

New Numerical Techniques to
Quantify and Predict the Effect of
Entrainment Defects, Applied to
High Pressure Die Casting

By

Robert Watson

A thesis submitted to
the University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

School of Metallurgy and Materials
Engineering and Physical Sciences
University of Birmingham
Edgbaston
Birmingham
B15 2TT
United Kingdom

October 2016

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

High Pressure Die Casting is an attractive option for automotive manufacturers, as it has a number of advantages over wrought process routes; it potentially allows stiff, lightweight aluminium structures to be designed with part consolidation, while also offering a lower piece cost than most other casting processes.

However, a lack of understanding of the defects which may result from the process leads to conservative design. An improved understanding of these defects could allow castings to be employed with greater confidence, leading to lighter vehicle structures.

Towards this goal, a novel algorithm was developed to predict the formation of entrainment defects, which may significantly limit the strength of cast parts. This model was integrated into FLOW-3D, a commercial transient fluid dynamics solver, and applied to a commercial High Pressure Die Casting.

Further theoretical advances were made, which offer a means of extrapolating a purely spatial distribution of predicted damage to a location specific statistical distribution, an improved way of statistically characterising the contribution of each defect type in mechanical data, and a means of finding correlation parameters for statistical distributions, such that the statistical variation in strength may be predicted at arbitrary locations within a casting.

A combination of casting experiments and numerical experiments were performed, to evaluate the fidelity of the new algorithms and underlying fluid flow solution, and also to test the influence of entrainment defects on the strength of High Pressure Die Castings.

Overall, it was found that defects formed by air entrainment can significantly limit the strength of High Pressure Die Cast parts. However, these defects are affected by other factors, and other defect types may also be significant. The methods and techniques explored in this work showed promise, but further advances will need to be made, before accurate *ab initio* strength prediction for High Pressure Die Cast parts can become a reality.

Acknowledgements

For most of the time I have spent on this research, I have worked in the Interdisciplinary Research Centre for Net-Shape Manufacturing. The words on the side of this building, a quote from the autobiography of Francois Jacob ¹, read:

Day science employs reasoning that meshes like gears... One admires its majestic arrangement as that of a da Vinci painting or a Bach fugue. One walks about it as in a French formal garden... Night science, on the other hand, wanders blindly. It hesitates, stumbles, falls back, sweats, wakes with a start. Doubting everything... It is a workshop of the possible...where thought proceeds along sensuous paths, tortuous streets, most often blind alleys.

There have been times when I have felt that my work belongs to the “Day science” category, and that I could now rationalise apparently conflicting concepts. There have been many more times when I have felt that “Night science” would be a more appropriate term- Certainly, I have doubted everything, and taken blind alleys.

There are many individuals and organisations who have provided me with support throughout these darker times. Firstly, I must thank my academic supervisor, Dr W. Giffiths, who has provided plenty of constructive criticism, but also helped me maintain focus and confidence through the harder times.

My industrial supervisor, Dr T. Zeguer, his team, and Jaguar Land Rover as a whole, have been invaluable in providing many of the tools and resources which made this project possible. I am also grateful to the EPSRC, for their share in supporting this project.

¹ *The statue within: An Autobiography, New York, 1988, p296*

I must also thank the technical staff at the University of Birmingham, namely Adrian Caden and Peter Crammer, for their time, their insights, and their humour. A special mention should also go to Simon Ruffle, of JVM castings, without whom much of the characterisation work presented would not have been possible, and also to Odd-Sture Hopperstad and Torrod Berstad of NTNU for useful conversations on the subject of damage mapping.

Aside from the technical aspects, I could not have completed this work without my friends, fellow researchers at the University, family, partner and cat, who have all helped to keep me sane over the years.

Table of Contents

List of Figures.....	x
List of Tables	xxii
List of Symbols and Acronyms	xxiv
List of Publications	xxv
1 Introduction	1
1.1 Motivation.....	1
1.1.1 High Pressure Die Casting.....	1
1.1.2 Failure Modelling	2
1.1.3 Casting Simulation	3
1.1.4 Summary.....	4
1.2 Objectives.....	5
1.2.1 Identification of Significant Defects in HPDC	5
1.2.2 Realistic Prediction of the Distribution of Entrainment Defects	6
1.2.3 Realistic Prediction of the Strength at any Location in a Casting	6
1.2.3.1 Relate the predicted distribution of entrainment defects to mechanical properties	6
1.2.3.2 Translate predicted defect distribution into a predicted strength distribution in a Finite Element model.....	6
1.2.3.3 Evaluate the predicted strength distribution using a component test.	7
1.2.3.4 Evaluate flow predictions using an analogue for High Pressure Die Casting	7
2 Literature Review.....	8
2.1 Entrainment Defects.....	8
2.1.1 Introduction.....	8
2.1.2 Entrainment Defects as Crack Initiators	9

2.1.3 Relationship to Other Types of Defect	14
2.1.4 Particularities of HPDC	16
2.2 Modelling entrainment defect formation	18
2.2.1 Non-CFD (Computational Fluid Dynamics) Approaches	18
2.2.2 Computational Fluid Dynamics Approaches	19
2.2.3 Lagrangian Particle Models	21
3 Theoretical Developments	24
3.1 Surface Area Entrainment Code	24
3.1.1 Introduction	24
3.1.2 Overview of the Formulation	25
3.1.3 Quantitative Surface Tracking	25
3.1.4 Surface Stability Assessment	30
3.1.5 Particle placement	35
3.2 Statistical Mapping Theory	37
3.2.1 Introduction	37
3.2.2 Overview of the Implementation	39
3.2.3 Definition of Entrainment Damage	39
3.2.4 Evaluation of Permutations	40
3.2.5 Statistical Characterisation	42
4 Method	45
4.1 Experimental	45
4.1.1 Characterisation of Bearing Mount Commercial Casting	45
4.1.2 Whole Casting Structural Impact Test	50
4.1.3 Controlled Partial Vacuum Casting	53
4.1.3.1 Description of the Equipment	53
4.1.3.2 Plate Casting Experiments	59

4.2	Simulation	64
4.2.1	Surface Area Entrainment Code Robustness Characterisation	64
4.2.2	Parametric Investigation of Partial Vacuum Casting	67
4.2.3	Controlled Partial Vacuum Casting Simulations	69
4.2.4	Simulation of Commercial HPDC Experiments.....	72
4.2.4.1	‘Short Shot’ Simulation Fidelity and Mesh Sensitivity Test	72
4.2.4.2	Full Shot Simulations	75
4.2.5	Mapping Predicted Damage Statistics to Casting FE Model.....	77
4.2.6	Whole Casting Structural Impact Test	79
4.3	Analytical	82
4.3.1	MLE Extension for Multiple Concurrent Defect Types with Classification	82
4.3.2	MLE Extension for Correlation of Strength and Predicted Damage	84
5	Results	88
5.1	Surface Area Code Robustness Characterisation	88
5.1.1	Case 1: Falling Droplet.....	89
5.1.2	Case 2: Top Gated Return Wave	93
5.1.3	Case 3: Plunging Jet.....	97
5.2	Controlled Partial Vacuum Casting.....	101
5.2.1	Parametric Investigation	101
5.2.2	Time-Value Data from Experiments	105
5.2.3	Comparative Frames from Experiment and Simulation	111
5.3	Commercial High Pressure Die Casting	121
5.3.1	Short Shot Investigation	121
5.3.1.1	Experimental	121
5.3.1.2	Simulation and Mesh Sensitivity Analysis.....	123
5.3.2	Full Shot Exercises	128

5.3.2.1	Simulation	128
5.3.2.2	Casting Experiments and Tensile Characterisation	136
5.3.2.3	Weibull MLE Analysis	149
5.3.2.4	Correlation of Predicted Entrainment with Fracture Stress	154
5.4	Commercial Casting Structural Impact Test	159
5.4.1	Drop Test Simulation	159
5.4.2	Drop Test Experiment	164
5.5	Summary	168
6	Discussion.....	169
6.1	Applicability of Flow Modelling	169
6.1.1	High Pressure Die Casting Analysis	169
6.1.2	Controlled Partial Vacuum Casting	172
6.2	Evaluation of SAEC Entrainment model	175
6.2.1	SAEC Robustness.....	175
6.2.2	Mesh Sensitivity	177
6.3	Significance of Entrainment Defects in HPDC Parts.....	178
6.3.1	Applicability of Entrainment	178
6.3.1.1	Fractography	178
6.3.1.2	Mechanical Variation.....	179
6.3.2	Multiple Defect Weibull Fitting.....	181
6.4	Entrainment Correlation.....	185
6.4.1	Descriptive Quality of the Correlation	185
6.4.2	Inferential Quality of the Correlation	186
7	Conclusion.....	189
7.1	Identification of Significant Defects	189
7.2	Realistic Prediction of the Distribution of Entrainment Defects	190

7.3	Realistic Prediction of the Strength at any Location in a Casting.....	191
7.4	Summary	191
8	Future Work	192
8.1	Prediction of the Distribution of Entrainment Defect Distribution.....	192
8.2	Experimental Methods	192
9	References	194
10	Appendices	198
10.1	HPDC Fractography Catalogue	198
10.1.1	Samples Classified as “Ductile Shear”	198
10.1.1.1	FGFS-7L-A	198
10.1.1.2	BGFS-2R-B.....	199
10.1.1.3	BGFS-15R-C.....	200
10.1.1.4	FGSS-3L-A	201
10.1.1.5	BGSS-1L-A.....	201
10.1.2	Samples Classified as “Porosity”	203
10.1.2.1	FGFS-11L-C.....	203
10.1.2.2	BGFS-1R-A.....	204
10.1.2.3	FGSS-6L-C	205
10.1.2.4	FGSS-10L-B	206
10.1.2.5	BGSS-13L-B	209
10.1.3	Samples Classified as “Carbon Film”	210
10.1.3.1	FGFS-9R-C	210
10.1.3.2	BGFS-10L-A.....	212
10.1.3.3	FGSS-10R-A	214
10.1.3.4	FGSS-15R-A	215
10.1.3.5	BGSS-11L-A	217

10.1.4	Samples Classified as “Thin Oxide”	218
10.1.4.1	FGSS-12R-C	218
10.1.5	Samples Classified as “Dark Pore”	220
10.1.5.1	FGSS-2R-C	220
10.1.5.2	FGSS-8R-B	220
10.1.5.3	BGSS-3R-A.....	221
10.1.6	Samples Classified as “Lap Defect”	222
10.1.6.1	BGSS-1R-B.....	222
10.1.6.2	BGSS-3R-B.....	224
10.1.7	Samples Classified as “Exogenous Inclusion”	226
10.1.7.1	BGFS-11R-A.....	226
10.2	Tabulated HPDC Sample Data	228
10.2.1	Full Gate, Fast Shot.....	228
10.2.2	Blanked Gate, Fast Shot.....	232
10.2.3	Full Gate, Slow Shot	236
10.2.4	Blanked Gate, Slow Shot	240
10.3	Tabulated Simulation Control Options.....	244
10.3.1	Entrainment Algorithm Characterisation Simulations.	244
10.3.2	Controlled Partial Vacuum Casting Simulations.	246
10.3.3	High Pressure Die Casting Simulations	248

List of Figures

Literature Review

Figure 2-1: Ellingham diagram for some metals and oxides, including aluminium and other commonly cast metals [9].	9
Figure 2-2: Data extrapolated from bottom filled cast Al-7Si-Mg samples, tested under a cyclic loading, illustrated the difference in the initial flaw size distribution for filtered and unfiltered samples [13].	11
Figure 2-3: Typical HPDC microstructure showing distinct banding, from Otarawanna et al. [30].....	17

Theoretical Developments

Figure 3-1: The result of the cell pairing stage applied to a bubble.	26
Figure 3-2: Process for defining donor and acceptor cells.	28
Figure 3-3: An illustration of why loss and entrainment are not necessarily equivalent, from Watson et. al. [58].	31
Figure 3-4: Illustration of theoretical effect of body force direction on a fluid surface.....	32
Figure 3-5: Illustration of entrainment in a case that may be resolved using the pressure gradient criterion.....	34

Method

Figure 4-1: The casting selected for analysis, showing large flat areas of constant section thickness, appropriate for tensile testing.....	46
Figure 4-2: Typical trace of nominal HPDC piston speed.....	48
Figure 4-3: Illustration of the inserts used in the BGFS and BGSS trials.	48

Figure 4-4: Test bar sampling locations for left hand castings (A and C), and for right hand castings (B and D).....	49
Figure 4-5: Two views of the configuration of the drop test rig.....	51
Figure 4-6: Impact location and dimensions of impactor.	52
Figure 4-7: Photographs of the main elements of the controlled partial vacuum casting system.	54
Figure 4-8: Block diagram outlining the main elements of the control system.	56
Figure 4-9: The pattern (A) and a preliminary casting (B) from the plate casting experiments.	59
Figure 4-10: Dimensioned drawing of the gate geometry.	60
Figure 4-11: Target pressure below atmospheric over time, for each experiment.....	63
Figure 4-12: The test cases used to demonstrate the SAEC particle algorithm.	65
Figure 4-13: The 20 parameter combinations used to investigate velocity and mesh sensitivity.	68
Figure 4-14: Labelled drawing of the model set-up.	70
Figure 4-15: Full (A) and blanked (B) gate configurations used in HPDC simulation.....	75
Figure 4-16: Illustration of constraints and loading points for drop-test Finite Element simulations.....	80
Figure 4-17: Illustration of virtual test bar gauge lengths used for damage correlation.	85
Figure 4-18: Parameter combinations “Random Walk Distance” and “Smoothing Radius” used.	86

Results

Figure 5-1 – Selected time-frames from the falling droplet simulation using SAEC.	91
--	----

Figure 5-2: Fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames from Figure 5-1 are labelled for context.	92
Figure 5-3: Time differential of fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames from Figure 5-1 are labelled for context.....	92
Figure 5-4: Selected time-frames from the top gated return wave using SAEC.	95
Figure 5-5: Fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames from.....	96
Figure 5-6: Time differential of fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames.....	96
Figure 5-7 – Selected time-frames from the plunging jet simulation using SAEC.....	99
Figure 5-8: Fluid surface area and particle count for plunging jet simulation using SAEC. The positions of time frames from Figure 5-7 are labelled for context.	100
Figure 5-9: Time differential of fluid surface area and particle count for plunging jet simulation using SAEC. The positions of time frames from Figure 5-7 are labelled for context.....	100
Figure 5-10: Response of particle count to changes in velocity.	102
Figure 5-11: Response of particle count to changes in mesh size.	102
Figure 5-12: “Corrected” velocity response, assuming power law effect from mesh size.....	103
Figure 5-13: Each of the 20 simulations, at the same fill fraction, demonstrating the effect of velocity on flow morphology.	104
Figure 5-14: Recorded pressure-time curve for the “Slow” controlled partial vacuum casting experiment.	106

Figure 5-15: Recorded pressure-time curve for the “Slow TC” controlled partial vacuum casting experiment, which was the only test instrumented with thermocouples.	106
Figure 5-16: Recorded pressure-time curve for the “Medium” controlled partial vacuum casting experiment.	107
Figure 5-17: Recorded pressure-time curve for the “Medium Fast” controlled partial vacuum casting experiment.	107
Figure 5-18: Recorded pressure-time curve for the “Fast” controlled partial vacuum casting experiment.	108
Figure 5-19: Recorded pressure-time curve for the “Dump” controlled partial vacuum casting experiment.	108
Figure 5-20: Thermocouple data recorded in the “Slow TC” experiment, with the abscissa scaled to show the filling of the mould (A), and with the abscissa scaled to show all recorded data (B).	110
Figure 5-21: Example time frames from the “Slow” controlled partial vacuum experiment.	112
Figure 5-22: Equivalent time frames from the simulation of the “Slow” experiment.	112
Figure 5-23: Example time frames from the “Slow TC” controlled partial vacuum experiment.	113
Figure 5-24: Equivalent time frames from the simulation of the “Slow TC” experiment.	113
Figure 5-25: Example time frames from the “Medium” controlled partial vacuum experiment.	114
Figure 5-26: Equivalent time frames from the simulation of the “Medium” experiment.	114
Figure 5-27: Example time frames from the “Medium Fast” controlled partial vacuum experiment.	116

Figure 5-28: Equivalent time frames from the simulation of the “Medium Fast” experiment.	116
Figure 5-29: Example time frames from the “Fast” controlled partial vacuum experiment ..	117
Figure 5-30: Equivalent time frames from the simulation of the “Fast” experiment.	117
Figure 5-31: Example time frames from the “Dump” controlled partial vacuum experiment	118
Figure 5-32: Equivalent time frames from the simulation of the “Dump” experiment.....	118
Figure 5-33: Photographs of the feed side (A) and ejector side (B) of the short shot casting chosen for simulation.....	121
Figure 5-34: X-ray photomontages of the left and right hand side (as shown in Figure 5-33 A) of the short shot casting chosen for simulation.....	122
Figure 5-35: The final frame of the short shot simulation, which was performed using the fine mesh, coloured by solid fraction.....	123
Figure 5-36: The final frame of the short shot simulation, which was performed using the intermediate mesh, coloured by solid fraction.....	124
Figure 5-37: The final frame of the short shot simulation, which was performed using the coarse mesh, coloured by solid fraction.....	124
Figure 5-38: Specific fluid kinetic energy over time, for each of the three short-shot simulations.....	125
Figure 5-39: Specific fluid kinetic energy over time, for each of the three short-shot simulations.....	126
Figure 5-40: Specific fluid kinetic energy over time, for each of the three short-shot simulations.....	126

Figure 5-41: Selected frames from FGFS (Full Gate, Fast Shot) simulation	129
Figure 5-42: Selected frames from BGFS (Blanked Gate, Fast Shot) simulation.....	130
Figure 5-43: Selected frames from FGSS (Full Gate, Slow Shot) simulation	132
Figure 5-44: Selected frames from BGSS (Blanked Gate, Slow Shot) simulation	133
Figure 5-45: Plot of particle history over time for the four full show HPDC simulations.....	135
Figure 5-46: Images from SEM fractography of sample FGSS-10L-B. A shows the transition between apparent interdendritic shrinkage porosity and a surface with a “crumpled bag” texture. B shows an EDS result from the “crumpled bag” region, as labelled in C.	137
Figure 5-47: Image from SEM fractography of sample BGSS-1R-B, showing a pore (right) with the “crumpled bag” surface texture.	139
Figure 5-48: Images from SEM fractography of sample FGSS-15R-A, showing a large region with an unusual texture, towards the left of image A. B compares wide area EDS spectra, taken from the left (red line) and right (blue line) sides of frame A.	139
Figure 5-49: Box-Plots summarising tensile data from all samples, according to trial and test bar position.	142
Figure 5-50: Cumulative UTS data from tensile testing of specimens from FGFS (Full Gate, Fast Shot) positions A, B and C, categorised by critical defect.....	144
Figure 5-51: Cumulative UTS data from tensile testing of specimens from BGFS (Blanked Gate, Fast Shot) positions A, B and C, categorised by critical defect.	145
Figure 5-52: Cumulative UTS data from tensile testing of specimens from FGSS (Full Gate, Slow Shot) positions A, B and C, categorised by critical defect.	146
Figure 5-53: Cumulative UTS data from tensile testing of specimens from BGSS (Blanked Gate, Slow Shot) positions A, B and C, categorised by critical defect.....	147

Figure 5-54: Data from “Fast Shot, Full Gate” experiment, after conversion to C-L fracture parameter, and showing the results of multi-defect MLE Weibull analysis.	151
Figure 5-55: Data from “Fast Shot, Blanked Gate” experiment, after conversion to C-L fracture parameter, and showing the results of multi-defect MLE Weibull analysis.....	151
Figure 5-56: Data from “Slow Shot, Full Gate” experiment, after conversion to true stress, and showing the results of multi-defect MLE Weibull analysis.	153
Figure 5-57: Data from “Fast Shot, Blanked Gate” experiment, after conversion to true stress, and showing the results of multi-defect MLE Weibull analysis.	153
Figure 5-58: Effect of “Initial Smoothing Radius” on the maximum likelihood output of second stage optimisation.	155
Figure 5-59: Effect of “Random Walk Distance” on the maximum likelihood output of second stage optimisation.	155
Figure 5-60: True stress fracture data derived from “Fast Shot, Full Gate” experiment, overlaid with multi-defect MLE Weibull analysis, and entrainment damage correlation.	157
Figure 5-61: True stress fracture data derived from “Slow Shot, Full Gate” experiment, overlaid with multi-defect MLE Weibull analysis, and entrainment damage correlation.	157
Figure 5-62: True stress fracture data derived from “Fast Shot, Blanked Gate” experiment, overlaid with multi-defect MLE Weibull analysis, and entrainment damage correlation.	158
Figure 5-63: The five mapped strength distributions, which were applied to the cast component in the drop test simulations.	160
Figure 5-64: A-E) Each of the five drop test simulations which used a stochastic strength map, F) Stress distribution without a stochastic strength distribution applied. Frames taken 0.001 s from the start of the simulation, at the onset of fracture. Each frame is coloured by “maximum principle stress”.....	161

Figure 5-65: Frames from drop test simulation, without fracture strength model.	163
Figure 5-66: Frames from 0.5 m drop test experiment, equivalent to those illustrated in Figure 5-64, with time normalised from the point of impact.	164
Figure 5-67: Frames from high speed video of the two drop tests carried out, with a drop height of 1 m.	165
Figure 5-68: Post-test photographs of the two castings tested with a drop height of 1 m.....	166
Figure 5-69: Fracture surfaces from castings with 1 m drop height.....	167

Discussion

Figure 6-1: Overlaid fluid surface shapes extracted from the “Medium” experiment (Figure 5-25), at $t_{\text{norm}}=0.08$ s (light grey), $t_{\text{norm}}=0.23$ s (dark grey) and $t_{\text{norm}}=0.28$ s (black).	175
Figure 6-2: Confidence region plots based on the Weibull fitted parameters for each entrainment data set, at test bar locations A, B and C.	183

Appendices

Figure 10-1: SEM SEI image showing the representative fracture surface for this sample...	198
Figure 10-2: SEM SEI image showing some apparent shrinkage porosity which was found in the sample.	199
Figure 10-3: High magnification SEM SEI image of what appears to be shrinkage porosity in this sample.	199
Figure 10-4: High magnification SEM SEI image of the Micro-Void Coalescence fracture, which was representative of the sample.	200
Figure 10-5: Low magnification optical image of the sample.	200
Figure 10-6: Low magnification SEM SEI image, giving an overview of the fracture face..	201

Figure 10-7: SEM SEI image of the Micro-Void Coalescence fracture, which was representative of the sample.	201
Figure 10-8: Low magnification SEM SEI image, giving an overview of the fracture face..	202
Figure 10-9: High magnification SEM SEI image, showing what appears to be shrinkage and trapped air porosity.	203
Figure 10-10: SEM SEI image showing a cluster of small pores.....	203
Figure 10-11: SEM SEI image giving an overview of the fracture face, and showing pores of various sizes, clustered away from the edges of the sample.....	204
Figure 10-12: Optical image providing an overview of the fracture surface.	204
Figure 10-13: Close-in SEM SEI image of the small pores scattered around the sample.....	205
Figure 10-14: SEM SEI image providing an overview of the fracture surface.....	205
Figure 10-15: SEM SEI image showing what appears to be an air pore around 1 mm in diameter in this sample	206
Figure 10-16: SEM SEI image showing a close up of the wrinkled texture present in this pore.	206
Figure 10-17: SEM SEI image showing a close up of the edge of the wrinkled region	207
Figure 10-18: SEM SEI image showing the transition between the wrinkled region and apparent shrinkage porosity.	207
Figure 10-19: SEM SEI image showing the regions of interest for the spectra in Figure 10-20yy, Spectrum 3 was taken from the edge of one of the wrinkles, and Spectrum 4 was taken from an apparent MVC region.	208
Figure 10-20: The EDS spectra acquired from the locations in Figure 10-19	208

Figure 10-21: SEM SEI overview of the sample, showing a cluster of porosity towards the through-thickness centre of the sample.	209
Figure 10-22: The central section of Figure 10-21, imaged at a greater magnification.	209
Figure 10-23: SEI SEM image of the sample, showing features where the metal appears to have delaminated at film-like features perpendicular to the applied stress. Also, the location of the spectra illustrated in Figure 10-24.	210
Figure 10-24: The EDS spectra, at the locations shown in Figure 10-23.	210
Figure 10-25: An SEI SEM image of a roughly planar defect.	211
Figure 10-26: A low magnification SEI SEM image providing an overview of the sample.	211
Figure 10-27: A higher magnification SEI SEM image of the feature in the centre of Figure 10-26.	212
Figure 10-28: SEM SEI image showing some apparent shrinkage porosity present in the sample.	212
Figure 10-29: SEM SEI image giving a representative view of the majority of the fracture surface.	213
Figure 10-30: SEM SEI image showing the planar film-like defect in this sample, together with the locations for the spectra shown in Figure 10-31	213
Figure 10-31: EDS spectra from the locations indicated in Figure 10-30.	214
Figure 10-32: An SEM SEI image, providing an overview of the sample.	214
Figure 10-33: An SEM SEI image, providing a higher magnification view of the planar defect towards the centre of Figure 10-32.	215
Figure 10-34: An SEI SEM image, showing a planar defect, with a feature that appears as a darker (smoother) vertical stripe towards the centre of the image.	215

Figure 10-35: Spectra taken from A) the darker central region in Figure 10-34 and B) the lighter normal fracture region to the left of the dark band in Figure 10-34.....	216
Figure 10-36: Low magnification SEI SEM image, showing a large planar defect (left), with a surface texture quite distinct from the normal fracture surface (right).....	216
Figure 10-37: Wide area EDS spectra taken from the Figure 10-36. The red line was sampled from the left hand region, and the blue curve was sampled from the right hand region. ..	217
Figure 10-38: Optical image showing a shiny planar defect on the right hand edge of the test bar.....	217
Figure 10-39: SEI SEM image showing an overview of the sample, with some porosity, and some fissure-like features visible.....	218
Figure 10-40: SEI SEM image showing a higher magnification view of the fissure-like feature towards the bottom of Figure 10-39.	218
Figure 10-41: SEI SEM image showing a cluster of small pores around a crevice.	219
Figure 10-42: Optical image showing dark, rounded, yet elongated features.....	220
Figure 10-43: Optical image showing dark elongated feature on the left, with rounded edges.	220
Figure 10-44: An SEI SEM image showing an optically dark feature.	221
Figure 10-45: An SEI SEM image showing a higher magnification view of the edge of the feature presented in Figure 10-44, suggesting some sort of compliant layer coating the dendrite network.	221
Figure 10-46: An SEI SEM image providing an overview of the sample, appearing to show a low level of porosity.	222
Figure 10-47: An SEI SEM image showing the edge classified as a lap defect. On this half of the test bar, the defect is continuous with the mould-wall face of the sample.....	222

Figure 10-48: An SEI SEM image showing a pore in the sample (right), with what appears to be oxide folds, and a possible trail of smaller pores extending to the left of the image....	223
Figure 10-49: An optical image of the sample, showing the lap-defect (top left) and a few pores on the order of 100 μm in size.	223
Figure 10-50: An SEI SEM image providing an overview of the sample, appearing to show a low level of porosity.	224
Figure 10-51: An SEI SEM image providing a close up of the lap defect. There appears to be a gradual transition between normal MVC microstructure at the left edge of the image, to a smooth morphology at the edge of the fracture surface.....	224
Figure 10-52: An SEI SEM image showing what appears to be a network of shrinkage porosity.....	225
Figure 10-53: An optical image of the sample, showing the lap-defect (top left) and a few pores on the order of 100 μm in size.	225
Figure 10-54: An optical image showing the inclusion towards the left of the image, together with some pores on the order of 100 μm in size.....	226
Figure 10-55: An SEI SEM image representative of the majority of the fracture surface.	226
Figure 10-56: An SEI SEM image showing normal fracture surface, and an inclusion with poor electrical conductivity in the bottom right quadrant. Also marked are the sampling locations for the EDS spectra in Figure 10-57.....	227
Figure 10-57: The EDS spectra sampled from Figure 10-56, indicating that the inclusion is calcium based mineral.	227

List of Tables

Method

Table 4-1: Specification for Castasil 37 alloy	45
Table 4-2: Process parameters which were constant for each trial.....	46
Table 4-3: Process parameters varied with each trial	47
Table 4-4: Parameters for Controlled Partial Vacuum Experiments.....	62

Results

Table 5-1: Effect of nominal mesh size on simulation run time.....	127
Table 5-2: Classification criteria for observed failure modes.....	141
Table 5-3: Proposed multi-component Weibull parameter estimates for “Fast Shot” Castings.	150
Table 5-4: Proposed multi-component Weibull parameter estimates for “Slow Shot” Castings.	152

Appendices

Table 10-1: Mechanical data for test bars from the full gate, fast shot trial.....	228
Table 10-2: Mechanical data for test bars from the blanked gate, fast shot trial.....	232
Table 10-3: Mechanical data for test bars from the full gate, slow shot trial.	236
Table 10-4: Mechanical data for test bars from the blanked gate, slow shot trial.	240
Table 10-5: Control options for simulations used for “Falling Droplet” case.....	244
Table 10-6: Control options for simulations used for “Top Gated Return Wave” and “Plunging Jet” cases.....	245
Table 10-7: Control options for parametric investigation.	246

Table 10-8: Control options for simulating the physical experiments.....	247
Table 10-9: Control options for the shot-sleeve simulations.	248
Table 10-10: Control options for the main filling simulations.	249
Table 10-11: Control options for the solidification simulations.....	250
Table 10-12: Control options for the short-shot simulations.	251

List of Symbols and Acronyms

Common symbols as used in equations

d	Variable describing arbitrary “damage” arising from the presence of defects
β	Shape parameter in Fréchet distribution
μ	Offset parameter in Fréchet or Weibull distributions
m	Shape parameter in Weibull distributions
σ	Stress, or strength in the context of a constant (as denoted by super/sub-script)
P	Probability of a given event or occurrence
ρ	Density
g	Acceleration due to gravity

Abbreviations used throughout the text

ADC	Analogue to Digital Converter
BGFS	Blanked Gate, Fast Shot
BGSS	Blanked Gate, Slow Shot
BOEC	Boolean Oxide Entrainment Code - an approach to free surface flow in an Eulerian mesh.
CDF	cumulative distribution function
CFD	Computational Fluid Dynamics
COTS	Commercial Off-The-Shelf
DCM	Die Casting machine
EDS	Energy Dispersive Spectroscopy
ESCs	externally solidified grains
FGFS	Full Gate, Fast Shot
FGSS	Full Gate, Slow Shot
HPDC	High Pressure Die Cast(ing)
MLE	Maximum Likelihood Estimation

MVC	Micro-Void Coalescence
NTNU	Norges Teknisk-Naturvitenskapelige Universitet (Norwegian University of Science and Technology)
OFEM	Oxide Film Entrainment Model
PDF	Probability distribution function
RPT	Reduced pressure Test
SAEC	Surface Area Entrainment Code
SEI	Secondary Electron Imaging
SEM	Scanning Electron Microscope
UTS	Ultimate Tensile Strength
VOF	Volume of Fluid

List of Publications

Watson, R., et al. *The Response of Entrainment Defect Distribution to Varied Fluid Flow Parameters in the High Pressure Die Casting of Al Alloys*. in *Materials Science Forum*. 2013. Trans Tech Publ.

Watson, R., et al. *Application of a Novel Entrainment Defect Model to a High Pressure Die Casting*. in *Advanced Materials Research*. 2014. Trans Tech Publ.

Watson, R., et al., *Towards Location Specific Statistical Fracture Prediction in High Pressure Die Castings*, in *10th European LS-DYNA Conference*. 2015, DYNAmore: Würzburg.

1 Introduction

1.1 Motivation

1.1.1 High Pressure Die Casting

High Pressure Die Casting (HPDC) is a permanent mould casting method, which provides a number of advantages over methods that use a non-permanent mould, such as sand or investment casting. The principle advantage of a permanent mould is that moulds do not need to be continuously refabricated, which contributes to the high dimensional accuracy of such parts, and can reduce the overall cost of the process.

Additionally, since the mould is a steel die, high heat transfer rates can be achieved. With no risk of the melt penetrating the mould, an intensification stage with pressures of ~100 MPa can be used to increase the mechanical properties of the part. In all, this allows a “shot”, comprising of one or more parts together with the running system and overflows, to be produced roughly once a minute.

The high fluid velocities, which may be obtained with the process, fill the mould cavity in tens of milliseconds, allowing relatively complex and thin walled (~3mm) geometries to be cast with a good surface finish.

These qualities make High Pressure Die Casting (HPDC) highly attractive for a number of industrial sectors, as high volumes of light alloy castings may be produced at a relatively low cost per unit. The process also allows geometric features such as ribs to be included, that could make the part lighter, stiffer, or easier to integrate into an assembly.

There are naturally disadvantages to HPDC. Firstly, there is a high initial cost in fabricating the permanent moulds, which means a relatively low overall cost can only be achieved with a high production volume. Secondly, the process has a reputation for being highly vulnerable to defects formed by the high filling speeds [1], supported by work by LaVelle [2].

This first factor also makes HPDC very difficult to experiment with, since it would be very expensive and time consuming to commission a die specifically for a set of experiments. Adapting an existing die for an experiment, perhaps by adding instrumentation, or making any changes to the geometry would need to be authorised by the owner of the die, since any changes are not easily undone.

Real time X-ray technology has allowed the fluid processes in ceramic mould castings to be better understood, however the technology does not yet exist that can image the flow of aluminium within a large steel die, particularly with a sufficiently high frame-rate and signal to noise ratio to be of value.

This would explain, at least in part, why the majority of literature concerning HPDC consists of industrial case studies, and there is very little research which asks significant questions about the underlying physical processes, which introduce variability to the castings.

It should be noted that the scope this work did not include the types of microstructural defect such as intermetallics, which are primarily driven by alloy chemistry.

1.1.2 Failure Modelling

For a casting to be successful, it must meet the customers' expectations. From the point of view of the foundry, this means producing a casting which is the correct shape, where the composition is within specifications, and which is free from defects (up to some agreed level). This is the *status quo*.

Alternatively, from the point of view of their customer – for example an automotive manufacturer – it means that the part should be designed so that the expected properties are achievable and realistic: for this to be the case, the simulations used to help design the part must accurately predict the behaviour of the material.

In the author's experience this is still an area for improvement, at least in the automotive industry. Aside from any inaccuracies introduced during structural analysis, it is far from

standard practice to account for the variability in the cast material, although some work by Dørum et. al. [3, 4] and Fagerholt et. al. [5] has demonstrated that it is possible.

This variability can arise from a number of sources. The yield stress, strain hardening behaviour and crack propagation resistance of the material will vary according to metallurgical factors such as composition and microstructure.

Moreover, the strength of the material will be limited, to a variable extent, by the presence of any defects or inclusions introduced during the casting process, or already present in the charge material before casting. Turbulent filling has been shown to be significant factor in sand castings [6], as will be expanded upon in the Literature Review.

1.1.3 Casting Simulation

The disadvantages of performing experiments with HPDC – time, cost and difficulty of extracting meaningful results – are all potential advantages of simulation. In principle, the design of a new die can be tested in simulation in a far shorter time and at much lower cost than constructing a new die.

More importantly, the simulation results will show the predicted fluid flow and temperature distribution, and can indicate where problems might occur. These problems could be occasional, such as gross porosity, or slowly damage the die, such as excessive thermal shock. In either case, it might take a lot of real world experimentation to detect the issue, which may lead to an unacceptable scrap rate.

However, these arguments pre-suppose that the simulation is sufficiently accurate to design a mould and process, which consistently produces good castings. As Ohnaka states in his review [7], there are aspects of the casting process which are difficult to predict using current technology. Some of these are the accurate prediction of bubbles, pores and entrainment defects, as will be discussed further in the Literature Review.

In general terms, modelling the flow of metal in a casting is a transient (time dependant) free surface problem (the fluid may or may not be anywhere in the domain), where drag is affected by solidification, which is affected by the temperature distribution, which is partially affected by the transport of heat by fluid flow. Surface tension and the mould cavity gas may also be modelled to some degree. Overall, the equations are complex and strongly interdependent. Nonetheless, multiple Commercial Off-The-Shelf (COTS) packages, such as FLOW-3D, MAGMAsoft and PRO-CAST, claim to provide accurate solutions, and foundries routinely depend on their results.

Perhaps the solutions are accepted because it is not possible to determine the same information experimentally. However, experiments must be performed which challenge the assumptions of the simulations, or help isolate their parameters, for the quality of the simulations to improve.

1.1.4 Summary

There are many reasons why High Pressure Die Casting is an appealing process for designers in high volume industries, such as the automotive sector: low unit cost, good dimensional accuracy, and good design freedom, compared to non-cast routes. This design freedom would allow a single cast part to be used in the place of an assembly of wrought parts, reducing complexity, and avoiding stress concentrations at the joints. Moreover, the form of the part can potentially be optimised to minimise weight, while retaining the required structural performance.

However, because the sources of the variability in their mechanical strength are not well understood, it is rare that HPDC parts are used in structural applications. Where HPDC parts are used in structural applications, the material properties are deliberately under-estimated, so that the actual properties of the part are almost guaranteed to exceed the design minimum. This can also be undesirable, because if the part performs better than this low expectation,

then it may place too much load onto other parts in the assembly, and also there will have been lost opportunity for saving weight.

Saving weight (light-weighting) is a key driving force in design at Jaguar Land Rover, as it enables the environmental credentials of the vehicles to be improved, without needing to compromise performance or passenger comfort.

If software tools were available which would allow the strength variation in HPDC parts to be minimised or accounted for, then this could lead to the better, lighter, and more robust designs.

1.2 Objectives

This work was directed towards that ultimate goal of being able to design better, lighter castings, through a better understanding of the defects which affect them. This technology must be built upon a number of smaller methods and understandings. The focus of this work was not to explore the details of any one of these, but to try to connect the pieces of the puzzle, in order to better understand which pieces are missing, and what shape they would need to be.

1.2.1 Identification of Significant Defects in HPDC

It was initially assumed that entrainment defects would be a significant factor limiting the reliability of HPDC parts, as will be discussed further in the literature review. This assumption was questioned:

- What variation in mechanical strength can be attributed to entrainment defects?
- How much variability in mechanical strength should be attributed to other types of defect, and what are these other defects?

1.2.2 Realistic Prediction of the Distribution of Entrainment Defects

It became apparent early in this work that an entrainment prediction algorithm would need to be devised or refined, which could provide a reliable basis for further investigation. Such an algorithm should have good physical proportionality, but it should also be easily applied to an arbitrary 3-D flow case, and not significantly add to the simulation run-time.

1.2.3 Realistic Prediction of the Strength at any Location in a Casting

The final objective, which was

1.2.3.1 Relate the predicted distribution of entrainment defects to mechanical properties

To better understand the above two lines of enquiry, their results needed to be compared. If a quantitative function can be found which links the predicted severity of entrainment to the mechanical strength in a sample, this would support the notion that entrainment defects are significant, and also indicate that the modelling is accurate.

Moreover, this exercise served as a key hypothesis test: if observed variation from entrainment could not be linked to the predicted entrainment variation, then this would strongly suggest a logical flaw in the analyses.

1.2.3.2 Translate predicted defect distribution into a predicted strength distribution in a Finite Element model

A key deliverable set by the industrial partner, Jaguar Land Rover, was to be able to use the predicted defect distribution as a basis for mapping a strength distribution into a finite element model, using the correlation found. The aim was to allow the predicted variability of the castings to be factored into structural simulations.

1.2.3.3 Evaluate the predicted strength distribution using a component test.

Forming a “predictive” strength model was also necessary in order to test the model’s predictions against reality. An experiment needed to be performed which involved loading a whole component. A finite element simulation of this experiment that used the strength mapping algorithm would then allow the methodology to be evaluated.

1.2.3.4 Evaluate flow predictions using an analogue for High Pressure Die Casting

A core part of the overall evaluation must also consider the amount of faith which may be placed in the flow predictions; without accurate flow predictions, the entrainment predictions would be potentially meaningless.

While it is not feasible to analyse the flow in a HPDC mould cavity directly, some level of validation can be performed by observing an analogue under real time X-ray. As such, a casting experiment was envisaged which would employ lower velocities than in HPDC, and use an X-ray transparent mould. However, to extend knowledge beyond what can be found in literature, these experiments were to involve greater velocities than a typical sand casting, and use geometry closer to the HPDC paradigm.

Simulations of those experiments could then be compared to experimental observations, and so allow the applicability of the CFD models to be evaluated.

2 Literature Review

2.1 Entrainment Defects

2.1.1 Introduction

Entrainment, in the present context, is the action of something being taken into the bulk of a liquid, as a result of the motion and shape of the fluid. It is probable the reader will have observed entrainment, when water is poured from sufficient height into a glass, air will be entrained into the fluid and form bubbles.

In this case, the bubbles would eventually float to the surface of the liquid, which would then be smoothed by surface tension. After a few seconds, there would be no observable difference between a glass of water that was filled in a way that introduced bubbles, and another glass of water which was filled without entrainment.

This is not necessarily the case for reactive liquid metals, such as aluminium, as will be explored throughout this section. Campbell [8] proposed that air would necessarily react with the surface of any exposed liquid aluminium, to form a skin of alumina (Al_2O_3) over a very short timescale. As the below Ellingham diagram (Figure 2-1) shows, the Gibbs free energy of formation for aluminium oxide is one of the most negative of any of the reactions listed. This demonstrates that alumina is extremely stable, and to somehow reverse the reaction would have to go against a strong chemical potential. Moreover, this supports the notion that liquid aluminium will oxidise almost immediately on contact with air, and also that such a film would be stable over the timescale of the casting process. This would mean that it is entirely plausible that an oxide film formed during filling could affect both the solidification process, and the finished casting.




Figure 2-1: Ellingham diagram for some metals and oxides, including aluminium and other commonly cast metals [9].

2.1.2 Entrainment Defects as Crack Initiators

It is highly likely any entrainment event would include both the oxidised surface of the liquid metal, and some volume of mould gas. This is an implication of the “no slip condition” [10], which simply states that fluid moving against a solid boundary does not slip against it; viscosity then dictates that the closer a particle of fluid is to a solid boundary, the closer its velocity will be to the velocity of that solid boundary.

Let it be temporarily assumed that an entrainment defect will contain some volume of gas, after the casting has solidified (the validity of this assumption will be explored in a later sub-section). If this is the case, the entrainment defect would form a discontinuity in the

metal: a small region which stress cannot be transmitted across. If a stress is applied across such a discontinuity, that stress must be transferred to adjacent load-bearing material, to satisfy force equilibrium. The fracture behaviour of an entrainment defect may, to some extent, be likened to that of a crack.

The below relation, derived from linear elastic fracture mechanics, relates the failure stress, σ_F , to size of the crack, a , and the fracture toughness of the material, K_{IC} . The remaining parameter, α , accounts for the shape and location of the crack, and would also be influenced by other nearby discontinuities [11, 12].

$$\sigma_F = \frac{K_{IC}}{\alpha\sqrt{\pi a}} \quad \text{Equation 2-1}$$

Nyahumwa [13] used a fatigue based approach to determine the effective initial flaw size distribution in several sets of cast Al-7Si-Mg samples; this metric effectively combines α and a (above) into a single parameter. The Paris law may be used to describe the rate that a crack grows in a metal under cyclic loading [14]. Given the critical crack size at fracture, and the crack growth parameters, this equation allowed Nyahumwa to infer the initial effective flaw size of the most damaging defect in a sample, based on the fatigue life of that sample. The results from two of those data sets are illustrated in Figure 2-2; for these two sets of data, it seems that the addition of a ceramic filter to the running system has significantly improved the quality of the cast material. This result may be attributed to a reduction in the damage caused by entrainment defects for the “filtered” tests, as filters are believed to reduce the formation of entrainment defects. It should be noted that the largest defect in either data set is significantly smaller than diameter of the pores of the filters.




Figure 2-2: Data extrapolated from bottom filled cast Al-7Si-Mg samples, tested under a cyclic loading, illustrated the difference in the initial flaw size distribution for filtered and unfiltered samples [13].

Tiryakioğlu [15] analysed Nyahumwa’s data sets, alongside those of other researchers, to determine which of three “extreme value distributions” is most appropriate to describe the observed effective flaw size distributions. In theory an extreme value distribution should be appropriate to model the size distribution of the defects which are observed to cause failure, because there is an inherent sampling bias towards the defects most likely to initiate fracture.

Statistical parameters for the generalised extreme value distribution were estimated for each data set using the Maximum Likelihood Estimation (MLE) approach, which is further detailed in the Method, section 4.3.1. This distribution includes each of the three types of extreme value distribution as a special case; in effect the ζ parameter “switches” between each case. By determining the confidence intervals for this ζ parameter, Tiryakioğlu was able to

assess how well each of the distributions describe each data set. It was found that the Type II (Gumbel) distribution was most representative.

This would suggest that the overall defect distribution – including those which did not initiate fracture – follows a log-normal distribution. This finding is supported by Nyahumwa's work to characterise the pore size distribution in metallographic sections of cast material [13]. However, the Gumbel distribution does not have upper or lower bounds [16]. This means that a random sample from the Gumbel distribution has some chance of being negative, which does not make physical sense. Tiryakioğlu's results [15] also showed that the Fréchet distribution could be applied to each studied data set within confidence limits. This may make more physical sense, since the Fréchet distribution does have a lower limit, when used to model upper extremes, such as the largest defect [16].

In engineering terms, there are a number of reasons why cast parts are not generally characterised using fracture mechanics terms. The first is simplicity: it is easier to use a single metric, such as failure stress, as opposed to evaluating Equation 2-1 for the most damaging defect, however its properties may be determined. More significantly, the material will yield before K_{IC} is reached for most castings of interest, meaning that linear elastic fracture mechanics would not apply. However, the fundamental principles remain valid- that a discontinuity in a material can initiate unstable crack growth (i.e. fracture), and that this can occur at a lower stress, if the equivalent flaw size of the defect is larger.

As indicated by Figure 2-2, the strength of the castings cannot properly be described as constant. Campbell and Green wrote one of the first works [6] which investigated how the variation in the UTS of cast aluminium may be linked to the formation of entrainment defects. In their results, the tensile test bars cast using a mould geometry that was intended to minimise entrainment showed a significantly reduced chance of failing at a low strength.

The authors further showed that the statistical distribution of the UTS of their samples could be most adequately described by the 2-parameter Weibull distribution [17], albeit using

a statistical analysis which was relatively rudimentary [18, 19]. The Weibull distribution is based on weakest link theory, which makes it well suited to model brittle fracture. The 2-parameter form of the distribution is given in Equation 2-2, where P is the probability that a sample would fail at a stress lower than σ , based on the distribution parameters, σ_0 and m . The position parameter (σ_0) is an average for the distribution, describing the expected 63rd percentile value for σ . The relative amount of variation in the distribution is defined by the Weibull modulus (m), where a lower value corresponds to greater scatter.

$$P = 1 - \exp \left[- \left(\frac{\sigma}{\sigma_0} \right)^m \right] \quad \text{Equation 2-2}$$

Campbell and Green found that the modulus parameter decreased with the supposed severity of entrainment. Reilly [20] later simulated these experiments of Campbell and Green using an Oxide Film Entrainment Model (OFEM), which will be discussed in its own right in a later section. Reilly's results supported Campbell and Green's assumptions for which of their three running systems would introduce a greater or lesser number of entrainment defects. Taken together, these works strongly suggest that the more entrainment defects which are introduced to cast material during filling, the more variable the mechanical strength of that material will be.

