspiel (D_a/uL < 1). As the number of cells increases, D_a/uL decreases, tending towards a plug-flow behaviour that would meet the assumption for small extents of dispersion ($D_a/ul < 0.01$).

Within this framework, it is now possible to design a reactor with a narrow RTD which ensures fluid elements entering the reactor spend similar times within it, minimising the presence of un-reacted materials and unwanted by-products in the reactor's output.

4.4 Effect of Baffle Offsetting

Once the numerical scheme was validated and implemented, simulations were carried out to exemplify and compare the flow behaviour in a standard COBR and in a COBR with offset baffles. In order to compare the effect of baffle offsetting, *L/D* and *e/D* were kept the same, with only *a/L* changing to include a certain amount of offset. This permitted comparison of the flow features between both configurations.

The fluid mechanics of oscillatory flow mixing have been investigated by numerous authors in previous studies, with the development of eddies induced by oscillatory flows in different geometries being first assessed by Mackley *et al.*,[304] before they elaborated on the flow patterns originated by sharp edges.[305] Figure 88 shows the flow features at 6 different instants within a cycle for the classic COBR design with opposing baffles facing each other, and under the flow conditions described in the previous section of this chapter. Due to the presence of the baffles, the accelerating flow between phases 0 and $\pi/2$ creates eddies that are allowed to expand into the centre of the cell as the flow decelerates between the $\pi/2$ and π phases, as depicted in Figure 88 (a). When the flow reverses, it is forced to flow around these eddies, creating two flow channels between them and the boundaries of the geometry. These free vortices detached from the walls of the cell are later swept one against each other (Figure 88 (b), (c) and (d)) coinciding with the generation of new eddies, this time in the opposite side of the baffles. These new vortices expand (Figure 88 (d), (e)) and are allowed to interact with the weakened previous ones (Figure 88 (f)) generated during the first half of the cycle. Finally, as the flow starts accelerating in the positive direction, all features are unravelled into the main bulk flow, arriving again at the situation depicted in Figure 88 (a).

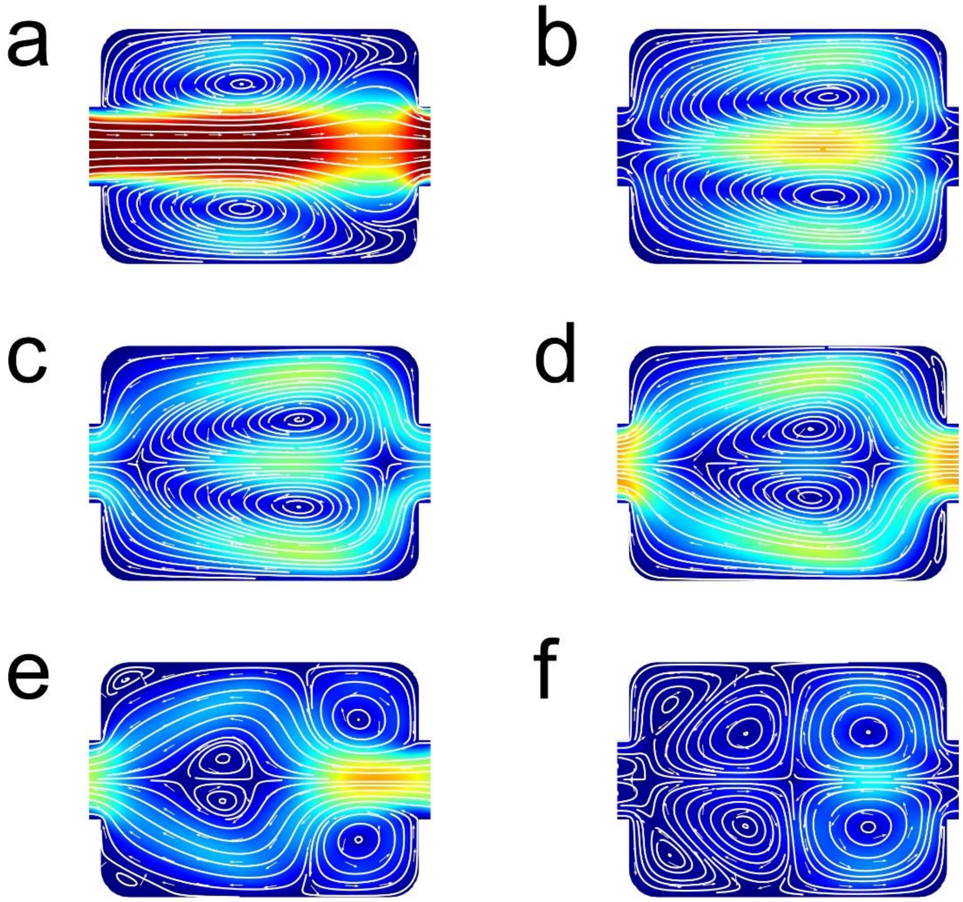

Figure 88. Velocity magnitude colour map, with streamlines and arrows depicting flow behaviour in a cell with (a/L=0) for six different oscillation phases. a) Φ = 3π/4 b) Φ = 11π/10 c) Φ = 6π/5 d) Φ = 27π/20 e) Φ = 7π/4 f) Φ = 48π/25.

With baffle offsetting, the flow dynamics differ significantly from this classic symmetric case. When the upstroke begins, the fluid accelerates and vortex formation starts on the downstream sides of the baffles. Unlike the case in which there is no offset between the baffles, one of the vortices becomes dominant as a result of the angle that the bulk flow forms with respect to the reactor's longitudinal axis when entering the cell, as shown in Figure 89 (a). This dominant vortex occupies almost 75% of the volume of the cell as the flow begins to decelerate, creating a strong recirculation, while the other vortex spreads slightly between adjacent cells (Figure 89 (b)). When flow inversion is initiated in the downstroke, the flow separates this dominant vortex from the wall pushing it towards the opposite side of the cell (Figure 89 (c)). As this vortex is pushed into the centre of the cell the reversed flow is forced to meander between the two vortices generated in the upstroke until eventually the minor vortex unravels into the main bulk flow (Figure 89 (d)). In this fashion, the dominant vortex has been formed and grown in one side of the reactor and then displaced towards the opposite side of the reactor. After the downstroke achieves its maximum velocity and starts to decelerate, two vortices are formed on the upstream side of the baffles. The lower vortex (shown in Figure 89 (e)) grows to create the dominant recirculation occupying the length of the cell; this is largely a mirror image of the vortex formation during the upstroke although slightly reduced in strength due to overall net flow through the reactor. The degree of offsetting has an influence on the formation of this vortex. If *a/L* is small, the dominant vortex is stronger than the vortex just generated by the backwards bulk flow upstream of the baffle. As *a/L* grows, the new vortex becomes stronger than the previous one. This determines the position of the centre of the new vortex, as they spin in the same direction and the weaker one disappears. As depicted in Figure 89 (f), as the backstroke continues decelerating, these vortices expand and occupy the area of the cell completely.

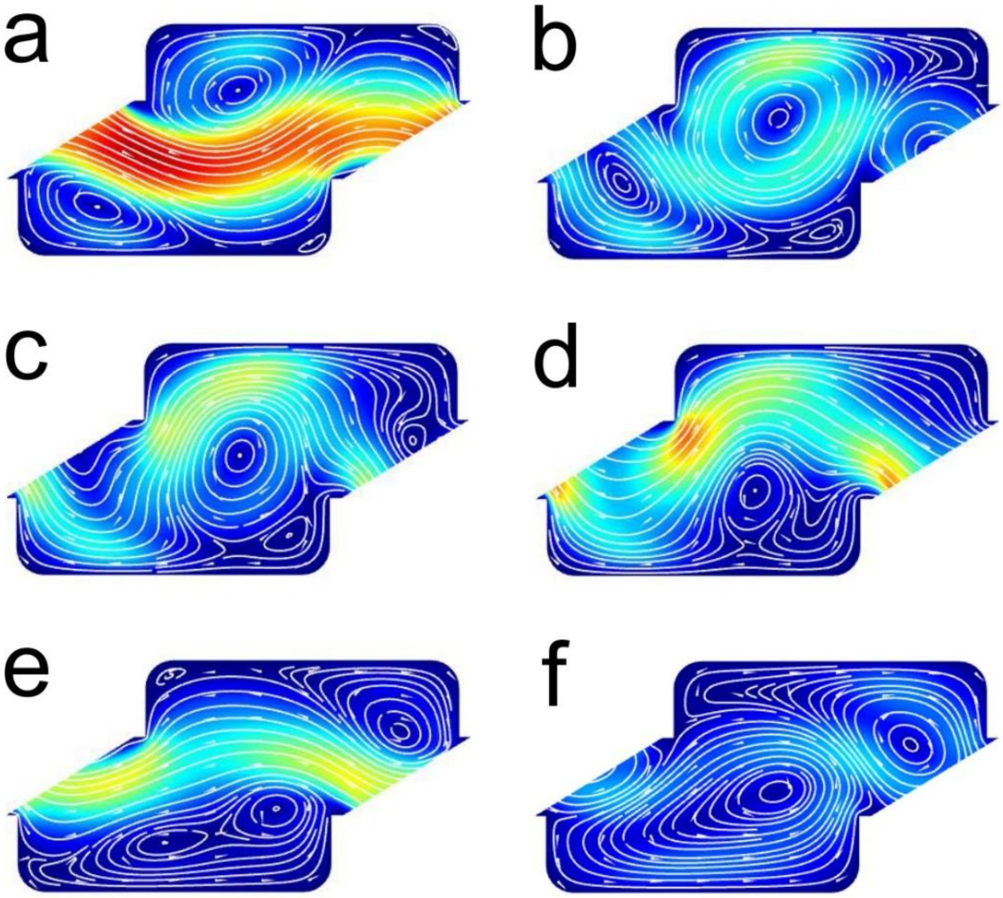

Figure 89. Velocity magnitude colour map, with streamlines and arrows depicting flow behaviour in a cell (*a/L=0.3794*) for six different oscillation phases. a) $\Phi = 3\pi/4$ b) $\Phi = 11\pi/10$ c) $\Phi = 6\pi/5$ d) $\Phi = 27\pi/20$ e) $\Phi = 7\pi/4$ f) $\Phi = 48\pi/25$.

To explore further the influence of baffle offsetting, two new simulations apart from the ones used to illustrate flow behaviour were carried out for *a/L=0.15*, and *a/L=0.5* whilst keeping the rest of the parameters the same as for the previous simulations. An additional value of *a/L=0.3794* was included as it will be demonstrated in section 4.5 that this is an optimum in terms of mixing performance. A series of parameters were evaluated to compare these three geometries quantitatively. The first two, instantaneous axial and transverse absolute averaged velocities, are defined as

$$U(t) = \frac{1}{A_{unit\ cell}} \iint |u(t)|\ dA \tag{100}$$

$$V(t) = \frac{1}{A_{unit\ cell}} \iint |v(t)|\ dA \tag{101}$$

respectively. They are the integrated absolute values of the axial and radial velocities in a reactor's cell, divided by its area. From these values, the velocity ratio can also be easily calculated as

$$R_v(t) = {U(t)}\big/{V(t)} \tag{102}$$

Of course, these quantities are time-dependent, so their mean and maximum values for 20 phases within an oscillation cycle were computed, and results presented in Table 18.

Table 18. Mean and maximum values for U, V and 1/R$_v$.

a/L		U	V	1/R_v
0	Mean	2.165×10^{-2}	0.616×10^{-2}	0.338
	Maximum	3.532×10^{-2}	0.805×10^{-2}	0.707
0.15	Mean	2.147×10^{-2}	0.897×10^{-2}	0.462
	Maximum	3.415×10^{-2}	1.099×10^{-2}	0.720
0.3794	Mean	1.989×10^{-2}	0.921×10^{-2}	0.513
	Maximum	3.306×10^{-2}	1.162×10^{-2}	0.935
0.5	Mean	1.965×10^{-2}	0.920×10^{-2}	0.523
	Maximum	3.299×10^{-2}	1.171×10^{-2}	0.983

From the results in Table 18, it can be observed how the values correlate to the simulation results. When no offset is present in the design, the lowest values for both V_{mean} and V_{max} are obtained, also coinciding with the highest U_{mean}

and U_{max}. This is an indicator of this geometry is suboptimal when creating transverse motion and mixing, especially when realising that it is this configuration the one that presents the lowest $1/R_v$, a parameter understood to characterise mixing in COBRs.[285] The values obtained for $1/R_v$ are similar to those reported in the literature,[127, 285, 286] falling well within the restriction proposed by Fitch *et al.*[127] that states that R_v <3.5 (or $1/R_v$ >0.285) is needed to ensure effective mixing in a COBR. In this regard, this geometry is clearly outperformed by the rest. V_{mean} is maximised for the case *a/L=0.3794*, making this design the best in terms of transversal mobility of the particles. However, it should be noted that this fact does not ensure that this particular geometry is the best in terms of RTD, since U_{mean} (which is correlated to the axial dispersion) has not been minimised in this case, being significantly lower when *a/L=0.5*, a geometry which also maximises $1/R_v$.

Summarising, even when the case *a/L=0.3794* presents higher axial velocities (a disadvantage in terms of RTD) and a lower $1/R_v$ than *a/L=0.5,* it maximises *V*, highlighting the trade-off between good mixing and low axial dispersion (good RTD) that will be explored in the next section of this chapter.

4.5 Optimisation Methodology

This section explores how COBR design can be optimised to maximise the mixing efficiency η and minimise the variance σ^2 of the RTD. The methodology will be demonstrated on the multi-objective optimisation problem defined in (103):

$$\begin{aligned} &\boldsymbol{minimise}: && \frac{1}{\eta}\left(\frac{L}{D},\frac{e}{D},\frac{a}{L}\right), \qquad \sigma^2\left(\frac{L}{D},\frac{e}{D},\frac{a}{L}\right) \\ &\boldsymbol{subject\ to}: && 1<\frac{L}{D}<1.75, \qquad 0.3<\frac{e}{D}<0.6, \qquad 0<\frac{a}{L}<0.5 \end{aligned} \tag{103}$$

4.5.1 Surrogate Modelling

The recent review by Haftka *et al*.[306] noted that for multi-objective problems with less than 100 design variables an effective approach is to use a gradient-free method where surrogate models of the system responses are used to create Pareto fronts which quantify the available compromises between competing objectives. In the present study, a 3^3 full factorial Design of Experiments (DoE) with *n=27* points is used to generate the CFD data for surrogate models of η and σ^2 throughout the design space. A nested 2^3 DoE is kept as a holdout dataset at the points maximising the distance to the previous DoE points for cross-validation of the surrogate models. Figure 90 shows the DoEs employed for surrogate modelling, where the blue dots represent the 3^3 full factorial DoE and the red dots represent the 2^3 full factorial that will be used as a holdout dataset.

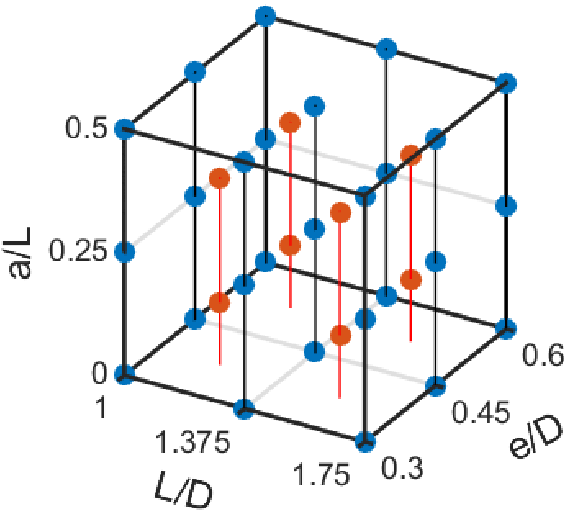

Figure 90. 3D representation of the design space, where blue points represent the 3^3 full factorial DoE used for the study and the red dots represent the holdout dataset.

The Moving Least Squares (MLS) method, already introduced in section 3.2.1 of this thesis, was used along with LOO and MCCV, and a golden search algorithm for model tuning. For MCCV, a random sub-sampling routine was implemented to carry out 2000 random sampling operations while ensuring that each DoE point has been removed the same number of times. The validation algorithm stores the RMSE over the k points left out at every iteration and calculates the mean RMSE over all 2000 iterations. A golden search algorithm is then used to identify the θ value that minimises the mean RMSE. The results of these cross-validations are summarised in Table 19.

Table 19 contains the following information. *CV RMSE* is the average RMSE over all $k_{\text{left out}}$ points that are removed from the full training set of n=27 DoE points, obtained using a surrogate model trained on the remaining (n- $k_{\text{left out}}$) DoE points. This value is minimised by the golden search algorithm that finds θ. *Training RMSE* represents the RMSE calculated over all 27 DoE points using the surrogate model already tuned and *Test RMSE* refers to the RMSE calculated over the nested DoE points held out during the LOO and MC cross-validation procedures. This means that the tuning algorithm does not have access to the holdout dataset and hence, *Test RMSE* is indicative of how the tuned model will perform for uncalibrated data. In each case, θ represents the optimal value that leads to the smallest associated *CV RMSE.*

Table 19. Cross-validation results. Sample size: 27 points.

	k **left out**	**θ**	**CV RMSE**	**Training RMSE**	**CV σ²**	**Test RMSE**
η	1 (LOO)	10	0.0533	0.0013	29.3405×10^{-4}	0.0683
	3	10	0.0535	0.0013	11.2370×10^{-4}	0.0683
	5	2.07	0.0559	0.0207	7.4275×10^{-4}	0.0627
	7	1.34	0.0589	0.0266	6.9889×10^{-4}	0.0606
	9	0.90	0.0646	0.0310	7.5544×10^{-4}	0.0593
		0		0.0419		0.05651
σ²	1 (LOO)	3.40	0.0277	0.0073	7.9538×10^{-4}	0.0171
	3	2.76	0.0273	0.0090	2.9719×10^{-4}	0.0175
	5	2.58	0.0283	0.0096	2.0285×10^{-4}	0.0176
	7	2.13	0.0318	0.0111	2.1892×10^{-4}	0.0180
	9	1.69	0.0367	0.0129	2.4812×10^{-4}	0.0185
		10		0.0006		0.0164

For the η surrogate model, Table 19 shows how, as $k_{\text{left out}}$ increases, both *CV RMSE* and *Training RMSE* also increase. This loss in precision is a result of the surrogate model being computed with a lower number of points, taking into account less information to generate its predictions. Conversely, *Test RMSE* decreases as the number of points left out grows. In other words, the model becomes better at completely unseen data as θ decreases and the model generalises. This is an indicator of the model overfitting the training set when only a small number of points are left out.[307, 308] Even when $k_{\text{left out}} = 9$, a substantial difference between *Test RMSE* and *Training RMSE* values indicates that the model is still over-fitting the data. In order to prevent this behaviour, and following the trend in decreasing values of θ, a surrogate model for $\theta=0$, which corresponds to the conventional least squares method, was computed. This yielded the lowest value of the *test RMSE* and also the minimum difference between this and the *training RMSE*. This result indicates either that a different choice of the polynomial base should be used for the metamodel or that including position dependent regression weights in this case is counter-productive. In consequence, the surrogate model chosen for η will be the one obtained when $\theta=0$.

