

A Novel Method for the Manipulation of Biological Cells using Fluid Flow and Realtime Control

by

Michael Curtis

A Thesis submitted to Monash University
for the degree of
Doctor of Philosophy

Department of Mechanical and Aerospace Engineering
Monash University

March 2016

Statement of Originality

I, Michael Curtis, declare that this thesis is my own work and contains no material that has been accepted for the award of a degree or diploma in this, or any other, university. To the best of my knowledge and belief, information derived from the published and unpublished work of others has been acknowledged in the text of the thesis and a list of references is provided in the bibliography.

Candidate: Michael Curtis

Submitted: March 2016

Copyright Statement

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Declaration for thesis based or partially based on jointly published or unpublished work

In accordance with Monash University Doctorate Regulation 17.2 Doctor of Philosophy and Research Master's regulations the following declarations are made:

I hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes one original paper published in peer reviewed journals. The core theme of the thesis is the Manipulation of Biological Cells using Fluid Flow. The ideas, development and writing up of all the papers in the thesis were the principal responsibility of myself, the candidate, working within the Department of Mechanical and Aerospace Engineering under the supervision of Prof. Andreas Fouras and A/Prof. Gregory Sheard.

The inclusion of co-authors reflects the fact that the work came from active collaboration between researchers and acknowledges input into team-based research.

In the case of chapter 3 my contribution to the work involved the following:

Thesis Chapter	Publication Title	Publication Status	Nature and extent of candidate's contribution
3	Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems	Published	Coding, Simulation, Analysis, Figures, Text

I have not renumbered sections of submitted or published papers in order to generate a consistent presentation within the thesis.

Sunday 6th March, 2016

Abstract

Cells are complex, with respect to both their biochemistry and mechanical properties. This complexity is multiplied further by the increasing body of evidence that the biochemical properties and mechanical properties may, in many cases, be intrinsically linked. Therefore, there is a clear unmet need for sophisticated and automated measurement tools that can produce quantitative measurements of cell deformation and cell responses to forces. In particular, the time-dependent response of cells to imposed force is of interest to the biological sciences, but there is a lack of tools that can measure this aspect of cell behaviour. Automation brings the capability to measure whole populations of cells and thereby provide both statistical relevance and repeatability. This is especially important in biological samples, which are known to vary widely in properties and behaviour on both an intra- and inter-sample basis.

In this thesis, the technology and capabilities have been developed that allow, for the first time, the study of cell mechanics in a non-contact, purely fluidic environment, in a fully automated fashion. To achieve this, a number of innovative techniques have been postulated, developed, proven and demonstrated. These include a novel control system simulator for the control of coupled fluid-particle systems, innovative approaches for realtime flow feedback through realtime image analysis, tracking cells in noisy conditions, and an advanced flow controller designed specifically for high bandwidth manipulation of flow in microfluidic systems.

The techniques developed and proven in this thesis allow the conception and construction of software-defined systems to manipulate not only cells, but also oil droplets, particles, and other suspended objects. These dynamic manipulation and measurement methods represent a significant advance over existing techniques for cell and droplet manipulation in diagnostics and research. In contrast to other methods, detailed multi-dimensional measurements (for example, both force-extension profiles and relaxation times) can be obtained using a single laboratory device, and on an individual cell basis, rather than through bulk measurement.

The technology has been successfully applied to the measurement of key mechanical properties of live red blood cells. In particular, the ability to trap cells in an automated fashion at high extensional rates (corresponding to a physiologically-relevant net force), and measure the resultant dynamic behaviour of the cells, greatly expands upon the existing tools available for cell analysis. Cells were obtained, subjected to a minimum of processing, and the dynamic changes in cell shape tracked as the cell was subjected to varying extensional forces. Unlike many of the preceding works, this was accomplished without the use of foreign manipulation devices — such as lasers or direct physical contact — and requires a minimum of operator intervention. By measuring both the steady-state and dynamic behaviours of cells, both the stiffness and time constant (viscosity) were obtained for individual cells. The work presented in this thesis enables researchers and clinicians to efficiently measure cell stiffness and viscosity simultaneously from single cells. This greatly improves the predictive and diagnostic power of cell mechanics measurements available to researchers and has the potential to revolutionise both our understanding of blood-related diseases (such as diabetes and malaria) and point-of-care diagnostics.

This thesis documents a number of the scientific and technical advances over the existing body of knowledge (Part I), across a number of fields (including imaging, fluid dynamics and control systems), that have arisen from the development of this technology (Part II). In addition, extensive validation work (Part III) will demonstrate the application of the work in Part II to the analysis of the properties of biological cells.

Acknowledgements

Firstly, I would like to acknowledge my supervisors, Prof Andreas Fouras, A/Prof Greg Sheard and Prof Kerry Hourigan, for their tireless support, encouragement, late night conversations and advice throughout this thesis. Without your help, this thesis would never have existed.

I also want acknowledge and thank my family for their support. My parents, Allan & Jenny, for sharing in my successes as well as disappointments. My wife, Rosemary, for being there to pick up the pieces, for sharing in all that I do. To Rosemary's family, for being my family away from home.

Of course, any large endeavour requires the help of a great many people. I would like to thank the present and past members of the Laboratory of Dynamic Imaging (Monash) for many fruitful discussions, for the opportunity to gain their advice, to learn from their mistakes (as they have no doubt learned from mine!), and for the many lunchroom conversations.

I would like to thank Dr Josie Carberry, Dr Elizabeth Gardiner (AMREP) for their advice in fluids and biology, respectively. Also thank you to Dr James Armitage (Monash/Deakin) for patiently explaining clinical practice and practical biology, and for assistance with experimental design, showing myself and my supervisors the clinical relevance of this technology, and for providing the inspiration to get it working.

I am indebted to the many talented tradespeople and technicians at Monash. Mr Hugh Venables, Mr Nat Derose and Mr Eric Wirth from the Mechanical Workshop have taught me a great many things about manufacturing and transformed my — sometimes unclear — drawings and the odd strange material or special order component into working parts. To Mr Mark Symonds and Mr Ken Sheridan, for their assistance with electrical assembly. To Mr Ian Reynolds, of the Department of Electrical and Computer Systems Engineering, who willingly helped out with the manufacture of microfluidic moulds, even though it was outside his department, and was using a machine that was intended for PCB prototyping.

List of Publications

Arising from the work reported in this thesis

- Michael D. Curtis, Gregory J. Sheard, and Andreas Fouras. “Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems”. In: *Lab Chip* 11.14 (June 28, 2011), pp. 2343–2351. ISSN: 1473-0189. DOI: 10.1039/C1LC20191C
- Andreas Fouras, Gregory J. Sheard, **Michael D. Curtis**, and James Armitage. “Method and device for application of fluid forces to cells”. Australian Patent AU2012/245075; US Patent 14,112,914. Licensed to 4Dx Pty Ltd.

Additional publications submitted during candidature

- Jordan Thurgood, Stephen Dubsy, Kentaro Uesugi, **Michael Curtis**, Rajeev Samarage, Bruce Thompson, Graeme Zosky, and Andreas Fouras. “Imaging lung tissue oscillations using high-speed X-ray velocimetry”. In: *Journal of Synchrotron Radiation* 23 (January, 2016), pp. 324-330. DOI: 10.1107/S1600577515021700

Contents

I. Background	1
1. Introduction and Review of Existing Literature	3
1.1. Concept	4
1.2. Physiological Importance	4
1.2.1. Active Mechanical Response in Cells	4
1.2.2. Cells as Transient Systems	6
1.3. Cell Manipulation Techniques	7
1.3.1. Micro-pipette Aspiration	8
1.3.2. Mechanical Manipulation	9
1.3.3. Magnetic Manipulation	9
1.3.4. Laser Tweezers	10
1.3.5. Laser Stretchers	10
1.4. Existing Microfluidic Systems for Manipulation	11
1.4.1. Fixed Microfluidic Channels	11
1.4.2. Rheoscopes and Cytometers	13
1.4.3. Saddle Point Flows	13
1.5. Similar Work	16
1.5.1. Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers	16
1.5.2. Hydrodynamic stretching of Single Cells for Large Population Mechanical Phenotyping	16
1.5.3. Cell Stretching Measurement Utilizing Viscoelastic Particle Focusing	17
1.5.4. Characterizing the Performance of the Hydrodynamic Trap Using a Control-Based Approach	17
1.6. Conclusions	17
1.7. Bibliography	17
2. System Overview and Microchannel Design	23
2.1. Preliminary Requirements	23
2.2. Design Parameters and Constraints	23
2.3. Prototype System	25
2.4. System Architecture	25
2.5. Microchannel Design and Construction	30
2.6. Conclusions	33
2.7. Bibliography	33
II. Advances toward Software-Defined Microfluidics	35
3. Simulation	37
3.1. CFD Simulation and Validation	37
3.1.1. Mesh Independence	39
3.1.2. Effect of a Cell on the Flow	39

3.1.3.	Effect of Aspect Ratio	42
3.1.4.	System Parameterisation	43
3.1.5.	Effect of Boundary Conditions	44
3.2.	Paper: Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems	48
3.3.	Tuning the System Geometry	58
3.4.	Conclusion	60
3.5.	Bibliography	60
4.	Flow Measurement	61
4.1.	Review of Flow Measurement Techniques	61
4.1.1.	Differential Pressure	61
4.1.2.	Thermal Techniques	61
4.1.3.	Ultrasound	62
4.1.4.	Particle Image Velocimetry	62
4.1.4.1.	Double-Pulsed PIV	64
4.2.	Realtime, GPU-accelerated, PIV	65
4.2.1.	Efficient Windowing	67
4.2.2.	Peak Detection	67
4.2.3.	Performance and Accuracy	69
4.2.4.	Accuracy and Robustness	70
4.2.5.	Performance	71
4.3.	Identification of the Saddle-Point	73
4.3.1.	Error and Accuracy Analysis	74
4.4.	Parameter Optimisation	77
4.5.	Adapting to Flow Velocity Changes	79
4.5.1.	Conventional Approaches	79
4.5.2.	Timing Modulation Approach	81
4.6.	Conclusion	84
4.7.	Bibliography	84
5.	Flow Control	87
5.1.	Existing Designs	87
5.1.1.	Syringe Pump	87
5.1.2.	Pressure Pumps	88
5.1.3.	On-Chip Multilayer Control Valves	88
5.1.4.	Off-Chip Diaphragm Valves	88
5.1.5.	Off-Chip Pinch Valves	89
5.1.6.	Summary	89
5.2.	Design and Development	89
5.2.1.	Actuator Type	89
5.2.2.	Challenges	91
5.2.2.1.	Relationship between Control Input and Flow	92
5.2.2.2.	Trapped Air and Fouling	92
5.2.2.3.	Hysteresis	92
5.2.2.4.	Pumping	94
5.2.3.	Final Design	96
5.2.3.1.	Bayonet Cap Design	96

5.3.	Closed-Loop Control	97
5.4.	Hybrid Closed-Open Loop Control with Calibration	98
5.4.1.	Calibration Procedure	101
5.4.2.	Calibration Analysis	102
5.4.3.	Forward Predictor Model	105
5.5.	The Motion Planner	107
5.5.1.	Tracking Valve State	107
5.5.2.	Kinematics Model	111
5.6.	Performance	112
5.7.	Conclusion	119
5.8.	Bibliography	119
6.	Cell Tracking	121
6.1.	Background Subtraction and Image Normalisation	122
6.2.	Segmentation	123
6.3.	Tracking	125
6.3.1.	Searching State	127
6.3.2.	Tracking State	128
6.3.3.	Recovery State	130
6.4.	Performance	131
6.5.	Validation	133
6.5.1.	Synthetic Image Benchmark	133
6.5.2.	Blood Cell Testing	135
6.6.	Conclusion	135
6.7.	Bibliography	136
III.	Application and Experimental Results	139
7.	Trapping Particles and Cells	141
7.1.	Trapping Particles	141
7.2.	Simulation	142
7.3.	Tuning the System for Higher Flow Velocities	148
7.4.	Extended Manipulation	151
7.5.	Conclusion	152
7.6.	Bibliography	154
8.	Clinical Pilot Study	155
8.1.	Experimental Procedure	155
8.2.	Measuring Cell Deformation	156
8.2.1.	Extracting Cell Shape	156
8.2.2.	Fitting the Ellipse	157
8.2.3.	Validation	158
8.3.	Time-Varying Behaviour of Cell Shape	159
8.3.1.	Minimum Deformation Forces	159
8.4.	Disease Model: Heat Treatment	161
8.4.1.	Steady-State Stiffness Changes Due to Heat Treatment	161
8.4.2.	Dynamic Deformation as a Sensitive Diagnostic	163

Contents

8.5. Conclusion	169
8.6. Bibliography	169
IV. Conclusion and Closing Remarks	171
9. Conclusions	173
9.1. Biomedical Relevance	173
9.2. Technical Advances	174
9.2.1. Coupled Fluid Control Simulation	174
9.2.2. Flow Measurement	174
9.2.3. High-Bandwidth Flow Control	174
9.2.4. Novel Realtime Cell Tracking	174
9.3. Closing Remarks	175
V. Appendices	177
A. Sample Configuration	179
B. Electronic Supplementary Information	183
B.3. Simulation	183
B.8. Clinical Pilot Study	183

List of Figures

1.1. Schematic of the proposed fluidic system, showing generalised flow directions and control principles	5
1.2. A red blood cell traversing the wall separating the splenic cord and sinus	7
1.3. Some of the observed shapes of a normal erythrocyte cell	8
1.4. Hyperbolic channel used to apply a constant elongational force to erythrocytes in suspension as they traverse the channel	11
1.5. Contracting microchannel designs for the study of platelet activation dynamics	13
1.6. One of the earliest rheoscopes for measuring cell deformation	14
1.7. Cross Slot device for observing polymer hysteresis	15
1.8. Eight way microfluidic cross, capable of producing rotational and extensional flow	15
2.1. Pressure-driven flow control concept.	24
2.2. Control system view.	26
2.3. CAD render of prototype system, showing the major components.	28
2.4. Component view of the major components in the complete system.	29
2.5. As in this cart-damper-cart analogy, a rigid system has a direct, lag-free, relationship between input and output	31
2.6. Microfluidic Chip to Tubing Interface, showing (a) a cross-section of one port assembly, (b) the finished chip, ready for use in the system	32
3.1. Geometries used for CFD simulations.	38
3.2. Mesh independence study (2D)	40
3.3. Mesh independence study (3D)	41
3.4. Effect of a stationary cell on the flow.	42
3.5. Extensional rate in the central region as a function of aspect ratio.	43
3.6. Visualisation of the velocity magnitude near the saddle-point, and the change in saddle-point position as a function of flow rate bias.	45
3.7. Effect of the flow rate bias on the position of the saddle-point and the flow structures.	46
3.8. Out-of-plane (3D) behaviour of the saddle-point position.	47
3.9. Velocity in the central region as a function of corner radius.	59
3.10. Saddle-point position as a function of flow rate bias for the ‘large’ channel	59
4.1. Schematic of a typical orifice plate flowmeter. A small, known, restriction is placed inline with the flow and the pressure drop across the restriction is measured. With this method, both the bulk flow rate and direction can be measured directly.	62
4.2. Schematic of process of obtaining flow velocities using Particle Image Velocimetry (PIV).	63
4.3. Example Double-Pulsed PIV Timing.	65
4.4. Expanded double-pulsed timing diagram, showing camera typical exposure and readout timing.	66
4.5. Double-pulsed PIV flowchart, showing the key steps in the efficient computation of velocity vectors.	68

List of Figures

4.6. Results of applying equation (4.3) in a situation where the spacing between windows is not a multiple of the image size. Window 3 extends past the end of the image array and, as 2D arrays are actually 1D arrays in memory, the remainder of the window “wraps” around the image and contains part of Window 1 (shaded region).	69
4.7. Synthetic flows used for validation of the GPU PIV code	70
4.8. Contour plot of the percentage velocity error of the analytical field estimate (4.7) as compared to a solution from a CFD solver. For much of the imageable area, the velocity error is less than 5 %, demonstrating that the analytical approximation produces a result close to the original. Errors increase further from the centre as the approximation does not account for the parabolic shape of the inlet and outlet velocities (Poiseuille flow).	75
4.9. Cell position (in the vertical, or outlet plane) over time for a number of saddle-point position accuracies.	76
4.10. Parameter variation study on the linearity and accuracy of the saddle-point estimation algorithm presented in section §4.3. Dashed lines in c) and d) indicate the region where the non-linearity error is ± 5 px.	78
4.11. Accuracy and Computational Performance of PIV and the saddle-point identification algorithm at various FFT Window Sizes and Vector Spacings.	80
4.12. Adaptive Gap Time Control block diagram.	81
4.13. Behaviour of Adaptive Gap Time Control during a flow rate ramp.	83
5.1. Air trapped in the system causes a severe lag in response.	93
5.2. Normalised area of a deformable silicone tube as a function of applied external pressure.	94
5.3. Generic model of a hysteretic system.	95
5.4. Peristaltic pumping effect as a result of rapid changes in valve actuator position. As the valve closes, the tubing wall is displaced by a certain amount. This results in the displacement of a corresponding volume of fluid (shaded region). Depending on the speed of valve closure, this may result in a significant, transient, change in flow rate (see section 5.2.2.4).	95
5.5. CAD renderings of the final pinch valve design; a) shows an exterior isometric view of the valve with bayonet cap removed and tubing inserted, b) shows a cross-section of the valve and its major components.	96
5.6. Exploded view of the Bayonet Cap, showing the major components.	97
5.7. PID feedback control of the saddle point position.	99
5.8. The inherent delays in the system that lead to a 15 ms delay from the initial change of flow conditions to the first measurement of that change.	100
5.9. Drift in the saddle-point as a result of system pressure changes.	100
5.10. Simplified control systems diagram, showing the essential elements of the hybrid approach, combining a fast open-loop controller with a slower closed loop compensation loop.	101
5.11. Results of initial calibration sweep for all four flow control valves.	103
5.12. Detailed calibration sweeps for the the left and bottom valves at two different channel sizes.	105
5.13. Sigmoidal fit to the valve calibration data, using data from bottom valve with a small (high flow rate) channel and the left valve with a large (low flow rate) channel.	107
5.14. Lookup table (LUT) formed from computing the sigmoid fits for both the positive-going and negative-going calibration curves, and then inverting these curves.	108
5.15. Finite State Machine model of the valve state tracking algorithm.	110

5.16. Traversal of the lookup table by the motion planner.	111
5.17. Step Response of the system with different shaping parameters applied by the motion planner.	114
5.18. Through the coordination of two opposing valves (in this case, the top and bottom valves), the performance of large steps can be further improved.	115
5.19. Step response of the flow control system at various net driving pressures.	116
5.20. Validation of calibration, open-loop estimation and motion planner.	118
6.1. Background removal and normalisation.	124
6.2. Region-growing image segmentation process.	125
6.3. Comparison of segmentation and thresholding algorithms.	126
6.4. Finite State Machine for the cell tracker.	127
6.5. Streamlines for a typical particle and the associated relationship with the distance deviation score p_D	130
6.6. Recovery of original track after interference due to a particle passing close to the tracked cell.	132
6.7. Thread occupancy diagram, showing the asynchronous nature of the streamline and histogram computations, relative to the tracker, and the asynchronous relationship between the main application and the tracker.	134
7.1. Comparison between microscopy images for a) a hollow glass sphere (red blood cell analogue); b) red blood cell	141
7.2. Particle and saddle-point positions during an active trap.	143
7.3. Typical processing pipeline, showing the data flow and inherent delays in the system that determine the minimum time that the system can respond to an event (change in position of the cell). Two frames are shown to indicate the continuous and alternating nature of the system control loop.	144
7.4. Typical flow (saddle-point position) noise while the system is idle.	145
7.5. Hysteresis and overshoot simulation, with experimental data as comparison.	147
7.6. Cell Positions in the y -plane from the simulation and corresponding experiments of a cell trap with an average flow velocity of $\bar{v} = 0.50 \text{ mms}^{-1}$ ($\dot{\epsilon} = 5 \text{ s}^{-1}$). Positions in simulation inverted for ease of comparison.	149
7.7. Trapping experiment and corresponding simulation, after tuning, with an average flow velocity of $\bar{v} = 0.75 \text{ mms}^{-1}$ ($\dot{\epsilon} = 10 \text{ s}^{-1}$).	151
7.8. Extended manipulation of a cell during trapping.	153
8.1. Illustration of the effect of the cost functions used in fitting an ellipse to sparse point data.	158
8.2. Deformation of a healthy red blood cell over time, as a series of images at key points in time.	160
8.3. A plot of Taylor Deformation Index (top) and cell orientation (bottom) against time, showing the effect of cell orientation on deformation behaviour.	162
8.4. Histogram of steady-state deformation of heat-treated and control red blood cells, where the y -axis is the normalised probability density. Normal distributions have been fitted to the distributions in both cases. The mean Taylor deformation index of the heat treated cells was significantly lower ($\text{DI} = 0.277$) than the non-treated cells ($\text{DI} = 0.372$, $P = 0.0003$, one-sided t -test).	163

List of Figures

8.5. Cell deformation over time for a number of red blood cells. (a) control (untreated) cells; (b) heat-treated cells. All lines are spline fits to the data; dotted lines indicate the final value for the heat treated cases.	166
8.6. Cell deformation over time for a number of red blood cells, showing the raw data points from the fitting algorithm and the associated spline fits. As is evident from the figure, this type of fit is effective at rejecting random noise and preserves the trends in the original data.	167
8.7. Time Constant and Final Deformation plotted as a 2-D map.	168

Nomenclature

A number of chapters in this thesis involve the discussion of mathematical concepts. For clarity, and as the subject matter spans a number of fields, each with their own particular conventions, the common nomenclature used in this thesis are briefly presented here:

Symbol	Meaning
$a, b, \text{etc.}$	Ordinary scalar
$\vec{a}, \vec{b}, \text{etc.}$	Vector quantity (German-style notation)
$\mathbf{a}, \mathbf{b}, \text{etc.}$	Matrix quantity
$f(x)$	Continuous function f on x
$f[x]$	Sampled, or discrete function on x (commonly used in signal processing literature)
$\mathcal{O}(N)$	“On the order of (N) .” An indicator of runtime scaling of a particular software algorithm. (e.g., $\mathcal{O}(N^2)$ means that the runtime is proportional to the square of the problem size.)

Part I.

Background

“Begin at the beginning,” the King said, very gravely, “and go on till you come to the end: then stop.”

— Lewis Carroll, *Alice in Wonderland*

1. Introduction and Review of Existing Literature

Cells are exposed to fluid forces in a variety of physiological contexts, including blood flow, air movement in lung alveoli,^{44,73} and shear stress on cardiomyocytes.^{10,31} Microfluidic devices, with similar length scales to blood vessels and other fluid transport channels in biological organisms, enable the construction of a customised bio-mimetic environment in the laboratory. By customising the flow geometry, and hence the fluidic environment of cells, forces can be applied to cells that mimic those encountered *in vivo*.^{41,58,70} In addition, microchannels can also include features such as constrictions,³³ precisely tuned outlet sizes^{30,68} and pillars⁴⁵ to direct and sort cells based on their size and behaviour in flow. Using these geometric structures, large numbers of cells can be processed rapidly, enabling high throughput analysis. However, the force profiles and sorting decisions are **static** — encoded in the geometry at the time of manufacture.

In contrast to these static devices, cells are dynamic. Cells have classical rheological characteristics, such as viscoelasticity,⁴³ and additionally exhibit active biochemical responses to applied force.³² The response by cells to externally-applied forces is complex and non-linear, due to interactions between membrane structure, proteins (such as haemoglobin) which may attach themselves to cell membranes,²⁸ internal organelles in the cell and viscous behaviour of the internal fluid. Pathological conditions modify the biochemical properties of cells, changing the force response. For example, it has been observed that patients with diabetes mellitus have decreased red blood cell stiffness,^{14,35,39} which may contribute to the secondary (co-morbid) conditions, such as impaired circulation and vascular retinopathy.^{22,54,55} The malaria parasite modifies both the membrane chemistry and the internal structure of red blood cells, causing a dramatic increase in red cell stiffness.^{42,46} Environmental forces also provoke active responses, such as directing the differentiation of stem cells;⁸ driving stiffness changes in blood vessel endothelial cells;⁴⁷ and causing otherwise healthy cells to become malignant.^{6,37}

Investigating dynamic behaviour, such as the response to varying forces applied to a cell in real-time, requires active control. In microfluidic systems, forces are usually applied to suspended cells in conjunction with a complementary technology, such as electric fields;^{5,25} magnetism;^{49,66} optical forces;²⁷ and surface acoustic waves.²⁴ These technologies require involved cell preparation or impose unnatural forces on the cells. Conversely, the use of the fluids to directly manipulate cells offers a preparation-free, bio-mimetic approach with a low risk of cell damage. In addition, the use of direct fluid manipulation allows multiple, independent, fluid control systems to be densely integrated on a single chip.⁵¹ In this thesis, significant advances have been made towards technology that enables a highly integrated, flexible device that makes use of active control to trap and manipulate red blood cells in a microfluidic environment.

Through the work conducted in this thesis, a new type of microfluidic system has been created — one that is substantially defined by the software systems that control it, rather than the physical geometry of the device. Using high-speed, automated control systems, the flow in the device can be adjusted in realtime to produce a wide variety of subtle and intelligent behaviours that would be impossible, or at least very difficult, to produce in a conventional microfluidic system. This work shows the development of the technology and its subsequent application to create a device that is able to steer; trap; stretch and sort biological cells in a way that is selective; comprehensive; automatic; and high-throughput.

1.1. Concept

The concept, developed extensively in this thesis, builds on an idea first popularised by Schroeder et al.⁶⁰ for the investigation of polymer hysteresis (both synthetic and natural — i.e., DNA — polymers were investigated). Instead of targeting a polymer, the goal of the studies presented in this thesis is to manipulate a particle or cell such that it remains trapped and stationary in the centre of an extensional flow (figure 1.1). While in this location, the cell experiences a combination of extensional (vertical direction in the example shown in the figure) and compressive (horizontal direction) forces, and undergoes a shape change. By simultaneously measuring this shape change and the fluid velocity, the dynamic stiffness behaviour of the cell can be inferred. This central location is also a region of zero fluid velocity — a stagnation point, or **saddle-point**. In theory, a particle that is perfectly placed at this saddle-point will remain there indefinitely, with the external forces acting on it perfectly balanced. However, any slight perturbation of the particle or the fluid conditions will result in an unbalanced force on the particle, accelerating it out of the central region. Hence, in practice, the system is unstable and a control system must be used to stabilise the system.

Active control is used to manipulate the flow in and out of the device and therefore modify the position of the saddle point in realtime. With reference to figure 1.1c, as the cell drifts towards Port C, moving the saddle-point a larger distance towards Port C will result in a net force that accelerates the cell back towards Port A. Conversely, moving the saddle-point towards Port A (figure 1.1d) will drive the cell back to Port C. By applying these simple principles, and extending this concept to 2D, a wide variety of behaviours can be reproduced with no modification to the system geometry.

Compared to a number of contemporary microfluidic systems, the chip geometry is very simple. The microfluidic chip consists solely of the four-way cross and associated external interconnection ports. Flow control is realised through external control devices, the four control valves. Decision support and control algorithms are provided by the imaging and computing sub-systems. While the fluidic system is very simple, the arrangement of the inlet and outlet ports allow sophisticated manipulation via control of the relative and absolute flow rates to the device.¹¹ This simplification of the mechanical system naturally implies an increase in complexity in another area. In this case, the complexity has been shifted to a sophisticated image-processing and control system running on a general-purpose computer system. Essentially, this means that the microfluidic system is *software-defined*. While the current hardware realisation uses a high performance server system, continuing advances in computer technology and the increasing power of portable, handheld devices, suggests that future software-defined microfluidic devices could target commodity, and highly portable, devices.

1.2. Physiological Importance

1.2.1. Active Mechanical Response in Cells

Many cells exhibit an active response to mechanical stimuli, a phenomenon observed in human endothelial cells over two decades ago.²³ One of the first articles to show experimentally that cells have a definitive mechanical response, Franke et al.²³ (1984), used a Couette cell to generate shear stresses on endothelial cells (found on the walls of the vasculature). After three hours of exposure to a shear stress of 0.2 Pa, actin stress fibres were found to increase in number and size. Although this type of mechanical response has an intuitive purpose — strengthening the vessel wall against increased flow induced stresses — the mechanisms involved in translating the sensed force into a biological response are complex. In the case of the endothelial cell stress response, over 100 genes are involved.²⁹ More recent research has shown that many cells exposed to mechanical forces exhibit similar properties.³² However, active response to mechanical forces is not limited to those cells with an obvious mechan-




Figure 1.1: Schematic of the proposed fluidic system, showing (a) the generalised flow directions and port definitions; (b) shape change of a deformable cell as a result of the forces imposed upon it in the central region; (c,d) active control of the saddle-point allows for manipulation of the cell position. (adapted from *Curtis, Sheard and Fouras¹¹*)

1. Introduction and Review of Existing Literature

ical role in physiology. Certain types of stem cells, such as those responsible for neuron and muscle formation, differentiate in different ways depending on the stiffness of the substrate on which they are grown.¹⁵ Forces induced on alveoli during lung development are critical for the formation of the branching tree structure typical of mammalian lungs.⁷³ These biochemical responses to applied force, and the converse — changes in exerted force as a result of biochemical processes — are collectively known as mechanotransduction.

Given the complex interactions that occur in regulating cell growth via mechanical feedback mechanisms and the prevalence of this behaviour across a wide variety of cell types, it should be unsurprising that many pathologies are a result of, or are influenced by, malfunctions in mechanotransduction. For example, cardiac muscle cells (cardiomyocytes) respond to increased cardiac demand by increasing in size as a response of increased mechanical load.³² In normal physiological conditions, cells grow in both length and thickness to accommodate this demand. However, in pathological conditions, cells can grow in only one of these directions, leading to cardiac abnormalities. What causes a cell to diverge from the physiological pathway into the pathological pathway is not yet completely understood.³²

Another cell that is strongly influenced by mechanical forces is the platelet. These cells are responsible for clotting and are the second most numerous cells in whole blood, accounting for approximately 5% of cells.⁵⁷ Once activated, the normally spheroidal platelet cells change shape and attach themselves to endothelium or other blood cells, eventually forming an aggregate. This physiological process plays an important role in wound healing and repair but can have disastrous consequences, particularly for patients with artificial heart assist devices (known as ventricular assist devices or VADs). Users of these devices are often required to take anti-clotting medication in order to prevent a dangerous build-up of aggregates, caused by a combination of VAD-induced shear and blood-material interactions.⁵⁶

The circumstances that cause platelets to activate and consequently cause blood clotting are presently not well defined. The biochemical cascade leading to platelet activation is complex, involving a number of molecules and can be easily disturbed, leading to excessive clotting or insufficient clotting. Furthermore, recent research^{52,70} shows that the shear history alone can lead to platelet activation, independent of biochemical triggers. Of particular note is the finding by Tovar-Lopez et al.⁷⁰ that the absolute shear is not sufficient to invoke activation and successful aggregation, but rather the history or transient changes to force applied to platelets are important to cause activation.

1.2.2. Cells as Transient Systems

Platelet cells are not the only cell type that exhibit transient behaviour. Indeed, the cells within complex organisms are in a constant state of flux, maintaining homeostasis (physiological equilibrium) and adapting to new and changed conditions (see the discussion on cardiomyocyte behaviour above). Timescales vary widely, with platelet activation occurring rapidly, whereas cell growth and cardiac remodelling is a significantly slower process. Force response in endothelial cells involves a number of very different timescales, with ion concentration changes apparent in less than a minute; biochemical signalling takes place over one hour; between one and six hours later, gene regulation and expression takes place, and adaptive response begins; finally, the mechanical stiffness changes become apparent at least six hours following the initial event.¹³

Cells have been modelled as rheological polymers, with viscoelastic characteristics such as relaxation times as well as non-linear frequency responses.^{29,43} A study involving embryonic fibroblast cells (NIH/3T3) found that the relaxation time for short (0 – 10 s) applications of force was in the order of 3 s, whereas force applied for an extended period (60 s) resulted in a relaxation time in the order of 5 minutes.⁷⁴ A review of a number of cell rheology experiments showed that frequency responses up to and beyond 10 kHz could be elicited from cells.²⁹




Figure 1.2: A red blood cell traversing the wall separating the splenic cord and sinus. The cell deforms to fit through an opening that is many times smaller than its usual diameter (reproduced from Figure 1, An and Mohandas.¹ Used with permission.)

A disadvantage of the rheological approach to cell mechanics is that it cannot easily differentiate between the contribution of the cell membrane and the contribution of the cell's internal structure and machinery to the overall response. An ideal cell for the study of the membrane is the red blood cell (erythrocyte), which has a very simple internal structure.

An erythrocyte can be considered to consist largely of a membrane (a lipid bilayer, potentially with attached haemoglobin molecules) which encapsulates a fluid solution that is similar in composition to blood plasma. This anucleic property is unusual. Most cells are vastly more complicated, including protein synthesis machinery, DNA, nuclei and incorporating myriad control and feedback mechanisms. Erythrocytes are also interesting and perplexing due to the large deformations that they can sustain without damage^{1,75} (see figure 1.2 for a particularly striking example), and the range of shapes which they can assume (see figure 1.3).⁶⁷ Additionally, erythrocyte membrane stiffness plays an important role in the behaviour of blood as it flows through arteries. The structure of blood flow in small capillaries is heavily influenced by the membrane stiffness — deformation of the erythrocyte gives rise to the formation of a parietal layer of blood plasma and other blood cells (e.g., platelets) around the circumference of the vessel, with erythrocytes in the central region of the vessel.⁴⁸ In addition, the elasticity of the membrane can also affect the behaviour of the cell in fluid flow. Barber et al.³ demonstrated that the membrane elasticity changes the streamline followed by a cell as it passes through a blood vessel bifurcation.

1.3. Cell Manipulation Techniques

Researchers have been attempting to quantify the properties and behaviour of biological cells for over half a century and, in the process, many techniques have been developed to subject cells to

1. Introduction and Review of Existing Literature




Figure 1.3: Some of the observed shapes of a normal erythrocyte cell (reproduced from Figure 1a, Svetina et al.⁶⁷ Used with permission.)

external manipulation. One of the earliest viable techniques to apply external forces to a cell, known as micro-pipette aspiration, was first proposed by Mitchison and Swann.⁵⁰ Since then, a number of other techniques have appeared, including direct mechanical manipulation; magnetic manipulation; atomic force microscopy (AFM); laser tweezers; laser stretchers and microfluidic techniques. These techniques are briefly reviewed below, with microfluidic techniques given a more thorough treatment in section 1.4.

1.3.1. Micro-pipette Aspiration

Micro-pipette aspiration remains one of the more popular techniques, perhaps due to its simplicity and wide applicability to a number of cell types. One of the earliest papers, Mitchison and Swann,⁵⁰ has received over 160 citations since publication, and a related paper, dealing with the application of the technique to erythrocytes, Evans,²¹ has received in excess of 200.

The technique involves placing a micro-pipette, filled with a working fluid such as water, close to the cell and using negative fluid pressure to draw the cell into the pipette. Evidently, the micro-pipette must have a diameter significantly less than that of the cell under test (Mitchison and Swann⁵⁰ manufactured a $50\ \mu\text{m}$ pipette, approximately half the size of the sea urchin eggs under consideration). Provided that a good seal is formed between the cell membrane and the micro-pipette, the cell is forced to deform into the pipette. The amount of deformation can be measured optically using standard microscopy techniques and the applied force is easily determined as a function of the suction pressure applied to the pipette.

Originally developed for the measurement of the membrane stiffness of sea urchin eggs, the micro-pipette aspiration technique has proved to be remarkably versatile, having been applied in a number of ways, including the measurement of erythrocyte stiffness²¹ and modified for measuring the forces required for cell-to-cell surface adhesion and tether forces with the use of two pipettes and an antigen-coated bead as a force sensor.⁶¹ Results presented in Evans,²¹ which estimated the shear modulus of a normal (healthy) erythrocyte to be $10\ \mu\text{Nm}^{-1}$, have been subsequently confirmed using other, more modern techniques^{12,27} which estimate the shear modulus to be in the range $9 - 18\ \mu\text{Nm}^{-1}$.

However, micro-pipette aspiration has serious limitations. Primarily, it is not a bio-mimetic process, and imposes forces on cells in ways that are non-physiological. As applied force increases, the cell is forced to deform to conform to the pipette. In the case of erythrocytes, when the radius of the deformed section matches the pipette radius, the remainder of the cell outside the pipette becomes spheroidal and no further extension is possible without causing the cell to neck or die.²¹

1.3.2. Mechanical Manipulation

A very simple method for manipulating cells involves using fine objects and micro-manipulators to directly press against the cells to determine stiffness and relocate the cells in the microscope. This is quite labour-intensive, but can be achieved with quite simple equipment. One of the earliest reported examples of using direct mechanical force to probe cell properties was reported in 1932 by Cole⁹. Small gold fibres $6\mu\text{m}$ thick were used to compress sea urchin eggs and determine cell surface stiffness. A similar technique was more recently (1991) reported in Thoumine et al.,⁶⁹ using what the author describes as “glass micro-plates” rather than gold beams. These micro-plates were about $2\mu\text{m}$ thick and were optically transparent, allowing observation and analysis from any direction in a microscope. The technique allows the use of standard commercial off the shelf micro-manipulators to manipulate cells, which significantly reduces the cost of the system and allows accurate movements. It is highly suited to adhered cells which can be made to attach to the micro-plates or a fixed surface such as a microscope slide. Blood cells, by contrast, are typically found in suspension, making this technique unsuitable for this class of cell.

However, suitable plates or beams for manipulating cells are inevitable also very flexible (the thin glass micro-plates using in Thoumine et al. had stiffnesses between 10^{-1} to $10^{-3}\mu\text{Nm}^{-1}$), making the apparatus very susceptible to mechanical vibration and external influences. Also, as the stiffness properties of the manipulators are highly variable, lengthy external calibration must be performed for each test. Finally, the low stiffness, and hence, low resonant frequency, combined with high fluid drag limits the maximum frequency of the forcing function to less than 10 Hz.

1.3.3. Magnetic Manipulation

Magnetism offers a unique pathway for non-contact “remote control” of cells.¹⁸ Electromagnets can generate large forces relative to the size of the actuator and can be controlled electronically. Unlike many other techniques, magnetic fields non-destructively penetrate the cell membrane and therefore magnetic manipulation can be used to manipulate internal cell machinery. Magnets have a long history of therapeutic and diagnostic use such as in magnetic resonance scanners. However, magnetic fields cannot necessarily be considered non-invasive, as certain cells, for example neurons, do respond to magnetic fields.²

A high throughput device was recently proposed by Spero et al.⁶⁶ for manipulating a large number of cells simultaneously using magnetic fields. Magnetic or paramagnetic beads, coated with appropriate antigens, were attached to the cell membranes, allowing remote manipulation via electromagnets embedded in a matrix of fluid cavities, each containing a single cell. As the driving magnets were designed to be placed outside the field of view of the microscope, optical access was unimpeded (however, the network of channels and magnets required to support such a system meant that cells could only be imaged in one plane). Forces of between 1pN and 1000pN could be applied with this system, where the actuators and cell under test fit into the space of four SBS wells[†] — approximately $8\text{mm} \times 8\text{mm}$.

Dobson¹⁸ reviews a number of magnetic manipulation and cytometry (twisting) techniques and applications. Recent work in magnetic nanoparticles has made it possible not just to probe the cell membrane, but also interact with receptors and ion channels on the membrane surface and exert external control on the cell’s internal behaviour. Dobson suggests that magnetic particles have a low inherent toxicity and could be used *in vivo* for non-invasive treatment strategies as well as *in vitro* investigations.

[†]SBS (Society of Biomolecular Screening) well plates are a standardised form factor that allow rapid and automated screening of a number of samples *via* an autosampling robot, or similar.

1. Introduction and Review of Existing Literature

Magnetic manipulators are limited to application of point loads on cells due to the need to bond magnetic beads to the cells in order to generate forces. In addition, magnetic beads must be specifically tailored with antigens to bond to a particular cell or portion of the cell membrane. As it would be time consuming to manually apply magnetic beads to individual cells, the cells are washed in a solution of beads. This randomises the location of the beads on the cell membrane, affecting the directions and magnitudes of force that can safely be applied to an individual cell in a random manner.

1.3.4. Laser Tweezers

Laser Tweezers, or Optical Traps are a versatile tool in cell biology and have been used for over 30 years in particle physics.⁶² For particle sizes much larger than the laser wavelength, an optical trap can be considered a geometric optics problem, and forces are imposed on the particle due to the change in momentum of light rays as a result of refraction at the interface between the medium and particle surface. While these forces are quite small, with relatively high laser powers (in the order of 200 mW–1500 mW) and small, low mass particles, the trapping forces are sufficient to retain a particle and control its movement.

Sleep et al.⁶⁵ applied this technique to the manipulation of red cell “ghosts” (erythrocytes which have been chemically treated to remove bound haemoglobin) which were then further treated via saponin lysis to make the cell membrane porous. This process was carried to simplify the later attachment of adhesive beads (coated with specific anti-gens for binding to the membrane), which were attached to opposite ends of the cells. Two laser traps were used to stretch the cells. One laser trap was used to fix the cell in place, while the other was able to be displaced by about $4\ \mu\text{m}$ in order to apply forces to the cell. Video microscopy was used to determine the force-extension profiles, with the position of the “fixed” bead used as an indicator of force applied (the effective laser trap stiffness was previously determined). However, using this method, Sleep et al.⁶⁵ calculated stiffness values that were somewhat different to those given in the existing literature. The forces applied in this experiment were much smaller than those reported in previous literature, about 25 pN, and the authors suggest this contributed to the much lower reported stiffness, as the cells have a non-linear force-extension response. Furthermore, the cells were exposed to a chemical or steroidal treatment, which may have modified their mechanical properties.

Dao, Lim and Suresh¹² used a similar technique to Sleep et al.,⁶⁵ however only one side of the cell had a bead attached to it, the other was rigidly attached to the coverslip. Using a high power laser the authors were able to exert a force of 600 pN on the cells. The authors found that the shear modulus of a typical erythrocyte varies with elongation (strain), with an initial value of $22.5\ \mu\text{Nm}^{-1}$ and reducing to $13.3\ \mu\text{Nm}^{-1}$ in the limit as the stretch ratio approached 3.

However, in common with the magnetic bead based manipulation techniques, the loads imposed on the cells are point loads. In addition, high laser powers contribute significant heating — between 1 K and 2 K temperature rise per 100 mW of laser power³⁸ — which has the potential to cause damage to the cells under test, especially for long running experiments. Furthermore, the use of higher laser powers require increasingly more complex optical filters to prevent the powerful laser light from interfering with imaging of the cell shape and position.

1.3.5. Laser Stretchers

The concept of laser tweezers can be modified to allow direct manipulation of a cell: rather than using optically induced forces to trap small beads attached to cells, the laser light can be used to directly apply force to a cell. By directing two divergent beams at opposites sides of the cell, the net force on the cell is zero, however there is a (relatively small) shear force as a result of the laser trap. This was




Figure 1.4: Hyperbolic channel used to apply a constant elongational force to erythrocytes in suspension as they traverse the channel (reproduced from Figure 2a, Lee et al.⁴¹ Used with permission.)

first proposed in Guck et al.²⁷ and shear forces of up to 630 pN were achievable with an 800 mW laser. The technique lends itself to high throughput applications as the cells do not need to be prepared with beads attached to them. Also, the force applied to the cell is distributed along the cell, rather than a point force, which better approximates physiological conditions for many cells.

However, the forces applied to the cell are dependent on the cell's refractive index, which varies between cell types and must be estimated with some accuracy. As the lower refractive index of cells is much closer to that of the working fluid (water) compared to polystyrene beads (used in laser traps), the forces that can be applied to cells are significantly lower than those that can be applied using laser tweezers.

In addition, the high power laser light directed at the cell may cause excess heating, phototoxicity (depending on the fluorophores and incident wavelength) or other damage to the cell, especially for long running experiments. Based on data presented by Kuo,³⁸ the temperature rise is expected to be in the order of ≥ 8 K. This is sufficient to change the stiffness of a red blood cell²⁰ and may result in lysis (death) of the cell.

1.4. Existing Microfluidic Systems for Manipulation

Microfluidics is a rapidly expanding field, especially in the biomedical sciences, where the potential for microchannels to act as bio-mimetic channels;⁷¹ miniaturised lab-on-chip devices and biological assay tools¹⁶ are being realised. As Vanapalli, Duits and Mugele⁷¹ point out, the length scales in microchannels can be made to be similar to those encountered *in vivo*. Microchannels are particularly suited for the study of blood diseases as the environment closely mimics a blood vessel in many respects. However, the technique is not limited to blood cells, as many cells are sensitive to mechanical forces and strain.

1.4.1. Fixed Microfluidic Channels

Lee et al.⁴¹ described a hyperbolic microfluidic channel (reproduced in figure 1.4) for imposing elongational flow on cells (red blood cells were used in the paper). The authors claim that the geometry chosen imposes a constant shear force on the cells moving through the channel in the hyperbolic section. CFD simulations in ANSYS FLUENT were used to predict the shear forces on a cell. However, the authors did not document whether these simulations were validated against similar real flows.

A suspension of cells was pumped through the channel using a syringe pump and imaged on a microscope. The cells are not tracked as they move through the field of view, hence the authors

1. Introduction and Review of Existing Literature

were careful to ensure that the image sampling rate was slow enough that no cells from one frame were present in any successive frames (an imaging frame rate of 10 frames/s or less was used). As a consequence, the deformation estimate was presumably based on a nominal initial size, rather than the initial size of the cell being imaged. As there is significant variation in size of red blood cells, this would have contributed to the relatively large experimental error reported by the authors (standard deviations were between 20 % and 35 % of the mean).

The device described is novel as it applies extensional flow to red blood cells, whereas many of the previous devices for measuring cell forces have applied pure shear. Red blood cells were found to behave differently in terms of their stiffness and deformation behaviour in extensional flow. This is an important result as most real physiological flows in vessels are not pure shear flows. Additionally, the authors found a non-linear relationship between applied shear and deformation, confirming the non-linearities reported by Dao, Lim and Suresh.¹²

Unfortunately, there are a number of problems with the method described by Lee et al.⁴¹ which limit its wider applicability. As the device is a fixed channel, there is no opportunity to modify the force profile that is applied to a cell — only the amplitude can be changed (by modifying the flow rate). As the channel is relatively short (and its length is necessarily constrained by the hyperbolic geometry), there is little opportunity to apply complex force profiles to a cell before it leaves the channel. Furthermore, the lack of particle tracking and the consequent use of single frame snapshots of cell deformation provides very little information about the transient cell behaviour or detailed investigation of variations in cell populations. By design the device is limited to a screening role, giving bulk trends for a population

A study by Tovar-Lopez et al.⁷⁰ investigated the mechanisms of platelet activation using a converging microchannel geometry. An abrupt change of depth was created in the microchannel, whereby the channel depth was reduced from 100 μm to 20 μm . Various approach angles (30°, 60° and 90°) were trialled to vary the shear history profiles applied to the platelet (see Figure 1.5). Local strain rates of around 40 000 s^{-1} were achieved with a Reynolds number of $\text{Re} = 1.2$ at the narrowest point. Strain rates were calculated using CFD and validation was performed using particle image velocimetry (PIV, a image-based flow measurement technique). However, due to imaging limitations, experimental results (PIV) were obtained at a much lower Reynolds number than the equivalent CFD. Simultaneous measurement of the velocity and platelet movement was therefore not carried out, which may have elucidated more information about the aggregation dynamics of platelets as they activate and flow downstream of the contraction.

The contracting microchannels used by Tovar-Lopez et al. lead to novel experiments into the observation of platelet activation and aggregation dynamics. In particular, they were able to observe the effect of differing strain gradients on both activation and aggregation. It was found that both the 60° and 90° inclines resulted in a large number of platelet activations. However, the aggregates formed were more stable in the 90° case. Tovar-Lopez et al. postulated that the lower flow region behind the 90° contraction allowed aggregates to form more reliably. Conversely, the aggregates in the 30° formed much further behind the contraction than in the other two cases.

Experiments on stenosed mouse mesenteric arteries, also presented in Tovar-Lopez et al., suggest that significant platelet activation occurs at strain rates around 40 000 s^{-1} but not at strain rates of 20 000 s^{-1} , indicating a step or switch-like activation behaviour. These experiments demonstrate that the microchannels used are good analogues of stenosed arteries, permitting further investigation into the behaviour of platelets in a stenosis with a controllable strain rate and improved optical access over live arteries. However, the authors only tested a single, fixed strain rate in the contraction (40 000 s^{-1}). While this guarantees that activation will occur *in vitro*, further experiments with differing geometries would provide more detailed information into the activation dynamics of platelet cells.




Figure 1.5: Contracting microchannel designs for the study of platelet activation dynamics (reproduced from Figure 1b, Tovar-Lopez et al.⁷⁰ Used with permission.)

1.4.2. Rheoscopes and Cytometers

The rheoscope and flow cytometer both measure the steady state deformation of cells in a well-defined shear flow. Schmid-Schönbein et al.⁵⁹ (figure 1.6) presented one of the first uses of a rheoscope for cell deformation measurement in 1973, and this technique has been used almost unchanged by Dobbe et al.¹⁷ in 2002. With reference to figure 1.6, the device involves two counter-rotating plates some distance apart, with the intervening cavity filled with fluid. By varying the rotation speed of the two plates, a predictable shear flow is setup, imposing a force on the suspended cells, thereby resulting in deformation. In a flow cytometer, the fluid shear is imposed on the flow by changing the flow rate through a thin tube. Laser diffraction techniques are used to measure the resultant deformation.^{34,36,64}

These techniques are capable of measuring a large number of cells efficiently and are a good physiological analogue to the deformation of cells in biological shear regimes, such as blood vessels. As the cell deformation is measured only as the cell(s) pass in front of the measuring optics (typically a camera for a rheoscope or a laser diffractometer for a flow cytometer), these devices typically only measure steady-state deformation (although a dynamic measurement is possible in a limited way — see, for example, Baskurt and Meiselman⁴).

1.4.3. Saddle Point Flows

Schroeder et al.⁶⁰ used extensional flow to observe the phenomenon of polymer extensional hysteresis. Fluid from opposing directions was pumped at a constant rate into a cross slot channel, and fluid exited at right angles. This produces pure extensional flow and exerts a force on a polymer in the centre of the fluid channel. A schematic of the channel design is shown in figure 1.7, where the channels were 7 mm wide and 150 μm deep. Some degree of control over the polymer position could be exerted by modifying the outlet impedance of one of the channels, thereby changing the relative flow rates. Extensional rates between 0.01 s^{-1} and 0.1 s^{-1} were achieved by Schroeder et al. However, at higher extensional rates (above 0.2 s^{-1}), feedback control of the particle position became difficult. However, the position of the polymer could be controlled adequately at low strain rates to allow polymers to be imaged for long periods of time (observation times of up to an hour were mentioned in the article).

The technique proposed by Schroeder et al. for observation of polymer dynamics is a promising technique for imaging of cells, as it has the potential for arbitrary extensional forces to be applied to the cell. Within the limitations of the fluid flow rates available and the system response time, any force waveform could be imposed on a cell using this technique. However, there are non-trivial stability issues with this microchannel design, especially in the case of biological cells. Polymers, once extended, form very long chains along the direction of flow out of the device, effectively anchoring the polymer, and reducing the effects of slight variations in polymer position. Conversely, biological cells are typically circular or near-circular and therefore are much more sensitive to variation in flow rates,

1. Introduction and Review of Existing Literature




Figure 1.6: One of the earliest rheoscopes for measuring cell deformation. The top and bottom plates rotate in opposite directions, creating a shear flow in the fluid between the plates. The rate of shear can be adjusted by changing the rotation speed. (reproduced from Figure 2, Schmid-Schönbein et al.⁵⁹ Used with permission)

as a small variation in the outlet flow rates could move the cell out of the saddle point region, at which point it will accelerate out of the channel due to fluid forces. These problems can be surmounted with sophisticated, active, realtime control.

This microfluidic technique is generating significant interest in the literature. Dylla-Spears, Sohn and Muller¹⁹ extended the technique for single molecule genotyping assays, as it allowed a DNA molecule to be extended and then imaged as a single long chain, allowing the location of genetic markers to be determined. Recent reviews into microfluidic manipulation of cells have mentioned the possibility of using microfluidic cross slots to exert forces on biological cells.^{71,72}

Computational Fluid Dynamics (CFD) simulations by Lee et al.⁴⁰ proposes a novel microchannel geometry that would allow a full range of flow types to be generated. Instead of the four way cross as used by Schroeder et al.,⁶⁰ eight channels are used (see Figure 1.8). Lee et al. define a flow-type parameter, ζ as

$$\zeta = \frac{|\underline{E}| - |\underline{Q}|}{|\underline{E}| + |\underline{Q}|}, \quad (1.1)$$

where $|\underline{E}|$ is the magnitude of the deformation tensor, and $|\underline{Q}|$ the magnitude of the vorticity tensor. Modelling performed by Lee et al.,⁴⁰ and supported by experimental data, demonstrate that the geometry used can generate flows with flow type parameters from $\zeta = -1$ (rotational flow) to $\zeta = 1$ (extensional flow) based on the ratios of the flow rates between the channels. The relationship between ζ and the flow rate ratio was highly non-linear but it appears that it could be collapsed into a linear relation after curve fitting.




Figure 1.7: Early cross-slot device by Schroeder et al. for observing polymer hysteresis (reproduced from Figure S1, Schroeder et al.⁶⁰ Used with permission.)




Figure 1.8: Eight way microfluidic cross, capable of producing rotational and extensional flow (reproduced from Figure 1b, Lee et al.⁴⁰ Used with permission.)

1.5. Similar Work

During the development and preparation of this thesis, a number of works have been published that are similar in concept to this work. These are Nève et al.⁵³ (2010), Gossett et al.²⁶ (2012), Cha et al.⁷ (2012), and Shenoy, Tanyeri and Schroeder⁶³ (2014). These will each be discussed in some detail, in chronological order, in the paragraphs to follow.

1.5.1. Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers (Nève et al., 2010)

Nève et al.⁵³ present a device that uses a combination of a hydrodynamic trap (of a design similar to that in presented by Schroeder et al.,⁶⁰ see figure 1.7) and an optical trap to simultaneously trap a cell and apply forces to it. Uniquely amongst the predicate works, Nève et al. measure the flow around the cell via particle image velocimetry simultaneously with trapping. This allows the authors to obtain quantitative measurements of cell deformation. Using this method, the authors present deformation and stress measurements for myoblast and chondroblast cells.

In contrast to the work presented in thesis, the mechanism of trapping cells in their work relies on a long wavelength laser (an infra-red laser of wavelength 1 064 nm), rather than fluids alone. This allows for a simpler control system — the microscope stage was manually positioned to target a particular cell in the (stationary) laser trap. Once a stable trap had been established, the flow rates were increased to apply forces to the cell. This is a manual (and undoubtedly slow) process, particularly as it requires the cell to be moving relatively slowly in order to allow the operator to position the microscope stage appropriately. In this thesis, the cells are trapped and positioned by an automatic control system, independent of operator intervention, and only fluid forces are used to achieve the trap. This avoids the potential complications of heating due to the application of high powered lasers directly to the cell. The authors note that little is known about the effects of laser trapping on cell viability, but that traps in excess of 20 s–35 s are known to result in cell death. In addition, the device presented in this work uses a much less complex optical setup, as a trapping laser is not required, resulting in a more compact and lower maintenance device.

1.5.2. Hydrodynamic stretching of Single Cells for Large Population Mechanical Phenotyping (Gossett et al., 2012)

In Gossett et al.,²⁶ a microfluidic cross-slot device is combined with a proprietary ‘inertial focussing’ device that ensures that all cells entering the cross-slot portion of the experimental setup lie exactly along the device centreline. As will be seen in chapter 3, this is an optimal configuration from a trapping perspective as all cells will — at least briefly — enter the device saddle-point. By employing extremely high flow rates (around 1 000 $\mu\text{L}/\text{min}$) and very high-speed camera acquisition rates (142 857 frames/s), significant deformations were observed. However, the individual cells were not trapped for long periods of time (without active control, the trap is unstable). Furthermore, due to the high transport velocities of the cells, the cells were only visible for a few acquisition frames — even at the high frame rates used. This is partly compounded by the small region of interest (ROI) and pixel binning settings required to achieve these frame rates — the effective imageable area was very small. As a consequence, dynamic measurements of cell deformation was not possible. Due to the small imageable area and high cell velocities, there was little to no change in cell deformation as the forces on the cell increased upon entering the centre of the device, and the cells did not remain trapped for long enough to measure the post-trap deformation. Hence, while impressive throughputs were obtained, only a static measure of cell deformation was presented.

1.5.3. Cell Stretching Measurement Utilizing Viscoelastic Particle Focusing (Cha et al., 2012)

Cha et al.⁷ presents a method for measuring the elongation of a large number of cells using a cross-slot device. To ensure that the cells travel along the centreline of the device, the cells were suspended in what the authors describe as as bio-compatible viscoelastic medium. This normalised the streamlines of the cells in the device such that 95 % of cells travelled along the device centreline. This viscoelastic medium also has the advantage that larger forces could be exerted on the cells relatively low flow rates, simplifying the imaging setup. However, as with the work presented by Gossett et al., the trap is purely passive - there is no active control to maintain the trap or select particular cells for analysis. This means that the applied force varies depending on the exact streamline followed by the cell. Additionally, no dynamic measurement was attempted — only final deformation was measured and reported.

1.5.4. Characterizing the Performance of the Hydrodynamic Trap Using a Control-Based Approach (Shenoy, Tanyeri and Schroeder, 2014)

This work is a successor to the pioneering work by Schroeder et al.,⁶⁰ and conducted by researchers in Charles Schroeder's research group. The Shenoy paper is similar to the work conducted and presented by the candidate in Curtis, Sheard and Fouras¹¹ (2011), but adds experimental validation of the control system. The authors appear not to have attempted to apply the device to biological cells, and instead have focussed on small, sub-micrometre beads. The paper presents an active trap, with good stability and long trap times. However, the extensional rates presented (a maximum extensional rate of $\dot{\epsilon} = 0.26 \text{ s}^{-1}$) do not significantly improve over the Schroeder et al.⁶⁰ paper. In Schroeder et al., DNA polymers were trapped at extensional rates of 0.1 s^{-1} with only manual control of the saddle-point position. In any case, these extensional rates are not sufficient for cell deformation measurements, and a more sophisticated control strategy, such as the one developed in this thesis, is required to scale the system to these higher extensional rates.

1.6. Conclusions

Biological cells are complex, both in respect of their biochemistry and mechanical properties. This complexity is multiplied further by the increasing body of evidence that the biochemical properties and mechanical properties may, in many cases, be intrinsically linked. Therefore, there is a clear unmet need for sophisticated and automated measurement tools that can produce quantitative measurements of cell deformation and cell responses to forces. In particular, the time-dependent response of cells to imposed force is of interest to the biological sciences, but there is a lack of tools that can measure this aspect of cell behaviour in an automated fashion. This importance of automation cannot be understated, as it is the ability to measure whole populations of cells that provide both statistical relevance and repeatability to this class of measurement, especially in biological samples, which are known to vary widely in properties and behaviour.

1.7. Bibliography

- [1] Xiuli An and Narla Mohandas. "Disorders of red cell membrane". In: *Br. J. Haematol.* 141.3 (2008), pp. 367–375. DOI: 10.1111/j.1365-2141.2008.07091.x.

