

Industrial Applications of Hybrid Modelling Techniques

Ben Dannatt

A thesis presented for the degree of
Doctor of Engineering

December 2018

Biopharmaceutical Bioprocessing Technology Centre
School of Chemical Engineering and Advanced Materials
Faculty of Science, Agriculture and Engineering



Abstract

In the present study, the application of hybrid modelling techniques is applied to industrial applications. Many of the studies currently known to the literature for the fields under examination are either purely model-based, theory-based or lab/pilot scale empirical studies. In this work, we present a hybrid approach whereby empirical data is used to form statistical models for relationships where no clear fundamental relationship can be described mathematically. Equally, first-principles models are employed where no suitable data can be gathered empirically. Finally, the process understanding, heuristics and recollections of plant operators, engineers and maintenance personnel can be integrated formally into the decision-making process of process design/optimisation.

The first half of this work is concerned with process development of a proprietary modular Gas-to-Liquids process, briefly comprised of a packed bed plate-fin 'mini-channel' Fischer-Tropsch reactor. Currently, little can be predicted about the flow or temperature performance of a complex reactor geometry in the design phase. Data-driven models provide a simplistic approximation with no added understanding. At commercially relevant scales, the parameters of interest are both costly and hazardous to iterate through empirical trial and improvement. By integrating offline analysis, online data and a novel temperature sensing scheme, we increase the spatiotemporal resolution of data while adding process understanding.

The second theme of this work is related to flue gas filtration in large-scale Biomass and Energy-from-Waste Power Generation plants. Flue gas filtration is overlooked as an opportunity for process improvement. We argue that a filtration system designed on the basis of lowest CAPEX, and operated at the lowest maintenance cost will not provide the lowest total cost of ownership. By integrating industrial historic data, maintenance records, commercial data and multivariate modelling methods, we produce a set of recommendations for improved operation. Commercially available solutions are benchmarked in predictive hybrid models on a ROI basis.

*Dedicated
to my Mam.*

Acknowledgements

The Author sincerely thanks the team at CompactGTL, particularly; Chris Bonner, Brian Keenan, Rob Knowles, Dr. Bob Peat, Dr. Hardi Koortzen, Dr. Joep Font Freide, Keith Bowen, John Shoyode, Asit Barchha, Lee Taylor and also Dr. Ross Morgan for his supervision. Thanks are also due to Svetlin Staykov at Silixia for his hard work. To the operations team at CompactGTL, I give you my deepest thanks and respect. I wish you good luck in your new challenges.

At Durham Filtration, I owe an immense debt of gratitude to Barry Goulden. I also thank Howard Nicholson for his guidance and support. At Lydall, thanks to Julian Connor and Helen Brain for their expertise and hard work. At Sembcorp, I thank Chris Plews, Dave Thompson, Gail Reilly and Nick Tann. I must also thank Vahid Mirsaiidi and Roger Taylor for their support.

At Newcastle University, The Author thanks Dr. Mark Willis for his supervision and Nikki Hawley for her pastoral care and support. Thanks to Dr. Chris O'Malley for his advice. Thanks also to Dr. Isabel Arce-Garcia and Richard Baron at ACMA for their assistance. I must also sincerely thank Prof. Gary Montague for 'sorting me out' on more than one occasion.

To the Master's students who worked with me; Hannah Eagle, Liam Day, Calum Taylor and Amelia Pettit, I congratulate you on your success, thank you for your contribution and wish you every success in your future careers.

The Author wishes to acknowledge the Engineering and Physical Sciences Research Council (EPSRC) for part funding this work (grant number EP/G037620/1).

The Author wishes to acknowledge Innovate UK for part funding this work (KTP number 010523).

Preface

This thesis describes research that was undertaken as part of an Engineering Doctorate in Process Development which was carried out in collaboration with CompactGTL Limited and Durham Filtration Limited and funded by the Engineering and Physical Sciences Research Council (EPSRC) and Innovate UK.

The thesis takes the format of a thesis by portfolio, which details a number of projects that are linked by the purpose of providing enhanced operability by implementing hybrid modelling methods.

Being an industrially focussed Engineering Doctorate, the projects reflect the research requirements of industry, and changed over the period of study to meet new research challenges within the organisations involved.

Contents

1	Introduction	1
1.1	Hybrid Modelling	1
1.1.1	Industrial Approach to Process Improvement	1
1.1.2	Illustrative Example - Industrial Approach	3
1.1.3	Illustrative Example - Hybrid Modelling Approach	4
1.1.4	The Role of Computation	5
1.2	Scope and Aims	6
1.2.1	Gas-To-liquids Process Development	6
1.2.2	Biomass Flue Gas Filtration Optimisation	7
1.3	Structure and Major Contributions of the Thesis	10
1.3.1	General Impact	10
1.3.2	Chapter synopsis and impact	11
2	An Overview of Modular Gas-to-liquids Technology	14
2.1	Chapter Abstract	14
2.2	Gas to Liquids Overview	14
2.2.1	Typical GTL Process	14
2.2.2	Fischer-Tropsch Synthesis	15
2.2.3	Small Scale Modular GTL	17
2.2.4	Compact Fischer-Tropsch Reactors and Process	19
3	Linking Operational History and Post-operational Reactor Condition	25
3.1	Chapter Abstract	25
3.2	Experimental Methods	26
3.2.1	Fischer-Tropsch Synthesis	26
3.2.2	Determination of pre- and post-operational differential pressure	28
3.2.3	Reactor Unloading	29
3.2.4	CHN Analysis	30
3.2.5	Scanning Electron Microscopy	32
3.3	Results and Discussion	32
3.3.1	FTS Operational History	32
3.3.2	Pre- and Post Operational Differential Pressure	35
3.3.3	Catalyst Unloading & Blockages	44

3.3.4	CHN Analysis	49
3.3.5	SEM analysis	58
3.3.6	In situ Endoscopy	63
3.4	Chapter Summary	67
4	Characterisation of Next-generation Compact Fischer-Tropsch-Tropsch Packed Beds Using X-Ray Tomography	68
4.1	Chapter Abstract	68
4.2	Materials and Methods	68
4.2.1	Mock Reactor	68
4.2.2	Image processing	69
4.3	Results & Discussion	72
4.3.1	Packing Parameters	72
4.3.2	Particle Geometric Properties	72
4.4	Chapter Summary	82
5	High Resolution Distributed Temperature Sensing of Modular Fischer-Tropsch-Tropsch Reactors	83
5.1	Chapter Abstract	83
5.2	Introduction	84
5.2.1	Distributed Temperature Sensing Principles	84
5.2.2	Fibre Optic Principles	86
5.3	Installation of Fibre Optics into a modular FT reactor	98
5.3.1	Pilot scale reactor	98
5.3.2	Lab-scale reactor	100
5.3.3	Software & Hardware Setup	102
5.4	Experimental Methods	110
5.4.1	Fischer-Tropsch Synthesis and Reactor Operation	110
5.4.2	Data Collection	111
5.5	Results and Discussion	112
5.5.1	The Scale and Shape of Data	112
5.5.2	Relating Linear Temperature to Real Space	115
5.5.3	End looping offsets and the implications on scale-up	118
5.6	Augmented Spatial Resolution	123
5.7	Chapter Summary	125
6	Biomass and EfW Flue Gas Filtration Overview	126
6.1	Chapter Abstract	126
6.2	Biomass Flue Gas filtration Overview	126
6.2.1	Typical Large Scale Biomass Power Generation Process	126
6.2.2	Flue Gas Filtration	131

7	Minimising TCO of Industrial Baghouses	138
7.1	Chapter Abstract	138
7.1.1	Master's Theses Summary	138
7.2	Experimental Methods	139
7.2.1	Biomass Power Generation and Filtration Trials	139
7.2.2	Historic Data Analysis	140
7.2.3	Offline Filter Analysis	144
7.2.4	Fly Ash Offline Analysis	147
7.3	Rapid Hybrid Computational Fluid Dynamic Design	148
7.3.1	Representing the design space	148
7.3.2	Inner Loop Design	149
7.3.3	Hybrid CFD	151
7.3.4	Cost Functions	152
7.4	Results & Discussion	153
7.4.1	Offline Filter Analysis	153
7.4.2	Missing data modelling	157
7.4.3	Compressed Air Consumption	160
7.4.4	Sorbent Dose	165
7.4.5	Pulse Sequencing	171
7.4.6	Cost Function Development	171
7.4.7	Hybrid Computational Fluid Dynamics Optimisation	175
7.4.8	Pleated Cartridge Conversion	180
7.4.9	Commercial Solutions	183
7.5	Predicted Commercial Impact	187
7.5.1	Plant Implementation	189
7.6	Chapter Summary	192
8	Conclusions and Recommendations to Industry	193
8.1	Chapter 2	193
8.2	Chapter 3	193
8.3	Chapter 4	195
8.4	Chapter 5	196
8.5	Chapter 6 & 7	196
	Appendix A Worldwide Patent WO 2016/001663 A1	198
	Appendix B Abstracts of the Master's theses	217
B.1	(Day et al., 2017)	217
B.2	(Eagle et al., 2017)	218
B.3	(Taylor et al., 2018)	218
B.4	(Pettitt et al., 2018)	219

List of Figures

2.1	Typical GTL processing route employing Fischer-Tropsch synthesis	15
2.2	Schematic graph to show windows of market applicability for several associated gas value recovery methods.	18
2.3	Schematic view of a simplified plate-fin reactor	20
2.4	Schematic to compare the layout and water partial pressure profiles of conventional single stage and CompactGTL's patented dual stage FT process	22
3.1	Schematic to show prototype reactor body arrangement	26
3.2	False colour photograph of reactor cross section	27
3.3	Schematic of δp measurement equipment	28
3.4	Schematic of vacuum suction apparatus	30
3.5	Graphical summary of a typical FTS run (SP09 shown)	34
3.6	DCS Screenshot to show typical operator and system response to a thermal event.	35
3.7	Scatterplot to show pre-operational catalyst loading versus post-operational δP , grouped by rows	40
3.8	Pre- and Post-operational analysis of SP09 grouped by column and row.	42
3.9	Contour plot to show interpolated $\Delta\delta P$ for SP09	43
3.10	Graphs to show location and extent of blocked regions in SP09 (mm).	45
3.11	3D view from the inlet of SP09 showing average location of runaway and average location and depth of blockage	48
3.12	Pie charts to show proportion of carbon and hydrogen (wt.%) present in each sample	50
3.13	Scatter graph to show δP of blockages vs length of blockage	53
3.14	Scatter graph to show δP of blockages vs length of blockage	54
3.15	Column graph to show comparison of carbon composition between samples sieved above and below 106 μm	56
3.16	Possible modes of carbon formation during FTS on cobalt catalysts	57
3.17	Electron micrograph (65 X magnification) of A) pre-operational catalyst. B) post-operational high δP catalyst (sieved to remove fines)	59
3.18	Electron micrograph (2500 X magnification) of A) pre-operational catalyst. B) post-operational catalyst.	60

3.19	Electron micrographs of A) sieved post-operational catalyst fines (350 X magnification). B) individual sieved post-operational catalyst fine particle (12000 X magnification).	61
3.20	Electron micrograph (50,000 X magnification) of individual post-operational catalyst particle superimposed with EDX elemental map for cobalt.	62
3.21	<i>in situ</i> free flowing catalyst (650 mm from inlet, Channel A1, C:H mol 0.22)	63
3.22	<i>in situ</i> photographs of radial blackening effect with superimposed channel dimensions and orientation. A) close-up view of the corner effect. B) view of majority of channel.	64
3.23	Schematic cross section of a single reactor channel showing the typical nature of the radial blackening effect observed by <i>in situ</i> endoscopy.	64
3.24	<i>in situ</i> photographs of 'Contact shielding' effect showing A) Close-up view of particles. B) wider view of particles showing both contact shielding and radial blackening.	65
3.25	Schematic diagram of two catalyst particles, illustrating the <i>contacting shielding</i> effect.	65
3.26	Photographs of blocked regions of a channel.	66
4.1	Schematic plan view of Perspex [®] block containing mock packed bed	69
4.2	Raw X-Ray Tomography image sections showing various artefacts	70
4.3	Comparison of raw and processed 2D x-ray tomography images	71
4.4	Horizontal and Vertical processed tomography planes for two particle sizes	72
4.5	Histograms to show geometric properties of two particle types derived from tomographic image analysis	73
4.6	Plots to show calculated shape parameter of two particle types derived from tomographic image analysis	75
4.7	Histogram and Tomographic surface projection showing tiered arrangement of particles for two diameters	76
4.8	Graphs and Tomographic Z-projection with average colour values (0 to 255) to show average bed void fraction as a function of channel width (μm for four cross sections. Channel 1 (600 μm particles) shown.	79
4.9	Graphs and Tomographic Z-projection with average colour values (0 to 255) to show average bed void fraction as a function of channel width (μm for four cross sections. Channel 2 (300 μm particles) shown.	80
4.10	3 Dimensional reconstruction of X-ray tomographic scan.	81
4.11	2D planar cross sections of tomographic scans for two particle sizes.	81
5.1	Schematic to illustrate differences in scale up / scale out behaviour for discrete and distributed temperature sensing schemes.	93
5.2	Schematic fibre optic element showing light pulse and corresponding backscattered signal due to external physical effects	94

5.3	Schematic to show three possible photon shifts following photon/molecule interaction in a fibre optic element	94
5.4	Schematic to show signal intensity vs frequency for backscattered portions of light, showing temperature dependent shifts	95
5.5	Schematic fibre optic waveguide demonstrating the critical angle, given two refractive indices.	95
5.6	Schematic cutaway of a generalised optic fibre section	95
5.7	Comparison of single mode and multi mode light transmission through a waveguide.	96
5.8	Comparison of core/clad refractive index profile for step index and graded index fibres	96
5.9	Graph to show effect of profile parameter, g , on refractive index profile of a graded index fibre	97
5.10	Mode propagation in a graded index fibre optic cable with a profile parameter of $g = 2$	97
5.11	Detail of tube plate showing spacing and header impingement	99
5.12	Schematic to show cross sectional view of a divider plate tube, showing a thermocouple with two passes of a fibre optic cable. All elements have close proximity to the tube wall.	100
5.13	Photograph to show insertion of thermocouple and a single pass of fibre optic cable into divider plate tubes of plate-fin Fischer-Tropsch reactor SP09.101	
5.14	Diagram of one of four band heaters used to encase lab scale FT reactors .	101
5.15	Schematic to show helical and axial mounting of a DTS element to a lab scale FT reactor	102
5.16	The Silixa Ultima-S DTS unit used in this study	104
5.17	Schematic flow chart of Silixa Ultima-S DTS components	104
5.18	Schematic to show fibre optic fusion splicing procedure	106
5.19	Various channel configurations possible with the DTS equipment	108
5.20	Graph to show achievable temperature resolutions at various sampling times as a function of fibre length, for fixed spatial resolution of 125 mm.	110
5.21	Graph to show (A) raw Stokes and Anti Stokes backscattered signal intensity as a function of length along the fibre optic cable and (B) signal processed temperature data for the same length. Data is from a single time point in a SP10 run. (Sampling time – 5 s, Spatial resolution, 125 mm)	113
5.22	Graph to show (a) trimmed Stokes and Anti Stokes signal and (b) Temperature as a function of length along the fibre.	115
5.23	Graph to show fibre temperature profile related to reactor geometry elements	116
5.24	Isometric cutaway to show SP10 internal divider plate structure highlighting key features of fibre optic path	117

5.25	Reactor temperature smoothing achieved as a result of spatial averaging of temperature data	118
5.26	Comparison of discrete thermocouples and sections of spatially averaged DTS fibre during an operational thermal event.	120
5.27	Plot of distributed temperature offset versus discrete temperature value for a thermal event	121
5.28	scaled distributed data, \hat{D} , compared to discrete data for a thermal excursion	122
5.29	Schematic to show relative size of lab, pilot and commercial scale reactors, showing the increase in tube length, w , at commercial scale.	123
5.30	Envisaged real time image of temperature data projected onto a commercial reactor geometry	123
5.31	Surface plot to show various features of DTS temperature data for a helically wound lab scale reactor.	125
6.1	Process flow chart of industrial biomass power generation plant	127
6.2	Transverse cross section of baghouse (axial midplane of compartment) . .	132
6.3	Schematic plan view of six-compartment baghouse lid arrangement and naming convention	132
6.4	Axial cross section of mid-plane of central inlet/outlet plenum in a six-compartment industrial baghouse in relation to the upstream duct.	133
6.5	3-Dimensional view of a six-compartment industrial baghouse in relation to the upstream duct.	133
6.6	Schematic cross section of filtration media structure	136
6.7	Schematic representation of the lime cycle	137
7.1	Raw historic process data for an extended period of operation (>2 years) A) Baghouse differential pressure (mbar). B) Volumetric Flow rate (m^3/s). C) Induced draught fan speed (% of max.).	141
7.2	Raw historic process data for an extended period of operation (>2 years) A) Stack exhaust particulate emissions (mg/m^3). B) Stack exhaust moisture content (wt.%). C) Stack exhaust HCl emissions (mg/m^3).	142
7.3	Raw historic process data for an extended period of operation (>2 years) A) Combustion, Boiler backpass and Stack temperatures ($^{\circ}C$). B) Solid biomass fuel flowrate (kg/s).	143
7.4	Schematic plan view of baghouse lids (A2 & B2) showing location installed filter brands and sampled locations.	145
7.5	Schematic view of a sampled filter bag showing locations of sub-sampled regions.	145
7.6	Flow chart to show general procedure of Rapid Hybrid Computational Fluid Dynamic Design	148

7.7	Three mechanical parts, generated by conventional and generative methods to perform the same duty	150
7.8	Main effects plots of air permeability ($dm^3/dm^2.min$) for offline sampled bags, in relation to geometric parameters.	154
7.9	Main effects plots of thickness (mm) for offline sampled bags, in relation to geometric parameters.	155
7.10	Main effects plots of weight (g/m^3) for offline sampled bags, in relation to geometric parameters.	156
7.11	Scatter plot of normalised pressure drop (mbar/MW) versus average age of filters (days) for two brands of filter.	157
7.12	Operationally linked bivariate plots showing plant-wide compressed air consumption against various plant KPIs	161
7.13	Schematic graphs to show system response to lime dose (D) (kg/s) reduction and enhanced mixing	168
7.14	CFD generated particle traces showing path of sorbent in the duct upstream of the baghouse for two mixing systems. Residence time, $t_{d,\mu}$ is shown.	169
7.15	Boosted air mixing nozzle designed for enhanced sorbent mixing.	169
7.16	Static mixing element designed for enhanced sorbent mixing.	170
7.17	CFD generated particle traces showing path of sorbent in the baghouse showing transverse mixing.	170
7.18	Comparison of sequential and skip pulse sequences showing path of re-entrained particles	172
7.19	Bivariate plot of measured baghouse differential pressure (mbar) versus parasitic load (kW).	174
7.20	Bivariate plot of predicted baghouse Dynamic DP reduction (% of mean) (less 5 mbar static DP) versus predicted yearly parasitic load savings (£/yr).	174
7.21	3-Dimensional views of steady-state CFD results of an industrial scale baghouse	175
7.22	Planar cross sectional view of CFD generated velocity magnitude (colour scale) (0-5 m/s) and vector (arrows) for 3 axial cross sections of Plant 1 Baghouse (filters removed).	176
7.23	Plan view of CFD generated velocity magnitude (colour scale) (0-5 m/s) and vector (arrows) for 2 vertical cross sections of Plant 1 Baghouse (filters removed).	177
7.24	Side elevation view of CFD generated velocity magnitude (colour scale) (0-10 m/s) and vector (arrows) for transverse mid-plane of the common inlet/outlet plenum.	178
7.25	Schematic cross section of single compartment baghouse showing sub-optimal air flow phenomena	181

7.26 3D schematic to show comparison between traditional bags and pleated cartridges	182
7.27 Optimised pulse cleaning benefits	184
7.28 Schematic comparison of diaphragm and spool valve pulse response to the same set-point change	185
7.29 Principle interactions in baghouse OPEX	188
7.30 Graphs to show predicted cumulative cash flow (CCF) (£'000s) for three proposed modifications.	191

List of Tables

2.1	Fischer-Tropsch synthesis reactions (de Klerk and Furimsky, 2010)	16
2.2	Response matrix to summarise general effects in FTS performance	16
3.1	SP08 post-operation δP (mbar)	39
3.2	SP09 sample location and quantity	44
3.3	Location of recorded thermal events during SP09 run	47
3.4	SP08 samples subjected to CHN elemental analysis	49
3.5	Raw results for CHN composition (wt.% \pm 0.3%). Results are in duplicate.	49
3.6	molar percentage of carbon and hydrogen in contamination and calculated C:H mol ratio. Values above \sim 0.48 are indicative of the presence of elemental carbon	51
3.7	Triplicate subsampled CHN analysis	52
3.8	Examples of various carbon species on cobalt FTS catalysts along with their hydrogenation temperatures. (Moodley, 2008)	57
4.1	Summary tomography parameters from two mock channels loaded with FTS catalyst of differing nominal particle diameter (channel height 4.06 mm).	72
4.2	Mean distance to wall (μm) for each tier (where Δ is the inter-tier distance, expressed in μm and radius units.)	77
5.1	Manufacturer supplied specifications for DTS element fibre portion. [†] Concentricity measured as (min. wall / max. wall). *Attenuation Measured on loose coil. (AFL Global, 2015)	92
5.2	Silixa Ultima range quoted maximum performance	103
5.3	Temperature resolution (\pm $^{\circ}\text{C}$) for various lengths of fibre at various sam- pling times achieved with the Silixa Ultima-S DTS equipment used in this study (spatial resolution fixed at 125 mm)	109
5.4	Sections of fibre joined by numbered nodes and their corresponding plate and vertical direction	116
5.5	Data regions to be averaged to give inlet-outlet orientated reactor tempera- ture profiles (letters with reference to 5.4, subscript ' _r ' denotes the data region is flipped horizontally	117
5.6	spatial resolution alignment with looping	119

5.7	Coefficient of determination (R^2) for various regression models applied to offset data shown in Figure 5.27. Good fit was obtained for most models.	121
5.8	Optimised parameters to fit scaled distributed temperature data onto discrete data and eliminate the offset	122
6.1	Summary of common fuel sources for Biomass and EfW (Korhaliller, 2010) (Popa, 2018)	128
6.2	Typical emissions limits for biomass combustion plants (European Commission, 2001) compared with actual emissions for a 33MWe plant (Wilson, 2012) (Day et al., 2017).	130
6.3	Table of common filter bag materials (Wang et al., 2004) (McKenna and Turner, 1989) (Mao, 2018).	134
6.4	Fabric pre-treatment options for filter bags. (McKenna and Turner, 1989)	135
6.5	baghouse cleaning systems	136
7.1	Average filter weights from offline analysis.	155
7.2	Summary of data-derived values for baghouse compressed air consumption	162
7.3	Summary of calculated operating expenditure for baghouse compressed air consumption	164
7.4	Economic estimates of sorbent cost	166
7.5	Summary of revenues and operating costs considered for Plant 1	173
7.6	Estimated cost and predicted benefits for Plant 1 implementation of various proposed modifications	190
7.7	Summary of predicted financial benefits. (median of range)	190

List of Publications

Patents

C. W. Bonner, B. Dannatt. (2016) *Catalytic Reactors Comprising Distributed Temperature Sensors*. World Intellectual Property Organization Patent no: 2016/001663 A1.

also published as

AU 2015282401 (2017), CA 2954214 (2016), CN 106999898 (2017),

EA 201790110 (2017), EP 3164206 (2017), GB 2527847 (2016), US 2017136432 (2017).

Master's Theses

supervised by the author

H. Eagle (2017) *Rapid Hybrid Computational Design of Industrial Energy Saving Flow Geometries - Biomass Fired Boiler Case Study*. Unpublished master's thesis, Newcastle University, United Kingdom.

L. Day (2017) *Cost Benefit Analysis of Filter Bag Performance Using Historic Operational Data*. Unpublished master's thesis, Newcastle University, United Kingdom.

C. Taylor (2018) *Industrial Baghouse Pulse-Jet Optimisation - A Hybrid Modelling Approach*. Unpublished master's thesis, Newcastle University, United Kingdom.

A. Pettitt (2018) *Modelling and Optimisation of The Pulse-Jet Cleaning Mechanism in a Filter Bag House*. Unpublished master's thesis, Newcastle University, United Kingdom.

Journal Articles

B. Dannatt, B. E. Goulden, M. J. Willis. (2018) *Optimisation of Pulse Jet Cleaning Frequency in a Biomass Combustion Baghouse - A Hybrid Modelling Approach*. Newcastle University, United Kingdom. (*manuscript in preparation*)

Conferences

B. Dannatt, M. J. Willis (2017) *Design optimisation using a hybrid modelling approach: an application to a biomass power plant*. [Poster] Exhibited at 10th World Congress of Chemical Engineering. Barcelona, Spain. 1 – 5th October 2017.

Chapter 1. Introduction

1.1 Introduction to Hybrid Modelling

Works presented in this thesis share the common theme of providing enhanced operability and profitability in industrial processes by applying hybrid modelling techniques. Hybrid Modelling is difficult to succinctly define. However, it can usefully be described as a set of approaches (empirical data gathering, first principles equations, statistical methods, simulation) which are combined to provide increased understanding of a given system, notwithstanding a lack of fundamental (first principles) understanding about one or more parts of the system. It is especially usefully employed in the process industries, where scale-up of well understood phenomena presents additional, often hidden, complexity.

Perhaps the most useful summary of the definition, history, development and applications of Hybrid Modelling is given by Glassey and von Stosch (2018). In their work, the authors begin with an intuitive example, allowing the reader to appreciate the role of Hybrid Modelling. In *this* work, where the focus is more specifically directed towards *industrial and commercial* applications, the author begins with the current attitudes and approaches taken in industry, before presenting a familiar example on which one can apply the concepts.

1.1.1 Industrial Approach to Process Improvement

While there are many examples of innovative and risk-tolerant industries and organisations, the following generalisation can be made about a typical large-scale industrial operation: they are averse to deviating from the *status quo*. For this reason, it follows that there is a fundamental contradiction in applying scientific methods to large, operational plants. Therefore the choices for carrying out research into process improvement options for large plants are as follows:

- *Do nothing, Change nothing.* - Whilst ostensibly flippant, this is the prevailing approach in large operational plants. The risk associated with change is seen as disproportionate to the gains of experimental process improvement. Consequently, significant effort and budget is committed to maintaining the plant at current or

nominal levels of various top-level KPIs (availability, uptime, unplanned shutdowns etc.)

- *Trust the salesman.* - Owner/operators of large plants are typically not expert in every unit operation and technology being utilised on site. This is especially true for ancillary equipment, utilities, and unit operations far up/down-stream of the *core* technology. Thus, Owner/Operators rely on the advice of third party specialists for the maintenance and specification of spares/repairs. These third parties are invariably 'salesmen', wishing to pursue their preferred methodology, product set or service.
- *Sub-set trial* - Often, Owner/operators will seek to reduce risk by installing or implementing a trial solution alongside the existing approach, with the existing approach taking the brunt of the duty. The perception is that the benefits observed for an implemented sub-set will scale *pro rata* when fully implemented. Unfortunately, this is hardly ever true. Furthermore, large plants tend to lack the high resolution spatiotemporal data required to detect the impact of such a small process change in the background of noise (both real and signal).
- *Pilot study* - A common approach to gaining confidence in a technology or solution is to perform a pilot study. Care must be taken to preserve important details in the scaled-down process. It is often difficult to decide what the scaling parameters should be, and therefore a reduced level of confidence is afforded at reduced scales. Conversely, the cost of the trial typically reduces proportionally to the reduction in scale. There must necessarily be a scale-up to commercial implementation, re-introducing risk and cost.
- *Develop a 1st principles model* - modelling by first principles allows fundamental understanding of a process to be employed in predicting a hypothetical scenario. While the 'cost' of a first principles model is very low, it is often an idealised version of events which involves many assumptions which are difficult to assess in terms of impact. These models, when robust, allow for intuitive and revealing recommendations as to the *direction* of change required to yield a particular effect. However, there is an understanding that the *exact* value predicted by a first principles model is almost never the true value in the real-world case (plant/model mismatch, PMM).

For each of the approaches listed above, there is a balance between cost and risk. This implies that the risk/reward relationship is known. i.e. there exists some cost function relating the increased profitability given by process improvement and the cost of implementing the improvement; a payback time and return on investment can be calculated. However, in practice, this is almost never the case.

1.1.2 Illustrative Example - Industrial Approach

Take a simple analogy from a familiar process: The hypothesis that "increasing car tyre pressure will reduce the operating cost (OPEX) of the car by increasing fuel economy". Now consider the analogous approaches that a large plant might take.

By *doing nothing*, there is zero risk of failure. However, tyre pressures will fluctuate and tend to decrease over time. Part of the OPEX of the car will be spent maintaining the pressure of the car between a high and low value (known as a dead-band). The car will also generally deteriorate over time, requiring an increased OPEX over time.

By *trusting a salesman*, one may be persuaded into purchasing high performance sports tyres which must be filled with high pressure nitrogen at regular intervals. The cost of this service may far outweigh the potential benefits of the improvement. Furthermore, the increased weight of the new tyres may *adversely* affect fuel economy, comfort, or cause increased wear on the suspension thus increasing OPEX.

By *performing a sub-set trial*, one could increase the pressure of only one tyre, keeping the other three the same. Should one expect that the benefits of doing this can be multiplied by four to give a prediction of the fully implemented benefit? Should one expect that the single inflated tyre will perform the same as four inflated tyres in operation? Intuitively, the answer to both is *no*. The non-uniform operation of the car will cause some unforeseen response. Increased localised tyre wear, vibration, handling issues, etc. The relatively small impact of increasing one tyre pressure may not be detected within the newly introduced noise, and is certainly not scalable as a predictor of uniformly implemented performance.

By *performing a pilot study*, we could cheaply assess the hypothesis by experimenting on the effect of differing tyre pressures on the battery life of a remote controlled car. This would give us some indication of the expected fuel economy of a full-size car. Intuitively, however, we sense that the scale-up implementation of the pilot-scale optimised solution would still be risky. What must be maintained between scales? How relevant are the various small-scale measurements we could take when compared to the full-size system? This varies for each process and across scales.

By *developing a 1st principles model*, we could use the well known laws of geometry and physics to predict that increased tyre pressure would reduce contact surface area and therefore friction between the tyre and road. We could use the chemical energy potential of the fuel mixture and the principles of internal combustion, combined with the mass of the car to calculate the distance over which a given volume of fuel could convey the vehicle under a given friction condition. We would have to make assumptions about wind direction and road surface and fuel price. We would make the simplification that the journey was uninterrupted, in a straight line and at constant speed.

Alternatively, we could perform a controlled empirical study whereby we operate a car at

different tyre pressures in a controlled environment and directly measure the impact on fuel consumption. We would be unable to control some of the conditions of the experiment (ambient temperature, pressure, wind direction, precipitation, traffic flow etc.) but we could control certain parameters as we perceive to be important (driving style, route etc.). The duration of the test would give us global data about the top-level KPI (fuel consumption), but it would be unfeasible to run the test for the whole lifetime of the car to assess the global impact on other responses (wear rates, etc.) since the car would not be providing value during the test (we can't use the car to go about our daily business during the experiment.)

Fundamentally, this is the reason that large plant operators are resistant to change. Seemingly none of the approaches in isolation gives both a satisfactory reduction in risk and a high level of confidence in the implemented solution. We cannot be sure of success, we do not have the correct metric for success and we do not have the sensors to measure success when it occurs.

1.1.3 Illustrative Example - Hybrid Modelling Approach

In Hybrid Modelling, we combine the approaches above to leverage the benefits of each, whilst negating the disadvantages. We rely on the following truths about large plants:

- Operation is inherently variable.
- This variation occurs over many timescales.
- This variation is often random and normally distributed. Where it is not, there is an underlying reason for the distribution of the variance which must relate back to some physical phenomena.
- Data is being captured in different forms, both by sensors and humans.
- This data encompasses the system response to variability.
- There will be 'gaps' in both the data (spatiotemporal resolution) and the first principles understanding of a given phenomena.
- With sufficient data and understanding, these 'gaps' tend not to overlap, and can be filled.
- Some system information is unknowable, given operational constraints (there is a window of operation).

Returning to the car analogy with these principles in mind we can devise a hybrid strategy:

We operate the car as normal. We collect all of the data we can about tyre pressures, weather conditions, road surface, driving style, carried weight etc. (in large plants, this

type of process data is measured and stored as a matter of course and years of historic operational can simply be downloaded). We also track important KPIs such as breakdown frequency, wear rates, maintenance requirements, fuel economy, etc. We may also take similar data from other cars operating normally (though differently). Through the daily routine operation of the car, all of the variables will be observed in combination with all other variables in a random manner, and the system response of this operating point will be the result of that particular operating point. For instance, tyre pressure will fluctuate over time, related to temperature, journey time, driving style and leak rate etc. and we will perform the normal dead-band maintenance cycle on tyre pressure.

Hybrid modelling allows us to use statistical tools to explain the variance in system response. We treat the historical data set as though we had designed a comprehensive experiment covering all variables and system responses within the window of operation. We can retrospectively append the data onto a Design of Experiments (DoE) structure to allow for statistically confident characterisation of the system. We can rule out data points where we have operational records to show that this data was not representative (When the car was in the garage, in a 5 hour traffic jam or on holiday in the mountains). We can perform simulations of known operating points and validate by characterising the mismatch between a first principles model and our data. We can use the offset between prediction and observed value to adjust our model and gain new fundamental understanding. We can use the adjusted model to predict the system response to operating points or combinations of variables not covered in the data set (we can validate this by reserving a sub-set of the data for validation).

Given enough time and data processing, we can arrive at a robust predictive description of the process such that we have confidence in not only the physical system response to change, but also a prediction of the global impact on OPEX. Our Hybrid Model not only answers the original question about tyre pressure, but is also able to make more general recommendations about reducing OPEX through a combination of process improvements.

1.1.4 The Role of Computation

At this point, there is still a 'leap of faith' required to make a process change. Our hybrid model was constructed based on an existing window of operation. If we wish to operate outside this window, because we predict a benefit to doing so, we must still generate a new data set to show that the models do not break down beyond the current window.

In recent years, the computational power of readily affordable hardware and cloudware has crossed an important threshold. It is now possible to *simulate* processes at industrially relevant scales in practical timescales. Commercially available software is now also advanced enough to exploit this power.

Computational Fluid Dynamics (CFD), once restricted to super-computers, is now a commercially available software tool in the same way that word processors and spreadsheet applications are. The underlying equations of CFD analysis, (Navier–Stokes equations) are well understood, robust and trusted. The benefit of increased computing power in the solving of these equations is the ability to dramatically increase the spatiotemporal resolution of the simulation. Advanced data structures and statistical methods reduce the simulation time and increase the accuracy and precision of the results.

Computational Fluid Dynamics therefore allows us to simulate outside of the window of operation in a risk-free sandbox, whilst still generating a highly representative data set. By iterating CFD and other simulations into our Hybrid Model, we are able to extend our model’s applicability beyond the current window of operation.

1.2 General Scope and Aims of This Thesis

1.2.1 Gas-To-liquids Process Development

The first half of this work is concerned with process development and commercial scale up of a modular Gas-to-Liquids process. The proprietary process of concern is briefly comprised of a packed bed plate-fin ‘mini-channel’ Fischer-Tropsch reactor. At the commencement of this work, the technology platform was approximately fixed in terms of fundamental design. However, the detailed and optimum design parameters for various elements of the process had yet to be determined. Of particular interest was the ability to understand, predict and control the thermal and flow characteristics of the reactors in order to avoid costly downtime of a commercial system.

Flow characteristics (such as differential pressure) are challenging to compute through first principles in real-world packed beds for a number of reasons. Namely the assumptions of first principles models in the current literature are idealised in a number of ways:

- Particles are mostly assumed to be spherical and mono-disperse.
- Channel geometries are simplified to circular or rectangular cross sections and limited in aspect ratio
- The models often rely on coefficients which have an origin in empirical studies, and are therefore not truly first principles and not guaranteed to be applicable to any given system.
- The properties of real fluids and beds are a function of time, whereas most models consider only steady-state scenarios.

The result of these limitations is that little can be predicted about the flow or temperature

performance of a complex reactor geometry in the design phase. Data-driven models provide a simple solution in that empirical data inherently encapsulates the deviations from ideal which limit the first principles models. The fluid and thermal performance is directly measured. However, the limitations of empirical models originate in their resolution. Measurements cannot be taken at infinitesimal resolution. A probe or sensor of some kind is required to be placed *in situ* for a representative measurement. The sheer presence of a sensor may in-fact alter the local performance (i.e. a thermocouple may act as a heat sink and alter the packing characteristics of the nearby bed, artificially reducing flow and temperature at a local level). This fact, coupled with the physical and data processing constraints of having many sensors *in situ* mean that spatial and temporal resolution of the empirical data is limited, often to a global level. This results in a similar idealisation to homogeneous performance as the first principles models.

Empirical modelling also comes at the expense of predictive design; a particular design iteration must be manufactured, commissioned and operated before it can be measured and modelled. At commercially relevant scales, the parameters of interest are both costly and hazardous to iterate through empirical trial and improvement. At laboratory and pilot scale, the costs and risks can be mitigated somewhat, but using pilot-scale empirical models to predict full-scale performance is also fraught. The same strength of data-driven modelling becomes a weakness; data-driven models ignore mechanisms, knowledge, heuristics, size and shape. One cannot be *certain* that an effect observed at small scale will be present at a larger scale, nor of the precise cause.

While many successful attempts have been made in scale-up and process development throughout history, there are surely untold failed attempts. Commercially, (and this EngD thesis necessarily addresses commercial concerns as a priority) we must seek to decrease the risk of failure to a minimum. A *hybrid* modelling approach, specifically targeting the limitations of the traditional methods outlined above, would allow for reduced risk of failure and greater understanding through the scales.

Specifically, can we increase the spacio –temporal resolution of data (without distorting the true undisturbed value)? Can we apply empirical data to modify first principles knowledge? Can we use high resolution retrospective analysis to explain and enhance low resolution empirical data?

1.2.2 Biomass Flue Gas Filtration Optimisation

The second theme of this work is related to flue gas filtration in large-scale Biomass and Energy-from-Waste Power Generation plants. The flue gas filtration technology utilised in these plants is a scaled-up version of traditional dust collector units. Dust control applications on the small scale rely on heuristics and experience to specify, with only very few design parameters *known* to be of importance in the design. Still, few objective and

widely accepted parameters are agreed upon. Consequently, the literature is dominated with commercial white-papers, with the claims of one contradicting the other.

End-of-Pipe filtration is regarded as a 'necessary evil' to achieve emissions targets, and is overlooked as an opportunity for process improvement, or even improved reliability. The upstream process receives far more resources (both in budget and manpower) for process understanding, optimisation and development. Consequently, the innovation in baghouse filtration technologies has been slow. Industrial baghouse technology has remained largely unchanged since its inception.

The result is that even world-leading projects (such as Wilton 10 Biomass - europe's first large scale biomass power generation facility) employ dated technology in the downstream flue gas filtration plant. During the procurement phase, budget and footprint constraints are likely to impact on the lower priority unit operations first, with flue gas filtration falling low on the priority list.

BATNEEC ("Best Available Techniques, Not Entailing Excessive Costs") is the European directive introduced in 1984 (84/360/EEC) and applies to air pollution emissions from large industrial installations. While this approach is taken to heart upstream, it may be invariably contrasted with the unofficial maxim CATNAP ("Cheapest Available Technology, Narrowly Avoiding Prosecution!") in the flue gas filtration stage.

A filtration system designed on the basis of lowest CAPEX, and then operated at the lowest possible maintenance cost sounds like a sound strategy for minimised "total cost of ownership" (TCO). However, combining initial purchase price and direct maintenance cost does not encompass the true TCO cost function. In simple example, a more efficient filter may be more expensive, but may protect downstream erosion or require less power to operate. With the inclusion of these wider 'costs' into a TCO model, it may be possible to arrive at a design or operating strategy with a far reduced total cost of ownership.

When applied to new plants, this methodology allows for the creation of a business case for applying BATNEEC to the previously overlooked filtration plant. In the case of existing plants, the same methods can be used to address operational challenges and perform upgrades or predictive maintenance on a "Return on Investment" (ROI) basis. This elevates the perception of downstream filtration from "necessary and costly" to "valuable and profitable".

The challenge of developing this minimum TCO framework for existing plants is discussed in this study. A large scale operational plant presents challenges which can be addressed by hybrid modelling.

- The cost and risk of experimentation is high and therefore there is a resistance to any change without sufficient confidence in the outcome.

- Online data collection is limited to the measurement devices installed at the time.
- Commercial output targets prevail, resulting in little or no experimental control over operational parameters during the study period.
- Operation is transient, with many hundreds of variables contributing to performance (only some of which are known, fewer are measured and even fewer are understood).

In this work, the development and application of hybrid modelling techniques is discussed. More specifically in the biomass flue gas filtration case study, the investigation proceeds by asking whether one can:

- Gather, process and summarise large amounts of historical data and relate it to intuitive and revealing metrics.
- Explain empirically observed phenomena in terms of a first principles mechanism or root cause, and quantify the tangible impact on key performance indicators.
- Populate a rich data picture by performing analysis of the sparse data points currently available.
- Establish the wider impacts of filtration design choices on the total cost of ownership.
- Provide a representative *in silico* simulation environment to allow for benchmarking of upgraded hardware in a risk and cost free manner.
- Combine data and understanding to develop a more comprehensive cost function for total cost of ownership.
- Predict the TCO savings (and ROI) for proposed solutions to allow for commercial implementation.

1.3 Structure and Major Contributions of the Thesis

1.3.1 General Impact

One of the key strengths of this work and contribution to the wider literature in both fields broadly covered in this work is the use of real operations and commercial data from industrially relevant sources. Many of the studies currently known to the literature are either purely model-based, theory-based or lab/pilot scale empirical studies. In this work, we present a hybrid approach whereby empirical data is used to form statistical models for relationships where no clear fundamental relationship can be described mathematically. Equally, first-principles models are employed where no suitable data can be gathered empirically. Finally, the process understanding, heuristics and recollections of plant operators, engineers and maintenance personnel can be integrated formally into the decision making process of process design/optimisation.

While there are certainly many indispensable references in the literature with regards to the theoretical basis for concepts relied upon in this study, it is hoped that the industrial data and discussion found in this work will provide the reader with a practical 'lens' with which to view the work of others:

- Real processes have high order interactions beyond our control.
- We can never measure the true value of a given parameter and the process of trying to measure it will distort the outcome.
- The data we do collect has hidden meaning and value.
- OPEX and CAPEX budgets are managed by different people, who don't always talk to each other.
- The 'optimum solution' will likely be decided on 'cost', even when the 'cost' is unknown.

Furthermore, the work sets out a framework for the so-called *Rapid Hybrid Computational Fluid Dynamic Design* whereby the overall computational demand (and therefore time) for optimising design parameters is minimised by the combination of systematically constraining the design space, statistically reducing the number of simulations and batch-to-batch iteratively (*in silico* or *in vivo*) validating the response.

The author has supervised and co-supervised four masters students to produce successful master's theses and attracted additional funding to continue this work commercially.

One of the key benefits of the industrially focussed EngD programme is the direct link to industrial need. Work carried out in this study has been done so in response to a particular

industrial challenge. For this reason, the impact of this work is largely measured in its industrial application and commercial value, some of which are outlined below.

1.3.2 Chapter synopsis and impact

Chapter 2

This chapter gives an overview of the proprietary modular Gas-to-liquids technology considered in the following chapters.

Chapter 3

This chapter forms the foundational understanding of the development of a next-generation technology platform for a commercial small-scale Gas-To-Liquids company. Discoveries, presented in this work, about the specific relationship between operational history and the pre/post-operational state of a reactor lead to a number of design and operational changes in the next-generation commercial reactors. By performing hybrid modelling, much more process understanding is gained. Specifically, the location of operational thermal events was causally linked to sintering of the catalytic bed and the presence of elemental carbon. The commercial implications of this understanding are significant, allowing the intellectual property holder to operate commercial reactors in a more safe and efficient manner.

Chapter 4

This chapter solves the issues of characterising a packed catalytic bed faced by the first principles approaches. X-Ray tomography is successfully applied to the novel next-generation reactor geometries through the use of image/data analysis. Characterisation of the fixed bed in this manner allows for a common-sense understanding of *localised* bed conditions, not described in data-driven or first principles models. The 3D computational representations of packed beds generated in this chapter provide the ability for further *in silico* flow studies to be performed.

Chapter 5

Work carried out in this chapter, presenting novel hardware and software configurations for enhanced understanding and operability of catalytic reactors, forms the basis of a worldwide patent for which the author is named as inventor. Fibre Optic Distributed Temperature Sensors are applied directly to the reactor geometry in a novel manner. The arrangement of the sensor is designed to achieve an apparent spatial resolution

beyond that of current hardware. An envisaged commercial control scheme is described in which the sensor is utilised to provide a dynamic level of detail, based on operational requirements. This allows for the safe and stable operation of a large number of reactor modules simultaneously.

The industrial applicability of this work was secured when the patent owner gained investment to build an £800 million Gas-to-Liquids plant based on the technology platform developed in part by the author, including implementation of the patented work discussed in this work.

Chapter 6

This chapter provides an overview of biomass and energy-from-waste power generation plants and more specifically the flue gas filtration unit operation considered in the following chapter.

Chapter 7

This body of work was carried out with access to two commercially operating large scale biomass power generation plants and co-funded by a filtration technology and solutions provider. Hybrid Models are presented with a view to reduce the total cost of ownership of the filtration system.

The impact of the study was two-fold. Firstly, the plant operators received recommendations for process improvements predicted to deliver upwards of £2.5 million in reduced OPEX over a 5 year period, with a payback period of ≤ 1 year. Implementation of these proposed improvements is currently in progress, with increased data capture ability providing a meaningful way to quantify the true impact. Secondly, the hybrid modelling methods, models and frameworks developed in this work now form a fully commercialised package of consultancy services offered by the co-funding filtration company. The increased technical capability is already generating additional revenue, predicted to reach approximately £1 million by 1 year after commercialisation.

Chapter 8

This chapter sets out the key conclusions of this work and the recommendations to industry and future academic studies in the field.

Appendix A

The published text of the worldwide patent is included in full.

Chapter 2. An Overview of Modular Gas-to-liquids Technology

2.1 Chapter Abstract

Modular Gas to Liquids may be an important component of the energy supply chain, given that new oilfields are becoming increasingly remote, smaller and variable in production. In this chapter, the fundamental process and design considerations for development of a proprietary modular gas to liquids (GTL) solution are set out. The Fischer-Tropsch process is described in detail with particular attention to the conditions required in commercial operations. The desired product of a GTL plant can be tuned by a subtle shift in operating conditions; temperature, pressure, syngas ratio etc. Conversely, slight undesirable disturbances in these conditions can hamper operations in terms of key performance indicators; Productivity, selectivity, ‘alpha’, conversion etc. A proprietary, two-stage braised plate-fin packed bed reactor is described. The particular design addresses several commercial requirements also set out in the chapter. Finally, the currently proposed mechanisms for Fischer-Tropsch catalyst deactivation are discussed.

2.2 Gas to Liquids Overview

Gas-to-Liquids (GTL) is the term given to a collection of refinery processes used to convert gaseous hydrocarbons (usually methane) into more useful liquid hydrocarbons. Liquid hydrocarbons are preferred over gaseous forms due to their increased energy density and transportability. Liquid hydrocarbons of varying chain lengths are usefully employed directly as liquid fuels or lubricants, or able to undergo various upgrading steps to convert to higher value chemicals. This is contrasted with the gas feedstock, which would otherwise be flared. (Compactgtl.com, 2015)

2.2.1 Typical GTL Process

A GTL plant typically comprises 4 key stages; feed gas treatment, synthesis gas generation, Fischer-Tropsch Synthesis (FTS) and an optional product upgrading section. The feed gas, following recovery of any Natural Gas Liquids (NGL’s), first undergoes a sequence of

treatment and conditioning steps. Steam is then added to the clean natural gas stream prior to heating and entry to a catalytic pre-reformer in which the higher hydrocarbons are reacted to give methane, hydrogen and carbon monoxide.

Fischer-Tropsch Optionally, the synthetic crude may be converted to diesel blendstock or premium fuel products by integrating mild hydro-processing, where the economic drivers dictate, including the ability to produce winter grade diesel blendstocks where appropriate. Figure 2.1 shows A simplified flow diagram of a GTL Plant. (de Klerk and Furimsky, 2010).

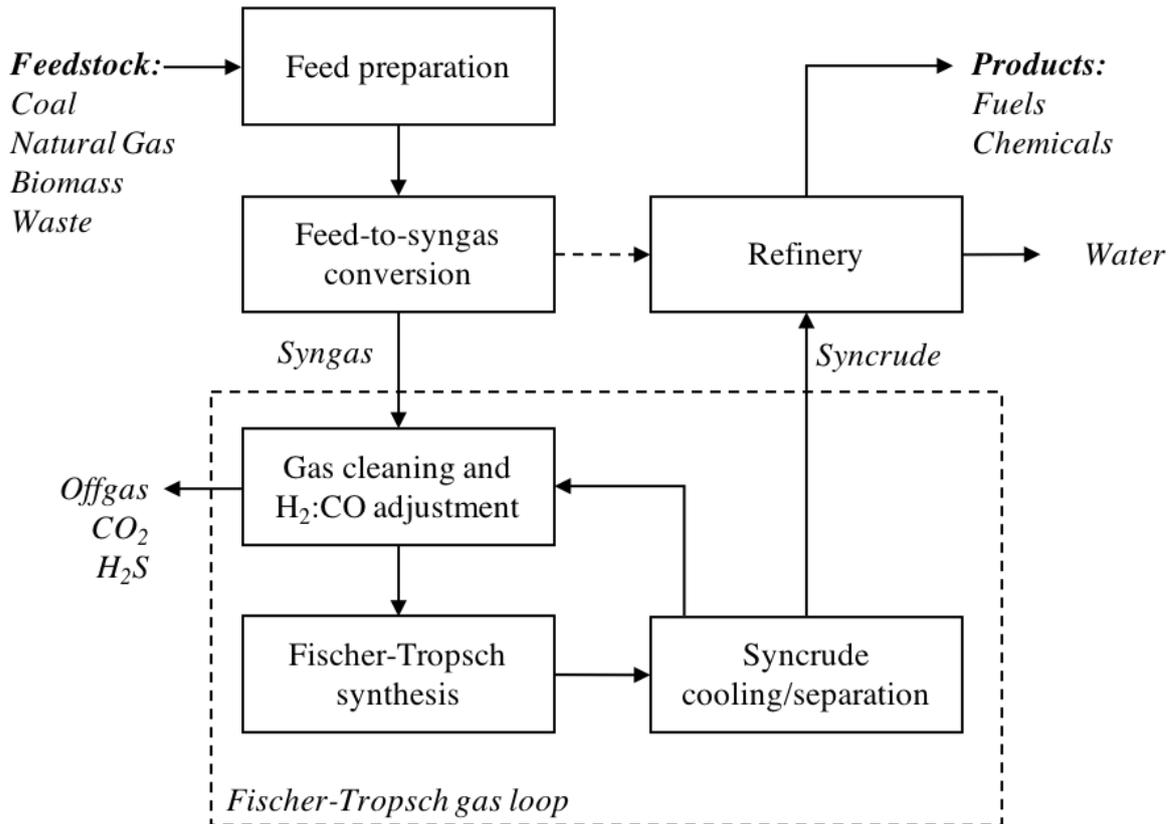


Figure 2.1: Typical GTL processing route employing Fischer-Tropsch synthesis. (de Klerk and Furimsky, 2010)

2.2.2 Fischer-Tropsch Synthesis

Table 1 shows the principal reactions taking place under normal operating conditions within a FT reactor (de Klerk, 2011). The principal goal in most FT applications is to produce a pure blend of olefins (alkenes) and paraffins (alkanes). There exist a number of side reactions which produce (usually) less desirable products. The choice of catalyst and reaction conditions allows selectivity to be tailored. Catalyst selection and composition is covered at length in previous work and will not be revisited here, save to say that cobalt-based catalysts doped with rare transition metals are preferred industrially (Moodley et al., 2009).

Table 2.1: Fischer-Tropsch synthesis reactions (de Klerk and Furimsky, 2010)

Type	Description	Reaction
Main Reactions	paraffins	$(2n + 1)H_2 + nCO \rightarrow C_nH_{2n+2} + nH_2O$
	olefins	$2nH_2 + nCO \rightarrow C_nH_{2n} + nH_2O$
	water gas shift (WGS)	$CO + H_2O \rightleftharpoons CO_2 + H_2$
Side Reactions	alcohols & esters	$2nH_2 + nCO \rightarrow C_nH_{2n+2}O + (n - 1)H_2O$
	aldehydes and ketones	$(2n + 1)H_2 + nCO \rightarrow C_nH_{2n}O_2 + (n - 2)H_2O$
	carboxylic acids and esters	$(2n + 1)H_2 + nCO \rightarrow C_nH_{2n}O_2 + (n - 2)H_2O$

Table 2.2: Response matrix to summarise general effects in FTS performance

Controlled Variable	Response Variable				
	Chain Length (α)	O:P ratio	Syngas Conversion	CH_4 Selectivity	Carbon Formation
Temperature	—	—*	+*	+	+
Pressure	+*	~	+*	—*	~
H_2 : CO ratio	—	—*	~	+	+
Space Velocity	~	+	—	—*	~

symbols indicate response to an *increase* in controlled variable level.

+: increase. —: Decrease *: desirable/favourable impact.

All Fischer–Tropsch reactions are highly exothermic; an average value for the heat of reaction is around 10 kJ/g of hydrocarbon product (Mousavi et al., 2015).

Key performance indicators in a commercial FTS reactor include syngas conversion, methane selectivity, usage ratio, olefin to paraffin ratio and product distribution, most commonly represented by the now ubiquitous Anderson-Schulz-Flory parameter, α (Schulz and Claeys, 1999) (Schulz, 1999). Typically, a high value of α (> 0.8) is desired, since this means that more of the methane is converted to higher molecular weight hydrocarbons. (in fact, α denotes the probability of chain growth in the FT reaction, and is only colloquially used to describe an achieved product distribution, for instance, “high alpha.”)

Most of the operating conditions for a given reactor and catalyst system vary the key performance indicators by intuitive equilibrium relationships. Table 2.2 summarises the key relationships in controlled versus observed parameters in FT reactions. An asterisk in Table 2.2 refers to a generally desirable effect in terms of increased profitability. For example, an increase in operating pressure will result in an desirable increase in α .

2.2.3 Small Scale Modular GTL

Current Market

Gas-to-liquids (GTL) has seen resurgence in research and development activity in recent years. Baxter (2012) notes that the commissioning of Shell's giant Pearl Qatar plant, plans by Sasol for large plants in North America and Uzbekistan, a record arbitrage between liquids and gas commodity prices and the emergence of recoverable shale-gas resources, many of them remote, have contributed to the debate.

There are currently five large scale GTL plants operating globally, with capacities ranging from 2,700 barrels per day (bbl/d) to 140,000 bbl/d. Shell operates two in Malaysia and one in Qatar, Sasol operates one in South Africa, and the fifth is a joint venture between Sasol and Chevron in Qatar. (Wood et al., 2012)

Implementation of further large-scale GTL plants in the near future seems almost certain. At capacities of 30-100 million barrels per day (bpd) and upwards, projects exhibit economies of scale which de-sensitise them to commodity price and project cost uncertainties. On the other hand, with oil and gas fields becoming smaller and more remote with each discovery, it is clear that small scale GTL will have much to answer for in future.

A number of small-scale GTL technologies have been conceptualised with varying degrees of success. However, none has yet been converted into a commercial project. One reason is the inherently high capital cost and complexity of GTL. High costs derive from exotic construction materials, high design temperatures, aggressive or corrosive environments and the use of specialised catalysts at all stages of the process. In the absence of economies of scale, small GTL attracts higher capital cost estimates per barrel of capacity. This can depress project internal rate of return (IRR) and increase sensitivity to market changes, making it difficult to realise attractive project proposals.

With capital costs marginalised, small scale GTL relies on incremental improvements at the development stage in order to emerge as a commercially viable venture. Bloomberg (2014) priced recently announced small-scale plants at \$100,000 for each barrel of capacity to produce liquid fuels. These 15,000 bpd plants are still an order of magnitude smaller in scale than Shell's Pearl venture in Qatar. Of this figure, the cost of catalyst is significant portion.

Window of Operation

With new oil fields becoming increasingly rare, small and remote, the typical economic model of large scale GTL plants breaks down. Coupled with the current low price of oil, this has triggered the economic viability of small scale GTL by making high value

wax lubricant. This is a dynamic situation, and a full assessment of market conditions is beyond the scope of this experimental report, but will be covered in future work. Figure 2.2 summarises the current options for handling associated gas. Small scale (modular) GTL is viable even in low productivity, remote locations.

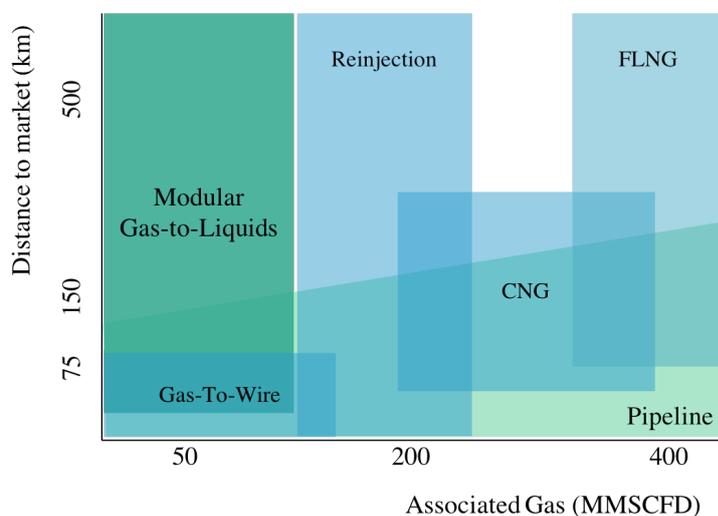


Figure 2.2: Schematic graph to show windows of market applicability for several associated gas value recovery methods in relation to the remoteness of the feedstock (km) and the flow-rate of associated gas (MMSCFD). (CompactGTL, 2015)

With the development of Small scale GTL come several challenges. Compared to large scale GTL, more highly active catalyst is used in a much smaller volume. The design of intensified modular FT reactors must therefore be dictated by their heat transfer capabilities. CompactGTL reactors are highly efficient at removing heat produced by the exothermic FT reaction. However, as has been presented in previous work, the reactors are occasionally subject to a local imbalance of heat generation and removal capacities. This results in a thermal excursion or thermal runaway; potentially damaging to the catalyst and reactor.

For this reason, high resolution temperature data is required to monitor and control compact FT reactors and guard against thermal events. With better control, Compact GTL reactors are able to operate closer to the ‘edge’ – higher productivity without causing a thermal runaway. Avoiding a runaway is crucial: modular GTL reactors should last 3-5 years on a single catalyst load if economic viability is to be maintained.

2.2.4 Compact Fischer-Tropsch Reactors and Process

Commercial Process Requirements

The drive towards small-scale GTL, as a result of the increased prevalence of remote, low yield, unconventional and variable productivity gas and oil fields has resulted in the development of intensified modular reactors. While several commercial efforts have resulted in a variety of designs, modern compact modular FT reactors are typically high surface area heat transfer structures filled with catalyst. The design is often limited by several external factors.

The requirements for a small scale GTL plant, operating remotely, differ from large scale operations in accessible regions. The design of the reactor system, as well as the plant as a whole, is largely dictated by these constraints. In particular:

- High availability - robust and reliable operation takes priority over process efficiency
- Suitable for oilfield conditions - ability to accommodate short and long term fluctuations in feed gas composition and flow rate, as well as fast shut-down and start-up capabilities
- Deployable in remote locations with poor logistics – modular systems constrained to shipping container dimensions and weight limits
- Maintainability - long service intervals in remote locations
- ASME VIII code compliant and certifiable equipment – process module manufacturing utilising established, certifiable manufacturing techniques
- Mass producible – utilising established companies to form a supply chain that ensures quality and volume delivery
- Suitable for onshore and offshore deployment and use
- CAPEX and OPEX – must fall within the ranges required for economic viability

Brazed plate-fin reactors

To meet the requirements outlined above, CompactGTL reactors, considered in this work, take the form of brazed plate-fin heat exchangers. Commonplace in the oil and gas industry, this design is a known quantity in terms of operation, maintenance and compliance.

Several hundred patents are available in the literature relating to the development of the proprietary reactors considered in this work. A most recent description can be found in (Bonner and Dannatt, 2015).

The novelty of the proprietary plate-fin reactors is the direct loading of catalyst within mini-channels: channels formed by the castellated fins with dimensions around 3 - 10 mm. The process of direct catalyst loading is made viable by the simple geometry of the exchanger and the cross-sectional dimensions of the channels; relatively large in relation to the competing technology; "micro-channel" reactors. Particulate catalyst with a mean diameter preferably less than 0.5 mm can be simply loaded into vertical channels on the fin side, while a heat transfer fluid is employed on the plate side to control the temperature.

This temperature control operates in two modes; firstly, during start-up, the temperature of the heat transfer fluid is increased by an external heater and this increases the temperature of the catalyst and syngas stream to a suitable temperature for FTS to occur. Secondly, during operation, the heat transfer fluid flow rate and temperature can be controlled to remove exothermic heat generated by the FTS reaction.

At commercial scale, a pressurised water/steam system is used to provide the heating and cooling duty for reactors. However, the heat transfer fluid used in pilot scale reactors in the present study is a eutectic mixture of 26.5 % biphenyl + 73.5 % biphenyl ether. Such a mixture is well-suited for heat transfer applications because of the relatively large temperature range of its liquid state.

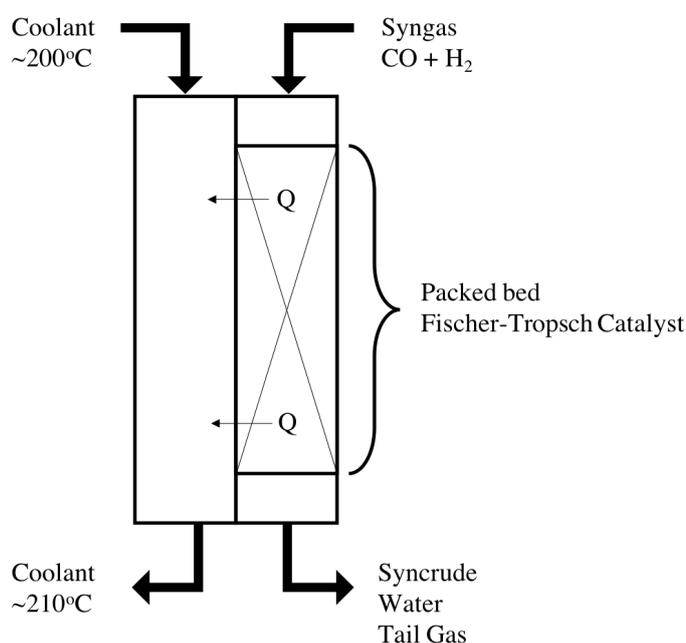


Figure 2.3: Schematic view of a simplified plate-fin reactor

The mini-channel configuration and plate-fin mechanical construction of the reactor units provide two critical characteristics: Firstly, very high specific heat transfer coefficients between the streams. Secondly, a relatively low metal inventory per unit of throughput. The high specific heat transfer coefficients between the streams in the reactor unit enable thermal stability to be achieved at reaction rates not attainable to date in conventional reactor designs. The low metal inventory, or high voidage, allows for an increased catalyst

loading per reactor volume. Together, these two aspects enable the compact reactor modules to meet the size, weight and operational constraints required for remote oilfield deployment, both onshore and offshore.

A further consideration regarding the supply chain for commercial implementation is the availability of an established, scalable and certifiable manufacturing route for plate-fin type reactors. Brazed plate-fin heat exchangers are extensively used in oil & gas and a wide range of industries as they provide a very good surface area to volume ratio compared to shell & tube or plate & frame units. Critically, ASME VIII pressure vessel code compliance protocols are also very well established for brazed plate-fin units.

Two Stage Modular FT Process

All FT processes create approximately one barrel of waste water for every barrel of synthetic crude product. This means that a partial pressure of water vapour develops in the process channels as the reaction progresses and this can cause hydro-thermal ageing of many FT catalysts (Moodley, 2008). To preserve catalyst life, the CompactGTL patented FT process is arranged in two stages, whereby conversion of the syngas to liquid product is limited in the first reactor. The resultant water vapour and condensable hydrocarbon product is then condensed from the tail gas from the first stage before the gas is passed into the second reactor.

A two stage FT system is capable of providing catalyst life of between 3-5 years compared to 1-2 years for an equivalent single stage system (2.4). The long catalyst life gives rise to high stability, availability, and low OPEX for a commercial plant.

In each FT stage, the syngas is pre-heated to ~ 220 °C before entering the reactor. The proprietary FT reactor units for the first and second stages are identical and comprise arrays of mini-channels constructed using brazed plate-fin construction, with two streams in a co-linear flow configuration. Typical operating conditions for the process stream are 220 °C to 235 °C at 25 bar. The FT process channels contain a particulate, cobalt based FT catalyst. As FT synthesis occurs down the length of each mini-channel, liquid hydrocarbons start to form and migrate under gravity and pressure to the outlet header after which they are cooled and collected. The FT reaction is highly exothermic so the adjacent coolant channels carry a pressurised water coolant to maintain acceptable thermal gradients in the radial and axial directions. This enables a more uniform overall rate of reaction profile and heat flux at every point in the reactor core. This maintains all the catalyst inventory at a more uniform temperature than would otherwise be possible, further promoting stable performance and long catalyst life.

Though it is possible to run FT reactors at a high conversion by applying more temperature to the process, later chapters in this work demonstrate that it is preferable to limit the

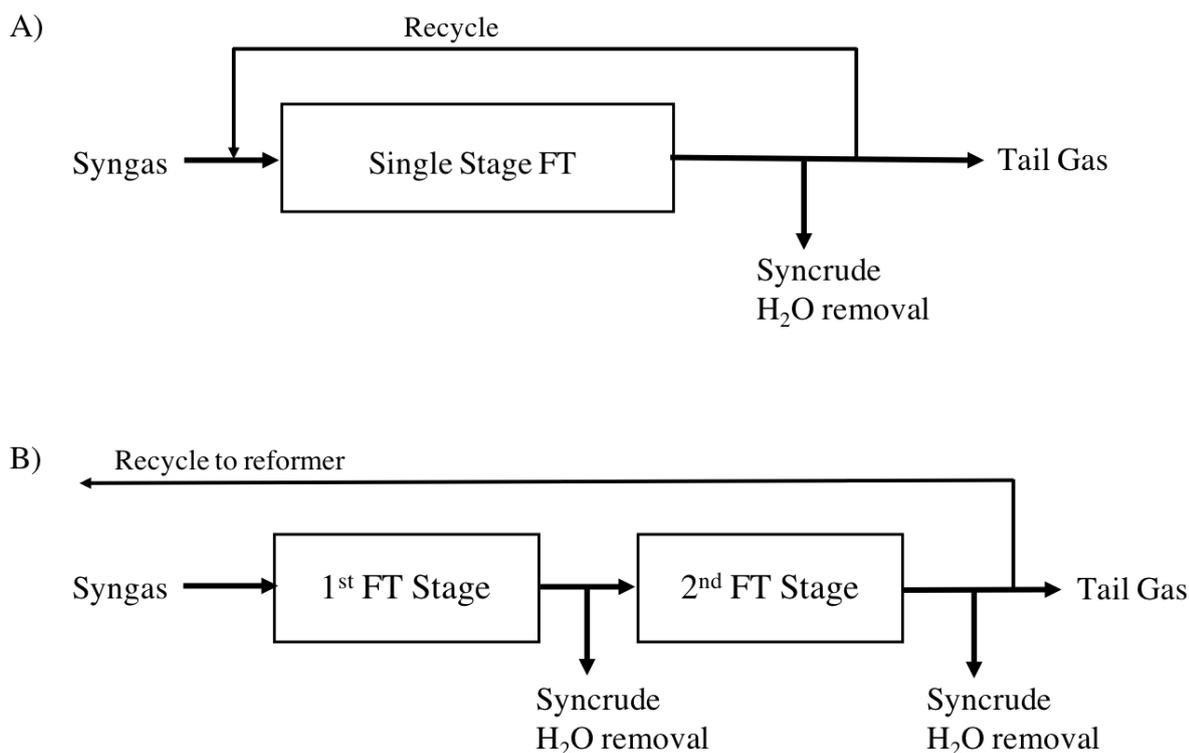


Figure 2.4: Schematic to compare the layout and water partial pressure profiles of conventional single stage and CompactGTL's patented dual stage FT process

conversion to below $\sim 60\%$ to ensure that a stable operation is maintained, without the risk of thermal runaways.

The use of liquid coolant also enables individual FT reactor modules on a plant to be operated at different temperature settings, whilst being supplied from a common, cost effective cooling system. Critically, for a complete plant operation, this accommodates operation of reactors with aged catalyst alongside reactors with virgin catalyst following catalyst replacement.

For commercial plants, the FT reactor modules will require a periodic removal and refurbishment cycle to replace aged catalyst with fresh catalyst in order to maintain the system performance. Catalyst replacement could be expected to occur every 3-5 years. Reactor modules are isolated in groups for removal and replacement with refurbished reactors containing fresh catalyst, whilst the remainder of the plant remains operational, maximising plant availability. The duration of this change-out is expected to be of the order of 4-5 days per group. For this reason, it is vital that change-out activities are well optimised. Future chapters of this work will discuss a number of reactor turnaround issues, including; unloading, reactor cleaning, post-operational analysis, catalyst loading and conditioning. These change-out activities can be phased in campaigns such that a given plant exhibits broadly the same overall output capacity at any given time, and the workload for reactor replacement and refurbishment activity is spread out evenly.

Fischer-Tropsch Catalysts

The lynchpin unit operation in GTL technology is Fischer-Tropsch Synthesis (FTS). Careful design of both catalyst and reactors are necessary in order to tailor conditions to the desired product. The anatomy of an FTS catalyst is fairly well known in industry, and only relatively minor variations exist. (Bartholomew, 2003), (Davis, 2007), (Iglesia, 1997). In general:

- Cobalt is typically used as the active Fisher Tropsch (FT) metal at around 10-30 wt%
- A second, usually noble, metal acts as a promoter comprising ≤ 1 wt% (Pt, Ru, Pd)
- A structural oxidic promoter (Ba, La, Zr)
- A high surface area, mechanically strong support (oxides of Si, Al, Ti are common)

Nanoscale cobalt particles are dispersed onto the support material in an optimal manner to avoid wasting exotic material. These particles need to be large enough to possess adequate surface architecture for the FTS activity to occur, but small enough for adequate dispersion and exposed surface to volume ratio. These competing factors seem to meet an optimum at a Co particle size in the order of 8 - 10 nm (Bezemer et al., 2006).

Cobalt has strong interaction with the, say, alumina support and the noble promoter is required in order to reduce the Cobalt metal to an active state for FTS. It has been argued that noble metals may have other benefits in terms of its influence on selectivity, poison abatement, activity regeneration and so on. (Eschemann et al., 2016)

Catalyst Deactivation

Despite the careful design of FT catalysts, it is true to say that in all cases, activity diminishes over time. This deactivation is thought to be the result of a combination of many deactivation mechanisms, principally:

Poisoning of the catalyst by sulphur or chloride is remedied by the inclusion of upstream guard beds in the process flow.

Cobalt oxidation or formation of compounds with the support, though likely to be slow at commercially relevant FTS conditions, is now widely regarded as a deactivation mechanism (van de Loosdrecht et al., 2007).

Sintering of cobalt nanoparticles to form agglomerates, thereby reducing effective surface area is also a likely mechanism. This could explain the approximately asymptotic nature of some observed activity curves. Perhaps cobalt nanoparticle sintering occurs early on in the time on stream (TOS) and reaches a limit when the apparent dispersion prohibits

further bridging (Phaahlamohlaka et al., 2018). Sintering is commonly discussed on the cobalt nanoparticle level. However, the present study also considers inter-pellet bridging in particulate catalyst. The point-to-point contact of each pellet reduces the effective surface area for the whole pellet on a μm scale.

Wax Holdup within the macro-pores of the cobalt-doped support may be responsible for reducing diffusion rates and inhibiting adsorption of the synthesis gas (syngas). With cobalt catalysed, low temperature Fischer-Tropsch (LTFT) the tendency is to produce longer chain wax. If the temperature is low enough, some of these waxes may maintain a degree of viscosity foul the pores.

Carbon lay-down, observed and discussed at length in the present work, has been identified as a potential cause of deactivation (Moodley, 2008). The present work identifies the operational cause and effect of carbon formation and its implications.

Long term operational data from commercially relevant micro-channel Fischer-Tropsch reactors is rare in the literature, Especially those focussing on the degradation of performance and catalyst deactivation. Most studies perform accelerated studies by externally increasing the temperature or spiking the inlet gas with contaminants. In a recent study, Steynberg et al. (2018) attempts this approach, operating for 100 day intervals, as compared to the 3-5 years of catalyst life required by micro-channel reactors for commercial viability.

Chapter 3. Linking Operational History and Post-operational Reactor Condition for Enhanced Operability of Modular Gas-to-Liquids Reactors

3.1 Chapter Abstract

Despite the careful design of Fischer-Tropsch catalysts, sustained catalyst life remains challenging. In order to achieve economic viability, recently announced small scale GTL plants must extend the life of catalysts by considering the overall reactor system. Deactivation mechanisms are found to be systemic to the reactor system as a whole and cannot be remediated simply by alternative catalyst design. Ultimately, end-of-life catalyst becomes more susceptible to unstable operation. The result is an increased likelihood of thermal runaway; manageable on a pilot scale, but potentially damaging at commercial scale. The effect is seen to be an acceleration in performance drop-off, sintering and carbon formation - all at the cost of catalyst life. Spent catalyst from end-of-run reactors must be rapidly unloaded for an economic turnaround with minimal asset damage.

Operating at the tipping point of catalyst performance and lifetime is found to present an additional challenge: the catalyst bed undergoes a mechanical change. Found to be enriched in carbonaceous fines, the particle structure of operationally unstable regions is all but erased. Instead, mechanically fast blockages present in the spent reactor are a hindrance to reactor turnaround and a liability to the costly reactors.

In this study, two similar prototype pilot scale Fischer-Tropsch reactors are benchmarked in extended runs. Their post-operational state is assessed by elemental analysis and imaging techniques. With a link to historic and online data, the analysis demonstrates that thermal events triggered by system-wide disturbances are correlated to the local presence of carbon-rich fragments of catalyst bed, restructured to form blocked portions, with increased resistance to flow. Methods developed in this work to measure and characterise the extent of the blockage show that both an axial and radial blockage profile are linked with operational history.

3.2 Experimental Methods

3.2.1 Fischer-Tropsch Synthesis

FTS was carried out at a dedicated pilot plant facility in prototype plate and fin fixed bed reactors with particulate 20% $Co(Ru) / Al_2O_3$ catalyst (Johnson Matthey, 500 μm diameter). In this study, two reactors are discussed; F-PRO-SP08 and F-PRO-SP09. (SP08 and SP09). Though ostensibly identical, the two reactors are subtly different owing to the inherent imperfections of the manufacturing process and catalyst loading techniques.

The prototype reactors used in this study are pilot scale representatives of CompactGTL's proprietary commercial reactor system, consisting of a multiplicity of process channels (fin side), 800 mm in axial length. The process channels are rectangular with face-on dimensions of 6 x 7 mm and are loaded with the catalyst previously described. These channels are arranged in four rows alternately with a multiplicity of coolant (plate side) channels. The coolant travels co-currently with the process gas.

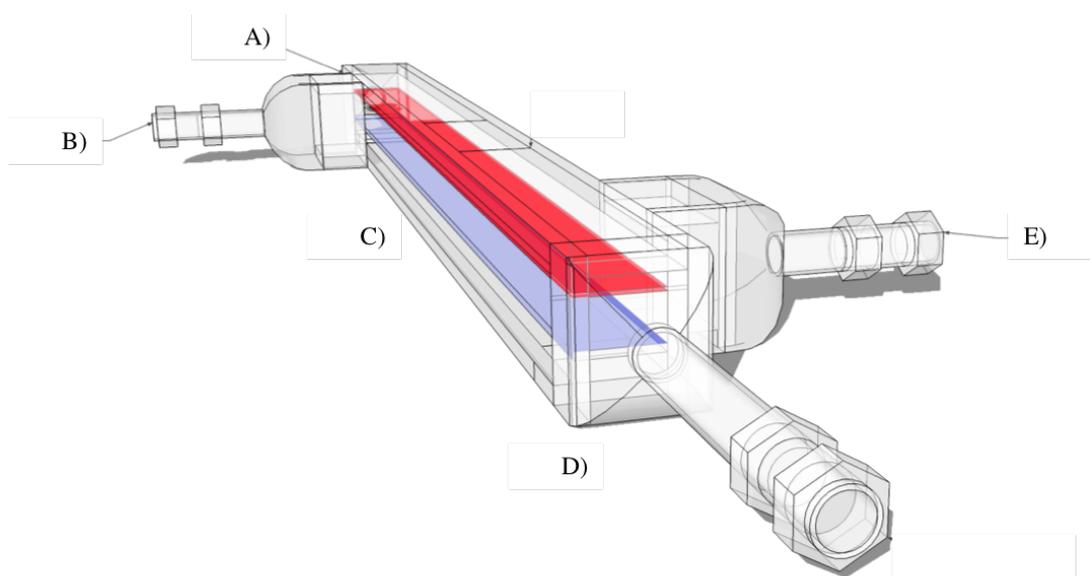


Figure 3.1: Schematic to show prototype reactor body arrangement with a) process inlet header, b) coolant inlet header, c) internal thermowell plates A&B (red and blue respectively), d) process outlet header, e) coolant outlet header.

The plate-side heating/cooling fluid is an eutectic mixture of 26.5% biphenyl and 73.5% biphenyl ether, held at appropriate temperature and pressure by means of an external heating and cooling loop. Both heating and cooling of the process side are supported depending on the mode of operation. The coolant flow-rate, for most of the study, was held at a high constant value (ΔT from coolant inlet to coolant outlet was not more than 0.2 $^{\circ}C$).

Mass flow controllers and pressure regulators along with gas pre-heaters control the process inlet gas conditions. Syngas is typically has a $H_2 : CO$ ratio of 1.9, though the effect of

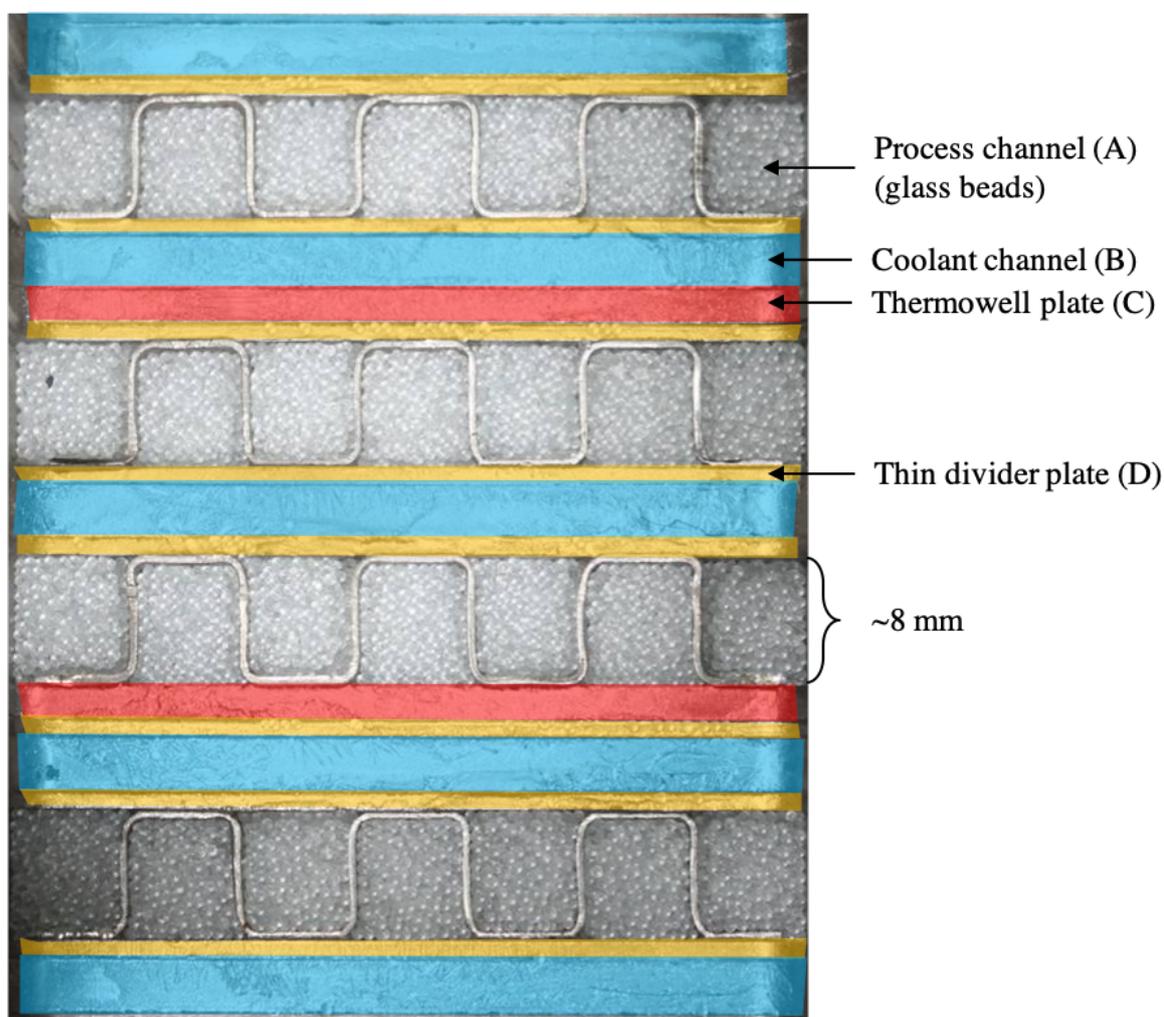


Figure 3.2: False colour photograph to show; A) Process channels filled with glass retention beads. B) Coolant layer. C) thick internal thermowell plates. D) thin plate to separate coolant and process. Overall reactor stack arrangement: BDADBCDADBDACDBDADB

varying composition is a controlled variable in some portions of the run. Process product is decanted in staged separation vessels to allow the segregation of water, oil, wax and tail gas. The composition of tail gas is routinely analysed by gas chromatography (GC) to close the mass balance and allow for the calculation of the FTS key performance indicators; productivity, selectivity, conversion etc. Offline physical and chemical analysis of the liquid and solid products is undertaken at time points throughout the run, in order to characterise the changing product profile during the catalyst lifetime.