Similar behaviour is seen for the σ^2 surrogate model, for which the cross-validation methods find decreasing values of theta as $k_{\text{left out}}$ grows. Apart from the cross-validated surrogates, an extra one is computed for $\theta=10$, producing the lowest values for both training and test RMSE. However, as discussed before, the enormous gap between these values suggests a severe over-fitting of the training data. Following ref.[259] the most appropriate $k_{\text{left out}}$ is chosen which minimises the ratio of the *Training RMSE* to the *Test RMSE*, so the surrogate model for σ^2 used in the following optimisation study use MCCV with $k_{left\ out}=9$ and $\theta=1.6898$.

4.5.2 Multi-Objective Optimisation

The multi-objective optimisation problem (103) is solved using Matlab to generate the Pareto front of non-dominated solutions shown in Figure 91. The genetic algorithm function *gamultiobj*, based on the algorithm NSGA-II [189] and included in Matlab's global optimisation toolbox is used to find the Pareto front. A custom function was developed that accepts as input a set of design variables and returns the values of $-\eta$ and σ^2 (as the algorithm aims to minimise both objectives, the sign of η is changed so it can be maximised while σ^2 is minimised). The function loads all options and parameters needed by the MLSM algorithm for the two different models selected and stores them as different structures whose fields can be accessed to provide the needed responses. The *gamultiobj* algorithm then makes use of these surrogate models to find the Pareto front.

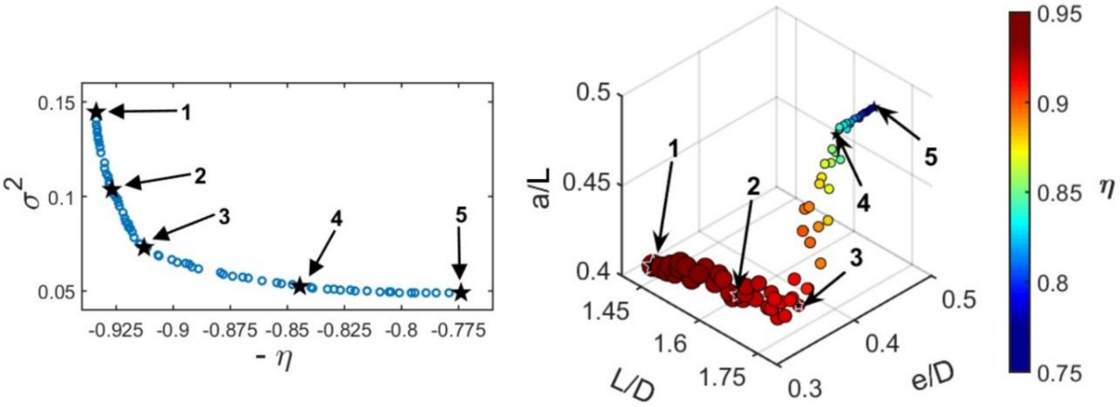

Figure 91. Pareto front (left) and its points represented in the design space (right) where the colour indicates the value of η and the area of the circles is proportional to σ^2. The validation points evaluated are indicated as black stars.

CFD results were obtained at five points along the Pareto front (marked as black stars in Figure 91) in order to provide a new test dataset for further validation of the surrogate models. These surrogate models will be computed carrying out MCCV over the 35 points resulting from the combination of the 27 points of the full factorial and 8 points of the nested holdout dataset with up to 11 DoE points left out to create the surrogate models and the results are shown in Table 20.

Table 20. Cross-validation results. Sample size: 35 points.

	$k_{\text{left out}}$	θ	**CV RMSE**	**Training RMSE**	**CV σ^2**	**Pareto front RMSE**
η	1 (LOO)	8.03	0.0575	0.0073	33.700×10^{-4}	0.0141
	3	3.58	0.0562	0.0244	11.751×10^{-4}	0.0321
	5	2.58	0.0548	0.0290	7.081×10^{-4}	0.0359
	7	1.30	0.0561	0.0359	5.510×10^{-4}	0.0396
	9	0.82	0.0569	0.0390	4.997×10^{-4}	0.0407
	11	0.70	0.0588	0.0399	4.975×10^{-4}	0.0410
σ^2	1 (LOO)	3.56	0.0243	0.0084	6.085×10^{-4}	0.0081
	3	3.41	0.0226	0.0087	1.991×10^{-4}	0.0081
	5	3.12	0.0234	0.0094	1.348×10^{-4}	0.0080
	7	3.01	0.0237	0.0097	1.061×10^{-4}	0.0080
	9	2.75	0.0255	0.0103	1.095×10^{-4}	0.0079

The updated surrogate model for η is obtained by using the θ value from cross-validation with 11 points removed, where training and test set errors are similar, whereas for the σ^2 model, LOO is used as this provides an appropriate balance between the training and test errors, in addition to representing the lowest value for the training error across all the cross-validation procedures. Figure 92 provides a visualisation of these updated surrogate models using isosurfaces.

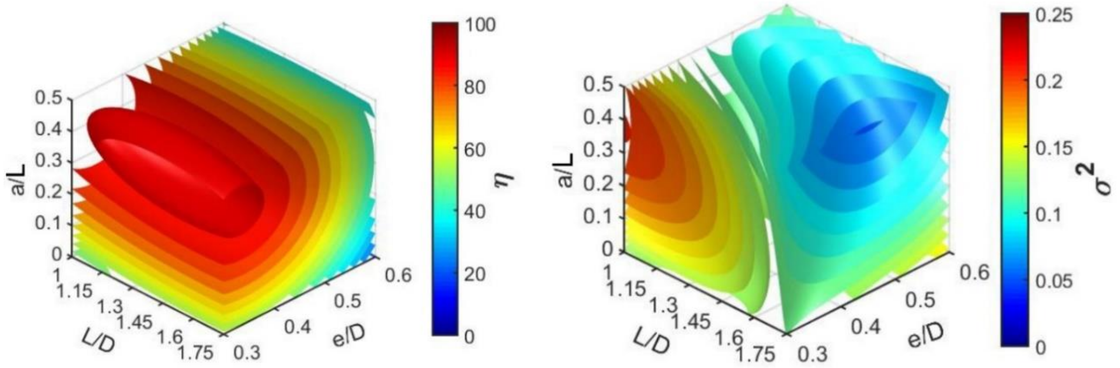

Figure 92. Isosurfaces representations of the models for η (left) and σ^2 (right).

For η, the region of interest comprising values higher than 0.9 is in an elongated ellipsoid in the *L/D* direction and the *L/D* parameter is the least influential design variable. There is clearly a maximum corresponding to the open area variable *e/D* ≈ 0.34 and *a/L* $\approx$ 1.37. More generally e/D and a/L have a significant influence on η. The surrogate model for σ^2 is more complex. Its lowest values are obtained for the maximum values of *a/D* and *L/D,* and for *e/D*≈0.45. This means that an elongated shape of the cell, as well as offsetting the baffles to the maximum, contributes to minimising axial dispersion.

An updated Pareto front, using the updated surrogate models for η and σ^2 which use CFD data from the nested hold out DoE in addition to the 27 points full factorial, is presented in Figure 93.

Table 21. Final solution.

x_i			f_i		$f(x_i)$		Absolute error		Relative error (%)	
L/D	**e/D**	**a/L**	**η**	**σ^2**	**η**	**σ^2**	**η**	**σ^2**	**η**	**σ^2**
1.452	0.338	0.379	0.872	0.13	0.927	0.136	0.0605	0.0065	6.98	5
1.749	0.350	0.389	0.867	0.081	0.912	0.076	0.0402	0.0054	4.61	6.61
1.750	0.450	0.500	0.741	0.057	0.774	0.049	0.0334	0.0084	4.51	14.6

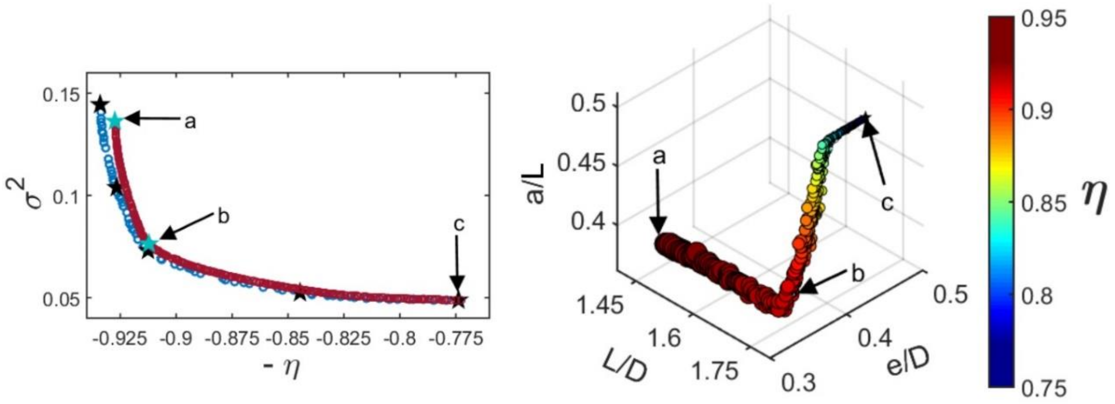

Figure 93. Revised Pareto front (left) and its representation in the design space (right) for the second pair of models used.

When the surrogate models are updated with more CFD data points, the main difference is that the new Pareto front predicts lower values for the mixing efficiency index, η. To assess the accuracy of the updated Pareto front, three validation points were selected along it. The first point was kept the same as the point selected for a minimum σ^2 in the previous front while the other two, shown as cyan blue stars in Figure 93, correspond to the position of highest curvature on the front and the point on the front predicting the highest η value.

Table 21 presents the results for these points. The geometry configuration calculated as to maximise the curvature of the front, yields values of 0.8721 for η, and 0.0817 for σ^2; hence providing a good trade-off between η and σ^2. In this case, the relative errors between the predicted results and the simulation results were of 4.6 and 6.6%, respectively. The highest error of 14.6%, obtained for the minimum σ^2 location, suggests the need for a bigger initial DoE or of an iterative approach generating new DoEs in the Pareto Optimality region of the design space. More precise models in this region can be computed, minimising the error over the hold out dataset (points for the previous Pareto front), but in order to come up with a reliable Pareto front, the shape of the metamodels needs to be globally correct and not just accurate in a particular region of the design space, even if that region is where the optimum could be. It is critical then to make sure that the models do not overfit the data. A good approach (and one exemplified here) to ensure this is to assess different cross-validation techniques and compare their relative performances. It can be concluded that the surrogate models created for both variables were reliable and a valid tool to perform a multi-objective genetic algorithm optimisation to ensure a good trade-off between the design objectives for the experimental plate mCOBR.

4.6 Summary

The first multi-objective optimisation strategy for COBRs is proposed in this chapter. The technique has been applied to the development of a small scale COBR but would be equally suited to the optimisation of more conventionally sized COBRs. Results presented here show the importance of analysing the temporal and spatial periodic flow field in order to identify the times and locations at which representative flow field data can be retrieved from CFD analyses. Performance-based metrics based on mixing efficiency index and the standard deviation of the residence time distribution are proposed and it is shown that, when creating surrogate models of these metrics, careful consideration of the balance between training and test data during the cross-validation stage is crucial in the determination of reliable optimal solution sets.

The surrogate models provide valuable insight into the influence of key design variables. For example, in the case of the mixing efficiency, the parameters *e/L* and *a/L* have the most significant effect over the performance of the device, achieving the maximum η values when $e/L \approx 0.34$ and $a/L \approx 0.38$, regardless of *L/D*. This configuration is proven to maximise the transversal mobility of the particles, as the new baffle offset feature induces asymmetries in the flow. The residence time distribution tends to be narrower for high values of *L/D*. This is explained because of the smaller width of the cells, reducing the contact area between axially adjacent volumes of fluid. The parameter a/L reveals itself crucial, as $a/L = 0.5$ makes for the narrowest RTD. This situation could be understood as a reactor with twice as many cells, separated by means of alternating baffles.

These surrogate models can be used within a multi-objective optimisation process to create Pareto fronts from which design decisions can be made, exploiting the trade-off between the competing objectives. Our proposed configuration (point b in Figure 93) improves the degree of mixing (η) at relatively little increase in the residence time distribution variance (σ^2) and outperforms any configuration that does not include baffle offsetting for both conflicting objectives.

The methodology introduced here can be extended to incorporate additional design variables and three-dimensional flow analyses, in which the challenge will be to constrain the additional computational costs within feasible timescales. Further work also includes the application of this optimisation to the wide range of other chemical reactor systems used during process development.

Finally, the work highlights how the vortex development induced by the inclusion of baffle offsetting is significantly different from the standard design in which no offset is included. The reactor is described paying attention to different flow features, and the effect of baffle offsetting quantified by means of the axial and transversal absolute averaged velocities.

Chapter 5
Conclusions and future work

5.1 Project outcomes

Lab-scale millifluidic reactors have been proven a useful tool for the automation of chemical experiments. However, these were constrained to single phase reactions, as generally their size limited the degree of mixing achievable. This thesis demonstrated the use of the fReactor, a miniature CSTR device enabling multiphasic reaction experimentation. In the same regard, and aiming to provide better RTDs, CFD was used in order to optimise the geometry of an mCOBR. Underpinning both of these studies is the use of optimisation methodologies. The outcomes of these studies will be discussed in comparison to the research gap identified in section 1.6, and the research aim outlined in section 1.7.1:

> *"The aim of this project is to develop flow reactors for automated exploration of reaction space. The work will focus on two platforms: (i) the fReactor CSTR and (ii) a continuous oscillatory baffled reactor. A key aim is to underpin both with optimisation techniques. In the case of the fReactor, such techniques will focus around the single and multi-objective optimisation of the output of multiphasic reactions; in the case of the COBR, they will be used to optimise the design of the device paying attention to different flow features".*

In addition, let us recall here the objectives of the research (originally introduced in section 1.7.2) for convenience:

> *"a) The design and development of a temperature-controlled version of the fReactor.*
>
> *b) Integration of said equipment into the automated rig at iPRD in order to enable automatic data gathering for multiphasic reactions.*
>
> *c) Analysis of the data obtained by the use of such equipment employing a variety of different metamodeling and cross-validation techniques, aiming to outperform the widely used least squares polynomial fitting.*
>
> *d) Investigation and quantification of the performance of a plate miniaturised continuous oscillatory baffled reactor via computational fluid dynamics.*
>
> *e) Computational fluid dynamics – enabled design optimisation of the plate mCOBR via surrogate modelling".*

The first and second objectives were achieved by the application of well-established engineering principles. The material of the fReactor body was selected meeting the safety requirements of the heating application. Apart from this, the hardware supporting the application was completely engineered to address its needs from its very conception, including the selection of sensors, actuators and high-power electronic elements. A PCB was designed providing appropriate connections for all the elements that integrated debugging components such as pull-down resistors and decoupling capacitors, resulting in a reliable and robust piece of equipment. Finally, programming of the microprocessor made use of data acquisition and signal filtering techniques, as well as providing a serial communication capability that allowed communication with the central computer of the laboratory. PI control was achieved by using an existent Arduino library, tuning it under different flow rate conditions and providing a gain scheduling that enables precise temperature control for flow rates up to 5 mL/min. A quick and stable temperature control was achieved, with an accuracy of ±0.3 °C. When compared with the Polar Bear tubular reactor for a specific temperature change, the control resulted 44% quicker, which translates into cutting waiting times and reagent consumption. In addition, extra hardware was developed for such system that enabled UV multiphasic chemistry [238] as well as robust stirring commercially available stirrers were not able to provide.

This device was successfully employed then in order to broaden the field of use of automated flow chemistry to multiphasic reactions. To exemplify this capability, a liquid-liquid multiphasic (aqueous/organic) Claisen-Schmidt condensation was carried out using the fReactors in an automated way under different conditions, according to a DoE sampling plan, where the products were separated via a liquid-liquid separator before injection into HPLC analysis.

In relation to objective c), metamodeling techniques were then successfully employed to generate response surface models for different metrics calculated from the HPLC results, where MLSM was used along with LOO and Monte Carlo cross-validation procedures. These were compared and different response surface models computed from the data according to the most appropriate selected models, that were then used to conduct both single and multi-objective optimisations of up to three competing objectives. Furthermore, an in-silico modelling and optimisation problem was carried out to demonstrate RBFs in addition to MLSM, for different basis functions and, again, different cross-validation approaches. It was demonstrated here that the use of the LOO and Monte Carlo cross-validation routines when compared to the frequently used least squares polynomial fitting leads to better surrogate models. A common misunderstanding when using response surfaces is to use the coefficient of determination R^2 as the indicator of the performance of a response surface. While this could be appropriate in some situations, in many cases this approach leads to models that greatly overfit the data, resulting in response surfaces with poor generalisation capability: oftentimes only able to predict responses with acceptable accuracy in the near surroundings of the data points. For example, using MLSM and LOO reduced the error by almost 5% when compared to a second order polynomial least squares fitting.

In order to further enhance the performance of the surrogate models employed, algorithms were developed for tuning the parameters of the models, that enabled identification of the radial basis function that minimised cross-validation RMSE for a LOO cross-validation with minimal user intervention. This metamodeling effort led to results that clearly outperformed least squares polynomial fitting: the error was reduced 6% with standard RBFs, and this result was surpassed when employing augmented RBFs, for which the error was reduced twice as much, slightly over 12%.

In addition to these experiments, the suitability of the fReactors was demonstrated for *in situ* XRD analysis. Using kapton tubing between fReactor modules enabled analysis of polymorph change due to solvent switching in the obtention of calcium sulphate. This would enable *in situ* assessment of multiphasic crystallisation processes in flow.

Finally, a CFD-enabled optimisation of the geometry of an mCOBR was carried out. A different kind of reactor to the CSTR, it allows the use of a greater number of cells, leading to much narrower residence time distributions than the fReactors in sensible cascade designs of up to five modules. The same metamodeling and cross-validation techniques employed for the achievement of the third objective were used in order to obtain the response surface models of the mixing efficiency index and the residence time distribution for the design, aiming to reach a compromise solution between the two. Insights about the inclusion of baffle offsetting in such plate configuration were given, with detailed description of the flows in a generic cell under different geometrical configurations, including quantification of the axial and transversal absolute averaged velocities in the cells. It was ascertained then that such plate COBR designs could benefit both in terms of maximising transverse mixing and minimising the variance of the RTD by including a baffle offset feature. In the case of RTD, the optimal design found outperformed the best performing opposing baffles design by reducing the RTD variance by 45%, while in the case of the mixing efficiency index it could be increased by 60% when selecting the optimal design.

A list with the key results and achievements of this work is provided here for clarity:

- A novel, temperature-controlled version of a miniaturised CSTR was developed that was able to cut the waiting times associating with heating by almost 45%, as well as increasing the resolution to ±0.3 ºC when compared to the commercially available Polar Bear Plus on which the setup relied on previously.
- In addition, the use of the fReactor active enabled multiphasic processing, as was demonstrated by carrying out an automated experiments for a Claisen-Schmidt reaction. This overcame one of the main limitations of the devices used previously.
- Multiple modelling techniques were used that outperformed the conventional LS fitting of polynomials, leading to reductions in error of up to 14% when using the same scattered data. These approaches relied in cross-validation based model tuning.
- Machine learning algorithms were developed to automatically compare the performance of different RBFs, for both their standard and augmented versions, minimising user intervention.
- The surrogate models calculated were then used in both simple and multi-objective optimisation in order to balance different conflicting objectives.
- The same techniques, coupled with CFD, were employed in the design optimisation of a plate mCOBR, for which the geometry of its cells was optimised parametrically.
- A new feature, namely baffle offsetting, was introduced in this study. The optimisation including this feature led to an outstanding 45% reduction in RTD variance and a 60% increase of the mixing efficiency index when compared to the best performing geometry without this feature (with opposing baffles).

