Imperial College of Science, Technology and Medicine
Department of Earth Science & Engineering

Investigating the effect of mesostructure on the shock response of granular materials through numerical modelling

James G. Derrick

Submitted in part fulfilment of the requirements for the degree of
PhD degree of Imperial College London and
the Diploma of Imperial College, July 2018

Abstract

Shock compaction of granular matter is of great interest to many fields such as planetary science, shock-chemistry, and ballistics. Often, in planetary science particularly, the scales are too large to perform experiments. Models are needed that parametrise the porosity. These are developed by investigating the response of granular systems whilst controlling influential properties. Granular matter, however, possesses many potentially influential properties and numerical models are the best method available to investigate granular matter whilst controlling its many properties. This thesis considers the effects of mesostructure on the shock response of granular matter. The number of contacts present in a granular system has a distinct effect on its shock response and it is observed that more contacts can artificially stiffen a granular material's shock response. The shock compaction of a bimodal powder, is an important problem in recent research into the formation of asteroids. A joint numerical-experimental approach was taken to investigate shock response in this system at the mesoscale. The large grains, embedded in the fine surrounding powder were found to 'shield' the fine powder as the shock wave passed. The porosity vector was introduced to quantify this effect. Finally, the thesis examines the efficacy of semi-explicitly resolving more complex granular materials in simulations, with the lunar regolith as a case study. The method was found to reproduce an explicitly resolved sample's response best when the parametrised material was engineered to possess the same porosity as the bulk porosity (through introduction of explicit pore space). Changing the grain morphology did not affect the system on the bulk scale, but local changes were observed and a critical angle a grain face could have with a shock wave was found that generated the highest local temperatures and pressures.

Acknowledgements

I would first like to thank my supervisor, Dr. Gareth Collins, for his help with the work presented in this thesis. He has been central to my research's progression and I can not thank him enough for his contributions. I would also like to thank Dr. Tom Davison for providing many hours of help and advice when I needed it. In addition, his work with pySALEPlot has been integral to every plot in this thesis, and inspired me to create pySALESetup in the first place.

I would like to thank all my collaborators, with whom I have worked on the material in Chapters 4, 5, and 6. First, I would like to thank Dr. Jeff LaJeunesse for his time and assistance with the benchmarking of iSALE and CTH, as well as Dr. John Borg for his assistance on the same project and his feedback on my first paper. Second I would like to thank Dr. Michael Rutherford for his endless patience with me, and constant diligence in working through the material we produced (Michael on the experimental side and I on the numerical side) in Chapter 5. Similarly, I would also like to thank Dr. Dan Eakins and Dr. Dave Chapman for their extensive help at the ESRF, as well as their professional guidance when compiling the results into various papers. Finally I would like to thank Dr. Phil Bland for his guidance on the work presented in Chapter 6. I also thank the iSALE developers (www.isale-code.de) and the ESRF for their providing of the facilities necessary to perform the experimental work.

I thank everyone in AMCG for their help and support when it was needed. I would also like to thank the many members of Imperial College Sinfonietta, Imperial College String Ensemble, Imperial College Fellwanderers, the Imperial College Union minibus team, and the Central London Orchestra for being great sources of catharsis, and distraction, throughout the PhD.

Finally I want to thank my friends and family. First are my wonderful parents (Caroline & Jeremy) for enduring many hours of stress and thesis discussion as well as their constant help and love. In particular, my Mum helped me by proof-reading much of my thesis (my Dad handling the rest) and I am eternally grateful for it! Next, I want to thank

my sisters (Catherine & Cecilia), and cousin (Natalie) for their support, and friendship, throughout, as well as providing somewhere for me to live in the final months. Next, I would like to thank the many friends I have been able to lean on every step of the way. There are too many who deserve recognition to list here but I will try to anyway: Laura, Katy, Ethan, Jas, Dave R., Rachel, Simon, Branwen, Zoe, Dan, Jonathan, Adam, Helena, Kirstin, Annabel, Piaf, Ben C., Maddy, Jason, Nick, Liyi, Tom W., Slavi, Gemma, Edward, Sally, Charlie, Ben Y. and Roan. I want to thank my excellent friend Sam for his unwavering confidence in my ability from start to finish and being a friend that I could always rely on. Finally, I thank my wonderful friend Hannah for helping me with many aspects of the work in this thesis, including proof-reading, presentation, and advice (both personal and academic) despite doing her own PhD at the same time. Hannah has been an indispensable source of support from beginning to end.

I dedicate this thesis to my Dad (a professor at Manchester), for telling me not to do it in the first place.

JGD gratefully acknowledges support from EPSRC EP/M506345/1.

Contents

Abstract **i**

Acknowledgements **iii**

Glossary **1**

1 Introduction **4**

1.1 Publications 8

1.2 Declaration of Originality 9

1.2.1 Instances of other's work in this thesis 9

1.3 Copyright 10

2 Introduction & Background Theory **12**

2.1 Introduction 12

2.2 Standard Definitions 12

2.2.1 Shock Waves 12

2.2.2 Granular Materials 15

2.3 A Problem of Scale 17

2.3.1 Shock Compaction at the Microscale 18

2.3.2 The Macroscale 18

2.3.3 The Mesoscale 20

2.4 Mesostructure in Granular Systems 24

2.4.1 Existing Mesostructural Investigations 24

2.5 Research Objectives 28

3 Methodology **30**

3.1 Introduction 30

3.2 iSALE: a Shock Physics code 30

3.2.1 Introduction & Background 30

3.2.2 Governing Equations 31

3.2.3 Equation of State 34

3.2.4 The Constitutive Relation 38

3.2.5 Discretisation 43

3.2.6 Shock Modelling 46

3.2.7 pySALESetup 46

3.3 Benchmarking iSALE and CTH 53

3.3.1 Calculating the Hugoniot 53

3.3.2 Results of the iSALE-CTH comparison 59

3.4 Resolution Test 61

3.5 Summary 63

4 Effect of grain-arrangement in mono-disperse systems. 65

4.1 Publication 65

4.2 Introduction 65

4.3 Set-up and Methods 68

4.3.1 Quantifying Mesostructure in the Particle Bed 72

4.4 Results 76

4.4.1 Sensitivity of the Hugoniot to Z 76

4.4.2 Sensitivity of the Hugoniot to the Strength Treatment of Cells Containing Multiple Materials 79

4.5 Discussion 80

4.6 Conclusions 85

5 Investigating compaction of bimodal powders: meteorite precursors 88

5.1 Publications 88

5.2 Introduction 89

5.3 Methods 92

5.3.1 Analog Materials Choice 92

5.3.2 Geometrical Set-up & Material Parameters 92

5.4 Results 97

5.4.1 Pure Matrix 97

5.4.2 Single Chondrule 107

5.4.3 Multi-chondrule mixtures 110

5.5 Discussion 119

5.5.1 The Porosity Vector, $\boldsymbol{C}$ 120

5.5.2 Implications for Meteorites 125

5.5.3 Implications for Shock Physics 126

5.6 Conclusions 127

6 Aspects of simulating the shock response of a complex powder 130

6.1 Introduction 130

6.1.1 The Lunar Regolith 136

6.2 Methods 137

6.2.1 Set-up & Materials 137

6.2.2 Resolving the Particle Size Distribution Semi-Explicitly 141

6.2.3 Varying Matrix Uniformity 144

6.2.4 Varying Grain Morphology 145

6.3 Results 148

6.3.1 Varying DE_{min} 148

6.3.2 Explicit Void 153

6.3.3 Morphologies 157

6.4 Discussion 161

6.4.1 Semi-explicit method suitability for modelling complex powders . . 161

6.4.2 Polygon Matrix-Shock Interaction 162

6.4.3 Implications for Regolith Studies 170

6.5 Conclusions 172

7 Conclusion **175**

7.1 Summary 175

7.2 Future Work 178

Bibliography **179**

A Chapter 4: additional material **188**

B Chapter 5: additional material **193**

List of Tables

4.1 Material parameters used in these simulations, predominantly taken from (Borg and Vogler 2008). 72

4.2 Spearman's rank coefficients of all fields vs Z 81

5.1 All material models used 95

5.2 Full Sipernat EoS & strength parameters 96

6.1 All material, & model, parameters used to model Basalt. 142

List of Figures

2.1 Example of a shock wave, and release wave, in pressure-position space . . . 14

2.2 Example of a shock in granular matter 16

2.3 Link between micro- macro- and mesoscales 20

3.1 $\epsilon - \alpha$ model compaction curve 39

3.2 Example of discretisation in iSALE 44

3.3 Effect of changing artificial viscosity 47

3.4 Material number (and velocities) of the MesoParticles2D iSALE example set up in pySALESetup 52

3.5 Benchmarking set-up geometry 54

3.6 Longitudinal stress experienced by the particle bed at an instant in time and the transversely averaged equivalent plot. 55

3.7 Position-time and u_p - position for TransAvg and Stress Contour methods. 56

3.8 Benchmarking particle bed 59

3.9 iSALE-CTH Hugoniot comparison and percentage differences breakdown . 59

3.10 iSALE resolution study 62

4.1 Contact vector example. 66

4.2 Study setup geometry 70

4.3 Demonstration of porosity gradients and $K - Z$ plot. 75

4.4 Example particle beds with different Z (= 0.4, 1.5) 75

4.5 $U_s - u_p$ and $U_F - u_p$ for $Z = 0.4 - 1.5$ 76

4.6 $-\sigma_{LL} - \rho$ for $Z = 0.4 - 1.5$ 77

4.7 Longitudinal variation in P, $\sigma_L L$, ρ 78

4.8 iSALE Hugoniots with and without stiction 80

4.9 Hugoniots in ϕ - u_p space 83

4.10 Hugoniots in $|\sigma_{LL}| - u_p$ (left) and $U_F - \rho$ space (right) 84

5.1 ESRF experiment & model geometry 94

5.2 iSALE Sipernat Hugoniots; EoS parameters comparison 97

5.3 Experiment radiographs alone 99

5.4 Qualitative radiograph/simulation comparison for pure Sipernat. 100

5.5 Diagram showing location of 'gap' and effect on shock front 102

5.6 On-axis shock position against time of pure Sipernat 103

5.7 104

5.8 (a-c) Porosity lineouts shockward and leeward of the chondrule in the single chondrule scenario for both experiments and models. (d) Comparison between numerical and experimental on-axis compaction at each time. . . . 104

5.9 Longitudinal porosity lineouts 106

5.10 Single chondrule visual results 108

5.11 Single chondrule transverse porosity lineouts 109

5.12 Undivided/divided radiographs example 112

5.13 Large chondrule visual results—mono-disperse bed 113

5.14 Large chondrule visual results—tomography-initialised beds 114

5.15 Large chondrule on-axis shock positions with time 115

5.16 Small chondrule visual results—mono-disperse bed 117

5.17 Small chondrule visual results—tomography-initialised beds 118

5.18 Small chondrules on-axis shock position with time 118

5.19 Numerical & experimental C and matrix porosity used to calculate it . . . 121

5.20 Angular distribution of C from a particle bed of the 'large' chondrules . . . 123

5.21 Angular distribution of C from a particle bed of the small chondrules . . . 124

6.1 Explicit - implicit resolution of a granular material demonstration 132

6.2 Matrix spatial uniformity assumption demonstration 134

6.3 Lunar sample PSD approximation 138

6.4 Bulk model set-up geometry 139

6.5 The PSD used to investigate the implicit vs explicit cut-off 144

6.6 Initial geometries of the simulations varying the implicit/explicit cutoff . . 145

6.7 The grain creation algorithm and example grains 147

6.8 The bulk responses of four simulations with varying explicit/implicit cut-offs150

6.9 The bulk, and GSC, responses of four simulations with varying explicit/implicit cut-offs . . . 152

6.10 All bulk responses of simulations with varying matrix non-uniformity. . . . 154

6.11 All GSC responses for simulations with varying matrix non-uniformity. . . 155

6.12 Peak temperatures and pressures of all samples, with varied matrix non-uniformity . . . 156

6.13 All bulk morphology variation responses. . . . 157

6.14 All GSC responses to morphology variation. . . . 159

6.15 Peak temperatures and pressures of all morphology variations. . . . 160

6.16 The bulk, and GSC, responses of three simulations with varying amounts of explicit void compared to the equivalent fully explicit results . . . 163

6.17 Grain shape construction demonstration . . . 164

6.18 Peak temperatures and pressures of three grain face orientations. . . . 166

6.19 Peak pressure in AoE variation with θ. . . . 167

6.20 Peak temperature in AoE variation with θ. . . . 168

A.1 $U_F - Z$ for all V_I . . . 189

A.2 $u_p - Z$ for all V_I . . . 190

A.3 $\rho - Z$ for all V_I . . . 191

A.4 $-\sigma_{LL} - Z$ for all V_I . . . 192

B.1 Analytical polycarbonate Hugoniot compared with experimental data . . . 194

B.2 Analytical copper Hugoniot compared with experimental data . . . 195

B.3 Analytical aluminium Hugoniot compared with experimental data 196

Glossary of Terms

This thesis uses a large amount of field-specific terminology and to avoid confusion, this glossary collects some of the more common terms and provides their definitions as a reference.

cppd Cells per particle diameter (see Chapter 3). This is a common metric that is used to refer to the number of cells used to resolve a grain. The grain in question is assumed to be circular and if it is non-circular then this applies to the equivalent circle with the same area as the grain.

stiction The intergranular resistance to shear. In Eulerian hydrocodes, when two grains come into contact their boundary will be manifested by a region of cells containing some of each grain. The grains do not experience friction in a typical sense; friction manifests as the resistance of the mixed cells to shear forces. This is 'stiction' (Borg and Vogler 2013).

PSD Particle Size Distribution. The distribution of particle frequency as a function of size. Frequency is commonly given in terms of the percentage weight of the whole sample. In Chapter 6, however, the material used is assumed to possess constant density, and simulations are 2D, so Area percentage is used instead.

Explicit resolution of a granular material When modelling a granular material and the grains are 'explicitly resolved', each grain exists in the simulation as an indepen-

dent object that can interact with other objects. Most numerical mesoscale studies resolve their granular materials explicitly. Borg and Vogler (2009), Derrick et al. (2018b), and Benson and Nellis (1994), for example.

Implicit resolution of a granular material When modelling a granular material and the simulation's cell size is greater than the grain size, a porosity model must be used, and any cells containing the granular material are treated as continuous. This is referred to as 'implicitly resolving' the granular material (Collins et al. 2011; Wünnemann, Collins, and Melosh 2006).

Semi-explicit resolution of a granular material When modelling a granular material that is sufficiently complex that implicitly resolving it would not capture the bulk response sufficiently and resolving it purely explicitly would be too costly, a third method can be used: resolving the material semi-explicitly. This resolves the large grains explicitly, and the small grains implicitly. The exact size range of the explicitly resolved grains is defined by the choice of DE_{min}.

$\boldsymbol{DE_{min}}$ Diameter of the smallest grain that is explicitly resolved when semi-explicitly resolving a granular material. This is an independent variable that is chosen before a simulation.

$\boldsymbol{DI_{max}}$ Diameter of the largest grain that is not explicitly resolved when semi-explicitly resolving a granular material. This is a dependent variable that depends on the particle size distribution as well as DE_{min}. In the bimodal mixture discussed in Chapter 5 $DE_{min} >> DI_{max}$, however, in a continuous PSD like that discussed in Chapter 6 $DI_{max} \approx DE_{min}$.

GSC 'Grain Size Category'. In Chapter 6 a grain size category refers to a selection of grains that fit into a certain size range. In a semi-explicitly resolved granular material, the implicitly resolved grains (the matrix) are considered to be in the

matrix GSC, although this is often referred to simply as 'the matrix'.

matrix In Chapter 5 'matrix' has a particular definition. In Chapter 6 matrix is used as the term for all implicitly resolved material, regardless of size.

Fines The smallest GSC in a simulation. In Chapter 6 'fines' are any grains that are implicitly resolved $D \leq DI_{max}$ (i.e. in the matrix).

Krumbein phi/ϕ_k A dimensionless, logarithmic, grain size metric commonly used in particle size distributions. $\phi_k = -\log_2\left(\frac{D}{D_0}\right)$, where D is the diameter of the sphere that has the same volume as the grain in mm[1], and D_0 is a reference length equal to 1 mm (so that the equation remains dimensionally consistent).

[1]In 2D it is the diameter of the circle with the same area.

Chapter 1

Introduction

High-velocity, collision, events are common throughout the known world and are present in a range of scenarios, from ballistics to asteroid impacts. These events are typically very destructive and have significant consequences for the materials involved. Asteroid impacts, for example, have been a key process in planetary formation and geological evolution of the Earth (Marchi et al. 2014). An understanding of the individual processes involved in a collision event is necessary to interpret the damage sustained by the colliding materials or to predict the outcome of a future event.

The primary process involved in an impact event is 'shock compression'. This occurs when a compression wave (e.g. sound wave) is driven through a material at a velocity greater than its sound speed. These are known as 'shock waves' and cause 'shock compression'. The passage of a shock wave through a material permanently deforms, compresses, and damages it. A material's response to shock is typically defined by its state immediately after the passing of the shock front; this is known as its 'shock state'. This state will be different for different shocks and by collating all possible shock states of a material its full response to shock is documented. This is often represented by a line plot in thermodynamic state variable space known as the 'Hugoniot' after P. H. Hugoniot, the

French engineer.

There are many studies into how specific materials are affected by shock compression and they broadly fall into three categories: numerical, experimental, and observational. Observational studies examine materials post-shock and infer the details of their shock compression from the end result. This approach is limited by the fact that shock compression is an inherently destructive process and the sample material can be challenging to recover for study. Impact crater, or meteorite, studies typically fall into this category, such as Marchi et al. (2014) or Forman et al. (2017), for example.

Experimental studies generally consider the shock response of small samples within well constrained scenarios. The shock response is typically extracted in shock-particle velocity space through the use of gauges inserted into the material prior to impact, or through surface-level velocimetry tests, such as VISAR (Velocity Interferometer System for Any Reflector). Experiments are able to examine the response of a material directly, which none of the other methods can do. They are limited, however, in their size and data extraction methods. It is typically unfeasible to perform shock experiments that are greater in size than a few metres because of the substantial energies, velocities, and pressures involved. Experiments are typically restricted to representative material samples, except in specific cases. In addition, directly extracting fields besides the shock and particle velocities within a sample during shock compression, is extremely challenging. In order to address these weaknesses, experiments are often performed in tandem with numerical modelling.

Numerical studies are unhindered by the size constraints of experiments or the accessibility of fields. They are capable of simulating small and large impact events with equal clarity. They are able to extract detailed information about the dynamic spatial distributions of fields (such as temperature) that are inaccessible by the other methods. Numerical models are therefore ideal for examining the dynamic response of a material to shock

compression, because they can observe the exact processes involved, as they occur.

Whilst numerical simulations are capable of resolving large-scale events (such as meteorite impacts) in full, it is challenging to simultaneously model processes on length scales that are several orders of magnitude apart. For example, when considering the effect of a 10 km diameter asteroid striking a desert of sand grains with a characteristic size of $\sim$ 1 mm, the simulation can not resolve every grain individually because it would be too computationally expensive. To perform simulations of large-scale impact events the porosity must be parametrised in cells containing porous matter. 'Porosity models' are needed that simulate the compaction of the porous material and ensure that cells in a simulation behave appropriately when they contain porous matter.

A porosity model is a continuum description of porosity compaction rather than a literal 'model of porosity', although this thesis will refer to it as a porosity model for ease. To develop porosity models, it is necessary to build an explicit understanding of the shock compaction process in granular matter and how properties of the granular materials affect it. For example, the stress at which grains fracture when exposed to shock compression, changes depending on the size of the grains involved (Neal, Chapman, and Proud 2012).

To investigate the effect of each property of a granular material (like grain size) on its shock compaction, an investigation must consider each property individually. The investigation can discern how much of an effect that property has on the shock compaction, linking the microscale property to a macroscale effect. To do so requires that all other potentially influential properties are restricted or controlled. The properties of granular materials, are extremely complex and it is almost impossible to control every property in an experiment. Numerical simulations, however, are inherently flexible and are able to do this. They are even able to simultaneously consider the bulk shock response of a material whilst resolving the shock compaction process at the microscale. These simulations resolve many individual grains explicitly (as clusters of cells), and thus consider how they interact as

a bulk (macroscale response) whilst simultaneously implementing microscale properties (e.g. grain fracture), and observing how these properties affect the response at each scale. These are known as 'mesoscale' studies.

Previous investigations into the shock compaction of granular materials have considered a variety of properties, such as grain size distribution (Borg and Vogler 2009), grain shape (Benson et al. 1994), and grain fracture (Lammi and Vogler 2012). Few studies, however, have considered how properties unique to the mesoscale might affect the shock response. One mesoscale aspect in particular is the presence of 'mesostructure'. Mesostructure is structure within a particle ensemble that consists of contacts between grains forming networks across the granular material. These networks may act to to resist compaction, or facilitate it, without altering the number of grains or their properties. This thesis is an investigation into how mesoscale properties influence the shock response of a granular material, with a particular emphasis on the mesostructure.

This thesis uses numerical models to examine several aspects of mesostructure and how they affect the shock response of a granular material. The primary code used is the semi-open source shock physics code, iSALE. Chapter 2 provides an overview of the literature relevant to this work. Chapter 3 presents the methodology used throughout the thesis as well as a benchmark of iSALE with the well-established shock physics code CTH. Chapter 4 aims to assess how the bulk response (shock and particle velocities) of a granular system is affected by the arrangement of grains. It does this by considering a monodisperse granular system of a well-studied granular material, tungsten carbide. Chapter 5 investigates how the bulk shock response is affected by a monodisperse, bimodal, granular system—large grains embedded in a matrix of fines. The chapter presents work interrogating the shock compression processes of bimodal powder mixtures in the context of asteroid formation. In particular it investigates the heterogeneity in compaction of this heterogeneous material. The study was performed as a joint numerical-experimental in-

vestigation and the chapter presents some of the first explicit mesoscale comparisons with a hydrocode in the field. Finally, Chapter 6 considers more complex granular systems that possess poly-disperse grains and complex morphologies. The chapter specifically considers the suitability of the simulation technique used in Chapter 5 (first used by Bland et al. (2014)) for modelling complex powders, with the lunar regolith as a case study. The chapter considers in particular how the set-up can change the temperatures and pressures experienced by the powder.

1.1 Publications

Papers published:

- *Derrick JG, LaJeunesse JW, Davison TM, et al.* Mesoscale simulations of shock compaction of a granular ceramic: effects of mesostructure and mixed-cell strength treatment, Modelling and Simulation in Materials Science and Engineering, ISSN:0965-0393 (Derrick et al. 2018b).

- *Derrick JG, Rutherford ME, Davison TM, et al.*, Interrogating Heterogeneous Compaction of Analogue Materials at the Mesoscale Through Numerical Modeling and Experiments, 20th Biennial Conference of the APS Topical Group on Shock Compression of Condensed Matter, Publisher: AIP Publishing, ISSN: 1551-7616 (Derrick et al. 2018a)

- *Davison TM, Derrick JG, Collins GS, et al.*, Impact-induced compaction of primitive solar system solids: The need for mesoscale modelling and experiments, Procedia Engineering, ISSN:1877-7058 (Davison et al. 2017)

- *Rutherford ME, Chapman DJ, Derrick JG, et al.*, 2017, Probing the early stages of shock-induced chondritic meteorite formation at the mesoscale, Scientific Reports,

Vol:7, ISSN:2045-2322 (Rutherford et al. 2017)

- *Rutherford ME, Chapman DJ, Eakins DE, Derrick JG, and Collins GS*, (accepted 1st December 2018) "Insights into local shockwave behavior and thermodynamics in granular materials from tomography-initialized mesoscale simulations", Journal of Applied Physics.

Conferences presented at:

- Pressure Energy Temperature & Extreme Rates; PETER (2015, 2016)
- Meteoritical Society (METSOC) workshop (2016)
- High-Velocity Impact Symposium; HVIS (2017; poster)
- Shock Compression Condensed Matter; SCCM (2017)

1.2 Declaration of Originality

I confirm that everything presented in this thesis is my own, except where clearly referenced, or otherwise noted.

1.2.1 Instances of other's work in this thesis

Gareth S Collins and Thomas M Davison have assisted in developing the themes, and ideas, presented throughout this thesis. In particular, Thomas M Davison created the plotting module pySALEPlot, which I have used in every chapter of this thesis and used as an inspiration for pySALESetup. The rest of the work, however, has been largely performed by myself, except where detailed below.

Chapter 3

The CTH simulations that were used as part of the benchmark were performed by Jeff LaJeunesse, and are published in Derrick et al. (2018b).

Chapter 4

The coordination number was a joint idea between myself and Gareth S Collins and is published in Derrick et al. (2018b).

Chapter 5

The experimental work in this Chapter was performed by myself, Michael E Rutherford, David J Chapman, Daniel E Eakins, Lukasz Farbaniec, and Joao P Duarte. The porosity vector was a joint creation between myself and Gareth S Collins. The work presented in this chapter is partly published in Derrick et al. (2018a), Rutherford et al. (2017), and a newly accepted paper by Rutherford *et al.*. An additional paper using this work is currently under review. This information is up to date as of December 2018.

Chapter 6

The work in this chapter is all my own, but involved significant direction and assistance from Gareth S Collins.

1.3 Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon

it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Chapter 2

Introduction & Background Theory

2.1 Introduction

The material in this Chapter is designed to give a reader the background required to understand the motivations for the research presented in later chapters as well as providing a brief review of the literature and explanations of some of the basic terms that are frequently used. Each subsequent Chapter will include a small introduction that covers concepts and motivations particular to it.

This Chapter begins by defining a series of terms, phrases, and concepts, common in granular shock physics and then goes on to detail where this thesis fits into the field.

2.2 Standard Definitions

2.2.1 Shock Waves

A shock wave is a discontinuous stress wave that progresses through a material at velocities greater than the sound speed. A shock wave carries energy, and its passage through

solid matter typically causes permanent damage to the material. Shocks are common throughout the known world, particularly in high-velocity events such as impacts, explosions, eruptions, etc.. Events such as these are often large in scale, and have significant consequences for their surroundings like building destruction, loss of life, or even total extinction (Marchi et al. 2014; Neal 2012; Schulte et al. 2010). This makes understanding these events something of great interest, and at the heart of the processes involved are shock waves.

Information travels in materials at the speed of sound in that material. A shock wave, however, travels faster than the sound speed. Material that is ahead of the shock wave can not react before the shock front arrives. This creates a build-up of material. Shock waves are characterised by this almost instantaneous rise in pressure and density, which persists until the material is 'released' (Figure 2.1). The material is 'released' upon its expansion post-shock compression. This typically occurs when the shock wave encounters a boundary with another material that the original material can expand into (air, for example), or a free surface. The expansion (and release) of the original material begins at the boundary and propagates inwards, releasing the pressure the original material was under (Figure 2.1). The material's state after the shock, but prior to its release, is known as a 'shock state'. This state is considered to be in equilibrium and does not undergo any other thermodynamic processes as part of the shock process (Petel and Jetté 2010). The shock response of a material is characterised by collating all its possible shock states. This collection is typically visualised by plotting each state as a point on a (usually 2D) state variable plot; pressure - volume, for example. A plot of all the loci of possible shock states of a material is known as a 'Hugoniot'. The Hugoniot is not a thermodynamic path, although it is typically plotted in thermodynamic spaces. Whilst it is common in thermodynamics to consider the standard state variables P, V, T, the Hugoniot is typically presented in shock velocity - particle velocity space because these two fields are the easiest to directly measure experimentally (Kerley 2013). Temperature and pressure

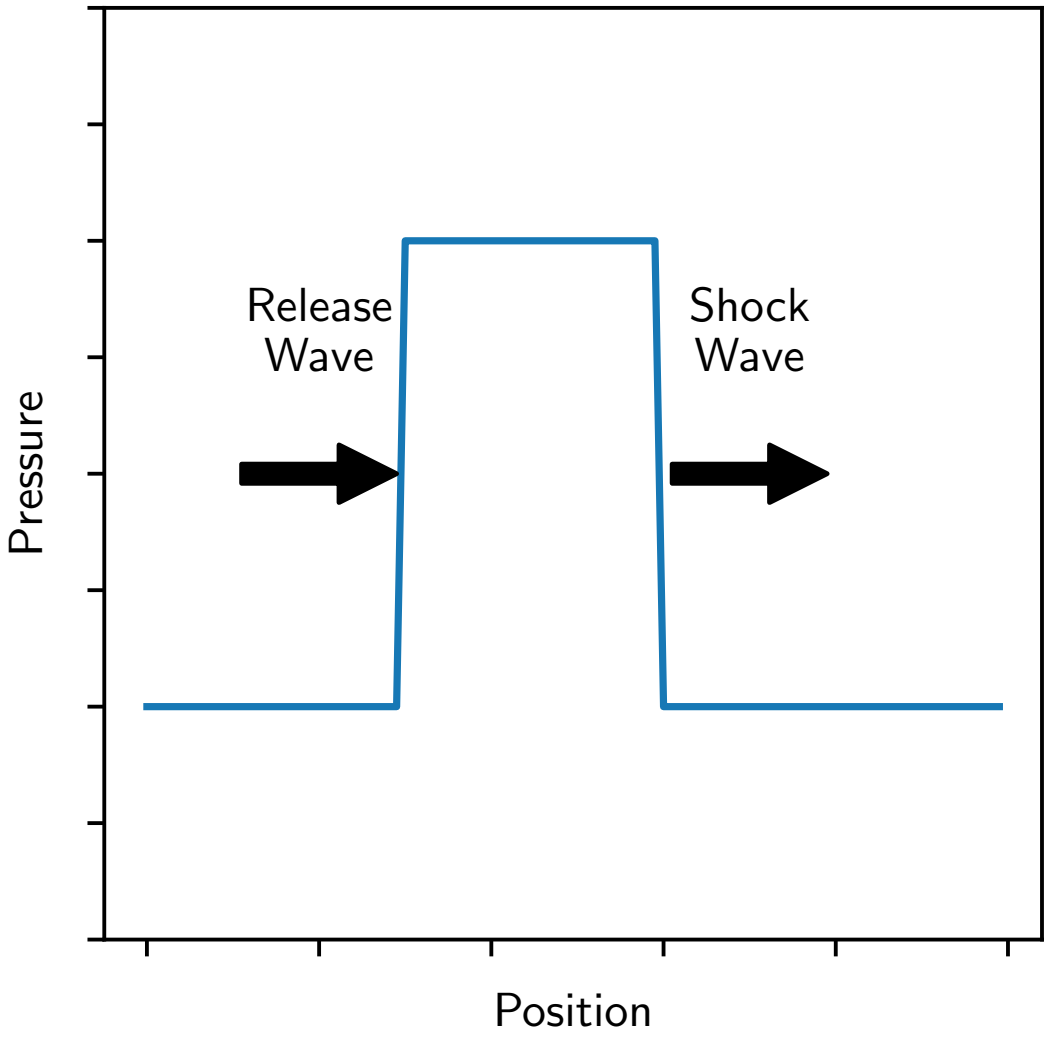


Figure 2.1: A basic example of a shock wave and its release wave plotted in pressure-position space. The shock wave is characterised by its discontinuous rise in pressure. The release wave is not necessarily sharp in real systems; its shape depends on the nature of the surface it is generated from and the material it is in.

must typically be calculated indirectly. 'Particle velocity' refers to the material velocity behind the shock and does not have a special meaning in a granular context.

2.2.2 Granular Materials

A granular material is typically characterised as an ensemble of solid, particles greater than one micron in size (P. de Gennes 1999). They are ubiquitous across Earth in beaches, deserts, soils etc., as well as throughout the Solar System in asteroids, comets and covering planetary surfaces (Bland et al. 2014; Cheng et al. 2016; Melosh 1989; Ormel, Cuzzi, and Tielens 2008; Wünnemann, Collins, and Melosh 2006). Their ubiquity has meant that they are inextricably involved in many vital processes across the world including construction (e.g. pilings) and defence (e.g. land mines and sand bags). This has made understanding their physical properties invaluable.

However, the general response of granular matter to shock compaction is still poorly understood (Borg et al. 2009; Fenton, Asay, and Dalton 2015). This lack of understanding is a result of significant complexity involved. For example, in real systems it is atypical that granular matter is made up of just one material. It is far more likely that it will contain grains of many shapes and sizes, with varying intra-granular porosities, and of a non-uniform density. For example, the lunar regolith's particle size distribution covers multiple orders of magnitude (Heiken, Vaniman, and French 1991). The number of variables (grain size, shape, composition, fracture, arrangement, number, etc.) in a single shock interaction is large and few are independent of all the others. A full understanding will come when the effects of each parameter are fully understood and can be used to analytically predict the outcomes of the shock compaction of an arbitrary granular material.

It is generally understood that the stresses brought on by shock compression force the grains in a powder together, crushing the pore spaces between them (Wünnemann, Collins, and Melosh 2006). The collapse of these spaces generates secondary shocks. The constructive interference of the secondary shock waves with the primary shock wave cause

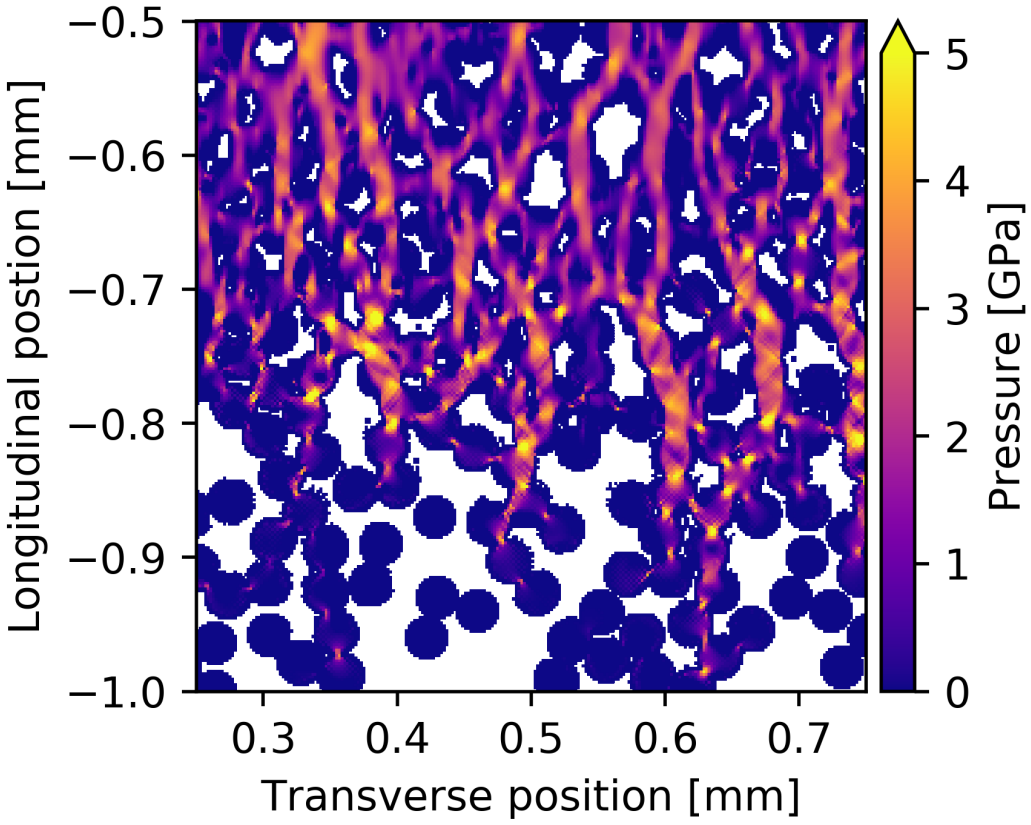


Figure 2.2: Example of shock propagation in a 2D granular material simulation (taken from Chapter 4); the shock is travelling down the image. The shock front is not linear and varies in position transversely. The network of stress bridges, or force chains, is clearly visible as 'fingers' of high pressure across the bed.

localised temperature and pressure spikes as a direct consequence (Bland et al. 2014). This compaction process is the root of many features of granular shock response. For example, pore collapse generates higher temperatures and pressures than the equivalent continuous material subjected to the same bulk stress (Derrick et al. 2018b; Kowitz et al. 2013).

A particularly important consequence of shock compaction in granular materials is the existence of stress-bridges, or 'force chains' at low stresses. Force chains are created because shock waves are transmitted across particle beds through grain-grain contacts. As the powder is compacted, its grains form a network of contacts, which the shock prop-

agates through. The shock propagates along the shortest available path in the network and can leave large portions of the grains relatively untouched (Borg and Vogler 2008; Dwivedi, Pei, and Teeter 2015). Some grains are not part of these networks at all and experience significantly lower stresses. The networks of stress are known as force chains (Borg and Vogler 2008; Dwivedi, Pei, and Teeter 2015; Zheng and Wang 2010). Force chains dictate the shock compaction process for stresses up to those required for complete grain deformation. Shock stresses above this point are sufficient to deform the grains as they pass, which becomes the dominant mechanic in shock compaction.

As a direct consequence of pore collapse and force chain formation, the shock wave is not as sharply defined as in continuous matter. The discontinuous pressure wave is spread out over several grains because it propagates along pathways made by grain contacts; some pathways will be longer than others and some will be shorter. The front is not always at the same longitudinal position for different transverse positions (to shock direction). It is typically several particle diameters thick, and distinctly non-planar because of this (Borg and Vogler 2008; Derrick et al. 2018b; Nesterenko 2001; Rutherford et al. 2017).

2.3 A Problem of Scale

In a typical shock compaction event (such as a meteorite impact) the materials involved are rarely 'perfect'. They are not uniformly dense, geometrically regular, or completely pure. They typically contain weaknesses, defects and impurities. These irregularities often have characteristic length scales that are orders of magnitude smaller than the event they are involved in and these irregularities dictate the bulk response of the event. For example, in the event of a meteorite (diameter > 1 km) striking a desert the characteristic length scale of the event will be on the order of the meteorite's size. The shock compaction of the sand, however, operates on a length scale that is approximately the size of a sand grain. Despite the disparity in length scales, the shock compaction processes in the sand

will have a direct effect on the meteorite's impact processes as the event unfolds. For example, the sand will respond differently if it is saturated with water than if it is dry (Perry et al. 2015; Perry et al. 2016). The combined effect of many microscale properties will have a direct consequence on the macroscale event.

2.3.1 Shock Compaction at the Microscale

In the context of granular materials the microscale is the scale at which single-grain effects are resolved; it is typically referred to as the 'particle-scale' in the literature (Li, Yu, and Li 2009). Microscale length scales depend on the grain sizes involved, however, its lower bound is of the order of microns ($L_{micro} \gtrsim 1\mu$m) and typical scales in the literature are on the order of mm. Below this scale, atomic features become relevant (nanoscale), such as Brownian motion (P. de Gennes 1999). The microscale properties have a direct impact on how the bulk granular material behaves. For example, fracture strength is a microscale property and the stress at which grains fracture will change the bulk response (Lammi and Vogler 2012). Force chains are another example of factors that determine the bulk response: the contact mechanics, which dictate how grains interact, are a microscale property of the powder (Sun, Wang, and Hu 2009) and changing them will alter how the force chains form, which will affect how the bulk material responds (Derrick et al. 2018b). The contact mechanics of two grains alone, however, are insufficient to examine how changing the mechanics affects the bulk response; a representative sample must be considered.

2.3.2 The Macroscale

The macroscale is the 'bulk' scale and whenever this thesis refers to the 'bulk' response it is equivalent to the macroscale response, unless stated otherwise. It is large compared to microscopic properties, which are unresolved ($L_{macro} >> L_{micro}$). In the context of

granular dynamics, a macroscale investigation treats a granular mixture as a continuum (Li, Yu, and Li 2009), often using a spatial resolution (through imaging or numerical models) that is not capable of explicitly resolving each grain. For example, a typical simulation of a 10 km diameter meteorite impact would use a cell-size on the order of 100 m, implying individual sand grains involved in that impact can not be resolved (Melosh 1989). In simulations where the granular material can not be resolved, it is common practice to treat cells containing the material as continuous and to employ a 'porosity model' to the cell that mimics the effects of compaction (Collins et al. 2011; Wünnemann, Collins, and Melosh 2006).

Porosity models are a continuum description of compaction; i.e. they ensure that a porous cell behaves as if it did contain a granular material (Wünnemann, Collins, and Melosh 2006). This is also known as 'implicitly' resolving a granular material. The $P - \alpha$, $\epsilon - \alpha$ and 'snowplough' models are examples of porosity models (Wünnemann, Collins, and Melosh 2006) (See Chapter 3 for more details on the porosity model used in this thesis: the $\epsilon - \alpha$ model). Porosity models are developed by observing the compaction of granular material, and empirically devising relationships that describe its compaction accurately (Wünnemann, Collins, and Melosh 2006). In order to further develop accurate porosity models, however, compaction must be observed at the microscale, as well as the macroscale. For example, Wünnemann, Collins, and Melosh (2006) observed a discrepancy between their model and experimental load-compaction data that could not be explained by considering the bulk exclusively. At low compaction, the sand exhibited a small upward kink in the load-compaction curve that was not seen in the $\epsilon - \alpha$ results. An examination of the particle bed at the microscale would be required to understand the processes that contributed to this effect.

Microscale effects are inherently unresolvable at the macroscale, but macroscale models require accurate continuum models to operate. A link between the two is required, and

to do this a new scale is introduced, the mesoscale (Figure 2.3).

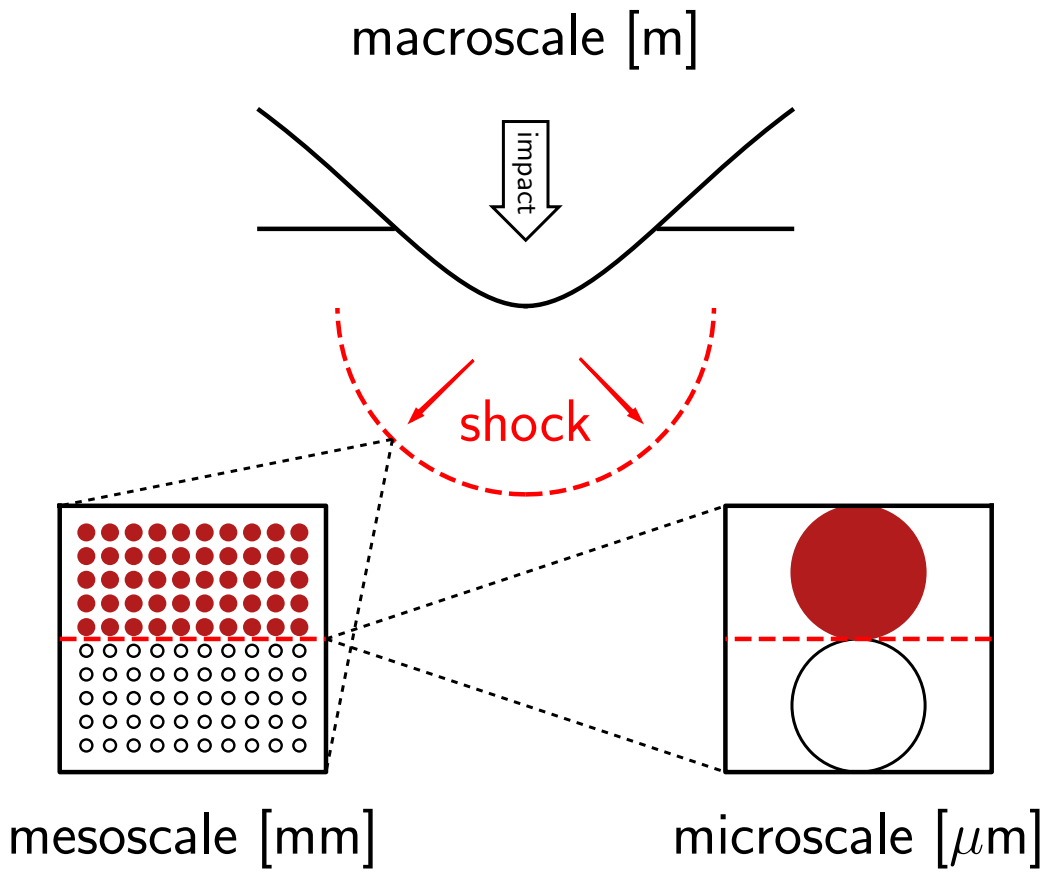


Figure 2.3: Explanation of each scale. Macroscale events are too large to resolve the microscale and the microscale is too small to predict the macroscale response. The mesoscale is introduced as a link between the two, although it is typically limited to simplified conditions in order to achieve the link.

2.3.3 The Mesoscale

The 'mesoscale' is an intermediate scale between macro and micro. It can be loosely defined as a level at which the micro- and macroscale responses can be observed simultaneously. In granular shock physics, a study at the mesoscale resolves a sufficient number of microscale systems (grains) such that the bulk response can also be considered. This is also known as 'explicitly' resolving a granular material. Simulating processes at the mesoscale can be computationally demanding, however, because the simulation must re-

solve spatial features at both the micro- and macroscales, which are often many orders of magnitude apart. To avoid extremely costly simulations, like resolving every grain of sand in a meteorite impact, mesoscale simulations typically examine simpler, smaller, systems (Rutherford et al. (2017) for example). They commonly operate on a scale that is large enough that the macroscale is resolved, but small enough that they do not need to resolve millions of grains at a time. The precise physical length scales will depend on the size of the grains involved, but the number of grains resolved is typically on the order of several hundred to a couple of thousand (Vogler (2015), for example).

Numerical mesoscale investigations into the shock compaction of granular materials use Eulerian (or ALE) codes, and typically avoid purely Lagrangian codes because of the additional computational cost required. CTH (McGlaun, Thompson, and Elrick 1990), ALEGRA (Summers et al. 1997) and iSALE (Wünnemann, Collins, and Melosh 2006), for example, are three codes that are commonly used to perform granular shock compaction simulations at the mesoscale. It is also common that these studies consider two dimensional simulations to avoid the additional cost of resolving a large number of grains in 3D (Borg et al. 2009; Borg and Vogler 2009; Derrick et al. 2018a).