A so-called operational ‘steady state’ is maintained by fixing conversion to a target figure. Changes in catalyst activity owing to deactivation and other deleterious changes in process conditions (e.g. changing syngas ratio, reaction pressure, inert partial pressure) are counterbalanced by increasing the applied temperature. (T_{app}). Thus, differential performance is often measured in terms of ‘loss of °C’; the amount by which the applied temperature must be increased to maintain conversion at a target notwithstanding the changes made.

3.2.2 Determination of pre- and post-operational differential pressure

The process side differential pressure is of particular interest since it may give some indication of the state of the bed at a given time during the extended run. An important subdivision of differential pressure should be made at the outset:

- Reactor differential pressure, or the pressure drop ‘through the headers’ is denoted by δp in units of bar. This value is measured online for a general metric on overall reactor state. This is achieved by subtraction of upstream and downstream pressure transmitter values. δp is closely correlated to the velocity of the gas- a variable that is often changed during an extended experimental campaign. For this reason, it is often more useful to refer to $\delta p/V$ in order to compare the state of the bed and not of the fluid.
- Single-channel differential pressure, denoted by δp in units of mbar, is the pressure drop across a single channel as given by purpose built measurement apparatus.
- Arising from work in this report, interesting results were discovered with regards to post-operational δp analysis. Further work, also discussed, included pre-operational δp . This allows for the calculation of the change in δp over time (particularly the difference between pre and post operational δp) and is referred to as $\delta p'$.

The δp measurement apparatus (Figure 3.3) consists of a digital differential pressure manometer (Extech Instruments 407910) attached to an impulse line consisting of a Swagelok 1/8" nozzle, Brooks mass flow controller and Swagelok back-pressure regulator, feeding air at a flow rate of 0.9 l min^{-1} and measuring the back-pressure generated. δp is sequentially measured in duplicate for each channel.

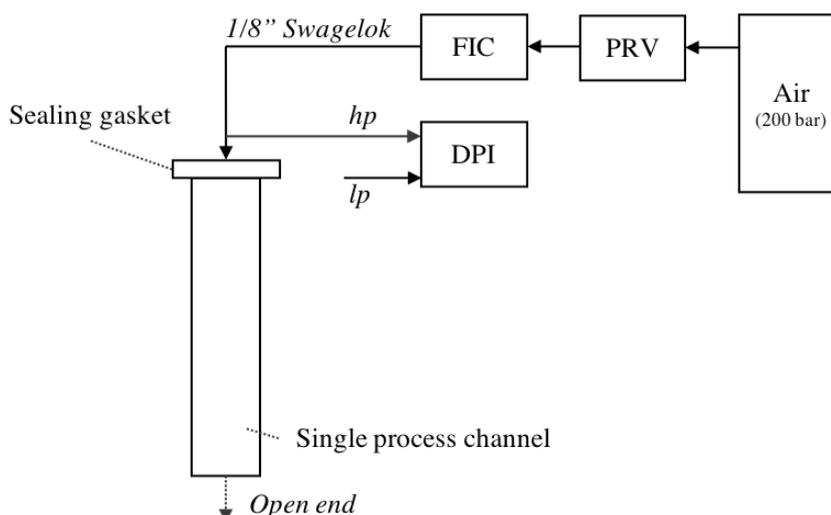


Figure 3.3: Schematic of δp measurement equipment. (PRV: Pressure Regulating Valve, FIC: Flow Indicator/Controller, DPI: Differential Pressure Indicator.)

3.2.3 Reactor Unloading

Reliably loading prototype reactors to deliver equal performance across each channel is the subject of much internal study and forms the basis of another body of work, save to say that there remain challenges in repeatedly achieving both optimum packing density and bed height. Challenges on the small scale are manageable, but small deviations at large scale equate to a significant financial impact. Remediation on commercial scale will also likely require capital-intensive solutions.

Unloading of spent catalyst from end-of-run reactors poses similar challenges. On the pilot scale, these challenges are mainly owing to the requirement to preserve sample integrity and separation for subsequent analysis. Methods developed in this work to systematically unload and record samples are applied with varying success depending on the degree of free-flowing particulate that remains in the channel. Commercially, the requirement to reliably and quickly unload tens of thousands of process channels containing hazardous material is a separate issue.

Reactor process side header manifolds are removed by cutting, exposing the channel structure. A retaining mesh and retaining glass beads are present on both the inlet and outlet face of the channels. The mesh is retained on the outlet side and the inlet mesh is removed. Retaining glass beads occupy typically the first 50mm of the channel depth from the inlet face. These beads are removed by vacuum suction using the vacuum collection system as shown in Figure 3.4.

The multiplicity of process channels is unloaded systematically, channel-by-channel in order to catalogue any findings. Undesired channels are first blocked by means of a flexible, inert quartz wool insert. This ensures that only material from the target channel is obtained or disturbed. Next, the following approaches are applied in order to extract samples.

‘Tipping and tapping’ is achieved by rotating the reactor from an ‘inlet up’ position through approximately 135-180° to an ‘inlet down’ position. Free flowing particulate is captured in a sample vessel and its mass is accurately determined on a calibrated microbalance (Sartorius R160P). Removal of internal thermocouples by pulling results in a slight loss of mechanical support offered by the presence of the thermocouple, and often yields a second free flowing sample. Further material may be collected by tapping the reactor body with a rubber mallet. The vibrations are often enough to allow yet more free flowing material to fall.

‘Gentle mechanical extraction’ employs the use of a thin, flat, 5 mm wide steel shim, which is slid between the fixed bed and the channel wall. This is often enough to break the wall support and may yield some fragments of the bed for direct sampling. This method is employed when a section of the fixed bed remains blocked and is not free flowing, but a sample of the blocked region is desired without too much mechanical disruption.

Mechanical disruption by ‘rodding’ the sample can quickly break up a blocked region. The disadvantage to this method is that it generates particulate ‘fines’. It is assumed that the rodding contributes to the break up of individual catalyst coated pellets and therefore distorts the sample before analysis. The advantage is the ability to quickly clear and segregate a whole channel for sampling; chemical composition of the sample is not altered, nor is depth information lost.

Vacuum suction of individual channels is also possible with the use of the aforementioned suction apparatus, shown in Figure 3.4.

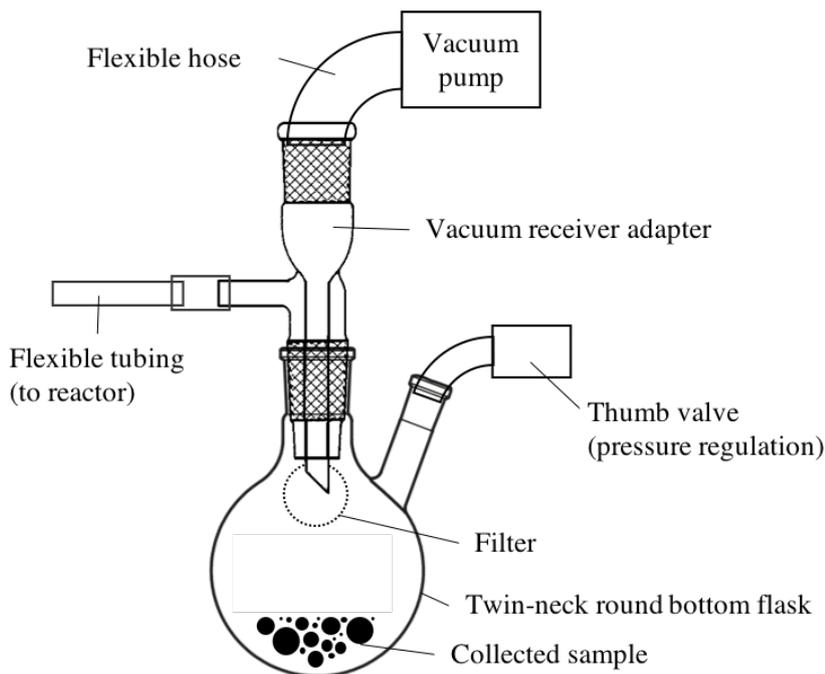


Figure 3.4: Schematic of vacuum suction apparatus

The advantage of vacuum suction is the lack of mechanical intervention. This method removes the doubt of fines created during the unloading procedure. Any fines present in the sample have originated in the operation of the reactor.

At each interval, the mass of the segregated sample is determined by the Sartorius R160P microbalance, along with depth information obtained by measuring the insertion depth of a narrow metal rod to the point at which contact is made with the remaining fixed bed. Periodically, images of the exposed bed surface are recorded by means of an Olympus iPlex TX industrial endoscope. Endoscopic images are a powerful method of analysing the undisturbed packed bed surface.

3.2.4 CHN Analysis

An intellectual property agreement prohibited the reverse engineering of the proprietary Johnson Matthey catalyst pellets in great detail. This limited the initial scope of the

investigation to a search for contamination and residue on the post-operational catalyst samples. Necessarily, a baseline for comparison allowed the same analysis to be carried out on pre-operational catalyst.

Short of quantifying the complete elemental composition of the samples through analytical techniques such as ICP-MS, a relatively inexpensive and quick method for scanning for contaminants was required. As discussed, the focus of this study was to determine the cause in loss of performance, or deactivation of catalyst. It has been noted that the cause of deactivation could be one of a combination of; poisoning, sintering, wax blocking, carbon laydown or oxidation. Poisoning of the catalyst is ruled out by the presence of guard beds upstream of the reactor. There are four guard beds in series:

- **Soda lime:** an alkaline bed to remove acidic gases (ie HCl and COS)
- **Alumina:** a second acid capture bed, which also acts as a drying column.
- **Zeolite 5A:** A high surface area, high temperature bed to capture $Fe(CO)_5$ and $Ni(CO)_4$. Carbonyl compounds are formed when the inlet syngas reacts with the metallic surface of the steel pipework in the system. These carbonyl contaminants are stable at room temperature but decompose on the hot zeolite surface to yield Fe and Ni metal and release CO into the gas stream.
- **SDS:** A commercial desulphuriser, comprising a CuO and ZnO mixture on an alumina support. The CuO is reduced to Cu using H_2 at temperature. This Cu and the ZnO capture trace sulphur compounds along with oxygen. The SDS bed operates at 50 °C to suppress the formation of CH_3OH from CO and H_2 reaction on a Cu catalyst.

Since poisoning is effectively negated, the remaining mechanisms – carbon laydown, wax holdup, oxidation, sintering would all change either the chemical composition or physical arrangement of the bed. Flash combustion of the extracted samples and analysis of the combustion products is a simple way to determine the molar composition of the samples.

All samples were tested on a Carlo Erba 1108 Elemental Analyser controlled with CE Eager 200 software, run in accordance with the manufacturer's instructions. The instrument is calibrated with acetanilide organic analytical standard. In brief, the method uses flash combustion of ground catalyst samples, which have been weighed into tin capsules using a calibrated microbalance (Mettler MX5). Samples are placed inside the auto-sampler, which is purged with helium. When the auto-sampler rotates, the samples are dropped into a vertical quartz combustion chamber, which is kept at a temperature upwards of 900 °C. When the sample is dropped, a pulse of oxygen gas is injected and the tin capsule containing the sample, which undergoes flash combustion. The helium carrier gas takes the gases produced from the flash combustion through the combustion reactor, which contains the oxidation catalysts copper (II) oxide and platinised alumina. All of the carbon present

in the sample is oxidised to CO_2 . Similarly, all the nitrogen is oxidised to NO_x . Hydrogen is oxidised to H_2O vapour. The helium carrier gas then takes the combustion products into a subsequent quartz reduction column containing copper wire, which acts as a reduction catalyst. The reduction column reduces the nitrogen oxides to Nitrogen gas. The resulting mixture of gases is carried by the helium carrier gas through a chromatographic column, which separates the CO_2 and N_2 gases. The separated gases and water vapour are detected by a thermal conductivity detector.

As a cursory examination of SP08, a reduced set of four samples was analysed in order to validate the method. As confidence in the technique was bolstered by the findings in this first experiment, SP09 was analysed more comprehensively; a set of 40 samples were analysed for CHN, targeting a range of sample depths and channel locations.

3.2.5 Scanning Electron Microscopy

In order to visually assess the catalyst surface and quantify the relative concentration and distribution of surface contaminants, a Hitachi S2400 Scanning Electron Microscope (SEM) was used, fitted with an Oxford Instruments Isis 200 ultra-thin window X-ray detector (EDX detector). Samples were loaded and sputter coated with a layer of atomic carbon. This was necessary to improve the conductivity of the samples and allow them to be seen by the microscope.

3.3 Results and Discussion

3.3.1 FTS Operational History

Reactor Performance

Two extended runs were carried out; the first one using the SP08 reactor and the following one using SP09. During the runs, key performance indicators (KPIs) are monitored and recorded. The SP08 run was interesting for a number of reasons. A series of thermal excursions throughout the campaign were suspected to be correlated to a gradual loss in performance of the reactor system. One theory was that thermal events were ‘knocking out’ whole channels, reducing the apparent available surface for reaction. The effect was not so straightforward but the ‘knock out’ effect was taken as common parlance. When the campaign was brought to an end, an attempt was made to explain the cause and effect relationship between the loss in performance and the post-run state of the reactor.

In the SP09 run, operation was much more cautious- temperature ramp rates were slower, gas flow rates were adjusted incrementally etc.- all in an attempt to avoid the thermal events seen in SP08. For the early part of the campaign, this was successful. However, when

pushing the maximum performance at high space-velocity, a thermal runaway occurred, which lead to a similar reduction in performance as seen in SP08.

Figure 3.5 shows a summary of KPIs achieved during the SP09 run. For most of the run, syngas ratio was not a controlled variable, and was held constant. SV was adjusted, the reactor allowed to settle, then temperature was added until a target conversion of 40% was achieved. This can be seen by the strong positive correlation between SV and T_{app} in Figure 6 for the first 1400 hours on stream. Naturally, as more gas is being fed, and conversion is being maintained at the cost of temperature input, the productivity of liquid hydrocarbons (PC_{5+}) is steadily increasing. Another important parameter to monitor, the methane selectivity (SCH_4) is shown to steadily increase throughout the run. Methane selectivity could be a good indicator of catalyst deactivation.

After the thermal runaway at approximately 1400 hours on stream, even a return to pre-trip SV could not recover performance. The applied temperature was increased significantly in an attempt to maintain conversion, resulting in a number of thermal runaways and further decreased performance.

As a general rule of thumb derived from parametric studies of previous reactor campaigns, the relationship between controlled variables and several key performance metrics can be described by 2.2.

Thermal Events

Five thermal excursions occurred during the latter half of the SP09 run. Owing to the rigorous safety integrity levels protecting the reactor and facility, typical thermal excursions are not a safety concern. Instead, a thermal runaway is an uncontrolled imbalance of heat removal and heat generation capacities of the system. Runaways can be allowed to go to completion if the exotherm is minor. In the event of a serious exotherm, the experiment is halted by the trip logic of the system and nitrogen purges the reactor. Figure 3.6 shows the anatomy of a typical thermal event. The sequence of events can be outlined with reference to the graph.

Firstly (white circle) the ΔP across the reactor (yellow trend) spikes due to a double dump of the downstream product section decanters (levels are shown by the blue trend). In response to the drop in reactor pressure, the temperature at the inlet of the reactor (taken from a thermocouple 50 mm into the bed, orange trend) drops sharply (orange circle). This is due to a sudden intake of gas from the feed of the reactor to equalise the pressure caused by the opening of product decanters. The reduction in temperature at the inlet side of the reactor means that the local conversion is reduced. This results in a subsequent higher partial pressure of unreacted syngas in the outlet side of the reactor. The higher than normal concentration of reactants in the downstream portion of the reactor gives rise

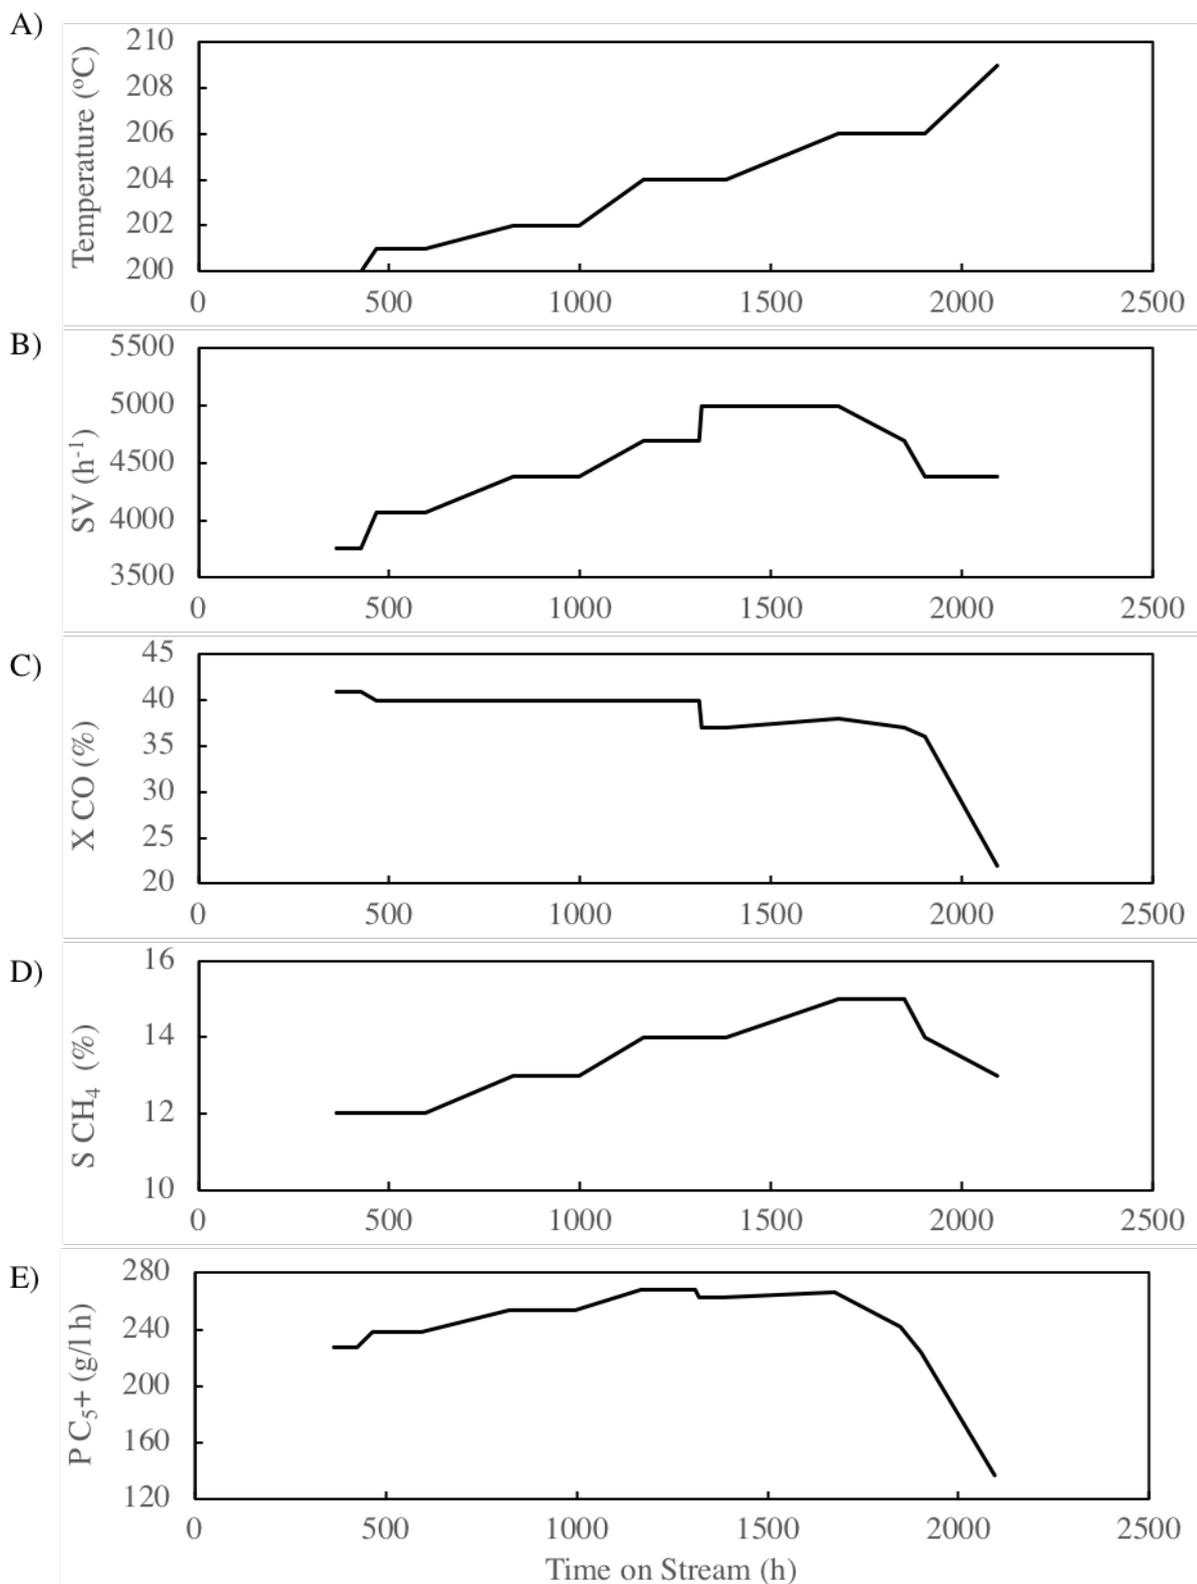


Figure 3.5: Graphical summary of a typical FTS run (SP09 shown)

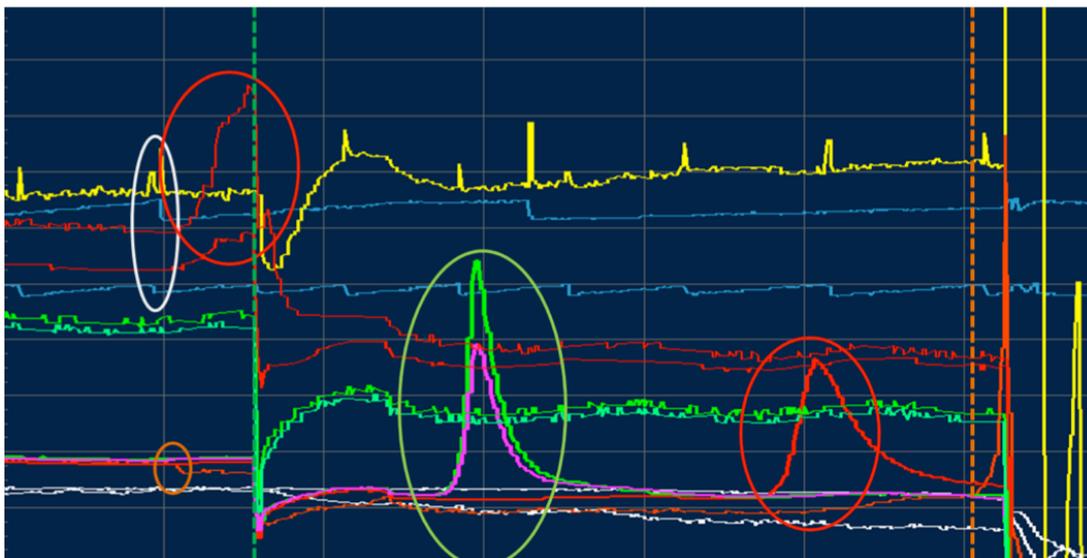


Figure 3.6: DCS Screenshot to show typical operator and system response to a thermal event. (Data is scaled)

to a higher than normal local reaction rate- the thermal equilibrium between the process and coolant side is disturbed. (First red circle) The outlet of the reactor (thermocouples targeting catalyst bed 40 mm from outlet) has a thermal excursion. Simultaneously, the reactor plate (metal) temperature in same location traces this rise. Next, operators reduce the set point temperature of the coolant, in order to cool the reactor. (The green trends show the coolant temperature, green dashed vertical line shows the time of the set point change) However, the temperature controller overshoot the larger than normal step change, but subsequently recovered.

After the control action and apparent recovery, (green circle) plate (metal) thermocouples (pink and green trends – 200 mm from inlet) have a large spike in temperature. No thermocouples are present in the bed at this depth, and so the local internal temperature is unknown but predicted to be significantly higher than the corresponding metal temperature. Following this, another plate thermocouple 100 mm from inlet, shown in the second red circle) has large spike in temperature. Again, no internal thermocouples were located at this depth; the internal temperature is predicted to be high. Finally, the thermocouple (orange trend) that experienced the initial front-end temperature drop has an exponential thermal runaway and trips plant.

3.3.2 Pre- and Post Operational Differential Pressure

Theoretical Determination of δP

Owing to the narrow channel geometry and particle size, pressure drop is heavily influenced by the effect of the wall. For this reason, an attempt at theoretically predicting single channel δP is problematic. The Ergun equation is commonly used in the design and

theoretical assessment of packed beds (Ergun, 1952). However, the universality of the constants has been questioned. Indeed, Foumeny et al. (1993) do away with the constants entirely and replace them with a consideration for the so called ‘diameter ratio’ – the ratio of tube to particle diameters. Foumeny et al. (1993) consolidate many previous attempts to consider this ratio as a parameter in the calculation, and generally achieve a better prediction to empirically measured values over a range of conditions.

Conditions considered in Foumeny et al. (1993) are an order of magnitude higher in flow velocity to the system considered in the empirical measurements of this report, and are limited to particles of a uniform size, packed into cylindrical channels. However, applicability of the Foumeny et al. model (referred to informally as ‘Ergun-Foumeny’ hereafter) is assessed in this work. If an Ergun-Foumeny value is close to an empirical value for the system in question, notwithstanding the limitations of application listed above, it might be a useful tool to predict how proposed future changes in the system (particle size, tube dimension etc.) would affect the δP . Furthermore, substitution of empirical results into the Ergun-Foumeny equation may be able to predict equivalent input parameters and help describe the difference between two systems. An example Ergun-Foumeny calculation applied to the loading of SP09 is shown below.

The starting point for Ergun-Foumeny is the determination of the tube to particle diameter ratio, d_r .

The particle diameter is limited by the fact that the system has a particle size distribution. Using a mean particle size estimated from measured the specification will allow for an initial calculation to be made. Since this is a strong limitation of the equation, it may be possible in future to compare measurements from a uniform particle obtained by sieving. Alternatively, substitution of the empirical value would allow calculation of an equivalent particle diameter which considers particle size distribution. In the initial case, however, a particle size of 600 μm is carried forward.

The tube diameter, d_t presents another limitation to the Ergun-Foumeny. The ‘rectangular’ channel must be summarised by a single parameter (in reality, manufacturing limitations mean that even a ‘rectangular’ description of the channel is not sufficient). Commonly used conversions are hydraulic diameter and equivalent diameter. Equivalent diameter is the diameter of a circular channel that gives the same pressure loss as an equivalent rectangular channel, which is calculated by the ratio of side dimensions:

$$d_t = 1.3 \left(\frac{(ab)^{0.625}}{(a+b)^{0.25}} \right) = 1.3 \left(\frac{(6 \times 7)^{0.625}}{(6+7)^{0.25}} \right) = 0.00708 \quad [\text{m}] \quad (3.1)$$

Therefore, the tube to particle diameter ratio is given by:

$$d_r = \frac{d_t}{d_p} = 11.80 \quad (3.2)$$

The d_r is significantly less than 50, where Ergun-Foumeny has shown to be most applicable. A rule of thumb for packed bed design is to have a minimum d_r of 10. At ratios of $50 > d_r > 12$, the Ergun equation tends to suffer by not considering wall effect. At $d_r < 12$, as observed in the current system, the common assumption of voidage for spherical particles becomes inaccurate. One strength of Foumeny et al. is the ability to predict the voidage using empirical models rather than rules of thumb based on assumptions.

Since $d_r < 1 < \sqrt{3/2}$, mean voidage is given by:

$$\varepsilon_m = 0.383 + \frac{0.25d_r^{-0.923}}{\sqrt{0.723d_r - 1}} = 0.3293 \quad (3.3)$$

This is seemingly close to the rule of thumb value of 0.4, but since the Ergun equation employs voidage terms raised to the power 3, it is crucial to accurately define voidage. Since voidage relies on the ratio of two parameters which have been identified as problematic in the current system, it may be possible in future to empirically measure voidage. In the current work, however, the Ergun-Foumeny voidage is carried forward as a best estimate.

The well known Ergun equation is given by (Ergun, 1952):

$$\frac{\Delta P}{L} = \frac{150\mu u(1 - \varepsilon_m)^2}{d_p^2 \varepsilon_m^3} + \frac{1.75\rho u^2(1 - \varepsilon_m)}{d_p \varepsilon_m^2} \quad (3.4)$$

An important step in Foumeny et al. (1993) was to include the rearrangement in (Macdonald et al., 1979):

$$\frac{\Delta P d_p^2 \varepsilon_m^3}{L \mu u (1 - \varepsilon_m)^2} = B \frac{\rho u d_p}{\mu (1 - \varepsilon_m)} + A \quad (3.5)$$

The advantage of the Macdonald et al. (1979) rearrangement is the linearisation of the equation into the form:

$$f' = B Re' + A \quad (3.6)$$

Where f' is the modified friction factor and Re' is the modified Reynolds number.

Foumeny et al. (1993) experimentally derives the constant, A , which is found to be 130. The constant B is replaced in Foumeny et al. (1993) by the ratio dependent regression coefficient given, for the current system by:

$$B = \frac{d_r}{0.335d_r + 2.28} = 1.893 \quad (3.7)$$

Substitution of A & B into the linear equation gives:

$$f' = 1.893 \frac{\rho u d_p}{\mu(1 - \varepsilon_m)} + 130 \quad (3.8)$$

For the current work, measurements of δP were taken in air at 0.9 l/min, therefore the velocity is given as:

$$u = \frac{Q}{A} = 0.357 \quad [\text{m s}^{-1}] \quad (3.9)$$

Note that velocity was calculated by taking the area of the rectangular channel rather than calculation of the equivalent area of the circular tube. Approximate values of density and viscosity are relied upon since accurate temperature measurement was not available for the test. Values taken at standard conditions have been used.

$$f' = 1.893 \frac{1.275 \times 0.357 \times 0.0006}{0.00001845(1 - 0.39)} + 130 \quad (3.10)$$

Now, replacement of the modified friction factor with the system parameters allows δP to be calculated by rearrangement:

$$f' = 176.1331 = \frac{\Delta P d_p^2 \varepsilon_m^3}{L \mu u (1 - \varepsilon_m)^2} \quad (3.11)$$

Rearrangement for the δP across a $6 \times 7 \times 800$ mm channel yields a predicted pressure drop of approximately 158 mbar:

$$\Delta P = \frac{176 L \mu u (1 - \varepsilon_m)^2}{d_p^2 \varepsilon_m^3} = 15769 \quad [\text{Pa}] \quad (3.12)$$

This compares to an average empirical δP of 214 mbar for SP09. Notwithstanding the limitations of the application of the equation listed above, this appears to be a significant under-prediction. Further work is necessary to predict the loaded δP for the current system.

Post Operational δP

Arising from the loss of performance in SP08, single channel δP became a new parameter of interest. Consequently, no pre-operational δP data is available for the SP08 run. This limits any strong interpretations being made since the change in δP ($\Delta \delta P$) is unknown. However, in hindsight, the values can be shown to be significantly higher than values

predicted by the Ergun-Foumeny equation or by comparison with later data gathered from SP09 (which has identical geometry). This shows that, in all likelihood, SP08 has a highly positive $\Delta\delta P$. Table 3.1 shows the post operational δP for SP08.

Table 3.1: SP08 post-operation δP (mbar)

Row	Post-operation δP per column (mbar)						
	A	B	C	D	E	F	G
1	822	830	825	850	926	830	706
2	402	490	476	452	470	450	330
3	674	830	812	780	845	880	746
4	406	404	350	412	400	382	326

Rather clearly, there exists a row-wise pattern in post operational δP . Rows 1 and 3 exhibit significantly higher δP (average = 811 mbar) than the alternate rows, 2 and 4 (average = 410 mbar). This surprising result implies that there is some systemic effect at work.

Some of the theories as to the origin of this effect can be assessed:

Differential local thermal performance

Recalling Figure 3.2, it is possible to notice a subtle lack of regularity between the arrangements of the various plates. The tube plate above row 2 is sandwiched between a thin plate and a cooling channel and is not in direct contact with the process channels. The tube plate below row 3 is not sandwiched in this way and is in direct contact with the FT Channels. This appears to be the only row-wise phenomenon. It is possible that the different layered arrangement alters the heat transfer gradient from process to cooling side. Poor heat transfer in some rows may contribute to an increased rate of deactivation in those rows by sintering, wax holdup or carbon deposition- all of which would result in a likely increase in δP . Though the lack of regularity is a row-wise effect, there appears to be no pattern by which it can be associated to the alternating δP pattern.

Differential loading

Since pre-operational δP was not recorded, it is not possible to show the $\Delta\delta P$. It is possible that alternating rows were loaded and packed differently, giving rise to a row-wise pattern of pre-operational δP . A comparison of loaded weights of catalyst to the post-operational δP is shown in Figure 3.7C.

Figure 3.7C clearly shows the groupings of high and low δP rows. Theoretically, an increased catalyst load within a channel of the same dimensions would indicate a better-settled bed (higher catalyst density), with reduced voidage. Lower voidage would result in a higher δP . However, this trend is not observed- indicating that there are subtle variations in geometry owing to imperfections in the manufacture of the reactors. Importantly, there

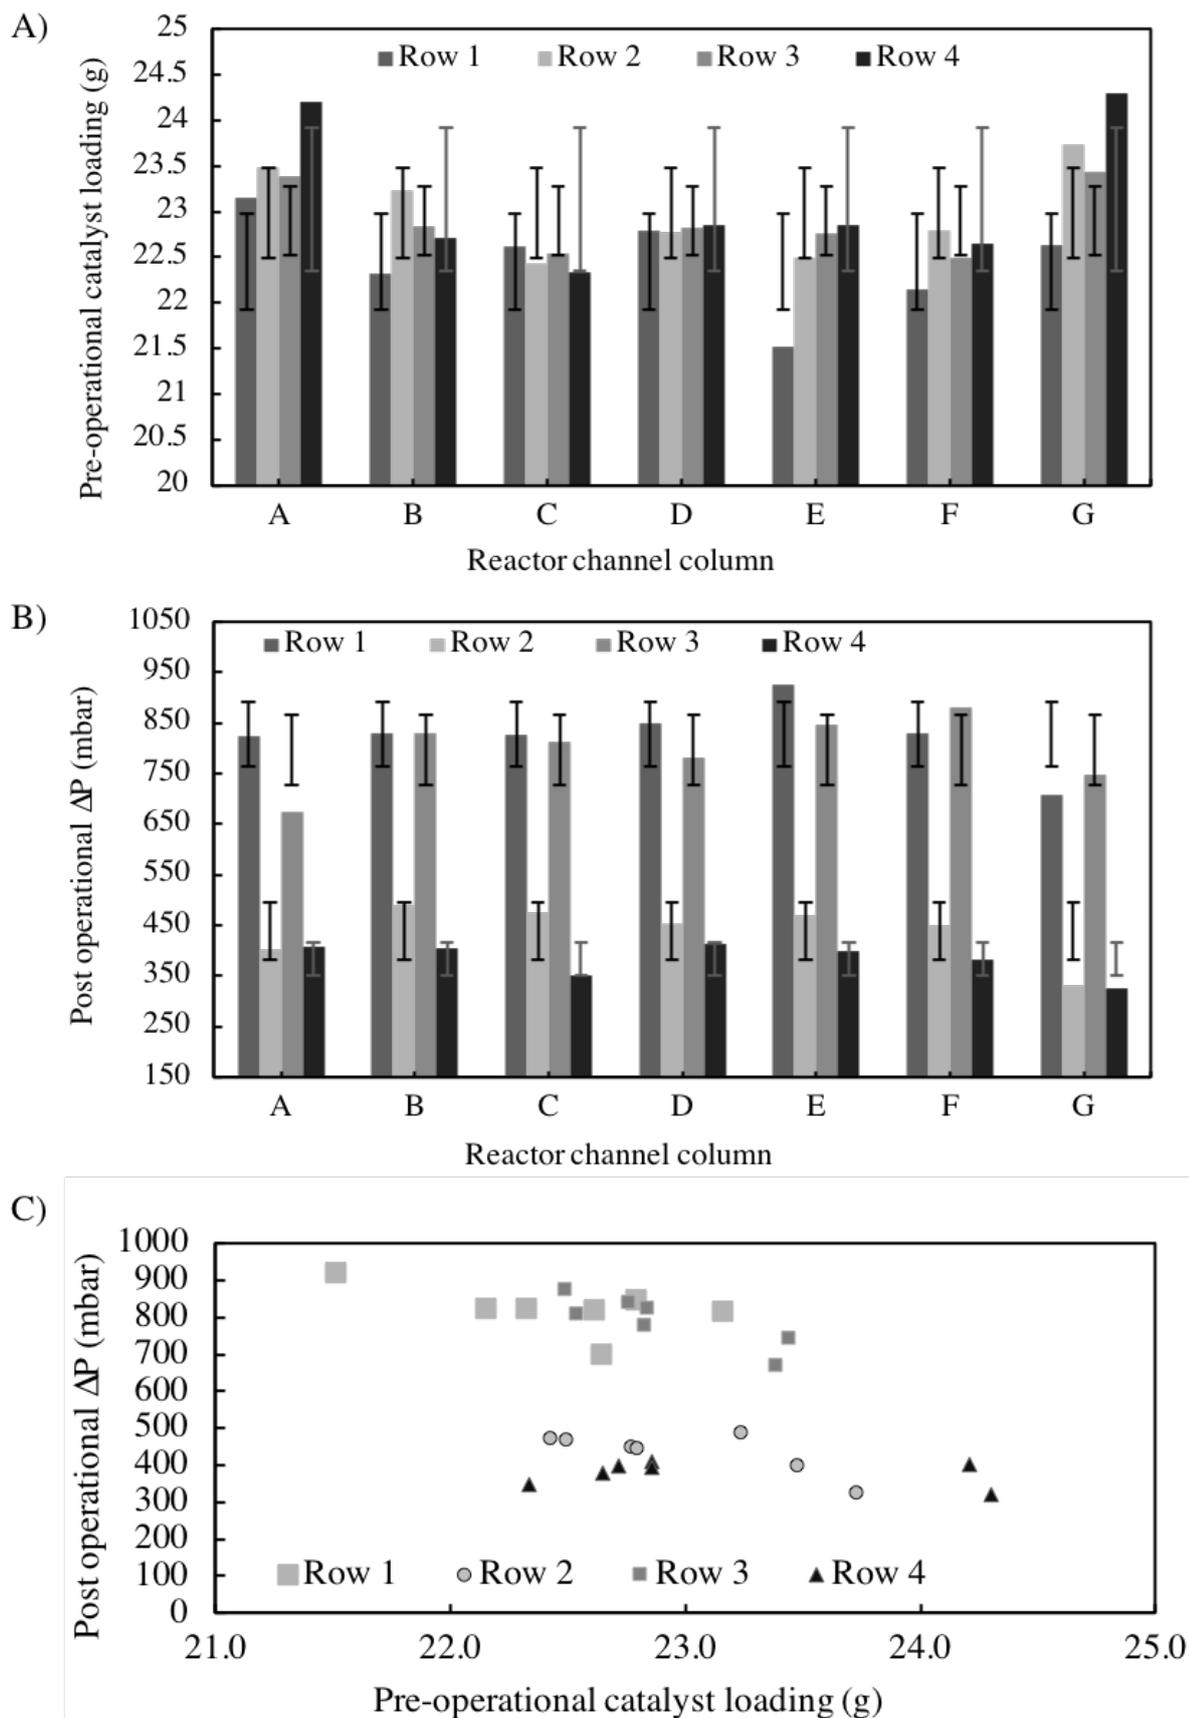


Figure 3.7: Scatterplot to show pre-operational catalyst loading versus post-operational δP , grouped by rows

is no grouping to suggest that higher δP channels 1 and 3 were loaded with more catalyst on average. In fact, no significant difference exists between the loading of any row (Figure 3.7A). Some potential outliers exist in terms of catalyst load, which can be explained with reference to Figure 3.7. A visual inspection of the reactor shows that edge channels (Columns A and G) are physically wider and show a lower δP on the whole. This is demonstrated by the subtle curvature for each grouping in Figure 3.7B. the increased width of edge channels also accounts for the significantly higher loading observed in channels A & G (Figure 3.7A)

Taken together, Figures 3.7A-C indicate that differential loading is not a dominant factor in post operational δP . Subtle variation in channel geometry does seem to influence δP and catalyst load, but this is not a row-wise effect.

Runaway propagation

Owing to the arrangement of the plate side coolant and channel dividing fins on the process side, it is likely in the event of a sustained thermal runaway that high temperatures are propagated preferentially in certain directions. Firstly, the flow direction of process gas would naturally force a heat source to spread axially downstream within the same channel. This propagation may also occur in the opposite direction, as discussed in the analysis of a runaway previously, owing to the differential local partial pressure of reactants. Next, fin-to-fin transfer would result in a local heat source spreading to adjacent channels within the same row. It is unlikely that heat transfer would occur from process row to row, since heat could not pass across a coolant row. This results in the potential propagation of thermal runaways or the damages caused by them to spread preferentially in a row-wise manner. The implication of this is that high δP rows should have seen a higher proportion of the thermal runaways during operation.

Comparison of pre and post operational δP in SP09

Following the results of the SP08 run, the opportunity was taken to obtain pre-operational δP data for the SP09 run, for later comparison.

Figure 3.8B shows that the pre-operational δP for SP09 is relatively uniform. With a mean δP of 214 mbar, and standard deviation of 15 mbar, the coefficient of variation (μ/SD) is 7%. Visual inspection of Figure 3.8B indicates that the sources of this variation are column A, and row 1. Since channels were loaded to an even height, rather than to an even weight of catalyst, this variance is assumed to be a product of differential settling and variations in geometry owing to imperfections in the manufacturing process. Notably, the values are all higher than the δP predicted by the Ergun-Foumeny equation described previously.

As with SP08, an online increase in reactor ΔP was observed. As expected this translated

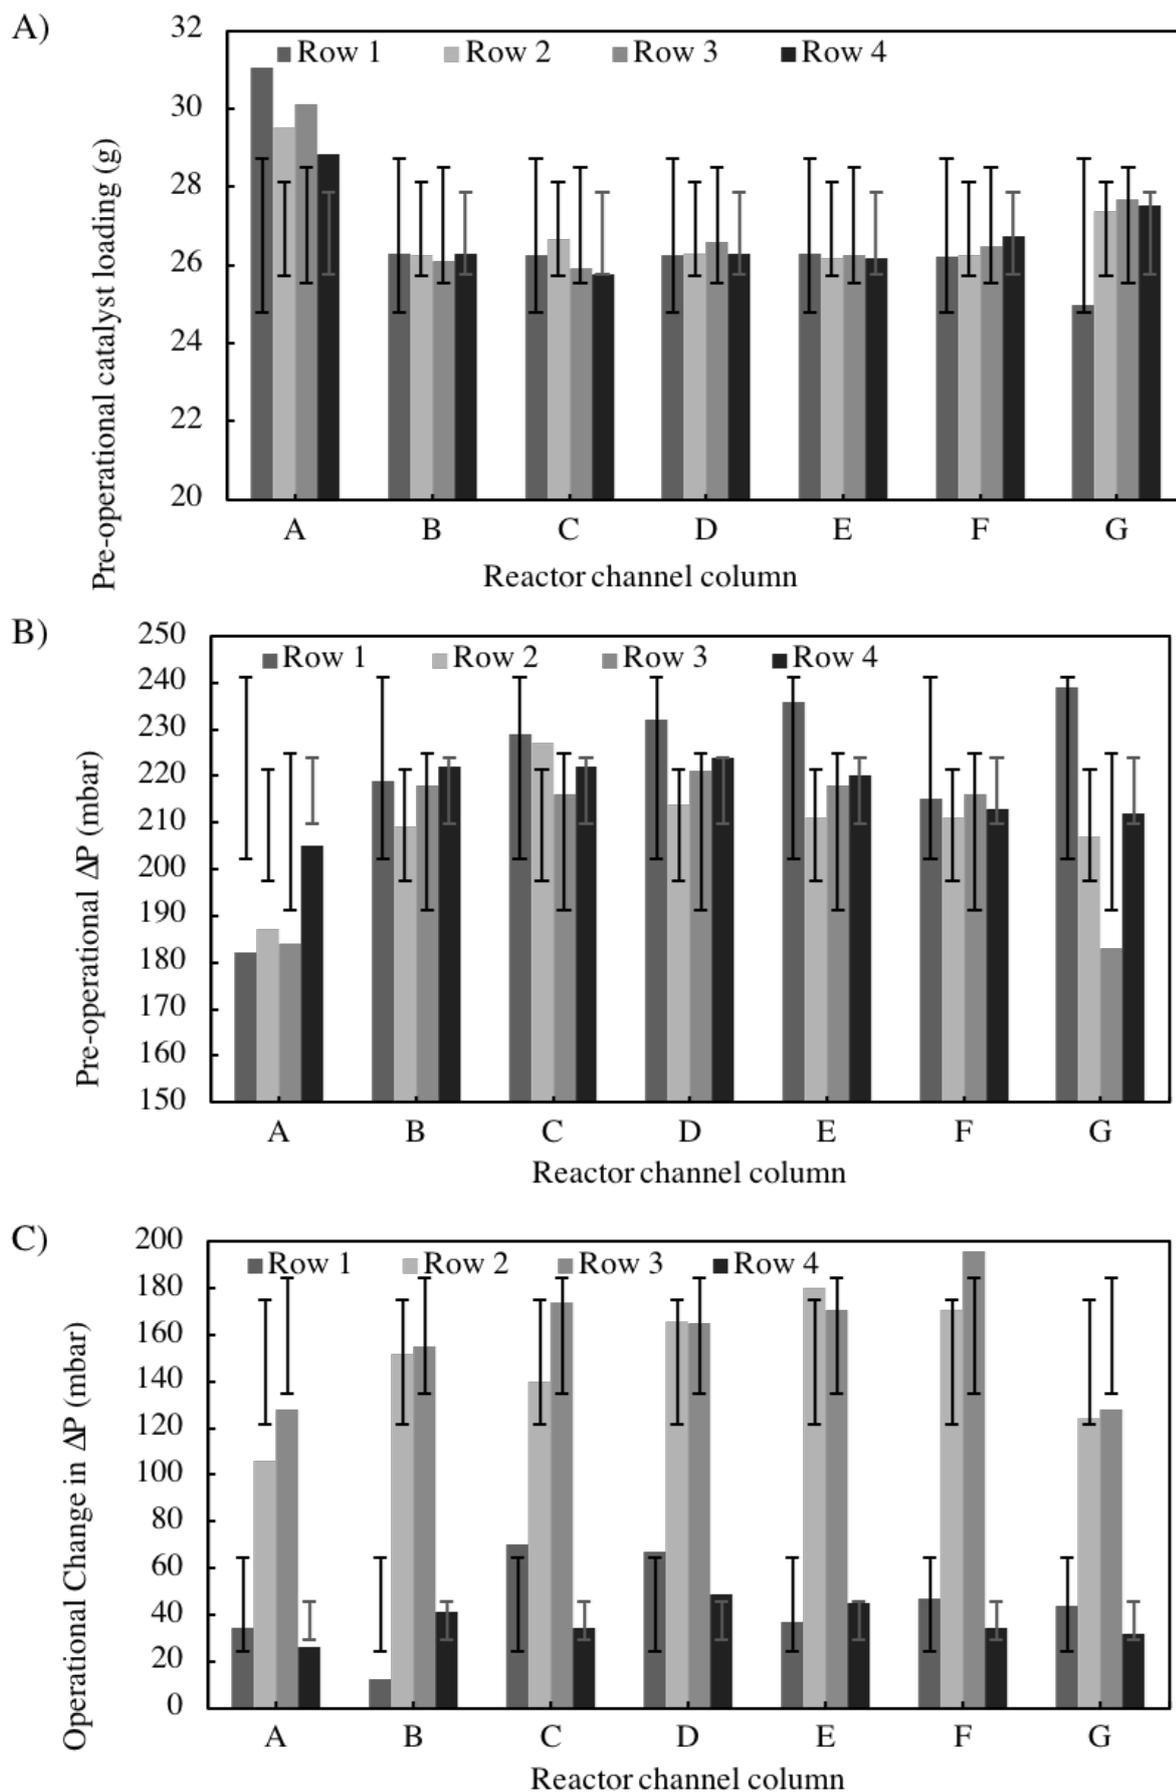


Figure 3.8: Pre- and Post-operational analysis of SP09 grouped by column and row. A) Pre-operational catalyst loading (g). B) Post-operational δP (mbar). C) Operational change, $\Delta\delta P$ (mbar). (Error bars show one row-wise standard deviation)

to an increase in post-operational δP for each channel. By subtracting, pre-operational δP , the $\Delta\delta P$ can be calculated. A positive value for each channel indicates that each channel suffered an increased pressure drop as shown in Figure 3.8C

Figure 3.8C ostensibly shows a row-wise pattern of $\Delta\delta P$. However, the pattern is not alternating as was seen with SP08. Instead, the two central rows show the highest increase in δP . Further analysis of the matrix is possible by means of a contour plot.

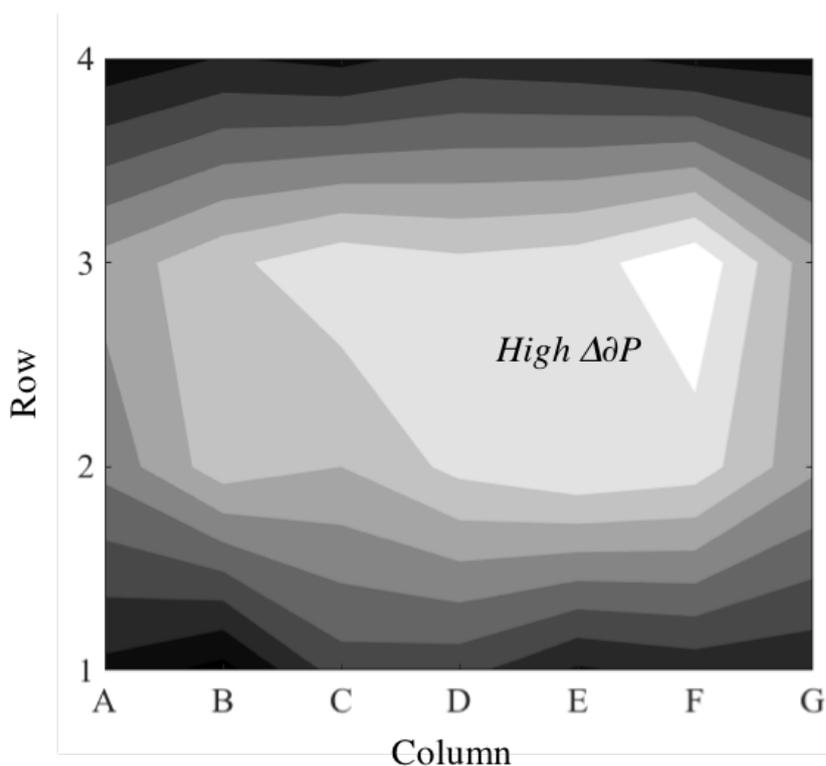


Figure 3.9: Contour plot to show interpolated $\Delta\delta P$ for SP09

Figure 3.9 shows an interpolated contour plot of the $\Delta\delta P$, which acts as a ‘heat map’ of $\Delta\delta P$ for the reactor. Clearly shown is the existence of a radial ‘hot spot’ relating to central channels- those not adjacent to an external wall.

The term ‘hot spot’ is convenient since theories shaped by the results in SP08 have involved temperature driven mechanisms for catalyst deactivation and damage which may also give rise to an increased δP .

For the first time, SP09 has hinted that heat transfer limitations causing thermal runaways are also leading to an increased back-pressure. A further explanation of a radial pattern of $\Delta\delta P$ is that the flow distribution is likely not uniform through the rectangular reactor header. Although a distribution plate and turbulent flow are present in the inlet header, it is likely that corner channels (and less so, edge channels) will receive a reduced flow compared to central channels. This reduced flow over the course of many hours on stream could result in a far fresher catalyst in these channels. The lack of FT activity comes with the lack of the proposed associated wax holdup, carbon laydown, sintering etc.- all

contributing to a lower $\Delta\delta P$ in corner and wall channels. However, sintering, carbon formation and product formation are all temperature driven. It may be no coincidence that edge channels, with greater ability to expend excess heat, are apparently less susceptible to these mechanisms.

3.3.3 Catalyst Unloading & Blockages

In SP08, it was noticed that channels were difficult to unload completely. There existed a mechanically tough blockage in most channels. Qualitatively, the extent of blockage was agreed to loosely conform to the row-wise δP pattern for SP08. This was the first hint that the increased δP was a result of some physical change in the bed other than the presence of solidified wax, as previously conjectured.

After initial δP measurements were taken for the post-operational SP09 reactor, it was necessary to attempt to unload the reactor of catalyst. This was carried out using the methods described previously. Some samples were retained for further analysis. In these channels, a complete axial profile of samples was collected. During unloading, it was clear that similar blockages existed in SP09. Due to further testing requirements, not all channels were completely unloaded. Instead, samples were taken from the channels indicated in Figure 3.2.

Table 3.2: SP09 sample location and quantity

Row	Column						
	A	B	C	D	E	F	G
1	4	2					
2	3		5		8		
3	2			2		4	
4	1						1

With all other channels, free-flowing catalyst was removed by vacuum suction, and the depth of the blockage was measured from the inlet and outlet of the reactor. The blocked portions of the reactor were left for further trials beyond the scope of the current work. Figure 3.10 summarises the blockage depth data for SP09.

Notably, no channels in row 4 of SP09 were blocked. Catalyst was completely removed by vacuum suction. Rows 1, 2 and 3 all had mechanically strong blockages in the regions shown in Figure 3.10.

Perhaps most strikingly, column G appears to have much larger blocked regions than other columns, with the caveat that column G was noted to be ‘tighter’ than other columns; the flexible rubber tubing used for vacuum extraction of the catalyst was compressed against the wall, making extraction difficult. The increased blockage depth in column G could

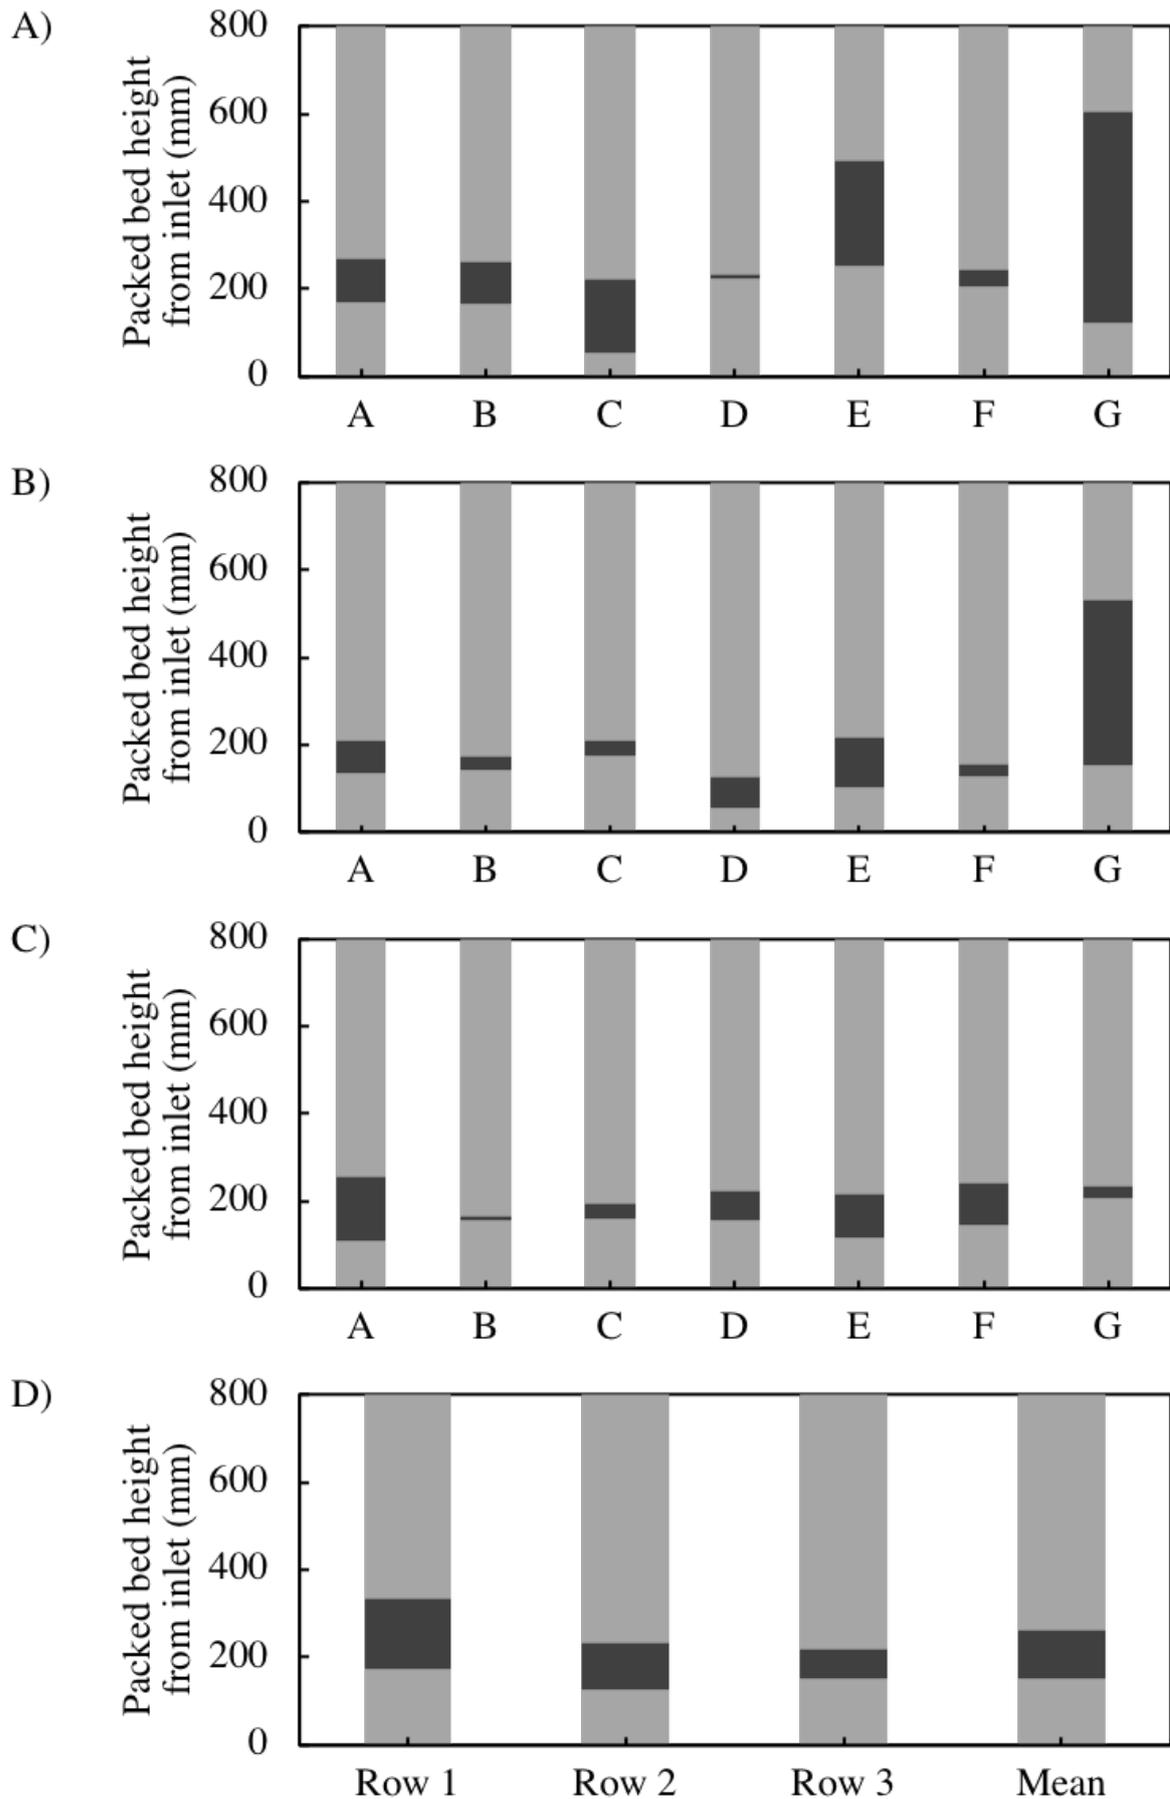


Figure 3.10: Graphs to show location and extent of blocked regions in SP09 (mm). A) Row 1. B) Row 2. C) Row 3. D) Row-wise average and global average.

therefore be for one of two reasons. Firstly, the difficulty in extraction meant that the figures obtained in column G were artificially high; Had there been effective extraction, the channels would have unloaded to similar depths. Secondly, the existence of ‘tight’ channels implies a reduced geometry- perhaps the reduction in channel dimensions had an adverse effect on the bed, forming more blockages. However, 3.8 does not seem to show a significant discrepancy in $\Delta\delta P$ between column G and other columns.

Figure 3.10 also demonstrates that, on the whole, blockages appear near the inlet end of the reactor. This can be further clarified by taking row-wise averages for each column, shown in Figure 3.10D.

As Figure 3.10D shows, the average depth from inlet of a blocked zone is 149 mm. The average length of a blocked zone is 113 mm. In order to assess the theory that thermal events are causing the blockage, it is necessary to assess the location of thermal events within the reactor. Table 3.3 shows the location of each recorded thermal excursion during the SP09 campaign. It is important to stress that temperature data is not available for every channel and at every depth. The location of thermocouples within the reactor is heavily influenced by requirements for tests not within the scope of this work. For this reason, an important comment may be made- the channel and depth that is recorded during a thermal event may not necessarily be the exclusive channel or depth at which the thermal event occurred. Table 3.3 demonstrates that this is true in two ways. Firstly, regarding the runaway event at 1386 HOS, a thermal excursion was observed on several thermocouples located in different channels at the same depth. This shows that thermal events can be radially distributed. Since this radial distribution transcends the row (and therefore coolant layer) boundary, it must be the case that the runaway is happening independently, rather than by conduction. The coincidence of location is therefore thought to systemic to the reactor, and not as a result of one event causing the other.

Secondly, during the runaway event at 1677 HOS, thermal excursions are observed on two thermocouples targeting different depths in the same channel. This shows that thermal events can be axially distributed. As shown in Figure 8, this axial distribution of a runaway is likely to be by causal propagation, analogous to a ‘fuse’ burning slowly upstream.

Since the recorded depth of runaways in Table 3.3 are limited by the spatial resolution of thermocouples, taking the average of all the data gives a more general indication of the general zone in which runaways tend to occur. Averaging depth, as well as X and Y coordinates for the runaway zone allows for the location of runaway events and blockages to be compared.

Figure 3.11 demonstrates that the average location of runaways and blockages are related. Figure 3.6 shows that runaways also tend to migrate axially in an upstream direction before they are detected by the nearest thermocouple within a channel. This may imply that there is an upstream skew on the average location of runaways when looking at the

Table 3.3: Location of recorded thermal events during SP09 run

TOS (hr)	Channel	Depth from inlet (mm)
1386	E1	170
	D2	170
	D2	290
	B3	170
	E3	170
1677	C4	50
	C4	110
1852	F2	50
1906	E1	170
2100	F2	50
Average		140

recorded thermocouple data. For any axially propagated runaway, it is possible that the area immediately downstream of the detected runaway has also undergone an undetected local runaway. This corresponds to the average blocked zone for the reactor being directly downstream of the average detected runaway. The local heat produced in a runaway might be the cause of the blockage. This might occur in one or both of the two heat driven mechanisms of catalyst deactivation discussed previously:

Sintering of the catalyst particles together to form agglomerates with reduced voidage and increased bulk mechanical strength. This would correspond to higher δP in the blocked zone. Necessarily, the temperatures required for sintering must be substantially higher than normal operating temperatures. Only channels that have experienced a local runaway would be capable of localised sintering.

Carbon laydown caused by higher temperature operation could occur as a result of cracking or alternative chemistry (e.g. Boudouard reaction, cracking). Carbon laydown could deposit carbon at the site of runaway or be carried a short way downstream before being deposited on nearby regions of the reactor. The deposition of carbon would block pores, coat particle surfaces and voidage, resulting in a higher local δP .

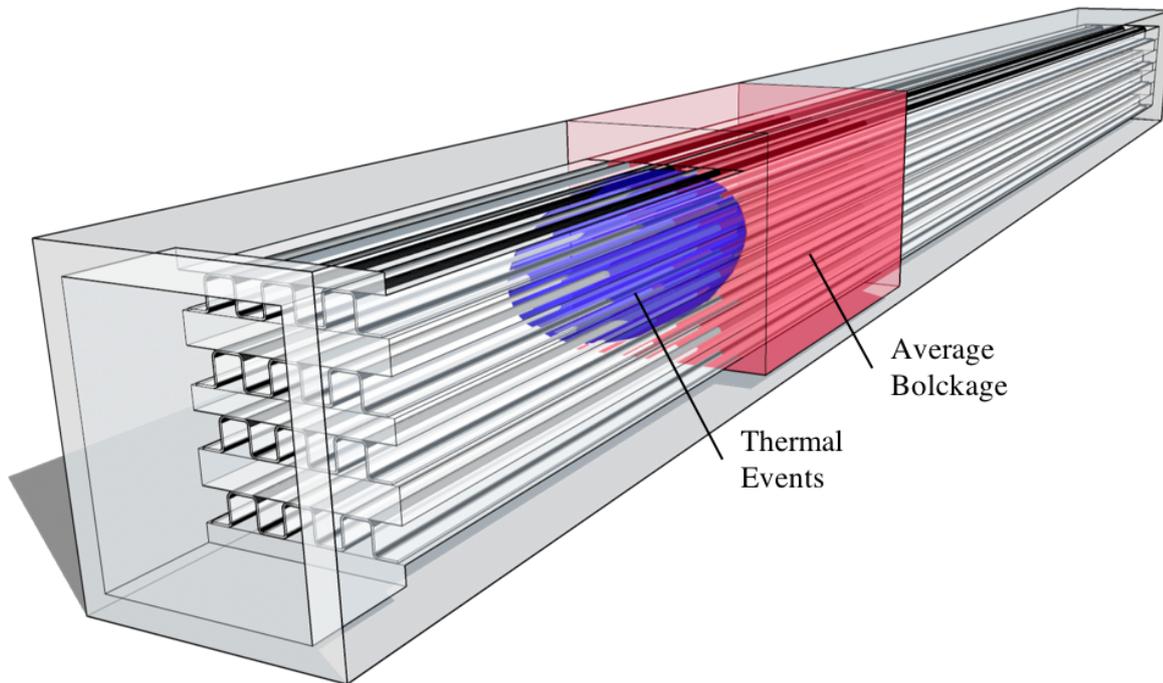


Figure 3.11: 3-Dimensional view from the inlet of SP09 showing average location of runaway (blue) and average location and depth of blockage (red)

Table 3.4: SP08 samples subjected to CHN elemental analysis

Sample	Description	Location	Extraction
1	Fresh Catalyst		
2	low δP free flow	E4	Vacuum from outlet
3	high δP blockage	E1	gentle mechanical from inlet
4	high δP fines	E1	106 μm sieve of sample 3

Table 3.5: Raw results for CHN composition (wt.% \pm 0.3%). Results are in duplicate.

Sample	N (wt.%)	C (wt.%)	H (wt.%)
1	0.26	0.07	0.93
	0.26	0.08	0.99
2	0.03	2.77	1.29
	0.04	3.13	1.35
3	0.02	12.16	1.21
	0.01	12.27	1.04
4	0.01	17.22	1.02
	0.03	16.97	1.02

3.3.4 CHN Analysis

Preliminary SP08 analysis

Following discovery of the blocked regions, four samples were systematically obtained from the post-operational reactor. Fresh stock of pre-operational catalyst was analysed for comparison. One sample was taken from both a low and high post-operational δP channel. Fine powder fragments in the high δP sample were sieved by a 106 μm mesh. This retained the intact particles and allowed a separate sample of fines to be analysed. Due to additional testing requirements, the number of samples for this experiment was limited to those listed in table 3.4

Raw results from the CHN determination for the four samples are presented in Figure 21. The results show good repeatability between the duplicates, and the average value for duplicates is carried forward for analysis.

It is possible to observe firstly that the background nitrogen is not a significant factor in this analysis. It remains roughly constant for all post-operational samples. Discounting Nitrogen, Figure 3.12 visually represents the results.

In Sample 1 - Fresh Catalyst - it is possible to observe traces of the three elements under scrutiny. Predictably, carbon is not present in the fresh catalyst in appreciable amounts. Also expected, hydrogen is present in the lowest amount of any of the 4 samples. While reduced species may be present in our fresh catalyst, we can conclude that no significant mass of hydrocarbons are present.

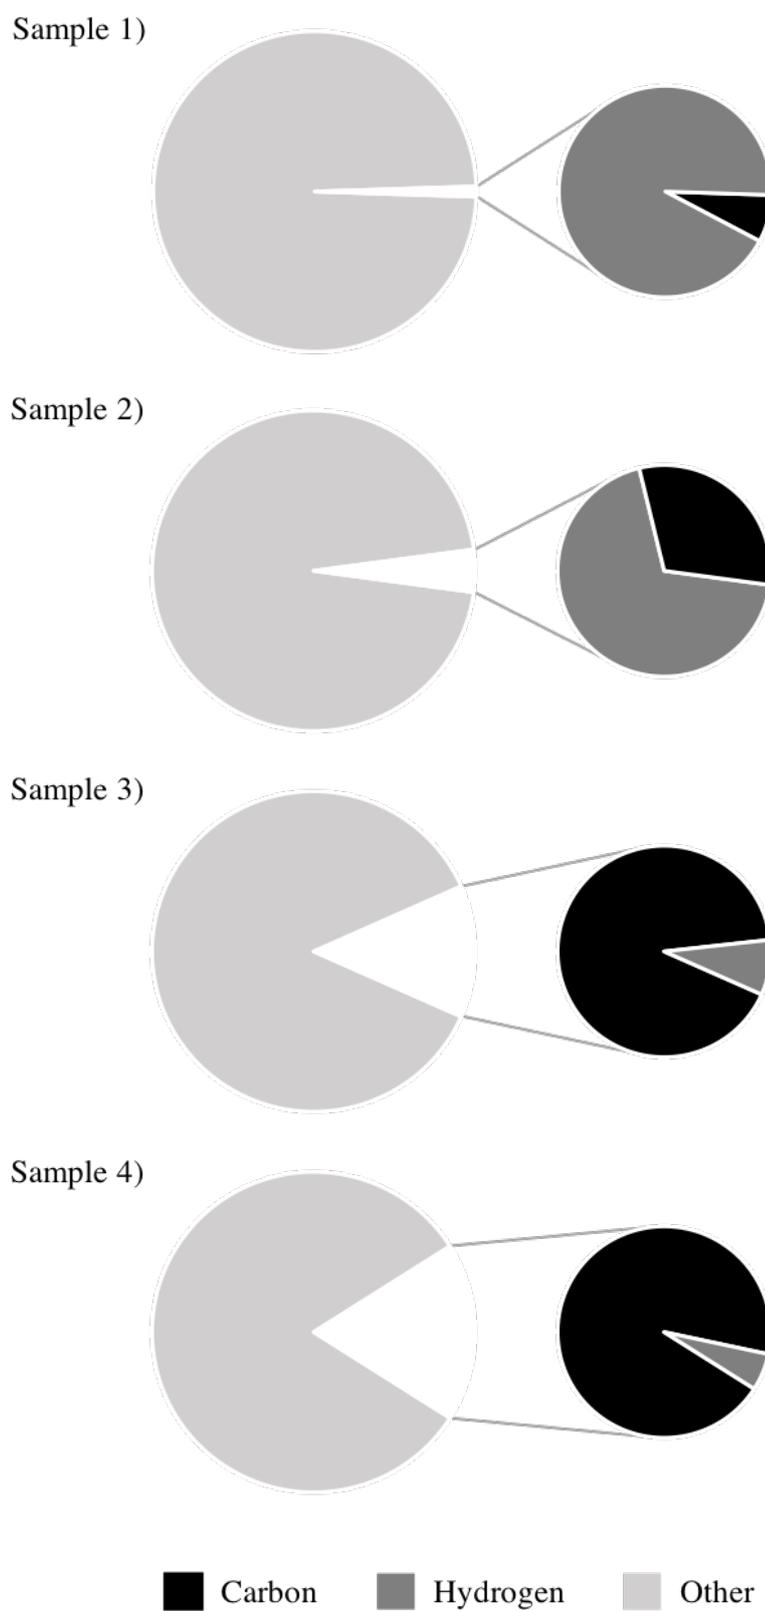


Figure 3.12: Pie charts to show proportion of carbon and hydrogen (wt.%) present in each sample (average from duplicates)

Table 3.6: molar percentage of carbon and hydrogen in contamination and calculated C:H mol ratio. Values above ~ 0.48 are indicative of the presence of elemental carbon

Sample	C (mol.%)	H (mol.%)	$C : H_{mol}$
1	0.65	99.35	0.01
2	15.7	84.3	0.19
3	47.5	52.5	0.90
4	58.28	41.72	1.40

In Sample 2 - vacuumed catalyst extracted easily from the outlet side of a low δP channel - we observe firstly that the background nitrogen is not a significant factor in this analysis. It remains roughly constant for all post-operational samples. Secondly, we can notice a significant increase in the amount of carbon and hydrogen in relation to the fresh catalyst. This can quickly and rationally be explained by the buildup of hydrocarbons in the macrostructure of the particles. However, closer inspection of the results may indicate some subtleties worth extracting.

The Carbon:Hydrogen mol% ratio may give an indication of potential sources of carbon contamination. This is calculated by converting the wt% to mol% for each species and dividing C/H. For sample 2, the C/H ratio is low. A crude assumption for a purely hydrocarbon based origin would be that the mol% of hydrogen would be $2_n + 2$ of that of carbon (the general formula for an alkane being C_nH_{2n+2}). Therefore a C:H mol ratio of ≤ 0.48 indicates that hydrogen is in molar excess compared to carbon. It is possible and plausible that most of the carbon present is in the form of hydrocarbons. However, beyond a C:H mol ratio of ≥ 0.48 , there is a molar excess of carbon compared to a purely hydrocarbon based source of contamination; it is likely that carbon exists as both hydrocarbon and elemental carbon, supporting the theory that carbon laydown is a source of deactivation. Table 3.6 shows the calculated molar ratio of carbon to hydrogen for each of the four SP08 samples.

In the case of sample 1 - Fresh catalyst - there exists a significant excess of hydrogen compared to carbon. As expected for fresh catalyst, virtually no contamination is present (1% total). Possible sources of apparent contamination are the equilibrium concentrations of moisture and other gasses present in the air around the porous structure of the sample. These traces amount to an insignificant background.

In sample 2, both carbon and hydrogen concentrations increase. Total C+H contamination is increased to approximately 4%. However, levels of contamination do not increase in a purely hydrocarbon-like ratio. Hydrogen is still the dominant of the two species. This indicates that the samples extracted from the downstream end of a low δP channel are effectively de-waxed and reduced during the shutdown procedure. Any wax holdup expected during operation has been removed. This correlates with the low δP - gas will preferentially flow down channels with low resistance to flow due to wax holdup. The

Table 3.7: Triplicate subsampled CHN analysis

Subsample	C (wt.%)	H (wt.%)
1	2.12	0.84
2	2.10	0.89
3	2.16	0.91
σ/μ	1.4 %	4.1 %

de-waxing procedure may be more effective in these channels, compounding the pressure effect.

In sample 3 - large particles from a high δP channel - there now exists much more carbon as a ratio to hydrogen. Compared to the ideal alkane ratio (0.48 C:H), there is an excess of carbon (0.9 C:H). This indicates that while hydrocarbon buildup may be a partial contributor, it is also probable that there exists some elemental carbon. This sample was extracted from the inlet of the reactor at the site of the blockage in one of the high δP channels. It gives a good indication that the nature of the blockage is, at least in part, elemental carbon. Absence of elemental carbon in the low δP channels, with reduced blockage, supports this.

Sample 4 - fine particles from the site of a blockage in a high δP channel - show a further increased wt.% carbon and the highest C:H mol ratio of all four samples (1.4 C:H) . The fine particles are therefore concluded to be fragments of catalytic material with a high buildup of predominantly elemental carbon. The fragments are not purely carbonaceous, since the mass% does not support this. As mentioned later in the report, the EDX analysis also confirms the presence of catalytic material even in the fine particles. This indicates that blocked regions are undergoing a physical and chemical change. Both elemental carbon laydown and the fragmentation of catalyst particles are occurring in blocked regions.

SP09 CHN Results

For SP09, Figure 3.2 denotes the location and number of samples taken (31 in total). Each of these samples were analysed for CHN composition. As with SP08, both blocked and free flowing regions were analysed. 106 μm mesh was also used to separate fine material from bulk where present. The primary advantage of the SP09 samples is that the location of the samples was more carefully recorded, allowing for graphical comparison.

Since the sample size was in the order of grams, and the amount required for CHN analysis was only in the order of micrograms, which amounts to only tens of individual particles, it was necessary to subsample the previously collected samples. In order to assess how representative the subsample would be, three subsamples from the same well-mixed bulk sample were first analysed.

Table 3.7 shows that the relative standard deviation (σ/μ) is relatively low ($< 5\%$) for both hydrogen and carbon composition. The subsampling of well-mixed bulk samples provides a reliable figure for determining bulk carbon and hydrogen composition. At this point it is useful to note that the nitrogen composition for every sample in the experiment was reported as being below the instrument detection limit. Nitrogen composition is disregarded in further analysis.

Next, with confidence in the method, results can be analysed by relating the carbon and hydrogen (CH) concentrations to location within the reactor. Considering with the average column-wise profile of an edge column (A), Weighted averages are used to account for the sample mass difference in the calculation. Taken on average, carbon is more prevalent in central rows- significantly in row 2, and less so in row 3. Row 4 was completely unblocked and it is unsurprising that CH concentrations are lowest there.

However, the results of unloading experiments show that Row 1 had a substantial blockage. This somewhat contradictory result seems to rule out the correlation between carbon and blockage. An attempt to characterise the blockage was made in order to assess this. With the post run reactor now partially unloaded of free flowing material, several blocked portions still remain. Assessing the relative contribution to δP of the blocked-only portions of the reactor compared to the whole bed has two advantages. Firstly, comparing the contribution to δP of the blockage and the free flowing bed will demonstrate conclusively that the blockage is the zone is the largest contributor to δP , secondly the blockage length may be used to plot $\delta P/L$ for the blocked regions.

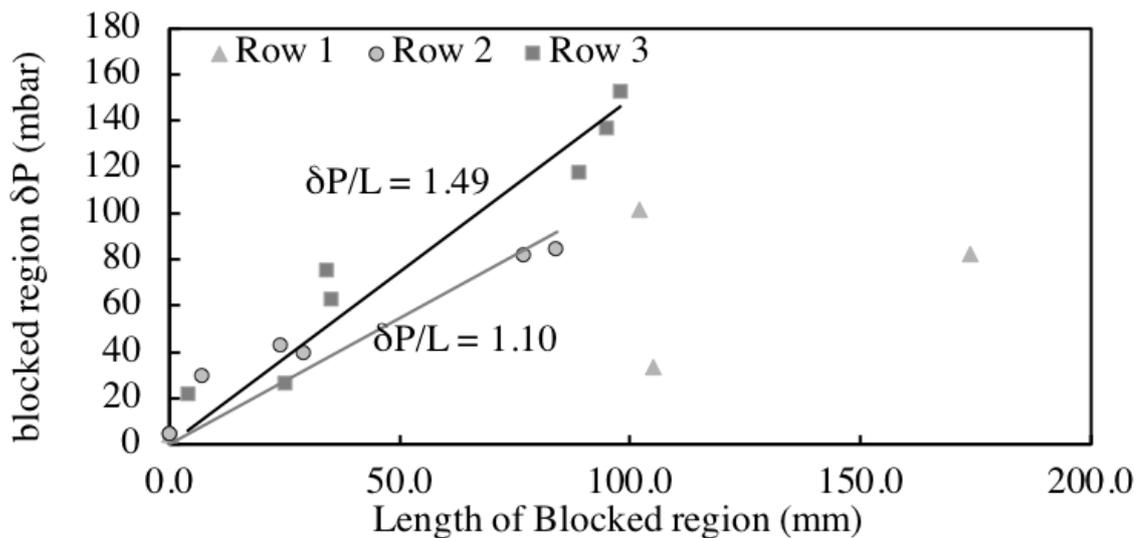


Figure 3.13: Scatter graph to show δP of blockages vs length of blockage, Grouped by rows. Linear trend line for rows 2 and 3 show the gradient, $\delta P/L$

Figure 3.13 demonstrates clearly that, as a rule, longer blockages offer more resistance to flow. This is shown by the positive correlation between δP and length of blockage. Furthermore, the row-wise groupings show that the characteristic resistance, given by

$\delta P/L$ is different for each row.