Yue and Green [21] performed a similar experiment, where cast plates were produced using three different running systems. Samples were then extracted from these plates and tested in bending. Yue also simulated these tests using Reilly's Oxide Film Entrainment Model (OFEM), and again found that the variation in mechanical strength correlated with the number of entrainment defects predicted to be introduced during mould filling.

In works from SIMLab (NTNU) [3-5], the authors suggested that a more appropriate measure for the strength of HPDC aluminium samples would be the Cockcroft-Latham fracture parameter [22]. The authors showed that, unlike UTS, the metric can adequately

account for failure in HPDC parts in stress states ranging from pure tension to pure shear. In their casting experiments [3], the authors' results indicated that the Weibull distribution of Cockcroft-Latham fracture parameter was affected by the severity of entrainment.

Timelli and Bonollo [23] used an empirical "Quality mapping" function (based on strength and failure strain), that to describe the strength of samples from a set of aluminium die castings. For the locations tested, the authors showed a correlation between the "quality function" and the area of defects on the fracture surface, and each of these varied from location to location.

Each of the experiments cited so far in this sub-section indicate a strong link between an increased chance of premature failure, and a supposedly more severe introduction of entrainment defects during mould filling. Across these works, however, there isn't enough evidence to categorically state that this link is one of direct cause and effect.

The results which most strongly support a direct cause and effect relationship are those where the defect that caused failure at a lower stress has been identified on the fracture surface using a Scanning Electron Microscope (SEM), and found to be an oxide of the parent material using a technique like Energy Dispersive Spectroscopy (EDS). Nyahumwa presented a number of examples of such observations [13].

Nevertheless, it is entirely possible that a defect positively identified as "oxide film" could have formed before mould filling, or that any of the other failures not observed in this way were falsely attributed to an entrainment defect. This raises the question of how (or whether) it is possible to easily identify the critical flaw as an entrainment defect.

2.1.3 Relationship to Other Types of Defect

The view of the entrainment defect commonly proposed, particularly in earlier works on the subject, is that of the so-called "oxide bi-film"; that is, two oxide films essentially in contact "dry side to dry side" [8]. This morphology is distinct from hydrogen and shrinkage

“porosity”, which are generally more rounded. When the oxide bi-film was first described scientifically, hydrogen and shrinkage porosity could be explained, and the oxide bi-film was a “new” form of defect, which needed to be explained in terms of entrainment. However, this point of view implies that an entrainment defect would always be formed with the “bi-film” morphology, and that it would stay that way throughout the casting process.

More recent research indicates that perhaps these three types of defect may not be as distinct as first thought. First, there is the theoretical perspective. The growth of a shrinkage pore can be understood as the need to supplement the volume change on solidification by the creation of “void”. Similarly, hydrogen porosity is driven by the chemical potential of super-saturated dissolved hydrogen, which can form when a casting solidifies. However, these mechanisms do not explain the nucleation of these defects.

Experiments have been conducted [24, 25], which used a Reduced Pressure Test (RPT) method along with X-ray equipment, to reveal film defects in the melt – presumed by the authors to be oxide films – and their growth into irregular bubbles. While the authors did not provide direct evidence that the observed bubbles were nucleated by bi-films, the second of these studies drew an alloy dependant correlation between the average number of pores and the total length of the films, indicating that the film was acting as a nucleant.

Raiszadeh and Griffiths [26] investigated the evolution of an oxide film defect under controlled experimental conditions, in terms of the gasses inside an entrainment defect. This research led to a description of the rate at which atmospheric oxygen and nitrogen are consumed by the melt. The authors developed a computer model based on his results [27, 28], which indicated that the lifetime of the internal atmosphere of an average defect could be around 3 minutes depending on the hydrogen concentration in the melt. In a later work, El-Sayed et. al. raised this estimate to between 5 and 30 minutes, depending on the alloy [29].

The results imply that the view of an entrainment defect strictly as a thin “bi-film”, is flawed, because they suggest that the “bi-film” morphology is a possible second stage in the

life-cycle of an entrainment defect. It is possible that with a sufficient solidification rate an entrainment defect could not take on the bi-film morphology before the casting solidifies.

As Raiszadeh and Griffiths [26] suspected, and later confirmed by El-Sayed et al. [29], the consumption of the oxygen and nitrogen within an entrainment defect occurs in two stages: first the oxygen is consumed, as oxygen is more reactive. Once the oxygen level is reduced to around 5%, nitrogen starts to react to form aluminium nitride. Both set of results showed that hydrogen can also diffuse through the film over time. This means that it is possible for an entrainment defect to evolve into something which may be identified as a hydrogen pore. In fact, it implies that an entrainment defect will either evolve into the bi-film morphology, or a hydrogen pore, depending on the hydrogen content of the melt.

In summary, a substantial body of research suggests that there is not necessarily a way to distinguish whether or not a given pore (shrinkage, hydrogen, etc.) observed in a cast part would exist without an entrainment event. It could be argued that if the pore would not exist without an entrainment event, then it should be classified as an entrainment defect. The simplest distinguishing feature of an entrainment defect is therefore an empirical link between the effect of the defect on the mechanical properties of a casting, and the severity of entrainment in the casting process. In physical terms, the link could be direct or indirect; however this line of reasoning requires a robust measure of the severity of entrainment.

2.1.4 Particularities of HPDC

Much of the above work was performed using sand castings, or other set-ups where the solidification rate is relatively slow. High Pressure Die Castings are particular in that the motion of the melt is significant throughout much of the solidification process. The below figure, from Otarawanna et al. [30], shows a full cross section of an HPDC sample, etched such that α -Al appears white, and Al-Si eutectic shows as black.



Figure 2-3: Typical HPDC microstructure showing distinct banding, from Otarawanna et al. [30]

A surface skin is visible in this image, marked by a solid white arrow. This is the first layer to solidify during filling of the mould, and is not always present in samples, as this skin may re-melt during filling. This layer is characterised by a fine grain structure, and a significant depletion of eutectic. Some authors [31, 32] have reported this layer to have improved mechanical properties.

Also visible in this image are occasional α -Al rosettes with a diameter on the order of 100 μm , substantially larger than average. These are commonly referred to as “Externally Solidified Crystals” (ESCs), and are believed to first form in the shot sleeve, before being transported to the die with the bulk of the fluid.

Two eutectic rich bands may be seen in the above figure, marked with dotted white arrows, and are referred to by Otarawanna et al. as “Defect bands”. The eutectic enrichment is believed to be formed by rheological phenomena in the semi-solid flow. The high eutectic

concentration would indicate these bands as the last to solidify, and would therefore be vulnerable to defects including shrinkage porosity.

2.2 Modelling entrainment defect formation

2.2.1 Non-CFD (Computational Fluid Dynamics) Approaches

One method of assessing the extent of entrainment during casting, would be to observe the casting process, perhaps using time real-time X-ray equipment, and pick out events that fit our understanding of how entrainment defects form: regions of surface turbulence, events where the liquid metal surface come together, or instances of obvious bubble entrainment. This method involves the application of a mental model, and is essentially the process used by Nyahumwa [13], Green and Campbell [6], among others.

However, this approach relies on the interpretation of the experimenter; it is not quantitative and it cannot be applied before a casting is produced. A mental model, which may be gradually gained from performing casting experiments can be applied to the design of a new mould, but without an experiment the new design could not be objectively verified.

Campbell has published a number of guidelines [33, 34] on the design of running systems, based on his experience and intuition. It should be noted that much of this experience formed whilst seeking to demonstrate the benefits of minimal entrainment, and so the findings were self-selecting to some extent. Additionally, many of the published experiments [6, 35-37] were based in the regime of low velocity sand castings, with low hydrogen content, where a difference could be resolved between entraining and non-entraining conditions. As such, it does not hold that Campbell's conclusions necessarily apply outside this regime, or that the mechanisms investigated are the only ones that may cause a significant effect.

One of the key elements of Campbell's advice is the "0.5 ms⁻¹ rule", a semi-quantitative entrainment criterion [38], which states that the melt velocity should not exceed 0.5 ms⁻¹.

This argument was derived by relating the kinetic energy of the fluid flow and the surface tension of the melt. The theory suggests that above 0.5 m^{-1} , surface tension cannot prevent the fluid from splashing, and so the fluid will splash and so form entrainment defects. In some circumstances this makes sense; if the fluid encounters a wall then its energy must be transferred into a splash. However, the way the rule is phrased suggests that it refers to a free flowing stream, and that surface tension might not control ripples formed by surface turbulence. In this case, one might suppose that the turbulent portion of overall kinetic energy would be more applicable than the macroscopic kinetic energy [39].

Furthermore, the theory does not consider the effect of gravity, or other body forces such as centripetal acceleration, which may stabilise the fluid surface more effectively at greater speeds. Perhaps the greatest weakness of the “ 0.5 ms^{-1} ” rule is that it does not give any quantification of the severity of the entrainment.

In an approach borrowed from hydraulic engineering [40], the flow may be analysed with respect to dimensionless numbers (e.g. Froude and Weber numbers) [41, 42]; in fact the Weber number is mathematically related to the derivation of the 0.5 ms^{-1} rule [20]. These illustrate the balance of forces in the fluid, and indicate the degree of the instability in the fluid surface. To some extent, these parameters can be evaluated by hand, but it is often more practical to involve Computational Fluid Dynamics (CFD).

2.2.2 Computational Fluid Dynamics Approaches

While the flow predictions of a Computational Fluid Dynamic (CFD) simulation are not necessarily accurate, the general capabilities of a CFD solver can be scientifically verified (or disproved), unlike a mental model. A numerical entrainment model can then be included, to identify and quantify entrainment based on the predicted flow patterns.

One of the simplest of these models is the “air contact time” model, available in commercial off-the-shelf packages such as MAGMAsoft and FLOW-3D. The premise of the

air contact time model is essentially to allow the simulated melt to oxidise, by continually adding to an “oxide concentration” scalar at the surface of the fluid [43]. This “oxide” will be carried through the simulation domain with the predicted fluid motion, using the numerical process of scalar advection[44]; entrainment is inherently predicted when the fluid flow takes the “oxidised” material below the surface.

The main advantage of scalar advection methods is that they are relatively easy to implement, since they follow the same equations as convective heat transfer, or the transport of species concentration.

Although Dørum et. al used the MAGMAsoft air contact time model to correlate entrainment with mechanical properties [3], there are a number of flaws with the air contact time model [44]. Perhaps the most significant is that the scalar accumulates whether or not entrainment is predicted. Additionally, the “oxide” may be taken below the fluid surface by benign fluid flow patterns. For example, an artefact of Eulerian scalar advection called “numerical diffusion” means that the oxide scalar can effectively diffuse into the bulk of the fluid, if the fluid travels in a direction normal to the fluid surface, although this can be mitigated using higher order algorithms. As a result, the oxide concentration scalar is more appropriately used as a tool to help visualise potentially entraining flow patterns, than as a quantitative measure of entrainment.

There is also the issue of proportionality; the model assumes that the severity of damage caused by oxides is directly proportional to the amount of time they have been in contact with the air. Such a simple relation seems unlikely. Furthermore, it goes against the notion that entrainment is reduced by slow, tranquil filling: The model would predict the same damage if one casting is filled steadily, and another is filled in half the time, with enough surface turbulence to double the average surface area.

In a review by Reilly et. al. [45], two scalar advection models are identified with a more believable basis for proportionality. The first is FLOW-3D’s “air entrainment” model, which

is closely related to the air contact time model, but which includes a factor to represent surface turbulence [46]. The intention with both of these models is to make the result more proportional to the volume of entrained air bubbles. However, since this FLOW-3D model is closely related to a standard air contact time model, it is likely to incorrectly identify entrainment in similar circumstances.

The second is MAGMA's "air entrapment model", which places a marker scalar to represent collapsed void regions [45, 47]. For entrainment to be detected, the formation and collapse of a void region must be resolved by the computational mesh; therefore, this method cannot practically predict the formation of bubbles smaller than ~5 mm. It could be argued that this is not a significant disadvantage, because mechanical strength would only be affected by the most significant defects.

One of the key flaws in the scalar advection method is that it cannot properly describe the behaviour of the defects after they have been entrained. Numerical diffusion is one part of this problem, but another is that the "defect concentration" cannot move relative to the fluid, without compromising the simplicity of a scalar advection algorithm. This means that the buoyancy of a large entrainment defect, such as those resolvable by the above "air entrapment" model, cannot be simulated. Ohnaka's review of challenges for casting simulation [7] supports this notion, stating that bubbles large enough to be observed in real time x-ray experiments [48] are strongly affected by buoyancy.

2.2.3 Lagrangian Particle Models

These problems can be avoided by treating the defects as individual particles [20], which follow a motion path dictated by a Lagrangian treatment of the fluid velocity field [49]. Buoyancy and drag forces can also be included in the motion equations.

By avoiding numerical diffusion, this approach could theoretically track an individual simulated entrainment defect from where it is created, to its final location, and vice versa.

This gives an impression of accuracy, in a similar way to listing a large number of decimal places; the true accuracy is limited by the validity of the underlying calculation.

Unlike scalar advection techniques, however, the physical basis of this approach means that the models could be experimentally validated, for instance using the Positron Emission Particle Tracking technique demonstrated by Griffiths et.al. [50].

With the exception of the JSCAST air entrapment model [51], entrainment models involving particle methods tend to be based in an academic environment.

Lin et. al. [52] and Dai et. al. [53] developed a detailed 2D model of the entrainment process, where free surface was described using a line joined by marker particles. Assuming instantaneous oxidation, this line also represents surface oxide. As such entrainment can be directly predicted by resolving the motion of this surface film, and allowing it to break when taken into the bulk fluid. This method follows a simple logic that relies on very few assumptions. However, the marker in cell method which the model relies on does not extend to 3D, which limits its applicability to real-world mould geometry. That said, the model could be used as a benchmark for the predictions of a more practical entrainment algorithm.

A common method of representing the shape of the fluid in 3D is the Volume of Fluid Method, described by Hirt et. al. [54]. The method as described by Hirt (later to be extended for use in FLOW-3D), and a variation described by Zhao et. al. [55], both define the fluid fraction in each cell, along with a description of the surface normal. This level of detail can be interpreted by an entrainment algorithm, to resolve “entraining” flow patterns with less ambiguity than with scalar advection methods, albeit less directly than with Dai’s model.

Such models have been developed by Ohnaka et. al. [56] and Reilly [20], which place Lagrangian particles to mark predicted entrainment events. Reilly’s model is implemented in the FLOW-3D commercial CFD package, and uses its internal VOF description of the fluid surface, along with the fluid velocity field, to identify entraining flows. The evidence presented in a number of works at the University of Birmingham [20, 21, 57] would suggest

that this model is able to reliably identify whether one mould geometry is more entraining than another closely related design.

However, it has not been demonstrated that the model is able to identify the areas in the casting which are most likely to be weakened by entrainment defects. Without this, it is difficult to say whether the model correctly distinguished the relative severity of entrainment events predicted at different times and places in the mould.

A critical assessment of the algorithm itself reveals a number of flaws. Reilly highlighted some circumstances in which the code can falsely identify entrainment, such as droplets or bubbles which move diagonally through the mesh [20]. This would be a significant problem in the highly turbulent flows of high pressure die casting.

The greatest weakness of Reilly's model is its Boolean nature [58, 59]; a particular flow pattern is either locally entraining, or it isn't. This means that the measure of damage – the number of particles placed – is dependent on the length of time that an entraining flow pattern persists, rather than a more physically meaningful quantity.

Moreover, the internal mechanism that maintains the approximate proportionality with time would be overridden if the point of entrainment moves to a different mesh cell. This was not a significant problem for Reilly's test cases, but could be a serious issue for more complex flow patterns, as it could lead to instances of runaway particle placement that are not linked to particularly severe entrainment events.

3 Theoretical Developments

3.1 Surface Area Entrainment Code

3.1.1 Introduction

In the Literature Review (Section 2.2), a number of methods of modelling entrainment were surveyed and criticised. Each of these was intended to indicate where in a casting a foundry might encounter problems with entrainment defects, and also indicate whether any changes to the design of the casting system might improve the casting. These algorithms are what might be called semi-quantitative; the information they predict is quantitative, and related to entrainment, but also to other factors, that might be impossible to separate. As such, their output is mainly useful to inform a foundry on qualitative decisions.

If the ambition is to take a new casting, and predict the mechanical properties at any given point, then a semi-quantitative algorithm is not sufficient, as there would be not be a prior design iteration – or some other known, closely related casting – to compare the results against.

Instead, the algorithm's output should be as quantitative as possible; ideally its output could be directly linked to mechanical properties, (possibly via a function that includes predicted parameters related to thermal history, microstructure, etc.), and not to factors that are unpredictable, or otherwise difficult to isolate.

This section details the development of the Surface Area Entrainment Code (SAEC), which was loosely based on Boolean Oxide Entrainment Code (BOEC) described previously [20]. This new algorithm was intended as another step towards such a fully quantitative Oxide Film Entrainment Model (OFEM). Firstly, it was intended to resolve whether or not a flow is entraining by calculation, rather than a pre-defined set of rules that may produce anomalous results for unexpected flow cases. Primarily though, the SAEC algorithm was intended to quantify the severity of entrainment, using the area of entrained surface film as a metric.

3.1.2 Overview of the Formulation

Like the Boolean Oxide Entrainment Code (BOEC), the present entrainment algorithm is implemented as a user subroutine in FLOW-3D, and places particles into the domain when entrainment is identified. These particles then move with the simulated fluid flow; each time-step their predicted motion over that time is tracked by FLOW-3D's Lagrangian Particle Tracking (LPT) algorithm.

Unlike the BOEC, it was intended that the Surface Area Entrainment Code (SAEC) would predict the area of surface film that was entrained, and place a number of particles proportionally. At a high level, the algorithm can be broken down into three stages:

1. Quantitatively assess the motion of the film assumed to be on the surface of the fluid, for each cell.
2. Assess the stability of the fluid surface, to determine whether it is possible for entrainment to occur in each cell.
3. For each cell identify whether surface film has been "lost" and also if entrainment could occur in that cell, place an integer number of particles dependent on that loss.

3.1.3 Quantitative Surface Tracking

The process of determining whether surface film has been created or lost in a cell is based on a continuity equation for area, conceptually related to conservation of mass or momentum. The details of the formulation will be expanded upon in due course. The approach relies on the following assumptions:

1. The surface of the fluid is coated in film, such that the area of that film is the same as the surface area of fluid in that cell.
2. The fluid surface area in each cell, as calculated by FLOW-3D, is sufficiently reliable to be used as the basis of these computations.
3. The surface film moves with the same velocity as the underlying fluid.

Under these assumptions, it is possible to calculate the motion of surface film into and out of each cell in the domain. For each cell, the “estimated surface area” can be extrapolated from one time-step to the next using the area continuity relation, which assumes that no surface area has been lost or created. At a later stage, the estimates for each cell will be compared to the calculated value for surface area. If the estimated value exceeds the calculated value, then surface film has been locally lost; conversely, if more surface area is found in a cell than expected, then surface film has been created there during the time-step.

In the work of Lin et. al. [52] and Dai et. al. [53], the surface of their 2-D fluid was represented directly by a 1-D line. In the present work the fluid surface is a 2-D manifold in 3-D space, but the surface is not directly modelled within the fluid simulation. Instead, it is implied by the cell fluid fraction number field, as defined by a Volume Of Fluid algorithm in FLOW-3D. That said, FLOW-3D’s VOF algorithms do define a relatively sharp fluid-void transition. The line indicating the fluid surface on the below figure was added by the post-processor, and is not part of the simulation.

However the liquid is discretised for computational solution, in topological terms, the shape of the melt at any instant can be considered as a 3D object, which has an infinitesimally thin 2D surface. While the surface may stretch and deform over time, points on the surface – perhaps representing an oxide film – may only move along the surface.




Figure 3-1: The result of the cell pairing stage applied to a bubble.

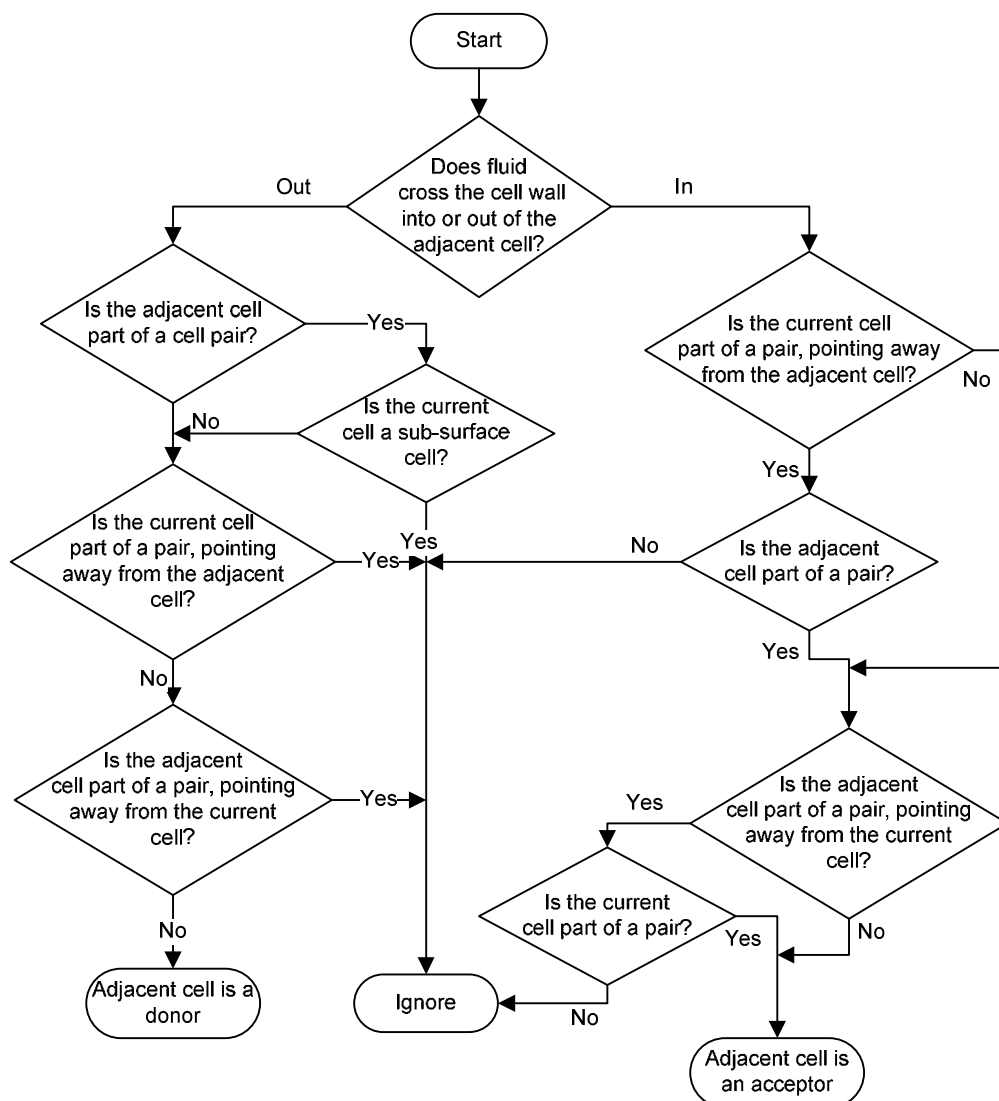
The area continuity relation was modified with logical conditions that constrain the flow of area to follow this behaviour, as opposed to the transport equations that generally apply to the VOF method in 3D. Central to these logical conditions is a subroutine which assigns a surface flag value to each cell. An example result of this flagging process is illustrated in Figure 3-1.

Cells with a fluid fraction >0.5 , which have a neighbouring cell that has a fluid fraction <0.5 , or vice versa; are coloured red, yellow, green and purple in the figure. Each of the cells flagged in this way “points” (+x, -x, +y, -y, +z and -z) at the adjacent cell which most sharply encloses the fluid surface; the pointing direction of the relevant cells in Figure 3-1 is indicated by their colour. Cells which point to one another are considered “paired”- these paired cells place logical bounds on the free surface. The network of cell pairs can be thought of as a virtual mesh, which reduces the complexity of the problem. Generally, cell pairs are consistent, and each cell in the pair points at the other, however this is not necessarily the case at corners (such as those above and below the purple cells in Figure 3-1).

A cell may therefore point to a cell which does not point back; such a situation is not ideal, but the negative effect of corners is believed to have been minimised by trial-and-error selection of the logical constraints. Some of these constraints are based on the fluid side neighbours of any paired cell, which will be referred to as “sub-surface cells”, and which are coloured dark blue in Figure 3-1. The light blue cells in this figure are “neutral” cells, which are not included by any other definition.

In some senses, the two members of a cell pair can be considered as a single virtual cell, which is theoretically guaranteed to enclose the fluid surface when the surface flags were assigned, the previous time-step and the following time-step. This is because FLOW-3D limits the time-step such that the fluid surface cannot travel more than half a cell in the direction normal to itself within a time-step, and the cell pairs are defined such that the fluid surface is between the centres of the two cells. From the perspective of each cell pair, the

motion of the surface film can therefore be reduced to a 2-D problem, with a domain roughly parallel to the fluid surface.



The intended result of these constraints on which cells are allowed to be a donor or an acceptor, is to largely restrict the motion of surface area, to emulate the behaviour of motion on a 2-D manifold. If these constraints were not imposed, numerical diffusion would spread the “surface” into the fluid, and erroneous entrainment would result.

In the present model, it is not necessarily the case that if cell A is a donor for cell B, then cell B is an acceptor for cell A. This is an example of the simplifications which make the result of the continuity relation an estimate, rather than a more exact solution, which would be more difficult to implement, and may significantly increase the run-time of the simulation.

Generally, the approximations and simplifications made were intended to over-estimate rather than under-estimate the amount of entrainment; a systematic false negative was considered to be more deleterious to simulation results than an equivalent false positive.

The approximate continuity relation is evaluated one time-step behind; at each time-step, information from the previous time-step is used to try to estimate the current state. This means that fewer assumptions need to be made. The set of equations used to evaluate the estimated change of surface area within a cell are given below, where \forall means “for all”:

$$A_D = \sum_{\forall \text{Donors}} \frac{A_i v_i}{x_i} \Delta t \quad \text{Equation 3-1}$$

$$K_A = \sum_{\forall \text{Acceptors}} \frac{v_i}{x_i} \Delta t \quad \text{Equation 3-2}$$

$$A_E = \begin{cases} A_D + A(1 - K_A), & \text{Constant} \\ \frac{A_D - (A_D - K_A A)e^{-K_A}}{K_A}, & \text{Depleting} \end{cases} \quad \text{Equation 3-3}$$

Equation 3-2 estimates the total area of surface film, A_D , carried into the current cell from donor cells, based on the area of surface film in each donor cell, A_i . This equation also

involves the component of fluid velocity from the donor cell to the current cell, v_i , the distance between the centres of the two cells, x_i , and the time-step, Δt .

In Equation 3-2, K_A is an advection constant, describing the flow of surface film out of the current cell; the other quantities in this equation have a similar meaning to above, but apply to acceptor cells. All of the quantities on the right hand side of these equations relate to the previous time step.

Equation 3-3 defines two different approximations used to estimate the surface area in the current cell, A_E , at the current time-step of the simulation. One function (“Constant”) assumes the area of surface film, which may be carried out to acceptors does not change over time; this form is used where K_A is small or where the area in the current cell increases over time. The other function (“Depleting”) is an integration, which assumes that the area of surface film available to be carried out of the current cell does change over time, but that A_D and K_A are constant; this approximation is only used if the time-step is not small, and more surface film is carried out than in. Choosing which of these functions to use based on a predicted increase or decrease of film area in the current cell, is equivalent to choosing the highest result from the two functions; this is done to try to ensure the estimated level of entrainment is an over-estimate.

3.1.4 Surface Stability Assessment

The second stage of the algorithm is to determine whether a modelled “loss” of fluid surface area is likely to equate to entrainment. Figure 3-3, from a previous published work by the author that summarised the algorithm [58], illustrates this concept. Both frames A and B illustrate a stream of fluid, which falls into a larger pool. Frame A shows an entraining flow, where the momentum of the fluid carries it beneath the surface of the pool, likely taking any surface film and mould gas with it.




Figure 3-3: An illustration of why loss and entrainment are not necessarily equivalent, from Watson et. al. [58].

Frame B of Figure 3-3 shows a different case, in which the fluid stream diverges when it hits the pool of fluid, perhaps because of an underlying obstacle; the key point is that this flow pattern is possible. In this case, the fluid at the surface also decelerates as it diverges. In reality, the surface film on such a flow would bunch up and form wrinkles, as illustrated.

Since the oxide film is a chemically stable solid, it is assumed that when the area the film needs to cover decreases, the film must wrinkle to accommodate that change, rather than thicken, or somehow dissipate; as such, surface film would be accumulated on the surface, and not entrained, or otherwise lost.

In a FLOW-3D simulation, these wrinkles could not form, because the surface film is not modelled, (except indirectly by the present algorithm). Since the entrainment estimate of the SAEC is based on the surface area calculated in each cell by FLOW-3D, the algorithm would indicate a “loss” of surface area in this flow, because more surface film would be carried into some cells than out, and the modelled surface area would remain constant.

Some entrainment criterion is needed to distinguish between the two types of flow shown in Figure 3-3: one where entrainment would occur in reality; and one where a model might indicate entrainment, but only because the underlying simulation does not resolve the fluid surface in sufficient detail.

In the SAEC, this is done by assessing the stability of the fluid surface. It was theorised that for any given fluid flow situation, the fluid surface may be stable (like a regular pendulum) or unstable (like an upside-down pendulum). If the surface is locally stable, then a small deviation in the fluid surface would shrink, and become a wrinkle. If the surface is locally unstable, then a small deviation is assumed to grow into an entrainment event.

The surface stability is assessed in two ways. Firstly, the macroscopic stability is assessed by checking the fluid surface for externally sharp corners, such as those illustrated in Figure 3-3 A. It was assumed that sharp corners such as these indicate the fluid surface is unstable at a scale that can be directly resolved by simulation, because otherwise the fluid flow simulation would smooth these corners under the action of gravity or surface tension.




Figure 3-4: Illustration of theoretical effect of body force direction on a fluid surface.

If a sharp corner is not detected in a surface cell, then the cell is checked for surface instability on a scale smaller than that explicitly resolved in the simulation, using an analytical stability relation. First, the direction normal to the fluid surface is determined, based on the “volume fraction fluid” number field. Then, the value of the pressure gradient in that direction is calculated, immediately below the surface. If the pressure increases towards the surface, then a small deviation in the fluid surface would grow into an entrainment event.

This can be understood by equating pressure gradient to a body force, such as gravity, as illustrated in Figure 3-4. Fluid mechanics dictates that the local pressure gradient and the body forces acting at a point are directly linked. If the pressure in both flows illustrated in this figure is assumed to vary linearly in the vertical direction (as a first approximation), then the highest points on the surface of both streams would have pressure below atmospheric, and the pressure at the lowest points at the surface of both streams would be above atmospheric. This cannot be the case, and so to resolve the force balance, the surface of the lower stream must be accelerating towards the dashed line, and the surface of the upper stream must be accelerating away from the dashed line. The dashed lines also represent the fluid surface, as it might be resolved using a mesh larger than the surface features.

Equivalently, the regions of mould atmosphere above the upper dashed line could be seen as buoyant proto-bubbles, which would rise into the upper stream, displacing fluid.

On the scale of several cells, surface tension is included in the stability assessment, because FLOW-3D includes surface tension in the pressure field calculation. However, at the smaller scale, surface tension is not included when determining whether surfaces such as those in Figure 3-4 are stable. This may seem counter-intuitive, but the small scale stability criterion is only relevant under potentially entraining conditions. As such, it is possible to use a logical “short-circuit” and formulate the calculation *given that* there is local surface “loss”, since the calculation will only be used if that is the case.

The purpose of the analysis is to resolve situations like the one illustrated in Figure 3-5. Here, two fluid streams meet, and accelerate downwards, as indicated by the arrows. The downward acceleration means that the fluid pressure decreases away from the surface, satisfying the “surface instability” criteria. During the initial formation of such an entrainment case, wrinkles of solid and chemically stable oxide film would start to pile up between the two flows. Surface tension is assumed not to be able to contract the surface film, and so surface tension could not counteract the growth of the entrainment event.




Figure 3-5: Illustration of entrainment in a case that may be resolved using the pressure gradient criterion.

Again, the dashed line represents the fluid surface as it might be resolved by the fluid flow solver. The entrainment algorithm is only one-way coupled to the fluid solver, and as such the solver would not account for the surface films and trapped gas separating the fluid streams.

3.1.5 Particle placement

The final stage of the algorithm estimates the area of surface film entrained in each cell, based on the calculations of the previous two stages, and the current state of the fluid surface. The particle placement subroutine analyses all cells, which were assigned to a cell pair, or contain more than 50% fluid. The area of surface film “lost” in a cell is evaluated as the current surface area in that cell, minus the “estimated surface area” in that cell.

Cells which are members of a pair are effectively evaluated together. This means that the fluid – and the associated surface film – can move between the members of a cell pair without generating anomalous results. This is implemented such that one member of the pair, the “primary cell”, performs the evaluation for both cells; the non-primary cell is skipped over. Whether or not a member of the cell pair network is a primary cell depends on its fill fraction; the primary cell will always be the fullest member of a pair.

When a primary cell is evaluated, the area of “lost” surface film calculated for that cell, then surface area “loss” for the cell which it point to is evaluated, and added on. This approach is tolerant of the case, which can arise at corners, where cells do not form consistent pairs. If the surface stability of the primary cell is assessed as “unstable”, then the subroutine continues, and the “lost” surface area is considered entrained. Any surface area which is lost, but not entrained is not carried forward to the next time-step. Given that the “estimated surface area” is generally overestimated, this would result in an accumulation of error.

It is also possible for a cell that is not in the surface layer of cell pairs to have a non-zero estimated surface area. This case generally arises when two surfaces collide and merge in the previous time-step, and so the cell was below the surface when the cell pairs were assigned. In this case, the cells are evaluated individually, and the surface stability criterion is skipped, since the film is essentially already entrained.

If an area of surface film was “entrained” in a cell, then that area is converted to a number of particles, using a proportionality coefficient. The entrained surface area is divided

by this proportionality coefficient, to calculate a real-valued (not integer) “particle number”. This “particle number” is randomly rounded up or down, using its fractional part as a bias. For example, if 1.3 mm^2 , of surface film is entrained in a cell, and the proportionality coefficient is $1 \text{ particle} \cdot \text{mm}^2$, then there is a 70% chance that 1 particle would be placed, and a 30% chance that 2 particles would be placed.

Particles are placed randomly within the fluid of the cell, as constrained by the cell’s fluid volume fraction, and orientation, and assigned a velocity equal to the velocity of the fluid at that point.

3.2 Statistical Mapping Theory

3.2.1 Introduction

Previous work on the reliability of castings [6, 21] found that it is not correct to describe the strength of cast material using a definite value. Instead, whatever the measure of strength, it must be described in terms of a statistical distribution, including a measure of the variation.

In some prior works [1, 6, 20], authors have hypothesised that an increase in the number of predicted entrainment defects in a region, should correlate with an increase in the variability in strength in that region, characterised by the Weibull modulus.

This may seem logical at first. If a test-bar has an entrainment defect on its fracture surface and a lower than average strength, then it follows that the defect caused that reduction in strength, and therefore increased the variation in strength for the data set. However, it does not then follow that an increase in the number of predicted entrainment defects would produce an increased variability in mechanical strength.

Let it be temporarily assumed that a perfect simulation were run, which accurately described the distribution of entrainment defects in each and every casting of a particular design. Each of the entrainment defects in those castings would have a counterpart in a very similar location in every other casting of that type. Given that mechanical strength variation was shown to exist in castings, the stress tolerance of each of these equivalent defects must vary from casting to casting. To avoid contradiction, this variation must be caused by something that would not also affect the location of these defects.

Fundamentally, casting is a fluid flow process, where that fluid has a low kinematic viscosity, meaning that turbulent eddies and vortices form easily. As a result, it is almost impossible to have a real casting process that is completely repeatable. For each real casting, small variations in the process conditions and the mould geometry will, to some extent, affect where and when entrainment defects form. Similarly, the precise pattern of eddies in the down-stream fluid flow will vary from casting to casting, and so even two entrainment defects

formed in an equivalent time and place are likely to be carried to different locations in the mould.

In contrast, whatever the accuracy of a simulation, it represents a single case, based on nominal or estimated process conditions. In theory, if a casting simulation involving some entrainment algorithm were run twice, then both runs should predict an identical or very similar distribution of defects.

One approach might be to run multiple simulations each with different initial or boundary conditions, to reflect the variability in the real casting process. However, this could be very computationally expensive, particularly given that each of these simulations must be sufficiently accurate to be meaningful.

The approach taken by Yue [21] was to consider the geometric distribution of entrainment defects within a simulated casting as a statistical distribution. Each simulated casting was split into divisions; this allowed a Weibull distribution to be fitted using the inverse of the defect count within each division as the random variable. For each mould geometry used, this could be directly compared the Weibull distribution of mechanical strengths; the test pieces were extracted from locations equivalent to the divisions of the simulation.

This comparison is logically consistent with the idea that no two castings are identical, since a correlation does not require the defects in each casting to be in equivalent locations to the defects predicted in simulation. That is to say, that the casting simulation may predict an elevated concentration of defects towards the centre of the casting, formed during a single large entrainment event. When the casting is performed experimentally, this large entrainment event would happen at slightly different times and locations, and slight variations in fluid flow pattern would carry the resulting region of concentrated defects to different locations in the mould. The location sampling statistical approach taken by Yue, would identify that some

samples would be significantly worse than others, both in experiment, and in simulation; the key feature is the elevated concentration of defects at some point in the casting.

However, assuming the entrainment model used was sufficiently reliable, the result of Yue's approach would be a predicted statistical distribution of mechanical strength for a whole casting. For larger and more complex castings, however, it may not be realistic to assume that each volume of fluid – carrying with it some level of entrainment damage – may end up at any location in the casting. Furthermore, from a design perspective, it would be desirable to know how the strength of the casting varies at different locations.

What is needed is a method of evaluating how the defect distribution might vary at each location in a casting, based on a single casting simulation, by characterising permutations of that result.

3.2.2 Overview of the Implementation

The implementation of the present location sampling statistical algorithm falls into three parts. Firstly, the measure of entrainment damage must be defined throughout the casting simulation result. The second stage is to evaluate the permutations of the damage value for any given point in the simulation. Lastly, the permutations at each point are characterised statistically. The algorithm was implemented as post-processor for FLOW-3D output files, which can read mesh, fluid and particle data using a library provided as part of the FLOW-3D installation.

3.2.3 Definition of Entrainment Damage

The measure of entrainment damage chosen was the smoothed particle concentration, with a smoothing radius of r_{smooth} . This was evaluated at the centre of each fluid containing mesh cell by dividing the smoothed particle count by the smoothed fluid volume. A smoothing function was used so as to reduce the noise introduced by random small scale

accumulations of particles, and to eliminate the step function associated with a particle being either in or out of a cell.

The smoothed particle count was calculated by counting each particle within $3r_{smooth}$ of the cell centre, with a bias based on Equation 3-4. For example, if a particle is $2r_{smooth}$ away from the centre of the cell where the smoothed particle concentration is being evaluated, that particle would be counted as 1.8% of a particle.

$$S(r) = e^{-\left(r/r_{smooth}\right)^2} \quad \text{Equation 3-4}$$

Similarly, the smoothed fluid volume is calculated by summing the volume of fluid in cells within $3r_{smooth}$, weighted according to Equation 3-4. For example, if the smoothed fluid volume is evaluated at cell A, cell B contains 1 mm^3 of fluid, and the centre of cell B is r_{smooth} away from the centre of cell A, then cell B's contribution to that smoothed fluid volume would be 0.37 mm^3 .

3.2.4 Evaluation of Permutations

Once the smoothed particle number density has been computed for each fluid containing mesh cell, the program reads in a list of 3-D Cartesian coordinates, which define where the permutations are to be evaluated. In this work, these coordinates were the centroids of elements from a Finite Element mesh, which had been geometrically mapped into the space of the casting simulation.

The model evaluates permutations using a random walk algorithm, which was adapted to work within the irregular Cartesian mesh used by the FLOW-3D simulation. The random walk starts at a cell near the given coordinates, and continually jumps from one cell to the next in a randomly selected direction (+x, -x, +y, -y, +z, -z), until the termination criterion is reached.

In principle, each time the random walk is run, the cell where the random walk ends contains a volume of fluid that could instead have ended up at the starting coordinates, if the simulation parameters had been slightly different. The random walk is evaluated a set number of times for each set of starting coordinates, to build up a statistical data-set.

One of the adaptations used to make the algorithm appropriate for an irregular mesh, was to use a total walk distance, r_{walk} (defined in metres), rather than a prescribed number of steps. Initially, a “distance to go” variable is defined as the square of r_{walk} . On each jump from one cell to another, the square of the cell-centre to cell-centre distance is subtracted from the “distance to go”, and the walk ends when this value falls below zero.

Additionally, the probability of jumping in each of the 3 ordinate directions was weighted by 2 to the power of the size of the cell in each of that direction. The probability of jumping in any given direction was further weighted by the inverse square of the separation of the cell centres in that direction. These adaptations were found by derivation, and verified by first implementing them in MATLAB, where the mesh structure could be arbitrarily varied, and the random walk result observed graphically.

It was also necessary to adapt the algorithm to ensure that the random walk would stay within fluid-containing cells. This was done by preventing a jump to a cell with less than 50% fluid volume, provided that cell has also the lowest fluid volume of any neighbour. It is therefore still possible for a jump to be made to a cell with less than 50% fluid volume, but if that is the case then the random walk must make its next jump to the fullest adjacent cell. This approach allows the random walk to function more realistically in cases where the casting is oriented diagonally to the mesh. Simply preventing a jump to any cell with a fluid volume fraction less than 50% was found to discourage the random walk from proceeding in directions where the surface of the casting was not aligned to the mesh planes.

3.2.5 Statistical Characterisation

For each set of starting coordinates, the random walk algorithm generates a set of geometric permutations, and a set of particle concentration values sampled from those locations. A statistical distribution is then fitted to those particle concentration values, and linked to the relevant set of starting coordinates. This means that the probable “damage” at each starting coordinate is defined mathematically, and can be manipulated using mathematical tools.

The 3-parameter Fréchet distribution, presented below, was used to characterise the particle concentration values. Equation 3-5 is a cumulative probability function for modelling the “damage”, or smoothed particle concentration, D , in a random sample of the casting.

$$P(D < d) = \begin{cases} 0 & , d < \mu \\ e^{-(d-\mu/d_0)^\beta} & , d > \mu \end{cases} \quad \text{Equation 3-5}$$

This distribution includes a threshold parameter, μ , which places a lower limit on the possible values of “damage” which may be sampled. The position parameter, d_0 , defines an “average” for the distribution, together with the threshold parameter; the 37th percentile of the distribution is defined as $\mu+d_0$. The variance of the distribution is a monotonically increasing function of $-d_0/\beta$.

Here, β is a shape parameter, closely related to the modulus parameter, m , in the Weibull distribution (Equation 2-1, Literature Review), however the value of β must be negative, whereas the value of m must be positive. This essentially reverses the direction of the distribution. The cumulative Weibull distribution begins with a “1 -” term, where the cumulative Fréchet distribution does not.

As it would seem from their respective cumulative probability functions, the Weibull and Fréchet distributions are very closely related. In fact, they are both “opposite” types of extreme value distribution, along with a third, the Gumbel distribution [15]. Weibull statistics

are widely used to describe failure, because they are suited to describing the low extremes, i.e. the weakest samples. Similarly, it is proposed that the Fréchet statistics are appropriate to describe defect damage, because they are suited to describing the high extremes, i.e. the most damaging defects; in fact, D in Equation 3-5 describes the most damaging of all the defects in a random sample.

Suppose that some stress is applied to a sample, and there is a relation which defines the critical level of “damage”, d_{crit} , that would cause the sample to fail at that stress. To use Equation 3-5 to model failure, d_{crit} may be used in place of d . For the sample to survive at a given stress, the most damaging defect must have a less than critical “damage” value. Therefore, for any arbitrary division of the sample into two parts, the most damaging defect in the first part, *and* the most damaging defect in the second part, must *both* be less damaging than the critical damage. In probability theory, the logical “and” is equivalent to multiplication, i.e. the probability that the result of two coin tosses are both tails is $1/4$.

$$e^{-(d_{crit}-\mu/d_0)^\beta} = e^{-q(d_{crit}-\mu/d_0)^\beta} \times e^{-(1-q)(d_{crit}-\mu/d_0)^\beta} \quad \text{Equation 3-6}$$

Equation 3-6 shows how splitting the sample is mathematically consistent, where q is the fraction of the sample attributed to the first part. For any q , both the left and right hand sides give the same probability that the sample would survive. This shows that for any splitting or recombination of parts, the Fréchet distribution can consistently describe the most damaging defect.

This analysis may be extended, such that an arbitrary number of elements, each with “damage” individually described by a Fréchet distribution, can be grouped together to form a virtual test piece, and the overall damage function for the test piece may be easily defined. This is the most important reason why the Fréchet distribution was used, because it allows the damage statistics to be correlated to test piece data, as described in the Method (section 4.3.2).

As stated in Section 3.2.4, element centroids were used to determine the random walk starting coordinates, and the random walk algorithm samples a set of damage values, for each of those starting points. The process of converting each set of damage values to Fréchet parameters for each element is as follows:

The cumulative probability function (Equation 3-5) was rearranged in terms of d , as shown in Equation 3-7, using x as a substitution that is defined in Equation 3-8. The first step in finding the parameters for this function is to sort the set of damage values in ascending order, so that each sample can be assigned a probability estimator P_i using Equation 3-9. It is then possible to re-define Equation 3-7 in terms of individual samples, by using P_i in place of the cumulative probability in Equation 3-8, resulting in Equation 3-10.

$$d = d_0 e^{(x/\beta)} + \mu \quad \text{Equation 3-7}$$

$$x = \ln[-\ln\{P(D < d)\}] \quad \text{Equation 3-8}$$

$$P(D < d) \approx P_i = (i + 0.5)/n \quad \text{Equation 3-9}$$

$$d_i = d_0 e^{(x_i/\beta)} + \mu + \text{err} \quad \text{Equation 3-10}$$

An iterative numerical procedure is then used to fit the parameters for the cumulative distribution function, based on a least squared error approach. Analytical functions for d_0 and μ were derived in β given that $\Sigma \text{err} = 0$, which allowed the overall solution to be found by solving $d(\Sigma \text{err}^2)/d\beta = 0$ using a bisection algorithm. This atypical solution method was primarily used because it is tolerant of zero-valued input data, but also because it is relatively fast and because it is believed to give greater weight to the upper section of the distribution, which is more important for subsequent analysis.