Experiments that directly examine the mesoscale are uncommon because they are fundamentally limited by the measuring techniques available (Rutherford 2017; Vogler and Clayton 2008). VISAR (Velocity Interferometer System for Any Reflector) and line-VISAR are two measurement techniques that are commonly used to extract shock states from granular shock experiments. They use laser interferometry to measure the free-surface velocities of the sample (Vogler 2015). The free-surface the speeds are measured at is usually the rear end of the sample; i.e. the surface the shock is propagating toward. VISAR uses a single laser and tracks the velocity of a single point on this free-surface as the shock wave reaches it. line-VISAR, however, uses a line of laser across the material, which means it is able to generate a spatial velocity distribution. Whilst the temporal

resolution is high for this technique, the spatial resolution is limited to a point, or 2D line. Both methods are common in shock physics experiments, however, they are only capable of measuring velocities at the end of the shock's propagation through the sample. VISAR is unable to extract information on the shock wave's progression within the material. For example, it can not directly calculate how the speed of the shock changed (if at all) during propagation. This limits its effectiveness as a tool when the shock front is distinctly non-planar; the front will possess a 3D spatial distribution of velocities and a single point, or line, is insufficient to extract a representative velocity, or velocities. Some studies have attempted to tackle this by improving measurement techniques, such as X-ray imaging (Rutherford et al. 2017; Vogler 2015). Although, X-ray imaging is not without weakness and in contrast to line-VISAR it is limited temporally as opposed to spatially (Vogler 2015). It can reveal details of the shock wave's passage through the particle bed that are unavailable to VISAR's surface-level insight, but only at a few select times; current work suggests that the limitation is approximately 4-5 frames per experiment (Rutherford 2017). As a consequence of the current difficulties in performing mesoscale experiments, many studies of the mesoscale response of granular matter to shock compaction are either wholly, or partially, numerical investigations (Benson and Nellis (1994) and Borg et al. (2009), for example).

Numerical simulations can be significantly more flexible in their set-up, than experiments, which facilitates restricting the number of potentially influential properties in a granular material. Microscopic properties can be restricted to the point that only one is considered at a time. It is because of this that mesoscale simulations usually adopt simple box geometries, rigid piston impacts, and systems involving 'small' numbers of particles (hundreds and thousands as opposed to millions or more). For example recent work by Lammi and Vogler (2012) investigated the effect of introducing grain fracture mechanics to an existing study into the shock compaction of a granular ceramic. This study considered 2D simulations using a relatively small number of grains ($\sim$ 1200) that was compacted

with a rigid, constant velocity, boundary condition. This allowed the study to consider the effects of grain fracture in isolation from any other potentially influential properties. Large volume fractions of completely pulverised material were observed using the new fracture model, which corroborated the previous recovery experiment results. Numerical mesoscale investigations into the shock compaction of granular materials commonly use Eulerian (or ALE) codes, and typically avoid purely Lagrangian codes because of the additional computational cost required. CTH (McGlaun, Thompson, and Elrick 1990), ALEGRA (Summers et al. 1997) and iSALE (Wünnemann, Collins, and Melosh 2006), for example, are three codes that are commonly used to perform granular shock compaction simulations at the mesoscale. It is also common that these studies consider two dimensional simulations to avoid the additional cost of resolving a large number of grains in 3D (Borg et al. 2009; Borg and Vogler 2009; Derrick et al. 2018a).

By restricting all potentially influential properties, and considering their effects on the shock response one at a time, it is possible to gradually reveal the relative importance of each microscale property in granular shock compaction. This, in turn, allows for the improvement of porosity models, where the features can be incorporated into the empirical relations. Some bulk features, however, have been observed that can not be attributed to the existing knowledge of microscale properties (Derrick et al. 2018b; Perry et al. 2016), and the number of properties that have been considered is small compared to the number that remain unexplored. In particular, 'mesostructure' is a relatively recently considered mesoscale property that has not been extensively examined. This is structure created by the arrangement, or inter-locking, of grains within a particle bed. The presence of certain mesostructures can cause a particle bed to resist or facilitate compression, and recent work has shown that it can have a noticeable effect on the shock response (Borg and Vogler 2009; De Cola et al. 2018; Perry et al. 2016).

2.4 Mesostructure in Granular Systems

2.4.1 Existing Mesostructural Investigations

The literature contains many studies that investigate the influences of micro and mesoscale granular properties on the shock compaction of granular materials. These investigations, however, are commonly particular to their respective contexts and do not examine the broader implications of the findings. Benson and Nellis 1994 for example, compared the shock compaction of two copper powders using two 2D numerical simulations. One powder contained exclusively circular grains, whilst the other consisted of rectangular grains with varying aspect ratios. Benson's results showed that the rectangular grains exhibited heating along vertical faces preferentially (parallel to shock direction), whereas circular grains exhibited no such preference. Benson concludes that the strength of the circular-grained powder would be more isotropic than the rectangular-grained powder. Why the vertical faces exhibit preferential heating is not examined, however, which might have implications for other morphology studies that consider rectangular grains, such as Borg and Vogler (2009). Eakins and Thadhani (2008) is another good example. Eakins *et al.* examine the shock compaction of Ni+Al powders using 2D mesoscale modelling with simulations that were initialised from SEM images of experimental samples. The authors observe that the grain morphology and configurations within their samples have a significant effect on their results. Their analysis, however, focused on two different set-ups and did not examine intermediary cases. Despite their direct consideration of morphological effects, it was unfeasible to draw broad conclusions about the affects of morphology from the examination of only two samples. This is another common feature in the literature: investigations typically consider two or three sample types when investigating a granular property, which limits the conclusions that can be drawn about the effect of the property on the shock response (Borg and Vogler 2009; De Cola et al. 2018; Fredenburg,

Dennis-Koller, and Dattelbaum 2012; Perry et al. 2016). Perry et al. (2016), for example, examine the effects of morphology on the quasi-static and shock compaction of three similar samples of sand. The samples differed in their surface morphologies and particle size distributions; one possessed rough sand grains, another possessed smooth grains, and the last (a sandy loam) possessed rough grains with many smaller grains in addition. They concluded that the morphological properties dominated at low strain rates, stiffening the material's response with increasing roughness; whereas, at high strain rates the mechanics of contacts between particles dominated the response. The study directly considers surface morphology and particle size distribution, however it only considers three samples. Nevertheless, the use of a small number of samples is not unjustified because it is unfeasible to restrict the properties of a powder in sample preparation to such a degree that a single property can be varied by incremental amounts between experiments. In addition, to be certain that the target property is the primary reason for difference in shock response, it is sensible to choose a few starkly different samples than many that are very similar. If, however, the understanding of granular compaction is to be furthered to the point where it can be fully parametrised in a porosity model, more investigations are needed that explore the effects of micro and mesoscopic properties on the shock response over broad ranges with large sample sizes.

There are some existing studies that specifically consider the effect of mesostructure on the shock response of a granular material, but few make direct reference to it. The most common approach is to examine it indirectly through the morphology. Fredenburg, Dennis-Koller, and Dattelbaum (2012), Borg and Vogler (2009) and Perry et al. (2016) state that the reason the morphology of their respective granular materials has an effect on the shock response is because changing the morphology affects the contact mechanics between grains, or the number of contacts themselves, which is the mesostructure of the particle bed by the definition used in this thesis. The core feature of mesostructure is the distribution of contacts between grains and their mechanics and these studies vary the

mesostructure by varying the morphology. In Borg and Vogler (2009), however, there is also a direct, numerical, investigation into mesostructure that is independent of the effects of morphology. The authors frame the mesostructure as 'disordered' versus 'ordered' grain arrangements. They varied the mesostructure through their particle bed-creation routines, which began with an ordered arrangement of grains, and applied a randomisation algorithm until the grains were suitably distributed. The randomisation algorithm was applied to two different initial ordered arrangements (square grid, and square gird rotated by 45°). Their results showed that the shock responses appeared to converge on their 'baseline' simulation results as disorder increased. Whilst this is one of the most direct investigations into mesostructure within the literature, the method used to parametrise it does not reveal the insight that is needed to be applied to a porosity model. The parametrisation method is a pre-creation method, as opposed to a post-creation method; it labels the bed based on the way in which it was created and not on what it created. To document the effect of mesostructure on the shock response of granular materials in sufficient detail that it can be incorporated into a porosity model, requires a parametrisation that can be applied to any given particle bed. The method introduced in Borg and Vogler (2009) can only be applied to particle beds created using the same method, which limits its general applicability.

There is a common feature across many of the studies discussed so far that has not yet been addressed: the shock responses of naturally occurring powders (other than sand), like soil, are rarely considered in the literature, despite their documented ubiquity. This is probably because naturally occurring powders are typically much more complex than artificially created powders and it is consequently more challenging to restrict all of their properties in a study. The former can have particle size distributions spanning multiple orders of magnitude, grains of multiple materials, and extremely angular morphologies. Modelling these powders is challenging because of the massive computational and temporal cost of simulating the large number of grains of many sizes and shapes.

Until recently granular materials have either been resolved 'explicitly' or 'implicitly' in numerical simulations. An explicitly resolved granular material resolves every grain as a separate object, and they are allowed to interact to produce the response of that material (almost all mesoscale simulations in the literature are 'explicitly' resolved granular material simulations). An implicitly resolved granular material, however, treats the material as continuous and parametrises the powder's compaction using a porosity model. Explicitly resolving a granular material will produce more accurate results, but there are many examples where resolving a granular material explicitly is not feasible (as discussed in Section 2.3.2). Simulations using porosity models have been shown to produce accurate results, when compared to explicit simulations, or experiments (Collins et al. 2011; Wünnemann, Collins, and Melosh 2006) and using porosity models drastically reduces the computational costs involved. Recent work, however, has begun to suggest that the response of some powders can not be captured by implicitly resolving granular matter and that they are also too complex to be resolved explicitly (Bland et al. 2014).

Bland et al. (2014) have introduced a new method of studying these complex powders computationally that avoids these problems. Bland et al. (2014) considers the shock compaction of precursor meteorite matter that consisted of large solidified melt droplets surrounded by a fine, porous, dust with grains that are orders of magnitude smaller than the droplets. The authors resolve the granular material 'semi-explicitly' by parametrising the dust and resolving the droplets (they resolve the smaller grains implicitly, whilst resolving the rest of it explicitly). Their results reveal that the droplets remain cold in comparison to the dust, which exhibits extreme temperature variations at small scales. This feature is unique to this strongly bimodal powder, but it also suggests that there might be other potentially influential mesoscale properties that have not yet been observed because previous capabilities were insufficient to measure them. This approach of semi-explicitly resolving a powder presents a potential method of modelling other complex powders that would otherwise be computationally unfeasible to simulate. Only a few

studies exist that use this approach (Bland et al. 2014; Davison, Collins, and Bland 2016; Davison et al. 2017; Rutherford et al. 2017) and its full capabilities have yet to be explored.

2.5 Research Objectives

The research presented in this thesis aims to address some of the flaws observed in the existing literature. In particular, this thesis aims to provide: a post-creation analysis of mesostructure in a particle bed and its effect on the shock response; a comparison of the semi-explicit modelling approach with experiment, using the same bimodal powder structure used in Bland et al. (2014), and an examination of the specific response of the particle bed at the mesoscale; and to examine the applicability of the semi-explicit approach when modelling complex powders with complex grain morphology and large size distributions, such as the lunar regolith.

Chapter 3 presents an initial benchmark of the iSALE code with another well-known and well-respected hydrocode, CTH, and a resolution analysis of iSALE. Good agreement is found between CTH and iSALE and the minimum recommended grain resolution was found to be 8 cells per particle diameter (circular grains).

Chapter 4 directly considers the mesostructure of a particle bed in a set-up based on the work presented in Borg and Vogler (2009). The chapter introduces the coordination number (average number of contacts per particle) as a means of quantifying the arrangement of particles in a bed. Increasing the coordination number caused the particle bed to stiffen its response, whereas lowering it weakened the material's response.

Chapter 5 presents a joint numerical-experimental analysis of the shock compaction of the bimodal mixture modelled by Bland et al. (2014) using analog materials. The investigation considers three systems numerically and experimentally: pure matrix (the fine dust), pure matrix with a single chondrule (a large grain), and multi-chondrule mixtures with

proportions analog to those used in Bland et al. (2014). The mechanics of bimodal compaction are analysed at the mesoscale and a dichotomy in compaction around the large grains is found; the large grains are shown to 'shield' the fine dust in their lee. To quantify this effect the chapter introduces the 'porosity vector', which points in the direction of least compaction across a grain and the chapter finds that the porosity vector's directional preference for the direction of the shock's passage is also present in multi-chondrule mixtures.

Finally, Chapter 6 presents an investigation into the efficacy of semi-explicitly resolving granular materials to model complex powders, with the lunar regolith as a case study. In particular the temperatures and pressures experienced by the material are considered. The chapter demonstrates that simulations resolving the regolith semi-explicitly can accurately model it and take into account its complex properties, provided the parametrised material possesses approximately the same porosity as the bulk porosity. In addition, it was found that grain morphology did not affect the bulk response of the material but had a measurable mesoscale effect. An optimum angle between the shock front and the shock-facing side of a grain was found that produces a hot spot in the matrix material along that face.

Chapter 7 summarises the findings in this thesis, presents any conclusions made over its course and outlines areas for future work.

Chapter 3

Methodology

3.1 Introduction

This Chapter introduces the common methods and techniques used across the whole thesis. It includes a background and brief overview of the principal code used in almost all the modelling presented (iSALE) as well as some details about the companion set-up routines developed by the author for this research. Finally, the principal methods behind calculating the Hugoniot (the loci of possible shock states) from mesoscale simulations of compaction of heterogeneous matter are discussed.

3.2 iSALE: a Shock Physics code

3.2.1 Introduction & Background

The main code used throughout this thesis is iSALE (impact Simplified Arbitrary Lagrangian Eulerian). iSALE has a long history that began in the 1980s with the SALE hydrocode (Amsden, Ruppel, and Hirt 1980). SALE was a simple FORTRAN code that could model compressible single-material Newtonian fluid flow at all speeds, including

shock waves. It solved the compressible Navier-Stokes equations, as well as conservation of mass and internal energy equations. From the early 90s, until the mid 00s, two versions of SALE were developed concurrently and independently, SALES-2 and SALEB, which added a variety of features. The Lagrangian part of the program was developed to take into account shear and tensile fractures and the Tillotson EoS (H. Tillotson 1962) amongst other features (Melosh, Ryan, and Asphaug 1992) (SALES-2). SALEB was developed by Ivanov *et al.* (Ivanov, Deniem, and Neukum 1997; Ivanov 2005) who added several features, which included the ability to have multiple materials in one cell; allowing advection of all variables in the Eulerian part; and additions of the semi-analytical EoS package, ANEOS (Thompson and Lauson 1974). Eventually these two were merged into one code, iSALE, appearing in the literature for the first time in Wünnemann, Collins, and Melosh (2006).

iSALE was originally designed to model impacts in the solar system, and its mesoscale modelling has required some modifications. The first published example of mesoscale modelling of granular materials, using iSALE, was in 2013 (Güldemeister et al. 2013).

3.2.2 Governing Equations

iSALE uses a finite difference method to solve the three principal equations of computational fluid dynamics. They are the continuity equation (conservation of mass); the Cauchy momentum equation (conservation of momentum/$\mathbf{F} = m\mathbf{a}$); and the conservation of energy equation.

The equations are typically three dimensional, however, only the 2D version of iSALE is used in this thesis, except in one special case in Chapter 5 which uses cylindrically symmetric geometry and is in 3D.

The Continuity Equation

The continuity equation arises when considering the mass fluxes into and out of a small, fixed, box of fluid and relates the change in density of that box with time, to the density, and velocity, gradients present.

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u} \cdot \nabla \rho = -\rho \nabla \cdot \boldsymbol{u} \tag{3.1}$$

Where the 3D velocity vector is $\boldsymbol{u} = (u, v, w)^T$, t is time and ρ is density.

The Cauchy Momentum Equation

The conservation of momentum equation is derived by applying Newton's 2nd law ($\mathbf{F} = m\mathbf{a}$) to the same box of material as in the previous section,

$$\rho \frac{\partial \rho}{\partial t} + \rho(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = \boldsymbol{f} + \nabla \cdot \overline{\overline{\sigma}} \tag{3.2}$$

where $\boldsymbol{f}$ is all external body forces per unit mass (e.g. gravity) and $\overline{\overline{\sigma}}$ is the stress tensor acting at a point, positive in tension. However, it is convenient to separate the isotropic (pressure) and deviatoric parts of of the stress tensor.

$$\overline{\overline{\sigma}} = \overline{\overline{s}} - p\overline{\overline{I}} \tag{3.3}$$

$$p = -\frac{1}{3}(\sigma_{xx} + \sigma_{yy} + \sigma_{zz}) \tag{3.4}$$

Where p is pressure (positive in compression), $\overline{\overline{s}}$ is the deviatoric stress and $\overline{\overline{I}}$ is the 3D identity matrix. This relation causes the $\nabla \cdot \overline{\overline{\sigma}}$ term to become:

$$\nabla \cdot \overline{\overline{\sigma}} = -\nabla p + \nabla \cdot \overline{\overline{s}}. \tag{3.5}$$

Thus, the general, Lagrangian, form of the Cauchy equation becomes:

$$\rho \frac{\partial \boldsymbol{u}}{\partial t} + \rho(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = \boldsymbol{f} - \nabla p + \nabla \cdot \overline{\overline{s}}. \tag{3.6}$$

This new form allows changes due to pressure to be calculated separately to stress, which simplifies the material models.

Conservation of Energy

The last continuity equation is that of (specific internal) Energy, E. It is found in a similar manner to the previous equations, by considering a small box of fluid. The surface forces will do work on the fluid and by summing their input with respect to time ($\Sigma \boldsymbol{F} \cdot \boldsymbol{u}$), a total rate of change of energy (E_T), with time, can be found.

$$\rho \frac{DE_T}{Dt} = \nabla \cdot (\overline{\overline{\sigma}} \cdot \boldsymbol{u}) \tag{3.7}$$

The equation, however, is needed in terms of internal energy E, and so the kinetic energy (E_K) must be subtracted from it.

$$\rho \frac{DE_K}{Dt} = (\nabla \cdot \overline{\overline{\sigma}}) \cdot \boldsymbol{u} \tag{3.8}$$

After expanding stress to deviatoric and pressure components and subtracting the kinetic energy terms, the full conservation of energy equation is found.

$$\rho \left(\frac{\partial E}{\partial t} + \boldsymbol{u} \cdot \nabla E \right) = -p \nabla \cdot \boldsymbol{u} + \mathrm{Tr} \left(\overline{\overline{s}} \cdot \frac{\partial \boldsymbol{u}}{\partial \boldsymbol{r}} \right) \tag{3.9}$$

Where $\boldsymbol{r} = (x, y, z)^T$. Despite being the last conservation equation, there are still too many unknowns for a solution. Two further equations are necessary to solve the system.

3.2.3 Equation of State

In addition to the three conservation equations an additional two equations are required to solve for all variables. These are not universal equations and different materials require different relations. One is the 'Equation of State' (EoS) which relates information about the compressibility (specific pressure p and density ρ) of the material to its specific internal energy.

$$p = p(\rho, E) \tag{3.10}$$

The EoSs used in this thesis include the Mie-Grüneisen EoS (most common), the Tillotson EoS, and ANEOS, a semi-analytical EoS package. ANEOS can represent phase changes in a thermodynamically consistent way (Thompson and Lauson 1974). This is the principle

advantage of its use over other EoSs like the Mie-Grüneisen EoS, which can not accurately represent phase changes (see Chapter 4 for details).

The Mie-Grüneisen (MG) EoS appears in this work more than any other and takes a very simple form that uses several standard material parameters as its input (initial density, sound speed, the linear $U_s - u_p$ gradient—S—, and specific heat capacity) and one EoS-specific variable: the Grüneisen gamma, Γ.

$$p = p_h + \Gamma\rho(E - E_h) \tag{3.11}$$

E_h and p_h are the Hugoniot states of the material. The MG EoS is favoured in the shock-physics field because of its simplicity, although it has a significant limitation. It has a singularity at a certain compression and because of this it is best for small compressions. To avoid this problem, in the simulations modelled here, the MG EoS is used for almost all materials that remain solid, and ANEOS is used for all others (principally in Chapter 6). Additionally iSALE uses the common approximation for the gamma coefficient $\Gamma\rho \approx \Gamma_0\rho_0$, which further simplifies equation 3.11.

The $\epsilon - \alpha$ Porosity Compaction Model

Compaction is inherently related to the compressibility of a material and it is because of this that porosity models are incorporated into the EoS. In iSALE simulations that require the parametrisation of porosity, the $\epsilon - \alpha$ model is used. The model relates the distension of a material ($\alpha = \frac{\rho_s}{\rho}$; ratio of solid, to bulk, density) to an applied mechanical volume strain ($\epsilon = ln(\frac{V_0}{V})$; logarithmic ratio of the original volume to current volume, positive in tension), separated from thermal volume strain. Distension is related to the more familiar porosity term ($\phi = \frac{V_{void}}{V_{total}}$; ratio of void space to total space) by a simple

expression:

$$\alpha = \frac{1}{1-\phi}. \tag{3.12}$$

The porosity model is implemented into the hydrocode through the EoS. In the EoS of a porous material the bulk pressure (p) will be a function of density, energy and distension. It is convenient, however, to calculate p by calculating the solid material pressure, p_s, first and converting to p by dividing by the distension (equation 3.13). The solid density (ρ_s) is required to calculate p_s and can be simply calculated by multiplying the bulk density (ρ, which is solved for by iSALE) by the distension.

$$p = f(\alpha, \rho, E) = \frac{1}{\alpha} p_s(\alpha\rho, E) = \frac{1}{\alpha} p_s(\rho_s, E) \tag{3.13}$$

Thus, the final step in implementing the porosity model is to calculate the distension from other state variables: the $\epsilon - \alpha$ model defines $\alpha = \alpha(\epsilon)$. Within iSALE, however, the differential form ($\frac{d\alpha}{d\epsilon} = f(\alpha, \epsilon)$) is used because α is updated with a simple iterative step (equation 3.14).

$$\alpha_{n+1} = \alpha_n + \frac{d\alpha}{d\epsilon}\frac{d\epsilon}{dt}\Delta t \tag{3.14}$$

The $\epsilon - \alpha$ model has four operating regimes: elastic compaction, exponential compaction, power-law compaction, and compression.

In the elastic regime the material's porosity decreases with increasing pressure and (negative) strain, however, the porosity is regained if the material is released. This regime continues until a threshold strain, ϵ_e, is reached. In this work $\epsilon_e = 0$ in all simulations, so the elastic regime is bypassed and the material enters the exponential compaction regime immediately.

In the exponential compaction regime, porosity is lost through pore collapse as grains rearrange themselves (Wünnemann, Collins, and Melosh 2006). The rate of change of distension with volumetric strain is assumed to be proportional to the distension (equation 3.15). Resistance to this process increases, however, as compaction increases because more grains are forced into close contact and the inter-granular friction forces rise. In order to account for this, the unit-less parameter $\kappa \lesssim 1$ is introduced to control the compaction rate. The final equation for this regime is shown below and when integrated produces the relationship $\alpha = \alpha_0 e^{\kappa(\epsilon - \epsilon_e)}$, which is the source of its name as the exponential compaction regime.

$$\frac{d\alpha}{d\epsilon} = \kappa\alpha \tag{3.15}$$

A critical strain (ϵ_x) or distension (α_x) marks the point at which the material can not be compacted any further by rearranging grains and thus grains begin to fracture. At this point the model shifts into the 'power-law regime', which lasts until total compaction of all pore space is achieved at the strain ϵ_c. In this regime an expression is needed to calculate the distension that follows three criteria: the gradient of the compaction curve must be smaller than in the previous regime because it should be harder to compress the material through grain fracture than rearrangement; the transition between the two regimes must be continuous; the compaction in this regime must approach $\alpha = 1$ ($\phi = 0\%$

smoothly. The $\epsilon - \alpha$ model uses the expression in equation 3.16 for this regime, which fulfils the three criteria specified (Wünnemann, Collins, and Melosh 2006).

$$\frac{d\alpha}{d\epsilon} = 2(1 - \alpha_x)\frac{\epsilon_c - \epsilon}{(\epsilon_c - \epsilon_e)^2} \tag{3.16}$$

Finally the compaction of the material follows the power-law regime, until the material is fully compacted. The distension aymptotes to 1, and the porosity to zero. Simultaneously, the rate of compaction smoothly approaches zero and once the material reaches total compaction it then begins to behave according to its continuous description.

By changing α_0, κ, ϵ_e, and α_x it is possible to model the compaction response of a wide range of porous materials and an example of the compaction curve for a fictitious material is shown in Figure 3.1 as a visual aid.

3.2.4 The Constitutive Relation

The last equation is the 'Constitutive Relation' which relates the strain ϵ, strain rate $\dot{\epsilon}$, pressure, energy E and damage D (see section 3.2.4 for details) to the deviatoric stress $\overline{\overline{s}}$.

$$\overline{\overline{s}} = g(\epsilon, \dot{\epsilon}, p, E, D) \tag{3.17}$$

The function $g(\dots)$ represents how the material resists deformation due to any forces applied to it. A classic example of a constitutive relation is Hooke's Law, which linearly relates stress to strain. This work assumes linear elasticity up to a yield stress that is defined by the strength model used.

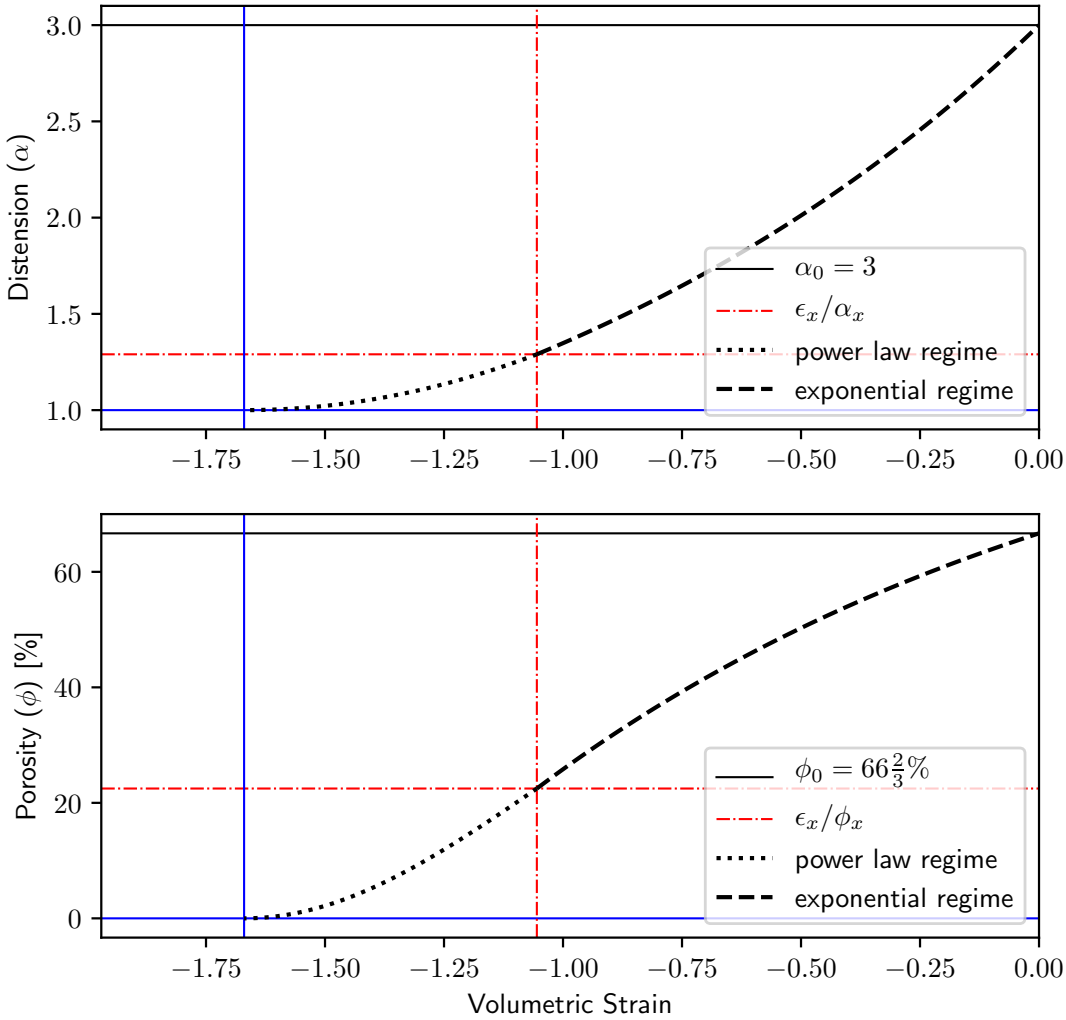


Figure 3.1: Example compaction curve for a fictitious material using the $\epsilon - \alpha$ model. This model was constructed using the parameter values: $\alpha_x = 1.29$, $\alpha_0 = 3$, $\epsilon_e = 0$, and $\kappa = 0.8$, which are representative of the materials considered in this thesis

Strength Models

'Strength models' refer to different methods of defining the stress at which a material ceases to respond elastically, as well as how this stress evolves with changes in strain and other variables. Strength models are slightly different from the constitutive relation,

which also includes how the material responds in an elastic regime. The strength model that is used predominantly in this thesis is the Von Mises Yield Criterion (Mises 1913). It is commonly referred to as an 'elastic-perfectly plastic' treatment and it defines the yield strength Y of a material as

$$Y = Y_0. \tag{3.18}$$

This is a scalar quantity, and to compare the stress tensor with it, the tensor must be condensed into a scalar form. This is done in iSALE by using the square root of the second invariant of the deviatoric stress tensor $\sqrt{J_2}$ (Collins et al. 2016).

$$J_2 = \frac{1}{6}((s_{xx} - s_{yy})^2 + (s_{yy} - s_{zz})^2 + (s_{zz} - s_{xx})^2) + s_{xy}^2 + s_{yz}^2 + s_{zx}^2. \tag{3.19}$$

It is common in hydrocodes to use other scalar measures of shear stress magnitude, such as the equivalent stress, which is used in CTH. This is a factor of $\sqrt{3}$ greater than $\sqrt{J_2}$. For example the Y_0 of a material in iSALE will be $\frac{1}{\sqrt{3}}$ times smaller than the equivalent Y_0 in CTH (see Chapter 4).

Strength models are often quite complex and Von Mises is very simplistic. A strength model commonly used for metals, and the second most common in this thesis, is the Johnson-Cook model (Johnson and Cook 1985). This relation additionally factors in other aspects that will affect the strength of a material, such as thermal softening

$$Y = (A + B\epsilon^N)(1 + C \ln \dot{\epsilon}) \left(1 - \left(\frac{T - T_{ref}}{T_m - T_{ref}}\right)^M\right) \tag{3.20}$$

A is the yield strength at the reference state ($T = T_{ref}$, $\epsilon = 0$, $\dot{\epsilon} = 1$) and T_m is the melt temperature. B, C, N, and M, are all material-specific user-inputted parameters.

Finally, two other strength models are used in this thesis and they are included here for completeness. The Drucker-Prager model[1] (Drucker and Prager 1952) and the 'Rock' model[2] (Collins, Melosh, and Ivanov 2004; Ivanov, Deniem, and Neukum 1997).

$$Y = \min(Y_0 + \mu p, Y_m) \qquad \text{Drucker-Prager} \tag{3.21}$$

$$Y = Y_d D + Y_i(1 - D) \qquad \text{Rock} \tag{3.22}$$

$$Y_d = \min(Y_{d0} + \mu_d p, Y_{dm}) \qquad \text{Rock; damaged strength} \tag{3.23}$$

$$Y_i = Y_{i0} + \frac{\mu_i p}{1 + \frac{\mu_i p}{Y_{im} - Y_{i0}}} \qquad \text{Rock; intact strength} \tag{3.24}$$

μ is the internal coefficient of friction, Y_m is the limiting strength at high pressure, and the i and d subscripts refer to 'intact' and 'damaged' respectively.

Damage Models

Damage is a scalar quantity $0 <= D <= 1$ that reflects the level of fragmentation of material within a cell (Collins, Melosh, and Ivanov 2004; Ivanov, Deniem, and Neukum 1997). One damage model is used in this work: the 'Ivanov' model. Under this model Damage is calculated using the expression shown in equation 3.25.

[1] Used for Sipernat in Chapter 5.
[2] Used for Basalt in Chapter 6.

$$D = \min\left(\frac{\epsilon_p}{\epsilon_f}, 1\right) \tag{3.25}$$

ϵ_p is an invariant measure of accumulated total plastic strain that is always positive. ϵ_f, however, is calculated using the following expression.

$$\epsilon_f = \max(\epsilon_{fb}, B(p - p_c)) \tag{3.26}$$

ϵ_{fb} is the minimum failure strain for low pressure states, p_c is a pressure above which failure is always compressional and B is a positive constant. Only models presented in Chapter 6 use this damage model.

Thermal Softening Models

When materials approach their melt temperature they typically become weaker, however, most of the scenarios considered in this thesis do not exhibit conditions where this occurs. In Chapter 6, however, there is sufficient heating of the material in some simulations for the material to approach (and exceed) the melt temperature. In this chapter a thermal softening model is needed. In addition, at high pressures the melt temperature will also be a function of pressure. In iSALE the pressure-dependency of the melt temperature is approximated using the common Simon-Glatzel equation (Simon and Glatzel 1929).

$$T_m = T_{m0}\left(\frac{p}{a} + 1\right)^{1/c} \tag{3.27}$$

Where T_m is the melt temperature at pressure p, T_{m0} is the melt temperature at zero pressure, and a and c are material constants.

The thermal softening of the material is approximated in iSALE using the following equation, which was originally published in Ohnaka (1995) and is referred to as the Ohnaka thermal softening model.

$$Y = Y_c \tanh\left(\xi\left(\frac{T_m}{T} - 1\right)\right) \tag{3.28}$$

Y is the thermally softened strength of the rock, Y_c is the cold yield strength (including damage), T_m is the melt temperature, T is the temperature of the rock and ξ is a material constant. Thermal softening is only used in Chapter 6 to model basalt.

3.2.5 Discretisation

To apply a finite-difference method all derivatives become differences and the independent variables (space and time) are discretised. When considering the simulation's domain it can not be described as a continuum and must be split into a grid of points, known as nodes. These are joined by straight lines to form cells. In iSALE vector fields such as velocity and position are defined at each node, whilst scalar fields, such as pressure or density, and the stress tensor are defined in cell centres (Figure 3.2). This set-up is known as a 'staggered mesh'. A first-order difference scheme is used to calculate solutions across the mesh (equation 3.29). Additionally, the points must be discretised in time, with time step Δt (equation 3.29).

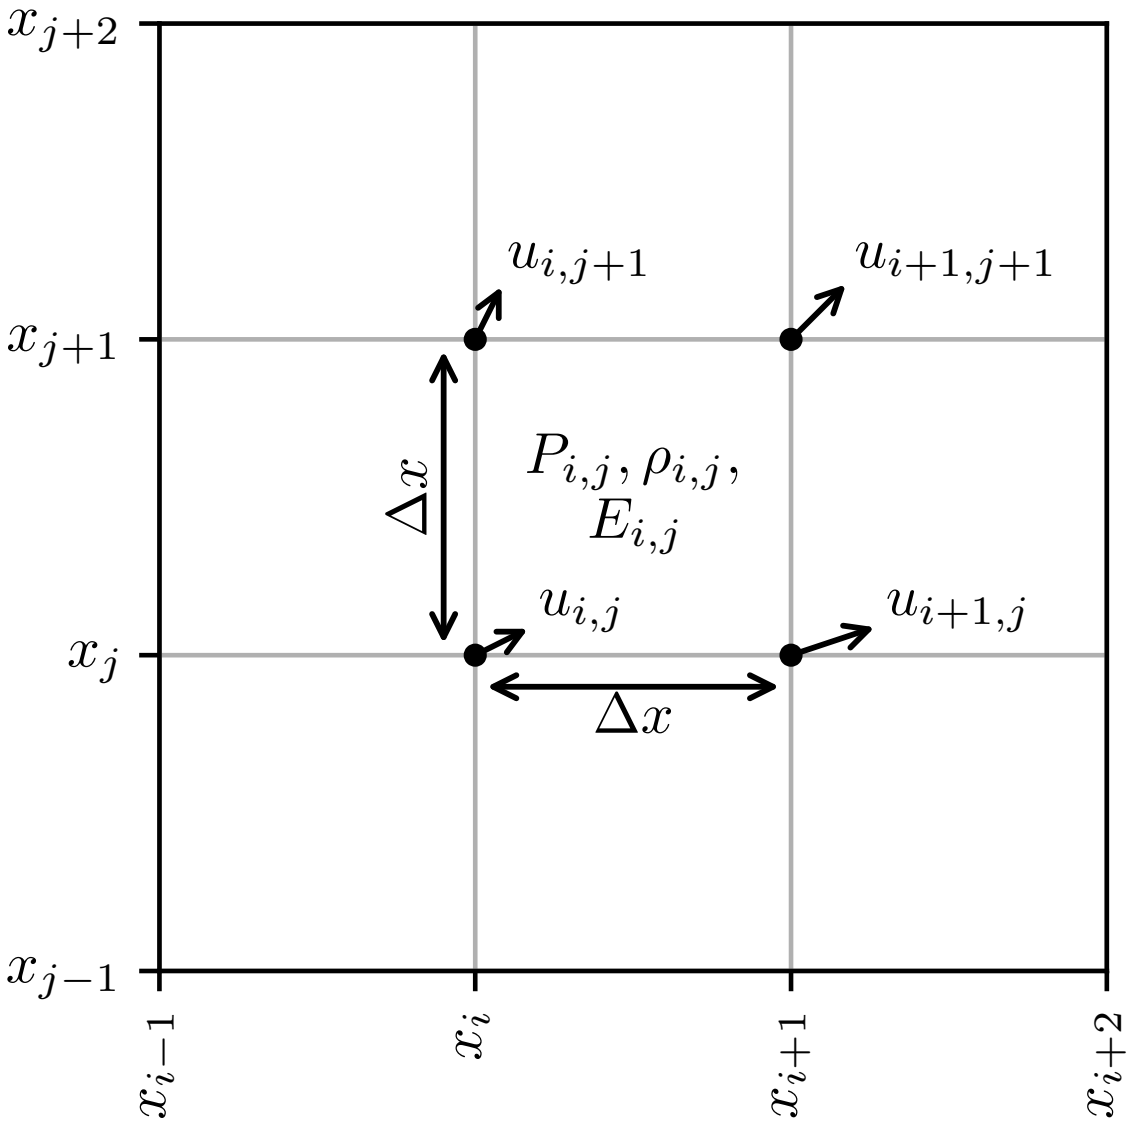


Figure 3.2: Space is discretised in iSALE into squares Δx by Δx in size. Position ($x_{i,j}$) and velocity ($u_{i,j}$) are defined at each node whilst scalar quantities (pressure—P—, density—ρ—, energy—E—) are defined in each cell centre. This is known as a 'staggered mesh'.

$$f'(x_i) = \frac{f(x_{i+1}) - f(x_i)}{\Delta x} + O(\Delta x) \qquad \text{Spatial Forward Difference Scheme} \qquad (3.29)$$

$$q^{n+1} = q^n + \frac{\Delta q}{\Delta t}\Delta t \qquad \text{Temporal Forward Difference Scheme} \qquad (3.30)$$

The simulations in this work, however, follow an ALE procedure. Velocities and positions are defined at each node, then at each time step these are updated by solving the Lagrangian conservation equations. The nodes are advected, as they would be in a Lagrangian calculation. At this stage, the deformed mesh is remapped onto its starting grid, and material is advected to new cells by calculating the flux into each cell. The new cell-centred parameters are calculated from the new contents of the cell, although, it is because of this method that cells may end up containing multiple materials. In this case, it is not clear how the bulk cell parameters required to solve the governing equations can be calculated.

Multi-material cell treatment

A natural consequence of using an Eulerian mesh of discretised cells is that each cell must be capable of containing multiple different materials (including void) at a time. iSALE tracks the interfaces between materials within each cell and uses them to constrain cell volume fractions by modifying the flux into and out of the cell. For example, this prevents material advecting into the next cell before the interface. The cell-centred parameters such as bulk pressure, temperature and density etc. are calculated by computing the average of each parameter, weighted by their respective material volume fractions. Calculating the strength of these mixed cells, however, is more complicated and can be done in three different ways.

The first method is called 'min': in this method the strength of the mixed cell is taken to be the same as the weakest material in that cell. The second method is 'const': in this method the cell strength of mixed cells is always equal to a user-specified constant (commonly zero). The third method is 'med': in this method the strength of each mixed cell is calculated as the volume-weighted average of all material strengths in that cell. When a cell contains some proportion of void, however, that cell always has a strength

(and pressure) of zero, regardless of which mixed-cell strength method has been chosen.

3.2.6 Shock Modelling

In the special case of shock modelling, it is problematic to model a discontinuous shock profile over a discrete mesh where the rise of a shock is occurring on a scale that is much smaller than the cell size. When this phenomenon is modelled in a discrete mesh significant instabilities (large oscillations) are generated at the shock front which have the potential to invalidate a simulation. One solution to this problem takes inspiration from the dissipative effects of viscosity (VonNeumann and Richtmyer 1950). The shock front in a viscous fluid is 'dissipated' and spread out, longitudinally. Thus, by introducing an artificial viscosity term (q) to the hydrocode, in addition to pressure, the shock is smoothed over several cells, reducing the numerical artifacts and maintaining stability (Anderson 1987).

$$\nabla p \rightarrow \nabla(p+q) \tag{3.31}$$

$$q = a_1 \rho A (\nabla \cdot \boldsymbol{u})^2 + a_2 \rho c \sqrt{A} \nabla \cdot \boldsymbol{u} \tag{3.32}$$

q is defined as being a function of user-input parameters a_1 and a_2, density ρ, velocity u, area A, and c, which is a dimensionless constant that is close to one. The inclusion of q maintains simulation stability, however, it has some limitations. In particular, it can cause artificial heating, with the greatest effects being along boundaries and interfaces.

3.2.7 pySALESetup

iSALE was not originally intended to perform mesoscale granular simulations in its inception and because of this its setup routines were very limited. Over the course of this

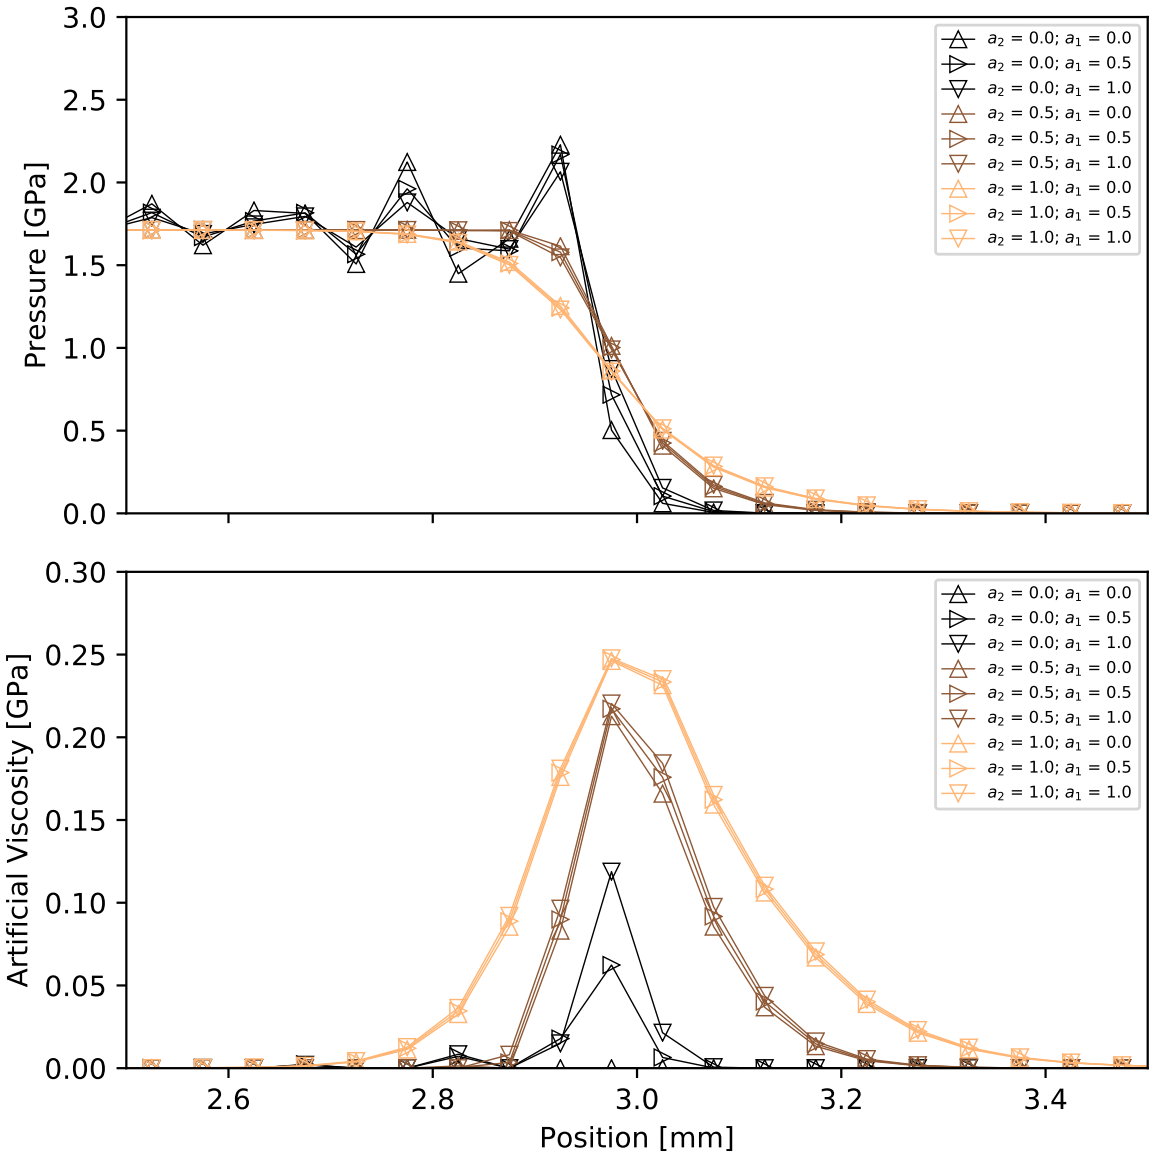


Figure 3.3: Plots of the pressure and artificial viscosity fields in the same 1D simulation with different values of a_1 and a_2. The linear term, a_2, has the greatest effect, significantly smoothing the field, whilst a_1 appears to fine-tune the response on a smaller magnitude. Artificial viscosity is zero away from the shock front and spikes at the discontinuity.

work the author has developed a series of simple routines that can be used to generate complex geometrical set-ups, which can initialise iSALE simulations. This capability is particularly useful in current mesoscale modelling.