In view of these, the impact of this thesis is twofold. Firstly, a new and responsive temperature controlled device enabling multiphasic processing is available now for chemists to carry out automated studies. The use of this reactor is very promising for process intensification of multiphasic reactions, where it presents clear advantages when compared to its tubular counterparts, maximising areas of contact between different fluids and mass transfer. In addition, the insert-based mCOBR design presented here leading to plug flow like behaviour brings with it a new degree of control over the residence time distribution, and such easy to manufacture and operate device could also be integrated into an automated rig, avoiding the broad RTD achieved with the fReactor while at the same time being able to process higher mass content and multiphasic products.

5.2 Recommended further research

The recommendations are presented here in three different areas: (i) metamodeling techniques in the field of automated flow chemistry experiments, and the use of (ii) fReactors and (iii) mCOBRs in the future. Process niches for the latest two are given in their corresponding sections.

5.2.1 Metamodeling techniques application

Introducing cross-validation would be beneficial, and generally can lead to better models that can help gaining understanding of reactions' landscapes. In the case of the frequent nonlinear responses obtained in chemistry, radial basis functions are recommended instead of polynomial models. These could also be implemented to be used iteratively in self-optimisation routines. Additionally, a Machine Learning algorithm similar to the one used in section 3.2.2 of this thesis could be implemented in order to detect the most appropriate RBF for a specific problem, without requiring user intervention. As there are two fundamental aims during a cross-validation process: minimising the objective error and also the relationship between test and training error, the algorithm can also include a capability to automatically select a model according to this bias-variance trade-off, trying to avoid the overfitting that is oftentimes incurred in by using a least squares fit with no cross-validation procedure.

5.2.2 fReactor automated experiments speed-up

Whilst the focus of this work has been on optimisation of flow chemistry, the multi-reactor arrangement of the fReactors could be used for batch chemistry based automated DoE or self-optimisation studies. With the fReactors in parallel and an appropriate system of distribution valves, reagents could be inserted into each single fReactor in different quotas and for different residence time and temperature conditions; in other words, as many experiments as fReactors present in such configuration could take place at the same time, resulting in great time savings. After the desired residence time, the resulting product would be sampled as it is flushed from the fReactor before injection of new reagents. This represents a full circle of returning to traditional batch chemistry, yet still capitalising on the advantages of flow chemistry presented in section 1.1.1. This could dramatically reduce the time for optimisation of batch processes.

5.2.3 Future mCOBR applications

The plate mCOBR design introduced in this thesis presents some characteristics that could make it useful in a variety of applications. The first one, following the trend of this thesis, would be its use in automated flow chemistry platform. Temperature control may be achieved via a controller using Peltier plates as actuators, combined with appropriate fans to dissipate the heat when using active cooling. Due to the modular nature of the controller developed for the fReactors, this is suitable for use along a mCOBR by simply providing the appropriate fittings to the actuators with relatively little work. This would include PID control tuning for the new device and the use of one of the outputs of the controller to activate the fans when required. This would lead to a much faster temperature control of the reactions, which translates in the reduction of waiting times and reagent consumption. The narrow RTD this design provides would allow to assess the output of the reactions for different mean residence times in a more accurate fashion than the fReactor used in continuous flow. However, the conditions under which this device may be used might be more restricted than those of the fReactor, especially when taking into account possible clogging problems due to the small cross-section of the mCOBR that might restrict it to lower solid mass percentages in the case of liquid-solid multiphasic reactions or other reactions in which precipitation may occur. In any case, having a range of lab-scale reactors in the lab benefits chemists, that will have to evaluate the trade-offs of each one for every specific application.

It would also be useful to explore the relationship between the mCOBR and larger COBRs, such as the mesoscale COBRs found in the research conducted by Harvey *et al.*[137] CFD analysis of such devices under different flow regimes and for different 3D baffle geometries could be used to formally optimise their performance according to different metrics. Among the geometries susceptible of being optimised in this way, and in connection with the introduction of asymmetries in the case of the mCOBR, a variable angle helicoidal baffle geometry could be assessed.

In relation with section 3.3 of this thesis, the mCOBR also poses an interesting tool for *in situ* crystallisation studies. Resembling the shape of the chip developed by Dr. Mark Levenstein,[271] and including windows that encompass the reactor's cells along with a combination of silicon and Kapton inserts, it would allow XRD analysis for multiphasic mixtures with a high degree of control over RTD that is obviously crucial for size control in crystallisation processes.

Bibliography

[1] R. L. Hartman, J. P. McMullen, and K. F. Jensen, "Deciding Whether To Go with the Flow: Evaluating the Merits of Flow Reactors for Synthesis," *Angewandte Chemie-International Edition,* vol. 50, pp. 7502-7519, 2011.

[2] T. Bergander, K. Nilsson-Välimaa, K. Öberg, and K. M. Lacki, "High-Throughput Process Development: Determination of Dynamic Binding Capacity Using Microtiter Filter Plates Filled with Chromatography Resin," *Biotechnology Progress,* vol. 24, pp. 632-639, 2008.

[3] C. Simms and J. Singh, "Rapid Process Development and Scale-Up Using A Multiple Reactor System," *Organic Process Research & Development,* vol. 4, pp. 554-562, 2000.

[4] J. R. McDonough, A. N. Phan, and A. P. Harvey, "Rapid process development using oscillatory baffled mesoreactors – A state-of-the-art review," *Chemical Engineering Journal,* vol. 265, pp. 110-121, 2015.

[5] R. Porta, M. Benaglia, and A. Puglisi, "Flow Chemistry: Recent Developments in the Synthesis of Pharmaceutical Products," *Organic Process Research & Development,* vol. 20, pp. 2-25, 2016.

[6] S. V. Ley, "On Being Green: Can Flow Chemistry Help?," *The Chemical Record,* vol. 12, pp. 378-390, 2012.

[7] N. Holmes, G. R. Akien, A. J. Blacker, R. L. Woodward, R. E. Meadows, and R. A. Bourne, "Self-optimisation of the final stage in the synthesis of EGFR kinase inhibitor AZD9291 using an automated flow reactor," *Reaction Chemistry & Engineering,* vol. 1, pp. 366-371, 2016.

[8] M. Rueping, T. Bootwicha, and E. Sugiono, "Continuous-flow catalytic asymmetric hydrogenations: Reaction optimization using FTIR inline analysis," *Beilstein Journal of Organic Chemistry,* vol. 8, pp. 300-307, 2012.

[9] S. G. Newman and K. F. Jensen, "The role of flow in green chemistry and engineering," *Green Chemistry,* vol. 15, pp. 1456-1472, 2013.

[10] J. Wegner, S. Ceylan, and A. Kirschning, "Flow Chemistry – A Key Enabling Technology for (Multistep) Organic Synthesis," *Advanced Synthesis & Catalysis,* vol. 354, pp. 17-57, 2012.

[11] V. Hessel, "Novel Process Windows – Gate to Maximizing Process Intensification via Flow Chemistry," *Chemical Engineering & Technology,* vol. 32, pp. 1655-1681, 2009.

[12] R. A. Bourne, R. A. Skilton, A. J. Parrott, D. J. Irvine, and M. Poliakoff, "Adaptive Process Optimization for Continuous Methylation of Alcohols in Supercritical Carbon Dioxide," *Organic Process Research & Development,* vol. 15, pp. 932-938, 2011.

[13] D. N. Jumbam, R. A. Skilton, A. J. Parrott, R. A. Bourne, and M. Poliakoff, "The Effect of Self-Optimisation Targets on the Methylation of Alcohols Using Dimethyl Carbonate in Supercritical CO2," *Journal of Flow Chemistry,* vol. 2, pp. 24-27, 2012.

[14] A. J. Parrott, R. A. Bourne, G. R. Akien, D. J. Irvine, and M. Poliakoff, "Self-Optimizing Continuous Reactions in Supercritical Carbon Dioxide," *Angewandte Chemie International Edition,* vol. 50, pp. 3788-3792, 2011.

[15] J. P. Knowles, L. D. Elliott, and K. I. Booker-Milburn, "Flow photochemistry: Old light through new windows," *Beilstein Journal of Organic Chemistry,* vol. 8, pp. 2025-2052, 2012.

[16] M. Oelgemöller and O. Shvydkiv, "Recent Advances in Microflow Photochemistry," *Molecules,* vol. 16, p. 7522, 2011.

[17] M. Oelgemoeller, "Highlights of Photochemical Reactions in Microflow Reactors," *Chemical Engineering & Technology,* vol. 35, pp. 1144-1152, 2012.

[18] X. Liu, B. Unal, and K. F. Jensen, "Heterogeneous catalysis with continuous flow microreactors," *Catalysis Science & Technology,* vol. 2, pp. 2134-2138, 2012.

[19] C. Jiménez-González, P. Poechlauer, Q. B. Broxterman, B.-S. Yang, D. am Ende, J. Baird*, et al.*, "Key Green Engineering Research Areas for Sustainable Manufacturing: A Perspective from Pharmaceutical and Fine Chemicals Manufacturers," *Organic Process Research & Development,* vol. 15, pp. 900-911, 2011.

[20] L. Saias, J. Autebert, L. Malaquin, and J.-L. Viovy, "Design, modeling and characterization of microfluidic architectures for high flow rate, small footprint microfluidic systems," *Lab on a Chip,* vol. 11, pp. 822-832, 2011.

[21] K. S. Elvira, X. C. I. Solvas, R. C. R. Wootton, and A. J. deMello, "The past, present and potential for microfluidic reactor technology in chemical synthesis," *Nature Chemistry,* vol. 5, pp. 905-915, 2013.

[22] R. Carlson and J. E. Carlson, "Design and Optimization in Organic Synthesis," in *Design and Optimization in Organic Synthesis*. vol. 24, ed Amsterdam: Elsevier Science Bv, 2005, pp. 1-574.

[23] P. L. Suryawanshi, S. P. Gumfekar, B. A. Bhanvase, S. H. Sonawane, and M. S. Pimplapure, "A review on microreactors: Reactor fabrication, design, and cutting-edge applications," *Chemical Engineering Science,* vol. 189, pp. 431-448, 2018.

[24] A. J. deMello, "Control and detection of chemical reactions in microfluidic systems," *Nature,* vol. 442, pp. 394-402, 2006.

[25] P. Watts and C. Wiles, "Recent advances in synthetic micro reaction technology," *Chemical Communications,* pp. 443-467, 2007.

[26] K. F. Jensen, "Microreaction Engineering—Is Small Better," *Chemical Engineering Science,* vol. 56, pp. 293-303, 2001.

[27] V. Hessel and H. Löwe, "Organic Synthesis with Microstructured Reactors," *Chemical Engineering & Technology,* vol. 28, pp. 267-284, 2005.

[28] D. M. Ratner, E. R. Murphy, M. Jhunjhunwala, D. A. Snyder, K. F. Jensen, and P. H. Seeberger, "Microreactor-based reaction optimization in organic chemistry—glycosylation as a challenge," *Chemical Communications,* pp. 578-580, 2005.

[29] S. Krishnadasan, R. J. C. Brown, A. J. deMello, and J. C. deMello, "Intelligent routes to the controlled synthesis of nanoparticles," *Lab on a Chip,* vol. 7, pp. 1434-1441, 2007.

[30] J. P. McMullen, M. T. Stone, S. L. Buchwald, and K. F. Jensen, "An Integrated Microreactor System for Self-Optimization of a Heck Reaction: From Micro- to Mesoscale Flow Systems," *Angewandte Chemie-International Edition,* vol. 49, pp. 7076-7080, 2010.

[31] J. P. McMullen and K. F. Jensen, "An Automated Microfluidic System for Online Optimization in Chemical Synthesis," *Organic Process Research & Development,* vol. 14, pp. 1169-1176, 2010.

[32] K.-J. Wu, V. Nappo, and S. Kuhn, "Hydrodynamic Study of Single- and Two-Phase Flow in an Advanced-Flow Reactor," *Industrial & Engineering Chemistry Research,* vol. 54, pp. 7554-7564, 2015.

[33] M. J. Nieves-Remacha, A. A. Kulkarni, and K. F. Jensen, "Gas–Liquid Flow and Mass Transfer in an Advanced-Flow Reactor," *Industrial & Engineering Chemistry Research,* vol. 52, pp. 8996-9010, 2013.

[34] J. P. McMullen and K. F. Jensen, "Rapid Determination of Reaction Kinetics with an Automated Microfluidic System," *Organic Process Research & Development,* vol. 15, pp. 398-407, 2011.

[35] B. J. Reizman and K. F. Jensen, "An Automated Continuous-Flow Platform for the Estimation of Multistep Reaction Kinetics," *Organic Process Research & Development,* vol. 16, pp. 1770-1782, 2012.

[36] C. J. Smith, N. Nikbin, S. V. Ley, H. Lange, and I. R. Baxendale, "A fully automated, multistep flow synthesis of 5-amino-4-cyano-1,2,3-triazoles," *Organic & Biomolecular Chemistry,* vol. 9, pp. 1938-1947, 2011.

[37] C. F. Carter, H. Lange, S. V. Ley, I. R. Baxendale, B. Wittkamp, J. G. Goode*, et al.*, "ReactIR Flow Cell: A New Analytical Tool for Continuous Flow Chemical Processing," *Organic Process Research & Development,* vol. 14, pp. 393-404, 2010.

[38] D. E. Fitzpatrick, C. Battilocchio, and S. V. Ley, "A Novel Internet-Based Reaction Monitoring, Control and Autonomous Self-Optimization Platform for Chemical Synthesis," *Organic Process Research & Development,* vol. 20, pp. 386-394, 2016.

[39] V. Hessel, C. Hofmann, P. Löb, J. Löhndorf, H. Löwe, and A. Ziogas, "Aqueous Kolbe−Schmitt Synthesis Using Resorcinol in a Microreactor Laboratory Rig under High-p,T Conditions," *Organic Process Research & Development,* vol. 9, pp. 479-489, 2005.

[40] M. K. S. Verma, S. R. Ganneboyina, Rakshith, and A. Ghatak, "Three-Dimensional Multihelical Microfluidic Mixers for Rapid Mixing of Liquids," *Langmuir,* vol. 24, pp. 2248-2251, 2008.

[41] M. K. S. Verma, A. Majumder, and A. Ghatak, "Embedded Template-Assisted Fabrication of Complex Microchannels in PDMS and Design of a Microfluidic Adhesive," *Langmuir,* vol. 22, pp. 10291-10295, 2006.

[42] F. Bally, C. A. Serra, C. Brochon, N. Anton, T. Vandamme, and G. Hadziioannou, "A Continuous-Flow Polymerization Microprocess with

Online GPC and Inline Polymer Recovery by Micromixer-Assisted Nanoprecipitation," *Macromolecular Reaction Engineering,* vol. 5, pp. 542-547, 2011.

[43] F. Bally, D. K. Garg, C. A. Serra, Y. Hoarau, N. Anton, C. Brochon*, et al.*, "Improved size-tunable preparation of polymeric nanoparticles by microfluidic nanoprecipitation," *Polymer,* vol. 53, pp. 5045-5051, 2012.

[44] V. Hessel, S. Hardt, H. Löwe, and F. Schönfeld, "Laminar mixing in different interdigital micromixers: I. Experimental characterization," *AIChE Journal,* vol. 49, pp. 566-577, 2003.

[45] S. Hardt and F. Schönfeld, "Laminar mixing in different interdigital micromixers: II. Numerical simulations," *AIChE Journal,* vol. 49, pp. 578-584, 2003.

[46] X. Yao, Y. Zhang, L. Du, J. Liu, and J. Yao, "Review of the applications of microreactors," *Renewable and Sustainable Energy Reviews,* vol. 47, pp. 519-539, 2015.

[47] B. Walsh, J. R. Hyde, P. Licence, and M. Poliakoff, "The automation of continuous reactions in supercritical CO2: the acid-catalysed etherification of short chain alcohols," *Green Chemistry,* vol. 7, pp. 456-463, 2005.

[48] R. A. Bourne, J. G. Stevens, J. Ke, and M. Poliakoff, "Maximising opportunities in supercritical chemistry: the continuous conversion of levulinic acid to γ-valerolactone in CO2," *Chemical Communications,* pp. 4632-4634, 2007.

[49] J. G. Stevens, R. A. Bourne, M. V. Twigg, and M. Poliakoff, "Real-Time Product Switching Using a Twin Catalyst System for the Hydrogenation of Furfural in Supercritical CO2," *Angewandte Chemie International Edition,* vol. 49, pp. 8856-8859, 2010.

[50] J. G. Stevens, R. A. Bourne, and M. Poliakoff, "The continuous self aldol condensation of propionaldehyde in supercritical carbon dioxide: a highly selective catalytic route to 2-methylpentenal," *Green Chemistry,* vol. 11, pp. 409-416, 2009.

[51] P. Licence, W. K. Gray, M. Sokolova, and M. Poliakoff, "Selective Monoprotection of 1,n-Terminal Diols in Supercritical Carbon Dioxide: A Striking Example of Solvent Tunable Desymmetrization," *Journal of the American Chemical Society,* vol. 127, pp. 293-298, 2005.

[52] X. Han and M. Poliakoff, "Continuous reactions in supercritical carbon dioxide: problems, solutions and possible ways forward," *Chemical Society Reviews,* vol. 41, pp. 1428-1436, 2012.

[53] P. N. Gooden, R. A. Bourne, A. J. Parrott, H. S. Bevinakatti, D. J. Irvine, and M. Poliakoff, "Continuous Acid-Catalyzed Methylations in Supercritical Carbon Dioxide: Comparison of Methanol, Dimethyl Ether and Dimethyl Carbonate as Methylating Agents," *Organic Process Research & Development,* vol. 14, pp. 411-416, 2010.

[54] J. P. McMullen and K. F. Jensen, "Integrated Microreactors for Reaction Automation: New Approaches to Reaction Development," *Annual Review of Analytical Chemistry,* vol. 3, pp. 19-42, 2010.

[55] R. A. Skilton, A. J. Parrott, M. W. George, M. Poliakoff, and R. A. Bourne, "Real-Time Feedback Control Using Online Attenuated Total Reflection Fourier Transform Infrared (ATR FT-IR) Spectroscopy for Continuous Flow Optimization and Process Knowledge," *Applied Spectroscopy,* vol. 67, pp. 1127-1131, 2013.

[56] R. A. Skilton, R. A. Bourne, Z. Amara, R. Horvath, J. Jin, M. J. Scully*, et al.*, "Remote-controlled experiments with cloud chemistry," *Nature Chemistry,* vol. 7, p. 1, 2014.