4 Method

4.1 Experimental

4.1.1 Characterisation of Bearing Mount Commercial Casting

Casting trials were performed at JVM Castings Ltd., to gather data on the defects that affect High Pressure Die Castings (HPDC) in a commercial setting, and to provide other qualitative information about the HPDC process. The casting selected, shown in Figure 4-1, was a mounting bracket for a prop-shaft bearing on a Jaguar XK sports car.

This part was chosen because it has large flat surfaces which facilitate the extraction of flat test bars, and also because it was one of the smallest castings that test bars could have been extracted from (~0.21 m in the x-direction as shown in Figure 4-1). The smaller size means that when the part is simulated, a practical run-time can be achieved whilst also resolving the flow in sufficient detail.

The alloy used for the castings was a recycled aluminium alloy to “Castasil 37” specification, as detailed in Table 4-1. The bulk alloy was melted in a gas powered furnace, then degassed using a rotary flux degasser immediately before being transferred to an electric holding furnace.

Table 4-2 details the key parameters and settings for the Die Casting Machine (DCM), which were fixed for each of the casting trials; these casting and machine parameters were based on the foundry’s standard practice, experience and capability.

Table 4-1: Specification for Castasil 37 alloy

Element	Cu	Mg	Si	Fe	Mn	Zn	Ti	Sr	Zr	Mo	Al
Minimum	-	-	8.5	-	0.36	-	-	0.010	-	-	Bal.
Maximum	0.05	0.06	10.5	0.15	0.60	0.07	0.15	0.025	0.3	0.3	Bal.




Figure 4-1: The casting selected for analysis, showing large flat areas of constant section thickness, appropriate for tensile testing.

Table 4-2: Process parameters which were constant for each trial

Parameter	Value
Shot Sleeve length	0.735 m
Shot Sleeve Diameter	0.09 m
First Phase Piston Speed	0.3 ms^{-1}
Shot Speed Transition Point	190 mm from die end
Casting Temperature	710°C
Intensification Pressure	32 MPa
Total Cast Part Volume, Both Sides (Nominal Mass)	$5.72 \times 10^{-4} \text{ m}^3$ (1.39 Kg)
Total Runner System Volume, Both Sides (Nominal Mass)	$3.9 \times 10^{-4} \text{ m}^3$ (0.95 Kg)
Total Overflow Volume, Both Sides (Nominal Mass)	$2.10 \times 10^{-4} \text{ m}^3$ (0.51 Kg)

Table 4-3: Process parameters varied with each trial

Trial Name	Third Phase Piston Speed	Gate Area (Blanking State)	Charge Mass
Short Shot	1.15 ms ⁻¹ Nominal	467 mm ² (Full)	0.4-2.1 kg in 0.1 kg steps
FGFS	3.35 ms ⁻¹	467 mm ² (Full)	2.9 Kg
BGFS	3.35 ms ⁻¹	280 mm ² (Blanked)	2.9 Kg
FGSS	1.12 ms ⁻¹	467 mm ² (Full)	2.9 Kg
BGSS	1.12 ms ⁻¹	280 mm ² (Blanked)	2.9 Kg

Two groups of casting trials were performed. The first were what are known as “Short Shot” trials, which have been used in prior works to help assess the pattern of fluid flow in the mould. In these tests, a less than nominal amount of liquid metal is dosed into the shot sleeve, as specified in Table 4-3. In these castings, the shape is not only determined by the shape of the mould, but also shows the pattern of fluid flow at the time the fluid solidified.

The other four trials listed in Table 4-3 were used to characterise the response of the castings’ tensile properties to changes in fluid flow parameters. For half of these trials, the full production standard shot speed was used. For the other half, a slower shot speed was used to push the metal into the mould; the value for this was based on the minimum speed they believed would make a “sound” casting. Previous works involving sand castings indicated that as filling speed increases, the level of entrainment related damage in a casting also increases.

Figure 4-2 illustrates the nominal piston velocity profile for one of the full speed shots, and also aims to clarify some of the terminology used in Tables 4-2 and 4-3.

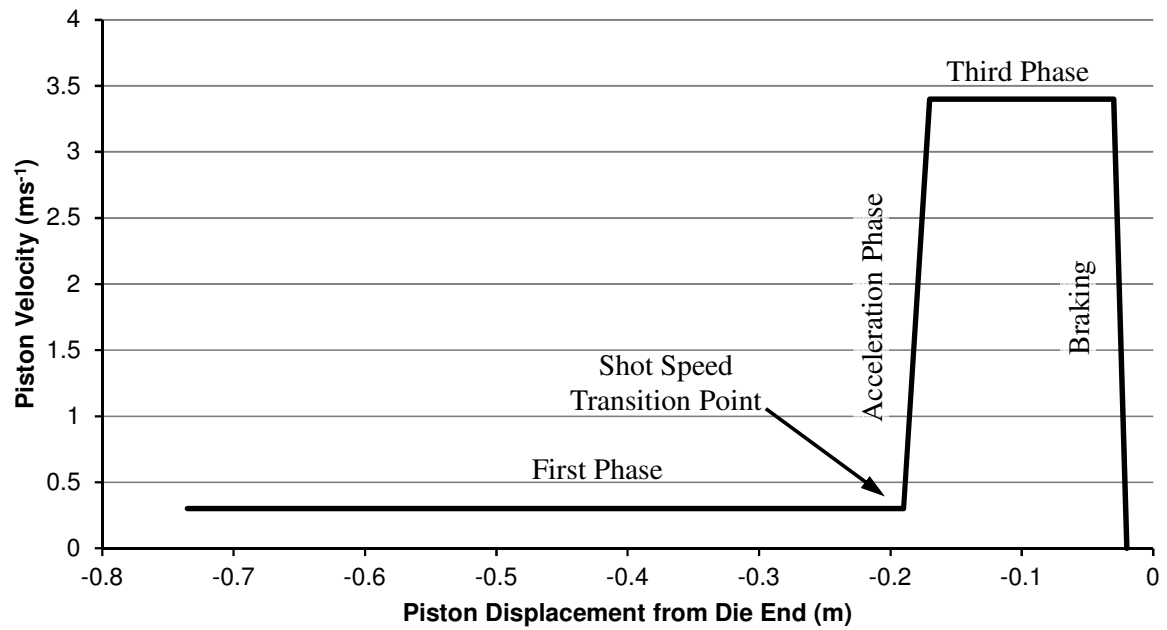


Figure 4-2: Typical trace of nominal HPDC piston speed.




Figure 4-3: Illustration of the inserts used in the BGFS and BGSS trials.

The other parameter that was varied in these trials was the state of the gate; for half of the trials, a blanking piece was inserted to block fluid flow through half of the gate, as shown in Figure 4-3. Preliminary simulations indicated that this blanked gate configuration would

strongly modify the pattern of fluid flow, and so significantly affect the final distribution of entrainment defects within the casting.

30 castings were made in each of the full shot trials. In a single shot, a casting was made in the left and right hand side of the mould, as pictured in Figure 4-3 A and B respectively. None of the castings in these trials were heat-treated in any way, although the foundry would normally send these parts for a stress relief heat treatment. Firstly, this was done to avoid increasing the ductility of the material, because less ductile material would be sensitive to crack initiation by smaller entrainment defects. Secondly, this removed a variable from the analysis, since the heat treatment may not have been completely consistent.




Figure 4-4: Test bar sampling locations for left hand castings (A and C), and for right hand castings (B and D).

Three test bars were punched from each of the castings, in the locations shown in Figure 4-4, one from near the gate, and two from the large flat region farthest from the gate. The test piece used was a JVM standard flat test piece, with a gauge length of 38.6 mm, and a nominal cross section area in the gauge length of 25.2 mm². A Zwick/Roell Proline was used to test these bars in uniaxial tension with a cross-head speed of 4 mm·min⁻¹, with a clip extensometer for measuring total strain.

The failed test-bars were sorted and classified according to their trial code, their shot number (1-15), the handedness of the casting (L or R), and their location (A, B or C). Each test bar's fracture surface was assessed visually, and notes were made on the nature of the fracture, and of any visible defects. This allowed the test bars to be additionally classified by the feature which was most likely to have initiated the fracture.

For a selection of the samples, the whole fracture surface was imaged optically at low magnification. This was principally to serve as a reference for future research.

Additionally, another subset of the samples was imaged using a Scanning Electron Microscope (SEM). Two SEMs were used in the investigation: the first was a Joel 6060, with an Oxford Inca EDS (Energy Dispersive Spectroscopy) system; the second was a Zeiss EVO 50, with a PGT Spirit EDS system.

4.1.2 Whole Casting Structural Impact Test

The aim of the characterisation work in the previous section was to help build a predictive model of the distribution of strength within a commercial High Pressure Die Cast (HPDC) part. The ambition is that this model would be able to predict the failure characteristics of a whole part, subject to arbitrary loading, rather than the uniform uniaxial tension applied in a tensile test. The part chosen for test was the same part used for the tensile characterisation exercise, because it was supposed that using another part would introduce too

much uncertainty, and reduce confidence in subsequent analysis. The castings for this test were made as per the FG-FS casting trials in the previous section.

As such, it was necessary to perform whole-part structural testing to assess the predictive capability of such a model. The test chosen was a drop test, in which the casting was secured to a pseudo-rigid structure, and an impactor is dropped from a prescribed height onto a prescribed location on the casting. The impactor is connected to a weighted trolley, which slides vertically on guide-rails.




Figure 4-5: Two views of the configuration of the drop test rig.

A dynamic test was chosen, because this more closely represents the type of load-case that the predictive model would be required to assess if it were applied in industry. As shown in Figure 4-5, the casting was clamped at an angle such that the large flat region of the casting is horizontal; in preliminary tests this was found to increase the energy transfer during impact, as the impactor would otherwise glance away from the casting.




Figure 4-6: Impact location and dimensions of impactor.

The mass of the trolley assembly used in the presented tests was 31.3 kg. In the first of these tests, the impactor was 0.5 m above the flat region of the casting when the trolley assembly was dropped. In the second and third tests, this distance was increased to 1 m.

The trolley was instrumented with accelerometers and a LASER displacement sensor, so that the changes in the momentum of the trolley assembly during the test could be calculated, and the forces within the system inferred.

Additionally, each test was filmed with a high speed camera, which was positioned side on (as in Figure 4-5 B) and recorded at $1000 \text{ frames} \cdot \text{s}^{-1}$. The position of the round crash-test style marker, which was affixed to the impactor, was determined in each of the recorded frames using in-house software. This provided a basis to determine the position of the impactor itself in two dimensions.

4.1.3 Controlled Partial Vacuum Casting


4.1.3.1 Description of the Equipment

One of the main aims of this work was to evaluate a commercial High Pressure Die Casting through simulation and through experimental work. To some extent, such an evaluation must consider which assumptions may be made about the behaviour of the liquid metal during the HPDC process. However, the author is not aware of any prior experimental work which observed the fluid flow behaviour of liquid aluminium, at fluid velocities significantly greater than the purported 0.5 ms^{-1} limit for “good quality” castings.

As such, an experiment was needed, which could challenge the quality of the fluid flow models for process conditions approaching those in High Pressure Die Casting. The process selected was a form of counter gravity casting, in which the liquid metal is sucked from a crucible and into the mould by a partial vacuum. This method was used, because it theoretically allows access to flow velocities up to 9 ms^{-1} without the safety concerns of other means of accelerating the metal, such as a pressurised furnace.

An advantage of a counter-gravity type process is that the initial conditions are far easier to account for in simulation, than would be the case for a gravity casting. This is important, because the fewer plausible reasons there are for any disagreement between the experiment and the predictions of the simulation, the more confidence there will be in the final evaluation.

Again, to allow a closer equivalence between a casting and its simulation, an electronic control system was implemented, which continuously monitored and recorded the partial vacuum level throughout each experiment. This data could then be used as a boundary condition in a simulation, in principle allowing each experiment to be simulated individually, without having to assume one of the most important boundary conditions.




A



B



C



D

Figure 4-7: Photographs of the main elements of the controlled partial vacuum casting system.

Furthermore, the electronic system provided closed loop feedback control over the partial vacuum during casting. This theoretically allowed the specified process parameters to be achieved far more reliably than if the system were under manual control.

Eliminating the need for direct human control also allowed the system to be used in an X-ray environment. This was necessary, because one of the principle results of the controlled partial vacuum experiments was the real-time X-ray footage of the fluid flow in the mould cavity, filmed at $60 \text{ frames}\cdot\text{s}^{-1}$. The shape of the fluid during casting is one of the key features which the simulation must accurately predict.

Again, to allow the mould cavity to be observed under X-rays, the partial vacuum chamber was constructed from two layers of 18 mm thick ply-wood, as pictured in Figure 4-7. This chamber is lined with a fibrous ceramic insulating material, for safety reasons, and sealed with a silicone based foam tape. The square metallic feature visible in the base of the chamber in Figure 4-7 A is where the mould assembly was bolted into the chamber. Beneath this, a large stainless steel block is set into the wood, to thermally separate the plywood from the mould assembly during casting.

The chamber sits on top of a small 16 kW resistance furnace, supported by a steel frame, as can be seen in Figure 4-7 B. The furnace was in turn mounted on a wheeled steel box-section frame, along with the other ancillaries for the controlled partial vacuum casting system, including the furnace controller. This allowed the whole system to be moved into and out of the University of Birmingham's real-time X-ray chamber as a single unit.

The wheeled frame also supports the vacuum pump, an Edwards model 18, and the vacuum reservoir, which was made by modifying a gas cylinder with vacuum fittings, and adding an additional inlet.

Frames C and D of Figure 4-7 show the pressure control box, along with the pipe-work and control valves. Pressure in the chamber can be controlled with four ball valves: manual vacuum shut-off, actuated vacuum control, actuated vent control, and manual vent by-pass.

One of the main design considerations for the system was to maximise the possible flow-rate of air out of the partial vacuum chamber, so that the greatest possible liquid metal velocities could be achieved. For this reason, 28 mm copper pipe was used, and the only angled connection was given a large radius.

Similarly, the decision was made to actuate the valves using 5W stepper motors, since these were considered to provide sufficient angular precision (0.9° per half-step), while allowing the valves to be actuated between the fully open and fully closed positions in 0.5 s. The motors needed to be able to react to changes in the chamber pressure as fast as possible, so that large changes in target pressure could be achieved and stabilised as quickly as possible, because this should allow control over high fluid velocities to be maximised.

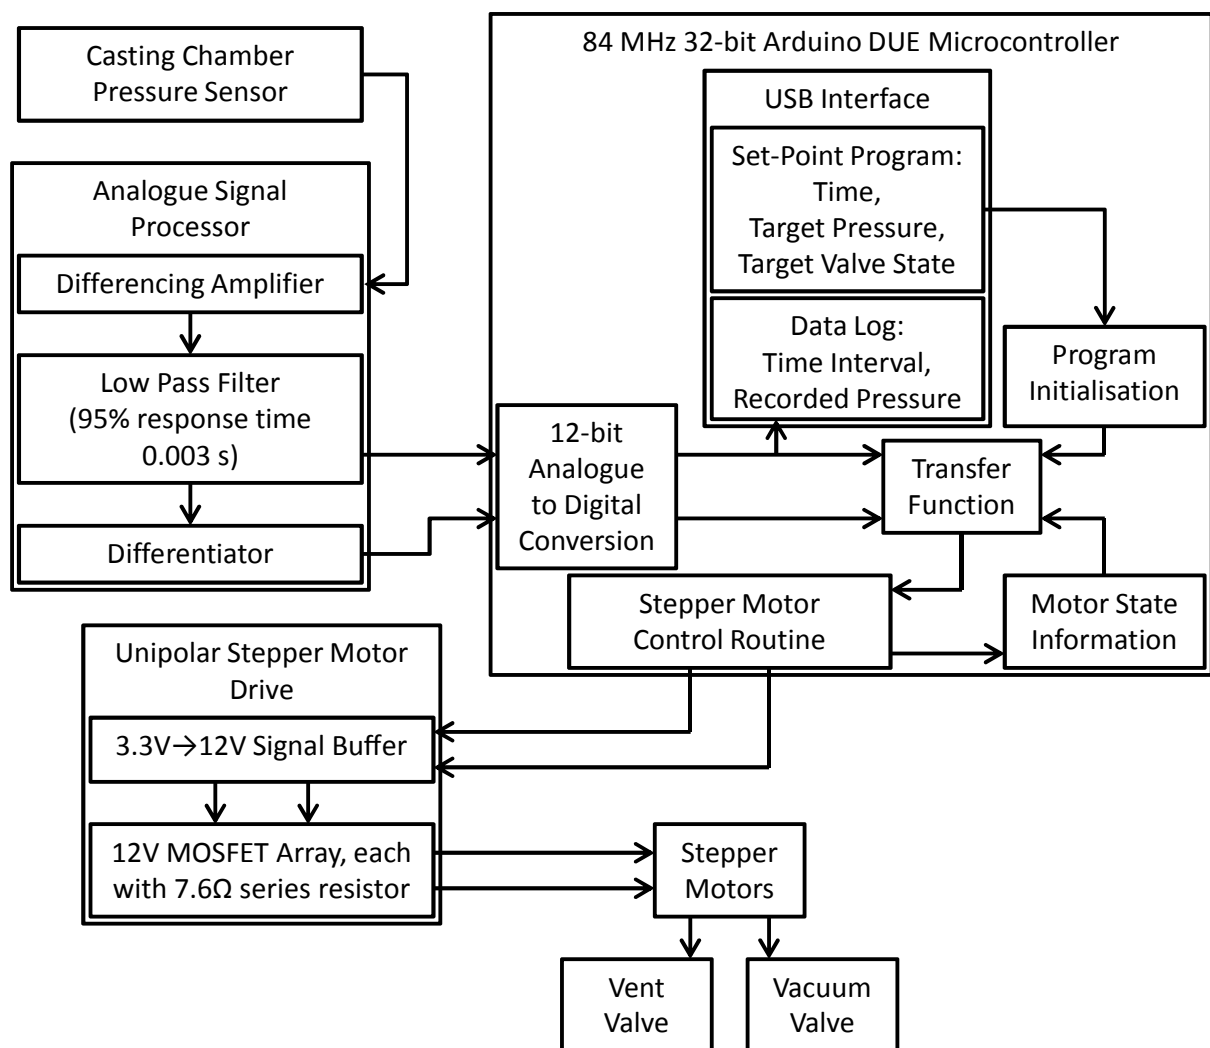


Figure 4-8: Block diagram outlining the main elements of the control system.

As illustrated in Figure 4-8, the parameters for an experiment, a list of set-points, were provided by an open-source interface to a USB memory stick. The set-point list is a series of lines, each specifying three values: the time from the start of the program that the set-point applies to, the target pressure below atmospheric, and the “target valve state”. This latter value controls and limits the how open both valves are in total. For example, a value of 50 might be specified: This would allow neither valve to open more than half way, but the controller would seek to control the pressure in the chamber such that one valve is $a\%$ open, and the other $b\%$ open, where $a+b=50$. For times between set-points, the target values are found by linear interpolation.

Equations 4-1 to 4-4 specify the functions used to calculate the motor control response, converted such that angles are specified in degrees, and pressures are gauge pressures in Pa. The motor control variables are target rotational speeds (in the open direction); the control functions for both vacuum and vent motors have a term which defines the pressure seeking behaviour, and a term which seeks the specified target valve state.

$$\dot{\theta}_{vacuum} = \dot{\theta}_{open} + \dot{\theta}_{pressure} \quad \text{Equation 4-1}$$

$$\dot{\theta}_{vent} = \dot{\theta}_{open} - \dot{\theta}_{pressure} \quad \text{Equation 4-2}$$

$$\dot{\theta}_{pressure} = 0.027\dot{P} + 0.18(P - P_{target}) \quad \text{Equation 4-3}$$

$$\dot{\theta}_{open} = 0.15 \left[\text{Target Valve State} - \frac{\theta_{vacuum} + \theta_{vent}}{0.9} \right] \quad \text{Equation 4-4}$$

The function that defines the pressure seeking behaviour, Equation 4-3, includes a term related to the difference between the target pressure and the measured pressure, and a term relating to the time differential of measured pressure. The form of this equation is based on a transfer function, which would prescribe a damped second order pressure response, if the rate of change of pressure in the casting chamber were directly proportional to valve angle. The

constants in this equation were based on the trial and error evaluation of a simplified numerical model of the system, which used loss coefficients calculated from a FLOW-3D simulation of the valves.

As indicated by Figure 4-8, the pressure signal is differentiated in the analogue signal processor, a network of operational amplifiers, resistors and capacitors. An analogue differentiator was used, because the numerical differentiation of high sample rate ($1000 \text{ samples} \cdot \text{s}^{-1}$) discretised data would introduce noise, which would then have to be digitally filtered.

The motor rotational speeds defined in Equations 4-1 and 4-2 are target speeds, which are used in the stepper motor control routine, to prescribe the actual speed of each motor. These speeds are firstly limited by the motor position; the motor should not turn the valve past the closed position, and the upper limit is defined by the “target valve state” as described above. The magnitude of angular acceleration is also limited to around $900 \text{ degrees} \cdot \text{s}^{-2}$; without such a limit, it is likely that the motor would skip steps, and the actual position would differ from the intended position.

The stepper motor control routine also translates the motor speeds into prescribed positions over time, and signals which coils should be energised by the motor drive circuit. A half-stepping scheme is used for positional accuracy. This signal is passed to two banks of four P-channel MOSFET transistors. Each of these transistors is wired in series with one of the four coils in each motor. Pairs of “opposite” coils (those which should not be energised at the same time) share a bank of 9 parallel resistors, with an equivalent series resistance of $\sim 7.6 \Omega$, and a combined power dissipation limit of 9 W. This, together with the addition of a cooling fan, allows 12 V drive circuitry to be safely used with 5 V motors. A magnetic field coil can be modelled with resistive and inductive elements. By using a resistive choke, the over-voltage is selectively applied to the inductive element of the coil, which should reduce

the rise-time of the magnetic field when the coil is first energised by ~60%, and therefore extend the motor torque to higher rotation speeds.

4.1.3.2 Plate Casting Experiments

The geometry chosen for these casting experiments was a 5 mm thick square plate, with a side length of 150 mm, as shown in Figure 4-9. A square plate was used because HPDC parts generally feature plate-like shapes, with a reasonably uniform section thickness that is significantly smaller than the overall dimensions of the part.




Figure 4-9: The pattern (A) and a preliminary casting (B) from the plate casting experiments.

The gate was situated in the centre of the plate's lower edge the base of the plate. The pattern for this, was made as a separate 3-D printed insert, the off-white part in Figure 4-9 A. This insert includes a semi-cylindrical feature, to allow the steel down-pipe to be glued into the mould as it is assembled. A pattern insert was used, because this would allow future

experiments to be carried out using different gate geometries, and pipe sizes, without modifying the pattern.

Figure 4-9 B shows a casting which was made during preliminary experiments to troubleshoot the casting system, and to gauge appropriate furnace temperatures. This photograph shows the steel dip tube, together with the plate which mounts the mould assembly to the casting chamber. The base of the pipe was coated with boron nitride to form a barrier between the steel and the liquid aluminium in the furnace. Also shown in this photograph are the “chill-vents” which were inserted into the mould above the plate. These were aluminium ~10 mm thick aluminium blocks, which had channels cut into them, such that air could flow freely through them, but liquid aluminium would freeze between them. This was necessary to ensure close equivalence between the air pressure in the mould and that in the casting chamber, and also to ensure that liquid aluminium would not spray out of the mould and into the wooden chamber. As an additional precaution, a sheet of fibrous ceramic material was glued to the outside of the mould, over the vent outlets.




Figure 4-10: Dimensioned drawing of the gate geometry.

Figure 4-10 details the geometry of the gate. The geometry is defined by smooth curves, to minimise the level of surface turbulence as the fluid enters the mould at slower speeds. The gate is also taller and wider than is usual in a sand casting. This is because the gate design was intended to approach the “fan” gate style sometimes used for High Pressure

Die Casting. These are designed to spread the metal throughout the mould, as it enters at high speed.

Other key dimensions and quantities related to these casting experiments were measured so that the experiments accurately simulated. A labelled drawing of these measurements is presented in section 4.2.3 as Figure 4-14.

Six experiments were performed using an Al-7Si alloy, and observed using a real-time X-ray unit, which recorded the flow patterns at $60 \text{ frames} \cdot \text{s}^{-1}$. The parameters for these experiments are listed in Table 4-4. Only one of the six experiments was instrumented with thermocouples, because it was believed that the thermocouples would disturb the fluid flow. This would be undesirable, as the fluid flow observations were the main objective of the experiments.

As such, the thermocouple instrumented test was a repeat of the “Slow” experiment. For this test, two type-K thermocouples were used. One was placed with its junction just above the gate, and the other with its junction near one of the upper corners of the mould.

These thermocouples were logged using a set of SCXI series National Instruments modules, which collectively interface to Lab-View, running on a desktop PC. The cold junction temperature was specified manually, based on a thermometer in the same room.

The furnace temperatures listed in Table 4-4 were measured by dipping a thermocouple into the liquid metal, immediately before mounting the casting chamber on the furnace.

The main parameter varied for each of the experiments was the rate of change of pressure in the casting chamber. This can be related to a rate of change of fluid height, as given in Equation 4-5, which should be approximately the rate that the fluid rises in the dip tube, and therefore the velocity of the metal as it reaches the gate. Another relation could have been used which links fluid velocity to pressure difference via kinetic energy but, the below formula was used because it is more applicable to a transient case involving rising fluid.

Table 4-4: Parameters for Controlled Partial Vacuum Experiments

Experiment Designation	Specified Pressure Ramp Rate (Pa·s⁻¹)	Equivalent Fluid Rise Rate (ms⁻¹)	Thermocouple Instrumented?	Furnace Temperature (°C)
Slow	-9,556	0.4	No	765
Slow TC	-9,556	0.4	Yes	764
Medium	-23,812	1	No	744
Medium Fast	-34,913	1.46	No	744
Fast	-45,848	1.92	No	734
Dump	-	-	No	732

$$\dot{h} = \dot{P} / \rho g \quad \text{Equation 4-5}$$

The full set pressure over time for each of the experiments is given in Figure 4-11, based on the parameters from Table 4-4 . The set-points used for each experiment specified the minimum target pressure in the casting chamber as 20 kPa below atmospheric, except for the “Dump” experiment. For this test, the minimum chamber pressure was set to full vacuum, so that the valves would only be controlled by the “target valve state”, and the system was directed to completely open the vacuum valve at full speed. The resulting liquid metal flow rate should therefore be as fast as possible, and have the most HPDC-like flow characteristics.

In each experiment, the pressure was directed to settle at 15 kPa below until 5 minutes after program start, when the vent valve would open, and the program would end. The equilibrium pressure required to fill the mould with liquid metal was approximately 10.6 kPa below atmospheric.

The “jump” shown in most of the set-point graphs when vacuum is first applied at 10 seconds after the start of the program is calculated to exactly mitigate the 0.15 s “lag” in the equilibrium response of Equation 4-3 to steadily changing pressure. For these tests, lines involving “target valve state” were used to direct vacuum valve to start opening before the 10

second mark, so that the valve would be just open enough to allow some air flow as soon as the partial vacuum was specified.




Figure 4-11: Target pressure below atmospheric over time, for each experiment.

The level of liquid metal in the crucible was maintained by adding a small quantity of pre-heated metal between each experiment, with a mass close to mass lost into casting and dip-tube during the previous test.

4.2 Simulation

4.2.1 Surface Area Entrainment Code Robustness Characterisation

In the literature review, the Boolean Oxide Entrainment Code (BOEC), developed by Reilly [20], was identified as the most promising Oxide Film Entrainment Model available outside of the present work. However, a number of issues were identified with this Boolean model, as presented in the Literature Review (Section 2.2.3), and in prior work by the author [58, 59].

With the aim of addressing these issues, a novel algorithm was developed, as described in Theoretical Developments, Section 3.1. This was called the Surface Area Entrainment Code (SAEC), because it employs mixture of logic and calculation to evaluate the entrainment of surface film based on the “flow” of surface area. This model identifies whether the free surface is unstable enough to be entrained, and also quantifies how much (if any) surface area can be entrained. In theory, these two measures would allow an Oxide Film Entrainment Model (OFEM) to determine when, where, and how much oxide film would be entrained with a reasonable accuracy.

However, this theory must be tested and evaluated. This provides an opportunity to demonstrate the functionality of the algorithm, but more importantly, the strengths and weaknesses of the algorithm must be identified, to support the evaluation of the other results based on that algorithm.




Figure 4-12: The test cases used to demonstrate the SAEC particle algorithm.

Three different test cases were simulated, to demonstrate a variety of flow conditions for the algorithm to evaluate. Each of these tests (as shown in Figure 4-12) are pseudo-2-D simulations, with symmetry boundary conditions enclosing a 6 mm thickness in the out of plane direction as pictured.

The simulation depicted in Figure 4-12 A is a “falling droplet” case. This was initialised with fluid already in the domain, where a circular droplet is given an initial velocity of 1 ms^{-1} downwards. For the other two models, fluid enters at a gate at a velocity of 1 ms^{-1} , as indicated by the arrows. A summary of the solution control parameters for these simulations are tabulated in the Appendices, section 10.3.1.

The qualitative results of these tests were the distribution of particles produced by the Oxide Film Entrainment Model (OFEM), together with the shape of the simulated fluid flow.

Quantitative results were also extracted from the simulation results files, to help evaluate the entrainment predictions; these were the particle count over time, and the total fluid surface area over time.

To simplify the equivalence between global surface area loss, and increased particle count, the proportionality constant for the SAEC algorithm was set such that one particle should represent 1 mm² of entrained surface film.

The fluid properties used in the simulation were those for A357 alloy, taken from the FLOW-3D database. Surface tension was included, but heat transfer and solidification were not modelled in these simulations.

The k - ε turbulence model [60] was used for each of the simulations presented in this work. shows the transport equation for ε_T (turbulent dissipation) as solved by FLOW-3D, as taken from the software's theory manual[61].

$$\frac{\partial \varepsilon_T}{\partial t} + \frac{1}{V_F} \left\{ u A_x \frac{\partial \varepsilon_T}{\partial x} + v A_y \frac{\partial \varepsilon_T}{\partial y} + w A_z \frac{\partial \varepsilon_T}{\partial z} \right\} = \frac{C1 \cdot \varepsilon_T}{k_T} (P_T + C3 + G_T) + \text{Diff}_\varepsilon - C2 \frac{\varepsilon_T^2}{k_T} \quad \text{Equation 4-6}$$

In the above, the left hand side of the equation amounts to the rate of change of ε_T for a given unit of moving fluid. On the right hand side, k_T is the specific turbulent kinetic energy, P_T and G_T are terms for the production of turbulence through shear or buoyancy respectively, and Diff_ε is short-hand for the inclusion of diffusion effects. $C1$, $C2$, and $C3$ are dimensionless parameters. In this work, their default values were used, which are 1.44, 1.92 and 0.2 respectively.

While these turbulence parameters themselves were not tuned for this work, preliminary investigations were undertaken which sought to reproduce published experimental work by Schmid and Klein [62], and by Sirrell, Holliday and Campbell [63]. This exercise found that the above parameters gave sufficient agreement with these experiments, if the surface roughness of the mould material was adjusted (coarse for sand casting, smooth for HPDC) to affect the turbulent production terms.

4.2.2 Parametric Investigation of Partial Vacuum Casting

To complement the partial vacuum casting experiments described in section 4.1.3, a parametric investigation was performed. The role of this investigation is firstly to characterise the general response of the flow pattern to increased nominal flow velocity. Secondly, it provides a simple flow case, with which to more rigorously characterise the response of the Surface Area Entrainment Code to changes in nominal velocity and mesh size.

For this investigation, a heavily simplified model of the plate casting described in section 4.1.3 was used, which largely neglected heat transfer in the dip tube, and the crucible was treated as a small volume supplied by a pressure boundary condition. A summary of the solution control parameters for these simulations are tabulated in the Appendices, section 10.3.2.

Instead of reducing the pressure in the mould over time, as was the case experimentally, an increasing pressure was applied as a boundary condition at the simulated crucible. Since the flow would be driven by pressure differences, rather than by absolute pressure, this approach was valid. As with the experimental work, a nominal flow rate may be equated to rate of change of fluid height, and thus to rate of change of pressure, using Equation 4-5 from Section 4.1.3.2.

Twenty different permutations of nominal velocity and mesh size are presented in this work: Nominal flow velocities were varied between 0.25 ms^{-1} and 1 ms^{-1} and the mesh size in the plate was varied between 0.65 mm and 1.25 mm, which corresponds to between 7 and 4 mesh cells across the 5 mm plate, respectively. Each permutation used a unique value both nominal velocity and mesh size, in what is known as “Latin square” sampling. This approach allows the largest number of unique values to be studied with the fewest experiments.

For both parameters, the values chosen would form a geometric sequence ($x_{n+1}=Ax_n$) if they were sorted. This allows the study to focus on lower values of mesh size and velocity,

and also makes it more correct to map the parameters to their logarithms in future analysis (i.e. power law regression).

The parameter combinations were chosen such as to maximise the determinant of their co-variance matrix, up to second order. This means that in later analysis, it is possible to fit an analytical description of the response (e.g. response surface), with confidence, and minimise interference from any second order interactions. The parameter combinations chosen are co-plotted in Figure 4-13.

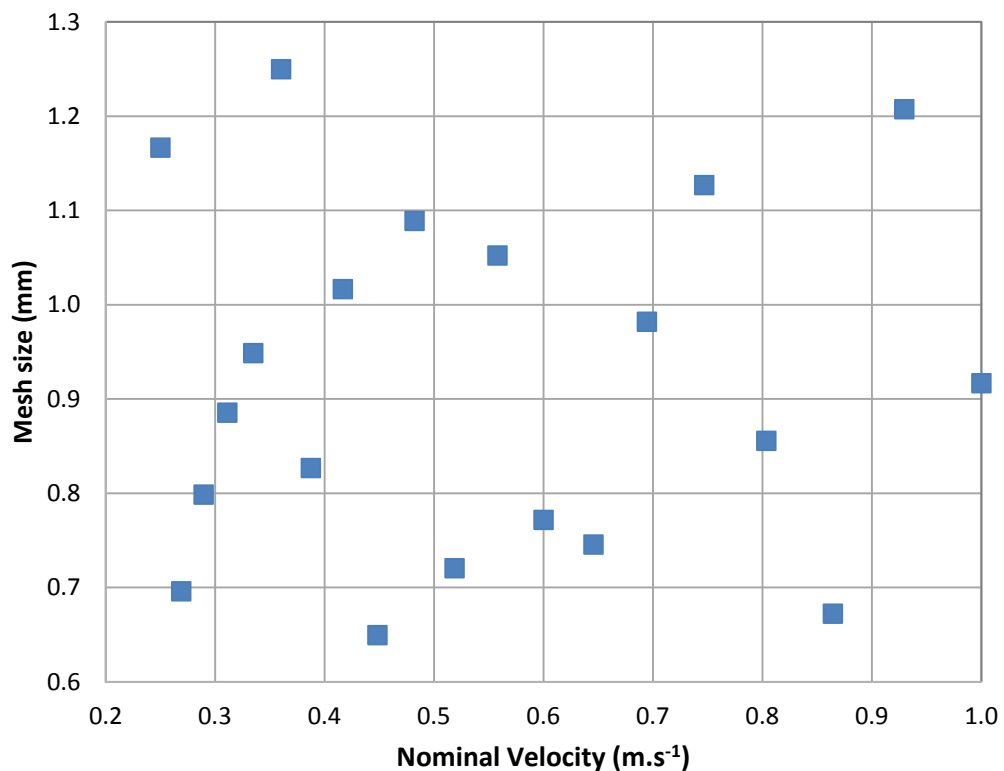


Figure 4-13: The 20 parameter combinations used to investigate velocity and mesh sensitivity.

4.2.3 Controlled Partial Vacuum Casting Simulations

The controlled partial vacuum casting experiments described in section 4.1.3.2, were translated into a simulation of each individual experiment, based on data gathered during that experiment. Thermal data and other information about the overall experimental set-up was translated into the initial conditions for the simulation. Crucially though, the pressure data-logs gathered by the Arduino DUE microcontroller were filtered and translated to become the time-dependant pressure boundary conditions used in each simulation.

In theory, this approach should provide a good opportunity to verify the capability of the simulation. Firstly, this is because “counter gravity filling” is hypothesised to significantly reduce the level of chaos in the casting process, because it starts from rest, and therefore minimises the introduction of turbulence. This means that the initial conditions of the casting simulation should be a close representation of reality. Secondly, by using the measured pressure curves as the main boundary condition for the simulations, the forces on the simulated fluid over time should closely match reality.

Overall, this approach should minimise the uncertainties in the assumptions made while building the models. As such, the cause of any discrepancies between the simulated flow patterns and the experimental real time x-ray footage could be isolated with more confidence. This increased confidence may be sufficient to challenge the physical approximations used in the simulations.

Each of these simulations was modelled in two stages, using the simulation restart feature in FLOW-3D. This allows a simulation to be initialised with the results of another simulation; this means that the final result of the first stage would become the initial conditions for the second stage. In this case, the first stage was a thermal-only simulation, where air was included as a second “phase”. This allowed the distribution of heat in the air, liquid metal and the dip tube to be initialised in a way which mimics that fact that the mould assembly was sitting in position for 20 minutes before the start of each experiment.

For the second stage of the simulation, the air around the dip tube was modelled as a “solid”, with the same thermal properties as air. This effectively allowed conductive heat transfer in air to be modelled, without invoking a two-fluid flow model; while FLOW-3D has that capability, simulations of that type were found to be very unstable in FLOW-3D.




Figure 4-14: Labelled drawing of the model set-up.

The geometry used for these simulations is presented in Figure 4-14, which was based on measurements taken from the casting experiments. This includes the placement of the virtual thermocouples, which allow detailed data to be extracted at a point. The steel dip-tube was modelled as 3.75 mm thick, as opposed to the actual value of 1.5 mm. This allowed the

dip-tube to be resolved by a coarser mesh, without introducing holes. The heat capacity and thermal conductivity of the mild steel dip tube were modified to preserve the “correct” heat flux based on 1-D heat flux approximations. Dominant heat flux was assumed to be axial in the first stage, and radial in the second stage.

The crucible was modelled with “wall” boundary conditions, with a specified temperature based on the furnace temperatures, as listed in Table 4-4.

The “chill-vents” were modelled as a region of “solidified fluid”, with 5 mm holes in. The heat transfer coefficient between liquid and “solidified fluid” is effectively infinite in the simulation.

In addition to the models of fluid flow with a k- ϵ turbulence model, heat transfer and solidification, the “adiabatic bubble model” was used. This allows mould gas to be accounted for, without using a completely coupled two phase simulation. Instead, separate void (mould gas) regions are treated with a uniform internal pressure which changes as the volume of each bubble changes. The model is based on ideal gas laws, but assumes no heat transfer between the bubble and its surroundings. A summary of the solution control parameters for these simulations are tabulated in the Appendices, section 10.3.2.

As with the parametric simulations, pressure was increased in the crucible, rather than reduced in the mould. This meant that the “Porous Mould” feature could be used to allow any bubbles trapped against the mould wall to vent to a consistent pressure.

4.2.4 Simulation of Commercial HPDC Experiments

4.2.4.1 'Short Shot' Simulation Fidelity and Mesh Sensitivity Test

As described in section 4.1.1, the characterisation of the commercial High Pressure Die Casting studied involved some "Short Shot" testing. These involved reducing the fill fraction of the shot sleeve, so that incomplete castings were produced. The piston speed must also be reduced (1.15 m.s^{-1} nominal piston speed), so that the piston braking system in the Die Casting Machine (DCM) can stop the piston, without relying on back-pressure from the liquid metal. These results can give valuable information about the characteristics of the flow, although there are severe difficulties related to their use as flow validation.

The lower the shot sleeve fill fraction, the less representative that test is of the standard casting; the quantity and distribution of fluid in the shot sleeve at the start of the shot is likely to have a strong effect on the dynamics of the fluid before it enters the mould. In particular, for lower fill fractions, the piston would accelerate to full speed before the runners are properly primed. However, a lower mass of metal would result in a more incomplete casting; the pattern of solidified fluid in the mould flowed into that shape, and the more incomplete the casting, the more that shape is determined by fluid flow forces, rather than the extent of the mould cavity. Therefore, to some degree, these solidified patterns can be equated to the flow patterns early in mould filling.

Only one of the experimental short shots was simulated; the one which used 1.9 kg of aluminium, corresponding to a shot sleeve fill fraction of 0.137. This test was chosen for simulation because it was one of the lowest mass castings to seem full from the outside. This meant that this charge mass represented a compromise between the above conflicting arguments; it was considered that simulated flow patterns for this test could be reasonably equated to the solidified fluid patterns in the lower mass short shot tests, while the parameters used in this short shot simulation would also be sufficiently relevant to the simulations of the full shot experiments.

As such, a number of preliminary short-shot simulations, not presented in this work, were used to determine plausible values for two key drag parameters, by comparing their flow patterns with the solidified fluid patterns in other short shot tests. These two drag parameters were the “Surface roughness”, which affects the generation term in FLOW-3D’s $k-\epsilon$ turbulence model, and a solidification drag parameter, which affects how quickly the kinetic energy of partially solidified metal is dissipated as the liquid flows through the solid. This latter parameter naturally depends on the heat transfer coefficient used in the simulation, which was $5 \times 10^5 \text{ W m}^{-1} \text{ K}^{-1}$, based on the work of Dour et. al. [64].

Three different short shot simulations are presented in this work; each is identical except that a different nominal average mesh size (0.5mm, 0.8mm and 1mm) was used for each one. This was to investigate how the simulation results varied with mesh size. The short-shot simulations were used for this investigation, because the short shot tests are not as numerically intensive as the full simulations described in section 4.2.4.2.

FLOW-3D has a feature to restart a simulation using data from a previous result, which allows the overall simulation to be broken down into stages. The mesh, domain, and other simulation parameters are not necessarily the same for each stage; in this case the mesh size was only varied in the second stage.

For these simulations, the domain of the first stage included the entire shot sleeve and the running system, but excluded the mould. The initially static fluid was set in motion directly by the moving piston (modelled using FLOW-3D’s moving object model), which was given a prescribed velocity according to the parameters specified Tables 4-2 and 4-3.

The transition to the second simulation stage was made immediately before the fluid reached the gate, which was the limit of the mesh used for the first stage. At this point, the shot sleeve was removed from the domain, and replaced by a mass flow boundary condition, and the rest of the mould cavity was added. The time when the correct amount of metal should have entered the mould was calculated, and the mass flow was shut off after this time.

The simulation was terminated 0.2 seconds of simulated time after the mass flow was shut off, to allow the fluid motion to cease, and form the pattern of solidified material.

In reality, the mould and running system produce two symmetrical castings. For each of the stages in these simulations only one of the symmetrical halves of the casting geometry is modelled, in the interests of simulation run-time. Where the two symmetrical halves meet (i.e. the shot sleeve) a symmetry boundary condition was used.

As with the other FLOW-3D simulations presented, surface tension was modelled ($\gamma=1 \text{ Nm}^{-1}$), and the FLOW-3D k- ϵ turbulence model was used. Heat transfer and solidification were also active, as described above, along with the adiabatic bubble model. A summary of the solution control parameters for these simulations are tabulated in the Appendices, section 10.3.3.

As well as the shape of the fluid at each recorded time-step, the outputs for this exercise include the kinetic and thermal energy over time for each case. These data will help isolate any mesh sensitivity effects, since the dissipation of energy to the mould can be strongly affected by how well the boundary layer is resolved.

4.2.4.2 Full Shot Simulations

The simulations of the full shot experiments were carried out using the same simulation parameters and model options which were used for the short shot simulations (section 4.2.4.1), unless specified otherwise in Table 4-3. In other words, a higher short-sleeve fill fraction was used, since a full casting was to be produced in each case, and the second phase shot speed and the gate condition were varied according to which of the casting trials the simulation related to. Figure 4-15 shows how the “Blanked” and “Full” gate trails were accounted for in simulation.




Figure 4-15: Full (A) and blanked (B) gate configurations used in HPDC simulation.

As with the short shot simulations, the simulations were divided into stages using FLOW-3D’s restart feature, and again, the domain for the first stage included the shot sleeve and excluded the majority of the mould cavity. As before, the transition to the second stage was made immediately before the liquid metal reached the gate.

In the second stage the remaining parts of the mould cavity were included in the simulation domain, and the shot sleeve was again removed, and replaced with a mass flow source. The second-stage mesh is the same as the nominally 0.8 mm mesh used in short-shot simulation. For these simulations, however, a modification to FLOW-3D’s mass source subroutine allowed the mass flow rate to be gradually reduced as the mould filled. This was found to be necessary to avoid pressure spikes towards the end of filling.

When the mould cavity was less than 95% filled, the mass source rate was determined by the second phase piston speed from Table 4-3. Above 95%, the source rate was reduced as a linear function of the remaining open volume, such that the mass source rate would be zero if the casting were completely full. If the casting were filled at a constant rate until it was completely full, then the pressure would be difficult to compute towards the end of filling, which would cause numerical instability, and spurious results. Furthermore, this approach smooths the transition between the simulation of filling, and the solidification only phase of the simulation. This allowed the fluid velocity field after filling to be carried into the solidification stage.

For the full shot simulations, a third stage was used to model the solidification, where the numerical options were less focussed on accurately resolving the fluid flow structure, so that the solution speed could be increased. One of the measures taken was to use the coarser (~1 mm) mesh from the short-shot study. The model, including fluid flow, was evaluated until 2.3 seconds of simulated time from the start of piston motion (~0.3-0.4 seconds after the casting is 95% full). This was to allow fluid flow in the plate sections to cease under the influence of solidification drag, so that the predicted entrainment defects to settle in their final positions, as they would in a real casting.

For the solidification stage, the piston was represented by a pressure boundary condition, set to the 100 MPa intensification pressure specified for the commercial casting process. Such a boundary condition is less likely to cause numerical instability, since any pressure spikes arising from numerical noise can theoretically be dissipated. Again to improve numerical stability, the parameters for cleaning up the fluid-void interface were altered to encourage any bubbles to collapse. A summary of the solution control parameters for these simulations are tabulated in the Appendices, section 10.3.3.

As with the short-shot simulations, only one symmetrical half of the casting geometry was simulated, using a symmetry boundary condition where applicable.

The history of particle count over time was extracted and compiled from each stage of each simulation. However, the principle results of these simulations are the fluid flow shape over time, and the particle distribution over time. The final distributions of solidified fluid and particles are some of the key inputs to the statistical damage mapping as detailed in section 3.2, and applied in section 4.3.2.

4.2.5 Mapping Predicted Damage Statistics to Casting FE Model

The entrainment damage prediction model used in this work is based on the theory presented in section 3.2.5, which allows a statistical distribution of “damage” to be calculated for an arbitrary “element” in the domain of the fluid flow simulation. “Damage” at a point is defined as the smoothed number density of particles (representing entrainment defects), centred on that point. The element-by-element statistical parameters output by the algorithm described in section 3.2.5 are for a Fréchet distribution, as defined by Equation 3-? in that section.