The setup routines were written in python as a module, to promote ease of use and

maintenance; they currently interface with iSALE by generating an additional input file that can be imported. In the future this module will likely become the primary simulation set-up method for iSALE, however, additional work will be required to reach this stage. The module is called 'pySALESetup' following the naming convention started by the comparative iSALE plotting module pySALEPlot (Collins et al. 2016).

pySALESetup primarily consists of one 'base' class, two 'object' classes and two 'extra' classes. The base class is the **Mesh** class. This defines the domain for the simulation to be set up in. It is initialised by providing its size and shape, as well as cell size. The class additionally contains several useful functions that allow a user to give materials velocities or fill the entire mesh with a certain material, for example. The class prioritises existing material when new material is added. If an object is inserted that overlaps with an existing one, the new object is cropped. This is the class, that when saved, produces the input file that can be read by iSALE.

```
import pySALESetup as pss
# create a Mesh instance, m
m = pss.Mesh(X=100,Y=500,cellsize=10.e-6)
# view the mesh materials
m.viewMats()
# view the mesh velocities
m.viewVels()
```

The object classes are '**Grain**' and '**Apparatus**'. These create physical 'objects' that can be subsequently inserted into a mesh. The **Grain** class is used to create instances of particles to be inserted into the simulation set-up. It currently supports circular, elliptical, and n-sided polygonal shapes (provided as a list of vertices). All shapes should be normalised onto a 2×2 Cartesian grid where $-1 <= x, y <= 1$. This makes grain morphology is physical-size independent; the exact vertices of a shape do not have to be specified by the user. The grain is scaled by the 'equivalent radius' parameter, which is the the radius of the hypothetical circle (in cells) that possesses the same area as the grain instance. This allows for easy scaling of particles; a particle can be created twice the size

of another by giving it double the equivalent radius, regardless of its morphology. When a **Grain** instance is created the resulting particle is created in a 'mini-mesh' (a square 2D array that has side length of equivalent diameter + 2) and stored within the instance for later use. This was chosen to avoid re-generating identical grains many times during set-up, particularly when they are complex shapes, which can be computationally costly.

Once created, **Grain** instances can be inserted into a target mesh using one of three in-built functions: *insertRandomly*, *insertRandomwalk* and *place*. *insertRandomly* selects an unoccupied region in a specified **Mesh** and places the grain there as specified material. *insertRandomwalk* places a grain in the mesh and has it undergo a random walk until it is in contact with other grains. 'In contact' is defined as overlapping with another object by no more than 100th the area of the placed grain, or 1 cell (whichever is larger). The third function, *place*, simply places a grain into the mesh at the specified coordinates and does not check if anything is already there.

```
import pySALESetup as pss
# create Grain instance; equivalent radius (eqr) = 5 cells
# no rotation (normally in radians)
# circular shape. Other parameters are required for other shapes
g = pss.Grain(eqr=5.,rot=0.,shape='circle')

volume_fraction = 0.
# insert grains until target volume fraction achieved
while volume_fraction <= 0.4:
        # insert randomly into a specified region as material 1
            for now
        g.insertRandomly(m,1,ybounds=[1.e-3,3.e-3])
        # calculate the new volume fraction in the particle
            region
        volume_fraction = m.vfrac(ybounds=[1.e-3,3.e-3])
```

The **Apparatus** class works in a similar manner, although it currently only supports n-sided polygons as input shapes (with vertices supplied as input). Additionally, a mini-mesh is not created for each instance because of their (typically) large size and less frequent occurrence.

```
import pySALESetup as pss
# create apparatus instances of the impactor and buffer plates
impactor = pss.Apparatus(xcoords=[0.,0.,1.e-3,1.e-3],
ycoords=[3.e-3,5.e-3,5.e-3,3.e-3])
bufferplate = pss.Apparatus(xcoords=[0.,0.,1.e-3,1.e-3],
ycoords=[2.e-4,1.e-3,1.e-3,2.e-4])

# place impactor into mesh m as material 2
impactor.place(m,1)
# assign impactor a velocity of -500 m/s to impactor
m.matVel(-500.,1)
# place buffer into mesh m also as material 1
bufferplate.place(m,1)

# delete top and bottom 10 cells to avoid boundary problems
m.top_and_tail(10)
```

The final two classes are not strictly necessary for set-up, however, they add a significant amount of functionality when creating particulate systems, which is particularly useful for mesoscale grain compaction simulations. These are the '**Ensemble**' and '**SizeDistribution**' classes. An **Ensemble** instance tracks a collection of **Grain** instances and facilitates the calculation of various multi-particle statistics such as size frequency distribution. However, its most useful function is to calculate the optimal material distribution between grains with the *optimise materials* function. This finds the optimum distribution of materials across the particles such that each particle is as far as possible (spatially) from another particle with the same material.

A typical script using pySALESetup will use all of these core classes. For example, the following code recreates the geometry of the iSALE example simulation MesoParticles2D using pySALESetup (Collins et al. 2016).

```
""" MesoParticles2D setup """

import pySALESetup as pss

# create a Mesh instance, m
m = pss.Mesh(X=100,Y=500,cellsize=10.e-6)
```

```
# create Grain instance; equivalent radius (eqr) = 5 cells
# no rotation (normally in radians)
# circular shape. Other parameters are required for other shapes
g = pss.Grain(eqr=5., rot=0., shape='circle')

volume_fraction = 0.
# insert grains until target volume fraction achieved
group = pss.Ensemble(m)
while volume_fraction <= 0.4:
        # insert randomly into a specified region as material 1
            for now
        g.insertRandomly(m,2,ybounds=[1.e-3,3.e-3])
        # add Grain instance to Ensemble
        group.add(g)
        # calculate the new volume fraction in the new region
        volume_fraction = m.vfrac(ybounds=[1.e-3,3.e-3])
# optimise materials using those not yet used
g.optimise_materials(np.array[3,4,5,6,7,8,9])

# create apparatus instances of the impactor and buffer plates
impactor = pss.Apparatus(xcoords=[0., 0., 1.e-3, 1.e-3],
ycoords=[3.e-3, 5.e-3, 5.e-3, 3.e-3])
bufferplate = pss.Apparatus(xcoords=[0., 0., 1.e-3, 1.e-3],
ycoords=[0., 1.e-3, 1.e-3, 0.])

# place impactor into mesh m as material 2
impactor.place(m, 1)
# assign impactor a velocity of -500 m/s to impactor
m.matVel(-500., 1)
# place buffer into mesh m also as material 1
bufferplate.place(m, 1)

# delete top and bottom 10 cells to avoid boundary problems
m.top_and_tail(10)

# view the mesh materials
m.viewMats()
# view the mesh velocities
m.viewVels()
# save the mesh as an input file for use with iSALE
m.save()
```

In general, when creating a mesh using pySALESetup, a **Mesh** instance should be created

first; this is followed by all **Grain** and **Ensemble** instances; finally, **Apparatus** instances are created and velocities assigned.

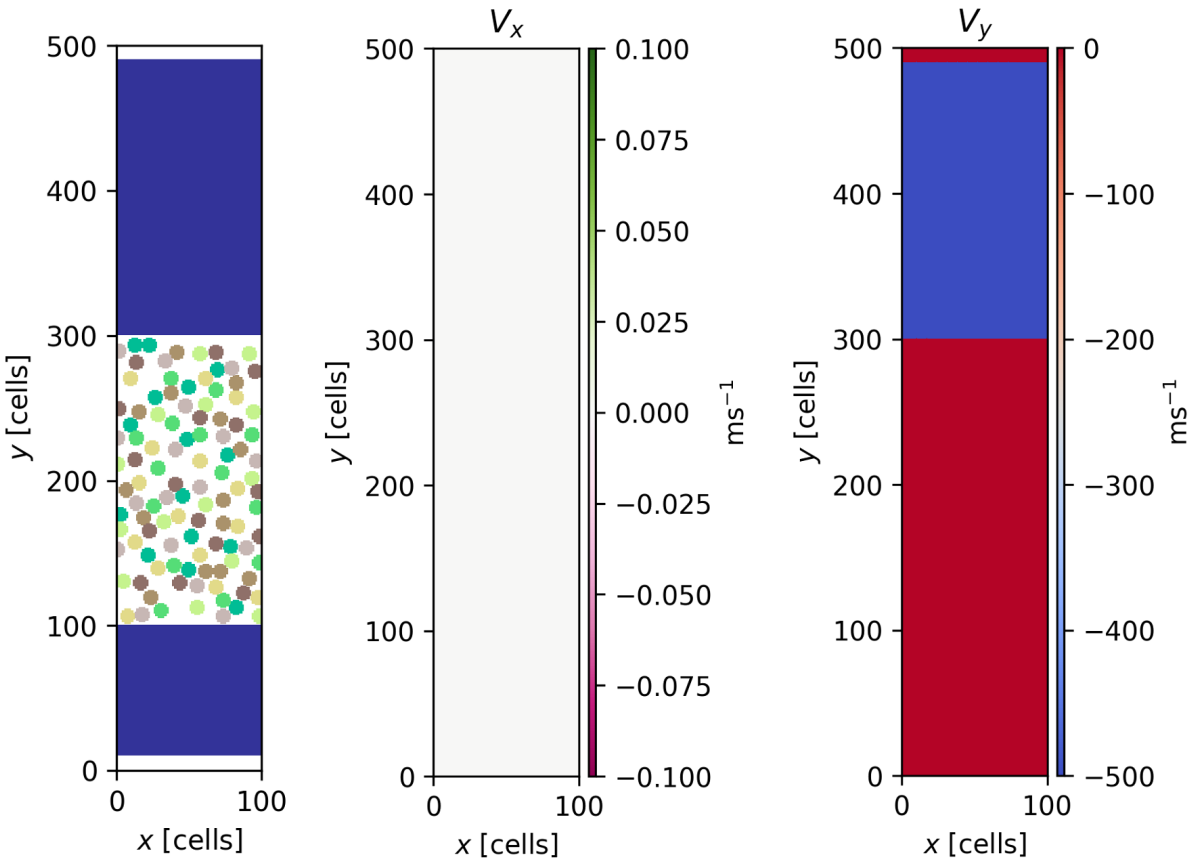


Figure 3.4: The MesoParticles2D iSALE example mesh, as generated by pySALESetup (Collins et al. 2016). The plots are (from left to right) of material number, horizontal velocity, and vertical velocities.

The last remaining class is the '**SizeDistribution**' class. This is an assistive tool, similar to **Ensemble**, that is not essential to the setup routines. It generates a particle size distribution based on either a user-defined function or one of four standard distributions. These are the uniform, normal, lognormal, and Weibull 2-parameter (Fang, Patterson, and Turner 1993) distributions. Then, given an interval of particle size, the class will return the predicted frequency of that size in the domain of the **Mesh**. This method is employed in Chapter 6.

3.3 Benchmarking iSALE and CTH

Before using iSALE to explore the shock compaction of particle beds, several simulations were run with both iSALE and the CTH shock physics code (McGlaun, Thompson, and Elrick 1990) to examine inter-code variability. Simulations to benchmark the two codes were based on numerical experiments described by Borg and Vogler (2008) and set up as close as possible to each other, using geometry and set-up shown in Figure 3.5. These simulations used the same material properties, geometries, boundary conditions (non-periodic), mixed cell strength treatments and the same particle arrangement (Figure 3.8 in Chapter 4). However, for this comparison a bulk porosity of $\phi = 45\%$ ($\equiv 55\%$ volume fraction) for the particle bed was used to allow an additional direct comparison with the experimental data (Borg and Vogler 2008). The benchmark simulations also used grains with 12 cells per particle diameter (cppd), and diameters of 24 μm because this was the preferred resolution of CTH. In order to construct a Hugoniot, simulations were run for a range of impact velocities. In order to reasonably compare the Hugoniots calculated from iSALE and CTH, it was necessary to devise a method for computing the Hugoniot states that could be applied equally to both sets of results.

3.3.1 Calculating the Hugoniot

In general, to calculate the state of a particle bed behind a shock for a given simulation it is necessary to extract a value of shock velocity (U_s), particle velocity (u_p), bulk density (ρ) and bulk longitudinal stress (σ_{LL}). The shock velocity is the speed of the shock wave's movement through the particle bed, which is generally calculated as the slope of a linear regression of shock front position against time. The other variables are calculated as averages over the particle bed behind the shock front of longitudinal material velocity, mass density and longitudinal stress. In this work, we used two different methods to identify the shock front position and calculate the shock variables for constructing the

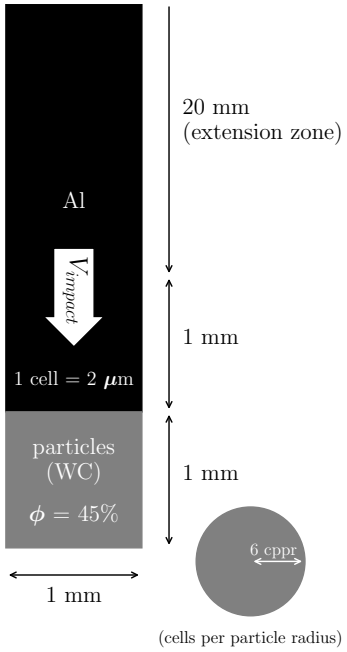


Figure 3.5: The simulation geometry used in the benchmarking. All material parameters and other details can be found in Chapter 4, upon which this test is based.

Hugoniot (the 'transverse average' method and the 'stress contour' method, Figure 3.6, right and left respectively).

The 'Transverse Averaging' Method

The 'Transverse Averaging' method (or 'TransAvg') produced a 1D wave profile in the longitudinal direction at each time step by transversely averaging all fields across the granular bed (Figure 3.6). The 1D profile was used to locate the instantaneous shock front position and find the shock state behind the shock front. The shock front position was identified as the longitudinal position at which the magnitude of the longitudinal stress exceeded half its value behind the shock front from the previous time step. (At time zero the longitudinal shock stress was taken to be zero). The instantaneous shock state (i.e., particle velocity, density, stress) at each time step was found by calculating the mean value of the relevant field between the shock front position and the impactor/particle bed interface position. Cells with no mass (void) were included to calculate the average

shock stress and mass density, but not particle velocity. This transverse averaging method has been used in multiple previous studies, including those using CTH (Borg and Vogler 2008; Borg and Vogler 2009).

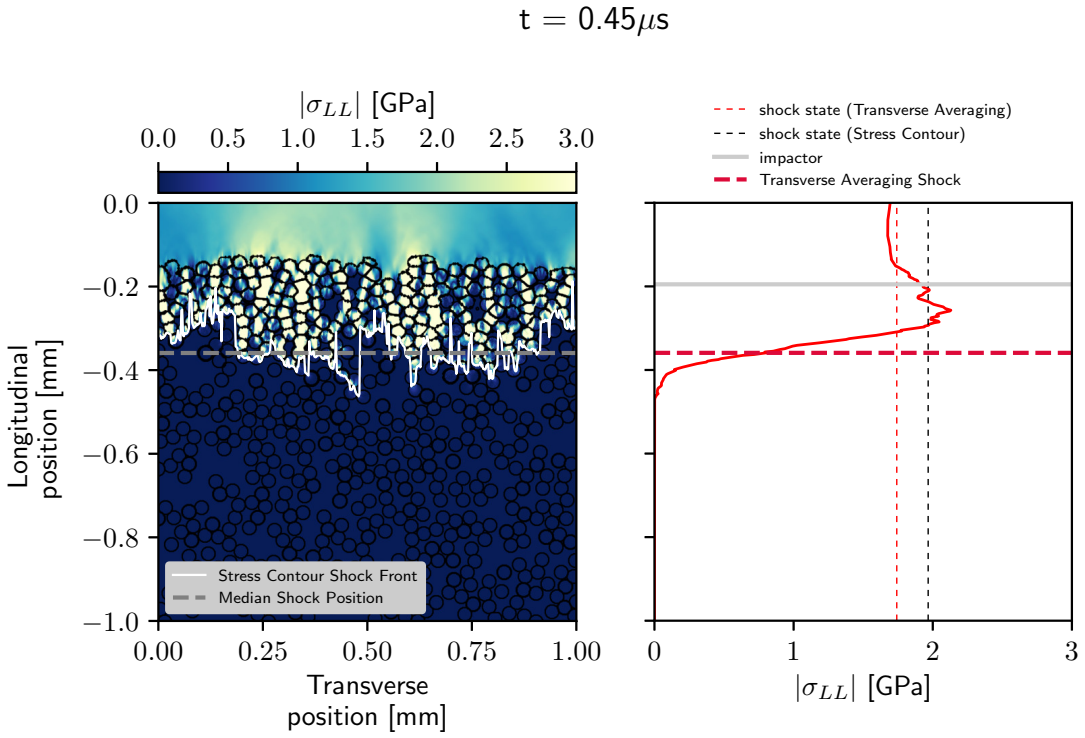


Figure 3.6: Comparison of the stress contour method (left) and transverse average method (right) at a single instance in time. (Left) A plot of longitudinal stress variation in the bed at one instance in time. The shock contour is shown in white whilst the particle positions are shown in black. (Right) The equivalent transversely averaged wave profile, used in the transverse average method at the same time. The impactor position was chosen as the position of the impactor/particle bed interface that was furthest along, longitudinally.

The 'Stress Contour' method

The TransAvg method is fast and simple to implement, however, the act of averaging all cells transversely before processing, will obfuscate features caused by heterogeneities. To address this, a second method was devised to locate the shock front using a contour of longitudinal stress (Figure 3.6) (Davison et al. 2017). this approach is referred to as the stress contour method. The selected stress contour was chosen to be equal to half

the value of the instantaneous longitudinal shock stress value from the previous time step. The contour is undefined at time zero. To track the shock front position with time, we defined the shock position as the median longitudinal position of the shock front contour. The instantaneous shock states of density, longitudinal stress, and particle velocity were found by taking the mean of their values between the shock front contour and the interface between the particle bed and the impactor. Again, empty cells were excluded when averaging particle velocity. As with the transverse average method, the time-averaged shock states were found by averaging the instantaneous shock states over the period that the shock wave was steady.

Comparison of Shock State Extraction Methods

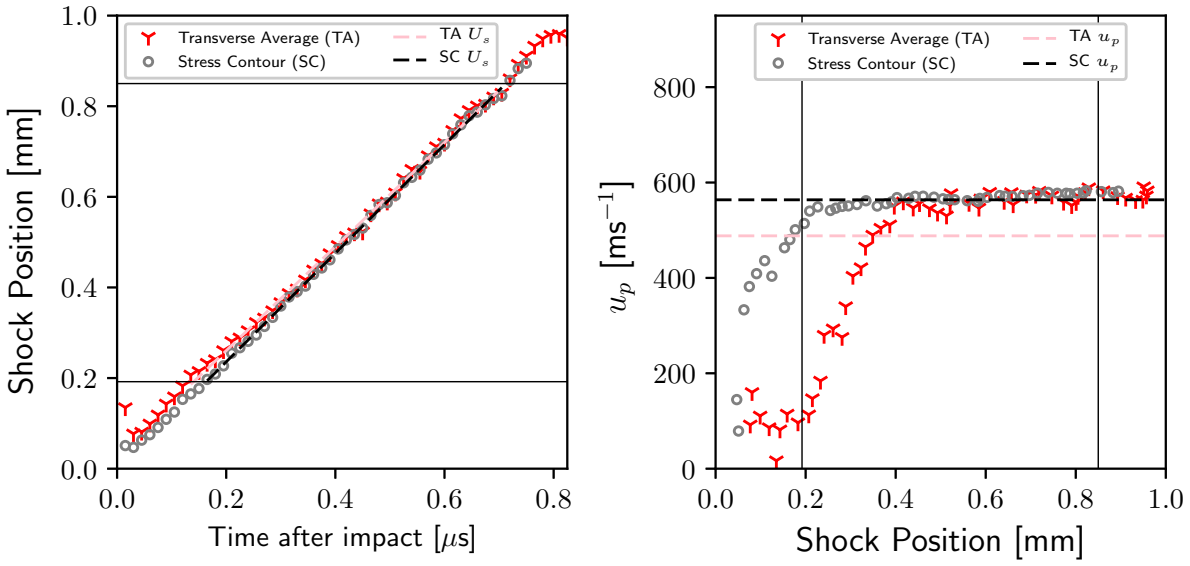


Figure 3.7: Shock position vs time (left) and particle velocity u_p vs shock position (right) for the 'Transverse Average' and 'Stress Contour' methods. Example simulations are taken from Chapter 4. Shock states were calculated by time averaging the fields during the period that the shock was steady. This was whilst the shock front was over six particle diameters into the bed and less than 0.85 mm from the end (solid horizontal & vertical lines respectively). The bulk values of shock and particle velocity in this region are shown (dashed pink and black lines, respectively).

To compare and contrast the two methods for extracting shock state variables from our

mesoscale simulation results, we calculated the instantaneous and time-averaged shock states for a number of simulations using both methods. The time-averaged shock state of a simulation was found by averaging the instantaneous shock states over all times whilst the shock remained 'steady'. A steady shock travels at a constant speed and the shocked material maintains an unchanging state. A shock in a granular material is not steady until it has travelled some way into the particle bed because there is significant variation in shock position along the transverse axis. Nesterenko (2001) states that generally the distance a shock needs to travel in a granular material before it can be considered steady is approximately two particle diameters. In this particular case, the thickness of the shock was ∼4-6 particle diameters so it was assumed that the shock would definitely be steady by the time it reached a position of 6 particle diameters into the bed. In addition, the shock does not remain steady indefinitely and will be affected by any reflections caused by parts of the shock front reaching the end of the bed first. In this case, the position of the shock at which the end of the bed began to affect it was determined to be 0.85 mm into the bed (approximately 4-6 particle diameters from the end).

The two methods differ in their approach to identify the shock front and in how they calculate instantaneous shock states behind the shock front. A consequence of their differences is that the transverse average method took almost twice as long to calculate the steady state compared to the stress contour method (Figure 3.7). The transverse average method did not correctly identify that the shock was steady until it was approximately 12 particle diameters into the bed; 6 particle diameters after it had actually stabilised. A distinctly lower particle velocity was calculated in the time-average as a consequence (≈ 15% difference, Figure 3.7). In contrast, the stress contour method correctly identified the steady nature of the shock by the time it had travelled 6 particle diameters into the bed and generally calculated more consistent shock states.

The difference in calculated shock velocity between the two methods was up to ≈ 10%

(Figure 3.7). In addition, the stress contour method shock states were more consistent between time steps than in the transverse average calculation, which exhibited more scatter (Figure 3.7). For this reason, the stress contour method produced more consistent time-averaged shock states and Hugoniots with significantly less 'scatter'.

The efficacy of the transverse average method depends on the 1D wave profile of the shock being an accurate representation of the actual shock, which is compromised when the shock wave is non-planar. In such cases, the 1D approximation will exclude some shocked material from the shock state calculation, and similarly, will include some unshocked material (Figure 3.6). The closer the shock front is to a planar wave, the smaller this effect will be and the better the transverse average method becomes. The second, stress contour method does not have this limitation because it accounts for the 2D variation in shock front position. It is able to better represent the heterogeneity of the shock front that is characteristic of shock propagation in particulate systems at the mesoscale. For these reasons, the stress contour method was used to calculate the shock states in the mesostructure investigation. The shocks investigated in this thesis often can not be interpreted as planar because of the inherent heterogeneity involved when examining the effects of mesostructure. The Stress Contour method is used throughout (except when otherwise stated) because of this.

The transverse average method was used to calculate the Hugoniot states for both the iSALE and the CTH simulation results because it is a native technique to CTH. However, all other simulations in this thesis make use of the stress contour method because of the extensive heterogeneity present.

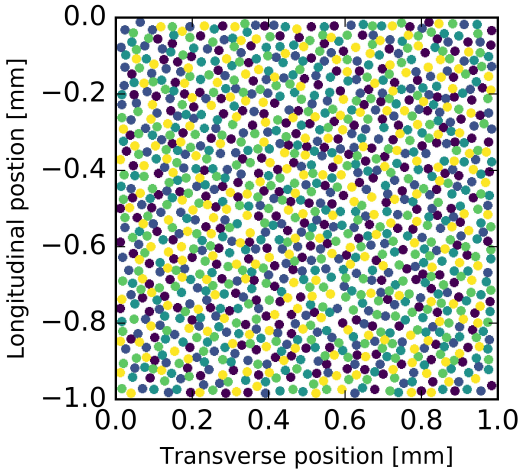


Figure 3.8: The particle bed used for the benchmarking (45% porous). Differently coloured/shaded particles correspond to different 'materials' and were kept constant between the codes.

3.3.2 Results of the iSALE-CTH comparison

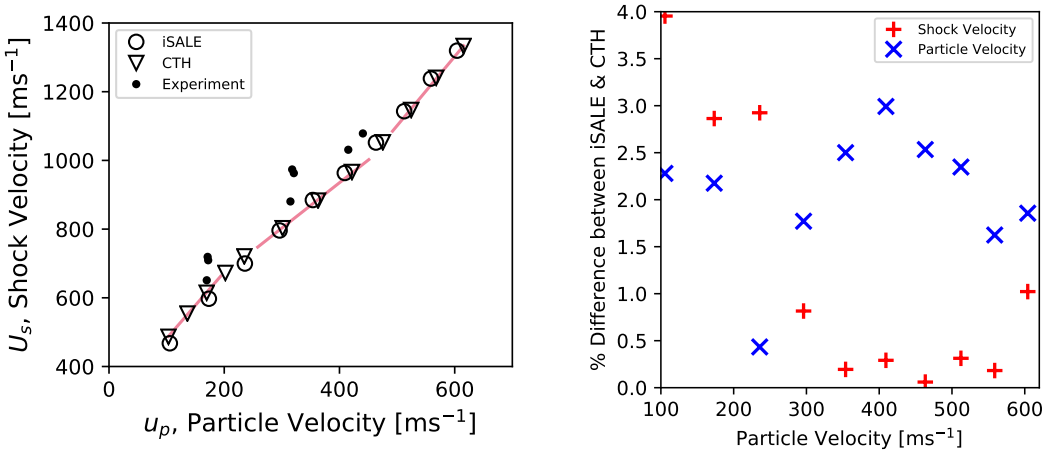


Figure 3.9: (Left) Comparison between Hugoniot responses in iSALE and CTH, for a weighted average mixed cell strength treatment. Lines (red) are included as a visual aid showing the linear 'sections' (Borg and Vogler 2008). (Right) Percentage difference between the iSALE & CTH shock & particle velocity results. Only comparable data is presented here. The two results never differ by more than 4% for either variable.

The Hugoniot response (in U_s-u_p space) inferred from the CTH and iSALE simulations using identical particle arrangements and mixed-cell strength treatments are in good agreement (Figure 3.9). Differences in shock and particle velocities between the simulations is consistently $\lesssim 4\%$ (Figure 3.9). Previous work has suggested that the granular tungsten carbide Hugoniot consists of approximately four linear 'sections' (in U_s-u_p space), possibly reflecting different regimes of behaviour (Borg and Vogler 2007; Borg and Vogler 2008; Borg and Vogler 2009). Although the data presented here do not extend to sufficiently low particle velocities to observe one of these sections, three approximately linear regions are apparent (Figure 3.9).

Figure 3.9 also includes published experimental Hugoniot data for a WC powder (Borg and Vogler 2008). The characteristic grain size (32μm) of the WC powder used in the experiments was larger than that used in our simulations and the grains were highly angular. In other regards, the experimental set-up was very similar to the simulation set-up used here. Over the limited range in particle velocity for which experimental data exist, the experimental and numerical data show a very consistent U_s-u_p slope (Experiment ≈ 1.46; iSALE ≈ 1.56; CTH ≈ 1.32), but the numerical data lie somewhat below the experimental data. This offset can be attributed to the 2D approximation used and a potential underestimation of the strength of individual WC particles. The bulk strength of the particle bed may be underestimated because of the lack of out-of-plane rigidity produced by a 3D grain network. Previous studies have found that fully 3D simulations, over-predict the shock response, shifting the Hugoniot too far up in U_s - u_p space (Borg and Vogler 2009). 3D simulations with 'sliding' (no resistance to intergranular motion), however, under-predict the shock response; they were found to be in good agreement with 2D simulations without sliding (Borg and Vogler 2009). These findings suggest that the use of 3D simulations would address the under-prediction of the shock response, however, they would not be sufficient to bring the results in line with the experimental data. Alternatively, the simulations may not accurately represent the strength of the WC

particles. The dynamic yield strength of WC is not well understood and is likely strain and strain rate dependent. In this work, the shear strength of WC is treated as a constant (see Table 4.1 in Chapter 4), with an estimated value derived from a measurement of the Hugoniot Elastic Limit (HEL) of 6.6 GPa (Dandekar and Grady 2002). Moreover, previous work using CTH found that increasing the yield strength of the individual grains from 5 GPa to 7.5 GPa was sufficient to bring the numerical results in line with the experimental Hugoniot (Borg and Vogler 2007; Borg and Vogler 2008; Borg and Vogler 2009). Hence, it is possible that uncertainty in the WC yield strength or an oversimplification of the WC strength behaviour account for the disparity between experiment and simulation.

3.4 Resolution Test

Finally, resolution choice is an important factor when setting up simulations; it is a primary factor in computational cost alongside domain size. The more cells involved in the simulation, the greater the length of time required for the simulation to complete. The resolution, ideally, should not affect a simulation's output, however, in practice numerical results are often sensitive to the spatial resolution. There is typically a minimum resolution, below which the simulation is insufficient to give meaningful results. At the other end of the scale, there will be an upper resolution beyond which the results converge. Above this resolution increases in resolution will result in increasingly small changes to the results, which will eventually asymptote to zero at a sufficiently high resolution (theoretically). When setting up simulations it is important to know the minimum resolution to achieve meaningful results, as well as an upper resolution above which no major changes are seen. In mesoscale particulate simulations resolution is parametrised as 'cells per particle diameter' (cppd) because the particles are typically the smallest objects in the simulation and the ones of interest. A resolution test using the benchmark simulation configuration was performed in iSALE to find the cppd operating range. This was done

by testing the variation between 2 and 20 cppd for circular grains (Figure 3.10). The bulk shock and particle velocities were used as test fields because of their status as the primary shock state variables in a shock wave experiment.

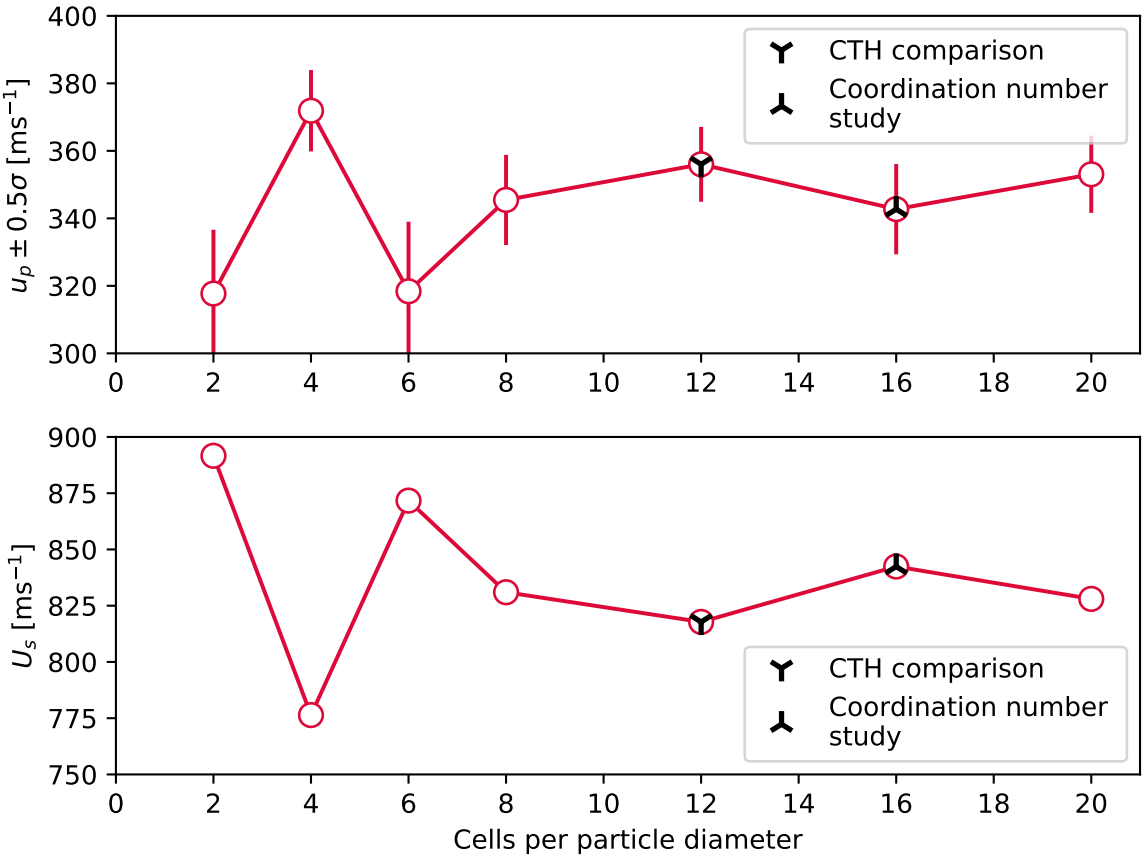


Figure 3.10: Resolution test for multiple cppd. Grain locations, and sizes, were kept constant between simulations; only resolution was altered.

The results demonstrate general agreement between almost all the resolutions considered; the results of the first two resolutions (2 and 4 cppd) differ by only $\sim 16\%$ and $\sim 14\%$ in u_p and U_s respectively. The differences between the simulations reduce significantly for ≥ 8 cppd to $\sim 3-4\%$ and the difference between 6 and 8 cppd is $\sim 5-8\%$. This suggests the results are converging after ~ 8 cppd. No other trend was observed and it was concluded that resolutions above 8 cppd would not exhibit differences $> 5\%$ caused by resolution. The results presented here suggest that the operating range of cppd is 8+;

simulations using less than 8 cppd are likely to possess larger uncertainties and 8 cppd should be the minimum resolution recommended when considering particulate ensembles with varying sizes.

3.5 Summary

iSALE is a shock physics code that solves the compressible Navier-Stokes equations through a first-order, finite-difference, method. In addition to these equations, an equation of state (EoS) must also be chosen to relate the pressure, internal energy, and density. This thesis uses the Mie-Grüneisen EoS in almost all cases, except in Chapter 6, where ANEOS is used (a semi-analytical EoS package). In simulations that use implicit porosity iSALE uses the $\epsilon - \alpha$ porosity model to parametrise porosity and compaction. The porosity model is implemented through the EoS by relating the bulk pressure of a cell to its solid material pressure. This is done by dividing the solid material pressure by distension, and multiplying bulk density by the distension to extract the solid material density.

iSALE was not originally intended to be used for granular compaction simulations and possessed limited set-up routines for this purpose. pySALESetup is a python module that was created to simplify and expand these set-up routines. It was written in python to emphasise ease of use and consists of five base classes **Mesh**, **Grain**, **Apparatus**, **Ensemble**, and **SizeDistribution**. The classes can loosely be categorised as a base class that contains objects (**Mesh**); two object classes that can be created and inserted into the **Mesh** class (**Grain** and **Apparatus**); and two distribution classes that work based on existing **Grain** instances (**Ensemble** and **SizeDistribution**).

iSALE was compared against a well-established shock physics code, CTH, in a benchmark test. The codes found the Hugoniot of a simple particle bed of WC grains by striking it with an Al impactor over a range of impact velocities. In order to make a meaningful comparison between the two codes, two methods used to calculate a Hugoniot from their

results were tested. Out of the two methods, the stress contour method was found to be the most suitable for the simulations presented in this work because it did not require the shock front to be approximately linear. This is predominantly the case throughout this thesis. The transverse average method, however, was used by iSALE in the benchmark, because it is CTH's native method. In general, close agreement was found between the simulation results ($< 4\%$ difference in u_p and U_s for all impact velocities).

Finally, a resolution test was performed for iSALE, using the benchmark set-up geometry. As resolution increased the bulk parameters u_p and U_s appeared to oscillate, and differences on the order of $\sim 15\%$ were exhibited at the lowest resolutions. For resolutions greater than 8 cells per particle diameter (cppd), however, the u_p and U_s did not vary by more than 5%. This suggests that a minimum resolution of 8 cppd for a grain is recommended for granular shock compaction investigations in iSALE.

Chapter 4

Mesoscale simulations of shock compaction of a granular ceramic: effects of mesostructure and mixed-cell strength treatment

4.1 Publication

The majority of the work in this chapter has been published as "Mesoscale simulations of shock compaction of a granular ceramic: effects of mesostructure and mixed-cell strength treatment" by *J. G. Derrick et al.* (Derrick et al. 2018b).

4.2 Introduction

Several recent mesoscale simulation studies have investigated the effects of specific microscale phenomena, such as sensitivity to inter-granular shear-resistance (Borg and Vogler

2013; Dwivedi, Pei, and Teeter 2015), grain morphology (Fredenburg et al. 2014) and grain fracture (Lammi and Vogler 2012), on the macroscale shock response of particle systems. However, as more microscale effects are implemented there is also a need to consider the effects of certain model assumptions that are poorly tested (Borg and Vogler 2009; Borg and Vogler 2013). In particular, it is common in mesoscale simulations to use either a single grain arrangement or multiple random arrangements when constructing a synthetic Hugoniot curve for the particle system. Yet, grain arrangement has been shown to have a discernible effect on shock response (Borg and Vogler 2013; Fredenburg and Thadhani 2013; Lomov et al. 2009; S and Basu 2015) with some studies directly investigating the effects of varying bed structure (Borg and Vogler 2009; S and Basu 2015). In general, these investigations quantify variation in structure between particle beds by the exposure to their respective bed-generation algorithms. They make no attempt to analyse the final bed structure and can not guarantee that their beds do not exhibit structurally similar features.

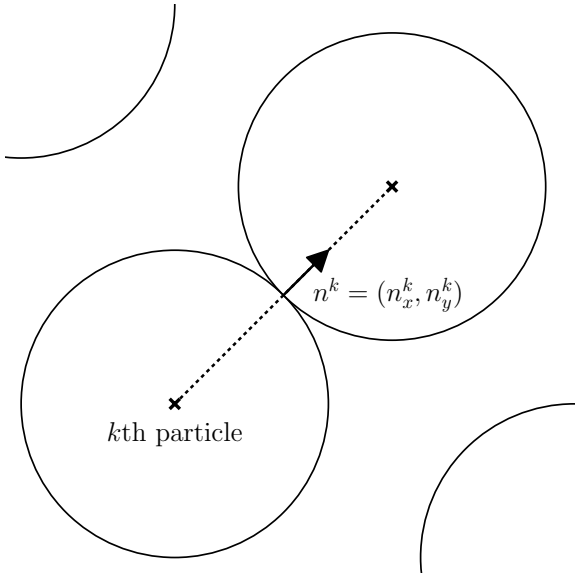


Figure 4.1: An example of a contact between two discs in a particle bed. The normalised contact vector is represented by n^k for the kth particle. The vector for the other particle in the contact would simply be $-n^k$.

By way of contrast, the effect of grain structure on quasi-steady state deformation is well studied and known to be significant (Fredenburg and Thadhani 2013; Kruyt 2012; Olney et al. 2011; Sun, Wang, and Hu 2009; Zhu, Nicot, and Darve 2016). A number of metrics, such as the 'Fabric Tensor' (Li and Li 2009), have been developed to quantify and track the mesostructure in a particle bed. The Fabric Tensor is a second order tensor that typically tracks all contacts between particles. Variations exist that alternatively track voids, or other features (Li and Li 2009). It has two invariances: 'The Coordination Number' (Z) and the 'Fabric Anisotropy' (A). The coordination number is the average number of contacts that an entity has with other entities in the system whereas the fabric anisotropy indicates the general directional preference of contacts in the ensemble (positive A implies contact vectors are generally in the x-direction) (Kruyt 2012; Li and Li 2009). This Chapter focuses on the coordination number Z. The full equations for a 2D particulate system are shown below (Kruyt 2012), where the fabric tensor is represented by $\overline{\overline{F}}$, the number of entities in the system by N, and the total number of contacts by M. In a system of 2D discs each contact vector (n^k) is equivalent to the normalised vector between particle centres (Figure 4.1). The calculation is more complicated in non-circular grain systems (Kruyt 2012). i and j refer to the components (x or y) of each contact vector for the kth contact.

$$\overline{\overline{F}} = \begin{bmatrix} F_{xx} & F_{xy} \\ F_{yx} & F_{yy} \end{bmatrix} \tag{4.1}$$

$$F_{ij} = \frac{2}{N} \sum_{k=1}^{M} n_i^k n_j^k \tag{4.2}$$

$$Z = F_{xx} + F_{yy} \tag{4.3}$$

$$A = F_{xx} - F_{yy} \tag{4.4}$$

Here, the effect of Z on the planar shock compaction of mono-disperse, 2D, particle beds using mesoscale numerical simulations is examined. All forms of mesostructure will depend on friction between particles to exist. The effect of removing this friction on the shock response of these materials is also considered. The simulations in this Chapter are based on previous laboratory and numerical experiments of the shock response of granular tungsten carbide (Borg et al. 2006; Borg and Vogler 2007; Borg and Vogler 2008) to provide a means of benchmarking iSALE simulation results with previous work. The porosity (and hence initial bulk density) of a particle bed will have a significant effect on both the shock response and the range of possible coordination numbers. Therefore porosity was kept constant between comparative simulations so that particle beds differed only in Z (additionally, A was kept as small as possible). The results presented show the mesostructure within the particle bed directly affects the shock response of the particle bed: an increase in coordination number ('more' structure) shifts the Hugoniot (locus of all shock states) to higher shock velocities, lower particle velocities and lower shocked densities for a given impact velocity. Sensitivity to structure implies a sensitivity to shear resistance between grains (inter-granular friction) and it is also shown that removing inter-granular resistance drastically reduces the resistance to compaction.

4.3 Set-up and Methods

Numerical mesoscale simulations were designed to simulate a 1D flyer plate impact experiment. They were loosely based on previous 2D planar mesoscale simulations performed by Borg & Vogler (Borg and Vogler 2008) that modelled the compaction of a granular tungsten carbide (WC) bed with a porosity of 45% (Vogler, Lee, and Grady 2007). This study considers an aluminium (Al) impactor, with an initial velocity (V_I), that strikes a particle bed of tungsten carbide (WC) with 50% porosity (ϕ, defined as the ratio of void space to occupied space in the particle bed). No driver plate was used. The full

set-up is presented in Figure 4.2. The impactor measured 21 × 1 mm, however only the millimetre of the impactor closest to the particle bed was resolved at high resolution; the remainder of the impactor was modelled with progressively larger cells that grew with distance away from the particle bed in order to minimise the computational cost of modelling a long impactor. A long impactor was necessary to delay the arrival of the release wave from its rear until after the compaction wave had traversed the particle bed. The particle bed measured 1 × 1 mm. A range of resolutions were considered, from 2 to 20 cells per particle diameter (cppd), and these were offset against their relative computational costs (see Figure 3.10). Consequently, a resolution of 2μm/cell was used and a particle resolution of 16 cppd employed. This gave the particles a characteristic diameter of 32 μm. Previous studies found that particle size and morphology variations did not affect the shock response of the simulated particle beds (Borg and Vogler 2009). Hence, a mono-disperse size variation was employed here to simplify and constrain the scope of the study, as was done in previous work (Borg and Vogler 2008). The boundaries of the simulation, perpendicular to impact, were frictionless, rigid and not periodic. Open, 'outflow', boundary conditions were used at the top and bottom (behind and ahead of the impact, respectively). The aim of these simulations was to quantify the shock response of the particle bed up to the pressure required for complete compaction. To generate a Hugoniot, a series of simulations were performed using a range of impactor velocities from 150 to 950 ms^{-1}, in intervals of 50 or 100 ms^{-1}, resulting in curves with approximately 10 points.