Row 1, shown to have large blockages by depth data alone, is revealed to have a wide distribution. Although the general trend of increased δP for longer blockages is apparent, no constant δP per unit length is attributable to the row as a whole. It is fair to say, however, that individual $\delta P/L$ values for row 1 fall below 1 mbar/mm.

In contrast, rows 2 and 3 show a much more regular pattern of $\delta P/L$. The gradients shown in Figure 3.13 are a good fit to the individual values. ($R^2 > 0.9$).

Using the gradient $\delta P/L$ as a coefficient to characterise the resistance to flow given by the blocked region shows that blockage *length* alone is not necessarily indicative of severity of blockage. The $\delta P/L$ for blockages correlates to the row-wise profile of carbon concentration observed previously. This affirms the relationship between carbon concentration and blockage severity.

Following on from the implications made by Figure 3.11, namely the relationship between runaways and blockage, a view of the axial CH profile might advance the theory that carbon is more prevalent in blocked or runaway zones of the reactor.

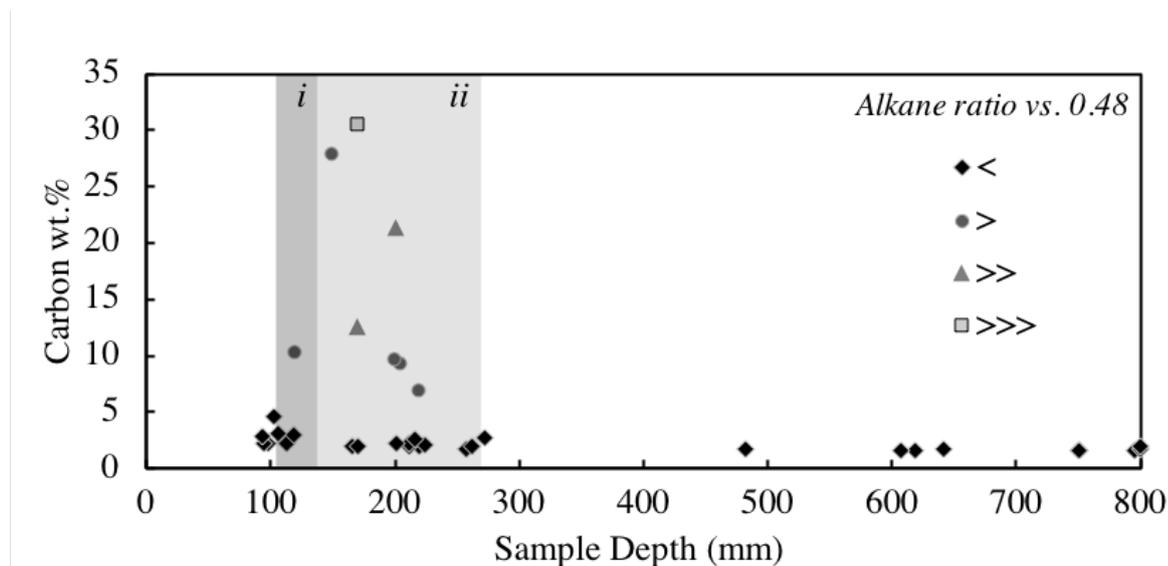


Figure 3.14: Scatterplot to show Carbon wt.% vs. depth from inlet in SP09. Groupings show alkane ratios less than and greater than the alkane ratio (0.48). Superimposed regions show (i) average runaway depth, and (ii) average blocked region

Figure 3.14 combines carbon wt.% data for all samples in the SP09 experiment plotted versus the depth at which they were obtained. One caveat to the depth information obtained is that the depth range at which a sample is taken is summarised by a single number; the final depth reached for the given sample. For instance, a sample taken from the depth range of 160-170 mm from the inlet will be denoted by a depth of 170 mm. From this, it is possible to observe the relative frequency at which samples are obtained along the depth seems to be higher in the region from 100 – 300 mm from the inlet. This is related to the fact that samples in this section were far less free-flowing and only small

amounts of sample were freed at a time. Between 300 – 800 mm from the inlet, it is clear that data points are less frequent. This is because free flowing sample was obtained and large fractions of the bed were removed at the same time. One data point serves as an average for a larger depth range. The extreme example is the data point at 800 mm, extracted as one free flowing sample from inlet to outlet, this data point serves as an average for the whole channel. This is closely matched, however, by the data point at 750 mm, which represents the first 50 mm from the outlet of every channel, extracted simultaneously and combined in a well mixed sample. Thus, axial and radial averages of free flowing material are shown to be similar in composition. Free flowing samples form a baseline of carbon of approximately 2 wt.%.

Figure 3.14 clearly shows the coincidence of a sharp peak in carbon composition corresponding to the zones of average runaway and blockage (shown in the shaded regions). That is to say that carbon is significantly more prevalent in blocked and runaway zones of the reactor. Furthermore, analysis of the groupings in Figure 27 show that C:H mol ratios above the idealised hydrocarbon ratio are exclusively present in the blocked and runaway zones, indicating the likely presence of elemental carbon in these zones. While this does not preclude the existence of elemental carbon in non-blocked regions, or indeed the presence of hydrocarbons in blocked regions, it suggests a significant shift in the relative amounts.

Though limited in scale, the SP08 experiments also seem to indicate the existence of carbon rich blocked regions. One key finding of the SP08 results were the presence of carbon rich fines in the blocked zone. Again, looking at SP09 data, a number of fines were observed to be present. This time, since careful vacuum suction of the channel was used to obtain the samples, there is less doubt as to whether the mechanical stresses of the sampling procedure were responsible for generating fines. Any fines present in the vacuumed sample are likely to have originated in operation of the reactor.

The relative carbon composition of samples and the fines contained within them are shown in Figure 3.15. For comparison, one free flowing outlet channel is compared with two samples taken from both the inlet and outlet side of a blocked zone. The free flowing outlet sample, taken from the region immediately downstream of the blockage in that channel is shown to have carbon rich fines. The larger particles from the same location contain carbon at approximately the baseline concentration discussed previously. By contrast, the large particles in the two blocked region samples are seen to have much higher carbon concentration than the baseline amount. The sieved fines are shown to have a still higher carbon concentration.

This seems to indicate that one of the following is occurring in blocked regions of the reactor: Either carbon is ‘dusting’ the surface of particles with a layer of fine carbon, which is loosely bound to the surface. This carbon coating could be enough to bind intact catalyst particles together to form a blockage. Subsequent disruption by extraction and

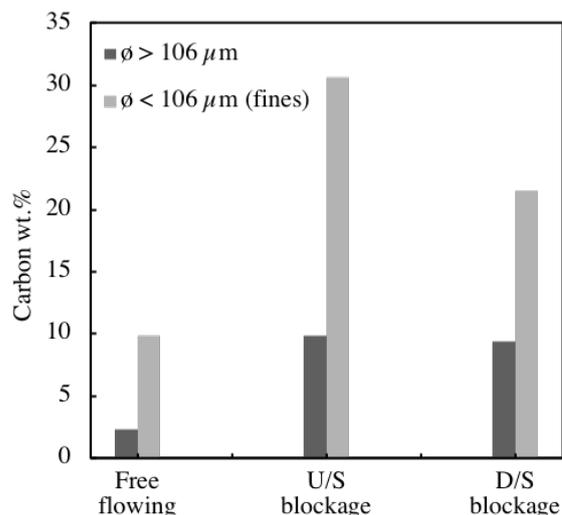


Figure 3.15: Column graph to show comparison of carbon composition between samples sieved above and below $106 \mu\text{m}$

sieving is possibly dislodging this carbon dust along with a significant quantity of other material ($>70\%$ of the fine material is not carbon). – It has been noted previously that fresh catalyst also has a tendency to produce fines as fragile portions of the cobalt-alumina surface are fragmented by abrasion. Else, blockages are formed by a parallel mechanism (e.g. sintering). The carbon laydown, fragmentation of the bed and blockages are related by cause but are not the same.

Further elemental analysis of the fines is required in order to establish the composition of the non-CHN component. This was not possible in the present study owing to the stringent intellectual property agreements in place. Establishing the form in which the carbon exists is also important in identifying the mechanism by which it was made. The latter may be answerable by looking at the operational history of the reactor. It is known that the reactor underwent several thermal runaways. The instantaneous local maximum temperature achieved is difficult to quantify due to the low spatial resolution of thermocouples, though it is potentially very high.

Moodley (2008) collates the results of several studies into the formation of carbon species, shown in figure 3.16.

The techniques used in this report are unable to identify carbon species directly – the scope for further analysis in this direction is clear. However, it will be noted that the de-waxing procedure for SP09 at the end of the run involved the passing of hydrogen at high temperatures in excess of $350\text{ }^\circ\text{C}$ for an extended period of time. This reducing environment is likely to have hydrogenated many of the possible forms of carbon present. The Moodley (2008) table of hydrogenation temperatures collated from several studies, referenced therein, is a useful guide. Based on these findings it is possible to predict that any remaining carbon in a post-operational reactor that has undergone the de-

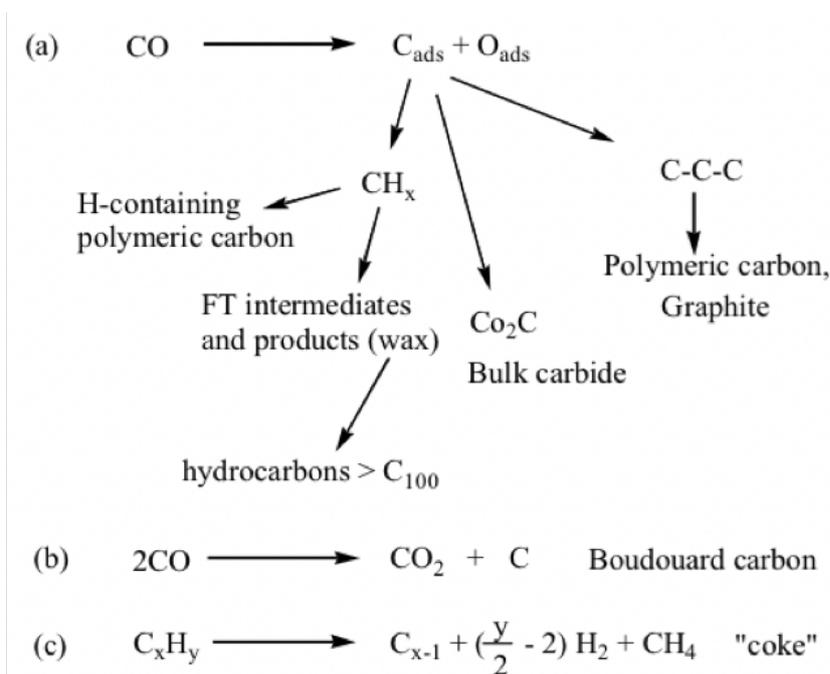


Figure 3.16: Possible modes of carbon formation during FTS on cobalt catalysts

Table 3.8: Examples of various carbon species on cobalt FTS catalysts along with their hydrogenation temperatures. (Moodley, 2008)

Carbon Species	$T_{\text{hyd}}(^{\circ}\text{C})$
CH_x fragments	<100
Surface Carbide	180-200
Bulk Carbide	<250
Paraffinic wax	250-350
Polymeric carbon	>350
Graphite/graphene	>620

waxing procedure is likely to contain only polymeric or graphitic carbon in addition to residual heavy hydrocarbons not removed by the de-waxing procedure. It is theorised that all components with a hydrogenation temperature lower than the de-waxing treatment temperature are all effectively reduced and not found in post-operational catalyst.

This theoretical hydrogenation of some carbon species does not preclude the prior formation of them during operation. Indeed, the formation and subsequent hydrogenation of said species could be responsible for the fragmentation and blockage of the bed observed in post operational catalyst. Further, more complex analysis of the forms and specific location of carbon are required in order to advance any theories as to the mechanism by which blockages are formed in carbon rich regions of the bed.

3.3.5 SEM analysis

Images taken at a variety of magnifications are presented in the following figures. Direct comparison or the physical appearance is useful to gain an understanding of the surface. Superimposition of the EDX backscattered elemental maps allows the composition of the samples to be qualitatively determined.

Fresh catalyst, shown in Figure 3.17A shows a relatively smooth surface appearance. At this magnification, macropores and smaller surface features are not visible. As a notable aside, catalyst is assumed to be approximately spherical. However, visual inspection of the particles shown in Figure 31 indicates that this assumption does not necessarily hold. Offsets between calculated and empirical values for δP derived from the Ergun-Fourmeny equations may arise from the observed lack of sphericity in addition to the particle size distribution and irregular channel geometry.

In comparison, figure 3.17B shows post operational catalyst from a high δP zone, which displays a different appearance at 65x magnification. It is possible to observe an increased particle size distribution. (Note that particles less than 106 μm have been removed by sieving) Some large particles remain; however, the majority of the sample is composed of smaller, more angular fragments. The degradation of the fixed beds into small fragments could lead to an increase in the δP by reducing voidage, however the bed was not observed to collapse and settle to a new level, indicating that the bed was held suspended. Figure 3.19A shows the sieved fines. Though sieved with 106 μm mesh, it is clear that the vast majority of the particles are much smaller ($<10 \mu\text{m}$) indicating that the particle distribution of the combined sample (sample 3 + sample 4) would be approximately binomial rather than normal as is the case for fresh catalyst.

Next, magnification was increased in order to observe the macrostructure of the particles. In each of the following images is shown the surface on one individual particle, typical of the sample, at 2500x magnification.

In Figure 3.18A, an open macroporous structure is clearly visible in the fresh catalyst. The surface appears to be relatively smooth. Surface pores appear to be in the order of 1-3 μm in size. By contrast, Figure 3.18B shows the much rougher surface of post-operational catalyst. There are many small granular particles of a range of sizes (1-5 μm), which rest on the larger surface of the intact particle. These small particles seem to block the macroporous structure of the larger. Figure 3.19B shows a close up of a fine particle. It is suspected that the small blocking fines shown to exist in sample 3 are indeed of the same nature as the particle in Figure 3.19B.

With a diameter of approximately 3 μm , the particle in Figure 3.19B is of the same size range as the macropores observed in the fresh catalyst and also of the same size range as the fine particles which appear to be blocking the pores of sample 3. No macroporous

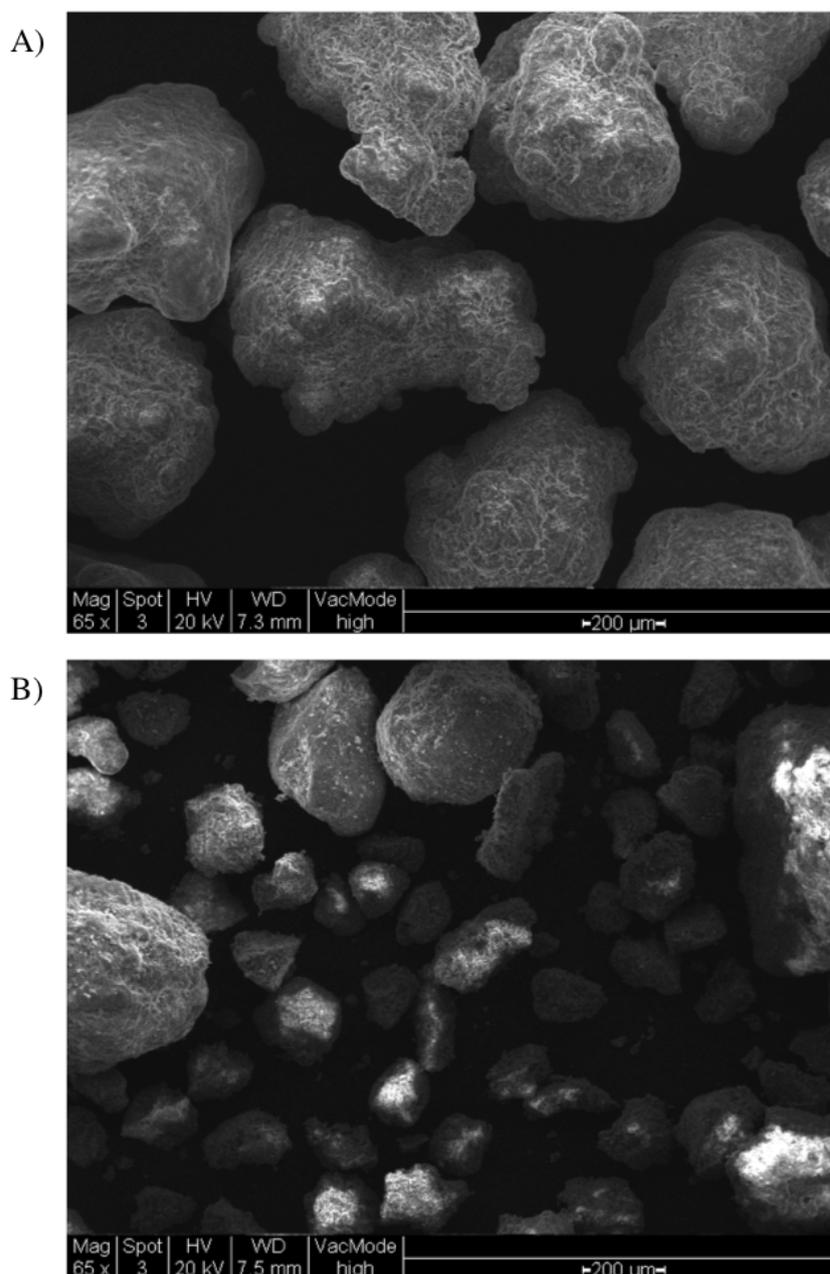


Figure 3.17: Electron micrograph (65 X magnification) of A) pre-operational catalyst. B) post-operational high δ P catalyst (sieved to remove fines)

structure is visible on the particle. CHN analysis reveals these fine particles to be highly contaminated with carbon. However, the composition was still made up of over 75% non-CHN species. This implies that the particle was not entirely formed in operation, but is a carbon rich particle of the catalyst and support.

The surface of a fine particle was viewed at 50,000 X magnification. The EDX map in Figure 3.20 shows the distribution of cobalt on the surface. The uniform presence of cobalt indicates that the fine particles are fragments of catalytic material, which have broken from the bulk support.

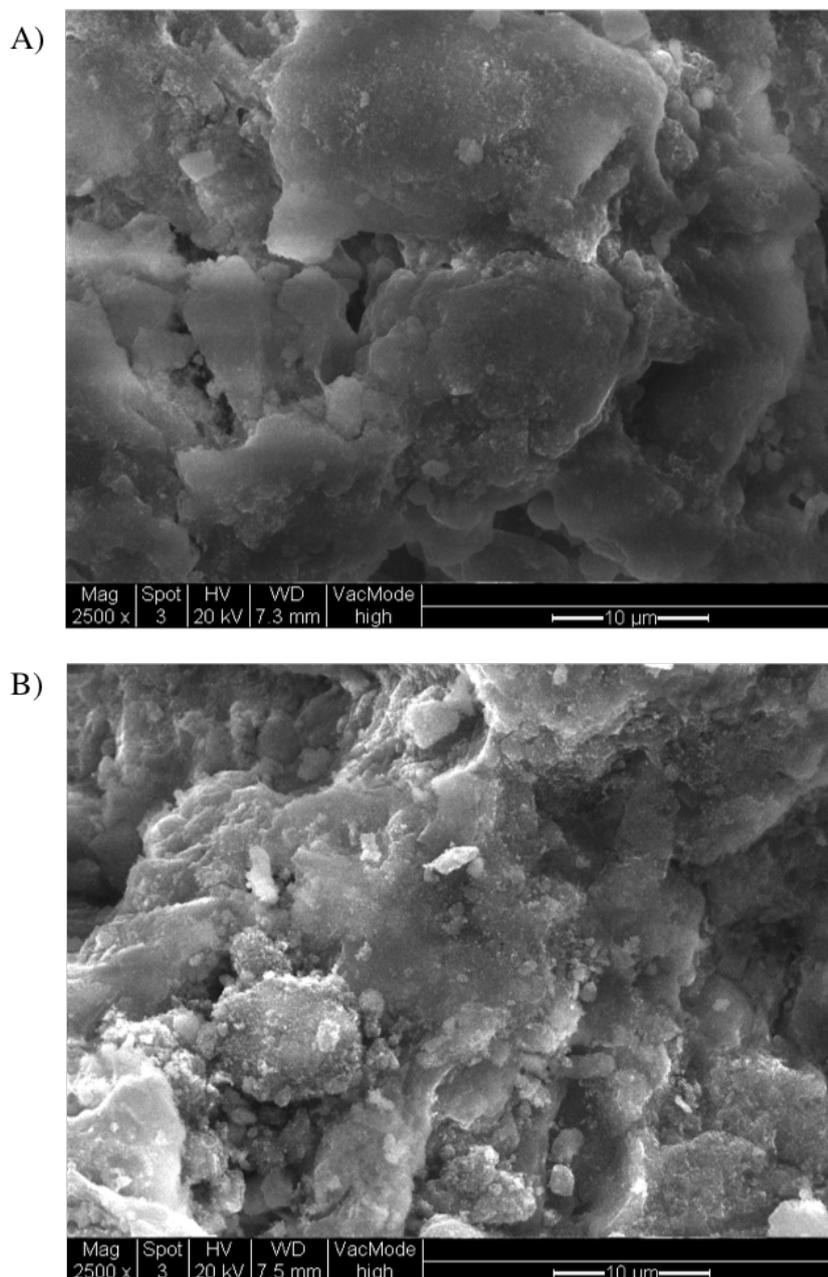


Figure 3.18: Electron micrograph (2500 X magnification) of A) pre-operational catalyst. B) post-operational catalyst.

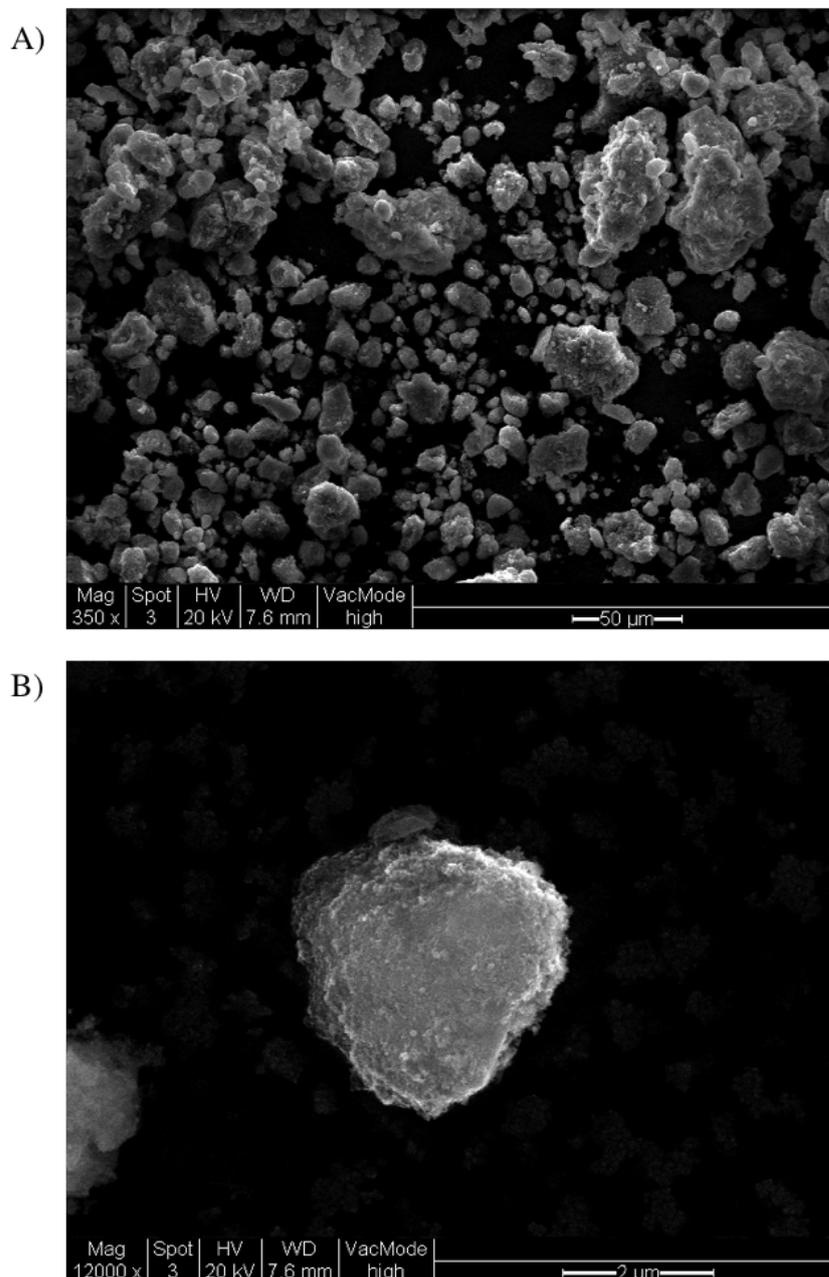


Figure 3.19: Electron micrographs of A) sieved post-operational catalyst fines (350 X magnification). B) individual sieved post-operational catalyst fine particle (12000 X magnification).

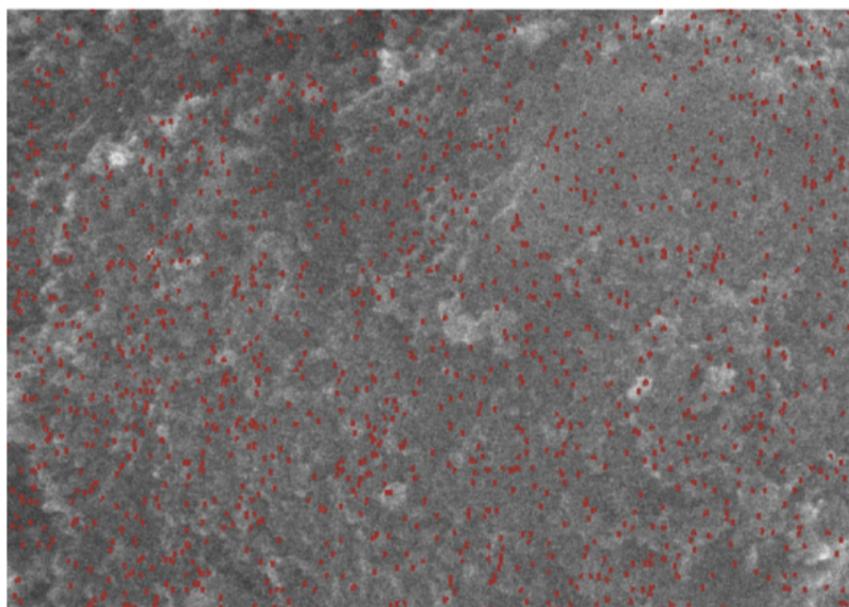


Figure 3.20: Electron micrograph (50,000 X magnification) of individual post-operational catalyst particle superimposed with EDX elemental map for cobalt.

3.3.6 In situ Endoscopy

During unloading and extraction of the post-run reactor, *in situ* endoscopic photography of the catalyst was undertaken. Initially to assist in the unloading process, endoscope images provided intriguing insights into the appearance of the bed. Previously reported results on blockage, particle fragmentation and the presence of carbon are observed in the images discussed in the following section.

Firstly, it is useful to note the appearance of free flowing catalyst. Figure 3.21 below shows a typical example. Some non-spherical eccentricities can be observed, supporting the observations of the SEM analysis discussed previously. In general, the bed exhibits a visible voidage and uniform particle colouration

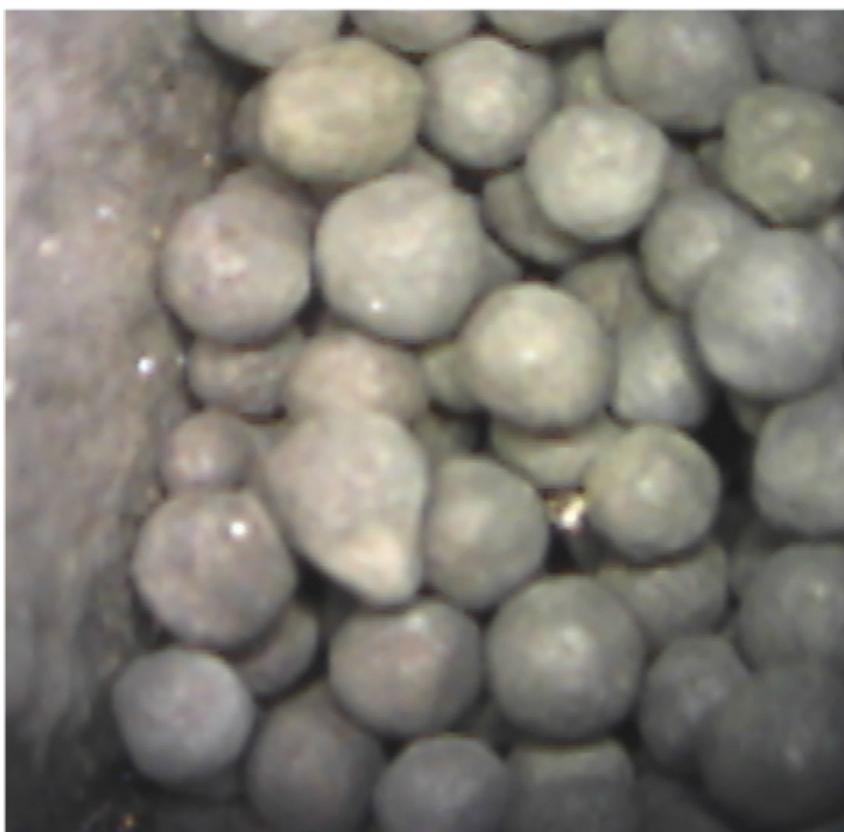


Figure 3.21: *in situ* free flowing catalyst (650 mm from inlet, Channel A1, C:H mol 0.22)

Next, two phenomena were observed in several channels during the unloading procedure, dubbed the following:

Radial blackening

Found throughout blocked regions of the channel, radial patterns of black material were found on the surface of the particles to form a ring of non-blackened particles adjacent to the wall of the channel. The effect was stronger in the corners.

Contact shielding

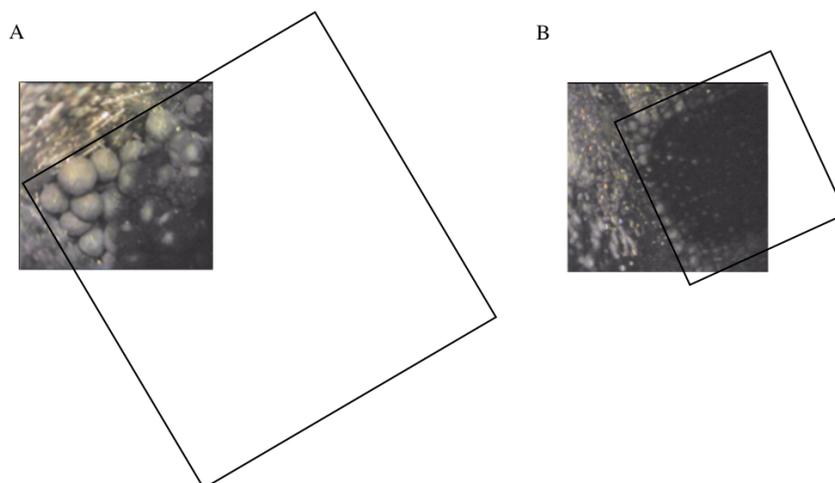


Figure 3.22: *in situ* photographs of radial blackening effect with superimposed channel dimensions and orientation. A) close-up view of the corner effect. B) view of majority of channel.

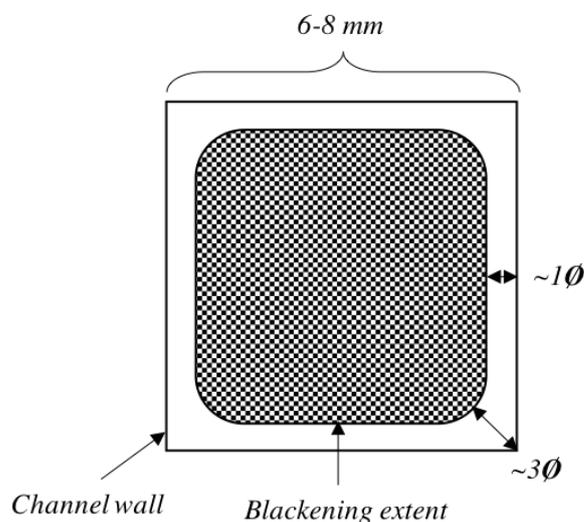


Figure 3.23: Schematic cross section of a single reactor channel showing the typical nature of the radial blackening effect observed by *in situ* endoscopy.

Observed on the surface of blackened particles, contact- shielding manifests as points with lighter colouration. Ostensibly, these points are areas of the surface which are in direct contact with a neighbouring particle. Since particles are closely packed, the result is a polka-dot appearance to particles.

Possible mechanisms

Radial blackening of the blocked regions is thought to be due to carbon laydown. This appears to be formed in the gas phase in the bed voidage and is laid down as a coating around particles. The radial effect may be explained by several potential mechanisms:

Perhaps the radial effect is a product of the flow distribution through the channel. Wall and corner regions within a channel do not receive the same gas flow and are therefore less susceptible to the gas phase laydown of contaminants.

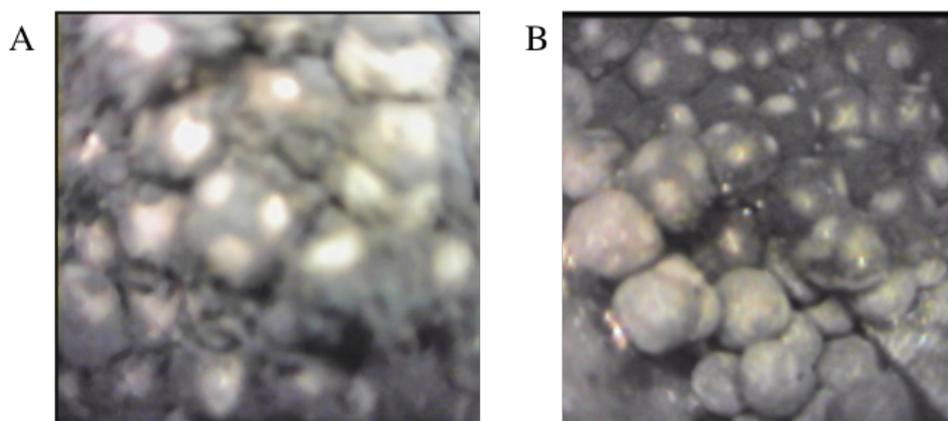


Figure 3.24: *in situ* photographs of 'Contact shielding' effect showing A) Close-up view of particles. B) wider view of particles showing both contact shielding and radial blackening.

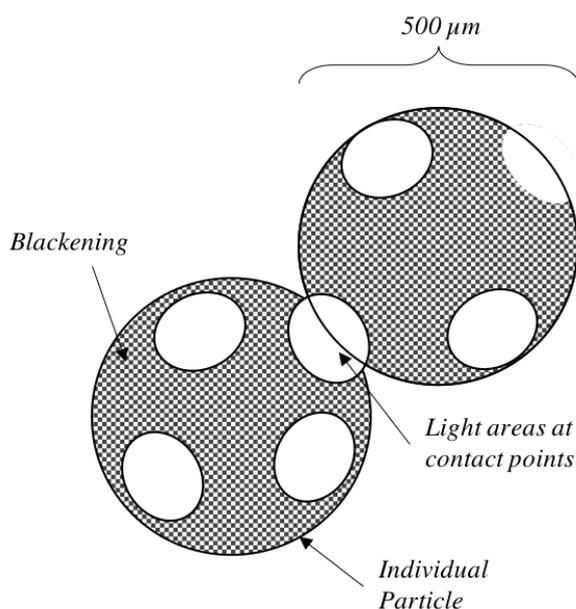


Figure 3.25: Schematic diagram of two catalyst particles, illustrating the *contacting shielding* effect.

The wall effect may be a profound influence; particles adjacent to a wall are more centrally aligned to each other, and form a regular packing structure aligned to the wall. Particles on the wall are effectively shielded by their upstream neighbour while in turn shielding their downstream neighbour. Contamination formed upstream is less likely to path towards shielded regions close to the wall. This is exacerbated by the low tube to particle diameter ratio, d_r , discussed previously.

Thirdly, proximity to the wall might afford those non-blackened particles the ability to more readily expend heat. If the contamination formation is a heat-driven process, effective heat transfer away from the particle could allow wall particles to remain below the heat of formation threshold of the contamination.

Contact shielding seems to be a much simpler effect- contamination formed in the gas phase is not being deposited where it cannot access. Particle point-to-point contacts are

shielding the surface from deposition of contamination. However, this is also compatible with a heat-driven mechanism of contamination formation. Particles experience a local heightened conductive heat transfer capability at particle point to point contacts. The local ability to dissipate heat in these regions may have a similar effect to the one proposed for particle-wall contact mentioned above.

A final observation of the *in situ* endoscopy is the appearance of the blocked regions within the channel. Notably, as the blocked region was exposed, the radial blackening effect gave way to a much more severe central amorphous zone within the channel. In this region, no individual particles were visible. Instead, agglomerated blocks of bed were present. Figure 3.26 shows *in situ* and *ex situ* examples of this effect, in contrast to the free flowing particles shown in Figure 3.21.

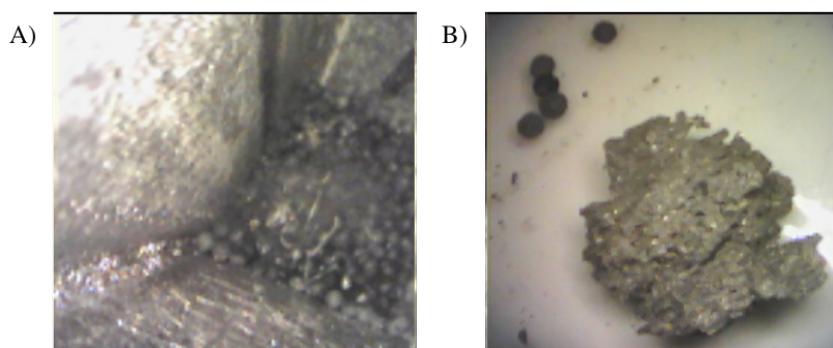


Figure 3.26: Photographs of blocked regions of a channel. A) *in situ* endoscopic image of blocked region. B) *ex situ* image of amorphous agglomerate next to intact particles for size comparison.

The agglomerate in Figure 3.26B was shown to have the highest wt.% Carbon (28.03 %) compared to all other samples, excluding the sieved fines of a blocked region which showed slightly more.

3.4 Chapter Summary

This chapter began by observing a subtle row-wise discrepancy in the construction of the modular FT reactor, SP08. Divider and cooling plates are stacked in an alternating sequence, giving rise to row-wise homogeneity of heat transfer performance. Upon inspection of the post-operational differential pressure, this was also found to observe a distinct row-wise pattern. During ‘post-mortem’ investigations of the reactor, this pattern of high pressure drop was found to correspond to the operational history of thermal events, and the presence of blocked regions of amorphous catalyst bed. Further examination found that these blocked regions were rich in elemental carbon thought to be laid down during thermal excursions during the run. The severity of blockages was quantified by normalising the local ΔP over the length of the blocked region. Two carbon laydown phenomena were explored, implying that the carbon was highly likely to originate in the fluid phase, since it follows patterns which correspond to fluid flow. An exploration of calculating and modelling pressure drop proved difficult, owing to the many flawed assumptions and simplifications in the literature. In future work, a method of capturing the *actual* 3D structure of the bed and modelling more directly will be explored.

Chapter 4. Characterisation of Next-generation Compact Fischer-Tropsch-Tropsch Packed Beds Using X-Ray Tomography

4.1 Chapter Abstract

Addressing previous limitations identified in this work regarding the applicability of the standard pressure drop models for packed beds, this chapter explores the possibility of characterising packed beds using 3D imaging. By means of X-Ray tomography of a mock reactor, packed with two different commercial catalysts, and by application of various image processing techniques, particle size distribution, voidage, tortuosity and other important fixed bed parameters can be captured not only as nominal values, but at a local level. The combined dataset can be reconstructed into *in silico* 3D models which may be explored further using computational simulation.

4.2 Materials and Methods

In order to quantify the packing characteristics of real catalyst particles in real process geometries, it was necessary to perform some form of imaging technique and subsequent data analysis. Rotating stage X-ray tomography is commonly employed in imaging complex structures. The technique allows for 3D models to be reconstructed and various bed characteristics to be accurately quantified.

X-Ray tomographic images of 2nd generation reactor channels (nominally 10 × 3 mm) was carried out for two FTS catalysts, varying principally in nominal particle diameter (300 μm & 600 μm). The differing particles represent two grades of catalyst which are candidates for commercial production. The smaller particles are specified to be more spherical and have a more precise particle size (narrower particle size distribution).

4.2.1 Mock Reactor

For this analysis, a 'mock', three-channel reactor was fabricated to the same dimensions as the pilot scale reactors considered in this case study, and therefore at the same channel

dimensions as a commercial modular reactor of the same design.

The mock reactor was embedded in Perspex[®] clear acrylic casing for the analysis (see Figure 4.1). Although the acrylic mounting is relatively transparent to high energy X-Rays, the irregular shape of the mounting from the perspective of the X-Ray source meant that a baseline disturbance value from the mounting would always be observed, but would vary with the angle of rotation. Therefore, a solid block of Perspex[®] was first used as a calibration standard. Further measurements of the mounted channel could be processed relative to this baseline.

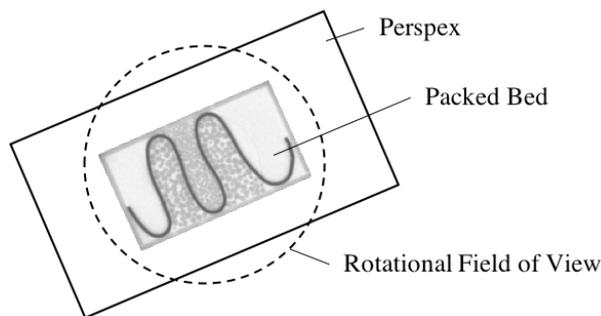


Figure 4.1: Schematic plan view of Perspex[®] block containing mock packed bed. Rotation around the indicated circle results in variable apparent thickness of perspex

As can be seen, entire sample assembly would not fit into the 16 mm field of view (FOV) due to the overhanging material. To overcome this, the tomographic scan was split into four overlapping scans, effectively increasing the FOV horizontally to 20 mm while also allowing for better resolution (5.8 μm in this case). Vertically, three measurements were made with a separation of 4 mm (overlapping 1 mm of the 5 mm FOV) such that potentially a 12 x 20 x 20 mm volume could be reconstructed. In practice, however, it was necessary to process the image in sections as the full volume would contain 25 billion voxels (volume pixels). Represented in 32-bit form this would correspond to a 100 GB image which is impractical to process. Due to the likelihood of noise and artefacts in the the measurement, each measurement was taken in five replicates.

4.2.2 Image processing

Figure 4.2 highlights some common artefacts within the imaging data. *ImageJ* software was used to apply image processing algorithms to the raw data to provide 'clean' images for further analysis.

3D median Filtering

Each voxel is set to the median value of its neighbourhood. The neighbourhood is defined by a radius. In this case, a radius of 2 voxels in the x-, y- & z-dimension was adequate to improve image quality (Schmid, 2014).

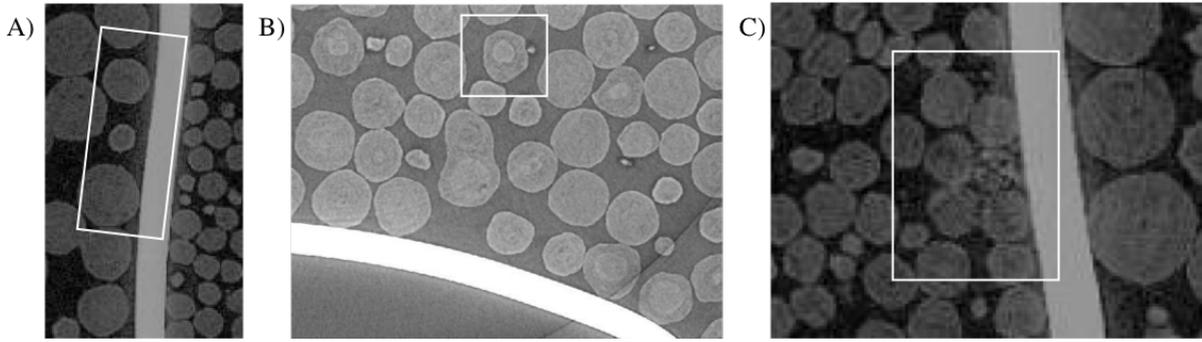


Figure 4.2: Raw X-Ray Tomography image sections showing various artefacts. A) High intensity regions near the channel walls. B) Particle inhomogeneity. C) Ring artefact due to image reconstruction.

2D Anisotropic filter

Anisotropic filters are a class of filter that reduces noise in an image while trying to preserve sharp edges.

ImageJ software contains an inbuilt implementation of the 2D Anisotropic filtering algorithm developed by Tschumperlé and Deriche (2005). Raw images are optimised by visually assessing the impact of the algorithm parameters:

Iterations - Raw images were processed in 20 iterations, with two smoothings per iteration.

a_1 (diffusion limiter along minimal variations) - a_1 influences the shape of the smoothing mask. $f_1(l_1, l_2) = (1.0 + l_1 + l_2)^{-a_1}$. The smoothing in each iteration is defined by a tensor (2x2 matrix), that is linear combination of tensors corresponding to minimal and maximal eigenvalue of structure tensor. f_1 and f_2 are weights of the tensors (Tschumperlé and Deriche, 2005). $a_1 = 0.5$.

a_2 (diffusion limiter along maximal variations) a_2 influences the shape of the smoothing mask: $f_2(l_1, l_2) = (1.0 + l_1 + l_2)^{-a_2}$. $a_2 = 0.9$.

dt (time step) - The result of the filter is proportional to the step, but too long time step yields numerical instability of the procedure. $dt = 20$.

h_{et} (edge threshold height) - Defines minimum "strength" of edges that will be preserved by the filter. $h_{et} = 5$.

The result of 3D median and 2D anisotropic filtering can be compared to the raw image in figure 4.3A&B

Thresholding & Segmenting

Next, the pre-processed images undergo a thresholding filter to classify regions into two categories; particle and void. Next, a watershed algorithm is applied to differentiate

individual particles. A size filter is applied to remove residual noise incorrectly evaluated as a small particle region. The result (Figure 4.3C) is an image map of individually identified particles which can be assessed for a number of parameters (Wall distance, diameter, roundness etc.).

In order to calculate particle parameters from the processed images, further image assessment relies on the reliability of the pre-processing steps. However, as can be seen in figure 4.3, some artefacts still remain. Most commonly, the thresholding process mis-identifies void space within a particle. this occurs both in the horizontal and vertical plane (figure 4.4) Since calculations such as the void fraction are calculated simply by calculating the area of processed images occupied by particles or void, any bias towards identifying void space caused by the thresholding algorithm will result in a slightly higher prediction of void fraction. This effect is unavoidable without causing unreasonable image distortion and is deemed to have little impact on the values for the purposes of this study.

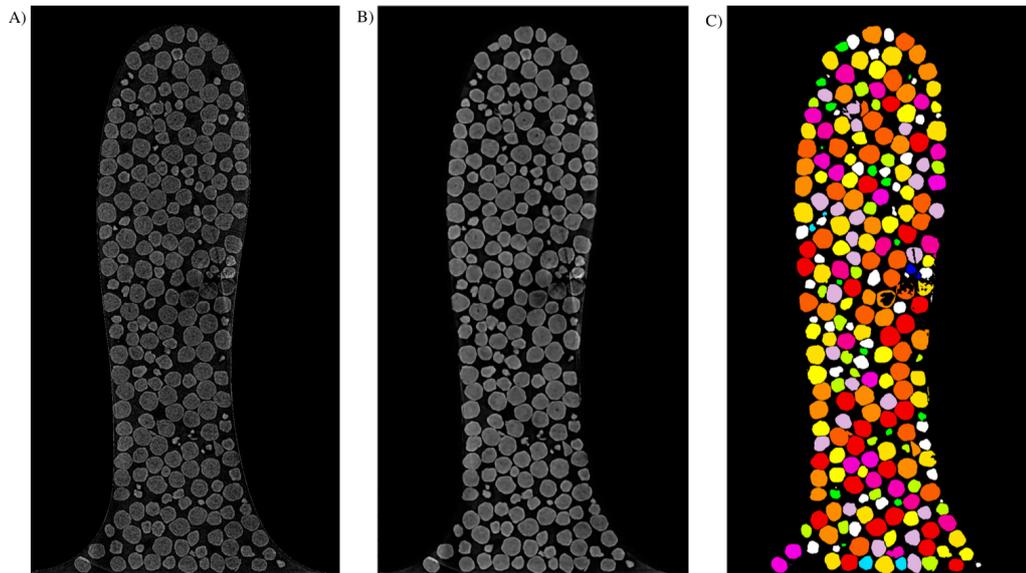


Figure 4.3: Comparison of raw and processed 2D x-ray tomography images. A) Unprocessed image of a single channel. B) Same data after 3D median Filtering and 2D Anisotropic filtering. C) Data after thresholding and segmenting. Random colour assignment identifies individual particles.

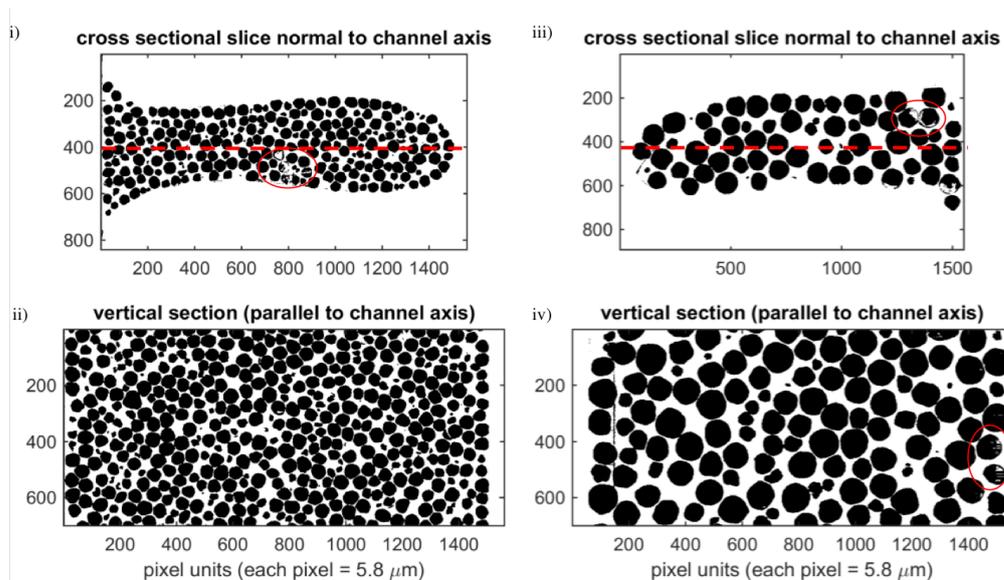


Figure 4.4: Horizontal and Vertical processed tomography planes for two particle sizes. (Dashed lines in upper images show location of vertical plane in lower images) Circles show regions where thresholding mis-identifies void space within a particle

4.3 Results & Discussion

4.3.1 Packing Parameters

Raw images were processed by the methods described earlier to allow for process relevant parameters to be calculated and reported. Table 4.1 summarises the global parameters for the packed beds.

Table 4.1: Summary tomography parameters from two mock channels loaded with FTS catalyst of differing nominal particle diameter (channel height 4.06 mm).

Property	Channel 1	Channel 2
Nominal Catalyst Diameter (μm)	600	300
Total Channel Volume (mm^3)	90.35	78.05
Total Particle Volume (mm^3)	52.81	47.61
Volume Fraction	0.58552	0.60992
Void Fraction	0.41682	0.39012
No. Particles	566	3418
Modal Particle Diameter (μm)	581	304
Modal Surface Area (μm^2)	1.16×10^6	3.03×10^5

4.3.2 Particle Geometric Properties

Characterisation of the different catalyst particle types was of interest. Figure 4.5 summarises the geometric property distribution for both particles.

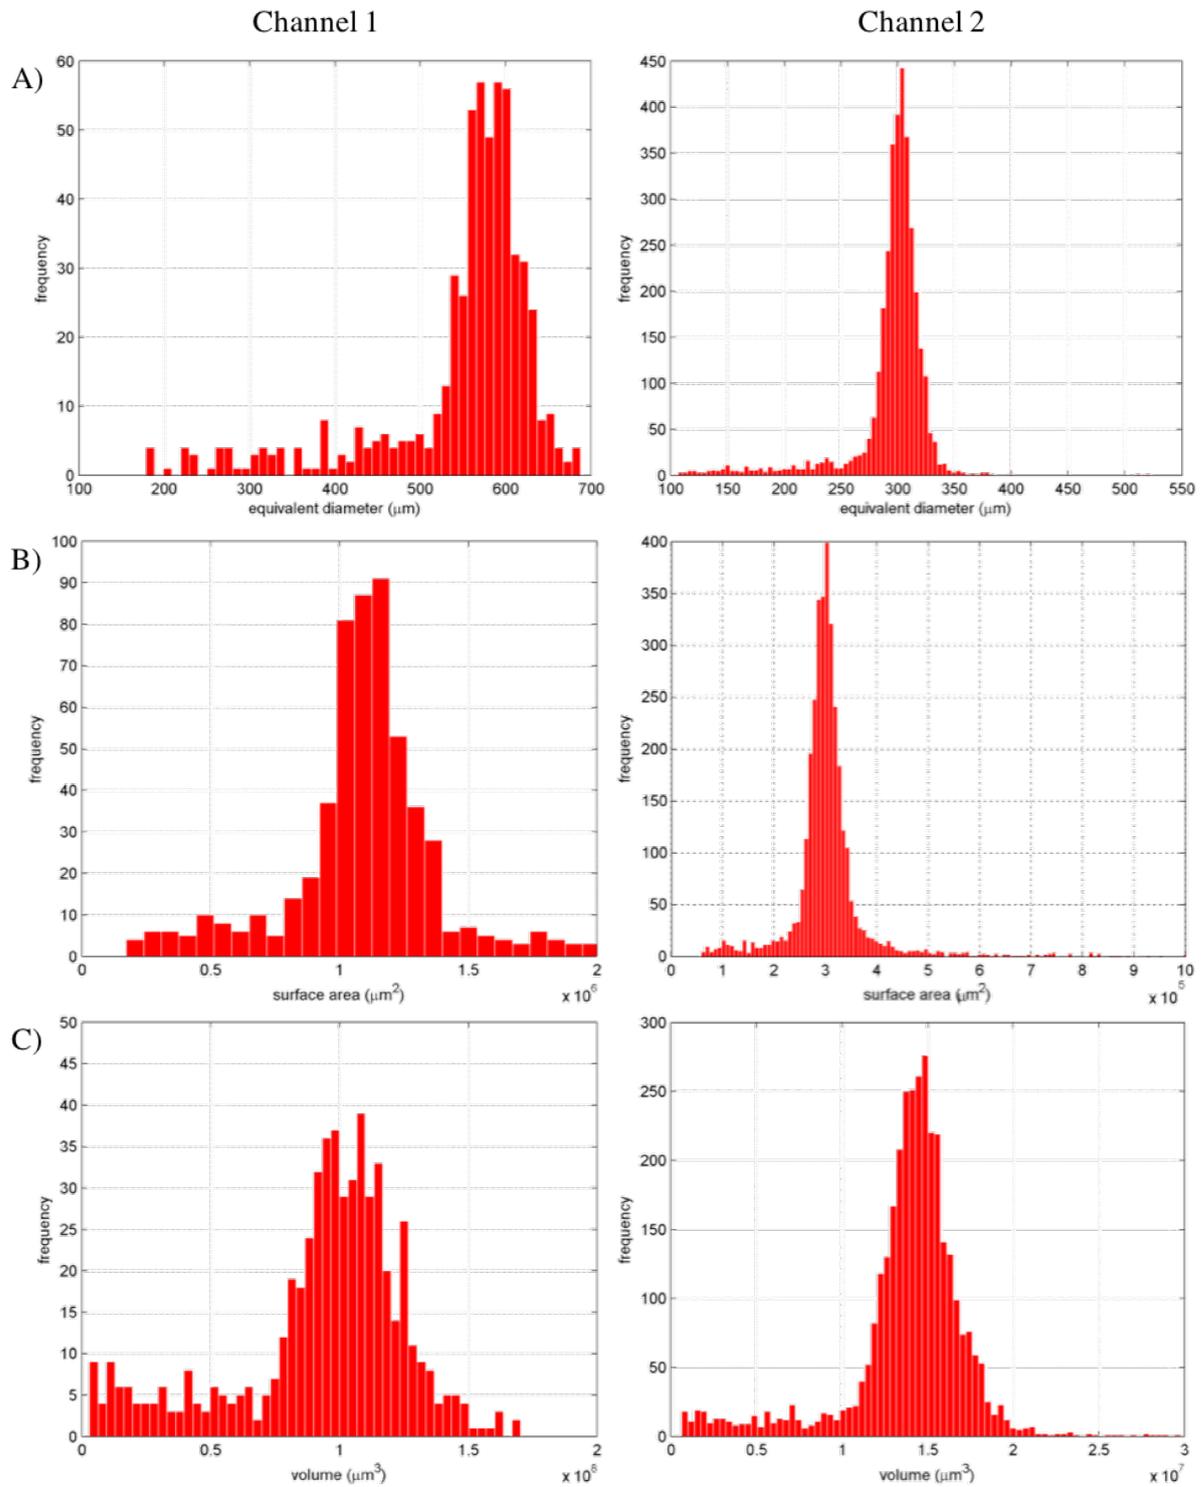


Figure 4.5: Histograms to show geometric properties of two particle types derived from tomographic image analysis.

Left) Channel 1 (600 μm particles) **Right)** Channel 2 (300 μm particles).

A) Equivalent Spherical Diameter (μm) **B)** Surface area (μm^2) **C)** Particle Volume (μm^3)

Diameter

The equivalent diameter of the particles is calculated by image analysis. (Figure 4.5A) Both channels show an approximately gaussian distribution of particle diameter, centred approximately on the nominal particle diameter specified for each. Channel 1 particles are slightly smaller on average than the nominal diameter, while Channel 2 shows slightly larger than nominal particles.

Both particle size distributions show a negative skew. There are three possible explanations for this. Perhaps the catalyst as-delivered contains a proportion of off-specification particles in the lower particle size range. Alternatively, the process of performing the experiment (through mechanical forces) the particles are fragmented into fines. Thirdly, it is possible that the image analysis is skewed by misidentified smaller particles.

Surface Area

With respect to surface area (Figure 4.5B), both types of catalyst observe a normal distribution. Clearly, and predictably the larger particles have a greater surface area. The nominal surface area is given by:

$$A = 4\pi r^2 \quad (4.1)$$

$$D = 300 \text{ }\mu\text{m}, A = 2.83 \times 10^5 \text{ }\mu\text{m}^2$$

$$D = 600 \text{ }\mu\text{m}, A = 1.13 \times 10^6 \text{ }\mu\text{m}^2$$

As can be seen by comparing these values to Figure 4.5B, calculated nominal values differ slightly from empirical values, especially for the 300 μm particles, which have a modal surface area of 3.03×10^5 .

Volume

When comparing volumes, the particles again show an approximately gaussian distribution, notwithstanding the previously mentioned negative skew.

The nominal volume is given by:

$$V = \frac{4}{3}\pi r^3 \quad (4.2)$$

$$D = 300 \text{ }\mu\text{m}, V = 1.41 \times 10^7 \text{ }\mu\text{m}^3$$

$$D = 600 \text{ }\mu\text{m}, V = 1.13 \times 10^8 \text{ }\mu\text{m}^3$$

Shape

A perfect sphere represents the minimum surface area for a given volume. All *real* shapes necessarily have a higher surface area to volume ratio.

Hakon Wadell defined sphericity as the surface area of a sphere of the same volume as the particle divided by the actual surface area of the particle.

By expressing a sphere's area in terms of the volume of a particle, as such:

$$A_s^3 = (4\pi r^2)^3 = 4^3 \pi^3 r^6 = 4\pi(4^2 \pi^2 r^6) = 4\pi \cdot 3^2 \left(\frac{4^2 \pi^2}{3^2} r^6 \right) = 36\pi \left(\frac{4\pi}{3} r^3 \right)^2 = 36\pi V_p^2 \quad (4.3)$$

One can arrive at a metric of shape such that:

$$S = \frac{A_p^3}{36\pi V_p^2} \quad (4.4)$$

Where $S = 1$ would indicate a perfect sphere and all real particles would have $S > 1$.

Figure 4.6A shows the calculated shape parameter, S , for each particle. It can be observed that the minimum shape factor lies between $1 < S < 1.1$ indicating that the particles are approximately spherical. Following the more stringent specification on sphericity of the 300 μm particles, the empirical results do show a slightly more spherical particle when compared to 600 μm particles.

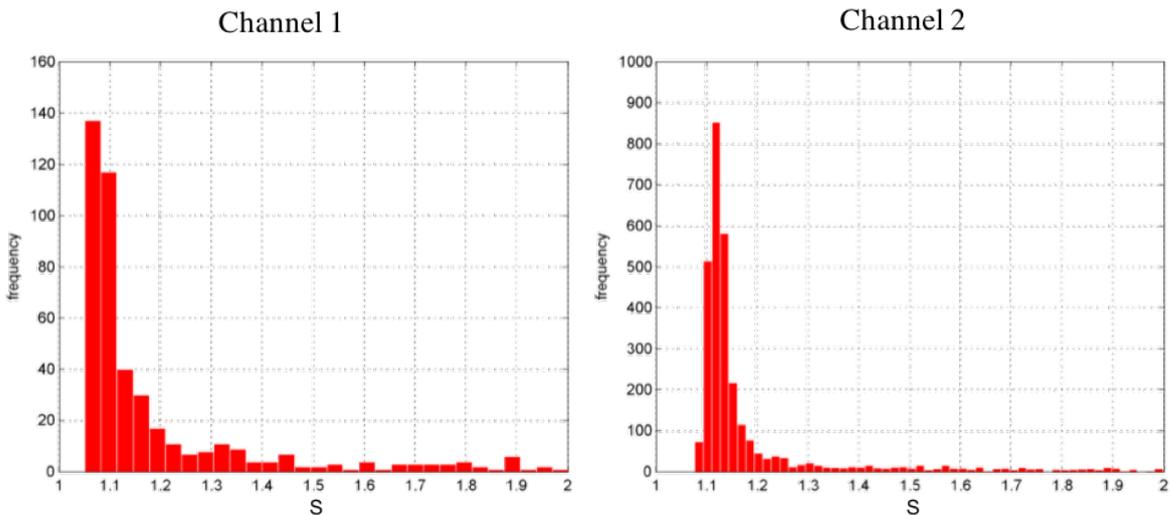


Figure 4.6: Plots to show calculated shape parameter of two particle types derived from tomographic image analysis.

Left) Channel 1 (600 μm particles) **Right)** Channel 2 (300 μm particles). $S = \frac{A_p^3}{36\pi V_p^2}$

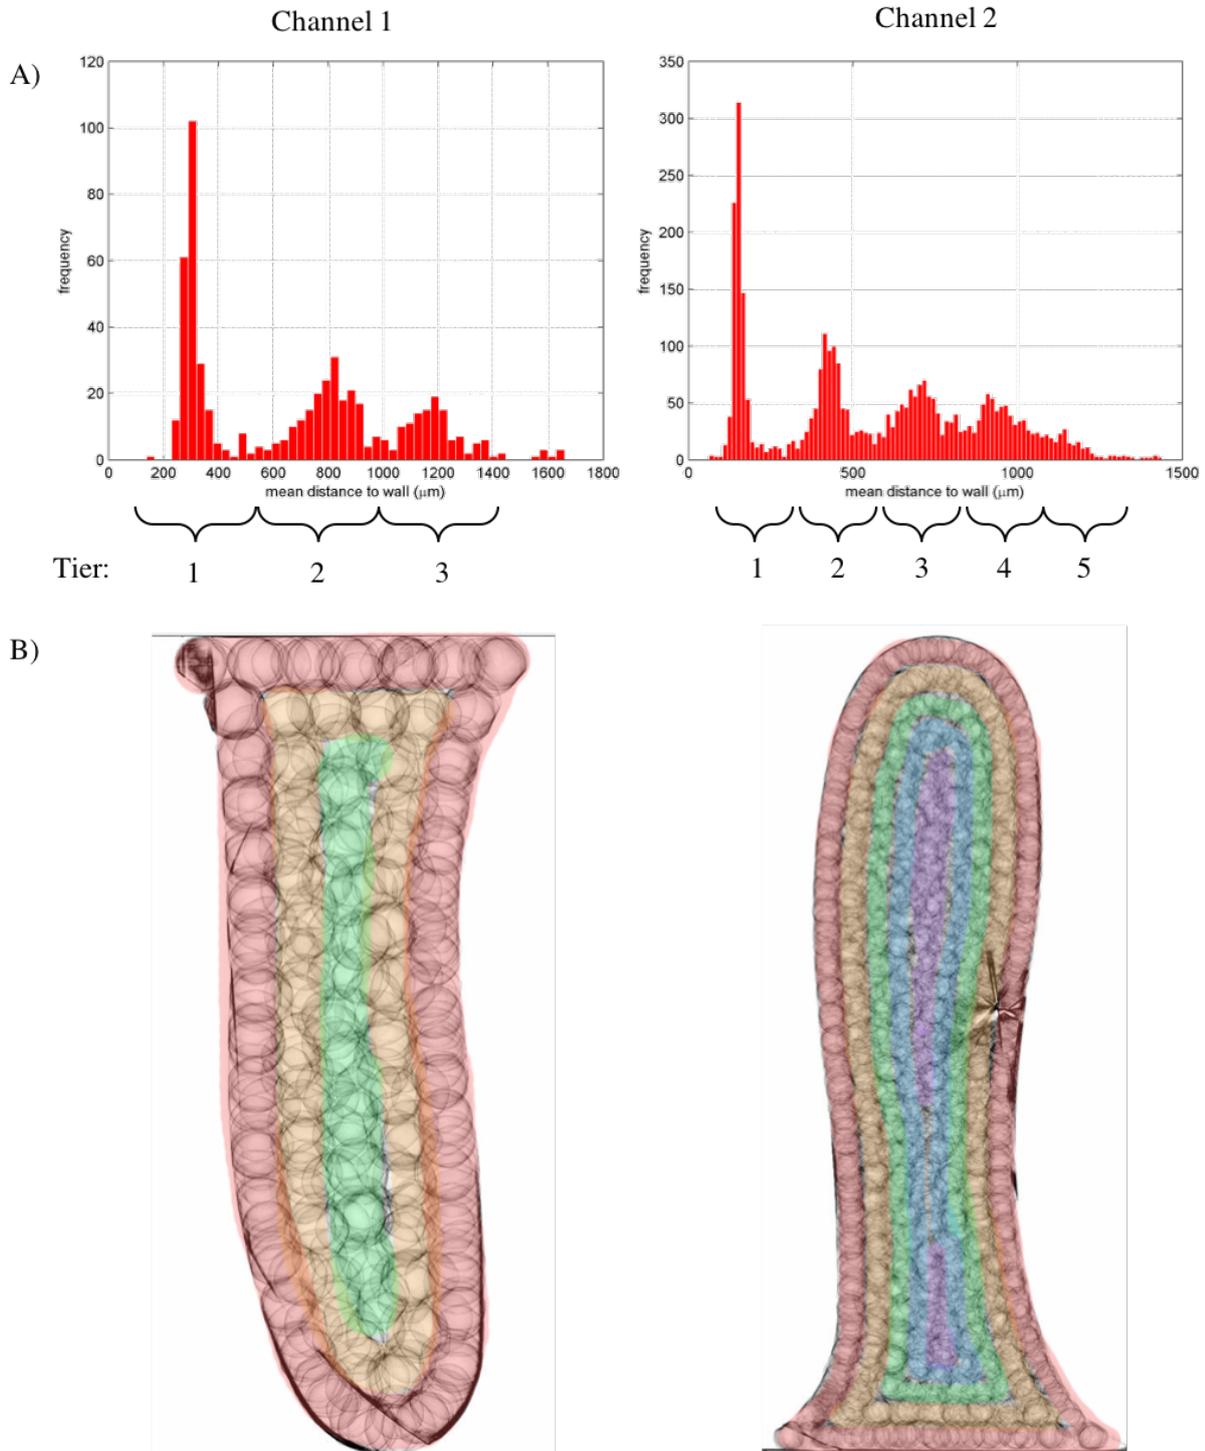


Figure 4.7: Histogram and Tomographic surface z-projection showing tiered arrangement of particles for two diameters

Left) Channel 1 (600 μm particles) **Right)** Channel 2 (300 μm particles).

A) Histogram of mean wall distance (μm) with groupings identified as tiers **B)** Plan view tomographic surface z-projection, coloured by approximate tier

Table 4.2: Mean distance to wall (μm) for each tier (where Δ is the inter-tier distance, expressed in μm and radius units.)

Tier	Channel 1 ($r_m = 300$)	Channel 2 ($r_m = 150$)
1	300	150
2	820 ($\Delta = 520 = 1.79r_m$)	430 ($\Delta = 280 = 1.84r_m$)
3	1190 ($\Delta = 370 = 1.22r_m$)	700 ($\Delta = 270 = 1.79r_m$)
4		940 ($\Delta = 240 = 1.58r_m$)
5		indeterminate

*

Voidage & Tortuosity

The commercial interest of this study is to provide high performance packed beds for carrying out Fischer-Tropsch synthesis to a high degree of conversion and stability for a low unit cost. Therefore, the volume of active catalyst per volume of reactor is a key parameter. The global volume ratio depends on the reactor design as a whole.

At an individual process channel level, the channel geometry and particle size are the main influencing factors contributing to the void fraction of the bed. While theoretical attempts to calculate voidage rely on these geometric parameters and consequently result in a nominal homogeneous voidage, tomographic analysis of a real packed bed allows for a much more localised representation of voidage.

Figures 4.8 & 4.9 show the local average voidage through the height of the bed versus the distance from the wall. The same tiered arrangement is observed as in 4.7. Voidage is higher in the inter-tier space. The mean void fraction for channels 1 & 2 are 0.42 & 0.39 respectively, with the larger voidage observed in the larger particles. This is intuitively the case, since the spaces between two larger contacting particles must necessarily increase. The local average void fraction (with respect to height) at the wall reaches as high as 0.9 and in the inter-tier voids as high as 0.65. However, for smaller particles the width of the inter-tier voids are much shorter (since table 4.2 shows the inter-tier period to be a function of particle radius).

This implies that any fluid flowing through the packed bed would preferentially flow through the void area, rather than traverse a tortuous path through many particles. This effect, known as 'channelling', results in shorter residence times in packed beds and consequently a lower reaction conversion. Smaller particles (or particle:bed aspect ratio) are therefore desirable for reducing the effect of channelling. The use of smaller particles also has the added benefit of increasing the surface area of active, accessible catalyst per unit volume of reactor. Though, perhaps at the expense of reactor stability.

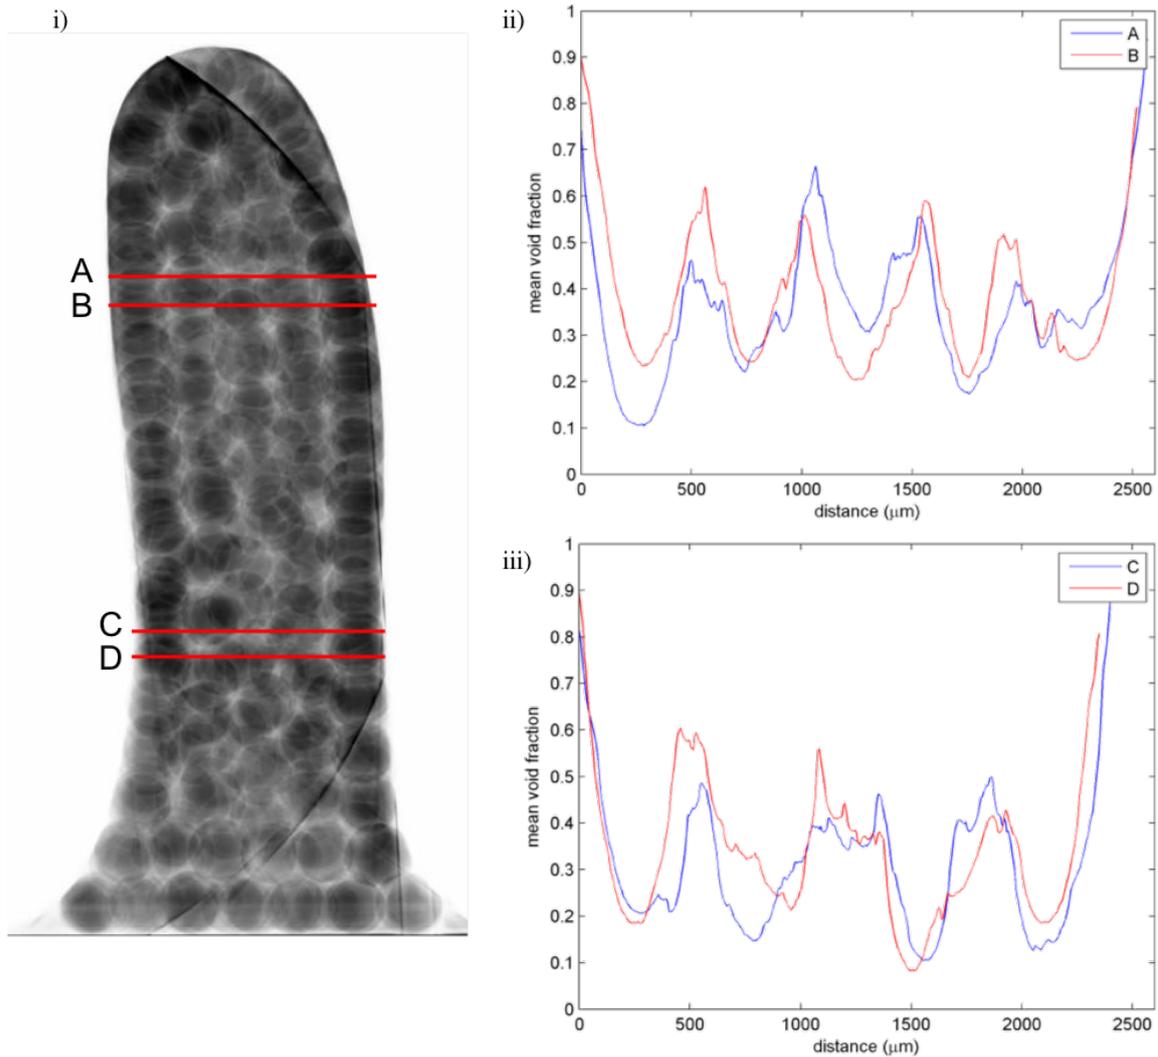


Figure 4.8: Representations of Mean Void Fraction for Channel 1

- i) Channel 1 (600 μm particles) tomographic Z-projection with average colour values (0 to 255) showing location of cross sections A, B, C & D.
- ii) Graph to show mean void fraction versus bed width for cross sections A & B.
- iii) Graph to show mean void fraction versus bed width for cross sections C & D.

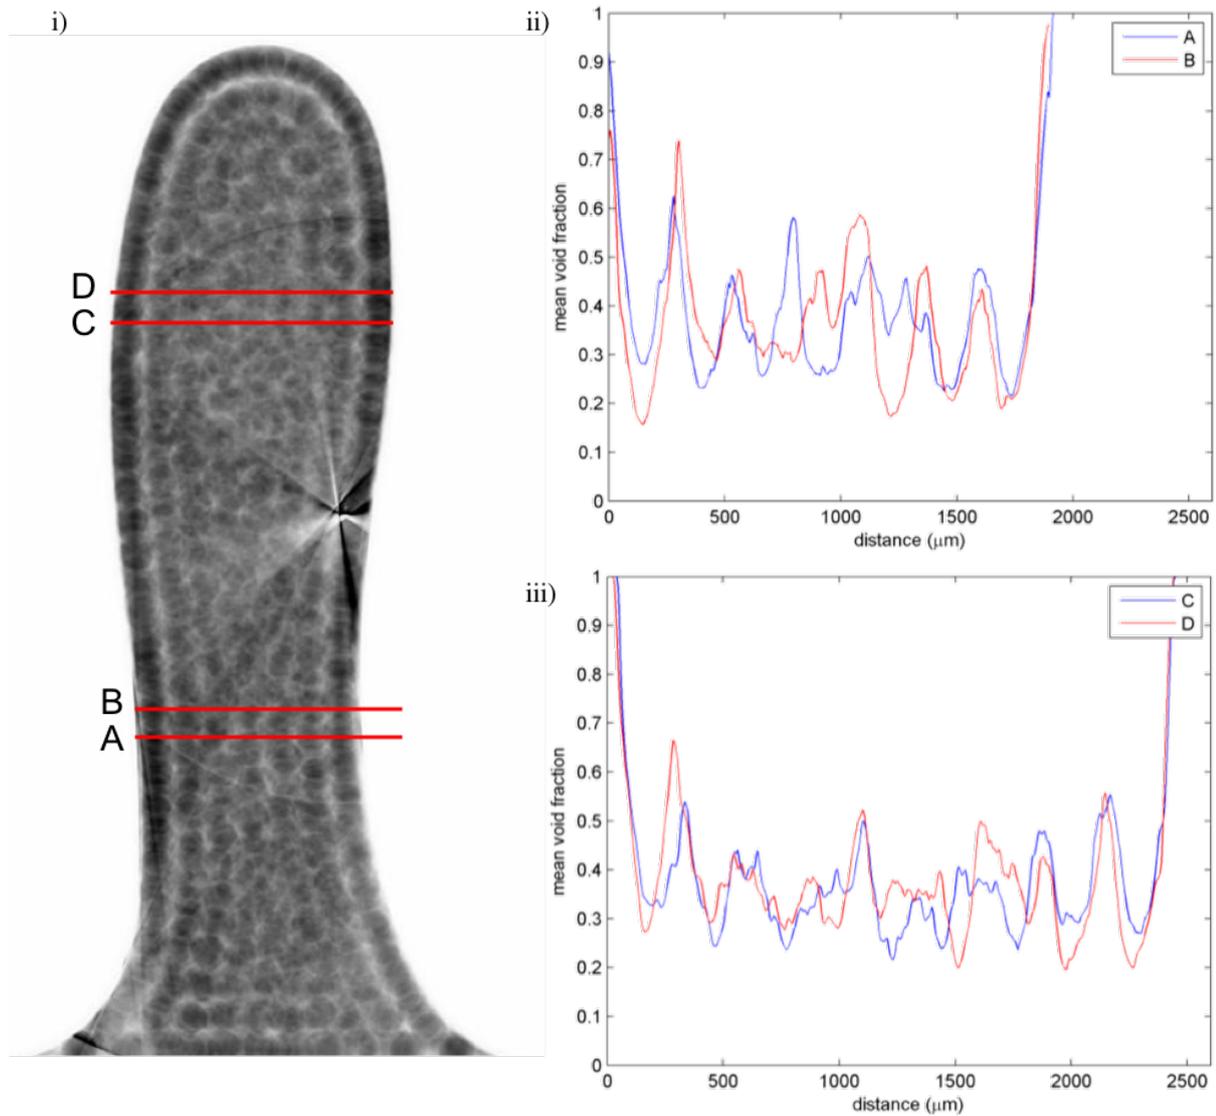


Figure 4.9: Representations of Mean Void Fraction for Channel 2

i) Channel 2 (300 μm particles) tomographic Z-projection with average colour values (0 to 255) showing location of cross sections A, B, C & D.

ii) Graph to show mean void fraction versus bed width for cross sections A & B.

iii) Graph to show mean void fraction versus bed width for cross sections C & D.