Equation 4-7 re-states the Fréchet distribution in terms of the cumulative probability, P , but including three additional parameters, d_{offset} , V and V_{eff} . The d_{offset} parameter is the “damage offset”, and is included to account for a constant background level of damage such as might result from defects formed during the filling of the shot-sleeve.

$$P = \exp \left\{ - \left(\frac{V}{V_{eff}} \right) \left(\frac{d + d_{offset} - \mu}{d_0} \right)^\beta \right\} \quad \text{Equation 4-7}$$

V and V_{eff} form a “volume scaling” term, where V is the volume of an element. Volume scaling allows an extreme value distribution to apply to different sample volumes; for example a larger sampling volume would be more likely to contain a critical defect. V_{eff} , the “effective volume”, is the volume of a test-bar gauge length, multiplied by a scaling factor, to account for the fact that the gauge volume of a test bar differs from the sampling volume used in the procedure described in section 3.2.5.

Both the damage offset and the effective volume parameters were found using the parameter estimation method described in section 4.3.2. This method was also used to find two other correlation parameters, A and B , are defined in Equation 4-8. These parameters define a power law relationship between the “damage”, d (as per Equation 4-7), and the true fracture stress of the element, σ^* .

$$\sigma^* = Ad^B \quad \text{Equation 4-8}$$

By sampling a value for P from a uniform distribution between 0 and 1, these two equations can be evaluated to sample a value for σ^* for each element, which has a defined Fréchet distribution for d .

However, there are practical complications to implementing this in the LS-DYNA Finite Elephant simulation, described in section 4.2.6. Firstly, the mesh used to represent the casting in the structural impact simulations was very large; to evaluate the whole statistical mapping procedure for each of these elements would have been too numerically intensive.

Instead, a coarser mesh, consisting of 36,142 tetrahedral elements was used to evaluate the element-by-element Fréchet statistics. Then each element in the finer mesh was assigned to its nearest neighbour in the coarser mesh; each of the “fine” elements assigned to the same “coarse” element form an element group. This group has known Fréchet parameters, inherited from the “coarse” element, and a known volume, which was the sum of the volumes of each “fine” element in the group. The fracture stress for that element group can then be evaluated, as above.

Secondly, LS-DYNA does not allow material properties to be specified on a per element basis, without either defining a material for each element, or a user material subroutine. The approach taken here was to take the evaluated fracture stress for each element group (as defined above), and to merge element groups with a similar fracture stress together, as a single “part”. In LS-DYNA syntax, elements assigned to the same “part” share a material definition.

The fracture stress of each part is the mean of the fracture stress for each element group assigned to it. The algorithm used to achieve this “part-binning” process was designed to minimise the least squares error between the sampled fracture stress of each group of elements in the “part”, and the fracture stress assigned to the part.

4.2.6 Whole Casting Structural Impact Test

As described in section 4.1.2, experiments were performed to investigate the behaviour of a whole cast part under arbitrary loading. The overall aim of this exercise was to investigate the ability of the strength prediction methodology developed during the course of this work to accurately predict the failure of cast parts, by loading the commercial cast part in a way which was not used for correlation.

To make the comparison between the experimental behaviour, and that predicted by simulation, the test which had the lowest impact energy was simulated a total of six times using a Finite Element model; only the strength of the cast material was changed in each run. One of these simulations placed no limit on the strength or ductility of the material. The remaining five simulations included a statistically sampled strength distribution, based on the FG-FS flow simulation, and the results of entrainment damage correlation. A “maximum principle stress” fracture criterion was used; if the maximum principle stress in an element exceeds a specified value, that element is deleted from the simulation. The method of correlating entrainment damage with strength, and the stochastic strength mapping method are described in sections 4.3.2 and 4.2.5 respectively.

These structural simulations were performed using the explicit non-linear Finite Element Analysis (FEA) software LS-DYNA. The parameters used for the material model were based on an average of the elastic and yield properties determined from the tensile testing of the two “Fast Shot” experiments (section 4.1.1). The yield curve used assumed Voce type strain hardening behaviour, [??] where the hardening parameters were fitted to the

data using a non-linear least squares approach. Strain-rate dependant effects were not accounted for.

In these simulations, the impactor was modelled as a solid rigid body, but the trolley was modelled as a point mass that was rigidly connected to the impactor. This rigid connection is shown as a blue dotted line in Figure 4-16. Frame B of this figure also shows the point mass used to represent the trolley; the value of this mass was calculated as the experimental mass for the trolley assembly, minus the modelled mass of the impactor.




Figure 4-16: Illustration of constraints and loading points for drop-test Finite Element simulations.

The green secondary axes, shown in Figure 4-16 frames A and B, were used to link to the horizontal (X and Y) and vertical (Z) directions in the experiment. This allowed the motion of the assembly to be defined simply, while allowing other aspects of the simulation to be defined in terms of the main axes. This simplified the defect mapping from the FLOW-3D simulation.

The centre of mass of the modelled trolley assembly was constrained in the X and Y directions of the coordinate axes shown in green in frames A and B of Figure 4-16; this corresponds to the horizontal constraint of the trolley in the experiment. The modelled trolley assembly was given an initial velocity of 3 ms^{-1} in the -Z direction of the green coordinate axes. Gravity was applied to all of the parts in the experiment in this same direction.

Preliminary simulations showed that the dynamics of the experiment could not be explained without accounting for flexibility in the system. As such, four parameters were added, which define a simplified model for this flexibility.

The first degree of flexibility is that the modelled trolley assembly was permitted to rotate around its centre of mass. A parameter was defined to describe an effective “spring”, which acts on the centre of the impactor in the horizontal directions to return the impactor to vertically below the centre of mass of the modelled trolley assembly. These “horizontal” and “vertical” directions are again in the orientation of the experiment, as defined by the green axes.

The component that the cast part is clamped to was also permitted some flexibility. The constraint of the cast part is defined as what is called a “nodal rigid body” in the LS-DYNA syntax; this means that a set of nodes from the casting, those labelled with blue crosses in frames C and D of Figure 4-16, are locked together such that they behave like a single rigid body. In this case, the nodal rigid body was constrained so that it could only rotate in the in-plane direction as depicted in Figure 4-16, around the node marked with black cross hairs in frames C and D. The dynamics of this nodal body were described by the remaining three

parameters, one for rotational inertia, one for damping, and one defining a torsional spring rate.

The values for each of these four flexibility parameters were estimated using a set of parametric simulations, which were not included in the thesis because they fall outside the core subject matter. In that tangential set of work, the simulated motion of the impactor was fitted to the experimental motion of the impactor, based on motion tracking from the high speed cameras, by varying these four parameters.

For the presented simulations, the mesh used to represent the casting was composed of 453,950 tetrahedral elements.

4.3 Analytical

4.3.1 MLE Extension for Multiple Concurrent Defect Types with Classification

This section describes some extensions to Maximum Likelihood Estimation (MLE), a numerical procedure used for fitting statistical data, and which is sometimes used for the analysis of failure data [18].

A probability distribution function defines the probability that an event would occur. The “likelihood” of this event can be defined as the logarithm of this probability. As such, an expression can be formed, which defines the total likelihood that each observed event (i.e. the failure of a test bar at a particular stress) occurred in an experiment, in terms of some statistical parameters. Then, a numerical procedure may be used to maximise the likelihood (and therefore the probability) that the observed events occurred in the experiment, by manipulating the parameters for the supposed statistical distribution. This is maximum likelihood estimation.

Logically, there is no value greater than the maximum. As a result, if a set of parameters are found by MLE – such as the position and scale parameters for a Weibull distribution – then there are no sets of parameters that better explain the observed result. One caveat is that a

local maximum is not necessarily the global maximum. Secondly, if a numerical procedure is used for maximisation, the result would not necessarily be the same as the analytical maximum (if that could be derived), which would be provably greater.

The expression for the total “likelihood”, L , as applied to the two parameter Weibull distribution, is given below. Here, σ_0 is the scale parameter, and m is the shape parameter or modulus, as per Equation 2-1 in the Literature Review.

$$L = \sum_{\forall i} \ln \left[\frac{d}{d\sigma_i} \left(1 - \exp \left\{ - \left(\frac{\sigma_i}{\sigma_0} \right)^m \right\} \right) \right] \quad \text{Equation 4-9}$$

The formulation of Equation 4-9 is based on the assumption that all observed failures are attributable to a single population of defects. Experimentally, this is often not the case. Suppose that two types of defects were observed to cause fracture, A and B. An analyst might choose to group the A samples together, and the B samples together, and analyse each set individually. It may seem logical to investigate the type A defects based only on group A samples, since it would be known that each group A sample failed because of a type A defect. However, there is a possibility that a type B defect was present in any or all of the group A samples, but did not cause failure. To be correct, the statistical approach should account for this possibility.

Suppose that group B samples generally have a higher failure stress than group A samples. It is entirely possible that unusually damaging group B defects could be masked by failure from a type A defect. As such, treating each group individually is likely to underestimate the variation in each set of defects to some extent.

The method presented in this work was also to perform two analyses, one for group A samples and one for group B samples. However, a modification to the likelihood equation was used to include group B (strictly, “Not A”) data in the group A analysis, and vice versa. Equation 4-10 expresses the modified likelihood relation for group A data.

$$L_A = \sum_{\forall i \in A} \ln \left[\frac{d}{d\sigma_i} \left(1 - \exp \left\{ - \left(\frac{\sigma_i}{\sigma_0} \right)^m \right\} \right) \right] + \sum_{\forall i \notin A} \ln \left[\exp \left\{ - \left(\frac{\sigma_i}{\sigma_0} \right)^m \right\} \right] \quad \text{Equation 4-10}$$

The additional term includes the probability that a group B sample survived failure from a type A defect. Analogously, this is the likelihood that a type A defect existed in each group B sample, but that the type A defect would have caused fracture at a higher stress than the type B defect which caused fracture.

To illustrate the effect of this term, suppose again that group A samples are generally weaker than group B, and that during the numerical maximisation of likelihood, a parameter combination is tried which under-estimates the strength of the group A samples. In this case, the likelihood that a sample would survive to fail from a type B defect would be too low, and so the second sum-term would act to increase the strength estimate for group A.

This approach was used to analyse the tensile data from the commercial HPDC experiments. The true stress at failure, σ^* , for each sample was used as σ_i in the above equations, where σ^* was defined as below, where ε^* is the engineering strain at failure:

$$\sigma^* = \text{UTS} \times [1 + \ln(1 + \varepsilon^*)] \quad \text{Equation 4-11}$$

4.3.2 MLE Extension for Correlation of Strength and Predicted Damage

As stated above, the MLE method can be used for fitting parameters of a statistical model to observed data. In pure terms, a statistical model is simply a function, which can be used to estimate the probability of an observation. Theoretically, any arbitrary function could be used for MLE, as long as it is possible to use that function to describe the likelihood of a set of observed events. The only practical requirement is that a maximisation algorithm can be successfully applied to the parameters of the resulting likelihood function.

As such, it is possible to add correlation parameters into a likelihood expression.

Suppose that the iterative fitting procedure is divided into an inner and an outer loop, and that

the outer loop handles the correlation parameters, where the inner loop finds the statistical parameter(s) to explain the variation that is not explained by the correlation. The lower the variation estimated by the inner loop, the greater the overall likelihood would be, because the probability density of each observation would be greater. This effect would drive the outer loop to find the correlation parameters that minimise the amount of residual variation.

In the present work, a power law relation was supposed between “strength” (true stress at fracture), and critical “damage” (particle number density), as defined in Equation 4-8. Using this expression, it is possible to link the observed mechanical strength to the statistical damage distributions for each test bar.

First, a finite element mesh representing the gauge length of each test bar was mapped into the domain of the fluid simulations, as illustrated in Figure 4-17. Each element was used as a starting point to estimate a statistical distribution of “particle number density” for that element, as described in section 3.2.5.




Figure 4-17: Illustration of virtual test bar gauge lengths used for damage correlation.




Figure 4-18: Parameter combinations “Random Walk Distance” and “Smoothing Radius” used.

The Fréchet damage statistics were mapped for each element in the gauge length of each test bar, based on each of the full-shot casting simulations, a total of 50 times. For each of these 50 runs, different values were used for the initial smoothing radius, used for calculating damage, and the random walk distance, used to sample this damage. The parameter combinations used are illustrated in Figure 4-18. As with the parameter selection for the parametric investigation of partial vacuum casting (section 4.2.2), “Latin square” sampling was used, and the determinant of the covariance matrix was maximised.

The data produced by each of these 50 statistical mapping runs was taken separately, to be passed into a Maximum Likelihood based correlation algorithm. The statistical data from

each run, covering each test-bar location in any or all of the full-shot casting experiments, is then correlated individually to the experimentally observed true fracture stress for each equivalent test bar.

For each i^{th} test bar the element-by-element damage distributions were combined to form a likelihood expression, Equation 4-12. In this expression, the subscript j refers to a property of an element, such as the volume of an element, V_j , or the Fréchet statistical parameters describing the probable damage in that element, $d_{0,j}$, β_j and μ_j . For correlation purposes, d is now a function of σ_i , in terms of A and B as defined by Equation 4-8. The remaining parameters are as defined in Equation 4-7, section 4.2.5.

$$L_i = \ln \left[\frac{d}{d d_i} \left(\exp \left\{ - \sum_j \left(\frac{V_j}{V_{eff}} \right) \left(\frac{d(\sigma_i) + d_{offset} - \mu_j}{d_{0,j}} \right)^{\beta_j} \right\} \right) \right] \quad \text{Equation 4-12}$$

Test bars where an entrainment defect was not deemed to have caused failure were treated with a modified likelihood expression, equivalent to the right hand term in Equation 4-10. The sum of these individual likelihood terms then becomes the total likelihood that each test bar failed at the observed stress, given that the probable fracture stress is described by the proposed function of “damage” (smoothed particle number density).

Overall, the correlation process has three levels, to find a total of 6 parameters. The innermost level is the solution of V_{eff} , which was achieved using a Newton-Raphson based function, given A , B and d_{offset} . This function is within a program, which uses a gradient ascent algorithm to seek the values for A , B and d_{offset} that maximise the overall likelihood of the observations, given all four of these parameters.

The final level of optimisation is to find which of the 50 parameter combinations for smoothing radius and random walk distance led to the greatest value of maximal likelihood, as calculated by the correlation program.

5 Results

5.1 Surface Area Code Robustness Characterisation

In the literature review, the Boolean Oxide Entrainment Code (BOEC), developed by Reilly [20], was identified as the most promising Oxide Film Entrainment Model available outside of the present work. However, this model was also criticised in the Literature Review (Section 2.2.3) for its insufficient robustness.

Since the Boolean model is based on a set of logical rules, which were empirically derived with reference to simple flow cases, it was hypothesised that this model could label a flow pattern as entraining when it is not, or fail to identify a flow which is entraining.

Section 3.1 in Theoretical Developments, presented a novel algorithm called the Surface Area Entrainment Code (SAEC), which was intended to solve each of these issues by using a more numerical approach.

To demonstrate these claims, a set of simulations are presented, which represent a variety of flow cases. Each of these models was run using the Surface Area Entrainment Code (SAEC). Additionally, an option called “particle ageing” in one of the user subroutines was activated, which changes the diameter of each particle over time, allowing the particles to be coloured based on their age.

For each of the simulation time-frames illustrated in this sub-section, a particle with a diameter of 0m (blue) has just been placed, and a particle with a diameter of 0.0005m (green) is 0.5s old. This was done to highlight where new particles are being created, and to provide a visual link between groups of particles that were created in a particular event.

In the accompanying graphs (Figures 5-2, 5-5 and 5-8), the particle count over time is presented alongside the fluid surface area over time, as calculated internally by FLOW-3D. The fluid surface area is provided as an independent measure of the severity of entrainment events, in particular for events where two volumes of fluid come together such that the surface between them is lost over a short time. These events should therefore show a step

jump in surface area and particle count. The magnitude of the jumps can be compared, and can also form a reference for other entrainment events.

In Figures 5-3, 5-6 and 5-9, these quantities are also graphed as time differentials, to better compare changes in the rate of particle placement and rate of surface area loss. Circular markers were added to each graph, at the positions relating to the illustrated time-frames of simulation, so as to put these frames into context.

The reader may find it useful to refer to Figure 4-12 in Section 4.2.1, where each of the flow cases in this section are illustrated diagrammatically.

5.1.1 Case 1: Falling Droplet

Figure 5-1 shows frames from a simulation, which begins with a “droplet” of fluid falling at 1ms^{-1} , into a static pool of fluid, as shown in frame A.

Between frames A and B, Figures 5-2 and 5-3 show a steep drop in the fluid surface area, as the droplet collides with the stationary fluid beneath it. Spikes in rate of surface area loss and particle placement rate can also be seen in Figure 5-3, showing that the surface area entrainment model are correctly identified this as an entrainment event.

Between frames A and B, the SAEC algorithm placed roughly 0.3 particles per mm^2 of global lost surface area, less than would be expected. However, Figure 5-1 B shows that the SAEC has correctly placed particles uniformly across the collision front, including at the edge, where the fluid collides at an oblique angle. This demonstrates the advantage of a flexible definition of surface collision; the Boolean entrainment criteria would not correctly treat a collision at an oblique angle.

In the circled region in frame C of Figure 5-1, there are two surfaces which come together at right angles, where the upper one moves to the left (towards the left hand surface), and the left hand surface flows upwards (towards to upper surface). This pattern would therefore satisfy one of the Boolean criteria for impinging flows. However, this is not an

impinging flow; the left hand surface is splash from the upper fluid stream, and moves upwards slower than the upper stream. The SAEC correctly detected this, with only a few particles shown within the circle, which were recently placed.

Between frames C and D, particle placement rate levelled off, as the fluid continued to travel up the walls and along the roof of the container. By frame E, this fluid has re-formed into a secondary set of droplets, which began to fall back towards the bulk of the fluid.

Between frames E and F there is another spike in the rate of surface area loss shown in Figure 5-3, as the first of these secondary droplets collides with the bulk of the fluid, and again between frames G and H when another secondary droplet collided with the bulk. Also between frames G and H, the bulk of the fluid collided with the fluid which remained at the right hand wall. In the time differential graph from the SAEC simulation (Figure 5-3), these events coincide with distinct spikes in particle production.

The total number particles produced during distinct surface collision events (A to B, E to F and G to H) represents 37.1% of the total particles produced by 0.4s. The initial collision immediately prior to frame B represents an isolated event, with no background signal; using this as a reference, the surface area code placed 5.9 times more particles by 0.4s of simulated time.

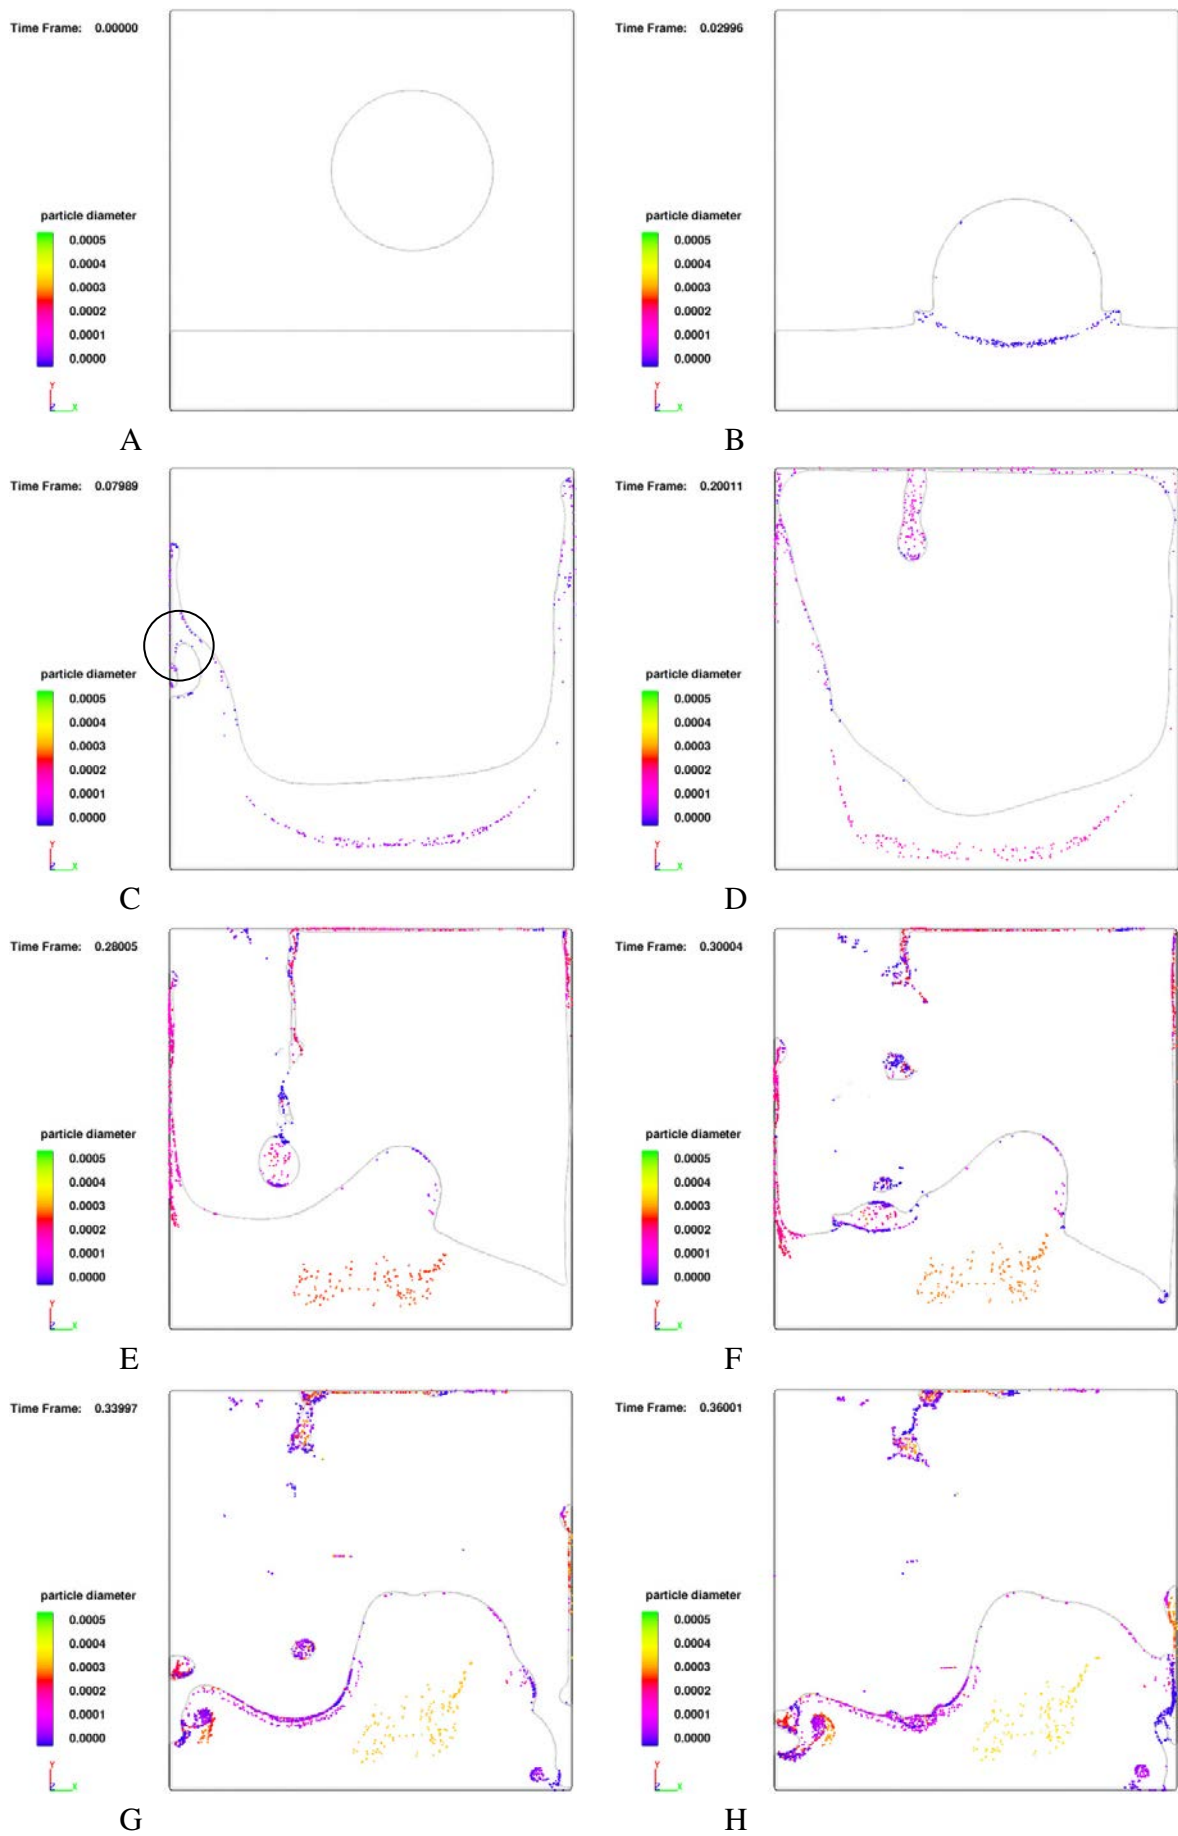


Figure 5-1 – Selected time-frames from the falling droplet simulation using SAEC.

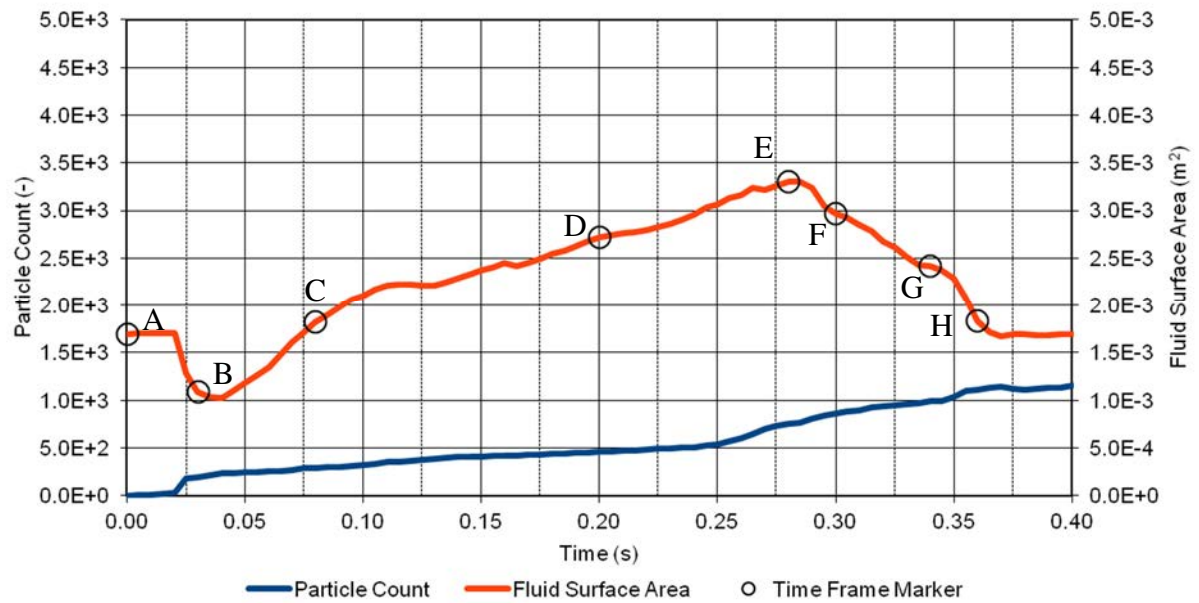


Figure 5-2: Fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames from Figure 5-1 are labelled for context.

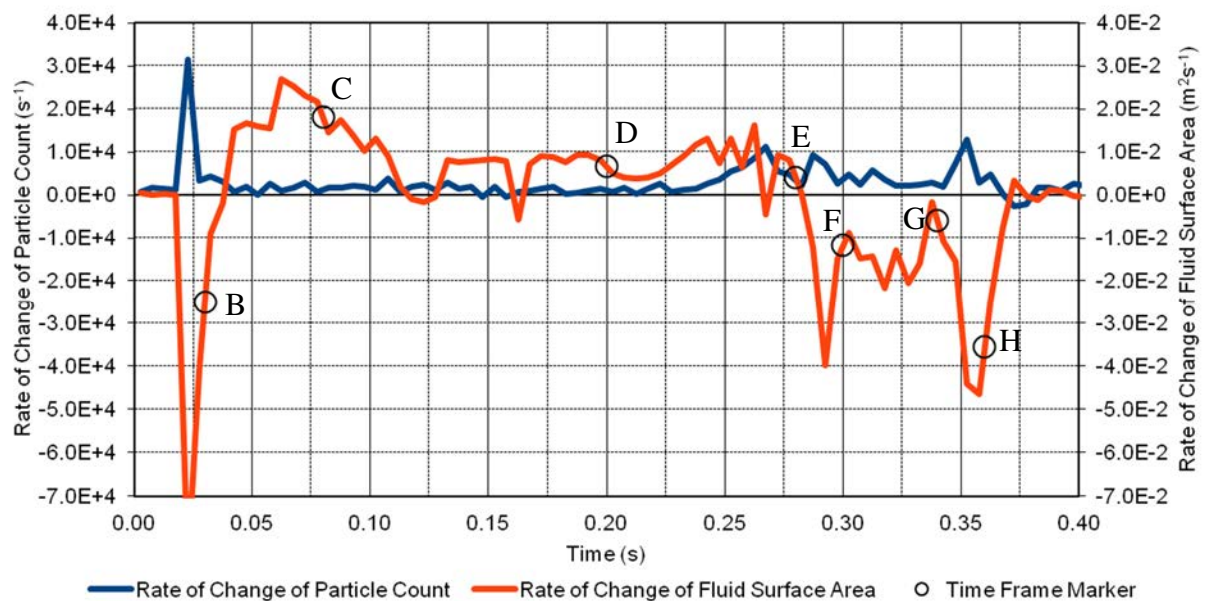


Figure 5-3: Time differential of fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames from Figure 5-1 are labelled for context.

5.1.2 Case 2: Top Gated Return Wave

In a “return wave” flow case, fluid enters the mould horizontally, and when it strikes the opposite wall, a wave is reflected that entrains air as it travels back across the fluid. The following simulation illustrates a slight variation, in which the fluid enters horizontally from the top of the cavity at 1ms^{-1} , and falls a short distance before continuing onward.

Figure 5-4 shows that the SAEC simulation places a significant number of particles on the lower surface of the stream from the start. This can be attributed to the two entrainment criteria for the surface area code:

Firstly, the stream is decelerating as it travels to the right; both because of friction with the mould wall, and also because gravity acts to increase the depth of the stream, meaning the fluid must slow to conserve volume. In the terms of the model, the surface of the slowing fluid is “lost”.

Secondly, the surface is defined by the model as “unstable”; because the fluid is above the void, there is a pressure gradient at the surface that should cause instabilities to grow. According to the models’ assumptions, this instability would cause any deviation in the surface – such as a wrinkle formed as the surface bunches up – to grow into a bubble and rise into the fluid as an entrainment defect. The particles in these simulations are modelled with negligible buoyancy however, and so are not taken into the bulk of the fluid.

As before, the circled region in frame B of

Figure 5-4, satisfies an impingement criterion for the Boolean code, since the left hand splash has rightward velocity, towards the point of “impingement”. In terms of relative velocity however, the flow is not entraining; the SAEC correctly identifies this, and does not place any new particles.

The particles which were lying on the surface in frame B have been taken beneath the surface by frame C, as the left hand jet rolls back on itself. This would imply there should be entraining conditions in the area circled in frame C. From the age of the particles in the

circled area it is apparent that since these “entraining” conditions formed, a small number were placed by the SAEC. Figure 5-6 does not show a significant change in the rate of particle placement by the surface area algorithm around the time of frame C.

As time progresses, particles continue to be placed at a slow rate in the region where the left hand jet rolls back on itself. By frame D, the right hand jet has rolled up around the right hand wall, and falls back towards the main stream. The entrainment algorithm has placed a significant number of particles on the leading edge of the splash, even though no loss of surface area is apparent.

Immediately prior to frame E, the right hand splash collides with the main falling stream; Figure 5-6 show a distinct spike in the rate of particle production coincident with this event.

Frame F shows a volume of fluid which has moved upwards from this collision, rolling against the flow of the main jet, creating a number of new blue particles in the circled region. Figure 5-6 also shows a spike in the particle production rate just before frame F.

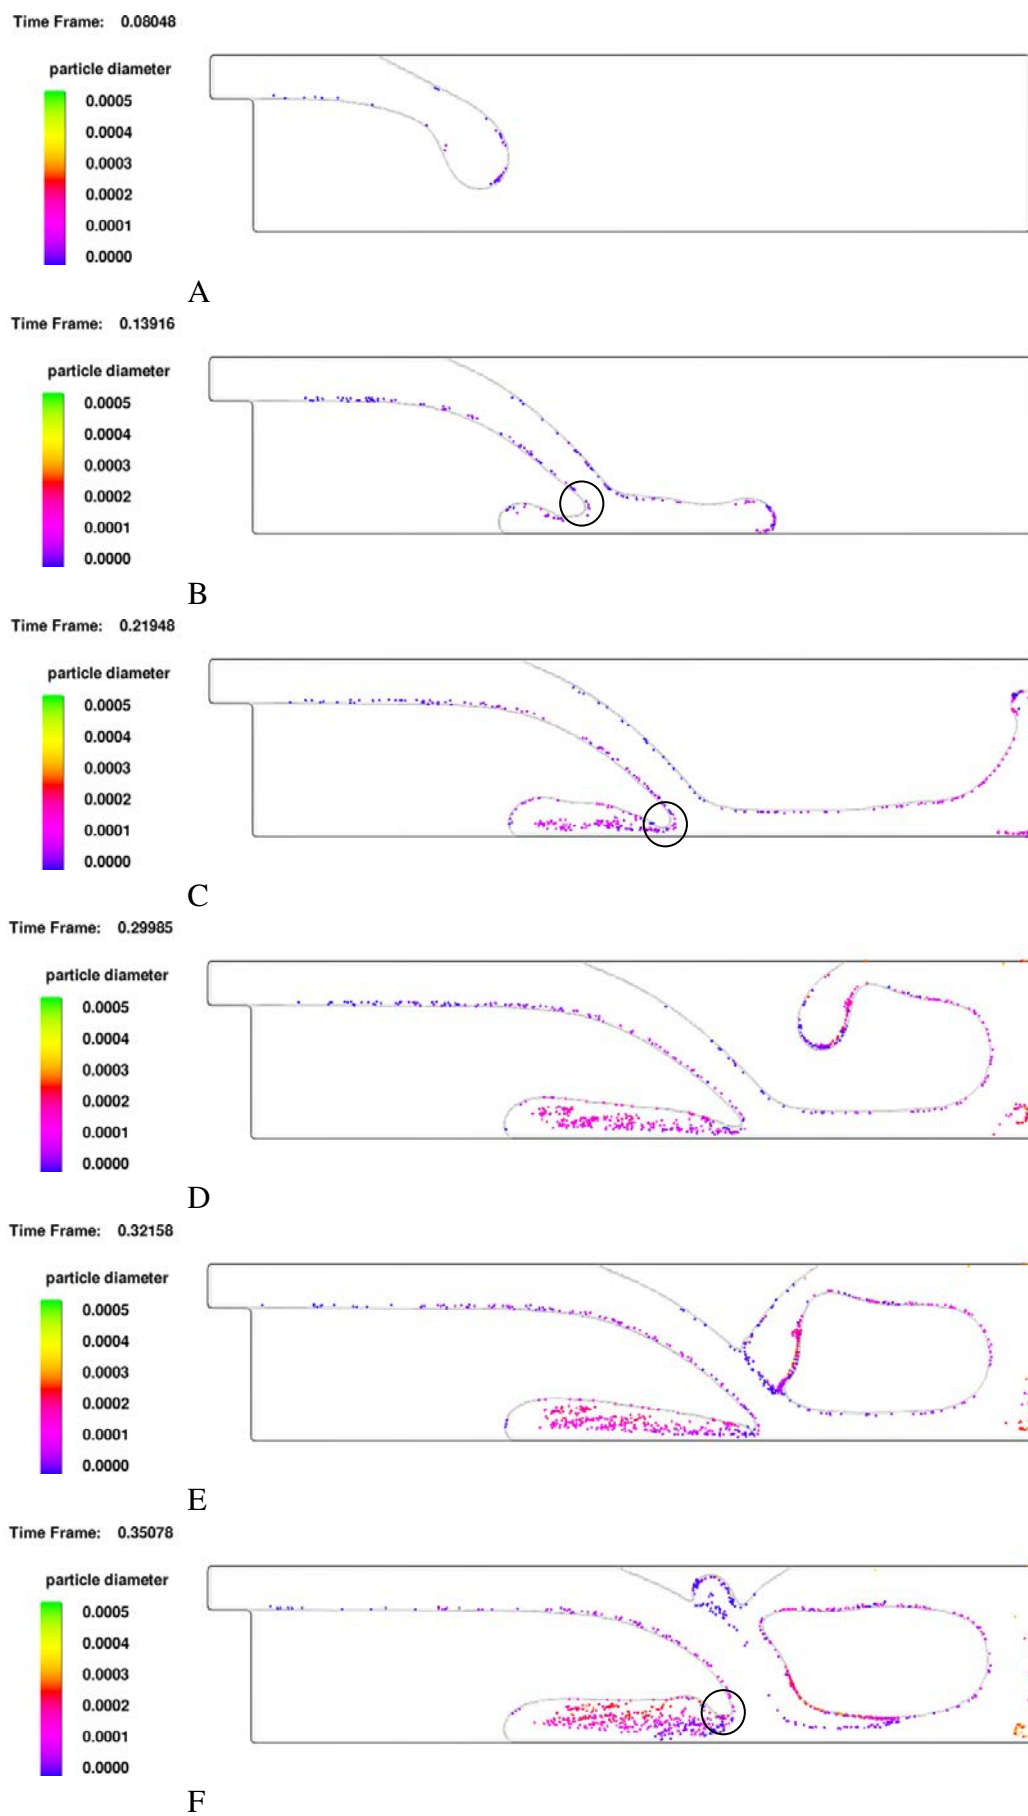


Figure 5-4: Selected time-frames from the top gated return wave using SAEC.

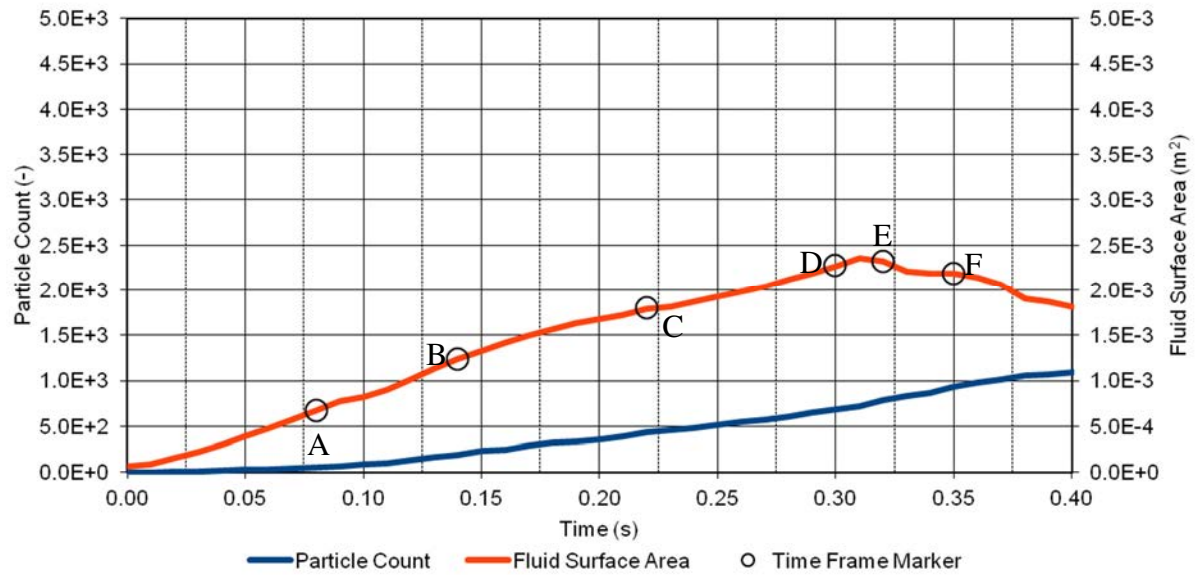


Figure 5-5: Fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames from

Figure 5-4 are labelled for context.

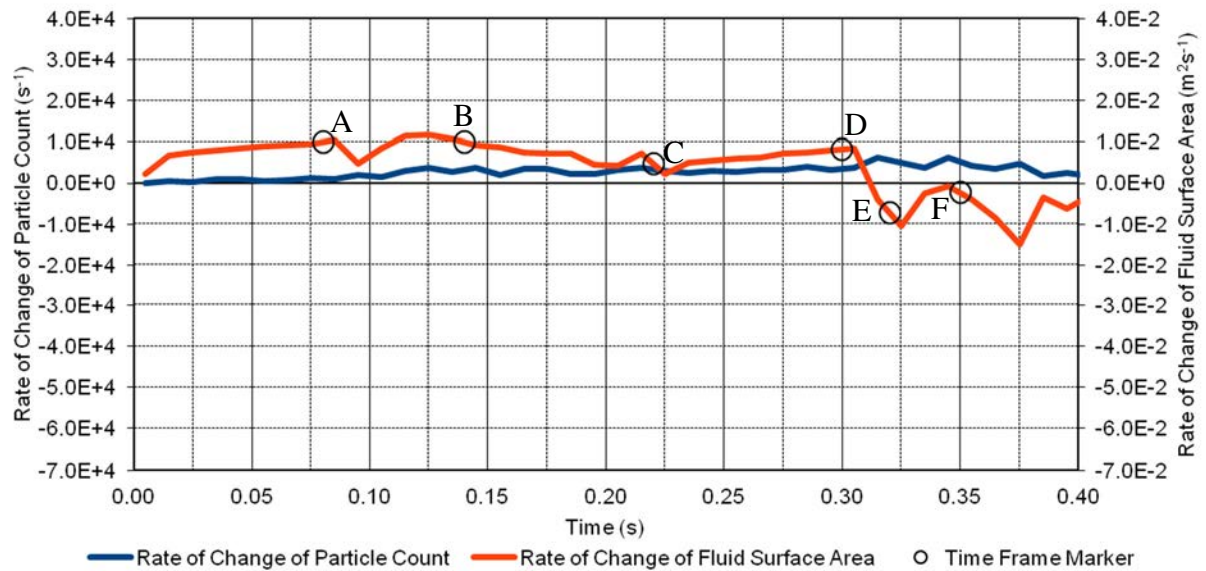


Figure 5-6: Time differential of fluid surface area and particle count for falling droplet simulation using SAEC. The positions of time frames

Figure 5-4 are labelled for context.

5.1.3 Case 3: Plunging Jet

The “plunging jet” flow case has been studied in prior works [20, 21], both experimentally and by simulation, and it is known to be highly entraining. In this scenario, a jet of fluid enters from the top of the mould cavity, and has enough energy to splash around the sides, before falling back, and impinging on the main jet. As such, it is a special case of top-gated filling.

As before, from the start of the simulation, the surface area algorithm placed particles on the surface of the fluid stream, with no obvious entrainment. Unlike the top gated return wave case, however, fluid accelerates as it falls under gravity, which should locally stretch the fluid surface. Thus, there should not have been any loss of surface area. It is possible that instantaneous small “losses” of surface area resulted from numerical noise, or as a by-product of the surface sharpening algorithm within FLOW-3D. Since “losses” are not reversible within the model, this would explain the placement of these particles as a systematic numerical error. If the average particle placement rate up to frame A is taken to be representative of the source of error (systematic or otherwise), and extrapolated to the duration of the simulation, then 711 particles, or 13.7% of the total, could be attributed to this error by 0.7s of simulated time.

Figure 5-7 shows the fluid stream falling and splashing against the bottom of the container (frame A), before rising up the sides (frame B). By frame B, the jets up the wall of the container have lost most of their upwards momentum. The model correctly identified entrainment in the circled area, where the right-hand jet rolled back on itself; Figure 5-9 shows a significant rise in particle placement rate prior to frame B.

The leftward fluid jet collided with the main jet in, as shown in frame C of Figure 5-7. In frame D, the rightward stream is shown falling to the base of the container, and colliding with itself. These were “collision” events, and Figure 5-9 shows a spike in the particle placement rate, which coincided with these events.

From its position in frame D, the right hand stream moved further to the left, as shown in frames E and F. Frame E shows the right hand stream impinging on itself very near the base of the main jet, creating what appears to be a double impingement event (circled), which may be assumed to be highly entraining. Figure 5-9 shows a distinct increase in the particle placement rate for the SAEC near the time of frame E.

Towards frame F, the point where the right hand stream is shown to fall back on the main stream moves upwards; the entrainment mechanism part impinging and part collision, and so is a non-intuitive mode, which may not be properly represented by boolean criteria. The surface area model correctly identifies the flow as highly entraining, and Figure 5-9 shows a significant spike in particle placement approximately 0.02 seconds before frame F.

Frame G shows the right hand void collapsing. During this collapse, which continued past frame H, the oxide film entrainment model placed a large number of particles; Figure 5-9 shows the particle placement rate peaked immediately prior to frame H, at the highest rate in the simulation. Frame H shows that at that point, the right hand volume of fluid came together with the main plunging jet, with their surfaces roughly parallel, which would explain the spike in particle placement.