A Mie-Grüneisen equation of state (EoS) was used for both the aluminium (Al) and the WC as was done in previous studies (Borg and Vogler 2008; Borg and Vogler 2009). The associated constants are presented in Table 4.1. A Johnson-Cook strength model was used for the Al (Johnson and Cook 1985), while the WC particles were treated as elastic, perfectly plastic with a constant yield strength, independent of any other variables such as strain and temperature. Neglecting thermal softening of the WC (and use of a simple

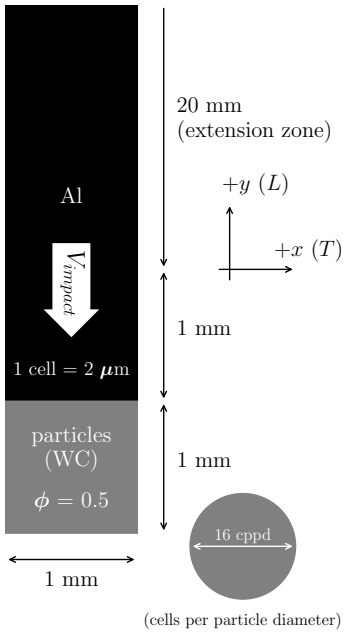


Figure 4.2: Figure showing the geometries used. Simulations were performed in the horizontal plane; the x and y directions are referred to as Longitudinal (L) and Transverse (T), respectively. The mesh is Cartesian, and structured and the simulations were not periodic in the T direction but used solid 'freeslip' boundary conditions. The extension zone used a decreasing resolution to minimise the computational cost of simulating a large impactor. The porosity is denoted by ϕ.

Mie-Grüneisen EoS) is justified given WC's high melting point (3143 K). Even at the highest impact velocity considered the bulk shock temperature of the WC was $\approx$ 1200 K, with a standard deviation of $\approx 400K$. Only a small percentage of cells within the particle bed ($< 0.1\%$) came within 500 K of the melt temperature, indicating that compaction of almost the entire bed was dominated by deformation of particles unaffected by thermal softening.

When particles come into contact, in iSALE, their shared interface is represented by so-called mixed-cells that are a mixture of each particle material (and void). Whilst iSALE does track the interfaces between materials in a mixed cell it does not track the interfaces between particles. Hence, the inter-granular resistance to shear is represented by the strength of these mixed cells: an approach referred to as 'stiction' (Borg and Vogler

2013). iSALE can distinguish between a number of different materials using volume fraction fields, regardless of whether the properties of these materials differ. These different material volume fraction fields are referred to as material numbers. In the simulations five different material numbers were used to represent particles in the bed, with assignments designed to minimise the number of contacts between grains of the same material number. Such instances result in unrealistic artificial 'welding' of particles. All particle beds that were used had zero contacts between particles of the same material number at time zero; five material numbers was the minimum required to achieve this. It is possible for the number of contacts a particle possesses to exceed four during a simulation; however, at such a point the distinction between particles becomes less important because of the removal of most of the porosity.

This chapter investigates the effect of two parameters on the granular WC Hugoniot: the initial arrangement of grains in the particle bed (Z), and the treatment of strength in cells containing multiple materials (grain-grain contacts). Mesostructure was quantified using the coordination number, an invariant of the fabric tensor of the initial bed. The initial density (and hence, porosity ϕ) was kept constant between beds comparing the effects of changing $\overline{\overline{F}}$; A was also kept approximately constant. In a particle bed with a high Z, the network of contacts typically covers more of the bed, and should provide a greater resistance to deformation. However, mesostructure relies on resistance between grains to have an effect. This resistance is approximated in (Eulerian) shock physics codes by stiction: the strength assigned to cells containing multiple materials. Two mixed-cell strength treatments were compared in simulations: in the first the strength of a cell was calculated as a weighted average of the individual strengths of each material occupying it (finite stiction); in the second (the 'friction-less' case) cells with multiple materials were automatically given a strength of zero (no stiction).

Table 4.1: Material parameters used in these simulations, predominantly taken from (Borg and Vogler 2008).

Parameter	WC	Al
Density, ρ (kgm^{-3})	15560	2703
Sound Speed, c_0 (ms^{-1})	5260	5288
Hugoniot Slope, S	1.15	1.3756
Grüneisen Gamma, Γ	1.0	2.14
Specific Heat, C_V (J/kgK)	172.3	861.7
Yield Strength, Y (GPa)	5.0	-
Poisson's Ratio, ν	0.2	0.28
Johnson-Cook model Parameters		
Strain coefficient, A (GPa)	-	0.26496
Strain coefficient, B (GPa)	-	0.42642
Strain rate coefficient, C	-	0.015
Thermal exponent, m	-	1.0
Strain exponent, n	-	0.34
Reference Temperature (K)	-	293

4.3.1 Quantifying Mesostructure in the Particle Bed

Multiple particle beds, containing randomised arrangements, were generated for the purposes of these simulations. Their bulk porosity ϕ and approximate number of particles were kept constant (to within $\approx$ 0.01%), along with the particle size and circular shape. Gravity was not included in simulations to allow for the generation of more varied values of the coordination number; that is, particles could be 'suspended' with zero contacts. Whilst this is an unlikely occurrence in a physical ensemble under terrestrial gravity, such arrangements can occur in very low-gravity environments. Such an arrangement also approximates the distribution of particles in a 2D slice through a 3D particle bed or in beds where the gravity vector is perpendicular to the impact direction.

The amount of mesostructure in the bed was controlled by varying the coordination number Z. Z was the chosen parameter to focus on because of its ease of calculation and interpretation. The average number of contact points in the particle bed is heavily dependent on the initial porosity of the particle bed, the definition of a 'contact', and the

algorithm used to place particles. Any one algorithm can have a multitude of potential particle arrangements, each with significant variation. A bed created at a given initial porosity with a fixed number of particles can only contain a finite number of possible Z values. A theoretical maximum value of $Z = 6$ is possible for a 2D bed of mono-disperse discs at the maximum packing density ($\approx 10\%$ porosity, $\approx 90\%$ volume fraction). This is for circles in an infinite plane; in a finite rectangle the true value will be slightly more porous and is equivalent to the theoretical maximum density (TMD) for that domain. The effect of mesostructure will be more visible in particle beds with initial densities significantly below maximum packing, which are capable of producing the full range co-ordination numbers, $0 \leq Z < 6$. Similarly, if Z has a maximum range of 6, A will have a range of 12: if $0 \leq F_{xx} + F_{yy} < 6$ it follows that $-6 < F_{xx} - F_{yy} = A < 6$ although physical ranges may differ. All chosen beds had $|A| \leq 0.05$, which was deemed sufficiently small.

To produce beds with a variety of Zs, a simple algorithm was employed using pySALE-Setup. Each bed was initialised by placing one particle into the domain at a random location. After this, particles continued to be placed at random until a threshold area fraction was achieved. The random locations were taken from unoccupied coordinates in the bed and particles were only placed if there was no overlap with any pre-existing particles. Once the area fraction threshold was reached the algorithm switched to a 'random walk' method. Particles were 'tentatively placed' (they could overlap other particles whilst in this state) into the bed and underwent a random walk around the bed. Once they were in contact with at least one other grain (and overlapped with any grains by no more than 1% of the area of one grain) they were fully inserted into the domain. This method allowed the algorithm to influence Z directly by varying the area fraction threshold. The threshold was varied over its full range (0 - target area fraction, 50%) to generate a library of particle beds. The coordination numbers of these beds were found and six were chosen for this study based on a pre-defined range (Figure 4.3). Only grain

centres were required to be within the boundaries of the bed and so partial particles were allowed on all boundaries.

Each particle was considered individually to calculate its contacts, and hence, the bed's fabric tensor (Figure 4.1). In the computational domain, each particle is represented by a circular region of cells. At time zero there are no mixed cells; all cells inside a particle are full. A 'contact' between two particles at time zero in the simulation was defined as two, or more, adjacent cells (diagonals were not considered adjacent) from two different grains with no void cells separating them.

Desired bed arrangements were those with variable Z but constant and approximately spatially uniform porosity, ϕ. However, for high coordination numbers some beds displayed significant porosity gradients (Figure 4.3 shows an exaggerated example of such a bed that was not used). To distinguish (and reject) such beds a third metric was used: the maximum porosity difference across the bed, K. This was calculated by first dividing the bed into 5 equal columns and calculating the porosity in each column. The maximum difference between the columns was recorded. This calculation was repeated along the orthogonal axis of the bed (rows in this case) and K was defined as the largest difference in porosity between any two rows or columns. Among the 272 particle beds generated, the maximum porosity difference across the bed was found to be small and relatively constant $K \lesssim 0.1$ for beds with a coordination number $Z \lesssim 1.9$ (Figure 4.3). Hence, only particle beds with $Z > 1.9$ were excluded from the study. Figure 4.4 shows two examples of particle beds that were chosen for this study. None of the generated beds exceed $Z = 2.5$ or come close to the maximum value $Z = 6$, because of the random nature of the particle placement algorithm and the interdependence of ϕ, Z and K. There is negligible chance of constructing a bed with $Z \approx 6$ at random, and with a bulk porosity of 50%, such a bed would necessarily have large K and be excluded.

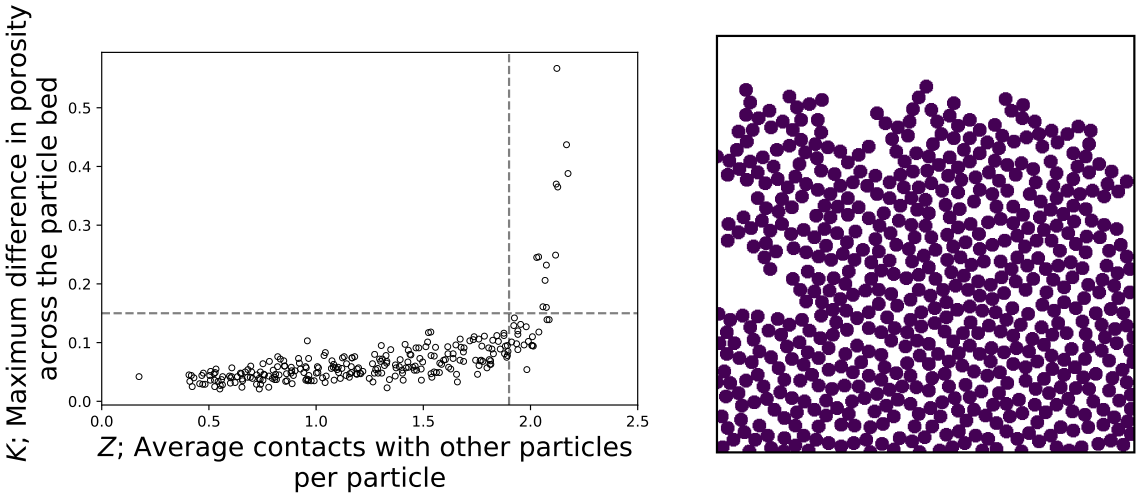
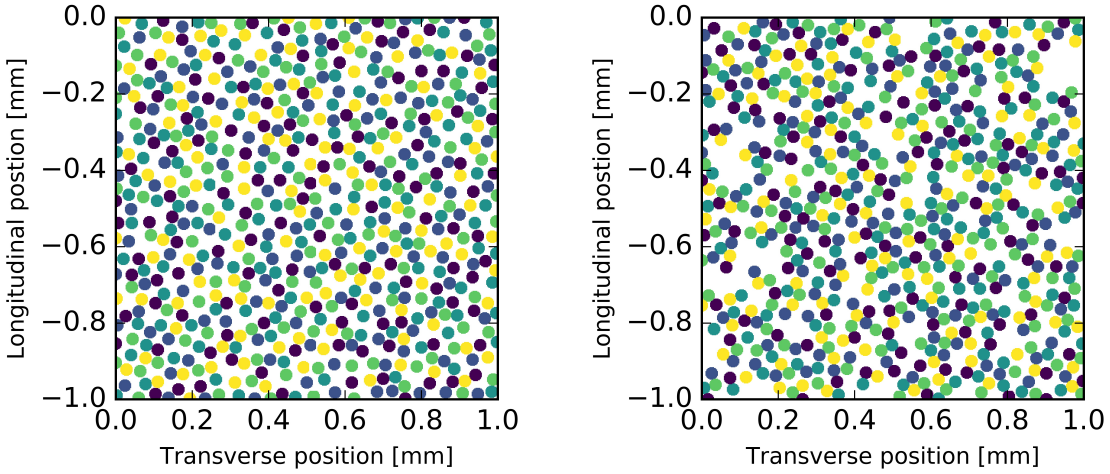


Figure 4.3: (Left) The porosity gradient metric K against the contact number Z for 272 different particle beds. The dashed lines display our guideline limit on K for use in this study and only beds from the bottom-left quadrant were used. (Right) Particle bed with a coordination number of $Z = 2.1$ and a $K = 0.567$ displaying significant porosity gradients. This bed is presented as an example of porosity gradients and was not otherwise used in this study.



Particle bed with a contact number of $Z = 0.4$ (50% porous).

Particle bed with a contact number of $Z = 1.5$ (50% porous).

Figure 4.4: Particle beds, with identical porosities but differing arrangements and hence, different contact numbers. The impact is considered to be travelling in the negative longitudinal direction as in Figure 4.2; the colours/shades indicate different material numbers.

4.4 Results

4.4.1 Sensitivity of the Hugoniot to Z

Hugoniot curves in U_s-u_p space were constructed for suites of simulations with the same initial grain arrangement and coordination number Z (Figure 4.5). As all post-shock parameters can be affected by changing Z, only the initial conditions (V_I, Z) are considered to be independent variables. A common trend in the results is for the U_s-u_p slope to be steep at low u_p, shallower at intermediate u_p and then steepen again at higher u_p. There is no sharp discontinuity between these regions and the Hugoniots appear to transition smoothly as u_p increases. The clearest variation with Z is observed at high particle velocities (> 400 ms^{-1}); as u_p decreases the existence of a monotonic relationship between U_s or u_p and Z becomes more unclear.

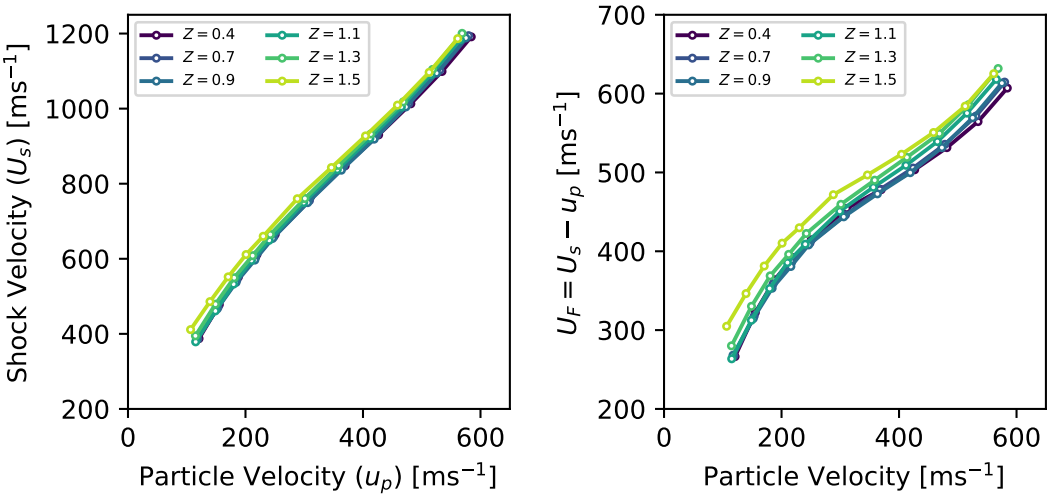


Figure 4.5: (Left) The difference in Hugoniot response of iSALE simulations over a range of coordination numbers (Z) in the $U_s - u_p$ plane with all other factors kept constant between them. Dark to light colour's represent increasing values of Z. (Right) The shock responses in $U_F (= U_s - u_p)$ - u_p space. This space amplifies differences between Hugoniots and presents a more accurate representation of the shock response (Kerley 2013). Standard errors were calculated for each Hugoniot point, however they were not included because they were too small to be visible.

Subtle differences in the Hugoniot curves are difficult to discern in U_s-u_p space because it

is not particularly sensitive to different material properties (Kerley 2013). Shock velocity in the moving frame, U_F ($= U_s - u_p$), as a function of u_p, is generally more sensitive to subtle variations in material properties (Kerley 2013). In this space, the differences and variations observed in U_s-u_p space are amplified and the change in Hugoniot shape with respect to Z is more pronounced (Figure 4.5). At particle velocities exceeding 400 ms^{-1} there is a systematic shift in the Hugoniot to lower u_p and higher U_s with increasing Z at constant V_I, however this trend seems to disappear at lower velocities.

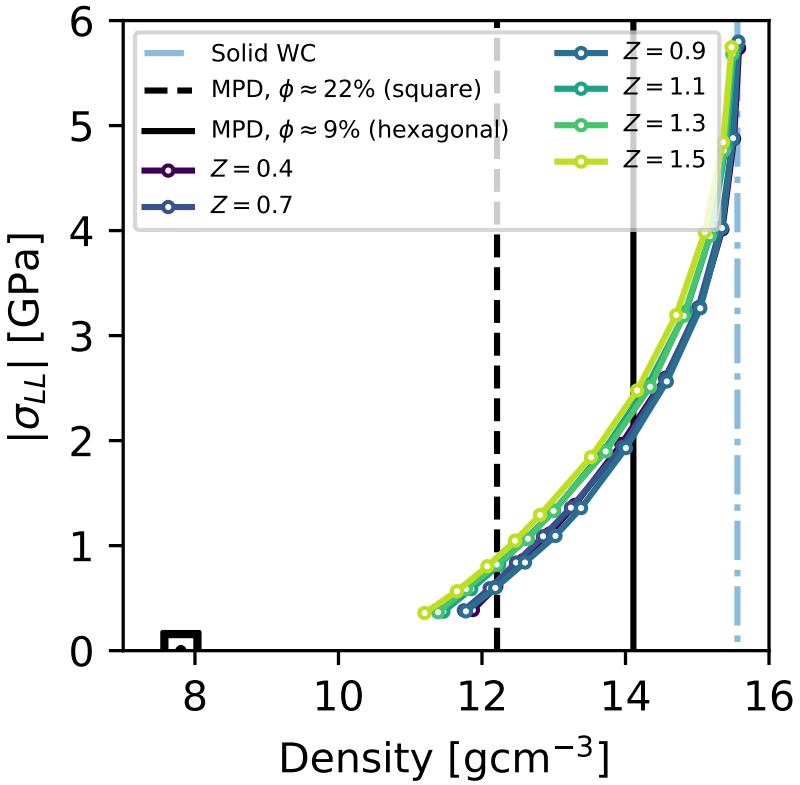


Figure 4.6: The shock response of the particle beds in Figure 4.5 in stress - density space. Dark to light colours represent increasing values of Z. The theoretical 'Maximum Packing Density' (MPD) used by Borg (Borg and Vogler 2009) (square packing) is indicated by a dashed grey line and the theoretical MPD for mono-disperse circles in a rectangle is indicated by a solid grey line (hexagonal packing). The starting density for all simulations is indicated by the $\times$ symbol that is half visible on the x-axis.

Values of σ_{LL} and ρ were extracted explicitly, as described in Section 3.3.1, rather than using the Rankine-Hugoniot Jump Conditions. However, unlike u_p, these fields exhibited

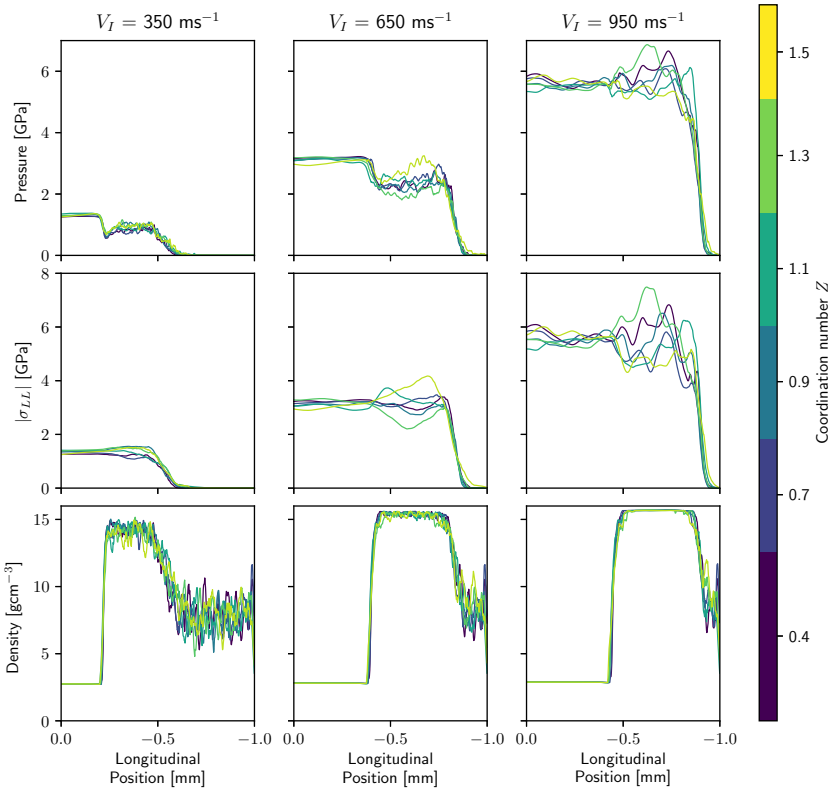


Figure 4.7: Variation in Pressure, longitudinal Stress ($|\sigma_{LL}|$), and Density along the longitudinal axis. The fields have been averaged transversely across the bed and are presented at the time the shock reaches a position of ≈ -0.8 mm. The shock is travelling left to right in all plots and the aluminium impactor is also shown; its location corresponds to the region of significantly lower density in each plot ($\rho \approx 3\text{gcm}^{-3}$).

notable variation across the particle bed behind the shock (Figure 4.7). In particular higher stresses were observed along force chains in the bed whilst grains that did not take part remained largely unaffected. The level of variation observed is consistent with that presented in the previous studies upon which this work is based (Borg and Vogler 2008).

The Hugoniots in stress-density space have less variation in the y-axis (stress) than in $U_s - u_p$ space, with all points achieving approximately the same stress for each impact velocity (Figure 4.6) although there appears to be a slight negative correlation of stress with Z for high V_I. In this space density is more sensitive to coordination number (mesostructure), and the Hugoniot shifts to lower densities as Z increases. However, the difference decreases for higher stresses to the point where it appears to vanish as they intersect with the fully

dense WC Hugoniot. Above this point the particle bed will achieve full compaction and removal of all void-space regardless of the initial mesostructure, and the bulk density will tend to approximately the solid density (≈ 15.56 gcm^{-3}).

4.4.2 Sensitivity of the Hugoniot to the Strength Treatment of Cells Containing Multiple Materials

Figure 4.8 compares two suites of iSALE simulations that use the same initial particle bed ($Z = 1.5$) but differ in their treatment of strength in mixed cells: 'with stiction'—a weighted average of the strengths of material in a cell (which was used for all other simulations so far presented in this work); and 'without stiction', *i.e.* 'frictionless' (any mixed cell has zero strength). Removing stiction for a simulation will make almost all structures involving particles 'unstable' and liable to break down until the particles reach their most stable arrangement possible in the time allowed.

The weighted average results appear as before (Figures 4.5 and 4.6), but removing stiction shows a single, approximately linear, trend in U_F-u_p space. The slope of this linear trend is steep throughout and is most similar to the slope of the high-u_p section of the Hugoniot for the weighted-average case, where the two intersect with each other (Figure 4.8). The removal of stiction allows particles to 'slip' past each other which removes the effects of mesostructure in the particle bed. In $|\sigma_{LL}|$-ρ space the particle bed is compacted to densities exceeding the maximum packing density (MPD) for all data points because particles are free to move until they reach a stable position. Similarly the stiction-less Hugoniot intersects with the weighted average Hugoniot, and the Hugoniot for initially solid WC, for high ρ.

The intersection between these two Hugoniots, and the solid WC Hugoniot, indicates that in this regime the shock response is insensitive to stiction between grains. This is as expected; when the V_I is large enough to compact all porosity out of the bed the effects of

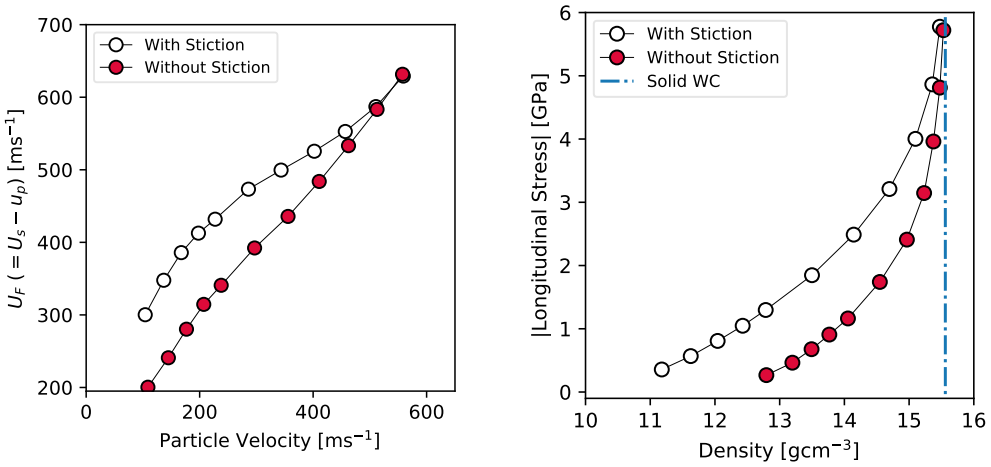


Figure 4.8: Difference between Hugoniot responses in iSALE, using treatment of a cell with multiple materials in it, with and without stiction. The $Z = 1.5$ bed is used in this comparison. The Hugoniots are shown in U_F - u_p space (left) and $|\sigma_{LL}|$ - ρ space (right).

grain-grain interactions on removal of void will no longer matter. Above this point there is no remaining porosity and the response of the homogeneous material dictates how the system evolves.

4.5 Discussion

Our results show that the Hugoniot is directly affected by variations in grain arrangement (Figures 4.5 and 4.6); as Z increases u_p, ρ and $-\sigma_{LL}$ appear to decrease and U_F increases. Spearman's rank correlation coefficient (represented by the symbol 'R_S' instead of ρ to avoid confusion with the symbol for density) can be used to quantify these progressions. Spearman's rank assesses how well the relationship between the two variables is represented by a monotonic function (Spearman 1904; Dodge 2008). R_S can vary from -1, strong decreasing trend, to $+1$, strong increasing trend; a value of 0 implies no correlation. U_F, u_p, rho and $-\sigma_{LL}$ have moderate to strong correlation with Z at all impact speeds. Correlation of u_p with Z is strong for all V_I ($|R_S| > 0.9$, Table 4.2) and the strength of correlation of U_F against Z increases sharply between $V_I = 450\text{ms}^{-1}$ and $V_I = 550\text{ms}^{-1}$

($u_p \approx 300 - 400 \text{ms}^{-1}$). Between these values it jumps from $R_S \approx 0.5 - 0.6$ (moderate) to $R_S \approx 0.8 - 0.9$ (strong), almost doubling (Table 4.2). Density similarly correlates strongly with Z, almost as strongly as u_p, but with more variation. Stress correlates strongly at low, and mid, impact velocities, however it decreases significantly for $V_I > 750$ ms^{-1}. At the highest V_I there is essentially no correlation visible. In this region there is full compaction of the particle bed (intersection with the solid WC Hugoniot) which implies particle deformation behind the shock. Z ceases to become a predictor of the post-shock stress because all stresses cause total deformation.

These correlations suggest that the coordination number Z is a useful predictor of the effect of mesostructure on shock response. However, further investigation is required to understand the jump in correlation for shock velocity. Shock velocity is intrinsically linked to the mechanisms of shock compaction as these dictate how the shock propagates, and an examination of whether this change in dependence of U_F on Z, corresponds to a specific state of compaction follows.

Table 4.2: Spearman's rank correlation coefficients (R_S) measuring the strength of correlation of U_F, u_p, ρ and $-\sigma_{LL}$ against Z. This table is a summary of the results presented in Figures A.2 A.1 A.3 and A.4 (Appendix)

Impact Velocity [ms^{-1}]	Spearman's Rank R_S			
	U_F - Z	u_p - Z	ρ - Z	$-\sigma_{LL}$ - Z
150	0.60	-1.00	-0.94	-1.00
200	0.43	-0.94	-0.94	-0.89
250	0.49	-0.94	-0.83	-0.94
300	0.60	-0.94	-0.77	-0.89
350	0.49	-0.94	-0.83	-1.00
450	0.60	-0.94	-0.71	-1.00
550	0.83	-0.94	-0.71	-1.00
650	0.83	-0.94	-0.83	-0.94
750	0.94	-0.94	-0.94	-0.66
850	0.89	-0.94	-1.00	-0.66
950	0.89	-0.94	-0.94	-0.14

Figure 4.9 presents the Hugoniots in ϕ - u_p space, and in U_F-u_p space for comparison. Porosity was calculated in a similar manner to u_p: $\phi = 1-$ the mean of the volume fraction

of cells in the shocked material. *Borg & Vogler* (Borg and Vogler 2009) proposed that the granular Hugoniot in the low u_p range undergoes a transition (change in U_s-u_p slope) at a porosity of $\phi \approx 21\%$, corresponding to the MPD for circles in a box, using square packing $MPD = \pi/4 \approx 0.79$. Maximum packing density can also be defined using a hexagonal arrangement of circles, which gives a greater solid volume fraction of $\frac{\pi\sqrt{3}}{6} \approx 0.91$ and a porosity of $\phi = 9\%$. Both values are included in Figure 4.9 as dashed and solid lines, respectively.

The Hugoniot curves presented here display the two distinct slopes in the low u_p region described in previous literature (Borg and Vogler 2009), but their transition does not occur as a sharp discontinuity at the square-packing MPD (Borg and Vogler 2009). However, the increase in correlation strength of U_F against Z, from moderately correlated to strongly correlated, occurs at approximately the hexagonal packing MPD ($\phi = 9\%$, $V_I = 450\text{ms}^{-1}$). This observation suggests the possibility that when compaction exceeds the MPD the pre-shock Z becomes a primary predictor for the shock velocity and that the hexagonal MPD is a key state of compaction in the compaction of a particle bed.

The effect of stiction on the shock response in numerical models has been discussed in previous literature, where removal of stiction in 2D simulations generally caused the simulations to underestimate the shock response of an experiment for a given V_I (Borg and Vogler 2013). Additional evidence for this is provided here. Z has a direct effect on the Hugoniot as it parametrises structure in the particle bed. It can be inferred that stiction between particles may have a significant effect as this would alter the ability of the particle networks to transmit the shock. Removing stiction from the models entirely could allow for greater inter-particle motion (total structure collapse), uninhibited by stiction between grains, and generally result in higher post-shock densities. Our results show that all previous features of the Hugoniot are absent, such as the three linear 'sections' (Figure 4.8), and that the post-shock densities achieved are greater than their stiction counterparts

in every case. This implies that variations in the Hugoniot from this relationship are due to bed structure.

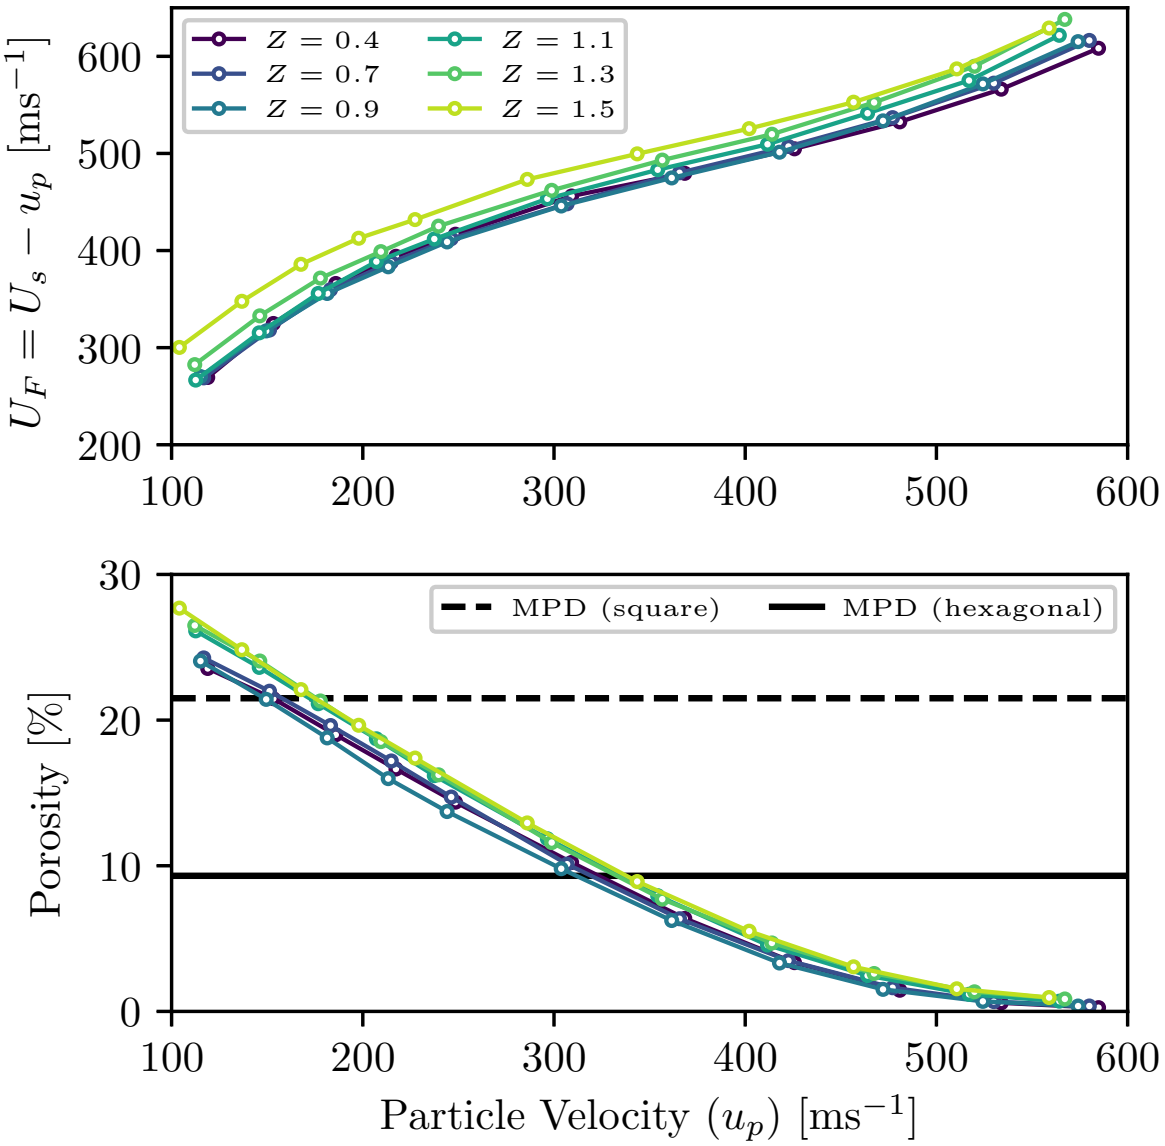


Figure 4.9: The Hugoniots from Figure 4.5 (top) with direct comparison to the Hugoniots in the corresponding ϕ - u_p space (bottom). The dashed and solid lines represent the equivalent porosity for the maximum packing density (MPD) of circles in a box for square and hexagonal packing respectively.

In all the spaces considered, the least sensitive state variables to the mesostructure are longitudinal stress and particle velocity, displaying minimal variation with Z. This insen-

sitivity is well-known (Kerley 2013) and it can be inferred that the least sensitive plane for comparing the shock response of particle beds with differing mesostructure will be the $-\sigma_{LL}$ - u_p plane (Figure 4.10). By minimising the effects of more sensitive variables it is possible to observe subtle changes in these insensitive variables. As V_I increases the range of u_p and $|\sigma_{LL}|$ increase from $\sim$ 14ms^{-1} and $\sim$ 0.03 GPa to $\sim$ 23 ms^{-1} and $\sim$ 0.1 GPa. However, the percentage range decreases from $\sim$ 13% and 8% to $\sim$ 4% and $\sim$ 2% respectively.

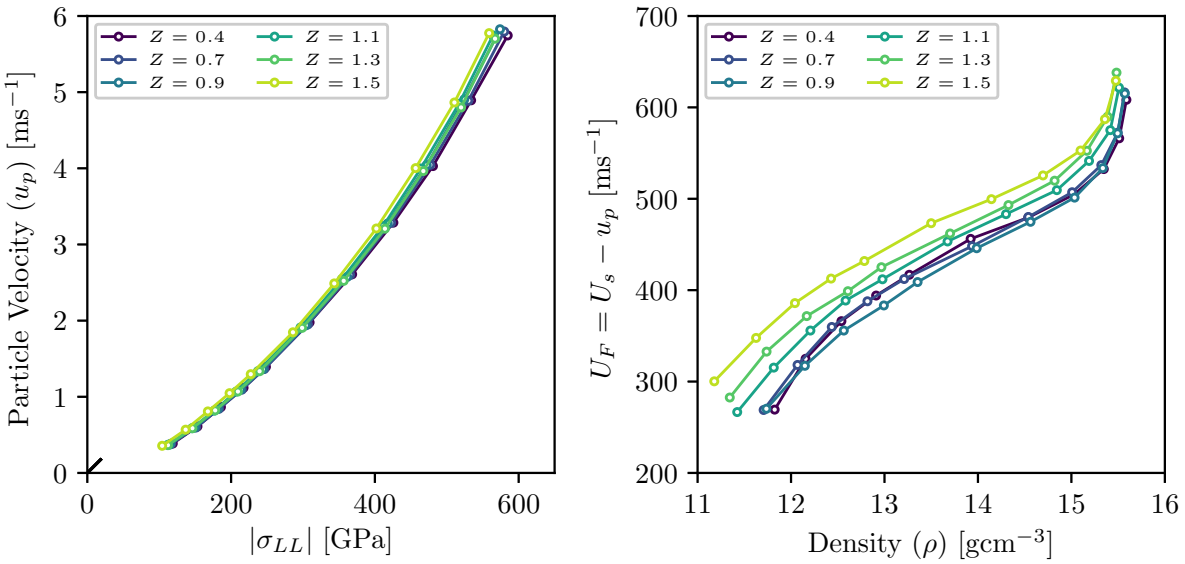


Figure 4.10: The Hugoniots from Figures 4.5 and 4.6 in $|\sigma_{LL}| - u_p$ (left) and $U_F - \rho$ space (right). (left) σ_{LL} and u_p are the two state variables that are least sensitive to mesostructure; Hugoniots in this space exhibit the least variation because of Z. (right) Density and U_F are the fields that are most sensitive to variations in Z therefore Hugoniots in this space present the effect of Z with the greatest clarity.

However, whilst u_p and $|\sigma_{LL}|$ have displayed the least effect caused by changing Z, as previously discussed, U_F and ρ have displayed the largest effect. The effect of Z on the Hugoniots will be plainest in $\rho - U_F$ space as presented in Figure 4.10. The trends are clearest in this space with the Hugoniots shifting to lower densities and higher shock velocities as Z increases.

The variation observed in U_s $(2-6\%)$ and u_p $(5-15\%)$ as a result of Z, is significant when compared to similar experiments. The original studies on which this investigation was based, present experimental Hugoniot data with small uncertainties, but a large scatter. The uncertainty on U_s was $\approx 0.5\%$ (Vogler, Lee, and Grady 2007), an order of magnitude less than the variation observed $(2-6\%)$. Additionally, for an approximately constant u_p, the studies displayed a spread in U_s of the order of $\approx 10\%$ (Figure 3.9) (Borg and Vogler 2008). The scatter is of a comparable order of magnitude to the variation observed due to Z but an order of magnitude larger than the reported uncertainty. This result suggests that the unquantified mesostructure could be a possible cause of this scatter. However, in studies of shock compaction of granular materials it is not common to quantify the mesostructure or grain arrangements of the pre-compacted samples, beyond measuring the bulk porosity. In experimental studies, it is practically impossible to use a common mesostructure for all experiments along a granular Hugoniot. Whilst it is possible to record the initial mesostructure in a particle bed before an experiment, it is typically difficult and costly (performing tomography on a sample before an experiment, for example). A simple method for dealing with the unknown variations in bed structure would be to combine results over multiple beds. In contrast, it is significantly more practical to record the pre-impact structure in mesoscale simulation studies. In the case of numerical studies, however, this is typically not presented and algorithms for generating initial grain arrangements used in different studies are typically not the same and are sometimes not specified at all.

4.6 Conclusions

This Chapter presents the results of a computational investigation into the effect of mesostructure on shock propagation through granular tungsten carbide. For the same bulk density, structure within the particle bed is defined by the average number of con-

tacts between particles (the coordination number, Z). The position of the Hugoniot in U_F-u_p space and $|\sigma_{LL}|$-ρ space is sensitive to the value of Z, and hence to the amount of mesostructure in the particle bed, for the same initial porosity/bulk density. A bed with larger Z (more mesostructure) generally has a faster shock speed, lower post-shock bulk density, and lower particle velocity. For a given impact speed the particle and shock velocity in the moving frame (U_F) correlate strongly with Z. However the strength of correlation of U_F with Z jumps from moderate to strong above an impact speed of 450 ms^{-1}. The transition corresponds approximately to the stress required to compact the bed to the equivalent hexagonal MPD for 2D circles in a box. Coordination number, therefore, appears to be a good predictor of shock response, and its predictive efficacy increases significantly for stresses that result in near complete compaction of the particle bed. This may be because there is insufficient time for particle rearrangement at higher stresses. At lower stresses, involving low-to-moderate compaction, shock speed and bulk density appear to also depend on aspects of mesostructure beyond the simple measure of number of contact points. A more nuanced metric may be required to fully describe the dynamic compaction of a particle bed. Future work should consider the evolution of the coordination number with time within a bed; this is non-trivial because of the presence of deformation at many V_I.

Our numerical results also show that the Hugoniot depends on stiction between grains. Stiction is an approximation of inter-particle shear resistance (friction and interlocking) that is used in Eulerian mesoscale simulations. The amount of stiction is varied by changing the strength of cells containing multiple materials, where two or more particles are in contact. Reducing this strength to zero, in all 'mixed' cells approximates removing inter-particle shear resistance entirely and produces a single linear U_s-u_p Hugoniot. The removal of stiction prevents particles locking together, making most existing contacts unstable structures, and encourages particle motion into more stable arrangements. Different methods for treating strength in cells containing multiple materials will likely result

in different Hugoniot curves, which should be accounted for when comparing results from mesoscale simulations.

This study has shown that varying the mesostructure within a mono-disperse particle bed has a distinct effect on its shock response. A simple method to account for this variation would be to combine results from multiple particle beds for the same model or experiment, but the coordination number, if tractable, provides a useful metric for estimating the effect. This work, however, focuses exclusively on the mono-disperse, circular, 2D scenario. Extending the investigation to 3D is necessary for developing a full understanding of the compaction mechanics. Studies have already shown that the transition is not trivial (Borg and Vogler 2013). In addition, it is important to examine whether the effect of mesostructure on shock response observed here extends to poly-disperse, complex morphology particle beds. The coordination number may be an ideal metric for studying such cases but a greater understanding of its full effect, such as evolution in time, is required.

Chapter 5

Interrogating the shock compaction of bimodal particle size distribution powders through numerical modelling and experiments: the shock compaction of precursor meteorite materials

5.1 Publications

Some of the work in this chapter has been published, and is referenced accordingly. Some is in the process of publication. The relevant publications are as follows: *Derrick JG, Rutherford ME, Davison TM, et al.*, Interrogating Heterogeneous Compaction of Analogue Materials at the Mesoscale Through Numerical Modeling and Experiments, 20th Biennial

Conference of the APS Topical Group on Shock Compression of Condensed Matter, Publisher: AIP Publishing, ISSN: 1551-7616 (Derrick et al. 2018a); *Davison TM, Derrick JG, Collins GS, et al.*, Impact-induced compaction of primitive solar system solids: The need for mesoscale modelling and experiments, Procedia Engineering, ISSN:1877-7058 (Davison et al. 2017); *Rutherford ME, Chapman DJ, Derrick JG, et al.*, 2017, Probing the early stages of shock-induced chondritic meteorite formation at the mesoscale, Scientific Reports, Vol:7, ISSN:2045-2322 (Rutherford et al. 2017).

5.2 Introduction

Granular mixtures are extremely common geologically (sand and soil are the most well-known examples) but they are often simplified in simulations as single-substance monodisperse, spherical-grained, powders (Borg and Vogler 2013, for example). This simplification is often justified, however, such as in terrestrial powders, which are subject to chemical erosion and may not have significant particle size distributions. In addition man-made, machined, powders are often chosen to possess spherical grains because it simplifies their study, as well as classification.

There is a common particulate system, however, for which monodispersity fundamentally does not apply, although grains can still be approximated as spherical. These are bimodal particle size distributions. These are powders that possess essentially two types of particle, one significantly smaller/larger than the other (typically by at least an order of magnitude). In these systems the smaller grain powder is often approximated as a continuum, in which the larger grains are embedded.

These bimodal systems are common when mixing man-made powders because they are the simplest particle size distribution available that is not mono-disperse. Investigations into the response of bimodal powders to shock loading has revealed a variety of interesting features unique to this set-up. Freim, McKittrick, and Nellis (1995) found that they

achieved final densities that were 10% higher than unimodal powders when shock loading a bimodal mixture of Al_2O_3-ZrO_2 (large grains $\sim 20\mu$m and small $\sim 2\mu$m) with the same conditions. Herbold et al. (2008) considered how a bimodal particle size distribution (Al and W particles) embedded in a PTFE continuum responded to shock compression. They found that when the Al and W particles were of comparable sizes (and fine, Al-2μm and W-1μm) the substance possessed an usually high compressive strength in comparison to a more stark bimodal system (Al-1μm and W-10μm).