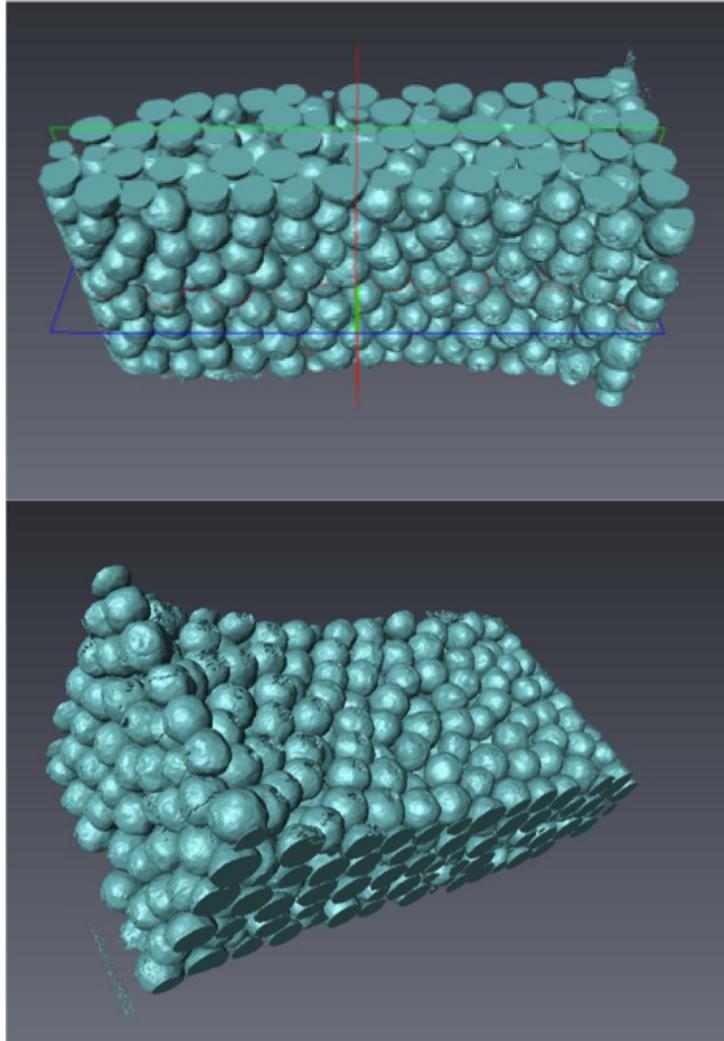


Figure 4.10: 3 Dimensional reconstruction of X-ray tomographic scan

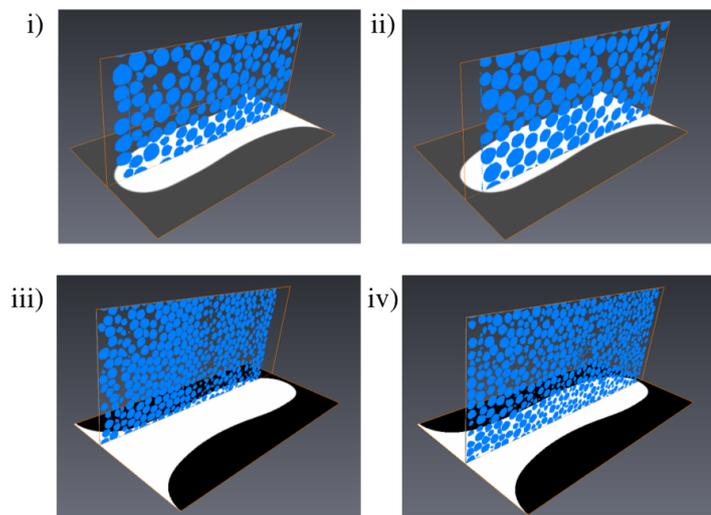


Figure 4.11: 2D planar cross sections of tomographic scans for two particle sizes. i) Central plane, 300 μm particles. ii) Wall proximal plane, 600 μm particles. iii) Wall proximal plane, 300 μm particles. iv) Central plane, 300 μm particles.

4.4 Chapter Summary

Owing to the need to more precisely quantify the structure of the catalyst packed bed, a 3D imaging technique, X-ray tomography was employed. A small mock reactor of the same channel dimensions was loaded with two generations of catalyst particle. After performing the X-Ray scan, several image processing algorithms were applied to the raw data to remove the effect of various artefacts. The fundamental fixed bed parameters and particle properties were compared for the two catalysts under investigation. A tiered pattern of catalyst packing was shown to correspond to the diameter of the particles, giving rise to low local tortuosity. This strongly supported the idea of a degree of fluid bypass and corroborated the observations of the previous chapter. Finally, a 3D model of the catalyst bed was constructed. The geometry file created by this process could be further developed to allow for direct simulation of properties such as pressure drop, which were shown to be elusive in the previous chapter. When taken together, chapters 3 & 4 show that the local performance and state of the reactor is difficult to know during operation. The lack of spatial resolution in the sensing equipment is a limiting factor in our understanding of the operational state. This resolution is set to worsen upon scale-up. In the next chapter, a novel high-resolution sensing method is explored, in an effort to overcome this issue.

Chapter 5. High Resolution Distributed Temperature Sensing of Modular Fischer-Tropsch-Tropsch Reactors

5.1 Chapter Abstract

Modularisation of a conventionally large-scale process presents unique operational challenges. The complex arrangement of modular reactors, whose interactions must be finely balanced to achieve desired productivity, requires a complex process control network. While scalability of the process has been verified, instrumentation and information systems have relied on conventional, full-scale technology. The requirement for high resolution data coupled with the modular arrangement of commercial plants results in an unfeasible requirement for instrumentation in terms of space, economics and operability.

Distributed temperature sensing (DTS) technology is commonplace in industry, applied to long distance pipelines and deep wells. By pulsing monochromatic light through fibre optic cables, a return signal given by Raman backscatter can be processed to give a one dimensional trace of temperature along the pipeline. Spatial resolutions in the order of several meters and temperature resolution in the order of degrees is adequate for traditional applications.

In this work, a novel application of DTS technology is explored. Fibre optic cable is arranged in a 3-dimensional configuration inside in a compact modular Fischer-Tropsch reactor. The DTS system is tuned to provide the necessary high spatial, temporal and temperature resolution for reactor operability, with the result that detailed reactor temperature profiles can be resolved in real time. The novel system is benchmarked against the conventional discrete temperature sensing elements and recommendations for potential commercial applications are discussed. The new method reduces a control problem requiring several thousand meters of cable, and hundreds of I/O channels to a scalable solution of a single channel per module.

5.2 Introduction

5.2.1 Distributed Temperature Sensing Principles

The Limitations of Discrete Temperature Sensing

Distributed Temperature Sensing is contrasted with the more commonplace discrete sensing method. Discrete sensing involves the use of electronic or mechanical sensors at defined points. When applied to reactor control, this typically means thermocouples wired physically into place at several key locations within a reactor. While thermocouples are relatively inexpensive and easily replaceable in conventional reactor systems, modular compact FT reactors present a unique challenge: High spatial resolution temperature data is required inside a complex 3D reactor architecture. To achieve sufficient spatial resolution to build an adequate profile of reactor temperature with discrete methods would require installing several thermocouples per reactor. Considering that even the most modest sized modular GTL plants would require upwards of 50 compact FT reactors, maintaining control by scale-out of the distributed sensing strategy would result in a challenging control system from an installation, maintenance and operations standpoint. Figure 4 shows the scale-out challenges of discrete temperature sensing compared with the much more saleable distributed temperature sensing solution.

In a distributed temperature sensing (DTS) system, the measurements do not rely on electronic or mechanical sensors at discrete locations along the line, but rather by observing the physical interactions between probing light energy and a fibre optic waveguide. Because the measurements do not rely on a multitude of discrete sensors, they can be taken at virtually all locations along the length of a cable and are inherently more reliable than former technology. The optical fibre can be thought of as a long thermocouple, along the entire length of which the distributed temperature profile can be determined (AFL Global, 2015).

How DTS Works

Distributed Temperature Sensing systems are optoelectronic devices which measure temperatures by means of optical fibres functioning as linear sensors. Temperatures are recorded along the optical sensor cable not at points, but as a continuous profile.

A short laser pulse is fired down a fibre optic cable. The length of the pulse is typically 10 ns or less (related to the spatial resolution, discussed later). Optical fibres are made from doped quartz glass. Quartz glass is a form of silicon dioxide (SiO_2) with amorphous solid structure. This light pulse propagates down the fibre through a region called the mode field, consisting of a core and cladding material (discussed later). Figure 5.2 shows

a general schematic of this principle.

The velocity of light is dependent on the medium through which it travels. For glass with an index of refraction of ~ 1.5 , the velocity of light is 2×10^8 m/s.

However, physical phenomena such as temperature or pressure and tensile forces, can affect glass fibres and locally change the characteristics of light transmission in the fibre. As a result of the attenuation of the light in the quartz glass fibres through scattering, the location of an external physical effect, such as a temperature event, can be determined so that the optical fibre can be employed as a linear sensor.

Thermal effects induce lattice oscillations within the glass. When light falls onto these thermally excited molecular oscillations, an interaction occurs between the incident photons and the electrons of the molecule. Light is scattered in all directions during this interaction. This scattered light undergoes a spectral shift by an amount equivalent to the resonant frequency of the lattice oscillation (Figure 5.3).

A fraction of the light is scattered back in the direction of incidence, and undergoes a spectral shift as outlined in Figure 5.3. The backscattered signal contains components of; Rayleigh scattering, with the same wavelength as the incident light, Stokes components with a higher wavelength than the incident light ('red shifted'), and Anti-Stokes components with a lower wavelength than the incident light ('blue shifted'). The shifted light is further subdivided into so called Brillouin and Raman peaks.

The intensity of the Raman Anti-Stokes peak is temperature-dependent, while the Raman Stokes peak is practically independent of temperature. The local temperature of the optical fibre can therefore be derived from the ratio (R_r) of the Anti-Stokes (λ_a) and Stokes (λ_s) light intensities by the following relation (Berthold, 2011):

$$R_r = \left(\frac{\lambda_s}{\lambda_a} \right)^4 \exp \left(-\frac{hcv}{kT} \right) \quad (5.1)$$

Where h is Planck's constant, c is the velocity of light, v is the optical frequency of the incident light, k is Boltzmann's constant and T is the absolute temperature.

The exploitation of this principle by filtering the backscattered light to extract Stokes and Anti-Stokes peaks is known as Optical Time Domain Reflectometry (OTDR), or Optical Frequency Domain Reflectometry (OFDR), depending on whether the incident light is modulated in the time or frequency domain.

Distributed Temperature Sensing is a rising technology. As the technical limitations and cost of DTS systems decrease the applications and prevalence will increase over former technologies providing greater accuracy and resolution of data across longer distances. This technology can be used in harsh environments where former technology would have

been prone to failure or unreasonable to deploy.

5.2.2 Fibre Optic Principles

Fibre Anatomy

A fibre optic cable is a form of waveguide employing total internal reflection to transmit a signal for long distances at high speed with minimal loss. The waveguide geometry in a fibre optic cable is formed by the differing refractive indices of two concentric materials, the core and cladding. Electromagnetic radiation enters the waveguide geometry at a given angle, θ . When θ is less than some critical value, θ_{max} , which depends on the properties of the fibre, all light is totally internally reflected. The angle at which some light is first able to escape the waveguide is the critical angle, θ_c (Kapron et al., 1970).

The anatomy of a general optic fibre (F/O) is shown Figure 5.6 to consist of a core, cladding, coating and jacket layers.

Fibre Modality

The functional aspects of the naked fibre in the DTS element are the core and cladding. The interaction between the two elements is not trivial. The core and cladding define the waveguide geometry.

For cores of small diameter (typically in the order of 10 μm) an incident light pulse may only find one possible path down the length of fibre. (that is to say, one solution to the Helmholtz equation for waves) (A. et al., 2011). For larger diameters, typically above 50 μm , it is possible for light entering the fibre to find multiple paths through the cable. The list of possible propagation paths is referred to as the ‘modes’. Naturally, depending on the incident wavelength and core diameter, it is possible to state whether a fibre is a single mode (SMF) or a multi mode fibre (MMF). It is also possible to calculate the number of modes in a MMF.

Firstly, the normalised frequency, V , is calculated:

$$V = \frac{2\pi\alpha}{\lambda} \sqrt{n_1^2 - n_2^2} \quad (5.2)$$

where α is the core radius, λ is the wavelength. n_1 and n_2 are the refractive indices of the core and clad respectively. The numerical aperture, NA, is often specified in place of the two separate indices and is defined as:

$$NA = \sqrt{n_1^2 - n_2^2} \quad (5.3)$$

The Numerical aperture of an optical system is a dimensionless number that characterises the range of angles over which the system can accept or emit light (P and K, 1982).

Thus, normalised frequency may be calculated from supplier specified parameters:

$$V = \frac{2\pi\alpha}{\lambda} NA \quad (5.4)$$

For the fibre used in this study, whose physical dimensions are given in Table 5.1, and discussed at length later, the calculation is as follows: For 1300 nm light and 25 μm core radius:

$$V = \frac{2 \cdot \pi \cdot 25 \times 10^{-6}}{1300 \times 10^{-9}} 0.2 = 24.1 \quad (5.5)$$

The criteria for a single mode fibre is that $V < 2.4048$, the first root of the Bessel Function (J_0) (Pal, 2015). In this case, V is significantly greater than the single mode cut off. It is verified to be multi-mode fibre.

Modal Dispersion

In the case of a single mode fibre, where light travels in only one path along the length, a signal input is perfectly represented by the output with a time delay equal to:

$$t_L = \frac{n_1 L}{c} \quad (5.6)$$

where L is the length of fibre traversed, n_1 is the refractive index of the core and c is the speed of light.

However, for a MMF, it is possible for light to take several paths through the fibre. Figure 5.7B illustrates the nature of multimode light transmission through an optic fibre. It can be appreciated that the length of each path taken through a multimode fibre will be slightly different. The shortest path through a given length of fibre is achieved in a single mode fibre; parallel to the axial direction of the fibre. The longest path in a multimode fibre is given by light traveling at the critical angle.

The range of possible path lengths (and therefore flight times) in MMF means that incident light is dispersed in the output signal, with different light arriving from different paths at a range of times. The effect is analogous to chromatography, where particles following a more torturous path through a chromatographic medium arrive later than those that pass

freely through. In fibre optics, this effect is amplified when polychromatic light is used, since different wavelengths of light are slowed by varying degrees through a medium.

Modal dispersion describes the peak spreading of output signals compared to input signals. For a single pulse of incident light, this does not present a significant problem. However, a series of initially well resolved pulses separated by fractions of a second may be unresolvable at the output. This problem is worsened with increasing distance.

To quantify the modal dispersion, it is possible to calculate the time window for an output signal from an instantaneous pulse of input light through a multimode fibre.

First, the time for light travelling at a given incident angle may be calculated:

$$t_L = \frac{n_1 L}{c \cos \theta} \quad (5.7)$$

Now, the minimum time t_{min} will occur at $\theta = 0$, which reduces the above equation to the single mode equation shown previously. The maximum time take will occur at the critical angle:

$$\theta_c = \cos^{-1} \frac{n_2}{n_1} \quad (5.8)$$

where n_2 is the refractive index of the cladding. Substituting this into the previous equation gives t_{max} .

$$t_{max} = \frac{n_1 L}{c \cos(\cos^{-1} \frac{n_2}{n_1})} \quad (5.9)$$

The simplified expression for t_{max} :

$$t_{max} = \frac{n_1^2 L}{c n_2} \quad (5.10)$$

So, for an instantaneous input pulse, the output signal would occupy a time interval, τ_i , given by:

$$\tau_i = t_{max} - t_{min} = \frac{n_1 L}{c} \left[\frac{n_1}{n_2} - 1 \right] \quad (5.11)$$

A common simplification in industry is to replace the multiple refractive indices with the numerical aperture, NA , defined previously. This is possible for core and cladding materials with suitably similar refractive indices. The measure of relative difference between the two indices is given by defining a parameter, Δ

$$\Delta \equiv \frac{n_1^2 - n_2^2}{2n_2^2} \quad (5.12)$$

And where $\Delta \ll 1$ (similar refractive indices):

$$\Delta \cong \frac{n_1 - n_2}{n_2} \quad (5.13)$$

Therefore, from the definition of numerical aperture, NA , given previously, the approximation follows:

$$\tau_i \approx \frac{n_1 L}{2c} \Delta \approx \frac{L}{2n_1 c} NA^2 \quad (5.14)$$

Importantly, the modal dispersion interval, τ_i , is proportional to the square of NA .

Modal dispersion is problematic for DTS applications for two main reasons. Firstly, several pulses per second (typically ~ 50 kHz) are required to give suitable sample resolution for accurately determining temperature. With sufficient modal dispersion, the output signal from separate pulses would merge into an unresolvable broad peak. Secondly, DTS does not rely on the transmission of the source signal to the end of the fibre, L , but instead on the return signal backscattered from each increment. Since dispersion is a function of length, each return signal is subject to a different modal dispersion. The effect is doubled since any given backscattered signal will have travelled both an outwards and inwards path to make the path back to the detector. Clearly Modal dispersion is to be minimised if clean, high frequency signals are to be achieved. Intuitively, reducing the numerical aperture is one way of achieving this. However, this reduces the range of angles over which the system can accept light from source emitter, so much finer lasers are required for low NA optics. A well designed fibre optics system must therefore compromise between the numerical aperture (and consequently, the data resolution) and the laser source (and therefore hardware cost).

Calculations presented so far have relied on the concept that the refractive index of the core and clad are distinct. The boundary between the two materials marks a sudden step change in refractive index. This class of fibres is referred to as ‘step index’ fibres. In contrast to step-index, there exists another class of fibre optic cables called ‘graded index’. In graded index fibres, the refractive index of the core is not uniform, and forms a gradient from the centre axis, $r = 0$, to the core/clad boundary, $r = a$.

For a graded index fibre, it will be appreciated that a number of potential refractive index profiles are possible. Refractive index gradients can range from linear to sharply curved. In order to quantify the characteristic shape of a given refractive index profile in relation to the physical properties of the system, a profile parameter, g , is included in a power-law

relation:

$$n(r) = \begin{cases} n_1 \sqrt{1 - 2\Delta \left(\frac{r}{\alpha}\right)^g}, & r \leq \alpha \\ n_1 \sqrt{1 - 2\Delta}, & r \geq \alpha \end{cases} \quad (5.15)$$

where n_1 is now re-defined as the refractive index of the central axis of the fibre core:

$$\lim_{r \rightarrow 0} n(r) = n_1 \quad (5.16)$$

Furthermore, n_1 is defined as the maximum refractive index in the system.

For the clad material ($r \geq \alpha$), the case above simplifies to an expression of uniform refractive index, this is the same as for step index fibres:

$$n(r) = n_2 \quad (5.17)$$

However, the refractive index profile of the core ($r \leq \alpha$) is given by:

$$n^2(r) = n_2^2 \left(\frac{r}{\alpha}\right)^g \quad (5.18)$$

This gives rise to a standard power-law relation. To explore the potential profiles produced by this expression, consider the two extremes of r in the expression; 0 and α .

At $r = 0$, $\frac{r}{\alpha} \rightarrow 0$, so $n(r) = n_1$

(by definition, the refractive index at the central axis)

At $r = \alpha$, $\frac{r}{\alpha} \rightarrow 1$, so $n(r) = n_2$

The above is true for all values of g . These are the boundary conditions with respect to r . Now, with respect to g :

At $g = 0$, no profile is present. $n(r) = n_2$. At $g = 1$, the profile is linear between the two boundary conditions. At $g = \infty$, the refractive index approaches the limit n_1 . At all other values of g , the curvature follows an increasing sharpness towards a step index profile as demonstrated in Figure 5.9.

For all practical purposes, any fibre optic cable where $g > 3$ is said to approach the behaviour of a step index fibre (Pask and Sammut, 1980). In order to find the optimum refractive index profile, it is necessary to consider the geometric principles.

The problem of modal dispersion is a result of light entering the system at different angles. These angles are reflected back at the same angle and therefore never deviate from their

path. Each path takes a different time to travel, since the paths are varying length but the speed of light through the medium is the same (since the refractive index is constant). In contrast, different modes in a graded index fibre are subjected to lower refractive indices close to the core/clad boundary. Consequently, high order modes, whose path is often close to the boundary, travel faster through the lower refractive index.

The optimum profile parameter would therefore be one which always refracts light back to the central axis, regardless of its entry angle. This property is found in parabolas, given by quadratic functions. This corresponds to a profile parameter of $g = 2$.

Figure 5.10 demonstrates how different modes are propagated through a graded index fibre. While it is clear that the path length for higher order modes is greater, the refractive index profile means that light travelling closer to the boundary are travelling faster than those in the centre. The difference in speed and path length are balanced for fibres with a parabolic refractive index profile and each mode remains in phase, eliminating modal dispersion. In real systems, the refractive index profile parameter is close to 2, and modal dispersion is minimised.

Now, calculation of the modal volume, M , may proceed. For a multi mode, parabolic graded index fibre (MMGI) ($g = 2$), with normalised frequency $V > 5$ (Pask et al., 1975):

$$M = \frac{V^2}{2} \left(\frac{g}{g+2} \right) = \frac{V^2}{4} = \frac{24.1^2}{4} = 145 \quad [\text{modes}] \quad (5.19)$$

This compares with the number of modes in a step index fibre ($g \rightarrow \infty$)

$$M = \frac{V^2}{2} \left(\frac{\infty}{\infty+2} \right) = \frac{V^2}{2} \quad (5.20)$$

Therefore, the number of modes in a graded index fibre is approximately half that of a step-index fibre of equal diameter, further reducing the impact of modal dispersion.

Fibre Properties

Given the considerations mentioned in the previous discussion, a multi mode graded index fibre was used in this study. The optical properties, while important, were not the only considerations in selection of the appropriate fibre. Physical properties and dimensions for the AFL VHM3000 series fibre are given in Table 5.1.

As discussed previously, the typical operating temperature for FTS is in the range of 180 – 250 °C. The exothermic reaction is liable to thermal excursion when performed with aged catalyst. Peak temperatures for a thermal excursion brought adequately under control should not exceed 250 °C, but temperatures up to 300 °C have been observed in some

Table 5.1: Manufacturer supplied specifications for DTS element fibre portion. [†]Concentricity measured as (min. wall / max. wall). *Attenuation Measured on loose coil. (AFL Global, 2015)

Properties		Value
Material	Hermetic coating	none
	Coating	Polyimide
	Max pressure (MPa)	≥ 690
Geometry	Core diameter (μm)	50 ± 2.5
	Clad diameter (μm)	125 ± 2
	Core non-circularity (%)	≤ 5
	Clad non-circularity (%)	≤ 1
	Core/Clad offset (μm)	≤ 1.5
	Coating diameter (μm)	245 ± 10
	[†] Coating concentricity (%)	≥ 80
	Optical	Numerical aperture (nominal)
*Attenuation:		
@ 850 nm (dB km^{-1})		≤ 3.0
@ 1300 nm (dB km^{-1})		≤ 3.0

cases. Any reactor achieving a temperature in this upper region would likely suffer from some form of catalyst degradation and significant carbon formation and should be taken offline for inspection and re-loading. However, a fibre optic cable observing the reactor should be suitably selected to survive such a thermal event, in order to safely control the reactor during the shut-down procedure.

The selected fibre in this study is specified as withstanding temperatures of $\geq 300^\circ\text{C}$. While no experimental effort was made to establish the true upper limit of temperature tolerance, guidance from the manufacturer suggested that this fibre, designed for ‘harsh environment’ was well suited to the scope of the study. Future studies should seek to establish the maximum operating temperature of the fibre, and also show both long and short term effects of operating at elevated temperatures.

Fischer Tropsch reactions are commercially operated at elevated pressures, typically above 20 bar. The fibre should also, therefore, be capable of withstanding operational pressures. Furthermore, shutdown strategies for tripped or runaway reactors may include stages of rapid compression or decompression. Though not specified by the manufacturer, the maximum tolerable rate of change of pressure should be established for the fibre.

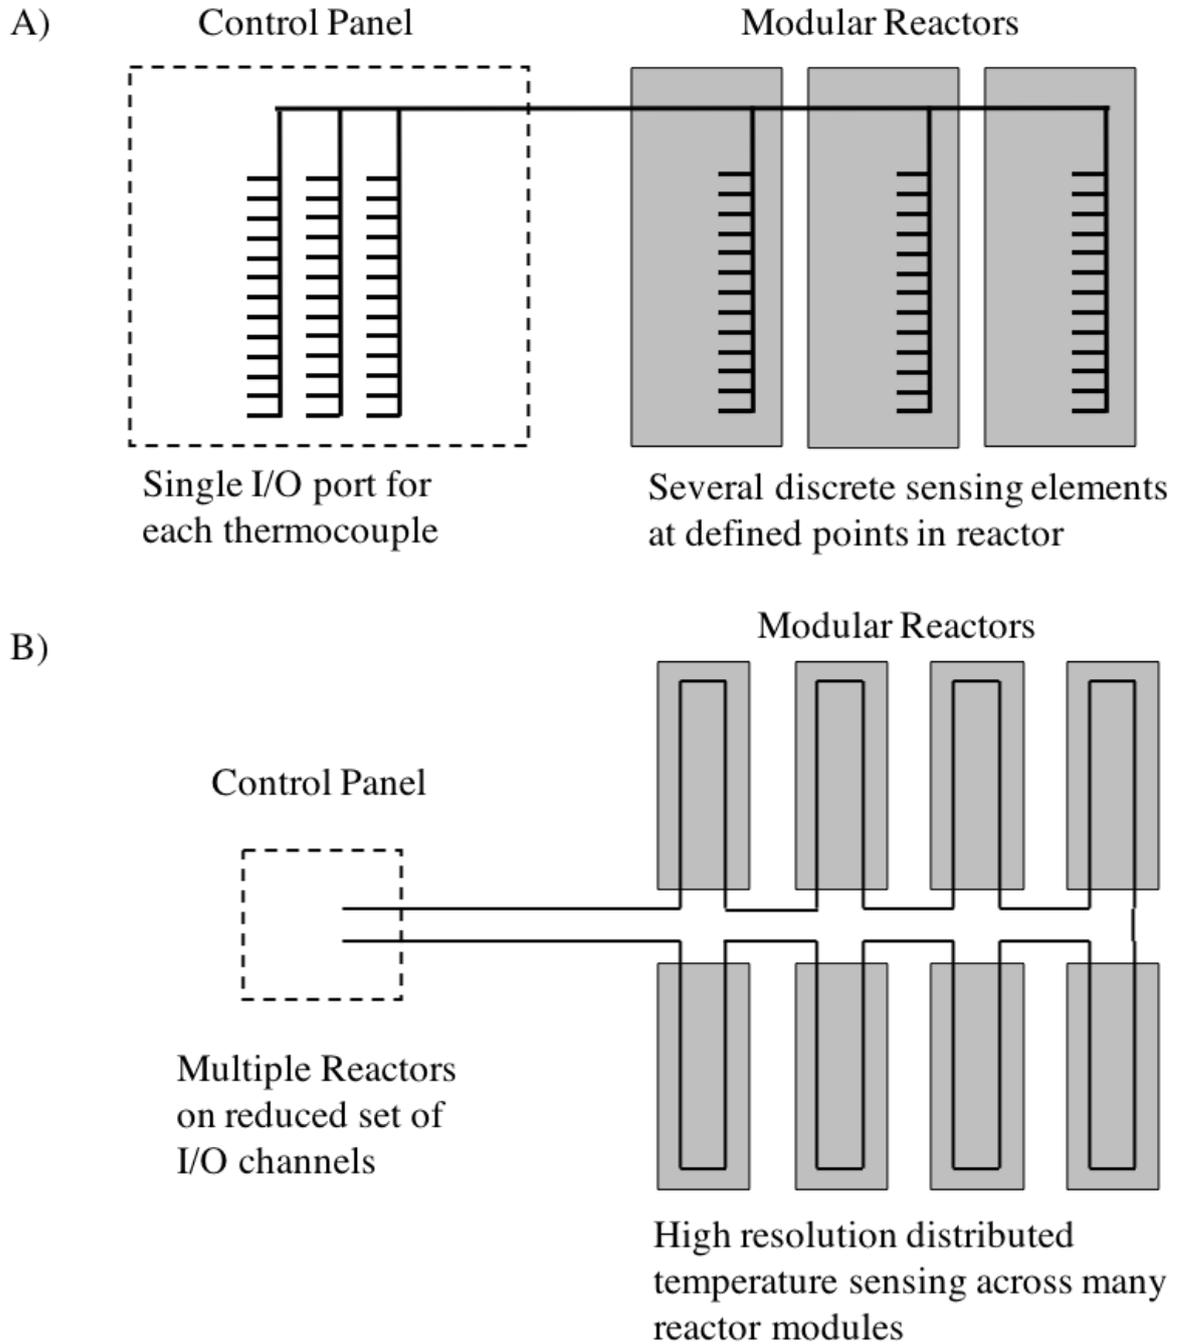


Figure 5.1: Schematic to illustrate differences in scale up / scale out behaviour for discrete (above) and distributed (below) temperature sensing schemes. Black lines show sensing elements.

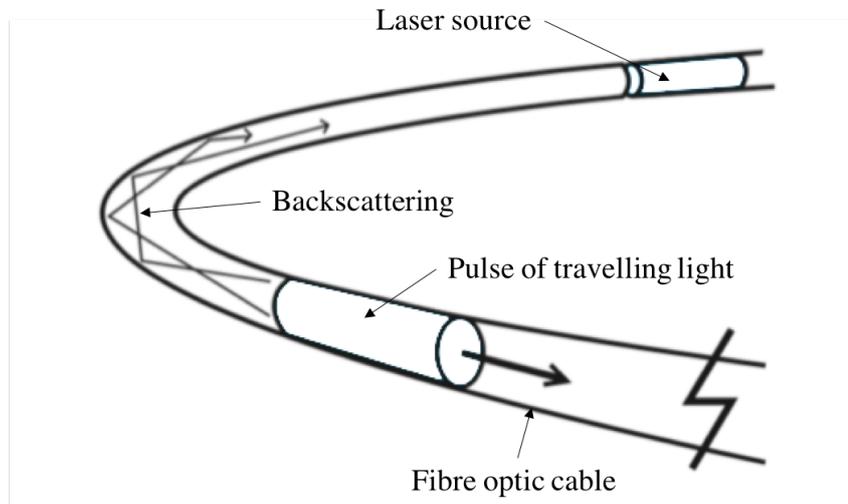


Figure 5.2: Schematic fibre optic element showing light pulse and corresponding backscattered signal due to external physical effects. (AFL Global, 2015)

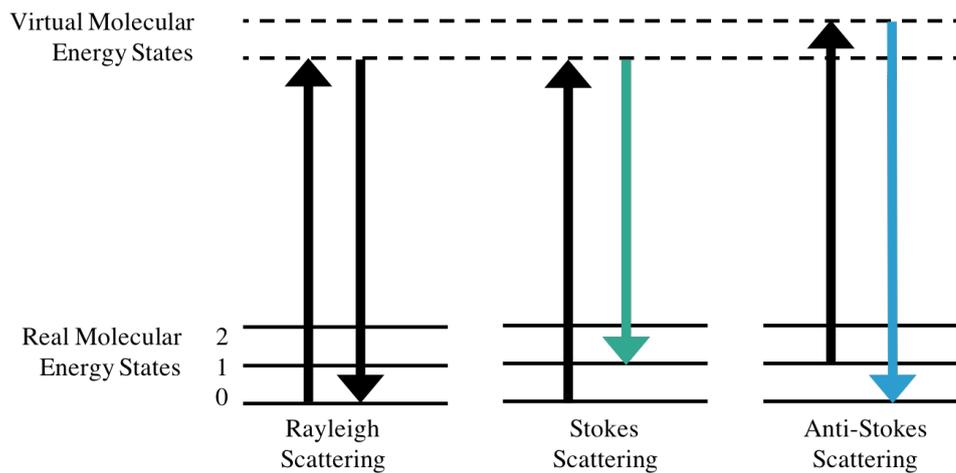


Figure 5.3: Schematic to show three possible photon shifts following photon/molecule interaction in a fibre optic element. (AFL Global, 2015)

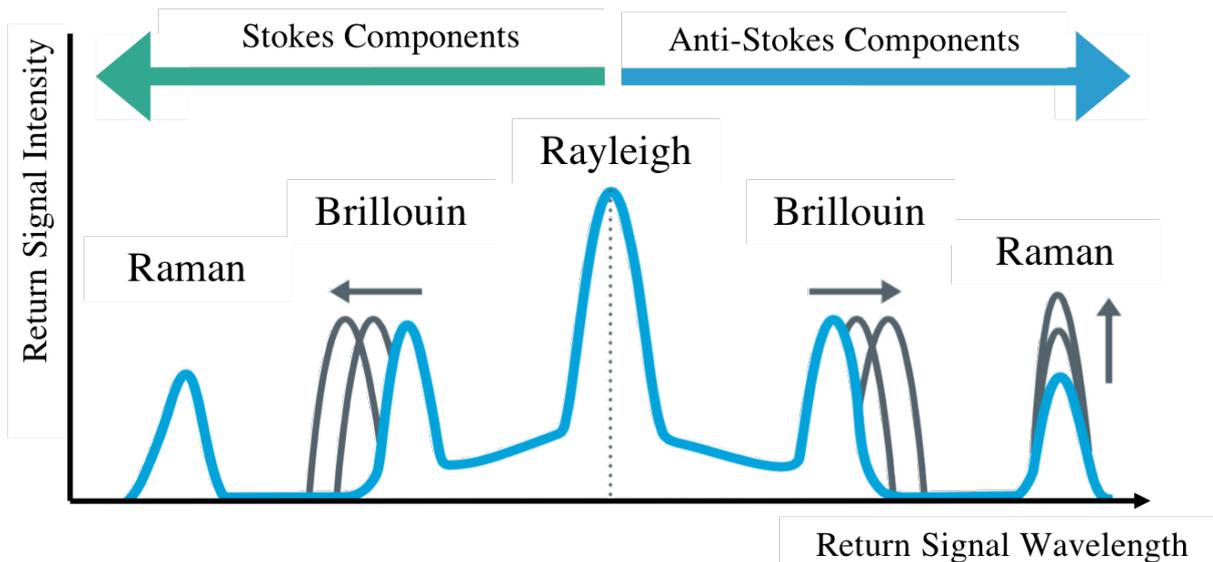


Figure 5.4: Schematic to show signal intensity vs frequency for backscattered portions of light, showing temperature dependent shifts. (AFL Global, 2015)

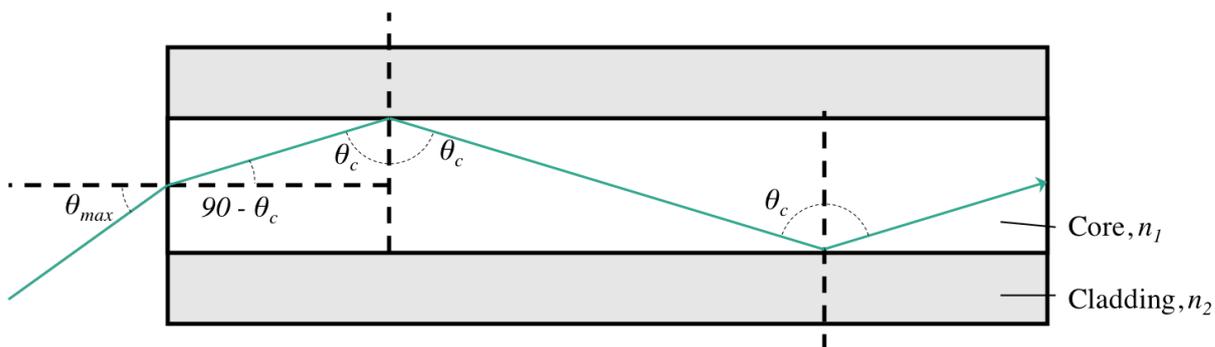


Figure 5.5: Schematic fibre optic waveguide demonstrating the critical angle, θ_c , given two refractive indices, n_1 & n_2 .

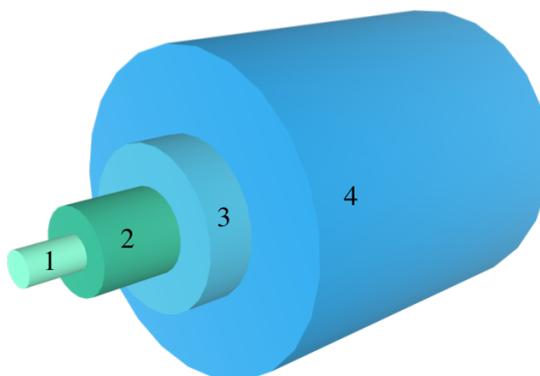


Figure 5.6: Schematic cutaway of a generalised optic fibre section showing: 1) Core, 2) Cladding, 3) Coating, 4) Jacket.

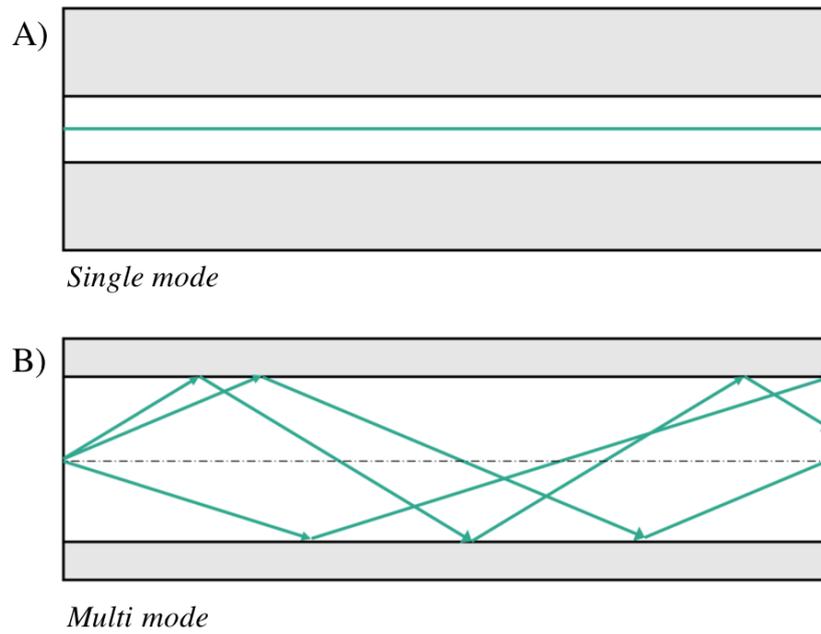


Figure 5.7: Comparison of single mode (A) and multi mode (B) light transmission through a waveguide.

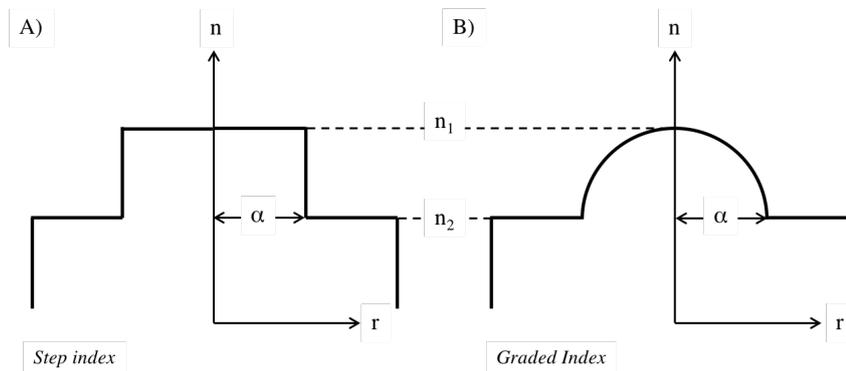


Figure 5.8: Comparison of core/clad refractive index profile for step index (a) and graded index (b) fibres. (Willingale, 2007)

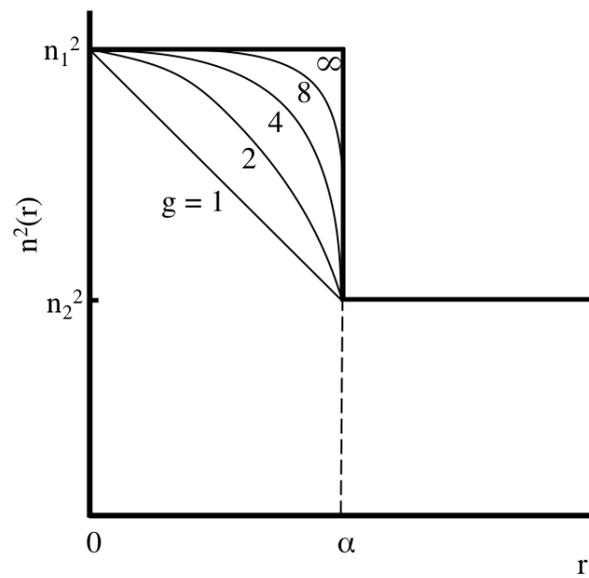


Figure 5.9: Graph to show effect of profile parameter, g , on refractive index profile of a graded index fibre

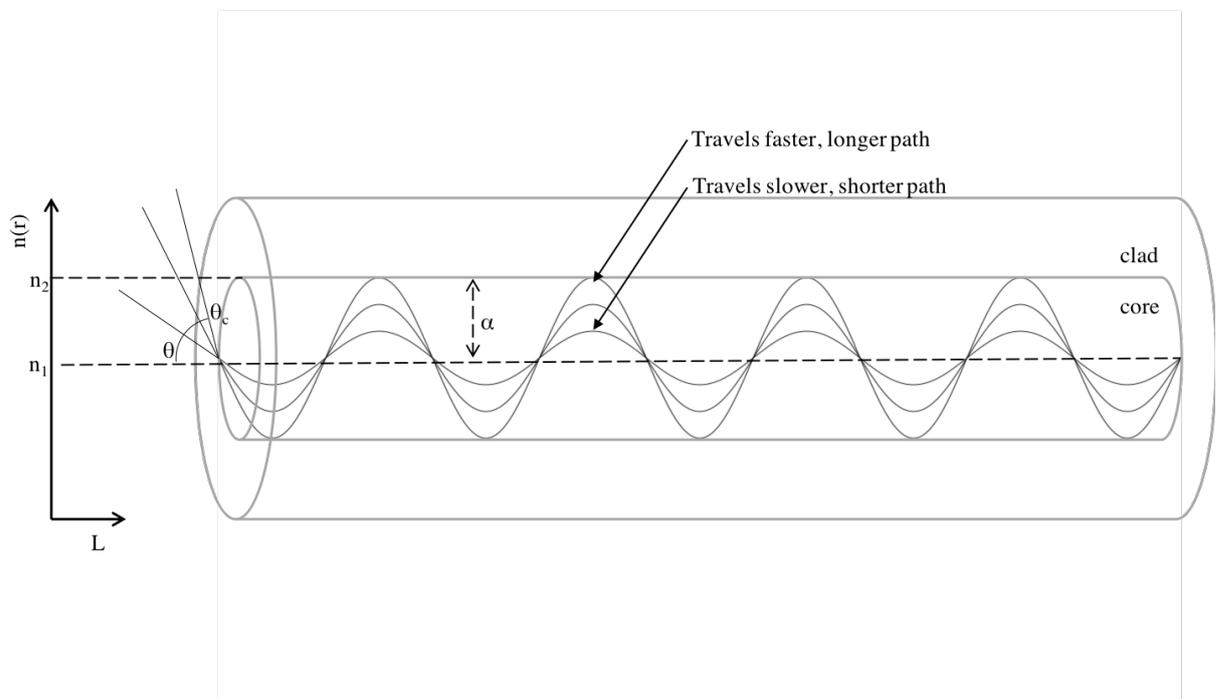


Figure 5.10: Mode propagation in a graded index fibre optic cable with a profile parameter of $g = 2$

5.3 Installation of Fibre Optics into a modular FT reactor

Installation of the DTS element presented many challenges, both physically and in terms of comparability to a full-scale application. Successful implementation on the small scale required several considerations to be made. Installation of the fibre optic cable was in itself an important study, since challenges faced and freely overcome at the small scale must be translated into a set of formal guidelines for installation on a commercial plant where ‘trial and error’ would not be tolerated.

5.3.1 Pilot scale reactor

Aside from the design elements of the reactors necessary for performing FTS reactions, already discussed, A number of aspects of the reactor are present to accommodate control hardware. ‘Divider’ or ‘tube’ plates present in the brazed, layered structure of the reactors consist of a complex arrangement of thin axial plates and several narrow tubes mounted at regular intervals in a transverse direction to the reactant flow direction.

These transverse tubes are designed to accommodate a discrete temperature sensing element to detect the reactor body temperature during operation. The tubes have a specified internal diameter of 1.12 mm. The thermocouples normally used within the tubes are 1 mm diameter. This ensures that the thermocouple is guaranteed to be in close proximity to the reactor body and likely to be touching the tube wall at the point where the temperature reading is made.

During the trial with the F/O cable *in situ*, it was necessary to also include the proven thermocouple technology. This would ensure that safe operation could be continued in the event that the F/O trial was a failure, but also allowed for benchmarking the current system against the DTS system. However, with a specified maximum diameter of 255 μm , the optical fibre would not fit in the same tube as the thermocouple, even under the best case scenario; a perfectly straight thermocouple making a uniform crescent-shaped space with the tube wall.

For this reason, the thermocouples used in the tube plates were re-specified as $\varnothing 0.75$ mm thermocouples, verified to give similar performance to the $\varnothing 1$ mm thermocouples at the expense of wall proximity or contact. A further improvement, not tested in this study, would be to fill the tubes with a suitable heat transfer paste after mounting thermocouples and F/O. This would overcome the issue of wall contact but may prove difficult to remove. The compatibility of the paste with the F/O would also need to be determined. Such a paste, when cured, might also concentrate areas of stress on the fibre at the insert and exit, making it more liable to breakage.

In the absence of such a modification, the F/O was able to be mounted alongside the

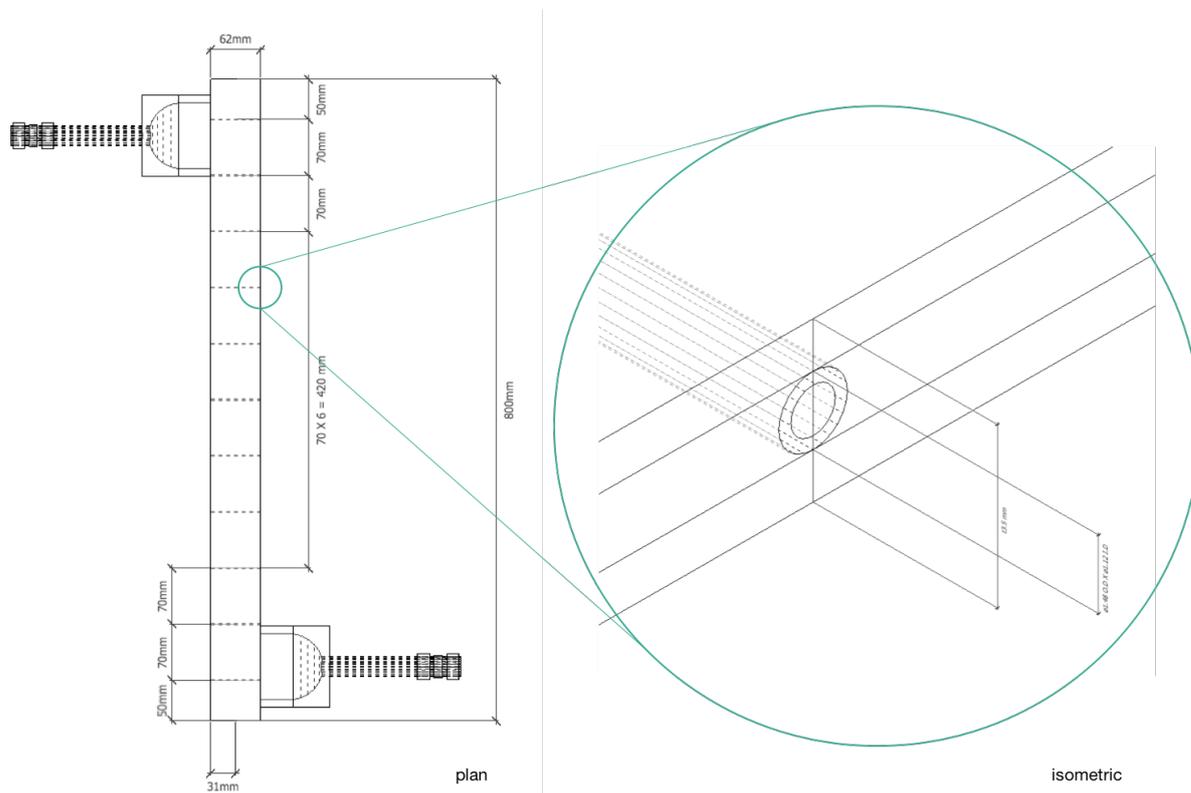


Figure 5.11: Detail of tube plate showing spacing and header impingement

thermocouple with little difficulty. Indeed, it was even possible to load two or more F/O in the same tube. This allowed for a number of possibilities to enhance the robustness of the DTS system.

- A single F/O could be mounted in several passes through the tube, giving repeat readings at different lengths along the fibre.
- Two or more separate cables could be mounted into the tube, allowing for redundancy or possibly the integration of a second measuring principle (e.g. distributed acoustic sensing, DAS) or second incident light wavelength.
- A thicker, more ruggedised fibre could be mounted to mitigate breakage or degradation (physical or chemical) of the fibre over time.

For this study, the same fibre was mounted with two passes through each transverse tube, allowing for a second measurement to be taken at each length within the reactor. Since two tube plates were present in the pilot scale reactor, and transverse tubes were located in the same axial location in each plate, further duplication of measurement for a given axial location was thus possible.

An important limitation of the tube plate design in the reactors can be observed in Figure 5.11. Both the first and last two tubes do not extend fully in the transverse direction, rather only to the mid-point of the reactor body. This is due to the presence of coolant inlet and outlet headers. For thermocouple mounting, this is not problematic since the

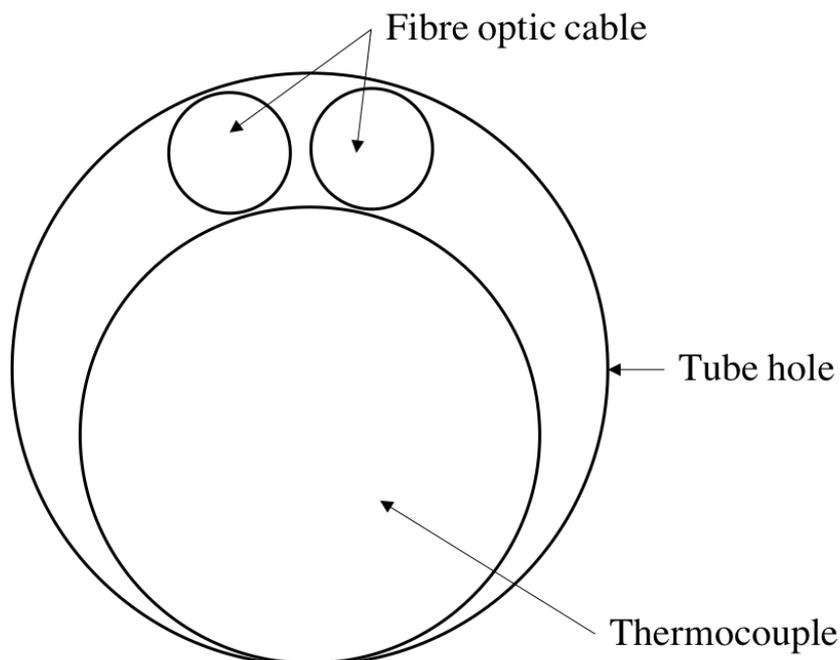


Figure 5.12: Schematic to show cross sectional view of a divider plate tube, showing a thermocouple with two passes of a fibre optic cable. All elements have close proximity to the tube wall.

thermocouple terminates inside the tube rather than being threaded completely through. However, It is not possible to thread fibre optic cables through these incomplete tubes. Instead, they are omitted. This results in a 190 mm ‘blind spot’ at each end of the reactor. The area covered by F/O is therefore only 420 mm. However, it has been demonstrated in previous work that the area of highest thermal activity is well within the measured zone.

The header impingement is not likely to be overcome in commercial-scale reactors, since this aspect of the reactors is retained in scale-up. Future reactor designs could include axial tubes instead of or in addition to transverse tubes. This would overcome the header impingement and resulting blind spots of the F/O. As discussed later, axial tubes may be seen to have additional benefits.

5.3.2 Lab-scale reactor

A second, single channelled reactor was also used in trials. While this reactor did not mimic the commercial scale reactor, it allowed for more flexibility in conducting experiments. (the pilot scale reactor was performing a parametric operations test for the duration of the fibre optic trials)

The lab-scale reactor consisted of a single process channel, encased in a large cylindrical metal block. Internal holes and two axial grooves allowed for the mounting of body thermocouples and in-channel thermocouples in a similar way to the pilot and full-scale reactors. However, the axial grooves did not extend the full length of the reactor. Instead



Figure 5.13: Photograph to show insertion of thermocouple and a single pass of fibre optic cable into divider plate tubes of plate-fin Fischer-Tropsch reactor SP09.

two grooves were separated by a mid point. This meant that it was not possible to mount a F/O axially in the groove.

The reactor block is enveloped by four tight fitting band heaters with independent set points (Figure 5.14). The axial grooves accommodate control and policeman thermocouples for the accurate control of the heaters using PID control software.

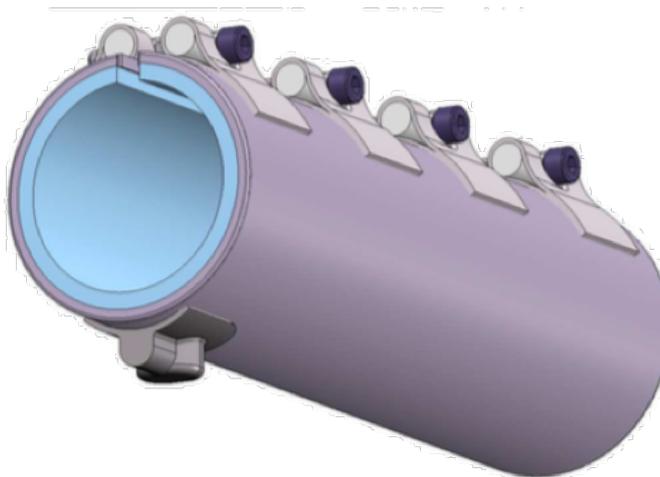


Figure 5.14: Diagram of one of four band heaters used to encase lab scale FT reactors

Since axial mounting in the groove was not possible, two alternative mounting options were explored- External helical winding and external axial mounting. F/O was attached directly to the external metal casing of the band heaters, first with an upwards cylindrical

helix and then with a downwards axial length of fibre.

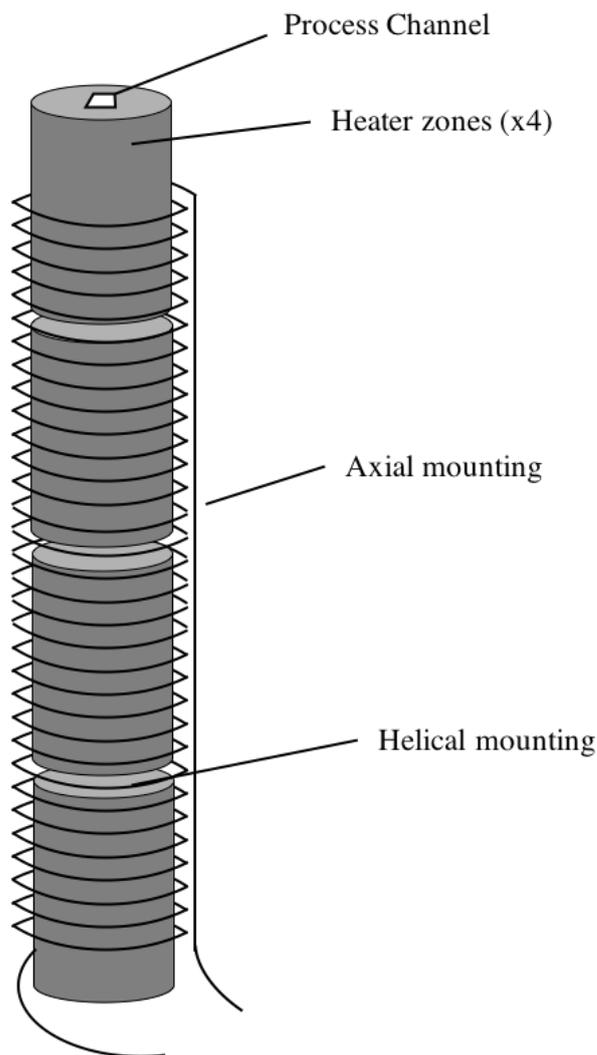


Figure 5.15: Schematic to show helical and axial mounting of a DTS element to a lab scale FT reactor

5.3.3 Software & Hardware Setup

Installation of the fibre optic DTS element and also the DTS equipment was undertaken with guidance from Silixa Ltd. Silixa hardware and software was used to process the DTS signal. This portion of the work, before on-line experimental readings could be gathered, was vital in ensuring that meaningful data could be collected. Through guided trial and improvement, a workable implementation was achieved. In this section, several of the key considerations of fibre hardware installation and software configuration are discussed, with particular emphasis on those issues relevant to implementation of a commercial scale system.

Table 5.2: Silixa Ultima range quoted maximum performance

	ULTIMA-S	ULTIMA-M	ULTIMA-L	ULTIMA-XT
Range (km)	5	10	35	10
Spatial resolution (mm)	125	250	1000	250
Temperature resolution (\pm K)	0.01	0.01	0.05	0.01
Minimum data frequency (s)	1	1	1	5
Fibre mode (Multi/Single)	Multi	Multi	Multi/Single	Multi

Hardware Specifics

Choice of DTS hardware in this pilot scale study was driven by trying to achieve the highest spatial resolution possible, while also achieving a suitable level of temporal and thermal resolution. A secondary criterion was the ability to attain suitable resolution across extended lengths of fibre. Conventional applications of DTS typically require spatial resolutions in the order of several meters, thermal resolution of $\pm 1^\circ\text{C}$ at a frequency of several seconds. Thus, it is possible to achieve the required precision over many kilometres. However, in the current application, much finer resolution is required, at the compromise of distance.

Silixa claim to have the worlds finest DTS system in terms of spatial resolution. Table 5.2 shows the quoted maximum performance figures associated with the four models of DTS equipment available from Silixa. It should be noted that “maximum” performance is given by the minimum value of resolution. Therefore, the Ultima-S hardware offers the maximum resolution with little compromise in terms of temperature and sampling frequency. It is important to note that the quoted maximum performances for each criterion are not necessarily achievable simultaneously. This performance ‘trade-off’ will be discussed later.

The Silixa Ultima-S DTS unit (Figure 5.16) was installed within the pilot reactor control system. The unit comprises a PC with graphical user interface operating system (Windows XP, in this case), networking ports to allow for integration with the plant control system (RJ45, Category 5), and the inbuilt optical system. The optical system is housed within the self contained unit and was not directly inspected during the experiments. However, Figure 5.17 shows a general schematic for the optical hardware components of the DTS system.

Splicing

Splicing is a key step in creating a workable optic system and involves the joining of optical fibres together. The method preferred by industry is known as fusion splicing, where electrodes melt two surfaces together to form a perfect seam. In this work, splicing was unconventional, since the portion of fibre forming the distributed temperature sensor



Figure 5.16: The Silixa Ultima-S DTS unit used in this study

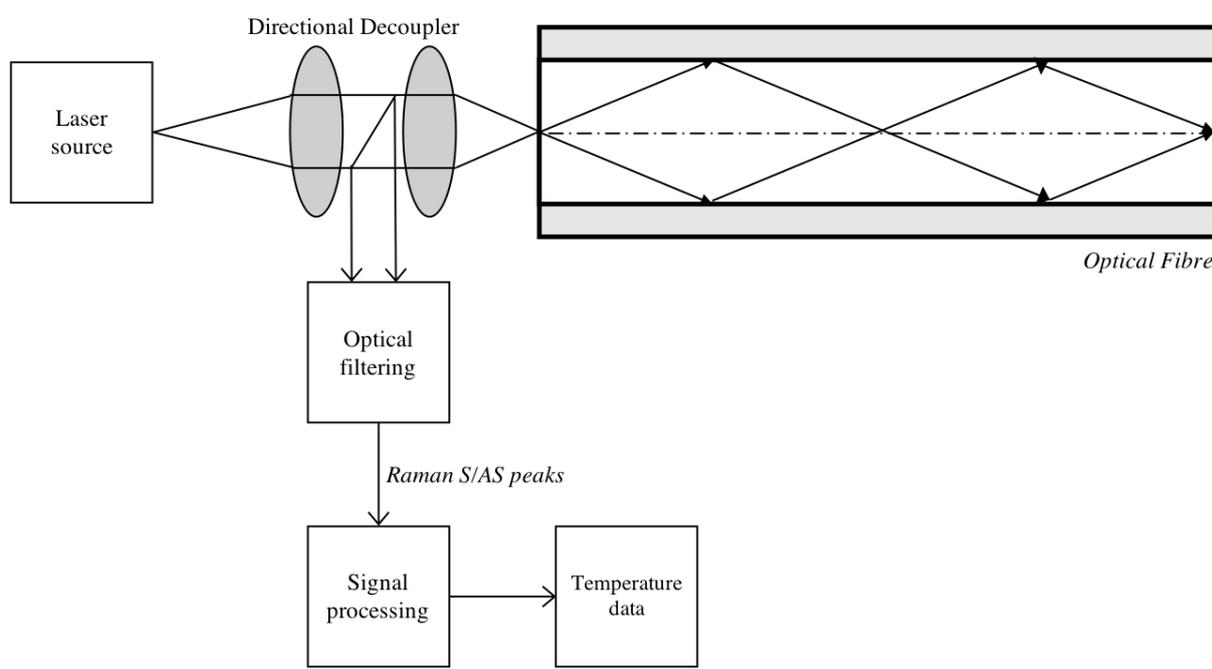


Figure 5.17: Schematic flow chart of Silixa Ultima-S DTS components

was naked fibre – with no buffer or jacket present. This naked fibre was to be spliced to a longer, jacketed fibre referred to as the surface cable. The surface cable’s optical properties are less important, save to say that they match closely to the DTS element cable and provide little in the way of optical losses. The function of the surface cable is to link the Ultima DTS equipment with the naked fibre inside the reactor. For this purpose, approximately 25 m of surface cable was required. Since dual core cable was used, it was necessary to run only one length of surface cable to to close the loop with both ends of the DTS element.

The splicing procedure itself is not strictly affected by the naked fibre, since fibre is stripped of its coatings down to the bare cladding before splicing occurs. The only difference in this case was that the fibre was much more delicate and fragile. Failure at this stage would result in having to re-load the fibre into a reactor. Failure at a commercial installation

might mean that several reactors in series would have to be re-loaded with fibre. For this reason, it is recommended that the naked fibre should be kept to a minimum length and surface cable should be spliced in a junction box at each opportunity. In a commercial scale reactor, this might mean that each reactor in series is linked by surface cable, or it may even be preferable to split the DTS element within the reactor, and have several junction boxes linking the fibre on a single reactor, for example: once per layer, or once per pass. This adaptation would further enhance the reliability of the DTS system, since any failures could be quickly identified and located, and that portion of fibre could be repaired without affecting neighbouring reactors or indeed working portions within the same reactor. This adaptation would require further testing, since it is unknown how the presence of such frequent splicing would affect the signal attenuation.

The splicing procedure is briefly outlined below and represented schematically in figure 5.18:

1. Strip fibre down to the clad material ($\varnothing 125 \mu\text{m}$) using a fibre optic stripping tool
2. Create 'window' in polyamide coating using heat treatment tool
3. Clean the fibre in the area to be spliced with isopropyl alcohol pads
4. Fit a splice protector sheath to the fibre, for use later
5. Cleave the fibre using the fibre optic cleaving tool – this tool guides a perfect 90° transverse cut across the fibre
6. Gently load the two ends to be spliced into the V-shaped grooves in the fusion splicer taking care not to touch the fibre with any surface
7. Tension the fibres to allow for a degree of slack
8. Inspect the viewfinder on the splice machine to verify alignment
9. Run automatic splice program to introduce tungsten electrodes and heat- weld the fibre
10. Slide the splice protector sheath over the spliced area
11. Splice integrity can be tested with optical time domain reflectometer (OTDR)

I/O Channel configuration

The Silixa Ultima-S DTS unit is equipped with up to eight optical I/O ports (the unit supplied for this study was limited to four channels). This means that up to eight fibre cores can be fitted. However, since there is only one light source, the light pulsing cycles between all active channels sequentially, reducing the effective sample frequency by a

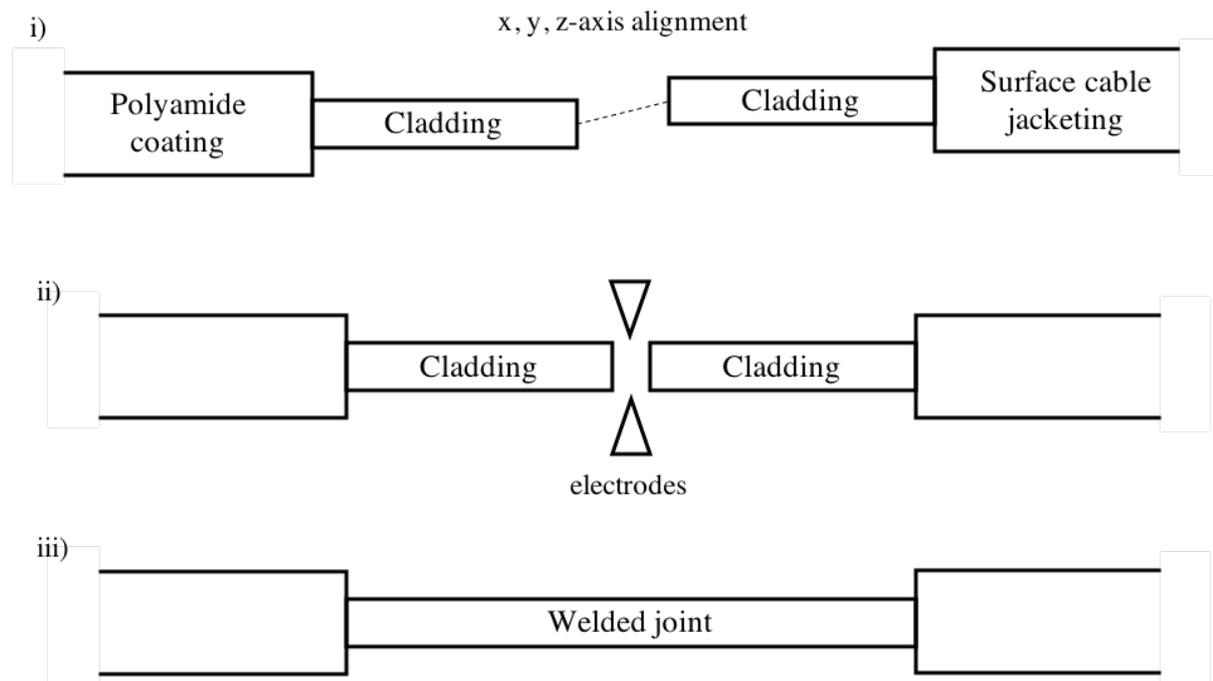


Figure 5.18: Schematic to show fibre optic fusion splicing procedure

factor of the number of channels. In a commercial installation, several channels could be employed to serve several groups of reactors, rather than having a single channel serve each group in series. This adaptation would reduce the length of fibre required in a single run and thus increase the robustness of the system.

However, since the goal of this study is to push the boundaries of resolution, a further adaptation was installed. Since the incident light and backscattered signal can be easily differentiated by filtering, it is possible to send pulses in both directions down the fibre. This opens the possibility of a single run of fibre occupying two channels, with pulses rapidly alternating between each channel (given the limitation that only one channel can be active at a time). Now, with two channels linking a single run, it is possible to have either a single or double pass of the reactor as the fibre makes a loop between the channel I/O ports. Figure 5.19 details the possible paths of fibre and configurations of channels for the pilot scale system. This set of configurations is easily scalable to incorporate several reactors in a commercial installation. The main characteristics of each configuration are outlined below, with reference to Figure 5.19:

- **a) Single channel, single pass** This is the simplest configuration, where a single channel sends incident light pulses in one direction through the fibre. A single distributed temperature sensor is positioned inside the reactor for a single measurement. This data requires the least interpretation since time of flight is directly related to the position in the reactor.
- **b) Single channel, dual pass** – More complex in terms of installation, this configuration allows for the DTS element to be looped back through the reactor for

a second reverse pass. (a second forward pass is also possible with some adjustment) Light is still sent in a single direction through the fibre, but two sets of temperature data are received in mirror image. This effectively increases the sampling frequency and spatial resolution, since more information is gathered per unit time, and mirrored data can be averaged to increase the accuracy of data.

- **c) Dual channel, single pass** – A fibre is not left open ended in this configuration. Instead, a loop is made and the fibre is terminated at another I/O port. This allows for light to be sent in two directions through the fibre. In contrast to single channel configurations, a full set of data can be gathered even if there is one break in the fibre at any length. Data can be sent back in different directions from either side of the break. This increases the robustness of the system. In another adaptation, the return position of the fibre may be mounted externally and axially on the reactor surface, giving further information about the state of the reactor.
- **d) Dual channel, dual pass** – In a combination of the two previous configurations, both the robustness and increased virtual sampling rate are achieved. In this configuration, mirror image data is gathered from two directions. The data can be combined in software to achieve a less complicated output. This arrangement would also allow temperature detection and reactor control to occur in the event of a breakage (or potentially two breakages, depending on location).

For the experimental portion of this work, a pilot scale plate-fin reactor was fitted with the dual channel, dual pass configuration shown in Figure 5.19(d). A single channel FT reactor was fitted with a dual channel, single pass configuration as per Figure 5.19(c), with the return pass being mounted axially to give an axial temperature profile.

Settings trade-off

With the fibre hardware installed, the software settings must be configured. Importantly, the quoted maximum performance figures for the DTS unit are not achievable simultaneously. The performance parameters under scrutiny in this work are as follows:

- **Length, L (km)**, – the maximum length over which a resolvable signal can be sustained.
- **Spatial resolution, Δz (mm)**, – the distance over which two distinct real temperatures can be resolved by the system
- **Sampling resolution (mm)** – the distance over which the software averages a set of spectra to give an average figure for that portion. (Related to the spatial resolution.)

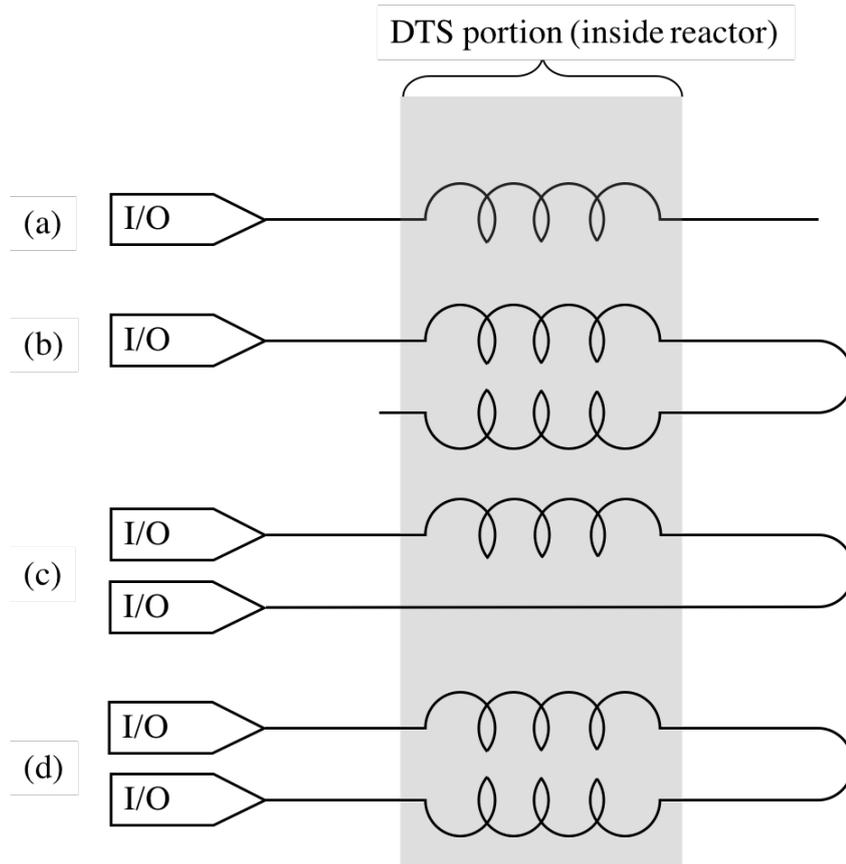


Figure 5.19: Various channel configurations possible with the DTS equipment

- **Measurement time, t (s)**, – the amount of time over which successive spectra for a given portion of fibre are averaged, to give a representative data point.
- **Thermal Resolution, ΔT ($^{\circ}\text{C}$)**, – the temperature difference required to practically differentiate between two distinct real temperatures

The relationship between each parameter is given by F_m , the so-called 'figure of merit', which describes the trade-off of parameters:

$$F_m = \frac{L}{\Delta z^2 \Delta T^2 t} \quad (5.21)$$

Given the physical limitations of the system, and the software limits in the Ultima-S settings, the following temperature resolutions are achievable under various conditions. (Table 5.3).

Since the best spatial resolution possible is 125 mm, this was held as fixed for the duration of the experiment, since enhanced spatial resolution is a key criterion for a commercial implementation. Table 5.3 shows that, for fixed optimum spatial resolution, temperature resolution is adversely affected by both an increase in fibre length and a decrease in sampling time. For temperature control of the reactors, the modular FT reactors require temperature readings as accurate as $\pm \sim 0.1$ $^{\circ}\text{C}$ at 1 s frequency. However, the goal of the

Table 5.3: Temperature resolution (\pm °C) for various lengths of fibre at various sampling times achieved with the Silixa Ultima-S DTS equipment used in this study (spatial resolution fixed at 125 mm)

Sampling time (s)	Temperature resolution (\pm °C)		
	0.5 km	1 km	1.8 km
1	0.31	0.37	0.50
10	0.12	0.13	0.19
60	0.05	0.06	0.07

distributed sensor is not to control the reactor under normal operation – This task would be left to highly accurate reference thermocouples (empirical models developed in this work have demonstrated the ability to accurately predict whole reactor temperature profiles based on a reduced set of parameters, such as inlet temperature and flow conditions). The commercial distributed temperature sensor would aim to provide a layer of protection against thermal events causing damage to the catalyst or indeed the reactor itself. For this reason, temperature resolution below $\pm \sim 0.5$ °C would provide adequate protection. Table 5.3 shows that even at the least optimal conditions within the scope of the study (1 s sampling time and 1.8 km fibre length), adequate temperature resolution was achieved.

Though unlikely, it may be possible for commercial implementations to require fibre lengths in excess of 1.8 km. At a maximum length of 5 km, temperature resolution below $\pm \sim 1.5$ °C is still achievable for a 1 second sampling interval. For sampling times of 10 seconds or above, an acceptable level of temperature resolution is maintained at 5 km. This relationship is shown in Figure 5.20.

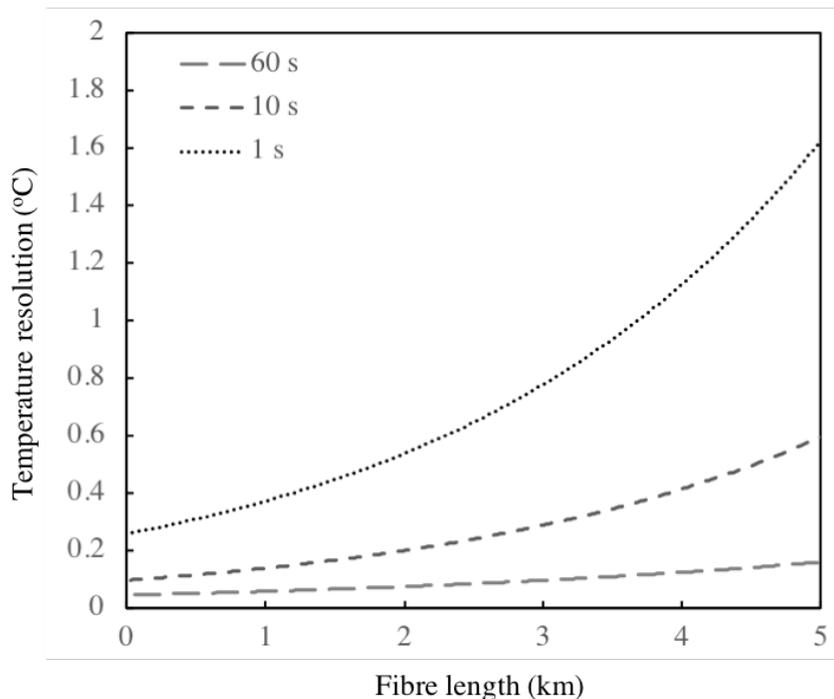


Figure 5.20: Graph to show achievable temperature resolutions at various sampling times as a function of fibre length, for fixed spatial resolution of 125 mm.

5.4 Experimental Methods

5.4.1 Fischer-Tropsch Synthesis and Reactor Operation

FTS was carried out at CompactGTL’s Wilton pilot plant facility in prototype plate and fin fixed bed reactors with particulate catalyst (20% $Co(Ru)$ on Al_2O_3 support.) (Johnson Matthey, $\varnothing \sim 300 \mu m$). In this study, two reactors are discussed; a 50-channel next generation pilot reactor, SP10) and single-channel lab scale reactor, RX4) (note that length and cross-sectional dimensions of individual channels is preserved across scales, and thus represent commercial conditions).

SP10 Run

The cooling fluid is held at an appropriate temperature and pressure by means of an external heating and cooling loop. Both heating and cooling of the process side are supported depending on the mode of operation. The coolant flow rate is controlled to maintain a ΔT of $\sim 5 \text{ }^\circ C$ across the reactor (this represents a change in operating philosophy from previously reported work, owing to further developments in the process, undisclosed in this work).

Mass flow controllers and pressure regulators along with gas pre-heaters control the process inlet gas conditions. Syngas is typically has a H_2/CO ratio of ~ 1.9 , though the effect of

varying composition is a controlled variable in some portions of the run. Process product is decanted in staged separation vessels to allow the segregation of water, oil, wax and tail gas. The composition of tail gas is routinely analysed by gas chromatography (GC) to close the mass balance and allow for the calculation of the FTS key performance indicators; productivity, selectivity, conversion etc. Offline physical and chemical analysis of the liquid and solid products is undertaken at time points throughout the run, in order to characterise the changing product profile during the catalyst lifetime.

A so-called operational ‘steady state’ is maintained by fixing conversion to a target figure. Changes in catalyst activity owing to deactivation and other deleterious changes in process conditions (e.g. changing syngas ratio, reaction pressure, inert partial pressure) are counterbalanced by increasing the applied temperature. (the temperature of the gas at the inlet, T_{app}). Thus, differential performance is often measured in terms of ‘loss of T’, or ΔT_{app} ; the amount by which the applied temperature must be increased to maintain conversion at a target notwithstanding the changes made to process conditions.

The pilot scale experimental work was carried out in parallel to another parametric study focussed on reaching stable operation at high conversion. The DTS experiment discussed in this report was merely able to observe changes in temperature, rather than recommend and enact changes in process conditions to observe the effects. In this regard the pilot scale DTS study was limited by the scope of the parametric study. The conditions observed by the DTS equipment were representative of a full campaign; start-up, ramping, steady operation, thermal events and shutdown.

RX4 Run

Since it wasn’t possible to alter any conditions in the SP10 run, but only to observe, the RX4 reactor was set up to allow thermal events to be simulated. No reactants were introduced into the reactor. Instead, the external band heaters were controlled through the distributed control system (DCS). Heaters were subjected to a range of set point changes in various sequences and rates, designed to simulate extreme cases of operational disturbances (start up, shut down, trips, thermal runaways).

5.4.2 Data Collection

Process data is collected every second and relayed to a supervisory control and data acquisition system (SCADA) for storage and operator manipulation. Aspen Process Explorer (AspenTech, 2007) was used to access the data from the server and view or export elsewhere.

DTS data was acquired using the ULTIMA Data Acquisition software (Silixa, 2014).

Optical signals processed by the instrument are time-averaged over the specified window (t) to give a trace of peak intensity over fibre length for both the stokes and anti - stokes Raman backscattered signal. This raw optical signal data is processed (by taking the ratio of stokes and anti-stokes signal intensity) to give a temperature value as a function of length for each time interval, t . Each time point yields an XML file containing the configuration parameters along with the stokes, anti-stokes and temperature vectors stored as comma separated variables (CSV). The XML file path can be loaded into the ULTIMA DTS Viewer software (Silixa, 2014) where the data is stitched together to allow playback of thermal events in real or accelerated time. Alternatively, the CSV files can be exported for analysis elsewhere.

5.5 Results and Discussion

5.5.1 The Scale and Shape of Data

Data size Challenge

The first challenge that implementation of a full scale DTS system poses is the storage and processing of the large amount of data that is generated. Indeed, during the pilot scale trials, two weeks of acquisition yielded several thousand files totalling > 50 GB of data.

In order to reduce the scale of the data, several adaptations to the system could be made. The most obvious solution is to reduce the sampling frequency. This would have the added benefit of increasing temperature resolution, as discussed previously. Another proposed solution explored experimentally is the implementation of boxcar or exception-based sampling, where a new data point would only be stored if it varied sufficiently from a previous data point. While the data size challenge is an important one, it's solution is perhaps outside the scope of this report – save to say that a commercial scale plant should be better equipped to store and process such large amounts of data.

Raw Data Processing

With masses of raw data being stored as separate files per time point, the first challenge is to understand the structure of a raw file and automate the extraction of relevant information. Raw files are XML documents, with object classes to define each configuration parameter. The raw measurement data is stored in a single object class as comma separated matrix, with a row vector, stored in a separate class, allocating structure to the matrix.

Extraction and delimitation of the CSV matrix and pairing with the transposed structure vector gives a common sense view of the raw data, at the expense of much greater file size.

Figure 5.21 shows graphically the extracted raw data for a single time point. This data was collected during steady state operation in the SP10 reactor, averaged over 5 seconds with a spatial resolution of 125 mm.