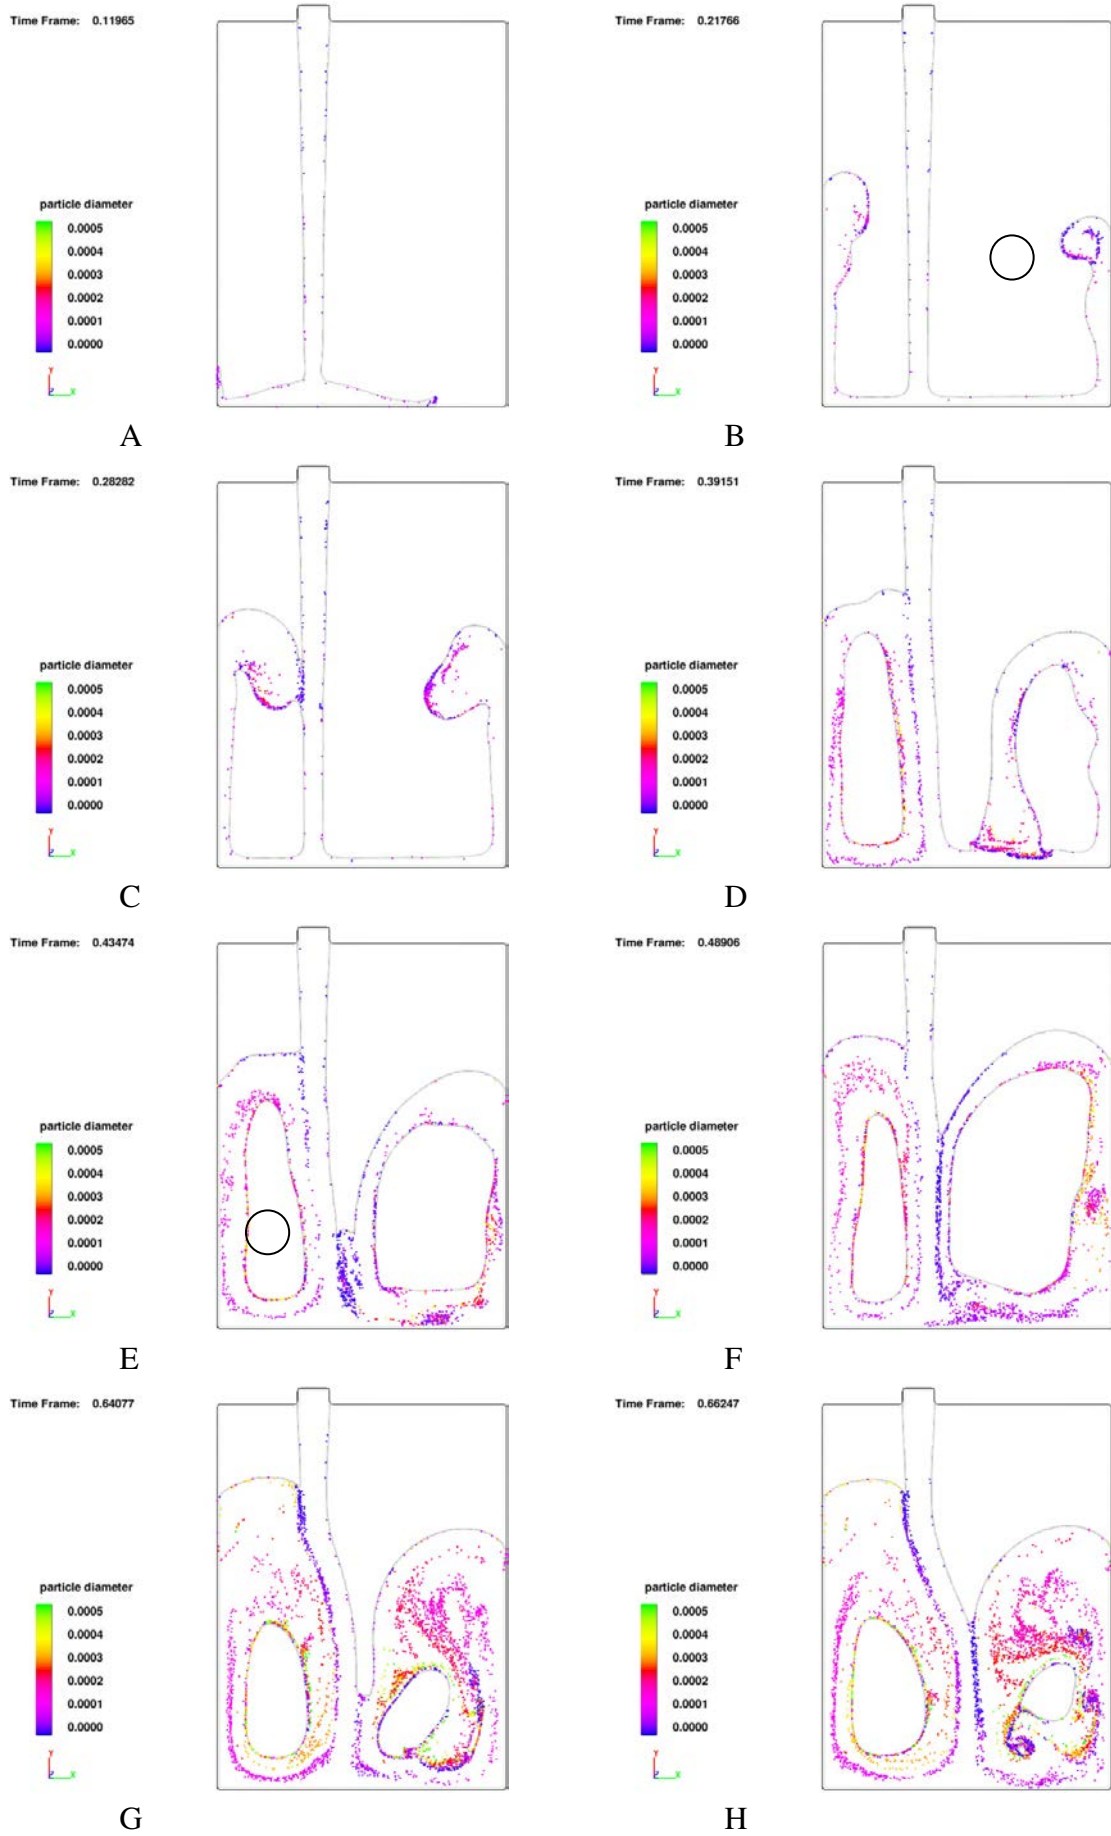


Figure 5-7 – Selected time-frames from the plunging jet simulation using SAEC.




Figure 5-8: Fluid surface area and particle count for plunging jet simulation using SAEC. The positions of time frames from Figure 5-7 are labelled for context.




Figure 5-9: Time differential of fluid surface area and particle count for plunging jet simulation using SAEC. The positions of time frames from Figure 5-7 are labelled for context.

5.2 Controlled Partial Vacuum Casting

5.2.1 Parametric Investigation

To complement the controlled vacuum casting experiments, and the simulations of those experiments, a parametric investigation was run, using a simplified model of the process, as described in the Method (Section 4.2.3).

In this investigation, the mesh size was varied between 0.65 mm and 1.25 mm, and the rate of change of pressure was varied such as to set the nominal flow speed between 0.25 ms⁻¹ and 1 ms⁻¹ for each simulation. The response variable was the number of particles placed by the surface area entrainment prediction algorithm, at the end of each simulation.

On the following page, Figures 5-10 and 5-11 show that the particle count – and by proxy the predicted area of entrained surface film – is sensitive to both mesh size, and the target flow velocity.

There is a lot of “noise” in both responses, however to some extent the noise in each plot can be explained by the influence of the un-plotted factor; that is to say that the influence of velocity appears as noise in the plot against mesh size, and vice versa.

Because of this, it is difficult to draw conclusions with confidence, except to say that the particle count increases with velocity, and decreases with mesh size. However, Figure 5-11 seems to show a roughly linear response of particle count to mesh size, when plotted on logarithmic axes. This would suggest that the effect of mesh size could be modelled using a power law.

Using that assumption, it is possible to try to minimise the noise (error from moving average) in the velocity plot, by subtracting a power law function of mesh size. The terms of this problem are defined as follows:

$$Noise = \sum \left[Y_i - \left(\frac{Y_{i-1} + Y_{i+1}}{2} \right) \right]^2 \quad \text{Eq. 5-1}$$

$$Y_i = \ln(P_i) + k(\ln(m_i) - \overline{\ln(m)}) \quad \text{Eq. 5-2}$$

Here, Y_i is the logarithm of the corrected particle count, m is the mesh size, and P_i is the particle count from simulation. When Eq. 5-2 was substituted into Eq. 5-1, and differentiated, an expression was found in k , to analytically minimise *Noise*. Using this expression, k was found to be -3.33. This means that if the linear mesh size is halved, the particle count can be expected to increase by roughly 10 times for a comparable simulation.




Figure 5-10: Response of particle count to changes in velocity.




Figure 5-11: Response of particle count to changes in mesh size.




Figure 5-12: “Corrected” velocity response, assuming power law effect from mesh size.

The response of corrected particle count, $\exp(Y_i)$, to velocity is graphed in Figure 5-12, where Y_i is as defined in Eq. 5-2. This analysis shows that when the influence of mesh size is subtracted for these data, the particle count stays relatively constant with nominal flow velocities below 0.5 ms^{-1} , but between 0.5 ms^{-1} and 1 ms^{-1} , the particle count increases by approximately 1 order of magnitude.

The investigation also provides an opportunity to qualitatively investigate the response of flow shape to increasing flow velocity. This will put the above results into a qualitative context, and help with the interpretation of experimental observations.

Figure 5-13 shows one frame with equivalent fill fraction from each of the parametric simulations, ordered by nominal flow velocity. Each frame is rendered using a custom post-processor, which can colour the fluid by smoothed particle number density; in this case red corresponds to $0.5 \text{ particles} \cdot \text{mm}^{-3}$.

Up to and including around 0.48 ms^{-1} , the flow front appears relatively quiescent at the selected fill fraction, although these flows are far from being absolutely stable. Some of these simulations show particle concentrations towards the centre of the flow front, presumably formed as the fluid splashed against the sides of the mould.

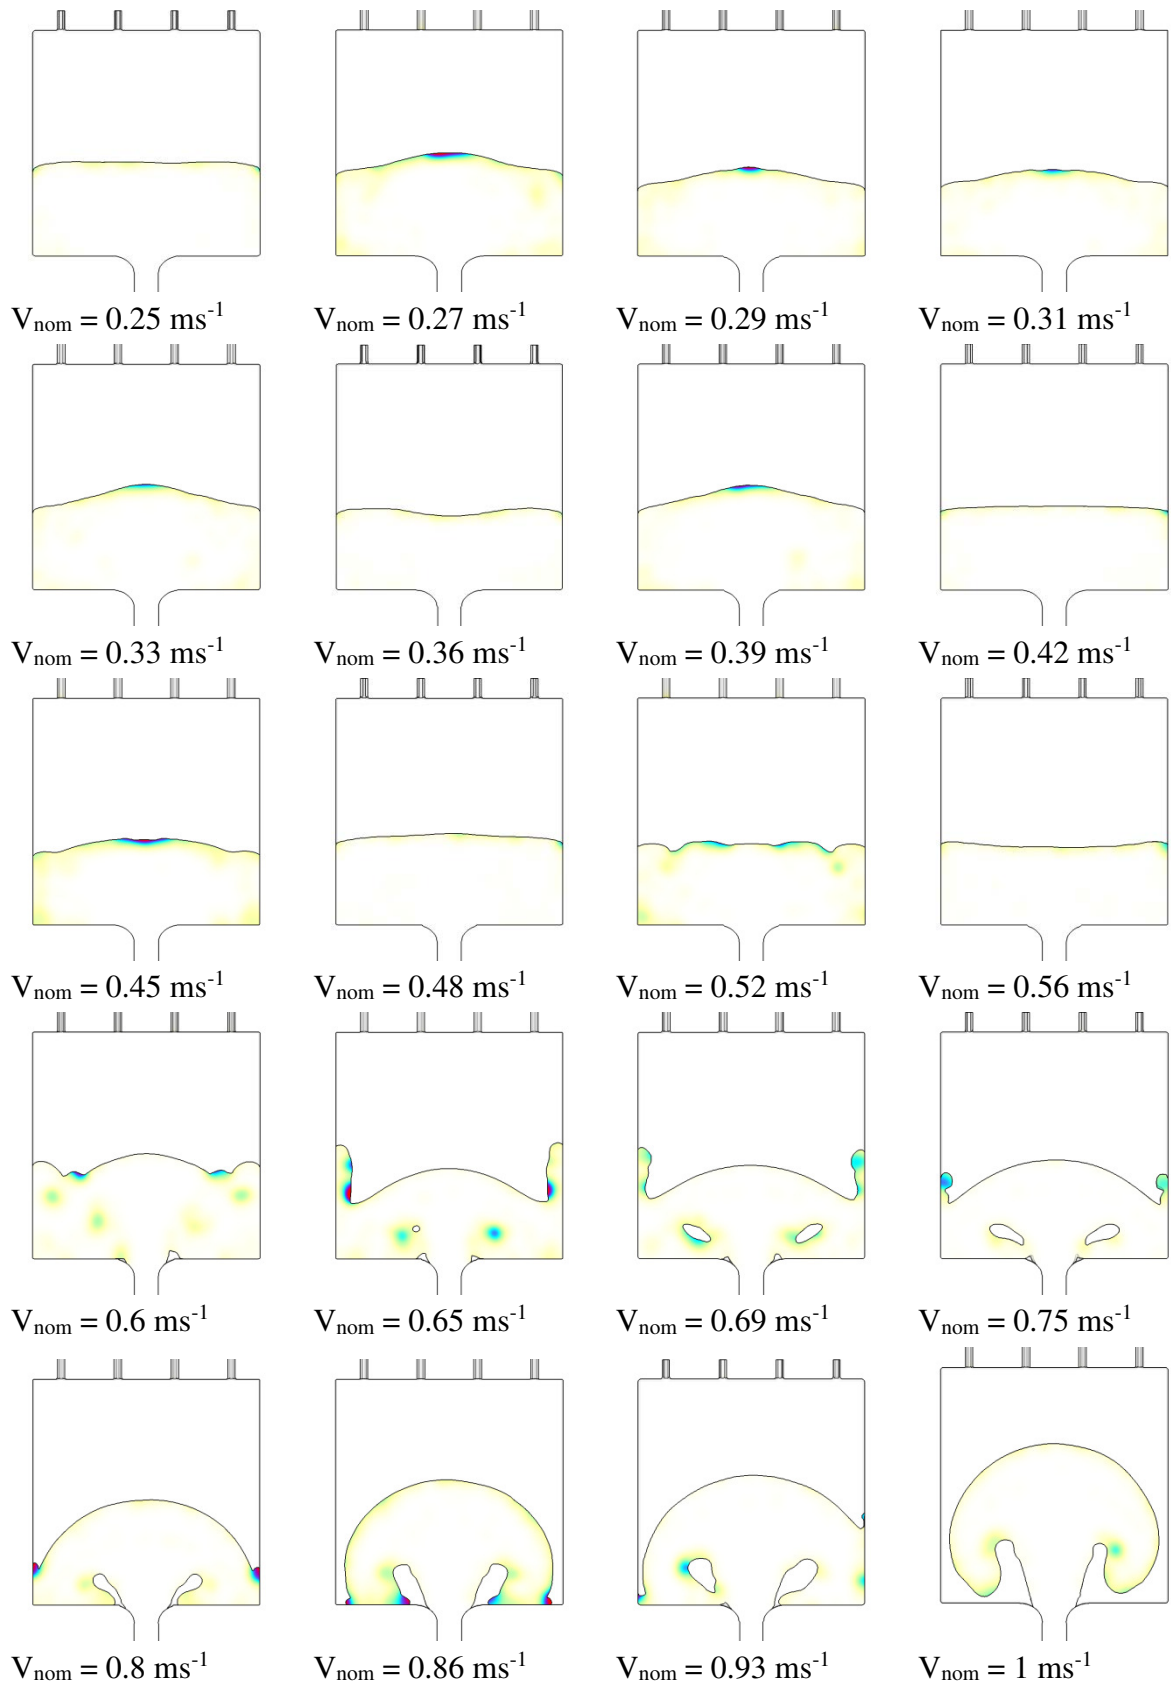


Figure 5-13: Each of the 20 simulations, at the same fill fraction, demonstrating the effect of velocity on flow morphology.

From 0.52 ms^{-1} to 1 ms^{-1} , the flow front is shown to evolve, as the vertical momentum of the flow increasingly overcomes gravitational acceleration. With increased velocity, the two “arms” of the main flow – most clearly seen in the 1 ms^{-1} case – fall farther from the gate. When these “arms” splash against the mould wall, a component of each splash travels towards the gate, forming a vortex, with a void at its core. These voids grow with increased nominal velocity.

5.2.2 Time-Value Data from Experiments

When the controlled partial vacuum casting experiments were performed, an Arduino DUE microcontroller was used to provide closed loop feedback control, to regulate the pressure in the partial vacuum chamber according to defined set-points. Details of the experimental set up were given in the Method (Section 4.1.2).

Inherent in closed loop feedback is the continuous measurement of the controlled variable, in this case pressure. The pressure curves were recorded at a rate of 100 samples/second (after signal processing), and are presented in this section.

In each of the following figures, the pressure curve from experiment is presented alongside the target pressure curve for each experiment. Each experimental run is referred to by the intended velocity, as prescribed by its target pressure set-points.

It is not clear why the recorded chamber pressure started decreasing at some time after the 10 second target, for the majority of these tests (with the exception of “Slow”, Figure 5-14). It does not appear to be a data synchronisation issue, because Figure 5-15 shows the recorded pressure curve transition to a negative gradient at 13 seconds after program start, when the valves are partially closed, and the target pressure steps down.

It is possible that the recorded values are too low, and that values which would have been negative have been truncated to 0.




Figure 5-14: Recorded pressure-time curve for the “Slow” controlled partial vacuum casting experiment.




Figure 5-15: Recorded pressure-time curve for the “Slow TC” controlled partial vacuum casting experiment, which was the only test instrumented with thermocouples.




Figure 5-16: Recorded pressure-time curve for the “Medium” controlled partial vacuum casting experiment.




Figure 5-17: Recorded pressure-time curve for the “Medium Fast” controlled partial vacuum casting experiment.




Figure 5-18: Recorded pressure-time curve for the “Fast” controlled partial vacuum casting experiment.





Figure 5-19: Recorded pressure-time curve for the “Dump” controlled partial vacuum casting experiment.


For the “Dump” experiment, the time-offset between the target pressure curve and the actual pressure curve can be attributed to the time take to actuate the valves. Unlike the other tests, there was no “pre-load” to allow the valve to open exactly on time, and the closed loop control was effectively overridden between 10 and 12 seconds after program start. At 10 seconds the vacuum valve started moving from fully closed to fully open, and at 12 seconds, the program parameters revert to closed loop control, and combined valve openness is set to 35%.

The thermocouple data from the “Slow TC” experiment is presented in Figure 5-20. These data show that as the mould filled, the thermocouples registered a maximum temperature of 332°C at the gate. This result is extremely unlikely to represent the real metal temperature at these times; these temperatures are approximately 200°C below the solidus temperature for the LM25 aluminium alloy used in these experiments.

Figure 5-20 B shows that the recorded temperature signal was overlaid with what appears to be a regular pattern of pulses in both signals, which would indicate some systemic error in the measurement system. At the start of the first step, the recorded temperature at the gate goes up from 271°C to 350°C. If the “high” level of these pulses were assumed to represent to “true” signal, and if 350/271 were taken as a multiplying factor to obtain the “true” curve, this would make the peak extrapolated temperature during filling 429°C. Since, this is still highly unlikely to be the true maximum temperature reached at the gate, this extrapolation would be invalid. As such, only the relative shapes of these curves, rather than their absolute values, are useful for analysis.



A



B

Figure 5-20: Thermocouple data recorded in the “Slow TC” experiment, with the abscissa scaled to show the filling of the mould (A), and with the abscissa scaled to show all recorded data (B).


5.2.3 Comparative Frames from Experiment and Simulation


While many of these casting experiments did not perform as expected, a good variety of pressure-time curves were recorded. As such, the real-time X-ray footage showed the response of the fluid flow pattern to a range of boundary conditions.

Moreover, the fact that the boundary conditions were recorded in real time, allowed each test to be simulated using that data. This negated the need to make assumptions about the boundary conditions, except for the assumption that the recorded pressure data was accurate. This provided an opportunity to test the fidelity of each of these simulations to the observed fluid flow shapes which arose from these different pressure histories.


In the following frames, the time (t_{norm}) has been normalised such that the metal first enters the mould cavity at $t_{\text{norm}}=0$. Figures 5-21, 5-23, 5-25, 5-27, 5-29, and 5-31, show example time frames from the real time X-ray footage of the experiments, which have been post processed to enhance the contrast between the liquid metal and the mould cavity. It should be noted that there is an exposure time; when the fluid moves during this exposure time, the image will experience motion blur.

Figures 5-22, 5-24, 5-26, 5-28, 5-30 and 5-32, show equivalent time frames from the simulation of these experiments. These have been rendered using a custom post processor, which was set to shade each pixel based on the thickness of fluid, mimicking the x-ray imaging process.


The first of the illustrated frames in Figures 5-21 and 5-22 show good agreement between the experiment and simulation. However, for the later time frames, the simulated fluid shape increasingly deviated from the experiment. The final illustrated time frames in each figure show that the  shaped flow structure persisted longer in the experiment, than in the simulation of that experiment.



$t_{\text{norm}}=0.22 \text{ s}$




$t_{\text{norm}}=0.32 \text{ s}$




$t_{\text{norm}}=0.55 \text{ s}$


Figure 5-21: Example time frames from the “Slow” controlled partial vacuum experiment



$t_{\text{norm}}=0.22 \text{ s}$




$t_{\text{norm}}=0.32 \text{ s}$




$t_{\text{norm}}=0.56 \text{ s}$


Figure 5-22: Equivalent time frames from the simulation of the “Slow” experiment.



$t_{\text{norm}}=0.15 \text{ s}$




$t_{\text{norm}}=0.25 \text{ s}$




$t_{\text{norm}}=0.52 \text{ s}$


Figure 5-23: Example time frames from the “Slow TC” controlled partial vacuum experiment.



$t_{\text{norm}}=0.14 \text{ s}$




$t_{\text{norm}}=0.24 \text{ s}$




$t_{\text{norm}}=0.52 \text{ s}$


Figure 5-24: Equivalent time frames from the simulation of the “Slow TC” experiment.



$t_{\text{norm}}=0.08 \text{ s}$




$t_{\text{norm}}=0.23 \text{ s}$




$t_{\text{norm}}=0.28 \text{ s}$


Figure 5-25: Example time frames from the “Medium” controlled partial vacuum experiment



$t_{\text{norm}}=0.08 \text{ s}$



$t_{\text{norm}}=0.24 \text{ s}$



$t_{\text{norm}}=0.28 \text{ s}$

Figure 5-26: Equivalent time frames from the simulation of the “Medium” experiment.


Figures 5-23 and 5-24 compare the “Slow TC” experiment and its simulation. It is apparent that the fluid velocities are higher in the experiment; the mould filled more over the same time, and the flow appears to have more momentum. A bubble, which has formed at the tip of the gate thermocouple, is visible in the final frame illustrated in Figure 5-23. The thermocouple can be seen to extend from this point, down and to the right. This thermocouple seems to have deflected the fluid flow towards the left hand side of the image. This is of course not reproduced in simulation.

The simulated thermocouple positioned immediately after the gate in the “Slow TC” simulation recorded that the temperature peaked twice, first at 695°C as the metal first entered the mould, and then again at 717°C 1.38 seconds later, during the final stages of filling, once the dip tube was extracting less heat from the incoming metal.


The “Medium” case is illustrated by Figures 5-25 and 5-26. Again, the flow in the experiment seems to have significantly more kinetic energy than the simulation, which was based on the pressure data recorded in that experiment. Interestingly, Figure 5-25 shows that the upper edge of the metal remained remarkably persistent between $t_{\text{norm}}=0.23$ and $t_{\text{norm}}=0.28$, as the fluid fell back down; it seems as if there was a solid structure holding some of the edges of the flow in place.

In Figures, 5-28 and 5-27, it is again apparent that the fluid had significantly more energy as it filled the mould in the “Medium Fast” experiment, than was reproduced in simulation.


Comparing Figures 5-29 and 5-30, it would seem not only that the simulation under-represented the vertical momentum of the fluid as it entered the mould, but also that in the experiment the gauge pressure in the chamber was of sufficient magnitude to hold fluid in the mould for longer. In the experiment the fluid did not finish draining out until $t_{\text{norm}}=1.63$, whereas in its simulation the fluid had finished draining by $t_{\text{norm}}=0.54$.



$t_{\text{norm}}=0.1 \text{ s}$




$t_{\text{norm}}=0.28 \text{ s}$




$t_{\text{norm}}=0.38 \text{ s}$


Figure 5-27: Example time frames from the “Medium Fast” controlled partial vacuum experiment



$t_{\text{norm}}=0.1 \text{ s}$




$t_{\text{norm}}=0.28 \text{ s}$




$t_{\text{norm}}=0.38 \text{ s}$


Figure 5-28: Equivalent time frames from the simulation of the “Medium Fast” experiment.



$t_{\text{norm}}=0.08 \text{ s}$




$t_{\text{norm}}=0.3 \text{ s}$




$t_{\text{norm}}=1.63 \text{ s}$


Figure 5-29: Example time frames from the “Fast” controlled partial vacuum experiment



$t_{\text{norm}}=0.08 \text{ s}$




$t_{\text{norm}}=0.3 \text{ s}$




$t_{\text{norm}}=0.54 \text{ s}$


Figure 5-30: Equivalent time frames from the simulation of the “Fast” experiment.



$t_{\text{norm}}=0.08 \text{ s}$




$t_{\text{norm}}=0.22 \text{ s}$



$t_{\text{norm}}=0.32 \text{ s}$

Figure 5-31: Example time frames from the “Dump” controlled partial vacuum experiment



$t_{\text{norm}}=0.08 \text{ s}$



$t_{\text{norm}}=0.22 \text{ s}$



$t_{\text{norm}}=0.32 \text{ s}$

Figure 5-32: Equivalent time frames from the simulation of the “Dump” experiment.

Figures 5-25, 5-27, 5-29 and 5-31 show that when the fluid was jetted upwards during the initial stages of filling for these experiments, a V-shaped structure formed which was thin in the x-ray direction across most of its section, but thick at its edges. The top of this structure is also thick. One hypothesis would be that is a sort of vertical hydraulic jump, where the thick upper section is supported by a thin high velocity jet.

As shown in Figure 5-32, the simulation better reproduced the flow patterns observed in experiment for the “Dump” case. In the first of the frames shown from Figures 5-31 and 5-32, the V-shaped “vertical hydraulic jump” feature is seen, and the angle of divergence in the jet is quite similar between both of these frames. However, in the simulation, the fluid formed a wide thick channel on the right hand side, as opposed to the more symmetrical structure observed experimentally. In addition, the simulation showed horizontally banded ripples. The apparent non-existence of these bands in the experimental observations is inconclusive; if they did occur, they would have been obscured by motion blur.

The peak volume flow rate into the mould region of the simulation was $3.57 \times 10^{-4} \text{ m}^3 \text{ s}^{-1}$, corresponding to an average velocity of 1.76 ms^{-1} over the cross section. The peak vertical component of the flow velocity immediately after the gate was 2.9 ms^{-1} .

The intermediate frames from both of these figures show that the fluid reached the bottom of the mould at a very similar time, in the simulation and in the experiment that is was based. However, in the final frame shown in Figure 5-32, the residual bubbles are substantially larger than those shown in the equivalent frame of Figure 5-31. It is possible that in the experiment the fluid had more energy, and so was able to compress the bubbles further. Additionally, in the experiment, the gas in the bubbles may have been partially consumed in an oxidation reaction.

It is unlikely that the gas escaped into the mould wall, as it is believed that the castings formed a solid skin with bubbles travelling within this skin. One of the castings (Medium Fast) had voids towards the top, although casting would appear solid at first glance.

Furthermore, the “Fast” casting, which partially drained out, did form a full square shape; it seemed that only the inner fluid drained out, leaving a solid deflated skin.

Other observations of the castings themselves, which may be of later relevance, are that there was significant sand penetration, and that there was some flash formation, even in the “Fast” experiment. Together, these suggest that there was more than enough superheat in the casting experiments. If the flash formation occurred because the forces in the mould were sufficient to force the two halves of the mould apart, then this may help explain some of the disparity between the experimental observations than their simulations, since the simulations assumed un-deformed mould geometry.

5.3 Commercial High Pressure Die Casting

5.3.1 Short Shot Investigation

5.3.1.1 Experimental

Figure 5-33 shows the short shot case that was selected for simulation. In this shot, 1.9kg of aluminium were used, with a shot sleeve fill fraction of 0.137. This fill fraction was selected, because it was one of the lowest to result in a casting which was apparently full. This made it more correct to use the patterns of solidified metal, which were observed using lower fill fractions, to inform the parameters used in the simulation of this casting.






Figure 5-33: Photographs of the feed side (A) and ejector side (B) of the short shot casting chosen for simulation.

While the above castings appear full, Figure 5-34 shows that this is not the case. In fact, there are several voids on the scale of 2 cm, and many smaller bubbles on with a diameter on the order of 2 mm. These smaller bubbles seem to be more prevalent in the horizontal centre of the casting, as pictured, whereas the larger voids generally lie around the edge of this region.



Left



Right

Figure 5-34: X-ray photomontages of the left and right hand side (as shown in Figure 5-33 A) of the short shot casting chosen for simulation.

5.3.1.2 Simulation and Mesh Sensitivity Analysis

As described in the Method, section 4.2.2.1, the simulation of the 1.9 kg short shot was performed using three different mesh sizes, nominally 0.5 mm, 0.8 mm and 1 mm. The 0.8 mm mesh was used in the selection of model parameters, and the 0.5 mm and 1 mm meshes were used to evaluate the sensitivity of the model to finer and coarser meshes respectively. Each simulation was terminated after 2.2 seconds of modelled time, from the start of piston motion in the shot sleeve.

Figures 5-35, 5-36 and 5-37 show the final frame of the simulations with the 0.5 mm, 0.8 mm and 1 mm nominal mesh sizes, respectively. The mesh size appears to have affected the tendency of the simulated fluid to go towards the inner or outer edge of the mould cavity.

The banding visible in these figures is an artefact of the rendering process. The colour is assigned based on the solid fraction at the centre of the cell, and since some of the “outer” cells’ centres are further into the casting than others, this results in the striped effect.

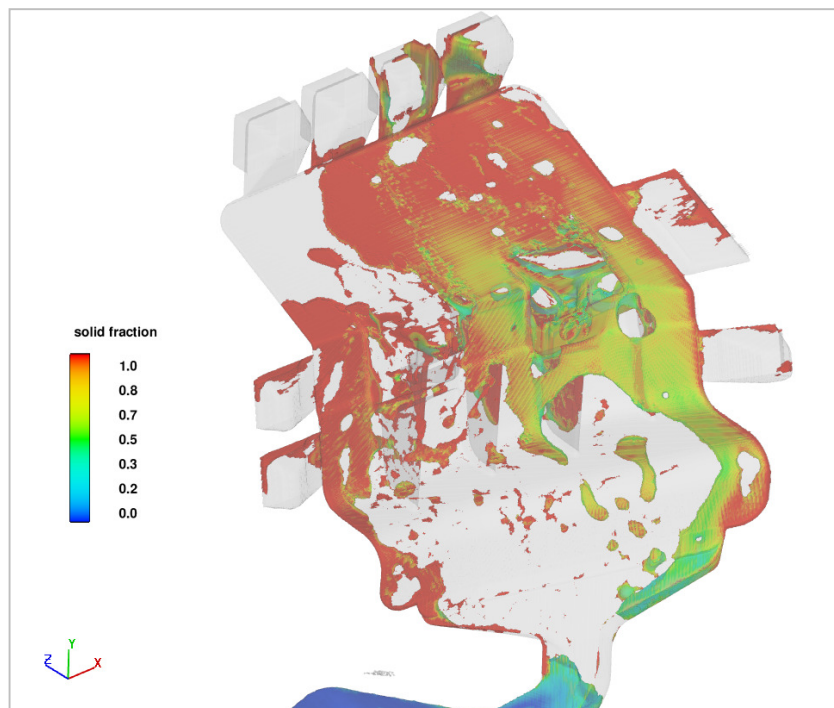


Figure 5-35: The final frame of the short shot simulation, which was performed using the fine mesh, coloured by solid fraction.

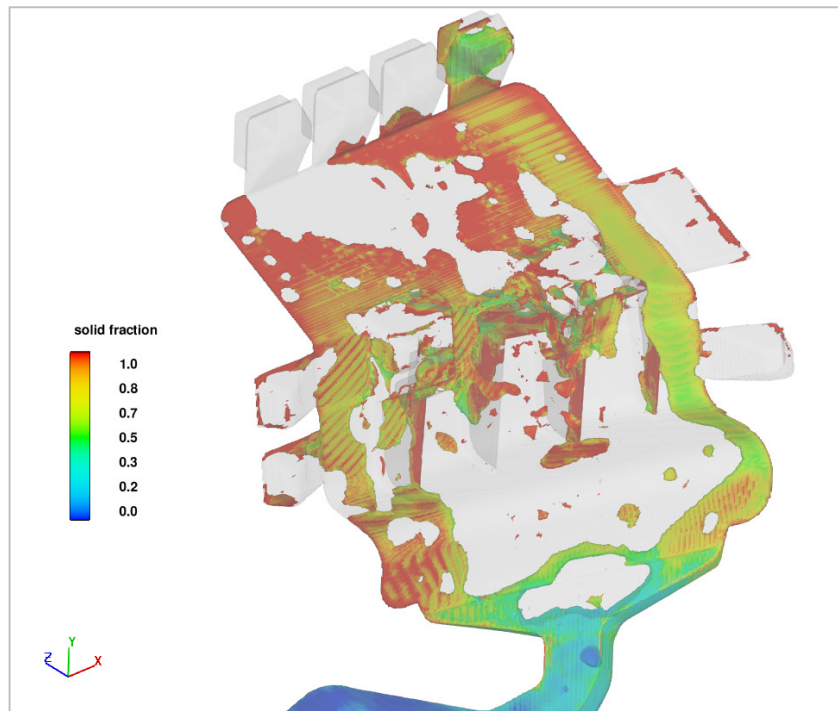


Figure 5-36: The final frame of the short shot simulation, which was performed using the intermediate mesh, coloured by solid fraction.

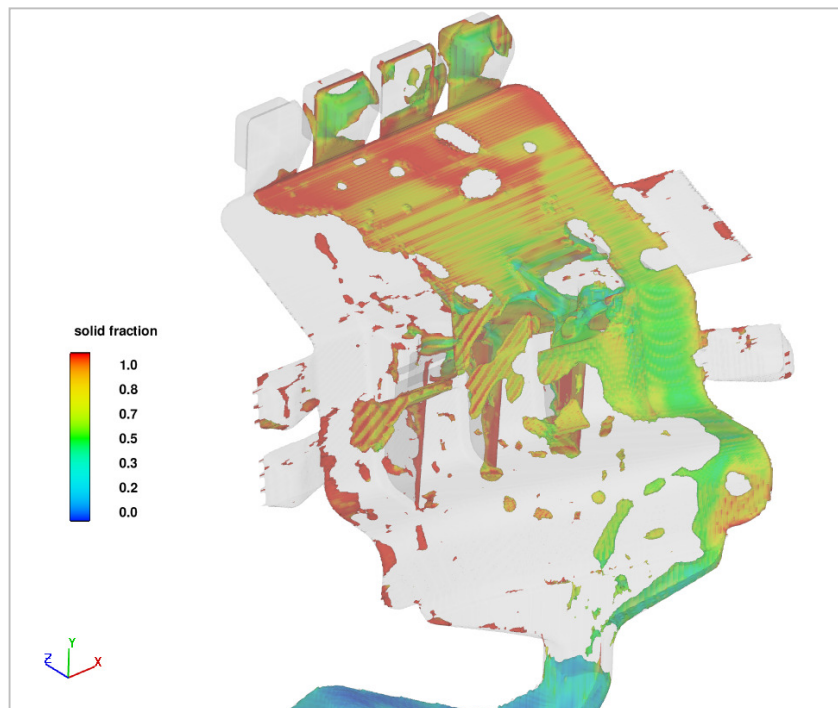


Figure 5-37: The final frame of the short shot simulation, which was performed using the coarse mesh, coloured by solid fraction.

Similarly, it is not clear from these images whether the region farthest from the gate was more solidified at the end of fine mesh simulation than it was in the others; it may simply be that with a finer mesh, the outer cells are closer to the mould wall.

The results shown in the above figures show a qualitative response to changes in the mesh size, but it is not clear from these whether the effect was the result of inherent mesh sensitivity in the CFD solver, or whether the response was to something small and random, but amplified by the chaotic nature of fluid flow.

The following graphs (Figures 5-38, 5-39 and 5-40) show the time history of some key variables, which were extracted from each simulation. Figures 5-38 and 5-39 demonstrate the changes in specific kinetic energy and fluid thermal energy over time, respectively. Both of these quantities might be expected to be mesh sensitive in some cases, since the mesh would affect the resolution of the turbulent boundary layer, which would in turn affect the transfer of heat and momentum from the fluid to the mould.




Figure 5-38: Specific fluid kinetic energy over time, for each of the three short-shot simulations.




Figure 5-39: Specific fluid kinetic energy over time, for each of the three short-shot simulations.




Figure 5-40: Specific fluid kinetic energy over time, for each of the three short-shot simulations.

However, for these simulations, the mesh size does not seem to strongly affect the rate of dissipation of thermal or kinetic energy, although there is a small difference in kinetic energy at the start of the coarse mesh simulation.

Figure 5-40, however, shows that the surface area of the simulated fluid was affected by mesh size, particularly for the fine mesh. This is likely to be because smaller features in the fluid flow can be resolved by a finer mesh, and in particular, there can be smaller bubbles and droplets, which would have a larger combined surface area. It is factors such as this, which may explain how the qualitative differences arise between the simulations, as shown in Figures 5-35, 5-36 and 5-37.

When selecting the mesh size, it is also important to consider whether the mesh will result in an unacceptably long simulation time. The effect of nominal mesh size on simulation time is tabulated below. Each simulation was run on the same computer, with a 3.4 GHz Intel i7 Quad-Core processor, with hyper-threading disabled.

Table 5-1: Effect of nominal mesh size on simulation run time.

Nominal Mesh Size (mm)	Total cells used for Flow Calculations	Run Time
1	1,140,200	12 hours 34 minutes
0.8	2,211,749	3 days 4 hours 59 minutes
0.5	8,776,074	26 days 5 hours 59 minutes

5.3.2 Full Shot Exercises

5.3.2.1 Simulation

The commercial high pressure die casting was simulated with its nominal fill fraction, as described in the Method, section 4.2.4.2. Each of the four different process variants were simulated; these were the combinations of “Fast” and “Slow” shot speed, and “Full” and “Blanked” in gate.

Selected frames from these simulations are presented in Figures 5-41, 5-42, 5-43 and 5-44. These frames were rendered using an in-house post processor, developed by the author, to colour the liquid metal according to particle number density. The particle number density was calculated as the smoothed local particle count, divided by the smoothed local fluid volume. Although this approach was not explicitly intended to account for porosity formation, where FLOW-3D models the formation of macroscopic porosity, the fluid volume decreases, and so the particle number density increases.

Figure 5-41 shows the simulation results for the FGFS (Full Gate, Fast Shot) case, which is the commercial process configuration. The simulated fluid behaved in a believable manner for a high speed flow; it fragmented and formed many smaller and larger bubbles, which were carried with the fluid until their collapse. As the HPDC foundry would hope, the simulation predicted that the material with the most entrainment damage was carried to the overflows by the end of the simulation.

For the FGFS simulation, the Surface Area Entrainment Algorithm (SAEC) also appears to have performed as expected. In the earlier stages of the simulation, it seems to have placed most of the particles at the leading surface of the flow front, which could be assumed to be highly unstable. There was no significant particle placement visible around moving bubbles, and particles were placed when the bubbles collapsed under pressure.

As stated in Section 4.2.4.2, the termination time of 2.3 s was chosen because this was the time when the simulated melt had solidified throughout all regions of interest.

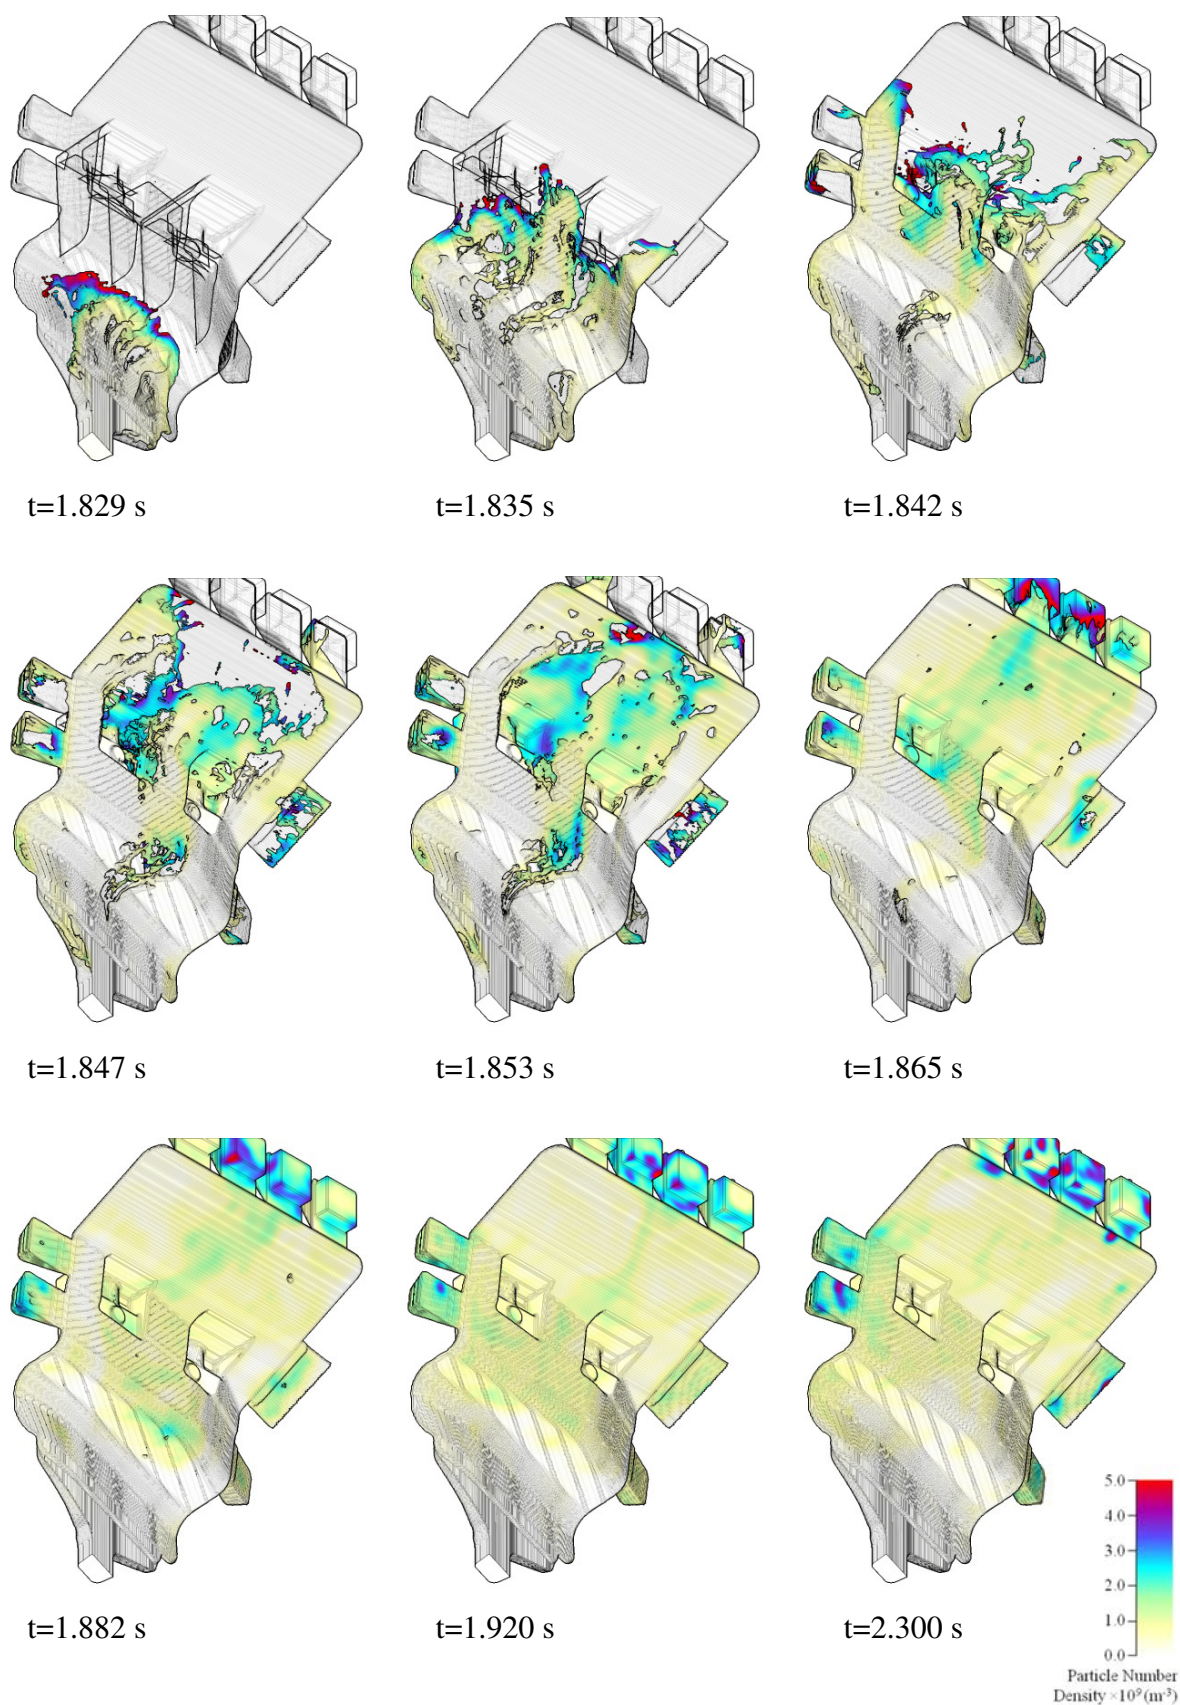


Figure 5-41: Selected frames from FGFS (Full Gate, Fast Shot) simulation

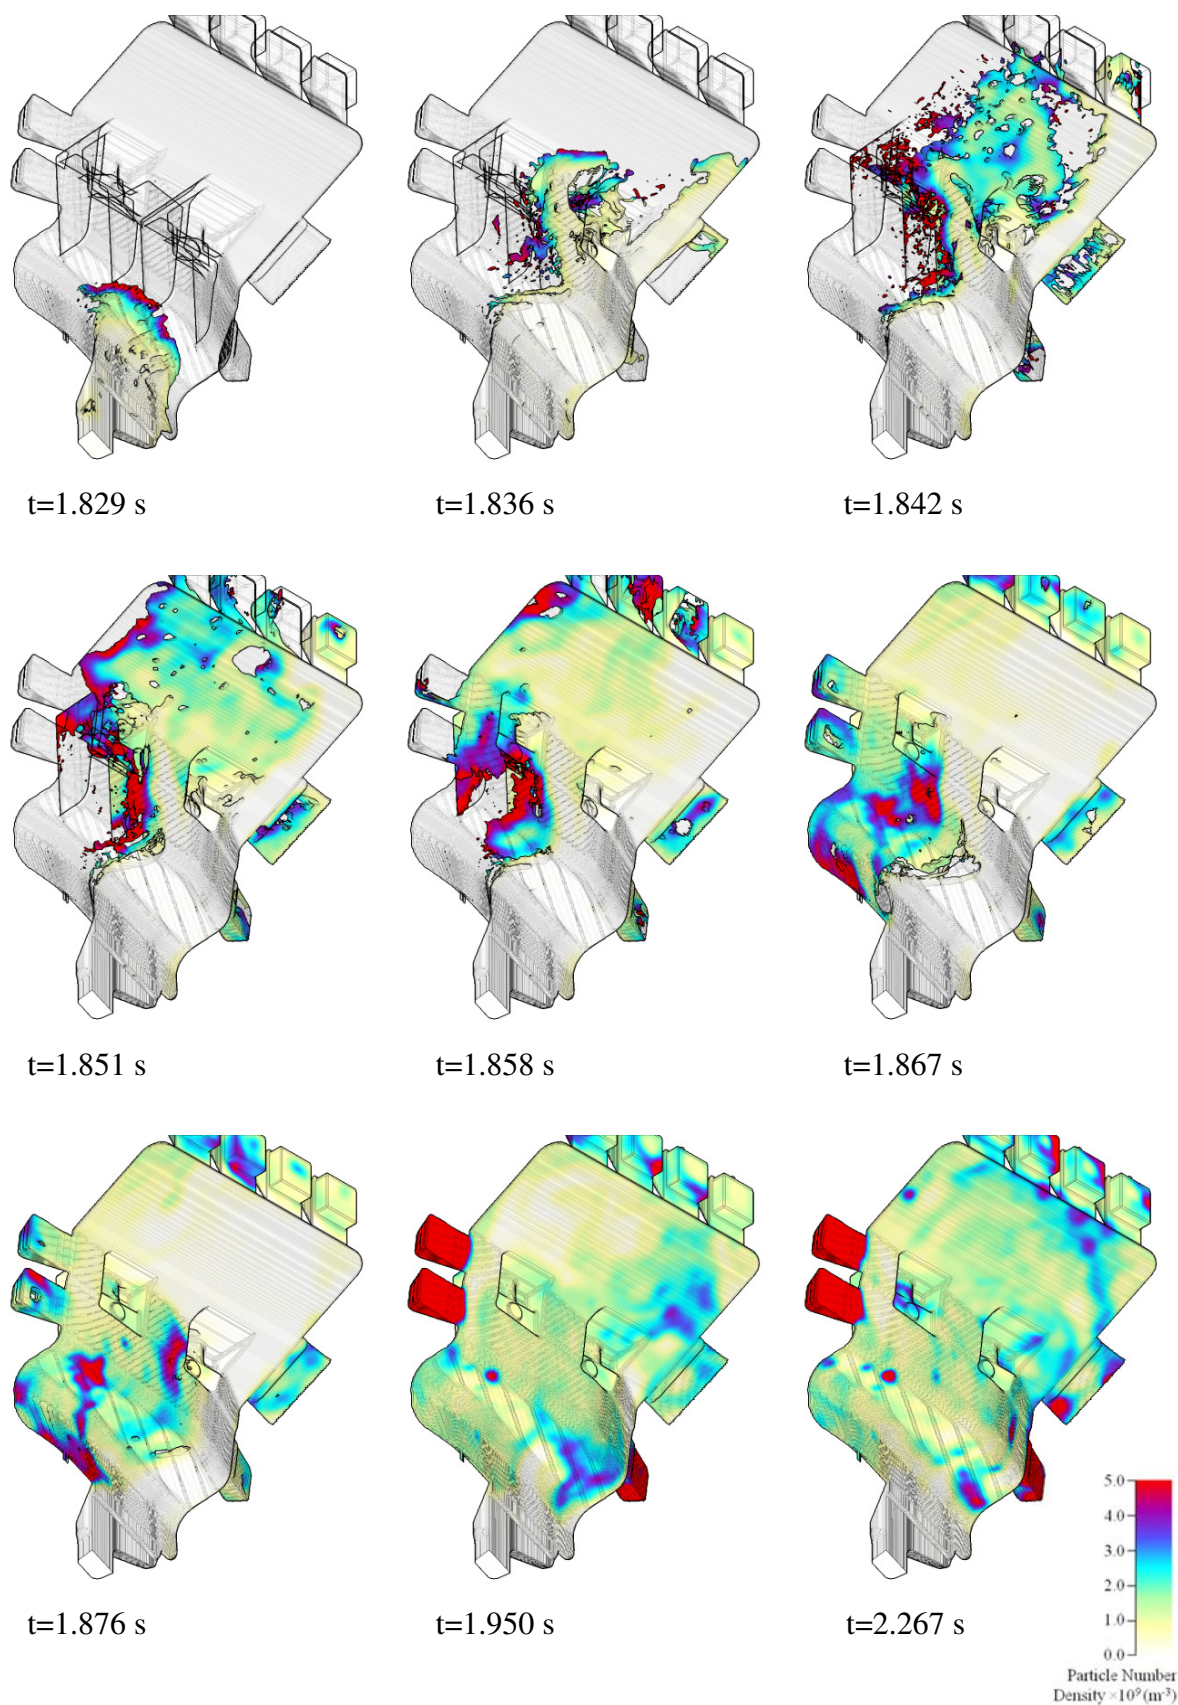


Figure 5-42: Selected frames from BGFS (Blanked Gate, Fast Shot) simulation

For the BGFS (Blanked Gate, Fast Shot) simulation, Figure 5-42, there is more cause to doubt the accuracy of the predicted entrainment pattern. Immediately prior to the $t=1.842s$ frame shown, a numerical instability in the pressure solution caused the fluid surface to fragment, particularly towards the left hand side. In the following frame shown, $t=1.851s$, these fragments have recombined, but the particle number density on this left hand edge is understandably high, because the surface area of the fluid fragments was lost as they recombined with the bulk fluid.