Despite the reputation of naturally occurring powders to not be as perfect as man-made powders, geometrically, there is a scientifically important situation which has become of great interest in recent years: meteorites. Meteorites are among the most studied rocks in existence because they are the primary source of information about the conditions present in the early solar system (Bland et al. 2014; Davison, Collins, and Bland 2016). By decoding the processes meteorites have been exposed to, and when this has happened, the conditions of the early solar system can be inferred. Precursor meteorite matter consisted of large, spherical, solidified melt droplets ('Chondrules' O(mm)) surrounded by an extremely fine ($O(\mu$m)), porous ($\phi > 50\%$), dust known as 'Matrix'. Meteorites are classified by the relative exposure of this precursor matter to three processes: aqueous alteration, thermal metamorphism, and impact processes. Impact processes include low velocity collisions (< 1 kms^{-1}) between asteroids in the early solar system (Bland et al. 2014). Shock classification of meteorites has traditionally been calibrated against relatively non-porous materials (Stöffler, Keil, and Edward R.D 1991). Bland et al. (2014) found that low-velocity collisions produced large temperature variations ($>$ 500K on the mm-scale) in the matrix whilst leaving chondrules relatively unaffected. Subsequent studies (Davison, Collins, and Bland 2016; Davison et al. 2017) have considered the shock response of these mixtures in greater detail by varying input parameters such as chondrule abundance and matrix porosity. Fundamentally, however, these studies have suggested that shock compaction was more prominent as a process than first thought.

The shock response of a chondrule-matrix mixture is a complex problem and it is difficult to extract detailed mesoscale information. Fundamentally, however, a bimodal mixture has a bimodal distribution in particle size and its response will depend on how a single chondrule-matrix system (one chondrule embedded in matrix) responds to shock compression. This chapter aims to consider this scenario in detail and also whether the conclusions drawn can be extended to multi-chondrule mixture systems. To investigate this, two linked investigations were devised. The first considered a pure matrix system and a single chondrule-matrix system, which allowed for explicit characterisation of the bimodal mixture dynamic in experiments and models. The second investigation considered multi-chondrule mixtures, analogous to precursor meteorite materials, and investigated how applicable the conclusions drawn in the single-chondrule scenario were to multi-chondrule mixtures.

In order to study these complex systems in detail a joint numerical-experimental approach was devised. The system was investigated using simple materials that were analog to precursor meteorite matter (Davison et al. 2017; Rutherford et al. 2017), to allow for repeatability and to enable easy simulation. The simulations were performed using iSALE as described in Chapter 3 and Bland et al. (2014) and Davison et al. (2017). Experiments were performed at the European Synchrotron Radiation Facility (ESRF) utilising state-of-the-art, *in-situ*, mesoscale X-ray imaging techniques. Agreement between experiments and simulations allowed for experimentally inaccessible features to be accessed through the models.

5.3 Methods

5.3.1 Analog Materials Choice

Chondritic mixtures consist of two materials: spherical, solidified, zero-porosity, melt droplets (chondrules) and a fine porous powder (matrix). Chondrules are typically between 0.1mm and 1 mm in diameter, whilst the matrix tends to consist of sub-μm monomers (Bland et al. 2014); despite variations in chondrule size there is still a gap of three orders of magnitude between the matrix grains and chondrules. This gap validates the bimodality assumption. The chondrules themselves rarely consist of a single material but are essentially silicates with varying mixtures of materials such as olivine and pyroxene (H. and Brigitte 2012). Chondrules have also been approximated as forsterite, olivine, and serpentine (Bland et al. 2014; Davison et al. 2017). Similarly the matrix is made up of broadly the same material as the chondrules, except it is significantly more porous, possessing an initial porosity of $\gtrsim 60\%$ (Bland et al. 2014). Given the necessary similarity between the chondrules and matrix, as well as their silicate nature, the analog material was chosen to be SiO_2. The chondrule analog was specifically chosen to be soda-lime glass beads because of their ease of acquisition and silicate nature. The matrix analog was chosen to be a fine-grained silica powder known as Sipernat silica 320-DS because of its small grain size, $\sim 7\mu$m, and high natural porosity ($\gtrsim 90\%$), which had to be pre-compressed before the experiments (Rutherford et al. 2017). A SEM image of Sipernat grains is present in Figure 5.1.

5.3.2 Geometrical Set-up & Material Parameters

The experiments and simulations were based on the same geometry. A 25 mm $\times$ 12.7 mm cylindrical sabot with a 2 mm $\times$ 12.7 mm copper or polycarbonate flyer—attached to one end— struck a copper or polycarbonate (PC) driver—$\approx$ 2.6 mm thick—with an

incident velocity of ≈ 600 ms^{-1}. The driver transferred the shock into the particle bed, which had dimensions of 9.8 mm $\times$ 6 mm (4 mm in the multi-chondrule mixtures) within an aluminium cell and with a PMMA window (full geometry in Figure 5.1). The Sipernat had a bulk porosity of $\approx 70\%$. In the single-chondrule investigation a single borosilicate glass rod (8 mm in length, 1 mm diameter), was used as the chondrule analog because of its comparable density to solid silica. When positioned along the X-ray axis it enhanced contrast in the radiographs and created a quasi-2D experiment. In the multi-chondrule mixture stage two sizes of spherical soda-lime glass beads were used as chondrule analogs, 200 μm and 500 μm diameter. These diameters were chosen to most closely match real chondrule length scales (Bland et al. 2014; Davison, Collins, and Bland 2016; Davison et al. 2017). In addition, the beads were mixed with the matrix such that they took up 30 vol%, the matrix occupied 70 %vol (with a porosity of 70 %) and the mixture had a bulk porosity of ≈ 50 % in accordance with predicted mixture make up based on meteoritic evidence of pre-compacted proportions of chondrites and matrix for a representative meteorite (Bland et al. 2014).

Experiments were performed at the ESRF on beam line ID19. Shots were performed using a 3 m, 12.7 mm bore, single stage light-gas gun aligned perpendicular to the X-ray beam. Four radiographs were captured indirectly per experiment using a single crystal LYSO scintillator (Rutherford et al. 2017) imaged by a pair of PI-MAX4 cameras, each capable of taking two images (although, in the multi-chondrule mixture experiments, only one camera was used). The intervals between each frame were 0.878μs, 0.708μs, and 0.878μs respectively. Their field of view was restricted to a 12.1 mm $\times$ 12.1 mm square that encompassed the particle bed and some of the surrounding apparatus (Figure 5.3) with an effective pixel size of 11.85 μm $\times$ 11.85 μm. In the single chondrule scenario a cylindrical rod oriented along the X-ray axis was used in place of a spherical bead to enhance contrast in the radiographs. It was a quasi-2D experiment.

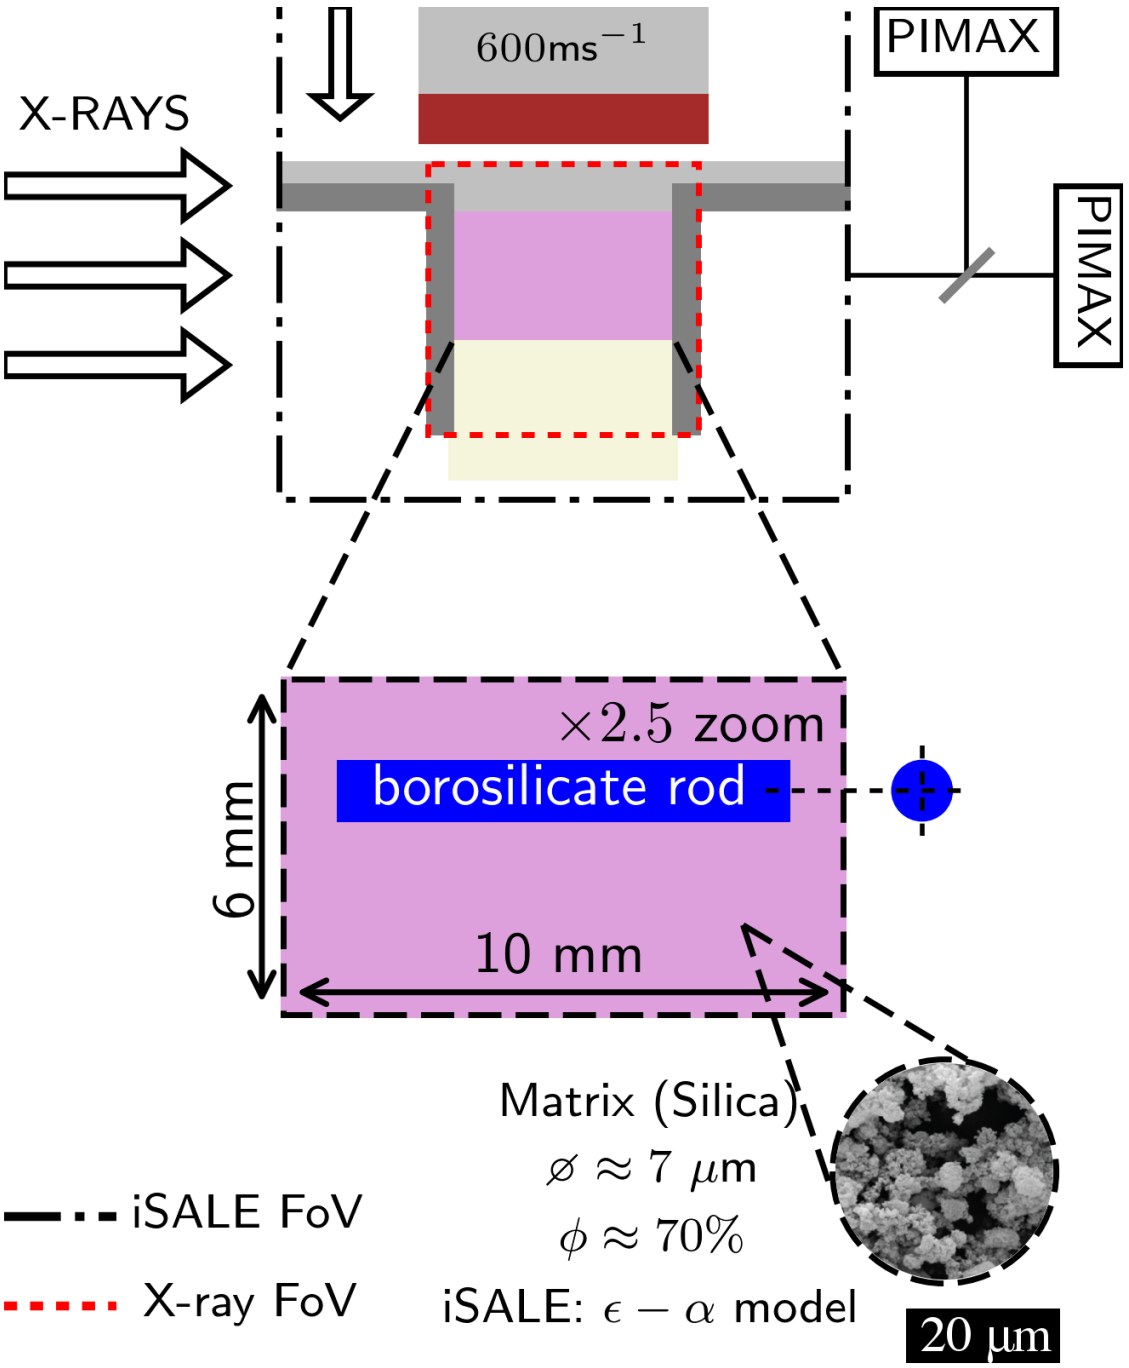


Figure 5.1: Full geometry (to-scale) used in the experiments and models based on previous work by *Rutherford et. al* (Rutherford et al. 2017). Only part of the impactor is shown.

The numerical models resolved the full cross-section geometry of the experiment in the simple 2D plane strain scenarios. When modelling pure matrix alone, however, cylindrical

Table 5.1: Full material models used.

Material	Equation of State	Strength Model
Polycarbonate	Mie-Grüneisen	Von-Mises
Aluminium	Mie-Grüneisen	Johnson-Cook
Copper	Mie-Grüneisen	Johnson-Cook
PMMA	Tillotson	Von-Mises
Boro-silicate glass	Mie-Grüneisen	Von-Mises
Soda Lime glass	Mie-Grüneisen	Von-Mises
Silica (matrix)	Mie-Grüneisen	Drucker-Prager
	$+ \epsilon - \alpha$ porosity model	

boundary conditions were used because of the set-up's axial symmetry. Whereas, only the central 10 mm was modelled when considering the multi-chondrule mixture particle beds. For the boundary conditions the models use 'freeslip' (friction-less, solid) boundaries on the edges perpendicular to the shock; and outflow (free flow out of the simulation) everywhere else. A 2D plane strain geometry was used for the single chondrule and multi-chondrule mixture scenarios. The matrix was modelled as a continuous material in which the porosity was parametrised. The $\epsilon - \alpha$ porosity model was used to describe compaction in this continuous material (Collins et al. 2011; Wünnemann, Collins, and Melosh 2006) and its parameters were tuned to match existing experimental $U_s - u_p$ and $P - \rho$ Hugoniot data (Borg et al. 2006; Chapman 2010) [1] The chondrules, however, were resolved explicitly. Most materials used the Von-Mises strength model, except for the aluminium and copper (Johnson-Cook strength model) and Sipernat (Drucker-Prager; see Chapter 3 for full details). The full parameters and models employed in the simulations are summarised in Tables 5.1 and 5.2.

[1] The EoS parameters for silica are contradictory across the literature. A Mie-Grüneisen EoS requires a linear $U_s - u_p$ relationship to operate appropriately and solid silica is only linear above $u_p \approx 2\text{kms}^{-1}$. Below this the Hugoniot curves upward as u_p decreases (Marsh 1980); other approaches have even assumed a quadratic relationship in the 'linear' section (McCoy et al. 2016). Values reported in the literature often use the (approximately) linear regime to derive values of S and c_0, or alternatively derive values from a porous Hugoniot. Neither approach is ideal, and low-velocity shock response studies of silica are rare. The parameters used here are sensible, and approximately consistent with previously published work (Lyzenga, Ahrens, and Mitchell 1983; Wackerle 1962). In addition, their apparent consistency with the experimental data presented here is sufficient evidence of their efficacy.

Table 5.2: Full details of all parameters used for Sipernat. Unreferenced parameters were chosen as sensible approximations; Sipernat is an unusual substance and this approach is necessary to generate a reasonable numerical model. This technique was employed because of the numerous contradictory accounts in the literature regarding fused silica's shock response (see Figure 5.2). However, Figure 5.2 shows that the chosen parameters are in good comparative agreement with existing experimental data.

EoS Parameters	
Solid density [gcm^{-3}]	2.203 (Wackerle 1962)
Sound speed, c_0, from $U_s - u_p$ relation [ms^{-1}]	1400 (Borg et al. 2006)
Grüneisen gamma, Γ	0.036 (Wackerle 1962)
$U_s - u_p$ constant, S	1.1 (Borg et al. 2006)
Specific heat capacity [JK^{-1}kg^{-1}]	740
Strength Parameters	
Cohesion [kPa]	10
Coefficient of internal friction	0.7
Limiting strength [MPa]	100
Poisson ratio	0.3
$\boldsymbol{\epsilon - \alpha}$ **Parameters**	
Elastic threshold, ϵ_0	0
Transition distension, α_x	1.29
Exponential coefficient, κ	0.8
Sound speed ratio, χ	0.33

In one suite of simulations the multi-chondrule mixtures were modelled as mono-disperse, randomly placed, chondrules embedded within matrix. In another suite of simulations, tomographical scans of the experimental mixtures were used to initialise the geometry in the models. Each tomograph was 'sliced' three times in 60° intervals, through the tomograph's centre, and parallel to the direction of impact. This produced three bitmap cross-sections of the particle bed that could be used to create simulation geometries. These 'slice-based' simulations provided the potential for a more nuanced analysis than standard methods.

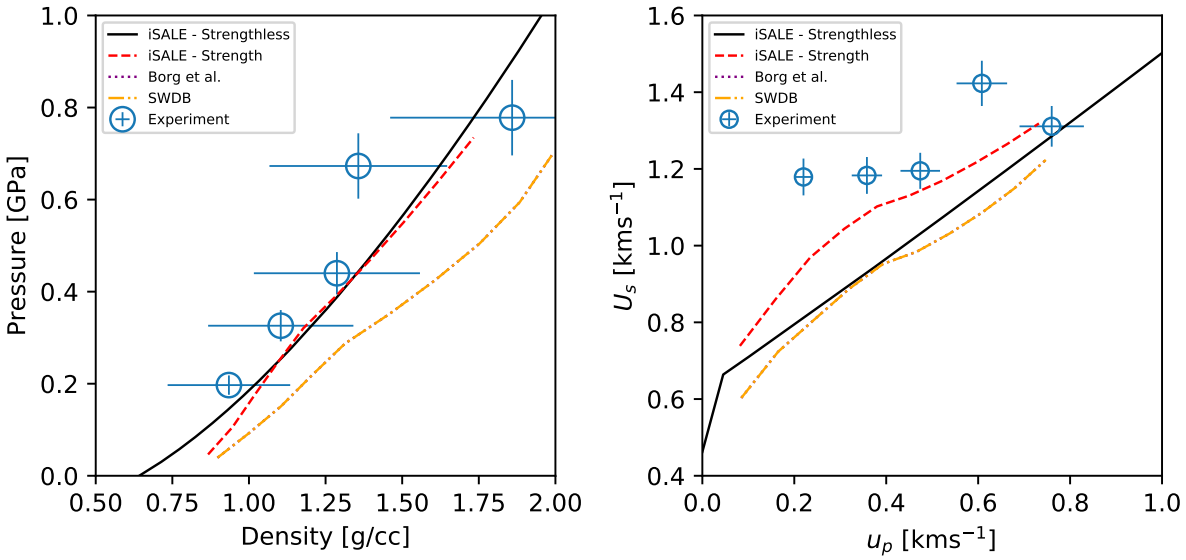


Figure 5.2: Hugoniots taken from iSALE using fitted $\epsilon - \alpha$ parameters compared with experimental Sipernat Hugoniot data (Chapman 2010). The Hugoniot using EoS parameters taken from Borg et al. (2006) is used for comparison (using the same strength model as the final parameters).

Due to a large experimental uncertainty on the time of the first frame after impact, but a low uncertainty on the inter-frame times, in the single-chondrule scenarios the best match between the first radiograph and the model was taken as a time zero for both. All times were measured relative to this 'zeroth' frame. This was not an issue in the multi-chondrule mixture impact experiments and time is measured from the moment of impact.

5.4 Results

5.4.1 Pure Matrix

The experimental radiographs display a clearly visible shock front propagating down the image (Figure 5.3). The shock front, however, is not planar and appears asymmetrically bowed in each frame. The convex-down shape unusually appears in every frame, which suggests that the shock front was bowed before entering the particle bed and that it

maintained this shape as it propagated. In addition, the shock front's asymmetry appears to be because it is rotated or 'tilted'. The first radiograph was designated T_0 and all subsequent frames were measured relative to it; this was because of a large experimental uncertainty on the impact time. The positions, however, were all measured relative to the pre-impact driver/powder interface position.

Between the first two frames (Figures 5.3 and 5.4) the experimental shock front travels approximately 113 pixels, or 1.3 mm (1 px = 11.85μm). This suggests an approximate shock speed of 1.5 kms^{-1}. The distances travelled between subsequent frames are ≈70 px/0.8 mm (frame 2 - 3) and ≈75 px/0.9 mm (frame 3 - 4), which produces inter-frame shock velocities of 1.2 kms^{-1} and 1.0 kms^{-1} respectively. A similar treatment of the driver's position produces driver velocities of 0.72 kms^{-1}, 0.44 kms^{-1}, and 0.52 kms^{-1} respectively. This result suggests that the shock and driver decelerated noticeably between frames 3 and 4.

The simulation frame that best agreed with the first radiograph's shock position was set as the simulation's T_0 to maintain consistency with the radiographs' timing convention (as shown in Figure 5.4). The simulation densities were artificially projected to give areal density for a better comparison with the radiographs. The simulation shock front does not appear to demonstrate the same level of bowing or tilt exhibited by the experiments. Instead, the shock front exhibits some bowing at late times but is largely planar at early times. The simulation was cylindrically symmetric so asymmetries like the tilt can not be replicated in this type of simulation. A direct comparison between the simulation and the radiographs is difficult because of the differing driver positions. In the first frame (in which the shock front positions were explicitly matched) the simulation's driver has propagated further than the experiment's driver. This difference in driver position persists throughout the subsequent frames. A similar, approximate, inter-frame shock and driver velocity analysis to that performed on the radiographs already shows that the

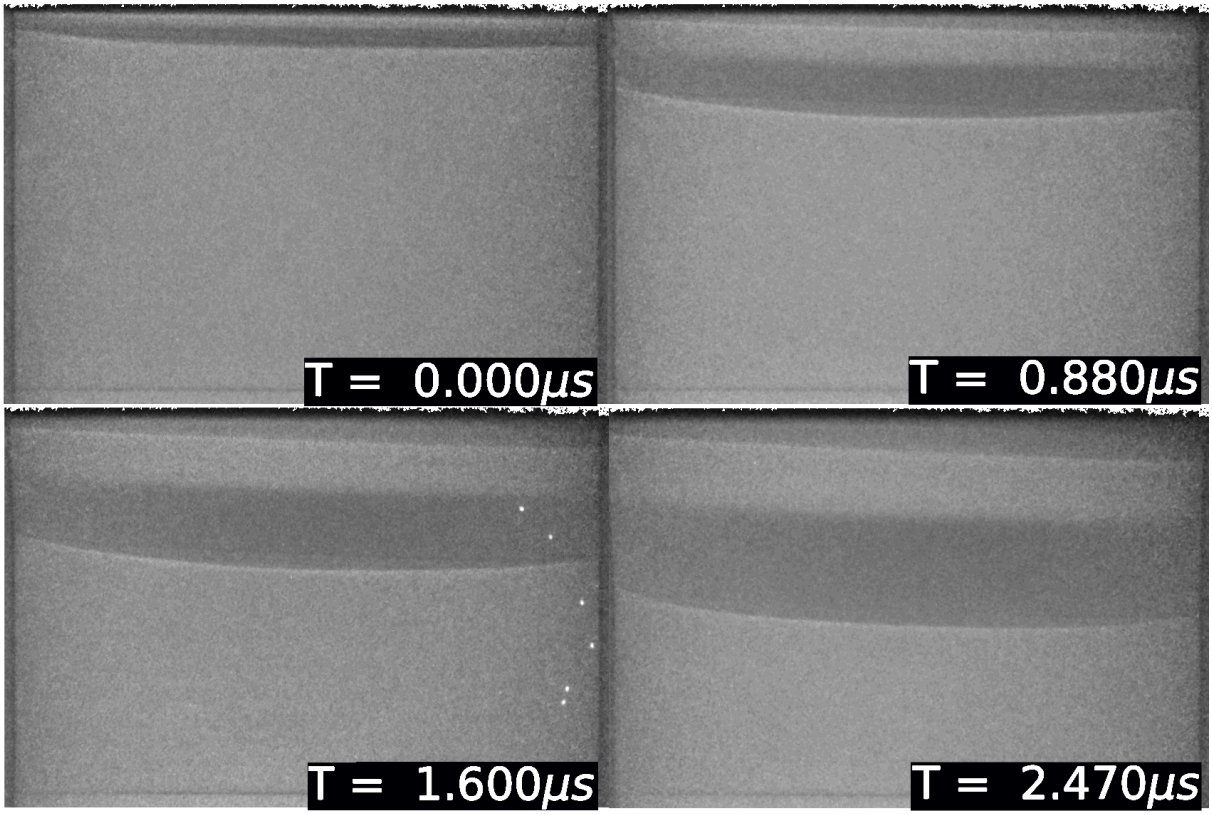


Figure 5.3: The four radiographs from Figure 5.4 in full. Non-horizontal impact on the driver in the experiments has caused a 'tilt' in the radiographs which was a source of uncertainty. All times were measured relative to the first frame because of a large experimental uncertainty on the impact time for this experiment.

simulation inter-frame velocities are 0.8 kms^{-1}, 0.7 kms^{-1}, and 0.5 kms^{-1} for the driver; 1.4 kms^{-1}, 1.3 kms^{-1}, and 1.3 kms^{-1} for the shock front. Both the driver and shock front decelerate after the first frame, but not to the same extent as exhibited by the radiographs. It is unknown how much this analysis is affected by the lack of bowing and tilt in the simulations, even if it is only approximate. Further analysis of the origins of these features is necessary.

In each radiograph the shock front appears to propagate toward the lower left corner of each image. In addition, the driver/powder interface is not perpendicular to the cell walls. This could be a result of the sabot striking the sample at an angle and causing the shock to propagate unevenly through the sample. This has produced asymmetric radiographs

and the tilt introduces ambiguity into driver and shock front position calculations This tilt was approximated to be $\approx 2.3°$ clockwise of the transverse plane.

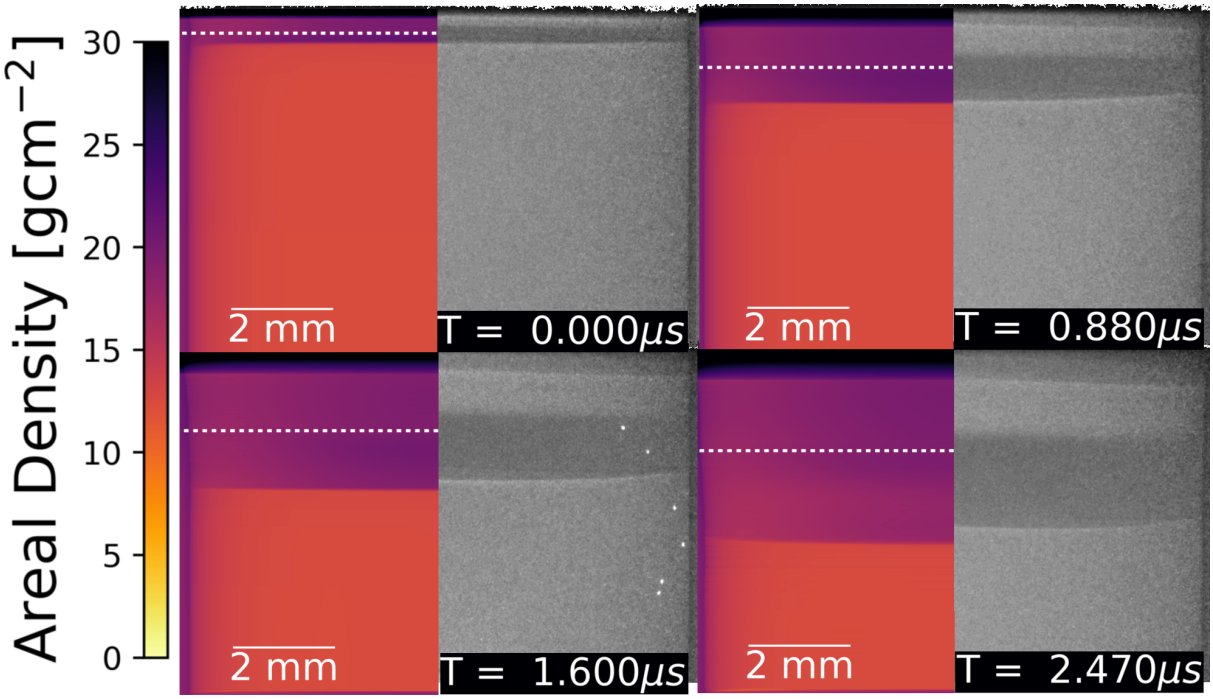


Figure 5.4: Direct radiograph/model comparison. Models are on the left and radiographs on the right; the times are calibrated against frame 1 which is 'Time 0' or T_0. Model data is projected into areal density to facilitate comparison with the radiographs. Impact direction is down the page.

The bowing of the shock front is present in the first frame and persists throughout the experiment, which suggests that the shock front was already bowed before entering the particle bed. Additional simulations were performed that replicated the general shape of this early bowing by introducing a gap between the driver 'plug' and the Al cell wall (Figure 5.5). The size was chosen to be 0.1 mm, which was approximately the uncertainty on the plug's machined diameter. The amount of bowing this introduced to the simulation shock front, however, was still insufficient to reproduce the shock front seen in experiments.

Another factor that may have affected the experiments is transverse porosity gradients within the pre-compacted particle bed (Rutherford 2017). Additional simulations that

implemented approximate transverse porosity gradients were able to recreate a reasonable approximation of the bowing of the shock at late times, however, in this case the shock was not bowed prior to entering the bed, and grew in amplitude with time. It is likely that a combination of the gap in the apparatus and porosity gradients caused the strong bowing seen in experiments. Unfortunately, the transverse porosity gradients are difficult to model in iSALE because the exact pre-impact gradients are unknown, and the additional simulations made sensible guesses.

The exact porosity gradients, set-up inaccuracies, and tilt effects within the experiments are unknown, preventing anything but a trial and error simulation approach to match the experiments exactly. In order to rule out inaccurate material models, their analytical Hugoniots were calculated from their EoSs and compared with published $U_s - u_p$ Hugoniot data (see Appendix for full comparisons). The material models were in good agreement for the field value ranges observed in these experiments, although, this method did not account for strength.

To avoid the substantial computational and temporal cost of taking tilt and bowing into account in a quantitative analysis, their effects were minimised by considering a centred region of width 750 μm ($\sim$ 100$\times$ the grain size - $\sim$ 7.5μm). This region was dubbed the 'on-axis' region. The shock position was calculated as the absolute distance travelled by the shock wave from the pre-impact driver/powder interface. The positions agree to within 10% in any frame and are in closest agreement at early times. The simulation shock front position is consistently further ahead of the radiograph shock front position, however, and the difference appears to increase with increasing time. This suggests a disagreement in shock velocity, even if the positions agree, as predicted by the earlier approximate analysis (Figure 5.6).

Unlike traditional shock experiment analysis techniques (stress gauges, or line VISAR, for example) the relative compaction of the matrix behind the shock is a feature that

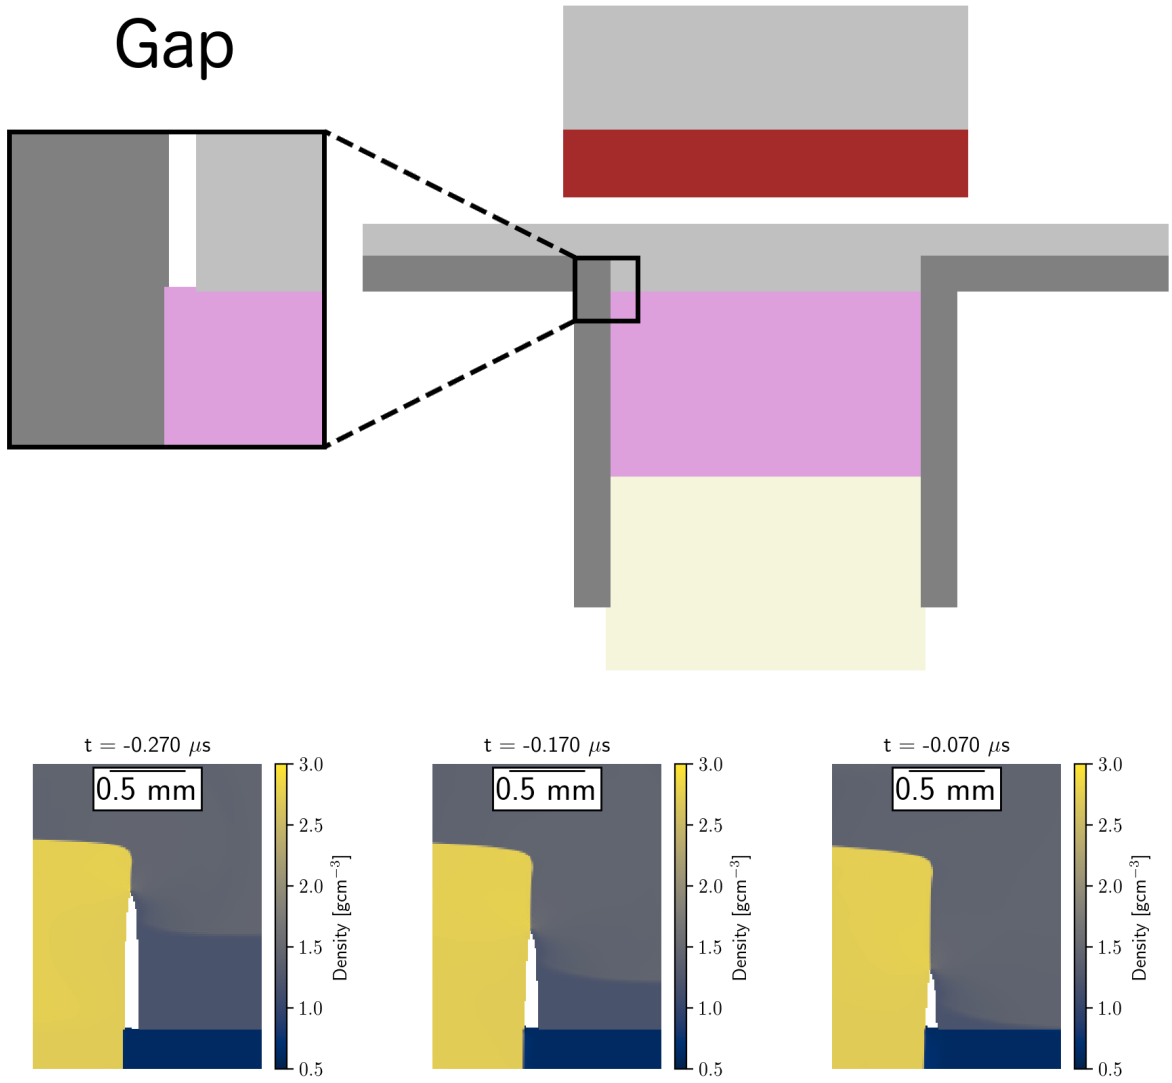


Figure 5.5: (Top) Diagram demonstrating the location of the gap between the driver and cell wall. (Bottom) Simulation results demonstrating how the shock front is bowed before entering the particle bed when the gap is included. Time is measured relative to the first frame of the experiments, which is time '0'. These frames occur before time 0 and thus have negative values..

can be examined directly using these mesoscale imaging techniques. The porosity of the matrix was calculated using an X-ray absorption model that used the radiographs as input (Rutherford 2017). Lines of porosity were extracted transversely and longitudinally across the bed. In order to compare porosities at equivalent longitudinal positions, all positions are plotted in their respective driver's frame; the driver position in each frame is the origin.

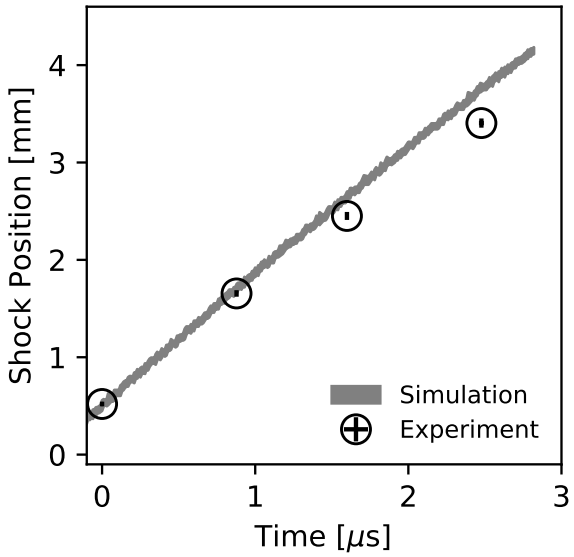


Figure 5.6: Shock position against time measured in the on-axis (750 μm) region. The 'band' of the model results is the region between the minimum and maximum positions of the shock at each time.

In the transverse case, lines of porosity ('lineouts') were considered 0.5 mm ahead of the on-axis driver position; at a fixed longitudinal position. A centred, transverse, 'working region' of 7.5 mm (ten times the size of the on-axis region) was considered for the lineouts, as opposed to the whole bed, because the absorption model used to extract porosity values was least accurate near the Al cell walls. The lineouts were also taken without correcting for the tilt. The longitudinal porosity was similarly compared at identical longitudinal positions, which is reflected in Figure 5.9.

In general there is good agreement between the experiment and simulation transverse porosity lineouts at each time (Figure 5.8). At the earliest time frame the compaction appears to be marginally greater in the models than experiments and the agreement improves as time progresses. This could be because of phase differences in the imaging caused by the proximity of an interface (Rutherford 2017) (Figure 5.7). The simulation exhibits a numerical artefact near the central boundary (Figure 5.8), and this was excluded in on-axis calculations. The simulation also shows a slight increase in porosity with

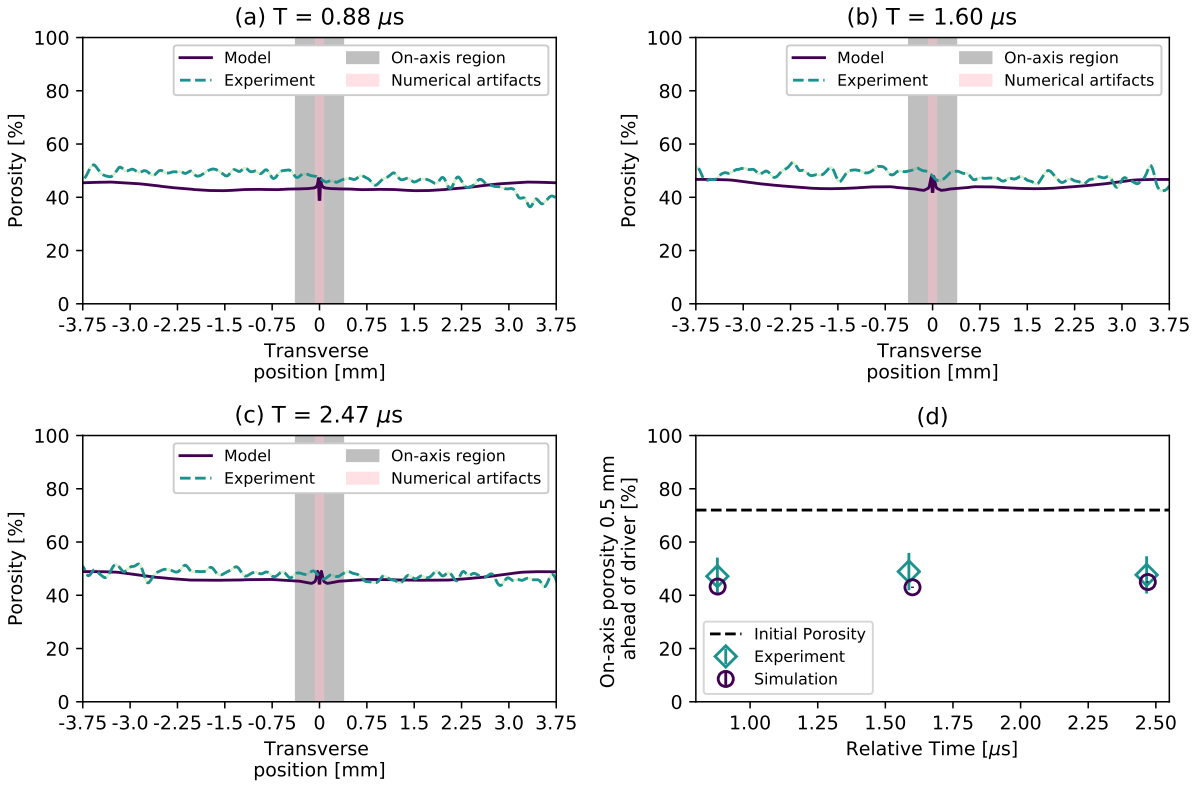


Figure 5.7

Figure 5.8: (a), (b), (c) Porosity of the post-shock powder for both experiment and simulation, 0.5 mm ahead of the driver. The shock had not advanced more than 0.5 mm at time 0, hence it is not included. The on-axis region is shaded in gray. The five cells closest to the centre of the cylindrically-symmetric simulation were not included in quantitative calculations because they are likely to be spurious. This is a common weakness of cylindrically symmetric models and the region of numerical artifacts is shaded in pink. The general gradient of increasing porosity toward the cell walls is because of lateral release in the simulations; this has been confirmed by additional plane strain simulations which exhibit the same feature. Otherwise, the asymmetric angle of the porosity in the experiment is in agreement with the tilt of the experiment, previously observed. (d) The on-axis average porosities at each time with uncertainties. Model results & experiment results are within uncertainty at all times. Experimental porosity was calculated using an X-ray absorption model which took the measured radiographic contrast as an input. (Rutherford 2017)

increasing radial distance. Additional simulations found that this is a symptom of lateral release and not a propagation of the artefact.

The longitudinal lineouts also show good agreement between experiment and simulation

(Figure 5.9) and exhibit non-uniform compaction behind the shock front. The results are compared at identical longitudinal positions (the 0 in each plot is the on-axis driver-powder interface). Both the simulation and experimental results exhibit the same general shape: a sharp drop in porosity across the shock front followed by a gradual decrease in porosity towards the driver. Very near the driver front, porosity increases slightly until it reaches the driver.

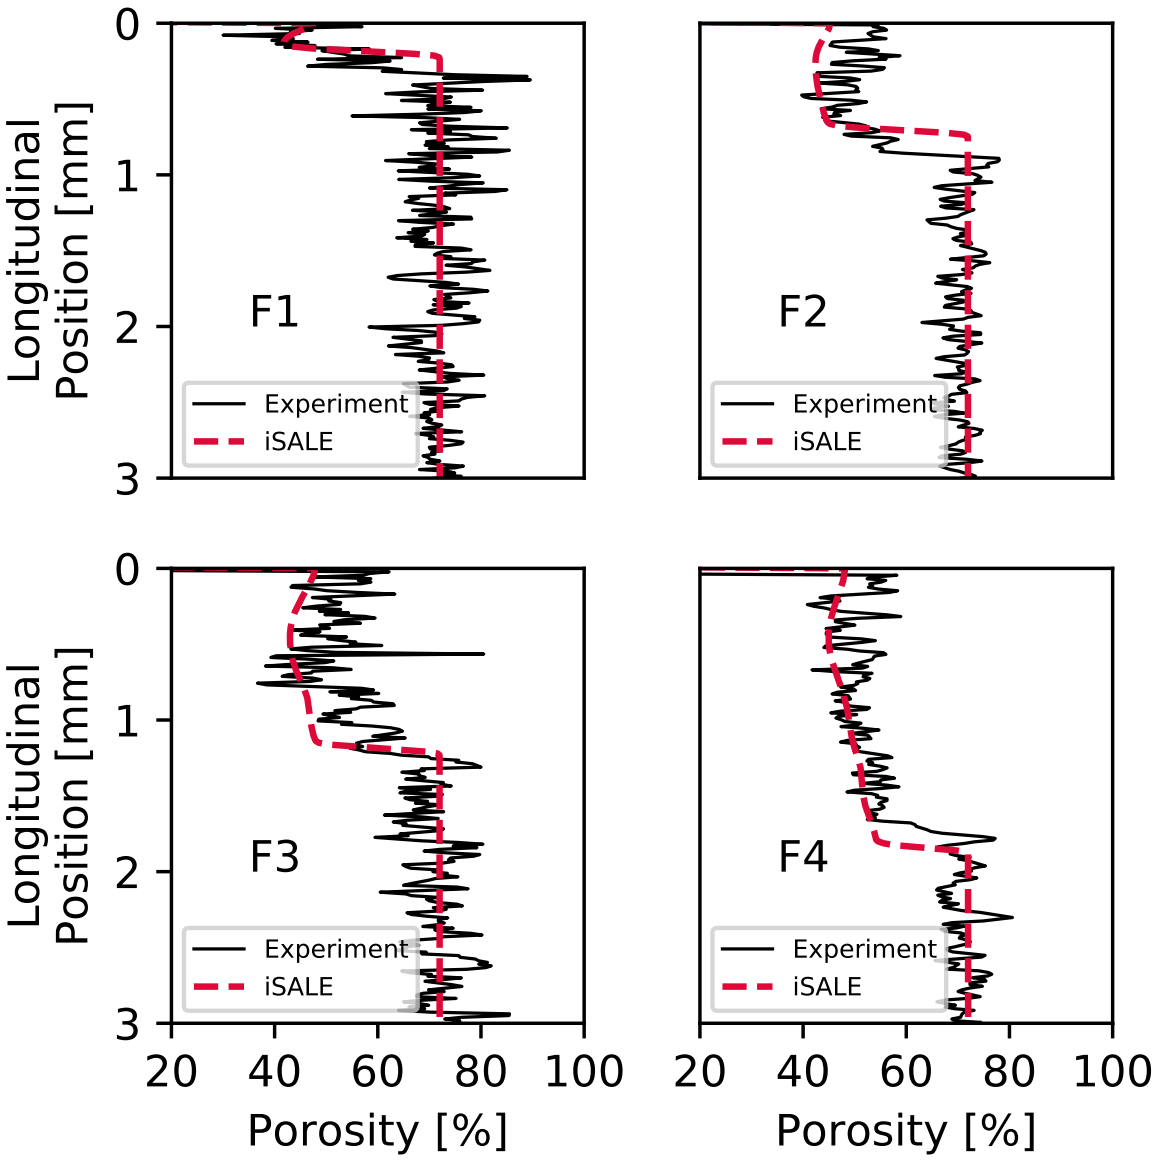


Figure 5.9: Longitudinal porosity line-outs of the on-axis porosity at each time for both experiments and simulations. Positions are relative to the driver/powder interface in each frame.

These results suggest that the simulations are a good representation of the experiments and provide support for the appropriateness of the material models used in the iSALE simulations. The largest differences between the experiment and the simulation are the bowing and tilt of the shock front, neither of which can be easily simulated. The conven-

tion of on-axis measurements, however, was devised to minimise these effects.