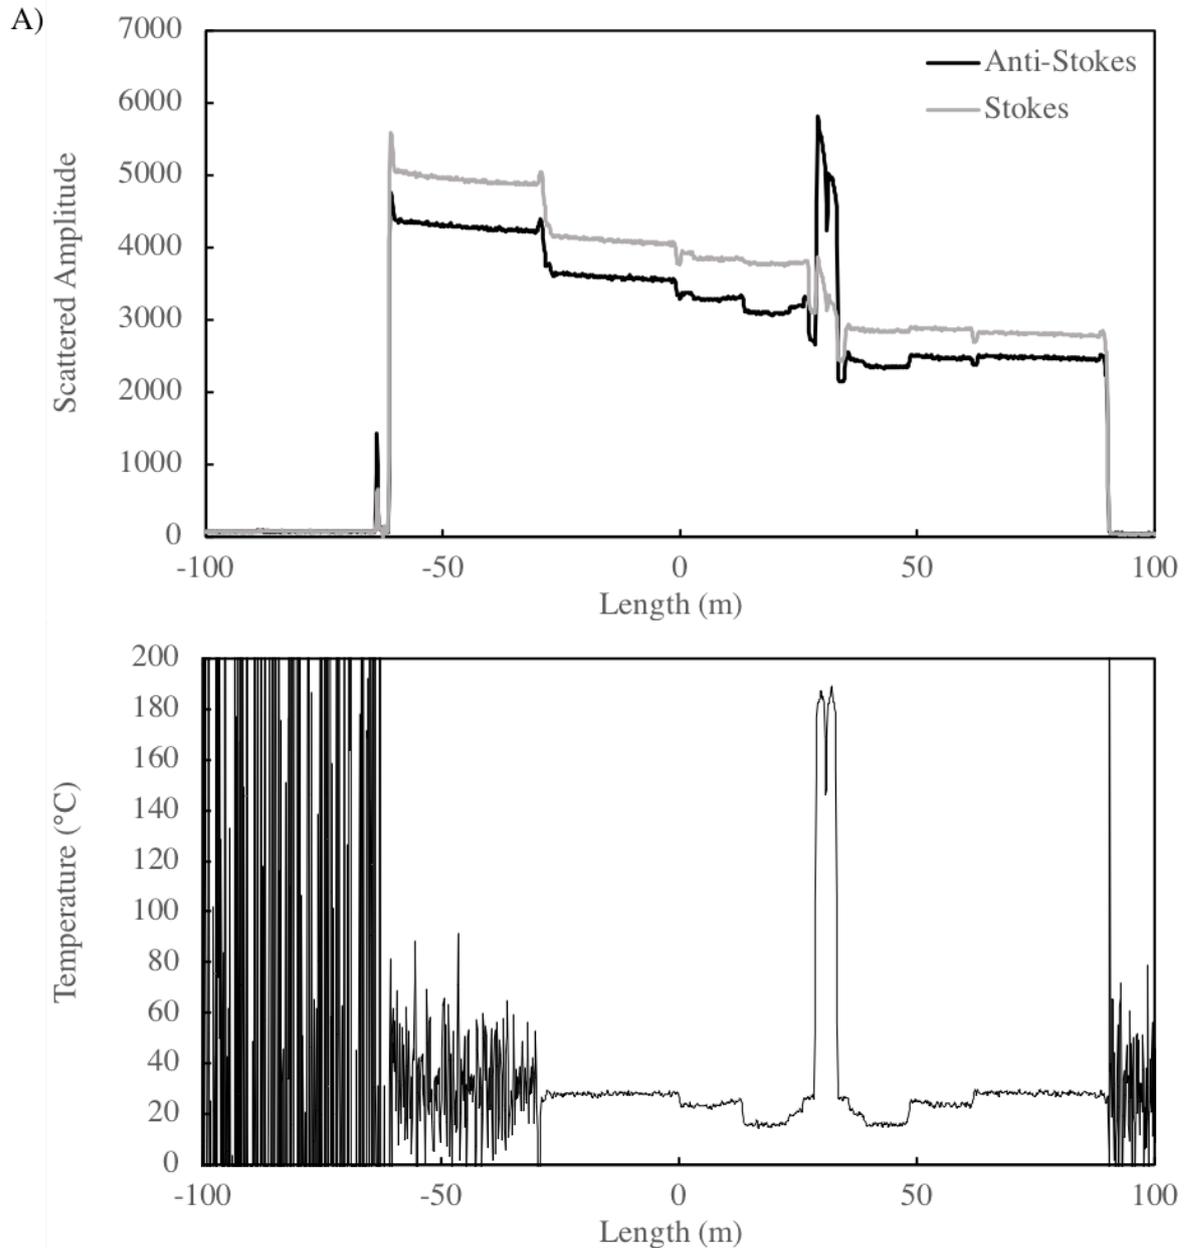


Figure 5.21: Graph to show (A) raw Stokes and Anti Stokes backscattered signal intensity as a function of length along the fibre optic cable and (B) signal processed temperature data for the same length. Data is from a single time point in in a SP10 run. (Sampling time – 5 s, Spatial resolution, 125 mm)

The x-axis in Figure 5.21, length (m), extends to ± 100 m. Storage of this data is superfluous, since 0 m marks the point at which the light exits the source emitter. The negative values are an artefact of the double ended data – Light is also transmitted in the opposite direction through the fibre. The ‘real’ data, excluding the end effects starts at 0 m and extends to a length of approximately 65 m. This includes 2 lengths of 25 m surface cable and the DTS portion inside the reactor. In a commercial implementation,

the data storage interval should be adjusted in the Ultima software to pre-cut the data to only include ‘real’ values from the fibre itself.

From Figure 5.21(a) it is also possible to appreciate the relationship between the backscattered signal and temperature. In heated regions, the Anti Stokes signal has a much stronger response. Both the Stokes and Anti Stokes signal intensities are diminished as a function of length, resulting in a gradual slope. The signals are also diminished as a strong function of temperature, giving rise to a much greater attenuation slope in high temperature regions. This seems to indicate that there would be a practical maximum exposure length for a DTS element, which would be a function of temperature. It is also important to note that attenuation is increased by bending of the fibre. The portion of the DTS element exposed to the high temperatures of the reactor is also the portion which takes the most torturous path, increasing the attenuation effect at this point. This limitation of temperature and bend induced attenuation would be overcome by changing fibre properties or increasing the incident laser power.

The attenuation of signal occurs in the opposite direction for the second channel signal in a double ended system. Thus, where one signal is most attenuated, the other is at its strongest. Averaging of the two signals allows for much greater lengths to be reached. The weakest signal from the averaged data therefore comes at the mid-point of the fibre loop. In this experiment, the mid point was the end loop of the DTS element. In a future adaptation, asymmetrical surface cable lengths could be used for the outward and inward surface cable runs. This would ensure that the central, attenuated signal does not coincide with the DTS region of interest.

After trimming of end effects, The DTS element of the system can be examined more closely. Figure 5.22 shows trimmed data (source data shown in Figure 5.21) highlighting only the DTS element.

Analysis of Figure 5.22 allows us to begin to locate key regions within the DTS element as a function of length along the fibre. Splices are often revealed by a sudden step change in Stokes and Anti Stokes intensity that is not accompanied by a change in temperature. The added optical resistance of the splice boundary affects the Stokes and Anti Stokes signals equally, and no change in the ratio is produced. In this way, the quality of a fusion splice can be assessed. If a splice is ‘bad’, an unacceptable step change in the backscattered signal intensity will be observed. It is recommended that pre-operational experiments are carried out when commissioning a commercial scale DTS installation to identify the location and quality of fusion splices and other optical defects (bends, kinks, compression etc. are also detectable in this way).

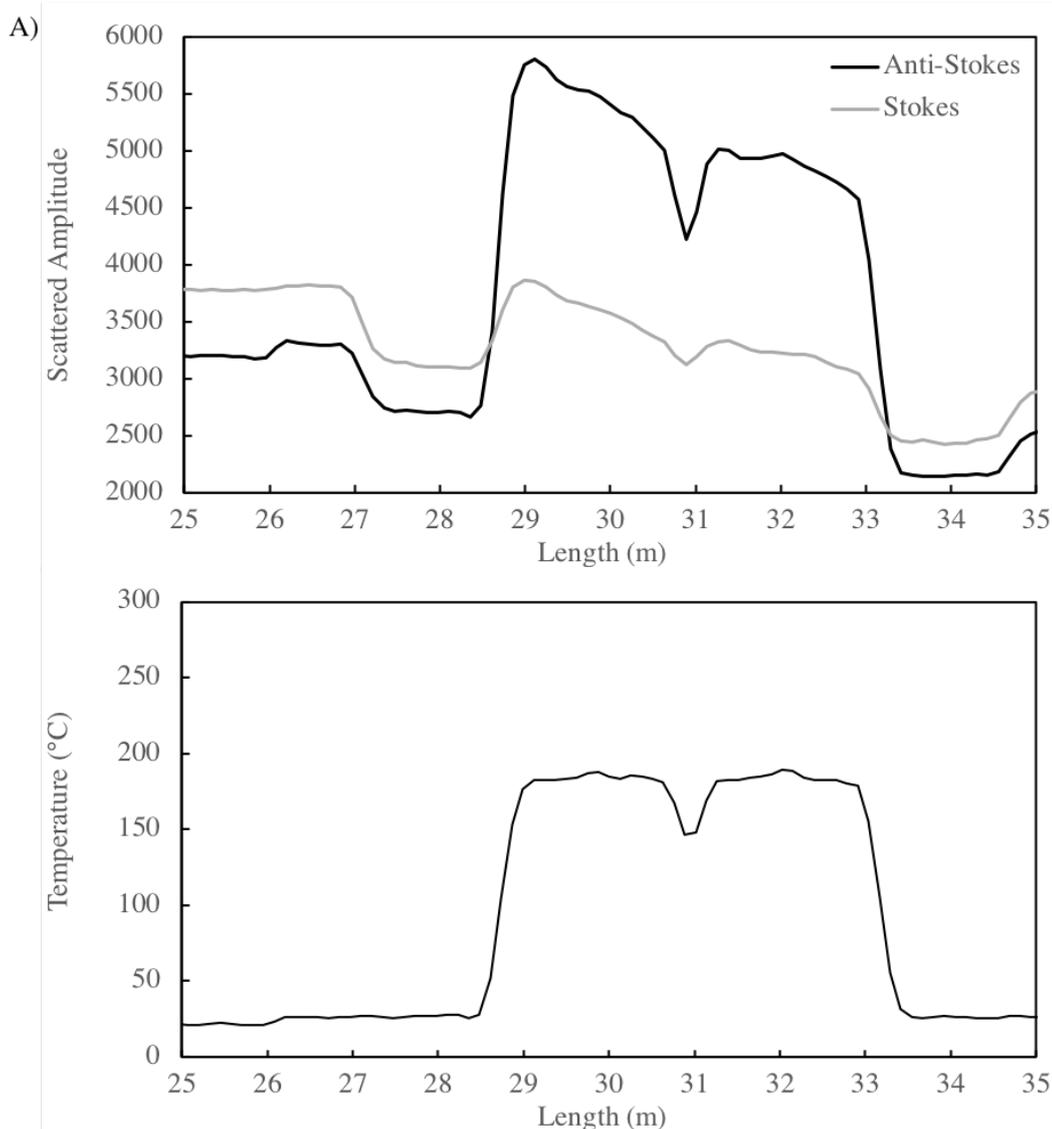


Figure 5.22: Graph to show (a) trimmed Stokes and Anti Stokes signal and (b) Temperature as a function of length along the fibre.

5.5.2 Relating Linear Temperature to Real Space

Measurement of the temperate regions along the fibre shows that approximately 1.64 m of fibre is threaded into each pass of the reactor, whose total length is considerably less than this. 420 mm of vertical reactor length is covered by the DTS element. Tortuous fibre paths are therefore shown to ostensibly increase the spatial resolution. A further and more extreme example of this is described later when considering the lab scale reactor. Relating the linear temperature data to ‘real’ space within the reactor is critical in extracting useful information from the DTS system.

One strategy employed in this work was the use of a cold air spray canister to spot-cool sections of the fibre. The temperature drop could be viewed in real time on the DTS software, and known locations in real space could be related to points in the linear trace. Figure 5.24 shows key features of the fibre optic path within the reactor. Spot cooling of

Table 5.4: Sections of fibre joined by numbered nodes and their corresponding plate and vertical direction

Region	Section	Plate	Direction
A	1 - 2	front	up
B	2 - 3	rear	down
C	3 - 4	rear	down
D	4 - 5	front	down

some of these points allowed for their location to be marked on the linear temperature trace. Cross referencing the temperature data with known locations in the reactor geometry means that a more intuitive understanding of the temperature data is possible.

Figure 5.23 shows the characteristic shape of temperature data within the DTS element during steady state operation. The linear profile can be explained in terms of real reactor geometry. Spot-cooling was used to verify the location of these points.

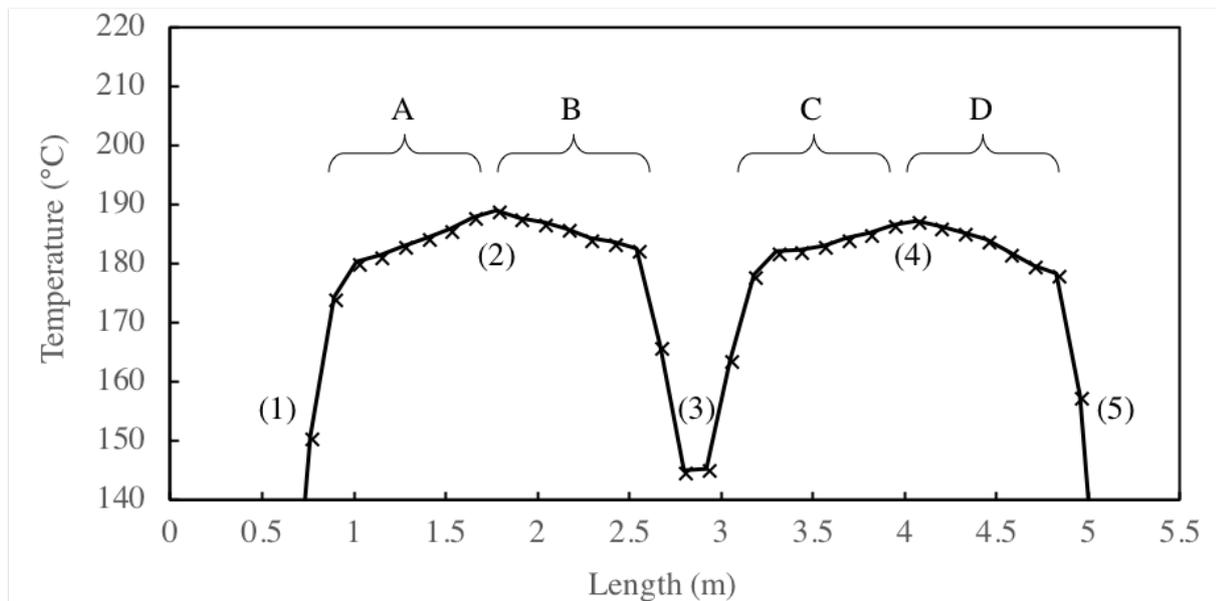


Figure 5.23: Graph to show fibre temperature profile related to reactor geometry elements. Numbers (1 - 5) refer to labelled geometry in Figure 5.24. Letters (a - d) refer to directional sections in Table 5.4

Together, Figures 5.23 & 5.23 illustrate one of the key subtleties of the DTS data. Common sense reactor temperature profiles are plotted from inlet to outlet. In the case of the DTS trace, two factors prevent this: The DTS element begins and terminates at the outlet side of the reactor (Figure 5.23 [1 & 5]) and the fibre optic passes travel both upwards and downwards with respect to the reactor. A further subtlety of the linear data is the fibre location within the two plates as well as vertical position.

Paying close attention to Figures 5.23 & 5.23, Table 5.4 resolves the spatial complexity of each main fibre section.

In order to take advantage of the repeat sampling offered by the looping of fibre, spatially

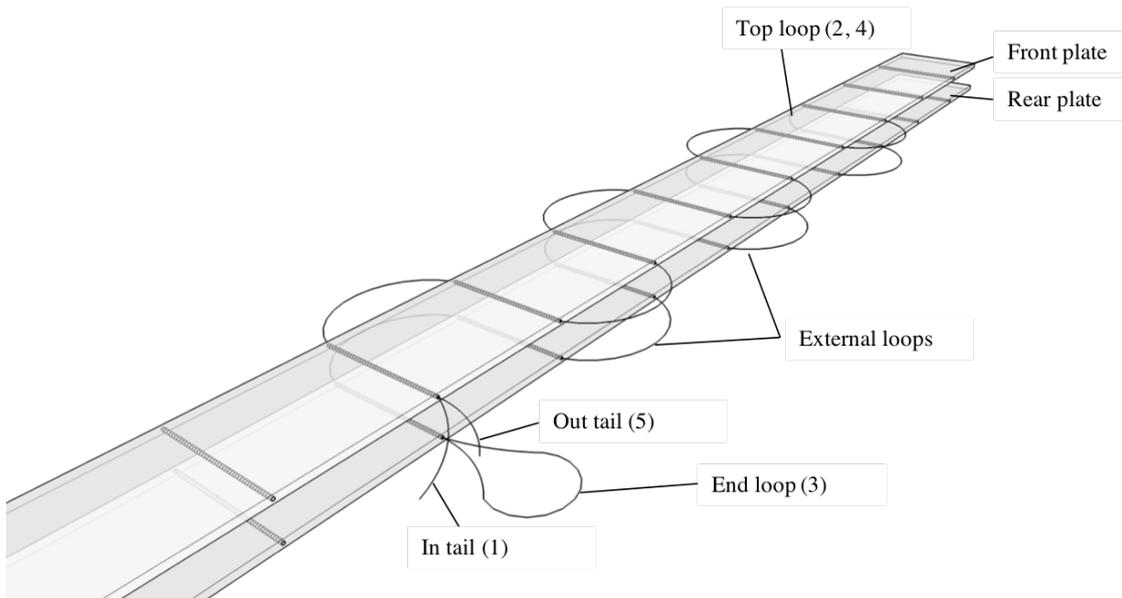


Figure 5.24: Isometric cutaway to show SP10 internal divider plate structure highlighting key features of fibre optic path; (1) In tail – where the fibre enters the reactor, (2,4) Top loop – the highest vertical position of the fibre, where fibre spans between tube plates, (3) End loop – fibre exits the reactor and is returned in the opposite direction through the same path. (5) Out tail – the fibre exits the reactor and is fusion spliced to return surface cable.

Table 5.5: Data regions to be averaged to give inlet-outlet orientated reactor temperature profiles (letters with reference to 5.4, subscript ' r ' denotes the data region is flipped horizontally)

Profile	Averaging
Front Plate	$A_r + D$
Rear Plate	$B + C_r$
Whole reactor	$A_r + B + C_r + D$

corresponding data can be averaged to give a more stable data point. Data can also be transposed in such a way that all data, after averaging is oriented from reactor inlet to reactor outlet. Table 5.5 shows the spatial averaging operations with reference to data matrices from data labelled in Table 5.5.

Flipping aligns the up and down directions of the data with respect to the reactor. In this way, averages can be taken to represent both divider plates individually, or the reactor as a whole. Figure 5.25 shows a plot of temperature data against reactor position, relative to the inlet of the reactor. The characteristic symmetry of the raw temperature data has been processed out and the data points are now more representative of real reactor space. Figure 5.25 demonstrates the temperature smoothing made possible by taking 'plate' and 'whole reactor' averages. The plate profiles in the pilot scale reactor are not expected to differ significantly, due to their proximity. However, on a commercial scale, comparison of several plate profiles could provide a further layer of detail to the temperature data. A plate whose temperature strays from the pack may be showing early signs of a thermal event. Control action, such as an increase in coolant flow rate, could be triggered by a

plate temperature profile deviating by a predetermined amount. Furthermore, once a rate of change in temperature is detected in a certain location, the DTS settings can be tuned to provide more detailed data, subject to requirements (e.g. by increasing temperature resolution during startup ramps or sampling frequency during thermal excursions.)

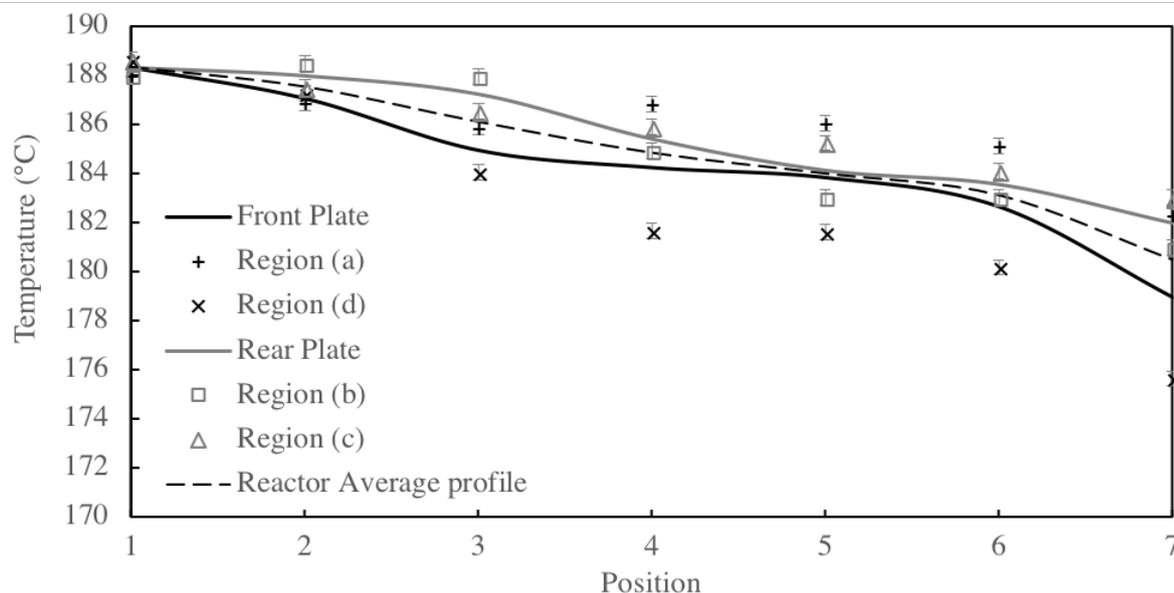


Figure 5.25: Reactor temperature smoothing achieved as a result of spatial averaging of temperature data

Figure 5.25 also shows the smoothed average reactor temperature profile. This profile is far less susceptible to noise, since 4 data points are averaged to produce each point on the average line. Control action triggered by such an averaged temperature trace would be significantly less susceptible to so called ‘spurious’ triggers – where noise or single data points exceed a threshold amount and triggers a control action unnecessarily. Exception-based sampling and spatial noise cancellation would allow for a high degree of confidence in any control action (e.g reactor shutdown).

The temperature profile shown in Figure 5.25 is characteristic of the steady state operation of a modular compact FT reactor of the type used in this study. A high inlet temperature is a result of high partial pressures of fresh syngas relative to water vapour and other products. This means that the reaction rate is fastest and therefore most exothermic at the inlet.

5.5.3 End looping offsets and the implications on scale-up

The data X-axis of Figure 5.25 (position 1-7) refer to seven sequential portions of 125 mm (the spatial resolution) of the fibre. Figure 5.25 therefore does not resolve temperature to vertical height of the reactor, but still to length of fibre (albeit much more aligned to the reactor geometry). Seven portions of 125 mm fibre totals 875 mm of fibre covering 420

Table 5.6: spatial resolution alignment with looping

Zone start	Zone end	% inside
Tube inlet	loop	49.6
Tube middle	loop	25.6
Tube outlet	tube	12.8
Loop Middle	Loop	49.6

mm of axial reactor length. Necessarily, since the reactor is only 62 mm wide, the 125 mm portions must encompass an element inside the tube, and an element outside, in the external loops (see Figure 5.24). Table 5.6 shows four boundary scenarios for where an averaging zone may fall within the fibre path. Naturally, a measuring zone may lie at any point between these extreme examples.

Table 5.6 shows that any given 125 mm fibre segment does not reside exclusively within the divider plate tubes. Instead, each spatial average is a combination of tube and loop temperature readings. In the most optimistic case, a 125 mm portion of fibre resides for a maximum of 62 mm within the tube. (Table 5.6 calculations are based on the external loops forming a semicircle with 70 mm diameter – the distance between successive tubes).

The coincidence of these dimensions; the spatial resolution ($\Delta z = 125$ mm), the tube length or reactor width ($w = 62$ mm) and the distance between tubes ($\Delta y = 70$ mm) gives rise to the approximate relationship that:

$$w \approx \Delta y < \Delta z \quad (5.22)$$

This means that the majority of the DTS element in the pilot scale reactors are in fact measuring from external portions of the reactor, rather than transverse sections through the tubes. This will inevitably skew the temperature data downwards, since it is certain that the temperature experienced by loops is much less than the temperature inside the reactor. Though insulation is present around the reactor during operation, this is primarily for operator safety and it is known that the temperature at the insulation wall is <70 °C. The temperature experienced by the loops must therefore be between 70°C and the reactor temperature (~ 200 °C).

Comparison to discrete thermocouples measuring single points inside the tubes gives some indication of the degree of offset experienced by the external looping effect. Figure 5.26 compares two point thermocouples with two spatially averaged sections of fibre situated physically closest to the two thermocouples. The time range encompasses a thermal excursion captured by the discrete temperature probes. In contrast, the distributed sensors did not register a significant change during the event, and only showed an appreciable difference when the reactor was cooled following a control action.

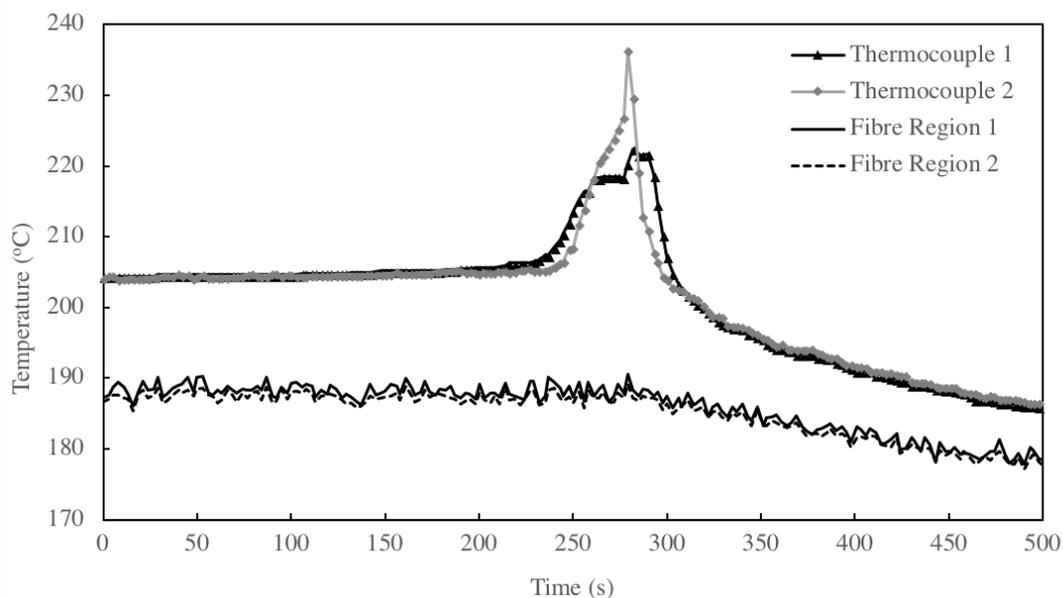


Figure 5.26: Comparison of discrete thermocouples and sections of spatially averaged DTS fibre during an operational thermal event.

Notably, there is a pronounced offset between the discrete and distributed temperature readings. Upon closer inspection of Figure 5.26, the offset is not constant. A plot of the offset for discrete versus distributed sensor averages is shown in Figure 5.27.

Figure 5.27 clearly shows that the offset is a strong function of temperature, which firmly supports the notion that the offset is caused by external loops exposed to colder temperatures than the reactor tubes. The offset is reduced at lower reactor temperatures, since the heat transfer gradient between the reactor and surrounding air is less severe. Figure 5.27 shows three distinct regions. In (a), at normal operation, the temperature offset is consistent, shown by a tight grouping. In region (b), during the thermal excursion, the offset is much more severe. The rate of exothermic temperature generation inside the reactor exceeds the capacity of the surroundings to remove heat. The surrounding temperature remains constant as the reactor rapidly increases in temperature. In region (c), during the post-control action cooling, the offset is lessened, as the heat removal capacity of the surroundings overtakes the heat generation within the reactor.

Simple regression of the offset versus temperature showed no clear descriptive model to be preferred (Table 5.8). The offset is almost certainly explained by several high order interactions in the dynamic system.

However, in an attempt to process this offset, the distributed data can be force-projected onto the discrete data by means of a transfer function. The transfer function parameters can subsequently be applied to all new data to account for the offset. The transfer function parameters are found by minimising the sum of squared offsets for a range of coefficients of differing order, removing those terms which have little impact on the overall fit.

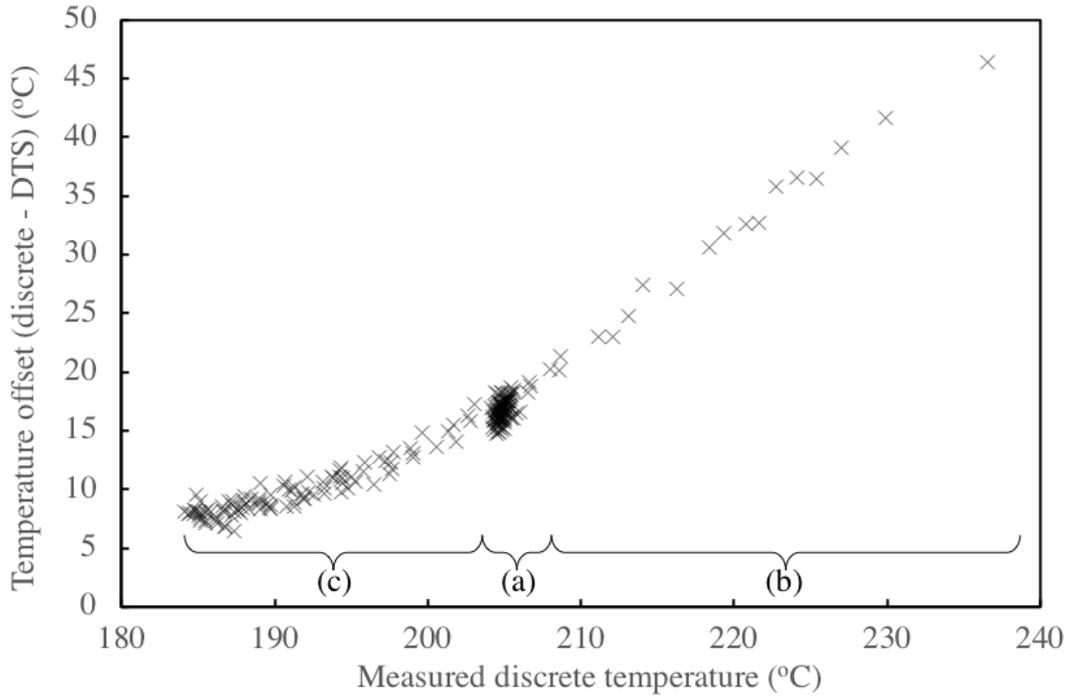


Figure 5.27: Plot of distributed temperature offset versus discrete temperature value for a thermal event

Minimise:

$$\sum_{i=1}^t (d_i - \hat{D}_i)^2 \quad (5.23)$$

Where:

$$\hat{D} = \alpha D^2 + \beta D + D^\gamma + D + \delta \quad (5.24)$$

and t is time, d_i is the discrete temperature at time i . D is the distributed temperature and α , β , γ , & δ , are model parameters. \hat{D} is the scaled distributed temperature accounting for the offset.

The scaled DTS data, shown in Figure 5.28, reduces the offset during steady state and

Table 5.7: Coefficient of determination (R^2) for various regression models applied to offset data shown in Figure 5.27. Good fit was obtained for most models.

Model	Fit (R^2)
logarithmic	0.884
exponential	0.970
power law	0.969
Linear	0.904
Quadratic	0.979
Cubic	0.983
4 th order	0.985
5 th order	0.985
6 th order	0.985

Table 5.8: Optimised parameters to fit scaled distributed temperature data onto discrete data and eliminate the offset

Coefficient	Value
α	5.933×10^{-3}
β	-1.018
γ	-5.777×10^{-2}
δ	-1.240×10^{-2}

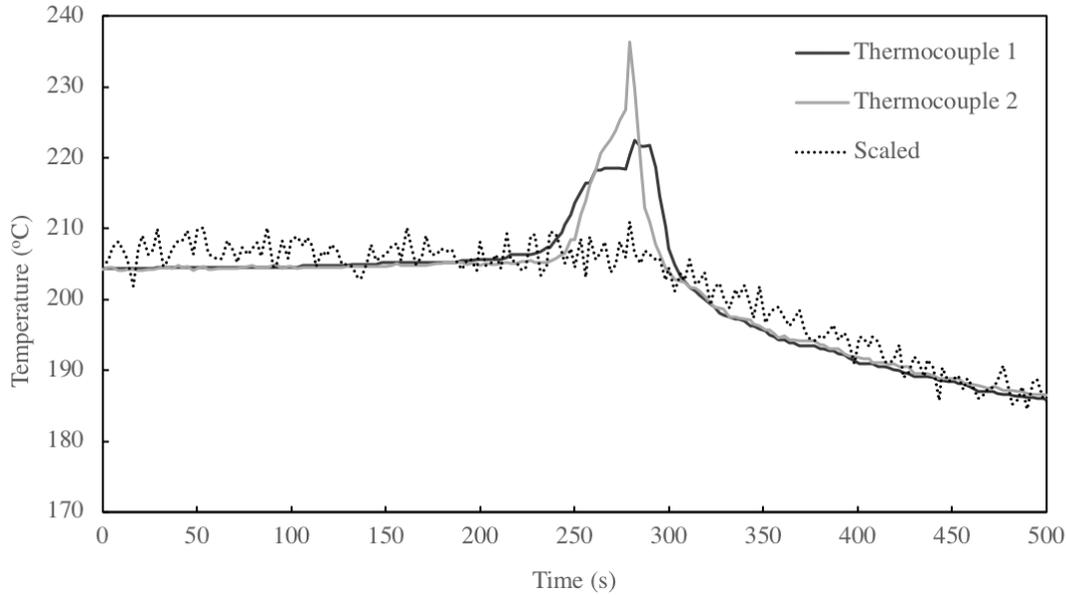


Figure 5.28: scaled distributed data, \hat{D} , compared to discrete data for a thermal excursion. The offset is minimised, but noise is amplified.

cooling, by increasing the gain of the signal. This has the deleterious effect of increasing the signal noise. However, even at high gain, the distributed temperature data does not succeed in capturing the thermal excursion. It seems, at this point, that any noise filtering transfer function would not yield better results. The effects of the external looping are thus concluded to be irrecoverable on the pilot scale reactor.

However, in contrast to pilot scale reactors, a commercial scale reactor increases in dimensions relative to the spatial resolution of the fibre. By significantly increasing the width of the reactor in the direction transverse to flow, and thus the tube length (w), the relationship between spatial resolution (ΔZ) and tube gaps (Δy) now becomes:

$$w \gg \Delta y \approx \Delta z \tag{5.25}$$

Therefore, at commercial scale, the effect of external looping will have a significantly reduced effect on temperature profiles. In commercial scale reactors, several sampling regions, (δz) will reside entirely within a tube, giving un-diluted temperature data in the transverse direction (Figure 5.29).

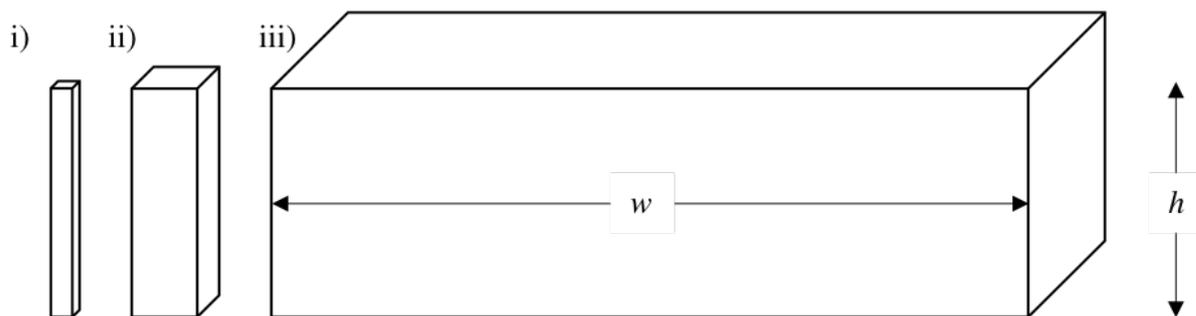


Figure 5.29: Schematic to show relative size of lab, pilot and commercial scale reactors, showing the increase in tube length, w , at commercial scale.

Commercial reactors fitted with a DTS element are therefore in a rare position of having greater 3D temperature data resolution than a pilot scale reactor. This is contrasted with the position at the start of this work, where maintaining a similar level of discrete temperature resolution was unfeasible.

With several data points per tube, several tubes per plate and several plates per reactor, temperature data from a DTS element in a commercial reactor, resolved to real space in a manner discussed in this report would allow for a real-time 3D map of temperature to be constructed. Operators of modular reactors could potentially view a 3D reactor model superimposed with rich temperature data in real time such as is envisaged in Figure 5.30.

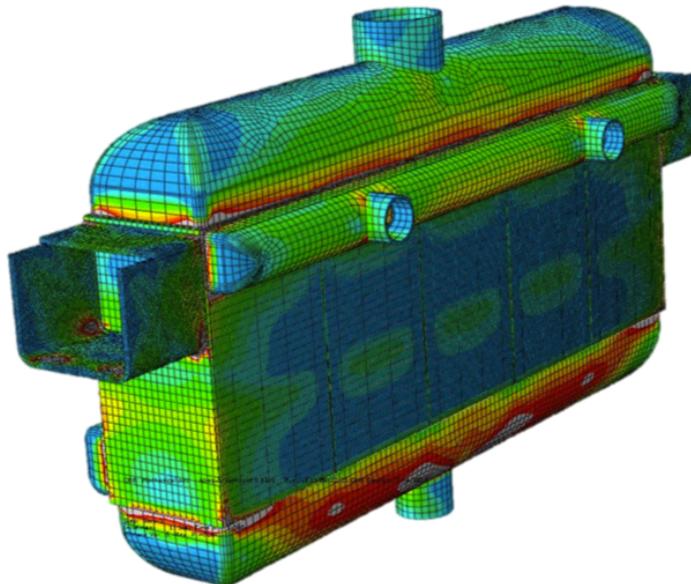


Figure 5.30: Envisaged real time image of temperature data projected onto a commercial reactor geometry

5.6 Augmented Spatial Resolution

In an attempt to further maximise the spatial resolution of the DTS probe, a lab scale reactor was helically wound (figure 5.15). In this manner, several sensing regions (126 mm

spans) of the DTS element cover only a very small axial length of a reactor (a function of its perimeter). Tightly wound fibre optic cable will approximately correspond to a single axial length (in increments of the fibre optic thickness). Thus, a helically wound DTS element can provide dramatically increased spatial resolution compared to thermocouples.

The four heater zones shown in figure 5.15 can be independently controlled. Figure 5.31 shows the response of the DTS element during heater step tests. The four heater zones are clearly identified during the step test, since they were independently controlled to various set points. Each heater zone is controlled by one traditional thermocouple positioned centrally in each zone, and thus the spatial resolution is no more than 200 mm axially, with 4 thermocouples across the 800 mm reactor. Such low resolution results in the assumption of uniform temperature across the entire heater zone. However, the winding applied to the DTS element allows for 10 sensing regions per heating zone. This effectively mimics the data of 40 thermocouples. In practice, the windings could have been made tighter if more sensing regions were desired. Axial mounting of the DTS element provided 5 sensing regions per reactor. Axial mounting may be suitable for much longer pipe runs, where low spatial resolution is required.

The high resolution temperature data from the helically wound DTS elements show some interesting subtleties in the data which are not appreciable at the low spatial resolution offered by traditional thermocouples. With reference to figure 5.31, region (i) shows a subtle oscillation over time for each heater zone. This is expected due to the PID heater controller operating within a dead band. Secondly, there is a subtle curvature to the temperature profile in the length dimension. That is to say that the DTS element provides data about the local temperature non-uniformity across the heater. This information was not possible to explore with a single thermocouple, and difficult to interpolate when more point thermocouples were used. Since the DTS data is effectively continuous, a smooth profile is obtained. Region (ii) shows the peak temperature after an aggressive ramp. Here, it is possible to observe that the non-uniformity effect observed in (i) at low temperatures is exaggerated at high temperatures. Previous assumptions about temperature uniformity across heater zones are shown to be lacking.

With the increased maximum spatial resolution, it is possible to exceed the required nominal spatial resolution. This allows for the DTS element to instead be tuned for enhanced temperature precision or temporal resolution during normal operation, and only when required, the high spatial resolution can be dialled in. Operating in this dynamic manner gives operators the freedom to focus on important events, while maintaining an overview of process areas which are behaving nominally. The added detail provided by high spatiotemporal resolution data allows for characterisation of a process in much greater depth than previously possible.

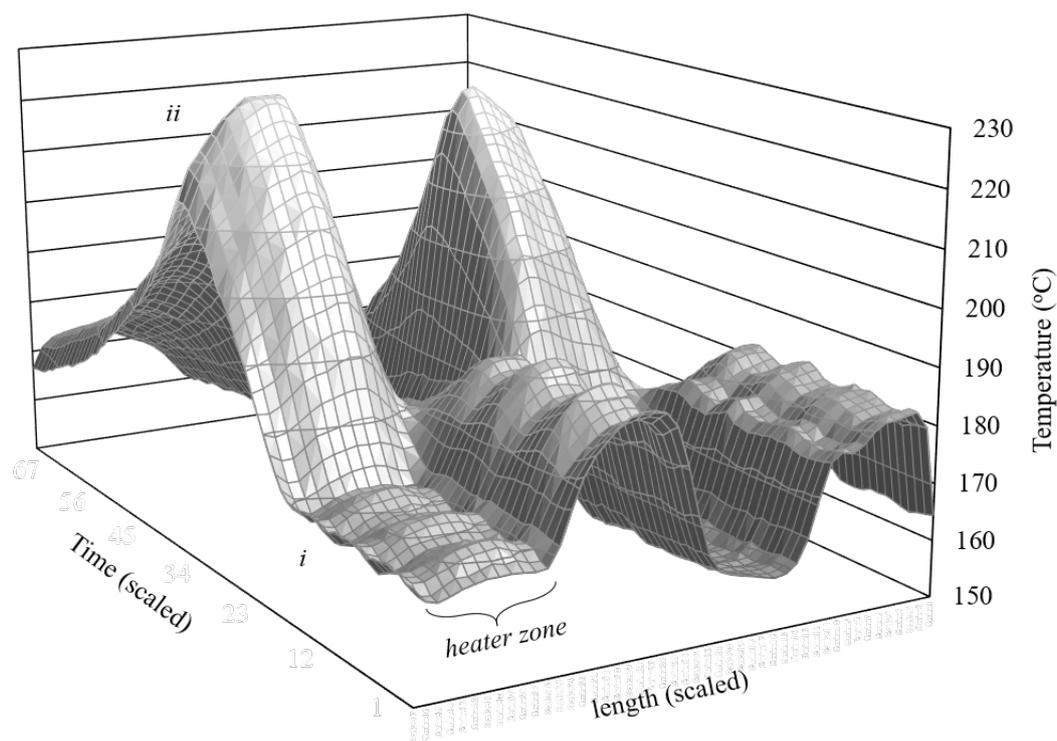


Figure 5.31: Surface plot to show various features of DTS temperature data for a helically wound lab scale reactor.

5.7 Chapter Summary

Fibre optic sensors are examined in terms of their optical properties and suitability as information carrying elements. Once a fibre specification has been established, the chapter continues to conceptualise the potential increase in spatial resolution and reduction in I/O requirements for a scaled up system. Next, the implementation of a fibre optic distributed temperature sensing system is described for the modular FT reactors under investigation. Several notes and recommendations on modifications to the reactor are made to improve this process in the future. Several mounting configurations were explored to bring about the potential for repeat measurements, robustness against breakage and augmented spatial resolution. Data was gathered for an extended run and processed using several data screening and processing techniques. It was observed that the temperature offset between distributed and discrete sensing elements was likely due to the effect of external looping in the DTS element. This effect would be eliminated in commercial scale reactors. In general, however, the DTS implementation was able to provide high levels of spatial resolution with little compromise in precision and accuracy. In a further enhancement, a helically wound lab scale reactor was shown to have an apparent spatial resolution much higher than the best possible with a linear DTS element. Further expansion of the novel concepts outlined in this thesis are suggested towards the implementation and processing of data to provide a flexible, dynamic resolution sensor to allow enhanced operability of a number of reactors simultaneously.

Chapter 6. Biomass and EfW Flue Gas Filtration Overview

6.1 Chapter Abstract

Biomass and waste are increasingly common feedstocks for thermal power generation. The technology employed in the combustion of these fuels is now mature; a derivative of long established coal power plants. This chapter first outlines the fundamental unit operations in a typical biomass or energy-from-waste power generation plant. Next, flue gas filtration technology is explored in greater detail, with particular attention being paid to the design parameters of industrial scale filter bag houses.

6.2 Biomass Flue Gas Filtration Overview

6.2.1 Typical Large Scale Biomass Power Generation Process

Biomass and Energy from Waste power generation plants are increasingly common in the UK, EU and world-wide. In general, the purpose of these plants is to produce either electrical power, steam or heat to supply an end user or customer. Though many processing routes exist (gasification, pyrolysis, anaerobic digestion), the most common is combustion. Plants range in scale from *ad hoc* < 50 kW units to > 250 MW of electrical output (MWe). For industrial power and utilities generation, the *typical* scale of Biomass and EfW combustion plants in the UK is currently ~10 - 50 MWe (UK Government, 2017).

Figure 6.1 shows a typical process flow of the key unit operations found in an industrial Biomass or EfW combustion plant. The process is briefly as follows (Storm, 2017) (Foster Wheeler, 2006):

Fuel Handling

A large proportion of the plant, by area, is concerned with the receipt, storage, preparation and conveying of fuel. Biomass in the form of pellets, sawdust, wood chips, logs, straw, etc. along with waste in the form of recycled wood, municipal solid waste (MSW) etc. are stored, sorted and processed to provide a fuel which is of suitable calorific content, moisture content and particle size distribution to allow for efficient combustion.

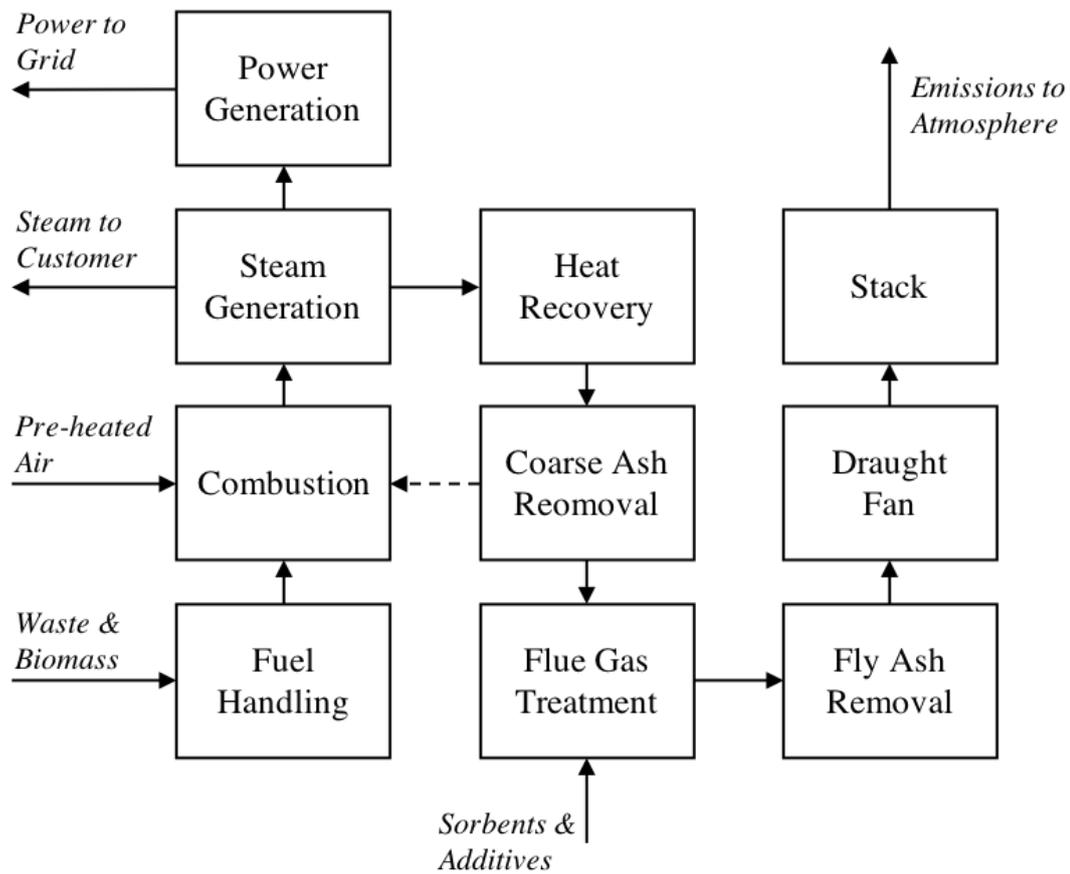


Figure 6.1: Process flow chart of industrial biomass power generation plant showing typical unit operations

Combustion

Fuel is mixed with air and ignited in a furnace. The combustion technologies employed are typically sub-divided into fixed-bed or fluidised bed, with fluidised beds being preferred at larger scales (Caillat and Vakkilainen, 2013).

Furnace temperatures of 800-900 °C are common, depending on the fuel being used. Temperature control can be achieved by varying the fuel:air mixture and also by controlling the oxygen partial pressure by means of a partial recycle of combusted gasses, depleted in oxygen. As well as fuel, refractories and additives (such as kaolin and olivine) are added to the combusting mixture. These have been shown to aid combustion, adsorb contamination and improve boiler performance. (Yang et al., 2018) (Hupa et al., 2017) (Abdullah et al., 2018) (Roberts et al., 2018) (Matúš et al., 2018) (Clery et al., 2018).

Steam generation

Steam is generated through a series of pre- and superheater coils present in the upper section of the furnace. The arrangement of the tubes with respect to the hot combustion gasses essentially forms a counter-current shell and tube heat exchanger which increases the temperature & pressure of the water/steam with high efficiency (Taler et al., 2018).

Table 6.1: Summary of common fuel sources for Biomass and EfW (Korhaliller, 2010) (Popa, 2018)

Fuel	Source	Example
Virgin wood	forestry industry	bark logs
	processing waste	sawdust pellets
Energy Crops	agricultural grasses	sugar, oil & starch crops miscanthus hemp
	short rotation crops	eucalyptus sycamore
Agricultural Residue	wet residue	manure
	dry residue	animal bedding
Food waste	catering industry	waste oils
Industrial waste	Industry	untreated wood paper pulp sewage slurry medical waste
Domestic Waste	Household	MSW

Power Generation

When electrical power generation is required, the steam generated by the process is passed through a steam turbine generator. The high pressure steam generated downstream is decompressed over a series of turbine blades, which are caused to rotate around a central shaft. This rotation induces an electric current which is transformed into useful electrical energy for export to a distribution grid (Tanuma, 2017).

Heat Recovery

Meanwhile, the hot combustion gasses exit the furnace into the 'backpass'. A further series of heat exchangers are positioned to recover energy from the combustion gasses to pre-heat water and air. This heat recovery increases the efficiency and profitability of the plant dramatically. For this reason, the assembly of heat recovery steps in the backpass is often referred to as the 'economiser' (Caillat and Vakkilainen, 2013).

Coarse Ash removal

The combustion gasses contain a mixture of combustion products and ash (Vassilev et al., 2017). Some ash is entrained into the backpass (known as fly ash). Separation of the more coarse particles of entrained ash can be achieved with relative ease by gravity or inertial

separation methods. In some cases, this coarse ash is recycled into the combustion section to extract further thermal energy and alter the combustion chemistry (Pissot et al., 2018).

Flue Gas Treatment

The combustion gas has now been cooled considerably, from ~ 900 °C to ~ 140 °C. It is still, however, laden with fine particulate ash known as 'fly ash'. Fly ash is not easily removed by inertial methods (such as cyclones) (Singh and Shukla, 2014). The flue gas is further contaminated by the products of combustion. Depending on fuel, this could comprise *HCl*, *HF*, Dioxins, Furans, *NO_x*, *SO_x*, *CO*, heavy metals (*Hg*, *Zn*, *Cd*, *Pb* etc.), Water vapour, *CO₂* and so on. Many of these emissions are strictly regulated by legislation, and a given plant will be obliged to reduce the concentration of these emissions to below an agreed level (given in the environmental permit)(European Commission, 2001) (European Commission, 2013) (European Commission, 2008).

In order to reduce the levels of emissions, flue gas additives, usually in the form of a fine powder or mist, are added. Good mixing and residence time of the additives allows for adsorption or chemical reaction to occur, reducing the concentration of emissions in the flue gas.

Most commonly in UK Biomass plants, dry sorbent injection (DSI) of particulate activated carbon (PAC) and Calcium Oxide powder (lime) is provided (Singh and Shukla, 2014).

Fly Ash Removal

The fine particulate ash entrained in the flue gas must be removed to meet particulate emissions limits. The two prevailing technologies for achieving this are electrostatic precipitation (ESP) and filtration (Hamilton et al., 2010). Of the filtration methods in use, Baghouse filtration is almost exclusively preferred (Singh and Shukla, 2014). The principles of baghouse filtration are covered in more detail later in this chapter.

Draught Fan

For flue gas flow to occur, there must be a pressure gradient between the combustion chamber and the atmosphere. This difference in pressure is the 'draught'. Draught may be generated by a blower or fan either upstream (Forced Draught) or downstream (Induced Draught) of the boiler. Induced Draught fans (ID fans) are more common, or where several fans are utilised for balancing, the ID fan performs the bulk of the duty.

ID fans require a significant amount of energy to operate. (often $\sim 5\%$ of the plants energy production is consumed by operating the ID fan.) The total amount of electrical energy *consumed* by a power generation plant is termed *parasitic load*. Over 50% of a plant's parasitic load can be devoted to operating the ID fan (Foster Wheeler, 2006).

As with all large industrial machines, the ID fan is costly and complex to maintain. In the

case of power generation, where the ID fan is in the flue gas stream, the condition of the flue gas (pH, moisture content, temperature, particulate etc.) has a direct impact on the performance and health of the ID fan.

Stack

The stack is the downstream outlet of the flue gas system to atmosphere. The stack contains critical measurement devices to monitor the various flue gas emissions as dictated by the environmental permit. The stack is almost invariably a tall, vertical duct (the height provides additional draught through the 'stack effect').

Table 6.2: Typical emissions limits for biomass combustion plants (European Commission, 2001) compared with actual emissions for a 33MWe plant (Wilson, 2012) (Day et al., 2017).

Species	Limit (mg/Nm ³)	Plant 1 Emissions (mg/Nm ³)
Carbon Monoxide	215	94
Particulate Matter	25	4.9
Hydrogen Chloride	25	5.0
Hydrogen Fluoride	2	<0.04
Sulphur Dioxide	155	32
Nitrogen Oxides	300	194
Cadmium and Thallium	0.05	<0.00081
Mercury	0.05	<0.00091
Group 9 Metals	0.5	0.15
Dioxins & Furans	0.01 (ng m ⁻³)	0.0024 (ng m ⁻³)

6.2.2 Flue Gas Filtration

Industrial Bag Filter Houses

Figure 6.2 shows a transverse cross section of the general construction of a multi-compartment baghouse. In general, flue gas flows axially through a central constant velocity plenum, where it is sub-divided into several separate compartments. Each compartment contains a number of long filter bags which are suspended from a 'tube sheet'. The bags are open at the top and held rigid by means of an internal metal cage. Alternatively, pleated filter cartridges of rigid construction are suspended without the need for a support cage (Beachler et al., 1995).

Figures 6.4 & 6.5 respectively show an axial cross section of the mid-plane of the central inlet/outlet plenum and a 3-Dimensional view of a six compartment baghouse in relation to an inlet duct. The cleaning system header tanks and compartment lids are also shown.

Flue gas entering the compartment is redirected towards the filter surface. The filter bags accumulate particulate material on the outer surface, whilst the now clean flue gas is allowed to pass through the porous filter media. The cleaned gas exits the filter bag through the open end of each bag, and is recombined in a common outlet plenum. The filter cake is periodically removed by means of a cleaning system, described more fully later.

Filters are periodically replaced to maintain the performance of the system (discussed in detail later). Replacement is carried out through access 'lids' above each compartment. Figure 6.3 shows a schematic plan view of the arrangement of filter compartments in a six compartment baghouse. The compartments are split symmetrically into streams A & B, and numbered axially. Compartment lids are numbered axially. Thus, the baghouse can be navigated by a coordinate system: "*B2-4*" is "*Side B, Compartment 2, lid 4*".

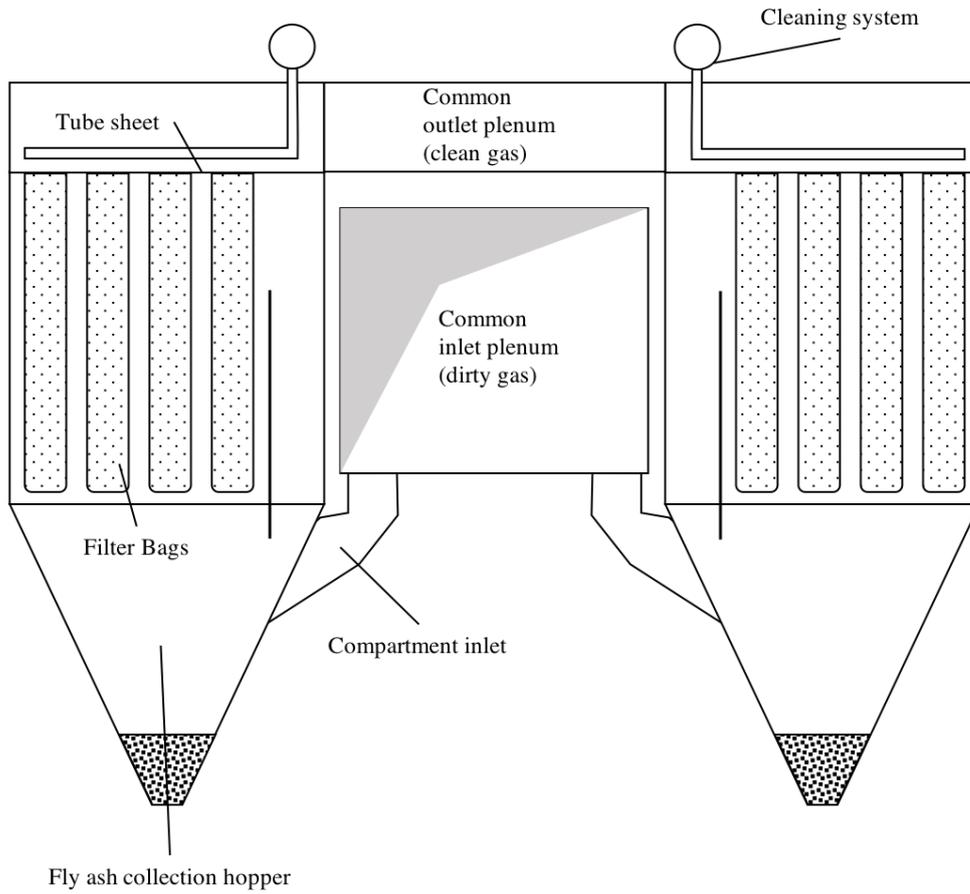


Figure 6.2: Transverse cross section of baghouse (axial midplane of compartment)

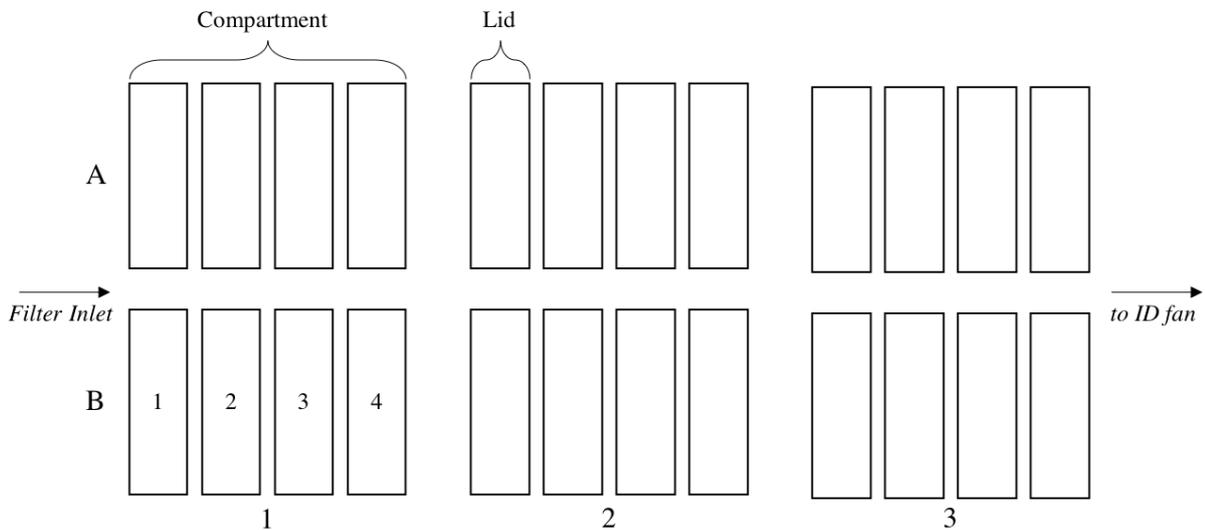


Figure 6.3: Schematic plan view of six-compartment baghouse lid arrangement and naming convention

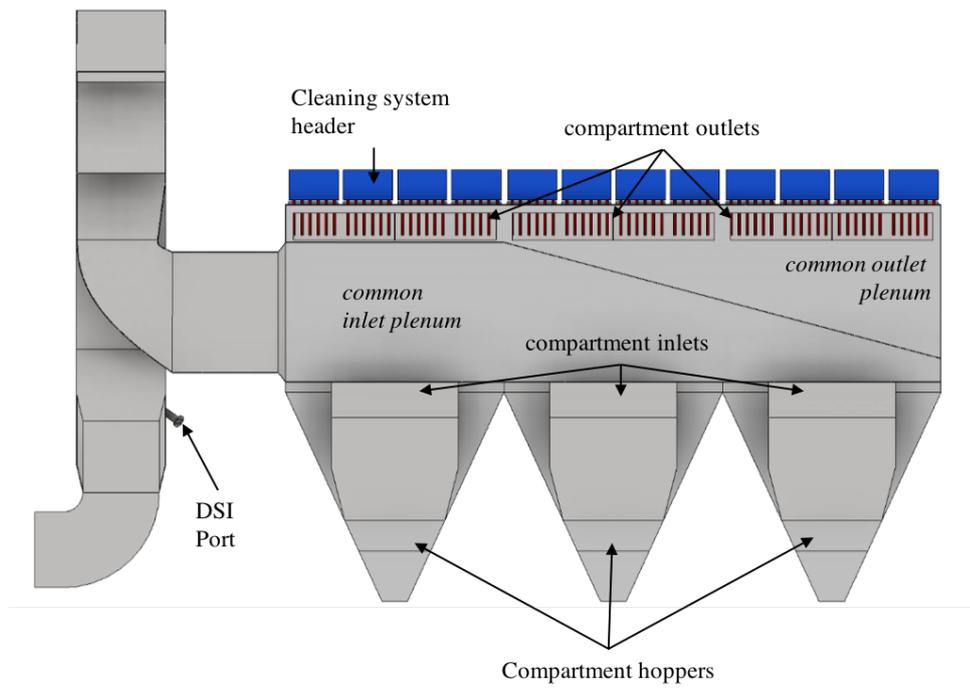


Figure 6.4: Axial cross section of mid-plane of central inlet/outlet plenum in a six-compartment industrial baghouse in relation to the upstream duct. Modified from Eagle et al. (2017).

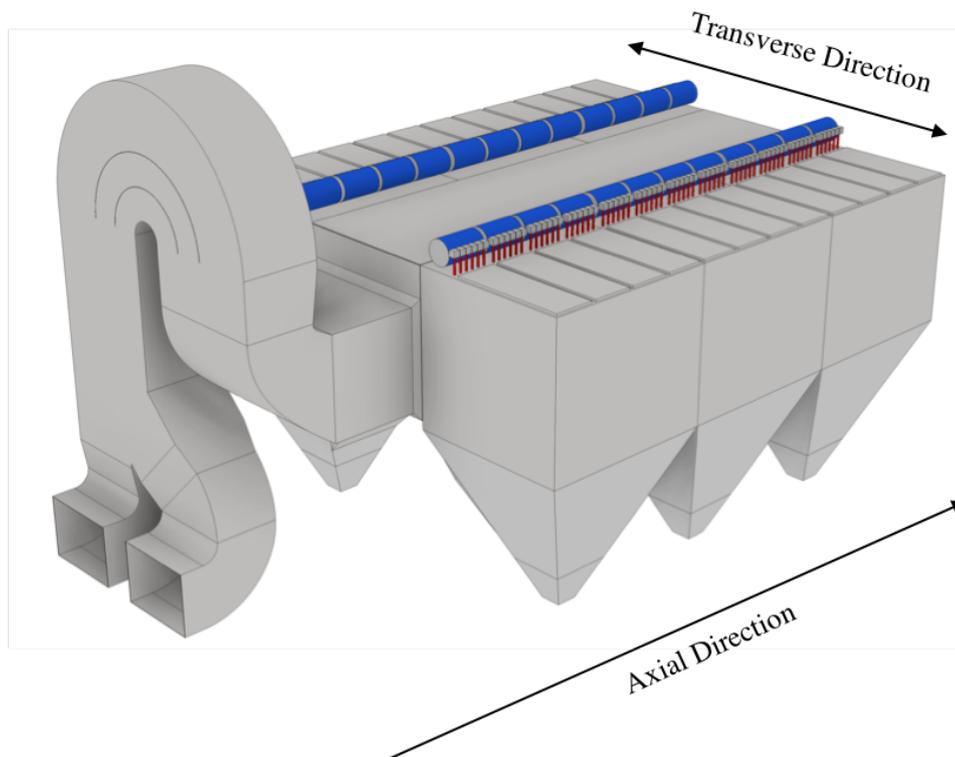


Figure 6.5: 3-Dimensional view of a six-compartment industrial baghouse in relation to the upstream duct. Axial and transverse naming convention shown. Modified from Eagle et al. (2017).

Fabric Filtration Media

The choice of filtration fabric (media) in a baghouse is critical for operation. Pre-treatment of the media with various physical and chemical processes can also enhance operation. Despite this, commercial media selection is an inexact science, often governed by unit cost. Table 6.3 lists various commercially available filtration media options. Aside from unit cost, the temperature resistance is of key importance in flue gas filtration. Typical process temperatures are $\sim 140^{\circ}\text{C}$ and therefore high temperature medias are required. This precludes the selection of some materials.

Table 6.3: Table of common filter bag materials (Wang et al., 2004) (McKenna and Turner, 1989) (Mao, 2018).

Material name	Media Construction	Max. Temp. ($^{\circ}\text{C}$)	Acid resist.	Alkali resist.	Mechanical strength
P84	Felted	232	Good	Fair	Fair
PTFE	Felted	232	Excellent	Excellent	Fair
Cotton	Woven	82	Poor	Excellent	Average
Modacrylic	Felted	71	Excellent	Excellent	Good
Polypropylene	Felted	93	Excellent	Excellent	Excellent
Aramid	Felted	204	Fair	Excellent	Excellent
PPS	Felted	200	Excellent	Excellent	Good
Acrylic (Orlon)	Felted	127	Good	Fair	Excellent
Dacron	Felted	121	Fair	Good	Good
Creslan	Felted	121	Good	Good	Good
Fiberglass	Woven	260	Good	Poor	Poor
Nextel 312	Ceramic	1300	Excellent	Excellent	Excellent

Global supply and pricing of the raw materials fluctuates greatly and is also likely to influence fabric selection (though only indirectly through the price/lead-time increase, since plant operators do not pay close attention to the global polymer commodity market). Indeed, many operators will maintain the *status quo* by repeatedly purchasing the Original Equipment Manufacturer (OEM) specification. The OEM specification was, in all likelihood, selected on the prevailing market conditions at the time of purchase (again, indirectly). As can be seen in table 6.3, felted media currently dominates the market. Currently, pricing for P84 and PTFE-based fabrics allow them to dominate, whilst Aramid fibres face global supply chain issues as a result of recent uptake by military customers.

In addition to the base media selection, there are various fabric modification steps which can take place, as outlined in table 6.4. Again, these treatments are not well linked with direct impact on performance, not least for a particular filtrate or process.

Table 6.4: Fabric pre-treatment options for filter bags. (McKenna and Turner, 1989)

Treatment	Method	Result
Calendaring	High pressure rollers	flattens & smooths fabric
Napping	Scraping metal needles across surface	raises surface fibres
Singeing	Passing over flame	removes straggling fibres
Glazing	High pressure & temperature pressing	fuses fibres
Coating	Impregnation with coating (e.g. PTFE)	enhanced properties

The general construction of felted filtration media is given in figure 6.6. A central woven support *scrim* forms a highly porous, mechanically strong substrate. The scrim does not offer any filtration efficiency itself, since the pores are *much* larger than that of the non-woven material. Either side of the scrim, a layer of felt is added. This felt is typically bonded to the scrim and to itself by needling. A series of densely packed needles is pressed or rolled into the surface to disrupt the fibre structure, forcing layers to mix. pre-treatment methods listed in table 6.4 can be used to fix the structure. The felt may be a single polymer or a blend. The blend may be different on the front (dust loading) face and the rear. For example, having a densely packed blend of P84 and PTFE on the dust loading face will give good dust release and filtration efficiency, while having a more loosely packed rear face will reduce pressure drop and increase depth-filtration dust holding capacity.

Filter cake forms on the surface of the filter media and performs the bulk of the filtration, as the size of particles is small, relative to the apparent pore size of the media. When particles enter the media cross section, they become irrecoverably embedded. This process is known as filter *blinding*, and is responsible for the irrecoverable increase in pressure drop over time for filtration media.

Cleaning Systems

In an attempt to increase the operational life of a filter, by periodically removing filter cake and, to a lesser extent blinding, cleaning systems are required. Table 6.5 lists the principle cleaning mechanisms employed by industrial bag houses.

Filtration Additives

Several commercially available filter additives can be dosed into a filtration system, usually by pneumatic conveyance of powder or aerosol spray of liquids. In Biomass and EfW applications, the predominant filter additives are PAC and Lime.

Powdered activated carbon (PAC) is applied to reduce the concentration of heavy metals,

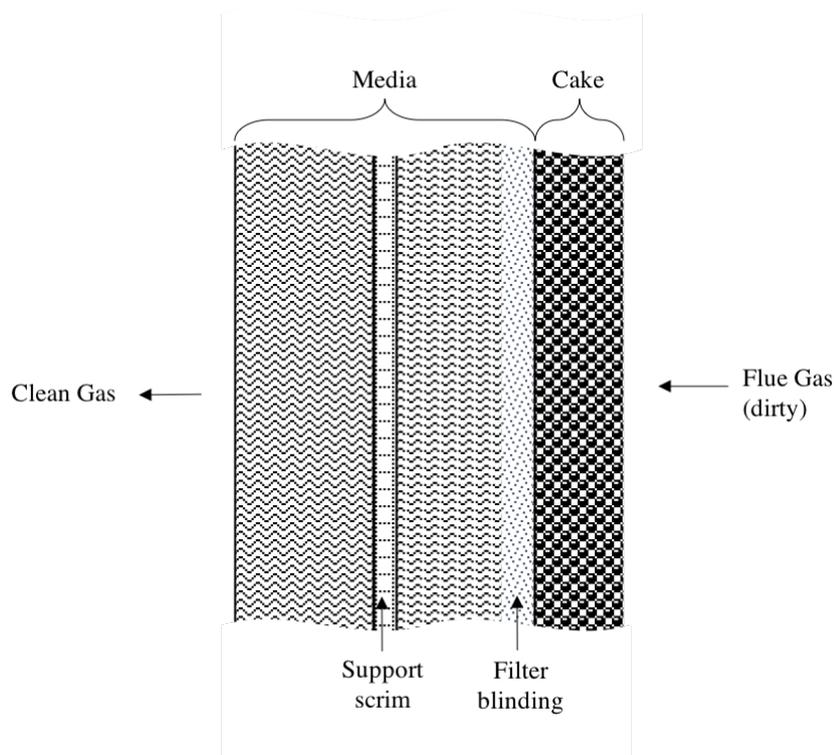


Figure 6.6: Schematic cross section of filtration media structure

Table 6.5: baghouse cleaning systems

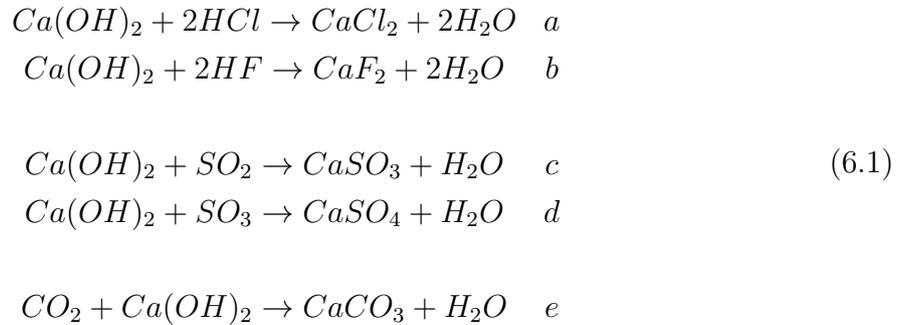
Type	Advantages	Disadvantages
Shaker	High collection efficiency Simple operation relatively low pressure drop	low air-to-cloth ratio low max. temperature high maintenance cost hopper human entry required
Reverse Air	High collection efficiency High temperature operation	low air-to-cloth ratio frequent offline cleaning cleaning air requires filtering high residual dust cake
Pulse Jet	High air-to-cloth ratio minimal residual dust smaller footprint No human entry required Low Pressure drop	Requires dry compressed air sensitive to moisture
Sonic Horn	Supplementary cleaning	

dioxins and furans at the stack. These species are adsorbed onto the macro-porous structure of the PAC. Being larger than the filtration pore size, the PAC particles are easily removed by filtration.

Lime (calcium oxide /calcium hydroxide) is dosed to neutralise acidic components of the

flue gas. If untreated, components dissolve in the moisture of the flue gas to produce acids such as sulphuric acid, which are highly corrosive. This can damage the infrastructure of the baghouse and erode the filter media creating holes in the fabric (Eagle et al., 2017) (Foster Wheeler, 2006).

The following reactions occur in the flue gas to remove the acidic components.



Notably, lime dosed into the flue gas stream is exposed to combustion products, H_2O and CO_2 , giving rise to the formation of limestone ($CaCO_3$) (figure 6.7). This limestone is not readily calcined in the baghouse environment (the required temperature is not reached) and therefore accumulates. This results in hard, brittle and impermeable deposits forming on the filter surface, resulting in high pressure drop and bag failures.

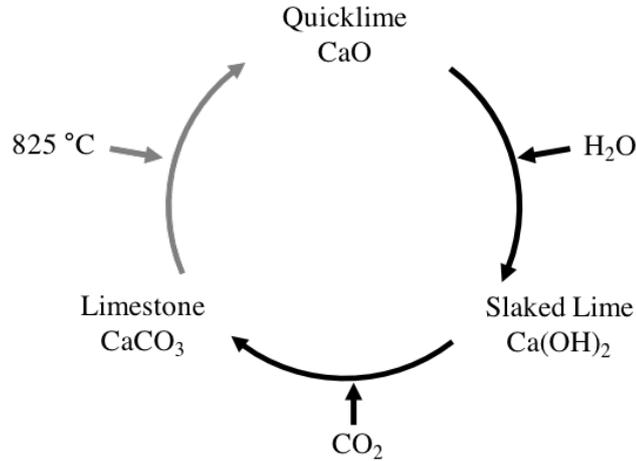


Figure 6.7: Schematic representation of the lime cycle

Chapter 7. Minimising Total Cost of Ownership of Industrial Baghouses

7.1 Chapter Abstract

This body of work was carried out in collaboration with two industrial biomass fired boilers (named here as Plant 1 and Plant 2), producing ~ 33 MWe & ~ 15 MWe respectively. In the current work, the focus is predominantly on Plant 1 (~ 33 MWe). Commercial and operational data was made available for the study. The author carried out this research with the assistance of four MEng chemical engineering master's students over the course of two years.

Plant 1, a 33 MWe biomass power generation plant is examined in detail. Real historic plant operations data is combined with commercial data to provide an economic decision model. A 3-dimensional model of key plant areas is constructed to perform computational fluid dynamics studies. Once a robust CFD environment is benchmarked to current operational data, the 3D models are then tuned to provide potential process improvements. This chapter explicates the developed approach of Rapid Hybrid Computational Fluid Dynamic Design, whereby CFD is combined with traditional response surface methodologies to arrive at constrained optima for a given geometry. Several case studies are given for the implementation of this approach, with the commercial benefits being predicted.

The overall aim of the study was twofold: Firstly, Plant 1 was, in essence, a customer of consultancy services aimed at improving operations and process economics. Secondly, the Industrial sponsor, Durham Filtration was concerned with developing process understanding and a set of approaches and tools to allow for ongoing consultancy services for other customers in the biomass and EfW sector.

7.1.1 Master's Theses Summary

In Day et al. (2017), we performed various multivariate methods on raw plant data. This work formed the foundation of understanding of baghouse performance in relation to plant KPIs. Key steps in this work were the normalisation of baghouse differential pressure to the MW output of plant, and a nonlinear multiple regression analysis, linking average bag

age to OPEX.

In Eagle et al. (2017) we developed a combined CFD & DoE approach which forms the pre-cursor to the Rapid Hybrid Computational Fluid Dynamic Design (RHCFDD) methodology formally stated in this work. In steady state simulations, various geometries were optimised. A CFD based minimisation of baghouse pressure drop was conducted in an extension of pilot studies performed by Chen (2001) and Chen and Cheng (2005). A key step in this work was to use industrially relevant geometries and scales, rather than the scale-down model used in previous work. Additionally, CFD was carried out to quantify the sorbent mixing performance and residence time provided by various commercially available mixers.

In Taylor et al. (2018) we applied RHCFDD to a transient simulation of an industrial pulse-jet cleaning system. Nozzle geometry was optimised to give enhanced cleaning of the filters. A key step was the relation of peak pressure at the bag inlet to the overall cleaning performance, allowing a computationally efficient metric of cleaning to be optimised.

In Pettitt et al. (2018) we modelled the cake formation and blinding of filter fabrics by combining mechanistic and data-driven models, and performed an optimisation of pulse frequency. The key step was the ability to separate a previously lumped parameter, cake resistance, into two parameters; one for that of the surface cake and one for the filter blinding effect. The result was a much more accurate prediction of the $\Delta P/t$ profile for baghouses, and a more sensible competing objective function for pulse frequency.

7.2 Experimental Methods

7.2.1 Biomass Power Generation and Filtration Trials

Biomass combustion and power generation was carried out in an industrial scale circulating fluidised bed (CFB) boiler (Foster-Wheeler) over a period of two years. A base-load electrical output of 33 MWe was produced. A corresponding flow-rate of high, medium and low pressure steam was also produced and supplied to nearby end users. The flue gas (approx. $50 \text{ m}^3 \text{ s}^{-1}$) was treated in accordance with the plant's operating procedures and environmental permit. Particulate Activated Carbon and dry powdered calcium oxide were supplied to reduce the concentration of acidic and toxic emissions. The resultant mixture was filtered in a six compartment industrial bag filter house (Lühr) containing 4608 individual flat-type filter bags. The boiler draught was induced by a large industrial variable speed fan. Some filtered flue gas was recirculated to provide combustion control, whilst the remainder was emitted to atmosphere through a 65 m stack (Foster Wheeler, 2006).

During the two year operating period considered in this study, several process disturbances

were observed, allowing for both steady and non-steady state operation to be characterised. The flue gas filters were changed a number of times and in varying fashion during the run, with the impact of these changes forming a large part of the discussion in this work. Changing of the filters is a costly and lengthy process which may be completed either during a planned maintenance shutdown or, to a more challenging extent, while the plant is operational - albeit at a lower capacity. Furthermore, the continued demands for power and steam meant that the degree of control over process parameters exerted by the author in this study were limited. However, access to both real-time and historic plant data allowed for the retrospective interrogation of the data.

7.2.2 Historic Data Analysis

Real time operational data was sampled for various industrially relevant performance indicators throughout the run, and stored on a database for periodic offloading. The data was periodically processed and analysed to provide specific process understanding and interventions. The full data-set was compiled at the end of the study period and processed to allow for various hybrid modelling techniques to be applied. Figures 7.1, 7.2 & 7.3 show downsampled raw process data for a number of measured process variables over the study period. As can be seen, the data has a number of artefacts relating to various operating points (startup, shutdown, steady state) as well as to process disturbances (variable fuel quality, atmospheric/weather conditions, equipment failure).

In total, > 1000 sensors are gathering data on site, with many values updating on a second by second basis. Over a two year period, this amounts to > 6×10^7 data points. Data pre-processing is fundamental to understanding such a large and complex data set. Pre-processing the data involves selecting only the measurements of interest (this step is often non-trivial and requires process understanding), filtering non-operational data points by cross-referencing production schedule information (e.g. cutting out data from plant shutdown or maintenance periods), or downsampling the data (reducing the spatio-temporal resolution of data by performing time averaging (e.g. Stack HCl emissions is level is measured and stored per second. This can be time-averaged over a window of, say, an hour) or spatial averaging (e.g. furnace temperature is given by 8 spatially distributed thermocouples. This can be reduced to processing, say, either the mean or maximum value) (Bivand and Krivoruchko, 2018) (Shrestha and Bhatta, 2018).

Screening, filtering and downsampling of process data without losing valuable process understanding or information is at the core of hybrid modelling, and is explored in many forms in this work.