Between the $t=1.876s$ and $t=1.95s$ frames, the remaining bubbles collapse and another large pressure instability occurs. The most apparent result of this is that the particle number density increased off the scale in some of the overflows. In fact, wherever there were voids, these voids expanded and contracted, and in that process, a large number of particles were placed, resulting in the “hot spots” visible towards the right of the casting in the $t=1.95s$ frame.

A feature of the metal flow, which occurred in each of the simulations, but which is most apparent in the BGFS simulation, is that the metal continues to flow for tens of milliseconds after the casting is full. This demonstrates the importance of modelling the fluid flow directly during solidification, particularly where the fluid has a lot of kinetic energy.

Figure 5-43 shows that the fluid modelled in the FGSS (Full Gate, Slow Shot) simulation behaved very similarly to that in the FGFS simulation (Figure 5-41). A key difference is that with a slower shot speed less fluid fragmentation was predicted, both at the flow front, and in the region just downstream of the gate.

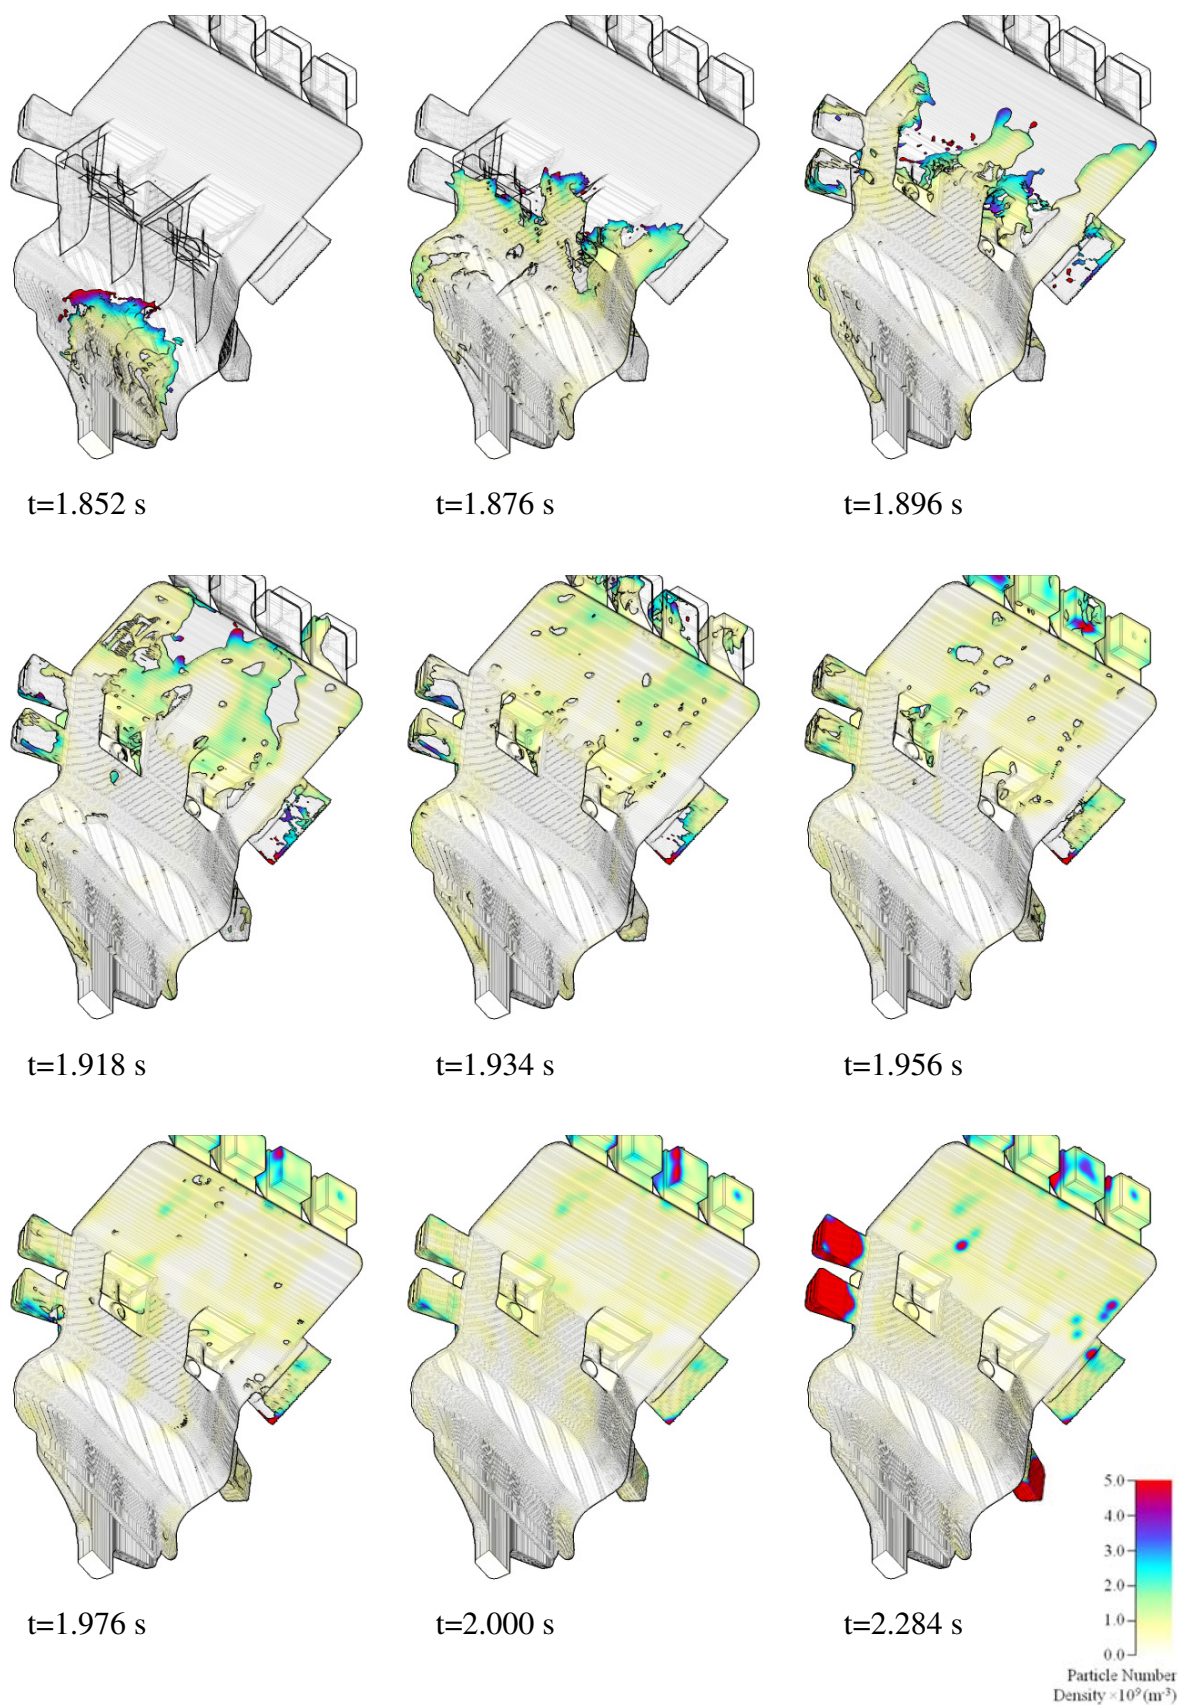


Figure 5-43: Selected frames from FGSS (Full Gate, Slow Shot) simulation

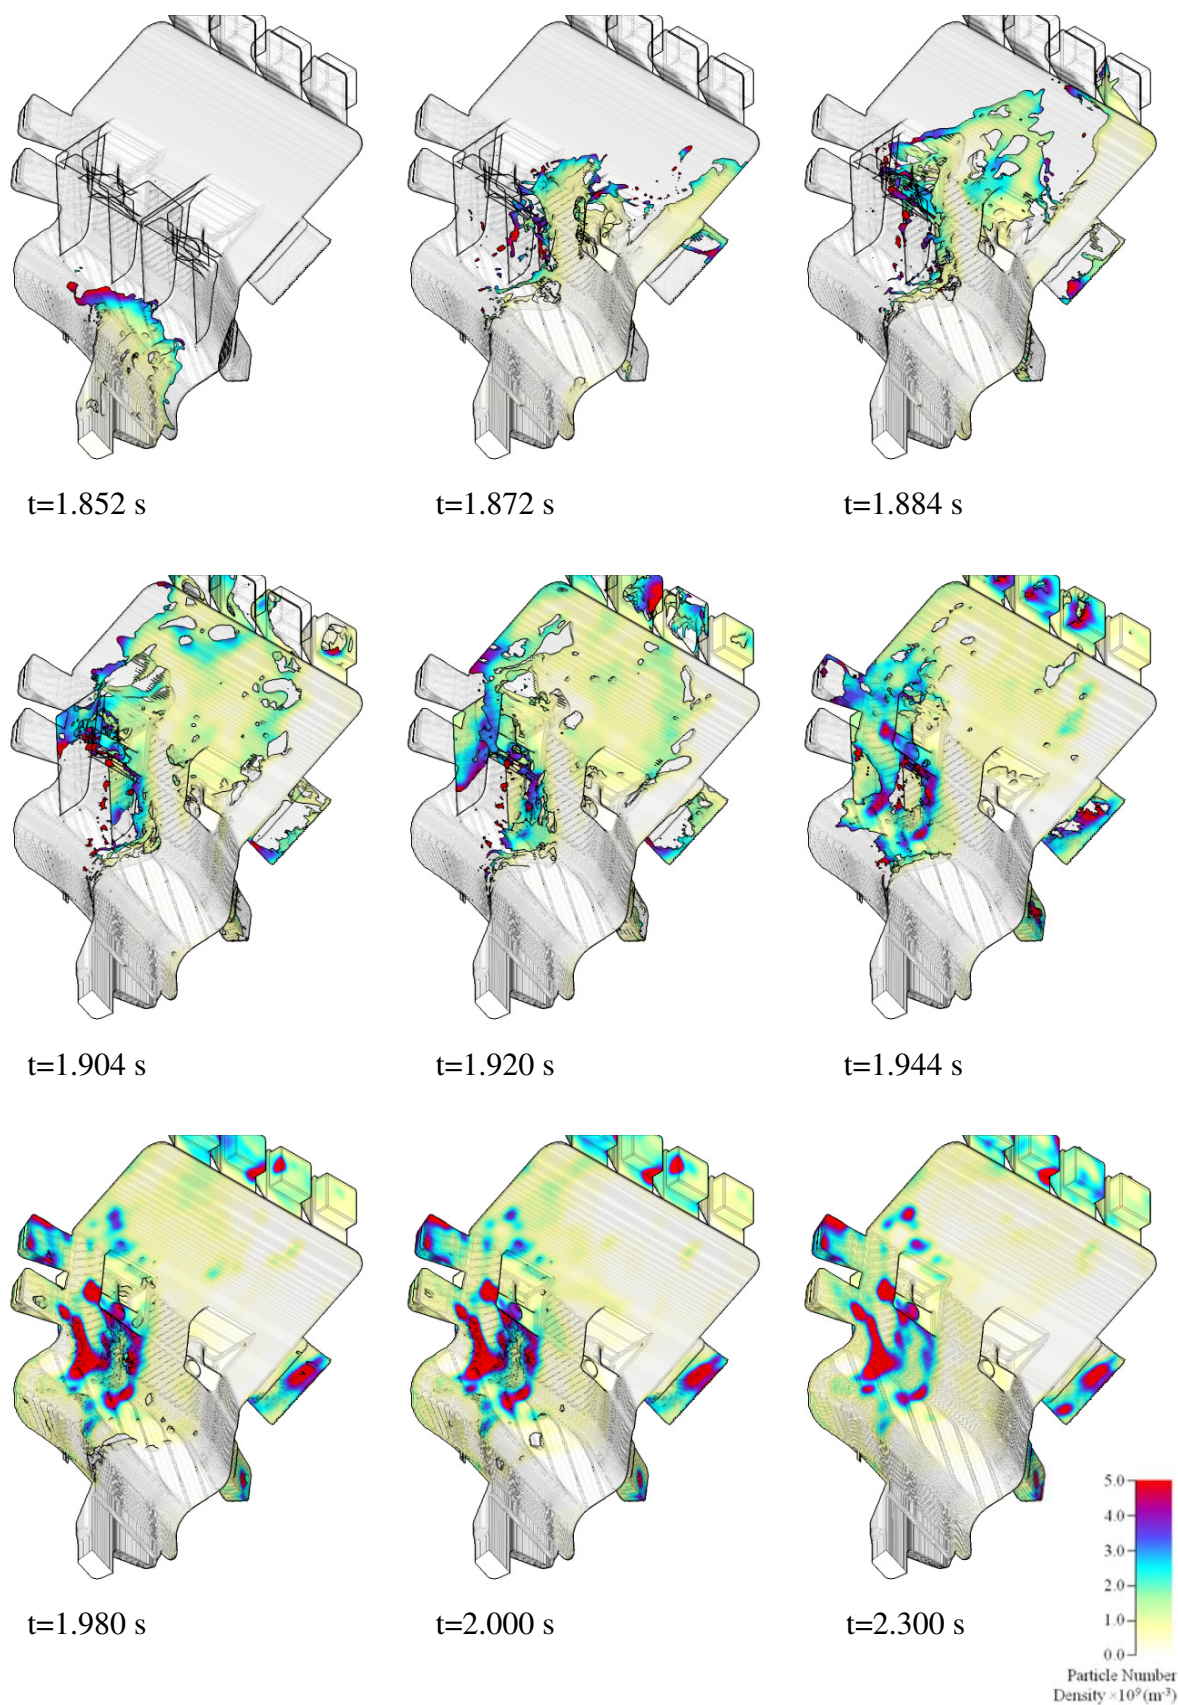


Figure 5-44: Selected frames from BGSS (Blanked Gate, Slow Shot) simulation

Again, the BGSS (Blanked Gate, Slow Shot) simulation (Figure 5-44) predicted significantly less fluid fragmentation than the “Fast Shot” equivalent. The BGSS simulation predicted that the lower left corner of the casting is one of the last places for the fluid to arrive, and was also one of the most entraining regions, as many fluid streams came together there in the final stages of filling. Interestingly, much of the fluid that eventually arrived in this region was relatively cold, and a significant proportion of it (perhaps only the skin) solidified soon after the predicted entrainment events, trapping the particles in place.

More quantitative information about the features and anomalies in the full shot simulations is provided by Figure 5-45. For the most part, the particle count over time increased relatively steadily in each simulation, until levelling off as the casting became full, with a number of exceptions.

The most significant anomaly is in the BGFS (Blanked Gate, Fast Shot) simulation, at $t \approx 1.89$ s, where the total particle count jumps from around 5×10^5 to around 8×10^5 . This coincides with a pressure instability in that simulation, as discussed above.

Towards the end of the FGSS (Full Gate, Slow Shot) simulation, Figure 5-45 shows that the overall particle count doubled, before the simulation terminated early at $t \approx 2.28$ s. It is not clear why such a large number of particles might be produced so late in the simulation, once the modelled fluid had largely solidified, but the reason is likely to be related to the cause of the early termination, such as a numerical instability. A similar, but smaller event can be seen in the curve for the BGFS simulation.

A number of small spikes in total particle count can be seen in Figure 5-45, which return to “normal” shortly afterwards. These are artefacts of when a simulation had to be restarted following an unexpected termination caused by a numerical instability.




Figure 5-45: Plot of particle history over time for the four full show HPDC simulations.

Figure 5-45 also shows downward steps in total particle count at $t=2s$, for the two “Slow Shot” simulations, and less obviously at $t=1.91s$ for the two “Fast Shot” simulations. These steps coincide with the times that each simulation transitions from the “filling” phase to the “solidification” phase. This is most likely explained by the fact that a slightly coarser mesh was used for the “solidification” phase. When the domain is re-divided into different cells, it is possible that some particles became “outside” the casting, and so were deleted.


5.3.2.2 Casting Experiments and Tensile Characterisation

For each of the simulated HPDC flow cases that were reported in the previous section, a casting experiment was performed, where the Die Casting Machine (DCM) parameters were altered to accommodate “Fast” or “Slow” third phase shot speeds, and where a blanking piece either was or wasn’t inserted into the gate at the relevant location.


15 shots were performed for each parameter case, and since each shot included one left and one right hand casting, a total of 120 parts were cast. 3 test bars were punched from each of these castings, at the locations as described in the Method, section 4.1.1, then tested in uniaxial tension, also as described in that section.

Some of the qualitative data that was gathered from these experiments is presented first, in Figures 5-46, 5-47 and 5-48. These show the results of SEM fractography of some of the samples, and were imaged using Secondary Electron Imaging (SEI). These figures are a selection of the imaging results taken throughout this research, which best represent the nature of the defects encountered. A full catalogue is included in the Appendices, section 10.1.


The sample references used in the following text and figure captions follow the format “trial code”-“shot number” Left/Right-“sample location”. For example, FGSS-10L-B would be from the 10th shot in the FGSS (Full-Gate, Slow Shot) trial, and would be sample B from the left hand casting. This nomenclature is consistent with the Appendices, both section 10.1 and also with section 10.2 which tabulates the complete set of tensile data.



A



B



C

Figure 5-46: Images from SEM fractography of sample FGSS-10L-B. A shows the transition between apparent interdentritic shrinkage porosity and a surface with a “crumpled bag” texture. B shows an EDS result from the “crumpled bag” region, as labelled in C.

The right hand side of Figure 5-46 A shows an example of a rounded pore, found in one of the test bars from an FGSS casting. Rounded pores are often attributed to excessive hydrogen in the liquid aluminium. The left hand side of Figure 5-46 A shows a different form of porosity, in which voids form between the growing dendrites. Evidence was highlighted in the Literature Review (Section 2.1.3), which suggests these forms of porosity may be initiated by an entrained surface oxide. Figure 5-46 B is an Energy Dispersive Spectroscopy (EDS) trace of part of the surface of the rounded pore, which shows a peak at the characteristic X-ray energy of oxygen. The site that the spectrum was taken at, shown in Figure 5-46 C, was selected to be at the upper edge of a wrinkle in what appeared to be a surface film, such as to maximise the amount of signal which was produced as electrons interacted with that film.

These results suggest that the surface of the “crumpled bag” texture visible towards the right of Figure 5-46 A is an oxide film. Since an oxide film can easily grow in area, but not easily shrink, it follows that as an oxygen containing pore grows, shrinks and otherwise changes shape, the film would form the kind of wrinkles observed in this micrograph.

Of further interest is the fact that the transition between the wrinkled region and the region of interdendritic porosity [65] is very gradual, with an intermediate smoother region. This would suggest that the region of interdendritic porosity did in fact grow from the rounded pore late in solidification.

Figure 5-47 shows another of these rounded pores, from one of the BGSS test bars. This pore also presented the crumpled texture which indicates the presence of a thin stable film on the fracture surface; a thick film would not crumple, and the wrinkles would not persist if the film were not stable (at least on the order of seconds at casting temperatures). A “young” aluminium oxide film – formed during casting, rather than being already present in the charge material – has these properties.




Figure 5-47: Image from SEM fractography of sample BGSS-1R-B, showing a pore (right) with the “crumpled bag” surface texture.




Figure 5-48: Images from SEM fractography of sample FGSS-15R-A, showing a large region with an unusual texture, towards the left of image A. B compares wide area EDS spectra, taken from the left (red line) and right (blue line) sides of frame A.

However, a satisfactory EDS spectrum was not acquired for the apparent film in Figure 5-47. This may be due to the fact that the pore was in a recess, so the characteristic x-rays could not reach the detector, that the film was too thin to show a signal in EDS analysis, or a combination of these factors.

The images in Figure 5-48, relate to a test bar sampled from near the gate, in one of the FGSS castings, which had the lowest tensile strength in this experiment. The film, towards the left of image A, was visible on both pieces of the test bar. Otherwise, the film did not show any of the characteristics of an oxide film, namely the crumpled texture illustrated in Figures 5-47 and 5-46. Wide area EDS analysis of the differently textured regions to the left and right hand side of image A – as illustrated by the red and blue lines in image B, respectively – showed a significant carbon response from the “film” region, where no carbon signal was returned in the region with a more normal fracture surface.

As mentioned previously, the fractographic evidence presented so far in this section, is a selection of the images gathered throughout the work, which best summarises the nature of the key defect types observed. Overall, a number of different failure modes were observed, and these were categorised into 7 groups, as detailed in Table 5-2. A complete image catalogue, which includes SEM and macroscope images, is presented in section 10.1 of the Appendices. This is arranged according to these classifications, and is intended as an auxiliary reference for the curious reader concerned with the specifics of each proposed grouping.

Microstructural evidence is not presented. As mentioned in the Introduction, the scope of this work did not include microstructure, although the influence of microstructure will be considered in the discussion.

Table 5-2: Classification criteria for observed failure modes.

Failure mode	Criteria	
Ductile Shear	“Smeared” surface texture, with no visible defects.	Fracture surface at ~45° angle to applied stress.
Porosity	At least one pore, generally with an uneven shape, and less than 1 mm in diameter.	Uneven, largely transgranular fracture surface, roughly normal to applied stress.
Carbon Film	Un-wrinkled, smooth (often planar) defect on both parts of the test bar.	Remaining fracture surface was uneven, and roughly normal to applied stress.
Thin Oxide	Flattened, crinkled feature, with a texture distinct from transgranular fracture.	Remaining fracture surface was uneven, and roughly normal to applied stress.
Dark Pore	At least one rounded bubble, sometimes elongated, with a dark inner surface, showing no sign of conforming to solidification shrinkage or wrinkling.	Remaining fracture surface was uneven, and roughly normal to applied stress.
Lap Defect	Small smooth region at the edge of the fracture surface, sometimes contiguous with a flat side of the sample.	Remaining fracture surface was uneven, and roughly normal to applied stress.
Exogenous Inclusion	Bulk non-metallic material embedded in the fracture surface.	Remaining fracture surface was uneven, and roughly normal to applied stress.

Figure 5-49 summarises the UTS data that was gathered from the commercial high pressure die castings, in the form of box-plots. The full data set, including other mechanical metrics, may be found in section 10.2 in the Appendices.

These box-plots show a high variability at some test bar locations, particularly position A, for most of the trials. Additionally, both the degree of variability, and the central values, change strongly from trial to trial, and from location to location. This would imply that the nature and severity of the defects also change from location to location, and from trial to trial.

Furthermore, most of these plots show an extended lower tail, and that the median lies above the mid-point of the range, suggesting that the Weibull distribution could appropriately model these data.

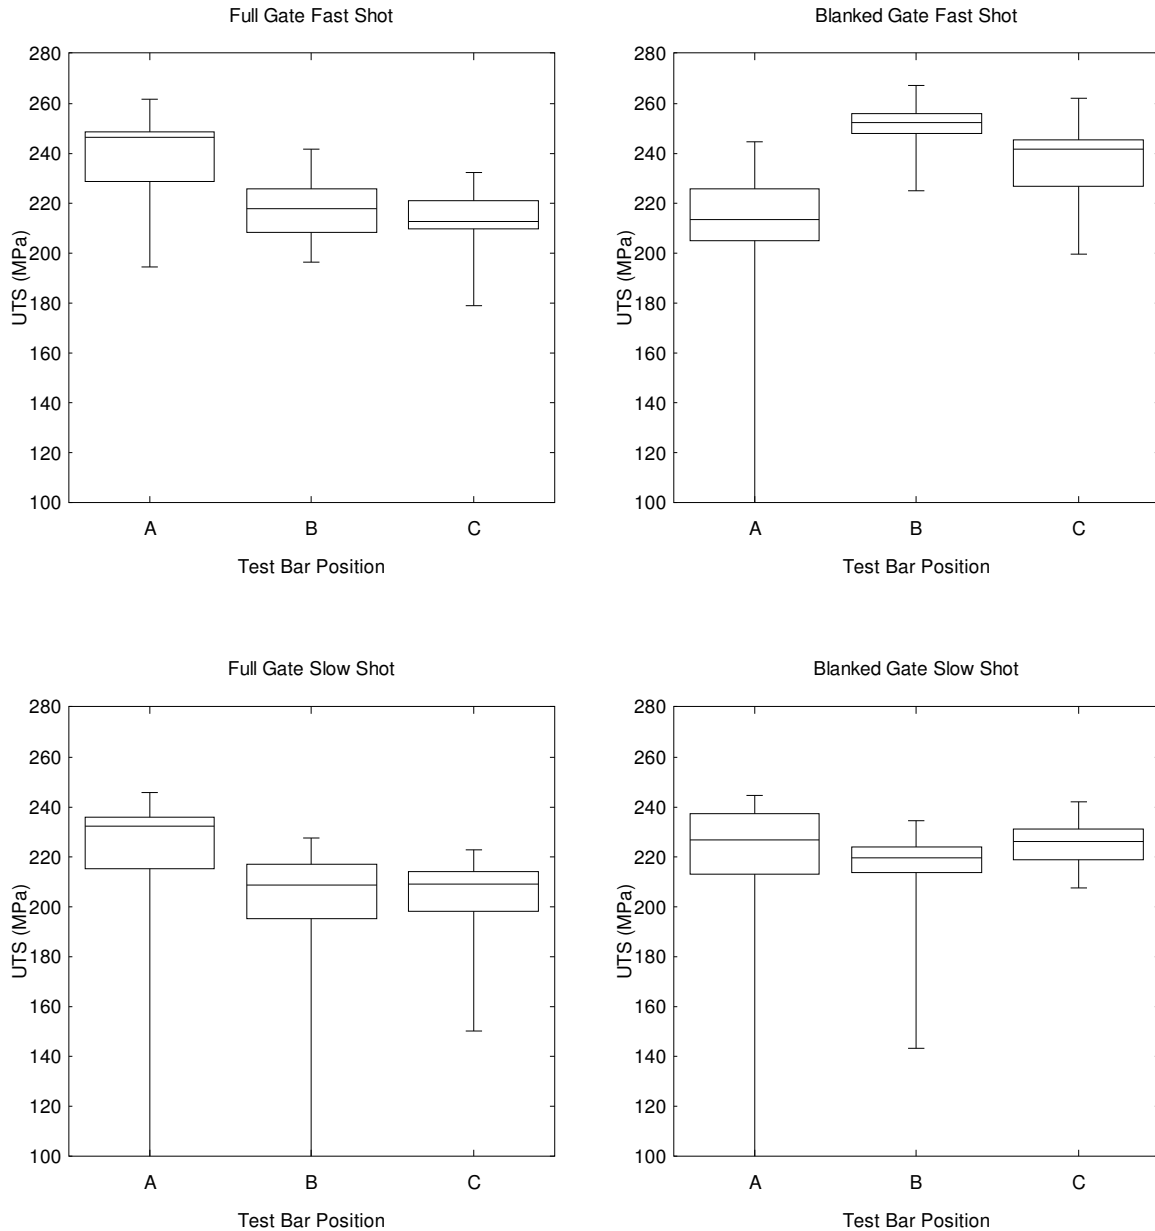



Figure 5-49: Box-Plots summarising tensile data from all samples, according to trial and test bar position.

In the following Figures, 5-50 to 5-53, the UTS results from tensile testing are presented as complete cumulative distributions, for each casting parameter set, and at each test bar location. In addition, each datum point is classified by the type of defect that was assumed to have initiated fracture, as per Table 5-2. If two types of defect were observed on the fracture surface, the largest was chosen as the initiating defect.


Perhaps unsurprisingly, the “Ductile Shear” failure mode was more dominant in the higher strength samples. It seems that in these experiments, the initiation of ductile shear placed the upper limit on the strength of test bars that did not fracture from a defect at a lower stress.

The so-called “Carbon Films” account for many of the lower strength samples, particularly for the slow shot tests. They were present in the fast shot castings, but they did not limit strength to a significantly greater extent than porosity. Curiously, their effect seems largely limited to samples from the A location, near the gate. This would suggest that their presence, or damaging nature, is somehow related to the last fluid to enter the mould, perhaps as the rest of the casting solidified.


“Lap Defects” and “Dark Pores” were most common in the BGSS (Blanked Gate, Slow Shot) and FGSS (Full Gate, Slow Shot) trials respectively, at the B and C locations. The fact that they occurred under slow shot conditions suggests that their occurrence is related to slower metal velocities. At lower velocities, defects are less likely to be broken up by fluid shear, and the metal would have lost more heat before reaching the far side of the mould.



A




B




C


Figure 5-50: Cumulative UTS data from tensile testing of specimens from FGFS (Full Gate, Fast Shot) positions A, B and C, categorised by critical defect.



A




B




C


Figure 5-51: Cumulative UTS data from tensile testing of specimens from BGFS (Blanked Gate, Fast Shot) positions A, B and C, categorised by critical defect.



A




B




C


Figure 5-52: Cumulative UTS data from tensile testing of specimens from FGSS (Full Gate, Slow Shot) positions A, B and C, categorised by critical defect.



A



B



C

Figure 5-53: Cumulative UTS data from tensile testing of specimens from BGSS (Blanked Gate, Slow Shot) positions A, B and C, categorised by critical defect.

Tests bars where failure was attributed to porosity make up a significant fraction of each of the data sets. As such, the effect of porosity damage can be seen to change significantly from dataset to dataset, based only on these cumulative distributions. Furthermore, the level of porosity damage has a significant effect on the overall distribution of strength in each dataset.

In general, the strength of samples from the A location (near the gate) was less limited by porosity than those extracted from the B or C locations (in the larger flat region). A notable exception to this was the data set for the B location of the “Blanked gate, Fast Shot” castings, which showed the highest strengths and the least variation in strength, of all the data sets.

The pattern of fluid flow was intended to be significantly altered by blanking the gate, as the above simulations predict. It seems that the distribution of porosity damage is strongly affected by whether or not the gate is blanked, which would suggest that the degree of porosity damage at each location is in turn strongly affected by fluid flow.

5.3.2.3 Weibull MLE Analysis

The tensile data were converted to true stress, and analysed using the multiple failure mode Maximum Likelihood Estimation (MLE) method described in section 4.3.1. In this analysis, samples that were categorised as porosity and thin oxide failures were treated together as “Entrainment”. Each of the other failure modes were treated individually.

The results of the Maximum Likelihood Estimation are provided in Tables 5-3 and 5-4. These results are illustrated graphically in Figures 5-54, 5-55, 5-56 and 5-57. Each of these figures show the data points attributed to “Entrainment” and “Not Entrainment”, the cumulative Weibull curves attributed to entrainment (thin lines). These figures also present combined cumulative curves, based on the parameter estimates for all failure modes (thick lines); this total cumulative curve was not directly fitted to the data.

To aid the interpretation of these results, the characteristics of the parameter fitting method will be briefly explained, with reference to data from the “Slow Shot, Blanked Gate” experiment, test bar location C. This data set failed mostly by entrainment, but there was one lap defect, and one so-called carbon film.

The parameter estimates for “lap defect” and “carbon film” both have scale parameters well beyond the nominal UTS for the alloy, and very low modulus parameters. This may seem incorrect, but the scale parameters estimated using the present method are both a measure of the expected value, and a measure of the probability that one defect type is observed, not another. If the reported parameter estimates with a high scale value for the carbon “lap defect” or “carbon film” are evaluated using the CDF for the Weibull distribution (Equation 2-2) for stresses around the average strength of the data-set, the probability of failure would be approximately 1 in 30 – the observed probability of occurrence for these two defect types.

If either of these samples were assessed individually using normal statistics, their standard deviation would be undefined, since there is not enough information about the distribution for the standard deviation to be computed. However, with the present MLE

method, the fact that most samples did not fail by these modes provides enough information for their shape and scale parameters to be balanced to optimally explain that result.

Naturally, since the estimates for the “Carbon” and “Lap Defect” groups are based on only one sample each, the confidence in these estimates is much lower than the confidence for the “Entrainment” parameter estimates.

The lack of confidence for the “Carbon Film” or the “Lap defect” groups does not affect the confidence in the “Entrainment” group. Similarly, if the supposed carbon film or lap defect were grouped together, this would not affect the parameter estimates for “Entrainment”.

Table 5-3: Proposed multi-component Weibull parameter estimates for “Fast Shot” Castings.

		Weibull Parameter Estimates at Test Bar Location					
		A		B		C	
Casting Parameter	Failure Mode	Scale, σ_0 (MPa)	Shape, m	Scale, σ_0 (MPa)	Shape, m	Scale, σ_0 (MPa)	Shape, m
Full Gate, Fast Shot	Ductile Shear	263.2	34.5	-	-	-	-
	Entrainment	262.4	25.8	228.3	18.0	224.3	23.3
	Carbon Film	320.7	6.2	-	-	576.8	3.8
	Lap Defect	-	-	768.4	5.7	-	-
Blanked Gate, Fast Shot	Ductile Shear	729.8	6.2	269.8	37.8	868.2	5.34
	Entrainment	230.2	15.2	267.5	24.8	250.2	19.0
	Carbon Film	257.9	11.3	-	-	860.4	5.4
	Exogenous Inclusion	4064.1	1.2	-	-	629.2	3.8




Figure 5-54: Data from “Fast Shot, Full Gate” experiment, after conversion to C-L fracture parameter, and showing the results of multi-defect MLE Weibull analysis.




Figure 5-55: Data from “Fast Shot, Blanked Gate” experiment, after conversion to C-L fracture parameter, and showing the results of multi-defect MLE Weibull analysis.

Table 5-4: Proposed multi-component Weibull parameter estimates for “Slow Shot” Castings.

		Weibull Parameter Estimates at Test Bar Location					
		A		B		C	
Casting Parameter	Failure Mode	Scale, σ_0 (MPa)	Shape, m	Scale, σ_0 (MPa)	Shape, m	Scale, σ_0 (MPa)	Shape, m
Full Gate, Slow Shot	Ductile Shear	252.9	50.44	-	-	-	-
	Entrainment	246.5	34.0	220.0	17.7	217.3	23.7
	Dark Pore	572.6	4.0	259.5	5.0	251.6	10.6
	Carbon Film	330.3	2.67	-	-	-	-
	Lap Defect	-	-	-	-	752.1	5.6
Blanked Gate, Slow Shot	Ductile Shear	252.8	97.2	-	-	-	-
	Entrainment	246.5	24.1	229.3	23.2	234.3	26.1
	Carbon Film	271.3	3.6	782.8	5.7	848.7	5.3
	Lap Defect	-	-	261.7	9.05	855.2	5.2




Figure 5-56: Data from “Slow Shot, Full Gate” experiment, after conversion to true stress, and showing the results of multi-defect MLE Weibull analysis.




Figure 5-57: Data from “Fast Shot, Blanked Gate” experiment, after conversion to true stress, and showing the results of multi-defect MLE Weibull analysis.

5.3.2.4 Correlation of Predicted Entrainment with Fracture Stress

To correlate the predicted entrainment damage with the observed mechanical properties, a total of 6 variables were optimised in three nested stages, as detailed in the method, section 4.3.2. Two of these variables, “Initial Smoothing Radius”, and “Random Walk Distance”, were used to generate the location dependant damage statistics from particle distribution data, at points within the gauge lengths of three virtual test bars. Particle distribution data was taken from the last time-frame of each of the full shot HPDC simulations, except for the BGFS (Blocked Gate, Fast Shot) run, because of the anomaly in that simulation. The process of generating the damage statistics from the particle distribution is slow, and is described in more detail in Theoretical Developments, section 3.2.5.

Each set of damage statistics was then used in an in house correlation program, to estimate the four parameters which link the damage statistics to the observed failure strength for each test bar. The program used a modified Maximum Likelihood Estimation method, as described in the method (section 4.3.2), to find the values of these four remaining parameters, which best explain the observed failure distribution. This program also reported the maximised likelihood that it was able to achieve by varying these parameters.

Figures 5-58 and 5-59 show that for lower values (~1 mm) of “Initial Smoothing Radius”, the maximised likelihood increased significantly. In general, as “Random Walk Distance” was increased, the maximised likelihood did increase, but the effect was far less significant.




Figure 5-58: Effect of “Initial Smoothing Radius” on the maximum likelihood output of second stage optimisation.




Figure 5-59: Effect of “Random Walk Distance” on the maximum likelihood output of second stage optimisation.

The optimal value pair was chosen simply as the one that gave the greatest value for maximum likelihood; this was a value of -845.8, from an “Initial Smoothing Radius” of 1 mm, and a “Random Walk Distance” of 58 mm.

For this case, the optimised parameters for transforming the damage statistics were $d_{offset} = -5.11 \times 10^8$, and $V_{eff} = 0.426 \text{ mm}^2$, relating to Equations 4-6 and 4-12 in sections 4.2.5 and 4.3.2 respectively. When the estimated values for the A and B parameters were substituted in Equation 4-7 (section 4.2.5), the function which maps a sample of damage, d , to fracture stress, σ^* , is given below:

$$\sigma^* = 1648.7 \times d^{-8.469 \times 10^{-2}} \quad \text{Equation 5-3}$$

The correlated strength distributions that result from these parameters are illustrated in Figures 5-60, 5-61 and 5-62 (thick lines). For reference, the cumulative failure data are also provided, along with the MLE Weibull fit for “Entrainment” for each data set (thin lines).

The fits seem reasonably good, given that they relate to statistical distributions, which inherently include variation and uncertainty. However, the strengths of the test bars in the A position seem consistently over-estimated. The algorithm used would compromise the quality of the estimates for the A position for a better overall fit, given failure modes other than “Entrainment” were most prevalent in the A location.




Figure 5-60: True stress fracture data derived from "Fast Shot, Full Gate" experiment, overlaid with multi-defect MLE Weibull analysis, and entrainment damage correlation.




Figure 5-61: True stress fracture data derived from "Slow Shot, Full Gate" experiment, overlaid with multi-defect MLE Weibull analysis, and entrainment damage correlation.




Figure 5-62: True stress fracture data derived from “Fast Shot, Blanked Gate” experiment, overlaid with multi-defect MLE Weibull analysis, and entrainment damage correlation.

5.4 Commercial Casting Structural Impact Test

5.4.1 Drop Test Simulation

Once the parameters were found that best account for the observed fracture strength of the tensile test bars, these can then form the basis of a predictive model to describe the probable strength distribution at any point in that casting. However, the quality of this prediction must be evaluated, as this should highlight any flaws in the process of forming the predictive model.

The strength prediction model was evaluated by performing a drop test experiment on the whole commercial High Pressure Die Casting investigated in section 5.3, as described in section 4.1.2 in the Method. This placed the casting under loading conditions that were not used for the calibration of the predictive model. One of these drop test experiments was then reproduced in LS-DYNA Finite Element simulations, as described in the Method, section 4.2.6. The drop test was simulated six times, once without any strength limitation applied to the material, and five times using a statistically mapped strength distribution, as described in the Method section 4.2.5.

The five strength distribution maps used are shown in Figure 5-63. This figure shows that each of these strength maps is different to some extent, as would be expected from a stochastic procedure, but that the overall pattern is consistent. In the region of the strength map around test bar location A, the strength is relatively high, and relatively uniform. Conversely, in the region that corresponds to test bar location B (farthest from the gate), the mapped strength is on average at an intermediate level, but a small fraction of the elements in this area were assigned a strength of around 200 MPa.







Figure 5-63: The five mapped strength distributions, which were applied to the cast component in the drop test simulations.




A




B




C



D



E



F

Figure 5-64: A-E) Each of the five drop test simulations which used a stochastic strength map, F) Stress distribution without a stochastic strength distribution applied. Frames taken 0.001 s from the start of the simulation, at the onset of fracture. Each frame is coloured by “maximum principle stress”.

In Figure 5-64, frames A-E show the same time frame (0.001 s) from each of these 5 simulations where a stochastic strength distribution was used. This was the first time frame which showed fracture in each of these simulations. For each of these simulations, a crack initiates towards the left of the stressed region, at the root of a curved feature.

There are small differences in the stress field around this crack in each of the simulations, at the time shown. Additionally, three of the simulations (A, C, and D) also show a second crack in the right hand side of each frame. The shape and starting position of each of these right-most cracks are different, as well as the shape of the stress field at each end of the crack. Again, this is as would be expected, given that each element in the modelled casting has a randomly assigned strength.

The equivalent frame from the run that did not invoke element failure is shown in frame F of Figure 5-64. Figure 5-65 illustrates the full course of that simulation. Two views are shown, one isometric view showing the stress field across the plate section, and one from the side, to be comparable with the experimental results presented in section 5.4.2.

From the side, the casting was predicted to rotate clockwise after the initial impact, before reaching an apparently constant rotated position. The impactor was predicted to have glanced away from the point of initial impact, down and to the right, aided by the rotation of the casting. The impactor was also predicted to have rotated slightly in its vertical axis.

The isometric view shows that for each of the illustrated time-frames, a maximum principle stress exceeding 225 MPa was predicted throughout large areas of the casting. The highest value of maximum principle stress at any integration point (element centre) was 262.47 MPa.