Bibliography

- [2] María J. Azanza, Ana C. Calvo and Agustín del Moral. “Evidence of Synchronization of Neuronal Activity of Molluscan Brain Ganglia Induced by Alternating 50 Hz Applied Magnetic Field”. In: *Electromagn. Biol. Med.* 21.3 (1 Jan. 2002), pp. 209–220. ISSN: 1536-8378. DOI: doi:10.1081/JBC-120015992.
- [3] Jared O. Barber et al. “Simulated Two-dimensional Red Blood Cell Motion, Deformation, and Partitioning in Microvessel Bifurcations”. In: *Ann. Biomed. Eng.* 36.10 (Aug. 2008), pp. 1690–1698. ISSN: 0090-6964. DOI: 10.1007/s10439-008-9546-4.
- [4] O. K. Baskurt and H. J. Meiselman. “Determination of red blood cell shape recovery time constant in a couette system by the analysis of light reflectance and ektacytometry”. In: *Biorheology* 33.6 (Nov. 1996), pp. 489–503. ISSN: 0006-355X. DOI: 10.1016/S0006-355X(97)00037-1.
- [5] Ali Bhagat et al. “Microfluidics for cell separation”. In: *Med. Biol. Eng. Comput.* 48.10 (1 Oct. 2010), pp. 999–1014. ISSN: 0140-0118.
- [6] Darci T. Butcher, Tamara Alliston and Valerie M. Weaver. “A tense situation: forcing tumour progression”. In: *Nat. Rev. Cancer* 9.2 (Feb. 2009), pp. 108–122. ISSN: 1474-175X. DOI: 10.1038/nrc2544.
- [7] Sukgyun Cha et al. “Cell Stretching Measurement Utilizing Viscoelastic Particle Focusing”. In: *Anal. Chem.* 84.23 (4 Dec. 2012), pp. 10471–10477. ISSN: 0003-2700, 1520-6882. DOI: 10.1021/ac302763n.
- [8] Kelly C. Clause, Li J. Liu and Kimimasa Tobita. “Directed Stem Cell Differentiation: The Role of Physical Forces”. In: *Cell Commun. Adhes.* 17.2 (Apr. 2010), pp. 48–54. ISSN: 1541-9061. DOI: 10.3109/15419061.2010.492535.
- [9] Kenneth S. Cole. “Surface forces of the arbacia egg”. In: *J. Cell. Comp. Physiol.* 1.1 (1932), pp. 1–9.
- [10] James C. Culver and Mary E. Dickinson. “The Effects of Hemodynamic Force on Embryonic Development”. In: *Microcirculation* 17.3 (2010), pp. 164–178. ISSN: 1549-8719.
- [11] Michael D. Curtis, Gregory J. Sheard and Andreas Fouras. “Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems”. In: *Lab Chip* 11.14 (28 June 2011), pp. 2343–2351. ISSN: 1473-0189. DOI: 10.1039/C1LC20191C.
- [12] M. Dao, C.T. Lim and S. Suresh. “Mechanics of the human red blood cell deformed by optical tweezers”. In: *Journal of the Mechanics and Physics of Solids* 51.11-12 (Nov. 2003), pp. 2259–2280. ISSN: 00225096. DOI: 10.1016/j.jmps.2003.09.019.
- [13] Peter F. Davies. “Flow-mediated endothelial mechanotransduction”. In: *Physiol Rev* 75.3 (1 July 1995), pp. 519–560.
- [14] E. J. Diamantopoulos et al. “Impaired Erythrocyte Deformability Precedes Vascular Changes in Experimental Diabetes Mellitus”. In: *Horm. Metab. Res.* 36.3 (Mar. 2004), pp. 142–147. ISSN: 0018-5043. DOI: 10.1055/s-2004-814337.
- [15] Dennis E. Discher, Paul Janmey and Yu-li Wang. “Tissue Cells Feel and Respond to the Stiffness of Their Substrate”. In: *Science* 310.5751 (18 Nov. 2005), pp. 1139–1143. DOI: 10.1126/science.1116995.
- [16] Petra S. Dittrich and Andreas Manz. “Lab-on-a-chip: microfluidics in drug discovery”. In: *Nat Rev Drug Discov* 5.3 (Mar. 2006), pp. 210–218. ISSN: 1474-1776. DOI: 10.1038/nrd1985.

- [17] J. G. G. Dobbe et al. “Measurement of the distribution of red blood cell deformability using an automated rheoscope”. In: *Cytometry* 50.6 (15 Dec. 2002), pp. 313–325. ISSN: 1097-0320. DOI: 10.1002/cyto.10171.
- [18] Jon Dobson. “Remote control of cellular behaviour with magnetic nanoparticles”. In: *Nat Nano* 3.3 (Mar. 2008), pp. 139–143. ISSN: 1748-3387. DOI: 10.1038/nnano.2008.39.
- [19] Rebecca Dylla-Spears, Lydia L. Sohn and Susan J. Muller. “Use of Stagnation Point Flows for DNA Trapping, Manipulation, and Target Sequence Detection”. In: *The XV International Congress on Rheology: The Society of Rheology 80th Annual Meeting*. Ed. by Albert Co et al. Vol. 1027. AIP, 2008, pp. 959–961.
- [20] H. Engelhardt and E. Sackmann. “On the measurement of shear elastic moduli and viscosities of erythrocyte plasma membranes by transient deformation in high frequency electric fields.” In: *Biophys J.* 54.3 (Sept. 1988). PMC1330348, pp. 495–508.
- [21] E.A. Evans. “New Membrane Concept Applied to the Analysis of Fluid Shear- and Micropipette-Deformed Red Blood Cells”. In: *Biophys. J.* 13.9 (Sept. 1973), pp. 941–954. ISSN: 0006-3495. DOI: 10.1016/S0006-3495(73)86036-9.
- [22] M. J. Fowler. “Microvascular and Macrovascular Complications of Diabetes”. In: *Clin. Diabetes* 26.2 (1 Apr. 2008), pp. 77–82. ISSN: 0891-8929, 0891-8929. DOI: 10.2337/diaclin.26.2.77.
- [23] Ralf-Peter Franke et al. “Induction of human vascular endothelial stress fibres by fluid shear stress”. In: *Nature* 307.16 (Feb. 1984), pp. 648–649.
- [24] T. Franke et al. “Surface acoustic wave actuated cell sorting (SAWACS)”. In: *Lab Chip* 10.6 (2010), p. 789. ISSN: 1473-0197. DOI: 10.1039/b915522h.
- [25] M. J. Fulwyler. “Electronic Separation of Biological Cells by Volume”. In: *Science* 150.3698 (12 Nov. 1965), pp. 910–911. DOI: 10.1126/science.150.3698.910.
- [26] Daniel R. Gossett et al. “Hydrodynamic stretching of single cells for large population mechanical phenotyping”. In: *Proc. Natl. Acad. Sci.* 109.20 (2012), pp. 7630–7635.
- [27] Jochen Guck et al. “The optical stretcher: a novel laser tool to micromanipulate cells.” In: *Biophys. J.* 81.2 (Aug. 2001). PMC1301552, pp. 767–784.
- [28] R. M. Hochmuth and R. E. Waugh. “Erythrocyte membrane elasticity and viscosity”. In: *Annu. Rev. Physiol.* 49.1 (1987), pp. 209–219.
- [29] Brenton D. Hoffman and John C. Crocker. “Cell Mechanics: Dissecting the Physical Responses of Cells to Force”. In: *Annu. Rev. Biomed. Eng.* 11.1 (Aug. 2009), pp. 259–288. ISSN: 1523-9829. DOI: 10.1146/annurev.bioeng.10.061807.160511.
- [30] Han Wei Hou et al. “Deformability based cell margination - A simple microfluidic design for malaria-infected erythrocyte separation”. In: *Lab Chip* 10.19 (2010), p. 2605. ISSN: 1473-0197. DOI: 10.1039/c003873c.
- [31] Jay R. Hove et al. “Intracardiac fluid forces are an essential epigenetic factor for embryonic cardiogenesis”. In: *Nature* 421.6919 (9 Jan. 2003), pp. 172–177. ISSN: 0028-0836. DOI: 10.1038/nature01282.
- [32] Diana E. Jaalouk and Jan Lammerding. “Mechanotransduction gone awry”. In: *Nat. Rev. Mol. Cell Biol.* 10.1 (Jan. 2009), pp. 63–73. ISSN: 1471-0072. DOI: 10.1038/nrm2597.
- [33] R. A. Jamison et al. “X-ray Velocimetry and haemodynamic forces within a stenosed aortic model at physiological flow rates”. In: *Ann. Biomed. Eng.* (2011).

Bibliography

- [34] C D Jennings and K A Foon. “Recent advances in flow cytometry: application to the diagnosis of hematologic malignancy”. In: *Blood* 90.8 (15 Oct. 1997), pp. 2863–2892. ISSN: 0006-4971. pmid: 9376567.
- [35] Hua Jin et al. “Detection of erythrocytes influenced by aging and type 2 diabetes using atomic force microscope”. In: *Biochem. Biophys. Res. Commun.* 391.4 (22 Jan. 2010), pp. 1698–1702. ISSN: 0006-291X. DOI: 10.1016/j.bbrc.2009.12.133.
- [36] Ferenc Kiss et al. “The power of slit-flow ektacytometry measurements for testing normal and heat treated red blood cells using various viscosity media in laboratory animals”. In: *Korea-Aust. Rheol. J.* 22.1 (2010), pp. 81–86.
- [37] Sanjay Kumar and Valerie Weaver. “Mechanics, malignancy, and metastasis: The force journey of a tumor cell”. In: *Cancer Metastasis Rev.* 28.1 (1 June 2009), pp. 113–127. DOI: 10.1007/s10555-008-9173-4.
- [38] Scot G. Kuo. “A Simple Assay for Local Heating by Optical Tweezers”. In: *Laser Tweezers in Cell Biology*. Ed. by Michael P. Scheetz. Methods in cell biology v. 55. Academic Press, 1998.
- [39] C. Le Dévéhat, M. Vimeux and T. Khodabandehlou. “Blood rheology in patients with diabetes mellitus”. In: *Clin. Hemorheol. Microcirc.* 30.3 (1 Jan. 2004), pp. 297–300.
- [40] Joo Sung Lee et al. “Microfluidic four-roll mill for all flow types”. In: *Appl. Phys. Lett.* 90.7 (2007), p. 074103. ISSN: 00036951. DOI: 10.1063/1.2472528.
- [41] Sung S. Lee et al. “Extensional flow-based assessment of red blood cell deformability using hyperbolic converging microchannel”. In: *Biomed. Microdevices* 11.5 (May 2009), pp. 1021–1027. ISSN: 1387-2176. DOI: 10.1007/s10544-009-9319-3.
- [42] Chwee Teck Lim. “Single Cell Mechanics Study of the Human Disease Malaria”. In: *J. Biomed. Sci. Eng.* 1.1 (2006), pp. 82–92. DOI: 10.1299/jbse.1.82.
- [43] C.T. Lim, E.H. Zhou and S.T. Quek. “Mechanical models for living cells—a review”. In: *J. Biomech.* 39.2 (2006), pp. 195–216. ISSN: 0021-9290. DOI: 10.1016/j.jbiomech.2004.12.008.
- [44] Mingyao Liu, A. Keith Tanswell and Martin Post. “Mechanical force-induced signal transduction in lung cells”. In: *Am. J. Physiol. Lung Cell. Mol. Physiol.* 277.4 (1 Oct. 1999), pp. L667–683.
- [45] Goher Mahmud et al. “Directing cell motions on micropatterned ratchets”. In: *Nat. Phys.* 5 (14 June 2009), pp. 616–612. ISSN: 1745-2481. DOI: 10.1038/nphys1306.
- [46] Alexander G. Maier et al. “Malaria parasite proteins that remodel the host erythrocyte”. In: *Nat. Rev. Micro.* 7.5 (May 2009), pp. 341–354. ISSN: 1740-1526. DOI: 10.1038/nrmicro2110.
- [47] Shannon McCue, Sabrena Noria and B. Lowell Langille. “Shear-Induced Reorganization of Endothelial Cell Cytoskeleton and Adhesion Complexes”. In: *Trends Cardiovasc. Med.* 14 (2004), pp. 143–151. DOI: 10.1016/j.tcm.2004.02.003.
- [48] George Mchedlishvili and Nobuji Maeda. “Blood Flow Structure Related to Red Cell Flow: Determinant of Blood Fluidity in Narrow Microvessels”. In: *Jpn. J. Physiol.* 51.1 (2001), pp. 19–30.
- [49] S. Miltenyi et al. “High gradient magnetic cell separation with MACS”. In: *Cytometry, A* 11.2 (1990), pp. 231–238. ISSN: 1097-0320.
- [50] J. M. Mitchison and M. M. Swann. “The mechanical properties of the cell surface: I. the cell elastimeter”. In: *J. Exp. Biol.* 31.3 (1954), p. 443.

- [51] Matthew S. Munson et al. “Image-based feedback control for real-time sorting of microspheres in a microfluidic device”. In: *Lab Chip* 10.18 (2010), p. 2402. ISSN: 1473-0197. DOI: 10.1039/c004708b.
- [52] Warwick S. Nesbitt et al. “A shear gradient-dependent platelet aggregation mechanism drives thrombus formation”. In: *Nat Med* 15.6 (June 2009), pp. 665–673. ISSN: 1078-8956. DOI: 10.1038/nm.1955.
- [53] Nathalie Nève et al. “Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers”. In: *Cel. Mol. Bioeng.* 3.3 (Mar. 2010), pp. 213–228. ISSN: 1865-5025. DOI: 10.1007/s12195-010-0113-3.
- [54] T. T. Nguyen, J. J. Wang and T. Y. Wong. “Retinal Vascular Changes in Pre-Diabetes and Prehypertension: New findings and their research and clinical implications”. In: *Diabetes Care* 30.10 (1 Oct. 2007), pp. 2708–2715. ISSN: 0149-5992, 1935-5548. DOI: 10.2337/dc07-0732.
- [55] Thanh T. Nguyen and Tien Yin Wong. “Retinal vascular changes and diabetic retinopathy”. In: *Curr. Diab. Rep.* 9.4 (Aug. 2009), pp. 277–283. ISSN: 1539-0829. pmid: 19640340.
- [56] Buddy D. Ratner. “Blood compatibility - a perspective”. In: *J. Biomater. Sci. Polym. Ed.* 11 (Dec. 2000), 1107–1119(13).
- [57] Michael Ross, Gordon Kaye and Wojciech Pawlina. *Histology : a text and atlas*. 4th ed. Philadelphia Pa.: Lippincott Williams & Wilkins, 2003. ISBN: 978-0-683-30242-4.
- [58] Massimiliano Rossi et al. “Tapered microfluidic chip for the study of biochemical and mechanical response at subcellular level of endothelial cells to shear flow”. In: *Lab Chip* 9.10 (2009), p. 1403. ISSN: 1473-0197. DOI: 10.1039/b822270n.
- [59] H. Schmid-Schönbein et al. “A counter-rotating “rheoscope chamber” for the study of the microrheology of blood cell aggregation by microscopic observation and microphotometry”. In: *Microvasc. Res.* 6.3 (Nov. 1973), pp. 366–376. ISSN: 00262862. DOI: 10.1016/0026-2862(73)90086-1.
- [60] C. M. Schroeder et al. “Observation of Polymer Conformation Hysteresis in Extensional Flow”. In: *Science* 301.5639 (Sept. 2003), pp. 1515–1519. ISSN: 0036-8075. DOI: 10.1126/science.1086070.
- [61] Jin-Yu Shao and Robert M. Hochmuth. “Micropipette suction for measuring piconewton forces of adhesion and tether formation from neutrophil membranes”. In: *Biophys. J.* 71.5 (Nov. 1996), pp. 2892–2901. ISSN: 0006-3495. DOI: 10.1016/S0006-3495(96)79486-9.
- [62] Michael P. Sheetz, ed. *Laser Tweezers in Cell Biology*. Methods in cell biology v. 55. San Diego: Academic Press, 1998. 228 pp. ISBN: 0-12-639370-2.
- [63] Anish Shenoy, Melikhan Tanyeri and Charles M. Schroeder. “Characterizing the performance of the hydrodynamic trap using a control-based approach”. In: *Microfluid Nanofluid* (9 Oct. 2014), pp. 1–12. ISSN: 1613-4982, 1613-4990. DOI: 10.1007/s10404-014-1495-7.
- [64] Sehyun Shin et al. “Slit-flow ektacytometry: Laser diffraction in a slit rheometer”. In: *Cytometry B Clin. Cytom.* 65B.1 (May 2005), pp. 6–13. ISSN: 1552-4949, 1552-4957. DOI: 10.1002/cyto.b.20048.
- [65] John Sleep et al. “Elasticity of the Red Cell Membrane and Its Relation to Hemolytic Disorders: An Optical Tweezers Study”. In: *Biophysical Journal* 77.6 (Dec. 1999), pp. 3085–3095. ISSN: 00063495. DOI: 10.1016/S0006-3495(99)77139-0.

Bibliography

- [66] Richard Chasen Spero et al. “High throughput system for magnetic manipulation of cells, polymers, and biomaterials”. In: *Rev. Sci. Instrum.* 79.8 (2008), pp. 083707–7. DOI: 10.1063/1.2976156.
- [67] Sasa Svetina et al. “The cooperative role of membrane skeleton and bilayer in the mechanical behaviour of red blood cells”. In: *Bioelectrochemistry* 62.2 (May 2004), pp. 107–113. ISSN: 1567-5394. DOI: 10.1016/j.bioelechem.2003.08.002.
- [68] Swee Tan et al. “Microdevice for the isolation and enumeration of cancer cells from blood”. In: *Biomed. Microdevices* 11.4 (2009), pp. 883–892. DOI: 10.1007/s10544-009-9305-9.
- [69] Olivier Thoumine et al. “Microplates: a new tool for manipulation and mechanical perturbation of individual cells”. In: *J. Biochem. Biophys. Methods* 39.1-2 (25 Feb. 1999), pp. 47–62. ISSN: 0165-022X. DOI: 10.1016/S0165-022X(98)00052-9.
- [70] Francisco Javier Tovar-Lopez et al. “A microfluidics device to monitor platelet aggregation dynamics in response to strain rate micro-gradients in flowing blood”. In: *Lab Chip* 10.3 (2010), pp. 291–302.
- [71] Siva A. Vanapalli, Michel H. G. Duits and Frieder Mugele. “Microfluidics as a functional tool for cell mechanics”. In: *Biomicrofluidics* 3.1 (2009), p. 012006. ISSN: 19321058. DOI: 10.1063/1.3067820.
- [72] Guilhem Velve-Casquillas et al. “Microfluidic tools for cell biological research”. In: *Nano Today* 5.1 (Feb. 2010), pp. 28–47. ISSN: 1748-0132. DOI: 10.1016/j.nantod.2009.12.001.
- [73] Hubert R. Wirtz and Leland G. Dobbs. “The effects of mechanical forces on lung functions”. In: *Respir. Physiol.* 119.1 (1 Jan. 2000), pp. 1–17. ISSN: 0034-5687. DOI: doi:DOI:10.1016/S0034-5687(99)00092-4.
- [74] Falk Wottawah et al. “Optical Rheology of Biological Cells”. In: *Phys. Rev. Lett.* 94.9 (11 Mar. 2005), p. 098103. DOI: 10.1103/PhysRevLett.94.098103.
- [75] Young-Zoon Yoon et al. “The nonlinear mechanical response of the red blood cell”. In: *Phys. Biol.* 5.3 (Aug. 2008), p. 036007. ISSN: 1478-3975. DOI: 10.1088/1478-3975/5/3/036007.

2. System Overview and Microchannel Design

The conception and implementation of any moderately complex system is, by nature, iterative. This is especially true where the techniques and operating principles need to be developed in addition to the system itself. Experiments yield new information that then dictates a modification, or possibly major revision, to the original system.

In order to orient the reader, the system, as it exists at the end of the thesis, will be briefly described in this chapter. Some of the major decisions that have impacted the rest of the system will be described, as these decisions have strongly influenced the end result and development of the techniques involved.

2.1. Preliminary Requirements

As stated in section §1.1, the goal of the system is to capture a cell, or particle, in an extensional flow, using an active control system. This capture event can then be analysed to determine deformation as a function of time and applied force. The analysis of the capture event is not necessarily time-critical and can be performed as a post-processing operation (some potential applications of the system may require realtime analysis, though these are outside the scope of this thesis). However, the process of capturing the cell is necessarily a realtime operation and requires a control system capable of processing data from the microfluidic system and updating control parameters in realtime. Conceptually, this requires careful control design (see chapter 3); a method of measuring the flow rates (see chapter 4); a method of regulating the flow into and out of the device (see chapter 5); and a method of imaging the cell position (see chapter 6). Additionally, it is desirable to control the overall shear rate, which allows trapping to take place at a number of different shear rates, as well as the potential for shear ramps to be imposed during a trap.

Measurement of the flow rates into and out of the device was considered to be an important requirement, as it allows for quantitative measurement of the environment the cell is exposed to throughout the time that the cell is imaged by the system. As the most practical method of measuring cell deformation is *via* optical inspection (i.e. imaging), it is convenient to use this same imaging system for flow measurement. In order to manage bio-compatibility and cleanliness concerns, the flow control valves selected were pinch valves, as these do not contact the working fluid. In order to pump the fluid around the system, a system based on low pressure air was devised (figure 2.1), as this is easy to control, provides natural damping (and therefore stability of the overall pressure, due to the compressibility of the air) and results in a minimum of contact between the pump and the fluid. Finally, it was clear that all of this would need to be driven by a sufficiently powerful computer system, capable of handling the image processing tasks.

2.2. Design Parameters and Constraints

These requirements lead to a control flow similar to that shown in figure 2.2. Due to the interrelationship between saddle-point position, flow rates, and the motion of the cell, the control system consists of two nested controllers. The inner control loop measures the flow in the microchannel and adjusts the control valves to achieve the desired flow rates and, therefore, the desired saddle-point position. The outer loop is responsible for determining the position of the cell and requesting a new saddle-point position (or flow rate bias - see chapter 3) to trap a cell or maintain a trapped state.

2. System Overview and Microchannel Design




Figure 2.1: Pressure-driven flow control concept (four of these are required for a complete system).

An important consideration in the development of any realtime control system is the speed of response required to achieve stability. In the case of red blood cells, deformation was reported in flow-cytometer devices at shear stresses between 0.5 Pa and 20 Pa.^{1,2,4} This implies a shear rate in water of between 500 s^{-1} and $20,000 \text{ s}^{-1}$ (by $\tau = \mu\dot{\gamma}$, where τ is shear stress, μ the kinematic viscosity — $\approx 1 \times 10^{-4}$ in water — and $\dot{\gamma}$ the shear rate). Reported extensional rates (the analogue of shear rate for extensional flows⁵) for a similar cross-junction device that uses optical trapping, rather than fluid flow to maintain cell location, reported deformation of myoblasts at extensional rates between 5 s^{-1} and 50 s^{-1} .⁷ By integrating the streamlines for a cell entering a device near the device inlet, and assuming a device characteristic length $w = 100 \mu\text{m}$, it is possible to estimate the time the cell dwells in the central region of the channel (an area $100 \mu\text{m} \times 100 \mu\text{m}$) in the absence of active control. At an extensional rate of 500 s^{-1} , the dwell time in the centre of the chip is 16 ms, of which 5 ms is the time that the cell is transported along the inlet centreline, 6 ms is spent dwelling within $5 \mu\text{m}$ of the saddle-point, and in the remainder of the time (4 ms), the cell is accelerated out of the central area. In this linear approximation, these times will scale linearly with the inverse of extension rate, *e.g.* the total dwell time at 5 s^{-1} will be 1600 ms. Using the control systems simulator developed in chapter 3, a number of idealised control simulations were carried out for a number of frame rates. At an extensional rate of 50 s^{-1} , and using the idealised control parameters (no flow response lag, infinite resolution), the minimum frame rate for stability was found to be 100 frames/s. Frame rates below this had trapping behaviours that were stable for a brief period, before becoming unstable. At 90 frames/s, the trap could be maintained for 650 ms and at 75 frames/s the trap was only stable for 350 ms. At frame rates below 75 frames/s, the system was unstable. These figures give upper bounds for achievable extensional rates and lower bounds for camera frame rates. Practical systems, due to the influence of noise, non-linearities and other complicating factors, may require higher frame rates to achieve stability at these extensional rates.

A review of available technologies at the time of purchase (2011), revealed two camera technologies that were capable of delivering frame rates in excess of 100 frames/s in realtime[†]. These were GigE-

[†]There are a number of cameras on the market that can image significantly faster than those presented here. However, they deliver their images to onboard (internal) storage, not directly to a computer in realtime.

Vision, of which the camera with the highest maximum framerate was the AVT GE680 (200 frames/s at $640 \text{ px} \times 480 \text{ px}$); and CameraLink, where the fastest camera was the Basler A504k (500 frames/s at $1280 \text{ px} \times 1024 \text{ px}$)[‡]. Due to a number of factors, including the ease of use of AVT’s software development kit, the ability for the camera to be located some distance from the computer system, and cost, the GE680 was selected.

Despite the selection of the slower of the two cameras, the data rates involved are nevertheless high. The GE680 camera produces a compressed (bit-packed) image stream at 1 Gbs^{-1} , which, after unpacking, equates to 117 MiBs^{-1} (the A504k produces a data stream at a rate of $1,200 \text{ MiBs}^{-1}$). In order to deal with the large data rates involved, a high performance computing system is needed.

2.3. Prototype System

A summary of the major components in the prototype platform is presented in table 2.1. A computing system with a combination of a graphics-processing unit (GPU) accelerator and multi-core CPU was selected in order to provide sufficient compute capacity for analysing the images from the camera in realtime, as well as managing the other parts of the control system. A custom microscope was constructed after initial experiments were affected by vibration and insufficient stiffness of the microscope stage and camera mounts. Use of a custom design allowed for tighter control over the rigidity over the system. Control valves and the custom lighting subsystem were able to be mounted close to the microscope, while maintaining mechanical isolation from the sensitive optics. The control valves in particular are mounted on a flexible platform in order to reduce the transmission of vibration resulting from the movement of the valves to the rest of the system.

A CAD model of the fluidic, optical and mechanical subsystems is shown in figure 2.3. Tubing is connected from the fluid reservoirs, is laid inside the control valves (by means of a removable cap) and is finally connected to the microfluidic chip with a hermetic bond. This design means that the tubing, fluid reservoirs and microchannel can be treated as a sealed unit and can be installed or removed, without disconnecting the tubing, simply by removing the valve caps. The final cost of the prototype is estimated to be under \$20,000 (AUD), as detailed in table 2.1.

2.4. System Architecture

In order to support the high framerates and interfacing requirements, a distributed system, consisting of software, electrical and mechanical subsystems was required. Figure 2.4 gives a component view of the complete control system. Arrows on the diagram indicate the direction of the predominant data flow. At the centre of this distributed system is the main control application (top-left), which communicates with the various subsystems (right hand side of the diagram); receives user input from the user interface application; as well as transmitting telemetry information to the user interface and storing the telemetry on disk for later analysis.

The components in the system are briefly described in the following paragraphs (from left to right in figure 2.4).

Control Software: A multithreaded application (known as ‘vulture’) was written by the author as part of this thesis, and is responsible for controlling all other devices and running the main control

[‡]This list is of course not exhaustive, but these two cameras are both capable of saturating their respective interfaces and therefore represent the state-of-the-art for GigE and 10-tap CameraLink Full (which uses two CameraLink cables for one camera).

2. System Overview and Microchannel Design

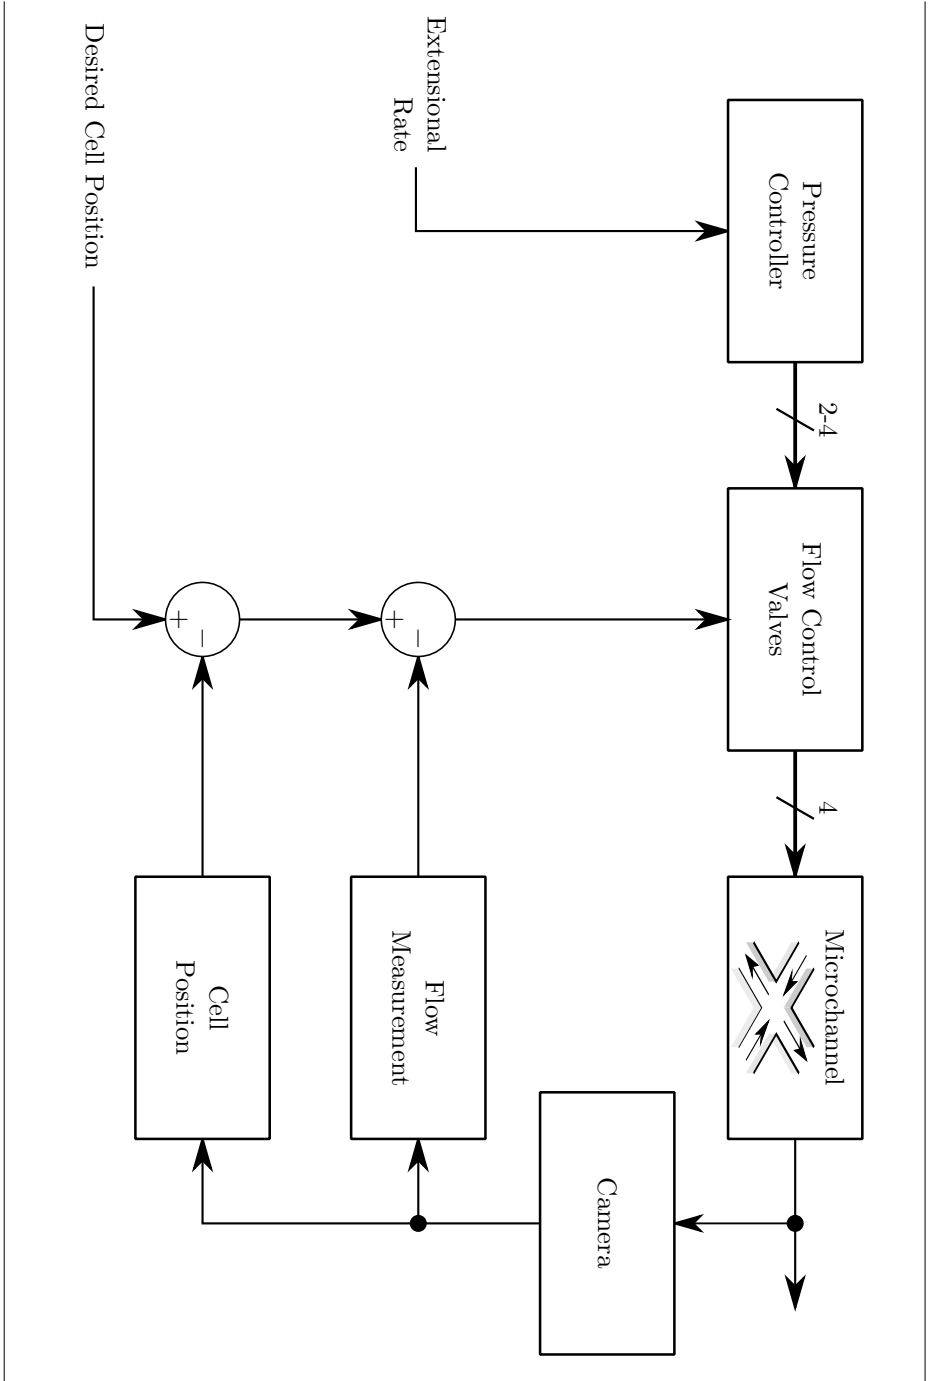


Figure 2.2: Control system view. The difference symbol, \ominus , is used here as an informal shorthand for a controller including its loop filter. Slashes on connecting lines indicate a bus of multiple signals or a group of fluid lines.

Component		Cost (AUD)
Computing System	Hewlett-Packard (HP) SL370s	\$ 6,370.00
CPU	2x Intel Xeon CPU E5620 A total of 8 cores @ 2.4 GHz	
GPU	NVIDIA Tesla M2050 2.5 GiB ECC memory; 448 CUDA cores	
System Memory	24 GiB ECC	
Camera	Allied Vision Technologies (AVT) GE680	\$ 2,661.50
Sensor	640 × 480 px, 7.4 μm square pixels	
Max. Frame Rate	205 frames/s	
Dynamic Range	12 bit A/D	
Communication	Gigabit Ethernet (GigEVision)	
Optics		
Microscope Objective	Leica ‘PL Fluotar’ 20× microscope objective (P/N: 506242)	\$ 1,277.60
Microscope	Custom assembly using Thorlabs 1” optical tubes and cage system	\$ 982.52
Illumination	Philips Luxeon LXML-PWC2 LED, LM3401-based driver, Thorlabs ACL5040 aspheric condensers	\$ 278.60
Flow Control		
Pressure Controller (EPRC)	Modified in-house (Laboratory for Dynamic Imaging) small animal ventilator.	\$ 2,175.33
Dynamic Timing Module	Addon to pressure controller by the author.	
Control Valves	Custom high-speed pinch-valves, designed by the author.	\$ 1,160.00
Miscellany		
Contract Manufacture	Machining of Mechanical Components and PCB Manufacture	\$ 4,000.00
Cabling, Cabinetry		\$ 823.03
Total		\$ 19,678.58

Table 2.1: Summary of the major components of the system and estimated final costs.

2. System Overview and Microchannel Design


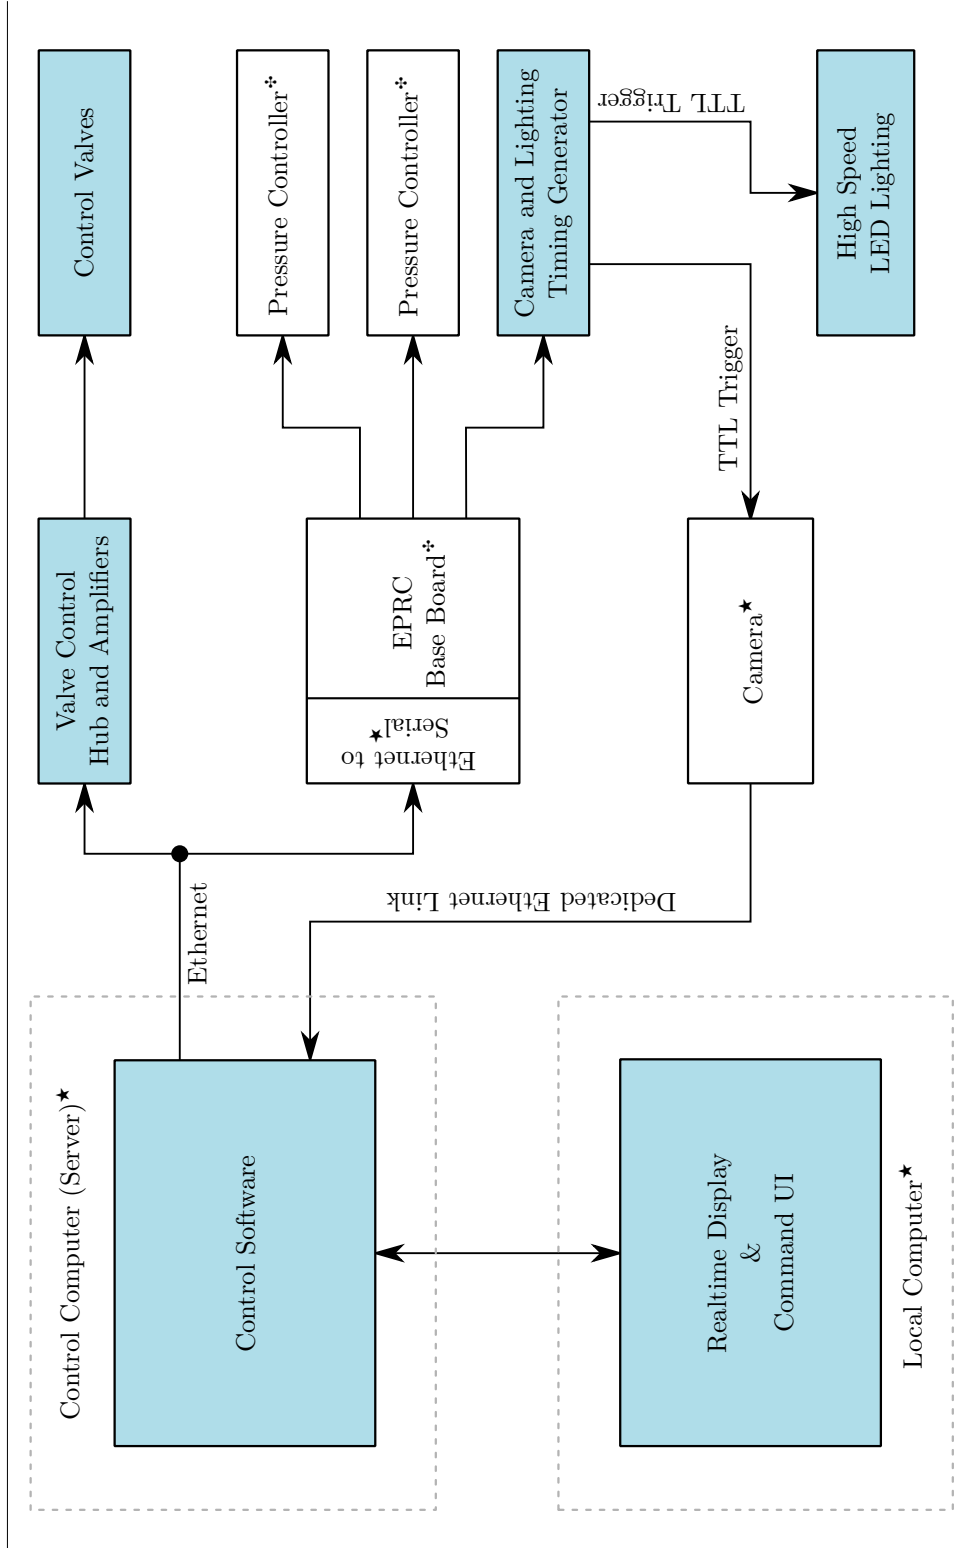


Figure 2.3: CAD render of prototype system, showing the major components.



★ Third-Party, off-the-shelf component.

♣ Component developed in collaboration with others in the laboratory as part of a small animal ventilator project (now commercialised)

Figure 2.4: Component view of the major components in the complete system, including third-party components. Shaded blocks indicate items developed as part of this thesis.

2. System Overview and Microchannel Design

loop. This is a large software application (over 12,000 lines as measured by SLOccount⁹) written from scratch in C++. Vulture is a long-running (daemon-style) process that runs in realtime (200 frame/s) and performs flow measurement (chapter 4), flow control (chapter 5), cell tracking and cell trapping control (chapter 6) functions. Additionally, the software handles a number of housekeeping functions such as logging, telemetry and the storage of images for later analysis. Many of the internal state variables and realtime measurements (known as telemetry data) are continuously logged to a file in HDF format³, which allows for detailed analysis of the system behaviour and has been used to supply data for many of the figures in this thesis.

Realtime Display and Command Interface: A graphical user interface for displaying realtime images and other data from the control system, as well as allowing a number of parameters to be modified in realtime and simple commands to be sent to the control software. As there are a large number of parameters in the system, and the system operates in realtime, the ability to change operating parameters in realtime has proven to be invaluable to the development, operation and general success of the project. This component is separated from the main control application and can run on a separate computer. This allows the control application to be long-running and independent of the user interface and allows the system to be monitored and controlled remotely. This separation also increases the stability of the system as UI crashes or hangs do not interfere with the more critical control software.

Valve Control Hub & Amplifier: An embedded ARM system (Philips LPC2378 on a development board by Olimex) running the third-party real-time kernel TNKernel.⁸ This generates dithered PWM pulses for precise generation of voltages for the system's flow control valves. The combination of high-frequency digital PWM and dithering results in a bandwidth of 2 kHz and a resolution of 16 bits ($\approx 180 \mu\text{V}$ for a 12 V bus voltage).

Expandable Pressure Regulation Controller (EPRC): An expandable system that hosts a number of daughter cards for pressure control. This was developed within the Laboratory for Dynamic Imaging and was extended for use with this project. A major addition was the timing card.

Timing Generator and Lighting Control: The timing generator is a programmable pulse generator for use with double-shutter PIV systems. It is unique in that it allows variation of the pulse timing in realtime without interruption to the pulse train. The specific use of this card for an innovate approach to PIV is developed in chapter 4.

2.5. Microchannel Design and Construction

While the microfluidic aspect of this system is comparatively simple, the success of the system is highly dependent on the propagation of small changes to flow rates to the rest of the fluid system in a short period of time (ideally, in less than the time it takes to acquire an imaging frame — i.e. $< 5 \text{ ms}$). Hence, the effective fluid damping of the microfluidic assembly is a critical parameter that requires optimisation. Small dimensional changes in the walls of the channel and tubing; air ingress / egress; and the unwanted presence of air bubbles all contribute to effective system damping, which increases the system characteristic delay and makes feedback control more difficult. A system which has no fluidic damping can be considered “rigid”, as there is a direct relationship between control input and system output, in the same way a rigid mechanical link directly transmits force (figure 2.5).




Figure 2.5: As in this cart-damper-cart analogy, a rigid system has a direct, lag-free, relationship between input and output

The most common substrate material in low-volume microfluidics manufacture is poly-dimethylsiloxane (PDMS), a soft, two-part silicone rubber that is easy to form using replica moulding. This produces a 'half-open' channel, where three-sides (nominally the top and sides) of a channel can be peeled away from a negative mould. In order to complete the channel, a second substrate material must be bonded to the remaining surface to produce a (usually, but not exclusively) rectangular channel. However, PDMS is hydrophobic and is therefore difficult to bond to other materials. Initial experiments were conducted with a PDMS-on-PDMS technique to seal the bottom of the microfluidic channel. In this method, adapted from Wu, Huang and Zare,¹⁰ an off-ratio mixture of cross-linker to polymer base is spin-coated onto a glass substrate, partially cured, and then brought into contact with the PDMS moulding and cured a second time. While careful application of this technique can result in an extremely good seal, it is time consuming and error prone, and small inconsistencies in the coating method or PDMS moulding can result in local de-lamination, which affects the effective seal. In the latter phases of the project, the manufacture of sealed channels was sub-contracted to an external company (uFluidix, Canada), which sealed the channels using a plasma-based surface-oxidisation technique.⁶

The techniques developed for sealing PDMS to the substrate have subsequently been successfully applied to the development of a unique, fluidically stiff, tubing interconnect (figure 2.6). As shown in figure 2.6a, the external tubing is press-fit into the moulded PDMS channel, then additional, off-ratio PDMS is poured into a mould form (a short length of $\frac{3}{8}$ " ID tubing) which allows the PDMS to be built up in a thin extrusion around the external tubing. The off-ratio PDMS bonds tightly to the existing material (as described above, PDMS-on-PDMS bonds can be very effective) and dramatically extends the contact area of the chip to the tubing. Additionally, if silicone tubing is used, the off-ratio PDMS also bonds to the tubing, forming a very tight seal. If polypropylene tubing is used, a short length of small-bore silicone tubing can be pressed over the PP tube to form a shim or sleeve. In this case, the PDMS bonds to the silicone, and the silicone tube shrinks around the PP tube due to the heat of curing, forming a tight press-fit. Both these methods have been found to be very effective and the resultant seal can withstand pressures well in excess of the operational pressures of the system. The high yield pressure of this interconnect is an indirect indication that the seal is rigid and adds very little to the overall system damping or air ingress/egress. Additionally, experimental data gathered from the valve calibration process (see section §5.4) shows a dramatic improvement in system response time with fluid interconnects sealed using this method, as compared to other methods examined.

2. System Overview and Microchannel Design




Figure 2.6: Microfluidic Chip to Tubing Interface, showing (a) a cross-section of one port assembly, (b) the finished chip, ready for use in the system

2.6. Conclusions

In this chapter, a detailed overview of the system has been provided. The complete system consists of a number of components, of which only a few will be discussed in detail in this thesis. The aspects of the system which are the focus of this thesis pertain to the development of the realtime control system. In developing this system, a number of significant contributions have been made to the field of active microfluidics, including: the development of simulation tools that can accurately model the behaviour of complex realtime fluid control systems; new techniques for measuring flow in microchannels in realtime in conjunction with novel feedback mechanisms for optimising the imaging process; accurate, high speed flow control systems that are non-contact, inexpensive and accurate; and tracking systems that can track the location of cells at very high framerates in realtime for the purposes of realtime control.

2.7. Bibliography

- [1] J. G. G. Dobbe et al. “Analyzing Red Blood Cell-Deformability Distributions”. In: *Blood Cells. Mol. Dis.* 28.3 (May 2002), pp. 373–384. ISSN: 10799796. DOI: 10.1006/bcmd.2002.0528.
- [2] J. G. G. Dobbe et al. “Measurement of the distribution of red blood cell deformability using an automated rheoscope”. In: *Cytometry* 50.6 (15 Dec. 2002), pp. 313–325. ISSN: 1097-0320. DOI: 10.1002/cyto.10171.
- [3] *Hierarchical Data Format*. version 5. (1997–2015). The HDF Group. URL: <http://www.hdfgroup.org/HDF5/>.
- [4] Ferenc Kiss et al. “The power of slit-flow ektacytometry measurements for testing normal and heat treated red blood cells using various viscosity media in laboratory animals”. In: *Korea-Aust. Rheol. J.* 22.1 (2010), pp. 81–86.
- [5] Sean S. Kohles et al. “Mechanical Stress Analysis of Microfluidic Environments Designed for Isolated Biological Cell Investigations”. In: *J Biomech Eng* 131.12 (Dec. 2009), p. 121006. ISSN: 0148-0731. DOI: 10.1115/1.4000121. pmid: 20524729.
- [6] M. P. MacDonald, G. C. Spalding and K. Dholakia. “Microfluidic sorting in an optical lattice”. In: *Nature* 426.6965 (27 Nov. 2003), pp. 421–424. ISSN: 0028-0836. DOI: 10.1038/nature02144.
- [7] Nathalie Nève et al. “Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers”. In: *Cel. Mol. Bioeng.* 3.3 (Mar. 2010), pp. 213–228. ISSN: 1865-5025. DOI: 10.1007/s12195-010-0113-3.
- [8] Yuri Tiomkin. *TNKernel real-time system*. version 2.5.1 (TN-Net 0.8.3). (2009). URL: <http://www.tnkernel.com/>.
- [9] David A. Wheeler. *SLOCCount*. version 2.26. (2004). URL: <http://www.dwheeler.com/sloccount/>.
- [10] Hongkai Wu, Bo Huang and Richard N. Zare. “Construction of microfluidic chips using polydimethylsiloxane for adhesive bonding”. In: *Lab Chip* 5.12 (2005), p. 1393. ISSN: 1473-0197. DOI: 10.1039/b510494g.

Part II.

Advances toward Software-Defined Microfluidics

I never did anything worth doing by accident, nor did any of my inventions come indirectly through accident, except the phonograph. No, when I have, fully decided that a result is worth getting, I go about it, and make trial after trial, until it comes.

— Thomas Edison

3. Simulation

Accurate system models form the basis of important tools that aid in both the design of the physical microchannels as well as the design of the control system and associated software. At the outset, Computational Fluid Dynamics (CFD) simulations were required to determine a number of important characteristics of the system, namely: a) the behaviour of the flow as the inlet and outlet conditions were varied, determining whether the relationship between these conditions and the flow observed in the central region was linear and/or easily describable; b) the effect of aspect ratio (that is, channel height) on the flow in the central region, and therefore the impact that aspect ratio changes would have on the cell trapping efficiency; and c) the effect of various geometries on the controllability of the system. Finally, using Computed Interpolated-Flow Hydrodynamics (CIFH), a unique technique developed as part of this thesis, the control system was simulated using pre-computed flow data from CFD simulations and a highly efficient hydrodynamics solver. This allowed for computationally-efficient, parametric studies to determine the stability of the system, impact of flow geometry, and to provide insights that informed the design of the controller. This simulation technique has been used extensively in the design, validation and simulation of the system. Further improvements to the simulator will be described in chapter 7.

3.1. CFD Simulation and Validation

A number of geometries were meshed in GAMBIT¹ and flows were simulated using an in-house spectral element computational fluid dynamics solver (VIPER). VIPER employs a nodal spectral element method for spatial discretisation and a third-order operator-splitting scheme for time integration based on backwards differentiation. The code solves the incompressible Navier–Stokes equations for a Newtonian fluid

$$\rho \frac{D\vec{u}}{Dt} = \rho \vec{g} - \nabla p + \nabla \cdot \boldsymbol{\tau}_{i,j}, \quad (3.1)$$

where ρ is the density of the fluid, \vec{g} the gravity acceleration vector, $\boldsymbol{\tau}_{i,j}$ the viscous stress tensor, and \vec{u} the velocity vector. The derivative of velocity, or substantial derivative, is defined as

$$\frac{D\vec{u}}{Dt} = \frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla) \vec{u}. \quad (3.2)$$

Two-dimensional⁷ and three-dimensional^{5,6} simulations discretize the flow domain using quadrilateral or hexahedral elements, respectively.

All simulations were based on a nominal central area of $100 \times 100 \mu\text{m}$ ($w = 100 \mu\text{m}$), at a Reynolds number of 1 (the characteristic length was chosen to be the channel width w). The selection of these nominal lengths allows easy scaling to a number of channel widths, without loss of generality. Geometries and meshes for the 2D and 3D simulations are presented in figure 3.1a-d, where h is the channel height for 3D simulations. The aspect ratio is therefore defined as the ratio between height and width, $AR = h/w$.

Inlet and outlet channels were extended to three times the characteristic length (i.e. $l = 3w = 300 \mu\text{m}$) to ensure that the flow was fully developed. The flow profile converged to the parabolic profile expected from Poiseuille flow within one diameter (w), indicating that $3w$ was adequate. Dirichlet

3. Simulation




Figure 3.1: Geometries used for the CFD simulations presented in this chapter. w defines the channel width; l inlet length ($l \geq 3w$); h channel height, hence aspect ratio $AR = h/w$. a) Geometry used for the 2D simulation, triple-headed arrows indicating the nominal flow directions; b) 3D geometry with two symmetry planes used for many of the 3D simulations; c) 3D geometry with one symmetry plane used for simulating the effects of different inlet flow rates on the 3D flow structures; and d) detailed view of the 2D mesh. In 2D, for computational efficiency, the grid is coarse in the inlet/outlet sections and is refined to a fine grid in the central region.

boundary conditions comprising of a unit average velocity were imposed on the inlet channels and a zero pressure condition was imposed on the outlet channels. The combination of the constant velocity (Dirichlet) and zero-pressure boundary conditions ensures that the solver is not over-constrained by directly specifying all inlet and outlet velocities. Finally, a no-slip (zero velocity) condition was imposed on all channel walls, and VIPER imposes a Neumann condition on the pressure to preserve third-order accuracy in time integration, following Karniadakis, Israeli and Orszag.²

The largest cell likely to be used with this device is the erythrocyte, with an average diameter of $\sim 8\mu\text{m}$. As the scale difference between the microchannel and biological cells is relatively large, the simplifying assumption has been made that the presence of a cell does not significantly affect the overall flow characteristics (the effect that the cell has on the flow in 2D is further explored in §3.1.2).

3.1.1. Mesh Independence

A coarse mesh was constructed with elements concentrated in regions where velocity gradients were anticipated. Mesh independence was verified by varying the order of the tensor product of Lagrangian polynomial shape functions within each element from 3 through 12. This is known as p -refinement² and typically results in exponential convergence of spatial errors. Grid refinement via p -refinement typically results in a more rapid convergence than increasing the density of spectral elements (h -refinement).

For the 2D solutions (figure 3.2), at $n = 8$, the percentage change in $\|L_2\|$ over the entire mesh was 0.2% (where $\|L_2\|$ is defined as the integral of the velocity magnitudes over the mesh); the percentage change in absolute velocity at a point $0.1w$ away from the centre was 0.4%. The velocity at the centre of the channel, which should be, in the ideal case, zero, was less than 3×10^{-3} , or three orders of magnitude below the base flow rate. Vorticity (figure 3.2b) around the corners of the channel is fully developed for $n \geq 8$.

In the 3D case, the trends are similar (figure 3.3). For both the $\|L_2\|$ and mean vorticity magnitude (x - y plane), there is an asymptotic relationship with increasing n , showing minimal change in both measures for $n \geq 6$. Hence for aspect ratios $AR \geq 0.62$, the solutions can be considered to be independent of the mesh density and geometry. At $AR = 0.20$, the solution required prohibitive computing resources for $n \geq 7$, and therefore the mesh independence of the solution at $n = 6$ could not be validated. A comparison with the trends at other aspect ratios suggests that this result is usable, especially for first-order measures, such as velocity.

3.1.2. Effect of a Cell on the Flow

The simulations presented so far in this chapter have not included cells in the modelling, and have rather assumed that the presence of the cell has a minimal effect on the flow characteristics. This is important to consider, as inclusion of a cell or multiple cells in the model would necessitate a coupled fluid-structure interaction model, which is more computationally expensive due to the need to solve both fluid and structural-mechanics models for each time-step. In order to verify that the presence of a cell has a minimal effect on the flow characteristics in this case, a cell-sized obstruction in the form of a rigid circular obstruction was inserted into the mesh at the centre, and the simulation was run with the same parameters as previously described. This is likely to result in a over-estimation of the change in flow conditions due to the rigidity of this obstruction when compared to a real cell, and its fixed, no-slip boundary condition. Figure 3.4a shows the flow in the central region both without the cell, and with the cell in two different locations (see caption). The general shape of the velocity is similar in the no-cell and cell at $(0,0)$ cases (Figure 3.4b), with the presence of the cell having a small effect on the local strain rate. A real cell is expected to deform subject to the local strain rate, lessening its effect on the local flow. In the final case, where the cell is offset to $(-10,0)$, there is

3. Simulation

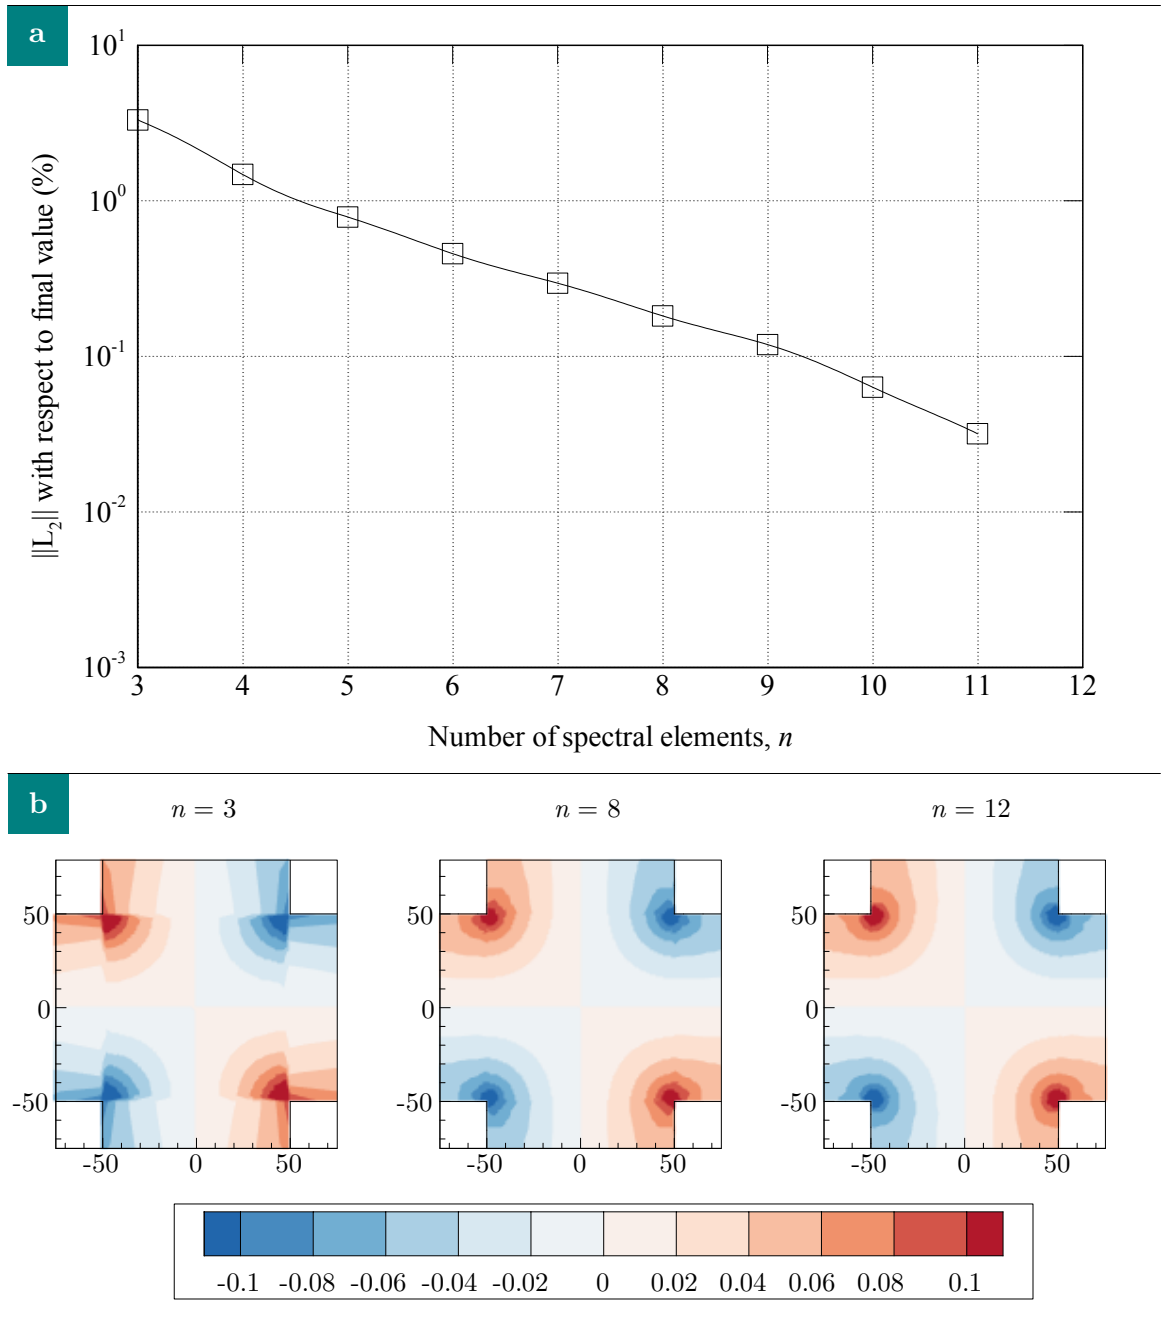


Figure 3.2: Mesh independence study for the 2D mesh, varying the number of spectral elements as an efficient way of uniformly refining the mesh. For $n \geq 8$, vorticity (b) is well-developed and the change in $\|L_2\|$ (a) is minimal.

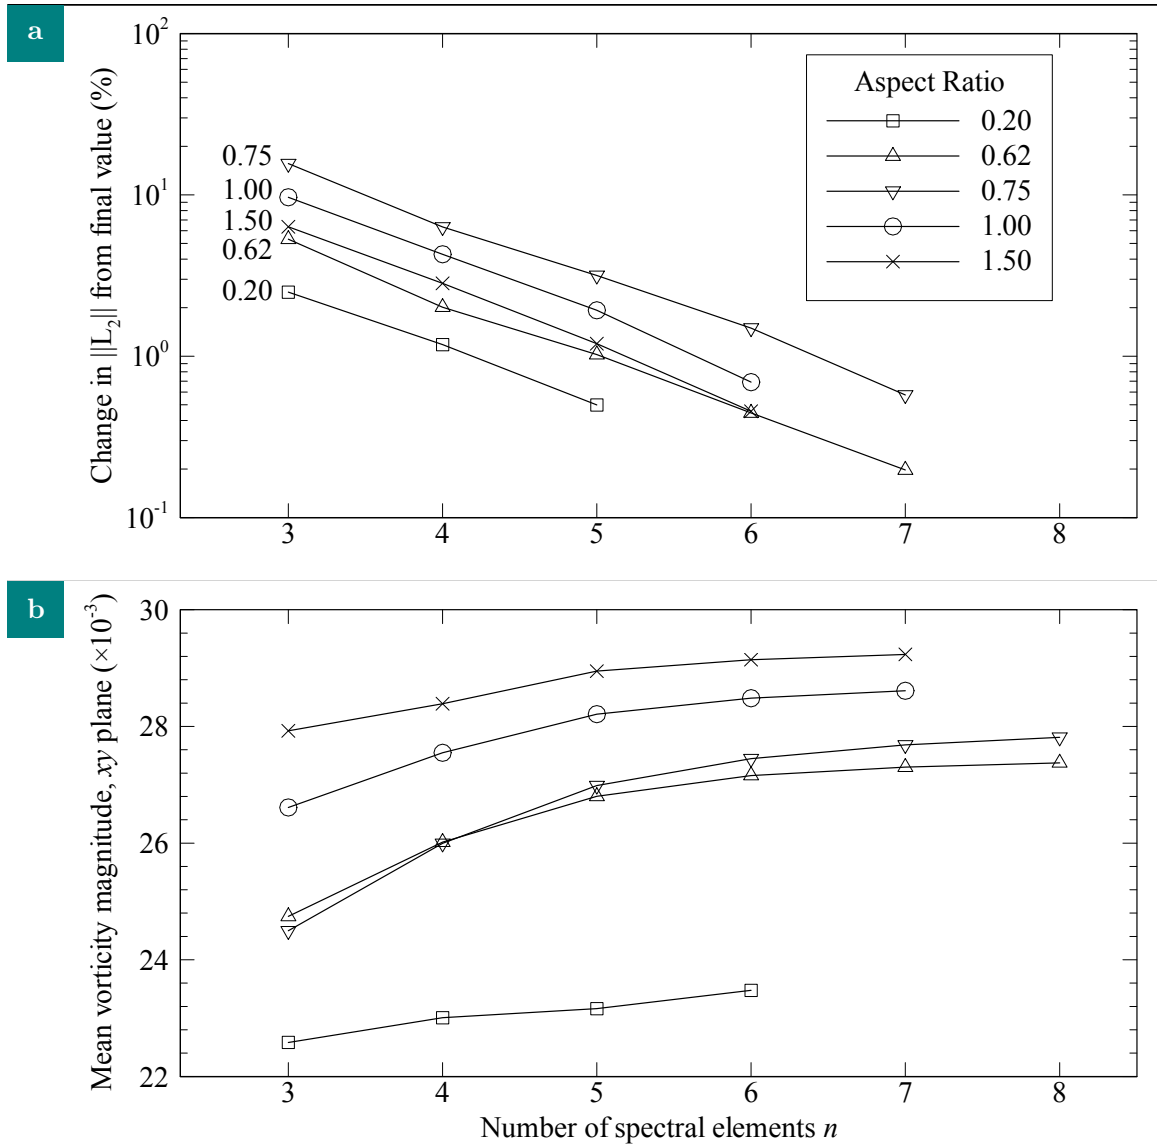


Figure 3.3: Mesh independence study for 3D meshes at a number of aspect ratios, again varying the number of spectral elements. a) Distance of $\|L_2\|$ from the final (maximum n) value; b) mean vorticity magnitude ($|\omega_{xy}|$) as a function of the number of spectral elements. In both cases there is an asymptotic trend as n increases, indicating that for $n \geq 6$, the flow is fully-developed.