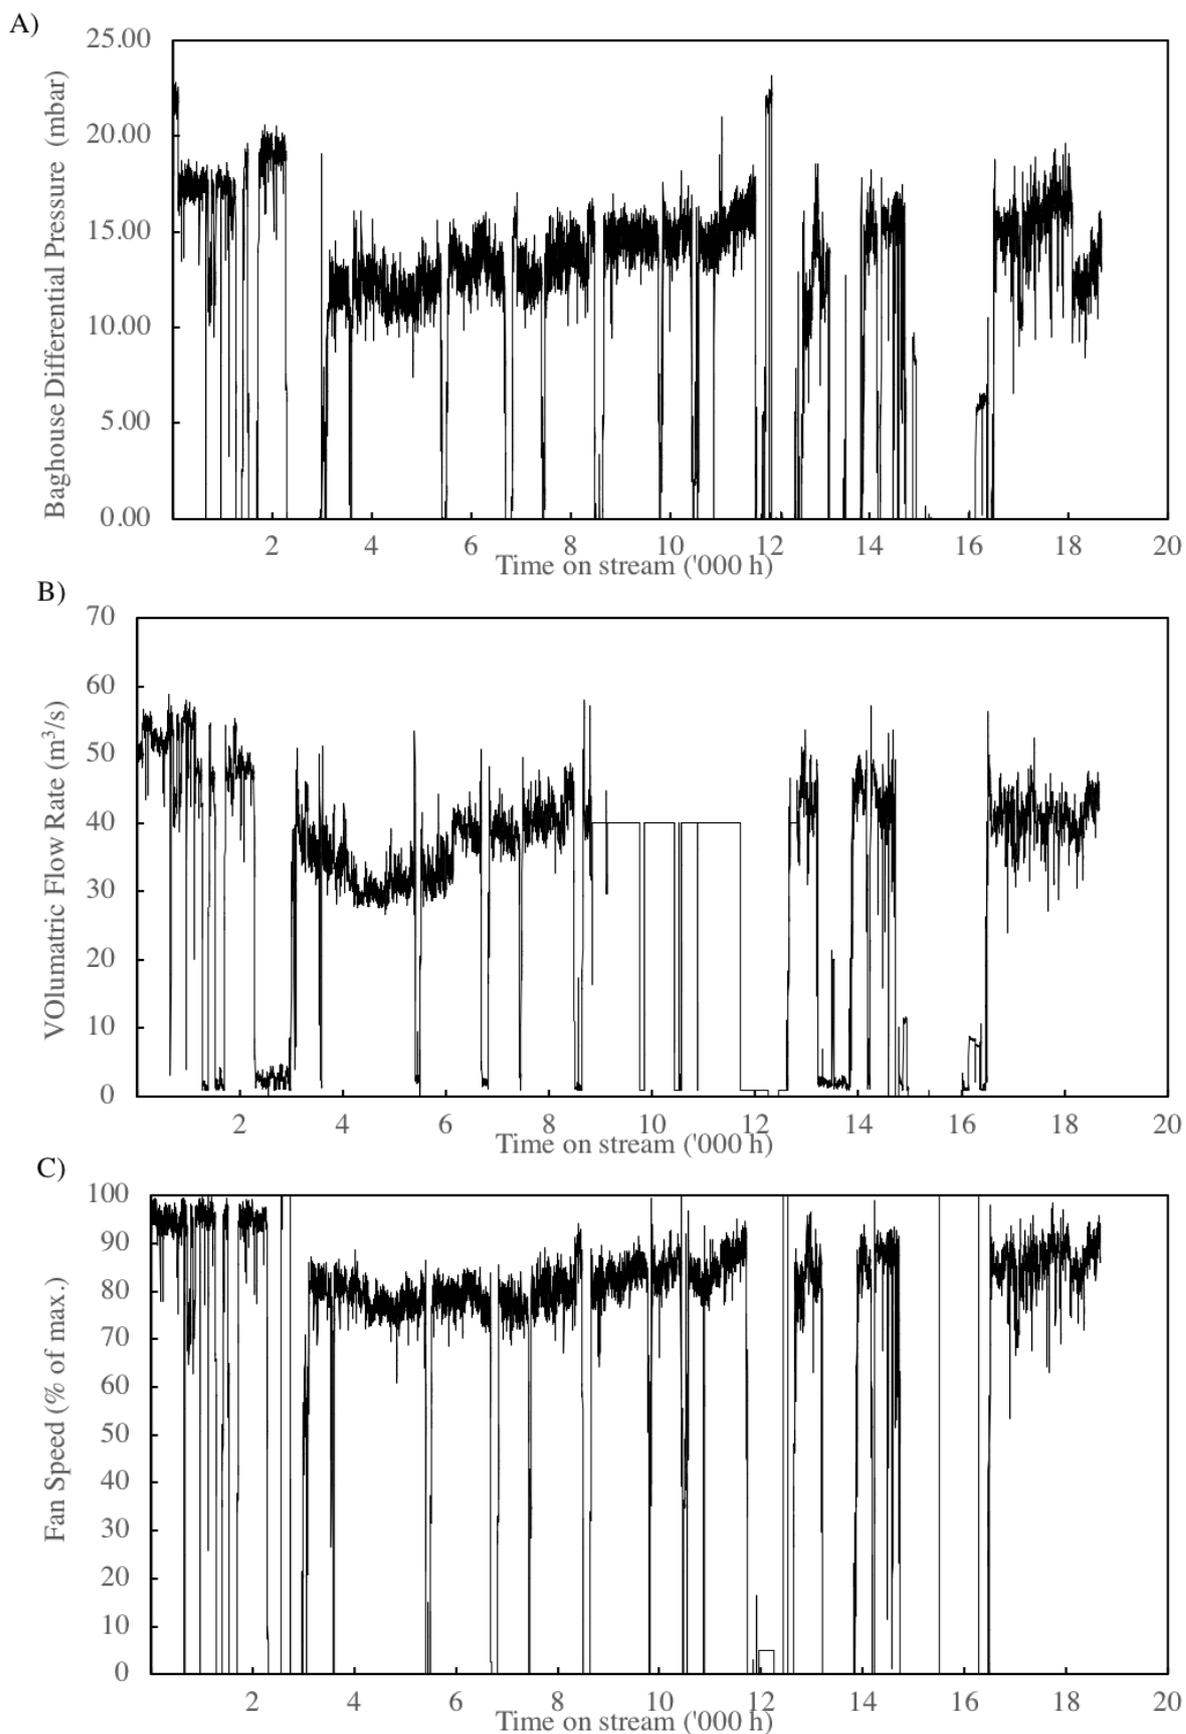


Figure 7.1: Raw historic process data for an extended period of operation (>2 years) A) Baghouse differential pressure (mbar). B) Volumetric Flow rate (m^3/s). C) Induced draught fan speed (% of max.).

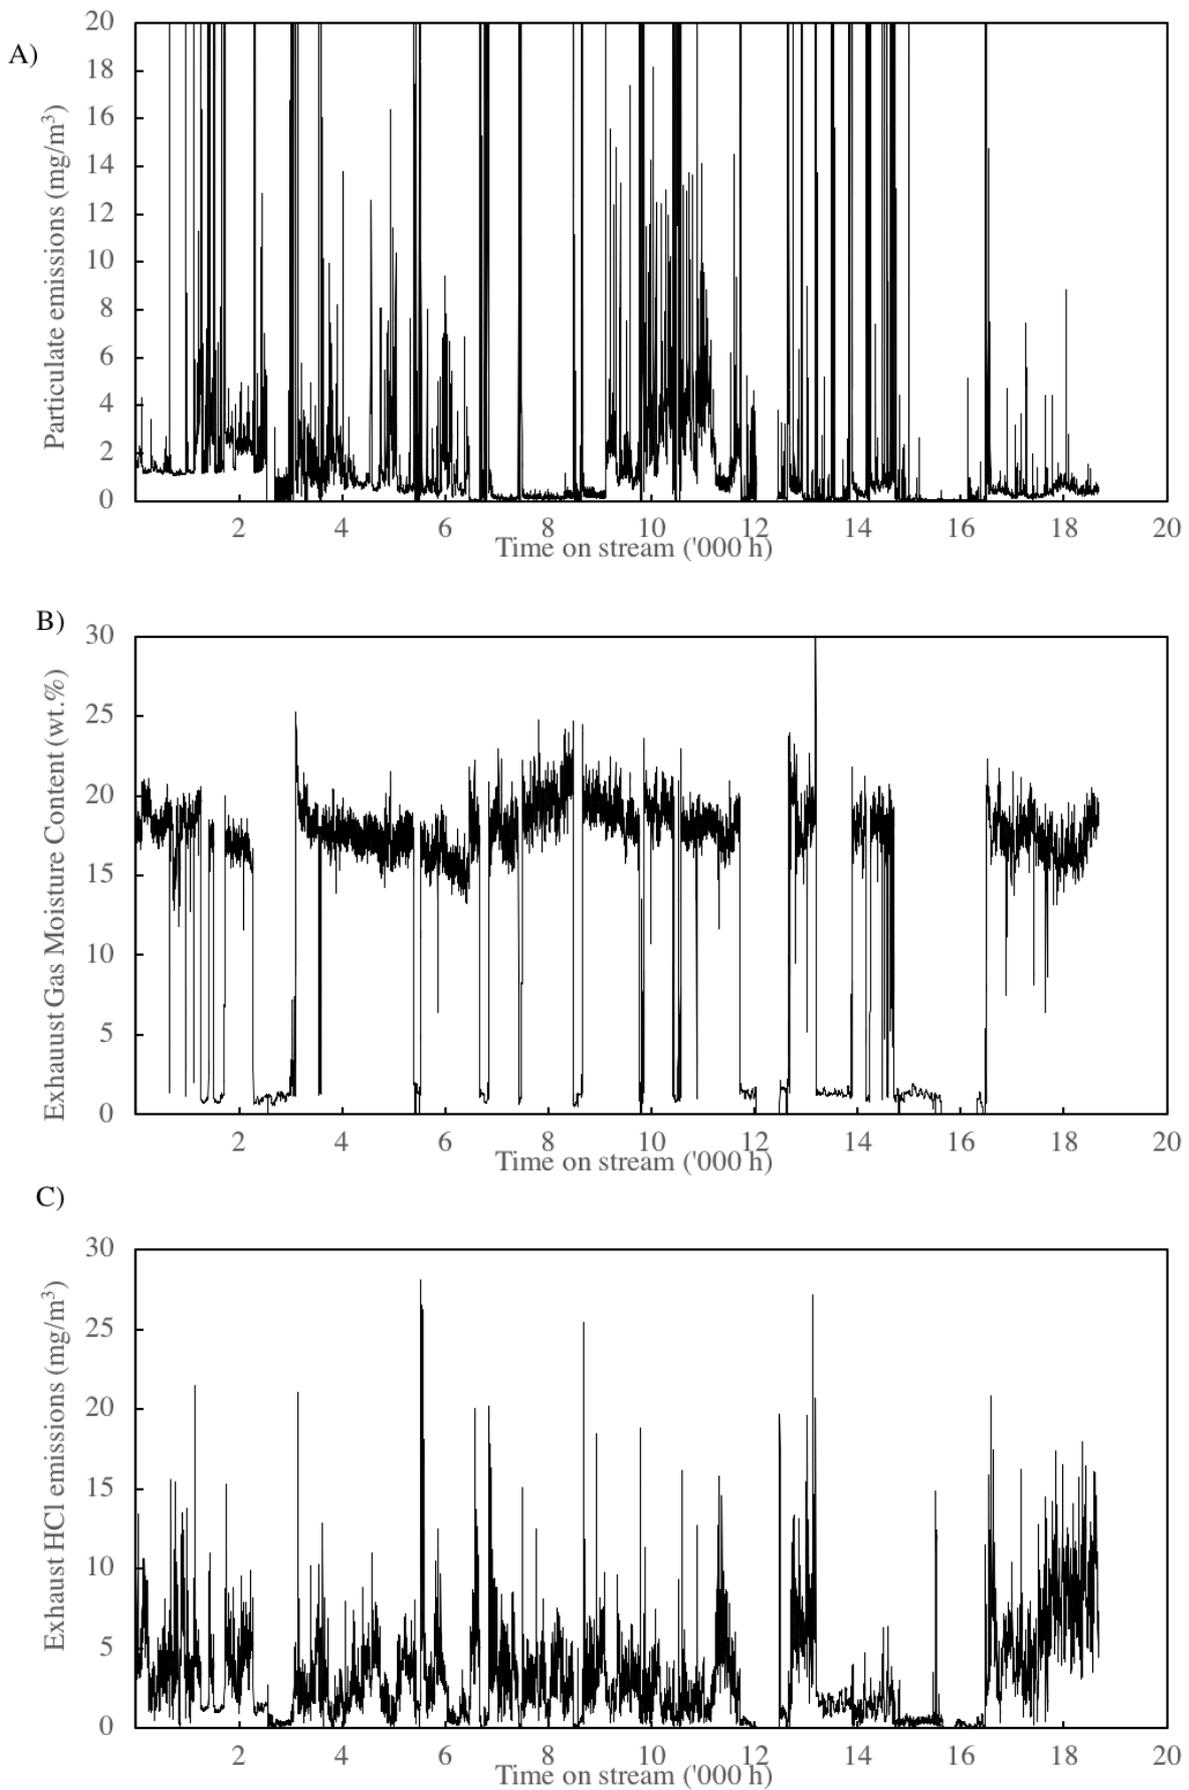


Figure 7.2: Raw historic process data for an extended period of operation (>2 years) A) Stack exhaust particulate emissions (mg/m^3). B) Stack exhaust moisture content (wt.%). C) Stack exhaust HCl emissions (mg/m^3).

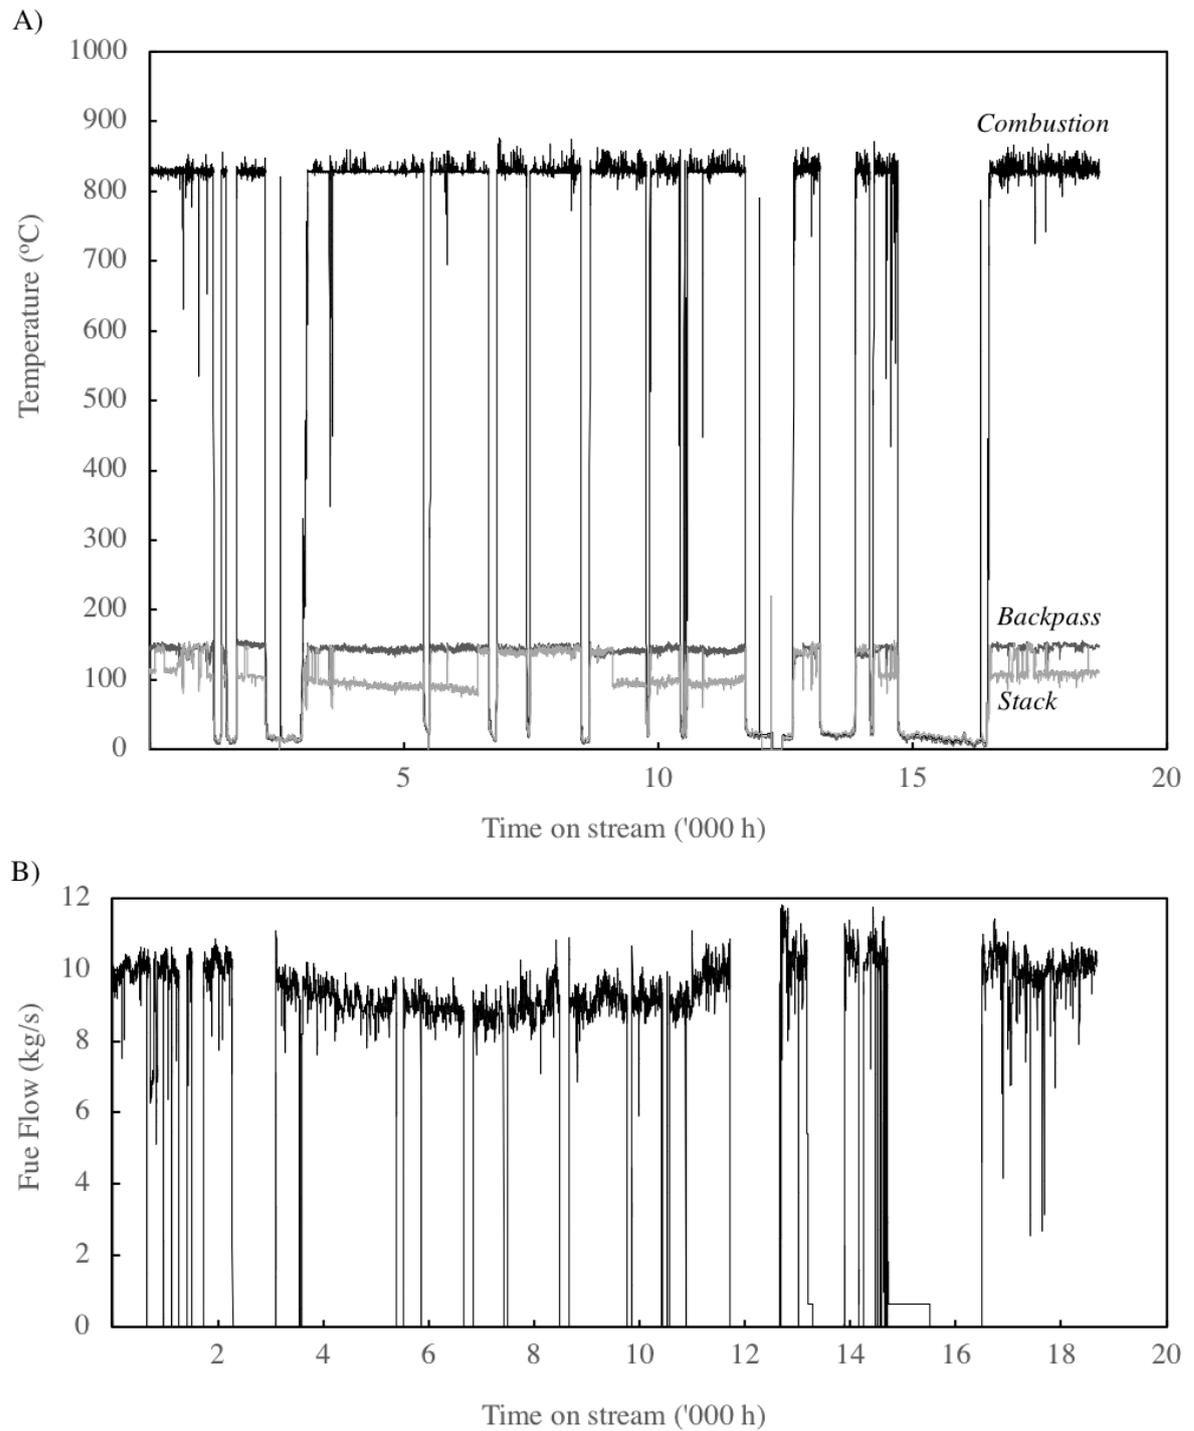


Figure 7.3: Raw historic process data for an extended period of operation (>2 years) A) Combustion, Boiler backpass and Stack temperatures ($^{\circ}\text{C}$). B) Solid biomass fuel flowrate (kg/s).

7.2.3 Offline Filter Analysis

A comparative experimental study of two brands of filter media was carried out at the industrial baghouse under investigation. The first filter media was manufactured in the UK by Tyne Tees Filtration, under the US company, Lydall. The second filter media was manufactured by the German company Gutsche. The Gutsche media alternative was of a slightly different construction, with finer fibres on the raw gas face, supposedly to improve dust release. The use of coarser support fibres also had an associated cost reduction. Other than those two attributes, the product behaviour would supposedly be the same. (Lydall-Gutsche, 2015)

Filter sample location

The two bag brands, Gutsche (G) and Tyne Tees Filtration (TTF), were installed in proximal locations of the baghouse to ensure the same duty. The bags were operated for nine months (approximately half of the predicted filter life) before a spot-sample was taken from strategic locations across four lids.

Shaded regions in figure 7.4 show the location of the differing brands of bags. In the middle compartments (A2 & B2). Crosses show the locations of samples obtained in this study. Six bags were taken from compartment B2 in symmetrical locations to allow for direct comparison. The sample locations also allow for some spatial effects to be quantified. It is proposed that a further study after a period of one additional year of operation would allow for further analysis of the performance over time. Unpublished commercial work by the author has indicated that for filters between 3 and 7 years of operation, a plateau of filter properties is reached. Long-term industrial data in this field is lacking in the literature.

Figure 7.5 shows the sub-sampling treatment applied to each sample. Horizontal and vertical strips, as well as a large square sample were required to complete cross- and machine-direction mechanical testing as well as physical performance. Samples from the top, middle and bottom of each bag were obtained. This procedure was also carried out for an unused OEM supplied bag (Lühr). In this way, the number of replicates for each sample is increased, whilst also allowing statistical evaluation of the performance of bags in three dimensions.

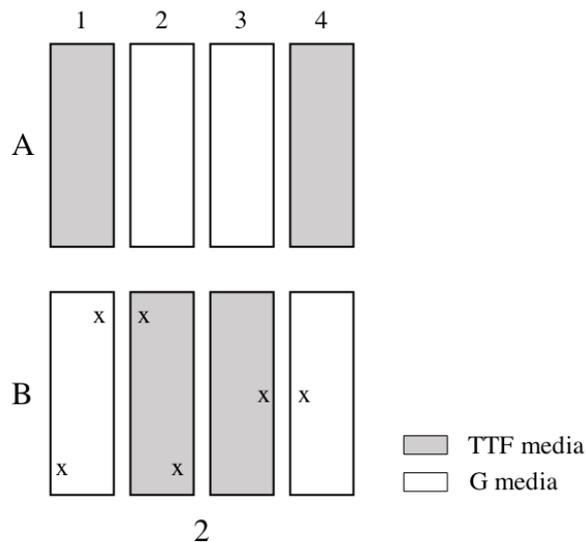


Figure 7.4: Schematic plan view of baghouse lids (A2 & B2) showing location installed filter brands. Crosses indicate locations of sampled bags.

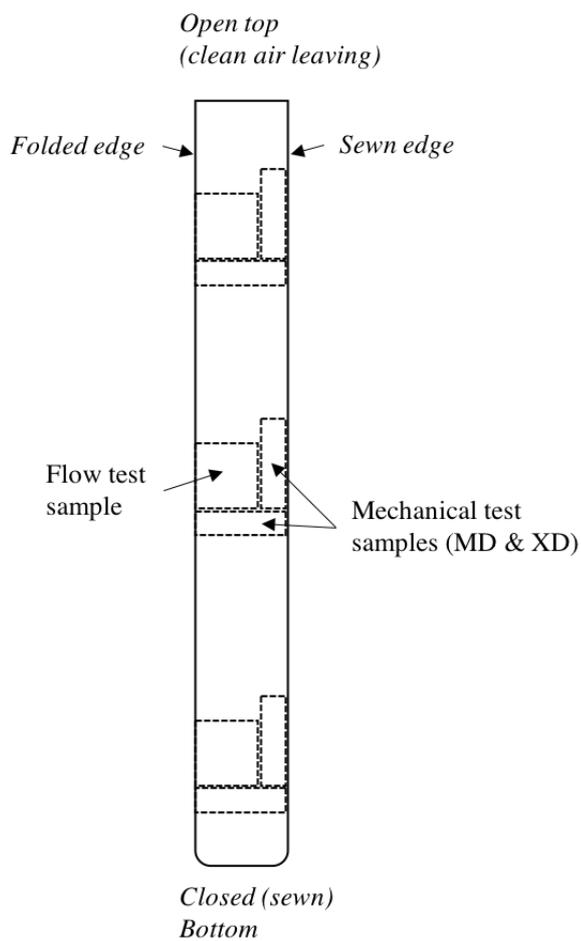


Figure 7.5: Schematic view of a sampled filter bag showing locations of sub-sampled regions.

Filter Sample Tests

Permeability

Defined as the velocity of an air flow passing perpendicularly through a test specimen under a prescribed air pressure differential over a certain time period, and having units of dm^3/dm^2min . The test is carried out according to BS EN ISO 9073-15:2008 (BSI, 2008b). In a further development, we obtain results for both 'as received' and 'cleaned' permeability. The cleaned permeability gives some indication of the degree of filter blinding and the potential performance recovery if filters were cleaned more effectively (e.g. offline clean or optimum pulse settings).

Fabric Weight

Defined as the mass per unit area of a filter media, and having the units g/m^2 . The test is carried out according to BS EN 29073-1:1992 (BSI, 1990). Again, we obtain results for 'as received' and 'cleaned' to give an indication of the mass of particles embedded within the filter media.

Thickness

The thickness of the filter media in units of mm, was measured by BS EN ISO 9073-2:1997 (BSI, 1997). Measurements were taken for 'as received' samples.

Mechanical Properties

Toughness, peak strength and elongation are measured in both the machine- and cross-direction for each sample according to BS EN ISO 9073-18:2008 (BSI, 2008a). These values, combined, give a picture of overall bag health and remaining life.

Machine and cross-direction refer to the axis on which the filter media is manufactured. In practice, however, the machine direction also refers to the length of the bag, and cross direction, it's circumference. Since the scrim is the most mechanically strong part of the bag, the mechanical strength is largely dictated by this layer.

Peak Load Refers to the pulling force required to rip apart a strip of filter material. This represents a considerable failure of the bags. While not representative of conditions found within the baghouse, it is a standardised relative measure of tensile strength.

Elongation at 50 N Refers to the stretchiness of filter media at 50 Newtons force. This is a modest force (as compared to the peak force in the high 100s or low 1000s.) For a 150 mm sample strip, a 1% elongation denotes a stretch of 1.5 mm. In this case, the ability to stretch in the machine direction shows a good response to pulsing (where the pulse propagates in the machine direction and must deform the bag sufficiently to allow dust release). Stretching in the cross-direction (cross-sectional circumference will loosen the bag relative to the cage). This will allow for more cage rubbing and potential breakages.

Expansion in the cross-direction will necessarily increase pore size, allowing more dust to enter and pass the filters in operation. *Elongation at Peak* refers to the final stretched length at the breaking point. A similar interpretation can be made as above, where machine direction stretch is favourable to cross direction stretch.

Toughness refers to the ability of the filter media to absorb energy and plastically deform without fracturing, or the amount of energy per unit volume that the media can absorb before rupturing.

7.2.4 Fly Ash Offline Analysis

Samples of fly ash were taken at various times throughout the run for offline analysis, both from the filter cake and hopper. Offline analysis was performed to characterise the physical properties of the fly ash:

Density

Though no specific standard exists for the determination of bulk, tapped and particle density of *fly ash*, the methods are well defined for various powders. In this work, bulk and tapped density were established by the methods described in BS ISO 18842:2015 (BSI, 2017a). While the standard makes specific reference to alumina, it was found to be applicable in the present study in the absence of a specific standard.

Free lime

The concentration of unreacted lime was determined in accordance with BS EN 451-1:2017 (BSI, 2017b).

SEM & EDX

A Hitachi S2400 Scanning Electron Microscope (SEM) was used, fitted with an Oxford Instruments Isis 200 ultra-thin window X-ray detector (EDX detector). Samples were loaded and sputter coated with a layer of atomic carbon. This was necessary to improve the conductivity of the samples and allow them to be seen by the microscope.

Electron Micrographs were inspected for particle morphology and qualitative particle size distribution.

7.3 Rapid Hybrid Computational Fluid Dynamic Design

A key approach developed in this work is that of so-called *Rapid Hybrid Computational Fluid Dynamic Design*. With reference to figure 7.6, the approach is summarised as follows:

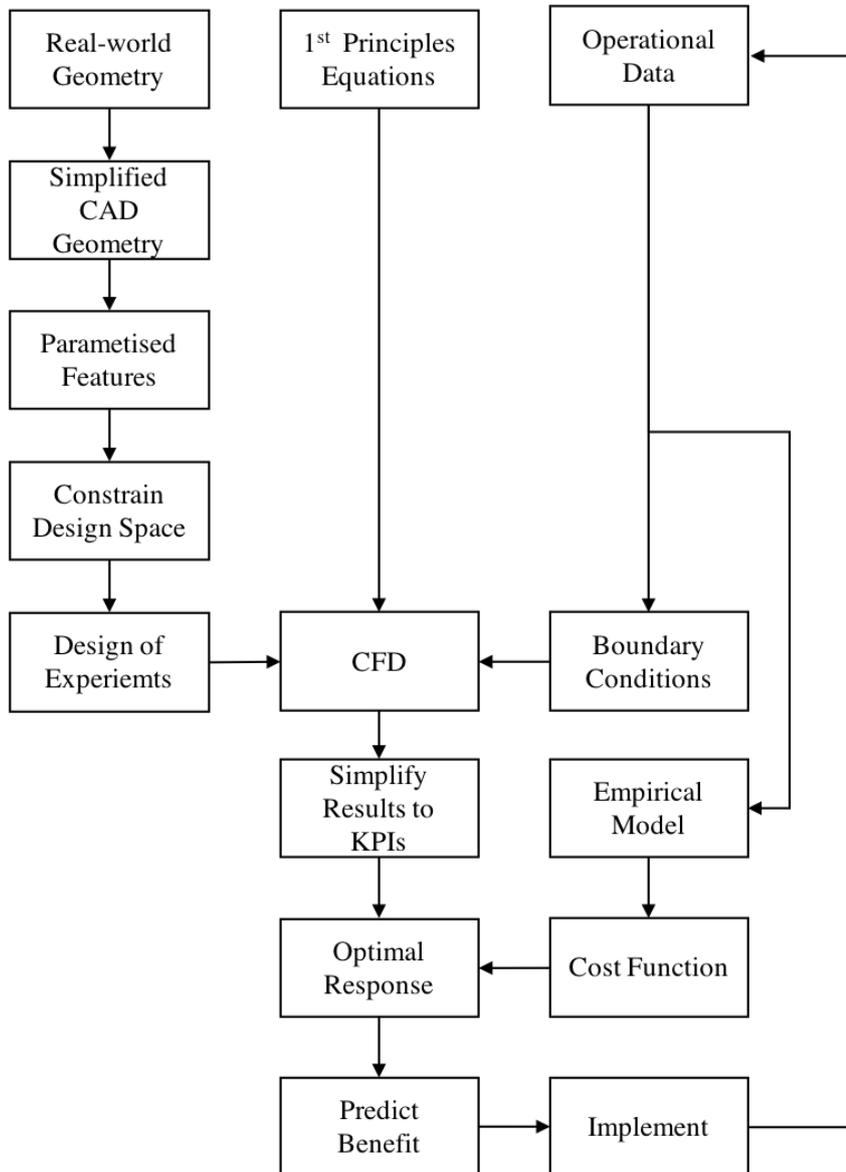


Figure 7.6: Flow chart to show general procedure of Rapid Hybrid Computational Fluid Dynamic Design

7.3.1 Representing the design space

Starting with real-world plant geometry, detailed 3-Dimensional Computer Aided Design (CAD) Models are produced. This *in silico* geometry can be constructed by physical inspection and measurement of the system, translation of the general arrangement (GA) drawings or, more recently, by Light Detection and Ranging (LiDAR) point cloud technology (Wu et al., 2018).

In any case, the Level of Detail (LOD) for the 3D model is an important consideration. For instance, it must be decided whether the accurate modelling of every bolt and rivet is required to provide an adequate representation. Increased LOD in the 3D model will decrease plant-model mismatch, at the expense of computational power requirements and solve time. Where *rapid* design iteration is required, the emphasis must be placed on simplification of the geometry in the model. Recent studies demonstrating the potential of this approach are provided in Rezeau et al. (2018) and Taghinia et al. (2018).

As well as LOD simplification, we can further simplify the model by reducing the model frame. For instance, where a repeating pattern of geometries exists, it may be possible to simulate only a single sub-unit of the pattern. To ensure that the scaled down model is representative, this may be done in stages.

1. Simulate the base case on the full geometry and validate against measured performance.
2. Create the sub-unit model
3. Use measured (if available) or simulated values from the base case simulation as the boundary conditions for the sub-unit model.
4. Perform rapid design iterations on the sub-unit model.
5. Implement the chosen sub-unit design in the full geometry model to validate the global response.

As can be seen, a rapid inner loop is created, whereby simulations of sub-unit design iterations remove the computational demands of simulating the whole system each time. By feeding the sub-unit with appropriate boundary conditions from an outer loop which can be validated against empirical data, it is possible to make these simplifications with minimum plant-model mismatch.

7.3.2 Inner Loop Design

Next, there must be a strategy by which we make intelligent sub-unit design iterations. There must also be a method to determine which of the sub-unit designs is 'optimum' (the concept of *optimum* is yet to be addressed within the framework).

Many approaches to design iteration have been demonstrated in the literature; Trial and improvement, Batch to Batch iterative learning, Design of Experiments (DoE), Deep Learning (DL), Artificial Neural Networks (ANN) are all common. However, many of these approaches are targeted at big data or process control challenges. Relatively few studies have applied this approach to *physical* parameter optimisation, and fewer still have applied these methods to industrial filtration.

One challenge faced when applying these methods to a retro-fit of an existing plant is that of physical, practical and financial constraints. Any design solution generated by an optimisation algorithm *must* be possible to manufacture, transport, assemble and maintain. It must also be affordable. While deep learning algorithms may produce the most intricate and precise geometries to optimise a given parameter, they pay no respect to the constraints listed above. The result is an optimum only in parameter response, but not in terms of a wider cost function. For example, Figure 7.7 shows three mechanical parts designed to perform a particular mechanical load. In 7.7A, the part is manufactured using standard stock materials (galvanised steel sheet and tube sections), whilst 7.7B&C are generated by deep neural networks and additive manufacture methods (Niehe, 2015). While the mechanical performance of these geometries is the same, there are clearly several practical limitations to the implementation of the neural network generated solutions in large industrial settings.

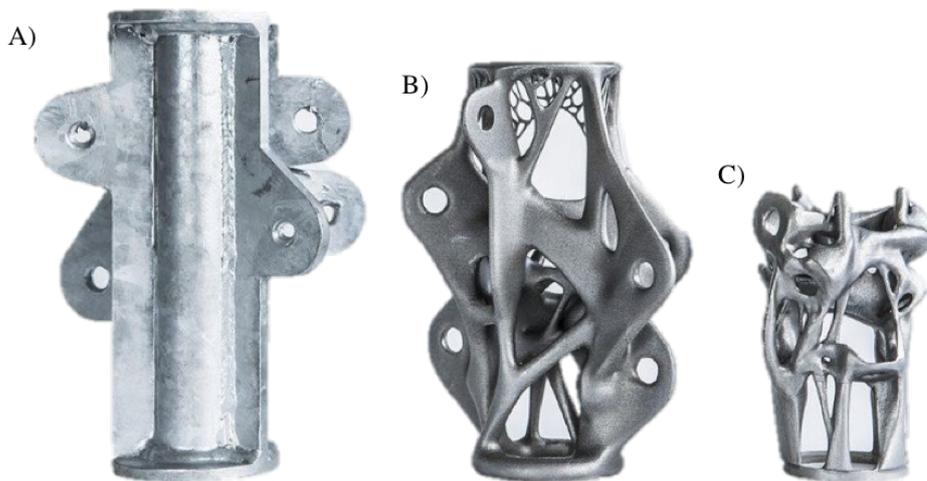


Figure 7.7: Three mechanical parts, generated by conventional and generative methods to perform the same duty (Niehe, 2015). A) Stock raw material and standard manufacture. B) Deep learning / additive manufacture geometry. C) Next-generation Deep learning geometry.

For this reason, the method developed here adopts a highly constrained approach. The use of Parametric CAD allows for dimensions and properties the modelled geometry to be changed efficiently. A well-defined parametric CAD model will respect the constraints of other fixed parameters when changing others. By limiting parameters to discrete levels, we can ensure that the solution is manufacturable at a reasonable cost. For example, if a parameter of interest is the thickness of a steel plate, the model should be constrained to use only 'standard' thicknesses of steel. Specifying 3.219747 mm steel sheets is impractical. Instead, the discrete values of 3.0mm or 3.25 mm should be chosen. If the model 'knows' this in advance, computational effort in simulating impractical geometries to unachievable precision is spared.

DoE is now a well recognised and robust approach to process development. More specifically, a Central Composite Response Surface Methodology (RSM) employed in this work has

been shown to adequately predict system response to non-simulated levels of system parameters by inter- or extrapolation (Zhu et al., 2016). A *significantly* reduced set of experiments (simulations) need be carried out to explain the variance across the design space when compared to a full-factorial or one-variable-at-a-time (OVAT) approach.

With a reduced set of simulations required for a reduced sub-unit of the global system at a reduced LOD, there is now a dramatic reduction in computational effort required to simulate the full design space with confidence.

7.3.3 Hybrid CFD

CFD is becoming an increasingly commonplace tool in research and industry. It is possible to simulate fluid and thermal performance of systems which would otherwise be too complex and time consuming to calculate by other means. The first principles basis of CFD is well understood. There are many reference works in the literature which are recommended reading for understanding the detailed partial differential equations employed. However, commercial software tools have been developed to allow for CFD to be applied in the same manner as spreadsheets or image editing. For the average user, the process is a *black box* where one must only have solid understanding of the required *inputs* and be capable of understanding the *outputs*.

Inputs

The *inputs* of a CFD model briefly comprise: the geometry, the mesh and the boundary conditions. In this hybrid framework, the geometry is provided by the DoE structure, the mesh is handled by software and the boundary conditions are delivered by empirical data.

For example, if an inlet flow rate boundary condition is to be set, historic plant data can provide the boundary condition in a number of ways:

- *Nominal value* - in the most simple boundary condition, the designed or operated flow rate is specified as a single value which does not vary over time.
- *Mean, mode or median value* - the observed variability in flow rate can be ignored and represented by a single value calculated from the distribution of measured values.
- *Monte Carlo value* - from the measured distribution of flow rates, the simulation can be iterated in a further inner loop to represent the probability distribution of outputs.
- *Dynamic value* - if transient simulation is used, the boundary condition can be set as the actual measured inlet flow rate as a function of time.

- *Empirical Model* - the raw data can be summarised by a model which is fed into the simulation. e.g. polynomial regression or piecewise linear.

Each of these approaches may be used, depending on the circumstance. Dynamic and Monte Carlo simulations come at the cost of increased computational demands but may be necessary to capture the response of highly dynamic or variable processes.

Outputs

Once the CFD simulation has been carried out, there is a further important step of extracting the outputs. One of the strengths of CFD is its ability to convey visual and qualitative understanding. However, the optimisation framework demands objective and quantitative metrics of performance. Modern CFD software packages have a number of tools to extract the simulated data in manageable forms. Once more, the required level of detail (resolution) must be decided. Modern CFD packages are able to solve industrially sized geometries with a resolution in the millions of vox els (volume pixels) and therefore data selection and downsampling becomes important.

For instance, the pressure at the outlet face of a large duct varies across two dimensions. For a reasonable resolution mesh, this could be characterised by a matrix consisting of 1000s of data points, or summarised by a single average value. One must decide not only on the property of interest (e.g. pressure, velocity, temperature, vorticity, kinetic energy, etc.), but the metric by which the property is reported (e.g. average, rate of change, variance, range, etc.). This will depend on the objective of the study and the current understanding of the system.

7.3.4 Cost Functions

In order to assess the impact of a change in system parameters to a meaningful extent, a cost function must be developed. The 'cost function' in this framework is quite literal; it is the *cost* of operating in a particular manner. *Total cost of ownership* is an extended cost function where both the direct, tangible costs and the indirect, intangible costs are combined.

For instance, tangible costs consist of materials & labour, energy, raw materials and utilities. Intangible costs could encompass reputation, environmental performance, occupational health.

7.4 Results & Discussion

7.4.1 Offline Filter Analysis

Permeability

Figure 7.8A shows a number of interesting results. Firstly, when comparing the two brands, it can be seen that Gutsche (Brand 1) bags have a significantly higher air permeability in operation (as received). Air permeability is directly proportional to pressure drop. (Indeed, the test method calculates air permeability by measuring pressure drop). Secondly, bags have markedly higher permeability at the top of the bag, compared with the bottom. This may be a result of poor cleaning - the pulse not being strong enough to clean the bag down its full length - as well as poor flow distribution at the flue gas inlet- the majority of the fly ash is exposed to the bottom of the bag.

Figure 7.8B shows the air permeability after the bags are air cleaned under vacuum. This laboratory cleaning method does not directly simulate the pulse cleaning regime at Plant 1, but represents the performance recovery possible under an 'ideal' cleaning scenario.

Interestingly, after laboratory cleaning, both brands return to a statistically identical permeability. This indicates that although the bags have similar properties in lab conditions, the Gutsche (brand 1) bags maintain a significantly lower pressure drop in operation. Good dust release and cleanability in operation is directly linked to lower parasitic load. Bags with better cleanability may confidently be pulsed less often. This reduces mechanical wear, compressed air usage and the enhanced filter cake also provides greater dust removal.

Both bags show a significant reduction in air permeability when compared to unused bags. The new bags have an average air permeability of $155 \text{ dm}^3/\text{dm}^2\text{min}$. Both sets of bags, as received, show an operational air permeability of just 10% of this value. Cleaned bags show a permeability half that of a new bag.

The data also suggests many spatial factors in the life and performance of the filter bags. This picture will become clearer after more evidence is discussed. This discussion will proceed by comparing the different bag brands. A report into spatial effects will be delivered once more samples can be gathered to complete the picture.

Thickness

Figure 7.9 shows that both brands of bag gain thickness compared with an unused bag (on average, >1.2 mm thicker). This thickness gain is likely due to a combination of pulse-cleaning expanding the felt (raising), and a gain of mass due to particle ingress (blinding). However, TTF bags gain appreciably more thickness compared to Gutsche

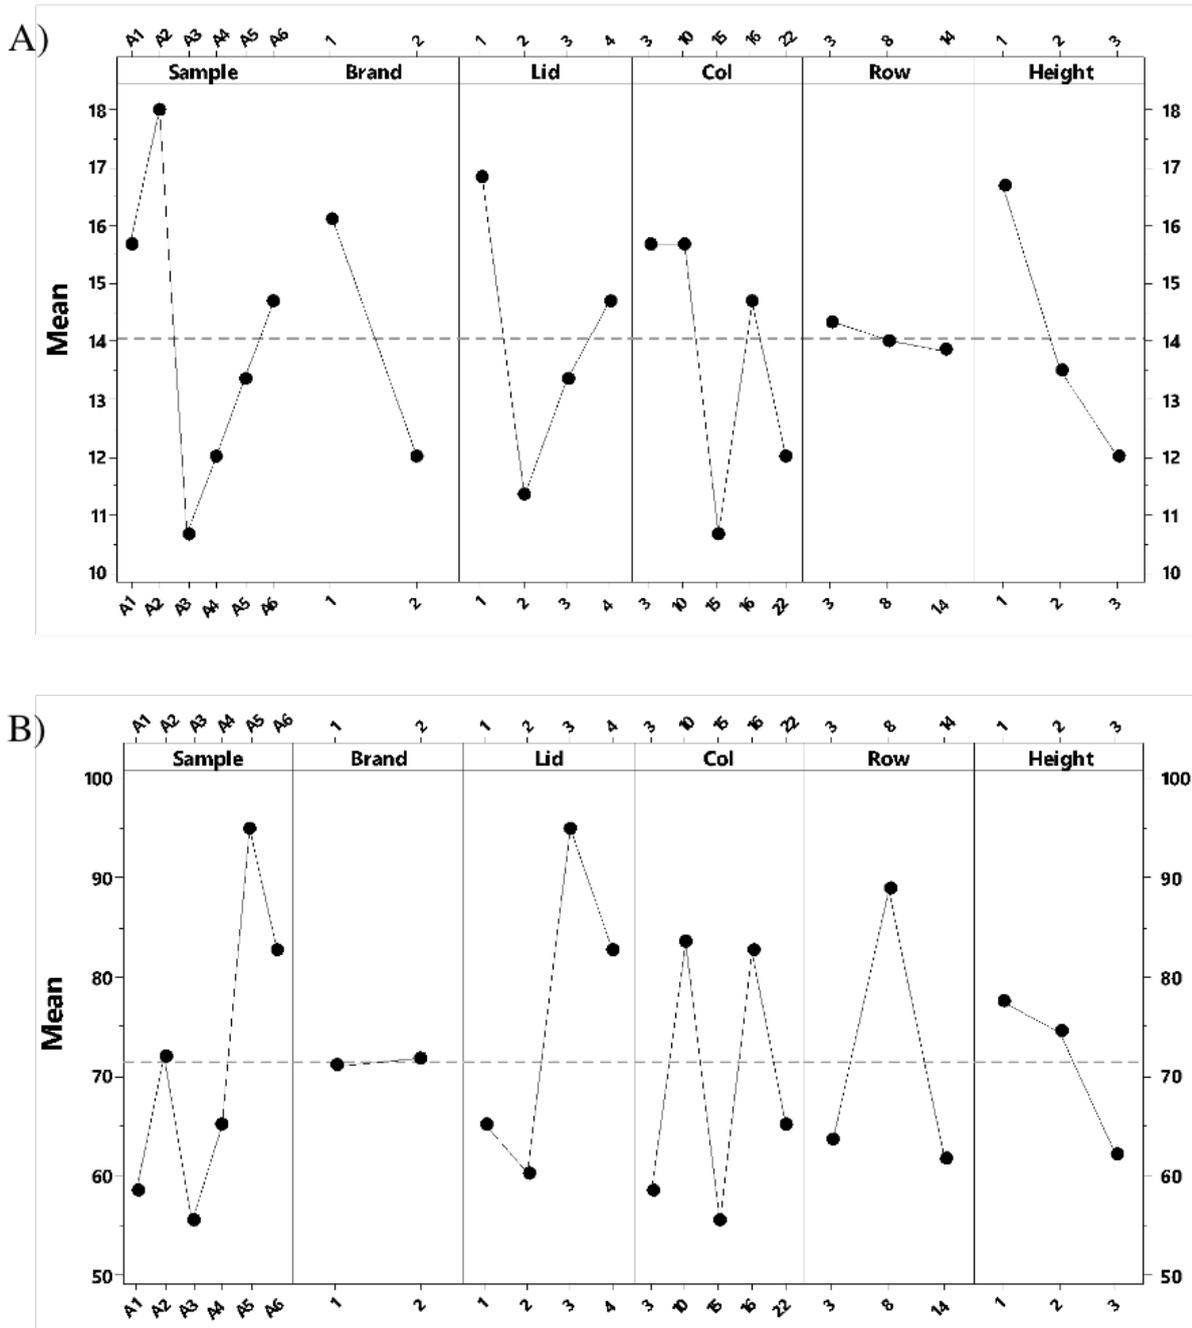


Figure 7.8: Main effects plots of air permeability ($dm^3/dm^2.min$) for offline sampled bags, in relation to geometric parameters. A) samples as received. B) Samples after vacuum cleaning. Brand: 1 = G, 2 = TTF. Height: 1,2,3 = top, middle, bottom

bags. This suggests a greater susceptibility to particle ingress and/or pulse-raised felt. The increased thickness of TTF bags likely results in a widening of pore diameter and therefore an increased rate of blinding and particulate emissions. Blinded fabrics will produce an unacceptable differential pressure much sooner in their lifetime, compared with a more robust fabric of consistent thickness.

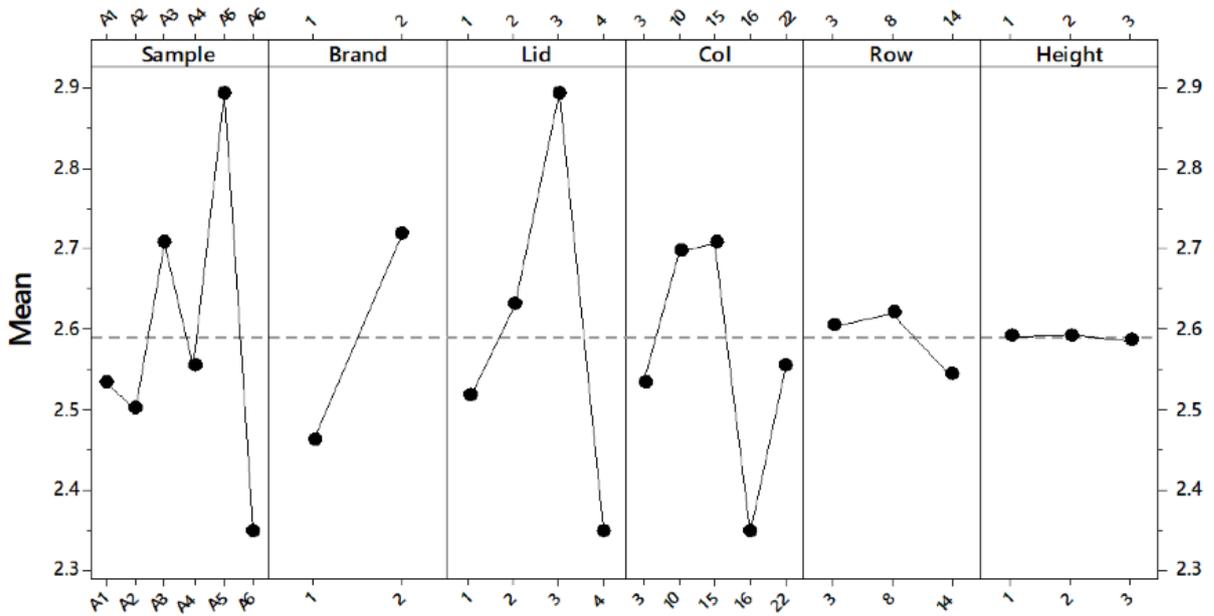


Figure 7.9: Main effects plots of thickness (mm) for offline sampled bags, in relation to geometric parameters (Samples as received). Brand: 1 = G, 2 = TTF. Height: 1,2,3 = top, middle, bottom.

Weight

With reference to Table 7.1 once the original fabric weight is considered, The Gutsche bags are shown to have accumulated 145 g m^{-2} of dust in operation compared with 30 g m^{-2} more in the TTF bags. This shows that the TTF bags have worse dust release properties in operation. Once cleaned, the Gutsche bags return to a lower level of residual dust, showing that less irreversible filter blinding has occurred.

Table 7.1: Average filter weights from offline analysis.

Brand	Filter Weights (g/m^2)				
	New	As Received	Accumulation	Cleaned	Residual
Gutsche	530	675	145	578	48
TTF	550	725	175	625	76
OEM	600				

Online Validation

In each of the standard offline tests undertaken, Gutsche bags perform better than TTF bags. They are shown to have better dust release and mechanical properties. In operation, this means a longer useful life in terms of both pressure drop and mechanical failure.

If the laboratory testing conclusions are true, there should be some appreciable change in operational parameters when changing to the Gutsche bags. In May 2016, 12 lids (half of the baghouse) was changed to the Gutsche bags. This significant change should show

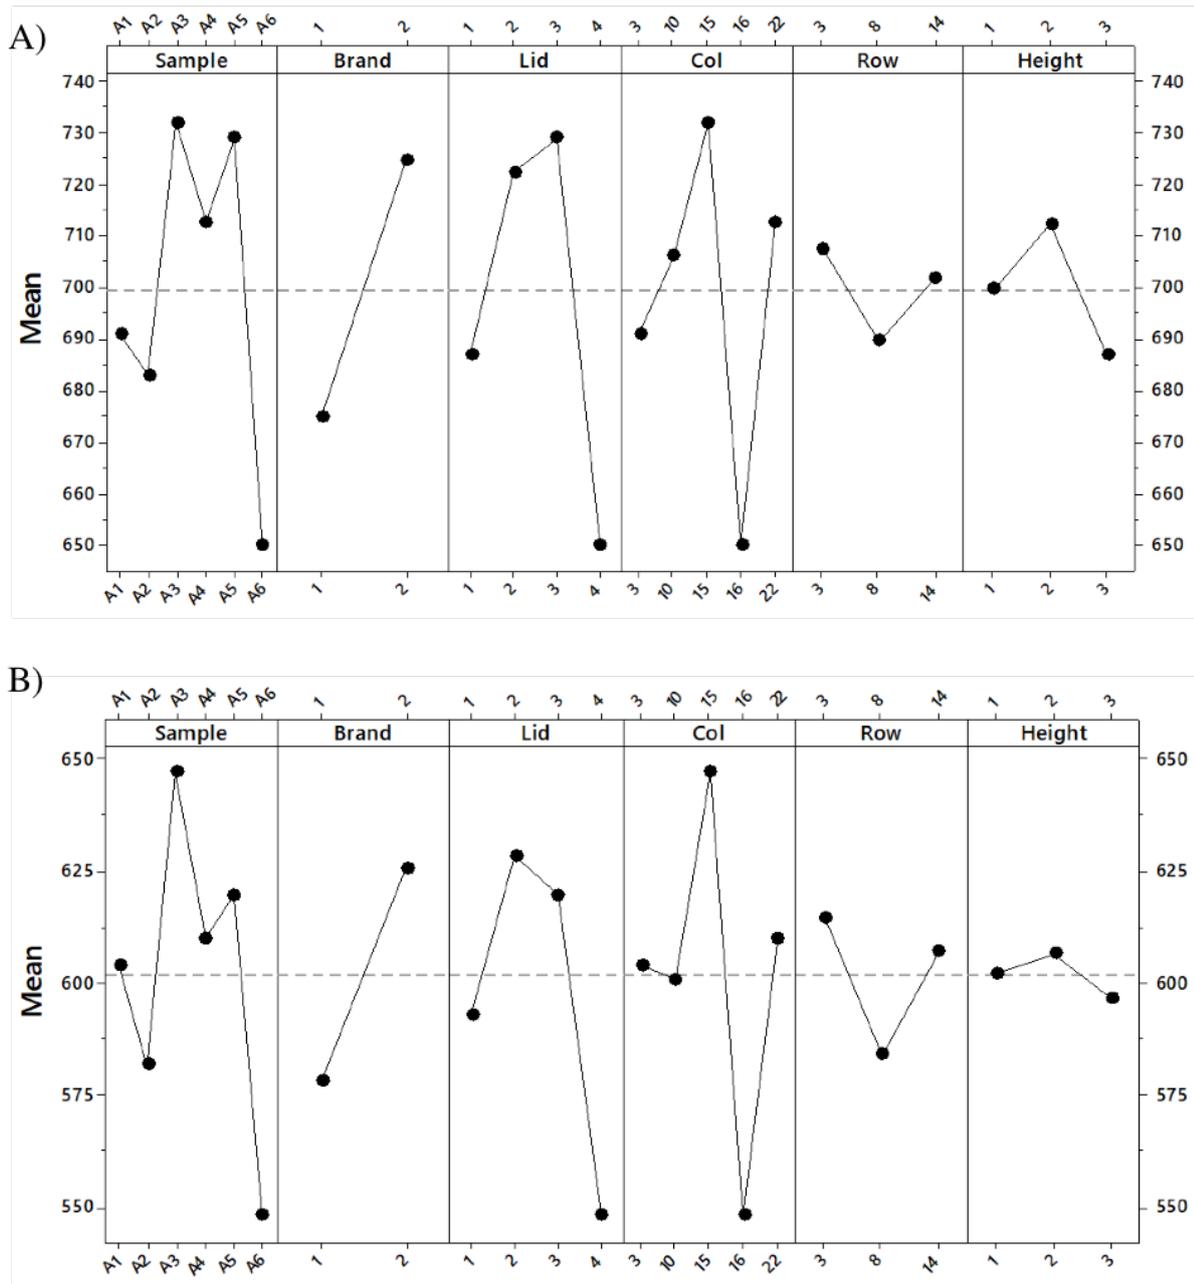


Figure 7.10: Main effects plots of weight (g/m^3) for offline sampled bags, in relation to geometric parameters. A) samples as received. B) Samples after vacuum cleaning. Brand: 1 = G, 2 = TTF. Height: 1,2,3 = top, middle, bottom

some step change in performance, over and above the usual temporary performance gain from refreshing bags, like-for-like.

However, during this time, Plant 1 also undertook some efficiency exercises and reduced turbine load significantly. Therefore, when calculating relative performance, it is important to consider this change in turbine load reduction (which necessarily brings about lower pressure drop). To correct for this, we present performance of the baghouse as differential pressure standardised against turbine load in units of mbar/MW. This allows the direct comparison of two data points regardless of turbine load.

As shown in Figure 7.11, even when normalised against reduced turbine load, the 50%

change to Gutsche shows a significant (24%) and sustained (Lines are parallel) reduction in baghouse pressure drop. This demonstrates that the laboratory interpretation translates into tangible operational savings. A 100% change to Gutsche is likely to increase the potential savings.

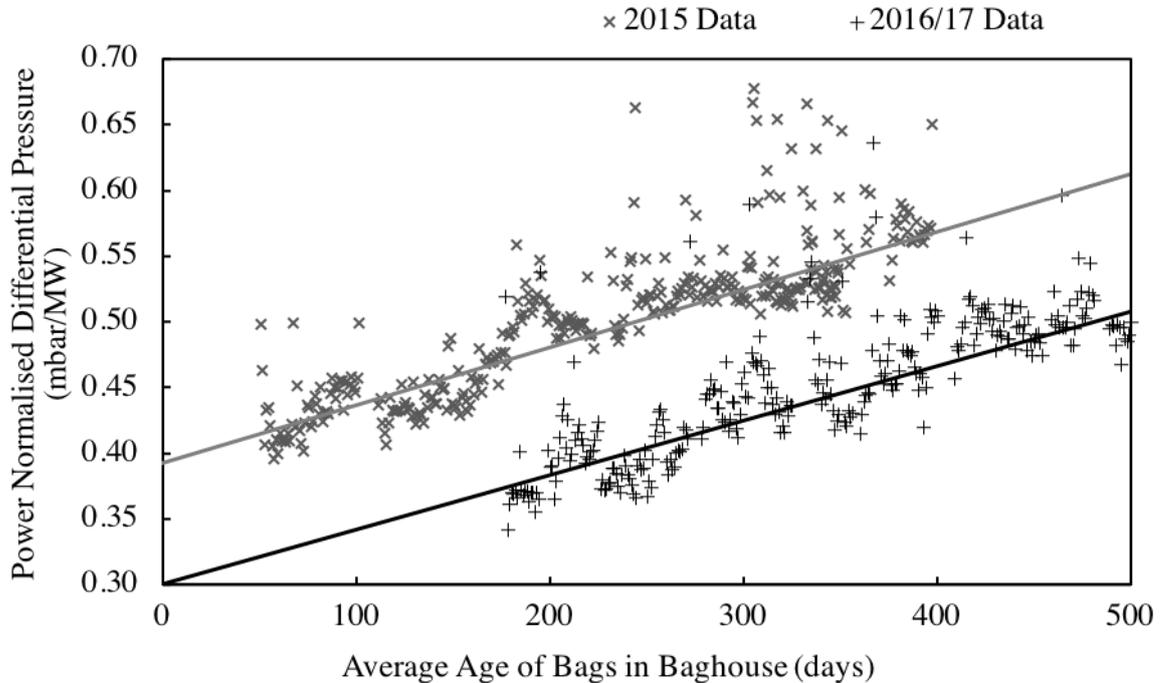


Figure 7.11: Scatter plot of normalised pressure drop (mbar/MW) versus average age of filters (days) for two brands of filter. 2015 Data: TTF bags. 2016/17 data: Partial change to G Bags.

7.4.2 Missing data modelling

Given the attitudes, budget and operational constraints discussed previously, filter baghouses are commonly controlled and monitored globally, giving rise to low spatial data resolution. Plant operators of large industrial power generation plants rely on one or two salient metrics to quickly assess the apparent health of a unit operation. However, if the goal of diagnosing baghouse performance, performing predictive and preventative maintenance and process optimisation is to be achieved, there is a clear need to heighten the level of detail to which data is collected and analysed.

Following the extended industrial campaign undertaken in the present work, several clear recommendations for increased data resolution are made below.

Differential Pressure

Differential pressure (DP) across the baghouse is a key parameter in diagnosing performance. Global DP (so-called ‘flange-to-flange’) is currently used to monitor the baghouse under

consideration in this study. High and low pressure impulse lines trace back to proximal locations above and below a central plenum, dividing the crude and clean gas streams. While this data is valuable in providing an overall impression, it relies on the sampling points being representative. Simulation has shown that local flow non-uniformities result in a non-uniform pressure distribution in practice.

The implication is that global DP is not sufficient to understand the degree of preferential flow between compartments. Increasing the spatial resolution of pressure measurements throughout the baghouse would allow for a more precise map of flow and local filter health to be established. Subtle operational changes could then be benchmarked in terms of improved performance.

Hopper Level

Hopper level is currently maintained between two hard-wired limits dictated by the fixed position of two level probes at the base of each hopper. While this system is adequate for controlling the level of the hopper through simple hi-lo deadband control, there is an opportunity for further process understanding to be gathered. The relative rate of increase in level for each hopper is indicative of the fly ash duty. Continuous, rather than discrete level monitoring would allow this revealing measurement to be made.

Results presented later in this work indicate that flue gas and fly ash duty are non-uniform not only in quantity, whether by mass or volume, but additionally the properties (density, particle size and composition) of the particulate differs across the compartments.

Additionally, a common design limitation found in vertical baghouses exists in the fact that the inlet geometry angles the crude gas stream towards the collected fly ash in the hopper with sufficient velocity that some of the already separated dust is re-entrained. This severely reduces the efficiency of the filters. This is worsened by the fact that the smaller, finer particles are more likely to be re-entrained; effectively distilling the retained particles to concentrate the fines, which is known to increase filter blinding, DP and parasitic load. Tighter level control may serve to reduce this effect. Continuous level monitoring would allow this behaviour to be diagnosed for each compartment, allowing for adjustments to re-distribute the burden more equally.

Particulate Monitoring

Currently, particulate emissions are monitored at the stack to meet environmental permitting regulations. The technology employed for this measure is perfectly adequate, and provides Plant 1 with a good raw signal of emissions. This raw signal must then be converted to mass flow by means of an isokinetic correlative model. The accuracy of the

reading is therefore reliant on the validity of the correlation. Furthermore, as with all optical detection methods, the measurement depends on a stream of particles passing the narrow beam.

As bags age and wear, small holes can quickly develop into large leaks. A leaky bag will concentrate flow to that area of the baghouse and therefore accelerates the wear of neighbouring bags. The result is that clusters of bags tend to wear in distinct patterns. Identification of these leaks is crucial in optimising the design, operation and maintenance of the baghouse.

Stack emissions spikes are currently used by the operators to identify leaks by roughly correlating the coincidence of a pulse from the cleaning system with a time delayed spike on particulate emissions at the stack. For the six-compartment baghouse, this means that any emission spikes are immediately diluted by a factor of six. The random variation of particulate across the compartments also multiplies the signal distortion by a factor of six. The result is a high signal-to-noise ratio and difficulty in spotting leaky bags early. Aside from the detection issue, transport delays mean that the exact location of the leak is never pinpointed. Enhanced local sensors situated at each compartment outlet would allow for far earlier detection of leaks, possibly allowing for a spot-change of bag long before it caused accelerated wear on neighbouring bags.

Pulse Cleaning System

The pulse cleaning system in the present study operates on a 2-mode sequencer; ‘normal’ mode and ‘over-pulse’ mode, where the pulse frequency is doubled by reducing the interval between sequential rows. The considered plant currently has no metering of baghouse compressed air usage. This is reportedly a common situation in industry. The compressed air is totalised further upstream and has many end-uses. The baghouse represents potentially the largest single use of compressed air in the plant and therefore carries a significant operating expense which is currently unquantified, and without recourse for optimisation. Air metering would allow for this gap in process understanding to be filled and give actionable metrics for optimisation.

Furthermore, there is currently no monitoring to allow the efficacy of the pulse cleaning to be monitored. Diaphragm valves and solenoids deteriorate over time. Pressure transmitters installed on each header tank would allow for the system response to solenoid and diaphragm impulse to be evaluated and deterioration to be diagnosed.

In the following section, this report aims to retrospectively quantify the compressed air consumption of the baghouse, notwithstanding the lack of data. Whilst a lengthy process, it serves to highlight several key areas of process understanding while also justifying the urgent need for compressed air metering to be installed on site.

7.4.3 Compressed Air Consumption

With interrogation of the plant data, it is possible to make an estimate of actual baghouse compressed air usage, compared to calculated values. It is noted that the units of compressed air consumption are given as '*normal thousand meters cubed per hour*' (NKm³/h).

Using plant-wide compressed air usage, along with time-aligned historic data, it is possible to isolate occurrences where the baghouse cleaning system was operational without other plant air users (e.g. during an offline clean at shutdown).

Figure 7.12A shows plant-wide compressed air usage plotted against ID Fan load (%). The fan load is used as a good indicator of an operational plant. Clearly, the highlighted cluster represents operational data - The ID Fan running at high load and compressed air usage in a high range. However, the lower left portion of the graph shows a small cluster of data points at 0% fan load and yet still a significant air use of ~ 0.42 NKm³/h. The sparse points in-between these two clusters represent transient operation (e.g. start-up/shutdown).

Cross referencing operational logs with the occurrence of the (0.42,0) cluster shows that the baghouse often underwent a period of extended offline cleaning (where the pulse cleaning system was allowed to continue post-shutdown in an effort to recover performance)

Figure 7.12B shows sorbent lime dose to the baghouse. Once more, operational values are clearly seen as a cluster. Transient values are also shown. In this case, highlighted region *i* indicates a plant operating at regular lime dose, and yet a compressed air demand below operational levels. This scenario is consistent with a baghouse in offline cleaning mode (after shutdown). The tail-off in lime dose at a sustained air duty is also observed (region *ii*). This is consistent with the shutdown of the lime dosing system and continuation of the baghouse pulsing. Regions *i* & *ii* show the baghouse cleaning air usage at ~ 0.42 NKm³/h. After lime addition is ceased, the tail-off shows consumption at ~ 0.4 NKm³/h.

Figure 7.12C shows operational baghouse differential pressure plotted against plant-wide compressed air consumption. The two groupings are divided by a differential pressure of 15 mbar. At this set point, the pulse cleaning system enters an 'over-pulse' mode (essentially, air consumption is doubled due to the increased pulse frequency). Although the data is noisy (due to the various other end users of compressed air), the increase in average compressed air consumption would be equal to the increased demand of the cleaning system as a result of over-pulsing. (since all other end users of air will vary by the same degree regardless of baghouse pressure). Since 'over-pulsing' doubles the air consumption of the baghouse relative to standard pulsing, the increase in average air consumption is equal to the standard consumption, and 2x the increase is equal to the over-pulsing consumption. The mean value of air consumption for the two groups (shown on the graph as lines) are 1.621 and 1.775 NKm³/h for below and above 15 mbar respectively. Thus, the

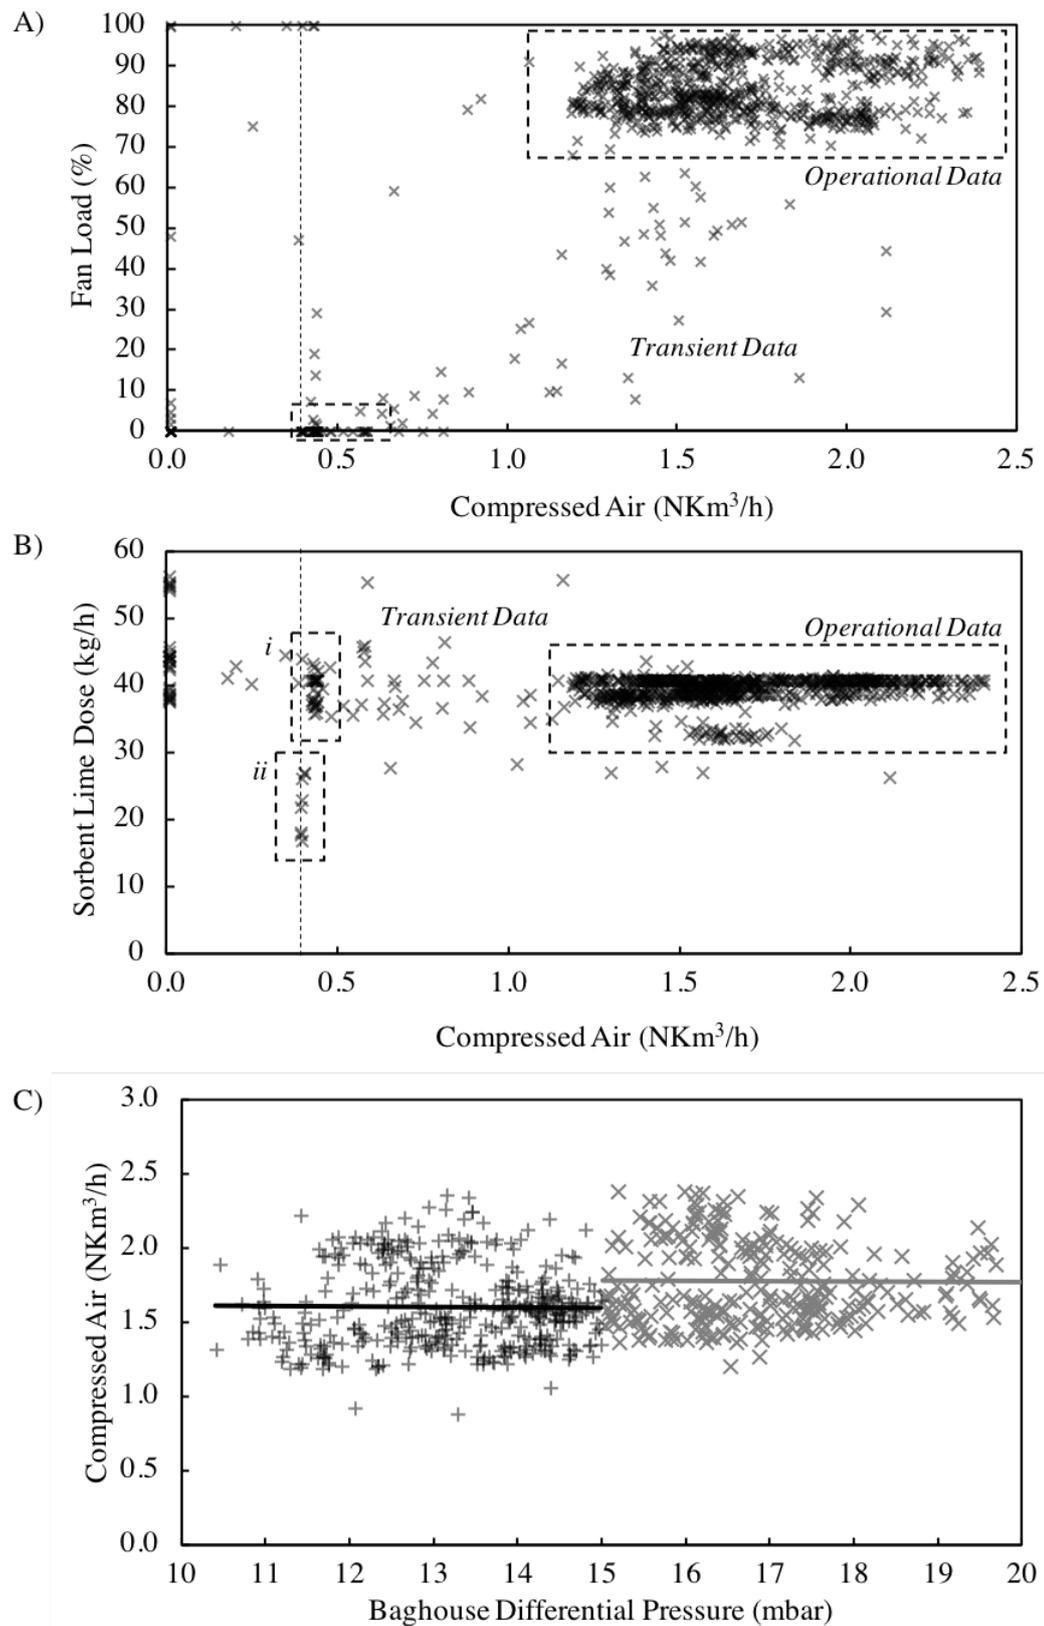


Figure 7.12: Operationally linked bivariate plots showing plant-wide compressed air consumption against various plant KPIs. **A)** Fan load (%). **B)** Sorbent Lime Dose (kg/h). **C)** Baghouse differential Pressure (mbar)

Table 7.2: Summary of data-derived values for baghouse compressed air consumption

Method	Value (NKm ³ /h)
Fan load	0.420
Lime dose	0.400
Pressure Drop (over-pulse)	0.154 (0.308)

difference in means for the two groups is 0.154 NKm³/h. Therefore, the air consumption for over-pulsing is approximately 0.306 NKm³/h.

So far, analysis of historic plant data has given three estimates of the compressed air consumption of the baghouse. The values are summarised in table 7.2. The difference between the first two estimates and the estimate given by the pressure drop calculation can be further explored. The first two estimates rely on the baghouse cleaning system data to be isolated from other end users based on absolute values of air consumption. Though the data is precise (all values within the cluster closely agree on the value being estimated) it suffers poor accuracy due to the many other destinations of compressed air which also correlate to the running of the plant in a given state. Furthermore, and importantly, these absolute values will also *include* any leakage of compressed air (i.e. from diaphragm valves) experienced by the system.

On the other hand, the pressure drop method arrives at an estimate by using relative, rather than absolute means. In this way, the result is accurate (the estimate is more specifically identifying baghouse air consumption) but lacks precision (there is a wide distribution either side of these mean values). However, since physical knowledge of the system allows us to explain this variance, more confidence in the means is afforded. Secondly, this relative measure does not include the leak rate (since the leak rate will be a constant function of system pressure, and not of pulse regime). Therefore, subtracting the relative from global estimates will also allow the leak rate to be estimated. All of these estimated values for air consumption and leak rate may then be compared with calculated values based on first principles.

In order to estimate the air consumption of the baghouse cleaning system from first principles, the following calculations can be made:

First by calculating the maximum possible volume of air per pulse, assuming a full header tank overpressure is drained:

$$V_{p,max} = V_h P_h = 0.05 \cdot 6 = 0.3 \quad [\text{Nm}^3] \quad (7.1)$$

where V_h is the header tank volume (0.05 m³) and P_h is the header tank pressure (6 bar(g)).

Now for a multi row pulse ($N_r = 2$), with a pulse interval, $T_i = 10\text{s}$, there are 720 pulses per hour ($F_p = 720$). Therefore the *maximum* air consumption per hour is given by:

$$Q_{max} = V_{p,max}F_p = 0.216 \quad [\text{NKm}^3/\text{h}] \quad (7.2)$$

This calculation gives a useful upper bound. In reality, however, the pulse valves introduce a choked flow regime which limits the possible flow rate.

n is the index of isentropic expansion. For an ideal gas undergoing an adiabatic process, n is given by the ratio of specific heats for constant pressure (C_p) and constant volume (C_v):

$$n = \frac{C_p}{C_v} = 1.4 \quad (7.3)$$

Given n , the ratio between critical pressure (P_c) and initial pressure (P_0) for a nozzle can be expressed as:

$$\frac{P_c}{P_0} = \left[\frac{2}{n+1} \right]^{\frac{n}{n-1}} = 0.528 \quad (7.4)$$

Now, for the given system with a valve flow factor (C) of $337.45 \text{ L}^3 \text{ s}^{-1} \text{ bar}$ ($K_v = 85 \text{ m}^3 \text{ h}^{-1}$) and a valve impulse time (T_o of 200 ms, the choked flow air consumption is calculated as:

$$Q_c = P_h C T_o \frac{P_c}{P_0} F_p = 0.154 \quad [\text{NKm}^3/\text{h}] \quad (7.5)$$

By following both data-driven and first-principles approaches, two values of air consumption have been derived. Firstly, the global data-driven approaches agree on the actual baghouse cleaning system compressed air consumption of $\sim 0.4 \text{ NKm}^3/\text{h}$. Next, both by relative and first-principles calculations, the pulse-cleaning compressed air consumption has been shown to be $0.154 \text{ NKm}^3/\text{h}$, when the baghouse differential pressure is below 15 mbar, and thus $0.308 \text{ NKm}^3/\text{h}$ during 'over-pulsing'. The difference between actual and pulse only consumption can be attributed to leaks.

Next, an economic calculation of compressed air consumption can be made. Compressed air in industry is typically generated on site by means of dedicated compressors. Thus, the *cost* of compressed air is principally the cost of power required to supply the required flow rate and pressure. For the system under consideration in this study, the cost for compressed air is given by the operator as $21 \text{ £}/\text{NKm}^3$. This represents a close approximation to the average found in industry, where typical values are given as $\sim 10\text{-}30 \text{ £}/\text{NKm}^3$.

Taking the global data for plant-wide compressed air consumption and assuming ~ 8000 hours of operation annually, the plant-wide compressed air consumption is calculated as

12023 NKm³/yr and £252,491/yr.

In a similar way the operating cost for baghouse compressed air can be calculated for the various cases calculated previously. These are summarised in table 7.3.

Table 7.3: Summary of calculated operating expenditure for baghouse compressed air consumption

Pulse mode	OPEX (£/yr)
Normal	25,872
over-pulse	51,744
Actual (data)	67,200

This calculation shows that the total cost of air leaks in the baghouse are in the region of £15,456 - £41,328 per annum, depending on the degree of over-pulsing for a given year.

This report has demonstrated that the compressed air consumption of the baghouse is a considerable operational expense for the given system. With this in mind, an air metering device should be installed immediately upstream of the baghouse to directly measure compressed air consumption.

The empirical and calculated approaches shown here have agreed upon values of air consumption, and have shown evidence for a significant compressed air leak. This is most likely to be one or more of the diaphragms, which are known throughout the industry to pose a significant failure point (144 failure points in the considered Baghouse).

Compressed air consumption plant-wide represents \sim £250,000 in OPEX. The baghouse cleaning system accounts for \sim £67,200 of this total (27%). Of this, only \sim £25k is usefully spent cleaning the bags (in normal pulse mode). The remainder is a result of the over-pulsing due to high baghouse pressure drop and leaks (likely due to ageing diaphragms.) Fixing air leaks, optimising pulse settings and maintaining \leq 15 mbar pressure drop is therefore able to save in the region of \sim £54,000 per year in compressed air costs alone. Coupled with the considerable parasitic load savings achieved by doing the same (quantified later), the benefits to optimising the cleaning system hardware and control software are substantial.

In terms of the wider work, this small calculation serves two purposes. Firstly, it highlights that a hybrid approach to calculations (empirical and first principles) is useful and, where there is missing data, often *necessary* to arrive at a reasonable level of confidence in an inferred measurement. Secondly, the calculation allows a further and wider exercise of hybrid modelling to progress in this work with the inclusion of air consumption into a cost function.

In order to quantify the benefit of reduced compressed air consumption gained by optimising pulse sequencing and fixing leaks, it is necessary that dedicated air metering is installed

on the Plant 1 baghouse cleaning system. The full extent of economic benefits will only be quantified with this metering in place.

In general, the solution must meet the following criteria:

Fast response time

Since the cleaning system does not demand constant flow, but very short pulses with rapid on/off times, the flow in the feed line will also be transient. It is important that any air flow meter is able to respond quickly enough to measure the whole pulse of air as it passes the meter. A slow meter will underestimate the total flow since air will pass before the instrument responds.

DCS interface

It would be preferential to have an instrument which can be linked to the control room for monitoring by the operators and data-logging for diagnostics and trending. Without this historical trend capability, the meter would require frequent manual visits to the local gauge in order to build even the most low-resolution trend. With high resolution, real time digital data collection, suboptimal performance can be identified.

Appropriate Range

Since the current true air consumption is unknown, we have made estimates outlined previously in this report. A meter ranged to 0 - 0.15 Nm³/s. is likely to encompass the relevant flow range, while still providing the required accuracy.

7.4.4 Sorbent Dose

Sorbent is currently dosed by a screw feed mixer, combining the output from two silos (one each for PAC and Lime).

The mass flow rate is therefore reliant on good correlation with screw speed. This relationship does not always hold if air gaps are present, or as the screw deteriorates by fouling or wear between services. The exact dosage of sorbent is therefore only known to within a margin of error. In any case, no routine measurements of unreacted sorbent are currently undertaken to assess the dose. Both a direct measure of sorbent mass flow and a routine measure of unreacted sorbent would enhance Plant 1's ability to optimise dosage. Optimum dose will minimise not only cost, but also emissions.

Dry Sorbent Injection (DSI) accounts for a considerable operating expense (OPEX) at the studied plant. Table 7.4 shows the total cost of sorbent material added (2015 budget year shown). In recent years, attempts had been made to minimise this cost by reducing the dose of DSI. The reduction in dose is also partially due to reduced plant utilisation over this time. Although the site did not see any appreciable difference in stack emissions of

HCl or heavy metals, a reported increased rate of corrosion in the downstream duct and ID fan could potentially be attributed to a dose reduction.

Table 7.4: Economic estimates of sorbent cost

Sorbent	Unit cost (£/T)	Dose (T/yr)	Annual cost (£/yr)
Lime	143	334	47,762
PAC	620	42	26,040
Total			73,802

The dose reduction exercise illustrates two important points; Firstly, the site was clearly overdosing DSI in the first instance. If DSI was added in a stoichiometric ratio to flue gas contaminants, a dose reduction would show an impact on stack emissions (assuming complete reaction). Secondly, the downstream corrosion impact shows that stack emissions alone are not a good measure of contaminant slip (contaminants evading the dry scrubber). Downstream metalwork may be acting as a ‘fine scrubber’ to remove any contaminant slip by metal corrosion. (e.g. HCl condensing or adsorbing onto steel).

From this, it is apparent that a systematic method for reduction of DSI dose is required, whilst also ensuring adequate dose to ensure that corrosion is minimised. One such approach is the so-called free lime determination method. Lime dose is reduced in small increments, and allowed to equilibrate between each change. After an equilibration time, the fly ash collected in the hopper is sampled and titrated with acid. The percentage of unreacted lime in the fly ash is calculated from the amount of acid consumed in the titration. This ‘free lime’ is present in the fly ash because it has not reacted with contaminants in the dry scrubber. Free lime therefore represents an overdose of DSI and therefore wasted OPEX.

Figure 7.13A(i) shows this response to a reduced lime dose (D). Minimisation of free lime (F) occurs by a step-wise reduction in D . At D_{min} , $F = 0$. An operational dose D_O should be chosen such that $D_O > D_{min}$ and $F \rightarrow F_{lim}$. ($D_O \leq D_{min}$ is not desired since this will cause incomplete scrubbing of contaminants.) D_{min} can intuitively be defined as the stoichiometric dose of DSI, plus some factor to correct for imperfect mixing and incomplete reaction. The mixing of DSI and consequent path length can be summarised by calculating the residence time, t_d , of DSI in the duct.