[57] Z. Amara, E. S. Streng, R. A. Skilton, J. Jin, M. W. George, and M. Poliakoff, "Automated Serendipity with Self-Optimizing Continuous-Flow Reactors," *European Journal of Organic Chemistry,* vol. 2015, pp. 6141-6145, 2015.

[58] V. Sans, L. Porwol, V. Dragone, and L. Cronin, "A self optimizing synthetic organic reactor system using real-time in-line NMR spectroscopy," *Chemical Science,* vol. 6, pp. 1258-1264, 2015.

[59] D. Cortés-Borda, K. V. Kutonova, C. Jamet, M. E. Trusova, F. Zammattio, C. Truchet*, et al.*, "Optimizing the Heck–Matsuda Reaction in Flow with a Constraint-Adapted Direct Search Algorithm," *Organic Process Research & Development,* vol. 20, pp. 1979-1987, 2016.

[60] B. J. Reizman and K. F. Jensen, "Simultaneous solvent screening and reaction optimization in microliter slugs," *Chemical Communications,* vol. 51, pp. 13290-13293, 2015.

[61] A.-C. Bédard, A. Adamo, K. C. Aroh, M. G. Russell, A. A. Bedermann, J. Torosian*, et al.*, "Reconfigurable system for automated optimization of diverse chemical reactions," *Science,* vol. 361, p. 1220, 2018.

[62] Y.-J. Hwang, C. W. Coley, M. Abolhasani, A. L. Marzinzik, G. Koch, C. Spanka*, et al.*, "A segmented flow platform for on-demand medicinal chemistry and compound synthesis in oscillating droplets," *Chemical Communications,* vol. 53, pp. 6649-6652, 2017.

[63] H.-W. Hsieh, C. W. Coley, L. M. Baumgartner, K. F. Jensen, and R. I. Robinson, "Photoredox Iridium–Nickel Dual-Catalyzed Decarboxylative Arylation Cross-Coupling: From Batch to Continuous Flow via Self-Optimizing Segmented Flow Reactor," *Organic Process Research & Development,* vol. 22, pp. 542-550, 2018.

[64] S. P. Sutera and R. Skalak, "The History of Poiseuille's Law," *Annual Review of Fluid Mechanics,* vol. 25, pp. 1-20, 1993.

[65] B. J. Kirby, *Micro- and Nanoscale Fluid Mechanics: Transport in Microfluidic Devices*. Cambridge: Cambridge University Press, 2010.

[66] G. I. Taylor, "Dispersion of soluble matter in solvent flowing slowly through a tube," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences,* vol. 219, pp. 186-203, 1953.

[67] O. Levenspiel, *Chemical reaction engineering*, 3 ed. New York: John Wiley & Sons, 1999.

[68] T. M. Squires and S. R. Quake, "Microfluidics: Fluid physics at the nanoliter scale," *Reviews of Modern Physics,* vol. 77, pp. 977-1026, 2005.

[69] M. Gonidec and J. Puigmartí-Luis, "Continuous- versus Segmented-Flow Microfluidic Synthesis in Materials Science," *Crystals,* vol. 9, p. 12, 2018.

[70] M. Adamo, A. S. Poulos, C. G. Lopez, A. Martel, L. Porcar, and J. T. Cabral, "Droplet microfluidic SANS," *Soft Matter,* vol. 14, pp. 1759-1770, 2018.

[71] X. Casadevall i Solvas and A. deMello, "Droplet microfluidics: recent developments and future applications," *Chemical Communications,* vol. 47, pp. 1936-1942, 2011.

[72] J. D. Tice, H. Song, A. D. Lyon, and R. F. Ismagilov, "Formation of Droplets and Mixing in Multiphase Microfluidics at Low Values of the Reynolds and the Capillary Numbers," *Langmuir,* vol. 19, pp. 9127-9133, 2003.

[73] C. E. Stanley, R. C. R. Wootton, and A. J. deMello, "Continuous and Segmented Flow Microfluidics: Applications in High-throughput Chemistry and Biology," *CHIMIA International Journal for Chemistry,* vol. 66, pp. 88-98, 2012.

[74] J. M. Ottino, S. R. Wiggins, M. R. Bringer, C. J. Gerdts, H. Song, J. D. Tice*, et al.*, "Microfluidic systems for chemical kinetics that rely on chaotic mixing in droplets," *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences,* vol. 362, pp. 1087-1104, 2004.

[75] A. M. Schweidtmann, A. D. Clayton, N. Holmes, E. Bradford, R. A. Bourne, and A. A. Lapkin, "Machine learning meets continuous flow chemistry: Automated optimization towards the Pareto front of multiple objectives," *Chemical Engineering Journal,* vol. 352, pp. 277-282, 2018.

[76] M. I. Jeraal, N. Holmes, G. R. Akien, and R. A. Bourne, "Enhanced process development using automated continuous reactors by self-optimisation algorithms and statistical empirical modelling," *Tetrahedron,* vol. 74, pp. 3158-3164, 2018.

[77] N. Holmes, G. R. Akien, R. J. D. Savage, C. Stanetty, I. R. Baxendale, A. J. Blacker*, et al.*, "Online quantitative mass spectrometry for the rapid adaptive optimisation of automated flow reactors," *Reaction Chemistry & Engineering,* vol. 1, pp. 96-100, 2016.

[78] A. Echtermeyer, Y. Amar, J. Zakrzewski, and A. Lapkin, "Self-optimisation and model-based design of experiments for developing a C–H activation flow process," *Beilstein Journal of Organic Chemistry,* vol. 13, pp. 150-163, 2017.

[79] J. S. Moore and K. F. Jensen, "“Batch” Kinetics in Flow: Online IR Analysis and Continuous Control," *Angewandte Chemie International Edition,* vol. 53, pp. 470-473, 2014.

[80] K. D. Nagy, B. Shen, T. F. Jamison, and K. F. Jensen, "Mixing and Dispersion in Small-Scale Flow Systems," *Organic Process Research & Development,* vol. 16, pp. 976-981, 2012.

[81] R. L. Hartman, "Managing Solids in Microreactors for the Upstream Continuous Processing of Fine Chemicals," *Organic Process Research & Development,* vol. 16, pp. 870-887, 2012.

[82] S. Pal and A. A. Kulkarni, "Interfacial precipitation and clogging in straight capillaries," *Chemical Engineering Science,* vol. 153, pp. 344-353, 2016.

[83] A. Suzuki, "Cross-Coupling Reactions Of Organoboranes: An Easy Way To Construct C□C Bonds (Nobel Lecture)," *Angewandte Chemie International Edition,* vol. 50, pp. 6722-6737, 2011.

[84] D. S. Surry and S. L. Buchwald, "Biaryl Phosphane Ligands in Palladium-Catalyzed Amination," *Angewandte Chemie International Edition,* vol. 47, pp. 6338-6361, 2008.

[85] H. Wakami and J.-i. Yoshida, "Grignard Exchange Reaction Using a Microflow System: From Bench to Pilot Plant," *Organic Process Research & Development,* vol. 9, pp. 787-791, 2005.

[86] K. W. Anderson and S. L. Buchwald, "General Catalysts for the Suzuki–Miyaura and Sonogashira Coupling Reactions of Aryl Chlorides and for the Coupling of Challenging Substrate Combinations in Water," *Angewandte Chemie International Edition,* vol. 44, pp. 6173-6177, 2005.

[87] I. P. Beletskaya and A. V. Cheprakov, "The Heck Reaction as a Sharpening Stone of Palladium Catalysis," *Chemical Reviews,* vol. 100, pp. 3009-3066, 2000.

[88] W. Cabri and I. Candiani, "Recent Developments and New Perspectives in the Heck Reaction," *Accounts of Chemical Research,* vol. 28, pp. 2-7, 1995.

[89] J. G. de Vries, "The Heck reaction in the production of fine chemicals," *Canadian Journal of Chemistry,* vol. 79, pp. 1086-1092, 2001.

[90] A. L. Hansen and T. Skrydstrup, "Fast and Regioselective Heck Couplings with N-Acyl-N-vinylamine Derivatives," *The Journal of Organic Chemistry,* vol. 70, pp. 5997-6003, 2005.

[91] T. W. Lyons and M. S. Sanford, "Palladium-Catalyzed Ligand-Directed C−H Functionalization Reactions," *Chemical Reviews,* vol. 110, pp. 1147-1169, 2010.

[92] J. Sedelmeier, S. V. Ley, I. R. Baxendale, and M. Baumann, "KMnO4-Mediated Oxidation as a Continuous Flow Process," *Organic Letters,* vol. 12, pp. 3618-3621, 2010.

[93] E. Quevedo, J. Steinbacher, and D. T. McQuade, "Interfacial Polymerization within a Simplified Microfluidic Device: Capturing Capsules," *Journal of the American Chemical Society,* vol. 127, pp. 10498-10499, 2005.

[94] W. Li, H. H. Pham, Z. Nie, B. MacDonald, A. Güenther, and E. Kumacheva, "Multi-Step Microfluidic Polymerization Reactions Conducted in Droplets: The Internal Trigger Approach," *Journal of the American Chemical Society,* vol. 130, pp. 9935-9941, 2008.

[95] M. Seo, Z. Nie, S. Xu, M. Mok, P. C. Lewis, R. Graham*, et al.*, "Continuous Microfluidic Reactors for Polymer Particles," *Langmuir,* vol. 21, pp. 11614-11622, 2005.

[96] S. L. Poe, M. A. Cummings, M. P. Haaf, and D. T. McQuade, "Solving the Clogging Problem: Precipitate-Forming Reactions in Flow," *Angewandte Chemie International Edition,* vol. 45, pp. 1544-1548, 2006.

[97] D. L. Chen, C. J. Gerdts, and R. F. Ismagilov, "Using Microfluidics to Observe the Effect of Mixing on Nucleation of Protein Crystals," *Journal of the American Chemical Society,* vol. 127, pp. 9672-9673, 2005.

[98] S. Xu, Z. Nie, M. Seo, P. Lewis, E. Kumacheva, H. A. Stone*, et al.*, "Generation of Monodisperse Particles by Using Microfluidics: Control over Size, Shape, and Composition," *Angewandte Chemie International Edition,* vol. 44, pp. 724-728, 2005.

[99] R. L. Hartman, J. R. Naber, N. Zaborenko, S. L. Buchwald, and K. F. Jensen, "Overcoming the Challenges of Solid Bridging and Constriction during Pd-Catalyzed C−N Bond Formation in Microreactors," *Organic Process Research & Development,* vol. 14, pp. 1347-1357, 2010.

[100] T. Noël, J. R. Naber, R. L. Hartman, J. P. McMullen, K. F. Jensen, and S. L. Buchwald, "Palladium-catalyzed amination reactions in flow: overcoming the challenges of clogging via acoustic irradiation," *Chemical Science,* vol. 2, pp. 287-290, 2011.

[101] S. Kuhn, T. Noël, L. Gu, P. L. Heider, and K. F. Jensen, "A Teflon microreactor with integrated piezoelectric actuator to handle solid forming reactions," *Lab on a Chip,* vol. 11, pp. 2488-2492, 2011.

[102] N. Pamme and C. Wilhelm, "Continuous sorting of magnetic cells via on-chip free-flow magnetophoresis," *Lab on a Chip,* vol. 6, pp. 974-980, 2006.

[103] A. I. Rodríguez-Villarreal, M. D. Tarn, L. A. Madden, J. B. Lutz, J. Greenman, J. Samitier*, et al.*, "Flow focussing of particles and cells based on their intrinsic properties using a simple diamagnetic repulsion setup," *Lab on a Chip,* vol. 11, pp. 1240-1248, 2011.

[104] N. Pamme, "Magnetism and microfluidics," *Lab on a Chip,* vol. 6, pp. 24-38, 2006.

[105] J. G. Kralj, M. T. W. Lis, M. A. Schmidt, and K. F. Jensen, "Continuous Dielectrophoretic Size-Based Particle Sorting," *Analytical Chemistry,* vol. 78, pp. 5019-5025, 2006.

[106] M. Dürr, J. Kentsch, T. Müller, T. Schnelle, and M. Stelzle, "Microdevices for manipulation and accumulation of micro- and nanoparticles by dielectrophoresis," *ELECTROPHORESIS,* vol. 24, pp. 722-731, 2003.

[107] A. Asthana, I. Zinovik, C. Weinmueller, and D. Poulikakos, "Significant Nusselt number increase in microchannels with a segmented flow of two immiscible liquids: An experimental study," *International Journal of Heat and Mass Transfer,* vol. 54, pp. 1456-1464, 2011.

[108] A. Günther and K. F. Jensen, "Multiphase microfluidics: from flow characteristics to chemical and materials synthesis," *Lab on a Chip,* vol. 6, pp. 1487-1503, 2006.

[109] M. W. Losey, R. J. Jackman, S. L. Firebaugh, M. A. Schmidt, and K. F. Jensen, "Design and fabrication of microfluidic devices for multiphase

mixing and reaction," *Journal of Microelectromechanical Systems,* vol. 11, pp. 709-717, 2002.

[110] C. Houben, N. Peremezhney, A. Zubov, J. Kosek, and A. A. Lapkin, "Closed-Loop Multitarget Optimization for Discovery of New Emulsion Polymerization Recipes," *Organic Process Research & Development,* vol. 19, pp. 1049-1053, 2015.

[111] D. L. Browne, B. J. Deadman, R. Ashe, I. R. Baxendale, and S. V. Ley, "Continuous Flow Processing of Slurries: Evaluation of an Agitated Cell Reactor," *Organic Process Research & Development,* vol. 15, pp. 693-697, 2011.

[112] P. Filipponi, A. Gioiello, and I. R. Baxendale, "Controlled Flow Precipitation as a Valuable Tool for Synthesis," *Organic Process Research & Development,* vol. 20, pp. 371-375, 2016.

[113] L. Liguori and H.-R. Bjørsvik, "Multijet Oscillating Disc Millireactor: A Novel Approach for Continuous Flow Organic Synthesis," *Organic Process Research & Development,* vol. 15, pp. 997-1009, 2011.

[114] A. D. Randolph, "The mixed suspension, mixed product removal crystallizer as a concept in crystallizer design," *AIChE Journal,* vol. 11, pp. 424-430, 1965.

[115] Y. Zhang, S. C. Born, and K. F. Jensen, "Scale-Up Investigation of the Continuous Phase-Transfer-Catalyzed Hypochlorite Oxidation of Alcohols and Aldehydes," *Organic Process Research & Development,* vol. 18, pp. 1476-1481, 2014.

[116] N. G. Anderson, "Practical Use of Continuous Processing in Developing and Scaling Up Laboratory Processes," *Organic Process Research & Development,* vol. 5, pp. 613-621, 2001.

[117] C. Wiles and P. Watts, "Continuous flow reactors: a perspective," *Green Chemistry,* vol. 14, pp. 38-54, 2012.

[118] Y. Mo and K. F. Jensen, "A miniature CSTR cascade for continuous flow of reactions containing solids," *Reaction Chemistry & Engineering,* 2016.

[119] Y. Mo, H. Lin, and K. F. Jensen, "High-performance miniature CSTR for biphasic C–C bond-forming reactions," *Chemical Engineering Journal,* vol. 335, pp. 936-944, 2018.

[120] K. E. Jolley, M. R. Chapman, and A. J. Blacker, "A general and atom-efficient continuous-flow approach to prepare amines, amides and imines via reactive N-chloramines," *Beilstein Journal of Organic Chemistry,* vol. 14, pp. 2220-2228, 2018.

[121] H. S. Fogler, *Elements of chemical reaction engineering*, 5 ed. Boston: Prentice Hall, 2016.

[122] M. R. Chapman, M. H. T. Kwan, G. King, K. E. Jolley, M. Hussain, S. Hussain*, et al.*, "Simple and Versatile Laboratory Scale CSTR for Multiphasic Continuous-Flow Chemistry and Long Residence Times," *Organic Process Research & Development,* vol. 21, pp. 1294-1301, 2017.

[123] T. Wirth, *Microreactors in organic synthesis and catalysis*, 2 ed. Weinheim; Chichester: Wiley-VCH ; John Wiley [distributor], 2013.

[124] E. L. Paul, V. A. Atiemo-Obeng, and S. M. Kresta, *Handbook of industrial mixing science and practice*, 18 ed. Hoboken, N.J: Wiley-Interscience, 2004.

[125] H. Wang, A. Mustaffar, A. N. Phan, V. Zivkovic, D. Reay, R. Law*, et al.*, "A review of process intensification applied to solids handling," *Chemical Engineering and Processing: Process Intensification,* vol. 118, pp. 78-107, 2017.

[126] T. McGlone, N. E. B. Briggs, C. A. Clark, C. J. Brown, J. Sefcik, and A. J. Florence, "Oscillatory Flow Reactors (OFRs) for Continuous Manufacturing and Crystallization," *Organic Process Research & Development,* vol. 19, pp. 1186-1202, 2015.

[127] A. W. Fitch, H. Jian, and X. Ni, "An investigation of the effect of viscosity on mixing in an oscillatory baffled column using digital particle image velocimetry and computational fluid dynamics simulation," *Chemical Engineering Journal,* vol. 112, pp. 197-210, 2005.

[128] M. S. R. Abbott, A. P. Harvey, G. V. Perez, and M. K. Theodorou, "Biological processing in oscillatory baffled reactors: operation, advantages and potential," *Interface focus,* vol. 3, p. 20120036, 2013.

[129] A. P. Harvey, M. R. Mackley, and T. Seliger, "Process intensification of biodiesel production using a continuous oscillatory flow reactor," *Journal of Chemical Technology & Biotechnology,* vol. 78, pp. 338-341, 2003.

[130] E. Lobry, T. Lasuye, C. Gourdon, and C. Xuereb, "Liquid–liquid dispersion in a continuous oscillatory baffled reactor – Application to suspension polymerization," *Chemical Engineering Journal,* vol. 259, pp. 505-518, 2015.

[131] S. Lawton, G. Steele, P. Shering, L. Zhao, I. Laird, and X.-W. Ni, "Continuous Crystallization of Pharmaceuticals Using a Continuous Oscillatory Baffled Crystallizer," *Organic Process Research & Development,* vol. 13, pp. 1357-1363, 2009.

[132] M. R. Mackley, K. B. Smith, and N. P. Wise, "Mixing and separation of particle suspensions using oscillatory flow in baffled tubes," *Chemical Engineering Research and Design,* vol. 71, pp. 649-656, 1993.

[133] A. Mazubert, J. Aubin, S. Elgue, and M. Poux, "Intensification of waste cooking oil transformation by transesterification and esterification reactions in oscillatory baffled and microstructured reactors for biodiesel production," *Green Processing and Synthesis,* vol. 3, pp. 419-429, 2014.

[134] M. Palma and R. Giudici, "Analysis of axial dispersion in an oscillatory-flow continuous reactor," *Chemical Engineering Journal,* vol. 94, pp. 189-198, 2003.

[135] A. P. Harvey, M. R. Mackley, and P. Stonestreet, "Operation and Optimization of an Oscillatory Flow Continuous Reactor," *Industrial & Engineering Chemistry Research,* vol. 40, pp. 5371-5377, 2001.

[136] N. Reis, A. A. Vicente, J. A. Teixeira, and M. R. Mackley, "Residence times and mixing of a novel continuous oscillatory flow screening reactor," *Chemical Engineering Science,* vol. 59, pp. 4967-4974, 2004.