3. Simulation

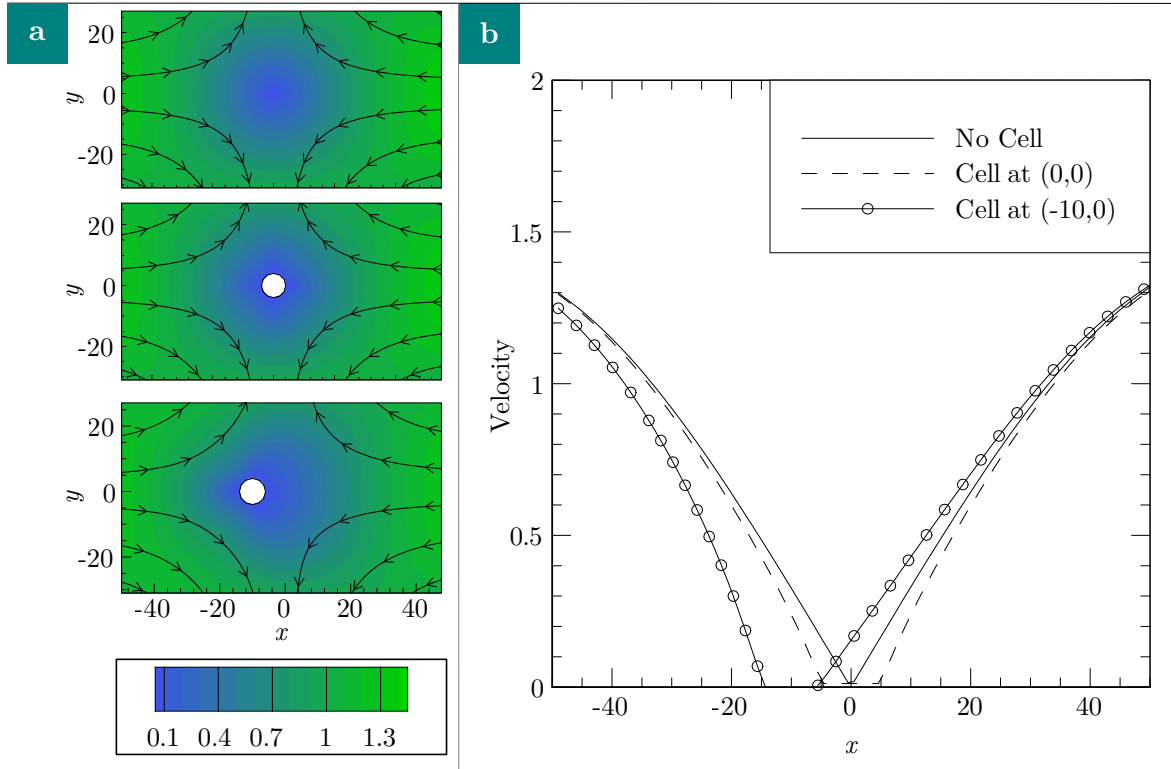


Figure 3.4: Effect of a stationary cell on the flow. a) velocity contours and streamlines for (from top): the no cell case, the cell at dead-centre (0,0), the cell offset at (-10,0); b) velocity in a slice through the centre ($y = 0$) as a function of x position.

a significant shift in the shape of the velocity profile to the left and right of the cell (Figure 3.4b). This is to be expected as the boundary conditions are such that the velocity is forced to zero close to the cell. In a real device, the cell will have a velocity much closer to the flow velocity as the high shear will cause it to accelerate toward the centre. Therefore, while there is a clear change in the characteristics of the flow with the introduction of a circular boundary in either of the two locations, the behaviour of a cell in these flows as compared to the no cell case will be very similar, allowing the model without the cell to be used as a basis for control system development. This has clear advantages in terms of solution efficiency, as including the circular boundary would require solution of a coupled fluid-structure interaction problem at each time step of the control system simulation as the flow rates are changed and the cell moves in response to the updated flow conditions.

3.1.3. Effect of Aspect Ratio

A number of 3D simulations were performed using the geometry in figure 3.1b to determine the effect of aspect ratio and the 3D behaviour of the flow. Out of plane velocity u_w was found to be negligible ($u_w/\|\vec{u}\| < 1.3 \times 10^{-3}$) for all aspect ratios, indicating that the flow is largely 2D. This has important implications for the control and trapping of cells in the device, as it means that out-of-plane motion of the cell/s as a direct result of the flow will be minimal (of course, the effects of gravity, tumbling and other dynamic behaviours may increase the out-of-plane motion beyond that predicted by the flow field).

In order to compare aspect ratios, the maximum extensional rate in the $z = h/2$ plane has been




Figure 3.5: Extensional rate in the central region as a function of aspect ratio (h/w). For aspect ratios ≥ 0.75 , the extensional rate is almost independent of aspect ratio (the variation in extensional rate from $AR = 0.75$ to 2D is less than 3.5 %).

considered. This is a measure of the maximum extensional rate that would be experienced by a cell travelling along the channel centreline and reaching the saddle-point (most likely via an active control system). These are plotted in figure 3.5, which shows an increase in extensional rate as aspect ratio is increased, with an asymptote of $19.4 s^{-1}$. Increasing extensional rate is a measure of the potential for higher deformation of cells in the trap for a given flow rate or inlet velocity, which is desirable as less working fluid is consumed in achieving a given extensional rate, and the cells are subjected to lower velocities outside the central region. However, lower aspect ratios will reduce out of focus effects in the imaging system, as the cells, if assumed to be randomly distributed in z , will be closer to the focal plane (assumed to be at $z/2$, but the same applies regardless of where the focal plane is located). In balancing these two concerns, an aspect ratio of 0.75 represents a good compromise, as the maximum extensional rate for $AR = 0.75$ is similar to the extensional rate at higher aspect ratios, and to that in the 2D simulations. This is convenient as, combined with the finding that out-of-plane velocity is small, it allows the behaviour of the system to be predominantly simulated using more computationally-efficient 2D simulations for $AR \geq 0.75$.

3.1.4. System Parameterisation

In devices based on the technology developed in this thesis, the control system affects the flow conditions in the central region by changing the flow rates into and out of the device. In order to maintain continuity and constant extensional rate, the sum of the flow rates into the device must be constant and equal to the sum of the flow rates out of the device. Hence, the system can be described by the combination of an average flow rate q_{avg} and the relative balance, or **bias**, between the two inlet and outlet flow rates, b_i and b_o , respectively. Hence the bulk flow rates for the first and second inlets, q_{i1} and q_{i2} are defined by

3. Simulation

$$q_{i1} = q_{avg}(1 + b_i), \quad (3.3)$$

$$q_{i2} = q_{avg}(1 - b_i), \quad (3.4)$$

where $-1 \leq b_i \leq 1$. The same relation is defined similarly for q_{o1} and q_{o2} , the bulk flow in first and second outlets, respectively, in terms of the outlet bias b_o . The averages of both $\{q_{i1}, q_{i2}\}$ and $\{q_{o1}, q_{o2}\}$ are q_{avg} for all b_i and b_o , fulfilling the continuity and constant extensional rate conditions.

Notational Conventions For notational convenience, the flow-rate bias can be expressed as a pseudo-vector $\mathbf{b} = (b_i, b_o)$. This notation will be used in the remainder of the text. Further, the pseudo-vector is defined as a Cartesian-style coordinate, where positive values for b_i result in a saddle-point position in the positive x direction, relative to the centre and positive values for b_o result in a saddle-point position in the positive y direction, relative to the centre. This requires that q_{i1} and q_{i2} correspond to the left-most and right-most inlet boundaries in figure 3.1a, respectively and that q_{o1} and q_{o2} correspond to the bottom-most and top-most boundaries, respectively.

3.1.5. Effect of Boundary Conditions

2D Model A parameter variation study was carried out to determine the linearity and controllability of the position of the saddle point, as a function of the input and output flow rates. A number of simulations were run based on the 2D geometry described above, with $n = 8$ (chosen as a balance between runtime and accuracy). Extending the analogy of the flow rate bias as a two-dimensional pseudo-vector, a mesh was formed of possible flow rate biases from $\mathbf{b} = (-0.9, -0.9)$ to $\mathbf{b} = (0.9, 0.9)$ with a grid resolution of 0.1. In the region from $\mathbf{b} = (-0.1, -0.1)$ to $\mathbf{b} = (0.1, 0.1)$ the grid resolution was increased to 0.02 to provide more detailed information in the critical central region.

In order to reliably identify saddle-point positions for all 900 CFD solutions, an automated method was required. It was observed that the velocity magnitude close to the saddle-point can be modelled by a cone (with the height of the cone representing velocity magnitude, see figure 3.6a), and therefore can be described by

$$\|\widehat{u_{x,y}}\| = \sqrt{\frac{(x - x_0)^2 + (y - y_0)^2}{h}} + c \quad (3.5)$$

where (x_0, y_0) is the centre of the cone and h and c are the height scale factor and minimum velocity, respectively. An optimal solution was identified with a non-linear solver using the Levenberg-Marquadt method,^{3,4} allowing an accurate, automated and largely mesh-independent evaluation of the saddle point position. The search area is iteratively determined, starting with the entire central area of the channel and reducing to a circle of radius $0.1w$, centred on the saddle point. The typical squared error is less than 1 over 2,000 mesh nodes.

As can be seen from Figure 3.6b, the relationship between b and saddle-point position s is cubic, and approximately linear in the region $-0.3 \leq b_i \leq 0.3$ (b_i shown, however the relationship for b_o is similar). Error bars indicate variation of the saddle-point position as a result of a coupling effect, where large values of b_o ($|b_o| > 0.3$) have an effect on the saddle-point position in the inlet direction s_i . This coupling effect is explored in more detail in figure 3.7: b_o is held constant in figure 3.7a-c, while b_i is increased. As b_i approaches 1.0, the saddle-point position in the outlet plane s_o , deviates from the uncoupled position at $b_i = 0$. This coupling effect, a consequence of the no-slip boundary condition on the channel wall, can be described generically in terms of the 2-norm of the flow rate

bias and the edge-related position change, or positional error

$$s_e = \|s_o(b_i, b_o) - s_o(0, b_o)\|. \quad (3.6)$$

This relationship is shown graphically in figure 3.7d. As the \mathbf{b} pseudo-vector approaches a channel corner ($\|\mathbf{b}\| \rightarrow \sqrt{2}$), the positional error increases. For values close to the centre, there is minimal coupling between the two planes, and it can be assumed that the relationships between b_i and s_i and b_o and s_o are both linear (figure 3.6b) and independent (figure 3.7d). For a positional error of less than $0.025w$ (2.5% of the channel width), $\|\mathbf{b}\|$ must be less than 0.80, a condition satisfied by $\{b_i < 0.5, b_o < 0.5\}$. This is a superset of the linear region identified earlier. In terms of saddle-point position, a flow rate bias range of $b = \pm 0.5$ corresponds to a saddle-point position $s \approx \pm 0.18w$. Effectively, this means that there is a square region in the centre of the channel of a size $0.36w$, where the system is both quasi-linear and independent (in terms of b_i and b_o).

3D Model The behaviour of the flow as a result of changing flow rate bias was also investigated in 3D. To effectively model the different inlet velocities, a geometry with a single plane of symmetry (figure 3.1c) and an aspect ratio of 1 was used. As this mesh is quite large (42,000 elements), a subset of flow rate biases were selected to observe the out-of-plane behaviour of the saddle-point as the flow rate bias is varied. To determine the location of the saddle-point, a line contour was computed using **TECPLLOT** (Tecplot, Inc, WA, USA) over the $y = 0$ plane (due to symmetry, no variation of b_o was possible, therefore the saddle-point location in y was fixed) for $\|u\| = 0$). For $|b| \leq 0.8$, the zero-velocity contours follow the expected pattern: at $b = 0$, the opposing flow rates are equal, and therefore the opposing velocities are equal for all z , hence the zero-velocity contour is a straight line

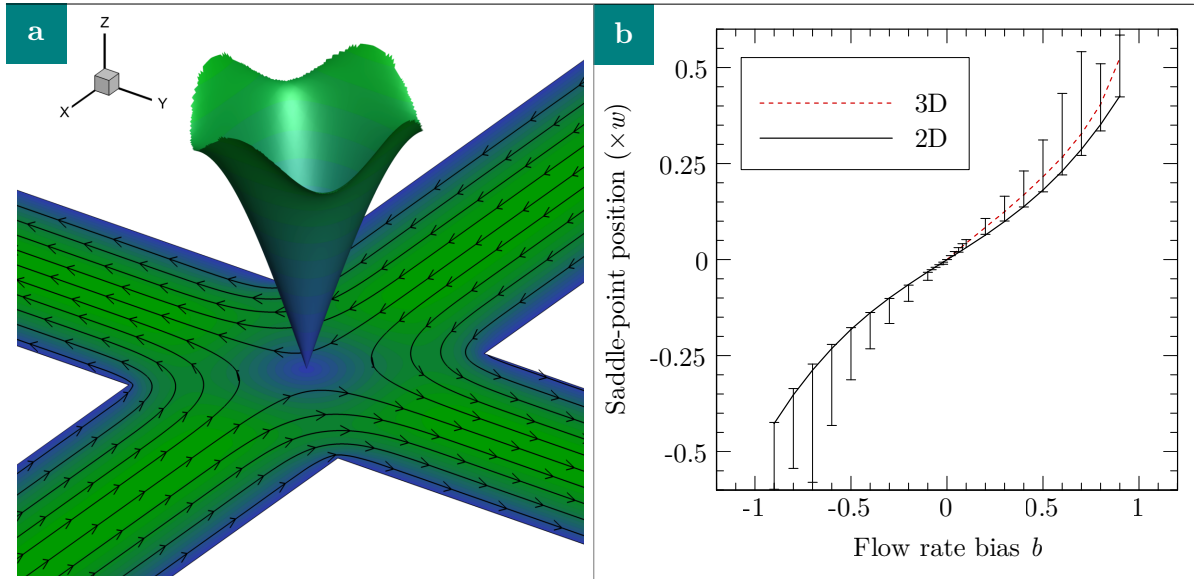


Figure 3.6: Saddle-point position measurement. a) Visualisation of velocity magnitude near the saddle-point as a 3D surface, showing the similarity to a conic section; b) Position of the saddle-point in the inlet plane as a function of the flow rate bias (b_i). Solid line indicates the saddle-point position where $b_o = 0$, error bars indicate the variation in saddle-point position as out-of-plane bias ($b_o \neq 0$) affects the position in the inlet plane. The curve is equivalent for saddle-point position in the outlet plane as a function of b_o .

3. Simulation

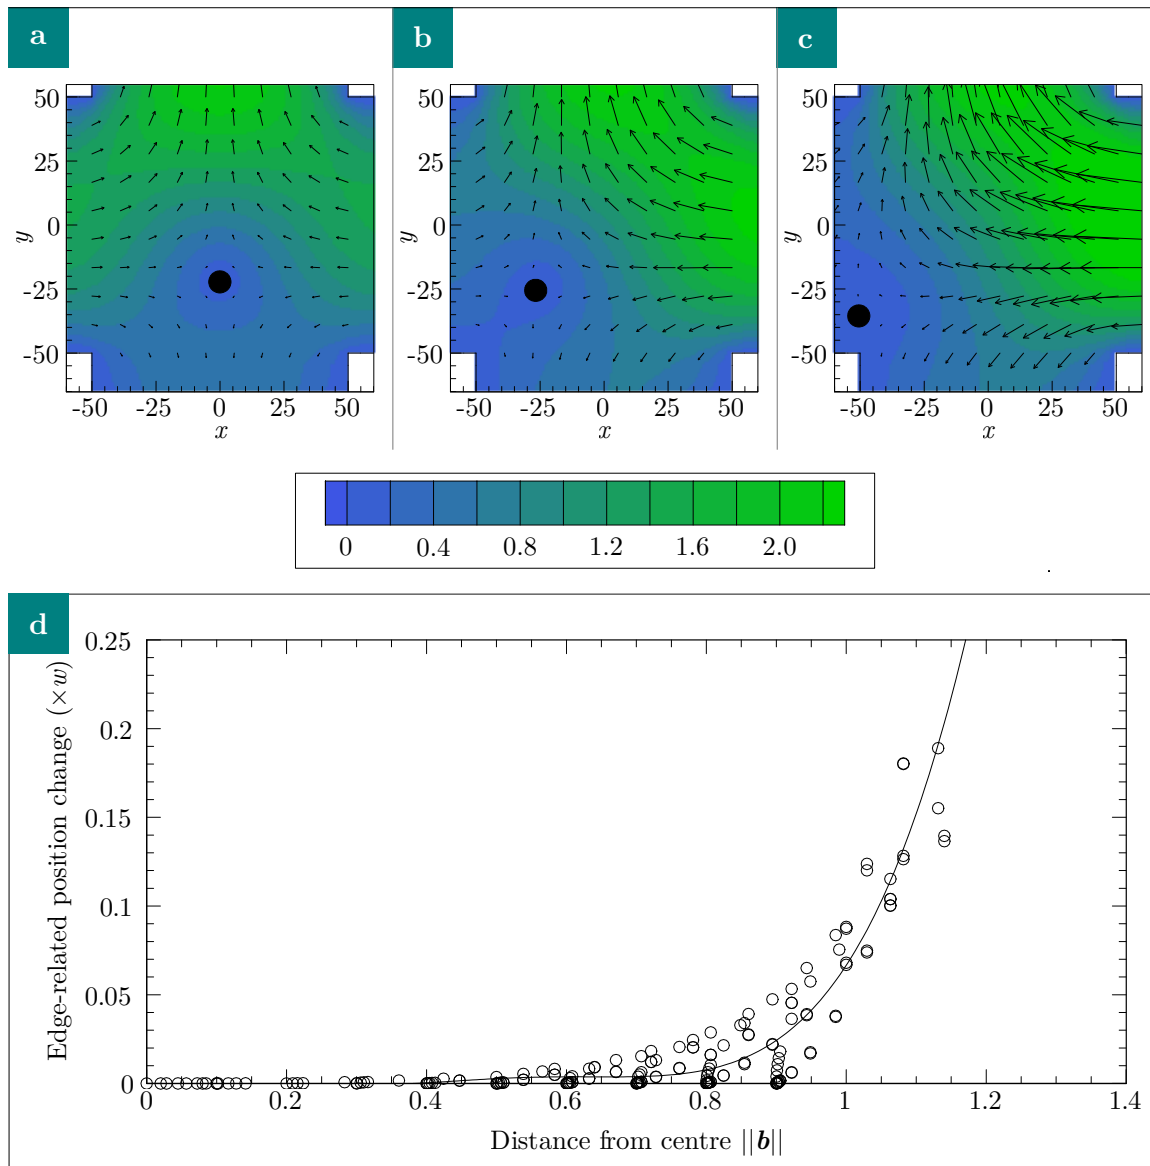


Figure 3.7: Effect of the flow rate bias on the position of the saddle-point and the flow structures. a-c) Velocity magnitude contours and vectors (overlaid) for a number of flow-rate biases. The location of the black circle indicates the position of the saddle-point. a) $\mathbf{b} = (0, -0.6)$; b) $\mathbf{b} = (-0.6, -0.6)$ c) $(-0.9, -0.6)$, showing the coupling between b_i and b_o as b_i is increased. d) Coupling in terms of the effect of the channel walls. As $\mathbf{b} \rightarrow (1, 1)$, *i.e.* the distance to the edge $\rightarrow 0$, the coupling between b_i and b_o increases, up to a positional error of $0.20w$ at $d_e = 0.28$ [$\mathbf{b} = (0.8, 0.8)$].




Figure 3.8: Out-of-plane (3D) behaviour of the device. Position of the saddle-point, as represented by zero velocity contours in the x - z plane (corresponding flow rate biases b labelled on each contour). In the central area, where $0.2h \leq z \leq 0.8h$, saddle-point position is height invariant through the channel, except in the extreme cases ($b = 0.9$, in particular). Dashed line at $x = 0.5w$ indicates the edge of the outlet channel (y direction), which causes the change in flow structure at $b = 0.9$.

at $x = 0w$. As b is increased, the zero-velocity contour shows a more parabolic shape as the velocity from one inlet is increased, hence the contour takes on the form $[1 + b]u(z) - [1 - b]u(z)$, where $u(z)$ has the form of a flat-topped parabola.

For flow rate biases $|b| \leq 0.8$ and in the region $0.2h \leq z \leq 0.8h$, the location of the saddle-point is independent of z (maximum variation of saddle-point position $< 0.1\%$). This means that the control system characteristics can be considered to be independent of z and therefore approach the 2D model, and furthermore means that identification of the z location of a cell is not required for accurate manipulation and trapping.

3. Simulation

3.2. Paper: Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems

In this paper, published in *Lab on a Chip*, the CFD simulations presented earlier in this chapter were used as an input into a novel system simulator, that allowed extensive exploration of the behaviour of the control system. By using flow data computed from CFD, along with an innovative technique to use the data from the parameter variation study presented earlier as a 2D lookup table, the response and stability characteristics of control system for trapping cells could be studied. As the model is a time-stepped non-linear model, system behaviours such as discretisation and noise could be simulated in some detail.

For direct links to the Electronic Supplementary information that accompanied this paper, please see section §B.3.

Declaration for the Paper “Feedback control system simulator ...”

Declaration by candidate

In the case of the paper “Feedback control system simulator for the control of biological cells in microfluidic cross-slots and integrated microfluidics systems,” the nature and extent of my contribution to the work was the following:

Nature of contribution	Extent of contribution (%)
Coding, Simulation, Analysis, Figures, Text	80 %

The following co-authors contributed to the work. If co-authors are students at Monash University, the extent of their contribution in percentage terms must be stated:

Name	Nature of contribution	Extent of contribution (%) for student co-authors only
Prof. Andreas Fouras	Editorial, Ideas	N/A
A/Prof. Gregory Sheard	Editorial, Ideas	N/A

The undersigned hereby certify that the above declaration correctly reflects the nature and extend of the candidate’s and co-author’s contributions to this work.

Candidate’s Signature		Date
----------------------------------	--	-------------

Main Supervisor’s Signature		Date
--	--	-------------

Lab on a Chip

Cite this: *Lab Chip*, 2011, **11**, 2343

www.rsc.org/loc

PAPER

Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems†

Michael D. Curtis,^a Gregory J. Sheard^b and Andreas Fouras^{*a}

Received 3rd March 2011, Accepted 9th May 2011

DOI: 10.1039/c1lc20191c

Control systems for lab on chip devices require careful characterisation and design for optimal performance. Traditionally, this involves either extremely computationally expensive simulations or lengthy iteration of laboratory experiments, prototype design, and manufacture. In this paper, an efficient control simulation technique, valid for typical microchannels, Computed Interpolated Flow Hydrodynamics (CIFH), is described that is over 500 times faster than conventional time integration techniques. CIFH is a hybrid approach, utilising a combination of pre-computed flows and hydrodynamic equations and allows the efficient simulation of dynamic control systems for the transport of cells through micro-fluidic devices. The speed-ups achieved by using pre-computed CFD solutions mapped to an n -dimensional control parameter space, significantly accelerate the evaluation and improvement of control strategies and chip design. Here, control strategies for a naturally unstable device geometry, the microfluidic cross-slot, have been simulated and optimal parameters have been found for proposed devices capable of trapping and sorting cells.

1 Introduction

Cells are exposed to fluid forces in a variety of physiological contexts,^{1,2} including blood flow, air movement in lung aveoli,^{3–5} and shear stress on cardiomyocytes^{6,7} and platelets.^{8,9} Microfluidic devices, with similar length scales to blood vessels and other biological fluid channels, enable the construction of a customised biomimetic environment in the laboratory. By customising the channel geometry, and hence the fluidic environment of cells, forces can be applied to cells that mimic those encountered *in vivo*.^{10–12} In addition, microchannels can also include features such as constrictions,¹³ precisely tuned outlet sizes,^{14,15} junctions¹⁶ and pillars¹⁷ to direct and sort cells based on their size and behaviour in flow. High throughputs can be obtained by making use of these geometric structures, allowing large numbers of cells to be processed rapidly. However, the force profiles and sorting decisions are static – encoded in the geometry at manufacturing time.

Cells are dynamic, exhibiting both rheological characteristics¹⁸ and active biochemical responses to applied force.¹⁹ The response to externally-applied forces by cells is complex and non-linear, due to interactions between membrane structure, internal

organelles in the cell and viscous behaviour of the internal fluid. Pathological conditions, such as diabetes, modify the biochemical properties of red blood cells, changing the force response.^{20–22} The malaria parasite modifies both the membrane chemistry and the internal structure of red blood cells, causing a dramatic increase in red cell stiffness.^{23,24} Environmental forces also provoke active responses, such as stem cells differentiating;²⁵ blood vessel endothelial cells changing stiffness;²⁶ and healthy cells becoming malignant.^{27,28}

Investigating dynamic behaviour, such as the response to varying forces applied to a cell in realtime, requires active control. Complementary technologies, such as electric fields;^{29,30} magnetism;^{31,32} optical forces;³³ and surface acoustic waves,³⁴ are often used to impose forces on suspended cells in microchannels. In contrast to these technologies, the use of the fluids to directly manipulate cells offers a preparation-free, biomimetic approach with a low risk of cell damage. In addition, the use of direct fluid manipulation allows multiple, independent, fluid control systems to be densely integrated on a single chip.³⁵ However, the interactions between actuators; the control algorithms; the cells under test and the dynamics of the fluid must be characterised in order to develop an efficient controller. Testing and characterisation of control algorithms experimentally is both time consuming and expensive. A large number of iterations is typically required to develop and tune an optimal controller. Not only is this time consuming experimentally, but may also be expensive in terms of reagent and sample costs. This ultimately limits the scope for optimisation and the complexity of systems that can be feasibly attempted. Flow geometry may require refinement

^aDivision of Biological Engineering, Monash University, Melbourne, Australia. E-mail: Andreas.Fouras@monash.edu

^bFluids Laboratory for Aeronautical and Industrial Research (FLAIR), Department of Mechanical and Aerospace Engineering, Monash University, Melbourne, Australia

† Electronic supplementary information (ESI) available: Experimental validation of CFD; movies of cells in a microfluidic device undergoing manipulation due to the control system. See DOI: 10.1039/c1lc20191c

and therefore the manufacture of new prototypes, with the attendant lead times and costs. Therefore, simulation of dynamic microfluidic systems is critical to the development of this field.

The greatest limitation of existing simulation techniques is that the computation time required is prohibitively long. Consequently, control simulations using these techniques are impractical and limit experimentation with control system design; control tuning parameters; and geometry. A common general approach for simulating problems of this nature formulates the system as a set of coupled systems that must be solved iteratively.³⁶ This is known as the *fluid-structure interaction* problem, where the structure in this case is a cell. Using this solution method, the fluid system is initially solved with an initial approximate cell location. The forces resultant on the cell are then calculated and cell deformation is estimated using a finite-element solid mechanics model of the cell. The fluid domain is then re-meshed based on the new cell deformation and a new solution to the fluid system is computed. This process repeats until the error between successive solutions has been reduced significantly.

After the error has reduced, net acceleration of the cell can be computed from the resultant forces and integrated to find an updated cell position. Using a standard numerical quadrature algorithm, such as a fourth-order Runge–Kutta algorithm, the iterative process described earlier must be repeated at least three further times to successfully integrate the acceleration to find an updated position. This process then repeats for every time step. Even for very efficient computational fluid dynamics (CFD) solvers and small solid mechanics models, each iteration step may take many hours on modern computer systems.

However, when the solid body can be described as a simple geometric shape and the flow can be considered to be minimally perturbed by the presence and motion of the solid body, a simplified approach based on the hydrodynamic drag equations can be used. Many cells can be approximated as spheres or ellipsoids and hence lend themselves well to this approach. Recently, this approach was used to model the motion of bacterial cells in a passive microfluidic trap.^{37,38} A background flow was computed for the microfluidic trap geometry and the bacterial motion was computed using hydrodynamics. However, in that study only a steady state flow with static boundary conditions was considered. In this paper, a hybrid approach, using a combination of precomputed flows and hydrodynamic equations is developed, allowing a dynamic control system for cells to be efficiently simulated.

In this study, it is demonstrated that the simulation technique can be used to characterise and design a control system to capture, trap and manipulate cells in a microfluidic cross slot (Fig. 1). First used to study polymer extensional hysteresis,³⁹ the microfluidic cross slot has recently been extended to study polymer beads and cells.⁴⁰ Due to the geometry and fluid boundary conditions in the device, a saddle point is formed in the centre of the channel, resulting in purely extensional flow. This exerts a force on the cell at the centre of the fluid channel, stretching it (Fig. 1a). Once the cell has been captured by the device, the cell can be redirected to one of the two outlets (Fig. 1b). Control over the cell position can be indirectly achieved by modifying the relative flow rates (Fig. 1c).

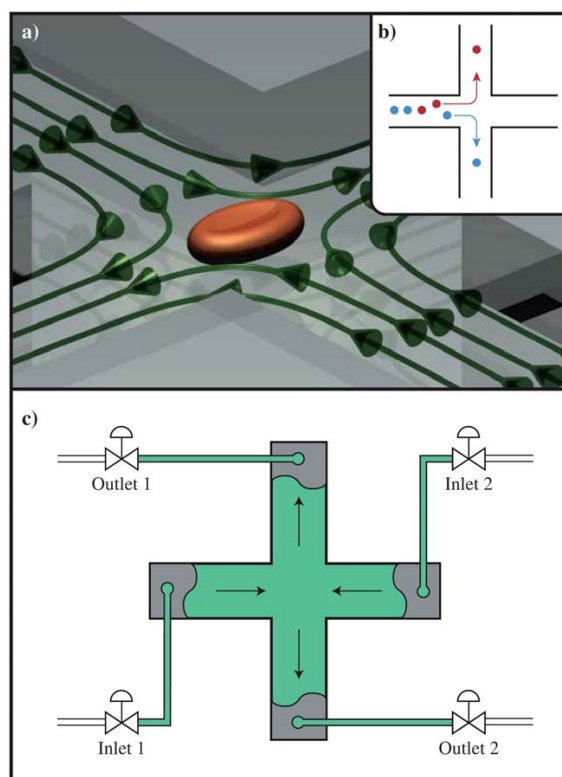


Fig. 1 Stretching (a) and sorting (b) cells using microfluidics. A red blood cell is shown (not to scale) undergoing stretching due to extensional flow. Fluid velocity, as indicated by the streamlines, results in shear stress on the cell; causing it to stretch. Sorting (b) can also be achieved using this device by directing cells toward one of the two outlets. Control of the shear, saddle point position and therefore position of the cell in the microfluidic cross can be effected by changing the relative flow rates of the inlets and outlets (c). For example, a higher relative flow rate through Inlet 1 will result in the saddle point moving to right, relative to the centre of the channel. Shear can be adjusted by modifying the overall flow rate.

2 Methodology

The method proposed in this paper, Computed Interpolated Flow Hydrodynamics (CIFH), significantly reduces the computation time compared to a conventional fluid-structure interaction approach, such that each step can be computed in less than a second. Three primary strategies are employed to achieve this: the motion of the cell is decoupled from the fluid, avoiding the expensive iteration between the fluid and solid mechanics models; the time evolution of the fluid system is treated as a Linear Time Invariant system, allowing flow solutions to be computed ahead of time; and flow solutions are interpolated, allowing the flow parameters to be varied continuously within the solution space.

Fluid-cell coupling

For small Reynolds numbers in the regimes of creeping and Stokes flow ($Re \lesssim 1$), the expensive computation of

fluid-structure interactions can be reduced to a closed-form equation known as the Basset-Boussinesq-Oseen equation (BBO).^{41,42} These equations use the hydrodynamic drag of the cell, along with its mass and viscosity to compute the acceleration of the cell in an ambient fluid. Drag correction factors for spherical and aspherical shapes are well-known, allowing great flexibility across a range of cell types without requiring direct computation of the solid mechanics.

Time independence of the fluid system

In the regime of Stokes flow, the inertia terms of the Navier–Stokes equations are negligible. Therefore, the flow solution can be completely described by the geometry and boundary conditions. As inertia is negligible, there is nothing to cause the fluid system to diverge from the time independent steady-state solution for a given input. This is analogous to a Linear Time-Invariant (LTI) system in signal processing theory – for a given input stimulus, the output is always identical, regardless of the time history of the system. As the fluid system is decoupled from the cell solution, the varying location of the cell over time does not change the time invariant properties of the flow. This allows flow solutions to be computed independently of the control simulation.

Interpolation of flow solutions

Finally, the LTI properties of the fluid system are utilised to reduce the number of flow solutions required by the simulation. As any combination of flow parameters (*e.g.*, boundary conditions) could be demanded by the fluid control system, it is necessary to efficiently find the solution to the flow for an arbitrary combination of flow parameters. This is achieved by discretising the flow parameters onto a grid of available solutions and interpolating between these available solutions, yielding an efficient method for generating the required flow solution. As for finite element methods, this method is accurate provided the discretisation is chosen to minimise abrupt changes in the solutions.

The overall process is summarised in Fig. 2. Firstly, a controller (Fig. 2a), which may be any multiple-input multiple-output (MIMO) discrete controller, demands a change to the flow parameters (in this example, flow rates) based on the cell location. Simulations of flow actuator noise and resolution are mixed into the flow parameters at this point. Secondly, the flow parameters are mapped to a grid of available flow solutions (Fig. 2b and inset). Next, the flow solutions closest to the demanded flow parameters are identified (Fig. 2c). These solutions are then interpolated to generate the flow solution for the flow parameters demanded by the controller (Fig. 2d). Finally, the equations describing the cell motion, the BBO equations are integrated to provide the location of the cell for the next time step (Fig. 2e). These processes are discussed in detail in the following sub-sections.

2.1 Flow field generation

Typically, in order to model the response to the changing fluid boundary conditions, a computational fluid dynamics (CFD) solver would be used to predict the flow field at every flow rate

requested by the controller. This is computationally expensive and time consuming, especially when simulating a large number of control scenarios. Even for a two-parameter system, the number of CFD runs required would be large in order to fully cover the space. Instead, in the present approach, the space is sampled and an interpolation approach is used to evaluate the flow field between these sample locations.

Flow solutions are obtained over an n -dimensional grid of the input parameters. The number of input parameters will be largely dictated by the complexity of the flow geometry and the degrees of freedom of the system. These flow solutions can be computed by any means, including a general-purpose CFD solver such as VIPER or FLUENT. As these flow solutions are computed once and then stored for future simulations, subsequent CIFH simulations are very efficient.

Each flow solution is treated as a node on a parameter space grid (see Fig. 2b for a two-dimensional example). In a process analogous to that used by the finite element method, the parameter nodes are meshed to form grid elements (in 2D, these are quadrilaterals). The meshed domain therefore encloses the space of input parameters for which there are solutions available. Any set of valid input parameters will be fully enclosed by exactly one parameter grid element[‡]. The parameter grid can then be efficiently and rapidly searched by an algorithm such as an Alternating Digital Tree.⁴³

During execution of the control simulation, flow parameters demanded by the controller are mapped onto the parameter mesh (see Fig. 2c). The element that encloses these parameters is then identified and linear interpolation is performed between the CFD solutions that comprise the element nodes (Fig. 2d). This is similar to the linear interpolation scheme used in the finite-element method, where the interpolation typically takes place between individual scalars rather than complete solutions.

Finally, to reduce the computational overhead associated with interpolating a number of large meshes, the interpolations are evaluated using a *lazy evaluation* scheme.⁴⁴ Instead of computing the interpolated values for each node in the resulting CFD solution, these values are only computed when needed by the simulation code and then cached in memory until the flow parameters are updated again.

2.2 Cell dynamics

A fourth-order Runge–Kutta integration scheme is used to integrate the cell velocity and hence simulate the evolution of its position over time (Fig. 2e). Provided the dimensions of the microchannel are large relative to the cells under consideration, cell dynamics can be decoupled from the fluid dynamics of the system. Therefore, the evolution of the cell position is independent of the fluid dynamics solutions and changes to the cell position, orientation or shape do not require new flow solutions to be obtained. However, cells will display inertial effects due to finite size and mass. These effects are described by integrating the

[‡] At element intersections, more than one element will enclose the set of input parameters. In this case one is arbitrarily chosen, as the nodes over which the interpolation takes place are identical in the plane of intersection.

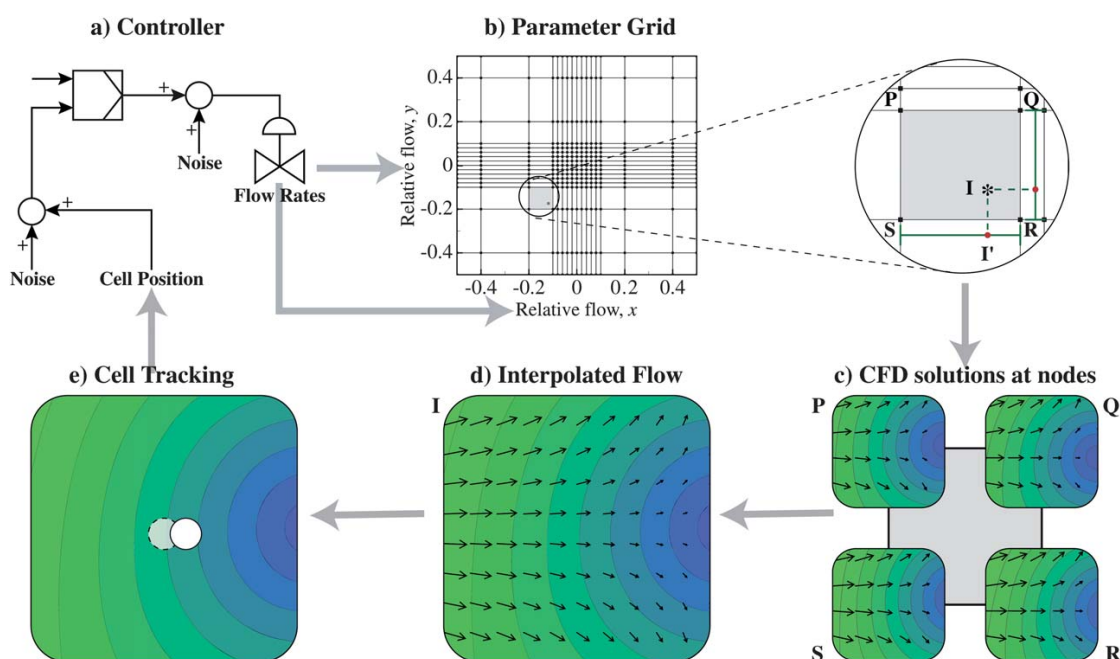


Fig. 2 CIFH control simulation process. A controller uses the cell position as an input to determine a new set of boundary conditions, in this case relative flow rates (a). Available solutions for these boundary conditions are mapped to a parameter grid, in this case with two dimensional, where each dimension relates to a flow rate parameter (b). The point marked with an asterisk (*) indicates the demanded combination of flow rate parameters (shown in the exploded view). The four nodes on the parameter grid forming the element boundary are identified (P,Q,R,S), where each node represents a complete CFD solution (b,c). The four solutions are then interpolated to find the resultant bilinear-interpolated flow (d). Finally, hydrodynamic equations are used to compute the motion of the cell and hence the updated cell position (e).

common hydrodynamic equation for particle drag in a flow, the Basset-Boussinesq-Oseen equation,⁴²

$$\begin{aligned}
 m_c \frac{d\mathbf{v}}{dt} = & -\frac{1}{2} m_f \frac{d}{dt} \left(\mathbf{v} - \mathbf{u} - \frac{\alpha^2}{10} \nabla^2 \mathbf{u} \right) \\
 & - 6\pi\alpha\mu_f \left(\mathbf{v} - \mathbf{u} - \frac{\alpha^2}{6} \nabla^2 \mathbf{u} \right) \\
 & - \frac{6\pi\alpha^2\mu_f}{\sqrt{\pi}\mu_f} \left(\int_0^t \frac{d}{dt} \left(\mathbf{v} - \mathbf{u} - \frac{\alpha^2}{6} \nabla^2 \mathbf{u} \right) \frac{d\tau}{\sqrt{t-\tau}} \right) \\
 & + (m_c - m_f) \mathbf{g} \\
 & + m_f \frac{D\mathbf{u}}{Dt},
 \end{aligned} \quad (1)$$

where m_c and m_f are the cell and fluid mass, respectively; μ_f is the fluid viscosity; α is the cell radius; \mathbf{g} is the acceleration due to gravity; \mathbf{v} is the cell velocity vector; and \mathbf{u} is the fluid velocity vector. The equation is reduced in complexity by taking into the account the following factors: the Faxen terms ($\nabla^2 \mathbf{u}$) can be neglected when $\alpha l \ll 1$, where l is the channel length scale, a condition which holds except for very small channels or large cells;⁴² and the integral term (known as the history term) approaches zero for Reynolds numbers less than 1.^{42,45} Additionally, the primary aim in this work is to model scenarios where the cell is already present in the flow and does not start from rest, hence reducing the importance of the history term in this case.

For this analysis, buoyancy is ignored as only the in-plane velocity of the cell within the channel is of interest.

Finally, the sphere drag used for the drag term in (1), *i.e.*

$$F = -6\pi\alpha\mu_f(\mathbf{v} - \mathbf{u}) \quad (2)$$

is substituted for the more general inertial drag equation⁴⁶

$$C_D = \frac{2F}{\rho_f \pi \alpha^2 \|\mathbf{v} - \mathbf{u}\|^2}. \quad (3)$$

This formulation is equivalent to that in (2), but lends itself more readily to describing aspherical shapes such as ellipsoids.⁴⁶ In the examples described in this paper, we employ $C_D = 24/Re_c$ for drag on a sphere at low Reynolds numbers.

Hence, expressing the masses in terms of areal density,

$$\begin{aligned}
 \frac{d\mathbf{v}}{dt} = & -\frac{1}{2} \frac{\rho_f}{\rho_s} \frac{d}{dt} (\mathbf{v} - \mathbf{u}) \\
 & - \frac{1}{2} \frac{\rho_f}{\rho_s} C_D \|\mathbf{v} - \mathbf{u}\| (\mathbf{v} - \mathbf{u}) \\
 & + \frac{\rho_f}{\rho_s} \frac{D\mathbf{u}}{Dt}.
 \end{aligned} \quad (4)$$

Where the cell is known to be undergoing deformation due to fluid forces, the drag coefficient C_D can be additionally varied to simulate the change in hydrodynamic behaviour due to the shape change. The drag correction factors for ellipsoids are well

known,^{42,47} allowing motion of aspherical cells such as red blood cells to be simulated. Additional behaviours are easily integrated into the simulation. For example, red blood cell tumbling,^{48,49} could be simulated by estimating the rotation rate from the capillary number⁴⁹ and applying the appropriate combination of the drag factors for axial- and normal-axis flow around an ellipsoid.⁴²

The BBO equations are also valid for simulations containing a number of cells by approximating the cells as an ensemble of single cells. For practical concentrations of cells, limited by the ability of the imaging system to reliably identify individual cells, this approximation is accurate. Any cell moving through a fluid produces a disturbance field in the fluid. For a single cell, the effect of this disturbance field is compensated for within the BBO equations. However, the effect of disturbance field from neighbouring cells cannot be compensated for. The disturbance field affects the base flow rate by just 10% at a distance of 5 diameters from the cell,⁴⁶ where the cell velocity is similar to the fluid velocity and the fluid and cell densities and viscosities are similar. Hence, the motion of cells for linear concentrations less than or equal to one cell per 5 diameters can be approximated by an ensemble of single cells moving through a fluid with minimal loss of accuracy.

For small cells, it may be important to model particle diffusion. This is a straightforward addition to the overall method, using the random-walk method of Ghoniem,⁵⁰ where particle positions are modified at each time step based on a Gaussian random distribution that is scaled based on the Schmidt number of the particles in the carrier fluid.

2.3 Control system

A simple control algorithm, based on PID (proportional integral derivative) control is used in this paper to control cell position (Fig. 2a). The control algorithm can be defined in terms of a pair of independent one-dimensional PID algorithms (one for each axis of control $-x$ and y), as

$$e = x_s - x_c \quad (5)$$

$$f = K(e + K_D \dot{e} + K_I \int e dt), \quad (6)$$

where e is the error, x_s and x_c are the desired and actual positions of the cell, respectively and K , K_I and K_D are the proportional, integral and differential gains, respectively.

CIFH is not limited to the simulation of PID control algorithms – any discrete control algorithm can be simulated using this method. Additionally, the simulation is efficient enough to be run in realtime inside the control loop for model-based control.

2.4 Actuator modelling

A limitation of the interpolation approach to modelling the change in boundary conditions is that it assumes that a change from one boundary condition to another can be undertaken instantaneously, regardless of the physical plausibility of this. However, for a non-turbulent, viscous flow, the response time of the fluid system will be dominated by the response time of the actuator controlling the fluid flow rate. Hence, fluid response

time is incorporated with the actuator response time to give an effective system response parameter. In this paper, actuator response has been modelled as a first order system, which will exponentially converge on a final value at a rate determined by the system time constant

$$\frac{dq}{dt} = \frac{1}{\tau}(q_{fv} - q), \quad (7)$$

where q is the flow rate, q_{fv} is the final value of the flow rate (the demand flow rate) and τ is the time constant. While a simple first order model has been used, the actuator model can be extended in the general case to a discrete transfer function with higher order responses and incorporate non-linear behaviours such as actuator hysteresis.

A typical practical actuator is also limited by positioning resolution. The sources of error in an actuator may include digitisation resolution and noise; sensor accuracy and the response of the positioning controller. The limiting response of the controller can be modelled by (7). Digitisation noise and resolution are modelled by first quantising the demanded value to the actuator resolution, followed by the addition of zero mean white noise (Fig. 2a), where the noise amplitude is equal to the half the resolution.

2.5 Observer modelling

The behaviour of a limited observer, for example a camera of finite resolution, is also modelled. Positional data into the controller will be limited by the resolution and update rate of the cell position sensor (in many cases, this would be a camera coupled with appropriate software). As there is little benefit in updating the controller parameters faster than input can be obtained from the sensors, the update rate of the sensors is modelled by adapting the controller update rate to be equal to or slower than the projected frame rate of the cell sensor. Resolution is modelled as a digitisation process, similar to the actuator model described previously, with the position obtained from the cell tracking model combined with a white noise process (Fig. 2a) and quantised to the imager resolution.

3 Results and discussion

3.1 Channel geometry

A computational fluid dynamics (CFD) model of the cross slot geometry was constructed, with a nominal channel width of 100 μm . The Reynolds number, as defined by the channel width and the average inlet flow velocity, was 1 for a working fluid of water. By scaling the magnitude of the solutions, flows with lower Reynolds numbers were also generated. A high-order spectral element solver^{51,52} was used to accurately compute the solutions to the flow. Two-dimensional and three-dimensional geometries were solved for, with 3D aspect ratios between 0.75 and 2. For aspect ratios above 1 (channel height equal to channel width), the flow approached the two-dimensional result and the out-of-plane velocity was minimal.

As stated in §2.1, a parameter space grid of the possible solutions is needed to efficiently generate flow fields for the control system. A non-dimensional parameter for the ratio of the flow velocities between the two opposing inlets was defined

$$f_x = \frac{v_1 - v_2}{v_1 + v_2}, \quad (8)$$

where v_1 and v_2 are the average inlet flow velocities for inlets 1 and 2, respectively (see Fig. 1c). The flow rate deviation can be defined similarly for the two opposing outlets.

CFD solutions for over 350 combinations of inlet and outlet parameters were computed. The two dimensional location of the saddle point was used to map the change in the characteristics of the flow with respect to the input parameters. In each case, the saddle point position was identified using an automated Levenberg-Marquardt optimisation scheme,^{53,54} whereby a conic section was fitted to the velocity magnitude field in the approximate vicinity of the saddle point. Validation was performed using an experimental flow model and PIV (particle image velocimetry).^{55,56} Flow velocity fields show excellent agreement at a number of saddle point locations (comparisons for saddle points at (0,0); and (0,20) are given in the ESI†). The relationships between the relative inlet flows and saddle point for CFD and PIV showed good qualitative agreement. In both cases, the change in saddle point position and flow structures had a piecewise linear relationship, demonstrating that bilinear interpolation can be used to accurately obtain solutions for any combination of f_x and f_y within the solution space (Fig. 2).

3.2 Trapping simulation

Using the CIFH method, the speed and repeatability of cell capture and the stability of a cell trap based on a cross-flow geometry can be assessed in a computationally efficient manner. Cells enter from one of the inlets, in this case inlet 1 (Fig. 1c), and are assumed to be suspended in solution. Injection of cells into the fluid is not simulated, rather it is assumed that the cells were previously prepared and suspended in the working fluid prior to device activation or mixed into the working fluid *in situ* using an upstream T- or oscillatory-mixer.¹⁶ In either case, the location of the cells within the channel as they are transported by the working fluid is randomised – only a very small percentage of the cells will be transported along the channel centreline. In the case of healthy red blood cells, a parabolic distribution is expected.⁵⁷ Hence, to investigate the ability of the system to successfully trap an arbitrary cell that is transported into the device, the control response was simulated with an initial cell position 15 μm away from the centreline.

Fig. 3 shows the control response of the cell as it moves to the saddle point in the centre of the channel for an idealised and two non-ideal cases (a short video of each case is also available in the ESI†). Simulation parameters were chosen to model a red blood cell in the flow, where an average red blood cell is assumed, filled predominately with solution similar to water and of radius 4 μm . Relevant parameters for simulation cases I–III are summarised in Table 1.

Increasing noise and decreasing observer resolution affects both the stability of the final trap and the time to achieve a stable trap. As resolution decreases and actuator noise increases, oscillatory behaviour becomes more prevalent. However, the overall noise remains low, due to the relative size of the cell and low flow velocities. Even in the low resolution and high noise case (Case III), position noise is much less than 1 μm . This is due to the presence of digitisation noise – the long-term average of the




Fig. 3 Controlling the position of the cell using fluid flow. By adjusting the relative flow rates in the device, and therefore moving the saddle point, the flow velocity can be used to direct the cell. Based on this principle, a feedback control system can be used to position the cell and stabilise the system. Response curves show the y (vertical) position of the cell vs. time for case I (solid line); case II (dashed line) and case III (grey line).

Table 1 Parameters chosen for the trapping simulations (§3.2)

Simulation Parameter	Case I	Case II	Case III
Imaging rate	200fps		
Cell radius	4 μm		
Cell:Fluid density ratio	1 : 1		
Initial y location of cell		15 μm above centre line	
Proportional gain, x	0.2		
Proportional gain, y	1.25		
Channel Width	100 μm		
Average flow velocity	1 $\mu\text{m s}^{-1}$		
Reynolds number	0.1×10^{-3}		
Observer resolution	Ideal	0.1 μm	1 μm
Actuator noise	None	None	1%
Actuator time constant	0	50ms	100ms

observed location after quantisation will approach the true value, resulting in a stable trap as long as the cell velocities are sufficiently small relative to the acquisition rate. Larger time constants increase the time required to trap the cell, largely due to the increased time required to initially change the flow rates when the control system is activated. An increase of the actuator time constant from 50 ms to 100 ms is sufficient to increase the overall system damping such that the response moves from an underdamped response and begins to approach an overdamped response, further lengthening the trapping time.

3.3 Gain optimisation

For a number of imaging applications, it is important to minimise oscillation of the cell position. Any small movement of the cell will raise the effective noise floor of the image analysis. Additionally, in situations where a secondary high-resolution camera is used to analyse the cell image, it may be necessary to

maintain the cell position within a small region of the channel, as the secondary camera may have a much smaller field of view than the imaging device used for control feedback. In the absence of a closed-form solution for the control response and therefore for the RMS error, non-linear optimisation techniques provide a route to optimisation of the control gains. In this paper, the simulation technique is used to map the space of potential solutions and therefore locate the region of lowest RMS error. This kind of optimisation could also be performed online using an adaptive gain estimation technique.⁵⁸

Over 180 simulations were run, varying the proportional gain in the x and y directions. Due to the computational efficiency of the CIFH technique, these simulations took less than 2 h to complete on a modern multi-core system, corresponding to around 10 CPU-minutes (the product of the number of CPUs and the total runtime of the simulation) per simulation. This contrasts with the time required to generate the CFD data for the parameter space, which was in excess of 42 CPU-days. Therefore, the CIFH technique decreases the simulation time by more than 500 times, once the initial dataset has been generated.

The control system in each case was configured to stably trap the cell while maintaining a mean flow velocity of $10 \mu\text{m s}^{-1}$ ($\text{Re} \approx 1 \times 10^{-3}$) in the inlet channels; after 20 s of trapping the cell, the overall flow rates were increased to a mean flow velocity of $100 \mu\text{m s}^{-1}$ ($\text{Re} \approx 10 \times 10^{-3}$); the cell was maintained in the trap for a further 20 s before the simulation was terminated. Fig. 4a shows a contour of the RMS error of the final 20 s ($v_{\text{avg}} = 100 \mu\text{m s}^{-1}$) of the simulation for each solution. Results for proportional gains in the y -direction less than 0.04 are not shown as they did not result in stable cell trap.

As expected, there is a combination of proportional gains where the error norm is minimised. Additionally, the error in the y axis is strongly dependent on the gain in the x direction, demonstrating that the control in the two axes cannot be considered linear and separable systems. Further improvement of the steady state error can be achieved with addition of differential gain. Fig. 4b shows the variation of the error with varying differential gains in both directions, with the proportional gains fixed at 0.06. By selecting gains that minimise the error, for example $K_{D_x} = 0.25$, $K_{D_y} = 0.2$ the steady state error norm reduces from $0.75 \mu\text{m}$ (proportional-only case) to $0.25 \mu\text{m}$ (proportional and derivative control).

3.4 Cell sorting

The modelling technique described in this paper can also be applied to a simple cell sorting device. The geometry developed previously is used to simulate sorting of cells into two groups. A hypothetical device, based on this geometry, with a camera located above the left hand inlet channel (see Fig. 1b) is presented. A number of metrics could be used to identify and categorise the cell, including diameter, shape, or fluorescent response. Using these metrics, the cell would be placed into one of two categories, corresponding to the fluid outlet to which the cell should be directed. For this scheme to be effective, the concentration of the cells in suspension must be low enough so that the fluid forces applied to one cell do not cause the next cell in suspension to flow towards the wrong outlet. The simulator was therefore used to determine the relationship between

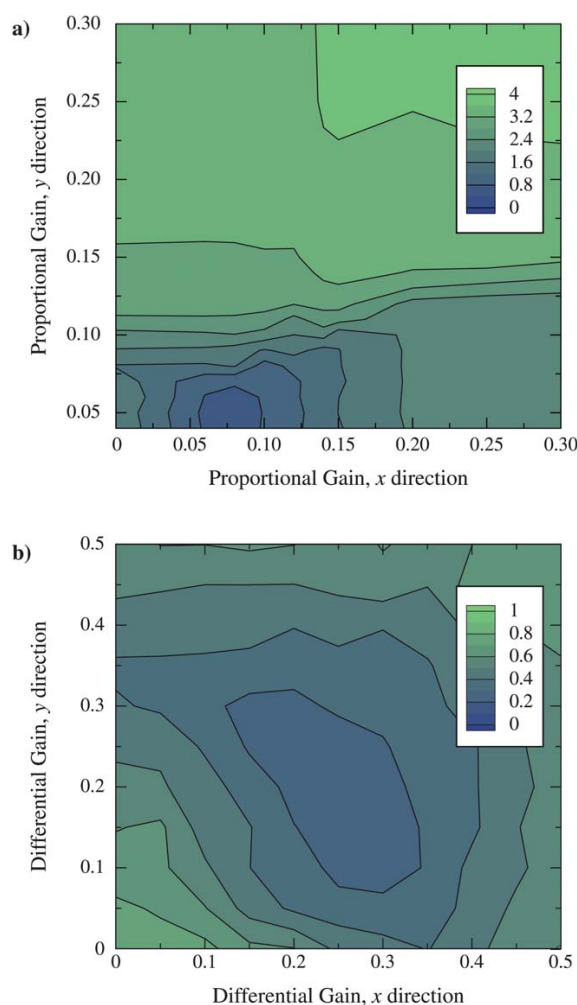


Fig. 4 The steady state error norm as a function of various PID gains. The control system sampling rate was fixed at 100Hz and the flow rate was $100 \mu\text{m s}^{-1}$. (a) Error as a function of proportional gains in the x and y directions. Data for $K_y < 0.04$ is not shown as the system was not stable. There is a clear region of stability and low error for $0.05 \leq K_x \leq 0.12$ and $0.04 \leq K_y \leq 0.08$. (b) Error as a function of differential gains in the x and y directions. For this study the proportional gains were fixed at $K_x = 0.06$ and $K_y = 0.06$.

concentration and accuracy. Parameters for this simulation are summarised in Table 2. Channel dimensions and geometry are the same as those discussed in §3.1. As with the previous simulations, a $4 \mu\text{m}$ red blood cell was simulated for both cell types. A practical sorting device based on these parameters could be used for sorting red blood cells based on their calcium-dependent fluorescence properties.⁵⁹

A number of simulation runs were performed, with 200 cells sorted in each run. Each cell was arbitrarily assigned a cell type, either alternating between 'type 0' and 'type 1' or randomly distributed. To simulate a fixed cell identification time, the cell type is not assigned until the cell is $50 \mu\text{m}$ from the centre of the channel. At the conclusion of each run, the cell type was

Table 2 Parameters chosen for the sorting simulations (§3.4)

Simulation Parameter	Value
Imaging rate	200fps
Cell radius	4 μm
Cell:Fluid density ratio	1 : 1
Initial y location of cell	Near centre line
Proportional gain, x	None (no control of x axis)
Proportional gain, y	0.8
Differential gain, y	0.05
Channel width	100 μm
Average flow velocity	10 $\mu\text{m s}^{-1}$
Reynolds number (water)	1×10^{-3}
Observer resolution	0.1 μm
Actuator noise	None
Actuator time constant	100ms

compared with the outlet port where the cell exited the fluid domain. The percentage of cells that exited through the correct outlet was used as a measure of the sorting success rate. Cells were inserted into the flow by the simulator at an number of different average injection rates. These injection rates were varied with a Gaussian distribution, with a standard deviation of 10% of the mean, to model the uncertainty present in a real dilution or injection scheme. A number of average rates, from one cell every 20 μm up to one cell every 80 μm , were simulated.