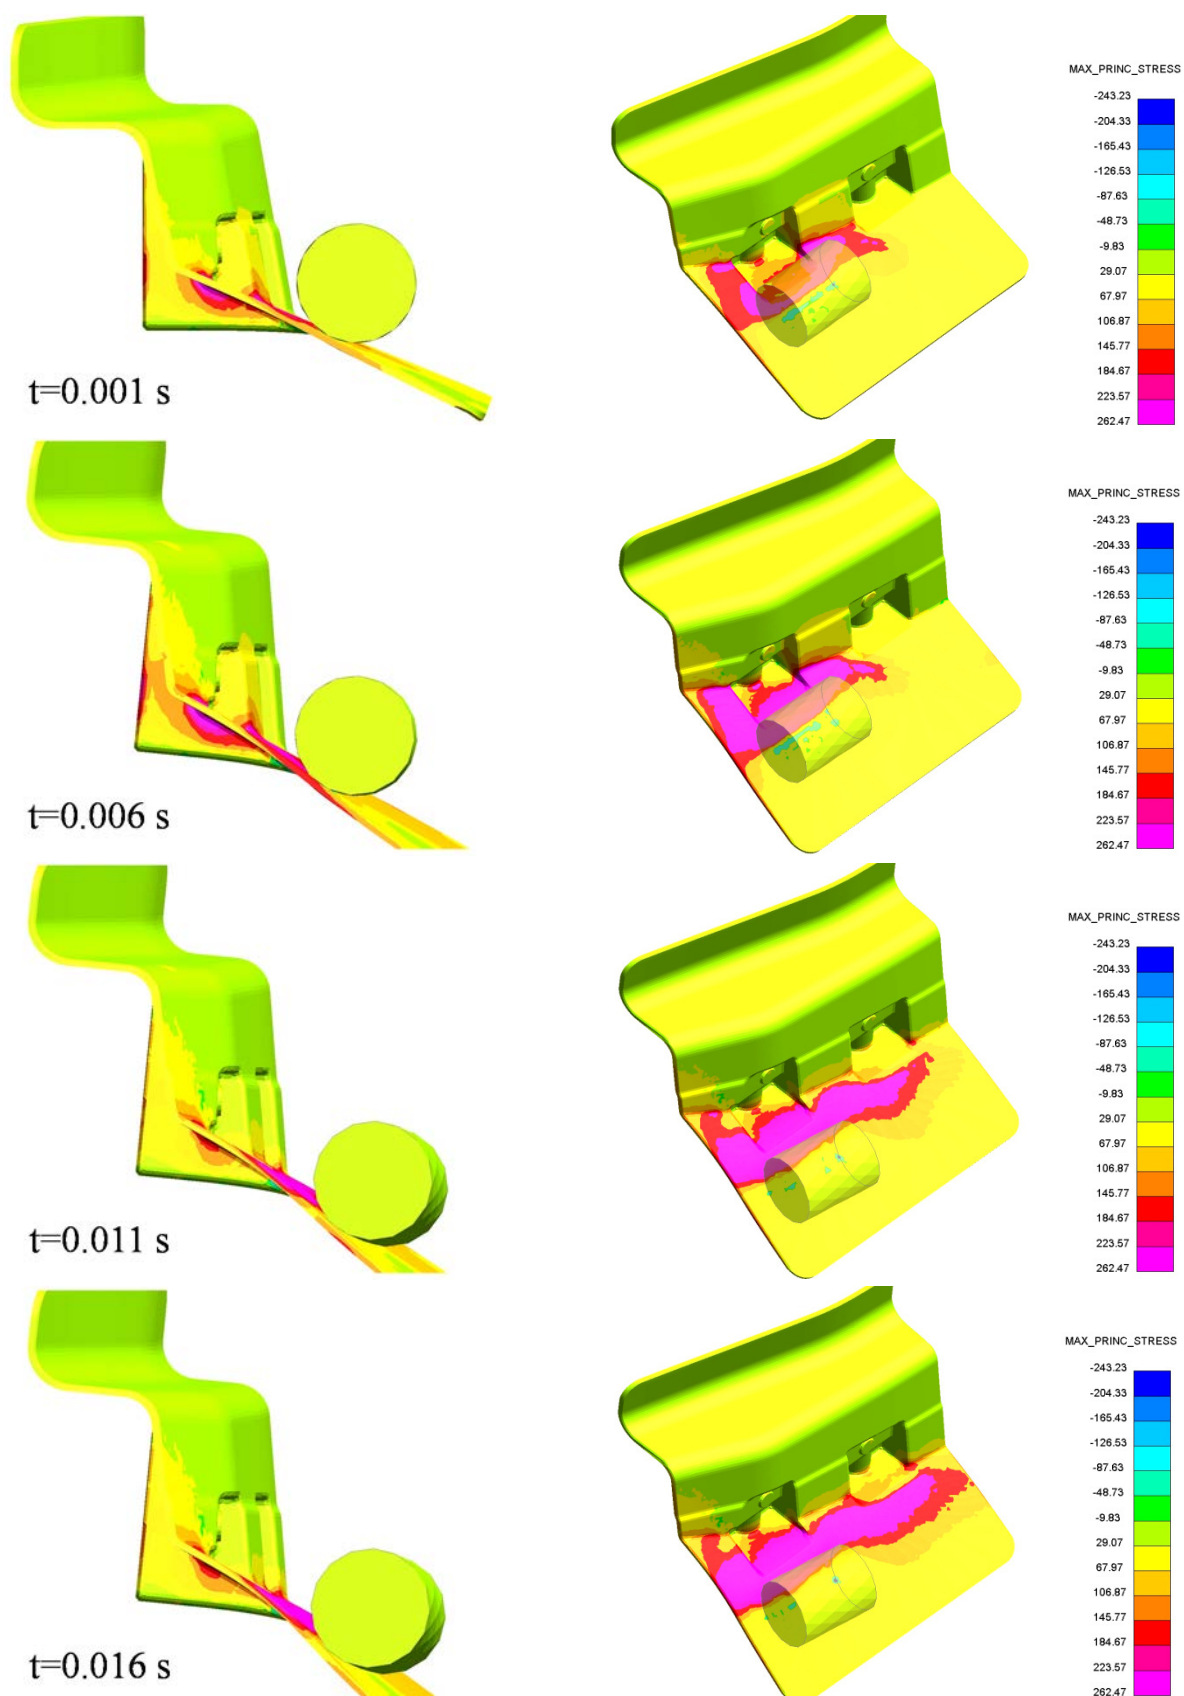


Figure 5-65: Frames from drop test simulation, without fracture strength model.

5.4.2 Drop Test Experiment

The results of three drop test experiments are presented in this section. The first of these is the case which was simulated, where the impactor was dropped from a height of 0.5 m above the casting, and struck the casting at approximately 3 ms^{-1} .

Frames from the high speed video of this test are shown in Figure 5-66. As with the simulation presented in Figure 5-65, the impactor can be seen to glance away from the casting to the left (in the orientation shown here). The casting was also shown to rotate slightly in the impact; the top right of the far side of the casting can be seen to move against the chequered background between the 0.001 s and 0.006 s timeframes. No fracture was observed.




Figure 5-66: Frames from 0.5 m drop test experiment, equivalent to those illustrated in Figure 5-65, with time normalised from the point of impact.

For the other two drop tests, the drop height was raised such that impactor was 1 m above the casting when the trolley assembly was dropped. Frames of high speed video from these two tests are shown in Figure 5-67, where the left hand frames show one test, and the right hand frames show the other. The time shown is normalised from the time of impact.




Figure 5-67: Frames from high speed video of the two drop tests carried out, with a drop height of 1 m.

As with the lower drop height, there was no fracture when the casting was first struck, and instead the impactor was deflected to the left in both of these tests. However, as the impactor was displaced further from its neutral position, a greater moment was placed at the base of the plate region. Eventually, fractures initiated here, and the plate folded downwards, taking the impactor with it. The final frames for both tests show slight differences, where the others are very consistent.

This can be attributed to differences in the fracture path taken by the crack in each test. As shown in Figure 5-68, the castings in the two tests did crack differently; the casting in the right hand frame showed a second crack towards its right hand side (as pictured), not present in the other casting.




Figure 5-68: Post-test photographs of the two castings tested with a drop height of 1 m.

Figure 5-69 shows macroscope images of the left hand fracture surface for each of the above castings. Each of these images appears to show a transgranular type fracture for most of the depth of the fracture, and below this a more shear dominated fracture. What appears to be distributed porosity may be seen throughout the transgranular region in both images. However, the pores in the leftmost image appear larger and more numerous than those in the right hand image.




Figure 5-69: Fracture surfaces from castings with 1 m drop height.

Overall, the results from the drop test experiment indicate that, quantitatively speaking, there is error in the predictive capability of the models used. That is, when the predicted strength distributions were applied, the finite element model predicted fracture at a significantly lower impact energy than was required experimentally.

However, qualitatively, the use of stochastic strength mapping produced different fracture patterns in simulation, in the same way that different fracture patterns were observed experimentally. Additionally, the experimental fractures showed evidence that damage from entrainment defects was indeed significant, and that this varied from casting to casting, thus supporting key assumptions.

5.5 Summary

The results presented in this section describe individual aspects of the experimental behaviour of High Pressure Die Castings, and the behaviour of individual numerical procedures, which aim to capture these behaviours.

The sand casting experiments which approached HPDC-like conditions demonstrated how fluid behaviour changes from lower to higher velocities, while the tensile analysis of the HPDC parts shows how the effect of each type of observed defect changes according to location and process variation. Together, the experimental results form a basis for a better understanding of how process variation may affect the strength of HPDC parts.

Besides the physical experiments, the results of trialling experimental numerical procedures were presented. Since these numerical models are implementations of the mental models of how various processes behave, they are a means of “bench testing” key assumptions, from the capabilities of fluid flow predictions, to the random character of entrainment defects.

In the following section, the results will be interpreted and cross-examined, in order to understand the confidence that may be placed in the various results and assumptions, and also to discuss what may or may not be concluded from the work.

6 Discussion

6.1 Applicability of Flow Modelling

6.1.1 High Pressure Die Casting Analysis

One of the main motivations for this work was to provide evidence to assess the quality of the results produced by a commercial CFD solver. The investigations into this line of enquiry were the Controlled Partial Vacuum Casting experiments and simulations, as well as the simulations and experiment of the “short shot” test. Some other insights were also provided by the “full shot” HPDC simulations.

An obvious discrepancy between the experimental results of the short shot test and the simulations was that the experimental casting was superficially “full”, where the simulation showed large voids. This can be attributed to a thin solid skin on the real casting, as shown by the X-ray view. The same effect was seen in the “Fast” Controlled Partial Vacuum experiment, where a less-than-nominal quantity of metal was also injected into the mould at a relatively high speed.

It is not possible for FLOW-3D to resolve the formation of a thin solidified skin that is left behind by fast moving fluid without using a significantly finer mesh, since the solid and liquid components of partially solidified fluid move together within a cell. A skin that is not modelled would undoubtedly have consequences for other aspects of the simulation, particularly where fluid returns to previously filled areas.

Furthermore, a significantly finer mesh would also be needed to account for the effect of complex rheological behaviours in the semi-solid state [30], which are believed to result in the microstructure shown in Figure 2-3 (Literature Review).

Another notable difference between the X-ray images of the casting, and the predictions of the simulations were in the size of the voids. Where the simulations predicted one or two voids which spanned a large area of the casting, the experiments showed some voids on the order of 2 cm, and many more on the order of 2 mm. The rounded shape of these smaller

voids indicates that they are either hydrogen porosity, or entrained bubbles. As discussed in the literature review, the inherently fast solidification rate favours the hypothesis that they are trapped air; there would not be time for hydrogen to diffuse out of the melt, or for the air to react with the liquid aluminium. The reduced size of the larger voids is likely to result from the bulking effect of the smaller voids.

FLOW-3D can model bulking effects in two ways, which were trialled in preliminary work, but not used in this work, as will be discussed shortly. The first method is by linking the density of the fluid to scalar concentration. That is, the software's scalar "Air Entrainment" model can be set to reduce the density of the fluid, and so increase its specific volume, as if it were filled with bubbles that are too small to resolve. The weakness of this approach is that it does not account for pressure; the fluid bulked in this way should be partially compressible, but the model assumes that it is not. This is a particular problem, given that the pressure in the fluid is greatly increased during the intensification stage. Additionally, the scalar treatment means that the bubbles cannot move independently of the fluid, meaning that their effect will be misrepresented where there are strong body forces (i.e. centrifugal accelerations).

The second is to adjust the numerical options so that FLOW-3D does attempt to resolve, and avoid artificially collapsing, bubbles down to the size of ~ 3 mesh cells. These bubbles can collapse and grow with changes in pressure, and are subject to buoyancy forces. However, this comes at the cost of decreased numerical stability, and increased runtime.

The simulations presented used a compromise between the default stable numerical options, and those which suppress unphysical bubble collapse, with the intention of improving the accuracy of the results.

However, pressure instabilities did occur during the presented simulations, as demonstrated during the "Fast Shot, Blanked Gate" simulation. These results (Figure 5-42) showed two severe pressure instabilities in the pressure solution, which resulted in the unrealistic fragmentation of the surface. These would have been caused by the collapse of

small bubbles, as described above. There are a number of forces acting on the surface of the bubble- its internal pressure, surface tension, and a virtual pressure which stabilises the bubble against the artificial “diffusion” of fluid volume into the bubble. When a bubble collapses, each of these forces will be removed, which will destabilise the pressure solution.

The nominal mesh size that was used for the “Full Shot” simulations was a compromise between run-time and supposed accuracy, which was deemed appropriate for an academic work. The total run time for each of these was on the order of weeks. As demonstrated in Table 5-1 the Short Shot simulations indicate that a two fold reduction in nominal mesh size (i.e. an eight-fold increase in mesh cells) corresponded to an increase in run time of around 50 times. This shows that the mesh size used was very close to the practical lower limit.

Using current technology, it would be certainly be prohibitive to attempt to refine the mesh to an extent where sharp variation in solid fraction or void fraction could be resolved directly.

Indeed, in an industrial setting, a run time on the order of hours (or at most days) would be required. This would necessitate fewer mesh cells, and so a less accurate representation of the heat and velocity fields. Certainly the *potential* accuracy of such a simulation is less than the *potential* accuracy of one that includes more detail.

For practical purposes it is possible that by using “calibration” of fluid properties and phenomenological models, a coarser mesh model could be adequately informative of the fluid flow dynamics in a mould, given that a finer mesh simulation would be slower and more prone to numerical noise. However, this “calibration” would have to be based on experimental work with few unknown variables.

6.1.2 Controlled Partial Vacuum Casting

The Controlled Partial Vacuum Casting work was intended to be such a set of experiments, where as many variables as possible were negated, controlled or measured. However, for each of the tests other than “Slow” and “Dump”, it was apparent that the fluid velocities in the simulation were significantly lower than in the experiments.

The most likely explanation would be a minor failure of the pressure measurement equipment after the “Slow” run, which meant that the pressure reading was offset low such that a partial vacuum below this offset recorded as zero. This would explain:

- Why the “slow” run shows a steady transition from atmospheric pressure to controlled pressure drop, and the others do not.
- Why only the “slow” results showed the chamber pressure start to drop at the target start time of 10s. The drop in vacuum at 13 s (the target time) for “Slow TC” indicates that the control timing was functioning correctly.
- Why, in the “Fast” case, where the vent valve stuck partially open, the pressure returned sharply to 0, rather than tending towards an equilibrium level of a few millibars.
- Why the “Fast” case was the most severely affected- since the vacuum was only briefly applied, this meant that the reading may have been trimmed to 0 for a large proportion of the time vacuum was actually applied.

This hypothesis assumes that the “Dump” test was affected, but that effect was not significant compared to the severity of the applied vacuum. As such, the majority of the simulations were not simulations of each exact experiment, as was intended.

Nevertheless, the experiments and the simulations provide some valuable insights regarding the nature of metal flow at the studied velocities, and the degree to which this may be simulated.

Firstly, the “Slow” experiment should compare well with its simulation (Figures 5-21 and 5-22). The shape of the flow is certainly similar for the first of the frames shown, although the shape of the fluid front appears “lumpier” in experiment than in simulation. This would indicate that simulation does not properly represent the details of the turbulent eddies beneath the surface in the experiment, which would cause these “lumps”.

For the second and third of the frames shown, the volume of fluid in the mould appears very similar in experiment and simulation, implying that the fluid flow rate is well matched. However, the experiment shows more fluid towards the centre of the mould than predicted by the simulation. This would suggest that the simulation could be over-predicting the effective viscosity of the fluid. The more momentum is exchanged between the incoming fluid and the fluid already in the mould, the more spread out the velocity field of the fluid reaching approaching the fluid front.

Another explanation is that, in reality, the gate geometry could have had more of a tendency to focus the incoming fluid towards the centre of the mould, either because of a weakness in the fluid model, or because the simulated gate geometry differed from that of the experiment. Observations of flash formation on the castings from these experiments would suggest that the mould halves were forced apart by the pressure difference. This would deform the gate, and may lead to a more focussed fluid stream.

In four of the experiments, and one of the simulations, a “V shaped” flow pattern was observed. The angle of the V should be a characteristic of the fluid’s response to the gate geometry. For the “Dump” test, the angles of the V in simulation and experiment were very similar. This indicates both that the gate geometry did not significantly differ between the simulation and the experiment, and that the fluid model used correctly represented the geometric effects. This supports the former explanation for the overly centralised distribution of fluid in the “Slow” simulation; that the fluid model over-estimates the effective viscosity of the melt. This could be an effect of the turbulence model used, or a result of numerical

viscosity; this is numerical artefact closely related to numerical diffusion, in that it over-predicts the diffusion of momentum.

In each of the real time X-ray results featuring the V shaped flow, the fluid across most of the central section of the V seemed quite thin. It is also possible that this fluid was composed of droplets which were too small or fast moving to be captured by the spatial or temporal resolution of the equipment, respectively.

In contrast, the simulation which showed the V shaped flow (“Dump”) incorrectly predicted that the flow would be biased towards one side of the mould. Fast moving small droplets would be very difficult to resolve using a Volume Of Fluid algorithm. As such it is possible that the spurious behaviour resulted from subroutines such as the “mist” region clean-up process, or artificial pressures intended to sharpen the fluid front. These functions are intended to correct the unphysical smearing of the fluid front, but could behave unexpectedly, given an unusual input fluid volume field.

The “Medium” experiment demonstrated a fluid behaviour that would be difficult for CFD to predict, as illustrated in Figure 6-1. This figure shows that from $t_{\text{norm}}=0.23$ s (dark grey) to $t_{\text{norm}}=0.28$ s (black), the upper front of the fluid did not drop down, or significantly change shape, although the shape of the sides of the “V” did change.

This would suggest that the top surface of the fluid was somehow being held in place across most of its width, between 0.23 s and 0.28 s. A possible explanation is that an oxide skin had formed on this fluid surface, between the two sides of the mould. Alternatively, the top surface of the fluid may have partially solidified. However, this seems unlikely, given the supposed high superheat in the melt, as indicated by the degree of sand penetration.




Figure 6-1: Overlaid fluid surface shapes extracted from the “Medium” experiment (Figure 5-25), at $t_{\text{norm}}=0.08$ s (light grey), $t_{\text{norm}}=0.23$ s (dark grey) and $t_{\text{norm}}=0.28$ s (black).

6.2 Evaluation of SAEC Entrainment model

6.2.1 SAEC Robustness

The numerical experiments to characterise the performance of the Surface Area Entrainment Code (SAEC) showed that the algorithm seems to positively identify entrainment events, and provide a believable indication of the relative severity of different entraining flows (colliding fronts, plunging jet, etc).

However, these numerical experiments also showed that the algorithm would place particles in situations which do not appear to be entraining. In the “Top Gated Return Wave” (

Figure 5-4), the algorithm placed particles on the underside of the incoming stream, to the extent that this masked the influence of more apparent entrainment events.

It is plausible that entrainment could occur in such a situation, given that the fluid stream would be decelerating (so surface film would bunch up), and also that gravity would act to destabilise the fluid front, so entrainment could occur. However, the simulation shows the underside of the fluid to be smooth.

If the assumptions of the entrainment code are correct, then the underside of this fluid would not be smooth in reality, and entrainment would occur, although perhaps not to the same extent. However, whether or not an entrainment algorithm should attempt to “second guess” the predictions of the underlying flow solver, is a larger and more esoteric question.

In the “Plunging Jet” example, Figure 5-7, the algorithm was found to place a limited number of particles on the surface of the downward flowing stream. This shows that the entrainment algorithm identified instances where fluid surface area was “lost”, although the macroscopic flow pattern suggests that the surface should be stretching.

The most likely explanation is that the algorithm was sensitive to noise arising from the various numerical processes. Within the model, a small and temporary “loss” due to noise would be considered entrainment, where in reality the surface film would first wrinkle, before breaking off and entering the bulk fluid.

If the wrinkles in the surface could be included in the model, and their growth into entrainment defects modelled using a less binary “stability” effect, then this could resolve both of the above issues.

Figure 5-45 illustrates the total particle count for each of the four “Full Shot” HPDC simulations. These traces show occasional steep increases in particle count, usually followed by a return to normality, when the simulation was re-started. This shows that the particle placement algorithm was sensitive to numerical instabilities in the underlying simulation, such as pressure oscillations that might cause bubbles to expand and collapse.

Overall, it seems that the SAEC is capable of interpreting the fluid flow patterns from the CFD simulation, but that it is overly sensitive to small numerical errors, which might arise from the flow solution, the entrainment algorithm or from the combination of the two. This sensitivity arises from the numerical, rather than Boolean, nature of the algorithm.

6.2.2 Mesh Sensitivity

Because the Surface Area Entrainment Code is based on a physical quantity, the results should not be inherently mesh sensitive; if two surfaces come together, the number of particles placed should depend on their area, and not on the mesh size of the simulation.

However, the results of the parametric study of mesh size and target in-gate velocity indicated that the particle count was sensitive to mesh size. The relationship of particle count to in-gate velocity was best explained if it was assumed that a factor of 2 reduction in nominal mesh size would result in around a 10 times increase in particle count.

If this effect is accounted for, the relationship between predicted entrainment and in-gate velocity (Figure 5-12) shows remarkable parallels with the results of Lai et. al. [66], who used a simple surface area based approach. Lai et. al. also found that the level of predicted entrainment was relatively constant below around 0.5 ms^{-1} in-gate velocity, and increased with in-gate velocity above this value, for a rising-jet flow. The agreement between these two sets of results suggests that both are qualitatively correct, and also supports the supposed proportionality of the SAEC to surface area.

It is not clear exactly what makes the SAEC mesh sensitive, but a likely candidate is that the fluid surface area predictions of the underlying flow solver are mesh sensitive. The results from the three “Short Shot” simulations with different mesh sizes showed that the FLOW-3D predictions were not significantly mesh sensitive in terms of energy transfer, but that the predicted surface areas in the fine and coarse meshed simulations differed by around 30% at 2 s from the start of piston motion (Figure 5-40).

The numbers in Figure 5-40 are instantaneous totals for a whole simulation. It is entirely possible that a finer mesh would result in increased numerical noise, since there would be more cells and more time-steps. The SAEC is believed to be sensitive to numerical noise, as suggested in the previous sub-section.

Additionally, the finer mesh would be better able to resolve small turbulent eddies. If these eddies result in more surface turbulence, it is likely that the algorithm would identify more entrainment.

6.3 Significance of Entrainment Defects in HPDC Parts

6.3.1 Applicability of Entrainment

6.3.1.1 Fractography

In the mechanical testing results, many different failure modes were found, which were attributed either to some defect in the test bar, or to ductile shear. Of the defects found on the fracture surfaces, only one was identified as an oxide film, fewer than were identified as exogenous inclusions.

However, over half of all the test bar fractures were attributed to a porosity defect. As was explained in the Literature Review (Section 2.1.3), it was believed that the solidification rates present in HPDC would give an entrainment defect a pore like morphology; that is, rounded rather than flattened. Such pores are pictured in Figures 5-46 and 5-47.

Using El-Sayed's estimate of $2.5 \times 10^{-6} \text{ mol} \cdot \text{m}^{-2} \cdot \text{s}^{-1}$ [29] for the reaction rate of oxygen with aluminium would indicate that in 0.3s (the approximate solidification time, based on the simulations) the oxide film on the surface of an air pore might grow approximately by 0.013 nm. Although one might expect the reaction rate to be faster at the higher pressures involved in HPDC, due to Le Chatelier's principle [67], this would explain why it was very difficult to positively identify a supposed oxide film using Energy Dispersive Spectroscopy (EDS).

Using the same estimate for reaction rate, the volume change over 0.3s can also be estimated for an air pore similar to the one pictured in Figure 5-47. If an estimated diameter of $200 \mu\text{m}$ is used, and the internal atmosphere is assumed to have a temperature and pressure of 900 K and 100 MPa respectively (based on casting parameters), the proportion of the original gas consumed after 0.3 s would be 2×10^{-6} . Again, the oxidation reaction might be expected to

be faster these higher pressures, but this figure supports the assertion that small air bubbles entrained during High Pressure Die Casting would not take on the bi-film morphology.

It was also suggested that an entrainment defect could potentially nucleate shrinkage or hydrogen porosity, although the high solidification rate would limit the potential growth of hydrogen pores. Figure 5-46 shows an entrainment defect, with a surface film formed by oxidation, as determined by EDS. In this figure, the entrainment defect transitions into a region of interdendritic shrinkage porosity.

While the above evidence does not conclusively show that *all* of the pores observed on the fracture surfaces were formed as a direct or indirect result of entrainment, it shows that the mechanism exists. It is possible that at least some of the pores could have formed without entrainment, but the mechanism for this has not been demonstrated.

An unexpected finding was the presence of flat defects, which were shown to have a carbon-rich surface, such as Figure 5-48, dubbed “Carbon film” defects. Although carbon is a common contaminant, it cannot be readily dismissed, since it was observed using two SEMs, and varied from site to site. While both the origin of the observed carbon signature, and their formation mechanism are unknown, they are unlikely to be related to the entrainment of gas by surface turbulence.

This is because they do not have the characteristic “crumpled bag” appearance, which arises from the fact that oxide films are free to crumple as they are transported through the melt, and additionally that the wrinkles are not smoothed out by surface tension or other forces. That the so-called carbon films do not show this appearance implies that they are not free to wrinkle, or there is a mechanism that smooths any wrinkling.

6.3.1.2 Mechanical Variation

Tensile testing showed that most of these “carbon film” defects were found in test bars from location A, near the gate. The results also found that some of the weakest samples in the

exercise failed because of these unexplained defects. This suggests that these defects could be a more significant limiting factor in the strength of HPDC parts than entrainment defects.

However, the tensile results also showed that these defects were most damaging in the slower speed tests; in the tests at full shot speed, the “carbon” defects caused a similar level of damage to the porosity defects.

Given the available evidence, these defects are consistent with the “Shot sleeve” defects, reported by Birch [68]. These defects were largely also found near the gate, and significantly reduced mechanical strength. Birch also found that raising the shot sleeve temperature helped to mitigate these defects. This supported his hypothesis that the defects were formed as a partially solidified layer in the shot sleeve, and were scraped off the inside of the shot sleeve by piston motion, before being carried into the mould by fluid flow.

It is widely understood that partially solidified material is carried from the shot sleeve into the mould, as this is held as the origin of Externally Solidified Crystals (ESCs) [30], as pictured in Figure 2-3 (Literature Review).

These “Shot sleeve” defects are consistent with the “carbon film”, because the substrate of solidified aluminium would prevent the defects from crumpling, both were found near the gate, and because the die-casters in the present investigation use a polymer based shot sleeve lubricant, which would form the carbon layer. It is also plausible that higher melt velocities would help to break up these pieces of solidified material, which would explain why they were less damaging in the fast shot trials.

The mechanical test results show that the reduction in strength attributed to porosity defects varied from location to location. In the “Fast-shot, Full Gate” trial, fractures attributed to porosity occurred between 200 and 225 MPa in the B location (farthest from the gate), and around 250 MPa at the A location (closest to the gate). The “Fast Shot, Blanked Gate” data showed the reverse result, with samples in the B location reaching the highest strength of all trials.

Irrespective of the results from the entrainment algorithm, the predicted fluid flow patterns were very different for these two cases. The “Fast Shot, Full Gate” simulation showed the mould filling from the gate outwards, which would suggest that the cleanest metal should end up near the gate, and the material most damaged by entrainment should solidify near location B. This pattern was predicted by the entrainment algorithm and the tensile test results support this assessment, if porosity defects are attributed to entrainment.

The “Fast Shot, Blanked Gate” simulation predicted that the fluid would sweep up around the far end of the mould, before returning towards the gate. This could reverse the distribution of entrainment defects, with more damage towards the gate, and cleaner material farther from the gate. Again, if porosity defects are attributed to entrainment, this would explain why the area weakened most by porosity changed when the gate was blanked.

The fact that the porosity defects weakened the material near the gate more in the “Fast Shot, Blanked Gate” tests shows that these porosity defects cannot solely be attributed to shrinkage. The runner and gate should provide the full intensification pressure, to counteract shrinkage, and their thermal mass should also give a favourable temperature gradient in the A location during solidification, which would again reduce shrinkage.

This does not mean that shrinkage did not contribute to the damage caused by porosity in the B and C regions; in fact Figure 5-46 shows an example of this. One of the weaknesses of the entrainment-strength correlation method used in this work is that it does not attempt to account for the interaction of entrainment defects with factors that are known to affect porosity, such as solidification shrinkage.

6.3.2 Multiple Defect Weibull Fitting

The graphs which show the cumulative failure distributions, Figures 5-54 to 5-57, have one line for the Weibull distribution fitted to “Entrainment” failures, and another for the total effect of each of the failure modes tabulated in Tables 5-3 and 5-4.


The line for total cumulative failure estimate matches the cumulative distributions of the failure data well. One of the reasons that axis transformations were not used to illustrate the data as a Weibull plot, was that it was felt that plain cumulative distribution plots illustrate the data in a more statistically relevant way; it is easy to interpret the deviations of the data points from the fitted lines as scatter in estimated probability. If the data were transformed to a Weibull plot, which linearises a perfect Weibull distribution, this picture would be distorted.

The good fit indicates that the methodology of splitting the overall failure distribution into multiple concurrent Weibull modes is mathematically sound, and also that the model adequately describes the present data-set. In particular, the method was able to account for the size and shape of the lower tails, for example in the slow speed data, where “carbon film” defects are prevalent. These could be accounted for using a 3-parameter Weibull distribution, but only by assuming some probability of failure at negative stress [18], and it would not have been possible to separate the effects of the different defect types.


The Weibull parameters presented in the results sections are estimates. As the sample size tends to infinity, the sample estimate of each parameter approaches the true population value for that parameter. Confidence testing provides an indication of how far a parameter estimated from data with a known sample size might deviate from its true value.

The disadvantage of the estimation method used is that the standard statistical confidence tests for the Weibull distribution are not strictly appropriate, except for the extreme case where all the failures may be attributed to a single mode. However, because this limiting case exists, the confidence tests can still be relevant, as a guideline.


Tirakioğlu and Hudak [69] published tables and formulae for calculating the confidence intervals for the shape parameter, m , and the position parameter, σ_0 . Figure 6-2 shows a “confidence region” plot of the Weibull analyses presented in Tables 5-3 and 5-4 based on these authors information, and a sample size of 30.



A



B



C

Figure 6-2: Confidence region plots based on the Weibull fitted parameters for each entrainment data set, at test bar locations A, B and C.

It should be noted that it was necessary to combine the equations provided in the aforementioned paper with a power law relation to transform between the present results, and the standardised distribution used by Hudak and Tirakioğlu. This allowed a confidence region to be derived for σ_0 which depends on the estimate for m , in the same way that the confidence region for the mean depends on the estimate for standard deviation.

As can be seen from Figure 6-2, the vertical error bars for the estimated modulus values overlap for many of the data sets. This indicates that the differences between the estimated modulus parameters could be attributable to statistical error, rather than a genuine difference in the underlying defect populations.

The same is not true for the position parameter; the overall variation between the estimates for this parameter is substantially larger than the confidence intervals for the estimates. This indicates that the changes in the value of this parameter, between different test conditions and sample locations, is not attributable to statistical error, and can instead be attributed to changes in the severity of entrainment.

The dominant view, purported by Campbell [1], is that a low Weibull modulus indicates a poor quality casting, and conversely that a good quality casting will have a high Weibull modulus. Figure 6-2 shows that it is the position parameter, and not the modulus parameter, which is most indicative of entrainment damage.

The above data also conclusively disproves Campbell's expectation that Weibull modulus values for High Pressure Die Castings should fall between 1 and 10; this range did not overlap with any of the presented confidence intervals for modulus. Of course, this analysis rejects the influence of the supposed shot sleeve defects, mentioned earlier.

6.4 Entrainment Correlation

6.4.1 Descriptive Quality of the Correlation

As shown in Figures 5-58 and 5-59, the correlation between mechanical strength and entrainment damage density was improved using a higher random walk distance, but was best where the initial smoothing radius was as small as possible.

This tells us that correlation of strength to predicted damage in castings can be improved by including some scattering effect, such as random walk. As the random walk distance increases, the result becomes less about the spatial distribution of the predicted defects, and is driven more by the statistical distribution of defect density, i.e. whether or not the defects form clumps.

This could indicate that the predicted spatial distribution of defects is wrong, and that adding scatter to the final position of the defects partially accounts for this. Given the known discrepancies between the fluid flow modelling and the actual behaviour of the melt, this is likely to be true to some extent.

However, the original logical basis for using a scattering function – such as random walk – was that the actual distribution of entrainment damage was expected to vary randomly from casting to casting. The improvement in correlation for a random walk distance of around 60 mm, could therefore indicate that this casting to casting variation in the location of entrainment damage was on the order of 60 mm.

Whatever the explanation for the response to random walk, the strongest response was to the initial smoothing radius. This parameter essentially sets the importance of damage localisation, and defines the values which are scattered by the random walk. The finding that an initial smoothing radius of 1 mm provided the best correlation suggests that variation in damage over small length scales has a strong effect on the variation in strength.

This could be attributed to the fact that a smaller smoothing radius captures the variations in defect density more realistically, since there is no smoothing radius in reality. In

physical terms, pockets of high defect concentration will be a more likely to initiate a crack than the same number of defects spread over a larger area. In addition, solidification shrinkage would make a cluster of air pores more damaging, if there aren't other air pores nearby to help feed the volume change.

When the best correlation parameters are used to try to infer the cumulative strength distribution for each test bar location, and for each casting trial (Figures 5-60, 5-61 and 5-62), the fit compares well with the fitted cumulative Weibull distributions, particularly for the B and C locations in the Slow-Shot trials. This indicates that the method of correlating entire distributions (rather than either σ_0 or m individually) using the modified Maximum Likelihood Estimation, is broadly valid.

However, the correlation consistently over-estimated the strength of the test bars in the A location. One possible explanation is that the model assumes that the level of “background damage”, which exists irrespective of the entrainment algorithm, is constant and uniform. The background damage is intended to represent defects formed as the metal was poured into the shot sleeve, for example. In reality, the background defects which reach the B and C locations are more likely to have been broken up by fluid forces, and so such defects at the A location, near the gate, would be proportionally more damaging.

6.4.2 Inferential Quality of the Correlation

Given a correlation between the simulated defect number density and entrainment damage, it is possible to predict the probable strength of the casting at any location. This final prediction is a product of the results from each individual stage of the investigation, and the assumptions made at each stage. The more valid the key assumptions, and the more accurate the results, the more robust the final prediction should be.

The randomised strength distributions shown in Figure 5-63 graphically demonstrate the qualities of the assumptions used to link the supposed defect number density and material

damage. It was assumed that for any given location in a casting, the severity of local entrainment damage varies randomly from one to the next, according to some distribution. This is mirrored in the way the predicted strength varies randomly on the scale of the finite element mesh used for the mapping process, with some areas generally stronger than others.

The length scale of strength variation in real castings is unknown. The assumption was made that volume scaling can be used to neutralise the effect of the scale of weakened regions. However, this is not necessarily the case, since volume scaling assumes that only the worst defect matters, i.e. there is no interaction between nearby defects.

In the LS-DYNA simulations of the drop test, which used the randomised strength maps, it can be seen that the fracture pattern varies between each casting. Similarly, in the 1 m drop tests, both runs showed a different cracking pattern. This is in line with theory, showing that a simulation with randomised fracture strength behaves as expected, and that this behaviour is qualitatively the same in reality.

However, there is a substantial quantitative difference between the predicted onset of cracking and the observed cracking behaviour. In each of the simulations that included material failure, cracks had appeared less than 1 ms after the impactor struck. Fracture did not occur in the physical experiment with the same impact energy (0.5 m drop height). With double the impact energy, cracking did occur, but not during the initial impact.

It is possible that the strain energy distribution in the simulation was not representative of reality, because the model did not properly account for the energy transfer between the impactor and the casting. However, as described in the method, four parameters were used to try to account for the dynamics of the test. If the 0.5 m drop experiment (Figure 5-66) is compared with the simulation that did not include material fracture (Figure 5-65), the overall motion is quite close. This suggests that the discrepancy between the experimental and the simulated cracking behaviour cannot be accounted for by an unrepresentative simulated stress

field. In hindsight, it may have been better to use a quasi-steady state experiment, so that the dynamic properties of the test rig would not need to be accounted for.

Another possible source of discrepancy, which would be mitigated by a quasi-steady state test, is that strain rate effects were not accounted for. Eivind [70] found that the yield strength of an A6060 aluminium alloy increased by approximately 40%, with a strain rate of 1 s^{-1} , compared to a strain rate of 0.01 s^{-1} . This would suggest that material should be more prone to cracking if strain rates were accounted for, since higher stresses would be reached before energy could be dissipated by plasticity.

Similarly, if the Weibull parameter estimates for the C location (closest to the loaded area) were used instead of the overall entrainment correlation, the fracture strength applied to the casting model would be around 15 MPa lower, according to Figure 5-60. This shows that a better correlation between the entrainment based strength prediction and the tensile test results would worsen the disagreement between the component level tests and their simulations.

Following this logic might suggest that the tensile testing did not properly quantify the strength of the material. Since the test bars were punched from the castings, they did show rough edges; the troughs in these rough edges would intensify the stresses at nearby defects, and so promote fracture. A solution to this could be to use a different technique to extract the test-bars, such as wire erosion. However, simply introducing a free edge means defects which intersect that free edge could be made more damaging than otherwise.

Another effect to consider is that in the drop test, the castings were loaded in bending, rather than in tension, and so a greater proportion of the load would have been borne by the outer skin. Given that castings solidify from the outside inwards, and that porosity is most prevalent in the last areas to solidify, it is possible that the outer skin of the casting was relatively free from defects, and so the casting was effectively stronger, as a structure. Microstructural differences in this outer skin layer have been reported in literature [30], which have also been reported to affect the mechanical properties of that layer [31, 32].

7 Conclusion

7.1 Identification of Significant Defects

A potentially novel variation of the Maximum Likelihood Estimation method, based on an extension of weakest link theory, was able to separate the statistical effects of different types of defects in tensile test data.

It was found that entrainment defects do have a significant effect in limiting the potential strength of High Pressure Die Cast parts. It was found that entrainment defects most significantly affected the average strength, rather than the variation in strength.

Additionally, porosity was found to fit within the definition of an entrainment defect, for HPDC. This was because: for short solidification times, trapped air bubbles would be expected to have the morphology of a pore; entrainment defects were found to be able to nucleate shrinkage porosity; and the distribution of porosity defects was found to be sensitive to fluid flow effects, in a way that is consistent with entrainment defects, and inconsistent with a purely thermal explanation.

It was presumed that the effects of trapped air, hydrogen and shrinkage would be almost impossible to separate, and therefore the three effects can logically be grouped together. No evidence was found to suggest that the porosity defects in the present work were unaffected by hydrogen or shrinkage, or that all of the observed pores can be attributed in some way to entrainment.

An unexpected finding was that a different class of defect was, in some cases, more significant than entrainment. These defects had a planar structure, and a carbon rich surface. While their origin was not conclusively determined, the evidence was consistent with the hypothesis that they were the remains of the solidified skin that would form in the shot sleeve, before being broken up by piston motion.

7.2 Realistic Prediction of the Distribution of Entrainment Defects

An improved entrainment algorithm was developed, which was shown to be able to detect entrainment with believable proportionality, for a wide variety of flow cases. However, this algorithm was also shown to be overly sensitive to numerical noise. The algorithm was also found to be very sensitive to unphysical surface break-up and bubble collapse that resulted from numerical instabilities in the flow models.

Because the entrainment algorithm was sensitive to both physical and unphysical fluid flow patterns, it may be concluded that the underlying CFD would have to be highly robust and accurate, for this type of entrainment model to be useful. It was shown that, for HPDC simulations, this was not the case for a current state-of-the-art flow solver.

It was also shown that there are physical effects in HPDC that are very difficult to resolve using known commercial CFD solvers, notably, the formation of a thin solidified skin layer, and the partially compressible bulking effect of entrained gas. Given that it is not possible to exactly account for every factor in a CFD simulation, the results will necessarily be inaccurate.

Under the assumption that the fluid flow patterns vary from casting to casting in reality, it is possible to accept a prediction with some inaccuracy, as one of many possible outcomes. A methodology was presented, based on a random walk algorithm, to statistically sample a spatial distribution of predicted defects. This allowed uncertainty and casting-to-casting variability to be accounted for, to some extent.

7.3 Realistic Prediction of the Strength at any Location in a Casting

This statistical sampling approach allowed a modified maximum likelihood method to be used for correlating the observed test bar failure statistics, to the predicted distribution of defects in the neighbourhood of each test bar.

In a descriptive capacity, the correlation was able to predict the cumulative probability of failure for most test bars with a reasonable accuracy, based on a function that maps damage statistics to strength statistics.

However, a randomised strength prediction based on this correlation was not able to account for the failure behaviour of a whole casting, under loading conditions that were not used to calibrate the model. While there are a number of possible explanations for the discrepancy are possible, it seems most likely that the bulk material behaviour was not appropriately described by test bar failure data, since the test bars would be weakened by extracting them from the castings.

7.4 Summary

Overall, the experiments in this work found that the fracture behaviour of High Pressure Die Castings can be extrapolated from our understanding of sand castings to some extent. However this extrapolation is not necessarily simple or intuitive and HPDC parts can be affected by additional defect types.

A number of new mathematical methods were developed to describe and predict the effects of defects in cast parts. Each of these methods showed promise and potential. However, it became apparent that the entrainment algorithm devised is prone to numerical error, and the underlying HPDC simulations were prone to numerical instability. Significant advances will need to be made in the speed and reliability of the predictions, before it will be practical to predict the strength distribution in High Pressure Die Castings *ab initio* in an industrial setting.

8 Future Work

8.1 Prediction of the Distribution of Entrainment Defect Distribution

The research approach presented in this work was based on the information that was available in the early stages of the project. However, the insights gained throughout this research indicated possible lines of enquiry, which may be pursued by future researchers.

The entrainment algorithm presented here has two main weaknesses. The first is its over-sensitivity to numerical error, and the second that realistic bubble dynamics (buoyancy and drag) should be included and calibrated.

Another potential area for improvement would be in the method of randomising the effective defect distribution. The random walk based scattering algorithm used in this work assumed that scatter was equally likely in all directions, whereas the scatter should be determined by fluid flow effects.

An entirely different form of entrainment prediction algorithm can be envisaged, where the formation of entrainment defects is not predicted directly by a subroutine of the flow solver, but as part of the post processing. This could allow the entrainment pattern to be re-assessed with different parameters (e.g. buoyancy), which would allow the effect of those parameters to be better understood, and also allow a more realistic randomisation process. Such an indirect approach would have less potential accuracy than a more direct method, but could be less sensitive to numerical error.

8.2 Experimental Methods

Although the results of the controlled partial vacuum casting system were useful and informative, the system did not perform as expected. It is recommended that future experiments based on this approach, use a PC based control system such as LabVIEW, and a commercial calibrated pressure sensor with on-board signal processing. This would add cost,

but also add to the value of the experiments. The stepper-motor based valve actuation used in the present work was prone to sticking; if possible, this should also be negated by design.

The tensile test results reported may under-represent the strength of the castings, since they were extracted using a punch. A better approach would be to cut the test-bar shapes from over-sized sections of the castings using wire erosion, or similar. However, the best way to avoid initiating fracture from a cut edge, while ensuring that the test is relevant, might be to test a U or V shaped section of a casting in bending, such that the cut edges are in compression. As a non-standard test, its results would however be more difficult to quantify.

In a similar vein, future whole casting mechanical tests should avoid error arising from unknown flexibility in the test equipment. A quasi-static test should be less prone to such errors than a dynamic test, but the deflection of the test rig should still be measured.

9 References

1. Campbell, J., *Complete Casting Handbook : Metal Casting Processes, Techniques and Design*. 2011, Butterworth-Heinemann.
2. LaVelle, D., *Aluminium die casting and factors affecting pressure tightness*. Transactions of American Foundrymen's Society, 1962. **70**: p. 641-647.
3. Dørum, C., H.I. Laukli, and O.S. Hopperstad, *Through-process numerical simulations of the structural behaviour of Al–Si die-castings*. Computational Materials Science, 2009. **46**(1): p. 100-111.
4. Dørum, C., et al., *Numerical modelling of magnesium die-castings using stochastic fracture parameters*. Engineering Fracture Mechanics, 2009. **76**(14): p. 2232-2248.
5. Fagerholt, E., et al., *Experimental and numerical investigation of fracture in a cast aluminium alloy*. International Journal of Solids and Structures, 2010. **47**(24): p. 3352-3365.
6. Green, N. and J. Campbell, *Statistical distributions of fracture strengths of cast Al • 7Si • Mg alloy*. Materials Science and Engineering: A, 1993. **173**(1): p. 261-266.
7. Ohnaka, I., et al., *Challenging issues in computer simulation of casting*. International Journal of Cast Metals Research, 2011. **24**(3-4): p. 133-138.
8. Campbell, J., *Chapter 2 - Entrainment*, in *Castings (Second Edition)*, J. Campbell, Editor. 2003, Butterworth-Heinemann: Oxford. p. 17-69.
9. Hasegawa, M., *Chapter 3.3 - Ellingham Diagram*, in *Treatise on Process Metallurgy*, S. Seetharaman, Editor. 2014, Elsevier: Boston. p. 507-516.
10. Batchelor, G.K., *An introduction to fluid dynamics*. 1967, London: Cambridge University Press.
11. Knott, J.F., *Fundamentals of fracture mechanics / (by) J.F. Knott*. 1973, Butterworth: London.
12. Zhuang, Z., et al., *Chapter 2 - Fundamental Linear Elastic Fracture Mechanics*, in *Extended Finite Element Method*. 2014, Academic Press: Oxford. p. 13-31.
13. Nyahumwa, C.W.M., *Influence of Oxide film filling defects on fatigue properties of cast Al-7Si-Mg Alloy / by Christian Wambura Mwita Nyahumwa*. 1997, Thesis (Ph.D) - University of Birmingham, School of Metallurgy and Materials, Faculty of Engineering.: Birmingham.
14. Paris, P. and F. Erdogan, *A Critical Analysis of Crack Propagation Laws*. Journal of Basic Engineering, 1963. **85**(4): p. 528-533.
15. Tiryakioğlu, M., *On the size distribution of fracture-initiating defects in Al- and Mg-alloy castings*. Mater. Sci. Eng. A-Struct. Mater. Prop. Microstruct. Process., 2007. **476**(1-2): p. 174-177.
16. Castillo, E., *Chapter 5 - The Gumbel, Weibull and Frechet Distributions*, in *Extreme Value Theory in Engineering*, E. Castillo, Editor. 1988, Academic Press: San Diego. p. 183-209.
17. Weibull, W., *A Statistical Distribution Of Wide Applicability*. Journal of applied mechanics, 1951. **103**.
18. Tiryakioğlu, M. and J. Campbell, *Weibull analysis of mechanical data for castings: A guide to the interpretation of probability plots*. Metallurgical and Materials Transactions A, 2010. **41**(12): p. 3121-3129.
19. Hudak, D. and M. Tiryakioğlu, *On comparing the shape parameters of two Weibull distributions*. Materials Science and Engineering: A, 2011. **528**(27): p. 8028-8030.
20. Reilly, C., *Development of quantitative casting quality assessment criteria using process modelling / by Carl Reilly*. 2010, Thesis (PhD.)--University of Birmingham,

- School of Mechanical Engineering, College of Engineering and Physical Sciences.: Birmingham.
21. Yue, Y. and N. Green. *Modelling of different entrainment mechanisms and their influences on the mechanical reliability of Al-Si castings*. in *IOP Conference Series: Materials Science and Engineering*. 2012. IOP Publishing.
 22. Cockcroft, M. and D. Latham, *Ductility and the workability of metals*. J Inst Metals, 1968. **96**(1): p. 33-39.
 23. Timelli, G. and F. Bonollo, *Quality mapping of aluminium alloy diecastings*. Metallurgical Science and Tecnology, 2008. **26**(1).
 24. Fox, S. and J. Campbell, *Visualisation of oxide film defects during solidification of aluminium alloys*. Scripta materialia, 2000. **43**(10): p. 881-886.
 25. Dispinar, D. and J. Campbell, *Porosity, hydrogen and bifilm content in Al alloy castings*. Materials Science and Engineering: A, 2011. **528**(10): p. 3860-3865.
 26. Raiszadeh, R. and W. Griffiths, *A method to study the history of a double oxide film defect in liquid aluminium alloys*. 2006.
 27. Raiszadeh, R. and W. Griffiths. *Estimation of the duration of the atmosphere within a double oxide film defect in pure aluminium*. in *World Foundry Congress, Istanbul, Turkey*. 2005.
 28. Raiszadeh, R. and W. Griffiths, *A semi-empirical mathematical model to estimate the duration of the atmosphere within a double oxide film defect in pure aluminum alloy*. Metallurgical and Materials Transactions B, 2008. **39**(2): p. 298-303.
 29. El-Sayed, M.A., et al., *Determination of the lifetime of a Double-Oxide Film in Al Castings*. Metallurgical and Materials Transactions B, 2014. **45**(4): p. 1398-1406.
 30. Otarawanna, S., et al., *Formation of the surface layer in hypoeutectic Al-alloy high-pressure die castings*. Materials Chemistry and Physics, 2011. **130**(1): p. 251-258.
 31. Cáceres, C.H., et al., *Microhardness mapping and the hardness-yield strength relationship in high-pressure diecast magnesium alloy AZ91*. Materials Science and Engineering: A, 2005. **402**(1-2): p. 258-268.
 32. Weiler, J.P., et al., *Variability of skin thickness in an AM60B magnesium alloy die-casting*. Materials Science and Engineering: A, 2006. **419**(1-2): p. 297-305.
 33. Campbell, J., *Castings practice : the 10 rules of castings / John Campbell*. 2004, Elsevier Butterworth-Heinemann: Oxford.
 34. Campbell, J.C., *Concise Castings: A Casting Workshop Lecture*. 2010: AFS. 104.
 35. Runyoro, J., *Design of the Gating System*. 1992, Thesis (PhD.)--University of Birmingham, School of Mechanical Engineering, College of Engineering and Physical Sciences.: Birmingham.
 36. Dai, X., et al., *Effects of runner system design on the mechanical strength of Al-7Si-Mg alloy castings*. Materials Science and Engineering: A, 2003. **354**(1-2): p. 315-325.
 37. Hsu, F.-Y., M.R. Jolly, and J. Campbell, *A multiple-gate runner system for gravity casting*. Journal of Materials Processing Technology, 2009. **209**(17): p. 5736-5750.
 38. Campbell, J., *Rule 2 - Avoid turbulent entrainment (the critical velocity requirement)*, in *Castings Practice*, J. Campbell, Editor. 2004, Butterworth-Heinemann: Oxford. p. 9-101.
 39. Cebeci, T., *Analysis of turbulent flows with computer programs*. 2013: Butterworth-Heinemann.
 40. Massey, B.S. and J. Ward-Smith, *Mechanics of fluids*. Vol. 1. 1998: CRC Press.
 41. Cuesta, R., et al., *Numerically modeling oxide entrainment in the filling of castings: The effect of the webber number*. JOM, 2006. **58**(11): p. 62-65.
 42. Isawa, T., *The control of the initial fall of liquid metal in gravity filled casting systems*. Department of Metallurgy and Materials, The University of Birmingham, Birmingham, 1994.