5.4.2 Single Chondrule

The addition of a single chondrule into the pure matrix resulted in some important differences in shock propagation within the sample bed. In particular, the chondrule's introduction affected the shock's speed and compaction in its vicinity. These effects did not extend across the whole particle bed, and were restricted to a region ($\sim$ 1.5 chondrule diameters) around the chondrule; this region will be referred to as its 'area-of-effect' or 'AoE'. Outside this region of influence, the response was very similar to the pure-matrix response. The shock wave was (approximately) planar and the matrix was compacted to $\approx$ 50% porosity. As the shock interacts with the chondrule it travels quickly through it, creating a protrusion in the shock front. This protrusion is most visible at early times and decays as the shock progresses. The protrusion appears to decay at a faster rate in the simulation than the experiment; it becomes almost invisible by the final frame in the simulation, while it is still clearly present in the final radiograph.

In a similar manner to the lineout comparisons made in Figure 5.7 transverse lineout comparisons were made between experiment and simulation. In this case, the rod served as the primary point of reference and lineouts were taken 0.25 mm 'leeward' (in the direction of shock propagation; the chondrule's lee) and 0.125 mm 'shockward' of the chondrule (analogous to 'windward'; in the direction that the shock has come from). When comparing relative compaction in the matrix in these locations we observe a distinct porosity anisotropy around the chondrule in both experiments and simulations (Figure 5.10). This effect appears to affect the matrix within $\approx$ one particle radius of the surface of the chondrule (its area of effect). There is less compaction in the lee of the chondrule compared to matrix at the same longitudinal position outside the area-of-effect. The reverse is true in the area-of-effect shockward of the chondrule; here there is greater

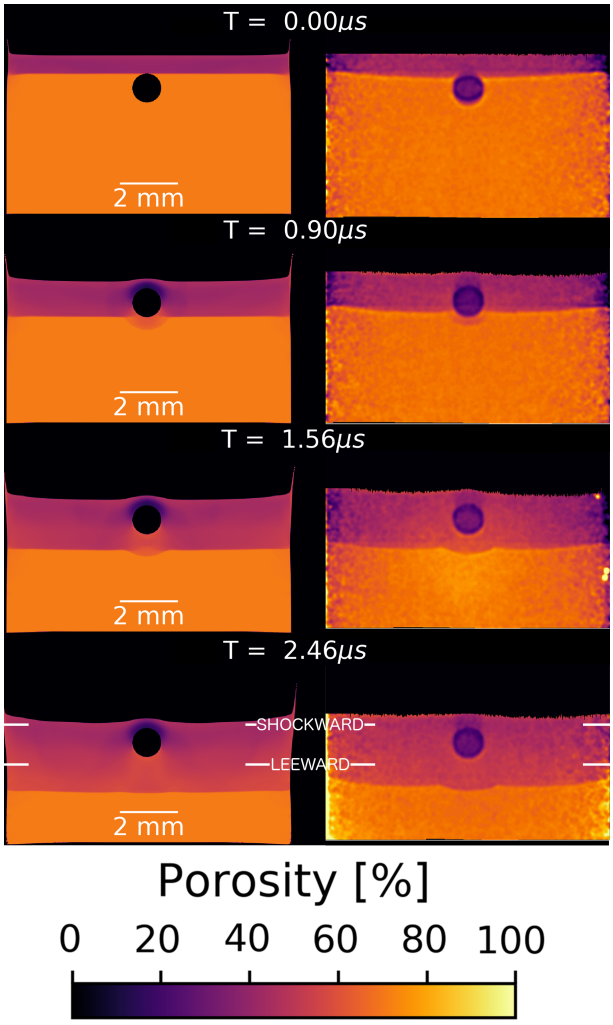


Figure 5.10: Porosity at four stages of shock propagation through samples that included a single chondrule (rod) embedded in the matrix. (Left hand column) Simulation results. (Right hand column) Experimental radiographs converted to porosity maps. Simulations were 2D and had plane-strain boundary conditions. Simulation and experiment position is calibrated by matching the position of the chondrule. Experimental porosities have been extracted using an absorption model (Rutherford 2017).

compaction shockward compared to matrix outside the AoE at the same longitudinal position. This feature occurs when the planar shock wave interacts with the chondrule. The shock wave is both reflected back into the shockward material, and transmitted through the chondrule. The reflected wave compacts the shockward material further, whilst the transmitted wave reaches the other side of the chondrule and propagates back into the matrix. The transmitted wave, however, is weaker because the shock wave lost

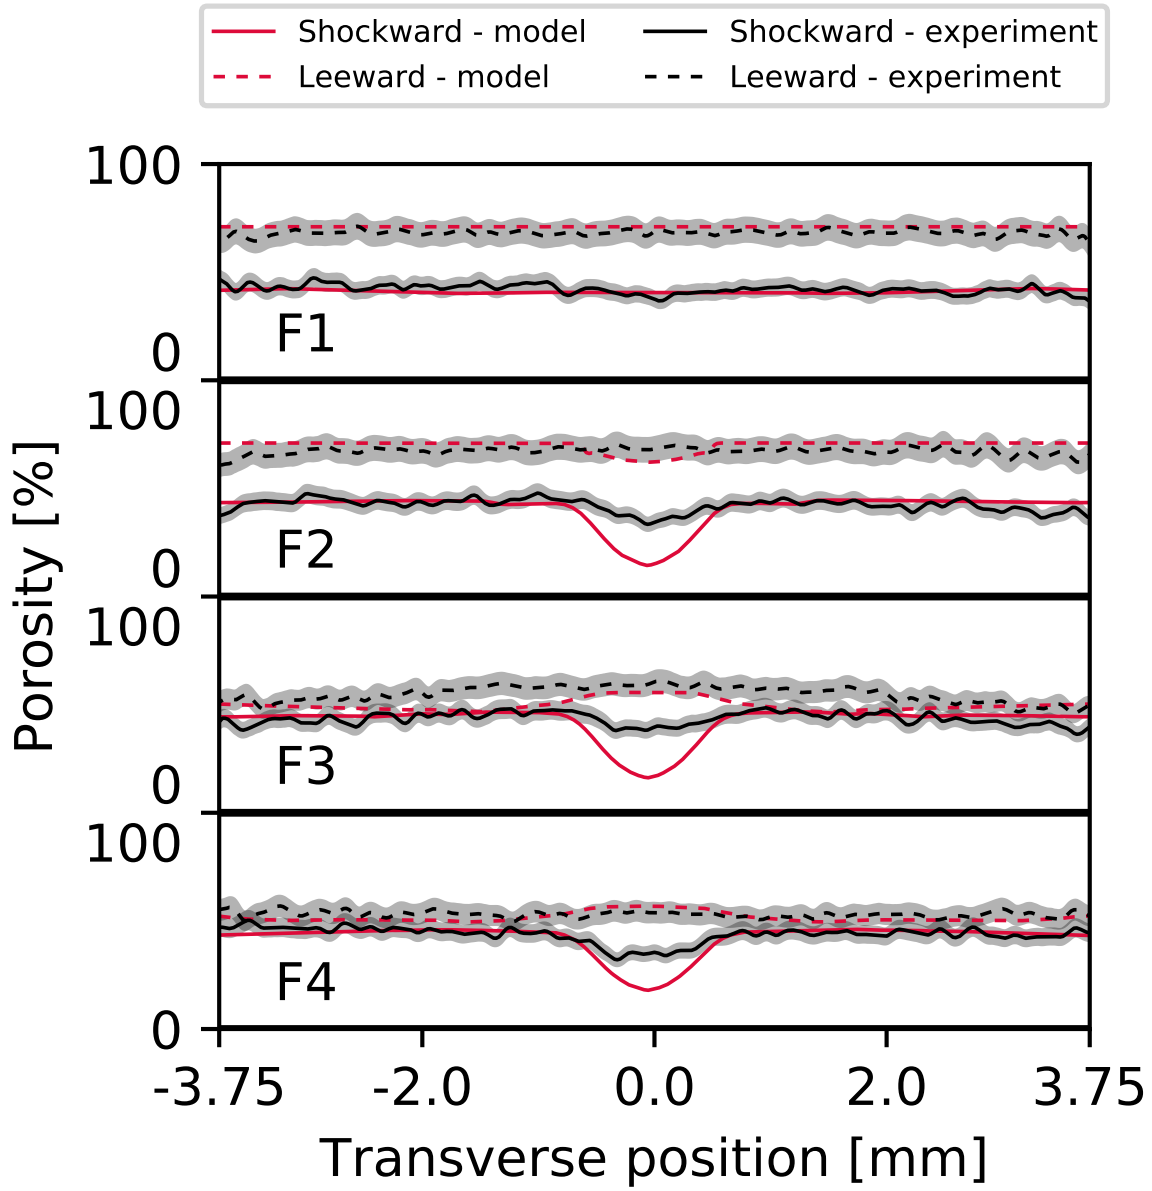


Figure 5.11: Line plots of final porosity in the experiments and simulations taken 250 μm leeward of the chondrule and 125 μm shockward of it. Only the central 7.5 mm (working region) has been considered to avoid effects of tilt or bowing on the results. Uncertainty on the experimental porosities was taken from inaccuracies in their calculation and is $\approx 7\%$) (Rutherford 2017).

energy in the transmission. This caused less compaction directly in the chondrule's lee because of the shock wave's reduced amplitude.

This feature is obvious in simulations but less clear in the experiments. This may be a

consequence of experiments being only approximately 2D: $\sim$ 1 mm of unaffected matrix at either end of the rod will contribute to the radiograph, which is a representation of integrated line density. Outside of the area of effect the porosities appear to be in good agreement, with simulation results lying well within experimental uncertainty (Figure 5.11). Experimental uncertainties come from noise in the radiographs and the simulation uncertainties principally come from uncertainty in the exact rod position (± 25 μm) owing to rod breakage and displacement during sample preparation. The model's rod position is based on the experiment's rod location, however, because the rod snapped in half during each sample preparation two rod 'ends' are visible in the radiographs and there is an uncertainty on the single position adopted in the models. This rod location disagreement is minimised after the shock passes because the matrix is reduced in volume, reducing the magnitude of variation in position. This makes it unobservable. It is possible, however, that the rod's final location in the experiments corresponded to a position further from the driver than was provided for the simulations. This would produce a region of lower shockward compaction because the shock would have had more time to dissipate before the re-shock, reducing its amplitude, and any further re-shock from the driver interface would also be reduced. In general the experiments and simulations demonstrate the same behaviour but to different extents. The spatial extent of the porosity anomaly, however, is consistent between the simulation and the experiment.

5.4.3 Multi-chondrule mixtures

The observations so far have been made in an isolated, simplified system of a single chondrule embedded in matrix. To understand the implications these results have for meteorites, an investigation into how the effects carry over to more realistic systems like those presented in Bland et al. (2014) is necessary. Multi-chondrule mixture experiments were conducted in order to extend the calibration exercise to more complex systems and explore whether the mesoscale effects observed around a single chondrule were also present in a

multi-chondrule scenario. This section is split into large and small chondrule sections for ease of presenting the results, however, there is no other difference between these set-ups than the size of the grains. There are two types of simulation within each section: particle beds that were initialised with monodisperse, randomly-positioned chondrules, and particle beds who's geometry was initialised by tomograph slices. The former are referred to as randomly placed simulations and the latter as tomograph initialised simulations. The experiments are simply referred to as large and small chondrule experiments.

The complexity of the mixtures made the radiographs difficult to interpret in their raw form (Figure 5.12). To enable a visual comparison between experimental and numerical results, radiographs from these experiments were divided by their pre-shot image (Rutherford 2017). All unchanged material appears monochrome; the driver is a region of lighter material; and the shock is visible as the wave of disturbance.

Large Chondrules

The shock propagation in the randomly-placed simulation is presented side-by-side with the equivalent radiographs for comparison (Figure 5.13). The mean shock front position appears to be consistent between the two, although, the experiment's shock front shape varies a lot between frames. In the first frame it is noticeably bowed, which is not unsurprising because the 'gap' between the driver and the Al cell wall that caused bowing in the pure-matrix experiment may still be present. In the last frame, however, the centre of the shock front develops a slight concavity and the convex-down shape has mostly disappeared. The mean experiment shock front position, however, remains consistent with the simulation shock front position. The simulation, however, does not account for this transverse variation in longitudinal position. This concavity is likely the result of a specific arrangement of chondrules. In the experiment, the chondrules were intended to be evenly spread within the matrix, but in practice this could not be guaranteed. Tomographi-

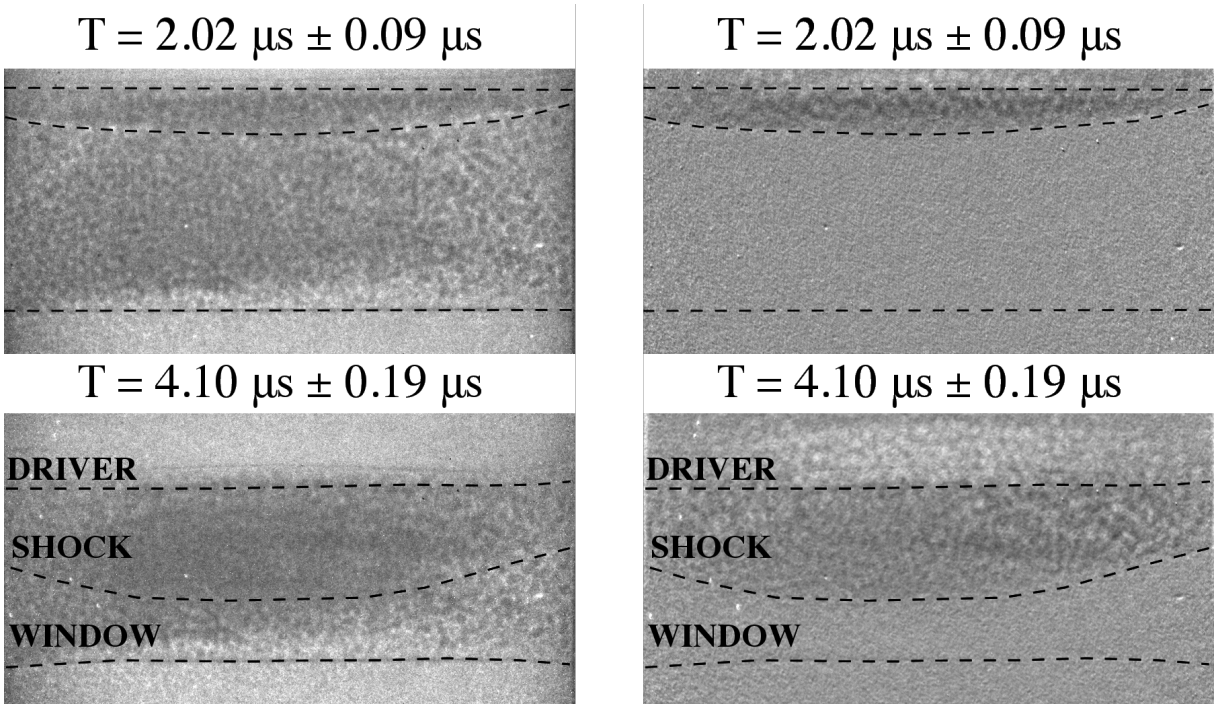


Figure 5.12: The two radiograph frames from a small-grain experiment: the shock is travelling down the image and the image width is exactly 10 mm. It is visually difficult to observe interfaces in the radiographs without the guides. To enable easy visual comparison with simulations, all other radiographs are divided by their pre-shot images and approximate visual guides are also included. (Left) The raw undivided radiographs. (Right) The same radiographs divided by their static versions; untouched material appears monochrome.

cal scans subsequently revealed strong heterogeneities in arrangement in some particle beds. Whilst these randomly-placed simulations appear to be relatively consistent with the radiographs' shock speed, they can not account for heterogeneities in the chondrule arrangement.

The tomography-initialised simulations, however, can account for some of the heterogeneous bed structure, caused by chondrule arrangement, by reproducing the arrangement in a slice of the particle bed. By performing multiple simulations, each using a different tomography slice, it is possible to examine the effects of the non-uniform chondrule arrangement. Three tomography-initialised simulations are presented in Figure 5.14, initialised from equi-spaced slices. Their mean shock front positions are similarly consistent with

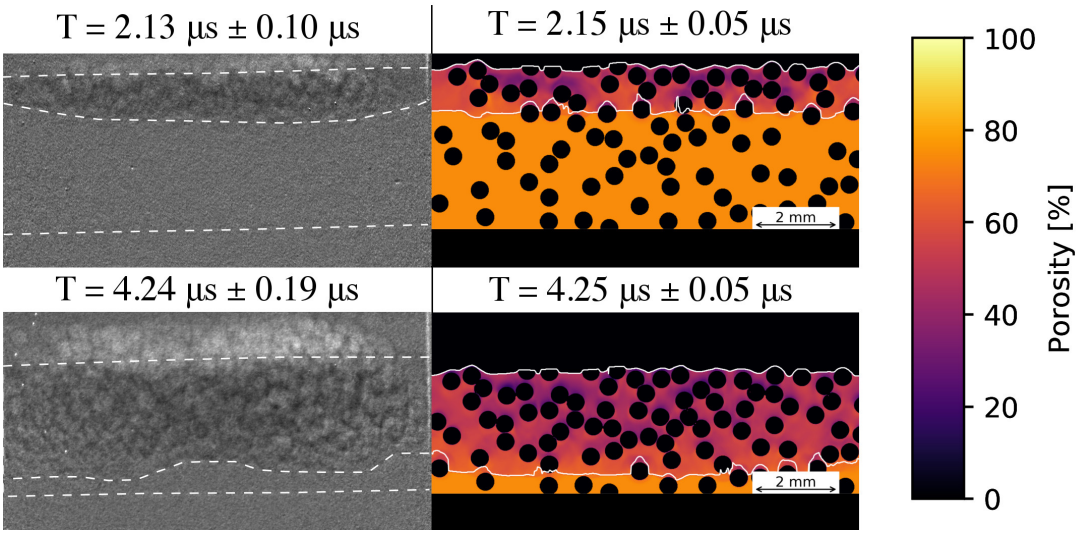


Figure 5.13: Large chondrule experimental results compared with the mono-disperse geometry using standard iSALE setup routines.

the radiographs' mean shock front positions (Figure 5.14) and additionally demonstrate some evidence of a slight concavity of the shock front observed in the slice 2-initialised simulation in the second frame. The other slice-initialised simulations do not demonstrate this effect. This is likely because the radiograph is a projection of the cylinder onto a 2D surface and for the concavity to be observed in the radiograph the chondrule distribution must exhibit the correct chondrule distribution in slices parallel to the radiograph. The slices all cross the centre of the sample radially, and may not be parallel to the radiograph. Only slice 2 demonstrates the effect which suggests this slice was the closest to parallel with the radiograph. The particles are clustered in an arch-like shape (Figure 5.14) which appears to be the primary cause of the unusual shock front geometry. At early times the shock front's centre accelerates through an initially dense layer of chondrules, forming the convex-down shape observed in frame one. In subsequent frames the central portion of the shock decelerates as it passes through a region consisting predominantly of matrix, whilst the outer portions accelerate through densely packed chondrules. This produces the slight concavity observed.

In order to compare quantitative shock front positions from the experimental and nu-

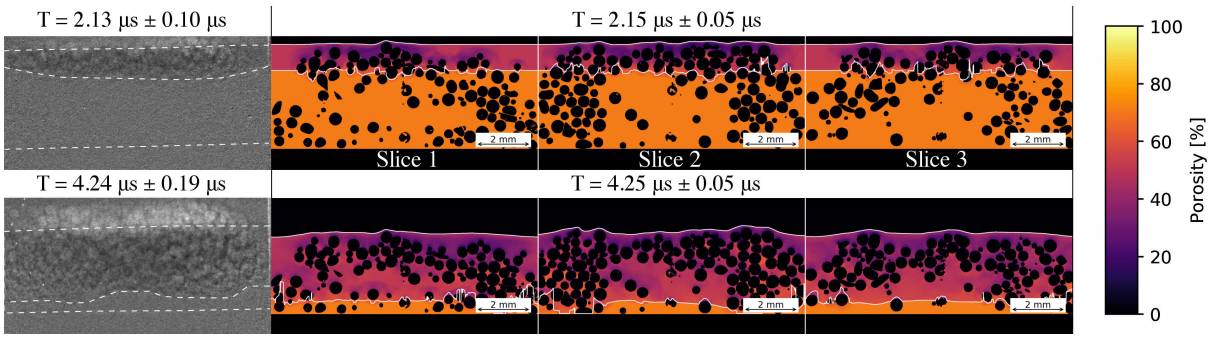


Figure 5.14: Large chondrule experimental radiograph compared with equivalent model frames from three simulations. Each simulation's geometry was initialised using a 2D 'slice' through the pre-shot particle bed. The white dashed lines are approximate and intended to be used as guides. The slices were taken in intervals of 120° and are referred to in subsequent plots by their angle: slice 1 = 0°; slice 2 = 120°; slice 3 = 240°.

merical results on-axis shock and driver positions were extracted from all experiments and simulations. The on-axis region is defined as being ten times the size of the largest grain, which is a chondrule in this case. This is a centred 4.25 mm region for the large chondrules and a centred 1.96 mm region for the small chondrules. Shock front position was measured relative to the initial driver/powder interface position. This allowed for a direct quantitative comparison of shock front positions between the experiments and both sets of simulations (randomly-placed and tomography-initialised). The shock front positions are presented in Figure 5.15. The shock positions recorded in the randomly-placed simulations, and in some of the tomography-initialised simulations, are within the uncertainty of the experimental positions. The gradients (i.e. shock speeds) also appear consistent, although, as the experimental data only consists of two points, a measure of the shock speed can not be more than an approximation. Interestingly, the simulation initialised from the second tomography slice ($\equiv$ 120°) is the one that is not consistent with the experiments, or the other simulation results. The slice 2-initialised simulation possesses a greater area fraction of chondrules than those initialised from other slices. It is probable that it is not a coincidence that this simulation exhibits atypical behaviour compared to the other simulations and it is likely the behaviour can be attributed to the

greater chondrule area fraction present. It has already been established that the chondrule distribution in the experimental beds is often non-uniform; this slice of the bed probably contains a larger area fraction of chondrules than would be there if the distribution was perfectly uniform. The slice 2-initialised simulation will experience a greater shock speed because of the increased chondrule fraction. The tomography-initialised simulations provide information about a particular experiment that would otherwise be unavailable when using randomly-placed simulations. The bulk response of the tomography-initialised simulations, however, may not agree with the experiment or the randomly-initialised simulations because the tomography slices used to initialise these simulations may posses too many, or too few, chondrules. This affects the bulk porosity of those simulations, which also affects the bulk response.

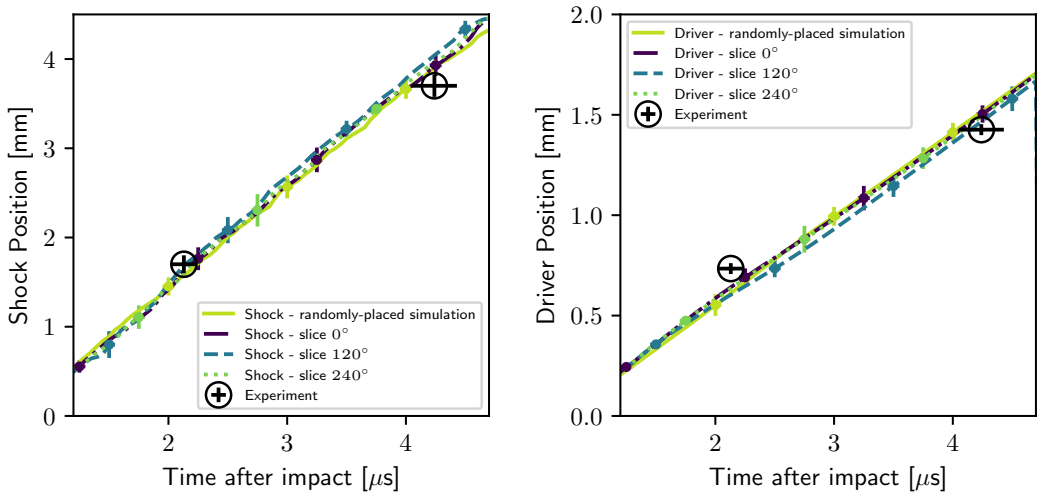


Figure 5.15: The on-axis shock and driver position for the three slices of the tomography-initialised simulations; the randomly-placed simulations; and experiments (hollow black circles). The experimental position uncertainty is $\pm 1\sigma$ on the on-axis positions and the temporal uncertainty comes from timing errors on the cameras. The simulation uncertainties are plotted on points at regular intervals; not all points are shown, however, because their number would obscure the plot.

The driver/powder interface positions agree in frame two, although there appears to be some disagreement in frame one. In general, the driver/powder interface speed appears faster in the simulations than the experiments. This suggests that the multi-chondrule

mixture in the experiments is not being compacted as much as in the simulations. This is consistent with the simulation results presented in the previous section, which show the simulation matrix consistently being more compacted compared to the experiments. In the previous section, however, the porosity observed in the simulations was still within the uncertainty of the porosity observed in the experiments. The discrepancy in shock front positions may be caused by differences between two and three dimensional studies. For example, there is out-of-plane rigidity present in 3D granular beds (Borg and Vogler 2009), which may be causing the experimental powder to resist compaction more than the 2D models.

Small Chondrules

The shock front position in the randomly-placed and tomography-initialised small chondrule simulations appears to be in agreement with the experimental shock front position in the first frame (Figure 5.16 & 5.17). Whilst the radiographs show that the experimental shock front was significantly bowed, it does not appear to be uniformly bowed; the bowing occurs near the cell wall and the central portion of the shock is approximately planar. The simulations' shock fronts appear to agree with the approximately planar section of the experiment's shock front. This is reflected in the quantitative results (Figure 5.18). This agreement occurs because the on-axis region is much smaller than in the large chondrule scenario (≈ 2 mm here) and does not reach the bowed sections in the experiment.

In the second frame, the tomography-initialised simulation shock front positions do not agree with either the experimental shock front position or the randomly-placed simulation shock front position. The tomography-initialised simulation shock fronts appear significantly further ahead than the other results (Figure 5.17 & 5.18). The randomly-placed simulations, however, do appear to be much closer to the experiments in terms of shock front position. This suggests that the shock speeds in the tomography-initialised

simulations are greater than those observed in the randomly-placed simulations and the experiments. This could be because the tomography slices were not representative of the area fraction, as discussed in the large chondrule section, and possessed a lower bulk porosity than the whole bed.

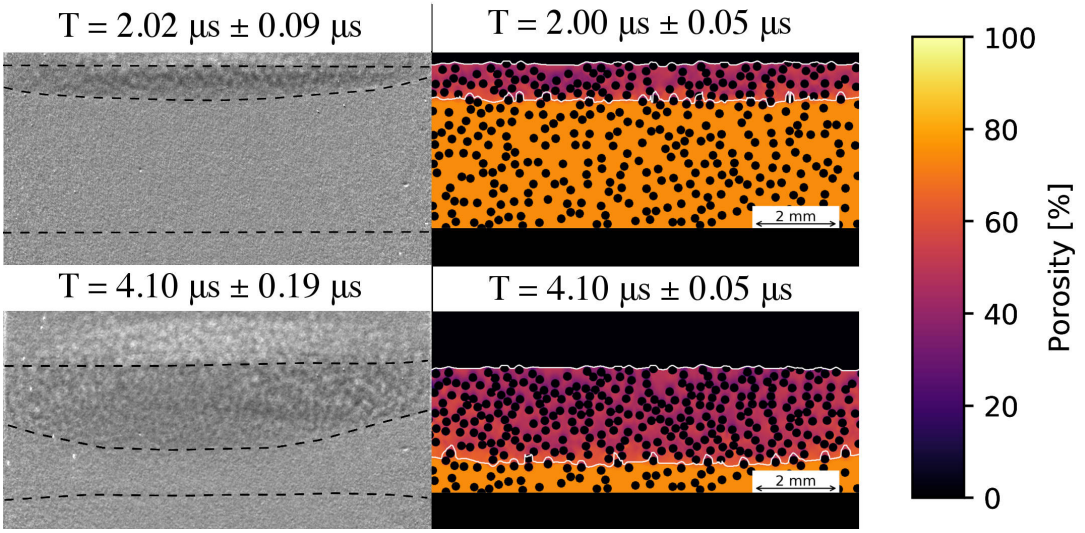


Figure 5.16: Small chondrule experimental results compared with the mono-disperse geometry using standard iSALE setup routines.

The tomography slices show that there was a dense cluster of chondrules in the centre of the sample. This is the likely cause of the significant bowing outside of this region observed in the radiographs (Figure 5.17). The slices may exhibit a greater chondrule area fraction than they would if the chondrules were distributed uniformly, however, it is possible this is not just caused by unpredictability in the experiment's set-up. The greater chondrule area-fraction could be a consequence of how the tomographs were converted to simulation input files. The tomographs had to be thresholded and reduced to just two colours (black—chondrules— and white—matrix—) to be converted into input files. If the tomographs were thresholded incorrectly, it could have produced input files with a disproportionately larger chondrule area fraction than each slice should have (they were all thresholded the same way). This would produce beds with a higher chondrule area fraction and higher shock speed. If this is the case, it would also effect the large chondrule

results, and may contribute to the reasons why those tomography-initialised simulations exhibited higher shock speeds than the experiments.

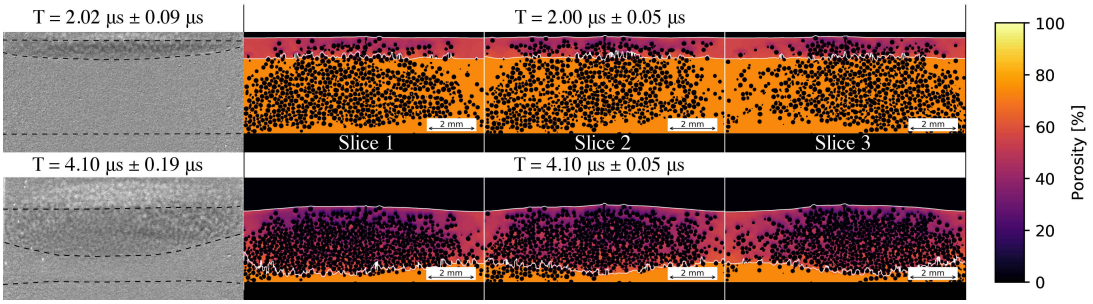


Figure 5.17: Small chondrule experimental radiographs compared with equivalent model frames from the three slice-based simulations. The white dashed lines are approximate and intended to be used as guides.

Whatever the cause, the on-axis shock and driver/powder interface positions are consistent with the hypothesis that the tomography slices are not representative of the whole bed and include a greater fraction of chondrules than matrix (Figure 5.18). The randomly-placed simulation, that was representative of the bulk mixture, is in much closer agreement with experiment.

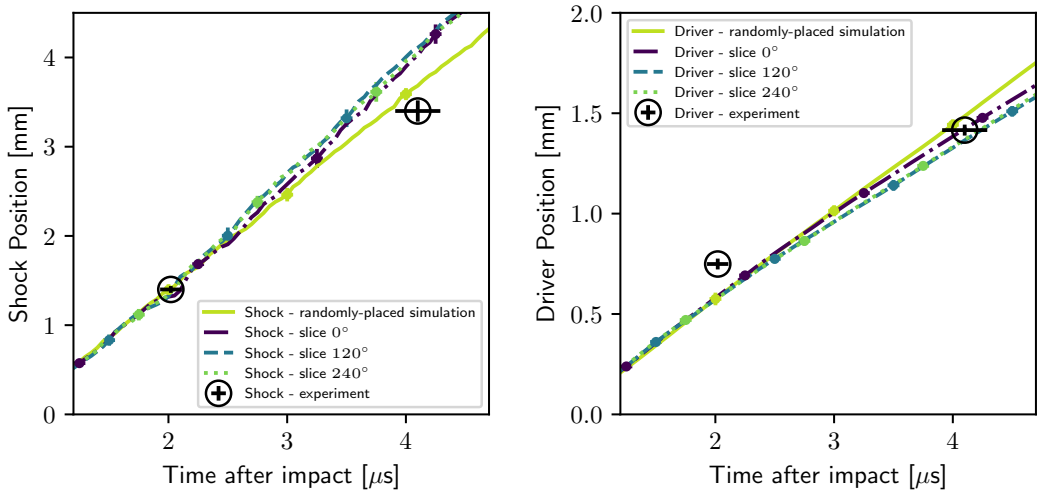


Figure 5.18: The on-axis shock and driver position for each slice; the normal set-up; and experiments (hollow black diamonds). The experimental position uncertainty is $\pm 1\sigma$ on the on-axis positions and the temporal uncertainty is as already stated.

The driver position results are the reverse of the shock position results: all the simulations

are consistent with the experiments in the second frame and disagree in the first frame. This result may appear as a contradiction to the results presented in Figure 5.18. This suggests that the shock entered the bed earlier in the experiments than it did in the simulations, and the faster shock speeds in the simulations allowed the simulated shocks to catch up to (and overtake) the experimental shock before they reached the end of the bed. Meanwhile, the simulations' driver speed was sufficient to catch up to, but not overtake, the experimental driver. This suggests that there is a timing uncertainty on when the shock entered the particle bed, which could also have a systematic effect on the results in the second frame. An additional source of inaccuracy could have been introduced when measuring the depth of the particle bed; it may have expanded after measurement but before the shot, raising initial porosity and bed-volume.

5.5 Discussion

The results presented in this chapter represent one of the first investigations into powder compaction of its kind, utilising X-ray radiography to explicitly resolve the dynamic compaction of the powder bed. Whilst some experiments have been performed that consider similar set-ups, such as Beitz et al. (2013) sample-recovery investigation into analog chondritic material compaction, there are no dynamic *in-situ* studies of this type. Beitz et al. (2013) investigated the impact velocity required to achieve final porosities that were consistent with known meteorites; the study did not consider mesoscale features. The primary investigations into the mesoscale features of bimodal powder compaction have been numerical. Amongst those, most have considered special cases of bimodality, such as Williamson (1990) who introduced bimodality to an ordered system of grains in order to vary the initial void fraction of the system. The work performed by Bland et al. (2014), Davison, Collins, and Bland (2016), Davison et al. (2017), and Rutherford et al. (2017) represents the main body of research that has considered this particular system. The

work presented in this chapter represents the logical extension to that by considering the system's shock response experimentally as well as numerically. The porosity dichotomy in particular has been proposed in existing studies (Davison, Collins, and Bland 2016; Davison et al. 2017) and this work presents explicit numerical and experimental evidence of its existence in an isolated system.

5.5.1 The Porosity Vector, C

In the single chondrule investigation experimental and numerical evidence was found for the existence of the porosity dichotomy predicted by Davison, Collins, and Bland (2016). The chondrule appeared to 'protect' the matrix in its lee resulting in a higher porosity on that side with a lower porosity on the opposite side. This result, however, does not indicate if the feature still exists in multi-chondrule systems. It is much more challenging to study compaction in these. In order to investigate the presence of this porosity dichotomy around individual chondrules, the 'porosity vector' ($\boldsymbol{C}$) was devised as a metric to quantify the magnitude, and direction, of this anisotropic compaction.

$\boldsymbol{C}$ is a polar vector (which shares its centre with the chondrule's centre) that points along the mean porosity gradient across a chondrule (Figure 5.19). $\boldsymbol{C}$'s angle is measured with respect to the direction of impact/intended shock direction: a value of 0° indicates that the porosity vector points in the same direction as the shock's propagation, whereas a value of ±180° indicates that the porosity vector points to where the shock has been. Positive and negative angles indicate directions left, and right, of the diagrams respectively and this choice is arbitrary. The porosity vector is calculated using the matrix within its area-of-effect—specifically an annulus that extends from the chondrule's rim by $\frac{1}{4}$ the chondrule diameter (250μm). This annulus of matrix is split into 20 'wedges' (annulus sectors) with equal arc lengths. The average porosity in each wedge was found. The relative difference in porosity on opposite sides of the grain was found by taking the difference in

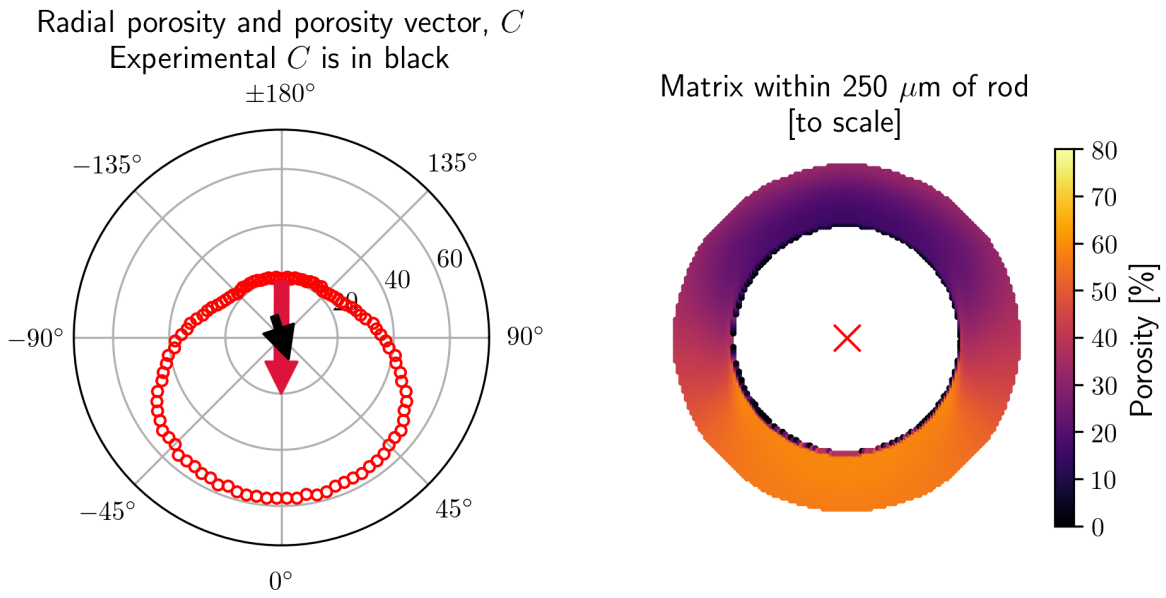


Figure 5.19: (Left) Experimental and numerical porosity vectors superimposed on the numerical radial porosity distribution around the rod shown in Figure 5.10. (Right) The porosity of the matrix surrounding the rod, which was used to calculate the porosity vector $\boldsymbol{C}$. An annulus half the radius in width was used.

porosity between opposite wedges. This produced 10 equi-spaced polar vectors that pass through the chondrule's centre. The weighted average of these 10 vectors (magnitude and angle separately) is the resulting porosity vector. In the single chondrule scenario the porosity vector is parallel to the shock direction ($\theta \approx 0°$) and has a magnitude of $\approx 20\%$ (Figure 5.19). This means that on average there is a positive porosity gradient across the chondrule demonstrating a difference of 20% porosity for an angle of $\theta \approx 0°$. The experimental $\boldsymbol{C}$ is mostly in directional agreement with the simulations. It points slightly to the right of 'downward', although this is approximately consistent with the angle of tilt, which suggests that the apparatus tilt causing the shock front to propagate at an

angle to its intended direction also affected the resulting porosity vector. The magnitude of the experiment's $\boldsymbol{C}$ is significantly less than the model's, but this is as expected, given the differences in porosity shockward and leeward of the grain that have already been discussed.

This porosity anisotropy is present in the isolated case, but it is not immediately clear if it is also present in more realistic scenarios, such as the multi-chondrule mixtures. The presence of the porosity anisotropy in multi-chondrule mixtures was tested by calculating values of $\boldsymbol{C}$ from both randomly-placed simulations and using an area-of-effect of $1.5\times$ the chondrule diameter. Not all chondrules were considered when calculating $\boldsymbol{C}$. Chondrules that had come into contact with each other were ignored because they had become 'welded' together, which hindered the algorithm used to calculate $\boldsymbol{C}$. In addition, chondrules that were partially within the domain were ignored because they did not have a clear centre, or annulus of surrounding matter with which to calculate $\boldsymbol{C}$.

The calculated porosity vectors were combined into two histograms (Figures 5.20 and 5.21) in order to compare their angles and distribution. They show that the majority of chondrules for which $\boldsymbol{C}$ was calculated in each bed, possessed a downward pointing vector. The porosity vector directions appear to be distributed about 0° (the direction of shock propagation). Although, in the small chondrule case, there is a slight skew to one side. Interestingly, the magnitudes of $\boldsymbol{C}$ are approximately half that of the idealised case; they lie in the 5-15% magnitude range compared to the 20% observed in the single chondrule case. They are more consistent with the experimental value of $\boldsymbol{C}$ for the single chondrule case, despite comparable impact velocities $\sim 600\text{ms}^{-1}$ and less lateral release. The shock is likely passing over chondrules and reflecting off those that are further ahead, similar to the re-shock observed shockward of the rod in Figure 5.11, producing lower than expected porosities. Not every porosity vector points in the direction of shock propagation, however. These divergent porosity vectors are likely caused by proximity to

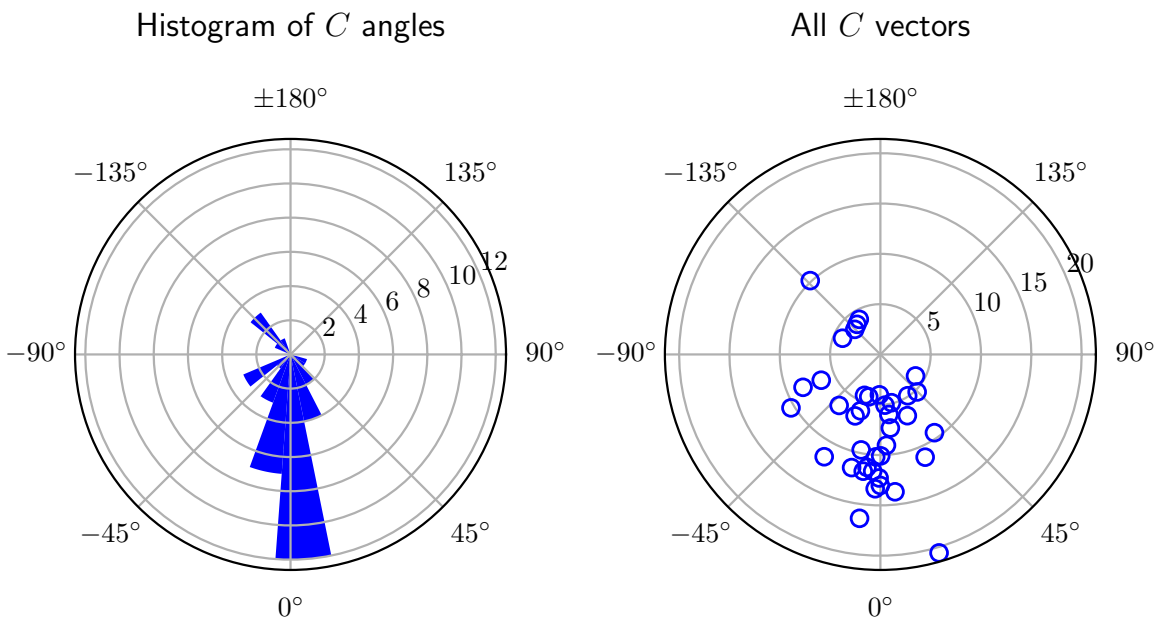


Figure 5.20: (Both) Details of all porosity vectors extracted from eligible large-grained simulations. (Left) Histogram of all porosity vector directions; frequency is radial. (Right) All porosity vectors as a polar scatter plot; porosity [%] is radial.

other chondrules. A closely-packed group of chondrules will compact the matrix to a much greater degree because of the many interfaces and motions of chondrules. For example, two chondrules positioned shockward and leeward of each other will significantly compact the matrix between them as the shock compresses them together. For the shockward grain this may be sufficient to give it a porosity vector that points shockward, whereas, the leeward grain would possess an unnaturally strong porosity vector leeward. The more grains there are, and the more densely packed those grains are, the more numerous this feature will be.