The resulting cell sorting success rates are shown in Fig. 5. For each cell concentration, 200 cells were simulated with an alternating distribution of cell types, representing the worst case, and a further 400 cells were simulated with a random distribution, representing the more common case. Excellent accuracy (95%) is achieved with a mean distance between cells of 75 μm and above. By factoring in the expected accuracy of the imaging system in correctly identifying a given cell type, the overall system accuracy

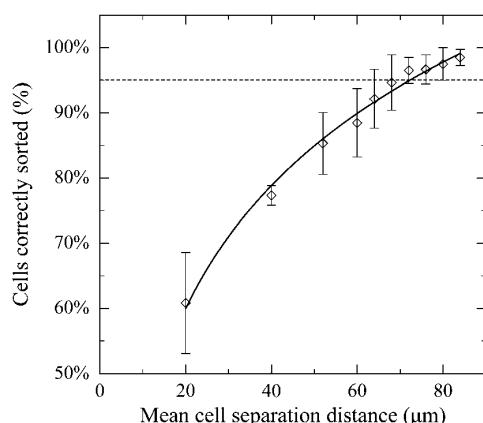


Fig. 5 Cell sorting using the cross geometry. Cells injected from the left inlet are arbitrarily and randomly assigned a type, and are then directed to the upper or lower outlet based on this type. Here the simulation is used to predict the probability that a given cell, assumed to be correctly identified, is directed to the correct outlet in the channel. Error bars indicate the range of values the tests performed, and the results have been fitted to a logarithmic function ($r^2 = 0.99$). The average flow velocity at the inlets is 10 $\mu\text{m s}^{-1}$. Over 95% accuracy is achieved if the cells are spaced at least 75 μm apart.

of a fluid-based cell sorting design can be estimated. This allows the system to be optimised for the throughput and accuracy demands of a given application.

4 Conclusion

An efficient, accurate and flexible method for simulating the control of suspended cells in microfluidic systems has been presented. It is expected that this method, CIFH, will have wide applicability in the design of new and sophisticated feedback control systems for microchannels. As the control simulation is based on a discrete time stepping method, a large number of existing control algorithms and techniques can be simulated using this technique. The method can take advantage of any flow solver, allowing the use of complex flow geometries. Careful choice of the parameter space and refinement of the parameter grid can be used to improve accuracy in cases where the CFD solutions become non-linear with respect to the control parameters. Additionally, the simple actuator models presented in this paper can be extended with measured transfer functions and hysteresis effects of real actuators to simulate the real world behaviour of complex systems.

Compared with earlier, conventional techniques, large improvements in the speed of the simulation have been demonstrated. These efficiency gains have been effected by employing three primary strategies: decoupling cell motion from the flow solver; using an LTI model to represent the change in flow boundary conditions over time; and interpolating between similar flow fields over multiple degrees of freedom. For the data presented in this paper, the generation of CFD flow fields took 42 CPU-days to complete, whereas a typical control simulation, such as those discussed in §3.3 takes under 10 CPU-minutes. Conservatively, the speed-up is well over 500 times when compared to a time-stepping CFD method.

The CIFH control simulation method has been demonstrated, by way of example, using a microfluidic cross slot. It has been shown that active PID control can be used to stabilise and capture a cell in the centre of the microchannel. Secondly, the efficiency of the technique allowed an exhaustive study of the control gain space to be performed to determine the optimal gains to minimise error. Finally, a proposed device for cell sorting has been investigated and the effectiveness of sorting quantified.

References

- 1 A. Fouras, M. J. Kitchen, S. Dubsky, R. A. Lewis, S. B. Hooper and K. Hourigan, *J. Appl. Phys.*, 2009, **105**, 102009.
- 2 R. A. Jamison, J. A. Armitage, J. Carberry, M. J. Kitchen, S. B. Hooper and A. Fouras, *Under Consideration for Current Pharmaceutical Biotechnology*, 2011.
- 3 M. Liu, A. K. Tanswell and M. Post, *Am. J. Physiol. Lung Cell. Mol. Physiol.*, 1999, **277**, L667–683.
- 4 H. R. Wirtz and L. G. Dobbs, *Respir. Physiol.*, 2000, **119**, 1–17.
- 5 M. L. Siew, M. J. Wallace, M. J. Kitchen, R. A. Lewis, A. Fouras, A. B. te Pas, N. Yagi, K. Uesugi, K. K. W. Siu and S. B. Hooper, *J. Appl. Physiol.*, 2009, **106**, 1888–1895.
- 6 J. R. Hove, R. W. Koster, A. S. Forouhar, G. Acevedo-Bolton, S. E. Fraser and M. Gharib, *Nature*, 2003, **421**, 172–177.
- 7 J. C. Culver and M. E. Dickinson, *Microcirculation*, 2010, **17**, 164–178.

- 8 W. S. Nesbitt, E. Westein, F. J. Tovar-Lopez, E. Tolouei, A. Mitchell, J. Fu, J. Carberry, A. Fouras and S. P. Jackson, *Nat. Med.*, 2009, **15**, 665–673.
- 9 E. Tolouei, C. J. Butler, A. Fouras, K. Ryan, G. J. Sheard and J. Carberry, *Ann. Biomed. Eng.*, 2011, **39**, 1403–1413.
- 10 S. S. Lee, Y. Yim, K. H. Ahn and S. J. Lee, *Biomed. Microdevices*, 2009, **11**, 1021–1027.
- 11 M. Rossi, R. Lindken, B. P. Hierck and J. Westerweel, *Lab Chip*, 2009, **9**, 1403.
- 12 F. J. Tovar-Lopez, G. Rosengarten, E. Westein, K. Khoshmanesh, S. P. Jackson, A. Mitchell and W. S. Nesbitt, *Lab Chip*, 2010, **10**, 291–302.
- 13 R. A. Jamison, S. Dubsy, K. K. W. Siu, K. Hourigan and A. Fouras, *Ann. Biomed. Eng.*, 2011, **39**(6), 1643–1653.
- 14 S. Tan, L. Yobas, G. Lee, C. Ong and C. Lim, *Biomed. Microdevices*, 2009, **11**, 883–892.
- 15 H. W. Hou, A. A. S. Bhagat, A. G. Lin Chong, P. Mao, K. S. Wei Tan, J. Han and C. T. Lim, *Lab Chip*, 2010, **10**, 2605.
- 16 A. Neild, T. W. Ng, G. J. Sheard, M. Powers and S. Oberti, *Sens. Actuators, B*, 2010, **150**, 811–818.
- 17 G. Mahmud, C. J. Campbell, K. J. M. Bishop, Y. A. Komarova, O. Chaga, S. Soh, S. Huda, K. Kandere-Grzybowska and B. A. Grzybowski, *Nat. Phys.*, 2009, **5**, 616–612.
- 18 C. Lim, E. Zhou and S. Quek, *J. Biomech.*, 2006, **39**, 195–216.
- 19 D. E. Jaalouk and J. Lammerding, *Nat. Rev. Mol. Cell Biol.*, 2009, **10**, 63–73.
- 20 C. Le Dévéhat, M. Vimeux and T. Khodabandehlou, *Clin. Hemorheol. Microcirc.*, 2004, **30**, 297–300.
- 21 E. J. Diamantopoulos, C. Kittas, D. Charitos, M. Grigoriadou, G. Ifanti and S. A. Raptis, *Horm. Metab. Res.*, 2004, **36**, 142–147.
- 22 H. Jin, X. Xing, H. Zhao, Y. Chen, X. Huang, S. Ma, H. Ye and J. Cai, *Biochem. Biophys. Res. Commun.*, 2010, **391**, 1698–1702.
- 23 C. T. Lim, *J. Biomech. Sci. Eng.*, 2006, **1**, 82–92.
- 24 A. G. Maier, B. M. Cooke, A. F. Cowman and L. Tilley, *Nat. Rev. Microbiol.*, 2009, **7**, 341–354.
- 25 K. C. Clause, L. J. Liu and K. Tobita, *Cell Commun. Adhes.*, 2010, **17**, 48–54.
- 26 S. McCue, S. Noria and B. L. Langille, *Trends Cardiovasc. Med.*, 2004, **14**, 143–151.
- 27 S. Kumar and V. Weaver, *Cancer Metastasis Rev.*, 2009, **28**, 113–127.
- 28 D. T. Butcher, T. Alliston and V. M. Weaver, *Nat. Rev. Cancer*, 2009, **9**, 108–122.
- 29 M. J. Fulwyler, *Science*, 1965, **150**, 910–911.
- 30 A. Bhagat, H. Bow, H. Hou, S. Tan, J. Han and C. Lim, *Med. Biol. Eng. Comput.*, 2010, **48**, 999–1014.
- 31 S. Miltenyi, W. Müller, W. Weichel and A. Radbruch, *Cytometry*, 1990, **11**, 231–238.
- 32 R. C. Spero, L. Vicci, J. Cribb, D. Bober, V. Swaminathan, E. T. O'Brien, S. L. Rogers and R. Superfine, *Rev. Sci. Instrum.*, 2008, **79**, 083707–7.
- 33 J. Guck, R. Ananthakrishnan, H. Mahmood, T. J. Moon, C. C. Cunningham and J. Käs, *Biophys. J.*, 2001, **81**, 767–784.
- 34 T. Franke, S. Braumüller, L. Schmid, A. Wixforth and D. A. Weitz, *Lab Chip*, 2010, **10**, 789.
- 35 M. S. Munson, J. M. Spotts, A. Niemistö, J. Selinummi, J. G. Kralj, M. L. Salit and A. Ozinsky, *Lab Chip*, 2010, **10**, 2402.
- 36 J. O. Barber, J. P. Alberding, J. M. Restrepo and T. W. Secomb, *Ann. Biomed. Eng.*, 2008, **36**, 1690–1698.
- 37 M. Kim, Z. Wang, R. H. W. Lam and T. Thorsen, *J. Appl. Phys.*, 2008, **103**, 044701.
- 38 M. Kim and C. Klapperich, *Lab Chip*, 2010, **10**, 2464.
- 39 C. M. Schroeder, H. P. Babcock, E. S. G. Shaqfeh and S. Chu, *Science*, 2003, **301**, 1515–1519.
- 40 M. Tanyeri, E. M. Johnson-Chavarria and C. M. Schroeder, *Appl. Phys. Lett.*, 2010, **96**, 224101.
- 41 M. R. Maxey, *Phys. Fluids*, 1983, **26**, 883.
- 42 E. E. Michaelides, *Particles, bubbles & drops: their motion, heat and mass transfer*, World Scientific, 2006.
- 43 J. Bonet and J. Peraire, *Int. J. Numer. Methods Eng.*, 1991, **31**, 1–17.
- 44 F. Burton and M. Huntbach, *IEEE Comput. Graphics Appl.*, 1984, **4**, 28–33.
- 45 D. J. Vojir and E. E. Michaelides, *Int. J. Multiphase Flow*, 1994, **20**, 547–556.
- 46 F. M. White, *Viscous fluid flow*, McGraw-Hill, 1991.
- 47 D. R. Breach, *J. Fluid Mech.*, 1961, **10**, 306–314.
- 48 Y. Sui, H. T. Low, Y. T. Chew and P. Roy, *Phys. Rev. E: Stat., Nonlinear, Soft Matter Phys.*, 2008, **77**, 016310.
- 49 K. Tsubota and S. Wada, *Phys. Rev. E: Stat., Nonlinear, Soft Matter Phys.*, 2010, **81**, 011910.
- 50 A. F. Ghoniem and F. S. Sherman, *J. Comput. Phys.*, 1985, **61**, 1–37.
- 51 G. J. Sheard, T. Leweke, M. C. Thompson and K. Hourigan, *Phys. Fluids*, 2007, **19**, 083601.
- 52 G. J. Sheard, *Journal of Fluids and Structures*, 2011, DOI: 10.1016/j.jfluidstructs.2011.02.005, **In Press**.
- 53 K. Levenberg, *Q. Appl. Math.*, 1944, **2**, 164–168.
- 54 D. W. Marquardt, *SIAM J. Appl. Math.*, 1963, **11**, 431–441.
- 55 A. Fouras, D. Lo Jacono and K. Hourigan, *Exp. Fluids*, 2007, **44**, 317–329.
- 56 A. Fouras, D. Lo Jacono, C. V. Nguyen and K. Hourigan, *Exp. Fluids*, 2009, **47**, 569–577.
- 57 G. Mchedlishvili and N. Maeda, *Jpn. J. Physiol.*, 2001, **51**, 19–30.
- 58 K. J. Aström and T. Hägglund, *Advanced PID Control, Instrumentation, Systems, and Automation Society*, 1st edn, 2006.
- 59 L. Kaestner, W. Tabellion, E. Weiss, I. Bernhardt and P. Lipp, *Cell Calcium*, 2006, **39**, 13–19.

3.3. Tuning the System Geometry

In the simulation work presented in this chapter, the intersection between the channels was modelled as a square intersection. In empirical studies, corner radii were introduced as a mechanism to change overall system dynamics and the effects these radii have on the system will be briefly discussed. In §5.2.2.4, it will be shown that with rapid changes in valve position, there is a peristaltic pumping effect as a result of the rapid change in tubing geometry. The magnitude of these pumping transients is largely independent of the system flow rate and is primarily related to the valve slew rate. Hence, as flow rate increases, the relative magnitude of the pumping transients decreases. By increasing the corner radius (r in figure 3.9b), the extensional rate (velocity) in the central region decreases for a given the inlet flow rate. This means that the system can be tuned to balance valve transients against the desired average extensional rate and system driving pressure. As r increases, the velocity in the central region and hence particle and cell motion decreases for a given flow rate, while also decreasing the relative magnitude of the valve transients. Hence, system dynamics can be tuned for high flow rates with minimal valve transients (with a large corner radius r) or efficient consumption of working fluid with larger perturbations as the flow is varied (lower r).

Figure 3.9 shows the effects of increasing corner radius r on the average velocity over a typical visible area (with respect to the camera and optics setup) in the centre of the channel. As this area is held constant, the effective size of the channel increases relative to the imaging area, resulting in a decrease in extensional rate and average velocity. For the experiments presented later in this thesis, two channel radii were chosen: the ‘small’ radius channel, with $r = 2.5w$ and a ‘large’ radius channel, $r = 10w$. The relative decreases in velocity for these two radii are $v|_{r=2.5w} = 0.178 \times v|_{r=0w}$ and $v|_{r=10w} = 0.017 \times v|_{r=0w}$, respectively. In terms of experimental parameters, compared to the ‘small’ channel, the ‘large’ channel will exhibit a valve pumping transient that is approximately $10\times$ smaller for a given average velocity.

The increase in corner radii also has an effect on the relationship between saddle-point position and flow rate bias. As the central region is much larger than the channel width, the saddle-point is displaced much further for a given flow-rate bias, relative to the channel width w . This effect can be seen in figure 3.10. A flow rate bias of $b = \pm 0.2$ produces a saddle-point position of $|s| = 0.5w$, as compared to $b = \pm 1.0$ for the $r = 0$ channel. For a fixed-size imaging region, this has the effect of reducing the sensitivity of the flow control valves. However, the coupling between planes is significantly reduced for the larger channel, due to the reduction in wall effects. In the $r = 0$ case, worst-case positional error due to coupling between planes was $0.5w$ ($|b| = 0.7$, $|s| = 0.57w$). In the $r = 10w$ case, worst-case error was 0.04 over the range $|s| \leq 0.5w$ and $0.12w$ at $|s| = 1.35w$. Furthermore, the relationship between saddle-point position over the region $|s| \leq 0.5$ is linear ($R^2 = 0.999$), rather than polynomial.

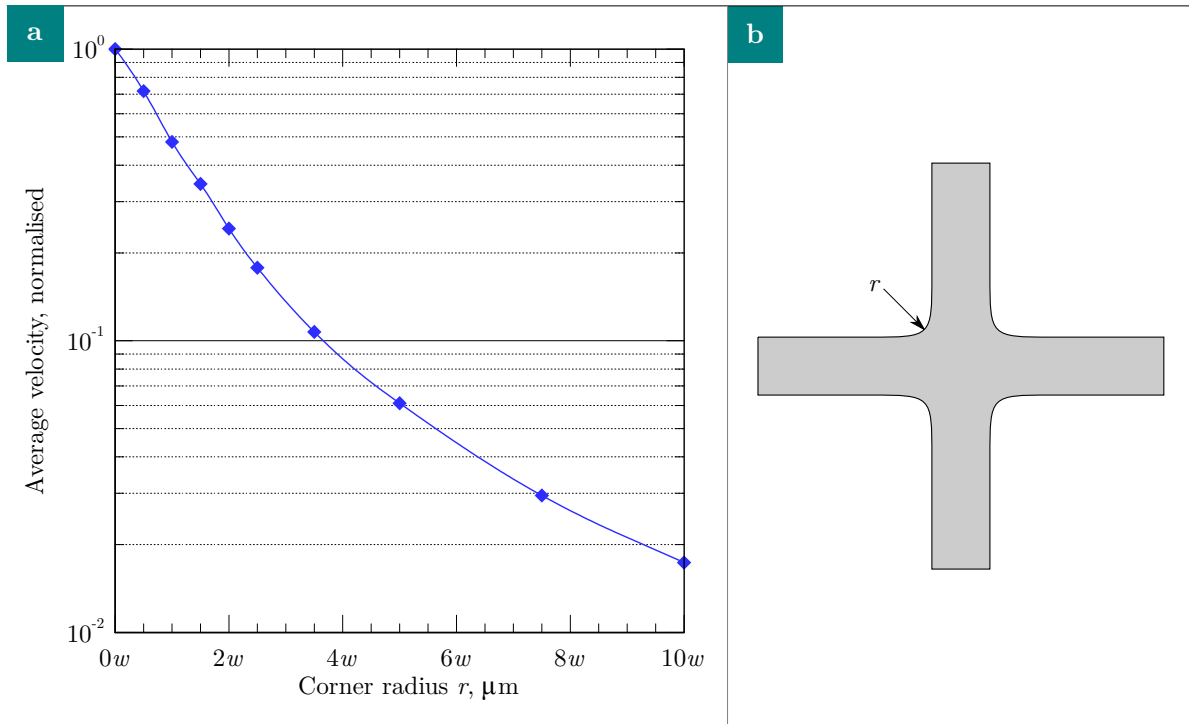


Figure 3.9: Velocity in the central region as a function of corner radius, normalised to the average velocity where $r = 0w$ (i.e., no radius, equivalent to the geometry in figure 3.1). As corner radius increases, average velocity decreases logarithmically.




Figure 3.10: Saddle-point position as a function of flow rate bias for the ‘large’ channel ($r = 10w$). Straight line and error bars are as before (the line plots the saddle-point position s_x for $b_y = 0$, and error bars indicate maximum deviation due to variation in b_y).

3.4. Conclusion

In this chapter, a detailed study of the flow behaviour in a cross-slot microchannel has been conducted. This led to development and presentation in *Lab on a Chip* of a new type of system simulator, that allowed extensive exploration of the behaviour of the control system. By using an innovative technique to use precomputed CFD flow data in a 2D lookup table, the response and stability characteristics of the control system for trapping cells could be studied in a computationally efficient manner. This simulation software has been used extensively throughout the development of the technology in order to better understand system behaviour — including stability limits, optimal parameters, and the effect of various flow geometries and actuator behaviours. Finally, the effect of corner radii in the central region have been extensively studied to show that both accuracy and undesirable transient behaviours can be optimised by carefully tuning the corner radius size to the expected flow velocity range.

3.5. Bibliography

- [1] *Gambit*. version 2.4.6. (2007). Ansys, Inc. URL: <http://www.ansys.com/>.
- [2] George Em Karniadakis, Moshe Israeli and Steven A Orszag. “High-order splitting methods for the incompressible Navier-Stokes equations”. In: *J. Comput. Phys.* 97.2 (Dec. 1991), pp. 414–443. ISSN: 00219991. DOI: 10.1016/0021-9991(91)90007-8.
- [3] K. Levenberg. “A method for the solution of certain non-linear problems in least squares”. In: *Q. Appl. Math.* 2.2 (July 1944), pp. 164–168.
- [4] D. W. Marquardt. “An Algorithm for the Least-Squares Estimation of Nonlinear Parameters”. In: *SIAM J. Appl. Math.* 11.2 (June 1963), pp. 431–441.
- [5] A. Neild et al. “Swirl mixing at microfluidic junctions due to low frequency side channel fluidic perturbations”. In: *Sens. Actuators B: Chem.* 150.2 (2010), pp. 811–818. ISSN: 0925-4005.
- [6] G. J. Sheard and K. Ryan. “Pressure-Driven Flow Past Spheres Moving in a Circular Tube”. In: *J. Fluid Mech.* 592.-1 (2007), pp. 233–262. DOI: 10.1017/S0022112007008543.
- [7] G. J. Sheard et al. “Flow around an impulsively arrested circular cylinder”. In: *Phys. Fluids* 19.8 (2007), p. 083601. ISSN: 10706631. DOI: 10.1063/1.2754346.

4. Flow Measurement

A critical feature of the system is the ability to measure the flow in realtime. This allows the control loop to be closed by the realtime feedback of flow information to the controller, as well as providing additional information to aid in cell tracking (as will be seen in chapter 6); valve calibration (chapter 5); and measurement of the shear-rate and distance of the cell from the saddle-point (chapter 8).

4.1. Review of Flow Measurement Techniques

There are a large number of techniques that can be used to measure the volumetric flow rate or flow velocity. Of these, a small number of techniques will be considered in detail with a view to selecting the technique most suitable to the applications under consideration. Of particular importance were the bio-compatibility of the technique, the cost and speed of measurement response.

4.1.1. Differential Pressure

Measurement of flow via differential pressure transducers is a simple, inexpensive technique that is used extensively across a wide variety of industries. The design of such flow meters is of sufficient importance that the ISO has published several standards on the topic (ISO 5167 parts 1 through 6). A simplified schematic of an orifice plate flow meter is shown in figure 4.1. An orifice plate flow meter is not the only type of differential pressure flow meter, but the principles are broadly similar regardless of whether a venturi, nozzle or orifice is used. The pressure drop across the restriction in the flow (in this case an orifice) can be related to flow rate in a repeatable fashion (for orifice plate flow meters, the relationship is simply $Q \propto \sqrt{\Delta p}$). This method is simple, reliable and well-documented.

However, most pressure sensors are not suitable for direct contact with liquids. This is often resolved with a gel-filled barrier[†], or a pocket of air (dark shaded section leading to P_1 and P_2 in figure 4.1). In either case, this presents an additional medium that will modify the frequency response of the final flow rate. Air, in particular, due to its compressibility will act as a low-pass filter on the flow rate signal. Furthermore, the constriction presents an additional surface for accumulation of biological material. The design of the orifice plate flow meter in particular is not unlike the constrictions introduced by Tovar-Lopez et al.²⁵ to induce platelet activation and aggregation.

4.1.2. Thermal Techniques

Sensors that use a thermal time-of-flight measurement have recently begun to reach commercial maturity in the microfluidics space.^{2,4,5} These sensors introduce a small increase in temperature to the flowing fluid and measure the time taken for this temperature increase to be conducted downstream. As the temperature sensing and heaters can be fitted externally to the wetted parts, these devices are non-invasive. However, this is a new technology in the microfluidics space and is still quite expensive (a particular Dolomite-branded sensor costs about AU\$3 000 per channel). Furthermore, it is unclear that the dynamic performance of these sensors is sufficient for realtime control (one manufacturer quoted a response time of 30 ms or $6\times$ slower than the target imaging framerate).

[†]For example, the MPX2300DT1 medical-grade pressure sensor from Freescale Semiconductor is manufactured with a gel barrier installed on the positive pressure port.

4. Flow Measurement




Figure 4.1: Schematic of a typical orifice plate flowmeter. A small, known, restriction is placed inline with the flow and the pressure drop across the restriction is measured. With this method, both the bulk flow rate and direction can be measured directly.

4.1.3. Ultrasound

In a similar fashion to the thermal sensors, ultrasound sensors are mounted externally to the wetted parts and are therefore non-contact. There are two basic types: a time-of-flight sensor, where an ultrasonic pulse is transmitted through the flowing liquid, and the time taken for the pulse to be received at the opposite side of the tube is measured; or an ultrasonic Doppler sensor, which measures the frequency shift in the transmitted pulses. In either case, these are expensive sensors (typically more expensive than the thermal sensors discussed above). Additionally, the typical measurement bandwidth is too low for this application. The ultrasonic sensor available to the author, the TRANSONIC T410 (Transonic Systems Inc.), appeared to have a frequency response on the order of 10 Hz in empirical studies.

4.1.4. Particle Image Velocimetry

Particle Image Velocimetry is an image-based technique that utilises the changes between successive imaging frames to measure the flow velocity. It has been applied to flows as diverse as those found in combustion chambers;⁹ helicopter rotor vortices;¹⁷ and microfluidic flows.¹¹ A simplified description of the technique follows in figure 4.2. The measured images are divided into small regions (interrogation windows) and a cross-correlation is performed between each window in the first imaging frame and the corresponding window in the second imaging frame. Formally, the direct cross-correlation is defined as

$$c[s_x, s_y] = \sum_{y=0}^h \sum_{x=0}^w a[x, y] \times b[x - s_x, y - s_y], \quad (4.1)$$

for $-\frac{w}{2} \leq s_x \leq \frac{w}{2}$ and,

$$-\frac{h}{2} \leq s_y \leq \frac{h}{2}$$

where $c[s_x, s_y]$ is the result of the correlation, w and h are the width and height of the correlation window, and $a[x, y]$ and $b[x, y]$ are the correlation windows for the first and second frames, respectively. For computational efficiency, the correlation is often performed in Fourier space

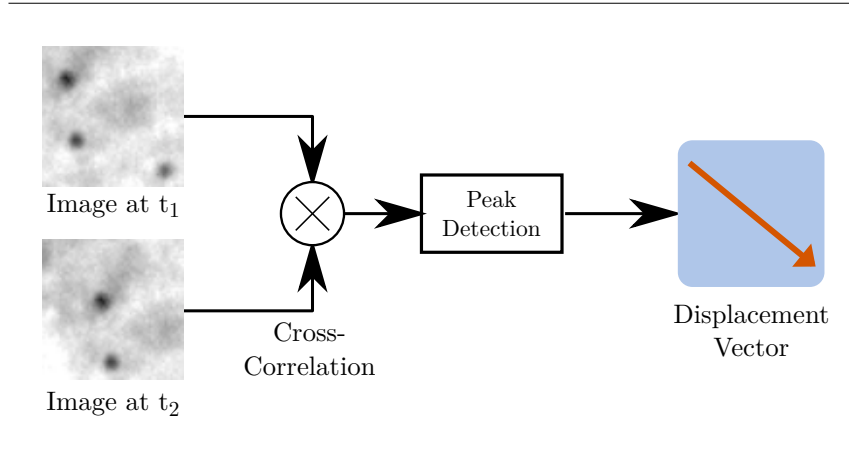


Figure 4.2: Schematic of process of obtaining flow velocities using Particle Image Velocimetry (PIV).

$$c[s_x, s_y] = \mathcal{F}^{-1} \{A(s) * B(s)\}, \quad (4.2)$$

where $A(s)$ and $B(s)$ are the Fourier transforms of a and b respectively and \mathcal{F}^{-1} is the inverse Fourier transform.

In either case, the result $C[s_x, s_y]$ can be considered to be a probability density function which is at its maximum at the most likely (modal) velocity for that imaging window.¹⁹ The location of this maximum is used to determine the velocity (displacement for one frame pair) for each imaging window (see figure 4.2). This process will be discussed in greater detail in section 4.2.2.

In order for PIV to be effective, the flow must be seeded with tracer particles to provide a moving pattern for the cross-correlations to match between frames (there are some flows that do not require the addition of tracer particles — for example some biological flows¹⁰ — but artificial seeding is the norm). This does introduce a bio-compatibility concern, due to potential interactions between the cells and tracer particles. However, a number of suitable particles are commercially produced for both laboratory and clinical use. PIV has also been successfully used in the past to analyse flow around biological cells.^{12,13}

The advantages of PIV are that it is inherently non-contact (as it is image based); and, as the system requires a high-speed imaging system for tracking the location of the trapped cell, does not require additional hardware for this application, which reduces the complexity (and cost) of the experimental setup. Additionally, high temporal resolution can be obtained, limited only by the speed at which the camera can image two successive frames. Finally, the flow can be measured much closer to the trapped cell and flow saddle-point than is possible with the techniques described earlier in this section, that all require a device to be inserted into the flow loop. This removes any measurement delays or other uncertainties that might be incurred by long tubing lengths between the flow meter and saddle-point.

However, PIV is typically employed as an analysis technique, rather than a realtime measurement technique. The computation of cross-correlations for each window takes quite some time, even when using the common implementation technique of computing the cross-correlations in Fourier space, rather than in the direct form of equation (4.1). A typical implementation on a 640×480 px image requires over 70 ms per frame to compute the cross-correlations and associated flow field. In section §4.2, a number of algorithmic improvements will be discussed that, combined with specialised computing hardware, reduce this computation time to below 5 ms per frame.

4. Flow Measurement

4.1.4.1. Double-Pulsed PIV

In the preceding section, it was mentioned that the measurement bandwidth of PIV is limited only by the time between acquisition frames. In this section, a common technique — double-pulsed (or double-shutter) PIV — will be discussed that allows the time between acquisition frames (the Δt) to be reduced to much less than might be implied by the camera framerate. In double-pulsed PIV, named for the type of laser traditionally used to provide illumination, frames are captured in pairs, where the time between exposures (i.e., the time over which incident light on the sensor is integrated, or simply *integration time*) within a pair is relatively short compared to the time between pairs. This is explained graphically in the timing diagram in figure 4.3, where the inter-pair time was chosen to be 8 ms, an effective integration time of 2 ms and an intra-pair time of 30 ms. The advantages of this style of acquisition are that the integration times are very short, reducing motion blur;¹ and fast flows, that would have large particle displacements between frames at the nominal acquisition rate, can be accurately captured by increasing the effective acquisition rate (and thereby reducing particle displacement) within the pair. This method of operation is only possible with cameras that are able to simultaneously read out the pixel intensities (which involves a time consuming conversion of each pixel from analogue to digital[‡]) while exposing a subsequent frame. Until recently, this has been only possible with *interline transfer* CCDs (charge-coupled devices), which have a secondary set of charge accumulators for each line (row) that can rapidly receive (transfer) charge from the primary set of charge accumulators that are exposed to light. It is these secondary charge accumulators which are read during the read out period, allowing the primary charge accumulators to receive light for a new exposure. Modern sCMOS (Scientific CMOS) imaging sensors now also allow this type of exposure in a mode known as *global shutter*.

It should be noted that although the illumination pulses are short and asynchronous, the imaging sensor has a constant (synchronous) exposure time. This has two important implications. From an experimental perspective, this means that the imaging sensor is integrating light over the full period even though the light source is pulsed at a fraction of this period. Thus, careful attention must be paid to extraneous sources of light in this type of imaging setup, as ambient light may introduce a secondary, undesirable, light source that increases the background noise in the image. Secondly, from a timing perspective, it should be noted that the asynchronous nature of the illumination pulses are in stark contrast to the synchronous readout behaviour of the imaging sensor. In the example shown in figure 4.4, the second illumination period of pair one, which begins at 18 ms and is 2 ms in duration, is associated with the second exposure period, which does not complete until 28 ms. Readout only begins at the end of the exposure period, due to the synchronous nature of the camera timing, and does not complete until 42 ms, or 24 ms after the start of illumination. The nominal frame rate in figure 4.4 is $1/15$ ms, and it follows from careful study of the timing, that the delay d between the start of integration and the completion of readout is $1/f < d \leq 2/f$ (where f is the frame rate). This delay is inherent to all interline transfer or global shutter acquisitions, and must be considered in any realtime imaging application.

[‡]A number of image sensors have multiple analogue to digital converters, but there are nevertheless many more pixels than there are A/D converters, so the conversion process is predominately a serial conversion.




Figure 4.3: Example Double-Pulsed PIV Timing. The time between exposures within a pair is 8 ms and the time between pairs is 30 ms.

4.2. Realtime, GPU-accelerated, PIV

Performing PIV on a graphics processing unit (GPU) has been of interest to a number of researchers.^{3,20,23,24} However, these works focussed on the speed gains that could be achieved in offline processing, and/or did not achieve the performance required to analyse the images in realtime at the high frame rates required for effective cell capture (≥ 200 frames/s).

In micro-PIV (that is, PIV performed on a microscope), high pixel displacements are commonly observed due to the relatively high magnification of the image.¹¹ With the $20\times$ magnification used in this device, the pixel displacements per frame are effectively magnified by the same amount, resulting in much larger pixel displacements than might be expected from the flow rates alone. Even at the relatively fast imaging rate of 200 frames/s, particle displacements much higher than the Nyquist-Shannon limit (that is, a displacement greater than half the window size) were encountered. Hence, a double-pulsed imaging setup was used to allow the particle displacements to be accurately captured. This had some important consequences to the implementation of PIV in the GPU. A naïve implementation would be to only perform analysis on every second image frame (i.e., at the end of each pair). However, the GPU can be used more efficiently by performing some of the processing ahead of time after the acquisition of the first frame in the pair. As discussed in section 4.1.4, cross-correlations for the purposes of PIV can be efficiently computed with the use of fast Fourier transforms. This transforms the expensive convolution operation into a multiplication in Fourier space, which can be therefore computed much more quickly than a spatial-domain convolution. In the first frame, the image can be windowed and the forward Fourier transform can be computed ahead of time. Once the second frame has been received, the same windowing and forward transform operation can be computed, and then multiplied by the previously-computed forward transform from the previous frame. This Fourier-space correlation can then be transformed back to the spatial domain by means of the backward transform and a displacement vector computed from the resulting correlation peak. As the Fourier transforms make up the bulk of the runtime, splitting the computational load between frames

4. Flow Measurement




Figure 4.4: Expanded double-pulsed timing diagram, showing camera typical exposure and readout timing. Using interline transfer (CCD) or global shutter (sCMOS), the integrated light captured during the exposure period can be transferred near-instantaneously to a readout buffer, where the pixels can be read out during a subsequent exposure.

in this way reduces the computation time in the second frame by approximately one third, as only one forward and one backward transform are required (a total of two transforms) rather than two forward and one backward (a total of three transforms).[§] This process is shown graphically in the flowchart in figure 4.5.

4.2.1. Efficient Windowing

In order to extract the maximum performance from graphics processors, work should be submitted in large batches.¹⁴ This allows the GPU to maintain high occupancy (*i.e.* a large proportion of the processing cores in the GPU are simultaneously performing work) especially on units that have a high number of cores (the NVIDIA M2050, purchased in 2011, possesses 448 cores; newer units have in excess of 1500 cores). The NVIDIA cuFFT API allows a large number of 2D transforms to be submitted in a single batch.¹⁵ However, the API requires that there is a consistent number of bytes between each window - the formula for computing the memory location of a window within the image array is simply

$$M_w = M_a + i_{stride}w, \quad (4.3)$$

where w is the window number, i_{stride} the number of bytes per window, and M_a and M_w the memory locations for the start of the array and start of the window, respectively. This formula only holds for pixel spacings between windows that divide exactly into the image width and height (see figure 4.6). However, this is often not the case, and therefore the image data must be first copied from the original image array into a series of windows. This reduces memory contention, and therefore improves performance, in the case of overlap (that is where the window spacing is less than the window size), as each FFT in the batch operates out of a separate region of GPU memory.

In order to implement the windowing routine on the GPU, the image data is first mapped to a texture. This allows efficient concurrent access of multiple processor cores to the same pixel location (which will occur with overlap) as well as built-in control of padding. The GPU texture hardware has an internal bounds checker which can be programmed to return a number of different values when pixel locations outside of the image are accessed. This relieves the GPU cores of performing these checks, which improves the performance of the windowing operation as conditionals incur a performance penalty on Single-Instruction-Multiple-Data (SIMD) architectures, such as GPUs, as all cores that are part of execution group (which NVIDIA calls a *warp*) must execute in lockstep. In this implementation, out-of-bounds pixels are set to zero, which is a common choice in PIV implementations^{16,18} (alternatives include padding the data with a mirrored copy of the image). Once the image data has been copied into a number of individual windows, the execution of the forward transforms can be queued on the GPU in one API call, which allows the CUDA API to schedule the transforms to achieve optimal occupancy and performance.^{14,15}

4.2.2. Peak Detection

After computation of the backward-transform, the result is a number of 2D correlation functions. These can be viewed as series of probability density functions (PDFs) of the most likely displacement/s of the particles in that window. Returning to the definition of the discrete 1D correlation function

[§]A similar optimisation can be applied to cinematic (that is, not double-pulsed) acquisition by using a buffer swapping technique - once the forward transform for the second frame has been computed, the buffers holding the forward transforms for frames 1 & 2 are swapped, and the forward-transform for frame 3 stored in the place of the forward-transform for frame 1. Hence the buffers are continually cycled through and only two Fourier transforms are required per frame.

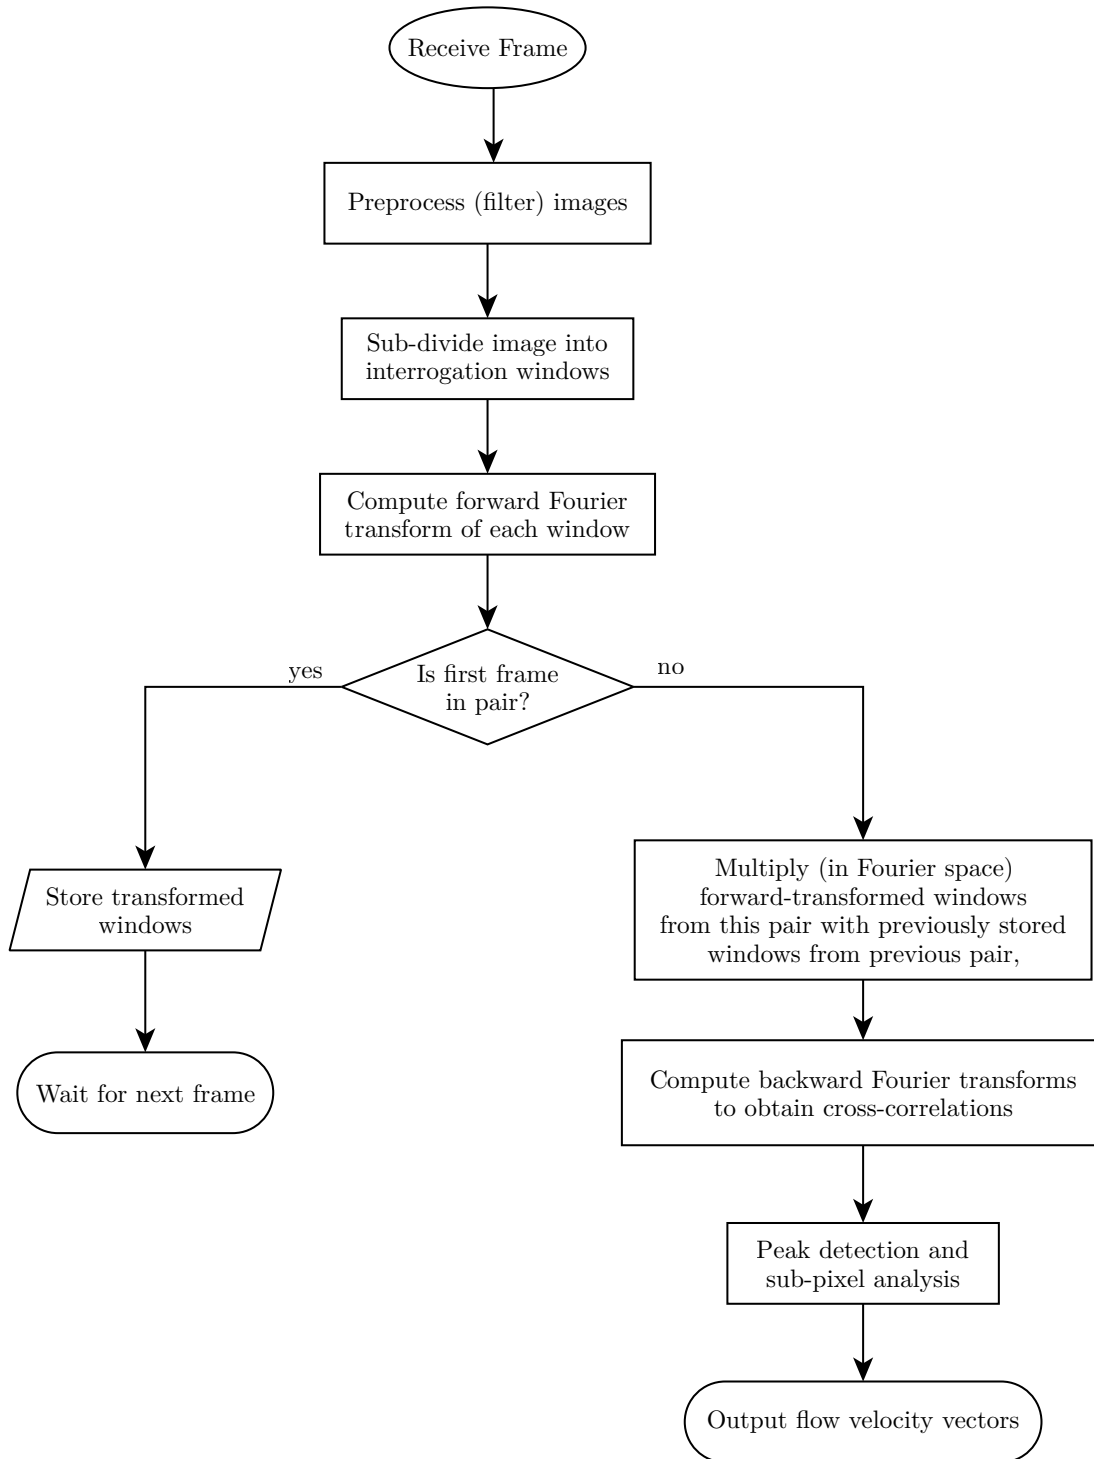


Figure 4.5: Double-pulsed PIV flowchart, showing the key steps in the efficient computation of velocity vectors.




Figure 4.6: Results of applying equation (4.3) in a situation where the spacing between windows is not a multiple of the image size. Window 3 extends past the end of the image array and, as 2D arrays are actually 1D arrays in memory, the remainder of the window “wraps” around the image and contains part of Window 1 (shaded region).

$$C[s] = \sum_{i=0}^N A[i+s]B[i], \quad (4.4)$$

where $A[i]$ and $B[i]$ are the two input datasets and s is the shift ($-N/2 \leq s \leq N/2$), it can be inferred that $C[s]$ is maximum where $A[i+s]$ (i.e., A , shifted by the shift s) is most similar (in a probabilistic sense) to $B[i]$. Extending this to 2D, the correlation function is maximised when the image \mathbf{A} , shifted by the shift vector \vec{s} results in an image that is most similar to the image \mathbf{B} . If all particles in the image have the displacement \vec{d} , then the correlation matrix \mathbf{C} will have a maximum at $\vec{s} = \vec{d}$.

This maximum location is found by a parallel reduction algorithm. The correlation matrices are first searched in blocks of 64 pixels. That is, each execution unit within the GPU is assigned to and loops through a block of 64 pixels in one of the correlation matrices. The location and value of the maxima within each block is recorded in shared memory. These local maxima are then compared in blocks of 32 to find the maximum within each block, and the location and maximum value are again recorded. This process repeats until there remains a single value for the maximum value and the location of that maximum value for each window. The locations of these maxima are the shift vectors \vec{s} that correspond to the most likely displacement of particles within the image. By performing the peak detection in stages, the number of GPU cores executing in parallel is maximised. A simplistic implementation might be to have each GPU execution unit loop through all pixels in a correlation matrix and record the maximum value. This implementation would be parallel across correlation matrices, but is not parallel within a matrix. The parallel reduction algorithm is parallel in both dimensions and is therefore superior.

4.2.3. Performance and Accuracy

To assess the accuracy and relative performance of the GPU PIV code developed, the output of the code was compared to a well-tested and proven code. The reference code selected was one developed by researchers at Monash (Fouras, Dusting, Lo Jacono, *et. al.*) using the FFTW Fourier transform libraries (executing on a single CPU core). This code has been validated extensively in a variety of flow regimes and imaging modalities.^{6,7} Validation was performed using a combination of synthetically generated images as well as real images taken from the device. Synthetic image generation was carried out using image generation software included in the reference code. Particles were initially randomly

4. Flow Measurement

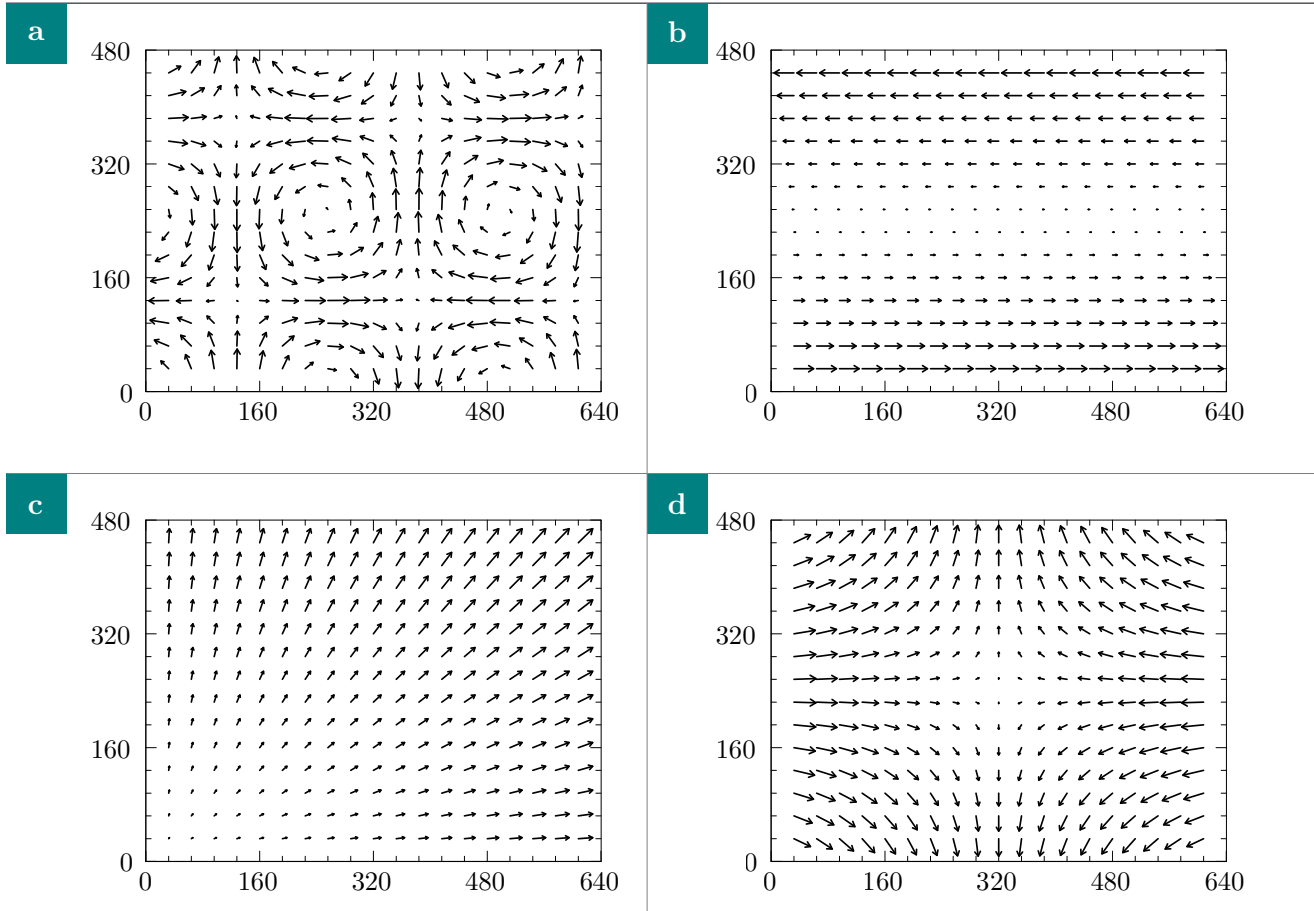


Figure 4.7: Synthetic flows used for validation of the GPU PIV code. a) Taylor vortices; b) simple shear flow; c) expanding flow designed to test sub-pixel resolution; d) extensional flow similar to the flow in the cross-slot device.

seeded (placed) in the image and displaced by an amount proportional to the flow vector in the respective particle locations. Reference flow fields for synthetic image generation were supplied from a mixture of numerically-generated flows (Taylor vortices, simple shear flow, expanding flow) and solutions from the CFD solver VIPER (cross-slot flow). Vector fields used to generate the synthetic images are shown in figure 4.7. For consistency and ease of comparison, all images, both synthetic and real, were 640×480 px and consisted of 20 image pairs (i.e. a total of 40 images). The seeding densities for synthetic cases a,b and c were 0.013 particles/px whereas the seeding density for synthetic case d was lowered to 0.003 particles/px to more closely correspond to the seeding density observed in the real images. Finally, image noise was added to the synthetic images in the form of white noise with mean intensities of 0% (that is, no noise); 12.5%; and 25%.

4.2.4. Accuracy and Robustness

After generating the synthetic images and acquiring images from the real device, all images were analysed with both the reference (CPU) code and GPU code. A window size of 64×64 px was chosen, with a spacing between each window of 32 px in each direction. Additionally, the reference code has a number of algorithms for sub-pixel interpolation (which improves upon the vector location found

via peak detection) - in all cases, the reference code was run with the most accurate algorithm (χ^2 fit) and a simple three-point interpolation. The three-point interpolation is significantly faster, yet potentially less accurate, and is also the algorithm used by the GPU PIV code.

To allow simple comparison between results, the error (in a least-squares sense) was then computed for each component of velocity

$$e_j = \sqrt{\frac{1}{N} \sum_{i=0}^N (\hat{\mathbf{x}}_{ij} - \mathbf{x}_{ij})^2}, \quad (4.5)$$

where E_j is the RMS error for the j^{th} component of velocity, N the total number of vectors in the solution, $\hat{\mathbf{x}}$ the result from PIV and \mathbf{x} the expected flow field. For the synthetic images, the expected flow field was that used to generate the synthetic images. For the experimental data, no such ground-truth or expected result exists. Instead, an approximation was generated by using the reference implementation with a multi-grid technique and vector rejection to achieve the maximum possible accuracy (the specific settings used were window sizes of 128×128 ; 64×64 ; and 32×32 , vector spacing of 16 px and a maximum vector error of 5 px).

To obtain an error estimate that is independent of the flow field, the error can be expressed as a percentage of the RMS average of the reference velocity magnitude

$$E = \frac{|e|}{\sqrt{\frac{1}{N} \sum_{i=0}^N |\mathbf{x}_i|^2}}. \quad (4.6)$$

The results of this error analysis are summarised in table 4.1. As can be seen from the table, the GPU implementation compares favourably with the reference implementation, achieving identical error scores when compared with the reference algorithm (in three-point mode). In most cases, the error in the results from the reference three-point and GPU implementations approach those of the reference χ^2 implementation, except notably in case c with 50% noise. This is a particularly stringent test, with very small sub-pixel velocities and demonstrates the limits of the three-point interpolation algorithm. Finally, an error estimate was computed (using equations 4.5 and 4.6) with the result from the reference implementation in three-point mode as the reference (\mathbf{x}) and the GPU result as the estimate ($\hat{\mathbf{x}}$). The error estimates obtained using this method are non-zero, while the error estimates as compared to the noise-free original estimate are identical. This indicates that the overall process is similar, with the same error characteristics, but that differences in implementation (particularly in the FFT implementation, which is FFTW in the reference implementation and cuFFT in the GPU implementation) generate results which are not identical.

4.2.5. Performance

Both the reference code and the GPU code were instrumented with timers that measured the computation time for each frame-pair. Care was taken in both cases to measure the time taken to read the images from disk independently of the computation time, allowing the separation of computation time from I/O time. In the application under consideration in this thesis, images originate directly from a camera, so the time taken to read an image from disk does not limit the performance of the system. Therefore, all figures presented below are exclusive of disk I/O time. To reduce the effects of noise in the timings due to the impact of other processes on the test system, the timings presented are an average of multiple runs.

4. Flow Measurement

Case	Image Noise	RMS Error (equation 4.6)			
		Reference (χ^2)	Reference (three-point)	GPU	GPU <i>c.f.</i> Reference
Synthetic Case a Taylor Flow	0 %	12.1 %	12.1 %	12.1 %	11.5 %
	12.5 %	11.8 %	11.8 %	11.8 %	11.0 %
	25 %	13.1 %	13.0 %	13.0 %	13.0 %
Synthetic Case b Shear Flow	0 %	4.7 %	4.7 %	4.7 %	4.7 %
	12.5 %	4.7 %	4.7 %	4.7 %	4.7 %
	25 %	5.7 %	5.7 %	5.7 %	6.6 %
Synthetic Case c Sub-Pixel Accuracy	0 %	45.3 %	45.3 %	45.3 %	45.1 %
	12.5 %	32.1 %	32.3 %	32.3 %	16.9 %
	25 %	71.9 %	100.1 %	100.1 %	137.8 %
Synthetic Case d Extensional Flow	0 %	11.1 %	11.0 %	11.0 %	11.4 %
	12.5 %	10.9 %	10.9 %	10.9 %	11.3 %
	25 %	11.1 %	12.2 %	12.2 %	12.9 %
Experimental Flow (Real Device)	N/A	N/A	10.6 %	10.6 %	12.5 %

Table 4.1: Comparison of RMS error in velocity as compared to the reference flow field. The final column, GPU *c.f.* Reference, compares the RMS error of the velocity computed by the GPU PIV code (developed in this thesis) with the Reference PIV code (in three-point mode). Shaded cells indicate errors which diverge significantly from the CPU (χ^2) result.

Implementation	Computation time (ms/frame)
Reference (χ^2)	171.0 ± 0.2
Reference (three-point)	74.5 ± 0.3
GPU (three-point)	2.2 ± 0.1

Table 4.2: Comparison of the computation times for various PIV implementations.

As can be seen from table 4.2, the computation time for the GPU implementation is significantly faster than the reference (three-point) version - over $30\times$ faster. While similar performance improvements may be realised with a multi-core CPU implementation, a 32 core machine would be required, assuming linear scaling, whereas the GPU implementation achieves this speed-up with a single lightly-loaded CPU core and a GPU, freeing any remaining CPU cores for other processing (for example, flow control — chapter 5; and tracking — chapter 6). Most importantly, the GPU implementation can process a frame-pair in much less than 5ms and is therefore well within the timing budget for realtime processing of images at 200 frames/s.

4.3. Identification of the Saddle-Point

In sections 3.1.5 and 3.2, it was shown that the location of the point in the flow where the velocity is zero (that is, the saddle-point) was an important determinant of the acceleration of the cell(s) undergoing trapping in the system, and is also an effective proxy for the relative flow rates. The method developed below is more accurate and less susceptible to noise than attempting to directly estimate the location within the flow field at which the velocity is zero, as it uses data from the whole field to determine the saddle-point location.

In chapter 3, a method was developed that fitted a modified conic section to the velocity magnitude of the flow field. This approach, while accurate and effective, involves the use of a non-linear solver and requires a number of iterations to achieve convergence. This was found to be too slow, and the time required for the solver to converge was variable. Therefore, this approach was inappropriate for a realtime system where the timing budget for this computation is 5 ms. While PIV results are only produced every 10 ms due to double-pulsed illumination, for good performance of the control system the delays incurred due to post-processing of the data should be minimised. Ideally, the analysis should be completed before the next frame. A simplification of this approach has been developed that uses the observation that the flow near to the saddle-point can be approximated by a pair of linear equations

$$\hat{\mathbf{v}} = \begin{pmatrix} u \\ v \end{pmatrix} = \alpha \begin{pmatrix} -[x - x_c] \\ +[y - y_c] \end{pmatrix}, \quad (4.7)$$

where (x_c, y_c) is a vector representing the saddle-point position, and α an arbitrary constant that is proportional to the average flow velocity. These equations can then be fitted to the measured velocities to determine (x_c, y_c) . As this system is a linear system of equations, the solution can be obtained via least squares, which is computationally very efficient as it only requires the computation of the singular value decomposition (SVD) of the system matrix (along with a few other matrix operations).

The suitability of (4.7) can be confirmed by plotting the velocity error $(\hat{\mathbf{v}} - \mathbf{v})$ of the analytical solution $\hat{\mathbf{v}}$ as compared to a vector field from a CFD solver \mathbf{v} . A contour plot of velocity error, expressed as a percentage, over the imageable area of the device is presented in figure 4.8. The mean error is 9.25% over the whole field and 4.5% close to the saddle-point. These error computations

4. Flow Measurement

do not imply that the saddle-point position (x_c, y_c) , computed by fitting these equations, has an associated error in the order of 10%. If the errors are symmetric and on opposite sides of the fitted line, the slope and R^2 value may be affected but there will be a minimal effect on the x -intercept of the two curves. For extremes of x_c and y_c , these errors will certainly not be symmetric within the imageable area, and more significant errors in (x_c, y_c) are to be expected. However, the accuracy of the saddle-point location is most critical for small values of (x_c, y_c) , where the linear approximation holds.

4.3.1. Error and Accuracy Analysis

In order to quantify the position error and accuracy characteristics of the saddle-point identification algorithm, a parameter study was conducted along similar lines to that presented in chapter 3.

Accuracy Requirements The accuracy of the saddle-point estimate will have a direct effect on the stability of the control system and its ability to trap a cell. Additionally, the stability of the system is more sensitive to the absolute accuracy of the saddle-point estimate when the cell is close to the centre of the imageable area (which, for this discussion, is assumed to be the control system's set-point, or goal). This can be verified by comparing the required response of a controller to trap a cell for two broad regions — one close to the centre of the imageable area, and one close to the edge of the imageable area. When the cell is far from the centre of the imageable area, the controller will demand a saddle-point position that is significantly further than the cell position, in order to rapidly accelerate the cell back towards the centre. As the cell moves closer to the centre, the controller will demand small changes in saddle-point position to maintain the cell position near the centre. A large absolute error, or uncertainty, in the position of the saddle-point close to the centre of the channel will result in the cell moving faster than the controller has been tuned for, or even in the wrong direction. Conversely, far from the centre, the same absolute error will be small, relative to the demand position, and have a lesser effect on the system stability.

The effect on the saddle-point error, modelled as Gaussian noise in the saddle-point position, has been simulated using the control simulator presented in chapter 3. Absolute saddle-point position accuracies were simulated for 0 px (ideal), ± 25 px, ± 35 px, ± 45 px and ± 55 px. Cell positions over time are plotted for each case in figure 4.9. The RMS of the cell position during trapping and the overshoot during capture are shown in table 4.3 as measures of the effect of accuracy on the 'near' and 'far' stability, respectively. As the other parameters in the simulation were chosen to be representative but arbitrary, the relative changes in RMS error and overshoot are of interest, rather than the absolute values. In terms of overshoot, accuracies of $\leq \pm 45$ px result in only a moderate increase in overshoot ($\leq 9\%$), whereas RMS error during trapping is affected significantly by saddle-point estimation accuracy, as predicted (30% at ± 25 px).

To aid in determining accuracy specifications, the near region n and far region f have been defined in terms of the imageable area A_I as follows

$$n \in \begin{pmatrix} -160 \\ -160 \end{pmatrix} \leq \begin{pmatrix} x \\ y \end{pmatrix} \leq \begin{pmatrix} 160 \\ 160 \end{pmatrix} \quad (4.8)$$

$$f : A_I \setminus n, \quad (4.9)$$

where (x, y) is a vector in pixels from the centre of the image, which corresponds to a square region that covers exactly half of the horizontal imageable area. Hence, the required accuracy can be defined in terms of the near region n and far region f (see table 4.4).




Figure 4.8: Contour plot of the percentage velocity error of the analytical field estimate (4.7) as compared to a solution from a CFD solver. For much of the imageable area, the velocity error is less than 5%, demonstrating that the analytical approximation produces a result close to the original. Errors increase further from the centre as the approximation does not account for the parabolic shape of the inlet and outlet velocities (Poiseuille flow).