43. Reilly, C., N.R. Green, and M.R. Jolly, *The present state of modeling entrainment defects in the shape casting process*. Applied Mathematical Modelling, 2013. **37**(3): p. 611-628.
44. Barkhudarov, M. and C. Hirt. *Tracking defects*. in *1st International Aluminum Casting Technology Symposium*. 1998.
45. Reilly, C., N. Green, and M. Jolly. *State of the Art Review of Modelling Entrainment Defects in the Shape Casting Process*. in *Shape Casting: Fourth International Symposium 2011 (in honor of Prof. John T. Berry)*. 2011. John Wiley & Sons.
46. Hirt, C., *Modelling Turbulent entrainment of air at a free surface*, in *Flow Science Technical Note FSI-03-TN61*. 2003, Flow Science Inc.
47. Reilly, C., *Review of the techniques available for modelling entrainment defects in the casting process*, in *Foundry Trade Journal*. 2012, ICME. p. 25-28.
48. Ohnaka, I., et al., *Mechanism and estimation of porosity defects in ductile cast iron*. International Journal of Cast Metals Research, 2008. **21**(1-4): p. 11-16.
49. Yeoh, G.H. and J. Tu, *Chapter 1 - Introduction*, in *Computational Techniques for Multiphase Flows*, G.H. Yeoh and J. Tu, Editors. 2010, Butterworth-Heinemann: Oxford. p. 1-20.
50. Griffiths, W., et al., *The Application of Positron Emission Particle Tracking (Pept) to Study the Movement of Inclusions in Shape Castings*. 2009.
51. Kimatsuka, A., et al., *Mold filling simulation of high pressure die casting for predicting gas porosity*. Modeling of Casting, Welding and Advanced Solidification Processes- X, 2003: p. 335-342.
52. Lin, J. and M. Sharif, *Numerical simulation of the movement, breakup and entrapment of oxide films during aluminum casting*. Aluminum Transactions, 1999. **1**(1): p. 71-78.
53. Dai, X., et al. *Modelling of liquid metal flow and oxide film defects in filling of aluminium alloy castings*. in *IOP Conference Series: Materials Science and Engineering*. 2012. IOP Publishing.
54. Hirt, C.W. and B.D. Nichols, *Volume of fluid (VOF) method for the dynamics of free boundaries*. Journal of computational physics, 1981. **39**(1): p. 201-225.
55. Zhao, H., I. Ohnaka, and J. Zhu, *Modeling of mold filling of Al gravity casting and validation with X-ray in-situ observation*. Applied Mathematical Modelling, 2008. **32**(2): p. 185-194.
56. Ohnaka, I., et al. *Porosity formation mechanism in Al and Mg alloy castings and its direct simulation*. in *Melting of Casting and Solidification Processes VI (6th pacific rim conference)*. 2004.
57. Hinton, E., *Development of a Casting Process for Automotive Austempered Ductile Iron Structural Components*. 2010, University of Birmingham.
58. Watson, R., et al. *Application of a Novel Entrainment Defect Model to a High Pressure Die Casting*. in *Advanced Materials Research*. 2014. Trans Tech Publ.
59. Watson, R., et al., *Towards Location Specific Statistical Fracture Prediction in High Pressure Die Castings*, in *10th European LS-DYNA Conference*. 2015, DYNAmore: Würzburg.
60. Harlow, F.H. and P.I. Nakayama, *Turbulence transport equations*. Physics of Fluids (1958-1988), 1967. **10**(11): p. 2323-2332.
61. *Theory: Turbulence Transport Models*, in *FLOW-3D 10.1.1 User Manual*. 2013, Flow Science Inc.
62. Schmid, M. and F. Klein. *Fluid flow in die cavities-experimental and numerical simulation, NADCA 18*. in *International Die Casting Congress and Exposition*. 1995.
63. Sirrell, B., M. Holliday, and J. Campbell, *Benchmark testing the flow and solidification modeling of Al castings*. JOM, 1996. **48**(3): p. 20-23.

64. Dour, G., et al., *Development of a non-intrusive heat transfer coefficient gauge and its application to high pressure die casting: effect of the process parameters*. Journal of Materials Processing Technology, 2005. **169**(2): p. 223-233.
65. Menard, G. and M. Richard, *Microfractographic Atlas of Cast Aluminium Alloys*. Centre Technique des Industries de la Fonderie, 188. 1992: Abe Books.
66. Lai, N., W. Griffiths, and J. Campbell, *Modelling of the potential for oxide film entrainment in light metal alloy castings*. Modeling of Casting, Welding and Advanced Solidification Processes- X, 2003. **2003**.
67. Le Chatelier, H. and O. Boudouard, *Limits of flammability of gaseous mixtures*. Bull. Soc. Chim.(Paris), 1898. **19**: p. 483-488.
68. Birch, J., *Shot Sleeve Defects*, in *Diecasting World*. 2001. p. 8-9.
69. Tiryakioğlu, M. and D. Hudak, *Unbiased estimates of the Weibull parameters by the linear regression method*. Journal of Materials Science, 2008. **43**(6): p. 1914-1919.
70. Semb, E., *Behavior of Aluminum at Elevated Strain Rates and Temperatures*. 2013.

10 Appendices

10.1 HPDC Fractography Catalogue

10.1.1 Samples Classified as “Ductile Shear”

10.1.1.1 FGFS-7L-A




Figure 10-1: SEM SEI image showing the representative fracture surface for this sample.




Figure 10-2: SEM SEI image showing some apparent shrinkage porosity which was found in the sample.

10.1.1.2 BGFS-2R-B




Figure 10-3: High magnification SEM SEI image of what appears to be shrinkage porosity in this sample.




Figure 10-4: High magnification SEM SEI image of the Micro-Void Coalescence fracture, which was representative of the sample.

10.1.1.3 BGFS-15R-C




Figure 10-5: Low magnification optical image of the sample.

10.1.1.4 FGSS-3L-A




Figure 10-6: Low magnification SEM SEI image, giving an overview of the fracture face.

10.1.1.5 BGSS-1L-A




Figure 10-7: SEM SEI image of the Micro-Void Coalescence fracture, which was representative of the sample.




Figure 10-8: Low magnification SEM SEI image, giving an overview of the fracture face.

10.1.2 Samples Classified as “Porosity”

10.1.2.1 FGFS-11L-C




Figure 10-9: High magnification SEM SEI image, showing what appears to be shrinkage and trapped air porosity.




Figure 10-10: SEM SEI image showing a cluster of small pores.




Figure 10-11: SEM SEI image giving an overview of the fracture face, and showing pores of various sizes, clustered away from the edges of the sample.

10.1.2.2 BGFS-1R-A




Figure 10-12: Optical image providing an overview of the fracture surface.

10.1.2.3 FGSS-6L-C




Figure 10-13: Close-in SEM SEI image of the small pores scattered around the sample.




Figure 10-14: SEM SEI image providing an overview of the fracture surface.

10.1.2.4 FGSS-10L-B




Figure 10-15: SEM SEI image showing what appears to be an air pore around 1 mm in diameter in this sample




Figure 10-16: SEM SEI image showing a close up of the wrinkled texture present in this pore.




Figure 10-17: SEM SEI image showing a close up of the edge of the wrinkled region




Figure 10-18: SEM SEI image showing the transition between the wrinkled region and apparent shrinkage porosity.




Figure 10-19: SEM SEI image showing the regions of interest for the spectra in Figure 10-20yy, Spectrum 3 was taken from the edge of one of the wrinkles, and Spectrum 4 was taken from an apparent MVC region.




Figure 10-20: The EDS spectra acquired from the locations in Figure 10-19

10.1.2.5 BGSS-13L-B




Figure 10-21: SEM SEI overview of the sample, showing a cluster of porosity towards the through-thickness centre of the sample.




Figure 10-22: The central section of Figure 10-21, imaged at a greater magnification.

10.1.3 Samples Classified as “Carbon Film”

10.1.3.1 FGFS-9R-C




Figure 10-23: SEI SEM image of the sample, showing features where the metal appears to have delaminated at film-like features perpendicular to the applied stress. Also, the location of the spectra illustrated in Figure 10-24.




Figure 10-24: The EDS spectra, at the locations shown in Figure 10-23.




Figure 10-25: An SEI SEM image of a roughly planar defect.




Figure 10-26: A low magnification SEI SEM image providing an overview of the sample.




Figure 10-27: A higher magnification SEI SEM image of the feature in the centre of Figure 10-26.

10.1.3.2 BGFS-10L-A




Figure 10-28: SEM SEI image showing some apparent shrinkage porosity present in the sample.




Figure 10-29: SEM SEI image giving a representative view of the majority of the fracture surface.




Figure 10-30: SEM SEI image showing the planar film-like defect in this sample, together with the locations for the spectra shown in Figure 10-31




Figure 10-31: EDS spectra from the locations indicated in Figure 10-30

10.1.3.3 FGSS-10R-A




Figure 10-32: An SEM SEI image, providing an overview of the sample.




Figure 10-33: An SEM SEI image, providing a higher magnification view of the planar defect towards the centre of Figure 10-32.

10.1.3.4 FGSS-15R-A




Figure 10-34: An SEI SEM image, showing a planar defect, with a feature that appears as a darker (smoother) vertical stripe towards the centre of the image.




Figure 10-35: Spectra taken from A) the darker central region in Figure 10-34 and B) the lighter normal fracture region to the left of the dark band in Figure 10-34.




Figure 10-36: Low magnification SEI SEM image, showing a large planar defect (left), with a surface texture quite distinct from the normal fracture surface (right).




Figure 10-37: Wide area EDS spectra taken from the Figure 10-36. The red line was sampled from the left hand region, and the blue curve was sampled from the right hand region.

10.1.3.5 BGSS-11L-A



Figure 10-38: Optical image showing a shiny planar defect on the right hand edge of the test bar.

10.1.4 Samples Classified as “Thin Oxide”

10.1.4.1 FGSS-12R-C




Figure 10-39: SEI SEM image showing an overview of the sample, with some porosity, and some fissure-like features visible.




Figure 10-40: SEI SEM image showing a higher magnification view of the fissure-like feature towards the bottom of Figure 10-39.




Figure 10-41: SEI SEM image showing a cluster of small pores around a crevice.

10.1.5 Samples Classified as “Dark Pore”

10.1.5.1 FGSS-2R-C



Figure 10-42: Optical image showing dark, rounded, yet elongated features.

10.1.5.2 FGSS-8R-B




Figure 10-43: Optical image showing dark elongated feature on the left, with rounded edges.

10.1.5.3 BGSS-3R-A




Figure 10-44: An SEI SEM image showing an optically dark feature.




Figure 10-45: An SEI SEM image showing a higher magnification view of the edge of the feature presented in Figure 10-44, suggesting some sort of compliant layer coating the dendrite network.

10.1.6 Samples Classified as “Lap Defect”

10.1.6.1 BGSS-1R-B




Figure 10-46: An SEI SEM image providing an overview of the sample, appearing to show a low level of porosity.




Figure 10-47: An SEI SEM image showing the edge classified as a lap defect. On this half of the test bar, the defect is continuous with the mould-wall face of the sample.




Figure 10-48: An SEI SEM image showing a pore in the sample (right), with what appears to be oxide folds, and a possible trail of smaller pores extending to the left of the image.




Figure 10-49: An optical image of the sample, showing the lap-defect (top left) and a few pores on the order of 100 μm in size.

10.1.6.2 BGSS-3R-B




Figure 10-50: An SEI SEM image providing an overview of the sample, appearing to show a low level of porosity.




Figure 10-51: An SEI SEM image providing a close up of the lap defect. There appears to be a gradual transition between normal MVC microstructure at the left edge of the image, to a smooth morphology at the edge of the fracture surface.




Figure 10-52: An SEI SEM image showing what appears to be a network of shrinkage porosity.




Figure 10-53: An optical image of the sample, showing the lap-defect (top left) and a few pores on the order of 100 μm in size.

10.1.7 Samples Classified as “Exogenous Inclusion”

10.1.7.1 BGFS-11R-A



Figure 10-54: An optical image showing the inclusion towards the left of the image, together with some pores on the order of 100 μm in size.




Figure 10-55: An SEI SEM image representative of the majority of the fracture surface.




Figure 10-56: An SEI SEM image showing normal fracture surface, and an inclusion with poor electrical conductivity in the bottom right quadrant. Also marked are the sampling locations for the EDS spectra in Figure 10-57




Figure 10-57: The EDS spectra sampled from Figure 10-56, indicating that the inclusion is calcium based mineral.

10.2 Tabulated HPDC Sample Data

10.2.1 Full Gate, Fast Shot

Table 10-1: Mechanical data for test bars from the full gate, fast shot trial.

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGFS-1L-A	246.9	143.0	71.09	2.4	Porosity
FGFS-1L-B	224.4	142.0	66.98	1.9	Porosity
FGFS-1L-C	221.2	139.0	64.83	2.0	Porosity
FGFS-1R-A	245.7	140.0	84.68	2.7	Porosity
FGFS-1R-B	237.2	138.0	96.53	2.2	Porosity
FGFS-1R-C	216.4	133.0	77.61	1.9	Porosity
FGFS-2L-A	247.0	135.0	79.28	3.0	Ductile Shear
FGFS-2L-B	199.4	130.0	65.7	1.6	Porosity
FGFS-2L-C	201.8	128.0	69.12	1.7	Porosity
FGFS-2R-A	257.4	138.0	85.64	3.9	Ductile Shear
FGFS-2R-B	241.6	141.0	71.67	2.5	Porosity
FGFS-2R-C	216.0	131.0	77.43	2.0	Porosity
FGFS-3L-A	250.6	139.0	71.43	2.8	Porosity
FGFS-3L-B	217.7	137.0	69.2	1.8	Porosity
FGFS-3L-C	209.8	131.0	70.72	1.7	Porosity
FGFS-3R-A	248.9	142.0	79.6	2.9	Porosity
FGFS-3R-B	234.9	140.0	79.59	2.3	Porosity
FGFS-3R-C	216.3	132.0	78.48	1.9	Porosity
FGFS-4L-A	248.1	133.0	76.96	3.3	Ductile Shear
FGFS-4L-B	205.0	130.0	72.98	1.7	Porosity
FGFS-4L-C	211.5	128.0	68.12	2.1	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGFS-4R-A	208.2	135.0	87.58	1.6	Carbon Film
FGFS-4R-B	221.2	134.0	80.9	2.1	Porosity
FGFS-4R-C	210.0	128.0	76.46	1.9	Porosity
FGFS-5L-A	227.4	133.0	79.05	2.2	Carbon Film
FGFS-5L-B	207.7	132.0	66.92	1.8	Porosity
FGFS-5L-C	204.0	127.0	67.39	1.8	Porosity
FGFS-5R-A	248.0	138.0	82.37	3.0	Porosity
FGFS-5R-B	220.7	139.0	80.85	1.9	Lap Defect
FGFS-5R-C	210.4	130.0	75.2	1.7	Porosity
FGFS-6L-A	247.3	137.0	75.62	3.1	Porosity
FGFS-6L-B	215.2	130.0	67.49	2.1	Porosity
FGFS-6L-C	213.0	129.0	64.3	2.0	Porosity
FGFS-6R-A	230.4	141.0	68.09	2.0	Porosity
FGFS-6R-B	223.4	137.0	79.75	1.9	Porosity
FGFS-6R-C	219.2	129.0	75.77	2.3	Porosity
FGFS-7L-A	261.6	143.0	75.55	3.3	Ductile Shear
FGFS-7L-B	210.7	139.0	69.3	1.5	Porosity
FGFS-7L-C	209.5	133.0	68.22	1.7	Porosity
FGFS-7R-A	225.6	138.0	84.4	1.8	Porosity
FGFS-7R-B	238.9	141.0	75.79	2.5	Porosity
FGFS-7R-C	231.6	130.0	85.04	2.8	Porosity
FGFS-8L-A	243.3	138.0	65.74	2.8	Ductile Shear
FGFS-8L-B	229.1	134.0	68.92	2.5	Porosity
FGFS-8L-C	223.5	133.0	64.83	2.3	Porosity
FGFS-8R-A	240.5	140.0	63.31	2.7	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGFS-8R-B	222.9	137.0	79.98	1.9	Porosity
FGFS-8R-C	230.5	131.0	77.21	2.7	Porosity
FGFS-9L-A	228.3	137.0	69.24	1.9	Ductile Shear
FGFS-9L-B	202.8	130.0	77.79	1.6	Porosity
FGFS-9L-C	211.7	128.0	70.21	2.0	Porosity
FGFS-9R-A	194.3	137.0	85.73	0.9	Carbon Film
FGFS-9R-B	224.9	139.0	81.33	1.9	Porosity
FGFS-9R-C	179.0	131.0	70.22	0.9	Carbon Film
FGFS-10L-A	229.0	135.0	76.52	2.2	Ductile Shear
FGFS-10L-B	201.5	133.0	61.84	1.6	Porosity
FGFS-10L-C	203.3	127.0	71.64	1.8	Porosity
FGFS-10R-A	253.1	138.0	82.06	3.6	Ductile Shear
FGFS-10R-B	217.8	133.0	86.41	1.8	Porosity
FGFS-10R-C	220.3	131.0	80.03	2.2	Porosity
FGFS-11L-A	233.4	136.0	70.53	2.3	Porosity
FGFS-11L-B	202.4	131.0	69.63	1.7	Porosity
FGFS-11L-C	196.4	129.0	66.41	1.4	Porosity
FGFS-11R-A	246.9	137.0	79.61	3.0	Porosity
FGFS-11R-B	225.7	140.0	76.08	2.1	Porosity
FGFS-11R-C	232.2	132.0	75.55	2.7	Porosity
FGFS-12L-A	251.9	135.0	65.63	3.6	Porosity
FGFS-12L-B	207.8	135.0	64.29	1.7	Porosity
FGFS-12L-C	214.6	128.0	67.36	2.2	Porosity
FGFS-12R-A	251.3	135.0	85.19	3.6	Ductile Shear
FGFS-12R-B	209.5	136.0	79.03	1.5	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGFS-12R-C	222.8	131.0	71.55	2.5	Porosity
FGFS-13L-A	225.7	137.0	67.63	2.0	Carbon Film
FGFS-13L-B	196.4	129.0	63.18	1.5	Porosity
FGFS-13L-C	211.4	130.0	65.73	2.0	Porosity
FGFS-13R-A	198.3	138.0	81.56	1.1	Carbon Film
FGFS-13R-B	226.4	136.0	80.86	2.1	Porosity
FGFS-13R-C	228.1	111.0	36.56	2.4	Porosity
FGFS-14L-A	247.6	140.0	72.02	2.6	Porosity
FGFS-14L-B	213.8	138.0	65.04	1.8	Porosity
FGFS-14L-C	209.7	133.0	63.4	1.8	Porosity
FGFS-14R-A	252.3	139.0	75.67	3.4	Ductile Shear
FGFS-14R-B	227.7	141.0	71.77	2.0	Porosity
FGFS-14R-C	211.9	131.0	78.6	1.7	Porosity
FGFS-15L-A	242.5	135.0	74.47	2.7	Ductile Shear
FGFS-15L-B	209.0	133.0	60.81	1.9	Porosity
FGFS-15L-C	211.7	126.0	66.4	2.1	Porosity
FGFS-15R-A	211.5	139.0	82.15	1.4	Carbon Film
FGFS-15R-B	212.9	141.0	71.64	1.5	Porosity
FGFS-15R-C	222.1	130.0	75.99	2.3	Porosity

10.2.2 Blanked Gate, Fast Shot

Table 10-2: Mechanical data for test bars from the blanked gate, fast shot trial.

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGFS-1L-A	215.0	137.0	70.78	1.5	Porosity
BGFS-1L-B	260.5	152.0	77.13	2.8	Ductile Shear
BGFS-1L-C	246.5	145.0	73.52	2.5	Carbon Film
BGFS-1R-A	227.8	146.0	79.98	1.5	Porosity
BGFS-1R-B	250.1	150.0	85.85	2.0	Ductile Shear
BGFS-1R-C	262.0	154.0	69.06	2.6	Porosity
BGFS-2L-A	233.5	135.0	77.67	2.5	Ductile Shear
BGFS-2L-B	247.4	143.0	73.61	2.2	Porosity
BGFS-2L-C	226.6	141.0	72.84	1.8	Porosity
BGFS-2R-A	154.7	139.0	74.18	0.4	Porosity
BGFS-2R-B	267.2	150.0	75.55	3.1	Ductile Shear
BGFS-2R-C	259.4	148.0	73.08	2.8	Porosity
BGFS-3L-A	220.9	133.0	70.33	1.8	Porosity
BGFS-3L-B	250.3	149.0	72.6	2.5	Porosity
BGFS-3L-C	210.9	137.0	71.54	1.2	Porosity
BGFS-3R-A	211.3	138.0	80.72	1.4	Porosity
BGFS-3R-B	261.5	148.0	80.22	2.8	Ductile Shear
BGFS-3R-C	250.2	145.0	70.76	2.6	Porosity
BGFS-4L-A	199.8	135.0	70.14	1.2	Porosity
BGFS-4L-B	246.0	141.0	79.62	2.7	Porosity
BGFS-4L-C	211.1	136.0	77.51	1.5	Porosity
BGFS-4R-A	201.9	134.0	78.75	1.3	Carbon Film

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGFS-4R-B	251.2	146.0	79.61	2.4	Porosity
BGFS-4R-C	245.3	143.0	72.15	2.6	Porosity
BGFS-5L-A	231.7	135.0	68.54	2.4	Porosity
BGFS-5L-B	252.0	146.0	75.76	2.8	Porosity
BGFS-5L-C	245.2	140.0	69.89	2.8	Porosity
BGFS-5R-A	226.8	142.0	69.91	1.8	Porosity
BGFS-5R-B	258.2	147.0	78.99	2.6	Porosity
BGFS-5R-C	240.4	145.0	70.2	2.1	Porosity
BGFS-6L-A	244.7	135.0	75.38	3.0	Porosity
BGFS-6L-B	224.9	145.0	71.62	1.7	Porosity
BGFS-6L-C	210.4	136.0	73.22	1.5	Porosity
BGFS-6R-A	225.2	140.0	72.34	1.7	Porosity
BGFS-6R-B	245.6	146.0	75.19	2.2	Ductile Shear
BGFS-6R-C	199.7	141.0	67.78	1.2	Exogenous Defect
BGFS-7L-A	219.7	135.0	72.33	1.8	Porosity
BGFS-7L-B	255.2	147.0	71.48	2.8	Ductile Shear
BGFS-7L-C	222.1	137.0	73.03	1.8	Porosity
BGFS-7R-A	225.8	141.0	70.57	1.8	Porosity
BGFS-7R-B	252.2	144.0	76.41	2.7	Porosity
BGFS-7R-C	222.7	143.0	72.99	1.6	Porosity
BGFS-8L-A	233.0	137.0	74.09	2.3	Carbon Film
BGFS-8L-B	234.5	145.0	78.34	2.0	Porosity
BGFS-8L-C	244.7	139.0	70.09	2.9	Porosity
BGFS-8R-A	234.3	138.0	81.12	2.1	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGFS-8R-B	257.6	147.0	80.76	2.7	Porosity
BGFS-8R-C	221.3	143.0	76.22	1.6	Porosity
BGFS-9L-A	220.8	136.0	76.54	1.8	Carbon Film
BGFS-9L-B	253.4	145.0	69.01	2.5	Porosity
BGFS-9L-C	233.1	138.0	67.78	2.4	Porosity
BGFS-9R-A	223.3	141.0	73.39	1.8	Porosity
BGFS-9R-B	247.2	148.0	77.42	2.2	Ductile Shear
BGFS-9R-C	242.6	141.0	75.95	2.4	Porosity
BGFS-10L-A	195.5	130.0	75.02	1.3	Carbon Film
BGFS-10L-B	252.3	144.0	81.19	3.1	Porosity
BGFS-10L-C	244.9	135.0	73.9	3.3	Porosity
BGFS-10R-A	211.1	139.0	70.64	1.4	Porosity
BGFS-10R-B	260.9	146.0	76.91	3.1	Ductile Shear
BGFS-10R-C	248.4	138.0	80.33	2.9	Porosity
BGFS-11L-A	204.8	133.0	67.05	1.6	Porosity
BGFS-11L-B	231.6	144.0	72.85	1.9	Ductile Shear
BGFS-11L-C	226.6	132.0	78.82	2.2	Porosity
BGFS-11R-A	91.3	0.0	70.98	0.0	Exogenous Defect
BGFS-11R-B	249.9	145.0	74.25	2.3	Porosity
BGFS-11R-C	231.1	139.0	73.87	2.2	Porosity
BGFS-12L-A	205.7	132.0	80.54	1.4	Porosity
BGFS-12L-B	254.1	141.0	84.34	3.6	Porosity
BGFS-12L-C	232.6	136.0	69.44	2.4	Porosity
BGFS-12R-A	209.3	136.0	73.75	1.5	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGFS-12R-B	255.8	144.0	78.91	2.8	Ductile Shear
BGFS-12R-C	243.9	140.0	73.4	2.6	Porosity
BGFS-13L-A	184.7	132.0	64.54	1.1	Carbon Film
BGFS-13L-B	248.6	142.0	70.09	3.1	Porosity
BGFS-13L-C	245.4	133.0	72.28	2.4	Porosity
BGFS-13R-A	210.3	136.0	73.82	1.5	Carbon Film
BGFS-13R-B	253.0	147.0	78.04	2.5	Ductile Shear
BGFS-13R-C	239.4	140.0	78.99	2.3	Porosity
BGFS-14L-A	225.0	134.0	69.45	2.1	Porosity
BGFS-14L-B	250.0	142.0	73.64	3.1	Ductile Shear
BGFS-14L-C	231.2	134.0	73.95	2.3	Porosity
BGFS-14R-A	202.7	136.0	78.48	1.3	Porosity
BGFS-14R-B	235.1	144.0	74.45	1.8	Porosity
BGFS-14R-C	252.1	140.0	74.77	3.1	Porosity
BGFS-15L-A	209.6	135.0	72.18	1.5	Porosity
BGFS-15L-B	253.6	142.0	73.19	3.3	Porosity
BGFS-15L-C	242.7	136.0	70.49	3.2	Porosity
BGFS-15R-A	208.1	139.0	67.62	1.4	Porosity
BGFS-15R-B	265.9	144.0	77.82	3.7	Ductile Shear
BGFS-15R-C	244.0	141.0	70.9	2.7	Ductile Shear

10.2.3 Full Gate, Slow Shot

Table 10-3: Mechanical data for test bars from the full gate, slow shot trial.

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGSS-1L-A	211.8	137.0	67.94	1.7	Carbon Film
FGSS-1L-B	217.2	141.0	65.66	1.8	Porosity
FGSS-1L-C	216.1	135.0	69.81	1.7	Porosity
FGSS-1R-A	236.0	137.0	66.16	2.4	Carbon Film
FGSS-1R-B	219.5	140.0	77.12	1.7	Dark Pore
FGSS-1R-C	211.7	136.0	69.38	1.6	Porosity
FGSS-2L-A	239.2	137.0	70.3	2.7	Porosity
FGSS-2L-B	216.7	138.0	72.9	1.8	Porosity
FGSS-2L-C	191.2	132.0	69.34	1.1	Porosity
FGSS-2R-A	218.2	134.0	79.64	2.0	Porosity
FGSS-2R-B	203.8	138.0	70.35	1.2	Dark Pore
FGSS-2R-C	196.0	133.0	73.34	1.2	Dark Pore
FGSS-3L-A	245.8	136.0	74.35	3.1	Ductile Shear
FGSS-3L-B	194.7	132.0	81.05	1.3	Porosity
FGSS-3L-C	208.6	132.0	65.26	1.8	Porosity
FGSS-3R-A	241.5	136.0	73.64	3.0	Porosity
FGSS-3R-B	206.3	137.0	70.15	1.3	Porosity
FGSS-3R-C	203.5	133.0	70.4	1.5	Porosity
FGSS-4L-A	118.7	0.0	69.63	0.1	Carbon Film
FGSS-4L-B	200.7	145.0	64.48	1.1	Porosity
FGSS-4L-C	211.2	142.0	68.86	1.4	Porosity
FGSS-4R-A	218.9	136.0	68.78	1.9	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGSS-4R-B	219.8	144.0	71.45	1.7	Dark Pore
FGSS-4R-C	193.7	143.0	66.98	0.9	Dark Pore
FGSS-5L-A	228.2	135.0	72.52	2.2	Porosity
FGSS-5L-B	219.1	145.0	68.59	1.5	Porosity
FGSS-5L-C	219.2	138.0	73.46	1.8	Porosity
FGSS-5R-A	214.2	139.0	73.38	1.6	Carbon Film
FGSS-5R-B	204.3	142.0	72.6	1.2	Dark Pore
FGSS-5R-C	210.5	139.0	71.74	1.6	Porosity
FGSS-6L-A	234.6	135.0	73.67	2.5	Carbon Film
FGSS-6L-B	197.3	137.0	71.54	1.2	Porosity
FGSS-6L-C	222.6	137.0	72.49	1.9	Porosity
FGSS-6R-A	227.9	134.0	75.63	2.3	Ductile Shear
FGSS-6R-B	212.8	144.0	71.66	1.4	Dark Pore
FGSS-6R-C	210.2	132.0	68.42	1.9	Porosity
FGSS-7L-A	242.1	136.0	71.24	3.0	Porosity
FGSS-7L-B	227.5	140.0	70.34	2.0	Porosity
FGSS-7L-C	205.4	135.0	69.5	1.5	Porosity
FGSS-7R-A	239.9	136.0	69.27	3.0	Ductile Shear
FGSS-7R-B	191.9	137.0	75.83	0.9	Dark Pore
FGSS-7R-C	207.7	135.0	69.78	1.6	Porosity
FGSS-8L-A	231.9	134.0	75.13	2.4	Porosity
FGSS-8L-B	196.4	135.0	72.22	1.2	Dark Pore
FGSS-8L-C	220.3	130.0	54.42	2.5	Porosity
FGSS-8R-A	233.4	132.0	83.38	2.6	Porosity
FGSS-8R-B	84.4	0.0	76.87	0.0	Dark Pore

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGSS-8R-C	207.0	136.0	68.23	1.5	Porosity
FGSS-9L-A	194.5	132.0	77.32	1.2	Dark Pore
FGSS-9L-B	216.4	139.0	66.64	1.8	Porosity
FGSS-9L-C	215.5	134.0	68.54	1.9	Porosity
FGSS-9R-A	231.4	136.0	75.09	2.4	Porosity
FGSS-9R-B	218.9	140.0	73.31	1.7	Porosity
FGSS-9R-C	187.8	133.0	77.87	1.1	Dark Pore
FGSS-10L-A	243.0	136.0	68.65	3.1	Porosity
FGSS-10L-B	208.5	140.0	66.33	1.4	Dark Pore
FGSS-10L-C	213.4	133.0	68.11	1.8	Porosity
FGSS-10R-A	146.6	133.0	80.38	0.3	Carbon Film
FGSS-10R-B	89.3	0.0	71.86	0.0	Dark Pore
FGSS-10R-C	189.3	134.0	68.89	1.0	Porosity
FGSS-11L-A	232.6	136.0	71.31	2.4	Porosity
FGSS-11L-B	214.4	137.0	70.91	1.8	Porosity
FGSS-11L-C	208.1	134.0	72.77	1.5	Porosity
FGSS-11R-A	206.3	132.0	81.01	1.5	Carbon Film
FGSS-11R-B	191.1	137.0	69.73	1.1	Porosity
FGSS-11R-C	196.4	132.0	71.98	1.3	Porosity
FGSS-12L-A	234.3	136.0	68.13	2.5	Carbon Film
FGSS-12L-B	220.8	137.0	68.93	2.0	Porosity
FGSS-12L-C	219.8	132.0	73.44	2.1	Porosity
FGSS-12R-A	233.1	137.0	69.67	2.4	Porosity
FGSS-12R-B	193.8	135.0	73.83	1.1	Porosity
FGSS-12R-C	150.0	132.0	68.83	0.4	Thin Oxide

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
FGSS-13L-A	235.4	135.0	73.12	2.6	Porosity
FGSS-13L-B	208.6	137.0	73.47	1.5	Porosity
FGSS-13L-C	220.2	133.0	68.94	2.1	Porosity
FGSS-13R-A	238.2	138.0	71.63	2.7	Porosity
FGSS-13R-B	211.8	137.0	74.88	1.6	Porosity
FGSS-13R-C	211.0	132.0	74.71	1.7	Dark Pore
FGSS-14L-A	234.2	136.0	66.94	2.6	Ductile Shear
FGSS-14L-B	214.6	137.0	68.95	1.8	Porosity
FGSS-14L-C	209.5	133.0	65.91	1.8	Porosity
FGSS-14R-A	201.8	133.0	82.44	1.4	Carbon Film
FGSS-14R-B	187.5	134.0	78.07	1.1	Porosity
FGSS-14R-C	214.2	131.0	72.36	2.0	Porosity
FGSS-15L-A	219.0	133.0	70.31	1.8	Carbon Film
FGSS-15L-B	220.3	137.0	70.01	2.1	Porosity
FGSS-15L-C	207.1	133.0	69.21	1.6	Lap Defect
FGSS-15R-A	16.8	0.0	0	0.0	Carbon Film
FGSS-15R-B	142.9	134.0	72.17	0.3	Porosity
FGSS-15R-C	182.2	130.0	74.89	1.0	Dark Pore

10.2.4 Blanked Gate, Slow Shot

Table 10-4: Mechanical data for test bars from the blanked gate, slow shot trial.

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGSS-1L-A	244.7	136.0	69.34	3.3	Ductile Shear
BGSS-1L-B	214.9	128.0	128.86	1.8	Lap Defect
BGSS-1L-C	232.2	136.0	81.06	2.3	Porosity
BGSS-1R-A	138.3	131.0	76.75	0.3	Carbon Film
BGSS-1R-B	202.2	136.0	112.32	1.1	Lap Defect
BGSS-1R-C	233.6	140.0	80.74	2.2	Porosity
BGSS-2L-A	239.8	134.0	75.88	2.6	Porosity
BGSS-2L-B	213.5	141.0	71.35	1.5	Lap Defect
BGSS-2L-C	226.5	134.0	71.04	2.1	Porosity
BGSS-2R-A	235.2	137.0	80.83	2.4	Carbon Film
BGSS-2R-B	221.4	141.0	81.98	1.6	Porosity
BGSS-2R-C	222.7	140.0	75.01	1.8	Porosity
BGSS-3L-A	213.0	132.0	73.86	1.8	Carbon Film
BGSS-3L-B	218.6	139.0	76.71	1.6	Porosity
BGSS-3L-C	219.5	135.0	74.76	1.9	Lap Defect
BGSS-3R-A	238.7	136.0	74.91	2.7	Thin Oxide
BGSS-3R-B	143.3	142.0	80.5	0.2	Lap Defect
BGSS-3R-C	239.9	140.0	80.14	2.5	Porosity
BGSS-4L-A	225.6	133.0	69.41	2.3	Porosity
BGSS-4L-B	226.0	140.0	73.49	2.0	Porosity
BGSS-4L-C	230.3	136.0	74.75	2.4	Porosity
BGSS-4R-A	230.6	135.0	85.29	2.4	Carbon Film

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGSS-4R-B	210.2	143.0	80.02	1.3	Lap Defect
BGSS-4R-C	226.5	141.0	80.35	1.9	Porosity
BGSS-5L-A	223.7	134.0	75.98	2.1	Porosity
BGSS-5L-B	227.4	142.0	65.09	2.0	Porosity
BGSS-5L-C	226.5	135.0	74.48	2.2	Porosity
BGSS-5R-A	218.7	136.0	83.35	1.7	Carbon Film
BGSS-5R-B	224.3	143.0	77.22	1.7	Porosity
BGSS-5R-C	214.9	139.0	85.89	1.6	Porosity
BGSS-6L-A	134.7	134.0	72.07	0.2	Carbon Film
BGSS-6L-B	213.7	139.0	66.17	1.4	Porosity
BGSS-6L-C	207.6	135.0	67.98	1.5	Porosity
BGSS-6R-A	227.8	135.0	75.39	2.2	Porosity
BGSS-6R-B	232.8	143.0	79.24	2.0	Porosity
BGSS-6R-C	227.6	140.0	79.56	2.0	Porosity
BGSS-7L-A	214.2	132.0	73.44	1.5	Carbon Film
BGSS-7L-B	220.4	141.0	72.36	1.8	Porosity
BGSS-7L-C	222.4	132.0	76.74	2.0	Porosity
BGSS-7R-A	242.6	137.0	75.04	2.9	Porosity
BGSS-7R-B	219.9	139.0	78.11	1.7	Porosity
BGSS-7R-C	231.6	143.0	79.27	1.9	Porosity
BGSS-8L-A	216.1	134.0	69.35	1.6	Carbon Film
BGSS-8L-B	231.5	139.0	67.45	2.2	Porosity
BGSS-8L-C	211.1	134.0	71.38	1.7	Porosity
BGSS-8R-A	213.2	139.0	74.61	1.5	Porosity
BGSS-8R-B	215.0	143.0	82.23	1.4	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGSS-8R-C	225.6	141.0	77.96	1.8	Porosity
BGSS-9L-A	214.2	134.0	68.97	1.7	Carbon Film
BGSS-9L-B	213.7	141.0	68.75	1.5	Porosity
BGSS-9L-C	221.4	134.0	69.74	2.0	Porosity
BGSS-9R-A	200.0	138.0	73.88	1.1	Porosity
BGSS-9R-B	228.8	145.0	71.15	1.7	Porosity
BGSS-9R-C	218.6	142.0	74.67	1.6	Porosity
BGSS-10L-A	227.8	130.0	75.75	2.5	Porosity
BGSS-10L-B	219.4	138.0	74.12	1.8	Carbon Film
BGSS-10L-C	230.3	137.0	70.6	2.3	Porosity
BGSS-10R-A	241.3	140.0	74.58	2.6	Porosity
BGSS-10R-B	143.8	143.0	79.25	0.2	Porosity
BGSS-10R-C	229.9	143.0	81.66	2.0	Porosity
BGSS-11L-A	106.9	0.0	76.18	0.1	Carbon Film
BGSS-11L-B	220.6	142.0	66.55	1.6	Porosity
BGSS-11L-C	215.1	136.0	69.48	1.7	Porosity
BGSS-11R-A	237.3	141.0	75.02	2.5	Carbon Film
BGSS-11R-B	219.4	145.0	84.85	1.4	Lap Defect
BGSS-11R-C	235.2	142.0	79.12	2.2	Porosity
BGSS-12L-A	231.7	133.0	68.85	2.8	Porosity
BGSS-12L-B	215.1	139.0	72.47	1.6	Porosity
BGSS-12L-C	214.0	135.0	69.09	1.7	Porosity
BGSS-12R-A	242.5	139.0	82.3	2.7	Porosity
BGSS-12R-B	220.2	143.0	79.09	1.7	Porosity
BGSS-12R-C	233.4	140.0	79.04	2.2	Porosity

Sample ID	UTS (MPa)	0.2% Proof Stress (MPa)	Young's Modulus (GPa)	% Elongation at Break	Classification
BGSS-13L-A	229.7	131.0	72.71	2.7	Porosity
BGSS-13L-B	234.4	139.0	72.04	2.5	Porosity
BGSS-13L-C	208.4	135.0	68.37	1.6	Porosity
BGSS-13R-A	203.2	132.0	88.29	1.3	Carbon Film
BGSS-13R-B	209.3	25.0	-7.22	1.2	Lap Defect
BGSS-13R-C	241.9	142.0	73.9	2.5	Porosity
BGSS-14L-A	43.8	0.0	76.87	0.0	Carbon Film
BGSS-14L-B	217.2	139.0	67.79	1.8	Porosity
BGSS-14L-C	222.7	136.0	64.26	1.7	Porosity
BGSS-14R-A	237.2	136.0	79.45	2.5	Ductile Shear
BGSS-14R-B	221.9	146.0	75.12	1.6	Porosity
BGSS-14R-C	236.5	142.0	76.82	2.3	Porosity
BGSS-15L-A	241.3	134.0	69.78	3.4	Porosity
BGSS-15L-B	201.0	139.0	67.21	1.3	Porosity
BGSS-15L-C	221.1	137.0	60.72	1.7	Carbon Film
BGSS-15R-A	159.1	137.0	75.33	0.4	Carbon Film
BGSS-15R-B	232.6	145.0	69.58	2.1	Porosity
BGSS-15R-C	209.2	141.0	74.37	1.3	Porosity

10.3 Tabulated Simulation Control Options

10.3.1 Entrainment Algorithm Characterisation Simulations.

Table 10-5: Control options for simulations used for “Falling Droplet” case

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Adiabatic
Unresolved Bubble Clean-up	Inactive
Moving Object Model	Inactive
Momentum Advection	First Order
Heat Conduction	Inactive
Heat Advection	Inactive
Solidification	Inactive
Solidification Shrinkage	Inactive

Table 10-6: Control options for simulations used for “Top Gated Return Wave” and “Plunging Jet” cases.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Inactive
Unresolved Bubble Clean-up	Inactive
Moving Object Model	Inactive
Momentum Advection	First Order
Heat Conduction	Inactive
Heat Advection	Inactive
Solidification	Inactive
Solidification Shrinkage	Inactive

10.3.2 Controlled Partial Vacuum Casting Simulations.

Table 10-7: Control options for parametric investigation.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Adiabatic
Unresolved Bubble Clean-up	Active
Moving Object Model	Inactive
Momentum Advection	Second Order Monotonic
Heat Conduction	Full Energy Equation
Heat Advection	First Order
Solidification	Active
Solidification Shrinkage	Inactive

Table 10-8: Control options for simulating the physical experiments.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Adiabatic
Unresolved Bubble Clean-up	Active
Moving Object Model	Inactive
Momentum Advection	Second Order Monotonic
Heat Conduction	Full Energy Equation
Heat Advection	First Order
Solidification	Active
Solidification Shrinkage	Inactive

10.3.3 High Pressure Die Casting Simulations

Table 10-9: Control options for the shot-sleeve simulations.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Unsplit Lagrangian method
Bubble Model	Adiabatic
Unresolved Bubble Clean-up	Active
Moving Object Model	Active
Momentum Advection	Second Order Monotonic
Heat Conduction	Full Energy Equation
Heat Advection	First Order
Solidification	Active
Solidification Shrinkage	Inactive

Table 10-10: Control options for the main filling simulations.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Adiabatic
Unresolved Bubble Clean-up	Active
Moving Object Model	Inactive
Momentum Advection	Second Order Monotonic
Heat Conduction	Full Energy Equation
Heat Advection	First Order
Solidification	Active
Solidification Shrinkage	Dynamic with fluid flow

Table 10-11: Control options for the solidification simulations.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Inactive
Unresolved Bubble Clean-up	Active
Moving Object Model	Inactive
Momentum Advection	Second Order Monotonic
Heat Conduction	Full Energy Equation
Heat Advection	First Order
Solidification	Active
Solidification Shrinkage	Dynamic with fluid flow

Table 10-12: Control options for the short-shot simulations.

Model	Option
Viscosity Model	(k-ε) turbulence model
Surface Tension	Active
VOF Method	Split Lagrangian method
Bubble Model	Adiabatic
Unresolved Bubble Clean-up	Active
Moving Object Model	Inactive
Momentum Advection	Second Order Monotonic
Heat Conduction	Full Energy Equation
Heat Advection	First Order
Solidification	Active
Solidification Shrinkage	Inactive