This analysis suggests that the porosity vector is a real consequence of bimodal powder compaction and that it commonly points in the direction of shock propagation when evaluated for a chondrule. This porosity anisotropy appears to resemble features seen

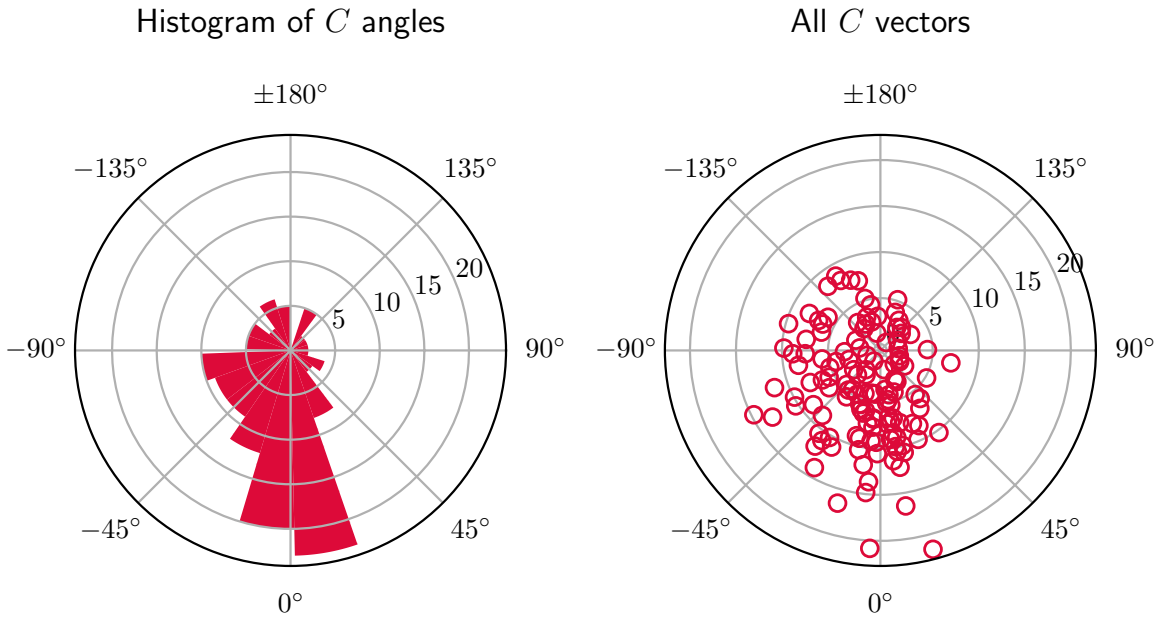


Figure 5.21: (Both) Details of all porosity vectors extracted from eligible small-grained simulations. (Left) Histogram of all porosity vector directions. (Right) All porosity vectors as a polar scatter plot.

around the chondrules and rock fragments in some meteorites. Meteorites such as NWA 5000 have large grains that appear to have 'shadows': these are dark and light sections that appear as if the grains have sheltered the smaller rocks behind them (light region). These features are consistent in their direction throughout the meteorite. It is possible that these 'shadows' are compaction anisotropies from a shock process in the past. If the 'shadows' and the porosity anomaly are related symptoms of a shock process, then this could present a new way to decode the direction and magnitude by which a meteorite was shocked. The meteorite, Leoville, is another example where the porosity vector could be of use (Nakamura, Tomeoka, and Takeda 1992). One study on Leoville suggested that deformation, and preferred orientation, of its chondrules could be linked to a preferred compaction direction (Nakamura, Tomeoka, and Takeda 1992). This study supposed a

preferred compaction direction may be related to pre-compaction porosity, and regular mild shock impacts. The work presented in this chapter, however, potentially presents a method of verifying if there was a preferred orientation. To test this hypothesis, future work needs to perform a more thorough analysis of the features, as well as shock recovery experiments to investigate final compaction distributions.

5.5.2 Implications for Meteorites

The existence of this porosity vector has significant implications for the shock categorisation of meteorites. It potentially provides a method by which the magnitude and direction of the shock that a meteorite has been exposed to can be calculated. Whilst porosity maps of meteorites do not currently exist, there has been significant work examining similar features. Forman et al. (2017), for example, considered the Allende meteorite and found that crystallographic preferred orientations of matrix surrounding the chondrules tended to point in a preferred direction. In addition, their calculations suggested that the temperatures and pressures experienced by the meteorite were consistent with the scales discussed in Bland et al. (2014). They concluded that the preferred orientations were a direct cause of shock propagation, similar to the conclusions drawn here. The similarity between that study and the porosity dichotomy discussed here is evident. The results of this chapter could potentially be used to categorise the shock stage of a meteorite using an analysis such as that presented in Forman et al. (2017). An agreement in the crystallographic preferred orientation, with porosity vector directions, in a meteorite would be solid evidence that both are symptoms of the same phenomenon. The magnitude of the porosity vector could additionally be used to assess the strength of the shock the meteorite was exposed to.

5.5.3 Implications for Shock Physics

The tomography-initialised method of modelling a particle bed is an example of considering the exact mesostructure of a granular mixture in an experimental setting, as discussed in Chapter 4. By investigating the 3D arrangement of grains within the particle bed, the reasons for particular geometric shock front features can be inferred. In Figure 5.14 for example, the chondrule density was predominantly at the sides of the cell and the slice 2-initialised simulation showed that at that point in the bed there were significantly more chondrules than in the other slice-initialised simulations. This suggested that the particle bed did not have axial symmetry for the chondrules, which could explain the existence of slight concavity in the shock front at late times. Whilst there are studies that use scans of the particle bed to initialise simulations, such as Eakins and Thadhani (2008), they are rarely used to explain quirks of the experimental response.

The set-up used in this research also has parallels with similar systems. For example, the compaction of composite granular materials (granular matter embedded within a continuous material) is a common problem in shock physics (Herbold et al. 2008). The matrix is not porous, but is still continuous. These materials are often created as explosives that do not detonate except under shock loading. Ensuring that the material releases energy at the optimum rate is the primary motivator. The mesostructure of the grains within the continuous material can be manipulated to optimise the rate of energy release with the explosive. The results of the mesoscale investigations presented here, the porosity vector in particular, has potential to be of use to the shock chemistry field. The results are in the optimum position to begin analysis of experimentally inaccessible fields like pressure and temperature. The effect of the porosity vector on the pressures and temperatures which a chondrule experiences has not been considered before and could have implications for explosive material structuring.

5.6 Conclusions

In conclusion, this chapter presents two joint numerical-experimental investigations into the shock compression processes involved in precursor meteorite material, through the use of mesoscale experiments and simulations. Experiments were conducted by driving shocks through analogous materials (Sipernat–fused silica powder– for the matrix and glass beads/rods for the chondrules) and recording real-time, *in-situ*, X-ray radiographs of the experiments at the ESRF (Rutherford et al. 2017). The first investigation considered the response of a single chondrule-matrix interaction and the second considered multi-chondrule analog mixtures.

In the first investigation, experiments and simulations were performed that investigated the compaction of a pure fine-grained, highly porous granular material analogous to the precursor of meteoritic matrix material (Sipernat 320-DS). Further simulations and experiments were performed using the same material, except now with a single non-porous chondrule of the same composition (borosilicate glass rod) embedded within it. Shock position and the relative compaction of the matrix were in agreement between the experiments and models. The compaction behind the shock front was in particularly good agreement between the experiments and simulations.

The central .75 mm of the bed was considered in quantitative analysis ($\sim 100\times$ grain size $\approx 7\mu$m in Sipernat, called the 'on-axis' region). In this region, the shock was deemed to be approximately planar in both the experiment and the simulation. Outside of this region the experiments were strongly affected by the tilt and shock front bowing.

In the single chondrule scenario the experiment and model were in good agreement in terms of shock position and compaction. The final porosity around the chondrule was strongly anisotropic in both; the chondrule appeared to 'protect' the matrix in its lee resulting in a higher porosity on that side with a lower porosity on the opposite side.

The magnitude of the contrast between the regions was greater in simulations than experiments, however the area-of-effect was comparable. This porosity distribution has similarities with features observed in known meteorites, such as NWA 5000. If these meteorites are formed by the same process, they potentially present a new way of decoding the shock magnitude and direction that meteorites were exposed to.

In the second scenario, fully analogous multi-chondrule mixtures were considered, consisting of 30 vol.% chondrules and 70 vol.% matrix, with a bulk porosity of 50%, which is within the typical range of values for precursor meteorite material (Davison, Collins, and Bland 2016). The projection of 3D set-ups onto 2D radiographs prevented the experiments from exporting meaningful spatial measures of density during compaction and comparison was predominantly in terms of shock and driver position. It was not possible to infer 3D spatial density distributions from the 2D radiographs in the multi-component mixtures, and comparisons were restricted to shock front and driver positions. Tomographical scans of the particle beds before each shot, however, allowed for numerical models to use bitmap 'slices' through the tomographs to initialise simulations. This provided an unprecedented insight into specific mesostructural aspects of compaction caused by irregularities in bead arrangement introduced in sample creation. Agreement between models and experiments was generally good, however, tomography-initialised models were not always representative of the bed because they were only a slice through the sample and subject to non-uniform packing effects, such as clumping. These slices did not represent the bulk response well, but provided key information about shock propagation in each experiment that would otherwise have been unavailable.

This work additionally presents some of the first results comparisons of mesoscale numerical models with explicitly mesoscale experiments. The results presented here demonstrate the efficacy of this method for interrogating the shock processes involved in a distinctly mesoscale phenomenon. This investigation leaves the work in a good position to begin ex-

tracting experimentally inaccessible fields from the simulation results, such as temperature and pressure. In addition, the porosity vector presents a potential new method of decoding the magnitude and direction by which a meteorite was shocked in the past. Future work would be able to investigate how the porosity vector translates to geological observations of real meteorites and make verifiable predictions about what should be observed. In particular an investigation that compares the porosity vector directly against crystallographic orientations of the matrix in a meteorite could provide the evidence needed to examine the efficacy of using this technique in meteorite categorisation. In addition, this work would benefit from extending the investigation to 3D, and considering how the porosity vector is affected by the transition.

Chapter 6

Aspects of simulating the shock response of a complex powder at the mesoscale: the lunar regolith

6.1 Introduction

It is common when modelling powders in the shock physics literature to approximate their geometries as 'ideal'. For example, a typical simulation will assume mono-disperse grain sizes, spherical grains, and an approximately uniform spatial particle distribution (Derrick et al. (2018b), for example). Previous studies, such as Borg and Vogler (2009), have found this assumption is mostly valid when considering the bulk response of the powders. In addition, it is not uncommon that the powder is engineered to possess the idealised properties, as is common in manufactured powders (Herbold et al. 2008), which removes the need for an approximation at all. Naturally-occurring powders, however, are rarely as simple. Materials such as the lunar regolith, terrestrial soil, construction gravel, etc., often possess particle size distributions that cover multiple orders of magnitude,

involve complex granular morphologies, and exhibit a non-uniform porosity distribution within the mixture. These powders are of great interest to the shock physics community because of their ubiquity (as discussed in Chapter 2) and recent studies have presented evidence that the mesoscale shock response of complex powders may be strongly influenced by their complicated micro- and mesocale properties (Bland et al. 2014; Davison, Collins, and Bland 2016; Derrick et al. 2018a; Rutherford et al. 2017). The complexity of these powders is directly attributable to their complex mesoscale responses, some of which were discussed in the previous chapter, and the results challenge the efficacy of using an idealised approximation when studying them.

The computational cost of modelling these complex powders explicitly, however, is typically large. For example, to extract meaningful results from a mesoscale iSALE simulation the smallest grain must be larger than the recommended minimum grain size (8 cppd in this thesis, see Chapter 3). If the physical domain size is large compared to the smallest grain within it ($\sim$ 25 mm to $\sim$ 7μm, as in the previous chapter, for example), then the minimum number of cells required to model that situation explicitly becomes prohibitively large.

In order to examine the mesoscale response of the precursor meteorite matter a semi-explicit resolution method was introduced, which avoided making an idealised powder approximation whilst simultaneously reducing the computational cost that would otherwise have been required. A semi-explicit resolution method resolves some of the grains in the granular material explicitly, whilst parametrising the remaining material as a porous continuum (Figure 6.1). This allows the simulation to implicitly resolve grains that it would otherwise be incapable of resolving explicitly. This semi-explicit resolution method was first introduced in Bland et al. (2014) and has been used in a handful of studies since, investigating the same heterogeneous powder (Bland et al. 2014; Davison, Collins, and Bland 2016; Davison et al. 2017; Derrick et al. 2018a; Rutherford et al. 2017) The semi-

explicit resolution method has significant potential to reduce the cost of modelling other complex powders, whilst simultaneously avoiding the flaws of using an idealised approximation that may misrepresent the powder's mesoscale response. Its accuracy, however, when simulating other powders with more complicated properties than the precursor meteorite material, has not been tested.

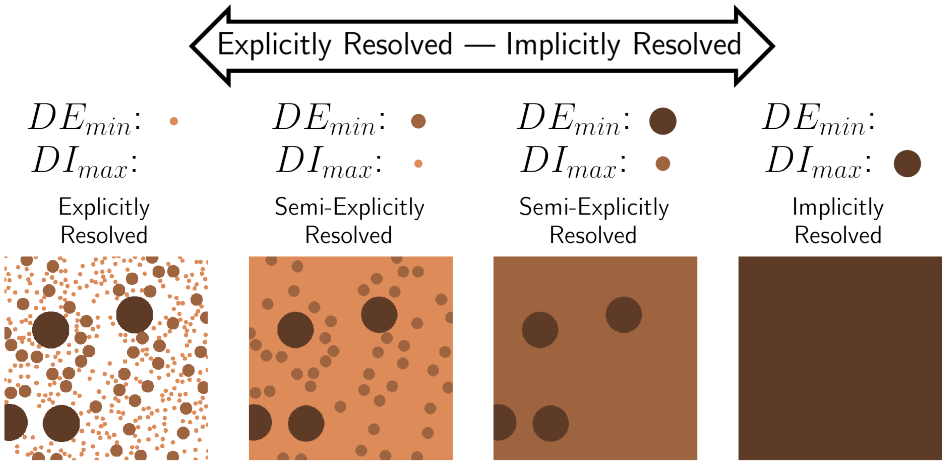


Figure 6.1: When resolving a granular material, either all the grains are resolved as individual objects (explicitly), or the grains are resolved as a continuum that uses a porosity model to behave appropriately (implicitly). The semi-explicit resolution approach resolves the large grains explicitly and parametrises the small grains.

When resolving a powder semi-explicitly, there will be a minimum explicitly resolved grain size (DE_{min}). I.e. in a simulated powder all grains resolved as individual objects will be greater than, or equal to, DE_{min} in size and any grains smaller than DE_{min} will be resolved implicitly. This parameter is chosen prior to the simulation. In the implicitly resolved material there will also be a maximum grain size for the granular matter that is resolved implicitly (DI_{max}). I.e. in the simulated powder all grains with sizes smaller than, or equal to, DI_{max} are parametrised in the porous, continuous, matrix. In general, the semi-explicit approach will be valid when $DE_{min} >> DI_{max}$, which implies that the smaller grains can be approximated as a continuum in comparison to the large grains.

This was the case in the previous chapter with the precursor meteorite matter. In a powder with a continuous particle size distribution, however, this condition is probably not true and it is more probable that $DE_{min} \approx DI_{max}$. It is not known how this affects a simulated powder's response when resolving the powder semi-explicitly. If resolving grains semi-explicitly produces meaningful results when being used to model a powder with a continuous particle size distribution, varying DE_{min} should have no effect on the powder's bulk response.

In Bland et al. (2014), Davison, Collins, and Bland (2016), and Davison et al. (2017) a large temperature dichotomy was observed between the matrix and the chondrules when resolving the granular material semi-explicitly. The matrix temperatures were orders of magnitude greater than the temperatures experienced by the chondrules. It is not clear if this result applies only when considering a powder where $DE_{min} >> DI_{max}$ is true, or if this is a general property of modelling a granular material using a semi-explicit resolution technique. In order to compare grains of the same size across simulations, they are divided into 'grain size categories' (GSCs) which are compared against each other. For example all grains in a simulation with diameters greater than 6 cppd, but less than 10 cppd could be considered a GSC. Using a small set of GSCs enables a simple comparison between the bulk responses of two equivalent GSCs in simulations with different values of DE_{min}, and avoids comparing individual grains which may not experience similar responses because of their location in the sample (this was not fixed). When varying DE_{min} a GSC will be explicitly resolved in one simulation but be included in the implicitly resolved grains in another. If the semi-explicit resolution method is able to produce meaningful results, it should also produce consistent responses in each GSC, independent of DE_{min}.

When parametrising material (resolving it implicitly) there is an underlying assumption that the spatial distribution of the grains being resolved implicitly is approximately uniform (Figure 6.2). This assumption, however, may not be true, and there may be instances

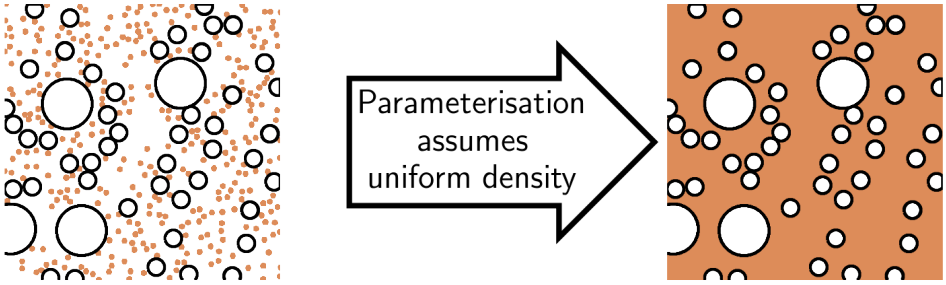


Figure 6.2: Matrix parametrisation assumption made in Figure 6.1 for the smallest grain size. When resolving grains implicitly, there is an underlying assumption that the implicitly resolved grains are uniformly distributed in space. The porosity model does not take clumping of grains into account. In some cases this is a reasonable assumption (Chapter 5), but when $DE_{min} \approx DI_{max}$ this may not be the case.

of 'clumping' (as discussed in Chapter 4). In situations where $DE_{min} >> DI_{max}$ the uniform spatial distribution assumption will be reasonable because the size difference means that the implicitly resolved grains can be reasonably approximated as a continuum. When $DE_{min} \approx DI_{max}$, however, it is more likely that the spatial distribution is non-uniform and large pockets of explicit void exist (Figure 6.1, 6.2). In order to assess the merits of resolving a granular material semi-explicitly, the validity of the assumption of uniformity must be tested.

Simulating powder beds that have been resolved semi-explicitly has not been considered in the past, and in particular, the effects of different grain morphologies (other than circular) have not been considered. If the semi-explicit resolution method can be used to

model complex granular materials, the bulk response of the granular material should be insensitive to morphology, as has been reported in previous studies considering explicitly resolved granular matter (Benson and Nellis 1994; Borg and Vogler 2009).

This chapter provides an exploration of the semi-explicit resolution method's ability to model other complex powders accurately. First, the effect of changing DE_{min} was investigated. Four simulations were performed that varied a particle size distribution (PSD) from fully explicitly resolved grains to fully-implicitly resolved grains, and the responses of four GSCs were compared between simulations. The GSC boundaries were chosen to coincide with the values of DE_{min} such that DE_{min} did not land in the middle of a GSC in any simulation. The effect of non-uniformity in the parametrised powder was examined by introducing varying amounts of explicit pores to a base simulation, and comparing the results with the results of the original. The matrix porosity was also kept constant between simulation (despite the addition of pores) by adding the equivalent number of fines from the PSD to balance out the explicit voids. Finally, the effects of altering the morphology of the explicit grains on the bulk and mesoscale shock response was examined by performing three simulations with fixed matrix porosity, grain positions, grain sizes, and orientations.

In order to consider these aspects the lunar regolith was chosen as a case study. This material is ideal for this purpose because of its extreme properties. In particular, it possesses a grain size distribution that crosses > 4 orders of magnitude, has extremely angular grain morphology and a high natural porosity ($\sim 50\%$). It is of great scientific importance to develop understanding of its shock compaction because it is a primary source of information about the Moon, and the Moon suffers continuous impact bombardment.

6.1.1 The Lunar Regolith

The Moon is the closest planetary body to the Earth and a scientifically important record of planetary formation. In particular it can provide details of impact mechanics, and the Earth's own history, through examination of its geology and make-up (Heiken, Vaniman, and French 1991). Unlike the Earth, the Moon's lack of an atmosphere means that impact craters do not erode away as they would on earth—the primary weathering system on the Moon is through additional impacts. Constant asteroid bombardment over millions of years, with no crust renewal (as would happen to terrestrial rock), means that nearly the entire surface of the Moon is covered by a thin layer (typically a few metres deep) of granular rock which is known as the 'lunar regolith' (Di et al. 2016; Heiken, Vaniman, and French 1991). It consists largely of basaltic rock which takes on several forms including breccia (rock containing pieces of other, older rock) and agglutinates (partially melted, and solidified, breccia). Additionally, recent studies have presented evidence that the Moon's crust is significantly less dense than originally thought (Han et al. 2014; Wieczorek et al. 2013), which could suggest that the regolith extends much deeper. The bulk porosity of the regolith has been estimated to vary from $\sim 44\% - 54\%$ (increasing with depth) although precise measurements are difficult because of the lack of *in-situ* samples (Heiken, Vaniman, and French 1991).

The particle size distribution of the regolith covers several orders of magnitude and can be approximated as a Gaussian distribution when considered as a function of Krumbein phi [1]. The particle size distribution of the regolith used here was taken from the Camelot crater rim (Heiken, McKay, and Brown (1974), sample 7508-1, see Figure 6.3).

The regolith samples taken on the Apollo missions, however, are the only lunar samples on Earth that have been taken directly from the Moon. This means they are some of

[1] Krumbein phi (ϕ_k) is a measure of grain size that is a common parameter in PSD research (Blott and Pye 2012). To avoid confusion with porosity (also 'ϕ') it is referred to as 'Krumbein phi/ϕ_k'

the most analysed materials on the planet, and the the regolith's grain morphologies have been extensively studied (Heiken, Vaniman, and French 1991). From this analysis, the ranges of various standard morphological parameters have been extracted, such as roundness, elongation, volume coefficient and others.

Three morphology parameters were considered in this chapter: elongation, roundness, and area ratio. Elongation (E) is the inverse of the aspect ratio, and regolith grains typically exhibit elongations in the range $1.32 - 1.38$, with an average of 1.35 (Heiken, Vaniman, and French 1991). They are referred to as 'somewhat elongated' (Heiken, Vaniman, and French 1991). The roundness of a grain is defined as $R = \frac{a}{q}$, where a = average radius of curvature of the grain's silhouette's circumference; and q = radius of largest circle that fits entirely within the grain's silhouette. Grains in the regolith have been found to possess average roundness in the range $0.19 - 0.26$ and are said to be 'angular to sub-angular' (Heiken, Vaniman, and French 1991). The area ratio is the area of the grain's silhouette, A, divided by the area of the smallest cuboid that encloses the grain's silhouette. The area ratio was found in lunar studies to be $\sim 0.7 - 0.74$ (Heiken, Vaniman, and French 1991).

6.2 Methods

6.2.1 Set-up & Materials

Regolith compaction simulations were set up in 2D plane strain geometry with two targets (1.25 mm $\times$ 3 mm), end-to-end, possessing opposing velocities (Figure 6.4). This produced a symmetric impact in which the two targets collide. Each target was given an initial velocity magnitude that was exactly half that of the intended impact velocity. Each target was divided longitudinally into three 'sections' of size 1.25 mm $\times$ 1 mm; the sections nearest, and furthest, from the collision site were designated 'buffer' regions (Figure 6.4).

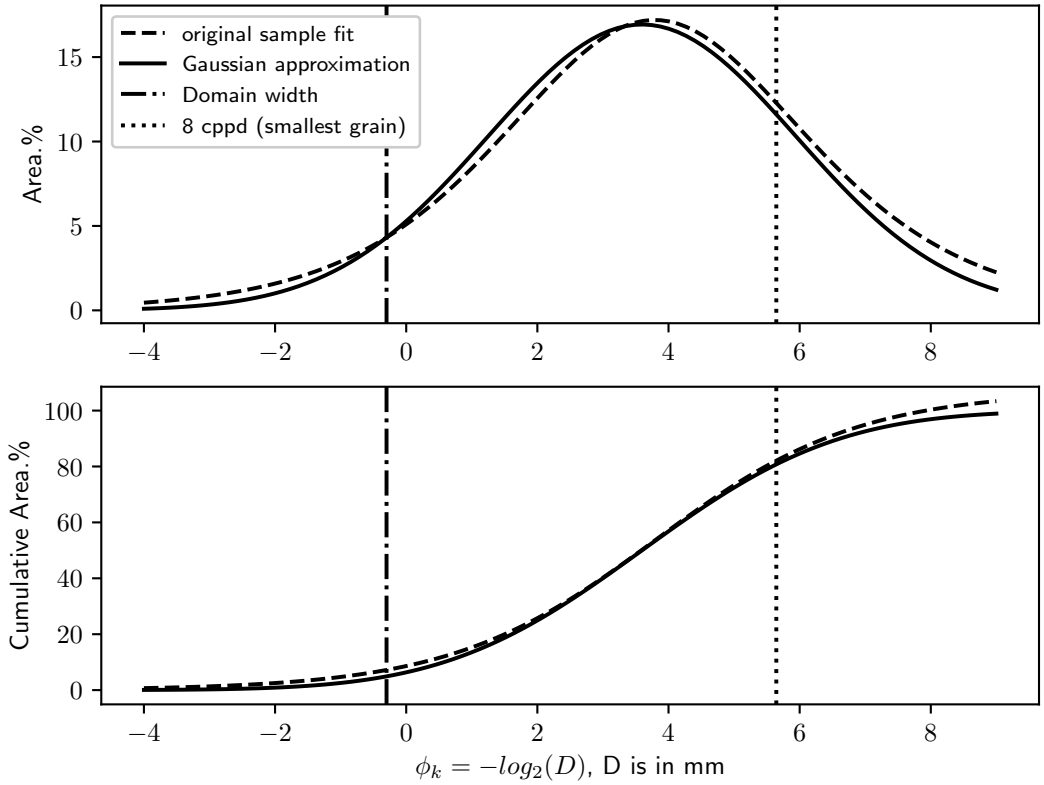


Figure 6.3: Approximate particle size distribution of a lunar sample taken from Heiken, McKay, and Brown (1974) and the additional approximation to a Gaussian distribution. This distribution was used as the base of all subsequent simulations.

The central sections of each target (surrounded by buffer zones) were the sample regions and all analysis was restricted to these sections. The buffer in front of the sample (closest to collision site) allowed the shock to stabilise before entering the sample region whilst the buffer behind it (furthest from the collision site) prevented the shock from releasing immediately off the rear face and into the sample. The targets only differed in terms of grain arrangement, meaning, each simulation was able to examine two samples at the cost of one.

The lunar regolith contains many different compounds and composites and it is not feasible to account for them all (Heiken, Vaniman, and French 1991). This investigation

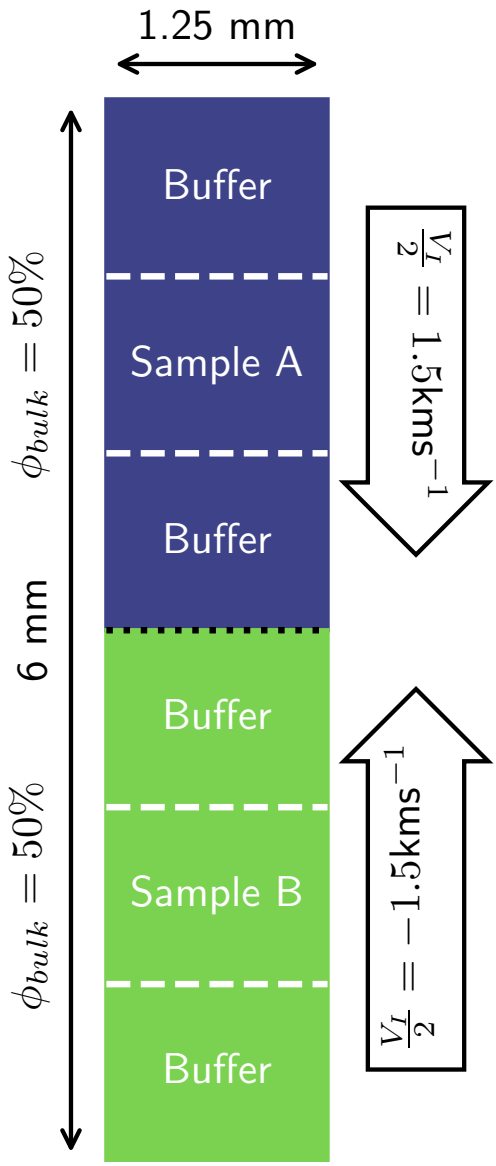


Figure 6.4: Initial geometry of regolith compaction simulations; simulations used a 'mirrored' geometry with buffer regions either side of each 1.25mm × 1mm sample.

approximated the regolith as a material that has been used to model the Moon's planetary crust in the past—basalt (Ivanov, Melosh, and Pierazzo 2010). The material parameters from Ivanov, Melosh, and Pierazzo (2010) were used to construct the material model for basalt, with some minor strength modifications to reflect the different material

scales. The simulation parameters from which this model is derived were originally used to model the lunar crust, however, the strength of this crust does not represent a homogeneous material. Rather, the crust would be full of discontinuities, fractures, and other weaknesses. This will be represented in the crust's intact strength. In a granular version of the same material the difference in scale is sufficient that most of these weaknesses won't exist. This is because the typical length scale of individual grains is much smaller than the typical length scales of the weaknesses. The intact strength of the grains will necessarily be much greater than the intact strength of the crust because of the absence of these size-dependent weaknesses.

Similarly, the intact strength of the parametrised material must represent the material as a whole, and will not be equal to the intact strength of a single grain. It should represent the whole ensemble. Thus, it will have a much weaker intact strength than the crust, because the material is not continuous and thus possesses many more 'weaknesses' than the crust, even if they are not the same ones. The intact strength of the lunar crust used in Ivanov, Melosh, and Pierazzo 2010 was 20 MPa, so the intact strength of the explicitly resolved grains was set to ten times this at 200 MPa, whilst the intact strength of the matrix of fine grained particles was set at 0.2 MPa, as a first approximation. Previous studies have found that the strength of soft rocks can decrease by as much as ten times as their size is increased (Brace 1981). In this case, however, the size range is many orders of magnitude greater (km to mm) than the range considered in Brace (1981) which was approximately a single order of magnitude (cm to m). The greater range will enhance the effect and suggests that an increase in strength of one order of magnitude for the grains is reasonable. The Rock strength model (see Chapter 3) and the $\epsilon - \alpha$ model were used to parametrise porosity in the implicitly resolved matrix material. The basalt (Ivanov, Melosh, and Pierazzo 2010) ANEOS (Thompson and Lauson 1974) EoS was used for the material; see Table 6.1 for full details.

Three impact velocities of 0.75 kms^{-1}, 1.5 kms^{-1}, and 3 kms^{-1} were considered (except where otherwise stated) and several bulk parameters were extracted from each sample region to examine how consistent the bulk response was between set-ups. Unlike previous chapters, the targets are in motion prior to impact and the buffer plates consist of the same granular material as the sample. Unfortunately, the standard stress-contour method can not easily be applied in this situation because of these features. To avoid this problem this chapter uses Lagrangian tracer particles (tracers) as an alternative. At the start of each simulation a tracer is placed in every non-void cell and these tracers are advected through the mesh with the piece of material they are assigned to, recording the various fields it is exposed to at each time. Tracer particles have constant mass and do not affect the simulation results. Unlike the normal stress-contour method, tracers can track the maximum temperatures and pressures they experience throughout the simulation. This provides previously unavailable information for analysis. Each tracer tracks the maximum temperature and pressure experienced by the piece of material it represents as well as the porosity. The bulk parameters were found for both the entirety of each sample and each GSC within them by averaging the respective fields across all relevant tracers.

6.2.2 Resolving the Particle Size Distribution Semi-Explicitly

The number of particles it is possible to resolve in a single simulation is a function of domain size and resolution. Consequently, there is a maximum area fraction that can be achieved for a given PSD. The largest grain that can be resolved can not exceed the size of domain, however, the smallest grain that can be resolved must have at least 8 cppd (see resolution test in Chapter 3). Reducing the size of a computational cell will reduce the physical size of the largest grain that can be resolved (for constant mesh size) and increasing the mesh size will drastically increase the computational cost of a model. The domain size was as shown in Figure 6.4, with a cell size of 2.5 μm $\times$ 2.5 μm, which produced a mesh size of 500 $\times$ 2400 cells. This mesh was size chosen as a compromise

Table 6.1: All material, & model, parameters used to model Basalt.

EoS	ANEOS (Thompson and Lauson 1974)
Strength Model (see Chapter 3)	Rock
Poisson ratio	0.25
Cohesion [kPa]	10
μ_i	1.4
μ_d	0.6
Y_{dm} [GPa]	2.5 (grains); 0.25 (matrix)
Y_{im} [GPa]	2.5 (grains); 0.25 (matrix)
Yield Strength [MPa]	200 (grains); 0.2 (matrix)
Damage Model (see Chapter 3)	Ivanov
ϵ_{fb}	10^{-4}
B	10^{-11}
p_c [GPa]	0.3
Thermal Softening Model (see Chapter 3)	Ohnaka (Ohnaka 1995)
Melt Temperature [K]	1473
Thermal softening constant (ξ)	0.7
a (Simon-Glatzel approximation)	4.5×10^9
c (Simon-Glatzel approximation)	3
Porosity Model (see Chapter 3)	$\epsilon - \alpha$
α_x	1
κ	0.98
χ	1
ϵ_e	0

between computational cost and the number of grains that could be resolved. It covered the region of the PSD with the largest area fraction of grains. The largest particle allowed was 1.25 mm in diameter ($\phi_k \approx -0.3$) and the smallest was 20 μm ($\phi_k \approx 5.6$). Estimates of the regolith's bulk porosity lie in the range $44 - 54\%$ (Heiken, Vaniman, and French 1991) so a bulk porosity of 50% was used for all simulations.

Varying DE_{min}

To provide evidence of the semi-explicit method's efficacy when modelling complex powders, two criteria must be met. First, all simulations should produce the same bulk response independent of DE_{min}. Second, the bulk response of grains in a GSC within the sample should be insensitive to DE_{min}. I.e. grains of a particular size should exhibit responses consistent with grains of the same size in the other simulations. The equivalence of samples with different DE_{min} was examined by comparing the mean peak shock temperature and pressure of each simulation (in bulk and in terms of the GSCs), as well as the final porosity of each simulation.

To examine the efficacy of resolving grains semi-explicitly, it was necessary to compare the semi-explicit resolution method against a fully-explicitly resolved simulation. Unfortunately, this was not possible using the domain and mesh sizes chosen for modelling the lunar regolith because of the regolith's large PSD width. To address this problem a second PSD was created for the implicit/explicit investigation with one third the standard deviation of the main PSD, but the same bulk porosity 50% (Figure 6.5). The implicit/explicit cut-off was defined by varying the smallest grain size that was resolved, and parametrising all grains below this size. Four scenarios were considered: a fully explicit simulation where the smallest grain size was 6 cppd[2]; two semi-explicit simulations with cut-offs of 20, and 40 cppd each; and a fully implicit simulation that parametrised the whole PSD (Figure

[2]The number of grains smaller than this size was non-zero but very small and a slightly smaller grain size (than the recommended minimum) was chosen such that the remaining material would require a parametrised matrix porosity as close to 100% as possible. 6 cppd required a matrix porosity of $\sim 97\%$.

6.6). In each case a different matrix porosity was required to maintain the same bulk porosity, $\phi = 50\%$. The matrix porosities were: 100% (no matrix) for 6cppd, 85% for 20 cppd, 65% for 40 cppd, and 50% for the fully implicit simulation. The implicit/explicit simulations used impact velocities of 3 kms^{-1} because this should produce the largest values for each bulk field and the most variation will be observed between simulations.

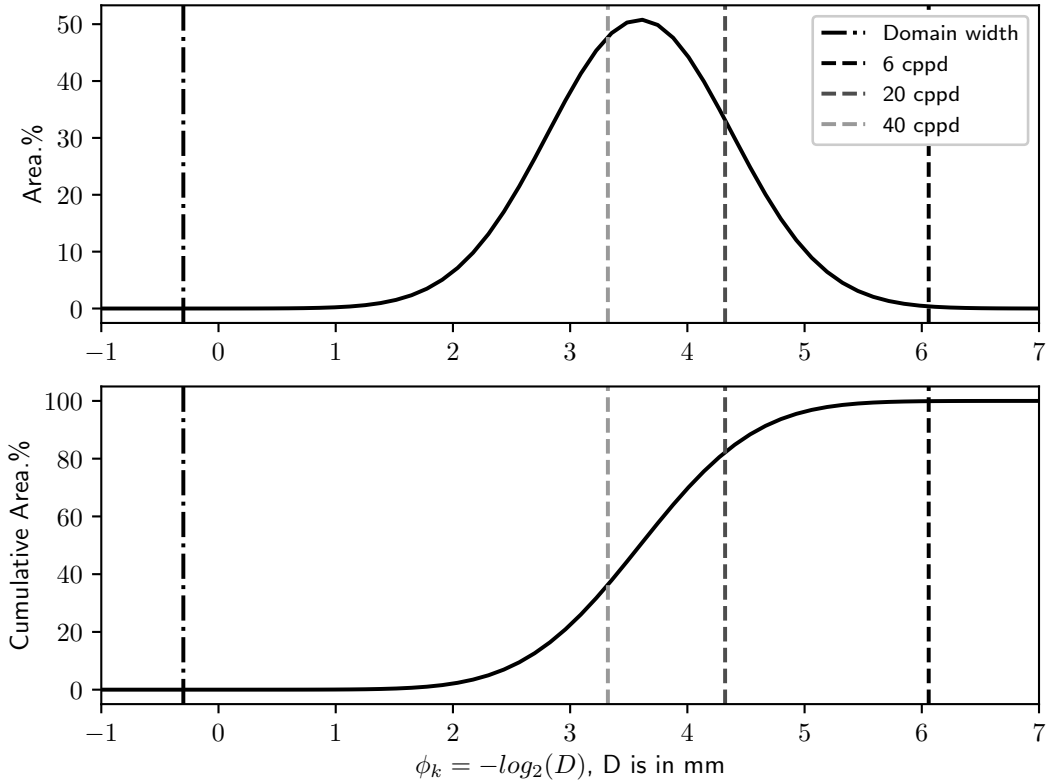


Figure 6.5: The PSD used in the explicit/implicit investigation. It possessed the same mean as the PSD in Figure 6.3 but a standard deviation that was three times smaller. The three main implicit/explicit cut-offs are shown as dashed lines.

6.2.3 Varying Matrix Uniformity

To investigate the effect of varying the spatial uniformity of the matrix, three different simulations were run. A baseline simulation was performed that possessed a uniform matrix distribution. Two subsequent simulations introduced non-uniformity to the matrix

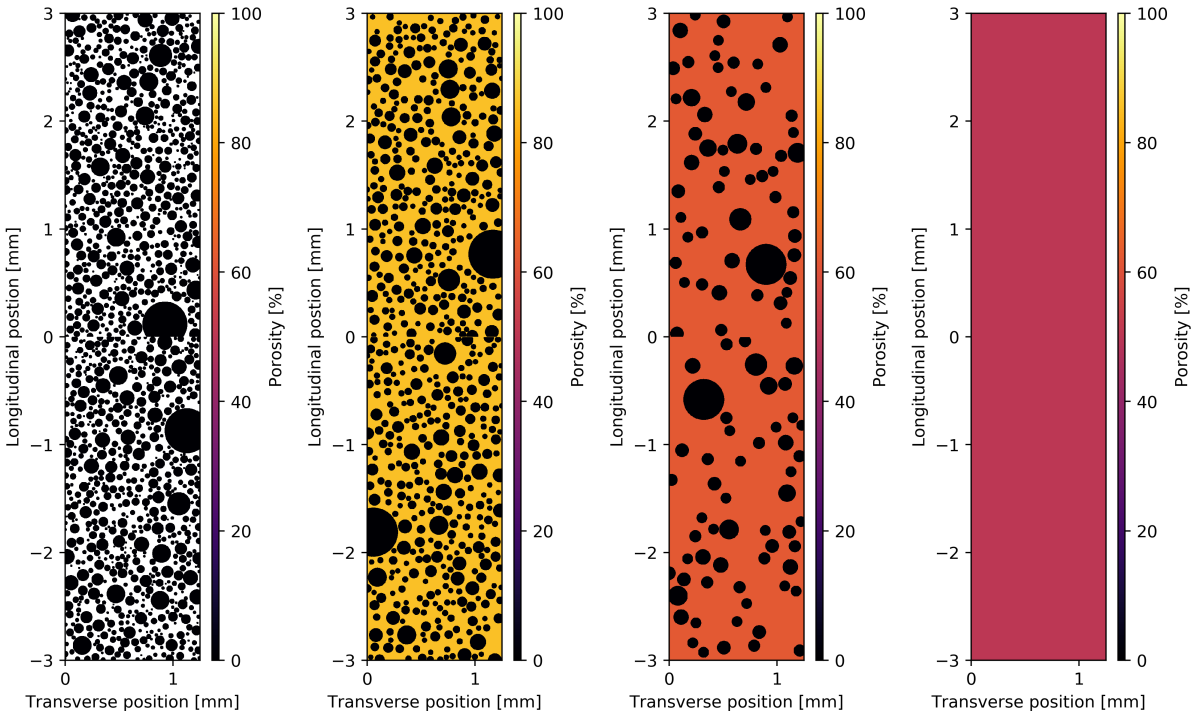


Figure 6.6: The starting geometries and porosities of the fours simulations that varied the explicit/implicit cut-off. From left to right, the cut-offs were: fully explicit, semi-explicit 20 cppd, semi-explicit 40 cppd, and fully implicit.

by adding explicit void pores (of volume fraction v_{fV}) and additional grains (small grains chosen from the PSD[3]) to keep the matrix porosity (and bulk porosity) constant. The pores were randomly placed in the matrix as circles and possessed the smallest allowed shape (8 cppd circle). The total area fraction of the explicit grains in the simulation was labelled as v_{fG}. The three simulations were $v_{fV} = 0\%$, $v_{fG} = 17\%$; $v_{fV} = 12\%$, $v_{fG} = 25\%$; and $v_{fV} = 25\%$, $v_{fG} = 33\%$ (see Figure 6.12).

6.2.4 Varying Grain Morphology

To investigate the effect of varying grain morphology on the shock response a single bed was set up using the base PSD and circular grains, as well as a fixed matrix porosity of 70%. The grain sizes and positions were recorded and three further simulations were

[3] $D \geq 8$ cppd

created using the same position but different grain shapes. Three grain shapes were considered in total: circles, squares, and 'realistic' grains that approximated common regolith morphologies. The area and orientation of each grain was fixed between simulations; occasionally this resulted in overlapping grains because of the new shapes, but this was allowed in order to minimise any other potential effects that would otherwise be introduced by changing the arrangement.

'Realistic' Grain Creation

'Realistic' grains were one of the morphologies considered, however, regolith grains are unlike most terrestrial grains: there is no atmosphere or water on the Moon, and so grains do not experience any chemical erosion, and the only effective erosion is through continuous meteorite bombardment (Heiken, Vaniman, and French 1991). This means that the majority of lunar grains are sharp, angular, and distinctly irregular in shape.

Realistic 2D grain models were constructed based on the regolith grains' irregular properties. Unfortunately not all of the parameters translated well for the purposes of reconstructing a grain. The volume coefficient, for example, is calculated as the 'volume of a particle divided by the cube of the diameter of the circle that encloses the same area as the particle profile' (Heiken, Vaniman, and French 1991). Whilst this provides valuable information about regolith grains it can not easily be used to construct a grain in reverse. As a first approximation to regolith grain morphology, three simple parameters were chosen to use in construction: elongation, roundness, and area ratio.

A simple algorithm was devised to create grains according to these parameters and their specified ranges. Grains were generated by creating a 2D rectangle with sides of length L and B. The first four vertices of the grain were generated at random positions along each side of the rectangle (one on each side). Four additional vertices were generated as midpoints between each adjacent pair of existing vertices. These four new vertices

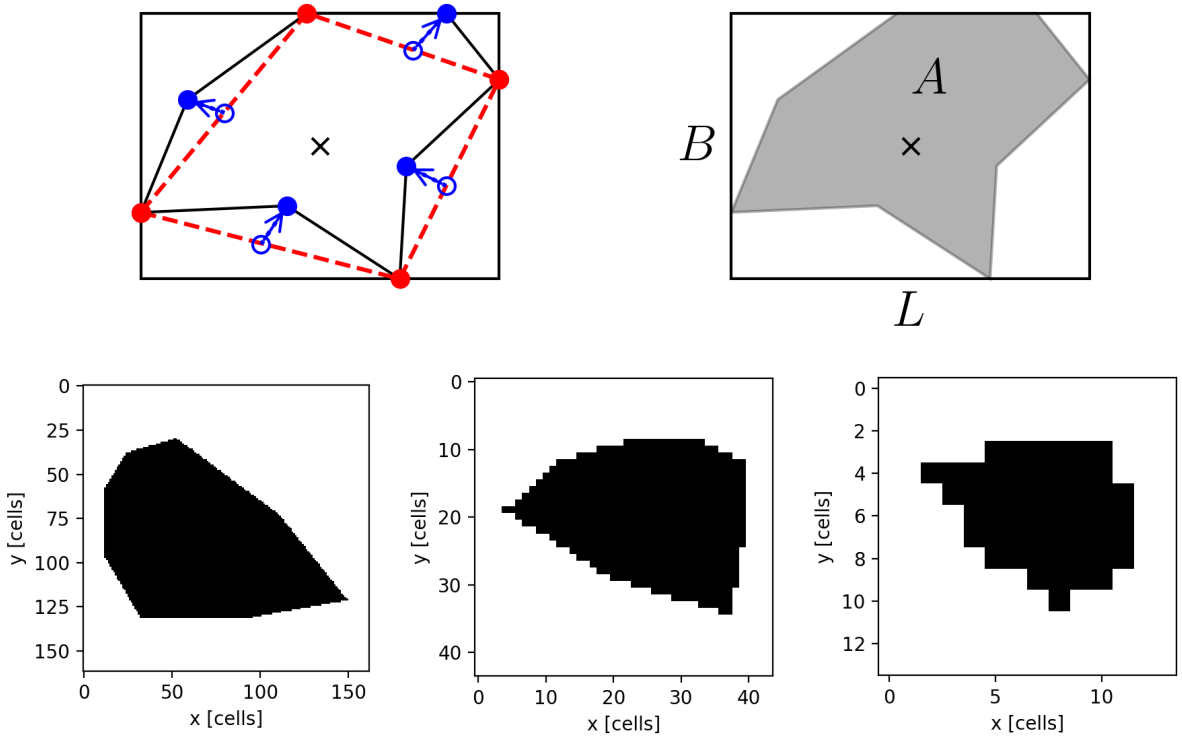


Figure 6.7: (Top left) Diagram demonstrating the grain creation algorithm: four vertices are chosen, one on each side of the rectangle. Next all midpoints are selected and moved radially out or in (at random) by $0.2q$, and $q \approx \frac{B}{2}$. If the result has an area ratio, $A_r = \frac{A}{L.B}$ between 0.65, and 0.75, the grain is added to the grain library. (Top right) the final example grain with area A, bounded by the box of shape $L \times B$. (Bottom row) Three grains of varying sizes, created by this algorithm for a simulation. The far left and right grains represent the extremities of grain sizes and are the largest and smallest grain sizes available (respectively).

were then shifted radially (relative to the centre of the rectangle) by $\pm 0.2q$ (direction was chosen at random and q was estimated to be $\sim B/2$). Finally, the area ratio of the shape was tested, and if it lay in a previously specified range it was deemed successful, and saved. If unsuccessful, the algorithm restarted and that grain was discarded. A library of grains was generated using this method.