4. Flow Measurement

Accuracy	RMS error	Increase in RMS error (<i>c.f.</i> ideal)	Over-shoot	Increase in overshoot (<i>c.f.</i> ideal)
Ideal	20.0 px	–	218.3 px	–
± 25 px	25.7 px	29 %	213.9 px	– 2 %
± 35 px	29.6 px	48 %	227.9 px	4 %
± 45 px	57.4 px	187 %	238.7 px	9 %
± 55 px	52.5 px	162 %	299.7 px	37 %

Table 4.3: Effect of saddle-point position accuracy on system stability.




Figure 4.9: Cell position (in the vertical, or outlet plane) over time for a number of saddle-point position accuracies.

Region	Accuracy
Near	$< \pm 25$ px
Far	$< \pm 45$ px

Table 4.4: Saddle-Point position accuracy specifications for control stability during capture and trapping.

Method Flow fields were computed using the CFD solver VIPER^{21,22} for a number of flow rate deviations (*bias*). The position of the saddle-point for each flow field was then confirmed using the conic-fit algorithm described in 3.1.5. These flow fields were then used as the basis for synthetic image generation. Parameters used for the synthetic image generator were the same as case d in section 4.2.4, with the exception that 100 image-pairs were generated to allow for a better estimate of standard deviation for each measurement point.

Once the synthetic images were generated, the images were analysed using the GPU PIV code with a window size of 64×64 px and spacing of 24 px (these parameters were chosen for consistency with the results presented in section 4.4). Finally, the velocity fields were analysed using the algorithm described above to determine saddle-point location.

Results In general, there is good agreement between the actual saddle-point position and the estimated position. In the horizontal (x) plane, the saddle-point estimate was linear and showed a small amount of underestimation (figure 4.10a). The slope of the line of best fit between x_c and \hat{x}_c was 0.97. Non-linearity error (deviation from the line of best fit, figure 4.10c) increases as the saddle-point is moved further away from the centre of the imageable area. This is to be expected as errors in the fit of the estimate to the data are likely to cancel (equal and opposite around the saddle-point) where the saddle-point is close the centre, and become more pronounced as the saddle-point is moved further away. In the vertical (y) plane, these non-linearity errors are more pronounced (figure 4.10b,d). This is due to the more limited field of view in the vertical direction, resulting in fewer velocity vectors to fit the saddle-point estimate. With fewer vectors over which the estimate error is effectively averaged by the process of least-squares fitting, the vertical direction is more sensitive to the non-linearity of the flow.

Overall, the RMS error over the near region in the horizontal plane is 2.4 px with a maximum non-linearity error of ± 5 px, similarly in the vertical plane, the RMS error was 10.8 px with a maximum non-linearity error of ± 5 px. Over the entire imageable area, the non-linearity error increases to ± 10 px and ± 40 px, respectively. In all cases, these specifications exceed those in table 4.4.

The non-linearity errors are highly dependent on the specific channel geometry used. For example, for a cross geometry with a larger radius ($R = 1000 \mu\text{m}$, rather than $R = 250 \mu\text{m}$), non-linearity and RMS error in the near region reduce to (1.2 px, 5.1 px) and ± 3 px, respectively. Over the entire imageable area the non-linearity is ± 5 px and ± 7 px in the x and y directions, respectively. This is because a much smaller area of the geometry is within the imageable area, relative to the flow geometry, as the radius increases. As can be seen in figure 4.8, as the imageable area decreases in size relative to the flow geometry, the error of the analytical field estimate decreases, improving the accuracy of the saddle-point identification algorithm.

4.4. Parameter Optimisation

Accuracy of PIV, and therefore of the saddle-point position, is closely related to the window size (which affects the velocity range that can be measured) and vector spacing (which affects the spatial resolution of the final result). Vector spacing \vec{s} should be no greater than the window size \vec{w} (i.e. $\vec{s} \leq \vec{w}$) as vector spacings greater than the window size leave pixel regions which do not contribute to the final result (under-sampling). Full coverage of the input image is achieved when $\vec{s} = \vec{w}$. As spacing decreases, the number of windows to be analysed increases, hence the computational burden of computing the vector field increases. Window size also has an effect on computation time, as it takes longer to compute to FFT of larger windows (typically of the order $\mathcal{O}(\log N^2)$, where N is the window size in each direction⁸).

4. Flow Measurement

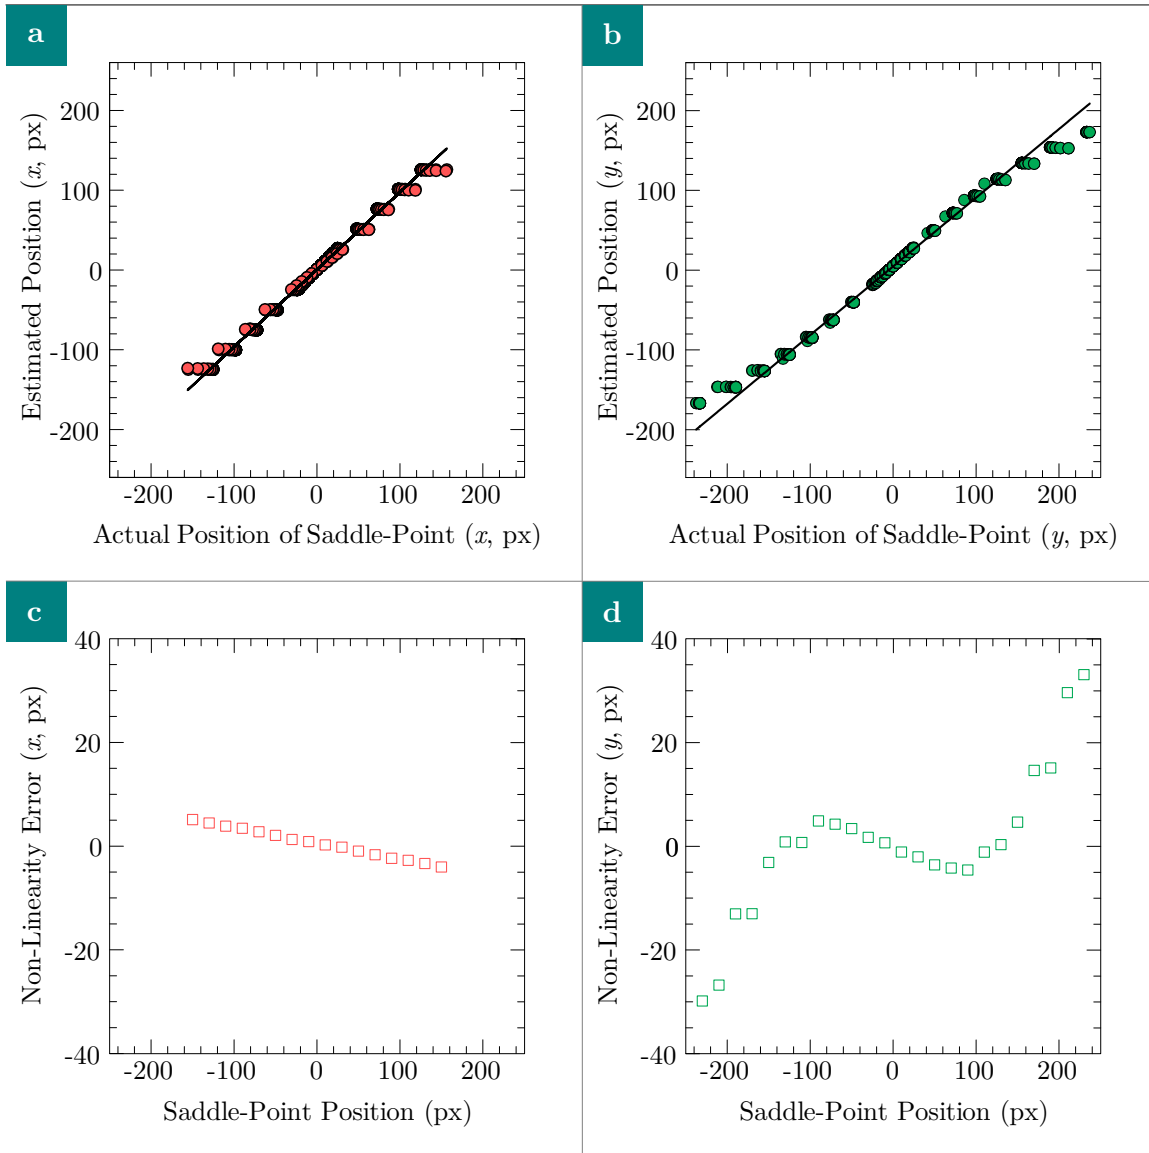


Figure 4.10: Parameter variation study on the linearity and accuracy of the saddle-point estimation algorithm presented in section §4.3. Dashed lines in c) and d) indicate the region where the non-linearity error is ± 5 px.

In order to quantify the effects of these trade-offs in terms of the saddle-point positional error, a number of parameters were compared using data obtained experimentally. 4000 images were recorded and analysed at a number of window sizes and spacings. The average seeding density was $0.001 \text{ particles/px}$. Vector spacings between 25 % and 75 % of the window size were chosen, as these often give optimal results.^{19,28} To quantify the measurement error, the saddle-point position was assumed to be steady, at a small but unknown offset from $(0, 0)$. Hence the standard deviation of the saddle point position $\|(\sigma_x, \sigma_y)\|$ was used as an indicator of measurement error. Figure 4.11 shows the results of this study for a number of window sizes, along with the time to compute a frame (this time includes other computations performed by the control software, and therefore is greater than the results in section 4.2.5 might imply). Of the window sizes analysed, only 64×24 , 64×32 , 128×64 and 32×16 meet the computation time criteria ($t \leq 5 \text{ ms}$) and, of those, a window size and spacing of 64×24 results in the lowest standard deviation in the saddle-point position.

4.5. Adapting to Flow Velocity Changes

Extensional flows are difficult to measure accurately with PIV due to the large velocity gradients present in the image. Close to the saddle-point, velocities are small (close to zero), which would imply that the optimal window size should be as small as practical in order to maximise resolution.²⁸ Further from the saddle-point, velocities increase rapidly, and are much larger than those close to the saddle-point. This would imply a much larger optimal window size to ensure that the displacements are smaller than $1/4$ of the window size. Displacements larger than this lead to increased measurement error.²⁶ These two requirements are in conflict and cannot be resolved without modifications to the measurement technique. In addition, the average velocity can increase dramatically as a result of shifts in the saddle-point location (a location that is close to the edge of the imageable area will result in a much larger average velocity) or changes to the overall flow rates that might take place for calibration, purging or shear rate changes during trapping.

4.5.1. Conventional Approaches

The problem of varying velocity scales, both spatially and temporally, has been well-studied in the literature, and is of particular interest to researchers in the field of turbulence, amongst many others. Brief descriptions of two of the more common techniques follow.

Irregular Grids Rather than enforcing a single window size across the image, the image is divided up into spatial regions based on the predominant velocity length scales. Areas with lower velocities are analysed with smaller windows, whereas for high velocities, much larger windows are chosen to accurately capture the motion.

This approach introduces a number of challenges in this application. Firstly, if the grid spacing is held constant, the computation time becomes variable and less predictable as the size of the analysis windows, and therefore the computational overhead, vary. Alternatively, the number and sizes of the analysis windows could be held constant and the grid spacing (i.e., the mesh) could be varied. This introduces complexity as the software must handle these irregular grid spacings efficiently (*in extremis*, this requires a structured auxiliary mesh approach such as that discussed in chapter 3). Furthermore, the regions of high and low velocities will change as the saddle-point position changes. This introduces a potentially unstable feedback loop into the system, as the division of the image into these sub-regions cannot be determined *a priori*. The reliability and accuracy of the saddle-point position estimate relies on accurate velocity measurements from PIV, and the accuracy of the PIV is

4. Flow Measurement




Figure 4.11: Accuracy and Computational Performance of PIV and the saddle-point identification algorithm at various FFT Window Sizes and Vector Spacings. The dashed line at 5ms on the time plot represents the maximum allowable computation time per frame (given an imaging frame rate of 200 fps). Only 64×24 , 64×32 , 128×64 and 32×16 meet the computation time criteria and, of those, a window size and spacing of 64×24 results in the lowest standard deviation ($\sigma = 2.9$ px) in the saddle-point position.

dependent on the flow characteristics remaining steady in a given location. This may result in periods of time where the system becomes unstable while the PIV and saddle-point estimates oscillate.

Multigrid In this approach, the PIV algorithm is run successively at progressively smaller window sizes (normally in powers of two, e.g. 128, 64, 32). For each iteration, the centre of the window in frame 2 is displaced by the velocity estimate from the previous iteration. This effectively means that all PIV measurements (except for the first) are a measure of the error in the previous measurement. The velocity is the sum of the window displacement and the calculated velocity from that iteration. In this way, PIV at the smaller window sizes is only measuring a small change in velocity, increasing accuracy. This approach does not suffer from the potential oscillation of the previous method, but is significantly more computationally expensive, requiring the computation of several PIV measurements. As can be inferred from the timings in 4.11, there is not sufficient time in the timing budget per frame to allow for multiple PIV computations.

4.5.2. Timing Modulation Approach

In contrast with the conventional techniques described above, this approach takes advantage of the unique ability of the PIV measurement system developed in this thesis to adapt to system changes in realtime. Recall that, in a double-pulsed PIV setup, the displacement per frame is dependent on the inter-frame (“gap”) time, or Δt (section 4.1.4.1). As the timing is generated by an on-board card in the EPRC (Expandable Pressure Regulation Controller, see section §2.4), the Δt can be changed in realtime by sending a remote command to the unit. This enables an adaptive approach to PIV timing, Adaptive Gap Time Control (AGTC), in which the Δt is varied based on characteristics of the flow.

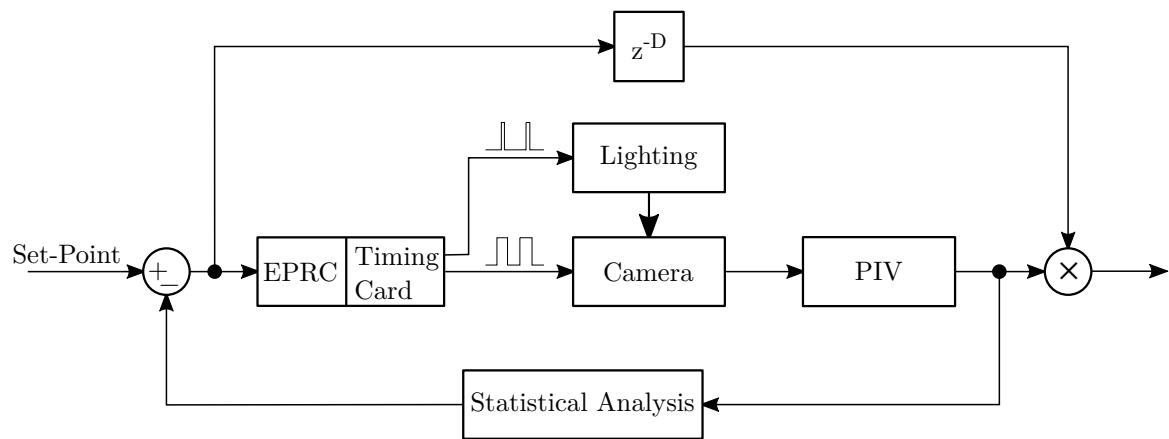


Figure 4.12: Adaptive Gap Time Control block diagram.

Computationally, this is very efficient, as the PIV algorithm is unchanged, and only requires a small amount of additional logic to compute the required statistics from the PIV vector field and update the Δt . However, this approach will not be as accurate at resolving high velocity gradients as the more conventional approaches outlined in section 4.5.1. Instead, the feedback loop will converge to an optimal solution that minimises the velocity error for the entire field, given the constraints imposed by the realtime PIV algorithm.

4. Flow Measurement

With reference to figure 4.12, timing pulses are generated by dedicated hardware on the EPRC with a resolution of 200 ns. The camera trigger pulses and first lighting pulse have a consistent phase relationship from cycle-to-cycle. The timing of the rising edge of the second lighting pulse, which determines the effective Δt , is varied based on commands received from the control software. The lighting timing pulses affect the effective time of the second frame captured by the camera and quantified by the PIV algorithm, hence modifying the effective displacement per frame on a global basis. The vectors output from PIV are analysed (Statistical Analysis) to determine a fitness function that gives a measure of the fitness of the chosen Δt for the current flow field. This fitness function is then compared with the Set-Point to determine a new Δt for the next frame pair.

As the Δt is constantly being updated, the raw velocity vectors from PIV cannot be used directly by the rest of the system. Therefore the velocity is scaled by the Δt and image resolution d_{px} , i.e.

$$\vec{u} = \vec{u}_{raw} \frac{d_{px}}{\Delta t_{delayed}}.$$

As there is an inherent delay D in the control loop, both in the relatively slow EPRC and in process of image acquisition and PIV analysis, the Δt used in the final velocity scale factor is delayed by a calibrated number of frames ($\Delta t_{delayed}$).

Specific Implementation For the particular flows under investigation, a simple mean of the velocity magnitudes ($f = \frac{1}{N} \sum \|\vec{u}_{raw}\|$) was found to be an effective measure of fitness for a specific timing interval. Other measures, such as the signal-to-noise ratio of the cross-correlation peak or the number of noisy vectors (as measured by an algorithm such as the Westerweel and Scarano ‘Universal’ outlier rejection technique²⁷), could be used either in conjunction with, or instead of the mean velocity magnitude. As there is a small amount of jitter (± 1 frame) in the timing update via the EPRC, which translates into a glitch in the corrected velocity \vec{u} , care was taken to minimise the rate of updates to reduce the likelihood of glitches. An update is only performed if the following relation does not hold

$$\frac{3}{4} f_s \leq f \leq \frac{4}{3} f_s, \quad (4.10)$$

where f is the fitness function and f_s the set-point. An updated Δt is computed with the following recursive formula

$$\Delta t = \frac{f_s}{f} \Delta t_{delayed}. \quad (4.11)$$

This updated Δt is rate-limited to $500 \mu s/\text{step}$ to ensure an under-damped response. Finally, to prevent rapid limit cycling, the updated timing information is only sent to the EPRC at most every 100 ms (systemic delay $D \approx 80$ ms).

Sample velocity and Δt traces from the system are shown in figure 4.13. As velocity increases, the system progressively decreases Δt from an initial final of $2056 \mu s$ and reaches a final value of $363 \mu s$ in order to maintain an average raw displacement within the allowable bounds. Without AGTC, the average velocity $\|\vec{u}_{raw}\|$ would have reached $28 \text{ px}/\text{frame}$, well outside the allowable bounds for reliable PIV measurement.




Figure 4.13: Behaviour of AGTC during a flow rate ramp. The Δt is progressively decreased from $2056 \mu\text{s}$ to $363 \mu\text{s}$ in order to maintain the average raw displacement within the limits in equation (4.10) (dashed line). The raw velocity shows a saw-tooth pattern during the velocity ramp, whereas the compensated velocity $\|u\|$ recovers the true behaviour of the system.

4.6. Conclusion

Using a variety of optimisations and modern graphics processing hardware, a highly-efficient, realtime flow measurement method has been developed that will be used extensively in the following chapters (in particular chapter 5 and chapter 6). Validation studies demonstrate that the developed methods and associated code has comparable accuracy to other, conventional, codes. In addition, a realtime method for analysing the flow information has been presented that can extract the saddle-point position in constant time while maintaining excellent accuracy in the critical region close to the centre of the channel. Finally, a novel method has been developed for realtime optimisation of PIV imaging parameters that allows the system to rapidly adapt to large variations in flow velocities without loss of accuracy or large computational overhead.

4.7. Bibliography

- [1] Ronald J. Adrian. “Particle-imaging techniques for experimental fluid mechanics”. In: *Annu. Rev. Fluid Mech.* 23.1 (1991), pp. 261–304.
- [2] Bronkhorst High-Tech B.V. *LIQUI-FLOW mini Micro Fluidic Mass Flow Meters for liquids*. 2012. URL: http://www.bronkhorst.com/en/products/liquid_flow_meters___controllers/liqui-flow_mini/.
- [3] F. Champagnat et al. “Fast and accurate PIV computation using highly parallel iterative correlation maximization”. In: *Exp Fluids* 50.4 (17 Mar. 2011), pp. 1169–1182. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-011-1054-x.
- [4] Dolomite Microfluidics. *Mitos Sensor Units*. 2010. URL: http://www.dolomite-microfluidics.com/images/stories/PDFs/datasheets/mitos_sensor_units_product_datasheet.pdf.
- [5] Elveflow. *Microfluidic flow sensors*. 2009. URL: <http://www.elveflow.com/microfluidic-flow-control-products/microfluidic-flow-control-module/microfluidic-liquid-mass-flow-sensors/>.
- [6] Andreas Fouras, Jonathan Dusting and Kerry Hourigan. “A simple calibration technique for stereoscopic particle image velocimetry”. In: *Exp Fluids* 42.5 (28 Mar. 2007), pp. 799–810. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-007-0293-3.
- [7] Andreas Fouras, David Lo Jacono and Kerry Hourigan. “Target-free Stereo PIV: a novel technique with inherent error estimation and improved accuracy”. In: *Exp. Fluids* 44.2 (Oct. 2007), pp. 317–329. ISSN: 0723-4864. DOI: 10.1007/s00348-007-0404-1.
- [8] Matteo Frigo and Steven G. Johnson. *benchFFT software and results (part of fftw)*. URL: <http://www.fftw.org/benchfft/>.
- [9] Bénédicte Galmiche et al. “Turbulence characterization of a high-pressure high-temperature fan-stirred combustion vessel using LDV, PIV and TR-PIV measurements”. In: *Exp Fluids* 55.1 (8 Dec. 2013), pp. 1–20. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-013-1636-x.
- [10] R. A. Jamison et al. “X-ray Velocimetry and haemodynamic forces within a stenosed aortic model at physiological flow rates”. In: *Ann. Biomed. Eng.* (2011).
- [11] Ralph Lindken et al. “Micro-Particle Image Velocimetry (PIV): Recent developments, applications, and guidelines”. In: *Lab Chip* 9.17 (2009), p. 2551. ISSN: 1473-0197. DOI: 10.1039/b906558j.

- [12] Warwick S. Nesbitt et al. “A shear gradient-dependent platelet aggregation mechanism drives thrombus formation”. In: *Nat Med* 15.6 (June 2009), pp. 665–673. ISSN: 1078-8956. DOI: 10.1038/nm.1955.
- [13] Nathalie Nève et al. “Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers”. In: *Cel. Mol. Bioeng.* 3.3 (Mar. 2010), pp. 213–228. ISSN: 1865-5025. DOI: 10.1007/s12195-010-0113-3.
- [14] NVIDIA Corporation. *CUDA Toolkit — Best Practices Guide*. URL: <http://docs.nvidia.com/cuda/cuda-c-best-practices-guide/>.
- [15] NVIDIA Corporation. *CUDA Toolkit — cuFFT*. URL: <http://docs.nvidia.com/cuda/cufft/>.
- [16] M. Raffel, C. E. Willert and J. Kompenhans. *Particle Image Velocimetry: A Practical Guide*. 3rd Edition. Springer-Verlag, 1998. 253 pp. ISBN: 3-540-63683-8.
- [17] M. Raffel et al. “Recording and evaluation methods of PIV investigations on a helicopter rotor model”. In: *Exp Fluids* 40.4 (1 Mar. 2006), pp. 665–665. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-006-0120-2.
- [18] T. Roesgen. “Optimal subpixel interpolation in particle image velocimetry”. In: *Exp Fluids* 35.3 (26 July 2003), pp. 252–256. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-003-0627-8.
- [19] Chaminda R. Samarage et al. “Polynomial element velocimetry (PEV): a technique for continuous in-plane velocity and velocity gradient measurements for low Reynolds number flows”. In: *Meas. Sci. Technol.* 23.10 (1 Oct. 2012), p. 105304. ISSN: 0957-0233. DOI: 10.1088/0957-0233/23/10/105304.
- [20] Thomas Schiwietz and Rüdiger Westermann. “GPU-PIV.” In: *VMV*. 2004, pp. 151–158.
- [21] G. J. Sheard and K. Ryan. “Pressure-Driven Flow Past Spheres Moving in a Circular Tube”. In: *J. Fluid Mech.* 592.-1 (2007), pp. 233–262. DOI: 10.1017/S0022112007008543.
- [22] G. J. Sheard et al. “Flow around an impulsively arrested circular cylinder”. In: *Phys. Fluids* 19.8 (2007), p. 083601. ISSN: 10706631. DOI: 10.1063/1.2754346.
- [23] Shuhei Tarashima, Satoshi Someya and Koji Okamoto. *Acceleration of Recursive Cross-Correlation Piv Using Multiple Gpus*. WOS:000319843000145. New York: Amer Soc Mechanical Engineers, 2011. ISBN: 978-0-7918-5454-9.
- [24] Shuhei Tarashima et al. “GPU accelerated direct cross-correlation PIV with window deformation”. In: 15th International Symposium on Applications of Laser Techniques to Fluid Mechanics. Lisbon, Portugal, 2010.
- [25] Francisco Javier Tovar-Lopez et al. “A microfluidics device to monitor platelet aggregation dynamics in response to strain rate micro-gradients in flowing blood”. In: *Lab Chip* 10.3 (2010), pp. 291–302.
- [26] J. Westerweel, D. Dabiri and M. Gharib. “The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings”. In: *Experiments in Fluids* 23.1 (1 May 1997), pp. 20–28. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s003480050082.
- [27] Jerry Westerweel and Fulvio Scarano. “Universal outlier detection for PIV data”. In: *Exp Fluids* 39.6 (Aug. 2005), pp. 1096–1100. ISSN: 0723-4864. DOI: 10.1007/s00348-005-0016-6.
- [28] C. E. Willert and M. Gharib. “Digital particle image velocimetry”. In: *Exp. Fluids* 10.4 (1 Jan. 1991), pp. 181–193. DOI: 10.1007/BF00190388.

5. Flow Control

An active microfluidic device needs an accurate, high-speed method of controlling flow into and out of the device. As the system has very little reliance on the characteristics of the chosen geometry, desired behaviours can only be imposed upon the system by regulating the flow of fluid into and out of the chip's various ports (hereinafter referred to as the chip's boundaries). In the design and construction of this system, a number of possible flow control devices were considered, which have been broken down into two broad categories: pumps and valves. As the intended application included manipulation of biological cells, particular attention has been given to bio-compatibility aspects and the ability to clean and sterilise, or completely dispose of, the fluid-contacting components.

5.1. Existing Designs

One method of changing the flow at a system boundary is to directly affect the transport of fluid into or out of the system (i.e. a pump). This has the advantage of providing very direct control. Alternatively, one can locally restrict the geometry of a system, hence increasing the resistance to flow and thereby reducing the flow rate (i.e. a valve). Here, a number of common flow control devices used in microfluidics are briefly reviewed. This is by no means an exhaustive list, but rather a list of technologies that were considered by the author for the purposes of this thesis.

5.1.1. Syringe Pump

A syringe pump is one of the simplest pumps and is favoured in the medical industry due to its simplicity and inherent hygienic properties (the fluid only contacts a disposable or sterilisable syringe). A linear actuator - often a stepper motor combined with a lead screw to give linear motion - drives the syringe, displacing the fluid contained within and therefore resulting in a pumping action. The primary advantage of this system is that it gives a very direct relationship between pump control input and rate and volume of fluid delivered. However, at low flow rates, such as those commonly encountered in microfluidic devices, system friction begins to overcome the imposed motion, resulting in a 'stick-slip' motion,^{7,11} whereby the syringe remains stationary ('stick') until the imposed displacement results in a sufficient force to overcome static friction, at which point the syringe barrel slides forward to the new position ('slip'). Additionally, stepper motor syringe pumps (which dominate the market) suffer from transmitted vibration due to small discontinuities at low flow rates, introduced by the discrete distances moved in each microstep. This manifests itself as pulsatility in the flow, which can be strongly related to syringe pump parameters such as flow rate and syringe diameter.^{13,14,21} Additionally, the syringe pump is a system with substantial rotational and linear inertia,[†] which severely limits the maximum flow control bandwidth - that is, the maximum rate of change of the system flow rate.²³ In Zhou et al.,²³ it was reported that the minimum flow rate rise time for a syringe pump with a glass syringe was $\tau_{0.5} = 2.1$ s into a PDMS microchannel. The minimum achievable rise time was $\tau_{0.5} = 0.3$ s when coupled with a PMMA (rigid) microchannel.

[†]There are some designs with specialised linear motors that only have linear inertia. However, they are rare, and nevertheless this does not affect the general principle detailed here.

5. Flow Control

5.1.2. Pressure Pumps

Pressure pumps can be thought of as consisting of two pumps: an upstream pump that supplies a secondary fluid (often air); which then exerts a pressure on the primary fluid. Providing that there is a pressure difference between one of the system boundaries and another, fluid will flow in the direction of the pressure gradient. For low flow rates and simple geometries, the flow can simply be predicted via $\dot{Q} = -\Delta P/R$ (where ΔP is the pressure drop and R the bulk resistance). With prudent system design, noise inherent to the upstream pump is mechanically filtered by the compressible secondary fluid, resulting in a very stable and consistent flow.¹⁰ However, this filtering does come at a disadvantage: any changes to flow or volume in the upstream pump will be bandwidth-limited by that filter.

5.1.3. On-Chip Multilayer Control Valves

These devices are constructed in a multilayer fashion, consisting of a conventional microfluidic fluids layer sandwiched on top of a gas control layer.²² At certain points, the lower channel (the fluid layer) intersects with a channel in the control layer, resulting in a thin, deformable, subsection which acts as a diaphragm valve. As gas pressure is increased, the force on the thin section results in deformation into the fluid layer, increasing system resistance and eventually completely occluding the fluid channel.

These devices have a number of attractive properties. Firstly, they have the same bio-compatibility properties as the chip itself, as the fluid only comes into contact with the microfluidic chip. Secondly, as they are directly integrated into the chip, the external hardware is limited to an external gas pressure controller.²² An interesting technique uses the valve membrane as a lens to provide positional feedback in realtime. However, it is unclear how sensitive this technique might be to dirt and optically dispersive fluids.

The primary disadvantage of this type of device for the present application is that the multilayer process dramatically increases the complexity, and therefore cost, of device manufacture. Secondly, the dynamics of this type of valve, as shown by a recent study,¹² are highly asymmetric, with the devices produced in the paper requiring 12 ms to partially close, and 16 ms to fully close, yet took 40 ms to return to an open state. The dynamics in the closing direction are attractive, however the slow (and asymmetric) relaxation time is a concern. Furthermore, the relationship between pressure and channel occlusion is highly non-linear, implying that calibration would be necessary to linearise the system.

5.1.4. Off-Chip Diaphragm Valves

An alternative approach to an on-chip multilayer control valve is to place the diaphragm off-chip, which allows for the use of commercial off-the-shelf diaphragm valves, and significantly increases the choices available in terms of actuation (for example, direct methods using solenoids and voice-coils can be considered). These devices have the advantage of mass-production and have well-defined flow characteristics.^{3,9,16,20} However, the diaphragms of these valves are not intended to be frequently replaced or cleaned, and are a significant part of the cost of the valve. Additionally, commercial diaphragm valves commonly have a relatively large dead-volume (a fluid-contacting volume that does not form part of the valve action) and the diaphragm mechanism contacts the fluid directly. All of these features (dead-space, complex geometries, components sliding past one another) increase the risk of fouling in biological settings due to the increased surface area of the diaphragm material, large contact areas, and potential for lysis at the points of inter-component contact.

5.1.5. Off-Chip Pinch Valves

Pinch Valves operate by using the system tubing as part of the valve itself. The tubing is clamped into the valve and an actuator impinges upon the tubing, thereby modifying its shape. Through this process, the effective pressure drop across the clamped section is increased and the flow through the tube can be reduced. The advantages of this approach are that the system bio-compatibility is not affected, and the tubing can be easily replaced if it becomes fouled. In addition, there are no additional interconnections to be made - points where dirt and air can accumulate - and the dead-volume is effectively zero. However, the characteristics of the valve becomes heavily dependent on the tubing characteristics, which is not normally designed as a valve membrane. This imposes additional challenges on the control of the system, especially when full proportional control of the flow rates are required.

5.1.6. Summary

The important characteristics of the flow control methods considered are summarised in table 5.1. The comparison focuses on two general areas: Consumable Cost and Bio-compatibility; and Control Behaviour (Static and Dynamic). Consumable Costs and Bio-compatibility are related - poor bio-compatibility will increase fouling, requiring increased consumable replacement or cleaning. Even with excellent bio-compatibility, some fouling is inevitable and cross-sample contamination concerns will require frequent consumable replacement. Microfluidic chips in particular, as a result of their very small channel dimensions and limited external access, are extremely prone to fouling that may be impossible or impractical to clean. Control behaviour has direct relevance to the system: the ideal flow controller would be predictable, repeatable and have bandwidth that exceeds that of the other system components. This is rarely achievable, but to maximise the capability of the system to modify flow behaviour in realtime, dynamic behaviour should be prioritised.

5.2. Design and Development

A detailed review of the options presented in table 5.1 led to the decision to use custom-designed pinch valves paired with a pressure pump for the further development of the technology presented in this thesis. The choice of a custom design may seem surprising, but a review of commercially-available designs found that most were oriented at simple on/off operation, and as such did not have the precision or sufficiently rigid tubing retention designs to support repeatable linear operation. In addition, all were paired with simple solenoid actuators that did not lend themselves to proportional operation.

To simplify further discussion, some of the design decisions were made on the basis of the dimensions of the tubing selected. This was the smallest readily-available silicone tubing that could be obtained commercially, which had an inner diameter of $800\ \mu\text{m}$ and corresponding outer diameter of $1600\ \mu\text{m}$ (effective wall thickness of $400\ \mu\text{m}$).

5.2.1. Actuator Type

A number of actuators were considered. Briefly, these were: lead-screw driven linear actuators; piezo-electric motors; and voice coil actuators. **Linear actuators** are a common solution to this type of problem and are used to drive a number of industrial diaphragm valves. While they can be extremely accurate, their dynamic performance is more limited. To traverse the distance from fully-open to fully-closed (in this case $800\ \mu\text{m}$), the motor must rotate through a certain number of full revolutions.

Technology	Consumable Cost	Bio-compatibility	Control Behaviour (Static)	Control Behaviour (Dynamic)
Syringe Pump	Low (depends on syringe quality)	Good	Direct. Syringe displacement directly corresponds to volume delivered.	Limited by relatively large mass of the motor and syringe barrel. Lower quality syringes exhibit much more pronounced 'stick-slip' behaviour.
Pressure Pump	Low	Good, although fluid exposed to potentially non-sterile gas.	Change of pressure results in a change of flow (dependent on the bulk flow resistance of the system).	Limited by rate of change of secondary pump and mechanical filtering due to compressibility of gas.
On-Chip Valves	High (complex manufacture)	Good	Relationship between fluid pressure drop and control gas input likely to vary between individual chips.	Depends on upstream gas pressure controller. Limited by compressibility of control gas.
Diaphragm Valves	May be high depending on diaphragm fouling	Poor	Pressure drop can be predicted from manufacturer's published data.	Lightweight diaphragm and common actuator characteristics lead to good dynamic behaviour.
Pinch Valves	Low (tubing)	Good	Relationship between fluid pressure drop and applied force or tubing displacement heavily dependent on tubing properties, and may vary considerably between tubing samples.	With flexible tubing and suitable actuator, potentially very good.

Table 5.1.: Comparison of flow control technologies.

	Linear actuator	Piezoelectric actuator	Voice coil actuator
	PhysikInstrumente MP-15	PhysikInstrumente U-264	MotiCont LVCM-025-029
Maximum Velocity	0.3 mms ⁻¹	250 mms ⁻¹	N/A
Position Resolution	0.2 μm	0.1 μm	
Maximum Displacement	6 mm	> 50 mm	16 mm
Time to travel 800 μm	> 2.5 s	< 3.2 ms	< 3 ms [‡]
Size	∅17 mm × 60 mm	57 × 63 × 10 mm	∅25 mm × 35 mm
Cost	~ \$2000	~ \$10 000	US\$250

Table 5.2: Small survey of commercial actuators of various types. The 800 μm distance used as a transit time benchmark is based on the inner diameter of the selected tubing.

As resolution is increased - that is the lead screw pitch is decreased, the number of revolutions required to move between fully-open and fully-closed increases. Hence there is an inherent trade-off between resolution and dynamic performance. **Piezoelectric actuators** are a relatively new technology that offer many advantages over traditional linear actuators. Specifically, piezoelectric actuators can achieve extremely high velocities and have very little moving mass, so can accelerate very quickly. Due to the complexity and relative infancy of this technology, piezoelectric actuators are significantly more expensive than their motor-driven counterparts. **Voice coil actuators** are used extensively in the disk drive industry, where they are used to drive the armature that contains the read/write heads. The advantages of these actuators are that they are very simple to produce - simply consisting of an electromagnet paired closely with a fixed magnet. As the current in the coil increases, the magnetic force increases, driving the fixed- and electro-magnets apart. These actuators can be made with very small moving masses and are suitable for applications where the range of motion is relatively small.

A comparison of the performance characteristics of some representative commercial actuators of each type is shown in table 5.2. In terms of both performance in this application, and unit cost, the voice coil actuator is clearly the most appropriate choice.

5.2.2. Challenges

Several potential problems with the design were identified and are discussed in the sub-sections to follow. In §5.4 and section §5.5, these challenges will be elaborated upon and solutions discussed.

[‡]Based on parameters (moving mass, effect of gravity) of final system, ignoring the effects of the tube.

5. Flow Control

5.2.2.1. Relationship between Control Input and Flow

In a pinch valve, there is no universally-applicable relationship between the control input (regardless of whether we define this as applied force or total tubing displacement) and pressure-drop across the valve. The primary reason for this is that the tubing behaviour is not tightly controlled as it would be in a commercially-built diaphragm valve. Variances in the manufacturing process may mean that the tubing wall thickness may vary slightly between batches of tubing or even along the length of a single piece of tubing. Heating from the valve actuator may result in changes to the shear modulus, or other mechanical properties of the tubing material (commonly silicone).¹⁹ Finally, a trade-off must be made between tubing flexibility (which would demand the thinnest wall thickness and a low durometer silicone - reducing the actuation force required); and tubing robustness (which would demand thicker walls and higher silicone durometers).

In common with all control valves considered, pinch valves do not actually change the flow rate. Instead, they change the pressure drop across the valve, which has a corresponding effect on the flow rate in their local branch of the fluid system. In order to build an accurate flow controller, the flow rates in the system must be continuously measured and the pressure drop adjusted in order to maintain the desired flow rate.

5.2.2.2. Trapped Air and Fouling

Air is very difficult to eliminate from microfluidic systems. Common chip-to-tubing interfaces provide ideal locations for air to accumulate in areas of low or no flow. Dissolved gases in the working fluid often lead to the formation of bubbles, even when there is no other route for air ingress. Predominant chip substrate materials such as PDMS are gas permeable, allowing air to diffuse through the substrate. In addition, the low flow rates used in microfluidic systems lead to very slow transport of gasses out of the system (in fact, surface tension often results in fluid flowing around a bubble rather than convecting it downstream).

The presence of air is particularly harmful to dynamic behaviour. As the upstream pressure is rapidly changed (due to an adjustment in valve position), the air acts as a spring, first compressing as a result of the increased pressure drop, then slowly expanding to restore the original shape of the bubble. As shown in figure 5.1, where the system was carefully purged of air before the experiment began, even small amounts of trapped air can significantly impair dynamic performance.

Fouling has an effect of the repeatability of the system over time. As dirt (for example lysed or clotted cells) accumulates in the system over time, the pressure drop increases, this leads to a change effective resistance. This can modify the relationship between valve pressure drop and system flow rate, which will cause an undesirable change in the system behaviour over time.

5.2.2.3. Hysteresis

Hysteresis, the dependence of the system output to both the input and history of past inputs, is a behaviour common to many systems. With respect to pinch valves, it means that the force required to achieve a certain pressure drop is different when the pressure drop is being increased compared to when the pressure drop is decreased. Hysteresis has been well-studied in the context of collapsible tubes,⁴⁻⁶ a field highly analogous to the deformation of silicone tubes via an external actuator. A number of studies^{4,5,8} have confirmed that silicone tubes have non-linear, and hysteretic flow-deformation characteristics. The findings presented in Bertram⁵ are of particular relevance to the application discussed in the sections to follow. In that study, detailed relationships between the applied external force (expressed in terms of pressure) and area (directly proportional to flow resistance, or flow rate




Figure 5.1: Air trapped in the system causes a severe lag in response. As the valve closes (increasing excitation), the saddle-point position increases (moves towards the inlet corresponding to the closing valve), as expected. However, the saddle-point position reaches an asymptote at $s_x = 50 \mu\text{m}$. Counter-intuitively, the saddle-point position then begins to increase again, even as the valve excitation voltage begins to decrease (shaded region). The system rapidly returns to steady state at $t = 62\text{s}$. This behaviour, where the response severely lags the input and additionally exhibits a stick-slip response ($t = 38\text{s}$) is a result of compression of air bubbles trapped in the flow loop, followed by sudden changes in response as the bubble shifts in the channel or reaches its maximum compression.




Figure 5.2: Normalised area of a deformable silicone tube as a function of applied external pressure (reproduced from Figure 4b in Bertram.⁵ Used with permission.). Significant hysteresis is observed in the central region where the area is changing most rapidly.

for a fixed pressure source) were developed. One result from Bertram is shown in figure 5.2 (note: the y -axis is the independent variable in this figure, and is therefore rotated compared to the similar figures presented in the present work), which exhibits similar pressure-area relationships to the excitation-flow rate relationships shown in the latter sections of this chapter.

In the remainder of this thesis, hysteresis will be discussed primarily in terms the positive-going and negative-going curves, and the region between them (the hysteretic region). This is a convenient model from a control perspective as it allows a direct mapping between the expected response and the current system input. These terms are described in the simplified diagram in figure 5.3.

5.2.2.4. Pumping

Pumping is a phenomenon that affects both pinch- and diaphragm-valves. As the overall flow rate is decreased, the small volume of fluid displaced by the valve as it changes position begins to have a significant effect on the flow through that branch. This is most easily explained in the simple 1D diagram (figure 5.4). For the purposes of this discussion, the worst case, where the valve moves from fully open to fully closed, will be considered. However, the principle is the same for all displacement changes. As the valve closes, the fluid in the shaded section is forced downstream. The resultant flow can be approximated by $\dot{Q}_{pumping} \approx \frac{\Delta V}{t_{close}}$, where ΔV is the displaced volume and t_{close} is the time required to close the valve. As the system flow rate decreases, $\dot{Q}_{pumping} \gg \dot{Q}$. Furthermore, this flow is symmetric and away from the valve in both directions, which means that flow in the affected branch may temporarily reverse against the system flow. This places an upper limit on the dynamic performance of the valve, and/or a lower limit on the system flow rate.




Figure 5.3: Generic model of a hysteretic system, showing the essential elements: a positive-going curve, traversed when the system input is increasing and a negative-going curve, traversed when the system input is decreasing. The area between the two curves, known in this text as the hysteretic region, must be traversed whenever the system input changes direction. In this simplified model, no change in the system response occurs while the system input is within the hysteretic region.




Figure 5.4: Peristaltic pumping effect as a result of rapid changes in valve actuator position. As the valve closes, the tubing wall is displaced by a certain amount. This results in the displacement of a corresponding volume of fluid (shaded region). Depending on the speed of valve closure, this may result in a significant, transient, change in flow rate (see section 5.2.2.4).




Figure 5.5: CAD renderings of the final pinch valve design; a) shows an exterior isometric view of the valve with bayonet cap removed and tubing inserted, b) shows a cross-section of the valve and its major components.

5.2.3. Final Design

The final valve is designed around a series of interlocking cylinder sections, as can be seen in figure 5.5b. The advantages of this design are that the various components are self-aligning along the central axis; the components can be almost exclusively machined on a lathe, simplifying manufacture. A bayonet-cap design has been employed to secure the tubing into the valve. This allows for easy and quick removal and replacement of tubing in the valve, while securing the tubing during normal operation.

5.2.3.1. Bayonet Cap Design

In the design of the bayonet cap, a number of conflicting requirements had to be overcome. In particular: the cap must rotate freely in order to facilitate easy removal and installation, yet the cap must not rock or move during operation; the cap must tightly press against the tubing to prevent undesirable movement as the actuator head is driven into the tubing, but the cap must not be so tight as to rip or crimp the tubing as it is rotated into place. In order to fulfil these requirements,




Figure 5.6: Exploded view of the Bayonet Cap, showing the major components.

a spring button was added to the cap, shown fully extended in figure 5.5a and partially retracted in figure 5.5b. A labelled exploded view of the cap is shown in figure 5.6. The spring button is free to rotate relative to the bayonet cap body and is forced to its maximum extension by the preloaded spring. As the bayonet cap is inserted onto the valve body, the spring button contacts the valve body first, pressing the tubing flush with the top of the valve body. The large metal-on-metal contact area between the valve body and the spring button results in substantial static friction, which prevents the spring button rotating as the bayonet cap is rotated into place. The preloaded spring ensures that the bottom surface of the spring button remains pressed flush against the valve body once the bayonet cap is locked in place.

5.3. Closed-Loop Control

To provide feedback as to the current state of the system, the methods developed in chapter 4 were used to measure the flow for a range of control inputs. As the control algorithms for cell capture can be naturally expressed in terms of saddle point location, and given that the saddle point location is a robust measurement as it is computed using a fit to a number of velocity vectors, it is convenient to express the flow control system in terms of saddle-point position. As demonstrated in chapter 3, saddle-point position is a suitable proxy for relative flow rate. Therefore, in the remainder of this chapter, the discussion will focus on the control of the saddle-point, rather than the flow itself, although this does not imply a lack of generality of the methods presented to flow control in general.

5. Flow Control

Early work conducted using the control valves used a traditional closed-loop proportional-integral-derivative (PID) controller to regulate the position of the control valves and therefore the flow rates into and out of the microfluidic chip. After some tuning, this can provide robust and stable control of the saddle-point position. In addition, the exact characteristics of each valve need not be precisely measured, as the control algorithm can continuously vary the valve output until the desired saddle-point position has been reached. Some typical response curves are shown in figure 5.7. In the ideal case (figure 5.7a), with no hysteresis, the saddle-point position changes almost instantly in response to the demanded position. The small delay observed (15ms) is inherent to the camera exposure, readout and analysis system and will be referred to subsequently as the system inherent delay (see for a graphical representation of the system inherent delay). However, when the system is in a hysteretic state, that is, when the valve must change direction, the behaviour is significantly different (figure 5.7b). As the span of the hysteretic region is significantly larger than the actuation region, the system expends significant time with no measurable change in the saddle-point position as the valve output is slowly increased through the hysteretic region. To resolve this, the proportional gain could be increased, or substantial integral gain could be increased. However, the gains required to rapidly traverse through the hysteretic region are many times greater than those that are appropriate for stability in the non-hysteretic state. The effects of gains that are too large are evident in figure 5.7c. The system is initially stable, and then a small perturbation drives the system into uncontrolled oscillation. Clearly, therefore, a more sophisticated approach is required.

5.4. Hybrid Closed-Open Loop Control with Calibration

During trapping, a system response that approaches the control loop rate is highly desirable, as the candidate cells/particles move very quickly relative to the imaging rate (see introduction to chapter 6). A system which does not model the hysteresis behaviour will require a number of time steps to transition the system through the hysteretic region, as the control regimes are very different (see figure 5.7). Additionally, because of the first requirement, the system must react stably at a rate close to or faster than the system inherent delay. This rules out a purely closed-loop system, as by definition, a closed-loop system needs at least one time-step, and typically several, to behave stably. However, without a closed loop, the system is susceptible to drift, where the steady-state saddle-point position, in the absence of a control input will slowly change. This is due to the build-up of debris, or relocation of small air bubbles, which cause small changes in the fluid resistance of various sections of the fluid circuit. Secondly, changes in system pressure will cause a change in the steady-state position of the saddle-point as the small variations in tubing length and resistance in each of the four sections of the fluid circuit act upon the steady-state flow rates. The effect of a change in system pressure on the steady-state saddle-point position can be observed in figure 5.9.

To achieve the desired system response and allow the system to compensate for drift, a hybrid closed-open approach has been developed, similar to that proposed by Shoukat Choudhury, Thornhill and Shah,¹⁸ but with a significantly expanded system model and auto-calibration capabilities. A simplified diagram, showing the essential elements of this approach is shown in figure 5.10. The top line, from command input to plant output, forms the open-loop portion of the model. The plant output is delayed by the inherent delays present in the system, and analysed using the flow measurement techniques described in chapter 4. The error between the command and the actual position is passed through a low-pass filter, then integrated, before used as an offset to the original command. Essentially, the closed-loop component of the system acts at a much slower rate than the open-loop components, and acts to determine the bias required to impose upon the command input to compensate for the long-term drift.

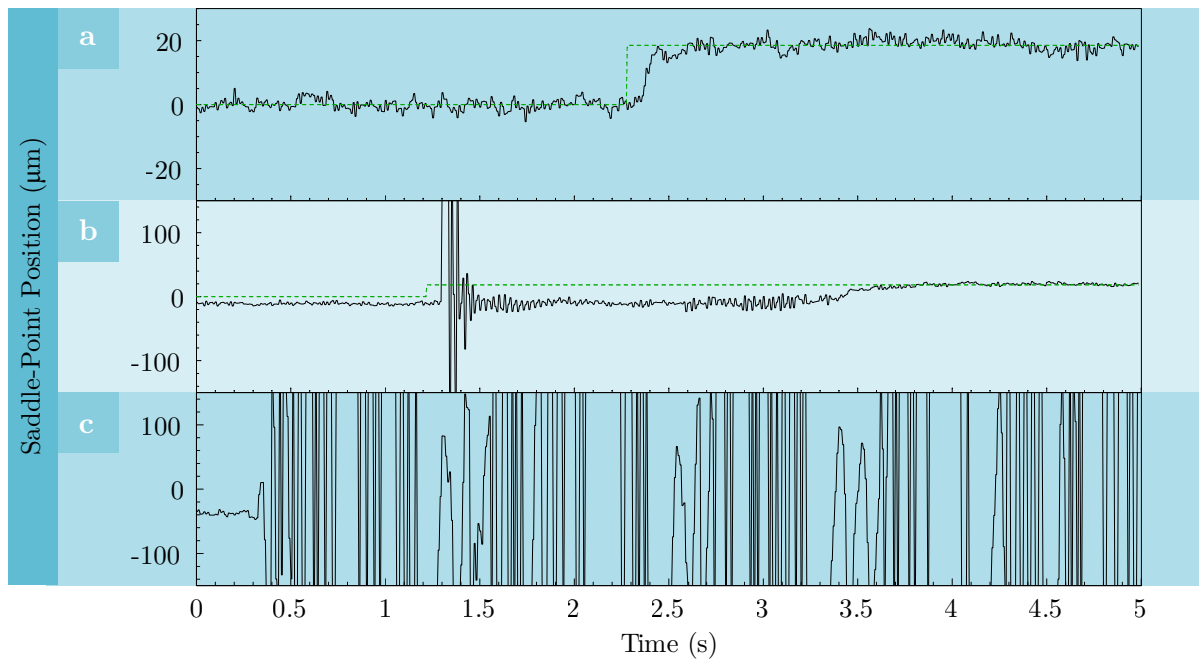


Figure 5.7: PID feedback control of the saddle point position. a) with no hysteresis (i.e. the system has already moved in the positive direction, and is directed to move in the same direction again), the system transitions to the new requested saddle-point position after a short delay; b) in a hysteretic condition (the system was previously moving in the negative direction, and has now been directed to move in the positive direction), there is a long delay and some transition noise as the valve moves through the hysteretic region; c) increasing the PID gains to decrease the transition time through the hysteretic region results in instability.

5. Flow Control

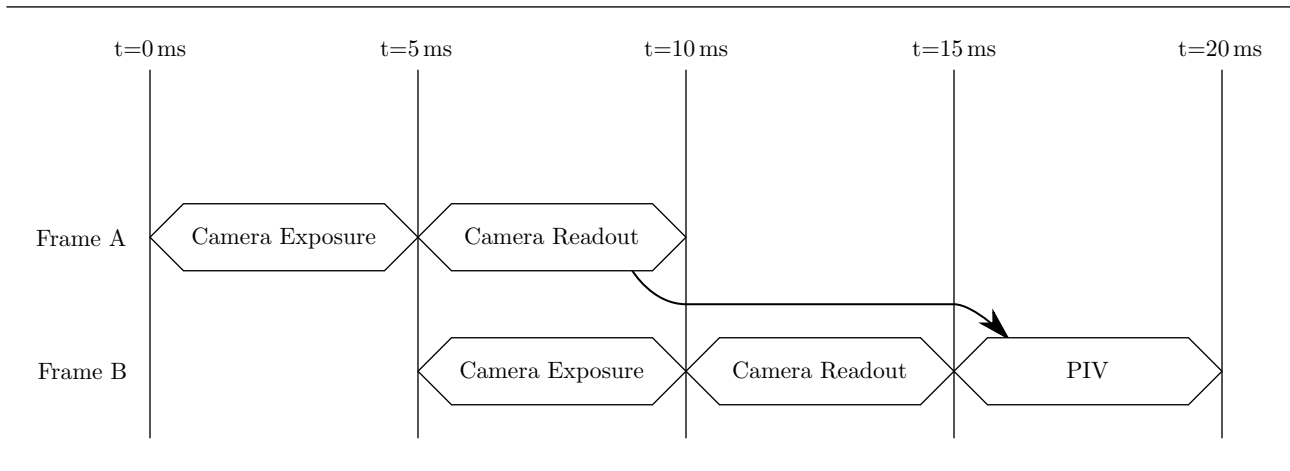


Figure 5.8: The inherent delays in the system that lead to a 15 ms delay from the initial change of flow conditions to the first measurement of that change. A change in flow conditions may occur at any time up to the beginning of the second exposure (Frame B), after which the information contained in the signal does not reach the rest of the system until after the PIV analysis is complete.




Figure 5.9: Drift in the saddle-point as a result of system pressure changes, showing the saddle-point position in the y -direction (solid-line) and the system pressure ($-\circ-$). The saddle-point position drifts 40 μm in the positive y -direction as a result of a 12 kPa decrease in system pressure.

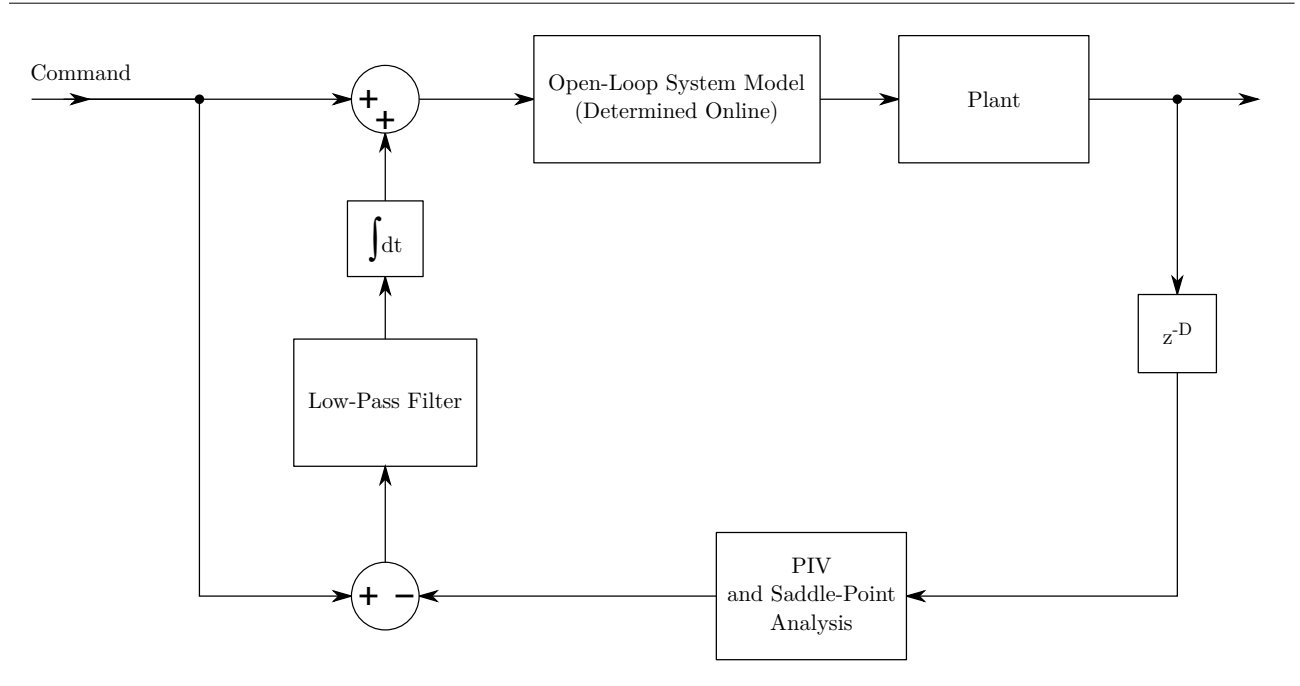


Figure 5.10: Simplified control systems diagram, showing the essential elements of the hybrid approach, combining a fast open-loop controller with a slower closed loop compensation loop.

In order to determine the parameters for the open-loop model, a calibration procedure must first be performed.

5.4.1. Calibration Procedure

As previously established in this chapter, the relationship between flow and control input and the hysteretic behaviour of the system are affected by a number of factors that can vary between experimental runs. For example, changing or cleaning the microfluidic channel used, replacing tubing, or simply moving the tubing in the valve (so that a different section of tubing is used), can all dramatically affect these relationships, and therefore the calibration. Hence, it was critical that automated methods of calibrating the valves and deriving suitable relationships between desired saddle point location and control input were developed, as manual calibration and analysis is time consuming and lacks repeatability.

When the software is initially started, no calibration data is available,[§] and the range of input voltages (the system *excitation*) that produce the desired change in saddle-point position (the system *response*) must first be determined for each valve. To achieve this, the valve input voltage for the first valve (the choice of the valve to calibrated first is arbitrary) is swept slowly from the minimum value to the maximum in discrete steps. The saddle-point position for each input voltage is averaged and recorded. If the saddle-point position exceeds the measurement range (that is, the saddle point is no longer in the field of view) for more than five voltage steps or the maximum voltage is reached, the valve is considered to have reached its maximum excursion and the negative-direction sweep is

[§]It is possible to provide an initial calibration file from a previous execution, but in this case it is assumed that no suitable data is available.

5. Flow Control

started. As for the positive sweep, the input voltage is swept slowly in discrete steps to the minimum value. As before, the saddle point position and valve input voltages are recorded. The measurement of both the positive and negative behaviours is necessary due to the presence of hysteresis. In a truly linear system, one would expect that the relationship between excitation and response in the positive and negative directions would be identical. However, this is not the case for the present system.

Once the measurement procedure has completed for all valves (in this system, there are four), a series of calibration curves can be generated, showing the relationship between valve input voltage and saddle point position in both the positive and negative directions. A sample initial calibration result can be seen in figure 5.11. This sample result highlights some of the peculiarities of this system. Firstly, over the wide voltage range swept, only a small range of excitation voltages produce a response in the form of a change the saddle-point position. This is due to a number of factors: the force developed by the actuator must first overcome gravity and then preload the tubing wall before deformation of the internal wall begins to take place; also, the field of view is limited - in this particular channel, the camera can see approximately one fifth of the total channel in each direction, hence limiting the range of voltages that have a measurable response. Secondly, there is substantial hysteresis present in the system. A larger excitation voltage is required in all cases to close the valve (reduce the flow rate and move the saddle-point away from the centre) than to open the valve (increase the flow rate and return the saddle-point closer to the centre). In many cases this voltage difference is greater than the voltage required to move the saddle-point from 0 (centre) to maximum once hysteresis has been overcome. This poses a challenge for the control system as small voltage changes are required when progressing linearly in a single direction, but much larger voltage changes are required as soon as a direction change is requested.

It should be noted that the resolution of the measurement in figure 5.11 is poor over the critical range of excitation values where the saddle-point position is changing. For efficiency, the initial sweep is made over a wide range of voltages in a coarse fashion, and analysis is performed on the results to provide a first estimate of the valve-tubing-system behaviour. Once this initial result has been obtained, a more thorough sweep can be performed for each valve using the unique calibration information obtained from the first estimate. Hence, the calibration estimate can be refined iteratively, producing a more accurate result with each pass. The analysis and subsequent calibration passes are explained in detail in the next section.