Figure 7.13A(ii) shows the response of DSI dose reduction to stack HCl emissions (H). As shown, for some dose reduction $D_1 \rightarrow D_2$, the response in H may be relatively subtle. In Figure 7.13A.iii, the increased corrosion rate due to a greater emissions slip is shown. Any systematic approach to DSI dose reduction must therefore also consider these factors. D_2 may produce $H < H_{lim}$ (environmental permit limit), yet yield an unacceptably high corrosion rate ($C_2 > C_{lim}$).

One issue identified with the current system is the poor mixing of sorbent in the duct.

CFD simulations of the DSI injection regime was carried out to understand the mixing and any potential for improvement.

Figure 7.14A shows the simulated path of sorbent under operational conditions. It demonstrates that poor mixing of the sorbent is currently achieved. Most of the sorbent particles follow a streamlined path with little chance to interact with contaminants in the duct. By the downstream end of the duct, most of the sorbent has lost all momentum and settles to the bottom right hand side of the outlet. In operation, this would imply a significantly non-uniform dose of DSI between compartments. Indeed, filter maintenance records show a dramatically shorter life in the first right-hand compartment, B1. A local overdose of lime in this compartment, as predicted by the CFD model, is plausibly responsible for accelerated wear or filter blinding.

An increased residence time can be achieved by increasing the degree of mixing in the duct. Residence time of the sorbent particles is a measure of the degree of mixing. High residence time indicates a thoroughly mixed stream, giving maximum time and distance for sorbent particles to physically interact with the flue gas stream. Therefore, by simply increasing mixing, the performance of the dry scrubber can be greatly enhanced. Furthermore, a more evenly distributed dose of sorbent will coat filter bags in each compartment more uniformly, yielding lower overall pressure drop and greater filter life.

Figure 7.14B shows the simulated path of sorbent when an enhanced boosted air dosing nozzle is used. One commercially available nozzle, shown in Figure 7.15, includes a motorised fan, which both increases air flow (and therefore velocity) and introduces rotational velocity into the DSI stream. Combined, these effects dramatically increase turbulence of the DSI. Figure 7.14B shows a significantly increased residence time of sorbent due to increased turbulence.

As an alternative to the boosted air dosing nozzle, the whole process stream may be mixed by inclusion of a static mixer. Several commercial designs of mixer exist, and some with the express purpose of mixing solids and gasses in stream. One such 'air blender' is shown in figure 7.16, the static mixer has concentric regions of counter-rotational baffles. this causes a 2-directional swirl . Within a given number of downstream duct diameters, the stream is completely mixed. The disadvantage of static mixers is the inevitable introduction of pressure loss, and therefore increased energy costs. Any pressure loss must be offset

Figure 7.17 Compares the mixing behaviour for the current system compared with two design options for static mixing in the inlet duct.

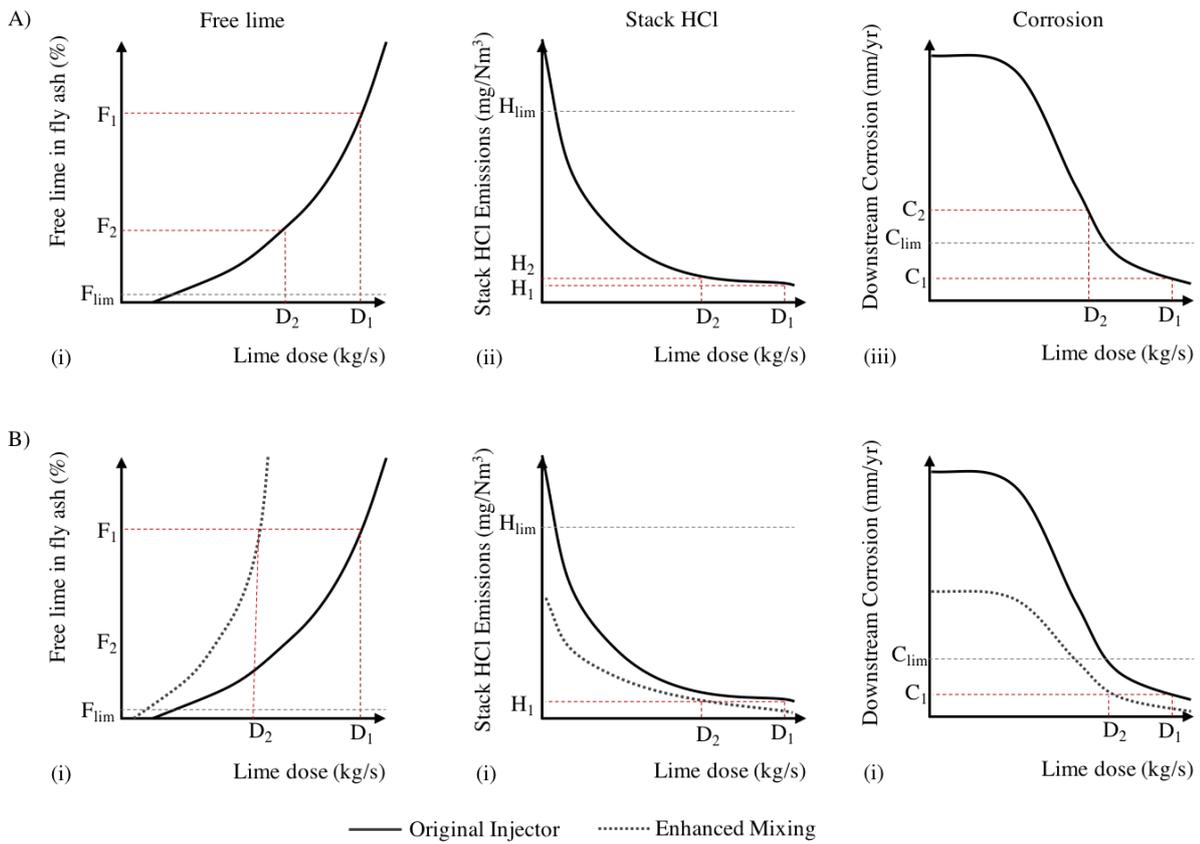


Figure 7.13: Schematic graphs to show system response to lime dose (D) (kg/s) reduction and enhanced mixing

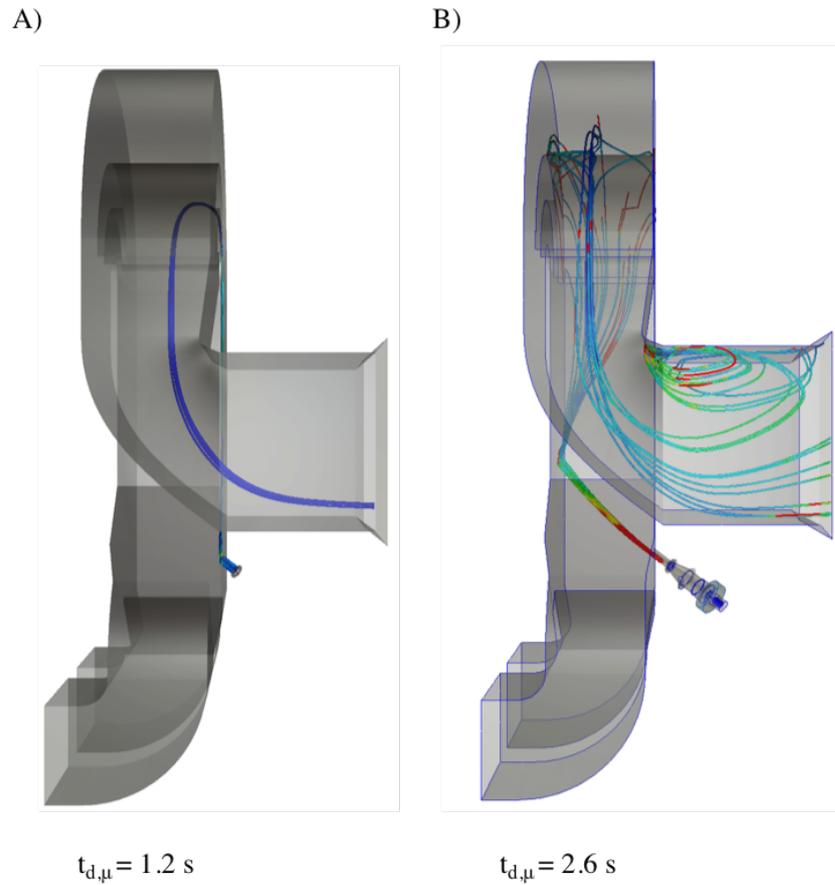


Figure 7.14: CFD generated particle traces showing path of sorbent in the duct upstream of the baghouse. A) The as-built system ($t_{d,\mu} = 1.2$ s). B) Enhanced mixing through boosted air nozzle ($t_{d,\mu} = 2.6$ s)

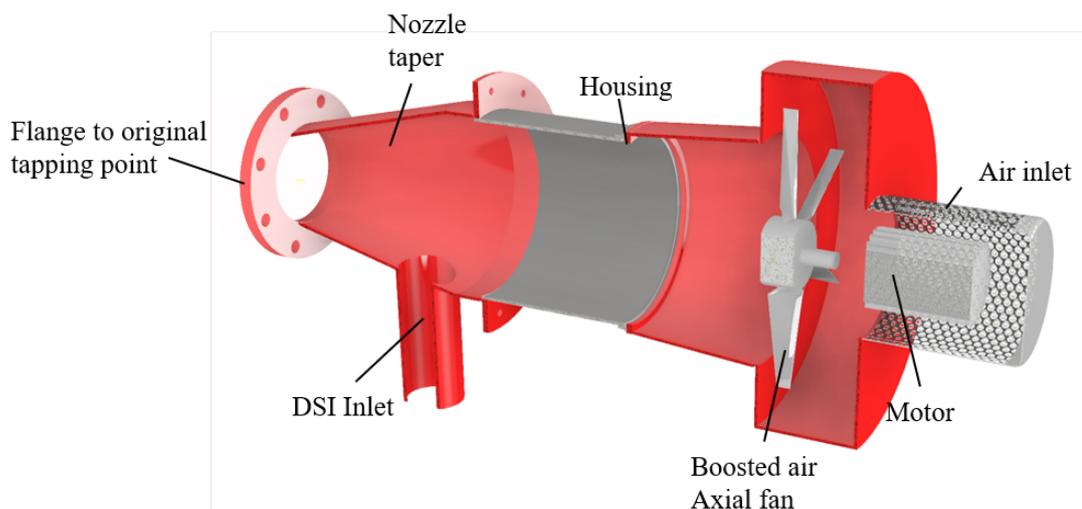


Figure 7.15: Boosted air mixing nozzle designed for enhanced sorbent mixing.

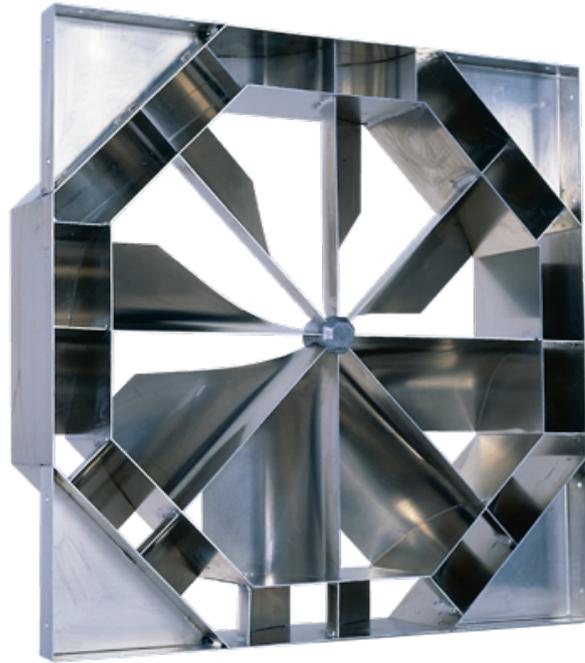


Figure 7.16: Static mixing element designed for enhanced sorbent mixing.

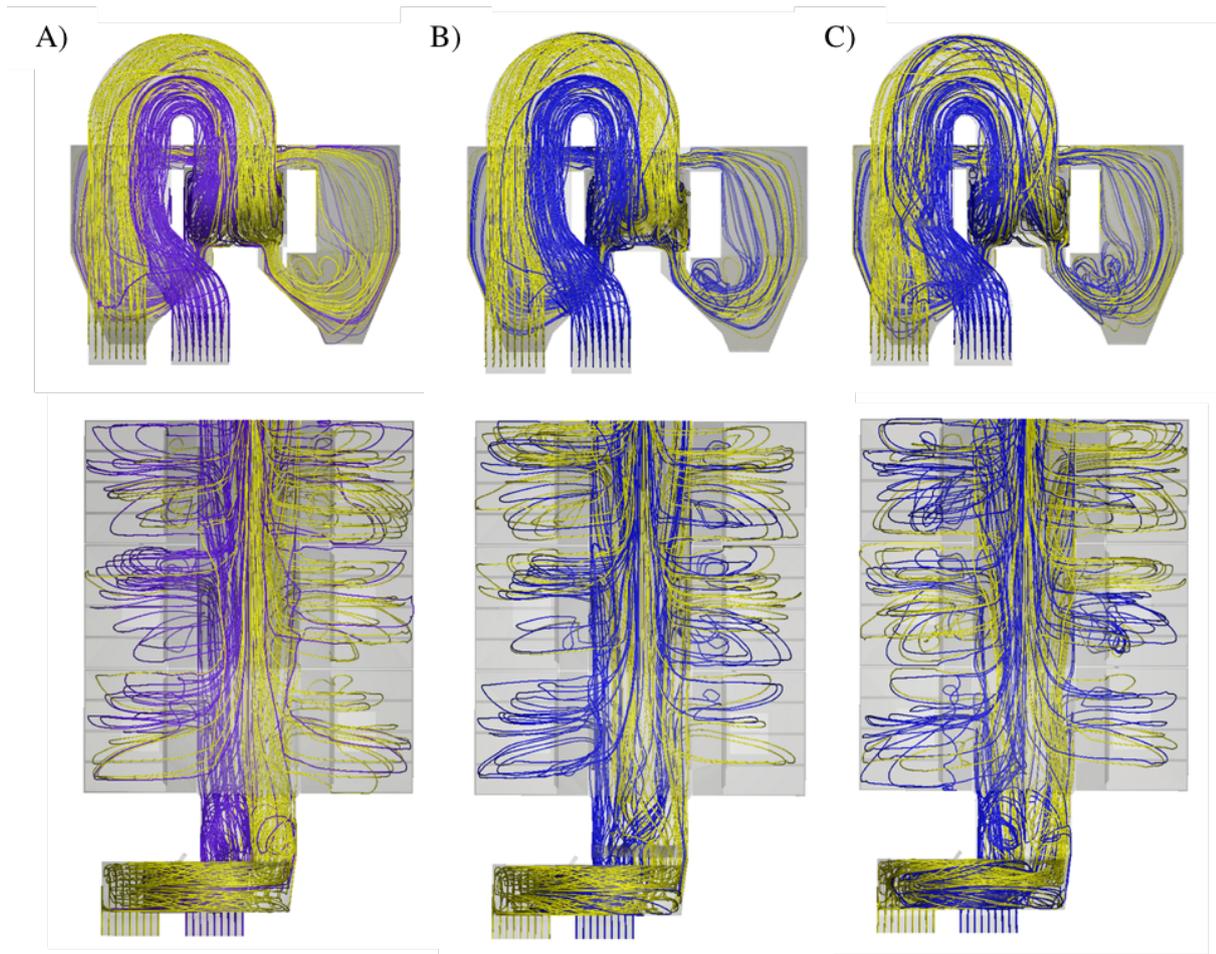


Figure 7.17: CFD generated particle traces showing path of sorbent in the baghouse showing transverse mixing. A) as-built system. B) Horizontally mounted static mixer in baghouse inlet plenum C) Vertically mounted static mixer in upstream duct.

7.4.5 Pulse Sequencing

Plant 1 currently pulses at factory default settings. Each diaphragm valve cleans two adjacent rows of bags. Then, after a fixed period of time, the adjacent valve pulses the next two rows of bags. This sequence continues down the length of the baghouse.

When a bag is pulsed, the bound dust is shaken loose. There is still a continuous upwards flow of gas, which opposes the pulsed dust in its fall to the hopper. Some of the pulsed dust is therefore entrained, and is immediately pulled back to the filter material. The dust most likely to stay entrained is the lighter, finer dust. Under the current sequence, the pulsed bag is directly next to the bag most previously cleaned. In this scenario, dust and air from the pulsed bag takes the path of least resistance and the most previously cleaned bag receives the majority of the pulsed bag's entrained dust. Since the previously cleaned bag is unprotected by a filter cake, the fine entrained particles are able to engrain themselves into the fabric and cause filter blinding. This effect is illustrated in Figure 7.18A.

In contrast, many improved pulse sequences can be developed. The simplest case is the staggered skip-2 sequence, where two rows are kept between the most recently cleaned row and the currently cleaned row. This arrangement is also shown in Figure 7.18B.

This simple staggered sequence is effective at reducing re-entrainment issues. However, the Plant 1 baghouse has the added complexity of being a multi-compartment with multi-row pulsing. Simple sequence modes introduce preferential flow dynamics which increase bag wear in certain locations. For stable operation, more advanced, dynamically adjusting pulse settings may prove an additional benefit.

7.4.6 Cost Function Development

In order to assess the impact of design and operational changes on the total cost of ownership, it is first necessary to understand both the revenue sources and *direct* operating cost within the scope of the filtration system. These are briefly summarised in table 7.5 and further explored below.

The energy required to power the plant equipment (fans, valves, sensors, heaters, screws etc.) is termed parasitic load. The Induced Draft (ID) fan is by far the largest contributor to parasitic load on plant. Because of this, parasitic load is a strong function of pressure drop across the baghouse- since the ID fan has to work harder to overcome the resistance to flow when pressure drop is high. The penalty for high parasitic load is considerable at Plant 1, due to the double Renewable Obligation Certificate (ROC) index. For each MWh of power sold to the grid, a further 2 ROCs are awarded. Current strike prices and ROC values mean that the value of a MWh is approximately tripled. Parasitic load is therefore

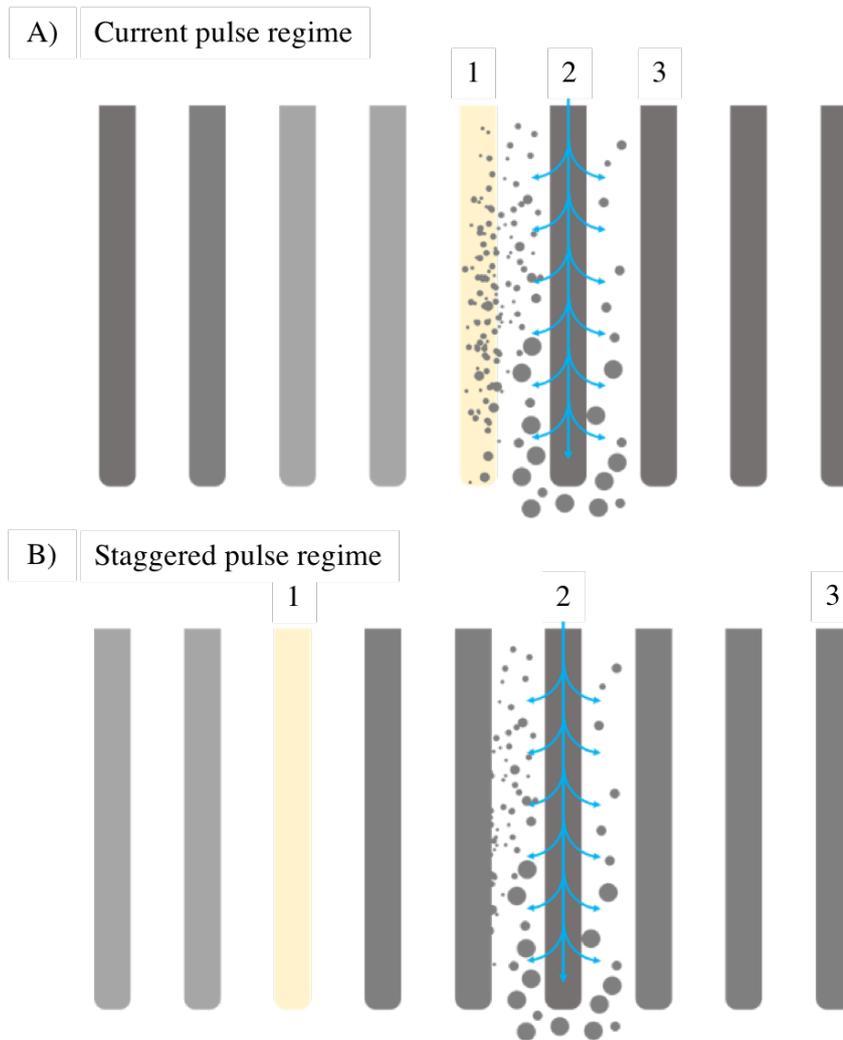


Figure 7.18: Comparison of sequential and skip pulse sequences showing path of re-entrained particles. **A)** Sequential (current) pulse regime. **B)** Skip-2 pulse regime

a loss in both MW value and 2x ROCs, as it is power which has been produced but not sold.

Parasitic load is clearly related to the lifecycle of filter bags, as well as the general flow of flue gas through the filter. As a bag filters fly ash and sorbent from the flue gas, a dust cake forms on the surface. The pulse cleaning system removes this cake periodically to expose a clean fabric surface for fresh cake formation. Over a sequence of many pulses over time, mechanical wear loosens the felt structure of the filter fabric, allowing fine particles to ingress into the felt. Pulse cleaning continues to remove some of the surface cake, but the fine particle ingress is difficult to remove. Over time, ever more particles work deeper into the felt and cause a significant resistance to flow. At this point the cleaning system becomes even less effective, since the pulse must overcome the resistance of the clogged fabric in order to remove the cake. This snowball effect is known as filter blinding, and results in high differential pressure across the filter fabric. The differential pressure over time (DP/t) profile results in a steady increase of parasitic load over time, as the ID fan works harder to overcome the pressure drop of the filter.

Table 7.5: Summary of revenues and operating costs considered for Plant 1

Type	Source	Impact (approx.)
Revenues	MWh Output	£40/MWh
	ROCs ($\times 2$)	£90/MWh
	Steam Output	customer dependent
OPEX	parasitic load	£130/MWh
	filter replacement	£100-200k/yr
	compressed air	£250k/yr
	sorbent	£70k/yr
	reliability	>100k/day shutdown cost

Whilst a large proportion of baghouse pressure drop is related to the fabric pressure drop, there is also a ‘static’ pressure drop across the system. The static pressure drop is the resistance to flow caused by the convoluted flow path the air must take to traverse the filters. This flow path is largely unchangeable within the current scope of improvements and must therefore be seen as a fixed, unavoidable loss.

In order to calculate the static pressure drop of the filtration system, it is simulated in 3D computational fluid dynamics software with real world flow conditions. The benefit of CFD simulations is that hypothetical measurements can be taken with high repeatability. For the static pressure drop, the operational volumetric flow-rate travelling through an ‘empty’ baghouse – with no filter bags is simulated.

The results of the simulation show a 5 mbar static pressure drop through the inlet duct and baghouse when empty. This 5 mbar static pressure drop can be included in all further calculations to ensure that possible improvements made by various modifications are not overestimated. For example, an improvement which is likely to achieve a 10% reduction in DP from 15 mbar would not yield a reduction of 1.5 mbar, but 1 mbar, since it is only possible to improve the dynamic pressure drop ($15 - 5 = 10$ mbar).

Figure 7.19 shows the relation between plant parasitic load and baghouse DP. Since the strike price and ROC price are known, the lost revenue due to parasitic load may be calculated directly. Therefore, the lost revenue due to increasing DP can be calculated. Figure 7.20 shows the substantial parasitic load savings achieved by reducing the dynamic pressure drop of the baghouse. For instance, a 10% reduction in dynamic pressure drop can save £36,674 per year. A 40% reduction saves £182,344 and so on.

These potential savings can now be used to benchmark the relative benefit of any proposed modification to the filtration system. For example, a modification which results in a 10% reduction in DP and costs £36k to complete will have a payback time of 1 year.

In terms of cost function, the OPEX associated with *raw materials* (sorbent, compressed air, filters), is assumed to observe a linear relationship with their consumption. Ignoring minor

fluctuations in commodity pricing and volume dependent pricing scales, this assumption holds true for Plant 1. Thus, if a modification was predicted to provide a 40% reduction in sorbent consumption, a 40% reduction in OPEX associated with sorbent would be incurred.

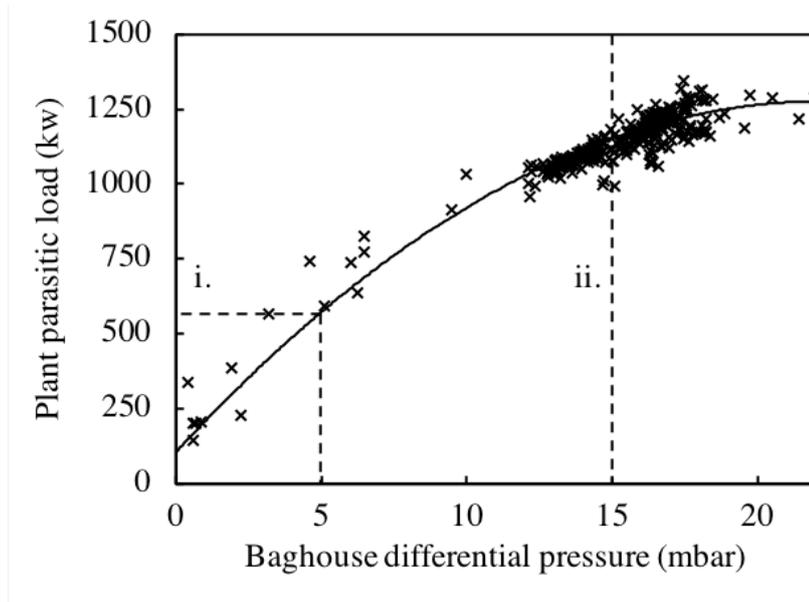


Figure 7.19: Bivariate plot of measured baghouse differential pressure (mbar) versus parasitic load (kW)(time averaging 24 h, shutdown data removed). i.) 5 mbar static pressure drop. ii.) 15 mbar pulse cleaning threshold.

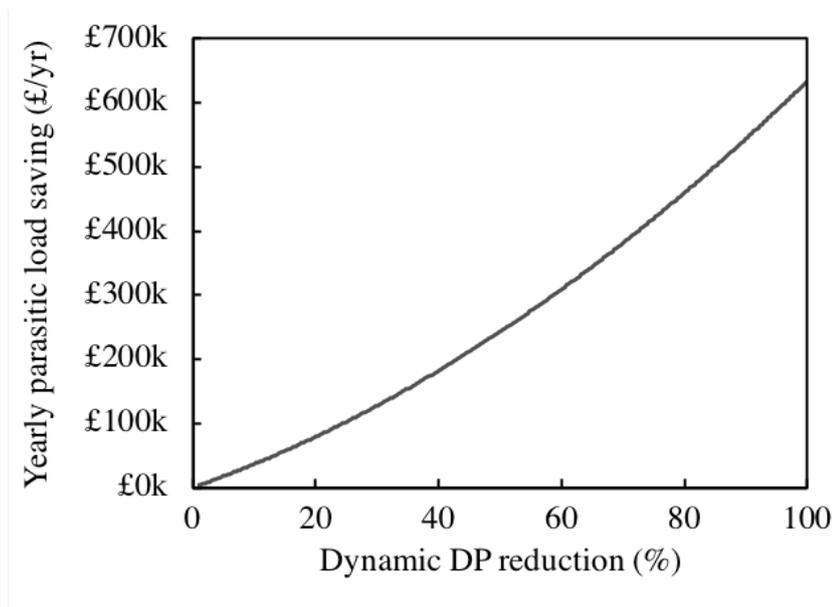


Figure 7.20: Bivariate plot of predicted baghouse Dynamic DP reduction (% of mean) (less 5 mbar static DP) versus predicted yearly parasitic load savings (£/yr).

7.4.7 Hybrid Computational Fluid Dynamics Optimisation

A fundamental issue in optimum baghouse design and operation is understanding the subtle dynamics of fluid flow through the system. Seemingly innocuous design decisions made by the OEM can have significant economic or operational impact on the real-world plant.

Through an initial CFD simulation of the as-built system in the current study, it is possible to characterise and design-out the various sub-optimal flow patterns present within the system. Figure 7.21 gives a visual representation of the air traces within the simulated baghouse geometry.

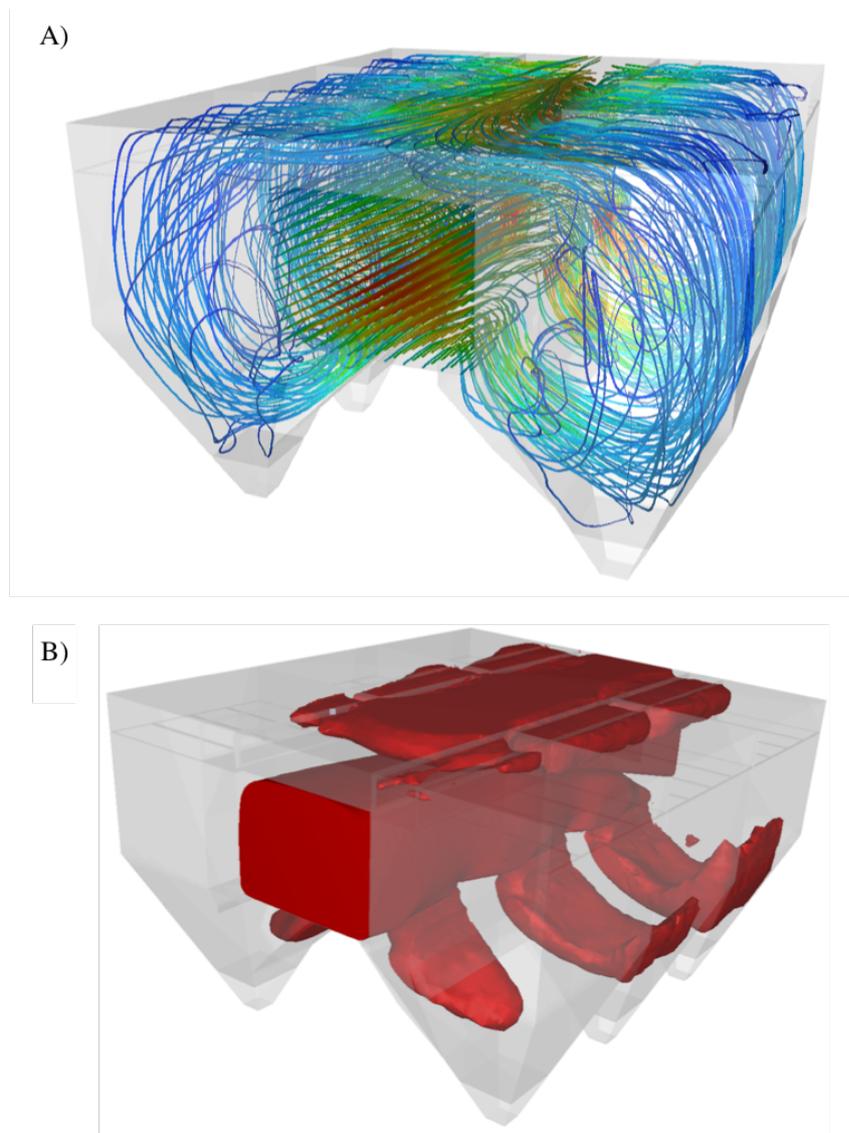


Figure 7.21: 3-Dimensional views of steady-state CFD results, showing **A)** Air traces coloured by velocity magnitude. **B)** Isometric volume with surface of constant velocity (>10 m/s).

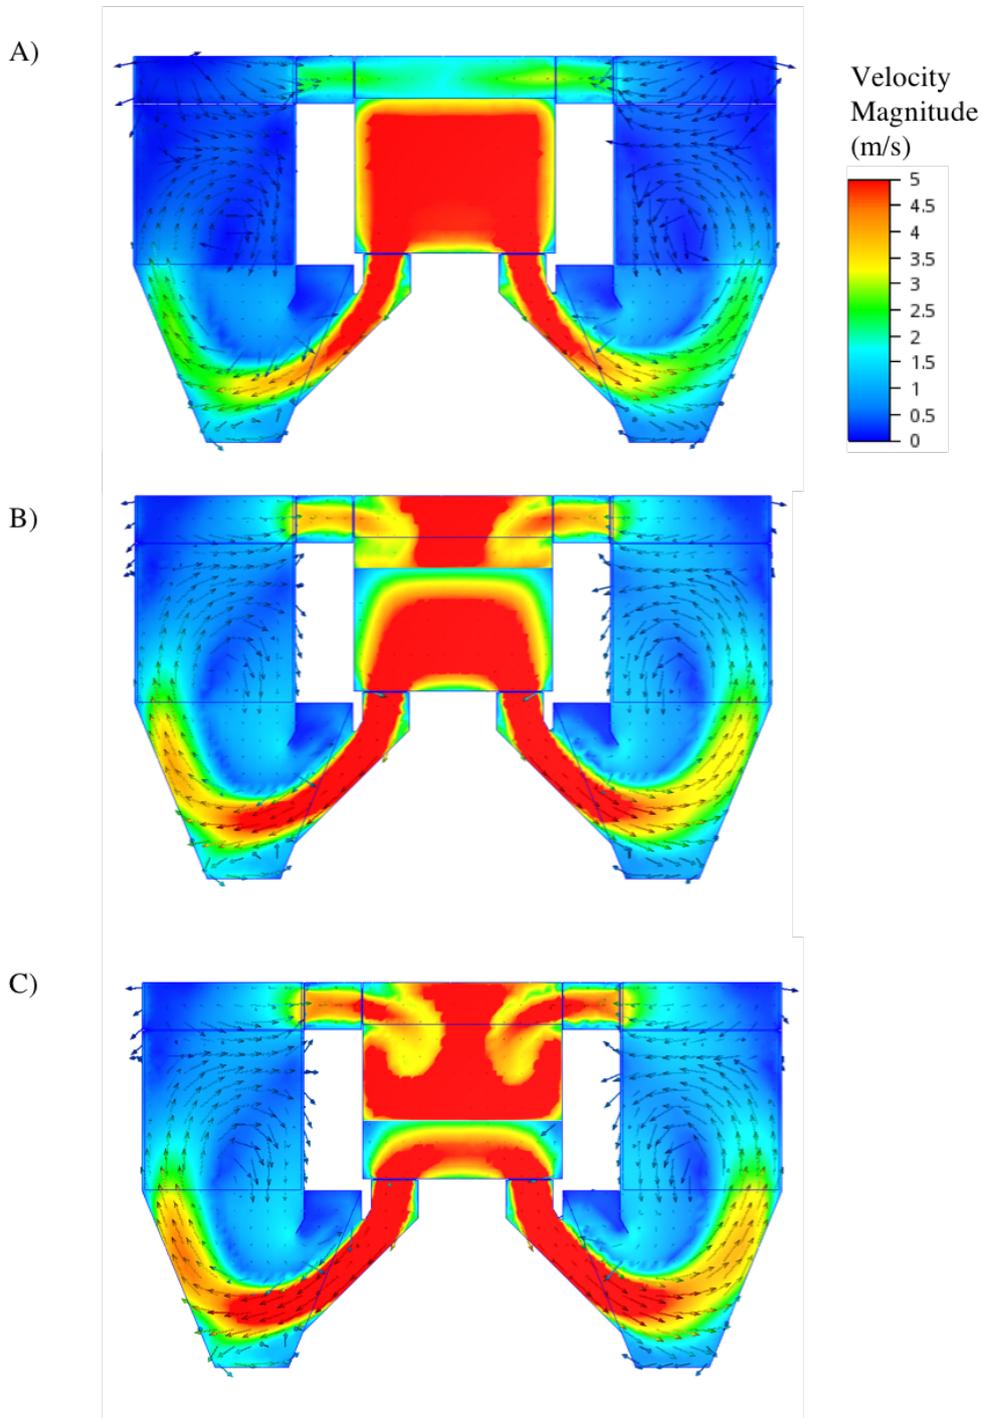


Figure 7.22: Planar cross sectional view of CFD generated velocity magnitude (colour scale) (0-5 m/s) and vector (arrows) for 3 axial cross sections of Plant 1 Baghouse (filters removed). A) Compartments A1 & B1. B) Compartments A2 & B2. C) Compartments A3 & B3.

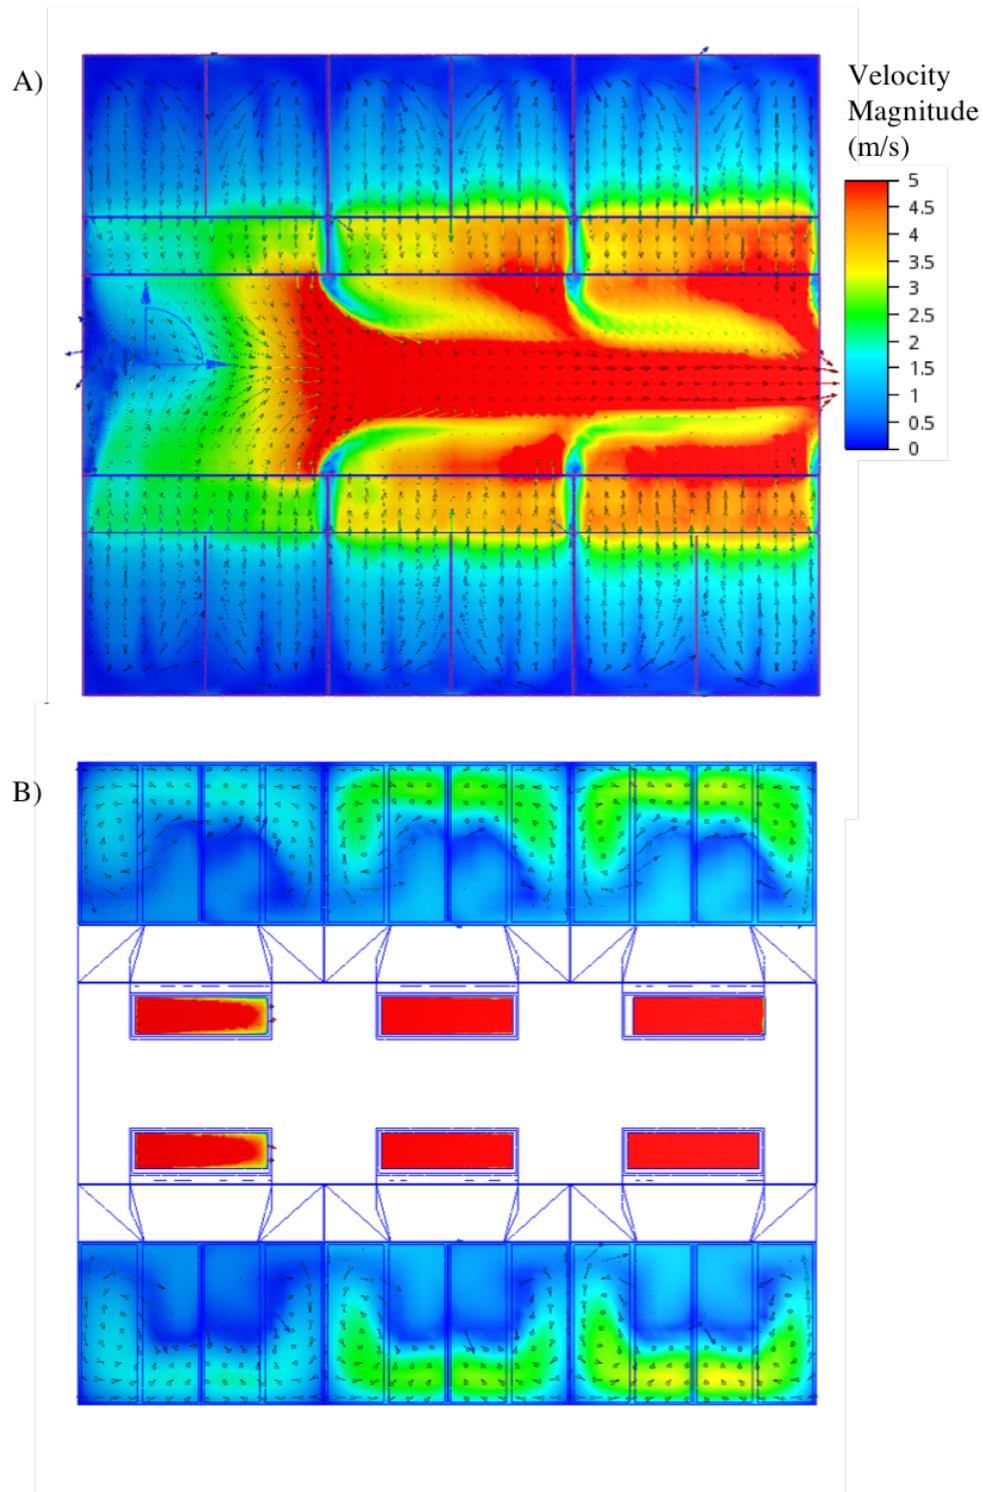


Figure 7.23: Plan view of CFD generated velocity magnitude (colour scale) (0-5 m/s) and vector (arrows) for 2 vertical cross sections of Plant 1 Baghouse (filters removed). A) Vertical mid-plane of compartment outlet plenum. B) Vertical plane of compartment inlet face.

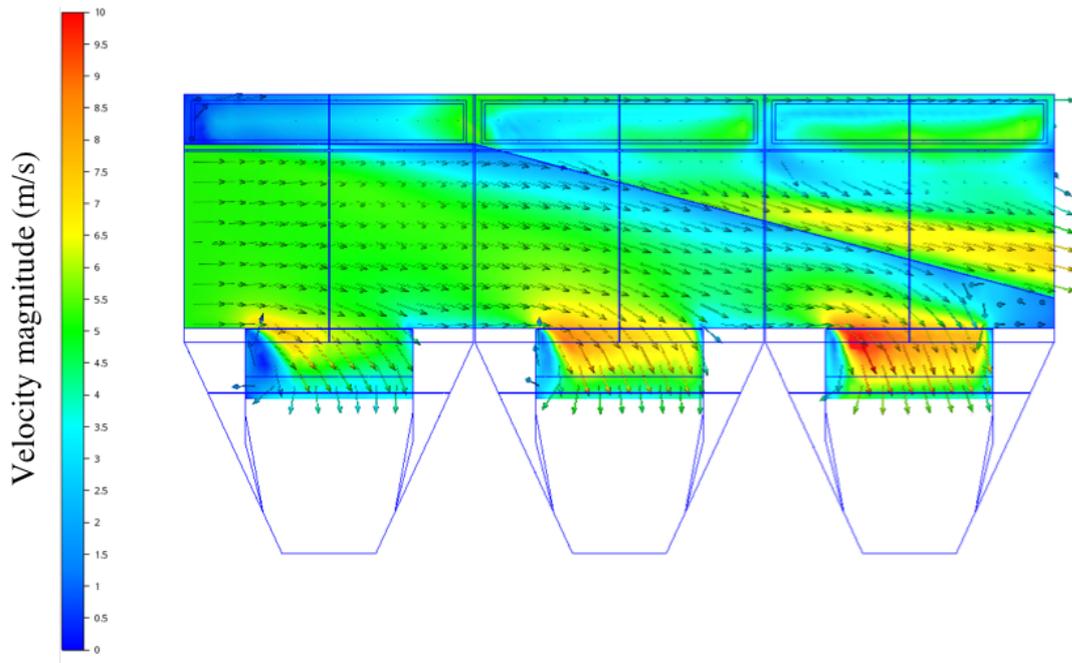


Figure 7.24: Side elevation view of CFD generated velocity magnitude (colour scale) (0-10 m/s) and vector (arrows) for transverse mid-plane of the common inlet/outlet plenum.

Flow Issues

Figure 7.25 shows a cross section of a single compartment of the baghouse under investigation. Issues identified by CFD analysis are schematically represented.

Flow uniformity is an important end goal for baghouse design. Flow uniformity occurs both system-wide and on a more localised level within the filtration system. Uniformity also changes with respect to time.

Compartment Duty

Proper duct geometry design and mixing gives rise to equal volumetric flow rate of flue gas for each compartment. This ensures that each compartment ages at the same rate. CFD analysis has shown a subtle imbalance of volumetric flowrate between compartments in the Plant 1 baghouse, giving rise to differential wear between compartments.

Sorbent Distribution

Good mixing and uniform flow ensures that each filter receives the same dose of sorbent. The sorbent performs the duty of protecting the filters from blinding, while also acting as a catalytic bed to remove trace acid gasses and heavy metals. Poor mixing has been shown to introduce significant non-uniformity both between compartments and within individual compartments. In short, while the dose of sorbent is high, a large percentage of bags receive far too little dose, while others receive far too much. This results in higher emissions and dramatically shorter bag life.

Lateral Uniformity

Referring to the flow uniformity within compartments, it is shown by CFD analysis that at operating conditions, filters at the far wall from the central duct observe a greater velocity and volume of flue gas, and consequently a higher duty of fly ash. Disadvantageously, the filters at this wide lateral position are also furthest from the cleaning system, resulting in diminished cleaning performance. The compound effect is reduced filter life, greater wear rate and increased pressure drop.

Time Dependent Uniformity (short term)

Over a short term, as the cleaning sequence processes through the axial length of the baghouse, there exists a continually shifting flow non-uniformity. Since a flowing gas will follow the path of least resistance, a greater volumetric flowrate of flue gas will pass through the most recently cleaned rows. Since this changes over time, there is a continual change of balance within the system.

Time Dependent Uniformity (long term)

Owing to the local non-uniformities mentioned above, particular areas of the baghouse are prone to suffer accelerated wear, which can lead to small failures in the filter fabric. At this point, a new path of least resistance is created, which increases the local flow of flue gas and fly ash. The increased local flow further accelerates wear rates in this zone, and the failures quickly spread to nearby filters. This problem has been shown to snowball at Plant 1, resulting in compartment B1 consistently showing bag failures, which have required spot-changes or whole compartment changes at half the usual expected filter life.

Due to flow non-uniformity, a number of issues are observed in the Plant 1 filtration system:

Bottom Bag Abrasion

Due to the widening of the inlet into the hopper, and the apparent narrowing of the hopper in both the downwards and outwards directions, crude gas is accelerated towards the far wall. This causes a high velocity cross flow along the bottom of the bags, resulting in premature erosion by the particulate laden gas stream. Considering the wear pattern and maintenance records for the Plant 1 Baghouse, this correlates with areas of high wear and bag failure.

Back wall impingement

As a further result of the inlet geometry, flue gas is forced to the outer wall, where it loses much of its energy to turbulence. This further increases erosion of the outermost bags, but also concentrates the accumulation of dust and sorbent. This disproportionate duty of the rear bags, coupled with the increased wear rate means that only a small proportion

of the baghouse is actually working at design conditions. Much of the baghouse is either far above or far below the design value and consequently operating far below the optimum in terms of pressure drop, parasitic load, cleaning performance, bag life and emissions performance.

Hopper re-entrainment

Due to the downwards direction of the inlet duct, along with the cross-sectional area reduction towards the hopper base, crude gas is accelerated towards the bottom of the hopper with sufficient velocity to re-entrain already separated particles. This effectively increases the duty of the filters as the same dust is in constant recirculation.

High Can Velocity

The upwards velocity of flue gas in the compartment. The cross-sectional area of the compartment minus the combined cross sectional area of the filter elements dictate the can velocity. At lower can velocity, separation efficiency is higher and dust entrainment is lower. Lower velocity also results in a lower pressure drop.

It is possible to reduce can velocity by three means; firstly, one can increase the cross-sectional area of the baghouse. Secondly, by reducing the cross-sectional area of the filter elements and finally by reducing the volumetric flow rate (i.e. By adding more compartments). In practice, however by improving flow uniformity, local high can-velocity areas can be eliminated.

7.4.8 Pleated Cartridge Conversion

As discussed previously, Plant 1 suffers from bottom bag abrasion, wall impingement, ash re-entrainment, low drop-out, high can velocity and high filtration velocity. Figure 7.26 shows a comparison between bag filters and pleated cartridges, highlighting the difference in flow behaviour between the two.

With reference to figure 7.26, the benefits of upgrading to pleated cartridges are as follows:

Larger Drop-out Zone

For the same filtration area, a pleated cartridge can be much shorter. For Plant 1, 1 meter of additional drop out zone can be achieved. This will allow more fly ash to disengage from the flow before it reaches the filters. In effect, the fly ash duty observed by the filters is reduced. The relative rate of filter blinding will therefore be reduced, extending filter life and reducing pressure drop.

Eliminate Bottom Bag Abrasion

The shorter cartridges will be completely withdrawn from the high velocity abrasion zone.

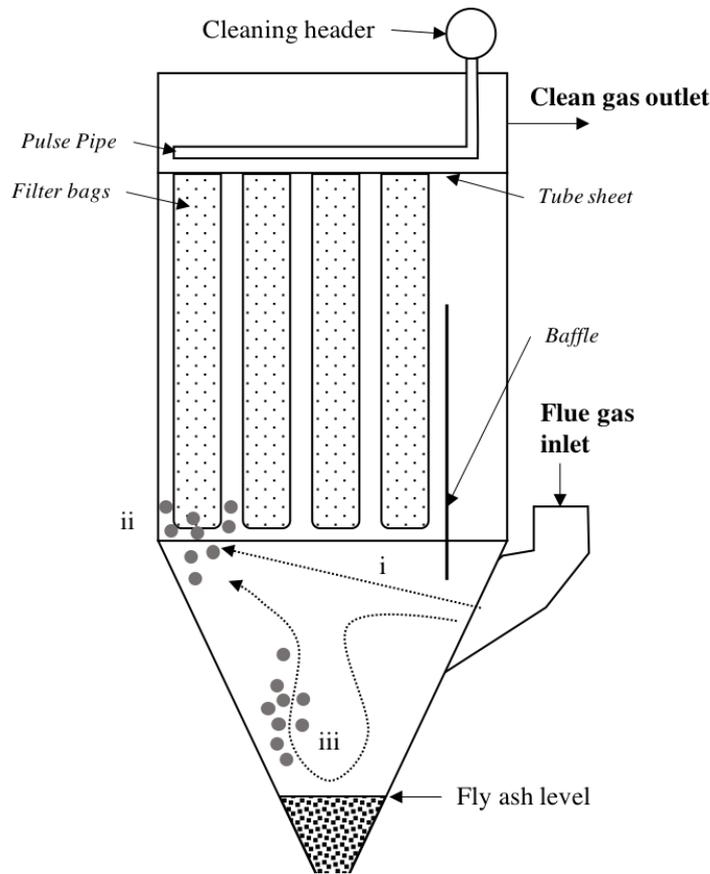


Figure 7.25: Schematic cross section of single compartment baghouse showing sub-optimal air flow phenomena. **i)** Bottom Bag abrasion. **ii)** Back wall impingement. **iii)** Hopper re-entrainment

In addition, the more robust construction of cartridges ensures that bottom abrasion is eliminated. This greatly increases filter life, as well as improving emissions performance and downstream erosion rates.

Lower Can Velocity

The increased surface area of cartridges means that they may be less densely packed to provide the same performance. Therefore, the space between filter elements is increased. This reduces the upwards velocity of the flue gas and therefore increases filtration efficiency and improves pressure drop.

Increased Air-to-cloth Ratio

The ultimate benefit of cartridge filtration is the ability to dramatically increase the capacity of the baghouse. 3D modelling of commercially available cartridge products has shown that a 70% increase in air to cloth ratio is achievable, while also achieving the benefits listed above. Air to cloth ratio is a key metric in defining a filter's efficiency and capacity. Essentially, more filtration area per unit volume of air means that each area of filter works less hard. This dramatically increases filter life, reduces pulse cleaning requirements and reduces pressure drop.

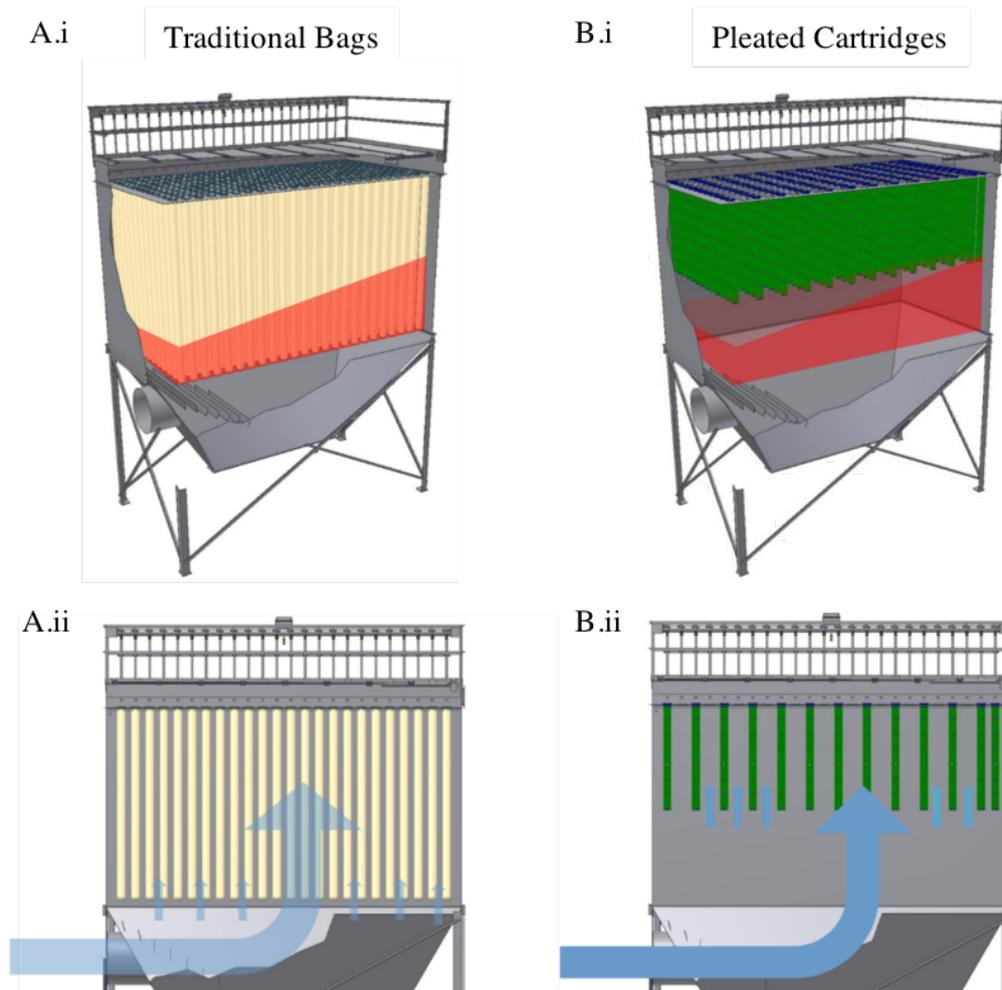


Figure 7.26: 3D schematic to show comparison between traditional bags and pleated cartridges. A) traditional bags. B) Pleated cartridges of identical air-to-cloth ratio as traditional bags in (A).

Greater Filter Life

Combining the factors mentioned above, typical cartridge filter life in biomass applications is five years. This reduces the downtime and labour cost of installing bags. Over this time, much more consistent emissions and pressure performance are maintained. The combined effect of increased filter life and reduced number of elements required means that, even before other benefits are included, the relative cost of ownership is lower for cartridges versus bags.

Reduced Pressure Drop

One of the most significant tangible savings comes in the form of the reduced pressure drop. Baghouse to cartridge conversions in power generation case studies have shown a pressure drop reduction of up to 55%. If this step change in performance were to be achieved at Plant 1, the resulting parasitic load reduction would be worth \geq £250,000 per year, or £1.25M over the lifetime of the filter.

Reduced Cleaning Cycle Demand

Due to their open structure, cartridges are pulsed less frequently and at a lower pressure. Case studies have shown that up to 50% less header pressure and 36x less frequent pulsing is required; resulting in up to 72x less compressed air usage compared with bags. This results in significant compressed air cost savings for Plant 1.

7.4.9 Commercial Solutions

FilterSense

Plant 1 currently operates the downstream filtration plant at 'factory default' settings. This study has identified several flaws in the current settings. Namely; pulse sequence, duration, interval and intensity. Additionally, the baghouse is controlled by a single differential pressure set point at 15 mbar. This means that the Plant 1 baghouse has no method for actively maintaining pressure drop by dynamically adjusting pulse settings.

In addition, bag leak detection is currently a laborious process and does not pinpoint leaks to beyond a single compartment. Stack particulate levels are slow to respond, and it could take several broken bags to detect and locate a leak. By this time, downstream damage has already occurred.

Thirdly, This study has identified a potentially significant compressed air leak through the diaphragms. If true, the pulse cleaning performance would be substantially reduced. Plant 1 currently has no method to identify poorly performing valves.

Furthermore, the current process for diagnosing these issues has been slow and laborious for operators and the author. The poor spatial resolution and access to only offline historic data means that data analysis is always slightly out of date, and slightly lacking in confidence. Plant 1 also has no means to *confidently* and *completely* quantify current baghouse performance, nor can improvements be benchmarked and economic benefits be accurately established.

A commercially available hardware solution, *FilterSense* addresses each of these issues. A *Filtersense* package was specified and costed for implementation in Plant 1:

IntelliPULSE cleaning system can maintain pressure drop at a set point by dynamically adjusting pulse settings (frequency, impulse time, sequence).

DynaCharge™ are locally mounted, high precision triboelectric particulate sensors which can pinpoint bag leaks and intelligently track the degradation of bags before stack particulate levels rise.

SP 10T probes diagnose valve performance to avoid poor pulse performance and leaks.

Lastly, *B-PAC™ G3 PRO* integrates the whole system and logs the high-resolution data to provide accurate and reliable data for quantifying the benefit of system modifications.

MAC valves

Plant 1 has 144 pulse cleaning valves which are vital to the performance of the whole plant. Poor filter cleaning leads to emissions non-compliance. A symptom of poor filter cleaning is high pressure build-up across the plant. This causes a high parasitic load and low boiler efficiency.

Parasitic load currently costs Plant 1 roughly £1,250,000 /year (1.2 MW + ROC basis). While most of this is unavoidable, a large proportion is due to sub-optimal filter cleaning. Low boiler efficiency (and backing off on fuel in order to recover efficiency) represents lost revenue to Plant 1 of approximately £4M pa (33 – 29 MW + ROC basis).

The filter cleaning system uses compressed air. Plant 1 consumes £ 250K /year in compressed air. Of this figure, £67,000 is due to the filter cleaning. This report has previously shown strong evidence that this figure is dramatically higher than required (although there is no air metering at the filter cleaning system). Between £15K - £ 42K in compressed air is wasted on leaks in the existing filter cleaning valves. £25 K pa. in compressed air is spent on ‘over-pulsing’ the filters in an effort to reduce baghouse differential pressure. An optimised pulse cleaning system delivers several tangible and indirect benefits, shown in Figure 7.27

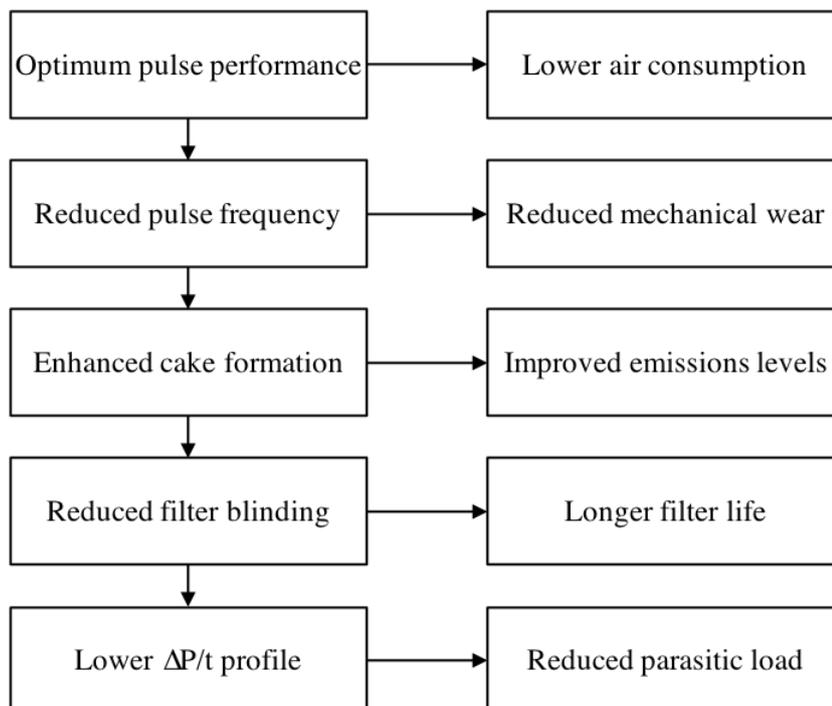


Figure 7.27: Optimised pulse cleaning benefits

Currently, the overwhelming majority of industrial pulse-jet baghouses use diaphragm

pulse valves to deliver the cleaning pulse. A recent commercial offering by *MAC valves, inc.* specifically designed for pulse-jet bag houses (MAC PV series) relies instead on a long life, bonded-spool design, designed to surpass the performance and reliability of the existing diaphragm valve technology.

In order to be effective, pulse valves need to deliver a precise and sharp pulse. It is not the volume of air, but the mechanical wave of pulse energy which shakes the dust from the filter bags. Stronger, more defined pulses therefore consume much less air to achieve the same cleaning performance or a much greater cleaning performance for the same compressed air consumption. Figure 7.28 compares the pulse profile in response to a control action for spool and diaphragm valves.

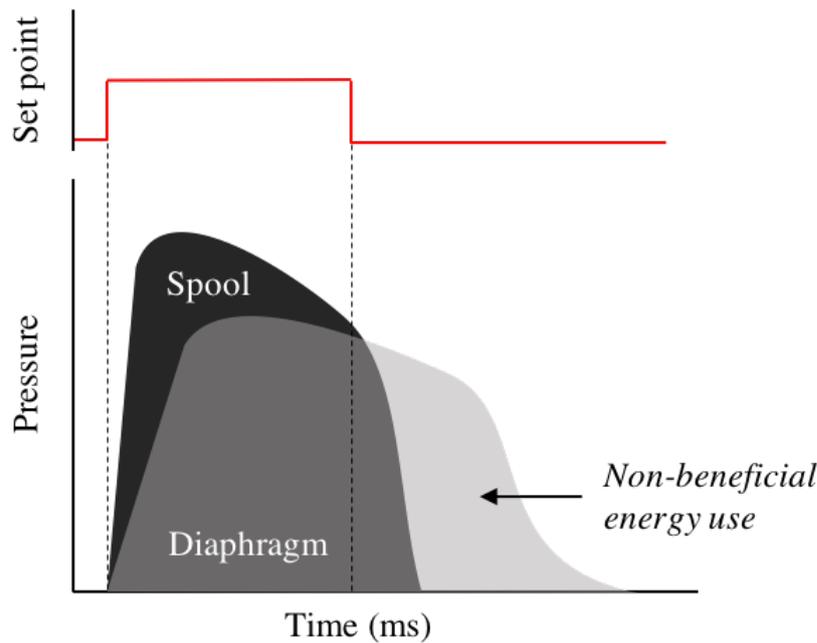


Figure 7.28: Schematic comparison of diaphragm and spool valve pulse response to the same set-point change

With reference to figure 7.28, several advantages of the spool valve technology can be explored. The *MAC* valve has been shown to have a shorter response time when compared to a competing diaphragm valve, resulting in a sharper increase in pressure over time, and achieving a higher peak and sustained pressure for the duration of the impulse time. When subject to the same control action, diaphragm valves are slow to respond owing to the dampening effect of the mechanical flexibility of the rubber diaphragm. Higher valve C_v results in lower peak pressures for a given header tank pressure in diaphragm valves. The same time delay is observed at the 'off' impulse, with diaphragm valves being relatively slower to close and fully seal. This delay results in non-useful compressed air passing the valve. The cumulative effect of this is wasted compressed air.

Given the performance difference outlined in figure 7.28, a change to spool valves affords the operator a number of ways to reduce OPEX, as follows:

Enhanced Cleaning - Given the *same* pulse parameters (header pressure, impulse time, frequency), the spool valves will clean the bags more effectively, allowing for an OPEX saving in parasitic load (Reduced ΔP).

Same cleaning - reduced settings - If current cleaning performance is adequate, the OPEX saving can be made by reducing compressed air consumption. Since spool valves give a higher peak pressure in response to a given header pressure, the header pressure may be reduced to give the same cleaning performance of existing valves. Since the time delay to both 'on' and 'off' set points is reduced, the on-time set-point of the pulse can be reduced accordingly.

Same cleaning - Reduced Frequency - Since each cleaning cycle is more effective, the frequency of cleaning cycles can be reduced. This linearly reduces the compressed air consumption and is also shown to increase bag life. OPEX savings on compressed air, bag replacement and parasitic load are all possible.

In practice, a combination of the effects listed above is shown to reduce OPEX, with the general advantage of spool valves being that they allow for such an optimisation to take place. Tiral baghouses fitted with MAC valves have shown a reduction in both pulse duration and pulse frequency giving rise to between 40-50% reduction in compressed air consumption in addition to achieving significantly better cleaning performance.

Optivel

Considering the results of the bag trial, discussed previously, A full upgrade to Gutsche fabric is recommended. The trial has shown a significant and sustained reduction in pressure drop with no impact on performance. The Gutsche fabric in the trial was *Microvel*, which represented a minimal change from the original filters. However, with confidence in the brand established, it is recommended that a further study of Gutche's *Optivel* fabric is undertaken.

Optivel blends the standard P84 fabric with PTFE, widely known as a non-stick material (Teflon). The addition of PTFE is predicted to result in enhanced dust release and therefore lower pressure drop in operation, and longer filter life. Furthermore, the PTFE blend adds additional chemical resistance.

While the effect of chemical coatings has not been discussed explicitly in this work, two commercially available coatings are recommended for trial. *Antafin*, an anti-adhesive treatment for 'sticky' dust offers resistance to acids, hydrolysis and other chemicals while at the same time improving the cleaning performance. *Membratex* is designed specifically for flue gas cleaning in biomass, incineration and waste-to-energy applications. A fine layer of PTFE particles coats the surface of the fabric, preventing particle ingress.

Sorbent Mixing

As discussed previously, sorbent mixing in the duct is currently poor. Most of the sorbent does not interact with the flue gas stream for a meaningful residence time, and most of the sorbent is filtered by the front right compartment, B1. This non-uniformity of sorbent distribution leads to sub-optimal emissions performance as well as non-uniform bag wear.

CFD simulations show that the greatest impact for the lowest cost can be achieved by installing a static mixer within the duct. These static mixers are designed to introduce counter-rotational concentric turbulence in order to maximise sorbent mixing and residence time, while minimising pressure drop. However, the static mixer inherently introduces some pressure drop.

CFD studies have shown that the optimum design is likely to introduce a 2 mbar pressure drop. This figure is currently unsustainable in the Plant 1 system from an operability point of view, but with the many proposed benefits of other modifications discussed in this report, it may be possible to support this increased pressure drop in the future.

7.5 Predicted Commercial Impact

Commercial implementation of the solutions identified in this work require that a business case be made. Tangible financial benefits are quantified and projected into the future, generating a predicted breakeven point (the time at which the additional investment is matched by the cumulative savings achieved). This is most commonly achieved by a cash flow analysis.

An attempt was made to quantify the tangible financial benefit of the proposed upgrades. This is achieved by passing predicted performance improvements through the economic model derived earlier. In this way, performance improvements in pressure drop, air usage, bag life, sorbent dose etc. are given a financial value. In addition to the financial benefits, an attempt is made to identify the more intangible, but still important benefits. For example, it is currently difficult to put a value on fan blade wear, corrosion, reduced emissions etc. Further studies in this direction for biomass & EfW plants may prove valuable in quantifying the tangible *cost* of these indirect effects.

The proposed upgrades address a whole series of issues within the filtration system, as discovered by CFD analysis and Knowledge Acquisition described in this work. The discussion so far has highlighted that many factors have direct and indirect influences on other factors. For instance, reduced pulse frequency has a direct impact on compressed air consumption, but also increases bag life and reduces particulate and acidic emissions by the mechanisms we have described. The interaction and compounding of benefits is therefore maximised when modifications are undertaken together. Indeed, some of the more

pronounced benefits of each modification are only achieved when taken in tandem with other modifications. Figure 7.29 schematically represents some of the principle interactions in relation to key tangible OPEX metrics.

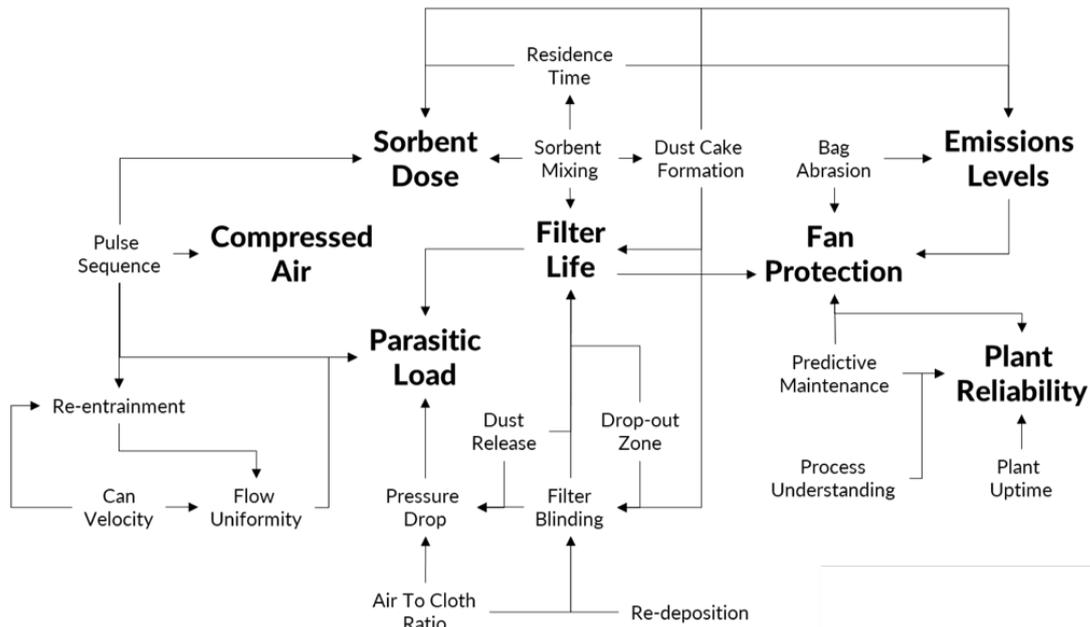


Figure 7.29: Principle interactions in baghouse OPEX

When estimating the tangible benefits for each solution, a combination of case studies, CFD simulations, performance data, trials and detailed supplier specifications are relied upon. By incorporating uncertainty and variability, the economic predictions are presented as a range between an upper and lower bound. With reference to table 7.7, the following benefit metrics are established.

a) Parasitic Load

A percentage value is presented. This refers to the percentage reduction in dynamic differential pressure, after the CFD modelled static pressure drop is subtracted. This percentage reduction is then passed through the data-driven model which correlates pressure drop to parasitic load, and therefore to revenue on a MW+ROC basis.

b) Compressed Air

A percentage value is presented. This value refers to the percentage reduction in normal-cycle compressed air consumption. The reduced air consumption attributed to lower over-pulse requirements and air leaks are also considered.

c) Sorbent dose

The percentage value refers to the combined reduction in PAC & Lime assuming a constant ratio. In actuality, it may be possible to tune the dose of each component independently, as their response to the various modifications is likely to differ slightly.

d) Filter Life

At Plant 1, the current average filter life is ~2 years, with 50% of the baghouse changed each year. A filter life increase therefore has the financial benefit of reducing the effective yearly cost (parts and labour) of replacing filters. A 3 year service life would entail changing a third of the baghouse (2 compartments) every year. Alternatively, 50% of the baghouse could be changed every 1.5 years. In either case, a saving is made on the parts and labour requirements relative to the current regime.

e) Plant Reliability

Based on plant data, it is possible to estimate the cost of plant uptime per day. An estimate is made as to the value of extra days uptime due to the modification. This uptime can be attributed to the ability to perform predictive/preventative maintenance, or the reduction in shutdown duration or start-up delays. This 'uptime' can also be attributed to higher electrical output.

Additionally, two intangible benefits are also expected, but not quantified in the current work.

Fan wear

Currently, no data is considered to allow a reduction in fan wear to be financially quantified. This benefit is expected to be significant where particulate emissions and filter leaks can be reduced or identified early. With access to further data on the maintenance cycle and costs for the ID Fan, this important element of OPEX should be included into the financial benefit prediction.

Emissions Performance

The direct emissions performance benefits we refer to are particulate emissions and Acid + Heavy metal emissions, since these are directly addressed by the modifications we make. Further to this, the increased efficiency of the whole system will also result in an effective reduction in CO₂ emissions on a per-MW basis. Again, there is no current understanding as to how Plant 1 can ascribe a financial value to these benefits, given that emissions are already below the permitted value.

7.5.1 Plant Implementation

Following on from the recommendations made to Plant 1 in the process of preparing this work, several of the proposed solutions were implemented. A systematic approach to sorbent dose reduction was commenced, resulting in an approximate halving.

PTFE blended fabric has replaced P84 in the baghouse, resulting in significant pressure drop savings. Currently, 50% of the baghouse has been upgraded, with the remainder

Table 7.6: Estimated cost and predicted benefits for Plant 1 implementation of various proposed modifications

Modification	Cost (£)	Annual Predicted Benefit (range)				
		Parasitic Load (%) ^a	Compressed Air (%) ^b	Sorbent Dose (%) ^c	Filter Life Life (years) ^d	Reliability (days) ^e
FilterSense	53,871	5 - 10	20 - 50	5 - 20	0 - 1	1 - 2
Gutsche Filters	0*	5 - 25	-	-	0.5 - 1	0 - 1
MAC PV	144,600	10 - 25	40 - 90	-	0.5 - 1	0 - 1

*Cost relative to current regime is £0. (Parts and labour cost for existing filters versus recommended is equal)

Table 7.7: Summary of predicted financial benefits. (median of range)

Modification	Payback time (months)	5-year NPV (£'000s)
FilterSense	6	635
Gutsche Filters	0	440
Mac PV	12	755

to be installed in early 2019. The benefits of the PTFE fabric are expected to result in further benefits. For example, once the baghouse is fully PTFE, it is expected that pulse cleaning frequency can be reduced, with no adverse effects.

Compressed air metering has been installed on site, allowing for predictive and preventative maintenance of the pulse valves. The compressed air consumption data is also providing a cost benefit justification for further cleaning system upgrades.

Perhaps most interestingly, Plant 1 has opted to install a full suite of baghouse monitoring and control equipment with the ability to remotely monitor and control the baghouse settings.

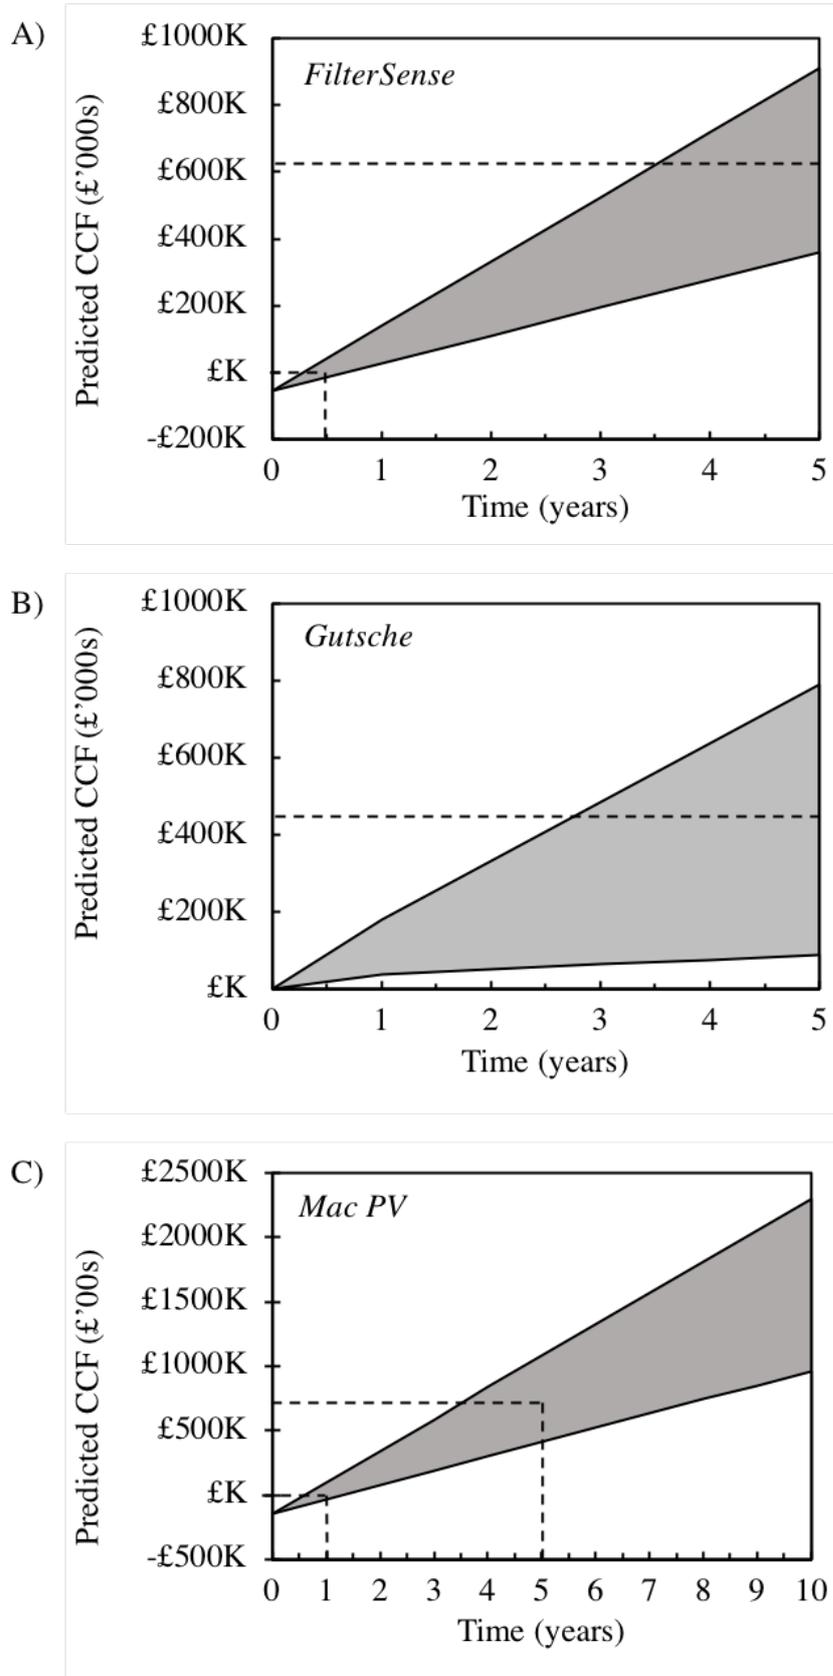


Figure 7.30: Graphs to show predicted cumulative cash flow (CCF) (£'000s) for three proposed modifications. A) FilterSense. B) Gutsche Filters. C) Mac pulse valves.

7.6 Chapter Summary

In this chapter, large amounts of online historic process data were presented and characterised. The impact of process variability and design variance was discussed at length with a particular focus on the flue gas filtration system. Next, several operational filter samples were obtained and subjected to various standard textile test methods. The results, coupled with online data and a 3D plant model were pushed forward into a novel modelling framework, outlined in this work, named Rapid Hybrid Computational Fluid Dynamic Design, whereby Design of Experiments methodologies are applied to Computational Fluid Dynamics to significantly reduce the computational demands and time requirements for constrained optimum designs to be generated. By using the modelling frameworks presented and developed in this work, several recommendations and improvements were proposed to the industrial collaborator. Where these proposals were implemented, this chapter presents the real economic and process improvement. Where the proposals were pre-implementation, this chapter discusses the predicted benefits as forecasted by the models. In all, proposals amount to a significant reduction in the total cost of ownership of the filtration system.

Chapter 8. Conclusions and Recommendations to Industry

8.1 Chapter 2

This chapter covered the basic GTL technology concepts required to navigate the proceeding chapters. At the time of writing, Commercial-scale and long-term operation studies of micro-channel FT reactors was limited. Therefore, the principles presented were necessarily based on the first hand industrial experience of the author (having assembled, commissioned and operated one of the two dominant micro-channel FT reactor technology platforms in existence). However, since the time of writing, new studies have been published which more specifically address the considerations of micro-FT reactors. While these studies are predominantly either lab-scale or *in silico* studies, there are several important concepts not specifically covered in this work. A refreshed outlook should be taken to integrate these developments into future designs of commercial Micro-FT reactors.

At the commencement of this body of work, the price of crude oil was $> \$100/\text{bbl}$. Therefore the focus of FT development was focussed on the production of syncrude from associated natural gas. However, at the conclusion of this work, the oil price was $\sim \$35/\text{bbl}$. This dramatic change in process economics drove a rapid reemphasis of technology development onto the production of high performance oils and waxes. Market fluctuation is likely to continue into the future. Focus should be given to designing processes, reactors and catalysts which can cope with changes in operating mode in response to market conditions.

8.2 Chapter 3

This chapter largely focussed on post-operational analysis of the SP09 reactor, driven by the results of previous extended runs which have indicated the presence of blockages in particular locations of the reactor.