[137] A. N. Phan and A. Harvey, "Development and evaluation of novel designs of continuous mesoscale oscillatory baffled reactors," *Chemical Engineering Journal,* vol. 159, pp. 212-219, 2010.

[138] A. N. Phan, A. Harvey, and J. Lavender, "Characterisation of fluid mixing in novel designs of mesoscale oscillatory baffled reactors operating at low flow rates (0.3–0.6ml/min)," *Chemical Engineering and Processing: Process Intensification,* vol. 50, pp. 254-263, 2011.

[139] A. N. Phan and A. P. Harvey, "Effect of geometrical parameters on fluid mixing in novel mesoscale oscillatory helical baffled designs," *Chemical Engineering Journal,* vol. 169, pp. 339-347, 2011.

[140] O. Okafor, A. Weilhard, J. A. Fernandes, E. Karjalainen, R. Goodridge, and V. Sans, "Advanced reactor engineering with 3D printing for the continuous-flow synthesis of silver nanoparticles," *Reaction Chemistry & Engineering,* vol. 2, pp. 129-136, 2017.

[141] F. Almeida, F. Rocha, and A. Ferreira, "Analysis of Liquid Flow and Mixing in an Oscillatory Flow Reactor Provided with 2D Smooth Periodic Constrictions," *U.Porto Journal of Engineering,* vol. 4, pp. 1-15, 2018.

[142] N. Reis, A. P. Harvey, M. R. Mackley, A. A. Vicente, and J. A. Teixeira, "Fluid Mechanics and Design Aspects of a Novel Oscillatory Flow Screening Mesoreactor," *Chemical Engineering Research and Design,* vol. 83, pp. 357-371, 2005.

[143] T. M. Floyd, M. A. Schmidt, and K. F. Jensen, "Silicon Micromixers with Infrared Detection for Studies of Liquid-Phase Reactions," *Industrial & Engineering Chemistry Research,* vol. 44, pp. 2351-2358, 2005.

[144] Z. Qian, I. R. Baxendale, and S. V. Ley, "A Continuous Flow Process Using a Sequence of Microreactors with In-line IR Analysis for the Preparation of N,N-Diethyl-4-(3-fluorophenylpiperidin-4-ylidenemethyl)benzamide as a Potent and Highly Selective δ-Opioid Receptor Agonist," *Chemistry – A European Journal,* vol. 16, pp. 12342-12348, 2010.

[145] P. Koos, U. Gross, A. Polyzos, M. O'Brien, I. Baxendale, and S. V. Ley, "Teflon AF-2400 mediated gas-liquid contact in continuous flow methoxycarbonylations and in-line FTIR measurement of CO concentration," *Organic & Biomolecular Chemistry,* vol. 9, pp. 6903-6908, 2011.

[146] T. Brodmann, P. Koos, A. Metzger, P. Knochel, and S. V. Ley, "Continuous Preparation of Arylmagnesium Reagents in Flow with Inline IR Monitoring," *Organic Process Research & Development,* vol. 16, pp. 1102-1113, 2012.

[147] H. Lange, C. F. Carter, M. D. Hopkin, A. Burke, J. G. Goode, I. R. Baxendale*, et al.*, "A breakthrough method for the accurate addition of reagents in multi-step segmented flow processing," *Chemical Science,* vol. 2, pp. 765-769, 2011.

[148] J. S. Moore and K. F. Jensen, "Automated Multitrajectory Method for Reaction Optimization in a Microfluidic System using Online IR Analysis," *Organic Process Research & Development,* vol. 16, pp. 1409-1415, 2012.

[149] K. Poscharny, D. C. Fabry, S. Heddrich, E. Sugiono, M. A. Liauw, and M. Rueping, "Machine assisted reaction optimization: A self-optimizing

reactor system for continuous-flow photochemical reactions," *Tetrahedron,* vol. 74, pp. 3171-3175, 2018.

[150] M. Adamczyk, J. Fishpaugh, J. Gebler, P. Mattingly, and K. Shreder, "Letter: Detection of reaction intermediates by flow injection electrospray ionization mass spectrometry: reaction of chemiluminescent N-sulfonylacridinium-9-carboxamides with hydrogen peroxide," *European Journal of Mass Spectrometry,* vol. 4, pp. 121-125, 1998.

[151] D. Fabris, "Mass spectrometric approaches for the investigation of dynamic processes in condensed phase," *Mass Spectrometry Reviews,* vol. 24, pp. 30-54, 2005.

[152] Y. Takahashi, R. Sakai, K. Sakamoto, Y. Yoshida, M. Kitaoka, and T. Kitamori, "On-Line High-throughput ESIMS Detection of a Reaction Product Using Synthesis and Extraction Microchips," *Journal of the Mass Spectrometry Society of Japan,* vol. 54, pp. 19-24, 2006.

[153] M. C. Mitchell, V. Spikmans, and A. J. d. Mello, "Microchip-based synthesis and analysis: Control of multicomponent reaction products and intermediates," *Analyst,* vol. 126, pp. 24-27, 2001.

[154] S. Meyer, R. Koch, and J. O. Metzger, "Investigation of Reactive Intermediates of Chemical Reactions in Solution by Electrospray Ionization Mass Spectrometry: Radical Cation Chain Reactions," *Angewandte Chemie International Edition,* vol. 42, pp. 4700-4703, 2003.

[155] M. Brivio, A. Liesener, R. E. Oosterbroek, W. Verboom, U. Karst, A. van den Berg*, et al.*, "Chip-Based On-Line Nanospray MS Method Enabling Study of the Kinetics of Isocyanate Derivatization Reactions," *Analytical Chemistry,* vol. 77, pp. 6852-6856, 2005.

[156] B. V. Silva, F. A. Violante, A. C. Pinto, and L. S. Santos, "The mechanism of Sandmeyer's cyclization reaction by electrospray ionization mass spectrometry," *Rapid Communications in Mass Spectrometry,* vol. 25, pp. 423-428, 2011.

[157] M. C. Mitchell, V. Spikmans, A. Manz, and A. J. de Mello, "Microchip-based synthesis and total analysis systems ([small micro]SYNTAS): chemical microprocessing for generation and analysis of compound libraries," *Journal of the Chemical Society, Perkin Transactions 1,* pp. 514-518, 2001.

[158] D. L. Browne, S. Wright, B. J. Deadman, S. Dunnage, I. R. Baxendale, R. M. Turner*, et al.*, "Continuous flow reaction monitoring using an on-line miniature mass spectrometer," *Rapid Communications in Mass Spectrometry,* vol. 26, pp. 1999-2010, 2012.

[159] Z. Zhou, X. Li, and R. N. Zare, "Optimizing Chemical Reactions with Deep Reinforcement Learning," *ACS Central Science,* vol. 3, pp. 1337-1344, 2017.

[160] A. F. Chrimes, K. Khoshmanesh, P. R. Stoddart, A. Mitchell, and K. Kalantar-zadeh, "Microfluidics and Raman microscopy: current applications and future challenges," *Chemical Society Reviews,* vol. 42, pp. 5880-5906, 2013.

[161] M. Lee, J. P. Lee, H. Rhee, J. Choo, Y. G. Chai, and E. K. Lee, "Applicability of laser-induced Raman microscopy for in situ monitoring of imine formation in a glass microfluidic chip," *Journal of Raman Spectroscopy,* vol. 34, pp. 737-742, 2003.

[162] P. D. I. Fletcher, S. J. Haswell, and X. Zhang, "Monitoring of chemical reactions within microreactors using an inverted Raman microscopic spectrometer," *ELECTROPHORESIS,* vol. 24, pp. 3239-3245, 2003.

[163] B.-H. Jun, M. S. Noh, G. Kim, H. Kang, J.-H. Kim, W.-J. Chung*, et al.*, "Protein separation and identification using magnetic beads encoded with surface-enhanced Raman spectroscopy," *Analytical Biochemistry,* vol. 391, pp. 24-30, 2009.

[164] S. Mozharov, A. Nordon, D. Littlejohn, C. Wiles, P. Watts, P. Dallin*, et al.*, "Improved Method for Kinetic Studies in Microreactors Using Flow Manipulation and Noninvasive Raman Spectrometry," *Journal of the American Chemical Society,* vol. 133, pp. 3601-3608, 2011.

[165] T. A. Hamlin and N. E. Leadbeater, "Raman spectroscopy as a tool for monitoring mesoscale continuous-flow organic synthesis: Equipment interface and assessment in four medicinally-relevant reactions," *Beilstein Journal of Organic Chemistry,* vol. 9, pp. 1843-1852, 2013.

[166] N. Cherkasov, Y. Bai, A. J. Expósito, and E. V. Rebrov, "OpenFlowChem – a platform for quick, robust and flexible automation and self-optimisation of flow chemistry," *Reaction Chemistry & Engineering,* vol. 3, pp. 769-780, 2018.

[167] C. Mateos, M. J. Nieves-Remacha, and J. A. Rincón, "Automated platforms for reaction self-optimization in flow," *Reaction Chemistry & Engineering,* 2019.

[168] C. Houben and A. A. Lapkin, "Automatic discovery and optimization of chemical processes," *Current Opinion in Chemical Engineering,* vol. 9, pp. 1-7, 2015.

[169] J. Y. Buser and A. D. McFarland, "Reaction characterization by flow NMR: quantitation and monitoring of dissolved H2via flow NMR at high pressure," *Chemical Communications,* vol. 50, pp. 4234-4237, 2014.

[170] J. Bart, A. J. Kolkman, A. J. Oosthoek-de Vries, K. Koch, P. J. Nieuwland, H. Janssen*, et al.*, "A Microfluidic High-Resolution NMR Flow Probe," *Journal of the American Chemical Society,* vol. 131, pp. 5014-5015, 2009.

[171] M. V. Gomez, A. M. Rodriguez, A. de la Hoz, F. Jimenez-Marquez, R. M. Fratila, P. A. Barneveld*, et al.*, "Determination of Kinetic Parameters within a Single Nonisothermal On-Flow Experiment by Nanoliter NMR Spectroscopy," *Analytical Chemistry,* vol. 87, pp. 10547-10555, 2015.

[172] F. Dalitz, M. Cudaj, M. Maiwald, and G. Guthausen, "Process and reaction monitoring by low-field NMR spectroscopy," *Progress in Nuclear Magnetic Resonance Spectroscopy,* vol. 60, pp. 52-70, 2012.

[173] A. Nordon, A. Diez-Lazaro, C. W. L. Wong, C. A. McGill, D. Littlejohn, M. Weerasinghe*, et al.*, "Consideration of some sampling problems in the on-line analysis of batch processes by low-field NMR spectrometry," *Analyst,* vol. 133, pp. 339-347, 2008.

[174] E. Danieli, J. Perlo, B. Blümich, and F. Casanova, "Small Magnets for Portable NMR Spectrometers," *Angewandte Chemie International Edition,* vol. 49, pp. 4133-4135, 2010.

[175] K. Halbach, "Design of permanent multipole magnets with oriented rare earth cobalt material," *Nuclear Instruments and Methods,* vol. 169, pp. 1-10, 1980.

[176] E. Danieli, J. Perlo, A. L. L. Duchateau, G. K. M. Verzijl, V. M. Litvinov, B. Blümich*, et al.*, "On-Line Monitoring of Chemical Reactions by using Bench-Top Nuclear Magnetic Resonance Spectroscopy," *ChemPhysChem,* vol. 15, pp. 3060-3066, 2014.

[177] D. Cortés-Borda, E. Wimmer, B. Gouilleux, E. Barré, N. Oger, L. Goulamaly*, et al.*, "An Autonomous Self-Optimizing Flow Reactor for the Synthesis of Natural Product Carpanone," *The Journal of Organic Chemistry,* vol. 83, pp. 14286-14299, 2018.

[178] M. Rubens, J. H. Vrijsen, J. Laun, and T. Junkers, "Precise Polymer Synthesis by Autonomous Self-Optimizing Flow Reactors," *Angewandte Chemie International Edition,* vol. 58, pp. 3183-3187, 2019.

[179] A. L. Cauchy, "Methode generale pour la resolution des systemes d'equations simultanees," *Comptes Rendus de l'Academie des Science,* vol. 25, pp. 536-538, 1847.

[180] J. Barzilai and J. M. Borwein, "2-POINT STEP SIZE GRADIENT METHODS," *Ima Journal of Numerical Analysis,* vol. 8, pp. 141-148, 1988.

[181] L. Armijo, "Minimization of functions having Lipschitz continuous first partial derivatives," *Pacific J. Math.,* vol. 16, pp. 1-3, 1966.

[182] W. Spendley, G. R. Hext, and F. R. Himsworth, "SEQUENTIAL APPLICATION OF SIMPLEX DESIGNS IN OPTIMISATION AND EVOLUTIONARY OPERATION," *Technometrics,* vol. 4, pp. 441-&, 1962.

[183] J. A. Nelder and R. Mead, "A SIMPLEX-METHOD FOR FUNCTION MINIMIZATION," *Computer Journal,* vol. 7, pp. 308-313, 1965.

[184] J. Yu Cheng and T. Mailund, *Ancestral population genomics using coalescence hidden Markov models and heuristic optimisation algorithms* vol. 360, 2015.

[185] M. W. Routh, P. A. Swartz, and M. B. Denton, "Performance of the Super Modified Simplex," *Analytical Chemistry,* vol. 49, pp. 1422-1428, 1977.

[186] W. Huyer and A. Neumaier, "SNOBFIT -- Stable Noisy Optimization by Branch and Fit," *ACM Trans. Math. Softw.,* vol. 35, pp. 1-25, 2008.

[187] Y. Censor, "Pareto optimality in multiobjective problems," *Applied Mathematics and Optimization,* vol. 4, pp. 41-59, 1977.

[188] M. T. M. Emmerich and A. H. Deutz, "A tutorial on multiobjective optimization: fundamentals and evolutionary methods," *Natural Computing,* vol. 17, pp. 585-609, 2018.

[189] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation,* vol. 6, pp. 182-197, 2002.

[190] J. Knowles, "ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems," *IEEE Transactions on Evolutionary Computation,* vol. 10, pp. 50-66, 2006.

[191] M. T. M. Emmerich, K. C. Giannakoglou, and B. Naujoks, "Single- and multiobjective evolutionary optimization assisted by Gaussian random field metamodels," *IEEE Transactions on Evolutionary Computation,* vol. 10, pp. 421-439, 2006.

[192] A. I. J. Forrester and A. J. Keane, "Recent advances in surrogate-based optimization," *Progress in Aerospace Sciences,* vol. 45, pp. 50-79, 2009.

[193] E. Bradford, A. M. Schweidtmann, and A. Lapkin, "Efficient multiobjective optimization employing Gaussian processes, spectral sampling and a genetic algorithm," *Journal of Global Optimization,* vol. 71, pp. 407-438, 2018.

[194] J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani, "Predictive entropy search for efficient global optimization of black-box functions," presented at the Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, Montreal, Canada, 2014.

[195] J. B. Edel, R. Fortt, J. C. deMello, and A. J. deMello, "Microfluidic routes to the controlled production of nanoparticles," *Chemical Communications,* pp. 1136-1137, 2002.

[196] J. R. Davis and A. I. H. Committee, "Part II. Irons, Steels, and High-Performance Alloys," in *Metals Handbook Desk Edition 2nd Edition*, ed: Taylor & Francis, 1998.

[197] *Arduino Uno Rev3.* Arduino.cc. Retrieved October 16, 2019, from https://store.arduino.cc/arduino-uno-rev3

[198] P. R. N. Childs, J. R. Greenwood, and C. A. Long, "Review of temperature measurement," *Review of Scientific Instruments,* vol. 71, pp. 2959-2978, 2000.

[199] J. X. J. Zhang and K. Hoshino, "Mechanical Transducers: Cantilevers, Acoustic Wave Sensors, and Thermal Sensors," in *Molecular Sensors and Nanodevices*, J. X. J. Zhang and K. Hoshino, Eds., ed Oxford: William Andrew Publishing, 2014, pp. 321-414.

[200] C. Hagart-Alexander, "Temperature Measurement," in *Instrumentation Reference Book (Fourth Edition)*, W. Boyes, Ed., ed Boston: Butterworth-Heinemann, 2010, pp. 269-326.

[201] *Precision Temperature Amplifiers with Cold Junction Compensation Rev.D.* Analog.com. Retrieved October 17, 2019, from https://www.analog.com/media/en/technical-documentation/data-sheets/ad8494_8495_8496_8497.pdf

[202] *TLVH431, TLVH432 Low-Voltage Adjustable Precision Shunt Regulators.* ti.com. Retrieved October 17, 2019, from http://www.ti.com/lit/ds/symlink/tlvh431b.pdf

[203] J. Blum, "Using Transistors and Driving Motors," in *Exploring Arduino: Tools and Techniques for Engineering Wizardry*, ed: John Wiley & Sons, Inc., 2013, pp. 63-68.

[204] J. Blum, "Digital Inputs, Outputs and Pulse-Width Modulation," in *Exploring Arduino: Tools and Techniques for Engineering Wizardry*, ed: John Wiley & Sons, Inc., 2013, pp. 29-31.

[205] T. Hirzel. *Arduino PWM foundations.* arduino.cc. Retrieved November 23, 2019, from https://www.arduino.cc/en/tutorial/PWM

[206] *FQP30N06L N-channel QFET MOSFET.* Onsemi.com. Retrieved October 17, 2019, from https://www.onsemi.com/pub/Collateral/FQP30N06L-D.pdf

[207] H. Zumbahlen, "Printed Circuit-Board Design Issues," in *Linear Circuit Design Handbook*, H. Zumbahlen, Ed., ed Burlington: Newnes, 2008, pp. 821-895.

[208] J. Blum, "Interfacing with Liquid Crystal Displays," in *Exploring Arduino: Tools and Techniques for Engineering Wizardry*, ed: John Wiley & Sons, Inc., 2013, pp. 199-219.

[209] J. Blum, "The 12C Bus," in *Exploring Arduino: Tools and Techniques for Engineering Wizardry*, ed: John Wiley & Sons, Inc., 2013, pp. 163-180.

[210] P. Horowitz and W. Hill, "Foundations," in *The art of electronics*, 3rd ed, 2018, pp. 1-68.

[211] M. Faraday, *Experimental researches in electricity*. London: Printed by Richard Taylor, 1832.

[212] A. Robbins and W. Miller, "Inductive Transients," in *Circuit analysis : theory and practice*, 5th ed Clifton Park, NY: Cengage Learning, 2013, pp. 457-478.

[213] "Electrodynamics," in *Electromagnetics Explained*, R. Schmitt, Ed., ed Burlington: Newnes, 2002, pp. 75-88.

[214] J. M. Jacob, "Power Switches," in *Power electronics : principles & applications*, ed Albany: Delmar Thomson Learning, 2002, pp. 269-321.

[215] *1N4001 thru 1N4007 Vishay General Semiconductor Datasheet.* Vishay.com. Retrieved 18 October, 2019, from https://www.vishay.com/docs/88503/1n4001.pdf

[216] *IPC-2221 Generic Standard on Printed Board Design.* Ipc.org. Retrieved 18 October, 2019, from http://www.ipc.org/TOC/IPC-2221.pdf

[217] P. Wilson, "Printed circuits," in *The Circuit Designer's Companion*, 3rd ed: Newnes, 2013, pp. 47-83.