5.4.2. Calibration Analysis

Analysis of the calibration curve proceeds in an automated fashion as soon as the software has finished the calibration sweeps of all four valves. An estimate of the 99th and 1st percentile saddle-point positions is obtained from the curves. These are used as estimates of the maximum and minimum responses, respectively. These two values are then averaged to find an estimate of the response expected at the centre of the calibration curve. This middle response value is mapped to the closest point on the calibration curve, and an iterative linear regression process is then used to find a first estimate of the shape of the curve.

This middle value is used to form an initial guess for the the lower excitation value x_1 and the upper excitation value x_2 , which are the excitation values over which the valve is considered to behave linearly. A third-order polynomial is fit to the points bounded by x_1 and x_2 by linear least squares. The error estimate can then be defined in terms of the response $r[x]$ and the response estimate $r'(x)$

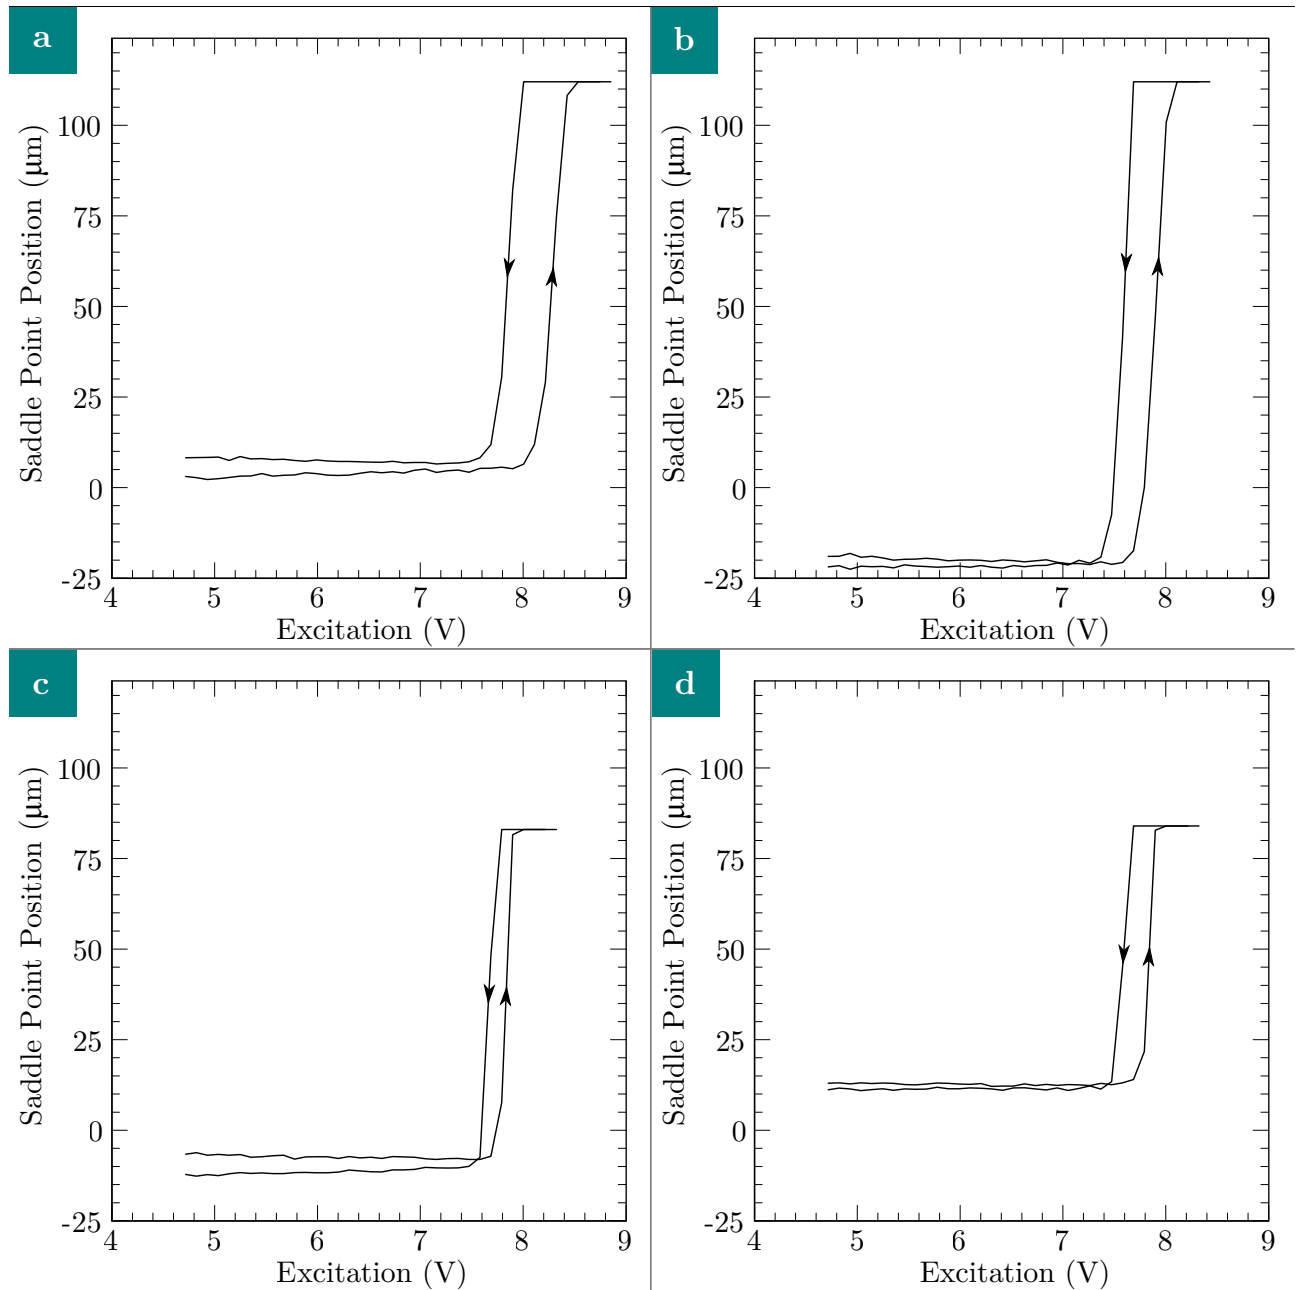


Figure 5.11: Results of initial calibration sweep for all four flow control valves: (a) left; (b) right; (c) top; (d) bottom. Of note is the relatively small range of excitation voltages over which control is effective; the substantial hysteresis that varies in magnitude and nature between valve-tubing systems; and finally, the initial (unperturbed) position of the saddle point can expand the measurable range of the valve, at the expense of the measurable range of the opposite valve.

5. Flow Control

	x_1	x_2
Positive-going	7.90 V	8.43 V
Negative-going	7.58 V	8.11 V

Table 5.3: Initial estimates for the minimum activation voltage (x_1) and maximum effective voltage (x_2) of the left valve, derived from calibration curves in figure 5.11. Voltages between these two limits result in changes to the flow rates observed in the system. A greater voltage (on average, 0.32 V higher) is required to produce the same flow rate in the positive direction, demonstrating the significant hysteresis present in the system.

(note that r is a discrete function as indicated by the $r[x]$ notation)

$$e = \sum_{x=0}^N (r[x] - r'(x))^2, \quad (5.1)$$

where the response estimate $r'(x)$ is defined as a piecewise function over $x_1 < x < x_2$, evaluating to a constant outside this region, and $f(x)$ is the usual form of a third-order polynomial, where a , b , c and d have previously been determined by the least squares estimate

$$r'(x) = \begin{cases} f(x_2) & x \geq x_2 \\ f(x) & x_1 < x < x_2 \\ f(x_1) & x \leq x_1, \end{cases} \quad (5.2)$$

$$f(x) = ax^3 + bx^2 + cx + d. \quad (5.3)$$

Equation 5.2 has the effect of emulating the saturation behaviour seen at the beginning and end of the calibration curves. This means that the error (5.1) is minimised when the section of the curve bounded by x_1 and x_2 closely matches a cubic polynomial, and the values outside the region bounded by x_1 and x_2 are constant and close to the initial / final values of the cubic over that range. The estimates for x_1 and x_2 are subsequently improved by incrementally decreasing (x_1) / increasing (x_2) the estimates until the slope of the error estimate (5.1) becomes positive. While simplistic, this minimisation scheme produces good results for a relatively small number of points (the total number of measurement points range from 50 – 200).

This estimation scheme provides robust and accurate estimates of both the initial and final values of excitation over which the valve produces measurable changes in response, but also allows a measure of the average, minimum and maximum hysteresis if both positive and negative curves are analysed. Applying this estimation scheme to the left valve in figure 5.11, we obtain values for (x_1, x_2) in both the positive and negative directions (table 5.3).

This results in a hysteresis estimate of 0.32 V (obtained by computing $x_{1_p} - x_{1_n}$; $x_{2_p} - x_{2_n}$ and averaging the results). These values also allow estimation of parameters for a second sweep: a starting value of 7.58 V and final value of 8.43 V. The software extends these values by 25% of the sweep range to reduce the risk that a noisy initial measurement will bias subsequent results. Hence the actual search parameters are a sweep of 7.36 V to 8.64 V. The results of this more detailed sweep can be observed in figure 5.12a. Note that the sweep range actually chosen by the software is slightly narrower




Figure 5.12: Detailed calibration sweep of (a) left valve from the same experimental run as shown in figure 5.11. The initial calibration pass allowed the sweep limits to be refined, providing a more detailed calibration curve in the second pass. Additionally, to highlight the full-closed behaviour of the valves, the sweep for the bottom valve (b) with a smaller channel has been plotted (symbols omitted from the plot so that shape of the end of the curves in both directions can be clearly seen).

than those shown in table 5.3. There are two reasons for this: the lower limit was computed by a slightly more advanced algorithm, which will be described in section 5.4.3, and resulted in a slightly higher estimate for x_1 ; the early termination condition described in section 5.4.1 resulted in the sweep terminating before reaching x_2 .

5.4.3. Forward Predictor Model

The simplified model presented in section 5.4.2 is sufficient for determining the limits of the valve's useful excursion and estimating hysteresis, but does not provide sufficient information to act as a forward predictor. In order to provide this more detailed model, a non-linear solver is employed (CERES SOLVER¹). However, non-linear solvers will, by design, seek a *local* minimum.¹⁵ This may not correspond to the desired solution, as many such local minima may exist. Hence, the algorithms described in section 5.4.2 are used to compute an initial guess for the parameters, ensuring that the minimisation problem converges toward the valve characterisation curve.

A visual and analytical study of a number of calibration curves has led to the observation that the curves follow the form of a sigmoidal curve. This is particularly apparent in the data presented in figure 5.12b, where the full extent of the excursion of the valve can be observed. This results in an estimator of the form

5. Flow Control

$$y'(x) \propto \text{erf}(x), \quad (5.4)$$

where $\text{erf}(x)$ is the Gauss error function $\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$.² This estimator can be parameterised in terms of the shift in x (x_c); the slope of the curve (or x stretch factor, x_s); the maximum height in y (y_{max}) and the y offset (y_{off}), as

$$y'(x) = \frac{\text{erf}\left(\frac{x-x_c}{x_s}\right) + 1}{2} y_{max} + y_{off}, \quad (5.5)$$

where $\frac{\text{erf}(x)+1}{2}$ simply shifts the y values of the error function to lie in the range $[0, 1]$. Initial estimates of the parameters can be computed from the parameter estimates obtained in the algorithm described in section 5.4.2:

$$\begin{aligned} y_{off} &= y[x_1], \\ y_{max} &= y[x_2] - y[x_1], \\ x_s &= \frac{x_2 - x_1}{4}, \\ x_c &= \frac{x_1 + x_2}{2}. \end{aligned} \quad (5.6)$$

Applying these parameter estimates to the positive-going curves in figure 5.12 results in the $y'(x)$ estimates shown in figure 5.13a (note, the algorithm equally applies to the negative-going curves, these have been omitted for the sake of clarity). These estimates show significant error compared to the original measured data, particularly in the case of the large channel - the estimation algorithm expects that the limits $[x_1, x_2]$ describe the bounds of a pure sigmoid, which is not the case where the field of view is limited, as with a larger channel. However, in both cases the guesses provide reasonable initial search parameters that are close to the local minimum of the cost function. After a small number of iterations of the non-linear solver, the fit converges (figure 5.13b) to the measured calibration curve. In both cases there is good agreement between the final fitted curve and the measured data.

These fitted curves are of sufficient quality, and are sufficiently robust in the presence of noise in the calibration measurement, to be used as forward predictors of the voltage required for a specific saddle point position. To achieve this, the curves are inverted, by interpolation, to build a numerical lookup table (LUT) of the voltages required, in both the positive-going and negative-going directions, to induce a particular deviation in the saddle-point position. A sample LUT is shown in figure 5.14 (the mapping between saddle-point position and voltage does not extend to zero due to the non-ideal initial conditions discussed earlier). Plotting the data in this way places the hysteresis and non-linearity of the system in stark relief. Of particular interest is the change in hysteresis with respect to saddle-point position. This is related to the deformation energy of the tubing as it is compressed. While these phenomena have not been investigated in detail, the change in elastic deformation of the tubing wall is greatest during the initial deformation of the tube, as the valve progresses from fully-open to closed. As the tubing approaches a fully-closed (fully-pinched) state, the change in deformation energy is lesser, which results in a lower observed hysteresis in this region.

With this automated calibration and characterisation method, a numerical forward model of excitation voltage for any desired saddle-point position (within the system limits) can be obtained for any valve in the system. However, the forward model does not provide information about how to transition between the positive-going and negative-going excitation curves. Furthermore, some means

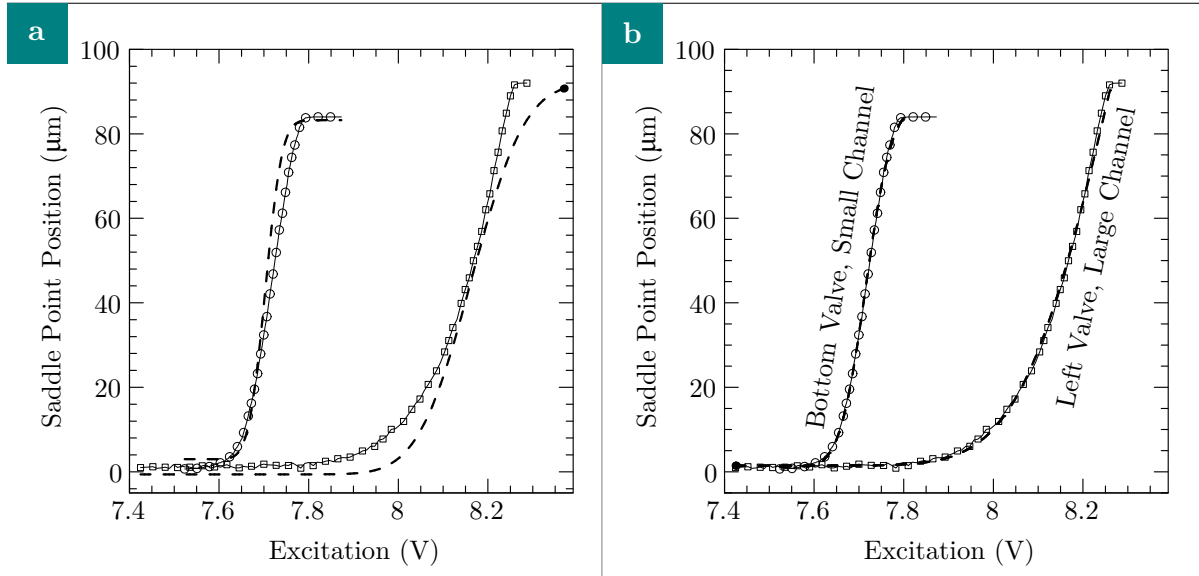


Figure 5.13: Sigmoidal fit to the valve calibration data, using data from bottom valve with a small (high flow rate) channel (—○—) and the left valve with a large (low flow rate) channel (—□—). (a) Initial guess (— — —) computed using the parameter estimates in equation 5.6; (b) final data fit (— — —) after 10 iterations of the non-linear solver. The calibration curves for the left valve have been offset by $-20 \mu\text{m}$ to facilitate comparison.

of tracking the current state of the system is required. For this, a further set of algorithms, collectively known as the **Motion Planner** are required.

5.5. The Motion Planner

The Motion Planner is responsible for managing the rate of change of the valve voltage outputs and tracking the state of the system, so that the correct excitation curves can be followed in all instances. To implement this, a finite state machine (FSM) is used to track valve state, and a kinematics model is used to control the rate of voltage change.

5.5.1. Tracking Valve State

The hysteretic state of the valve-tubing system is unobservable - that is, there is no system parameter or measurement that can directly provide information as to whether the valve-tubing system is in a positive- or negative- hysteresis state. It would be possible to infer this state based on the saddle-point changes with valve position, but this would remove many of the dynamic performance advantages of the partially-open-loop predictor model. Instead, the software tracks the hysteresis of the valve-tubing system by tracking the history of the commands sent to the motion planner. The valve is initially reset to a known state (the minimum voltage on the negative-going curve), which causes the tubing to relax to the full-open state. Starting from this known initial state, all commands are subsequently tracked and the internal state updated.

The state tracking algorithm must keep track of a number of external and internal state variables. The important state variables and constants are defined in table 5.4.




Figure 5.14: Lookup table (LUT) formed from computing the sigmoid fits for both the positive-going and negative-going calibration curves for the left valve described above, and then inverting these curves. Of note is the large hysteresis and non-linear nature of the relationship between saddle-point position and excitation voltage. Net hysteresis (blue|grey line) has been computed by subtracting the positive-going excitation from the negative-going excitation voltage. There is a trend of decreasing hysteresis as the valve progresses to fully closed, which gives insight into the underlying valve mechanics.

Variable	Type	Definition
c	External	Current position command to the motion planner, in terms of saddle-point position.
p	Internal	Current valve position, in terms of saddle-point position.
E_d	Internal	Demanded excitation voltage - the voltage requested by the state tracker.
E_a	Internal	Current excitation voltage, demand voltage after filtering by the kinematic model.
ϵ_a	Constant	Approach tolerance. Minimum distance $ c - p $ below which the state tracker considers that the valve has reached the commanded position.
ϵ_h	Constant	Hysteresis tolerance. Value to compare with $c - p$ to determine whether slope is positive or negative. This value is non-zero to prevent limit-cycling.
s_h	Internal	Hysteresis State: positive or negative .
s_m	Internal	Motion State: rapid or approach .

Table 5.4: Nomenclature used in the state tracker.

A formal state transition diagram for the valve state tracker is shown in figure 5.15. To simplify the diagram, *goalReached* is defined to be true iff $|c - p| < \epsilon_a$

Output behaviour of the system can be described in terms of the two transition states: approach and rapid. The **approach** state is the simplest and is represented by the two states on the right hand side of figure 5.15. In these states, the hysteresis state of the valve-tubing system is unchanged and the excitation proceeds along one of the two excitation curves (positive or negative, depending on the current hysteresis state). Increases (where $s_h = \mathbf{positive}$) or decreases (where $s_h = \mathbf{negative}$) to the command variable care rate-limited, as appropriate, by the kinematic model and a new excitation voltage is calculated via the lookup table. The internal state variable p , representing the current valve position, is updated as appropriate.

The **rapid** state is transitioned to if the difference between the command position and the current position becomes positive when the hysteresis is negative, or conversely, if the difference is negative when the hysteresis state is positive. In this state, the output behaviour is more complex. The hysteresis state is changed to reflect the expected hysteresis state once the transition is complete — for example, if the system were previously traversing the negative hysteresis curve, and the command input becomes more positive than the current position, the hysteresis state will switch to **positive**, and the motion state will switch to **rapid**. A new excitation voltage, E_d , is then computed via the same lookup table as before and passed to the kinematic model. As the actual excitation voltage E_a is changed by the kinematic model, the equivalent saddle-point position is back-calculated via the lookup table ($p_{inverse}$). At the start of the transition, there will be no valid saddle-point position for the output excitation voltage E_a as the system will be traversing the transition region (see figure 5.16 for an example of a transition from a **negative** hysteresis state to **positive**). However, as the voltage continues to change, the voltage will reach a point where it passes into a valid point of the relevant excitation curve (dashed line, $p_{inverse}$ valid on figure 5.16). Once the back-calculated position $p_{inverse}$ reaches the current position p , the back-calculated position is used to update p (final section on figure 5.16). Finally, once p is sufficiently close to c , that is the condition $|c - p| < \epsilon_a$ is satisfied, the motion state transitions to **approach**.

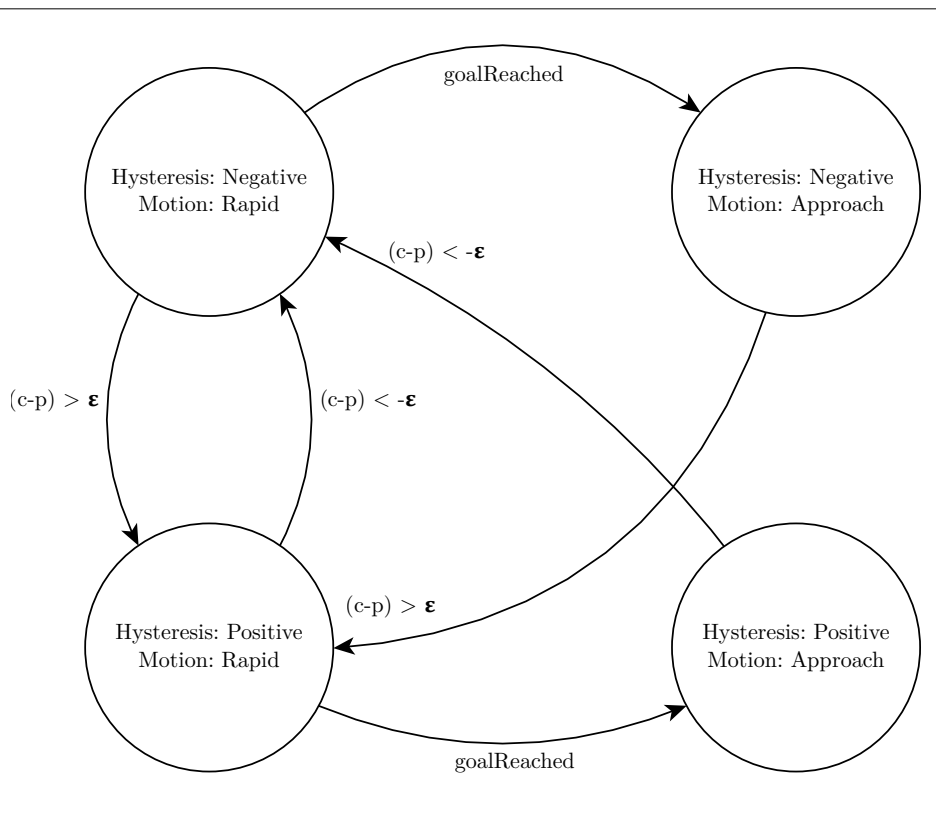


Figure 5.15: Finite State Machine model of the valve state tracking algorithm.

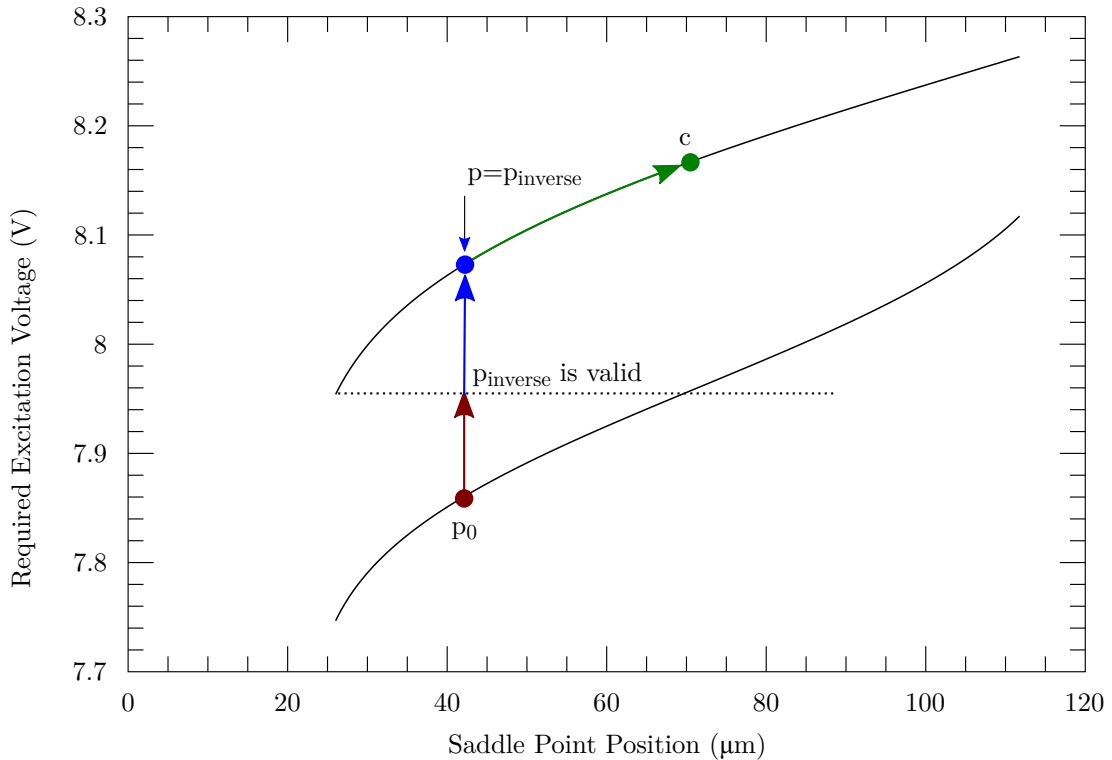


Figure 5.16: Traversal of the lookup table by the motion planner. The initial position is p_0 and the hysteresis state is negative. A new command is given (c) which causes a transition to the **rapid, positive** state. The back-calculated position, $p_{inverse}$ first becomes valid as the voltage crosses the start of the positive excitation curve. The current valve position p is first updated with $p_{inverse}$ once $p_{inverse}$ exceeds p_0 .

The continuous back-calculation to update p ensures correct behaviour if the command value c is rapidly changed. For example, if the system is transitioning to a positive state, and c is decreased below the original value, but above p , then no state changes are necessary and the system continues to approach the new, lesser c position. Conversely, if the system is transitioning to a positive state and the command input c is decreased (by the upstream controller) to a value above the previous value of c , but the back-calculated valve position p has already exceeded this new value of c , the system will correctly transition to a negative hysteresis state. A less sophisticated system model would not behave correctly in these circumstances.

5.5.2. Kinematics Model

The Kinematics Model is responsible for limiting the rate of change of the voltage output to the valves, and therefore controls the kinematic behaviour of the valve-tubing system. This is necessary as the valve should have the maximum possible velocity during the transition between hysteresis curves (the distance between p_0 and $p = p_{inverse}$ in figure 5.16), but should slow down and traverse much more slowly once the desired position has been reached. The control loop for the kinematic update is much

5. Flow Control

faster than the imaging control loop ($t_{step} = 500 \mu\text{s}$) allowing very fine control over acceleration and traversal velocities.

For the purposes of this discussion, the transition from a lower voltage to a higher voltage will be used as an example, but the same principles apply to the reverse direction. For this model, the system motion is considered to consist of three phases: acceleration from the current velocity (v_0) to the maximum velocity (v_{max}); a constant-velocity phase at v_{max} ; followed finally by deceleration from v_{max} to the final velocity (v_f). The rates of acceleration and deceleration may be different and are given by the constants a_+ and a_- , respectively. The distances (i.e., voltage) travelled during the acceleration and deceleration phases are

$$d_{accel} = \int_0^{t_{accel}} a_+ t + v_0 dt, \quad (5.7)$$

$$d_{decel} = \int_0^{t_{decel}} -a_- t + v_{accel} dt, \quad (5.8)$$

assuming that the distance travelled during the constant-velocity phase is 0. The distance travelled in the constant-velocity phase can be easily calculated as $d_{cv} = d_{total} - (d_{accel} + d_{decel})$. Solving for the minimum time required to accelerate from v_0 and subsequently decelerate to v_f allows for the computation of the distance travelled in all three phases (d_1 , d_2 and d_3 , respectively).

Once these distances have been computed, a simple discrete-time integration scheme can be used to time-step the kinematics model:

$$\begin{cases} v[t] = v[t-1] + a_+ t_{step} & d[t] < d_1, v[t] < v_{max} \\ v[t] = v[t-1] - a_- t_{step} & d[t] > d_2, v[t] > 0, \end{cases} \quad (5.9)$$

$$\begin{cases} d[t] = d[t-1] + v[t] t_{step} & d[t] < d_3 \\ d[t] = d[t-1] & d[t] \geq d_3. \end{cases} \quad (5.10)$$

5.6. Performance

The performance of the flow control system is critical to the success of the final system. Inaccuracies or slow response to command inputs will result in unstable trapping behaviour. Poor accuracy will require larger control gains, and therefore saddle-point excursions in order to guarantee that the saddle-point is above or below the cell position, as required to direct the cell in the correct direction. This will result in unstable behaviour. Similarly, a response that is too slow will result in the cell accelerating out of the field of view before the system can correct for the behaviour.

Dynamic Performance Assessments of the flow controller step response were used to measure the dynamic performance of the system. The step response is also a good model of the system behaviour during the initial trap. In the idealised case, a cell will be identified by the system when it is close to the y -centreline and close to the edge of the field of view. The cell will then begin to rapidly accelerate to one of the outlets, depending on the location of the cell relative to the centreline. The control system will then need to reposition the saddle-point such that the cell is driven back towards the centreline before the cell reaches the new saddle-point position — once the cell has passed this position, the net force will only serve to accelerate the cell further out of the field of view. Therefore, the speed at which the saddle-point can be repositioned has a direct effect on the maximum velocity

of the cell, and therefore the maximum extensional rate (see section §2.2).

In figure 5.17, a selection of step responses are presented for different motion planner parameters. All parameters have been normalised and expressed in terms of the percentage of full-scale travel (% FS). Use of this non-dimensionalisation allows the specification of these parameters for all four valves independently of calibration. As can be seen in the figure, adjusting the final velocity and deceleration has a dramatic effect on the speed of response and overshoot. In terms of the time to reach the set-point, a deceleration of $a_- = 20\% \text{ FS/s}^2$ and final velocity of $v_f = 20\% \text{ FS/s}$ results in a slow response of $t_{ss} = 800 \text{ ms}$ with minimal overshoot. With some optimisation, this can be improved to $t_{ss} = 375 \text{ ms}$ with minimal overshoot ($< 15\%$ of the step size). In terms of likely cell trapping response, this compares favourably with the dwell time computed in section §2.2 for an extensional rate of 5 s^{-1} (1600 ms). Based on dwell times, the response rate would be sufficient for initial capture at extensional rates up to 15 s^{-1} (with a corresponding dwell time of 530 ms).

The system response can be further improved by actuating both outlet (or inlet) valves simultaneously. For example, if the saddle-point is at the maximum positive direction (which corresponds to the top valve at its maximum excursion), the saddle-point can be rapidly repositioned to the maximum negative direction by simultaneously opening the top valve and closing the bottom valve. This principle is demonstrated in figure 5.18, where the saddle-point position is rapidly transitioned from $+60 \mu\text{m}$ to $-40 \mu\text{m}$ (a step of $100 \mu\text{m}$). This transition takes 100 ms (shaded region in figure 5.18), a significant improvement over the response time in figure 5.17. Furthermore, the overshoot that would occur as a result of such a rapid transition is significantly reduced, as there is an equal and opposite overshoot due to the motion of the opposing valve.

Effect of System Pressure on Dynamic Response In section §3.3, it was alluded that the system overall flow rate, which is determined by the system driving pressure in this experimental setup, has an important effect on overshoot and therefore dynamic response. It is important to make a distinction here between the system overall flow rate and the extensional rate. As demonstrated in section §3.3, there is a predictable relationship between the central corner radii and the extensional rate as a function of overall flow rate. Therefore, the relationship between dynamic response and system pressure shown in figure 5.19 demonstrates the trade-off that must be made between desired extensional rate, dynamic response, overshoot and channel geometry. Excessive overshoot on the order of twice the set-point (such as that exhibited at low driving pressures, e.g. $P = 0.5 \text{ kPa}$) has an adverse effect on stability. However, a slower system response (for example that shown at $P = 12.0 \text{ kPa}$) also has the effect of reducing system stability by lengthening the time it takes to respond to new cell positions. Hence figure 5.19 shows the family of dynamic response behaviours that can be chosen from by varying the general tubing and channel wall geometry (in this case tubing ID = $800 \mu\text{m}$ and the channel walls were $100 \mu\text{m}$ wide and $75 \mu\text{m}$ high, as discussed in chapter 2). Based on the selected dynamic behaviour, the channel corner radii can be chosen to optimise for the desired extensional rate. For the results presented in subsequent chapters, the underdamped behaviour at $P \geq 8.0 \text{ kPa}$ was preferred, and a corner radius of $10w$ ($r = 1 \text{ mm}$) was chosen.

Accuracy To assess accuracy, a similar process to the initial calibration sweeps was employed. However, instead of sweeping the raw voltage, the controller set-point is swept from $0 \mu\text{m}$ to the maximum saddle-point position ($90 \mu\text{m}$ in the outlet direction). This provides an assessment of the repeatability and accuracy of the controller, including the open-loop estimator, state predictor and motion planner. figure 5.20 shows the results of this process plotted for the left and top valves (figure 5.20a and b, respectively). Dashed lines indicate the ideal response, where the actual measured position of the saddle-point is exactly the same as the requested position in each traversal direction. As can be seen




Figure 5.17: Step Response of the system with different shaping parameters applied by the motion planner. The system was preconditioned to ensure that the step transitioned through the hysteretic region by ensuring that the previous command was a positive step. The parameters chosen for (a, —□—) results in a stable, but slow response to the step. Increasing the deceleration rate (b, —○—) results in a faster response but with significant overshoot. Finally, an increased deceleration, coupled with a slower final velocity (c, —◇—), results in a faster response with minimal overshoot.




Figure 5.18: Through the coordination of two opposing valves (in this case, the top and bottom valves), the performance of large steps can be further improved. The target saddle-point position ($-40 \mu\text{m}$) is reached within 100 ms (shaded region), and the overshoot that occurs as a result of such a rapid transition is significantly reduced, as there is an equal and opposite overshoot due to the motion of the opposing valve.




Figure 5.19: Step response of the flow control system at various net driving pressures. To facilitate comparison, both saddle-point position and valve position have been normalised such that the initial position is 0 and the final, steady-state position is 1. Saddle-point position has additionally been filtered to reduce high-frequency flow-related noise and therefore increase readability. As pressure is decreased, the system moves from an overdamped response to an underdamped response, as the effects of the rapid change in valve position become more prominent.

from the figure, the accuracy in the positive-going direction is very good, with the response curves for both valves closely following the ideal curve in both cases. In the negative-going direction, there is an average error of $< 10 \mu\text{m}$. This is due to the non-linear behaviour of the valve hysteresis, which varies depending on the speed of approach. Hence, the hysteresis measured during the relatively slow sweep differs slightly from the actual hysteresis when the valve position is transitioned rapidly from the positive-going to negative-going states by the motion planner. Nevertheless, this small residual error represents a significant improvement in linearity compared to the native behaviour of the valve-tubing system, where the uncompensated error is greater than the full-scale excursion of the valve.

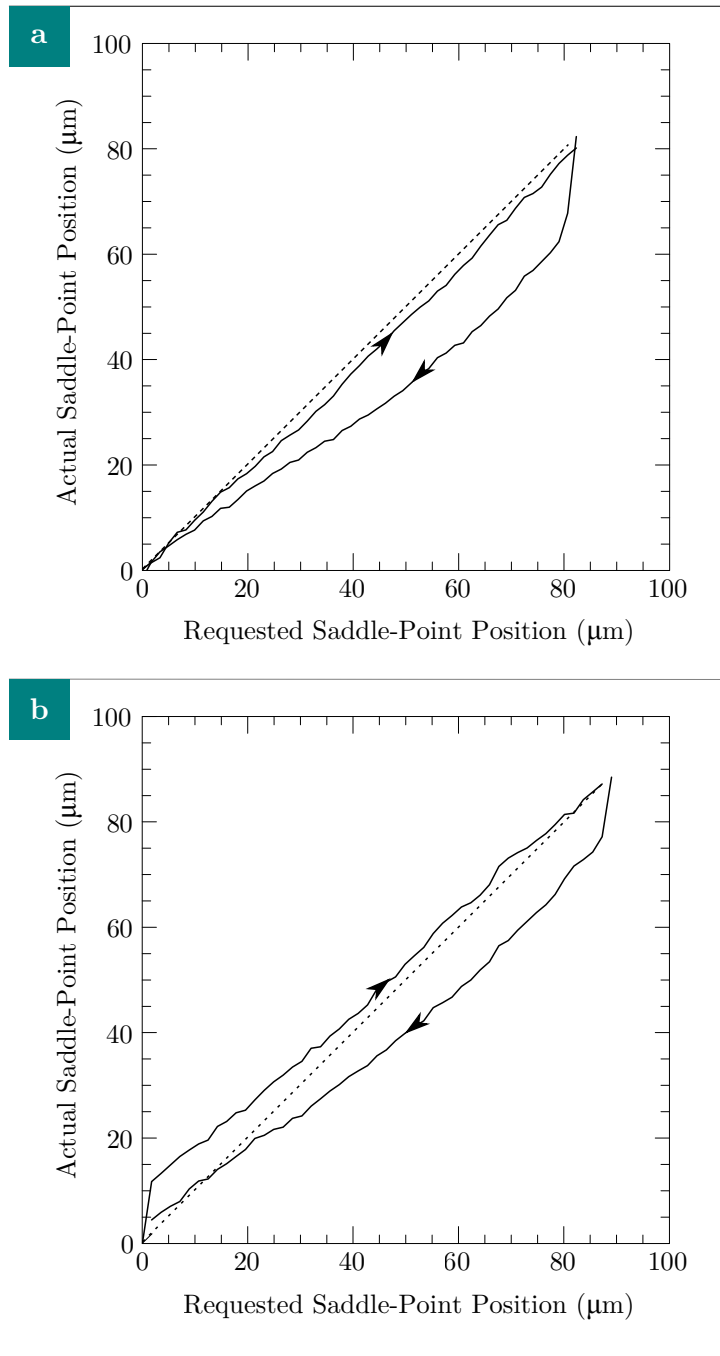


Figure 5.20: Validation of calibration, open-loop estimation and motion planner for (a) left valve and (b) top valve. In a similar fashion to the calibration sweep, a number of saddle point positions are requested sequentially in both the positive-going and negative-going directions. The dashed line in both sub-figures represents the ideal response — whereby the actual saddle-point position is exactly the same as that requested. There is a small amount of hysteresis remaining after correction, but the linearity of the valve-tubing system is significantly improved.

5.7. Conclusion

Control systems can radically transform the performance characteristics of dynamic systems. In this chapter, the strongly non-linear, hysteretic system consisting of the valves, silicone tubing and the fluidic system has been transformed into a linear, high-bandwidth flow controller. Transient performance, which is of paramount importance to the successful trapping of cells, was found to be excellent, and a significant improvement over control with state-of-the-art on-chip polymer valves,¹⁷ where fluid response times 250 ms have been reported (in this thesis, 100 ms response times have been achieved in similar tests). This performance could be further improved by further miniaturising the off-chip valves presented in this thesis, which would allow the valves to be moved closer to the chip and reduce dead volume. Additionally, the automatic calibration methods presented enable very precise control of flow rates in fluid systems, independently of geometry and, unlike other systems, do not require lengthy expert setup or precise manufacture.

Finally, the use of active control in software to stabilise and modify system behaviour shifts complexity away from the relatively expensive and consumable microfluidic chip and into a standardised and durable hardware platform and associated software. In so doing, the flow control system developed in this thesis enables software-controlled microfluidic platforms which can have behaviours added, changed and tuned *via* software updates. This enables rapid assays that reduce unit cost through a shared platform, can be upgraded in the field, and take advantage of sophisticated software for flow control and image processing.

5.8. Bibliography

- [1] Sameer Agarwal and Keir Mierle. *Ceres Solver*. version 1.4.0. (2012). Google, Inc. URL: <http://ceres-solver.org/>.
- [2] Larry C. Andrews. *Special functions of mathematics for engineers*. In collab. with Society of Photo-optical Instrumentation Engineers. 2nd. SPIE, 1998. 275 pp. ISBN: 978-0-8194-7846-7.
- [3] *ASCO 202 Series Proportional Valve Datasheet*. 2014. URL: <http://www.ascovalve.com/applications/products/ASCO-ProportionalValves.aspx>.
- [4] C. D. Bertram. “Experimental Studies of Collapsible Tubes”. In: *Flow Past Highly Compliant Boundaries and in Collapsible Tubes*. Ed. by Peter W. Carpenter and Timothy J. Pedley. Fluid Mechanics and Its Applications 72. Springer Netherlands, 1 Jan. 2003, pp. 51–65. ISBN: 978-90-481-6235-2 978-94-017-0415-1.
- [5] C. D. Bertram. “The effects of wall thickness, axial strain and end proximity on the pressure-area relation of collapsible tubes”. In: *Journal of Biomechanics* 20.9 (1987), pp. 863–876. ISSN: 0021-9290. DOI: 10.1016/0021-9290(87)90146-1.
- [6] C. D. Bertram, C. J. Raymond and T. J. Pedley. “Mapping of instabilities for flow through collapsed tubes of differing length”. In: *J. Fluids Struct.* 4.2 (1990), pp. 125–153.
- [7] Liyu Cao and H.M. Schwartz. “Stick-slip friction compensation for PID position control”. In: *Proceedings of the 2000 American Control Conference*. Proceedings of the 2000 American Control Conference. Vol. 2. 2000, pp. 1078–1082. DOI: 10.1109/ACC.2000.876666.
- [8] Debadi Chakraborty et al. “Fluid-structure interaction in deformable microchannels”. In: *Phys. Fluids 1994-Present* 24.10 (1 Oct. 2012), p. 102002. ISSN: 1070-6631, 1089-7666. DOI: 10.1063/1.4759493.

Bibliography

- [9] *Clippard EVP Series Proportional Control Valves Datasheet*. 2003. URL: http://www.clippard.com/downloads/PDF_Documents/Product%20Data%20Sheets/EVP_Proportional_Valve_Data_Sheet.pdf.
- [10] Yang Jun Kang and Sung Yang. “Fluidic low pass filter for hydrodynamic flow stabilization in microfluidic environments”. In: *Lab Chip* 12.10 (24 Apr. 2012), pp. 1881–1889. ISSN: 1473-0189. DOI: 10.1039/C2LC21163G.
- [11] Dean Karnopp. “Computer Simulation of Stick-Slip Friction in Mechanical Dynamic Systems”. In: *J. Dyn. Sys., Meas., Control* 107.1 (1 Mar. 1985), pp. 100–103. ISSN: 0022-0434. DOI: 10.1115/1.3140698.
- [12] Alan T. H. Lau et al. “Dynamics of Microvalve Operations in Integrated Microfluidics”. In: *Micromachines* 5.1 (10 Feb. 2014), pp. 50–65. DOI: 10.3390/mi5010050.
- [13] Zida Li et al. “Syringe-pump-induced fluctuation in all-aqueous microfluidic system implications for flow rate accuracy”. In: *Lab Chip* 14.4 (21 Jan. 2014), pp. 744–749. ISSN: 1473-0189. DOI: 10.1039/C3LC51176F.
- [14] R.H. Liu et al. “Passive mixing in a three-dimensional serpentine microchannel”. In: *J. Microelectromechanical Syst.* 9.2 (June 2000), pp. 190–197. ISSN: 1057-7157. DOI: 10.1109/84.846699.
- [15] George Nocedal and Stephen Wright. *Numerical Optimization*. 2nd. Springer Series in Operations Research and Financial Engineering. Springer, 2006. 664 pp.
- [16] *Parker - VSO-MI - Miniature Proportional Valve Datasheet*. URL: <http://ph.parker.com/us/12051/en/vso-mi-miniature-proportional-valve>.
- [17] Harrison Prentice-Mott, Mehmet Toner and Daniel Irimia. “Microfluidic proportional flow controller”. In: *J Micromech Microeng* 20.11 (15 Oct. 2010), pp. 1–8. ISSN: 0960-1317. DOI: 10.1088/0960-1317/20/11/115020. pmid: 21874096.
- [18] M.A.A. Shoukat Choudhury, N.F. Thornhill and S.L. Shah. “Modelling valve stiction”. In: *Control Engineering Practice* 13.5 (May 2005), pp. 641–658. ISSN: 0967-0661. DOI: 10.1016/j.conengprac.2004.05.005.
- [19] C. R. Siviour et al. “High strain rate compressive behaviour of a Silicone Elastomer”. In: *2008 SEM XI International Congress & Exposition on Experimental & Applied Mechanics*. 2008 SEM XI International Congress & Exposition on Experimental & Applied Mechanics. 2008.
- [20] *Takasago NPV Series Proportional Diaphragm Valve Datasheet*. 2011. URL: <http://www.takasago-fluidics.com/news/2011/04/000089.html>.
- [21] Elham Tolouei et al. “Effect of Hemodynamic Forces on Platelet Aggregation Geometry”. In: *Ann Biomed Eng* 39.5 (Jan. 2011), pp. 1403–1413. ISSN: 0090-6964. DOI: 10.1007/s10439-010-0239-4.
- [22] Marc A. Unger et al. “Monolithic Microfabricated Valves and Pumps by Multilayer Soft Lithography”. In: *Science* 288.5463 (4 July 2000), pp. 113–116. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.288.5463.113. pmid: 10753110.
- [23] Jianhua Zhou et al. “Pumping-induced perturbation of flow in microfluidic channels and its implications for on-chip cell culture”. In: *Lab Chip* 11.13 (13 June 2011), pp. 2288–2294. ISSN: 1473-0189. DOI: 10.1039/C0LC00466A.

6. Cell Tracking

A necessary prerequisite for trapping a particle or cell in a software-defined system is a mechanism for automatically identifying and tracking it. For the applications considered in this thesis, these particles are cells, and the techniques developed herein are optimised for red blood cells. Algorithms for tracking cells are usually intended for microscopy post-processing, where speed is not a concern^{1,8,9,13,22,26} or operate over relatively long time-scales.^{12,28} Realtime tracking systems have been developed by a number of researchers.^{3,11,14,17,19,24,29} However, these algorithms are optimised for low-noise images where typically only one cell is in the field of view at a time^{3,11,14,19,29} and may place particular constraints on the system, such as very low bit-depth¹⁷ or small image dimensions.²⁴

In this work, algorithms have been developed that are robust in the presence of noise (due to imaging constraints and the presence of tracer particles) and are fast enough to track cells in realtime. A particular challenge was the speed at which the cells enter and then subsequently leave the field-of-view when not redirected by the trapping control system. In the ideal case, for a cell entering close to the system centreline, and with a relatively slow flow, the transit time for entry to exit is < 500 ms. This means that the tracking system cannot rely on user input to define seed points or assist in the initial identification of candidate cells — it must be fully automated.

In this application, there is a strict time-budget, and therefore the tracking algorithms must be focused on efficiency as well as accuracy and autonomy. Frames are received from the camera at 200 frames/s, and the main control loop also updates at the same rate. However, due to the double-pulsed acquisition (section 4.1.4.1), it is only necessary to process one of the frames in each pair for cell tracking. This is because the intra-pair motion of the tracked cells is significantly smaller than the inter-pair motion and therefore processing both images in each pair does not add significantly to the information obtained. Therefore, the tracker is only invoked on every second iteration of the control loop (an effective rate of 100 frames/s or a time budget of 10 ms). However, the control loop would ideally make use of the updated cell position on the next iteration of the loop - a time budget of 5 ms. Hence the cell tracker should ideally have updated the cell position after 5 ms, but only needs to be ready for another frame every 10 ms.

The algorithms presented in this chapter fit within the available time-budget, and makes full use of the rich data available from this unique system and extensive use of the multi-threading capabilities of the target system. As there are potentially several cells within the field-of-view in a any given frame, as well as additional noise from the tracer particles, accurate tracking is required to feed consistent information to the cell capturing layer of the control system.

Tracking reliability has a direct effect on the stability of the trap and the control system in general. To demonstrate this, two potential instability modes will be briefly discussed. The first, and perhaps most intuitive, is a loss of tracking. Poor identification or inter-frame matching may lead to the algorithm ‘losing’ the cell track - that is, it can no longer find a valid match in the current frame. At this point, in the absence of a restoring force from the control system, the cell will rapidly drift out of the trap due to the flow velocities in the outlet direction. If the cell is outside the trap (in the process of being trapped), the control system will likewise fail to make the necessary adjustments to the flow velocities to guide the cell into the trap. This results in the first mode of instability. The second mode is a result of poor inter-frame matching of cells. Consider two cells in the flow of similar size and orientation being conducted towards the centre of the device close to, and on opposing sides of, the inlet centreline (orientation at this point in the flow is strongly related to the flow velocity

6. Cell Tracking

gradient, and the size distribution of red blood cells is relatively narrow,²³ so this occurrence is not as unlikely as it may appear). As the two cells appear similar (from an imaging perspective) the matching algorithm may, on a frame-by-frame basis, match either one to the similar cell in the previous frame. This oscillation may occur until one or both of the cells have been accelerated further from the system centreline. The end result of this instability will be that the control system will rapidly oscillate between trapping one cell or the other, ultimately becoming unstable.

In general, tracking algorithms can be separated into two general phases: identification (segmentation) and tracking (matching segments across subsequent frames). Due to the additional background noise in this system, arising from the choice of imaging modality and presence of tracer particles, an additional step to reduce the background noise is also required.

6.1. Background Subtraction and Image Normalisation

In order to present a normalised image to the segmentation algorithm, the input images are first flattened (intensity-normalised) and the background removed. This improves the reliability of segmentation and tracking as the algorithms are better able to identify regions of similar pixel intensities when the intensity of the background noise has been reduced.

The presence of a significant image background is a side-effect of the transmission microscopy setup, where the lighting is focused through an imperfect and uneven lens surface - the top surface of the microchannel (see chapter 2). Additionally, the lighting from an LED source, when projected onto a flat imaging plane such as a camera, follows a Gaussian intensity pattern due to the uneven emission pattern of an LED.¹⁶ In order to compensate for the effects of uneven lighting, as well as the background-related noise, an estimated background image is computed and used to compute a background-subtracted and intensity-normalised image:

$$\mathbf{i}_{c,i,j} = \frac{\mathbf{i}_{i,j} - \hat{\mathbf{b}}_{i,j}}{\widehat{\mathbf{b}}_{i,j}(\bar{\mathbf{b}})^{-1}}, \quad (6.1)$$

where $\mathbf{i}_{i,j}$ is the intensity of the input image at i, j ; $\hat{\mathbf{b}}_{i,j}$ the estimate of the background average at i, j ; $\bar{\mathbf{b}}$ the average of all pixels in \mathbf{b} ; and, $\mathbf{i}_{c,i,j}$ the compensated image.

The background image is estimated from a long-term recursive average⁶ of the input pixels of the form:

$$\mathbf{b}_{i,j}[t] \simeq \alpha \mathbf{x}_{i,j}[t] + (1 - \alpha) \mathbf{b}_{i,j}[t - 1], \quad (6.2)$$

where the exponential averaging coefficient α is chosen so that the average $\bar{x}_{i,j}$ approximates an average of 25 s of frame data. By induction, it follows that α can be computed in terms of the equivalent length of a moving average filter N by the relation:

$$\alpha = 10^{\log_{10}(fv)/N}, \quad (6.3)$$

where fv is the final value of the contribution of the frame at time $t - N$ to the average at time t . Empirically, a final value of 0.20 gives similar results to the equivalent moving average filter,

$$\mathbf{b}_{i,j}[t] = \frac{1}{N} \sum_{i=1}^{N+1} x_{i,j}[t - i], \quad (6.4)$$

but is computationally much less intensive as only two multiplications and one addition is required per pixel (equation (6.2)), rather than several hundred per pixel (where $N \approx O(100)$, equation (6.4)).

As the images are constantly changing due to the motion of particles within the flow, there is the potential for the background averaging process to result in a blurred image of the particle motion. To mitigate against this, two averages are computed, where one is based on a decimated timebase (the ‘slow’ average) and the other is updated for every frame (the ‘fast’ average). These two averages are then simply combined using a weighted average to form the current background average

$$\mathbf{b}_{i,j} = \beta \mathbf{b}_{\text{slow}_{i,j}} + (1 - \beta) \mathbf{b}_{\text{fast}_{i,j}}. \quad (6.5)$$

The slow average is given a much higher weight ($\beta = 0.85$), which allows the system to compensate for small intensity or background fluctuations, without the effects of blurring dominating the background image. This is an example of a multisample technique, that to the best of the author’s knowledge, has not previously been applied in this way to background recovery.

A sample result is shown in figure 6.1. As is apparent from the background image (figure 6.1b), there is an intensity gradient due to optical imperfections (misalignment, lensing effects, uneven LED emission). By normalising the image by the average intensity, this gradient is substantially removed, as can be observed from comparison of the respective line plots in (e,f). At $x = 80$ px, the intensity-normalised image (c) demonstrates a sharper transition between the dark interior of the RBC and the edge of the cell. Similarly, the two particles close to the edge of the image have the same intensity in (c), as compared to (d) where the particle at $x = 625$ px has approximately half the intensity of its neighbour. Intensity normalisation assists the segmentation and tracking processes by making edges more distinct and reducing spurious intensity variations that may interfere with the inter-frame matching process.

6.2. Segmentation

The segmentation algorithm uses a combination of thresholding and level set techniques, known as region growing.^{2,10,15} In a departure from typical region growing algorithms, this approach is focused on extracting small regions of interest from a noisy background, rather than subdividing a large region into labelled segments. The approach is broadly similar, though the choice of seed points and termination conditions must be more carefully considered.

The algorithm first identifies bright pixels as possible seed points (thresholding) and then expands each seed point to find an area of constant (level) gradient and intensity (level sets). Due to a number of factors (the effects of Fourier optics due to the effective numerical aperture of the imaging system, image noise, lensing effects of the particles and cells), the intensities within a cell or particle are not constant. Hence, the level set is defined as a pair of inequalities,

$$L_{\text{particle}}(I, \nabla I) = \left\{ (\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n) \left| \begin{array}{l} I[\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n] > I_{\text{min}} \\ \nabla I[\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n] < m_{\text{max}} \end{array} \right. \right\}, \quad (6.6)$$

where I_{min} is the minimum allowable intensity, m_{max} the maximum allowable gradient, and $I[\vec{x}]$ the intensity of the background-compensated input image at coordinate vector \vec{x} . m_{max} is a system tunable parameter and I_{min} is a composite of two parameters

$$I_{\text{min}} = \max(I_{\text{min}_{\text{local}}}, I_{\text{min}_{\text{global}}}), \quad (6.7)$$

where $I_{\text{min}_{\text{global}}}$ is a system tunable, and $I_{\text{min}_{\text{local}}}$ is the local minimum. Finally, the segmentation is rejected if the effective radius of the level set is less than r_{min} ,

$$\| \max(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n) - \min(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n) \| < 2r_{\text{min}}. \quad (6.8)$$

6. Cell Tracking




Figure 6.1: Background removal and normalisation. a) Sample frame from original image sequence; b) background image computed using the recursive average filter as described in the text (equations 6.2 and 6.5); c) image from a) after background subtraction and normalisation (equation 6.1); d) image from a) using a standard moving average filter and background subtraction (implemented using the ImageJ software²⁵). e,f) show the intensity profile for the line $y = 100$ px in images c) and d), respectively (as indicated by the pointers on the images).




Figure 6.2: Region-growing image segmentation process. a) an initial seed point is identified, and the surrounding pixels are interrogated in each direction; b) pixels that are added to the region are then treated as new seed points and the surrounding pixels are again interrogated; c) the segmentation is complete when no further pixels are added to the region.

This reduces the number of small tracer particles and noise in the final list of segmented particles.

To determine the initial seed points, a further threshold, I_{start} is introduced, which determines the minimum intensity required for a pixel to be considered as a potential seed point for the region growing algorithm. Once a seed point is found (via an exhaustive search of the image), the neighbouring pixels are interrogated in an attempt to expand the level set (region growing). Pixels successfully added to the level set are also added to a global mask. This global mask excludes a successful level set from subsequent seed point searches, so that the same region is not segmented multiple times. This process is shown graphically in figure 6.2.

In order to reduce the potential effects of varying lighting intensities, and to make the thresholds easier to specify, a histogram is computed for each frame and the thresholds $I_{min_{global}}$ and I_{start} are specified in terms of percentiles in the image histogram. Additionally, the cell is expected to be significantly brighter than the surrounding background, so a second limit, the local minimum can be computed from the local average intensity

$$I_{min_{local}} = \frac{1}{N} \int_{\Omega} I(\vec{x}) d\vec{x}, \quad (6.9)$$

$$\Omega = \left\{ (\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n) \mid \begin{pmatrix} x_0 - 2r_{min} \\ y_0 - 2r_{min} \end{pmatrix} \leq \vec{x} \leq \begin{pmatrix} x_0 + 2r_{min} \\ y_0 + 2r_{min} \end{pmatrix} \right\},$$

where Ω is the integration domain, consisting of a region that is $\pm 2r_{min}$ from the seed point. r_{min} , another system tunable parameter, is the minimum radius of a particle for a successful segmentation.

This level-set approach, coupled with appropriately-chosen termination conditions for intensity gradient and minimum intensity thresholds, results in much better shape-recovery than simple thresholding methods and avoids many of the noise problems that can be encountered with thresholding. Figure 6.3 demonstrates this. The original shape of the red blood cell is recovered accurately in figure 6.3b, whereas the thresholding techniques in figure 6.3c and figure 6.3d suffer from loss of shape information or significant amounts of pseudo-random noise, respectively.

6.3. Tracking

Tracking is achieved with a small finite state machine (FSM) with three basic states: **searching**; **tracking**; and **recovery** (figure 6.4). The FSM is initialised in the **searching** state. Once a suitable

6. Cell Tracking




Figure 6.3: Comparison of segmentation and thresholding algorithms. a) Original image, after background subtracted. Inset shows $2\times$ zoomed section of the top left corner (as shaded); b) Final aggregate mask image from segmentation showing the connected regions identified by the algorithm. Segmented particles with an equivalent radius less than $r_{min} = 5$ px have been rejected; c) thresholded image via the maximum entropy method of Kapur, Sahoo & Wong, which chose a threshold of 47%; d) thresholded image manually adjusted to 35% to recover the original shape of the cell (as shown in the centre of the inset). The level set techniques used in the segmentation algorithm recover the shapes of the objects in the field of view without the extraneous noise suffered by the thresholding techniques.