Elemental analysis was limited in scope to examining the concentration of carbon and hydrogen, due to intellectual property constraints. Future efforts to overcome this constraint should be attempted. The molar ratio of carbon and hydrogen residues indicates that elemental carbon is more prevalent in blocked channels. The presence of other forms of carbon is not ruled out and further examination is required to this end. Perhaps the

use of X-Ray Diffraction techniques would allow the definitive characterisation of the carbon species present. However, inspection of the operational history with reference to the hydrogenation temperature of potential carbon sources seems to imply the presence of polymeric or graphitic carbon, rather than carbonyls or hydrocarbons. This is supported by the ratio of carbon to hydrogen.

In situ endoscopy allowed for the blockages to be seen clearly for the first time; the radial blackening and contact shielding observed certainly indicate the likely mechanism of carbon laydown to be related to flow distribution and possibly inter-particle and particle-wall heat transfer. Endoscopy and scanning electron microscopy also indicate that an amorphous, fragmented bed at the site of blockage replaces the particle structure found in free-flowing regions. These blockages are brittle and composed of fines. SEM images also show the presence of fines within the macropores of individual particles, hinting at a possible cause for the loss in performance observed in end-of-run reactors. The fines are demonstrated to be enriched in elemental carbon, but still containing catalytic support material, indicating that the particles themselves are undergoing a mechanical change in the formation of a blockage.

The correlation between increased reactor ΔP , individual channel δP and the presence of blockages is clear. The individual contribution of blockages to the overall channel δP is dominant. However, a row-wise pattern of severity is characterised by the δP per unit length. Central rows, more prevalent in carbon, are also shown to be more severe in terms of blockage characteristics. Further work is needed in the modelling of accurate pressure drop predictions for fixed beds at the low diameter ratio. Attempts made in this work are limited by the use of a mean particle diameter in place of particle size distribution, an equivalent diameter equating to a cylindrical channel and the inability to account for particle porosity.

The general picture has emerged that operating at the limit of catalyst life and performance leads to the risk of thermal runaway, which in turn produces extreme local temperatures systemic to the reactor geometry and conditions. In SP09, thermal events more common at the inlet of the reactor are seen to have an axial and radial propagation. It is proposed that the alternative chemistry at these uncontrolled temperatures, along with any physical effect on the bed are responsible for forming mechanically resistive blockages in the regions surrounding the peak thermal excursion. Edge channels may be less susceptible to this effect, owing to their enhanced ability to dissipate heat into the surrounding metal.

It has emerged that the desire of long-life catalyst and high productivity operation are competing objectives. The reactor temperature profile should be carefully managed. The current method of reactor temperature measurement and control is limited in terms of spatial resolution. This work has shown that the propagation of runaways occurs in unexpected ways from the point of view of the operator, when faced with such low

resolution data. Further work in Chapter 5 of this work addresses this challenge directly, by development of a novel reactor temperature monitoring and control solution. It is hoped that implementation of this solution at the commercial scale will provide the necessary resolution for operators (or control algorithms) to operate at the limit of thermal performance, without the risk of the catalyst damage and deactivation discovered in this chapter.

Additional work by the author, not disclosed in this report, has addressed the issues relating to the post-operational turn-around of blocked reactors, with a focus on solutions which are practical for remote locations and minimal downtime. It is highly recommended that commercial ventures employing micro-FT technology develop similar methods and approaches for the eventuality that a commercial scale reactor become blocked in the manner described in this work.

8.3 Chapter 4

This chapter solves the issues of characterising a packed catalytic bed faced by the first principles approaches. X-Ray tomography is successfully applied to the novel next-generation reactor geometries through the use of image/data analysis. Characterisation of the fixed bed in this manner allows for a common-sense understanding of *localised* bed conditions, not described in data-driven or first principles models. The 3D computational representations of packed beds generated in this chapter provide the ability for further *in silico* flow studies to be performed.

Computational Fluid Dynamics methodologies developed in later chapters of this work are recommended to be re-applied to the micro-FT reactor design by parametric modelling of reactor channel geometries, given the detailed bed parameters obtained in this work.

Additional work by the author, not disclosed in this report has addressed the challenges of pre-operational catalyst loading of micro-FT reactors. However, this work is likely to show poor scalability to commercial scale reactors. For instance, at lab-scale, the catalyst loadings were meticulously recorded to $\pm 0.0001g$ and individually vibrated to allow for bed settling. In a commercial scale reactor, where $\geq 50,000$ channels may be present, this is not a practical method. Significantly reduced packing efficiency is likely to be obtained, with less knowledge of the local distribution of catalyst. Although work in this chapter has characterised a small sub-unit of a single channel, further understanding of the global distribution of catalyst in a commercial scale reactor is required. If the previous chapter's conclusions are followed, non-uniformity in reactor loading is likely to result in 'hot-spots' where thermal events are more likely to occur.

8.4 Chapter 5

This chapter presented the results of a trial implementation of distributed temperature sensing technology to provide high resolution temperature data from operational compact Fischer-Tropsch reactors.

Successful physical installation of the DTS element allowed for an extended operational run to be conducted under DTS supervision on the pilot scale reactor. The DTS equipment provided suitable thermal, spatial and temporal resolution to capture steady state operation. However, the configuration of the fibre in the small-scale reactor meant that a temperature offset caused by external cooling was observed. This dilution of the spatial averaging of DTS data meant that peak temperatures during a thermal excursion were not closely followed by the DTS system compared to discrete thermocouples mounted locally. This is concluded to be a phenomenon caused by the coincidence of pilot scale reactor dimensions with the spatial resolution of the DTS element, and therefore would not be an issue at commercial scale.

Several recommendations are made in this work to facilitate scale-up of this technology to commercial scale reactors:

The sampling and handling of large amounts of data generated should be considered carefully at commercial scale. Exception-based, boxcar or rate of change sampling methods should be developed to reduce the bandwidth and storage capacity requirements.

Enhancing the robustness of the DTS system by employing dual-pass and dual-channel installations has been shown to safeguard the system against one or even two fibre optic breakages. The attenuation of signal at greater distances than those covered in this work should be assessed. Commercial installations are likely to require several kilometres of fibre optic cable to provide adequate coverage.

Commercial implementation of this work will occur in the near future and commercial-scale results and data will need careful interpretation in order to establish a reliable control scheme based on the readings provided by the DTS unit.

8.5 Chapter 6 & 7

This body of work was carried out with access to two commercially operating large scale biomass power generation plants and co-funded by a filtration technology and solutions provider. Hybrid Models are presented with a view to reduce the total cost of ownership of the filtration system.

Several conclusions relating to specific areas of operation are reported. Most notably, the report shows that there is scope for drastic improvement in the process economics

of bag house filters. Even the most basic regression analysis of historic plant data will reveal to operators the hidden relationships between operational set-points and economic performance.

This work has developed a framework for performing geometrically constrained optimisations of plant geometries in a computationally efficient manner, whilst also showing good relation to the physical system. This is made possible by integrating process data at the boundary condition to computational simulations.

Low spatiotemporal data resolution in key areas of baghouse performance are limiting the confident and accurate determination of OPEX. The resulting lack of process understanding translates to a disinterest in the economic performance of the filtration system. We recommend a commercially available hardware package designed to address the missing data problem. It is hoped that implementation of this solution will allow for the approaches developed in this work to be revisited with increased confidence.

The impact of the study was two-fold. Firstly, the plant operators received recommendations for process improvements predicted to deliver upwards of £2.5 million in reduced OPEX over a 5 year period, with a payback period of ≤ 1 year. Implementation of these proposed improvements is currently in progress, with increased data capture ability providing a meaningful way to quantify the true impact. Secondly, the hybrid modelling methods, models and frameworks developed in this work now form a fully commercialised package of consultancy services offered by the co-funding filtration company. The increased technical capability is already generating revenue, predicted to reach approximately £1 million by 1 year after commercialisation.

**Appendix A. Patent WO 2016/001663 A1: Catalytic Reactors
Comprising Distributed Temperature Sensors**



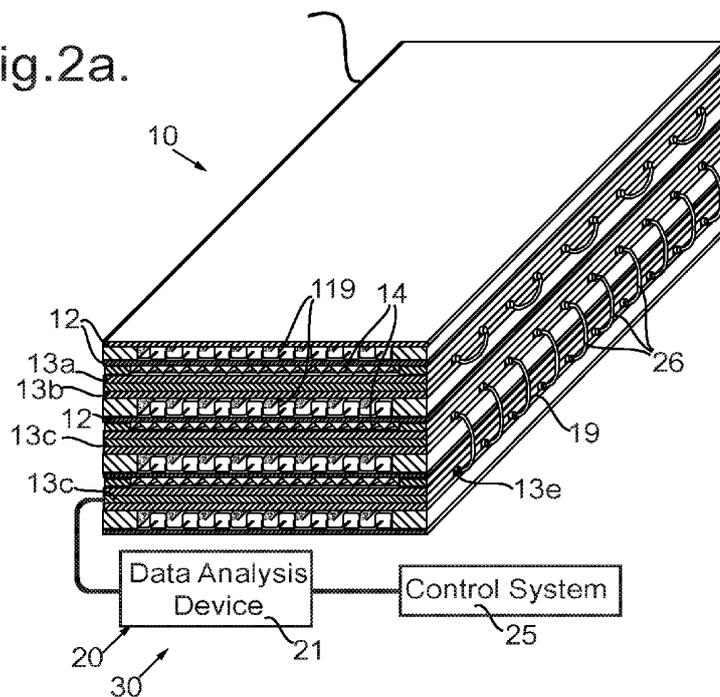
- (51) **International Patent Classification:**
B01J 19/00 (2006.01) *B01J 19/24* (2006.01)
- (21) **International Application Number:**
PCT/GB2015/051922
- (22) **International Filing Date:**
1 July 2015 (01.07.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
1412000.0 4 July 2014 (04.07.2014) GB
- (71) **Applicant:** COMPACT GTL PLC [GB/GB]; 19 Blacklands Way, Blacklands Way, Abingdon OX14 1SY (GB).
- (72) **Inventors:** BONNER, Christopher William; 1 Sherbrooke Avenue, Wilnecote, Tamworth Staffordshire B77 5EJ (GB). DANNATT, Benjamin; 1 Sycamore Avenue, Grimsby, North East Lincolnshire DN33 2AY (GB).
- (74) **Agent:** HUFTON, Victoria Alice; Stratagem Intellectual Property Management Limited, Meridian Court, Comberton Road, Toft, Cambridge CB23 2RY (GB).

- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:
— with international search report (Art. 21(3))

(54) **Title:** CATALYTIC REACTORS COMPRISING DISTRIBUTED TEMPERATURE SENSORS

Fig.2a.



(57) **Abstract:** A catalytic reactor is provided comprising a plurality of first flow channels including a catalyst for a first reaction; a plurality of second flow channels arranged alternately with the first flow channels; adjacent first and second flow channels being separated by a divider plate (13a, 13b), and a distributed temperature sensor such as an optical fibre cable (19). The distributed temperature sensor may be located within the divider plate, or within one or 10 more of the flow channels.

WO 2016/001663 A1

CATALYTIC REACTORS COMPRISING DISTRIBUTED TEMPERATURE SENSORS

The present invention relates to a chemical reactor, particularly but not exclusively a chemical reactor in which an exothermic reaction occurs. The invention also relates to the control of a reaction taking place in such a chemical reactor. It would be relevant, for example, to a plant and a process for treating natural gas to produce a liquid product, particularly but not exclusively for performing steam methane reforming followed by Fischer-Tropsch synthesis.

A process is described in WO 01/51194 and WO 03/048034 (Accentus plc) in which methane is reacted with steam, to generate carbon monoxide and hydrogen in a first catalytic reactor; the resulting gas mixture is then used to perform Fischer-Tropsch synthesis in a second catalytic reactor. The overall result is to convert methane to hydrocarbons of higher molecular weight, which are usually liquids or waxes under ambient conditions. The two stages of the process, steam/methane reforming and Fischer-Tropsch synthesis, require different catalysts, and heat to be transferred to or from the reacting gases, respectively, as the reactions are respectively endothermic and exothermic. Reactors for these reactions may be formed as a stack of plates, with flow channels defined between the plates, the flow channels for the different fluids alternating in the stack. Suitable catalysts for the reactions taking place within the flow channels may be provided.

As described in WO 01/51194 the catalyst may be supported on an insert which may take the form of a foil. Alternatively, catalyst may be applied directly to the walls of the flow channel. In a further alternative, the catalyst may be provided as catalyst particles which may form a bed of particles. The term "particles" means, for example, pellets, spheres, extrudates, trilobes, powders, granules, fibres, particulates, particulate solids or any solid components comprising the catalyst material that are suitable for use in the catalytic reactor. The particles may be arranged in a bed of particles.

One advantage of this reactor configuration, with alternate flow channels carrying different fluids, is that heat can be transferred efficiently between first and second flow channels. As a result, very active catalysts can be used resulting in a high level of process intensification. However, if heat is not transferred away sufficiently rapidly a thermal runaway may occur which may result in damage to the catalyst. Repeated thermal runaways may even compromise the integrity of the reactor as a whole. As a result, control of the reactions taking place within these reactors is key to ensuring the longevity of the catalyst and also the reactor.

Thermocouples are typically employed to monitor temperature. WO 2014/096779 discloses a reactor block having a plurality of alternately arranged flow channels, wherein the plurality of flow channels are arranged such that fluids in at least two adjacent flow channels can exchange heat through an intervening wall, and wherein the intervening wall defines at least one channel extending within the reactor block and communicating with an outside surface of the block and is dimensioned to accommodate a temperature sensor.

The temperature sensors are thermocouples of around 1 mm diameter and they are disposed between the channels within the intervening wall.

According to an aspect of the present invention there is provided a catalytic reactor comprising a plurality of first flow channels including a catalyst for a first reaction; a plurality of second flow channels arranged alternately with the first flow channels; adjacent first and second flow channels being separated by a divider plate; and a distributed temperature sensor provided within the divider plate. The provision of a distributed temperature sensor provides a much greater degree of flexibility in relation to the frequency of measurement and also the location of measurement in comparison with the thermocouples used in the art.

The distributed temperature sensor may be a fibre optic cable. The fibre optic cable may follow a tortuous path through the divider plates in order to provide temperature data for all of the adjacent first and second flow channels. A fibre optic cable may be provided in each divider plate. The distributed temperature sensor does not provide point temperature measurements, but provides average temperatures over successive portions of its length. It may be preferable to perform such an average in a direction where there is little temperature variation. The temperature will usually vary more significantly in the flow direction than transverse to the flow direction, so it may therefore be beneficial to arrange the path of the distributed temperature sensor so it is transverse to the flow direction in the first flow channels.

Alternatively, in another aspect of the invention, the distributed temperature sensor (such as a fibre optic cable) may be provided within one or more first flow channels and/or one or more second flow channels. This configuration provides more direct measurement of the temperature of the process fluids. However, the environment within the flow channel will be more hostile than within the divider plate, which may limit the useful lifetime of the fibre optic cable.

The first flow channels may include a Fischer-Tropsch catalyst and the second channels carry a heat transfer fluid, in use. Alternatively, the first flow channels include a steam methane reforming catalyst and the second channels further include a combustion catalyst.

- 5 The reactor of the present invention may be included in a plant for performing Fischer-Tropsch synthesis, together with a data analysis device. The plant may further comprise a control system operatively linked to the data analysis device.

10 The control system may be configured to shut in the reactor if the data analysis device detects a rate of change of temperature in excess of a predetermined value. Temperature measurements obtained from optical fibre cables are much less susceptible to noise than thermocouples described in WO 2014/096779.

15 The reactor may comprise a stack of plates. For example, the flow channels may be defined by thin metal sheets that are castellated and stacked alternately with flat sheets; the edges of the flow channels may be defined by sealing strips. Flow channels may alternatively be defined between flat plates by spacer bars. The stack of plates forming the reactor is bonded together for example by diffusion bonding, brazing, or hot isostatic pressing.

20 The invention will now be further and more particularly described, by way of example only, and with reference to the accompanying drawings, in which:

Figure 1a shows a sectional view of part of a reactor suitable for Fischer-Tropsch synthesis;

Figure 1b shows a sectional view of a modification to the reactor of Figure 1a;

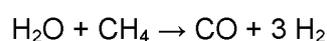
Figure 2a shows a perspective view of a reactor block according to the present invention;

25 Figure 2b shows a cut-away view of the reactor block of Figure 2a;

Figures 3a and 3b are graphs showing the temperature profile along the length of the reactor of Figure 2a under different circumstances;

Figure 4 shows a modification to the reactor block of Figure 2a.

30 The invention is of relevance to a chemical process for converting natural gas (primarily methane) to longer chain hydrocarbons. The first stage of this process may use steam reforming to form synthesis gas, that is to say the reaction of the type:

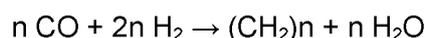


35

This reaction is endothermic, and may be catalysed by a rhodium or platinum/rhodium catalyst in a flow channel. The heat required to cause this reaction may be provided by

combustion of an inflammable gas such as methane or hydrogen, which is exothermic and may be catalysed by a platinum/palladium catalyst in an adjacent second gas flow channel.

The synthesis gas is then used to perform Fischer-Tropsch synthesis to generate a longer
5 chain hydrocarbon, that is to say:



which is an exothermic reaction, occurring at an elevated temperature, typically between
10 190°C and 280°C, and an elevated pressure typically between 1.5 MPa and 2.5 MPa (absolute values), in the presence of a catalyst such as iron, cobalt or fused magnetite. A suitable catalyst for the Fischer-Tropsch synthesis comprises a coating of gamma-alumina of specific surface area 140-230 m²/g with about 10-40% cobalt (by weight compared to the alumina), and with a promoter such as ruthenium, platinum or gadolinium which is less than
15 or equal to 10% the weight of the cobalt, and a basicity promoter such as lanthanum oxide.

The synthesis gas stream is cooled and compressed to the elevated pressure, say 2.0 MPa, and is then fed to a catalytic Fischer-Tropsch reactor. The Fischer-Tropsch reactor may be a compact catalytic reactor formed from a stack of plates as herein described; the reactant
20 mixture flows through one set of channels, while a coolant flows through the other set.

Referring now to Figure 1a there is shown a reactor block 10a suitable for use in performing Fischer-Tropsch synthesis, the reactor block 10a being shown in section and only in part. The reactor block 10a consists of a stack of flat plates 12 of thickness 1 mm spaced apart so
25 as to define channels 15 for a coolant fluid alternating with channels 17 for the Fischer-Tropsch synthesis. The coolant channels 15 are defined by sheets 14 of thickness 0.75 mm shaped into flat-topped sawtooth corrugations. The height of the corrugations (typically in the range 1 to 4 mm) is 2 mm in this example, and correspondingly thick solid edge strips 16 are provided along the sides, and the wavelength of the corrugations is 12 mm (the arrangement
30 being described in more detail below). The channels 17 for the Fischer-Tropsch synthesis are of height 5 mm (typically within a range of 1 mm to 10 mm), being defined by bars 18 of square or rectangular cross-section, 5 mm high, spaced apart by 80 mm (the spacing typically being in a range of 20 - 100 mm) and so defining straight through channels.

35 Referring now to Figure 1b there is shown an alternative reactor block 10b suitable for use in performing Fischer-Tropsch synthesis, the reactor block 10b being shown in section and only in part. In many respects the reactor block 10b resembles the reactor block 10a,

identical components being referred to by the same reference numerals. The reactor block 10b consists of a stack of flat plates 12 of thickness 1 mm spaced apart so as to define channels 15 for a coolant fluid alternating with channels 117 for the Fischer-Tropsch synthesis. The coolant channels 15 are defined in addition by sheets 14 of thickness 0.75 mm shaped into flat-topped sawtooth corrugations as described above, with solid edge strips 16. The channels 117 for the Fischer-Tropsch synthesis are sealed by solid edge bars 18 and are defined in addition by sheets 119 of thickness 1.0 mm shaped into castellations of height in the range of 4 mm to 12 mm, preferably 5 mm. In the preferred example the resulting channels 117 are of width 7 mm and of height 6 mm and extend straight through the stack from one face to the opposite face. As with the channels 15, 17 within the reactor block 10a, the channels 15 and 117 in reactor block 10b extend in parallel.

The reactor blocks 10a or 10b may be made by stacking the components that define the channels 15 and 17 or 117, and then bonding them together for example by brazing or by diffusion bonding. If the reactor block 10a or 10b is to be used for Fischer-Tropsch synthesis, the reactor block 10a or 10b is then turned through 90° so that the channels 15 and 17 are upright in use. If the reactor block 10a or 10b is to be used for steam methane reforming or other reactions, it may not be necessary to turn it through 90° as it is less critical that the channels 15 and 17 are upright for these reactions. In each case the first flow channels 17 or 117 would be then provided with a catalyst (not shown) for the Fischer-Tropsch reaction.

Each plate 12 may for example be 1.3 m by 1.3 m, or 1.2 m by 0.8 m, so the channels 17 would be 1.3 m long or 0.8 m long, respectively. Preferably the channels 17 are no more than 1.5 m long, and preferably at least 0.3 m long. Although only a few layers of the stack are shown, the reactor block 10a or 10b might in practice have ten, twenty or thirty layers containing the reaction channels 17, or as many as 100 such layers. The flat plates 12, the bars 18, and the castellated sheets 14 and 119 may be of a high-temperature steel alloy, or of an aluminium alloy, for example 3003 grade (aluminium with about 1.2% manganese and 0.1% copper).

Referring now to figures 2a and 2b, these show perspective views of portions of a reactor block 10 of the present invention. The reactor block 10 is shown as being constructed as described above for the reactor block 10b, but alternatively it might be constructed as in the reactor block 10a. It differs from the reactor block 10b only in that at least one of the flat plates 12 is formed of two divider plates 13a, 13b that are held apart by multiple rectangular plates 13c arranged side-by-side, but spaced slightly apart so as to define narrow gaps 13d between successive plates 13c. Within each such narrow gap 13d is a tube 13e, the external

diameter of the tube 13e being the same as the thickness of the rectangular plates 13c. All of these components 13a-13e are bonded together during assembly of the reactor block 10, as described above. It will be appreciated that an alternative reactor block of the invention would differ from the reactor block 10a only in that at least one of the flat plates 12 would be formed by divider plates 13a, 13b along with the requisite rectangular plates 13c and the tubes 13e as described above.

A fibre optic cable 19 is passed through each of the tubes 13e between the divider plates 13a, 13b. The fibre optic cable 19 forms a distributed temperature sensor 20 which is configured to measure the temperature at multiple positions along the fibre optic cable 19 using a data analysis device 21. The fibre optic cable 19 has a diameter of less than 1 mm, which ensures that it can pass through the tubes 13e in the reactor block 10. The fibre optic cable 19 may have multiple cores which provide the capability of providing multiple temperature measurements, for example providing the facility for making measurements from opposite ends of the fibre optic cable 19, by sending optical pulses in opposite directions along separate cores or fibres within the fibre optic cable 19.

As illustrated schematically in Figures 2a and 2b, the fibre optic cable 19 makes a number of passes through the reactor block 10 resulting, overall, in a tortuous path. At the point where the fibre optic cable 19 leaves the reactor block 10, it forms a loop 26 taking it to the next selected entry point into the reactor block 10. The positioning of the fibre optic cable 19 is dictated by the tubes 13e provided between the divider plates 13a, 13b. The fibre optic cable 19 makes a substantially straight line path from the entrance to the exit of each tube 13e. The tubes 13e are sized to conform closely to the diameter of the fibre optic cable 19 so that the position of the fibre optic cable 19 is known and movement of the cable 19 is minimised. It will be appreciated that the tubes 13e and therefore the corresponding portions of the fibre-optic cable 19 extend in directions that are transverse to the flow direction through the reaction channels 17.

The positions of the tubes 13e are selected to provide an even coverage of the reactor volume in order to ensure reliable and safe temperature monitoring throughout the life of the reactor block 10. However, under some circumstances (not illustrated), it may be preferable to intensify the monitoring of temperature in certain sections of the reactor block 10. This may be effected by providing additional narrow gaps 13d and tubes 13e in some parts of the reactor block 10 so that the distance between consecutive passes of the optical fibre cable 19 through the reactor block 10 is reduced. This may be applicable in part of a reactor block 10 which is known to be susceptible to thermal excursions.

Overall, the reactor block 10 may have over one thousand such tubes 13e for temperature measurements. For example, there may be ten, twenty or fifty tubes 13e in each divider plate 13a, 13b. The reactor block 10 may be provided with divider plates 13a, 13b between each castellated sheet 14 so that a fibre optic cable 19 can be provided to obtain temperature data pertaining to the temperature in the immediately adjacent flow channels 15, 17 for each of the ten, twenty or one hundred layers of the reactor block 10.

The temperature measurements obtained from the distributed temperature sensor 20 are averages over successive short lengths of the cable 19 (for example over successive 1 m, 25 cm or 12 cm lengths). This ensures that any temperature excursion will be captured by the data analysis device 21. The temperature is measured at regular intervals, throughout the length of the cable 19. Data is then extracted by the data analysis device 21 by analysing signals from along the length of the fibre optic cable 19. The length of successive data point locations is currently limited by the data processing capacity of the data analysis device 21, and may for example be successive 12 cm lengths of the optical fibre 19. The data analysis device uses the Boxcar method to minimise repetition in data collected. In this context, the Boxcar method allows an exception to be set at, for example, 0.5°C, and then data is only recorded if it differs from the previously recorded value by more than this amount, which may be called the exception quantum. This reduces the amount of data to be analysed, without reducing the accuracy or frequency of the measurement.

The measurement frequency is a trade-off between the time interval and the accuracy of data required. For example, for a time interval of 10 s, a 1°C change in temperature may be detected to an accuracy of 0.2°C. For a 5°C change in temperature, the accuracy of measurement is 0.4°C for a 60 s time interval, but 0.18°C for a 5 minute time interval. The accuracy increases the longer the time interval, that is to say the time between successive measurements. During particularly sensitive operations, such as the start-up of a reactor block 10, it may be preferable to increase the frequency of measurement.

The data analysis device 21 may be an Optical Time Domain Reflectometer or a Optical Frequency Domain Reflectometer or it may be a combined device comprising the functionality of both time and frequency domain reflectometry. The data analysis device 21 is configured to analyse all data collected from the fibre optic cable 19. The extent and nature of this analysis will change according to the reactor conditions and the nature of the procedures being carried out. For example, during normal operation of the reactor block 10, the temperature measurements from the entire reactor block 10 may be analysed in order to

provide maximum temperature, minimum temperature, average temperature and rate of change of temperature across the reactor block 10 as a whole.

In addition, the full temperature profile of the reactor block 10 may be calculated. This additional data may be accessed periodically during normal operation, for example on a daily basis. This may provide an early indication of degradation of the catalyst in part of the reactor block 10 if the operating temperature in that part of the reactor block 10 differs from the average temperature calculated across the reactor block 10 as a whole.

Moreover, access to the full profile will be especially useful during sensitive operations such as start-up of the reactor block 10 and during such operations the full profile may be calculated continuously. In addition the data analysis device 21 may be configured to respond to reactor conditions, for example, providing full profile data when the rate of change of temperature exceeds a predetermined threshold which may be indicative of a thermal runaway.

The data analysis device 21 is linked to a control system 25 used to control the operation of the reactor block 10. The data analysis device 21, control system 25 and reactor block 10 may collectively be referred to as the plant 30. The data analysis device 21 is configured to provide some level of automatic control of the reactor block 10. If the rate of change of temperature exceeds a predetermined value, the data analysis device 21 is configured to send a signal to the control system 25 to effect either a change in the composition of the feed gas or a partial or total shutdown of the reactor block 10. For example, when a first predetermined threshold is reached, the percentage of inert gases in the feed may be increased to slow the reaction. If this does not stabilise the temperature within the reactor block 10 and a second predetermined threshold value is reached for the rate of change of temperature, then the reactor may be shut in under 100% inert gases. The data analysis device 21 may have multiple channels, each of which obtains data from one fibre optic cable 19. There may be 2, 4, 6, 8 or more channels. Each fibre optic cable 19 may be up to 2 km long. For example it may be 1.5 m, 50 m, 500 m or 1 km long.

Figure 3 shows two examples of temperature profiles of a reactor block 10. The flow channels 17 through which the reagents for the Fischer-Tropsch synthesis reaction flow run parallel to the Y-axis. In each case the full temperature profile showed no significant variation in temperature, T, across the width of the reactor block 10; in Figure 3 the temperature profiles are shown along the Y-axis, that is to say in the direction of flow in the channels 17. Figure 3a shows the temperature profile of the reactor block 10 in a steady

state of a Fischer-Tropsch synthesis reaction. The temperature T rises initially in the first 20% of the length of the reactor block 10. The temperature T then falls throughout the remaining length of the reactor block 10. The temperature T varies between a maximum of 205°C and a minimum of 200°C, averaging 202°C. The maximum rate of change of temperature is no more than 0.1°C/min (and may in practice be zero).

Figure 3b shows an instantaneous example of a temperature profile of the reactor block 10 when a thermal runaway is taking place. The temperature in the first quarter of the length of the reactor block 10 rises very steeply before dropping quite steeply. The high temperatures in the first part of the reactor block 10 cause all of the reagents to be used and therefore the reaction rate in the remainder of the reactor block 10 is lower than it would be under steady state conditions. The maximum temperature is 220°C and the minimum is 200°C, averaging 205°C. The maximum rate of change of temperature is 0.8°C/min.

It will be appreciated that since measurements are taken for example over successive 12 cm lengths along the optical fibre 19, with a wide reactor block there may be several measurements taken at different positions across the width of the reactor block 10. So not only can temperature variations along the length of the reactor block 10 be monitored, but temperature variations across the width can also be monitored. The increased level of detail in temperature throughout the reactor block 10 (for example as shown in Figures 3a and 3b), enables the process to be further optimised and thereby operated closer to operating boundaries, reducing the required safety margins and thereby increasing yield.

The provision of the fibre optic cable 19 within the divider plates 13a, 13b enables the fibre optic cable 19 to be replaced without opening up the reactor block 10. Whilst the fibre optic cable 19 should be capable of withstanding a thermal runaway, there may still be circumstances in which replacing the fibre optic cable 19 may be preferable. Although the fibre-optic cable 19 is shown only shown as passing through a single reactor block 10, it will be appreciated that a continuous fibre-optic cable 19 may either be provided with an end termination (not shown) at the end remote from the data analysis device 21, or may be passed in a similar way through tubes 13e in another such reactor block 10.

Referring now to Figure 4, there is shown a modification to the reactor block 10 suitable for use in performing Fischer-Tropsch synthesis. The reactor block 10 is shown in section and only in part. The reactor block 10 is the same as the reactor block 10 of Figure 2a in that it consists of a stack of flat plates 12 of thickness 1 mm spaced apart so as to define channels 15 for a coolant fluid alternating with channels 117 for the Fischer-Tropsch synthesis. In this

modification a fibre optic cable 19 is arranged within the second flow channels 15, through which the coolant flows, in use. In another alternative, the fibre-optic cable 19 may instead extend through the first flow channels 117 in which the reagents for Fischer-Tropsch synthesis flow in use. The provision of the fibre optic cable 19 in the process channel, rather than within the divider plates, makes it more sensitive to temperature excursions, but it is a more hostile environment in the process channel in comparison with the divider plate and therefore the operational lifetime of the fibre optic cable 19 may be compromised. In a further alternative one fibre optic cable 19 may extend through the flow channels 15, as shown in figure 4, or through the flow channels 117, and another fibre-optic cable 19 may extend through the tubes 13e, as shown in figure 2a.

It will be appreciated that the size of the reaction channel 17 or 117 may differ from that described above. However, the reaction channels 17 or 117 are preferably at least 1 mm deep, preferably at least 2 mm deep, to provide adequate space for catalyst; and are preferably no more than 20 mm deep, more preferably no more than 10 mm deep, as it is difficult to ensure substantially uniform temperature throughout such a deep channel. The smallest dimension of the cross section of the first flow channel may be greater than or equal to 1 mm and less than or equal to 50 mm. The smallest dimension of the cross section of the first flow channel may be greater than or equal to 2 or 4 or 6 or 8 or 10 mm and less than or equal to 50 mm. The smallest dimension of the cross section of the first flow channel may be greater than or equal to 2 or 4 or 6 or 8 or 10 mm and less than or equal to 20 mm. The smallest dimension of the cross section of the first flow channel may be greater than or equal to 2 or 4 or 6 or 8 mm and less than or equal to 10 mm. The smallest dimension of the cross section of the first flow channel may be greater than or equal to 2 or 4 or 6 mm and less than or equal to 8 mm.

The smallest dimension of the cross section of the first flow channel may be the same or similar to the orthogonal dimension of the cross section of the first flow channel. A reactor provided with at least first flow channels that are substantially square in cross section may be particularly advantageous where an insert structure is used in each first flow channel, either a catalyst-carrying insert or an insert solely for enhancing heat transfer.

Conversely, a reactor provided with elongate channels, in which a first dimension orthogonal to the sheets considerably exceeds the extent of the channel in the direction 12 parallel to the sheets 12, may provide advantages when a highly active catalyst is used. The elongation of the channel in the first direction reduces the total number of layers provided within a reactor of a given size and therefore the amount of metal used may also be reduced. This

may result in considerable cost savings. This elongate dimension is balanced by the smaller extent of the channel in the orthogonal direction. This provides a reduced heat transfer distance to the fins on either side of the channel, and aids the structural integrity of the reactor.

5

Where the chemical reaction is endothermic, then the heat transfer fluid in the second flow channels may be a hot fluid, or may be a fluid which undergoes an exothermic reaction such as combustion. Where the chemical reaction is exothermic, then the heat transfer fluid in the second flow channels may be a coolant fluid. By way of example the catalytic reactor may be for performing Fischer-Tropsch synthesis.

10

It should be understood that the Fischer-Tropsch reaction is a comparatively slow reaction. The purpose of the Fischer-Tropsch synthesis is to generate hydrocarbons in which the carbon chain is longer than that of methane, and indeed preferably at least C5 and so are normally liquids or waxes. A practical reactor must therefore generate a significant quantity of such longer-chain hydrocarbons per unit time, and should be selective towards the formation of such longer-chain hydrocarbons rather than methane. It has been found that providing a greater volumetric loading of active catalyst material can enhance both the conversion of CO to hydrocarbons, and also the productivity of the desired hydrocarbons.

15

20

The Fischer-Tropsch reaction is typically carried out at a temperature less than or equal to 300°C and typically about 200°C, so a wide range of materials may be selected for the reactor. For example the reactor may be made of an aluminium alloy, stainless steel, high-nickel alloys, or other steel alloys. It should be realized that the foregoing example embodiments should not be construed as limiting. Other variations and modifications will be apparent to persons skilled in the art upon reading the present application.

25

Moreover, the disclosure of the present application should be understood to include any novel features or any novel combination of features either explicitly or implicitly disclosed herein or any generalization thereof, and during the prosecution of the present application or of any application derived therefrom, new claims may be formulated to cover any such features and/or combination of such features.

30

CLAIMS

1. A catalytic reactor comprising
a plurality of first flow channels including a catalyst for a first reaction
5 a plurality of second flow channels arranged alternately with the first flow channels;
adjacent first and second flow channels being separated by a divider plate ; and
a distributed temperature sensor provided within the divider plate and/or within one or more
first flow channels and/or one or more second flow channels.
- 10 2. The catalytic reactor according to claim 1, wherein the distributed temperature sensor is a
fibre optic cable.
3. The catalytic reactor according to claim 2, wherein the fibre optic cable follows a tortuous
path through divider plates in order to provide temperature data for all of the adjacent first
15 and second flow channels.
4. The catalytic reactor according to claim 2 or claim 3, wherein a fibre optic cable is
provided in each divider plate.
- 20 5. The catalytic reactor according to claim 2, wherein the fibre optic cable is provided within
one or more first flow channels.
6. The catalytic reactor according to claim 2 or claim 5, wherein the fibre optic cable is
provided within one or more second flow channels.
- 25 7. The catalytic reactor according to any one of the preceding claims, wherein the first flow
channels include a Fischer-Tropsch catalyst and the second channels carry a heat transfer
fluid, in use.
- 30 8. The catalytic reactor according to any one of claims 1 to 6, wherein the first flow channels
include a steam methane reforming catalyst and the second channels further include a
combustion catalyst.
- 35 9. A plant for performing Fischer-Tropsch synthesis, the plant comprising a catalytic reactor
according to any one of the preceding claims, and a data analysis device.

10. The plant according to claim 9, further comprising a control system operatively linked to the data analysis device.

11. The plant according to claim 10, wherein the control system is configured to shut in the
5 reactor if the data analysis device detects a rate of change of temperature in excess of a predetermined value.

Fig. 1a.

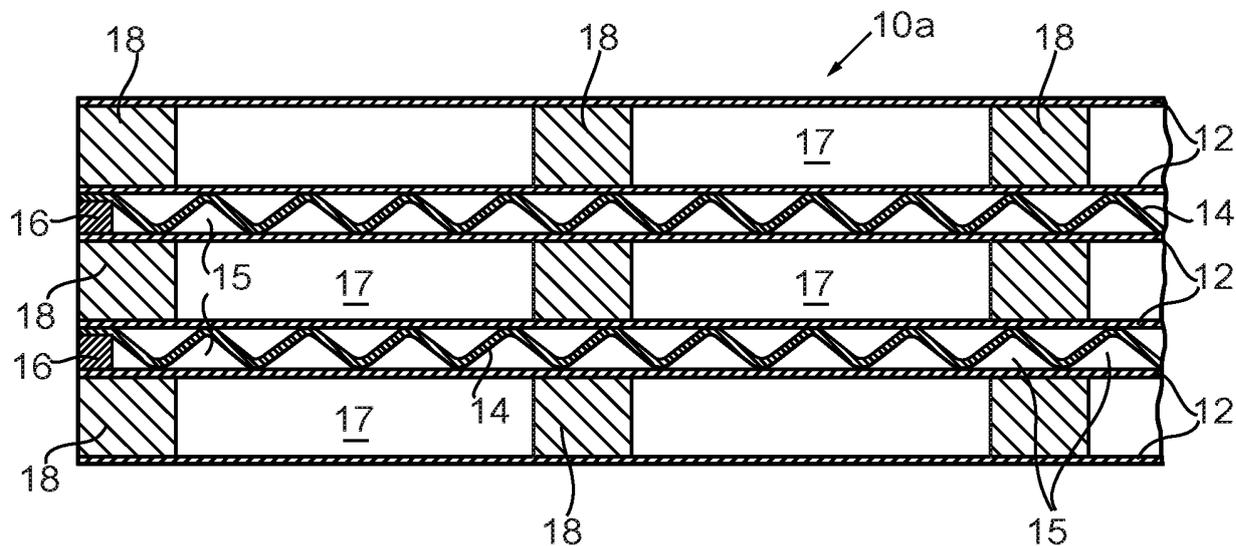
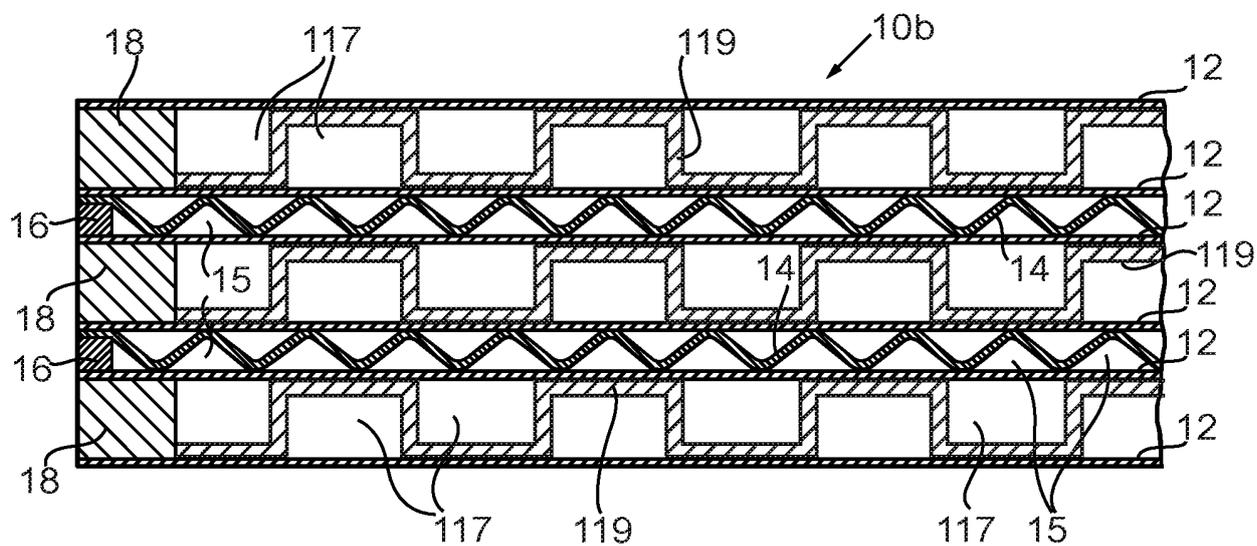


Fig. 1b.



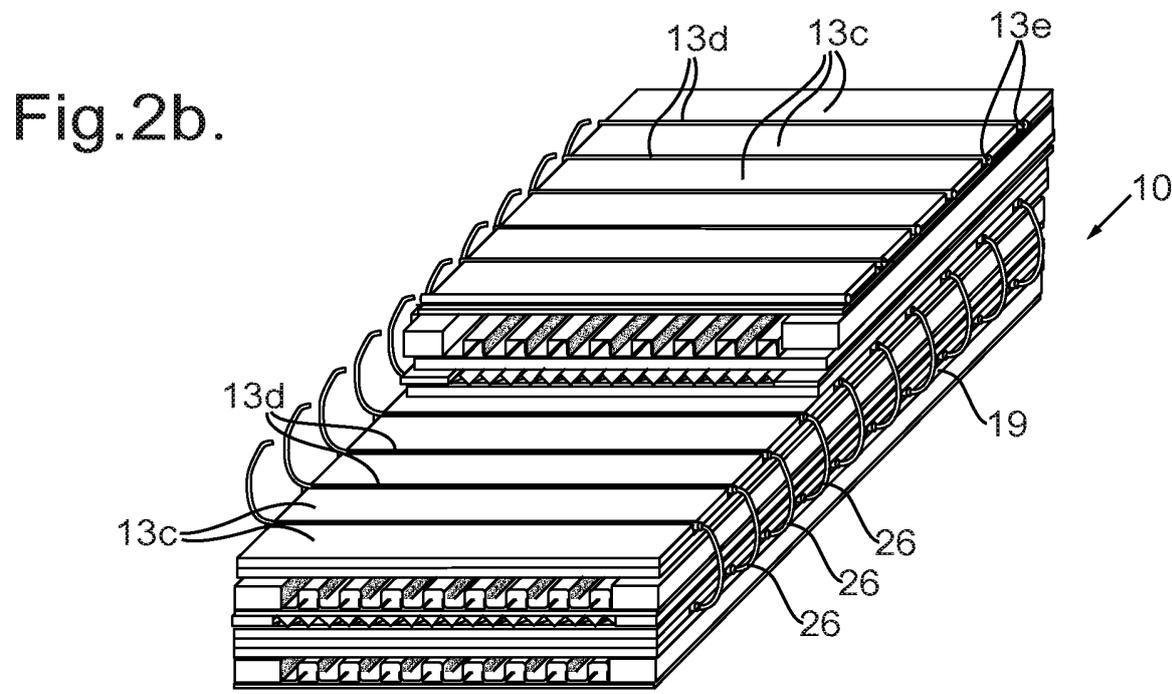
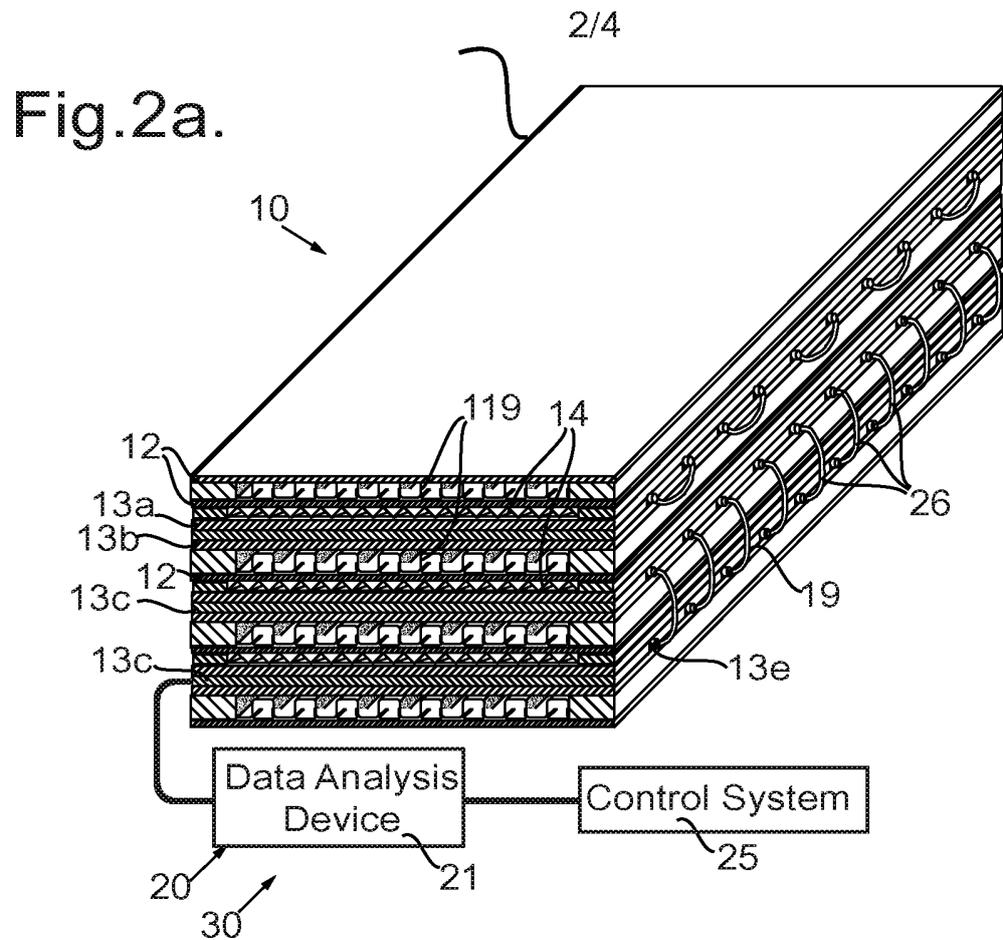


Fig.3a.

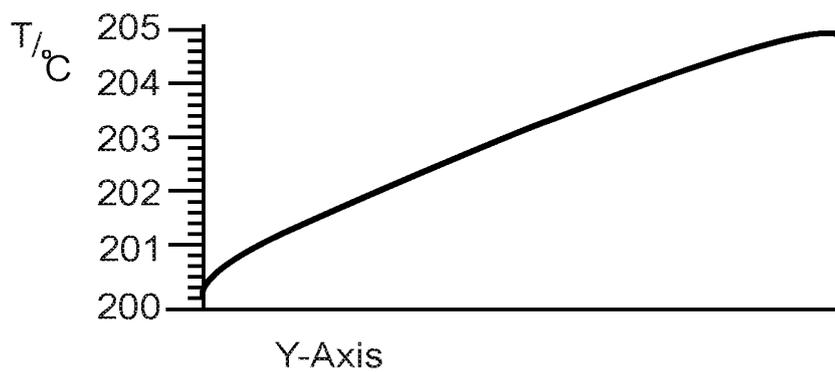


Fig.3b.

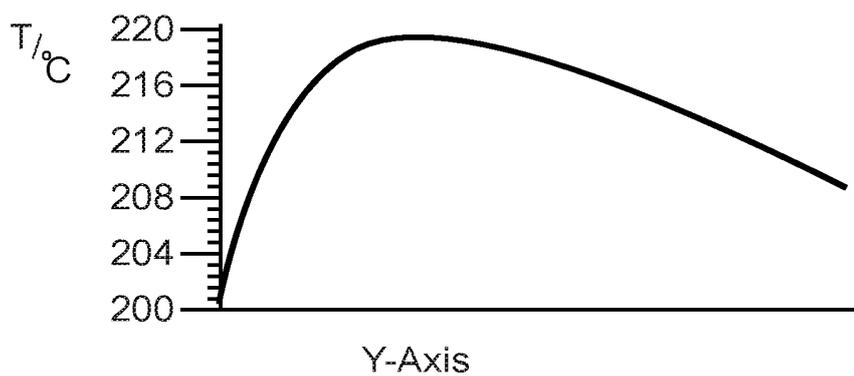
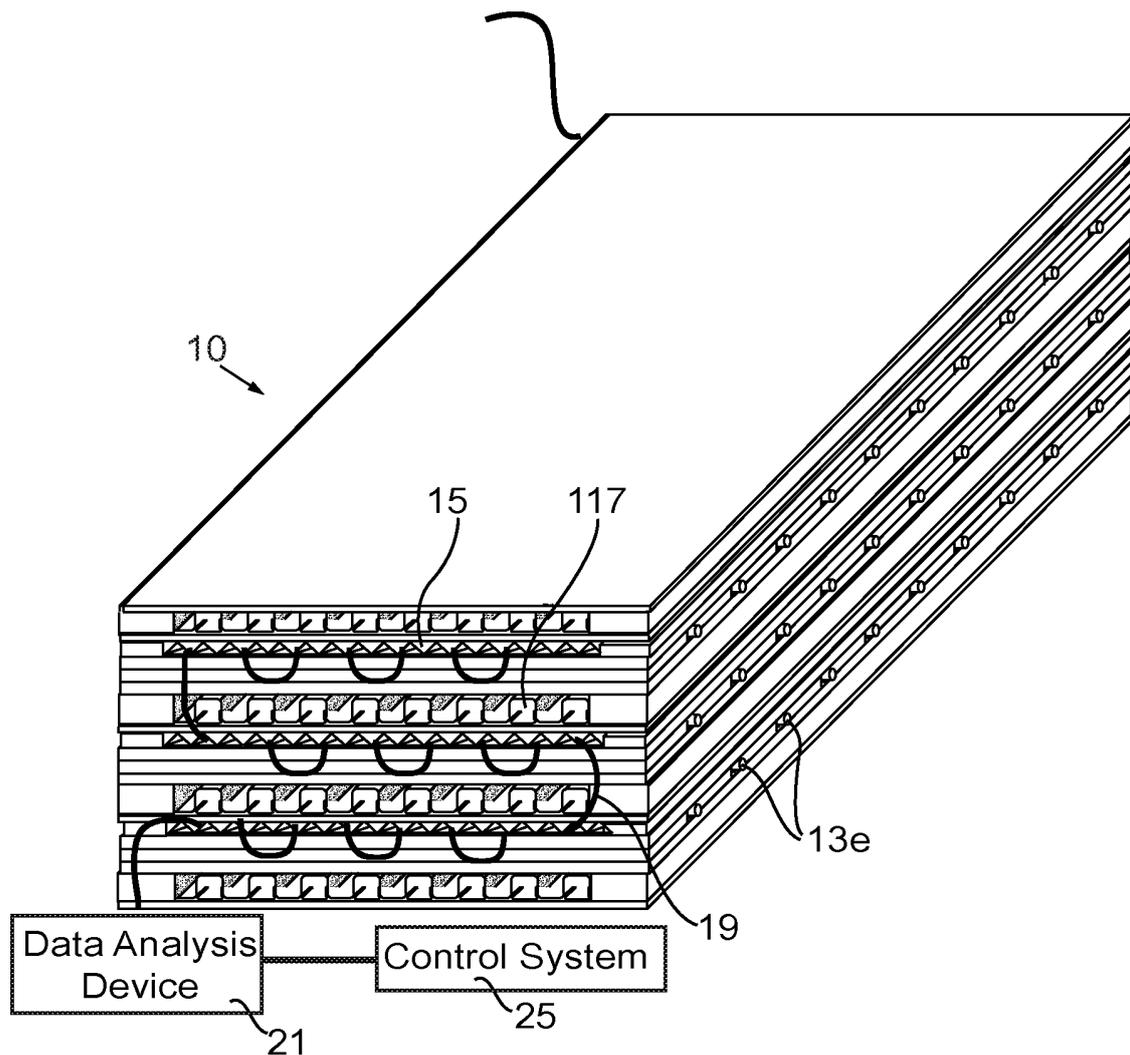


Fig.4.



Appendix B. Abstracts of the Master's theses

B.1 (Day et al., 2017)

As biomass boilers become more common across the UK and Europe, the need to understand how to maximise their efficiency and output increases. There has been extensive research into most properties of a biomass energy system, ranging from fuel type to the waste gas treatment leaving the boilers.

In this work, we discuss an often-overlooked issue relating to the flue gas filtration. End-of-pipe clean-up is often seen as a necessary evil to meet legal and environmental requirements, and therefore gains little attention as a potential source of significant efficiency savings. All energy producing sites in the UK must keep emissions within the environmental permit granted to them by the local authorities. To stay within these limits, biomass power generation sites use a variety of methods, including spraying additives to remove acidic and heavy metal contaminants, as well as using a fabric filter to remove solid particulate matter.

In this study, the efficiency of the flue gas filtration system at a UK-based 33MW biomass fired boiler is investigated. The filter media and cleaning methods are discussed in detail and plant historic data analysed with the aim of establishing enhanced operating strategies to reduce total cost of ownership of biomass flue gas filtration systems. This approach is rare in academic literature, and has been achieved by collaborating with a major UK filtration supplier, Durham Filtration, along with a multinational utilities company and a global filter media manufacturer.

In this work, we present power plant performance data as a function of filter time on stream, normalised against turbine load. This step allows for like for like comparisons to be made between several operating points over many years. Furthermore, we combine maintenance records, offline analysis, historic plant data, commercial power and raw material pricing and multivariate statistics to give economic models of filtration performance.

We find that optimal life span of a filter bag is two years, after which, bag wear and filter blinding result in high pressure drop and increased stack emissions. We present a methodology for benchmarking different grades and brands of filter media in operation, allowing for predictions of bag life and OPEX to be made. When comparing two brands of

filter media, we conclude that there was a significant difference in operating and mechanical qualities of the brands under investigation.

Finally, future work on the site under investigation is recommended, with work to the dry sorbent addition rate required, after the 30% correlation between lime addition and HCl concentration highlighted the excessive addition rate currently in operation.

B.2 (Eagle et al., 2017)

End-of-pipe filtration is chronically under-optimised in industry; it is viewed as a necessary operating expense to meet air quality standards. In power generation plants, flue gas filters present a significant pressure drop and consequently, compressor duty. Uneven flow distribution causes disproportionate wear within a filtration plant and can result in shorter filter life and therefore higher maintenance costs.

In this work we demonstrate a rapid and computationally efficient methodology for the optimal design of flow enhancing geometries to reduce the operating expense of a 33 MW biomass power plant. We propose a novel design optimisation approach that combines Computational Fluid Dynamics (CFD), Design of Experiments (DoE) and real plant data.

Dry Sorbent Injection (DSI) testing compared the mixing abilities of two technologies in order to reduce DSI dosage. Sorb-tech nozzle incorporated boosted air during injection to increase DSI residence time by more than double the original technology. Findings support a reduction in DSI dosage could be made due to enhanced mixing.

The multi-compartment, pulse jet baghouse for a 33 MW biomass power plant is simulated in three dimensional CFD software. Flow enhancing geometries (ladder vanes) were parameterised (angle, length, thickness etc.), and carried forward into a DoE structure. Regression modelling of real plant historical data was used to provide a correlation between various plant parameters (differential pressure, velocity profile etc.). Optimal geometry was found ($\theta_A = 58^\circ$, $\theta_v = 96^\circ$) and re-simulated to quantify Model-Model Mismatch (MMM).

B.3 (Taylor et al., 2018)

This report focuses on the optimisation of the pulse cleaning system within the baghouse filtration unit of Wilton 10 (W10), a biomass energy plant in Teesside. To conform to environmental regulations, impurities must be removed from the process stream. Subsequently, build-up of impurities occurs on the surface of the filter bags with in the baghouse. Pulse cleaning is utilised as a means of extending bag life and reducing the pressure drop across the baghouse. Using computational fluid dynamics to model geometric improvements to

the W10 pulse line, peak pressure is used as the primary metric within this report to quantify improvements to cleaning efficiency. Demonstration of correlation between peak pressure and operational costs provide scope for further work to quantify the savings in operation costs from this study.

Results of a batch-to-batch geometric improvements yielded that including a basic nozzle geometry (10mm diameter by 10mm height) on the pulse holes of the system resulted in a 33.8% increase in peak pressure ($p < 0.01$) adding isolation baffles to this design extended this to a 36.4% increase. Due to this, a Design of experiment analysis was performed to achieve a more optimised nozzle geometry which could be re-implemented into W10 for benchmarking. Response surface analysis showed local optima for the three dimensional factors tested (nozzle diameter 9 mm, height 5 mm and length from filter bag 115 mm). The optimised nozzle yielded a resultant increase in peak pressure of 16.8% ($p < 0.01$) in comparison to the current system configuration.

By demonstrating a methodology of how geometric improvements can systematically increase peak pressure the cleaning efficiency of any industrial plant can be optimised. As well as quantifying this improvement through reduced operating costs.

B.4 (Pettitt et al., 2018)

The optimal operating parameters of pulse jet cleaning systems are poorly understood, however, impact significantly on the expense of bag filter house operation. The optimal pulse interval is investigated for a 33 MWe industrial biomass plant; Wilton 10. The purpose of the cleaning system is to prevent an accumulation of fabric dust loading which would increase the filtration resistance. However, evidence that the pulse jet cleaning mechanism results in additional resistance due to dust incorporated into the bag fibres themselves (blinding) has been identified.

The optimal pulse interval is proposed to be a compromise between minimising surface dust loading whilst also that of blinding. Filter bags from the plant were taken offline periodically and analysed to determine a profile for the surface dust loading and blinding separately. Increasing the pulse interval from 15 minutes to 60 minutes was found to correspond to reduced differential pressure across the bag filter house, reducing operating costs of the induced draft fan and pulse jet system by £152,000/year for 3 μm diameter fly ash. Additional cost savings of increased bag lifetime and reduced maintenance were also identified. Increasing the particle diameter of the fly ash also reduced differential pressure across the bag filter house and operating costs due to an increased fraction of solids falling directly into the hoppers.

References

- N. D. A., B. P. E., and U. Eric. *Optical Fibers*, chapter 2, pages 9–33. Wiley-Blackwell, 2011. ISBN 9781118014103. doi: 10.1002/9781118014103.ch2. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118014103.ch2>.
- A. H. Abdullah, R. Mat, S. Somderam, A. S. A. Aziz, and A. Mohamed. Hydrogen Sulfide Adsorption by Zinc Oxide-Impregnated Zeolite (Synthesized from Malaysian Kaolin) for Biogas Desulfurization. *Journal of Industrial and Engineering Chemistry*, 2018. ISSN 1226-086X. doi: <https://doi.org/10.1016/j.jiec.2018.05.003>. URL <http://www.sciencedirect.com/science/article/pii/S1226086X18302211>.
- AFL Global. VHM3000 Series Harsh Environment Fibers., 2015. URL <http://www.aflglobal.com/Products/Specialty-Optical-Fiber/Harsh-Environments-Products/VHM3000-Series-Harsh-Environment-Fibers.aspx>.
- C. H. Bartholomew. History of Cobalt Catalyst Design for Fischer-Tropsch Synthesis. In *Proceedings of the American Institute of Chemical Engineers Spring National Meeting*, New Orleans, 2003.
- D. S. Beachler, J. Joseph, and M. Pompelia. Fabric Filter Operation Overview, 1995.
- J. W. Berthold. *Industrial Applications of Fiber Optic Sensors*, chapter 13, pages 349–371. Wiley-Blackwell, 2011. ISBN 9781118014103. doi: 10.1002/9781118014103.ch13. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118014103.ch13>.
- G. L. Bezemer, J. H. Bitter, H. P. C. E. Kuipers, H. Oosterbeek, J. E. Holewijn, X. Xu, F. Kapteijn, A. J. van Dillen, and K. P. de Jong. Cobalt Particle Size Effects in the Fischer-Tropsch-Tropsch Reaction Studied with Carbon Nanofiber Supported Catalysts. *Journal of the American Chemical Society*, 128(12):3956–3964, mar 2006. ISSN 0002-7863. doi: 10.1021/ja058282w. URL <https://doi.org/10.1021/ja058282w>.
- R. Bivand and K. Krivoruchko. Big data sampling and spatial analysis: “which of the two ladles, of fig-wood or gold, is appropriate to the soup and the pot?”. *Statistics & Probability Letters*, 2018. ISSN 0167-7152. doi: <https://doi.org/10.1016/j.spl.2018.02.012>. URL <http://www.sciencedirect.com/science/article/pii/S0167715218300579>.

-
- C. W. Bonner and B. Dannatt. Catalytic Reactors Comprising Distributed Temperature Sensors. World Intellectual Property Organization Patent no: 2016/001663 A1., 2015.
- BSI. BS EN 29073-1:1992 Methods of test for nonwovens. Determination of mass per unit area. Technical report, British Standards Institute, London, 1990. URL <https://doi.org/10.3403/00373938>.
- BSI. BS EN ISO 9073-2:1997 Textiles. Test methods for nonwovens. Determination of thickness. Technical report, British Standards Institute, London, 1997. URL <https://doi.org/10.3403/00970517>.
- BSI. BS EN ISO 9073-18:2008 Textiles. Test methods for nonwovens. Determination of breaking strength and elongation of nonwoven materials using the grab tensile test. Technical report, British Standards Institute, London, 2008a. URL <https://doi.org/10.3403/30148436>.
- BSI. BS EN ISO 9073-15:2008 Textiles. Test methods for nonwovens. Determination of air permeability. Technical report, British Standards Institute, London, 2008b. URL <https://doi.org/10.3403/30148424>.
- BSI. BS ISO 18842:2015 Aluminium oxide primarily used for the production of aluminium. Method for the determination of tapped and untapped density. Technical report, British Standards Institute, London, 2017a. URL <https://doi.org/10.3403/30293302>.
- BSI. BS EN 451-1:2017 Method of testing fly ash. Determination of free calcium oxide content. Technical report, British Standards Institute, London, 2017b. URL <https://doi.org/10.3403/30323496>.
- S. Caillat and E. Vakkilainen. 9 - Large-scale biomass combustion plants: an overview. In L. Rosendahl, editor, *Biomass Combustion Science, Technology and Engineering*, Woodhead Publishing Series in Energy, pages 189–224. Woodhead Publishing, 2013. ISBN 978-0-85709-131-4. doi: <https://doi.org/10.1533/9780857097439.3.189>. URL <https://www.sciencedirect.com/science/article/pii/B9780857091314500091>.
- C.-J. Chen. Reducing Pressure Drop in a Baghouse Using Flow Distributors. *Journal of the Air & Waste Management Association*, 51(10):1471–1475, 2001. doi: 10.1080/10473289.2001.10464364. URL <https://doi.org/10.1080/10473289.2001.10464364>.
- C.-J. Chen and M.-T. Cheng. Effect of Flow Distributors on Uniformity of Velocity Profile in a Baghouse. *Journal of the Air & Waste Management Association*, 55(7): 886–892, 2005. doi: 10.1080/10473289.2005.10464678. URL <https://doi.org/10.1080/10473289.2005.10464678>.

-
- D. S. Clery, P. E. Mason, C. M. Rayner, and J. M. Jones. The effects of an additive on the release of potassium in biomass combustion. *Fuel*, 214:647–655, 2018. ISSN 0016-2361. doi: <https://doi.org/10.1016/j.fuel.2017.11.040>. URL <http://www.sciencedirect.com/science/article/pii/S0016236117314424>.
- CompactGTL. CompactGTL website, 2015. URL compactgtl.com.
- B. H. Davis. Fischer-Tropsch Synthesis: Comparison of Performances of Iron and Cobalt Catalysts. *Industrial & Engineering Chemistry Research*, 46(26):8938–8945, 2007. doi: 10.1021/ie0712434. URL <https://doi.org/10.1021/ie0712434>.
- L. Day, B. Dannatt, B. Goulden, and M. Willis. Cost Benefit Analysis of Filter Bag Performance Using Historic Operational Data. Master’s thesis, Newcastle University, 2017.
- A. de Klerk. *Fischer-Tropsch Synthesis*, chapter 4, pages 73–103. Wiley-Blackwell, 2011. ISBN 9783527635603. doi: 10.1002/9783527635603.ch4. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9783527635603.ch4>.
- A. de Klerk and E. Furimsky. *Catalysis in the Refining of Fischer-Tropsch-Tropsch Syncrude*. RSC Catalysis Series. The Royal Society of Chemistry, 2010. ISBN 978-1-84973-080-8. doi: 10.1039/9781849732017. URL <http://dx.doi.org/10.1039/9781849732017>.
- H. Eagle, B. Dannatt, B. Goulden, and M. Willis. Rapid Hybrid Computational Design of Industrial Energy Saving Flow Geometries - Biomass Fired Boiler Case Study. Master’s thesis, Newcastle University, 2017.
- S. Ergun. Fluid flow through packed columns. *Chemical Engineering Science*, 48(2):89–94, 1952. URL <https://ci.nii.ac.jp/naid/10003393451/en/>.
- T. O. Eschemann, J. Oenema, and K. P. de Jong. Effects of noble metal promotion for Co/TiO₂ Fischer-Tropsch catalysts. *Catalysis Today*, 261:60–66, 2016. ISSN 0920-5861. doi: <https://doi.org/10.1016/j.cattod.2015.06.016>. URL <http://www.sciencedirect.com/science/article/pii/S0920586115003818>.
- European Commission. Directive 2008/1/EC concerning integrated pollution prevention and control (Codified version). *Official Journal of the European Union*, L 24/8:8–29, 2008.
- European Commission. Directive 2001/80/EC of the European Parliament and of the Council of 23 October 2001 on the limitation of emissions of certain pollutants into the air from large combustion plants. *Official Journal of the European Union*, (309):1–21, 2001. ISSN 0144-557X. doi: 10.1039/ap9842100196.
- European Commission. DIRECTIVE OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the limitation of emissions of certain pollutants into the air from

-
- medium combustion plants . *Official Journal of the European Union*, 0442(1386):19, 2013.
- Foster Wheeler. Biomass Plant Operation Manual. Technical report, Varkaus, Finland, 2006.
- E. A. Foumeny, F. Benyahia, J. A. A. Castro, H. A. Moallemi, and S. Roshani. Correlations of pressure drop in packed beds taking into account the effect of confining wall. *International Journal of Heat and Mass Transfer*, 36(2):536–540, 1993. ISSN 0017-9310. doi: [https://doi.org/10.1016/0017-9310\(93\)80028-S](https://doi.org/10.1016/0017-9310(93)80028-S). URL <http://www.sciencedirect.com/science/article/pii/001793109380028S>.
- J. Glassey and M. von Stosch. *Hybrid Modeling in Process Industries*. CRC Press, 2018. ISBN 9781351184366. URL <https://books.google.co.uk/books?id=TQpQDwAAQBAJ>.
- S. Hamilton, S. Fleming, and R. Stewart. The assessment of flue gas particulate abatement in wood burning boilers. *Report for Forestry Commission Scotland ED56285*, 3:1–39, 2010.
- M. Hupa, O. Karlström, and E. Vainio. Biomass combustion technology development – It is all about chemical details. *Proceedings of the Combustion Institute*, 36(1):113–134, 2017. ISSN 1540-7489. doi: <https://doi.org/10.1016/j.proci.2016.06.152>. URL <http://www.sciencedirect.com/science/article/pii/S1540748916302103>.
- E. Iglesia. Design, synthesis, and use of cobalt-based Fischer-Tropsch-Tropsch synthesis catalysts. *Applied Catalysis A: General*, 161(1):59–78, 1997. ISSN 0926-860X. doi: [https://doi.org/10.1016/S0926-860X\(97\)00186-5](https://doi.org/10.1016/S0926-860X(97)00186-5). URL <http://www.sciencedirect.com/science/article/pii/S0926860X97001865>.
- F. P. Kapron, D. B. Keck, and R. D. Maurer. Radiation Losses in Optical Waveguides. *Applied Physics Letters*, 17(10):423–425, 1970. doi: 10.1063/1.1653255. URL <https://doi.org/10.1063/1.1653255>.
- S. Korhaliller. The UK’s Biomass Energy Development Path. (November):1–45, 2010. URL <http://pubs.iied.org/pdfs/G02921.pdf>.
- Lydall-Gutsche. Leading the way to zero emissions. Website,, 2015. URL <http://www.lydall-gutsche.com/en/high-performance-filtration-2/gutsche-microvel/>.
- I. F. Macdonald, M. S. El-Sayed, K. Mow, and F. A. L. Dullien. Flow through Porous Media-the Ergun Equation Revisited. *Industrial & Engineering Chemistry Fundamentals*, 18(3):199–208, 1979. doi: 10.1021/i160071a001. URL <https://doi.org/10.1021/i160071a001>.

-
- N. Mao. 17 - Engineering design of high-performance filter fabrics. In M. Miao and J. H. Xin, editors, *Engineering of High-Performance Textiles*, The Textile Institute Book Series, pages 435–488. Woodhead Publishing, 2018. ISBN 978-0-08-101273-4. doi: <https://doi.org/10.1016/B978-0-08-101273-4.00018-4>. URL <https://www.sciencedirect.com/science/article/pii/B9780081012734000184>.
- M. Matúš, P. Križan, L. Šooš, and J. Beniak. The effect of papermaking sludge as an additive to biomass pellets on the final quality of the fuel. *Fuel*, 219:196–204, 2018. ISSN 0016-2361. doi: <https://doi.org/10.1016/j.fuel.2018.01.089>. URL <http://www.sciencedirect.com/science/article/pii/S0016236118300978>.
- J. D. McKenna and J. H. Turner. *Fabric Filter– Baghouses I: Theory, Design, and Selection : (a Reference Text)*. Number v. 1. ETS, 1989. URL <https://books.google.co.uk/books?id=cGtwHAAACAAJ>.
- D. J. Moodley. *On the Deactivation of Cobalt-based Fischer-Tropsch-Tropsch Synthesis Catalysts door*. PhD thesis, Eindhoven University of Technology, 2008. URL <http://alexandria.tue.nl/extra2/200811712.pdf>.
- D. J. Moodley, J. Loosdrecht, van de, A. M. Saib, M. J. Overett, A. K. Datye, and J. W. Niemantsverdriet. Carbon deposition as a deactivation mechanism of cobalt-based Fischer-Tropsch-Tropsch synthesis catalysts under realistic conditions. *Applied Catalysis. A, General*, 354(1-2):102–110, 2009. ISSN 0926-860X. doi: 10.1016/j.apcata.2008.11.015.
- S. Mousavi, A. Zamaniyan, M. Irani, and M. Rashidzadeh. Generalized kinetic model for iron and cobalt based Fischer-Tropsch synthesis catalysts: Review and model evaluation. *Applied Catalysis A, General*, 506(Complete):57–66, 2015. doi: 10.1016/j.apcata.2015.08.020.
- P. Niehe. 3D makeover for hyper-efficient metalwork, 2015. URL <https://www.arup.com/news-and-events/news/3d-makeover-for-hyperefficient-metalwork>.
- F. P and T. K. Optimizing the Numerical Aperture in Graded Index Fibers , 1982. URL <https://www.degruyter.com/view/j/joc.1982.3.3/joc.1982.3.3.101/joc.1982.3.3.101.xml>.
- B. P. Pal. *Fundamentals of Fibre Optics in Telecommunication and Sensor Systems*. New Age International (P) Limited, New Delhi, India, India, 2015. ISBN 8122436625, 9788122436624.
- C. Pask and R. A. Sammut. Experimental characterisation of graded-index single-mode fibres. *Electronics Letters*, 16(9):310–311, 1980. ISSN 0013-5194 VO - 16. doi: 10.1049/el:19800223.

-
- C. Pask, A. W. Snyder, and D. J. Mitchell. Number of modes on optical waveguides. *J. Opt. Soc. Am.*, 65(3):356–357, mar 1975. doi: 10.1364/JOSA.65.000356. URL <http://www.osapublishing.org/abstract.cfm?URI=josa-65-3-356>.
- L. Pettitt, B. Dannatt, B. Goulden, and M. Willis. Modelling and Optimisation of The Pulse-Jet Cleaning Mechanism in the Filter Bag House of a 33 MWe Industrial Biomass Combustion Plant. Master's thesis, Newcastle University, 2018.
- T. N. Phaahlamohlaka, M. W. Dlamini, M. W. Mogodi, D. O. Kumi, L. L. Jewell, D. G. Billing, and N. J. Coville. A sinter resistant Co Fischer-Tropsch catalyst promoted with Ru and supported on titania encapsulated by mesoporous silica. *Applied Catalysis A: General*, 552:129–137, 2018. ISSN 0926-860X. doi: <https://doi.org/10.1016/j.apcata.2017.12.015>. URL <http://www.sciencedirect.com/science/article/pii/S0926860X17305707>.
- S. Pissot, T. B. Vilches, H. Thunman, and M. Seemann. Effect of ash circulation on the performance of a dual fluidized bed gasification system. *Biomass and Bioenergy*, 115: 45–55, 2018. ISSN 0961-9534. doi: <https://doi.org/10.1016/j.biombioe.2018.04.010>. URL <http://www.sciencedirect.com/science/article/pii/S0961953418300990>.
- V. I. Popa. 1 - Biomass for Fuels and Biomaterials. In V. Popa and I. Volf, editors, *Biomass as Renewable Raw Material to Obtain Bioproducts of High-Tech Value*, pages 1–37. Elsevier, 2018. ISBN 978-0-444-63774-1. doi: <https://doi.org/10.1016/B978-0-444-63774-1.00001-6>. URL <https://www.sciencedirect.com/science/article/pii/B9780444637741000016>.
- A. Rezeau, L. I. Díez, J. Royo, and M. Díaz-Ramírez. Efficient diagnosis of grate-fired biomass boilers by a simplified CFD-based approach. *Fuel Processing Technology*, 171: 318–329, 2018. ISSN 0378-3820. doi: <https://doi.org/10.1016/j.fuproc.2017.11.024>. URL <http://www.sciencedirect.com/science/article/pii/S0378382017312900>.
- L. J. Roberts, P. E. Mason, J. M. Jones, W. F. Gale, A. Williams, and C. Ellul. Investigating the impact of an Al-Si additive on the resistivity of biomass ashes. *Fuel Processing Technology*, 178:13–23, 2018. ISSN 0378-3820. doi: <https://doi.org/10.1016/j.fuproc.2018.05.018>. URL <http://www.sciencedirect.com/science/article/pii/S0378382018301644>.
- B. Schmid. VIB.jar - A 3D version of the minimum, maximum and median filter., 2014.
- H. Schulz. Short history and present trends of Fischer-Tropsch synthesis. *Applied Catalysis A: General*, 186(1):3–12, 1999. ISSN 0926-860X. doi: [https://doi.org/10.1016/S0926-860X\(99\)00160-X](https://doi.org/10.1016/S0926-860X(99)00160-X). URL <http://www.sciencedirect.com/science/article/pii/S0926860X9900160X>.

-
- H. Schulz and M. Claeys. Kinetic modelling of Fischer-Tropsch product distributions. *Applied Catalysis A: General*, 186(1):91–107, 1999. ISSN 0926-860X. doi: [https://doi.org/10.1016/S0926-860X\(99\)00166-0](https://doi.org/10.1016/S0926-860X(99)00166-0). URL <http://www.sciencedirect.com/science/article/pii/S0926860X99001660>.
- M. B. Shrestha and G. R. Bhatta. Selecting appropriate methodological framework for time series data analysis. *The Journal of Finance and Data Science*, 4(2):71–89, 2018. ISSN 2405-9188. doi: <https://doi.org/10.1016/j.jfds.2017.11.001>. URL <http://www.sciencedirect.com/science/article/pii/S2405918817300405>.
- R. Singh and A. Shukla. A review on methods of flue gas cleaning from combustion of biomass. *Renewable and Sustainable Energy Reviews*, 29:854–864, 2014. ISSN 1364-0321. doi: <https://doi.org/10.1016/j.rser.2013.09.005>. URL <http://www.sciencedirect.com/science/article/pii/S1364032113006643>.
- A. P. Steynberg, S. R. Deshmukh, and H. J. Robota. Fischer-Tropsch catalyst deactivation in commercial microchannel reactor operation. *Catalysis Today*, 299:10–13, 2018. ISSN 0920-5861. doi: <https://doi.org/10.1016/j.cattod.2017.05.064>. URL <http://www.sciencedirect.com/science/article/pii/S0920586117303863>.
- K. Storm. 6 - Biomass Plant Equipment. In K. Storm, editor, *Industrial Piping and Equipment Estimating Manual*, pages 195–235. Gulf Professional Publishing, 2017. ISBN 978-0-12-813946-2. doi: <https://doi.org/10.1016/B978-0-12-813946-2.00006-X>. URL <https://www.sciencedirect.com/science/article/pii/B978012813946200006X>.
- J. H. Taghinia, M. M. Rahman, and X. Lu. Effects of different CFD modeling approaches and simplification of shape on prediction of flow field around manikin. *Energy and Buildings*, 170:47–60, 2018. ISSN 0378-7788. doi: <https://doi.org/10.1016/j.enbuild.2018.03.075>. URL <http://www.sciencedirect.com/science/article/pii/S0378778817335648>.
- D. Taler, M. Trojan, P. Dzierwa, K. Kaczmarek, and J. Taler. Numerical simulation of convective superheaters in steam boilers. *International Journal of Thermal Sciences*, 129:320–333, 2018. ISSN 1290-0729. doi: <https://doi.org/10.1016/j.ijthermalsci.2018.03.005>. URL <http://www.sciencedirect.com/science/article/pii/S1290072918300243>.
- T. Tanuma. 1 - Introduction to steam turbines for power plants. In T. Tanuma, editor, *Advances in Steam Turbines for Modern Power Plants*, pages 3–9. Woodhead Publishing, 2017. ISBN 978-0-08-100314-5. doi: <https://doi.org/10.1016/B978-0-08-100314-5.00001-4>. URL <https://www.sciencedirect.com/science/article/pii/B9780081003145000014>.

-
- C. Taylor, B. Dannatt, B. Goulden, and M. Willis. Industrial Baghouse Pulse-Jet Optimisation - A Hybrid Modelling Approach. Master's thesis, Newcastle University, 2018.
- D. Tschumperlé and R. Deriche. Vector-valued image regularization with PDEs: a common framework for different applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):506–517, apr 2005. ISSN 0162-8828. doi: 10.1109/TPAMI.2005.87.
- UK Government. Digest of UK Energy Statistics (DUKES): renewable sources of energy. Technical report, 2017. URL https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/633782/Chapter_6.pdf.
- J. van de Loosdrecht, B. Balzhinimaev, J.-A. Dalmon, J. W. Niemantsverdriet, S. V. Tsybulya, A. M. Saib, P. J. van Berge, and J. L. Visagie. Cobalt Fischer-Tropsch synthesis: Deactivation by oxidation? *Catalysis Today*, 123(1):293–302, 2007. ISSN 0920-5861. doi: <https://doi.org/10.1016/j.cattod.2007.02.032>. URL <http://www.sciencedirect.com/science/article/pii/S092058610700140X>.
- S. V. Vassilev, C. G. Vassileva, Y.-C. Song, W.-Y. Li, and J. Feng. Ash contents and ash-forming elements of biomass and their significance for solid biofuel combustion. *Fuel*, 208:377–409, 2017. ISSN 0016-2361. doi: <https://doi.org/10.1016/j.fuel.2017.07.036>. URL <http://www.sciencedirect.com/science/article/pii/S0016236117308906>.
- L. Wang, N. Pereire, and Y. Hung. *Handbook of environmental engineering: Air pollution control engineering*. Humana Press, 2004. ISBN 1617373974. URL <https://link.springer.com/book/10.1007%2F978-1-59259-778-9>.
- R. Willingale. University of Leicester Department of Physics and Astronomy Lecture Notes, 2007. URL <http://www.star.le.ac.uk/zrw/courses/lect4313.html>.
- G. Wilson. Comparison of Emissions Monitoring from existing biomass plants. Technical report, SKM ENVIROS, Shrewsbury, 2012. URL [https://infrastructure.planninginspectorate.gov.uk/wp-content/ipc/uploads/project/Comparison_of_emissions_from_existing_biomass_plants_\(Appendix_7.6_of_Environmental_Statement\).pdf](https://infrastructure.planninginspectorate.gov.uk/wp-content/ipc/uploads/project/Comparison_of_emissions_from_existing_biomass_plants_(Appendix_7.6_of_Environmental_Statement).pdf).
- D. A. Wood, C. Nwaoha, and B. F. Towler. Gas-to-liquids (GTL): A review of an industry offering several routes for monetizing natural gas. *Journal of Natural Gas Science and Engineering*, 9:196–208, 2012. ISSN 1875-5100. doi: <https://doi.org/10.1016/j.jngse.2012.07.001>. URL <http://www.sciencedirect.com/science/article/pii/S1875510012000947>.

-
- Q. Wu, H. Yang, M. Wei, O. Remil, B. Wang, and J. Wang. Automatic 3D reconstruction of electrical substation scene from LiDAR point cloud. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2018. ISSN 0924-2716. doi: <https://doi.org/10.1016/j.isprsjprs.2018.04.024>. URL <http://www.sciencedirect.com/science/article/pii/S0924271618301321>.
- W. Yang, Y. Zhu, W. Cheng, H. Sang, H. Xu, H. Yang, and H. Chen. Effect of minerals and binders on particulate matter emission from biomass pellets combustion. *Applied Energy*, 215:106–115, 2018. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2018.01.093>. URL <http://www.sciencedirect.com/science/article/pii/S0306261918301077>.
- L. Zhu, M. Li, and R. R. Martin. Direct simulation for CAD models undergoing parametric modifications. *Computer-Aided Design*, 78:3–13, 2016. ISSN 0010-4485. doi: <https://doi.org/10.1016/j.cad.2016.05.006>. URL <http://www.sciencedirect.com/science/article/pii/S0010448516300264>.