[218] P. Wilson, "Ground and power distribution," in *The Circuit Designer's Companion*, 3rd ed: Newnes, 2013, pp. 2-42.

[219] *Miniature Rectangular Thermocouple Connector Fascia Socket IM-K-FF Type K IEC Datasheet.* Labfacility.com. Retrieved October 21, 2019, from https://www.labfacility.com/product_pdf/download/pdf/id/378

[220] *Miniature Thermocouple Connector Plug IM-K-M Type K IEC Datasheet.* Labfacility.com. Retrieved October 22, 2019, from https://www.labfacility.com/product_pdf/download/pdf/id/344

[221] *1614 21 Lumberg Snap-In DC Power Socket Datasheet.* Retrieved October 21, 2019, from https://downloads.lumberg.com/datenblaetter/en/1614_21.pdf

[222] *1636 07 Lumberg DC Power Plug Datasheet.* Lumberg.com. Retrieved October 21, 2019, from https://downloads.lumberg.com/datenblaetter/en/1636_07.pdf

[223] R. G. Lyons, *Understanding digital signal processing*. Upper Saddle River [etc.]: Pearson Education International, 2013.

[224] *AN118 - Improving ADC Resolution by Oversampling and Averaging.* Silabs.com. Retrieved October 21, 2019, from https://www.silabs.com/documents/public/application-notes/an118.pdf

[225] *AVR121: Enhancing ADC resolution by oversampling* Microchip.com. Retrieved from http://ww1.microchip.com/downloads/en/Appnotes/doc8003.pdf

[226] M. Pelgrom, *ANALOG-TO-DIGITAL CONVERSION*. [Place of publication not identified]: SPRINGER, 2018.

[227] K. Ogata, *Modern control engineering*. Boston [Mass.]; London: Pearson, 2011.

[228] B. Beauregard. *Arduino PID Library.* playground.arduino.cc. Retrieved October 23, 2019, from https://playground.arduino.cc/Code/PIDLibrary/

[229] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," 1942.

[230] G. H. Cohen and G. A. Coon, "Theoretical considerations of retarded control," *Transactions of ASME,* vol. 75, pp. 827-834, 1953.

[231] E. Dahlin, "Designing and tuning digital controllers," *Instruments and Control systems,* vol. 41, pp. 77-83, 1968.

[232] M. T. Coughran. *Lambda Tuning - The Universal Method for PID controllers in Process Control.* Controlglobal.com. Retrieved October 28, 2019, from https://www.controlglobal.com/assets/13WPpdf/131022-coughran-controllers.pdf

[233] D. J. Leith and W. E. Leithead, "Survey of gain-scheduling analysis and design," *International Journal of Control,* vol. 73, pp. 1001-1025, 2000.

[234] D. Cambié, C. Bottecchia, N. J. W. Straathof, V. Hessel, and T. Noël, "Applications of Continuous-Flow Photochemistry in Organic Synthesis, Material Science, and Water Treatment," *Chemical Reviews,* vol. 116, pp. 10276-10341, 2016.

[235] C. W. Coley, M. Abolhasani, H. Lin, and K. F. Jensen, "Material-Efficient Microfluidic Platform for Exploratory Studies of Visible-Light Photoredox Catalysis," *Angewandte Chemie International Edition,* vol. 56, pp. 9847-9850, 2017.

[236] K. P. L. Kuijpers, C. Bottecchia, D. Cambié, K. Drummen, N. J. König, and T. Noël, "A Fully Automated Continuous-Flow Platform for Fluorescence Quenching Studies and Stern–Volmer Analysis," *Angewandte Chemie International Edition,* vol. 57, pp. 11278-11282, 2018.

[237] C. Sambiagio and T. Noël, "Flow Photochemistry: Shine Some Light on Those Tubes!," *Trends in Chemistry,* 2019.

[238] J. A. Manson, A. D. Clayton, C. G. Niño, R. Labes, T. W. Chamberlain, A. J. Blacker*, et al.*, "A Hybridised Optimisation of an Automated Photochemical Continuous Flow Reactor," *CHIMIA International Journal for Chemistry,* vol. 73, pp. 817-822, 2019.

[239] A. Pomberger, Y. Mo, K. Y. Nandiwale, V. L. Schultz, R. Duvadie, R. I. Robinson*, et al.*, "A Continuous Stirred-Tank Reactor (CSTR) Cascade for Handling Solid-Containing Photochemical Reactions," *Organic Process Research & Development,* 2019.

[240] M. Smith and J. March, *March's advanced organic chemistry : reactions, mechanisms, and structure*. Hoboken, N.J.: Wiley-Interscience, 2007.

[241] B. L. Hawbecker, D. W. Kurtz, T. D. Putnam, P. A. Ahlers, and G. D. Gerber, "Aldol condensation: A simple teaching model for organic laboratory," *Journal of Chemical Education,* vol. 55, p. 540, 1978.

[242] L. Claisen and A. Claparède, "Condensationen von Ketonen mit Aldehyden," *Berichte der deutschen chemischen Gesellschaft,* vol. 14, pp. 2460-2468, 1881.

[243] J. G. Schmidt, "Ueber die Einwirkung von Aceton auf Furfurol und auf Bittermandelöl bei Gegenwart von Alkalilauge," *Berichte der deutschen chemischen Gesellschaft,* vol. 14, pp. 1459-1461, 1881.

[244] C. P. Jasperse. *Aldol synthesis of dibenzalacetone.* web.mnstate.edu. Retrieved December 8, 2019, from http://web.mnstate.edu/jasperse/Chem365/Aldol%20Reaction.doc.pdf

[245] *Zaiput Flow Technologies Liquid-Liquid/Liquid-Gas Separators.* Zaiput.com. Retrieved December 18, 2019, from https://www.zaiput.com/technical_notes/Separators-TechnicalSpecification.pdf

[246] G. J. Janz and S. C. W. Jr., "Space ‐ Time Yield and Reaction Rate," *The Journal of Chemical Physics,* vol. 23, pp. 1550-1551, 1955.

[247] R. K. Henderson, J. Kindervater, and J. Manley, "Lessons learned through measuring green chemistry performance-The pharmaceutical experience," in *11th Annual Green Chemistry and Engineering Conference, Washington, DC*, 2007.

[248] *Sustainable Industry Classificaton System™ (SICS™) #HC0102,* S. A. S. Board.

[249] C. Jimenez-Gonzalez, C. S. Ponder, Q. B. Broxterman, and J. B. Manley, "Using the Right Green Yardstick: Why Process Mass Intensity Is Used in the Pharmaceutical Industry To Drive More Sustainable Processes," *Organic Process Research & Development,* vol. 15, pp. 912-917, 2011.

[250] K. Choi, B. D. Youn, and R.-J. Yang, *Moving least square method for reliability-based design optimization*, 2001.

[251] E. Loweth, G. de Boer, and V. Toropov, "Practical recommendations on the use of moving least squares metamodel building," in *13th International Conference on Civil, Structural and Environmental Engineering Computing*, Crete, Greece, 2011.

[252] Q.-S. Xu and Y.-Z. Liang, "Monte Carlo cross validation," *Chemometrics and Intelligent Laboratory Systems,* vol. 56, pp. 1-11, 2001.

[253] R. R. Picard and R. D. Cook, "Cross-Validation of Regression Models," *Journal of the American Statistical Association,* vol. 79, pp. 575-583, 1984.

[254] P. Zhang, "Model Selection Via Multifold Cross Validation," *The Annals of Statistics,* vol. 21, pp. 299-313, 1993.

[255] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of ICNN'95 - International Conference on Neural Networks*, 1995, pp. 1942-1948 vol.4.

[256] Z. Wang and G. P. Rangaiah, "Application and Analysis of Methods for Selecting an Optimal Solution from the Pareto-Optimal Front obtained by Multiobjective Optimization," *Industrial & Engineering Chemistry Research,* vol. 56, pp. 560-574, 2017.

[257] M. D. Morris and T. J. Mitchell, "Exploratory designs for computational experiments," *Journal of Statistical Planning and Inference,* vol. 43, pp. 381-402, 1995.

[258] M. E. Hosea and L. F. Shampine, "Analysis and implementation of TR-BDF2," *Applied Numerical Mathematics,* vol. 20, pp. 21-37, 1996.

[259] "Bias Variance Decomposition," in *Encyclopedia of Machine Learning*, C. Sammut and G. I. Webb, Eds., ed Boston, MA: Springer US, 2010, pp. 100-101.

[260] H. Rocha, "On the selection of the most adequate radial basis function," *Applied Mathematical Modelling,* vol. 33, pp. 1573-1583, 2009.

[261] "Constructing a Surrogate," in *Engineering Design via Surrogate Modelling*, ed, pp. 33-76.

[262] Z. Wu, "Compactly supported positive definite radial functions," *Advances in Computational Mathematics,* vol. 4, p. 283, 1995.

[263] H. Fang and M. F. Horstemeyer, "Global response approximation with radial basis functions," *Engineering Optimization,* vol. 38, pp. 407-424, 2006.

[264] A. J. Lewry and J. Williamson, "The setting of gypsum plaster," *Journal of Materials Science,* vol. 29, pp. 6085-6090, 1994.

[265] A. E. S. Van Driessche, L. G. Benning, J. D. Rodriguez-Blanco, M. Ossorio, P. Bots, and J. M. García-Ruiz, "The Role and Implications of Bassanite as a Stable Precursor Phase to Gypsum Precipitation," *Science,* vol. 336, p. 69, 2012.

[266] H. Weiss and M. F. Bräu, "How Much Water Does Calcined Gypsum Contain?," *Angewandte Chemie International Edition,* vol. 48, pp. 3520-3524, 2009.

[267] J. P. Ingham, "8 - Mortar, plaster, and render," in *Geomaterials Under the Microscope*, J. P. Ingham, Ed., ed Boston: Academic Press, 2013, pp. 137-162.

[268] B. Guan, G. Jiang, Z. Wu, J. Mao, and B. Kong, "Preparation of α-Calcium Sulfate Hemihydrate from Calcium Sulfate Dihydrate in Methanol–Water Solution under Mild Conditions," *Journal of the American Ceramic Society,* vol. 94, pp. 3261-3266, 2011.

[269] N. B. Singh and B. Middendorf, "Calcium sulphate hemihydrate hydration leading to gypsum crystallization," *Progress in Crystal Growth and Characterization of Materials,* vol. 53, pp. 57-77, 2007.

[270] U. Tritschler, M. Kellermeier, C. Debus, A. Kempter, and H. Cölfen, "A simple strategy for the synthesis of well-defined bassanite nanorods," *CrystEngComm,* vol. 17, pp. 3772-3776, 2015.

[271] M. A. Levenstein, C. Anduix-Canto, Y.-Y. Kim, M. A. Holden, C. González Niño, D. C. Green*, et al.*, "Droplet Microfluidics XRD Identifies Effective Nucleating Agents for Calcium Carbonate," *Advanced Functional Materials,* vol. 29, p. 1808172, 2019.

[272] P. Kraft, A. Bergamaschi, C. Broennimann, R. Dinapoli, E. F. Eikenberry, B. Henrich*, et al.*, "Performance of single-photon-counting PILATUS detector modules," *Journal of synchrotron radiation,* vol. 16, pp. 368-375, 2009.

[273] S. Shian and K. H. Sandhage, "A gas-tight Cu Kα x-ray transparent reaction chamber for high-temperature x-ray diffraction analyses of halide gas/solid reactions," *Review of Scientific Instruments,* vol. 80, p. 115108, 2009.

[274] M. Basham, J. Filik, M. T. Wharmby, P. C. Y. Chang, B. El Kassaby, M. Gerring*, et al.*, "Data Analysis WorkbeNch (DAWN)," *Journal of synchrotron radiation,* vol. 22, pp. 853-858, 2015.

[275] J. Filik, A. W. Ashton, P. C. Y. Chang, P. A. Chater, S. J. Day, M. Drakopoulos*, et al.*, "Processing two-dimensional X-ray diffraction and small-angle scattering data in DAWN 2," *Journal of Applied Crystallography,* vol. 50, pp. 959-966, 2017.

[276] R. T. Downs and M. Hall-Wallace, "The American Mineralogist crystal structure database," *American Mineralogist,* vol. 88, pp. 247-250, 2003.

[277] W. H. Bragg and W. L. Bragg, "The reflection of X-rays by crystals," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character,* vol. 88, pp. 428-438, 1913.

[278] P. F. Schofield, K. S. Knight, and I. C. Stretton, "Thermal expansion of gypsum investigated by neutron powder diffraction," *American Mineralogist,* vol. 81, pp. 847-851, 1996.

[279] P. Ballirano, A. Maras, S. Meloni, and R. Caminiti, "The monoclinic I2 structure of bassanite, calcium sulphate hemihydrate (CaSO4 · 0.5H2O)," *European Journal of Mineralogy,* vol. 13, pp. 985-993, 2001.

[280] X. Ni, H. Jian, and A. W. Fitch, *CFD modelling of flow patterns in an oscillatory baffled column* vol. 57, 2002.

[281] P. Gough, X. Ni, and K. C. Symes, "Experimental Flow Visualisation in a Modified Pulsed Baffled Reactor," *Journal of Chemical Technology & Biotechnology,* vol. 69, pp. 321-328, 1997.

[282] D. González-Juárez, J. P. Solano, R. Herrero-Martín, and A. P. Harvey, "Residence time distribution in multiorifice baffled tubes: A numerical study," *Chemical Engineering Research and Design,* vol. 118, pp. 259-269, 2017.

[283] A. N. Phan and A. P. Harvey, "Characterisation of mesoscale oscillatory helical baffled reactor—Experimental approach," *Chemical Engineering Journal,* vol. 180, pp. 229-236, 2012.

[284] X. Ni, M. R. Mackley, A. P. Harvey, P. Stonestreet, M. H. I. Baird, and N. V. R. Rao, "Mixing through oscillations and pulsations - A guide to achieving process enhancements in the chemical and process industries," *Chemical Engineering Research & Design,* vol. 81, pp. 373-383, 2003.

[285] M. Manninen, E. Gorshkova, K. Immonen, and X. W. Ni, "Evaluation of axial dispersion and mixing performance in oscillatory baffled reactors using CFD," *Journal of Chemical Technology & Biotechnology,* vol. 88, pp. 553-562, 2013.

[286] A. Mazubert, D. F. Fletcher, M. Poux, and J. Aubin, "Hydrodynamics and mixing in continuous oscillatory flow reactors—Part I: Effect of baffle geometry," *Chemical Engineering and Processing: Process Intensification,* vol. 108, pp. 78-92, 2016.

[287] A. Mazubert, D. F. Fletcher, M. Poux, and J. Aubin, "Hydrodynamics and mixing in continuous oscillatory flow reactors—Part II: Characterisation methods," *Chemical Engineering and Processing: Process Intensification,* vol. 102, pp. 102-116, 2016.

[288] D. Reay, C. Ramshaw, and A. Harvey, "Compact and Micro-heat Exchangers," in *Process Intensification (Second Edition)*, D. Reay, C. Ramshaw, and A. Harvey, Eds., ed Oxford: Butterworth-Heinemann, 2013, pp. 91-120.

[289] M. R. Chapman, Y. M. Shafi, N. Kapur, B. N. Nguyen, and C. E. Willans, "Electrochemical flow-reactor for expedient synthesis of copper–N-heterocyclic carbene complexes," *Chemical Communications,* vol. 51, pp. 1282-1284, 2015.

[290] M. R. Chapman, S. C. Cosgrove, N. J. Turner, N. Kapur, and A. J. Blacker, "Highly Productive Oxidative Biocatalysis in Continuous Flow by Enhancing the Aqueous Equilibrium Solubility of Oxygen," *Angewandte Chemie International Edition,* vol. 57, pp. 10535-10539, 2018.

[291] X. Ni, G. Brogan, A. Struthers, D. C. Bennett, and S. F. Wilson, "A Systematic Study of the Effect of Geometrical Parameters on Mixing Time

in Oscillatory Baffled Columns," *Chemical Engineering Research and Design,* vol. 76, pp. 635-642, 1998.

[292] N. M. F. Reis, "Novel Oscillatory Flow Reactors for Biotechnological Applications," PhD Thesis, Escola de Engenharia, Universidade do Minho, 2006.

[293] B. Finlayson, A. Aditya, V. Brasher, L. Dahl, H. Quan Dinh, A. Field*, et al.*, "Mixing of Liquids in Microfluidic Devices," presented at the COMSOL Conference, Boston, USA, 2008.

[294] P. M. Doran, *Bioprocess engineering principles / Pauline M. Doran*: London ; San Diego : Academic Press, [1995] ©1995, 1995.

[295] E. L. Cussler, *Diffusion: Mass Transfer in Fluid Systems*, 3 ed. Cambridge: Cambridge University Press, 2009.

[296] B. Armaly, Durst, F., Pereira, J., & Schönung, B., "Experimental and theoretical investigation of backward-facing step flow," *Journal of Fluid Mechanics,* vol. 127, pp. 473-496, 1983.

[297] P. Sung-Jin, K. Jung Kyung, P. Junha, C. Seok, C. Chanil, and C. Jun Keun, "Rapid three-dimensional passive rotation micromixer using the breakup process," *Journal of Micromechanics and Microengineering,* vol. 14, p. 6, 2004.

[298] A. D. Stroock, S. K. W. Dertinger, A. Ajdari, I. Mezić, H. A. Stone, and G. M. Whitesides, "Chaotic Mixer for Microchannels," *Science,* vol. 295, pp. 647-651, 2002.

[299] R. H. Liu, M. A. Stremler, K. V. Sharp, M. G. Olsen, J. G. Santiago, R. J. Adrian*, et al.*, "Passive mixing in a three-dimensional serpentine microchannel," *Journal of Microelectromechanical Systems,* vol. 9, pp. 190-197, 2000.

[300] B. Stoeber, D. Liepmann, and S. J. Muller, "Strategy for active mixing in microdevices," *Physical Review E,* vol. 75, p. 066314, 2007.

[301] Y. Li, R. K. Reddy, C. S. S. R. Kumar, and K. Nandakumar, "Computational investigations of the mixing performance inside liquid slugs generated by a microfluidic T-junction," *Biomicrofluidics,* vol. 8, p. 054125, 2014.

[302] A. Hashmi and J. Xu, "On the Quantification of Mixing in Microfluidics," *Journal of Laboratory Automation,* vol. 19, pp. 488-491, 2014.

[303] O. Levenspiel and W. K. Smith, "Notes on the diffusion-type model for the longitudinal mixing of fluids in flow," *Chemical Engineering Science,* vol. 6, pp. 227-235, 1957.

[304] G. F. Knott, M. R. Mackley, and M. J. Lighthill, "On eddy motions near plates and ducts, induced by water waves and periodic flows," *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences,* vol. 294, pp. 599-623, 1980.

[305] C. R. Brunold, J. C. B. Hunns, M. R. Mackley, and J. W. Thompson, "Experimental observations on flow patterns and energy losses for oscillatory flow in ducts containing sharp edges," *Chemical Engineering Science,* vol. 44, pp. 1227-1244, 1989.