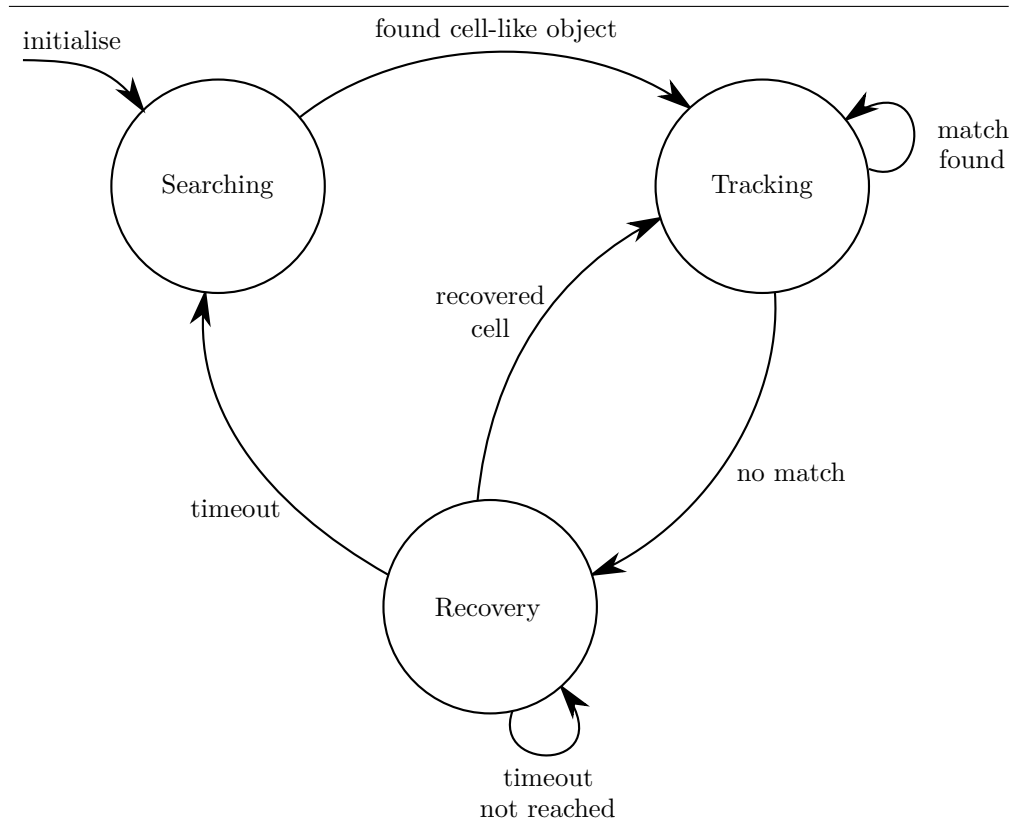


Figure 6.4: Finite State Machine for the cell tracker.

cell has been detected, the software transitions to the **tracking** state. The software remains in the **tracking** state while a valid match can be found between the reference cell and one of the cells in the current frame. If a valid match cannot be found, the software transitions to the **recovery** state. The software remains in the **recovery** state until the cell is located again (and therefore transitions back to the **tracking** state) or the **recovery** state times out, in which point the tracked cell is marked as lost and the software returns to the **searching** state. These three states are elaborated in the sub-sections to follow.

6.3.1. Searching State

In the searching state, the software will automatically select a candidate cell to track. This decision can be made based on various criteria, including size and initial position. In the current implementation, the software searches for a cell which is greater than a defined minimum size and whose centroid is within the region S (equation (6.10)). Additionally, cells that are closest to the centre of the region S are preferred.

$$\left\{ S(x, y) \mid \begin{array}{l} 0 < x < w/2 \\ \wedge \quad -h/6 < y < h/6 \end{array} \right\} \quad (6.10)$$

This region has been selected to maximise the success of the trap. Cells that enter the system outside of the vertical band (corresponding to $1/3$ of the total image height) will be rapidly accelerated towards one of the two outlets. Hence the system strongly prefers cells which are closer to the

6. Cell Tracking

centreline. Secondly, cells enter from the inlet closest to $x = 0$ in the current configuration, therefore the region captures those cells that have recently entered the field of view.

6.3.2. Tracking State

Once a candidate cell has been selected, the software transitions from the **searching** state to the **tracking** state. An ideal tracking function must correctly match the tracked cell from the previous frame (the *reference cell*) to one of a number of candidate segmented cells in the current frame. This is a problem explored extensively in the field of Particle Tracking Velocimetry, where probabilistic methods are common.^{4,5,18} Particle Image Velocimetry, explored in detail in chapter 4, can also be used to assist the matching process.⁷ In contrast to the methods presented in the literature, which use an iterative technique to find the best match between each particle in the reference frame to a corresponding particle in the current frame, the algorithm presented in this thesis searches for a single particle (cell) in the current frame that best matches the reference particle (cell). While the iterative process is more accurate, and requires less tuning, it can be computationally expensive as the matching process scales with $\mathcal{O}(n^2)$ (where n is proportional to the number of particles). This is particularly concerning in a realtime algorithm, as the number of particles in a given frame may vary depending on the segmentation input, therefore introducing large variations in the time required to complete the matching process. Initial testing showed that this iterative process would work acceptably in small-scale tests, but then would fail in real-time testing as the number of input particles varied on a frame-by-frame basis.

An algorithm has therefore been developed to score a match between two particles that can be used as an absolute (rather than relative) predictor. That is, scores below a certain threshold should be considered too weak for an acceptable match, and the best score for a given particle-pair is considered to be the correct pairing. To do this, a stochastic scoring mechanism, using a number of objective criteria, has been implemented. Each criteria is scored and transformed to a probability-like variable using a modified Gaussian function,

$$p^*(x, \sigma) = e^{-x^2/2\sigma^2}, \quad (6.11)$$

where x is the variable of interest and σ^2 the allowable variance in the variable. Note that this differs from the standard definition of a Gaussian probability distribution as the integral from $-\infty$ to ∞ is not 1. Instead, the value of the function is 1 when the likelihood is maximised ($|x| \simeq 0$). While not formally correct in terms of probability theory, this formulation is convenient as the product of multiple such probability-like variables (*scores*) is 1 when all scores are maximised. This has the corollary that a minimum acceptable score can be chosen that does not need to be adjusted as criteria are added or removed from the algorithm. Hence in all cases the match score $M(C|R)$ can be computed as

$$M(C|R) = \prod_{i=0}^N p_i^*(C|R), \quad (6.12)$$

where N is the number of criteria and the $M(C|R)$ notation has the same meaning as in probability theory (the expectation of a match of the candidate cell C given the reference R).

A suite of four criteria have been developed to estimate the likelihood that a candidate cell is the same cell as the reference cell. These are: size deviation; intensity deviation; distance deviation; and time deviation. These criteria are defined in table 6.1.

Criterion	Description
Size Deviation	The deviation between the area of the candidate cell (a_c) and the reference cell (a_r).
	$p_s^* = p^* \left(\frac{a_c - a_r}{a_r}, \sigma_s \right) \quad (6.13)$
Intensity Deviation	The deviation in average intensity (grey levels) between the candidate and reference cell.
	$p_I^* = p^* \left(\frac{\bar{I}_c - \bar{I}_r}{\bar{I}_r}, \sigma_I \right) \quad (6.14)$
Distance Deviation	The deviation in mean distance to the estimated cell streamline (motion path) between successive frames.
	$p_D^* = p^* (d, \sigma_D) \quad (6.15)$
Time Deviation	The deviation in effective transit time along the streamline between successive frames.
	$p_t^* = 1 + p^* \left(\frac{t_c - t_r}{t_r}, \sigma_t \right) \quad (6.16)$

Table 6.1: Scoring criteria for matching cells between frames.

The final two criteria (distance and time) require further explanation. These criteria are based on the computation of fluid streamlines. For each frame in which a cell was successfully identified and tracked, the path of that cell is estimated by integrating the velocity vectors computed by the Flow Measurement code (chapter 4) to obtain a streamline. A fourth-order Runge-Kutta integration of the velocity is performed with initial conditions set to the cell centroid position. As the time between frames is accurately known, the position of the cell can be estimated for each successive frame:

$$\vec{l}(t_1) = \int_{t=0}^{t_1} \vec{u}(x, y) dt. \quad (6.17)$$

However, as this is a simple forward prediction for a massless particle, and assumes a steady-state flow, the accuracy of this streamline will reduce as the time since the streamline was last computed increases. In particular, the flow may be constantly changing as part of an on-going cell trap in response to control commands. Hence, the streamline is re-computed at every timestep while the FSM is in the **tracking** state.

Distance deviation and time deviation measure two components of the deviation from the estimated centreline (streamline). Distance deviation measures the distance, normal to the streamline, that the cell has strayed from the estimated streamline. A graphical representation of the distance deviation criteria is shown in figure 6.5, where $\sigma_D = 30$ px. Conceptually, a line is constructed between the streamline and the candidate cell that minimises the distance from the cell to the streamline, and is

6. Cell Tracking

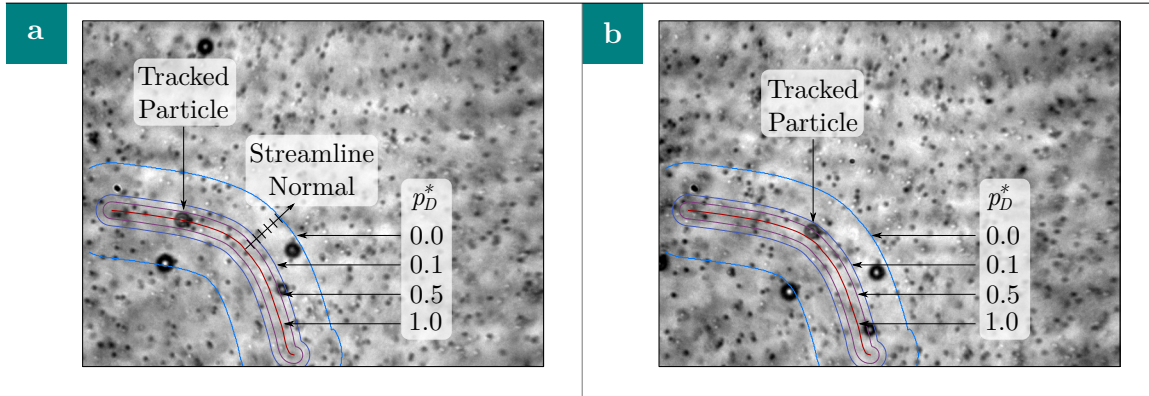


Figure 6.5: Streamlines for a typical particle and the associated relationship with the distance deviation score p_D . a) Initial position of a tracked particle, which originated at the far right of the image (at the start of the $p_D = 1.0$ contour), showing contour lines representing regions of equal p_D and a normal from the streamline (tick marks indicate 10 px distance increments along the normal). The centroid of the tracked particle (as indicated) lies on on the $p_D = 1.0$ contour, and hence is assigned the maximum possible score, representing a high probability of a match with the original particle. In subfigure b), time was advanced by 20 frames (0.1 s), at which point the particle is now on the $p_D = 0.5$ contour. The remaining particles in the image are further from the streamline, and are assigned lower scores. In normal operation, the streamline is computed at the new centroid after each successful match, which, in this case, would result in a score for the particle in b) of $p_D = 0.99$.

therefore normal to the streamline. Applying equation (6.15) to this minimum distance results in the distance deviation score.

On the other hand, time deviation measures the change in distance along the arc of the streamline. The streamline is a vector-valued function of the form $(x, y) = \vec{l}(t)$, hence increasing distance along the arc is directly related to time. The change in distance travelled along the streamline between frames, $t_c - t_r / t_r$, is therefore a measure of the match in the velocities of the reference cell and the candidate cell. As this quantity is indirectly based on a numerical derivative and is therefore based on the assumptions of a massless particle and steady flow, the time deviation is potentially noisy. To combat this, the time deviation score is biased, or offset, by 1.0 to ensure that a low score will not result in a complete rejection of a potential match. However, a higher score does indicate that a match between two particles is more likely.

6.3.3. Recovery State

The **recovery** state is designed to mitigate against the first instability mode discussed in the preface to this chapter: loss of tracking, leading to control system instability or complete loss of the cell. The software enters the recovery state from the tracking state if no successful match was found in the current frame. While in this state, there are two possible outcomes: the simplest is the software waits for a few frames, finds no further match, and transitions back to the **searching** state. At this point, the original cell is lost and a new one is sought. The second outcome is a recovery of the original cell and a continuation of tracking. To achieve this, the last known characteristics of the cell at the point where tracking was lost are stored, or frozen in time (normally, they are continually updated at each tracking step). The candidate cells are then compared with this frozen reference cell. The streamline is

considered to remain valid during this time, and the time deviation parameter is appropriately scaled based on the number of frames that have elapsed since tracking was lost:

$$p_{t_{recovery}}^* = 1 + p^* \left(\frac{t_c / N_{lost} - t_r}{t_r}, \sigma_t \right), \quad (6.18)$$

where N_{lost} is the number of elapsed frames. Effectively, this means that the new position of the particle is forward-projected based on its last known velocity. If a match is found that exceeds the minimum score, then the software transitions back to the **tracking** state and tracking resumes as before.

While this recovery step is very simple in concept, it is very effective at recovering from a number of common tracking errors. Early in the software development process, it was found that many of the tracks were aborted due to apparent dramatic changes in shape and intensity caused by other cells or tracer particles drifting in front of or behind the tracked cell. This caused an apparent rapid shape and intensity change, leading to loss of track. These abnormalities were only present in the image for a few frames — the flow velocities vary slightly in the out-of-plane direction, therefore causing these particles/cells to move past one another. Hence, by projecting the new position of the cell and assuming a constant shape and intensity, the system recovers from the temporary distortion in appearance and the original cell track can be correctly resumed. An example of this process can be seen in figure 6.6.

6.4. Performance

In order to ensure that each frame is processed within the available time budget, the problem needed to be parallelised across multiple CPU cores. In this instance, use of the GPU was not an option as all available GPU resources are dedicated to background subtraction (§6.1) and PIV (§4.2). Region-growing segmentation in a data-parallel algorithm is not straightforward²¹ as the image cannot be cleanly partitioned into independent sections. In a naïve approach, the image could be partitioned into a number of small sub-windows, and the cells that are wholly in each sub-window could then be segmented by separate threads of execution. However, where cells lie on the boundary between one or more sub-windows, the solution is less straightforward. One approach would be to allow the cell to be segmented multiple times and then the duplicate cells finally removed from the list of candidate cells in a post-processing step. This is potentially time-consuming, as depending on the locations of the cells in the image, several cells may need to be processed multiple times, and then duplicates identified. The duplicate identification step is potentially non-trivial as the segmented shape of the cell may differ slightly between sub-windows due to different choices of seed-points.

In the work conducted in this thesis, with the image sizes involved, it was found that the segmentation could be completed within the 5 ms time budget on a single thread, but some of the other processing could cause the algorithm to exceed its time budget, depending on the complexity of the image. table 6.2 shows typical timings extracted from a number of runs of the software. The critical processing, namely segmentation and tracking, is completed within 5 ms, but the timing budget is exceeded if the histogram and streamline computations are performed in series with the segmentation. Additionally, the segmentation & tracking code is dependent on the histogram and streamline processing, which appears to preclude parallelisation or re-ordering of the execution. However, it is acceptable for the streamline computation to be delayed by one or two frames, as the cell should not deviate significantly from the streamline over 10 ms. Secondly, the direct dependence on the histogram computation can be mitigated by using a scaling technique. The histogram, along with the average image intensity, can be computed for the preceding frame. In order to provide the segmentation and

6. Cell Tracking




Figure 6.6: Recovery of original track after interference due to a particle passing close to the tracked cell. a) Tracked cell (marked with shaded circle); b) loss of track as the secondary particle now appears to be connected to the tracked cell (the secondary particle is not physically connected to the cell, it only appears to be as a result of the imaging modality), \times marks the position of the new, combined segment, shaded circle marks the last known position of the cell; c) the cell moves vertically downward, but the two particles are still connected and no valid match can be found; d) the secondary particle is now sufficiently distant from the tracked cell that the track is recovered. Note that the centre of mass is offset slightly in this frame, as there is still some 'bleed' of the cell segmentation into the secondary particle. This is resolved in subsequent image frames as the secondary particle moves further away.

	Average Time (ms)	Maximum (99th percentile) (ms)
Segmentation & Tracking	2.4 ms	3.8 ms
Histogram Computation	2.7 ms	3.2 ms
Streamline	1.1 ms	1.9 ms
Total (serial processing)	6.2 ms	8.9 ms

Table 6.2: Typical timing information for the cell tracker code (on the reference system)

tracking code with the image intensity percentiles, the average for current frame is first computed (which is much quicker than a full histogram computation) and then used to scale the percentiles from the previous frame,

$$P[t] \approx P[t - 1] \frac{\bar{x}[t - 1]}{x[t]}, \quad (6.19)$$

where P is the percentile to be estimated and \bar{x} the image intensity average. This approximation holds as long as the shape of the distribution of intensities remains constant between images - i.e., the image content does not significantly change. Where the time between the current frame and previous frame is small, this assumption is valid and the scaling approach will effectively compensate for variations in the overall lighting intensity of the image.

With these changes to the algorithm dependencies, the entire process can be pipelined and therefore parallelised. The processing pipeline is illustrated graphically in figure 6.7, showing the interactions between the main application thread, the GPU and the tracker, histogram and streamline computation threads. As shown in the diagram, the tracking result is transmitted back to the main thread as soon as it is available (hence well within the 5 ms window), while the histogram and streamline computations can be performed in the background on other threads, and are not required until the next tracking frame. Therefore, by pipelining the algorithm, the various sub-steps can be split across multiple CPU- and GPU-cores, allowing the system to meet the 5 ms deadline imposed by the camera framerate and system specifications. Additionally, the main application thread is able to perform other operations, the most important of which is the realtime control loop and camera image timing while computation takes place in other threads.

6.5. Validation

Cell particle tracking is difficult to validate as a ground truth measurement is often unavailable or requires manual cross-checking. A number of papers reviewed in the development of these methods use a manual double-blind or single-blind style study to validate their methods^{1,8,27} or pre-existing methods,⁹ if they are validated at all. In validating this tracking code, a combination of synthetic image methods (common amongst PIV practitioners) and manual cross-checking was used.

6.5.1. Synthetic Image Benchmark

In order to test the ability of the cell tracking software to accurately reproduce a tracked path, a series of synthetic images were generated. The synthetic image generation software used was similar to that described in chapter 4, but with the following modifications. Particles were seeded randomly in the flow and integrated within the fluid domain using a method similar to the solution of the streamline integral described in section 6.3.2. These particles were randomly sized and composited

6. Cell Tracking

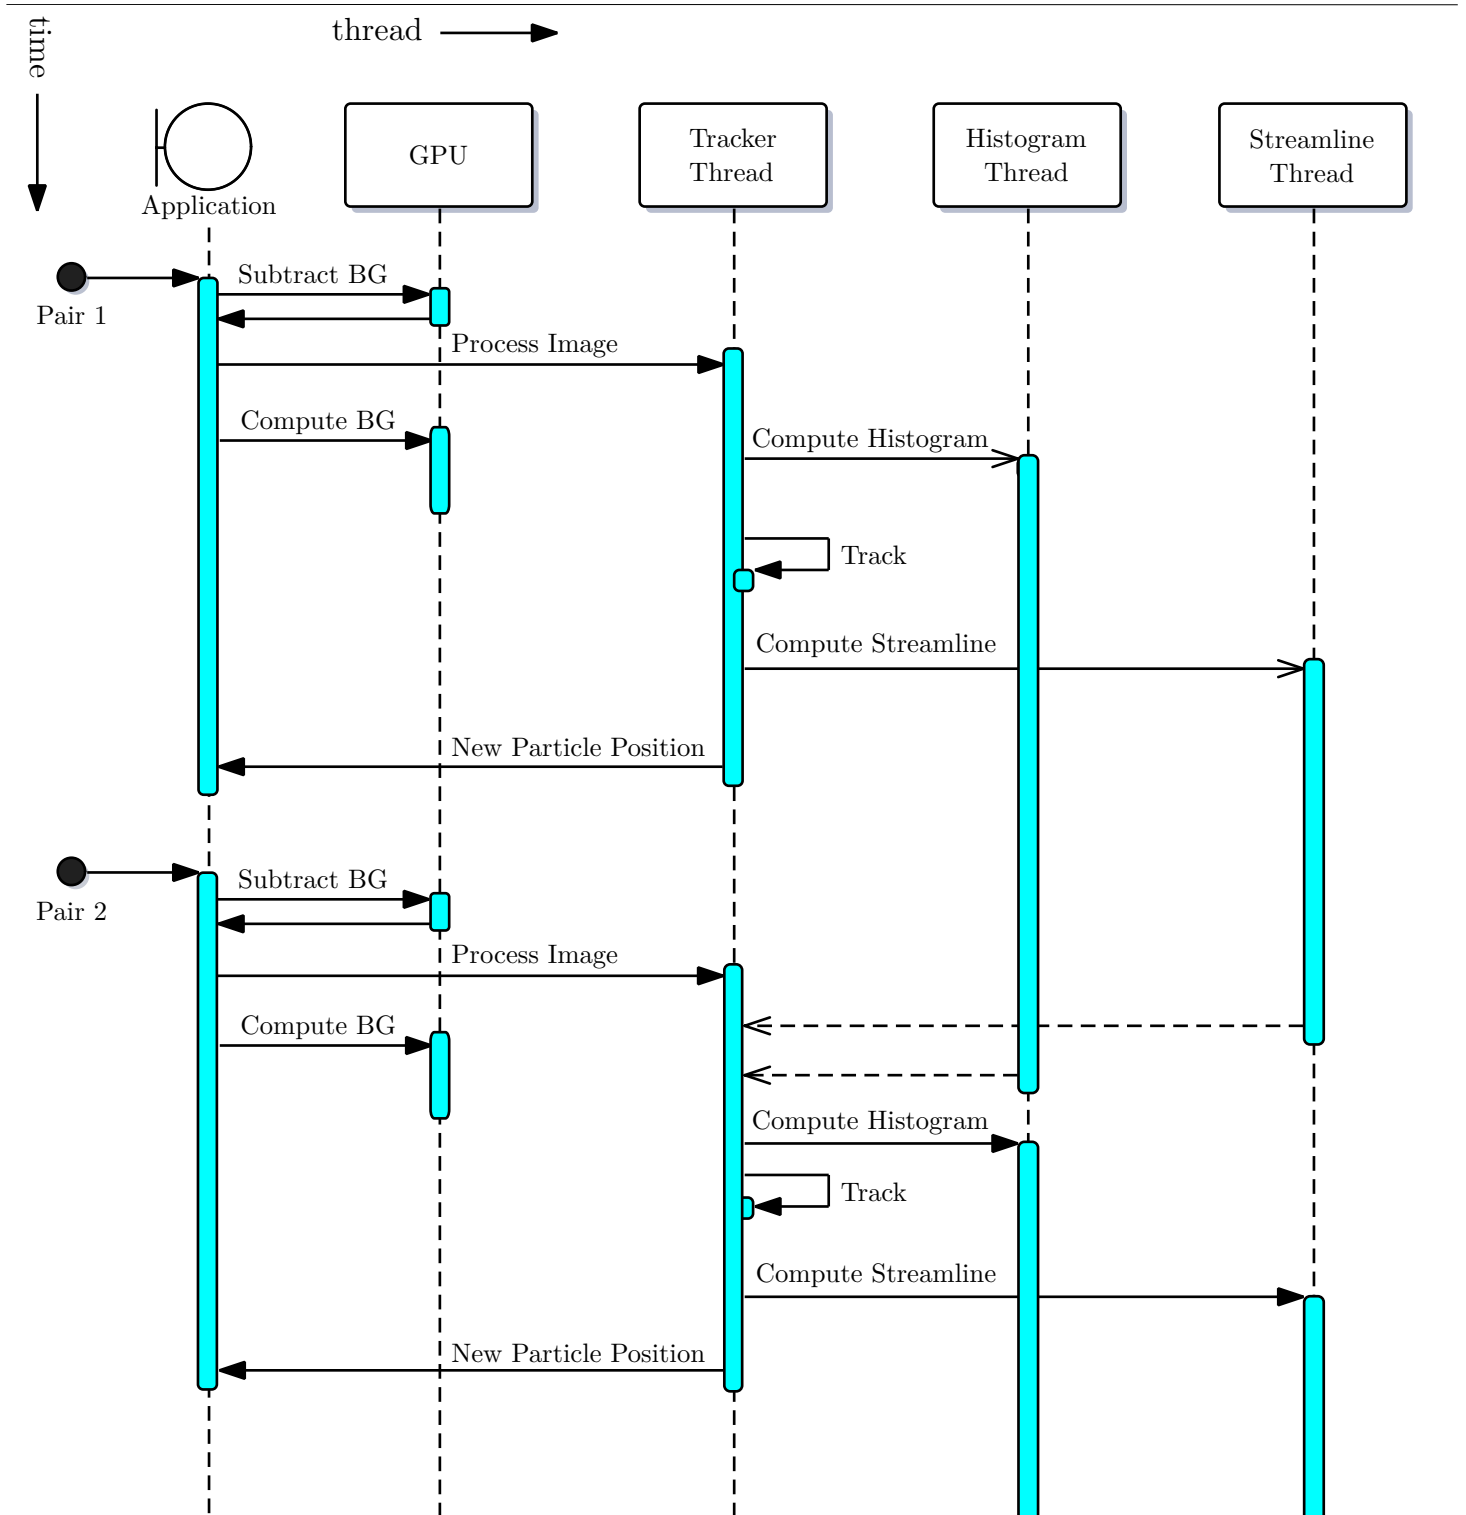


Figure 6.7: Thread occupancy diagram, showing the asynchronous nature of the streamline and histogram computations, relative to the tracker, and the asynchronous relationship between the main application and the tracker.

into the final image using a method similar to Olsen and Adrian.²⁰ A list of particle locations at each frame were output to act as the ground truth for later analysis and comparison. To simulate red blood cells in the flow, this process was carried out with a seeding density of 3×10^{-5} particles/px and a radius range of 13 px – 17 px. A seeding density of 1×10^{-3} particles/px and a radius range of 1.3 px – 2.9 px was used to simulate tracer particles. The resulting images were then composited with an experimentally-obtained background image (similar to figure 6.1b) using a linear combination.

This benchmark is an extremely difficult test for the tracking code, as the particles are largely of uniform size and intensity. Hence, the distance and time deviation terms become more important in the matching algorithms. Excluding failed tracks, the RMS average position error over all 250 particle tracks was < 0.82 px and the median of the errors for each track, treated separately, was 0.22 px. Approximately 48 % of particle tracks in the synthetic case resulted in a incorrect match - this demonstrates the importance of the size and intensity terms in differentiating between two nearby particles.

6.5.2. Blood Cell Testing

In order to assess the accuracy of the system with real images, recorded images of blood cells were replayed into the program. These images were of moderate quality and a number of cells were polluted with adherent tracer particles. 4,000 images were used for this test. The locations of the tracked cells were compared to the original images and assessed for accuracy. In total, 50 cell tracks were observed (more cells were present in the data, but only one cell is tracked at a time, so a number were ignored by the system). Of these, 9 cells merged or split with other cells or particle blobs (where a merge or split is purely a visual phenomenon and does not indicate a physical merge or split of the cells and particles) or drifted out of focus, and were excluded from the analysis. 7 cell tracks were improperly terminated before the cells left the field of view, and in 2 cases, the cell tracking algorithm incorrectly switched from one tracked cell to another. This translates to 32 cells out of 41 cells that were successfully and correctly tracked throughout the field of view, giving a correct-match rate of **78 %**. This compares favourably with other, non-realtime (which are more sophisticated but significantly slower) algorithms, where success rates of 80 % – 95 % are typically achieved.^{5,18}

6.6. Conclusion

In this chapter, a novel automated tracking system has been presented and evaluated in terms of both performance and accuracy. The tracking algorithms presented are uniquely capable of tracking biological cells against a noisy background in the presence of other cell-shaped objects (tracer particles). Additionally, this tracking is carried out fully automatically, with cells of interested automatically selected, segmented and tracked without human intervention. This is necessary, as at an imaging rate of 200 frames/s, with cells entering and leaving the field of view in less than a second, it is impractical for a human operator to interact with the system in realtime. Finally, using a combination of CPU and GPU processing, and multithreaded, pipelined analysis, the cell tracking system was able to accurately segment and track cells at a rate of 100 frames/s. This capability is necessary for reliable capture of cells, as will be seen in Part III.

6.7. Bibliography

- [1] Scott T. Acton, Klaus Wethmar and Klaus Ley. “Automatic Tracking of Rolling Leukocytes in Vivo”. In: *Microvasc. Res.* 63.1 (Jan. 2002), pp. 139–148. ISSN: 0026-2862. DOI: 10.1006/mvres.2001.2373.
- [2] Rolf Adams and Leanne Bischof. “Seeded region growing”. In: *IEEE Trans. Pattern Anal. Mach. Intell.* 16.6 (1994), pp. 641–647. ISSN: 0162-8828. DOI: 10.1109/34.295913.
- [3] B. Ahmad et al. “On-line tracking and stimulation of swimming microorganism by on-chip microrobot”. In: *2014 International Symposium on Micro-NanoMechatronics and Human Science (MHS)*. 2014 International Symposium on Micro-NanoMechatronics and Human Science (MHS). Nov. 2014, pp. 1–3. DOI: 10.1109/MHS.2014.7006078.
- [4] S. J. Baek and S. J. Lee. “A new two-frame particle tracking algorithm using match probability”. In: *Experiments in Fluids* 22.1 (Nov. 1996), pp. 23–32. ISSN: 0723-4864. DOI: 10.1007/BF01893303.
- [5] Nicholas D Cardwell, Pavlos P Vlachos and Karen A Thole. “A multi-parametric particle-pairing algorithm for particle tracking in single and multiphase flows”. In: *Meas. Sci. Technol.* 22.10 (1 Oct. 2011), p. 105406. ISSN: 0957-0233, 1361-6501. DOI: 10.1088/0957-0233/22/10/105406.
- [6] C. W. Clifford and K. Langley. “Recursive implementations of temporal filters for image motion computation”. In: *Biol Cybern* 82.5 (May 2000), pp. 383–390. ISSN: 0340-1200. pmid: 10836584.
- [7] E. A. Cowen and S. G. Monismith. “A hybrid digital particle tracking velocimetry technique”. In: *Exp. Fluids* 22.3 (1997), pp. 199–211.
- [8] O. Debeir et al. “Tracking of migrating cells under phase-contrast video microscopy with combined mean-shift processes”. In: *IEEE Trans. Med. Imaging* 24.6 (June 2005), pp. 697–711. ISSN: 0278-0062. DOI: 10.1109/TMI.2005.846851.
- [9] O. Dzyubachyk et al. “Advanced Level-Set-Based Cell Tracking in Time-Lapse Fluorescence Microscopy”. In: *Medical Imaging, IEEE Transactions on* 29.3 (2010), pp. 852–867. ISSN: 0278-0062. DOI: 10.1109/TMI.2009.2038693.
- [10] Jianping Fan et al. “Seeded region growing: an extensive and comparative study”. In: *Pattern Recognit. Lett.* 26.8 (June 2005), pp. 1139–1156. ISSN: 01678655. DOI: 10.1016/j.patrec.2004.10.010.
- [11] T. Hasegawa et al. “A new framework for microrobotic control of motile cells based on high-speed tracking and focusing”. In: *IEEE International Conference on Robotics and Automation, 2008. ICRA 2008*. IEEE International Conference on Robotics and Automation, 2008. ICRA 2008. May 2008, pp. 3964–3969. DOI: 10.1109/ROBOT.2008.4543820.
- [12] Yeon Hwa Kwak, Sung Min Hong and Soon Sup Park. “A single cell tracking system in real-time”. In: *Cell. Immunol.* 265.1 (2010), pp. 44–49. ISSN: 0008-8749. DOI: 10.1016/j.cellimm.2010.07.001.
- [13] Rui Lima et al. “Measurement of Individual Red Blood Cell Motions Under High Hematocrit Conditions Using a Confocal Micro-PTV System”. In: *Ann. Biomed. Eng.* 37.8 (2009), pp. 1546–1559. DOI: 10.1007/s10439-009-9732-z.
- [14] Xinyu Liu and Yu Sun. “Visually Servoed Orientation Control of Biological Cells in Microrobotic Cell Manipulation”. In: *Experimental Robotics*. 2009, pp. 179–187.
- [15] Andrew Mehnert and Paul Jackway. “An improved seeded region growing algorithm”. In: *Pattern Recognit. Lett.* 18.10 (1997), pp. 1065–1071.

- [16] Ivan Moreno and Ching-Cherng Sun. “Modeling the radiation pattern of LEDs”. In: *Opt. Express* 16.3 (4 Feb. 2008), pp. 1808–1819. DOI: 10.1364/OE.16.001808.
- [17] N. Ogawa et al. “Microrobotic visual control of motile cells using high-speed tracking system”. In: *IEEE Trans. Robot.* 21.4 (Aug. 2005), pp. 704–712. ISSN: 1552-3098. DOI: 10.1109/TR0.2005.844686.
- [18] Kazuo Ohmi and Hang-Yu Li. “Particle-tracking velocimetry with new algorithms”. In: *Meas. Sci. Technol.* 11.6 (2000), p. 603.
- [19] H. Oku, I. Ishii and M. Ishikawa. “Tracking a protozoon using high-speed visual feedback”. In: *Microtechnologies in Medicine and Biology, 1st Annual International, Conference On. 2000*. Microtechnologies in Medicine and Biology, 1st Annual International, Conference On. 2000. 2000, pp. 156–159. DOI: 10.1109/MMB.2000.893761.
- [20] M. G. Olsen and R. J. Adrian. “Out-of-focus effects on particle image visibility and correlation in microscopic particle image velocimetry”. In: *Experiments in Fluids* 29.7 (Dec. 2000), S166–S174. ISSN: 0723-4864. DOI: 10.1007/s003480070018.
- [21] M.A. Palomera-Perez et al. “Parallel Multiscale Feature Extraction and Region Growing: Application in Retinal Blood Vessel Detection”. In: *IEEE Trans. Inf. Technol. Biomed.* 14.2 (Mar. 2010), pp. 500–506. ISSN: 1089-7771. DOI: 10.1109/TITB.2009.2036604.
- [22] E Perez-Careta et al. “Cell Tracking by Normalized Cross-Correlation with Image Processing”. In: *2008 Digest of the IEEE/LEOS Summer Topical Meetings*. 21 July 2008. DOI: 10.1109/LEOSST.2008.4590480.
- [23] Todd S Perlstein et al. “Red blood cell distribution width and mortality risk in a community-based prospective cohort: NHANES III”. In: *Arch Intern Med* 169.6 (23 Mar. 2009), pp. 588–594. ISSN: 0003-9926. DOI: 10.1001/archinternmed.2009.55.
- [24] Roel Pieters, Pieter Jonker and Henk Nijmeijer. “Trajectory Generation for 1000 fps Direct Visual Servoing.” In: *MVA*. Citeseer, 2011, pp. 39–42.
- [25] Wayne S. Rasband. *ImageJ*. (1997–2015). U. URL: <http://imagej.nih.gov/ij/>.
- [26] C. C. Reyes-Aldasoro, S. Akerman and G. M. Tozer. “Measuring the velocity of fluorescently labelled red blood cells with a keyhole tracking algorithm”. In: *J. Microsc.* 229.1 (2007), pp. 162–173.
- [27] Y.-N. Sun et al. “Live cell tracking based on cellular state recognition from microscopic images”. In: *J. Microsc.* 235.1 (2009), pp. 94–105. DOI: 10.1111/j.1365-2818.2009.03186.x.
- [28] Z. Wang et al. “Cell detection and tracking for micromanipulation vision system of cell-operation robot”. In: *Systems, Man, and Cybernetics, 2000 IEEE International Conference on*. Systems, Man, and Cybernetics, 2000 IEEE International Conference on. Vol. 3. 2000, 1592–1597 vol.3. ISBN: 1062-922X. DOI: 10.1109/ICSMC.2000.886249.
- [29] Y. Watanabe, T. Komuro and M. Ishikawa. “955-fps Real-time Shape Measurement of a Moving/Deforming Object using High-speed Vision for Numerous-point Analysis”. In: *2007 IEEE International Conference on Robotics and Automation*. 2007 IEEE International Conference on Robotics and Automation. Apr. 2007, pp. 3192–3197. DOI: 10.1109/ROBOT.2007.363965.

Part III.

Application and Experimental Results

Eureka!

I have found it!

— Archimedes

7. Trapping Particles and Cells

In this chapter, the techniques, algorithms and research described in chapters 2–6 are deployed in support of the final goal of trapping a cell. To this, the realtime flow measurement techniques (chapter 4) are used for system calibration and as inputs to the trapping and control algorithms; the flow control research (chapter 5) performs an integral role in steering the saddle-point to create and maintain a stable trap; the cell tracking algorithms (chapter 6) are used to determine the cell location in realtime — necessary prerequisite for trapping; and the simulation (chapter 3) system has been used to develop and inform the control system. Finally, the entire system is supported by innovative mechanical systems; optical systems; electrical subsystems; and software that have been developed during the course of researching this thesis (chapter 2).

7.1. Trapping Particles

Initial work involved $7\ \mu\text{m}$ hollow glass spheres, chosen as red blood cell analogues, in terms of size and projected 2D shape. A red blood cell is a biconcave disc of average diameter $7.4\ \mu\text{m}$, and when undeformed and oriented with the long axis parallel to the imaging plane, appears as a circular disc in 2D microscopy (*i.e.*, in 2D, the appearance of an undeformed, non-tumbling red blood cell is the same as that of a sphere of equivalent radius). Sample microscopy images are shown in figure 7.1. The use of hollow glass spheres is convenient as they are readily available, have a very long shelf-life, require a minimum of handling precautions and simple clean up.

As explained in section §1.1, trapping is achieved by continuously adjusting the flow rates into and out of the device to maintain a restoring force on the particle (cell) that drives the cell toward the centre of the field of view (and by extension, the centre of the device). This principle is seen clearly in the following experiment.

The $7\ \mu\text{m}$ hollow glass spheres are mixed in a aqueous solution of $d = 1\ \mu\text{m}$ polystyrene tracer particles and a mild detergent to enhance particle solubility. The working fluid was transferred to the system inlet reservoirs and the flow loop was purged of all remaining air. An automated flow control calibration was performed and the flow control code enabled. Once a steady stream of large particles




Figure 7.1: Comparison between microscopy images for a) a hollow glass sphere (red blood cell analogue); b) red blood cell

7. Trapping Particles and Cells

was visible in the images, the cell trapping control code was enabled. The system began to identify and subsequently trap these large particles and a number of images and telemetry data were saved for further analysis. Telemetry and images (with superimposed PIV streamlines) for one such trap are shown in figure 7.2. The control principles involved in forming and maintaining a stable trap are evident in the images in the upper part of figure 7.2. At $t = 0.60$ s, the particle is located $15\ \mu\text{m}$ above the centre of the field of view. To direct the particle back towards the centre, the saddle-point position is modified, by means of adjusting the relative flow rates into and out of the device, to $100\ \mu\text{m}$ above the centreline. This results in a net force on the particle that causes it to accelerate back towards the centre (as indicated by the fluid streamlines). Similarly, at $t = 0.75$ s, the cell is directed in the opposite direction, back towards the centre, by moving the saddle-point beyond the particle location. Finally, a stable trap is achieved and the saddle-point position and the particle remains stable at the centre of the device ($t = 4.40$ s).

The trapping result in figure 7.2 is typical for a captured particle at an average velocity of $\bar{v} = 0.15\ \text{mms}^{-1}$ and was achieved with a minimal amount of tuning of the cell control loop. A simple proportional controller is sufficient to stably trap a cell at these flow velocities and, as can be seen from the figure, the trap can be maintained for quite some time. The extensional rate $\dot{\epsilon}$ at this velocity is $\dot{\epsilon} < 2\ \text{s}^{-1}$. By way of comparison to a recent paper on a similar topic (Shenoy, Tanyeri and Schroeder⁸), their presented experimental results indicate a maximum reported extensional rate of $\dot{\epsilon} = 0.26\ \text{s}^{-1}$. However, empirical evidence suggests that this extensional rate is not sufficient to cause significant deformation of cells.⁷ Hence, higher velocities are required, which will decrease the stability of the system as the required response time approaches the control loop rate. These faster velocities therefore require a much more detailed treatment of the control system tuning.

7.2. Simulation

The system, as built, exhibits a number of behaviours which are not completely modelled by the simulation software developed in chapter 3. In order to obtain accurate simulations of the system behaviour, and therefore iterate more quickly to determine optimal parameters and test control strategies, the simulator was augmented with a number of new capabilities. These are: a model of the system delay, more realistic flow and noise models, and simulation of the hysteresis inherent in the flow control (see chapter 5). These will be described in more detail in the following sections.

System Delay There are a number of unavoidable delays in most digital systems, and the cross-slot is no exception. The typical delays from the initial event (cell motion) to response (valve motion) are summarised in figure 7.3. The camera exposure time is fixed and takes 5 ms; once exposure is complete, readout and transmission from camera to the computer begins, which also takes about $5\ \text{ms}^\dagger$; the image is then available in the control software's memory and the cell tracking algorithm starts processing the frame - in most cases, this is complete within 5 ms (see section §6.4) and the control system parameters (including valve positions) are updated close to the beginning of the next control loop. In the ideal case, the control system response is computed based on an event that occurred 15 ms earlier — a delay of three frames. As the cell tracking algorithm is only executed on every second frame (the double-pulsed illumination means that little additional information would be gained in executing

[†]In fact, exposure times vary depending on whether the frame is the first or second in the imaging pair, and the readout time is slightly less than 5 ms. However, the overall repetition rate between the start of the first frame in a pair, and the start of the first frame in the next pair is 10 ms. Hence, for the purposes of this analysis, it is sufficient to approximate the exposure and readout times as 5 ms each.

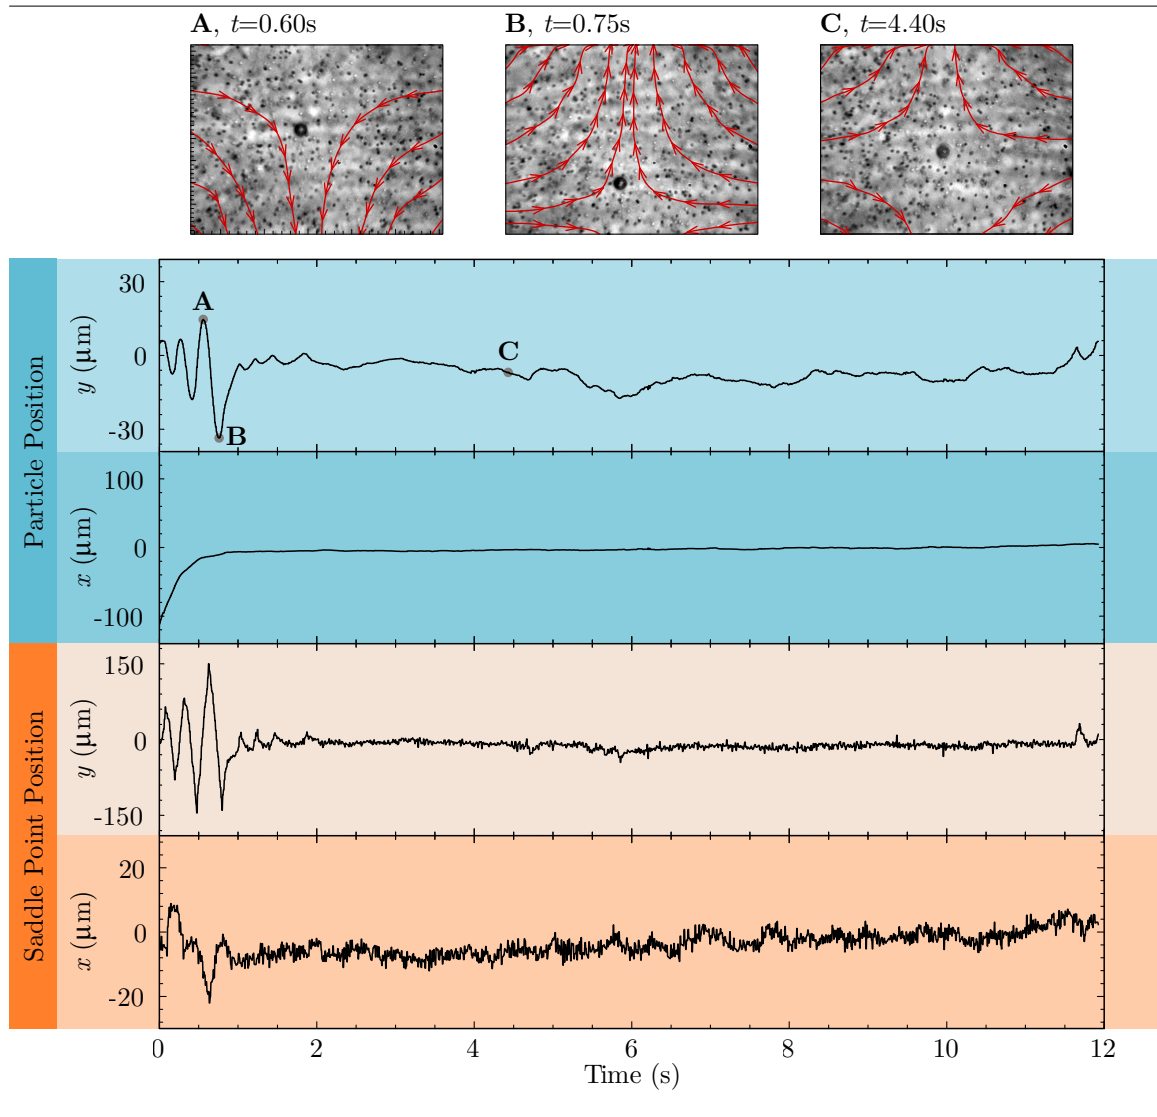


Figure 7.2: Particle and saddle-point positions during an active trap with an average flow velocity of $\bar{v} = 0.15 \text{ mm s}^{-1}$. There is some initial oscillation due to the inherent lag of the system, after which the particle is trapped in a stiff trap in the centre of the microchannel, with some small corrections required to maintain the position. Images and fluid velocity streamlines (top) show the position of the particle relative to the flow field after 0.60 s, 0.75 s and 4.40 s, respectively.

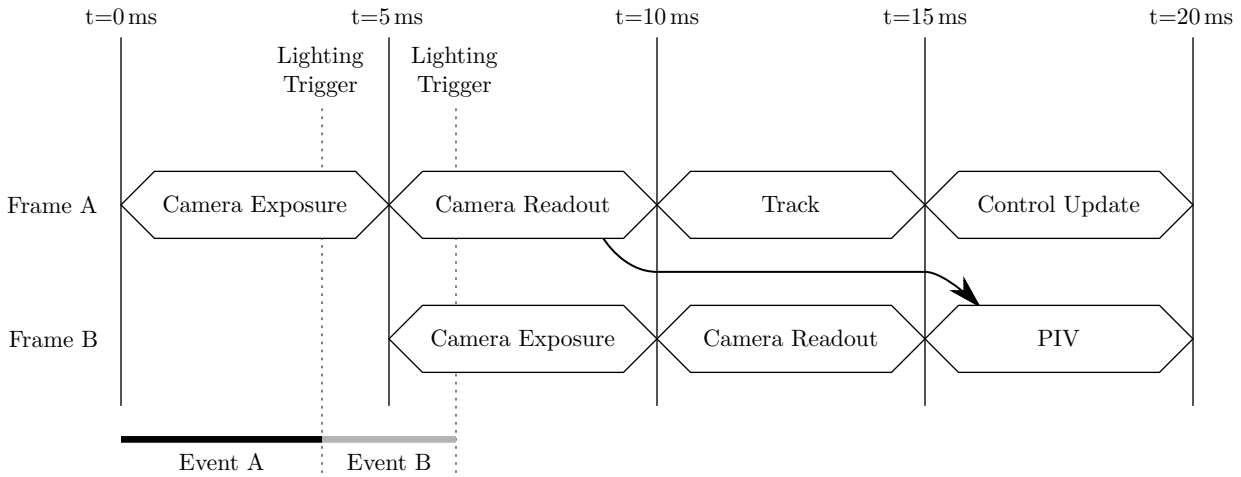


Figure 7.3: Typical processing pipeline, showing the data flow and inherent delays in the system that determine the minimum time that the system can respond to an event (change in position of the cell). Two frames are shown to indicate the continuous and alternating nature of the system control loop.

the tracking algorithm for every frame), an additional delay of 2 frames is introduced, resulting in best-case delays of between 15 ms and 25 ms.

This delay can be simply modelled in the form of a small ring buffer between the simulated cell position and the controller,

$$cell_{controller}[t] = cell_{simulated}[t]Z^{-D}, \quad (7.1)$$

where D is chosen to represent the average the frame delay. For this analysis, a conservative value of $D = 3$ was chosen.

Flow and Noise Models Due to a number of factors, including: variations in effective seeding density between frames; systemic errors in the PIV and saddle-point identification; and variations in the fluid system, the measured position of the saddle-point varies considerably over time. Typical experimental saddle-point position noise is shown in figure 7.4. For the data presented in figure 7.4, the standard deviation is $\sigma = 2.2 \mu\text{m}$ (or $\sigma = 5.9 \text{ px}$), which is significantly greater than the standard deviation that can be attributed to the measurement method alone (where experiments with synthetic data suggest the standard deviation due to measurement error is $\sigma = 2.9 \text{ px}$ — see section 4.4). This noise can be primarily attributed to noise in the pressure controller propagating into the system (as seen in section §5.4, changes in pressure have a direct effect on the saddle-point position) and variations in seeding density on a frame-by-frame basis.

To simulate the noise in the measured saddle-point position, Gaussian-distributed white noise was generated and inserted in the simulation in addition to the actuator digitisation noise. This noise was bandwidth-limited by decimating the output noise generator and interpolating between output points (a simple low-pass filter). Imposing noise at this point in the simulation means that this random noise affects the position of the simulated cell and, as such, simulates the pressure-controller induced noise. The remaining components of the noise can be adequately simulated as observer (rather than flow)




Figure 7.4: Typical flow (saddle-point position) noise while the system is idle.

noise by the unmodified simulator from chapter 3.

Hysteresis and Overshoot As discussed extensively in chapter 5, the valves exhibit hysteresis. Furthermore, the transition through the hysteretic region results in an amount of overshoot, which is related to the phenomenon of pumping due to valve motion (see section 5.2.2.4). Hysteresis can be simply modelled with a backlash-like hidden (or unobservable) state variable, which is updated each time step

$$h[t] = h[t - 1] + \frac{T}{k_r} \text{sign}(\Delta I[t]), \quad (7.2)$$

where $h[t]$ is the state variable, $I[t]$ the input vector, T the update period of the model, and k_r a rate variable that determines the time taken to transition through the hysteretic region. $h[t]$ is clamped between $-1.0 \dots 1.0$, where -1.0 indicates that the system is following the negative-going hysteresis curve, and 1.0 the positive-going curve. Values between these extremes indicate that the system is in the hysteretic region, and there are no observable changes to the output vector $O[t]$

$$O[t] = p[t] + \begin{cases} I[t] & h[t] = 1.0 \\ I[t] & h[t] = -1.0, \\ \text{unchanged} & \text{otherwise} \end{cases}, \quad (7.3)$$

where $p[t]$ is an internal variable which accounts for the disturbance of the saddle-point position as a result of pumping.

To model overshoot, or pumping, $p[t]$ is initialised to a constant k_p whenever $h[t]$ reaches either of its extremes (*i.e.*, whenever a change in direction has occurred). In all other cases, it is exponentially decayed

7. Trapping Particles and Cells

$$p[t] = \begin{cases} k_p & h[t] = 1.0 \text{ and } h[t-1] \neq 1.0 \\ -k_p & h[t] = -1.0 \text{ and } h[t-1] \neq -1.0 \\ p[t-1] & \text{time since last change to } p < k_h \\ 0 & \text{otherwise} \end{cases} \quad (7.4)$$

where k_h effectively becomes a hold-time constant, and models the time taken for the system to return to steady-state after the imposed impulse. In general, this will depend on the rate of change of the valves and the size of the excursion, but a single constant was found to be sufficient for modelling purposes in this case.

Validation To validate the model, a series of step changes were imposed on the saddle-point position by means of manual control override functionality built into the software. As these step changes are executed by the control software, these step change experiments measure the response of the physical system *as modified by* the flow controller. By determining the system parameters from this data, the physical and control systems are effectively modelled as a lumped system. As the purpose of these simulations is to optimise the cell trapping controller, assuming that the flow controller is already optimal, this simplifies the modelling with no loss of fidelity.

As shown in figure 7.5, a $10 \mu\text{m}$ step change was introduced to the experimental system in both the positive direction and the negative direction (that is, the sequence of commands was $0 \mu\text{m}$, $10 \mu\text{m}$, $0 \mu\text{m}$). The choice of step size was chosen to be small enough to avoid a saturation condition at low pressures, where the saddle-point position moves out of the field of view due to overshoot. Analysis of larger step sizes suggest that doubling the step (and hence the excursion of the saddle-point position) results in a 5% increase in overshoot at the pressures of interest, so the choice of step size can be made somewhat arbitrarily. To match the conditions that are typically used during cell trapping, the net driving pressure was set to 2 kPa, which resulted in an average flow velocity of $\bar{v} = 0.55 \text{ mms}^{-1}$.

Careful measurement was made of the noise, inherent delay, rise times and hysteresis overshoot from several experimental datasets and this was then translated into simulation parameters. Some empirical tuning of the parameters was then necessary to obtain a good match between observed and simulated results. The final parameters used in the simulation are shown in table 7.1, along with the original parameters used in Curtis, Sheard and Fouras² for comparison. The simulation was also run with an identical step-change in the saddle-point position imposed on the control system, the results of which are shown in the latter half of figure 7.5, which demonstrates excellent agreement between the experimental and simulated steps.

The benefits of this process are twofold: these parameters can now be used in subsequent simulations to optimise the controller; and, the whole process is a form of system identification for non-linear systems, where the resultant parameters can be used to infer characteristics of the underlying system. Of particular interest are the overshoot (k_p) and hold time (k_h) parameters (the rate parameter k_r is of lesser interest, as it is primarily a function of the width of the hysteretic region and the acceleration and maximum velocity parameters of the motion planner section §5.5). These parameters can be used inside the flow controller to compensate (and eliminate) the overshoot due to the hysteresis transition by making a smaller initial step (reduced by k_p) and then making a second step once the system has moved out of the hysteretic step (once $t > t_0 + k_h$, where t_0 is the time at which the original command was given). This is effective, but requires that these parameters are known *a priori* by the controller. As these parameters are weakly dependent on operating conditions (in particular, system pressure), k_p and k_h would need to be determined automatically and online.




Figure 7.5: Hysteresis and overshoot simulation, with experimental data as comparison. The experimental data was gathered at a net driving pressure of 2 kPa, and both positive and negative steps were imposed on the control system. Data from the simulation shows good agreement with the experimental system behaviour, indicating that the simulation is a good model of system behaviour.

7. Trapping Particles and Cells

Simulation Parameter		New Value	Typical value from §3.2
Imaging (Control Update) Rate		10 ms	5 ms (200 fps)
First-Order Time Constant	τ_f	0.01	10
	(§3.2 eqn 7)		
Observer Resolution		$< 0.37 \mu\text{m}$ ($< 1 \text{ px}$)	$0.1 \mu\text{m} - 1.0 \mu\text{m}$
Actuator Resolution		$0.5 \mu\text{m}$	$0.5 \mu\text{m}$
Flow Noise	σ_n	$0.2 \mu\text{m}$	—
System Delay	D	25 ms	—
Hysteresis Rate	k_h	25 ms	—
Hysteresis Overshoot	k_p	$4.0 \mu\text{m}$	—
Overshoot Hold Time	k_d	75 ms	—
Noise Bandwidth	B_n	200 Hz	—

Table 7.1: Simulation parameters chosen based on the experimentally observed behaviour, with the typical values from Curtis, Sheard and Fouras² included for reference. Dashed values indicate parameters which were not implemented in the original simulator.

7.3. Tuning the System for Higher Flow Velocities

Using the enhanced control systems simulator developed in section §7.2, the control system can now be optimised *in silico* for higher flow velocities. For this work, the range of velocities of interest were between $\bar{v} = 0.50 \text{ mms}^{-1}$ and $\bar{v} = 1.00 \text{ mms}^{-1}$. This corresponds to extensional rates of between $\dot{\epsilon} = 5 \text{ s}^{-1}$ and $\dot{\epsilon} = 12 \text{ s}^{-1}$.

The results of increasing the flow velocity, while making no other changes to the control system parameters, are shown in figure 7.6. Only the y -component of the cell positions are shown, as this direction is the least stable and most important to optimise. As can be seen in both the experimental and simulation results, the system traps the cell for $< 1 \text{ s}$ and then becomes unstable due to oscillations that grow rapidly, eventually resulting in loss of the cell.

One of the difficulties encountered at higher flow rates is that the cell can move a significant distance in the time between the initial image capture, subsequent analysis and the eventual control update. This, combined with the maximum stable rates of change enforced by the flow controller, means that the cell-trap system can quickly grow unstable as the control system does not respond fast enough to the rapidly accelerating cell. Simply increasing the control gains does not necessarily improve matters, as this means that the system will respond to small oscillations during a stable trap with a large control response, driving the system into instability. This is a classic non-linear problem caused by the extensional nature of the flow in the y -direction. To mitigate this, the system can be partially linearised by using a simple forward prediction. In chapter 6, a method was developed that used fluid streamlines to forward-predict the position of the cell to aid in track matching. This type of forward prediction can be also employed by the cell trapping controller to determine the likely position of the cell at the moment the control update will be applied

$$\vec{c}_c = \vec{l}(t_0 + t_d), \quad (7.5)$$

where \vec{c}_c is the compensated position of the cell, $\vec{l}(t)$ is a vector-valued function representing the fluid streamline, t_0 is the time at which the most recent image was taken and t_d the measurement delay




Figure 7.6: Cell Positions in the y -plane from the simulation and corresponding experiments of a cell trap with an average flow velocity of $\bar{v} = 0.50 \text{ mms}^{-1}$ ($\dot{\epsilon} = 5 \text{ s}^{-1}$). Positions in simulation inverted for ease of comparison.

7. Trapping Particles and Cells

Parameter		Original Simulation (chapter 3)	Sub-optimal (Figure 7.6)	Stable (Figure 7.7)
<i>x</i>-direction				
Proportional Gain	K	6.0	—	2.0
Differential Gain	K_D	0.25	—	0.1
Integral Gain	K_I	0	—	2.0
<i>y</i>-direction				
Proportional Gain	K	6.0	1.1	2.0
Differential Gain	K_D	0.2	0	0.1
Integral Gain	K_I	0	0	2.0

Table 7.2: Summary of control gains for the original simulation work (chapter 3, Curtis, Sheard and Fouras²); the sub-optimal trapping configuration (Figure 7.6); and the stable trapping configuration (Figure 7.7). For consistency with the later work, the proportional gains from the paper (which defined the gains in terms of flow rates) have been scaled to match the later definition (where the the gains are defined in terms of saddle-point position).

estimate. As can be inferred from figure 7.3, a suitable value for t_d is on the order of 25 ms. By applying this correction, the controller will exhibit small control responses during a stable trap, and much larger responses to cell motion further from the saddle-point, as required.

By combining forward-prediction with further optimisation using the enhanced simulator, the stability of the cell-trap system can be dramatically improved. Figure 7.7 shows the control response of the cell trap at $\bar{v} = 0.75 \text{ mms}^{-1}$ after optimisation. The simulated and experimental control responses are similar in the amplitude of oscillations, but differ in a number of other respects. The simulated response is more consistent, and underdamped compared to the small oscillations in the experimental response. However, the experimental system exhibits a number of periods of drift, where the cell moves away from the set-point and the system becomes more oscillatory in nature, before the integral gain component of the control equations restore the cell position to the set-point. The simulation also exhibits this type of behaviour (in particular around $t = 2 \text{ s}$ and $t = 6 \text{ s}$), and the RMS errors in the simulation and experimental cases are identical ($12.36 \mu\text{m}$). This demonstrates that, while there is a number of unmodelled behaviours that could be added to the simulator, the predictive power of the simulation is sufficient to guide optimisation.

Of particular note is the maximum amplitude in the experimental case is $\approx 30 \mu\text{m}$, which is similar to the maximum amplitude in the case in figure 7.2, where the average velocity was $5\times$ slower than that in figure 7.7 ($\bar{v} = 0.15 \text{ mms}^{-1}$ as compared to $\bar{v} = 0.75 \text{ mms}^{-1}$). This shows the importance of optimisation to maximise system stability — similar, although more oscillatory, responses can be obtained at much higher velocities after gain optimisation.

The final control gains chosen for the cell-trap system are shown in table 7.2, along with the optimal gains from the original simulation work (chapter 3) and the earlier unstable trap (figure 7.6). It should be noted that the gains chosen in the unstable case were extremely conservative — this is because initial optimisation work was performed online during experiments, which is extremely expensive (both in terms of time and consumables). Further development work on the simulator allowed rapid convergence on optimal gains (figure 7.7).