[306] R. T. Haftka, D. Villanueva, and A. Chaudhuri, "Parallel surrogate-assisted global optimization with expensive functions – a survey," *Structural and Multidisciplinary Optimization,* vol. 54, pp. 3-13, 2016.

[307] M. Kuhn and K. Johnson, *Applied predictive modeling*. New York: Springer, 2013.

[308] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An introduction to statistical learning : with applications in R*, 2013.

Appendix A
Technical Drawing of a fReactor 1 Inlet 1 Outlet

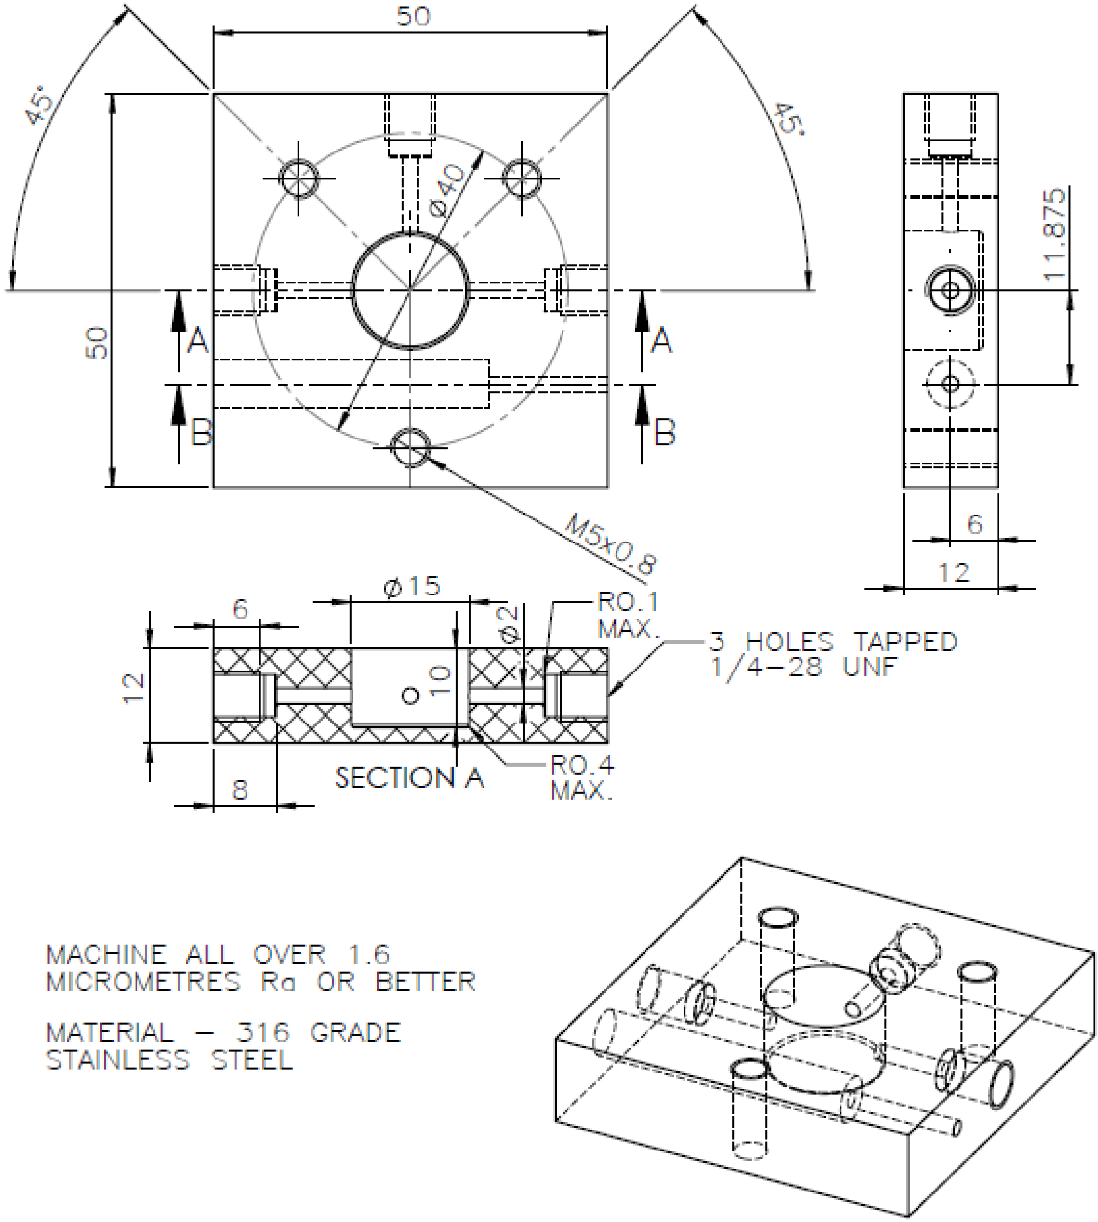

MACHINE ALL OVER 1.6
MICROMETRES Ra OR BETTER

MATERIAL — 316 GRADE
STAINLESS STEEL

Appendix B
Technical Drawing of a fReactor 2 Inlets 1 Outlet

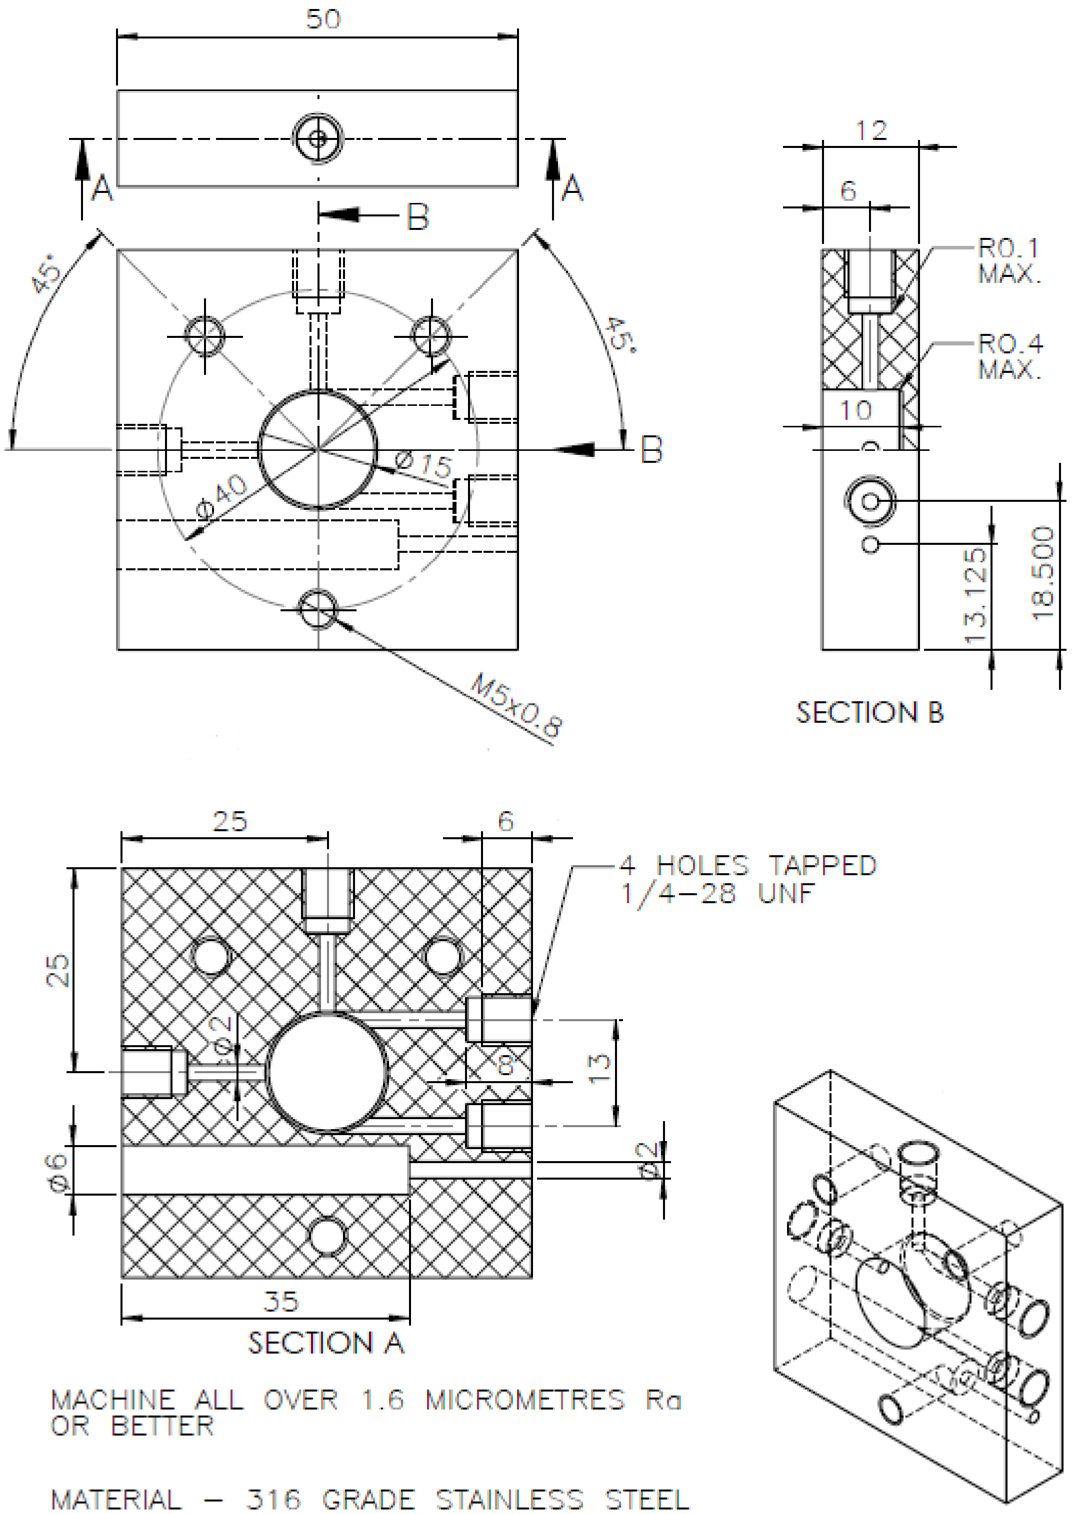

MACHINE ALL OVER 1.6 MICROMETRES Ra
OR BETTER

MATERIAL – 316 GRADE STAINLESS STEEL

Appendix C
Arduino Controller's Code

A flow chart of the programme is shown here that illustrates all actions carried out in order to implement the temperature controller, including all the features commented in the Chapter 3 of this thesis. The chart is divided into six different sections, for which excerptions from the code will be commented in the corresponding sections of this appendix.

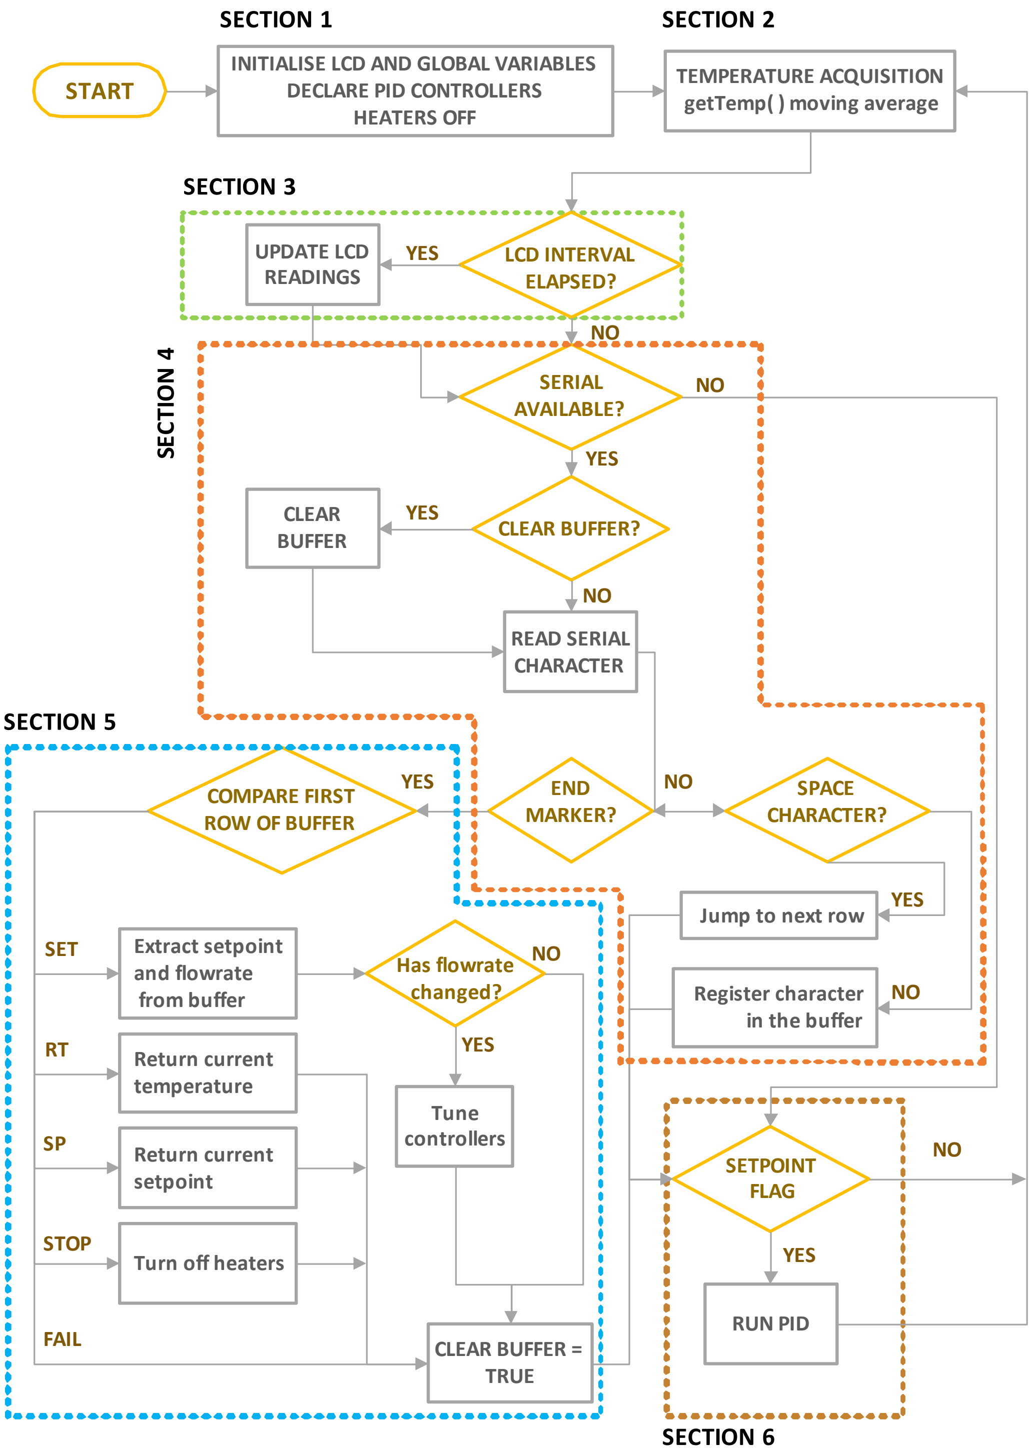

1. INITIALISATION

A number of libraries apart from Beauregard's PID library were used that enabled control over the LCD screen. The required number of pins needed to run the LCD screen is of 6 for screen control and an extra 3 for RGB backlight control, amounting to a total of 9 pins. However, the number of pins on the Arduino UNO board is quite constrained. In consequence, 1^{2}C communication protocol is used, which enables communication over two wires. In order to enable LCD control, three different libraries are needed:

- Wire.h: enables inter-integrated circuit 1^{2}C communication, a widely used communication protocol. This kind of communication only uses two lines, namely SDA and SCL, which are the serial data and serial clock, respectively. The data is transferred over the serial data line, while the clock line is used to synchronize the data transfer by sharing the same clock signal between the master and the slave devices. It will be used to communicate with the MCP2307 port expander chip of the LCD shield.
- Adafruit_MCP2307.h: used to run the MCP2307 port expander chip included in the LCD shield by using 1^{2}C communication. It is a 16-bit 1^{2}C I/O port expander. It makes use of the previously defined Wire.h.
- Adafruit_RGBLCDShield.h: This library includes the commands to manipulate the content and appearance of the LCD screen. It makes use of the two previous libraries to enable commands from Arduino to make use of 12C communication with the port expander to run the LCD screen.

In addition to these libraries, and as mentioned previously, an extra one is needed in order to implement the PID control:

- PID_v1: intensively used PID control library developed by Brett Beauregard. This library includes relevant features for the controller such as on-the fly tuning changes and reset window mitigation.

In the next lines, the libraries are included, the LCD screen is declared, and the white colour is defined for the screen. The interval at which the screen is intended to update with new temperature values is also defined as 1 s, since updating it in every loop will make the numbers in the screen change constantly, making it unreadable.

```
#include <Wire.h>                                 //12C COM with screen
#include <Adafruit_MCP23017.h>                    //LCD CHIP
#include <Adafruit_RGBLCDShield.h>                //LCD functions
#include  <PID_v1.h>                                //Brett  Beauregard
PID library
Adafruit_RGBLCDShield lcd = Adafruit_RGBLCDShield();
#define WHITE 0x7
long interval = 1000;                    //Interval between LCD updates
```

Next, the global variables relative to the serial communication are declared. *MAX_SIZE_COMMAND* and *MAX_NUM_PAREMETERS* are the number of elements in the *command* array that is used to store the incoming command string. They are also the number of rows and columns of the *commands_char* matrix that is used to split the received command into its different parts as explained in section 2.4.5. The integers count and command serve as indexes for *commands_char* when the characters are being split. Finally, the Boolean variable *clearBuffer* is initialized as 'true'. The state of this variable will be changed in the code whenever the flush of *commands_char* is needed.

```
#define MAX_SIZE_COMMAND 12
#define MAX_NUM_PARAMETERS 5
int count = 0;
char command[MAX_SIZE_COMMAND];
char commands_char[MAX_NUM_PARAMETERS][MAX_SIZE_COMMAND];
int ncommand = 0;
bool clearBuffer = true;
```

Later, the PID controllers for the three heated reactors are declared, and a boolean variable *SetPointFlag* will be used to track if the setpoint has already been introduced by the user to minimise computational consumption of resources and avoid running through undesired parts of code as long as the setpoint is remains unchanged.

```
double Setpoint, Input, Output;
double Input2, Output2;
double Input3, Output3;

PID Temp_PID(&Input, &Output, &Setpoint, 115, 0.04, 0, DIRECT);
PID Temp_PID2(&Input2, &Output2, &Setpoint, 115, 0.04, 0, DIRECT);
PID Temp_PID3(&Input3, &Output3, &Setpoint, 115, 0.04, 0, DIRECT);

bool Setpointflag = false;
```

The next section consists in the declaration of the variables needed in order to implement the oversampling and averaging routine. The times between both LCD and the serial communication updates is controlled through the variables *previousMillis* and *previousMillisSerial*, respectively. The unsigned long type is used to store large numerical data. With a size of 32 bits, it allows storing numbers up to 2^{32}. Since the time variables are expressed in milliseconds, this format was used to allow manipulation of the large numbers that arise during the continuous operation of the device.

```
#define RASIZE 75                       // Number of samples for o&a
double Ti;
int A[RASIZE][4];                       // Matrix to hold readings
bool start[4]={true,true,true,true};    // Flags for getTemp
double total[4] = {0,0,0,0};            // Sum of the temperature values
int index[4] = {0,0,0,0};               // Array index
int stindex[4] = {0,0,0,0};             // Array index until it is full

unsigned long previousMillis = 0;       // Time between LCD updates
unsigned long previousMillisSERIAL = 0; //Time between COM updates

unsigned long currentMillis;            // Time
unsigned long Tzero;                    // Beginning time

float previousfr = 0;                   // fr of previous experiment
```

The unsigned long variable *currentMillis* is the time value updated for each loop. The float variable *previousfr* is used to store the values of the previous flow rate fed into the reactors in order to compare it with the current one to check if changes in the PID tuning need to be carried out in order to cope with changing flowrates.