Figure 7.7: Trapping experiment and corresponding simulation, after tuning, with an average flow velocity of $\bar{v} = 0.75 \text{ mms}^{-1}$ ($\dot{\epsilon} = 10 \text{ s}^{-1}$)

7.4. Extended Manipulation

Much of this thesis has focussed on trapping cells — that is, the deformation, or other behaviour exhibited by the cells as a result of them remaining trapped near the stagnation point is the primary characteristic of interest. However, the creation of a stable trap opens up a number of other possibilities in terms of manipulation that go beyond the static observation of changes during trapping. As described in the simulation paper (section §3.2), a trapped cell could be analysed and subsequently sorted to one of the two outlets, or the trap used to stabilise the cell for a secondary, high resolution, imaging device. This concept can be further developed to enable other, more dynamic behaviours. The system, as built, has complete control over the fluid environment. The net pressure on the system can be manipulated in realtime, in addition to the relative flow rates. By increasing the net system pressure after the trap has stabilised, the extensional forces on the cell can be increased in a controlled manner: both the rate of change, and the magnitude of the change, can be controlled by the system. This ability to introduce a controlled change in extensional forces in a software-defined, realtime manner, enables new types of investigations, such as: studying the behaviour of cells as a result of shear histories (hypothesised by some to play a role in platelet activation⁶); measuring cell viscosities via forced oscillation or step changes (system identification⁴); and studying the non-linear relationship between extensional force and deformation on a cell-by-cell basis (a phenomena investigated in a bulk context in cell rheology studies, e.g. Dobbe et al., Kiss et al.^{3,5}).

An example of extended manipulation during trapping can be seen in figure 7.8. In this experiment, a red blood cell was trapped and subsequently analysed for shape changes using techniques and code that will be described in more detail in chapter 8. For the purposes of this discussion, the red blood

7. Trapping Particles and Cells

cell was suspended in the flow in a phosphate-buffered saline solution, and shape measurements were acquired by fitting an ellipsoid to the detected edges points on the cell. This yields both cell orientation and major and minor radii at each time step. In the literature on red blood cell deformation, the relationship between the two radii is commonly expressed as a dimensionless number known as the Taylor deformation index (after G.I. Taylor⁹)

$$\text{DI} = \frac{r_{\text{major}} - r_{\text{minor}}}{r_{\text{major}} + r_{\text{minor}}}, \quad (7.6)$$

which has the convenient property that a perfectly circular cell or object has a deformation index of 0, and has a maximum value of 1 (which corresponds to a degenerate ellipse — an infinitely thin line).

The trap in figure 7.8 can be divided into five main phases. In phase **A**, the cell re-orientates such that the major axis rotates from parallel to the inlet ($\theta = \pi$) to parallel to the outlet ($\theta = \pi/2$). In phase **B**, the cell is held steadily trapped and is oriented with the long axis parallel to the axis of maximum extensional force, and hence the cell begins to deform. In phase **C**, no appreciable change in deformation takes place, as the cell oscillates in the trap but is not exposed to increased forces. In phase **D**, a small amount of deformation is observed as the trap simultaneously becomes stable, and the average velocity increases slightly (this velocity increase is a side-effect of the controller rapidly changing the relative flow rates). Finally, in phase **E**, a velocity ramp is imposed on the cell and begins to further deform. This final deformation is limited by the instability of the trap at these higher velocities.

This simple experiment reveals a number of interesting properties about the red blood cell. Of note is that even though the red blood cell effectively undergoes relaxation in phase **C**, due to the decreased stability of the trap, there is no decrease in deformation (one exception is the small decrease at $t = 2$ s, but the cell quickly returns to the steady-state deformation for phase **C**). Additionally, there is substantial deformation in phase **D** even though the forces are applied in a pulsed fashion. These are all indicators that the cell shape is hysteretic with respect to applied force — the force required to cause the initial deformation of the cell is much greater than the force at which relaxation, or return to the original shape, occurs. This behaviour has also been observed by others in atomic force microscopy experiments with red blood cells.¹ Further, the relatively slow rate of change of deformation in phase **B** and the fact that the change in deformation appears to lag behind the application of the deforming force suggests that the cell is indeed viscoelastic (as confirmed in a number of other experimental modalities^{1,10,11}), and that these viscoelastic properties are measurable using this technique.

7.5. Conclusion

In this chapter, it has been shown how the various components of the system developed in this thesis have been combined to form a greater whole. In so doing, the mechanics of cells have been studied for the first time in automated, purely fluidic traps at physiologically-comparable extensional rates. The realtime image analysis system, comprised of flow measurement and cell tracking, coupled with an advanced flow controller designed specifically for microfluidics at low flow rates and high response rate, allows the conception and construction of software-defined systems to manipulate cells as well as oil droplets, particles, and other suspended objects. The types of experiments and systems enabled by this research include, but are by no means limited to, trapping for subsequent deformation measurement; trapping for assessment followed by sorting; and high-speed software-controlled sorting systems. By combining several control mechanisms, complex behaviours can be created. By coordinating the system overall flow rate and flow control subsystems, time-varying shear can be imposed on cells to assess the response of cells to varying environmental stimuli. As these behaviours are implemented




Figure 7.8: Extended manipulation of a cell during trapping. Deformation (solid line shows a smoothed spline fit to the data; dots show the raw data) and orientation measurements were obtained using code that will be described in detail in section §8.2 . Cell Position is shown decomposed into the high frequency (—) and low frequency components (- - -) to highlight the change in controller stability over time. Average flow velocity (solid line and dots as for deformation) is an indicator of the change in extensional rate over time. Extensional rate at the end of the velocity ramp is $\dot{\epsilon} = 17 \text{ s}^{-1}$.

in software, new test protocols can be rapidly developed, validated, and deployed on a large scale. Additionally, the basic hardware platform can support a number of potential uses, with the software defining the type of experiment performed.

7.6. Bibliography

- [1] Kristen E. Bremmell, Allan Evans and Clive A. Prestidge. “Deformation and nano-rheology of red blood cells: An AFM investigation”. In: *Colloids and Surfaces B: Biointerfaces* 50.1 (1 June 2006), pp. 43–48. ISSN: 0927-7765. DOI: 10.1016/j.colsurfb.2006.03.002.
- [2] Michael D. Curtis, Gregory J. Sheard and Andreas Fouras. “Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems”. In: *Lab Chip* 11.14 (28 June 2011), pp. 2343–2351. ISSN: 1473-0189. DOI: 10.1039/C1LC20191C.
- [3] J. G. G. Dobbe et al. “Analyzing Red Blood Cell-Deformability Distributions”. In: *Blood Cells. Mol. Dis.* 28.3 (May 2002), pp. 373–384. ISSN: 10799796. DOI: 10.1006/bcmd.2002.0528.
- [4] Gene F Franklin, J. David Powell and Michael L Workman. *Digital control of dynamic systems*. Half-Moon Bay, CA: Ellis-Kagle Press, 1998. ISBN: 978-0-9791226-1-3 0-9791226-1-9.
- [5] Ferenc Kiss et al. “The power of slit-flow ektacytometry measurements for testing normal and heat treated red blood cells using various viscosity media in laboratory animals”. In: *Korea-Aust. Rheol. J.* 22.1 (2010), pp. 81–86.
- [6] Warwick S. Nesbitt et al. “A shear gradient-dependent platelet aggregation mechanism drives thrombus formation”. In: *Nat Med* 15.6 (June 2009), pp. 665–673. ISSN: 1078-8956. DOI: 10.1038/nm.1955.
- [7] Nathalie Nève et al. “Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers”. In: *Cel. Mol. Bioeng.* 3.3 (Mar. 2010), pp. 213–228. ISSN: 1865-5025. DOI: 10.1007/s12195-010-0113-3.
- [8] Anish Shenoy, Melikhan Tanyeri and Charles M. Schroeder. “Characterizing the performance of the hydrodynamic trap using a control-based approach”. In: *Microfluid Nanofluid* (9 Oct. 2014), pp. 1–12. ISSN: 1613-4982, 1613-4990. DOI: 10.1007/s10404-014-1495-7.
- [9] G. I. Taylor. “The Formation of Emulsions in Definable Fields of Flow”. In: *Proc. R. Soc. Lond. Ser. Contain. Pap. Math. Phys. Character* 146.858 (1 Oct. 1934), pp. 501–523. ISSN: 09501207.
- [10] Nobuo Watanabe et al. “Dynamic Deformation and Recovery Response of Red Blood Cells to a Cyclically Reversing Shear Flow: Effects of Frequency of Cyclically Reversing Shear Flow and Shear Stress Level”. In: *Biophys. J.* 91.5 (1 Sept. 2006), pp. 1984–1998. ISSN: 0006-3495. DOI: 10.1529/biophysj.105.060236.
- [11] Young-Zoon Yoon et al. “The nonlinear mechanical response of the red blood cell”. In: *Phys. Biol.* 5.3 (Aug. 2008), p. 036007. ISSN: 1478-3975. DOI: 10.1088/1478-3975/5/3/036007.

8. Clinical Pilot Study

In this chapter, the technology developed in thesis has been applied to the measurement of mechanical properties of live red blood cells. In particular, the prototype device has been used to trap cells in an automated fashion at physiologically-relevant extensional rates, and measure the resultant dynamic behaviour of the cells. First, the experimental procedure for the preparation of the cells and system will be briefly explained, followed by a short explanation of the image processing techniques applied to extract shape information from the resultant images. Finally, the results of some clinical pilot studies will be presented and discussed.

8.1. Experimental Procedure

The chip was attached to the mounting platform and secured against rotation and translation with double-sided tape (non-foam type, Sellotape brand). Four lengths of 0.8 mm ID, 1.6 mm OD silicone tubing (Gecko Optical), each about 170 mm long were attached to the fluid reservoirs (Nordson Optimus 55 cc syringe barrels, Nordson Corp.) via Luer lock fittings. Approximately 10 mL of working fluid solution was prepared for each experiment. This consisted of a 1 % solution of 1 μm suspended particles (PolySciences Polybead) in clean analytical grade phosphate buffered saline (from Phosphate buffer saline powder, SigmaAldrich and ultra-pure water) and 1 mmolL^{-1} of the anti-coagulant K_2EDTA . To minimise clumping and reduce the functional properties of the particles, 150 mg of suspended particles were first dispensed by weight and suspended in a rinse solution, consisting of 1 % Triton X-100 detergent (SigmaAldrich) in ultra-pure water. This rinse solution was agitated in a sonication bath for approximately 60 min before repeated centrifugation and re-suspension at a relative centrifugal force (rcf) of 6000g for three minutes, up to a maximum of three times. The precipitate was then suspended in the PBS buffer to form the working fluid.

Approximately 4 mL of solution was initially added to the two inlet reservoirs, with 2 mL kept aside for later suspension of the red blood cells (RBCs). The system was then initially purged at a net driving pressure of 16 kPa until the trapped air in the tubing and microchannel was expelled. Once the air was completely purged from the system, the working pressure was reduced to between 1 kPa – 4 kPa and the online valve calibration procedure initiated (see section 5.4.2). Once the calibration had completed and the system was stable, blood was sampled from a willing volunteer. Blood was obtained via lancet prick (Owen Mumford Unistik 3 Extra) to either the ring or middle fingers of the volunteer's non-dominant hand. The first drops of blood were wiped off with clean, sterile gauze and 20 μL of blood was drawn up into a sterilised EDTA-dipped pipette tip. The sampled whole blood was then suspended in 1 mL of fresh PBS and EDTA solution and centrifuged for 2 min at 3600g. The precipitant was then pipetted into the remaining 2 mL of working fluid solution and agitated to ensure suspension. The RBC and working fluid solution was then transferred to the appropriate fluid reservoir and the system re-pressurised. The system was briefly returned to the purge pressure to ensure that a sufficient number of RBCs were conducted downstream into the chip, before the pressure was reduced to that desired for trapping and deformation measurement.

The composition of the working fluid, as well as those of the particle and cell washes, are summarised in table 8.1. While this procedure appears quite involved, it is not unusual for experimental preparation in the life sciences. In a commercial, or semi-commercial version of the device, one could envisage that the working fluid could be supplied pre-prepared, substantially shortening the procedure.

Component	Quantity or Concentration
Working Fluid	
Phosphate Buffered Saline Solution	10 mL
Rinsed Polybead Particles	≈ 150 mg
Triton X-100	trace
K ₂ EDTA	1 mmolL ⁻¹
Particle Rinse	
Ultrapure Water	10 mL (3 batches)
Triton X-100	1 %
Polybead Particles	150 mg
Cell Wash Solution	
Phosphate Buffered Saline Solution	2 mL
K ₂ EDTA	1 mmolL ⁻¹

Table 8.1: Summary of the compositions of the fluids used in the preparation of blood solutions for the clinical pilot.

8.2. Measuring Cell Deformation

Cell deformation is computed in an offline processing step after the experiment is complete. During the experiment, the automated macro code (described in chapter 7) initiates trapping, manipulates the system pressures and triggers an image export once the cell trap has finished. A large image buffer is used (by default 4000 frames, or 20 s worth of images) so that the complete trap and release process can be captured on disk. To maintain throughput, only the part of the buffer that contains a cell trap is exported as the relatively slow speed of the disks limits the number of cells that can be analysed in an experiment.

8.2.1. Extracting Cell Shape

Due to the relatively low resolution of the camera relative to the dimensions of a red blood cell (the effective pixel size is $0.37 \mu\text{m}/\text{px}$, giving an average diameter of an undeformed RBC of ≈ 21 px), extracting shape data is a challenge. The effective pixel size of the system is a function of camera pixel density and system magnification. The magnification could not easily be increased as this would have placed higher demands on the cell trapping and flow measurement parts of the system (an increase in magnification increases the effective particle and cell displacements per frame) as the cells would leave the field of view more rapidly before the system had time to stabilise the trap. This could be mitigated against by switching to a camera with a larger number of pixels. However, the current system is already close to the limit of available computing hardware and camera technology. A dual-camera setup, such as that described in Curtis, Sheard and Fouras,⁶ is a viable solution, but was not attempted. Hence, a solution needed to be devised to extract the maximum amount of information from the available data.

Background subtraction is first applied to the input images, following a similar procedure to that described in section §6.1. To determine the locations of the cell edges, Canny edge detection is used.⁵ A Gaussian blur is applied to the entire image with a blur radius $\sigma = 5$ px to reduce the amount of noise present in the derivative. Next, Canny edge detection is performed by taking the first derivative

of the image in both directions (Sobel filter) and computing both the magnitudes and angles of the edge gradients for all pixels. These gradients are binarised using the Canny thin-line thresholding algorithm to retrieve a thin edge from the image. In a modification from the standard algorithm, the binarised image is then multiplied by the gradient magnitudes. This results in thin lines, as per the standard Canny edge detection algorithm, with the addition of an edge-strength or confidence measure. This edge-strength measure is used later as part of the shape fitting algorithms.

A simple template-sliding based matching algorithm is then used to find an initial guess for the cell shape and optical parameters. Due to the edge enhancement effect of the optical setup (the illumination is partially collimated via a pinhole and therefore there is a weak holographic effect²¹), many cells appear as a bright central region surrounded by a darker exterior. The apparent size of the bright interior varies depending on the distance to the focal plane and shape of the cell. Therefore, the template-sliding algorithm is run several times with different interior sizes to find both an approximate location and interior radius. This interior radius is used to mask the internal edges of the cell from the later ellipse-fitting stages, improving the both the accuracy and noise rejection characteristics of the fit.

8.2.2. Fitting the Ellipse

The series of edge points collected by the modified Canny edge-detection algorithm are collected into an array of input points to be fitted to parametric ellipsoid. The CERES non-linear solver¹ is used to minimise a set of cost functions to determine the best-fit ellipsoid to the data. This ellipsoid is defined by the standard set of parametric equations

$$r(\theta) = \frac{r_m r_M}{\sqrt{[r_m \sin(\theta - \omega)]^2 + [r_M \cos(\theta - \omega)]^2}}, \quad (8.1)$$

$$x(\theta) = r(\theta) \cos(\theta) + x_0, \quad (8.2)$$

$$y(\theta) = r(\theta) \sin(\theta) + y_0, \quad (8.3)$$

where r_m and r_M are the minor and major radii, ω the rotation angle of the cell, and (x_0, y_0) the position of the cell centre. Expressing the ellipsoid in polar coordinates allows a pseudo-linear least squares approach to the solution, whereby the cost functions are defined in terms of the distance to the nearest edge for a given angle θ . Firstly, a distance-based minimiser $f_1[P]$ is defined

$$D_P = \left\| l_P - \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} \right\| - r(\theta), \quad (8.4)$$

$$f_1[P] = s_P^2 L(r(\theta))^2 D_P, \quad (8.5)$$

where l_P is the location of the edge P , D_P the difference between the radius estimate for the point P and the computed radius $r(\theta)$, s_P the edge strength measure (as described in section 8.2.1) and L a loss function that reduces the relative importance of edge points far from the initial radius estimate.

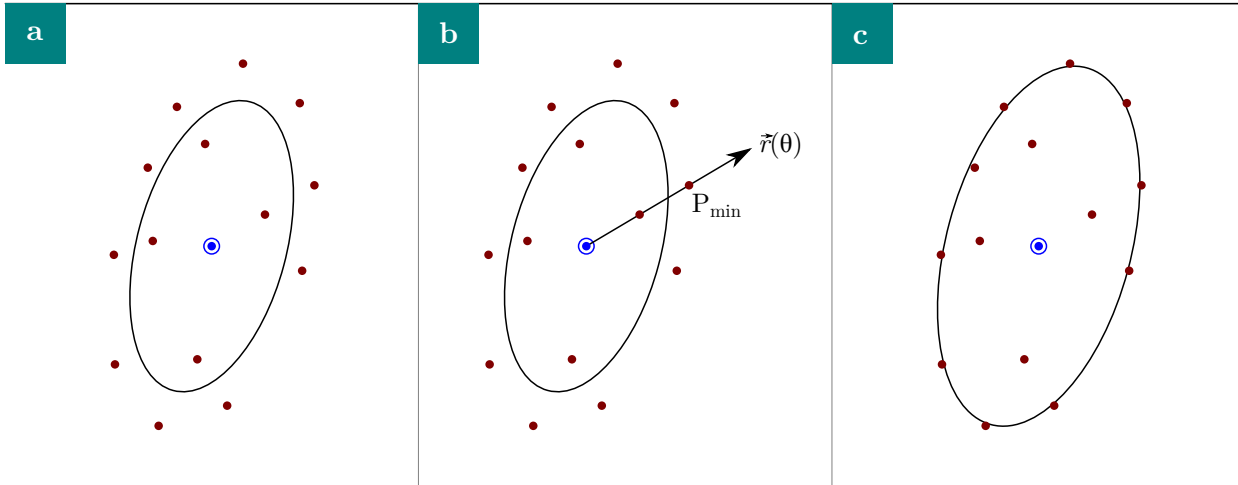


Figure 8.1: Illustration of the effect of the cost functions used in fitting an ellipse to sparse point data, where the extracted edges are represented by dots (\bullet) and the centre by a ringed dot (\odot). (a) Representation of the likely outcome from a fit using the least-squares type minimiser f_1 alone. The ellipse edge minimises the distance to all points, but does not pass through any of them; (b) graphical illustration of equation (8.6), where there are two edge points along the line formed by the ellipse centre at the radius vector $\vec{r}(\theta)$. The point closest to the curve of the fitted ellipse is chosen to contribute to the cost function; (c) representation of a likely fit using only the fitness function f_2 , all points chosen lie perfectly along the curve of the fitted ellipse, minimising f_2 .

Secondly, a best-match minimiser $f_2(\theta)$ is defined

$$P_{min}(\theta) = \operatorname{argmin}(P, D_P(\theta)), \quad (8.6)$$

$$\text{where } D_P(\theta) \text{ satisfies } \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} + \lambda \begin{pmatrix} x(\theta) \\ y(\theta) \end{pmatrix} = D_P(\theta)$$

$$f_2(\theta) = s_{P_{min}}^2 D_{P_{min}}, \quad (8.7)$$

where $P_{min}(\theta)$ selects the edge point that is closest to the estimated ellipse radius. f_1 is a conventional least-squares minimiser, where the optimal fitted function is one which minimises the squared distance to all edge points. On the other hand, f_2 is a fitness function that measures the degree to which the fitted function follows a pre-existing curve in the data. The effects of these function on a sparse set of edge points is shown in figure 8.1, in which the least-squares minimiser f_1 is biased by spurious edge points (figure 8.1a), whereas the fitness function f_2 is minimised when a majority of points lie along the curve of the fitted ellipse (figure 8.1b). However, as f_2 is discontinuous, it cannot be used as the sole cost function in a non-linear solver without risk of instability. Hence f_1 is used to find a series of likely solutions, and f_2 biases the solution toward one that fits a majority of edge points.

8.2.3. Validation

In order to characterise the noise rejection characteristics of the fit and the ability of the code to extract shape information from noisy, sparse image data, the code was first validated against a series of known inputs. The earlier trapping work, conducted using $6\ \mu\text{m}$ polystyrene particles (THERMO

SCIENTIFIC 7505A), provides an excellent source of data for this purpose. As these particles are mass manufactured, it is expected that the shape of the particles should be very close to perfectly spherical. Additionally, the nominal particle sizes is known, and it can be expected that the average diameter will converge to this value. Additionally, no (measurable) deformation is expected due to the material properties of the particles.

The average diameter of a small, random sample of particles, as measured by the algorithms described in the preceding section was $7.35 \mu\text{m}$ with a standard deviation of $\sigma = 0.87 \mu\text{m}$, after compensating for out-of-focus effects due to the numerical aperture of the system. This is well within the size distribution given by the manufacturer ($\bar{d} = 6 \mu\text{m}$, $\sigma = 1.5 \mu\text{m}$) and therefore indicates that the particle diameters are correctly recovered by this analysis. Secondly, the average Taylor deformation index was computed for the ensemble of sampled particles. The deformation index is highly sensitive to small changes in the ratios between the major and minor radii, and is therefore a good measure of the noise inherent in the analysis. After performing simple outlier rejection to remove cases where bright tracer particles have interfered with the measurement, the average Taylor deformation index was found to be $DI = 0.039$. This approaches the expected value ($DI = 0$ for perfect circles).

8.3. Time-Varying Behaviour of Cell Shape

By applying the ellipse-fitting techniques to the images captured from live cell trapping events, the deformation of the cell as a result of the forces exerted upon it by the fluid flow can be measured. A typical series of events is shown in figure 8.2, whereby the images have been cropped by a fixed amount around the cell centroid, in order to highlight the deformation of the cell. Figure 8.2b shows the fitted shape, as produced by the measurement software, which exhibits good agreement with the edge contours of the cell (within the bounds of the three parameter ellipse model). Of particular note is the ability of the least-squares fit to reject the spurious edges caused by the close proximity of a tracer particle (bottom left).

The tracer particles are made from a polystyrene copolymer that is adsorbed by many proteins.^{4,13,14} While the particles were thoroughly washed and treated to reduce this effect, some particles are still adsorbed by proteins on the surface of the cell. However, the presence of the particle in figure 8.2, while normally undesirable from a measurement standpoint, can be used to infer some important properties of the system.

Red blood cells are often observed to exhibit a property known as ‘tank treading’ when exposed to high shear rates.^{11,16,18–20} This is a phenomenon whereby the cell membrane rotates around the cell interior. As the particle location is static throughout the measurement in figure 8.2, the blood cell can be assumed not to be undergoing tank-treading in this extensional flow. Another commonly observed behaviour, ‘tumbling’, can also be excluded.^{11,18,19} Tumbling, also observed in high shear flows, is where the cell rotates about an axis parallel to the flow. This would result in the particle moving behind or in-front of the cell in the image sequence under discussion. As this is not taking place, we can assume that the cell is undergoing pure deformation as a result of the applied force, and not an apparent deformation due to an out of plane rotation of the cell.

8.3.1. Minimum Deformation Forces

While many of the cells enter the camera field of view with the long (major) axis oriented towards the device outlets ($\theta = 90^\circ$), a small number enter oriented towards the inlets ($\theta = 0^\circ$). As this is an unstable condition, the cells rapidly rotate toward $\theta = 90^\circ$. This behaviour is evident in the three deformation and angle histories presented in figure 8.3. In all three cases shown in the figure, there is

8. Clinical Pilot Study




Figure 8.2: Deformation of a healthy red blood cell over time, as a series of images at key points in time. (a) Raw image data, showing gradual extension of the cell due to the forces exerted on it by the drag forces in the flow; (b) elliptical shapes as fitted by the algorithms described in section 8.2.1 (light coloured outlines). The particle (bottom left of the cell in each frame) became attached to the cell due to random mixing. The presence of this particle, while normally undesirable from a measurement standpoint, can be used to infer some important properties of the system. Crucially, the relative location of the particle (a) on the cell remains static throughout the trap. This indicates that neither tumbling, nor tank-treading is taking place. Secondly the fitted cell shape (b) is not biased by the presence of the particle and is therefore an indicator of the noise rejection performance of the fitting algorithms.

a threshold angle — between 60° and 80° for the cases shown — at which the cell undergoes a rapid change in shape. As the cell rotates in the device, the angle between the resultant force vector (due to fluid flow) and the cell major axis becomes smaller and the elongation component of the forces on the cell thereby increases. Hence, this rapid change in deformation at a particular rotation angle suggests that there is a minimum force required to overcome the internal membrane stiffness. This behaviour is not uncommon amongst non-linear materials, but it is nevertheless interesting to observe in a complex mechanical system such as a red blood cell.

8.4. Disease Model: Heat Treatment

Heat treatment is a method of modifying the stiffness of live red blood cells. The raised temperature is thought to cause denaturation of the spectrin proteins that form part of the cell membrane, changing the mechanical properties of the cells.¹⁵ This effect is easy to control using a water bath, and the stiffness change can be adjusted by varying the length of time that the cells are exposed to the elevated temperature.^{7,8} In contrast to other artificial methods of stiffness manipulation, such as chemical treatment with glutaraldehyde, the stiffening process is easily terminated by returning the blood cells to a normal temperature (i.e. $25^\circ\text{C} \leq T \leq 37^\circ\text{C}$). Hence the heat treatment methods are an effective and reproducible method for benchmarking the sensitivity of the technique.

For the experiments to follow, healthy red blood cells were obtained using the method described in 8.1. The cells were suspended in a phosphate-buffered saline solution in a small centrifuge tube and then exposed to a temperature of $49.1 \pm 0.1^\circ\text{C}$ by immersing the tube in a water bath for a period of 2 min. The cell solution was then rapidly cooled by immersing the tube in a water bath at room temperature (26.8°C). The suspension time of 2 min was chosen to produce the minimum change in stiffness, as documented by Dobbe et al.,⁸ and therefore provides an indication of the minimum measurable change in stiffness.

8.4.1. Steady-State Stiffness Changes Due to Heat Treatment

While most suited to dynamic measurement of cell behaviour, the steady-state stiffness of cells can be measured using the techniques developed in this thesis. This measurement is an important benchmark as steady-state stiffness is commonly measured and reported in the existing literature,^{3,8,10,17,23} as it can be measured with a range of techniques (for example, optical tweezers,¹⁷ micro-pipette aspiration,¹⁰ and atomic force microscopy³). A particular extensive study of stiffness change, in terms of aspect ratios, as a function of heat treatment duration was reported in Dobbe et al.⁸ Using a flow exaktocytometer, they report a normal, or control aspect ratio of $\text{AR} = 3.01$, and the average aspect ratio after 2 min of heat treatment to be $\text{AR} = 2.71$. These experiments were carried out at relatively high shear rates, and so the deformation is greater than those obtained in the extensional flow device under discussion. However, the ratios in stiffnesses between the control and heat-treated samples is expected to be similar.

The average steady-state deformation was obtained from two experiments, conducted within hours of one another, where one sample was heat-treated as described in section §8.4 and the other was left untreated. Over 30 independent measurements of individual cells were obtained for each sample. For the control sample, the average aspect ratio was $\text{AR} = 2.27$ and for the heat treated sample, the average aspect ratio was $\text{AR} = 1.84$ ($P = 0.0003$). This represents a 24 % decrease in deformation, compared to the 11 % decrease reported in Dobbe et al.⁸ The relationship between cell stiffness and applied force is known to be non-linear,^{12,17,22} and is expected to have an S-curve relationship¹⁷ (almost no deformation at low forces, rapid change in deformation above a minimum activation force,




Figure 8.3: A plot of Taylor Deformation Index (top) and cell orientation (bottom) against time, showing the effect of cell orientation on deformation behaviour. 90° (horizontal dashed line) is the angle at which the cell major axis is oriented along the direction of maximum extensional force. In all three cases shown in the figure, there is a threshold angle — between 60° and 80° for the cases shown — at which the cell undergoes a rapid change in shape. The darker shaded region on the left indicates boundary between the initial deformation and the onset of the rapid shape transition as a result of the change in orientation.




Figure 8.4: Histogram of steady-state deformation of heat-treated and control red blood cells, where the y -axis is the normalised probability density. Normal distributions have been fitted to the distributions in both cases. The mean Taylor deformation index of the heat treated cells was significantly lower ($DI = 0.277$) than the non-treated cells ($DI = 0.372$, $P = 0.0003$, one-sided t -test).

followed by an asymptotic approach the maximum value). Hence, the deformation change as a result of increased inherent stiffness will be more pronounced at a lower applied force, as shown by these measurements.

Figure 8.4 shows the distribution steady state deformations, expressed in terms of the Taylor Deformation Index, instead of aspect ratio, for consistency with the other results presented in this thesis. The Taylor Deformation Index is a more sensitive measure (showing a 34% decrease in deformation for the same dataset) and therefore better highlights the changes in stiffness. As can be seen in the figure, there is a clear decrease in average deformation as a result of heat treatment.

8.4.2. Dynamic Deformation as a Sensitive Diagnostic

A key advantage of trapping and analysing the deformation of individual cells is the ability to measure both the final deformation (a measure of Hookean stiffness) and the rate of change of deformation (a measure of the effective viscosity of the cell). Unlike devices such as cytometers, which measure the bulk deformation of a large number of cells, or micro-pipette aspiration, which measures the deformation of a single cell in a labour-intensive process, this device measures the time-dependent deformation of individual cells in an automated process.

8. Clinical Pilot Study

Cell deformation as a function of time has been plotted for a number of control (untreated) and heat-treated cells in figures 8.5 and 8.6. In all cases, these deformations present the behaviour of a cell over a period of 2 s, starting from when it enters the field of view. Time bases for each individual deformation history have been aligned such that $t = 0$ s represents a cell position of $120 \mu\text{m}$ from the saddle-point (in the x -plane) to facilitate easy comparison. Figure 8.5 shows a least-squares spline fit to the measured deformations, while figure 8.6 shows the raw data points along with the associated spline fits. Figure 8.6 demonstrates that this type of fit is effective at rejecting random noise, while preserving the original trends in the data. In the heat-treated cases, the cells exhibited a greater propensity to acquire attached particles over time, and the measurement became noisy. However, the final steady-state deformations of the cells, after noise-rejection are represented by the dashed lines in figures 8.5 and 8.6.

The overall deformation trends, clearly seen in figure 8.5, reveal some interesting properties of the cells. In all cases, there is some non-zero initial deformation. It is expected that an unstressed cell would have a deformation index of close to zero, and so this indicates that the relatively small amount of force exerted on the cells at the entry to the microfluidic cross is sufficient to produce some deformation. The degree to which the cell is deformed at the entry-point differs between each cell, and is indicative of the natural inter-sample variation in mechanical properties between cells. The control and heat-treated samples were each obtained from a single donor and one lancet prick. This variation indicates that red blood cell deformation behaviour is too complex to be described by a single, one-dimensional, quantity and highlights the natural variability of red blood cells within the bloodstream. More detailed information is required.

For the control cells (figure 8.5a), there is a slow increase in deformation over time, indicating that the cell is viscoelastic. Cells C-1 and C-4 undergo an initial rapid change in deformation — due to the rotation shown in more detail in figure 8.3 — followed by a much slower increase to the final value. In the heat-treated case, the change in deformation is much faster, with the steady-state, or final value showing no further increase in deformation. This is a strong indication that cell viscosity is modified due to heat treatment, in addition to the commonly-reported stiffness changes. This finding is supported by micro-pipette aspiration and electric field experiments,^{2,9} where a decrease in viscosity was reported to accompany an increase in stiffness once the temperature of the cells was increased above 45°C .

Using this observation, the simple one-dimensional measurement can be significantly expanded to a two-dimensional measurement by including the rate-of-change of the deformation in terms of a time constant. The time constant can be obtained by modelling the deformation in terms of a simple RC model

$$x(t) = Ae^{-\tau t} + x_0, \quad (8.8)$$

where x_0 is the initial value, A the amplitude of the step-change and τ the time constant. Hence, the mechanical behaviour of the cell can be represented in terms of the final deformation

$$x_{fv} = \lim_{x \rightarrow \infty} [Ae^{-\tau t} + x_0],$$

and the time constant τ . The time constants and final values for each of the cells shown in figure 8.5 are given in table 8.2. The standard approximation of τ has been used ($\tau = t_1$, where t_1 satisfies $x(t_1) - x_0 = 2/3A$) and $t = 0$ (and therefore x_0) has been defined to be either the earliest available measurement point or, in the case of the rotating cells (C-1, C-4 and HT-1), the first measurement point where $\theta \geq 90^\circ$. As is evident from the table, there is a clear separation in the time constants for the heat-treated and control samples. This is further emphasised in the scatter plot (figure 8.7), where

Cell	Time Constant τ (s)	Final Deformation (Taylor DI)
C-1	0.675	0.294
C-2	0.550	0.297
C-3	0.420	0.320
C-4	0.580	0.238
C-5	0.335	0.348
C-6	0.405	0.394
C-7	0.400	0.372
HT-1	< 0.005	0.243
HT-2	0.115	0.166
HT-3	< 0.005	0.150
HT-4	0.210	0.315
HT-5	0.140	0.280

Table 8.2: Time constants and final deformations for both the control (C-1 through to C-7) and heat treated cells (HT-1 through to HT-5) in the study. There is a clear, and dramatic, decrease in time constant (and therefore viscosity) in the heat treated cells, in addition to the expected increase in stiffness (resulting in lower final deformations).

the plot can be divided into two clear regions, with the control cells occupying the high-deformation and long time constant region, and heat-treated cells occupying a low-deformation and short time constant region. It is expected that particular diseases would have a unique effect on the viscosity and steady-state stiffness (final deformation), allowing cells that are affected by these diseases to be plotted in different regions on this map. A one-dimensional measurement of stiffness alone (see figure 8.4) does not allow identification of disease states on a single-cell basis. The use of this additional information on single cells may enable more accurate diagnosis and/or earlier intervention in a clinical application of this tool. For example, as red blood cells are replenished at a finite rate, the percentage of diseased cells may be indicative of the history and duration of disease symptoms. This may have applications in the management of diabetes — the frequency and severity of hypo- or hyperglycaemic events over the last 3 months could be inferred, using the cell positions on the stiffness-viscosity map to determine the percentage of damaged cells and the severity of the damage.




Figure 8.5: Cell deformation over time for a number of red blood cells. (a) control (untreated) cells; (b) heat-treated cells. All lines are spline fits to the data; dotted lines indicate the final value for the heat treated cases.




Figure 8.6: Cell deformation over time for a number of red blood cells, showing the raw data points from the fitting algorithm and the associated spline fits. As is evident from the figure, this type of fit is effective at rejecting random noise and preserves the trends in the original data.




Figure 8.7: Time Constant and Final Deformation plotted as a 2-D map. The control and heat treated cells occupy two distinct regions on the map. It is expected that particular diseases would have a unique effect on the viscosity and steady-state stiffness (final deformation), allowing cells that are affected by these diseases to be plotted in different regions on this map. By comparison, the one-dimensional measurement of stiffness alone (see figure 8.4) does not allow identification of disease states on a single-cell basis. The use of the additional information afforded by this analysis method may enable more accurate diagnosis and/or earlier intervention in a clinical application of this tool.

8.5. Conclusion

In this chapter, the technology and associated device developed in this thesis have been successfully applied to the measurement of key mechanical properties of live red blood cells. In particular, the ability to trap cells in an automated fashion at high extensional rates, and measure the resultant dynamic behaviour of the cells, greatly expands upon the existing tools available for cell analysis. It has been shown that both the cell stiffness and cell viscosity are important measures in determining the mechanical properties of cells and thereby more accurately inferring a disease state. The potential of this dynamic measurement method for diagnostics and research extends beyond these two properties. For example, by modifying extensional rate in real-time, as explored in section §7.4, detailed measurements can be made of the force-extension profile and relaxation times of individual cells. In contrast to other methods, these measurements can be obtained from a single laboratory device, and on an individual cell basis, rather than a bulk measurement. Furthermore, the technique is capable of high throughput operation. With a fast storage system (based on one or more solid state drives, for example), cell throughput rates of over 30 cells/min are achievable.

8.6. Bibliography

- [1] Sameer Agarwal and Keir Mierle. *Ceres Solver*. version 1.4.0. (2012). Google, Inc. URL: <http://ceres-solver.org/>.
- [2] O. K. Baskurt and H. J. Meiselman. “Determination of red blood cell shape recovery time constant in a couette system by the analysis of light reflectance and ektacytometry”. In: *Biorheology* 33.6 (Nov. 1996), pp. 489–503. ISSN: 0006-355X. DOI: 10.1016/S0006-355X(97)00037-1.
- [3] Kristen E. Bremmell, Allan Evans and Clive A. Prestidge. “Deformation and nano-rheology of red blood cells: An AFM investigation”. In: *Colloids and Surfaces B: Biointerfaces* 50.1 (1 June 2006), pp. 43–48. ISSN: 0927-7765. DOI: 10.1016/j.colsurfb.2006.03.002.
- [4] M. M. Browne et al. “Protein adsorption onto polystyrene surfaces studied by XPS and AFM”. In: *Surface Science* 553.1–3 (20 Mar. 2004), pp. 155–167. ISSN: 0039-6028. DOI: 10.1016/j.susc.2004.01.046.
- [5] John Canny. “A Computational Approach to Edge Detection”. In: *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-8.6 (Nov. 1986), pp. 679–698. ISSN: 0162-8828. DOI: 10.1109/TPAMI.1986.4767851.
- [6] Michael D. Curtis, Gregory J. Sheard and Andreas Fouras. “Feedback control system simulator for the control of biological cells in microfluidic cross slots and integrated microfluidic systems”. In: *Lab Chip* 11.14 (28 June 2011), pp. 2343–2351. ISSN: 1473-0189. DOI: 10.1039/C1LC20191C.
- [7] J. G. G. Dobbe et al. “Analyzing Red Blood Cell-Deformability Distributions”. In: *Blood Cells. Mol. Dis.* 28.3 (May 2002), pp. 373–384. ISSN: 10799796. DOI: 10.1006/bcmd.2002.0528.
- [8] J. G. G. Dobbe et al. “Measurement of the distribution of red blood cell deformability using an automated rheoscope”. In: *Cytometry* 50.6 (15 Dec. 2002), pp. 313–325. ISSN: 1097-0320. DOI: 10.1002/cyto.10171.
- [9] H. Engelhardt and E. Sackmann. “On the measurement of shear elastic moduli and viscosities of erythrocyte plasma membranes by transient deformation in high frequency electric fields.” In: *Biophys J.* 54.3 (Sept. 1988). PMC1330348, pp. 495–508.

Bibliography

- [10] E.A. Evans. “New Membrane Concept Applied to the Analysis of Fluid Shear- and Micropipette-Deformed Red Blood Cells”. In: *Biophys. J.* 13.9 (Sept. 1973), pp. 941–954. ISSN: 0006-3495. DOI: 10.1016/S0006-3495(73)86036-9.
- [11] Badr Kaoui, George Biros and Chaouqi Misbah. “Why Do Red Blood Cells Have Asymmetric Shapes Even in a Symmetric Flow?” In: *Phys. Rev. Lett.* 103.18 (30 Oct. 2009), pp. 188101–4. DOI: 10.1103/PhysRevLett.103.188101.
- [12] Nathalie Nève et al. “Manipulation of Suspended Single Cells by Microfluidics and Optical Tweezers”. In: *Cel. Mol. Bioeng.* 3.3 (Mar. 2010), pp. 213–228. ISSN: 1865-5025. DOI: 10.1007/s12195-010-0113-3.
- [13] Polysciences, Inc. *Polymer Microspheres Datasheet*.
- [14] Polysciences, Inc. *Protocol for Adsorbing Protein on Polystyrene Microspheres*. Technical Data-sheet 238E.
- [15] Allen L. Rakow and Robert M. Hochmuth. “Effect of heat treatment on the elasticity of human erythrocyte membrane.” In: *Biophys J* 15.11 (Nov. 1975), pp. 1095–1100. ISSN: 0006-3495. pmid: 1201326.
- [16] T. W. Secomb et al. “Flow of Axisymmetric Red Blood Cells in Narrow Capillaries”. In: *J. Fluid Mech. Digit. Arch.* 163.-1 (1986), pp. 405–423. DOI: 10.1017/S0022112086002355.
- [17] John Sleep et al. “Elasticity of the Red Cell Membrane and Its Relation to Hemolytic Disorders: An Optical Tweezers Study”. In: *Biophysical Journal* 77.6 (Dec. 1999), pp. 3085–3095. ISSN: 00063495. DOI: 10.1016/S0006-3495(99)77139-0.
- [18] Y. Sui et al. “Tank-treading, swinging, and tumbling of liquid-filled elastic capsules in shear flow”. In: *Phys. Rev. E* 77.1 (31 Jan. 2008), p. 016310. DOI: 10.1103/PhysRevE.77.016310.
- [19] Ken-ichi Tsubota and Shigeo Wada. “Effect of the natural state of an elastic cellular membrane on tank-treading and tumbling motions of a single red blood cell”. In: *Phys. Rev. E* 81.1 (20 Jan. 2010), p. 011910. DOI: 10.1103/PhysRevE.81.011910.
- [20] Nobuo Watanabe et al. “Dynamic Deformation and Recovery Response of Red Blood Cells to a Cyclically Reversing Shear Flow: Effects of Frequency of Cyclically Reversing Shear Flow and Shear Stress Level”. In: *Biophys. J.* 91.5 (1 Sept. 2006), pp. 1984–1998. ISSN: 0006-3495. DOI: 10.1529/biophysj.105.060236.
- [21] Wenbo Xu et al. “Digital in-line holography for biological applications”. In: *Proc. Natl. Acad. Sci.* 98.20 (2001), pp. 11301–11305.
- [22] Young-Zoon Yoon et al. “The nonlinear mechanical response of the red blood cell”. In: *Phys. Biol.* 5.3 (Aug. 2008), p. 036007. ISSN: 1478-3975. DOI: 10.1088/1478-3975/5/3/036007.
- [23] Hu Zhang et al. “An optical-manipulation technique for cells in physiological flows”. In: *J. Biol. Phys.* 36.2 (1 Mar. 2010), pp. 135–143. DOI: 10.1007/s10867-009-9176-6.

Part IV.

Conclusion and Closing Remarks

9. Conclusions

In this work, a new type of microfluidic system, substantially defined by its software, rather than the characteristics imparted during manufacture (such as its geometry), has been developed and applied to red blood cell capture and analysis. In support of this goal, a number of innovations were required. A few of the more important technological innovations are: efficient air-free microfluidic interconnects; a novel control system simulator for the control of coupled fluid-particle systems; development of off-chip valves and control strategies that support them; lighting control and timing schemes that allow high speed, highly variable control of lighting to maximise the accuracy of image processing; and finally, efficient, real-time particle and cell tracking algorithms.

The technology developed and proven in this thesis enables the conception and construction of software-defined systems to manipulate not only cells, but also oil droplets, particles, and other suspended objects. These dynamic manipulation and measurement methods represent a significant advance over existing techniques for cell and droplet manipulation in diagnostics and research. In contrast to other methods, detailed multi-dimensional measurements (for example, both force-extension profiles and relaxation times) can be obtained using a single laboratory device, and on an individual cell basis, rather than through bulk measurement.

9.1. Biomedical Relevance

An ongoing theme in this thesis is the need for the capacity to make detailed measurements of key mechanical properties of live red blood cells. In particular, the technology developed in this thesis greatly expands upon the existing tools available for cell analysis by enabling devices that can trap cells in an automated fashion at high extensional rates, and measure the resultant dynamic behaviour of the cells. Cells were obtained, subjected to a minimum of processing, and the dynamic changes in cell shape tracked as the cell was subjected to varying extensional forces. This was accomplished without the introduction of foreign manipulation devices involving lasers or physical contact and requires a minimum of operator intervention. By measuring both the steady-state and dynamic behaviours of cells, both the stiffness and time constant (viscosity) were obtained for individual cells. This simultaneous measurement of two properties is significantly more powerful than a simple stiffness measurement. This is evident in the data presented in figure 8.7, where the heat-treated and control cells are clearly delineated into two groups, whereas using only stiffness measurement, this distinction is only evident on a bulk, statistical basis (figure 8.4). Thus, the measurement of two or more properties, enabled in an automated and high-throughput fashion by the device developed in this thesis, can identify the disease state of individual red blood cells. This allows more specific measurements of cell behaviour and may enable earlier diagnosis of disease. For example, once the stiffness and viscosity profile of glycosylated haemoglobin has been determined, the stages and likely physiological effects of diabetes can be rapidly and inexpensively assessed, improving patient care and outcomes.

The work presented in this thesis enables researchers and clinicians to efficiently measure cell stiffness and viscosity simultaneously from single cells. This greatly improves the predictive and diagnostic power of cell mechanics measurements available to researchers and has the potential to revolutionise both our understanding of blood-related diseases (such as diabetes and malaria) and point-of-care diagnostics.

9.2. Technical Advances

In order to meet the challenge of trapping cells at physiologically-relevant extensional rates, it was necessary to innovate across a number of fields which lead to a number of important contributions to the field, which have been published, patented or are under consideration for further patent protection. A few of the more important innovations are detailed in the following paragraphs.

9.2.1. Coupled Fluid Control Simulation

By using an innovative technique to use precomputed CFD flow data in a 2D lookup table, the response and stability characteristics of control system for trapping cells could be studied in a computationally efficient manner. This simulation software has been used extensively throughout the development of the device in order to better understand system behaviour — including stability limits, optimal parameters, and the effect of various flow geometries and actuator behaviours. This led to development and presentation in *Lab on a Chip* of a new type of control system simulator.

9.2.2. Flow Measurement

A highly-efficient, realtime flow measurement method using PIV has been presented alongside a realtime method for analysing the flow information that, taken together, can extract the relevant flow information in realtime while maintaining excellent accuracy. Additionally, a novel method has been developed for realtime optimisation of PIV imaging parameters, by dynamically tuning the illumination timing, that allows the system to rapidly adapt to large variations in flow velocities without loss of accuracy or large computational overhead.

9.2.3. High-Bandwidth Flow Control

Through innovations in control strategies for microfluidic valves system, the strongly non-linear, hysteretic valve-tubing-fluid system has been transformed into a linear, high-bandwidth flow controller. Transient performance was found to be excellent with response times in the order of 100 ms. Additionally, the automatic calibration methods presented enable very precise control of flow rates in fluid systems, independently of geometry and, unlike other systems, do not require lengthy expert setup or precise manufacture. Finally, the use of active control in software to stabilise and modify system behaviour shifts complexity away from the relatively expensive and consumable microfluidic chip and into a standardised and durable hardware platform and associated software. In so doing, the flow control system developed in this thesis enables purely software-controlled microfluidic platforms that can be employed for rapid and low-cost assays that can be field upgradable *via* firmware updates.

9.2.4. Novel Realtime Cell Tracking

The tracking algorithms (chapter 6) are uniquely capable of tracking biological cells against a noisy background in the presence of other cell-shaped objects (tracer particles). Additionally, this tracking is carried out fully automatically, with cells of interest automatically selected, segmented and tracked without human intervention. This is necessary, as cells enter and leave the field of view in less than a second, which means that it is impractical for a human operator to interact with the system in realtime. Finally, using a combination of CPU and GPU processing, and multi-threaded, pipelined analysis, the cell tracking system was able to accurately segment and track cells at a rate of 100 frames/s. Tracking and segmentation accuracies, as assessed by automated synthetic tests and detailed inspection were found to be in excess of 78%.

9.3. Closing Remarks

In this work, a series of new technologies have been presented that allow the construction of new, powerful tools that provide greater insights into cell mechanics than was previously possible. Furthermore, the ever-decreasing costs of computer hardware means that this technology has the potential to revolutionise patient care through the production of general-purpose, field reprogrammable devices that can measure the physical characteristics of a large number of cells for rapid diagnostic assays. As the technology in this area matures and becomes more broadly available, the depth of understanding of cell mechanics will rapidly accelerate. This may lead to the development of sophisticated diagnostic tools that can provide early and inexpensive diagnostics for a number of diseases. In pharmacology, the ability to assess the unwanted or intended effects of a particular drug on cell mechanics and behaviour has the potential to accelerate the development and approval for new drugs. The expanded ability to measure stiffness and viscosity changes in cells may lead to a greater understanding of diseases that target the red blood cell, such as malaria. Furthermore, the ability to impose forces on cells and measure the response in real time allows for the study of cell behaviours that are extremely difficult to measure by traditional means — such as the dynamic, and poorly understood conditions that lead to platelet activation. Finally, the ability to both measure and sort cells allows for applications outside traditional diagnostic treatment - such as the classification and sorting of viable cells for IVF.

The possibilities are endless.

Part V.
Appendices

A. Sample Configuration

In a number of places in this thesis, reference has been made to various system parameters or options that control the system, particularly in the control chapter (chapter 5) and tracking chapter (chapter 6). The following is an excerpt from the XML based configuration file that controls VULTURE, showing the typical settings and parameters used to obtain the results presented in Part III.

Listing A.1: Sample Configuration File for Cell Trapping

```
<?xml version="1.0" encoding="utf-8"?>
<Configuration xmlns="http://ldi.monash.edu.au/schemas/AvianConfiguration">
  <Program name="vulture">
    <Group name="OutputBuffers">
      <!-- Maximum size of the ring buffer used to store images in memory,
           this allows the last <N> frames to be saved to disk from a manual
           or automatic trigger -->
      <Number name="Images">4000</Number>
      <!-- Maximum size (in frames) of each PIV file. Larger numbers create
           fewer files, but increase the risk of lost data in rare case
           of a program crash or deadlock -->
      <Number name="PIV">4000</Number>
    </Group>
    <Group name="Camera">
      <String name="IP">192.168.1.41</String>
      <!-- Camera exposure time -->
      <Number name="exposure">4860</Number>
      <!-- Whether the light source is a double-shutter source. This also enables
           the trigger inputs in the camera -->
      <Boolean name="DoubleShutter">true</Boolean>
      <!-- The effective size, in micrometers of each pixel.
           (native pixel size / total mag) -->
      <Number name="EffectivePixelSize">0.35</Number>
    </Group>
    <Group name="ImageProcessing">
      <!-- Add a frame to the 'slow' ring buffer every <SlowDecimation> frames.
           See the Tracking chapter for details -->
      <Number name="SlowDecimation">500</Number>
      <!-- The number of frames in the 'fast' ring buffer -->
      <Number name="FastFrames">10</Number>
      <!-- The relative weight of the slow buffer -->
      <Number name="SlowWeight">0.85</Number>
    </Group>
    <Group name="VGT">
      <!-- Inherent delay (in frames) between a change to the illumination timing and
           the change becoming observable in the PIV analysis -->
      <Number name="Delay">8</Number>
      <!-- Minimum number of frames to wait after changing the timing before another
           change can be made. Higher numbers can damp oscillations at the expense of slower
           response -->
      <Number name="ChangeLockoutDelay">10</Number>
      <!-- The average (full field) velocity, in pixels/frame that the control
           loop should target. This is a tradeoff between the limits of subpixel
           accuracy (low numbers) and the number of aliased vectors (high numbers). -->
      <Number name="TargetVelocity">4</Number>
    </Group>
  </Program>
</Configuration>
```

A. Sample Configuration

```
<!-- The minimum time (microseconds) between pulses that can be supported
  by the hardware -->
<Number name="MinGapTime">130</Number>
<!-- The maximum time between pulses that is supported. This a function of
  the camera exposure time and the time within the first exposure that the
  first pulse is fired -->
<Number name="MaxGapTime">4300</Number>
</Group>
<Group name="ValveHub">
  <!-- IP Address of the valve communications hub -->
  <String name="IP">192.168.1.45</String>
  <!-- Mapping of amplifiers (left to right on panel) to valve locations
    on the chip. Locations (left,right,top,bottom) are as seen by the camera -->
  <String name="Map">L,T,B,R</String>
  <!-- Minimum and maximum valve positions. This gives a rough initial
    calibration to speed up the first calibration pass -->
  <Number name="Offset">14480</Number>
  <Number name="Scale">12880</Number>
</Group>
<Group name="ValveControl">
  <!-- Time (in microseconds) between each loop in the controller thread -->
  <Number name="InternalLoopRate">495</Number>
  <!-- Valve dynamics parameters -see Chapter 4 -->
  <Number name="ApproachVelocity">50</Number>
  <Number name="DeadbandVelocity">1500</Number>
  <Number name="Acceleration">60000</Number>
  <Number name="Deceleration">1000</Number>
  <Number name="Undershoot">1</Number>
</Group>
<!-- Parameters for the flow controller -->
<Group name="FlowPID">
  <!-- Standard PID parameters for x and y axes -->
  <Number name="Imax">2.5</Number>
  <Number name="DriveLimit">5</Number>
  <Number name="DifferentialHistory">5</Number>
  <Group name="GainsX">
    <Number name="P">-0.0001</Number>
    <Number name="I">0.001</Number>
    <Number name="D">0.05</Number>
  </Group>
  <Group name="GainsY">
    <Number name="P">-0.001</Number>
    <Number name="I">0.001</Number>
    <Number name="D">0.05</Number>
  </Group>
  <!-- Pressure - in cmH2O - below which the control system will switch off
    to prevent instabilities -->
  <Number name="LowPressureLockout">12</Number>
  <Group name="HystereticFilter">
    <Number name="Hysteresis">1</Number>
  </Group>
</Group>
<!-- Parameters for the cell trapping controller -->
<Group name="CellPID">
  <Number name="Imax">50</Number>
  <Number name="DriveLimit">320</Number>
  <Number name="DifferentialHistory">5</Number>
  <Group name="GainsX">
```

```

    <Number name="P">0.1</Number>
    <Number name="I">2</Number>
    <Number name="D">0.1</Number>
  </Group>
  <Group name="GainsY">
    <Number name="P">1.5</Number>
    <Number name="I">1</Number>
    <Number name="D">0.1</Number>
  </Group>
  <Group name="HystereticFilter">
    <!-- Minimum change in cell position before the controller will change
         the commanded position. Prevents oscillations as very small changes
         can introduce a much larger transient into the system due to
         valve hysteresis -->
    <Number name="Hysteresis">2</Number>
  </Group>
</Group>
<Group name="FlowCalibration">
  <Enum name="CalibrationMode">SaddlePoint</Enum>
  <Enum name="BalanceMode">Absolute</Enum>
  <Number name="StableWait">1000</Number>
  <Number name="Averages">50</Number>
</Group>
<Group name="FlowAnalysis">
  <Number name="Reject">5</Number>
  <Number name="NumPasses">2</Number>
  <Boolean name="WholeField">false</Boolean>
  <Number name="Recursive">0.5</Number>
  <NumberList name="CrossSize">640,480</NumberList>
  <Number name="WindowSize">64</Number>
  <Number name="Spacing">24</Number>
  <Number name="Averages">100</Number>
</Group>
<Group name="Tracker">
  <!-- See Tracking chapter -->
  <!-- Segmentation parameters -->
  <Number name="startThreshold">0.995</Number>
  <Number name="stopThreshold">0.75</Number>
  <Number name="MinIntensity">0.71</Number>
  <Number name="MaxGradient">0.25</Number>
  <Number name="SegmentIterations">1</Number>
  <!-- Minimum effective diameter - in pixels - for a candidate segment -->
  <Number name="minSize">20</Number>
  <!-- Matching probabilities -->
  <Number name="maxDistance">75</Number>
  <Number name="SizeDeviation">0.45</Number>
  <Number name="IntensityDeviation">0.3</Number>
  <Number name="DistanceDeviation">30</Number>
  <Number name="TimeDeviation">0.12</Number>
  <Number name="MinimumProbability">0.5</Number>
  <Number name="PIVRejectLevel">10</Number>
  <!-- Minimum velocity for streamline integration. This should be set to a small,
       but non-zero number. Streamlines that pass close to the saddle-point will
       generate a large number of integration points as the velocity trends to zero.
       This is expensive in terms of CPU time. Setting a minimum value here allows
       the integrator to skip over some of these points with a small loss in accuracy
       close to the saddle-point -->
  <Number name="MinStreakVelocity">0.1</Number>

```

A. Sample Configuration

```
<Number name="MaxRetries">6</Number>
<!-- Number of frames used to forward-project the cell location. This
      should be related to the total inherent delay of the system -->
<Number name="ForwardProjectFrames">7</Number>
<Boolean name="UseForwardProject">true</Boolean>
</Group>
<Group name="Macro">
  <!-- Number of frames to wait after the cell has reached the stability
        thresholds (see StabilityThreshold X & Y) before switching to
        low shear dwell -->
  <Number name="StableFrames">100</Number>
  <!-- Number of frames to dwell at each pressure -->
  <Number name="DwellFrames">80</Number>
  <Number name="HighShearPressure">25</Number>
  <Number name="LowShearPressure">25</Number>
  <!-- Thresholds, in terms of RMS error in pixels. Cell traps with RMS
        errors below these both thresholds in the respective directions are
        considered stable -->
  <Number name="StabilityThresholdX">80</Number>
  <Number name="StabilityThresholdY">30</Number>
  <!-- Pressure ramp rate in cmH2O/s -->
  <Number name="RampRate">0.5</Number>
  <!-- If true, only save the image frames from time of initial capture
        until release. If false, the number of image frames saved is controlled
        by the size of the image buffer -->
  <Boolean name="MinimalImageDump">>false</Boolean>
  <!-- The direction (up or down) to eject the cell at the end of the
        high shear dwell -->
  <Enum name="ExitDirection">up</Enum>
</Group>
<Group name="Ventilator">
  <!-- IP address of the ventilator huub -->
  <String name="IPAddress">192.168.1.42</String>
  <!-- Default pressure at system boot -->
  <Number name="DefaultPressure">25</Number>
  <!-- Gauge pressure to use for the negative gas supply. This may reduce
        if the requested net pressure difference results in the positive gas
        supply pressure exceeding the maximum pressure. e.g. if the net pressure
        demanded is 160cmH2O, then the negative gas supply will drop to -80cmH2O,
        as the maximum positive gas supply pressure is 80cmH2O -->
  <Number name="OffsetPressure">31</Number>
  <!-- Swap the negative and positive chambers in software.
        The hardware is identical. This allows more flexible tubing setups -->
  <Boolean name="SwapChambers">true</Boolean>
</Group>
</Program>
</Configuration>
```

B. Electronic Supplementary Information

Some of the results presented in this thesis are also available in video format. These are listed below in the order in which they relate to the thesis chapters. With the exception of those videos available directly from the journal, the complete set of videos can be accessed from Vimeo at vimeo.com/album/3391047.

The password for all videos is AppendixB.

B.3. Simulation

- B.3.1. Electronic Supplementary Information for the LabChip paper is available from the Royal Society of Chemistry website at:
<http://pubs.rsc.org/en/Content/ArticleLanding/2011/LC/c1lc20191c>
Three videos are available through the publisher — these correspond to the three simulation cases, where the resolution and noise parameters are varied (see the paper for further details).
- B.3.2. A video showing a potential sorting application of the device
Direct Link: vimeo.com/mdcurtis/b-3-2
In this video, it has been assumed that the cells have been identified and classified into two groups as soon as they enter the field of view. Some examples of methods that might be used to do this include: fluorophore tagging on the cells (similar to FACS sorting); shape measurement *via* image analysis; magnetic or electrical characteristics. Alternatively, the cell could be first be trapped, analysed for dynamic characteristics and sorted, although this last case is not specifically addressed by this video.

B.8. Clinical Pilot Study

- B.8.1. A short video demonstrating a cell entering the channel and becoming trapped
Direct Links: vimeo.com/mdcurtis/b-8-1s (slowed down); vimeo.com/mdcurtis/b-8-1r (realtime)
In this video, shown first slowed down, and then at realtime, a live red blood cell is trapped in the cross-slot as a result of active control, as described in chapter 7.
- B.8.2. A video demonstrating live manipulation of system parameters in the user interface developed as part of this thesis
Direct Link: vimeo.com/mdcurtis/b-8-2)
Here, the system was first purged by increasing the pressure to 16 kPa (the system is calibrated in cmH₂O, hence in the video 160 cmH₂O is shown), before reducing the system pressure to 2.5 kPa for trapping. After the pressure has stabilised, the automatic trapping macro is enabled and the system begins to search for cells to trap. Cells identified by the system are shown on screen as red markers. Once identified, the cells are trapped for a short period while data is recorded, then released.