In the language developed for Arduino IDE, setup() is a section of code that is run once. In this case the section is used to start the serial communication and place the thermocouple numbers in position in the LCD screen. Updating the LCD is time consuming, and as these indexes do not change their position during operation, it is preferable to do it once in this section instead of iteratively rewriting them. In addition to this operation, the PID controllers are set to automatic mode, and turned off.

```
void setup()
{
  analogReference(EXTERNAL);           //External reference 3.3 V
  Serial.begin(9600);                  //Baud rate

  lcd.begin( 16, 2 ); lcd.print("1:");
  lcd.setCursor(8,0); lcd.print("2:");
  lcd.setCursor(0,1); lcd.print("3:");
  lcd.setCursor(8,1); lcd.print("4:");

  Temp_PID.SetMode(AUTOMATIC);
  Temp_PID2.SetMode(AUTOMATIC);
  Temp_PID3.SetMode(AUTOMATIC);
  Tzero = millis();

  analogWrite(3,0);
  analogWrite(5,0);
  analogWrite(6,0);
  analogWrite(9,0);
}
```

2. TEMPERATURE ACQUISITION

The next section of the code written under loop() is iterative. The loop starts by acquiring temperature values from the thermocouples using *getTemp*, a function developed to this end and previously described in section 0. These values are rounded to the closest half for their display on the LCD screen.

```
void loop()
{

  double Temp[4];
  float rounded[4];

  for(int i=0; i<4; i++)
   {
      Temp[i] = getTemp(i);    //getTemp 0 or 1
      rounded[i] = (int)((Temp[i]*2.0)+0.5)/2.0;
   }
  currentMillis = millis();
```

3. LCD UPDATING

Next, *currentMillis* is updated with the current time since the board was turned on and compared with *previousMillis* to ascertain it the difference between these is equal or higher than the interval between LCD updates defined previously by the user. If this is the case, the LCD is updated with the rounded values and *previousMillis* equalled to *currentMillis* to continue quantifying the time elapsed in the next loop.

```
while(currentMillis-previousMillis >= interval)
 {
    lcd.setCursor(2,0);  lcd.print(rounded[0]);
    lcd.setCursor(10,0); lcd.print(rounded[1]);
    lcd.setCursor(2,1);  lcd.print(rounded[2]);
    lcd.setCursor(10,1)  lcd.print(rounded[3]);

    previousMillis = currentMillis;
  }
```

4. SERIAL COMMUNICATION

The next part of the code enables serial communication between Arduino and the master computer. The character *endMarker* is set to Matlab's fprintf terminator '\n' and *rc* is the variable used to contain the incoming character, which are read one at a time. If there are characters in the serial input buffer and the *clearbuffer* flag is true, the initial matrix used to store information sent over the USB port is cleared, overwriting every location with Arduino's null character '\0'. This is done after receiving the serial data completely, to leave a clear commands buffer for the next command. The variables *count* and *ncommand*, which are the column and row indexes of the commands matrix are then restored to zero and the state of *clearbuffer* changed to false to register the data when the next command arrives.

```
char endMarker = '\n';
char rc;

if (Serial.available() > 0)
 {
    if (clearBuffer == true)
     {
       for( int i = 0; i < MAX_NUM_PARAMETERS; i++ )
          {
            for( int j = 0; j < MAX_SIZE_COMMAND; j++ )
            commands_char[i][j] = '\0';
          }

        count = 0;
        ncommand = 0;
        clearBuffer = false;
     }
```

The expected command structure consists of a first identifier of the task; SET implies a new setpoint will be received, RT requests the temperature, SP requests the current setpoint and STOP turns off all the heaters. Of these, only SET comes followed by other characters in the string. After a blank space, the new setpoint temperature and again after another blank space, the flow rate of the experiment to be carried out at this temperature. At the end of the string the character ‘\n’ informs that the transmission has finished. In every iteration, while there are still characters in the incoming buffer, *rc* is updated with the next character in the buffer. If this is different from Matlab’s terminator and also different from a space character, the command matrix is updated with this character and its position within the same command incremented by one (*count++*). If it is a space character, the terminator ‘\0’ is written, the column counter is set to zero (beginning of new row in the matrix) and the row counter is incremented by one (*ncommand ++*) in order to move down to the next row of the buffer matrix.

```
rc = Serial.read();

if (rc != endMarker)
 {
   if(rc != ' ')
    {
      commands_char[ncommand][count] = rc;
      count++;
     }
    else
     {
       commands_char[ncommand][count] = '\0';
       count = 0;
       ncommand++;
      }
}      \\ end of if(rc!=endMarker){
```

5. STATE MACHINE

If, on the other hand, the character received is Matlab’s terminator ‘\n’, meaning that the whole command has been received, the command in the first row of the buffer matrix is compared against the four possibilities contemplated here. The first row determines the kind of action that is to be taken upon *commands_char[0]*. SET changes the *setpointflag* to true (this is a flag used to ensure PID computations do not take place until the user or master sends a SET command). The three controllers are set to manual in order to change the tuning parameters; the set temperature can be accessed in *commands_char[1]* and the flowrate of the experiment in *commands_char[2]* sent from the computer and Matlab. Once the data has been successfully read a string is sent back to the master to let it know the communication took place successfully. Then the flowrate is compared against a series of values to enable dynamic tuning, setting the appropriate K_p and K_i values and returning the PID mode to automatic. The variable *previousflowrate* is updated with the flowrate read from the commands char to enable its comparison with a new flow rate the next time a SET command is received.

```
else
   {

    if( !strcmp("SET", commands_char[0]))
      {
         Setpointflag = true;
         double TKp = 0;
         double TKi = 0;

         Temp_PID.SetMode(MANUAL);
         Temp_PID2.SetMode(MANUAL);
         Temp_PID3.SetMode(MANUAL);
         Setpoint = atof (commands_char[1]);
         float flowrate = atof (commands_char[2]);
         Serial.println("FIN");

          if (previousfr!=flowrate || previousfr==flowrate)
            {
              if(0 <= flowrate && flowrate< 0.25)
                {
                   TKp = 20*flowrate+65;
                   TKi = 0.02;
                }
              else if (0.25 <= flowrate && flowrate < 0.5)
                {
                    TKp = 20*(flowrate-0.25)+70;
                    TKi = 0.028*(flowrate-0.25)+0.02;
                }
              else if (0.5 <= flowrate && flowrate < 1)
                {
                    TKp = (flowrate-0.5)+75;
                    TKi = 0.006*(flowrate-0.5)+0.027;
                }
              else if (1 <= flowrate && flowrate < 2)
                {
                     TKp = 10*(flowrate-1)+75.5;
                     TKi = 0.01*(flowrate-1)+0.03;
                }
              else if (2 <= flowrate && flowrate < 4)
                {
                     TKp = 14.75*(flowrate-2)+85.5;
                     TKi = 0.005*(flowrate-2)+0.04;
                }
              else if (flowrate >= 4)
                {
                     TKp = 115;
                     TKi = 0.05;
                }

              Temp_PID.SetTunings(TKp, TKi, 0);
              Temp_PID2.SetTunings(TKp, TKi, 0);
              Temp_PID3.SetTunings(TKp, TKi, 0);

              Temp_PID.SetMode(AUTOMATIC);
              Temp_PID2.SetMode(AUTOMATIC);
              Temp_PID3.SetMode(AUTOMATIC);
              previousfr = flowrate;

              }
```

RT is a request of the temperature by the computer, so an array of the four temperatures is sent back. SP returns the current setpoint. Finally, STOP turns the controllers to manual and turns off all heaters. If the received command does not correspond to any of the latter, the command 'Fail' is sent back to Matlab to let it know the communication was not successful. Then, the buffer is set to be cleared in the next loop (*clearbuffer=true*).

```
              else if( !strcmp("RT", commands_char[0]))
              {
                Serial.print(Temp[0]); Serial.print(" ");
                Serial.print(Temp[1]); Serial.print(" ");
                Serial.print(Temp[2]); Serial.print(" ");
                Serial.println(Temp[3]);
               }
               else if( !strcmp("SP", commands_char[0]))
               {
               Serial.println(Setpoint);
               }
               else if( !strcmp("STOP", commands_char[0]))
               {
                Serial.println("FIN");
                Setpointflag = false;
                Temp_PID.SetMode(MANUAL);
                Temp_PID2.SetMode(MANUAL);
                Temp_PID3.SetMode(MANUAL);
                Output = 0;  analogWrite(3,Output);
                Output2 = 0; analogWrite(5,Output2);
                Output3 = 0; analogWrite(6,Output3);
               previousfr = 0;
              }
      else
      {
      Serial.println("Fail");
      }


    clearBuffer = true;
    }    \\ end of else{ -> last character received was\n

}      \\ end of if(Serial.available>0){
```

6. PID COMPUTING

If *setpointflag* is true, this is, if the first SET command has been sent over the time of use of the device, the PID outputs are calculated and fed to the output pins. These signals, as commented in the previous section, will be amplified with the use of MOSFETS and fed the heaters for effective temperature control.

```
if (Setpointflag)
  {                //Only computes when setpoint available, stops after
STOP
    Input = Temp[0];
    Temp_PID.Compute();
    analogWrite(3,Output);

    Input2 = Temp[1];
    Temp_PID2.Compute();
    analogWrite(5,Output2);

    Input3 = Temp[2];
    Temp_PID3.Compute();
    analogWrite(6,Output3);
  }

}      // end of loop(){
```

Appendix D
Matlab Communication Routines Code

In this appendix, the most important functions developed for Matlab's GUI are reproduced and explained. It includes different functionalities that the main program can access and run automatically without human operation.

The first of these functionalities is used to create Matlab's serial object for Arduino. This is done by the functionality name *objectCommand*, which in this case points to the function *carlosobj* developed specially for the Arduino controller. The function returns the serial object accepting a COM port number as a parameter.

```
function serialObject = carlosobj(com, userData)

% checks the number of arguments
error(nargchk(1, 2, nargin))

% error handling
if ~iscomport(com)
    % COM ports 1-255 allowed only
    error('com must be an unsigned integer from 1 to 255')
end

% initialises variables
serialPort = sprintf('COM%d', com);

% if serial object does not exist, defines serial object and opens it
if isempty(instrfindall('Port', serialPort))
    % creates serial object
    serialObject = serial(  serialPort,...
        'BaudRate', 9600,...
        'DataBits', 8,...
        'Parity', 'none',...
        'StopBits', 1,...
        'Terminator', 'LF',...
        'TimeOut', 2);

    % if userData was provided, modify the object
    if nargin >= 2
        % changes the user data
        serialObject.UserData = userData;
    end

else
    % errors
    error('Object(s) already exist on this serial port')
end
```

The function simply checks that the COM port assigned to Arduino does not already contain another device, and then creates and configures a serial object for it. The port in which to create the serial object (that should be the port to which Arduino is connected) has to be specified by the user. The baud rate is set to 9600 bits/s, matching that of the specified in Arduino IDE, as well as the data bits 8 shared in both devices. A minimum of 7 bits are required to transmit ASCII characters, but the default

8 bits (1 byte) is conserved here. The terminator character is set to ‘LF’, or line feed, which corresponds to the newline character ‘\n’. This indicates the end of a series of bytes that form a command. Finally, the time out is set to two seconds. This is the time Matlab stays trying to read from the buffer without receiving any bytes.

The next command is used to connect the device and initialise communication with the serial object created by the previous function, using it as its argument.

```
function carlosobjconnect(serialObject)

% checks the number of input arguments
error(nargchk(1, 1, nargin))

% error handling
if ~isserial(serialObject) || isrunning(serialObject)
    % errors
    error('serialObject must be a valid closed serial object.')
end

    % opens communication
    fopen(serialObject)

% pause for Arduino
pause(2)
```

This simple function is limited to check the number of arguments and, if the device is not a serial object or if it is already running, displays an error message before opening the communication using *fopen*, which connects the serial port object to the device. After that, it gives two seconds to Arduino to initialise before sending any commands.

The function developed to read the current temperature of the device is *carlosobjcurrenttemp*. It reads the current temperature in Celsius from Arduino provided the serial port object to which it is connected.

```
function currentTemp = carlosobjcurrenttemp(tempObj)
% error handling
if ~nargin
    % errors
    error('Insufficient arguments')
end

while (true)
     fprintf(tempObj, '%s\n', 'RT');
     Tempvect = fscanf(tempObj, '%f');
        if isvector(Tempvect)
            break
        else
            pause(0.1);
            flushinput(tempObj);
        end
end
currentTemp = Tempvect(1);

if currentTemp > 300 || currentTemp < -100
    currentTemp = NaN;
end
```

After checking the number of arguments, a while loop sends the command RT to Arduino and reads its response until an array is sent back. If the communication has failed it pauses for 0.1 seconds and flushes the input buffer in order to avoid reading bits from a previous failed communication. Then it requests again the temperature array containing the temperatures for all the thermocouples. The temperature for the first reactor is selected as the current temperature, since this is the most adverse value and takes the longer to reach a setpoint. Then, if the value is over 300 or below -100 °C, symptomatic of a broken connection in the thermocouple wiring, the value is discarded and set to NaN.

Finally, a function is needed to provide new temperature set points to the controller. This function is called *carlosobjwritesettemp*. It counts with two arguments: the first for the serial object and the second for the set point value. However, in order to enable dynamic control of the fReactors, a value for the flow rate must be provided to modify the PID tuning according to it. This is achieved by means of accessing a txt file (*FLOWRATE.txt*) in which the total flow rate provided by the pumps is saved periodically.

```
function carlosobjwritesettemp(tempObj, setTemp)

% checks the number of input arguments
error(nargchk(2, 2, nargin))

% error handling
if setTemp < 20 || setTemp > 150
    % errors
    error('Set temperature must be between room temperature and 150 °C')
end

 %imports flow rate from txt
 flowrate = importdata('FLOWRATE.txt');
 %checks flowrate is in double format
    if ~isa(flowrate,'double')
        error('The number in FLOWRATE.txt is not in double format');
    end

while(true)
    % sends command to Arduino
    fprintf(tempObj,'%s',   sprintf('SET   %.4f   %.4f\n',   setTemp,
flowrate));
    res = fscanf(tempObj,'%s');
      if strcmp(res,'FIN')
          break
      else
          pause(0.1);
          flushinput(tempObj)
      end
end
```

Here, the number of arguments is checked as usual. The temperature must be set between the limits of 20 and 150 °C for it to work. The current flow rate through the system is enquired from *FLOWRATE.txt*, checking that the number is in double format.

Then, a command is sent to Arduino that includes both the setpoint and the flowrate in the way specified previously in this sub-chapter. A response is expected from Arduino every time the set point is adjusted successfully. If this corresponds to the expected successful command (*FIN*) the function exits, while if this has not been received a delay is introduced before flushing the input buffer and repeating again the operation.

The same structure is used to require the current setpoint or stop the controller from heating, although these commands do not require any arguments after the identifier of the operation that are *SP* and *STOP*, respectively.

Appendix E
Aldol DoE Data

NaOH eq.	Residence time (min)	T (ºC)	YIELD (%)	STY (g/(L×min))	Cost (£/kg)	PMI
0,3	2,5000	50	0,0	0,0	∞	∞
0,75	2,5000	50	2,4	0,6	1356,7	651,6
0,525	6,2500	50	28,5	1,9	198,6	78,0
0,3	10,0000	50	3,0	0,2	1808,1	491,2
0,75	10,0000	50	2,3	0,1	1320,9	634,4
0,4125	4,3750	62,5	11,5	1,5	421,0	142,1
0,6375	4,3750	62,5	1,7	0,6	788,3	346,7
0,4125	8,1250	62,5	33,0	2,2	151,0	51,0
0,6375	8,1250	62,5	9,1	0,6	411,6	181,0
0,525	2,5000	75	19,8	4,5	204,1	80,2
0,3	6,2500	75	13,7	1,6	347,8	94,5
0,525	6,2500	75	15,3	1,5	241,0	94,7
0,75	6,2500	75	68,2	6,7	45,3	21,8
0,525	10,0000	75	84,6	6,1	38,3	15,0
0,4125	4,3750	87,5	63,9	10,4	59,7	20,1
0,6375	4,3750	87,5	51,9	7,8	60,6	26,6
0,4125	8,1250	87,5	82,5	7,8	43,0	14,5
0,6375	8,1250	87,5	58,3	4,7	54,1	23,8
0,3	2,5000	100	27,3	8,5	159,5	43,3
0,75	2,5000	100	76,8	22,4	33,6	16,2
0,525	6,2500	100	94,0	12,7	29,3	11,5
0,3	10,0000	100	60,7	4,9	70,1	19,0
0,75	10,0000	100	64,5	4,4	43,1	20,7

Appendix F
S$_N$Ar DoE Data

Residence time (min)	Morpholine eq.	Temperature ºC	Yield %
17,77	0,56	60,57	46,32
4,96	1,07	88,00	78,64
6,63	1,41	74,29	81,78
2,73	1,24	94,86	79,37
10,53	0,87	91,43	70,53
19,44	0,61	89,71	48,10
20,00	0,73	52,00	61,71
15,54	1,36	82,86	81,41
14,43	0,70	81,14	56,57
7,19	1,10	41,71	75,06
16,66	0,76	98,29	62,72
8,86	0,67	55,43	55,41
14,99	1,50	64,00	81,64
9,97	1,13	96,57	79,01
11,09	1,04	76,00	80,79
7,74	0,79	79,43	64,65
0,50	1,30	77,71	56,56
9,41	1,19	65,71	82,18
12,76	0,81	69,14	67,79
3,29	0,84	43,43	57,64
1,06	0,59	67,43	36,60
13,87	1,39	40,00	84,47
5,51	0,93	62,29	72,02
1,61	1,16	58,86	63,99
16,10	0,96	45,14	77,27
17,21	0,90	84,57	74,14
8,30	1,44	100,00	80,56
6,07	1,33	53,71	79,55
12,20	0,53	48,57	45,71
18,33	1,21	93,14	81,79
3,84	0,64	86,29	51,59
11,64	1,47	50,29	83,55
13,31	1,27	57,14	83,08
18,89	0,99	70,86	80,64
4,40	0,50	46,86	40,43
2,17	1,01	72,57	69,92