To further simplify the grain creation algorithm, the roundness and elongation were fixed, whilst the area ratio was allowed to vary. Elongation was defined as $E = \frac{L}{B}$, where L & B

refer to the length and breadth of the smallest cuboid that totally encloses the grain (Figure 6.7). Elongation was set as 1.35 in the creation algorithm. The roundness of each grain was set at 0.2. The area ratio, A_r, was defined as $A_r = \frac{A}{L.B}$ in this investigation and allowed to vary in the range $0.65 - 0.75$. This was a slightly larger range than observed in regolith investigations. This choice was made such that the algorithm would generate a greater variety of shapes whilst all other parameters were fixed.

The above method was suitable for this application, despite the approximations and simplifications, for two reasons. First, the simulations possessed a minimum resolution allowed and most grains were small. This meant that most grains had a small number of pixels to represent a complex geometric shape. The more complex shape a grain possessed the less likely it was to be appropriately resolved as a small grain (Figure 6.7). By keeping the shapes allowed limited to eight vertices, inaccuracies introduced by unresolvable complexity were reduced. Second, the majority of parameters used to classify regolith grains were only intended to be used on 3D shapes, and using them to create 2D shapes is already an approximation. An approximation of grain geometry was sufficiently accurate in this case.

6.3 Results

6.3.1 Varying DE_{min}

Bulk Response

The bulk responses of each simulation of the regolith are in reasonable agreement (Figure 6.8). To avoid confusion, and repetition, pressures and temperatures of the individual simulations will be referred to by the subscripts: E and I, for the explicitly resolved and implicitly resolved regolith simulations; and 20 and 40 for $D \geq 20$ and $D \geq 40$, respectively. Standard deviations will be represented by σ with the subscript P, T, or ϕ

to indicate which field they are related to. $\sigma_{P,E}$, for example, is the standard deviation on the mean peak bulk pressure of the explicitly resolved regolith simulation. The mass-weighted averages of the peak temperatures are consistent across all simulations, and the bulk peak temperatures of the $D \geq 20$ cppd (T_{20}), and $D \geq 40$ cppd (T_{40}) semi-explicit simulations differ from the fully explicit simulation (T_E) by $\sim 0.5\%$ and $\sim 12\%$ respectively.

The bulk average peak pressures, however, decrease as more and more of the material is parametrised. Interestingly the two semi-explicit simulations are approximately consistent with each other in terms of bulk response, and they differ by less than 1%. Their standard deviations are also similar: $\sigma_{P,20} \approx 2.1$[GPa] and $\sigma_{P,40} \approx 2.4$[GPa]. P_{20} & P_{40} differ from P_E by $\sim 10\%$ and are within σ_{PE}. They are in good agreement, but still follow the downward trend slightly as more material is parametrised. The lowest pressure value is exhibited by P_I, which is $\sim 45\%$ less than the P_E. This discrepancy can be attributed to the lack of mesoscale detail in the implicitly resolved regolith simulations; the mesoscale features cause large variations in pressure (spatially and temporally) and the peak pressures only capture the highest pressures. This means that comparing P_I with P_E is not a fair comparison. Logically P_I should be consistent with the shock state of the material, which will be an average of all the variation, as opposed to a maximum. When the post-shock pressures are considered (calculated using a modified stress-contour method) a more reasonable agreement is observed between all simulations (Figure 6.9).

The semi-explicitly resolved regolith simulations do not appear to achieve total compaction of the particle bed $\phi_{20}, \phi_{40} \gtrsim 0$. This is particularly significant because total compaction is achieved in the fully explicit simulation, which is simulating the same scenario. This difference is a result of the higher initial matrix porosities in the semi-explicit simulations. The impact conditions were sufficient to achieve the necessary stresses to crush all pore space in the implicitly, and explicitly resolved simulations. However, the

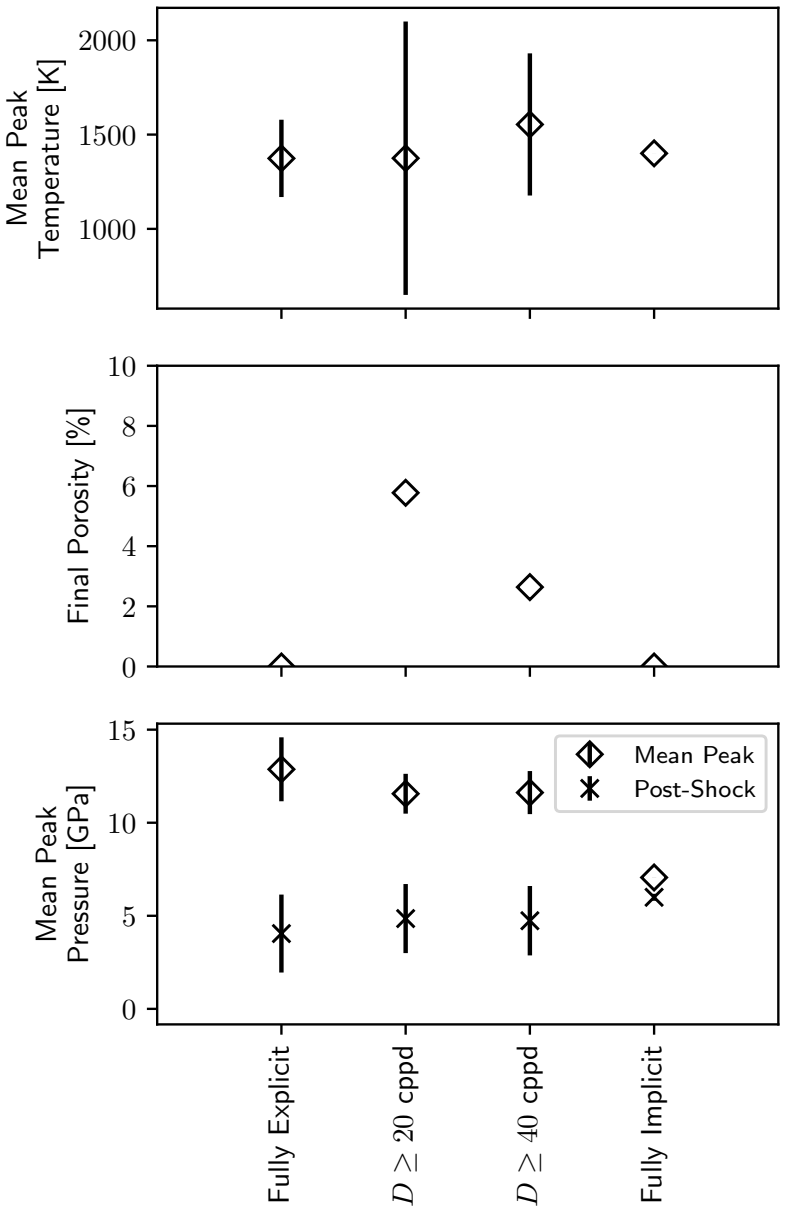


Figure 6.8: Graphical representation of the bulk responses of the four explicit/implicit simulations; error bars represent ± half a standard deviation. All bulk response averages were weighted by the mass of the tracers.

semi-explicitly resolved simulations possessed initial matrix porosities much greater than the bulk 50%. In the $\epsilon-\alpha$ model a greater strain is required to crush a powder by the same amount, as porosity decreases (see Chapter 3), and eventually the porosity asymptotes to zero. Increasing the amount of explicitly resolved grains is a trade-off with the effects

of increasing the matrix porosity. This effect appears to dominate the response of the simulation more and more as DE_{min} is decreased, until the matrix becomes entirely void, and the effect disappears. This is a weakness of the semi-explicitly resolved simulations and evidence that the parametrised material in them is not an accurate representation of its behaviour. The final porosities of the semi-explicitly resolved simulations are $\phi_{20} \approx 5\%$ and $\phi_{40} \approx 3\%$.

GSC Responses

To compare the response of specific GSCs between simulations the tracers within each sample region were divided into four categories depending on the object they started in. Tracers in grains were divided into three grain size categories: $6 \leq D < 20$ cppd, $20 \leq D < 40$ cppd, and $D \geq 40$ cppd. The fourth category consisted of any tracers that were present in the matrix; i.e. implicitly resolved grains. No simulation contained all four categories because the matrix was used instead of at least one category in each simulation, except in the fully explicit simulation, which had no matrix. The parametrised classes are shown in Figure 6.9 by the parentheses added to their label on the x-axis (e.g. '($D \geq 40$ cppd)').

In the semi-explicit simulations the mean peak temperature matrix response does not appear to agree with the explicit counterpart, which it is parametrising. Only the fully implicit simulation matrix broadly agrees with the temperatures exhibited by the fully explicit categories (Figure 6.8). The matrix temperature appears to decrease as more of the material is parametrised, and this correlates with the matrix porosity that each simulation started with. The highest matrix temperatures are observed in the simulation with the highest matrix porosity (85%) and the matrix temperatures decrease as the matrix porosity decreases (65% and 50%). This is not unsurprising because a more porous material experiences greater heating than a less porous one, and mass-weighting

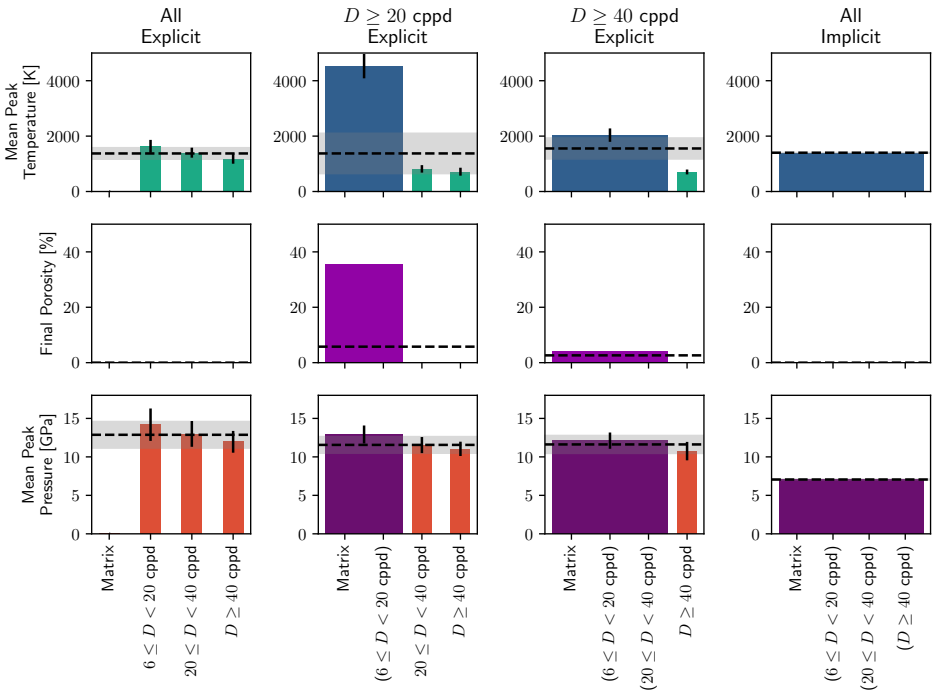


Figure 6.9: Graphical representation of the bulk (and GSC) responses of the four simulations varying DE_{min}; error bars represent $\pm$ half a standard deviation. The horizontal, black, dashed line in each plot represents that simulation's bulk response (mass-weighted average) and the grey shading around it represents $\pm$ half a mass-weighted standard deviation. The GSCs are ranges of particle diameters (D), in cells.

is only applied when calculating the bulk response. The explicit grains also exhibit lower temperatures in the semi-explicit simulations when compared to their counterparts in the fully-explicit simulation. The largest GSC's ($D \geq 40$ cppd) mean peak temperature is $\sim 50\%$ lower in the semi-explicit simulations than the fully-explicit one. The second largest GSC is also $\sim 50\%$ lower in the $D \geq 20$ cppd semi-explicit simulation.

The final matrix porosities exhibit a similar pattern to the final bulk porosities: the $D \geq 20$ cppd simulation exhibits a high final porosity, and the matrix of the $D \geq 40$ cppd simulation exhibits a lower porosity that is more consistent with the fully explicit simulation's final porosity.

By way of contrast, the mean peak pressures of grains in each GSC are broadly consistent

across simulations. All semi-explicit mean peak pressures, of each category, are within half a standard deviation of their category's result in the fully-explicit simulation, except for the implicitly resolved regolith simulation. This, however, is not a fair comparison (as discussed in the previous section) and is only included for completeness.

6.3.2 Explicit Void

Bulk & GSC Response

Three simulations were performed investigating the effect of increasing non-uniformity in the matrix. The grains of the reference simulation are referred to as 'Grain set 1' ($G1$). The grains introduced for the second simulation are referred to as 'Grain set 2' ($G2$) and the last lot of grains introduced for the third simulation are referred to as 'Grain set 3' ($G3$). Increasing the non-uniformity of the matrix generally appears to increase the bulk pressures, and decrease the bulk temperatures (Figure 6.10). Although, all the bulk temperatures and pressures are consistent across simulations with the same V_I. These trends are not easy to observe at the lowest V_I, although, the results suggest that most effects were small at this velocity because the peak pressures experienced by the grains and matrix were only slightly greater than zero. In general, introduction of non-uniformity has a minimal effect on the bulk response of the powder. The largest effect present is in the bulk porosity, which decreases significantly as non-uniformity increases.

In terms of the GSC responses, each set of grains appears to experience higher temperatures and pressures as the non-uniformity increases (Figure 6.11). The matrix exhibits the clearest trend, demonstrating increasing pressures and temperatures as non-uniformity increases. Interestingly the final porosity of the matrix appears unchanged by the introduction of non-uniformity, although, the value shown in Figure 6.11 is the average porosity of matrix tracers only, and does not include fines. The increase in the pressure and temperature response, with increasing numbers of pores, is not unexpected. Previous studies

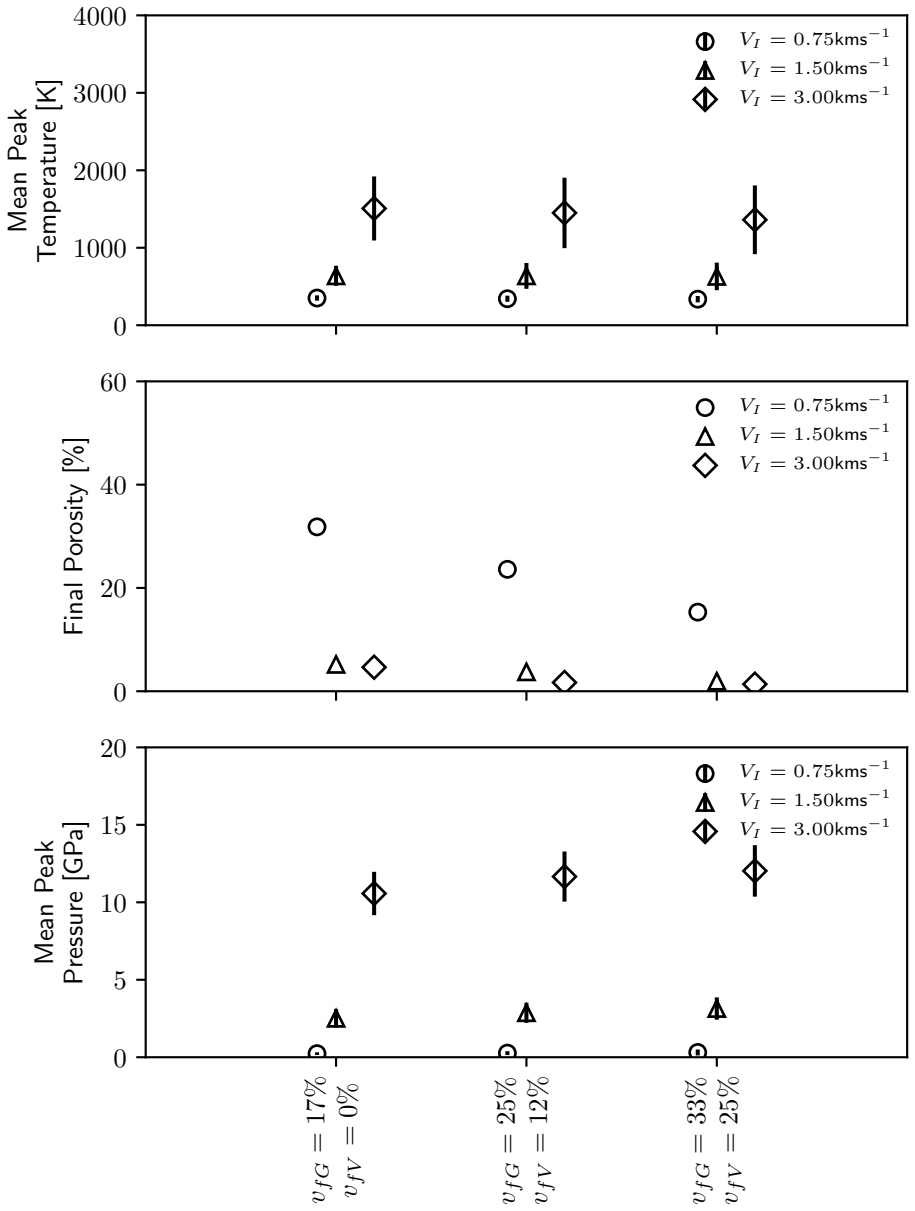


Figure 6.10: Graphical representation of the bulk responses of all samples using varied amounts of matrix non-uniformity; error bars represent ± half a standard deviation.

of porous media have shown that this is a characteristic of pore response (Kowitz et al. 2013).

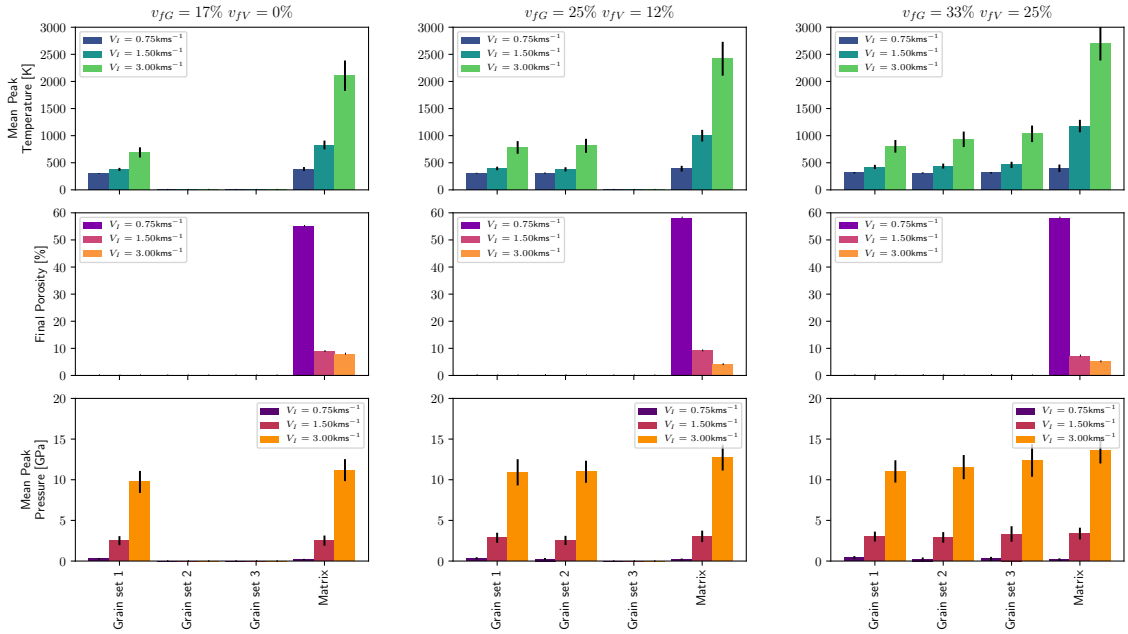


Figure 6.11: Graphical representation of the GSC responses of all samples using varied amounts of matrix non-uniformity; error bars represent ± half a standard deviation.

Mesoscale Response

The mesoscale results are shown in Figure 6.12. Generally the three simulations exhibit the same temperature distribution, independent of the amount of void explicitly resolved. The distribution is simply 'missing' where voids have been added. This is an interesting result because the tracers surrounding each pore are not exhibiting temperature spikes. It appears that the three simulations experienced the same temperature response and the addition of voids obscured it. The higher temperatures observed in the previous section (in simulations with large matrix non-uniformity) may be caused by the fact that there are fewer tracers to average over rather than being caused by hot-spots generated by pore collapse. Additionally, the grains that were added to increase v_{fG} were all small and probably experienced the greatest heating/peak temperatures. This observation, in addition to the fact that a higher v_{fG} results in more grain tracers, suggests that the addition of explicit pores does not necessarily have a strong effect on the temperatures experienced by a simulation. Figure 6.12 suggests that the particle arrangement matters

most.

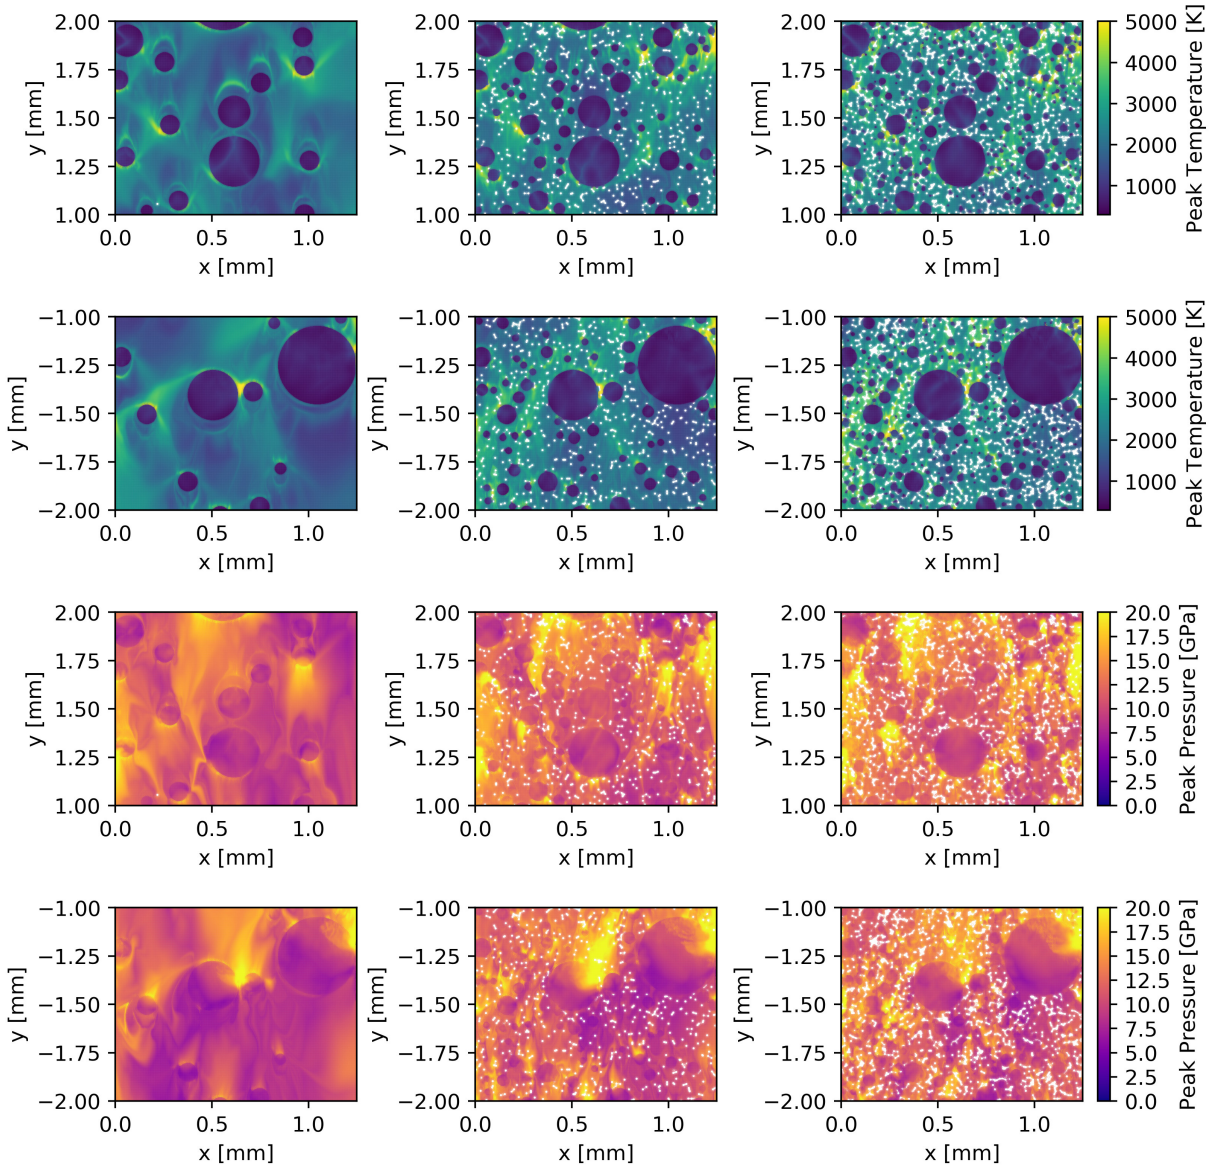


Figure 6.12: Graphical representation of the mesoscale response of each simulation that used a different quantity of explicit void/explicit grains (for a V_I of 3 kms^{-1}). Only the sample regions are shown because the pores are sufficiently small that they can not be resolved graphically when the entire simulation is plotted. In each pair of rows, the top image is the top sample, and the bottom is the lower sample. (Top two rows) The tracer peak temperature distribution, plotted in the particle bed's starting state. (Bottom two rows) The tracer peak pressure distribution, plotted in the particle bed's starting state. (Far left) Explicit void fraction, $v_{fV} = 0\%$, explicit grain fraction, $v_{fG} = 17\%$; (middle) $v_{fV} = 12\%$, $v_{fG} = 25\%$; (right) $v_{fV} = 25\%$, $v_{fG} = 33\%$.

The pressure results, however, suggest a different conclusion. The pressure distribution is

broadly similar between the three simulations but, as v_{fV} increases, so does the amount of material experiencing pressure spikes. The existence of explicit pores appears to have caused localised pressure spikes and exacerbated the effects in existing regions of high pressure.

6.3.3 Morphologies

Bulk Response

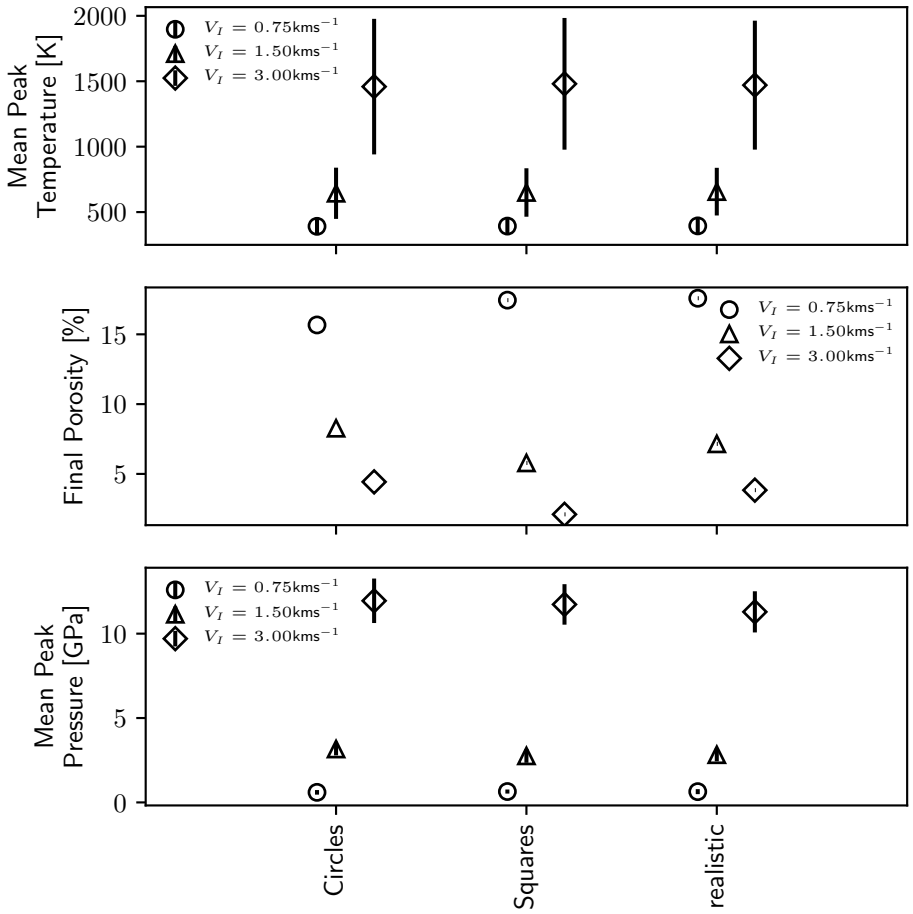


Figure 6.13: Graphical representation of the bulk response of all PSD samples; error bars represent ± half a standard deviation.

The bulk response of regolith with varying grain shapes appears to be independent of the grain shape because all bulk values are in reasonable agreement at each V_I (Figure

6.13). The results (whilst consistent) are not identical, however, and some small variations are present. For example, the circular grain simulations show the most compaction for $V_I = 3$ kms^{-1}, but the least for $V_I = 0.75$ kms^{-1}. The reverse is true of the square grain simulations (least compaction for $V_I = 3$ kms^{-1} and most for $V_I = 0.75$ kms^{-1}) and the realistic grain simulation results lie approximately in the centre. Otherwise, differences between the bulk responses are challenging to observe.

To investigate and compare the GSC responses in the morphology simulation results, the powder bed was divided into the same GSCs as those used in Section 6.3.1; the GSC responses do not exhibit many differences that can be attributed to the differing morphologies (Figure 6.14). Generally grains in each GSC appear to be consistent with the equivalent grains in other simulations. Morphology does not have a noticeable effect on the bulk shock response of the powder.

Mesoscale Response

Considering the samples' response at the mesoscale, however, reveals that whilst the bulk and GSC responses do not differ greatly, local differences are present. Figure 6.15 shows the spatial distributions of peak tracer pressures and temperatures in their starting locations for all three simulations. The temperatures are largely consistent between simulations: the spikes and lows occur in approximately the same locations, however, the circular, and realistic grains appear to exhibit more spikes in temperature on their shockward surfaces than their counterparts in the square grain simulation. The matrix surrounding one large grain on the left hand side at $y \approx 0.5$ mm makes this clear. In the square-grain simulation the matrix that lies on its shockward face does not exhibit temperature variations greater than those experienced by surrounding matrix. In the circle grain simulation it experiences two temperature spikes on opposite transverse sides of its shockward face. In the realistic grains simulation the matrix on this grain's shockward

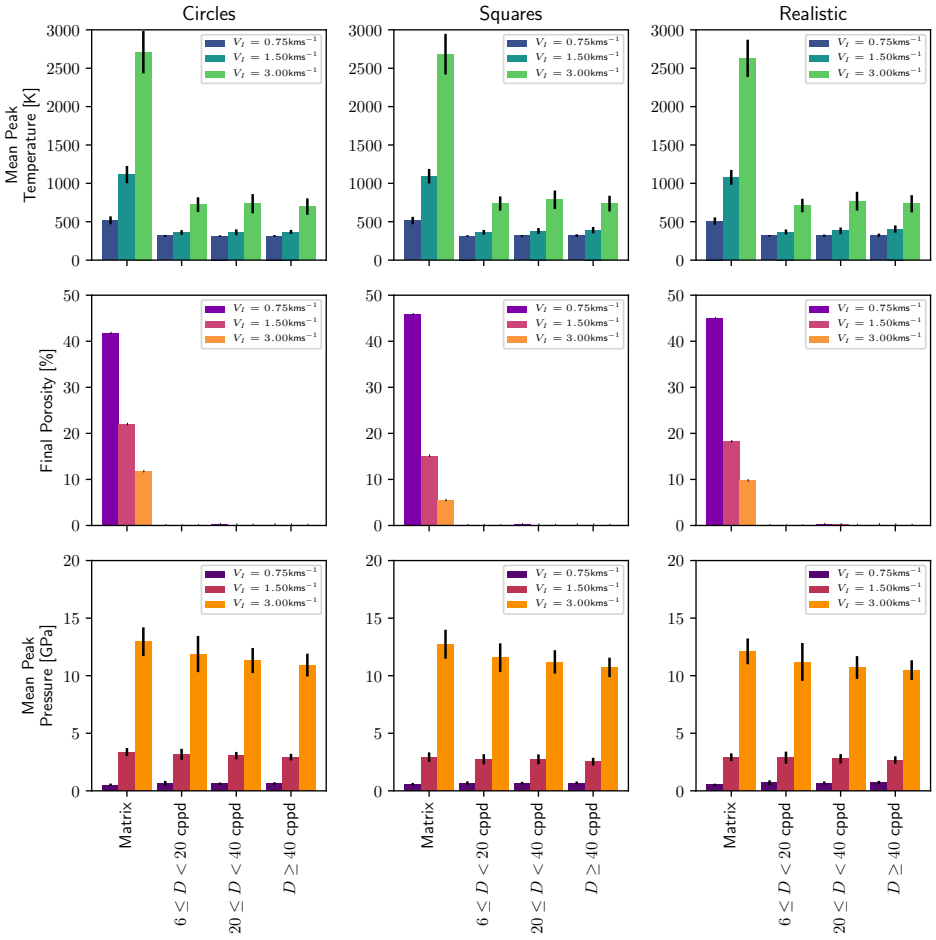


Figure 6.14: Graphical representation of the GSC responses of simulations using different morphologies; error bars represent ± half a standard deviation.

face exhibits a large spike in temperature compared to the surrounding matrix all along the face. Given that there are few grains shockward of this grain, it is probable that the grain experienced a similar shock wave in each simulation. It is plausible that the grain morphology is the cause of this difference, even if it is not clear why this would be the case.

Unlike the temperatures, the pressure distributions are too chaotic to interpret meaningfully. The responses appear very similar and do not exhibit any noticeable mesoscale differences between grains. The realistic grains may exhibit slightly lower bulk pressures than the other simulations but this is hard to decipher from these data (Figure 6.15).

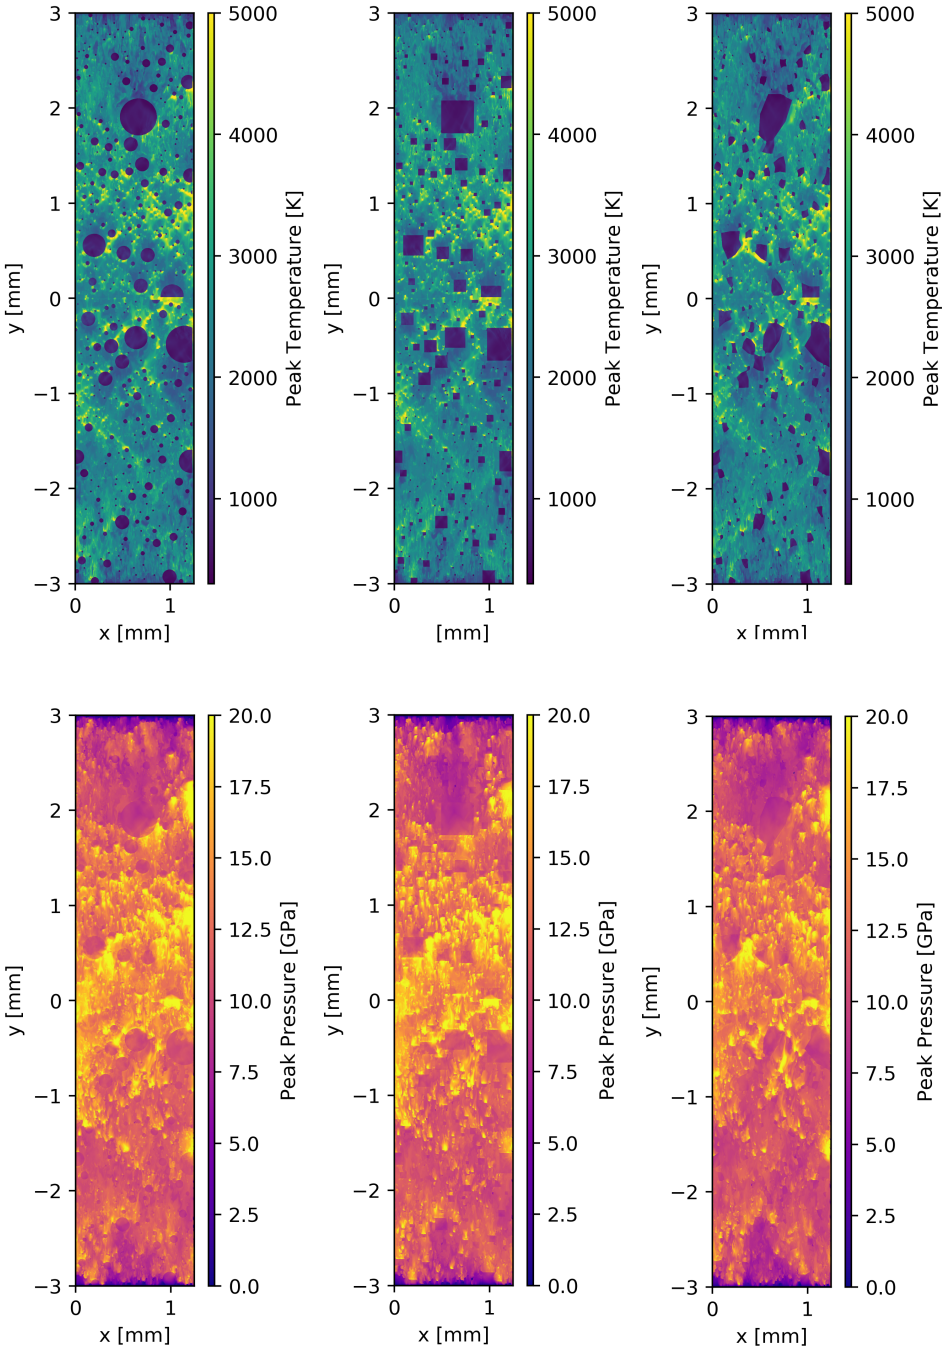


Figure 6.15: Graphical representation of the mesoscale response of each simulation employing different morphologies. (Top row) The tracer peak temperature distribution, plotted in the particle bed's starting state. (Bottom row) The tracer peak pressure distribution, plotted in the particle bed's starting state. (Left) Circular grain morphology; (middle) square grain morphology; (right) realistic grain morphology.

6.4 Discussion

6.4.1 Semi-explicit method suitability for modelling complex powders

Overall, the results presented here suggest that resolving a granular material semi-explicitly reproduces the bulk response of the explicitly resolved powder accurately. There are, however, some weaknesses to the method. In particular, the final bulk porosity is sensitive to the initial porosity of the matrix, and the higher it is, the higher the final bulk porosity. In addition, whilst the bulk response may be consistent, the larger GSCs tended to exhibit lower temperatures and pressures as DE_{min} increased. In particular, the porosities and peak temperatures in the matrix are the most inconsistent, but the pressures remain largely in agreement. Neither the implicitly resolved, nor the semi-explicitly resolved, regolith simulations are able to reproduce the explicitly resolved regolith simulation results completely, although it is reasonable to say that the semi-explicitly resolved regolith simulations come closest.

The results from simulations varying DE_{min} also suggest that there is a small temperature dependency on grain size. Explicitly resolved grains in GSCs covering smaller grains typically experience higher temperatures than grains in GSCs covering larger grains. This result appears to be insensitive to DE_{min}, even if the absolute values are not completely consistent between simulations, the relative variation remains consistent. This suggests that the temperature dichotomy between matrix and chondrules, documented by Bland et al. (2014), Davison, Collins, and Bland (2016), and Davison et al. (2017), is a physical effect.

The assumption of the uniformity of the matrix made by the semi-explicit, and implicit, methods may be the cause of the remaining discrepancy between the results. Section 6.3.2

showed how introducing explicit pores (as well as grains) increased the temperatures and pressures experienced by the matrix, whilst reducing its final bulk porosity. In addition, Section 6.3.1 noted how the high porosities of the matrix in the semi-explicitly resolved regolith simulations may be the cause of the high temperatures exhibited. Additional simulations were performed with the intention of combining these two effects. Explicit pore space was introduced to the $D \leq 20$ cppd simulation to lower the high matrix porosity, raise the pressures experienced and raise grain temperatures. Three simulations were performed, each with a different matrix porosity ($\phi = 65\%$, 50%, and 35%) and the equivalent number of explicit pores required to achieve them.

The results of these simulations show that introducing explicit porosity reduces the matrix temperatures, whilst keeping the bulk response consistent with the explicit simulation results (Figure 6.16). In addition, the pressures have been raised, and are now more consistent with the explicit simulation (in terms of the GSCs and the bulk). The final porosities have also been reduced, and now the bulk porosity is always zero, or very close to zero. The matrix, however, still possesses some final porosity in the simulations where a higher starting porosity is used. Overall, the most optimal simulation set-up, and the one that is most consistent with the fully explicit results, is the simulation that has a matrix porosity equal to the total bulk porosity of the entire sample. These observations suggest that the semi-explicit method can be used to simulate complex powders, such as the lunar regolith, but that the parametrised material should possess a bulk porosity that is equal to the total bulk porosity of the sample.

6.4.2 Polygon Matrix-Shock Interaction

In terms of the morphology, it is difficult to say how well the results shown here agree or disagree with existing literature. All investigations (so far) that have semi-explicitly resolved powders have (justifiably) considered spherical/circular explicit grains, and none

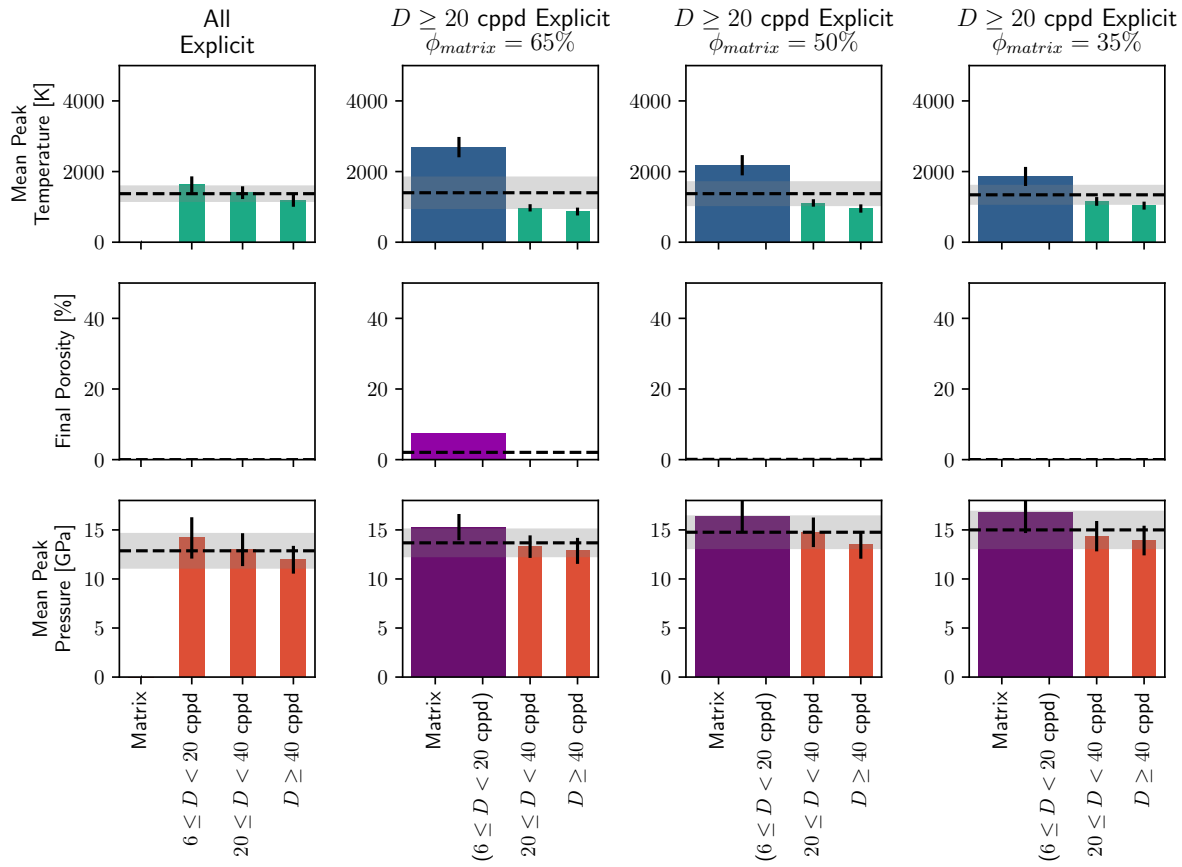


Figure 6.16: Graphical representation of the bulk (and GSC) responses of the three semi-explicit simulations with varying amounts of explicit void, compared with the fully explicit simulation results. Error bars represent ± half a standard deviation. The horizontal, black, dashed line in each plot represents that simulation's bulk response (mass-weighted average) and the grey shading around it represents ± half a mass-weighted standard deviation. The GSCs are in ranges of particle diameters (D), in cells.

have directly considered the effect of varying particle morphology. In this regard there are no other studies available for direct comparison, but the results suggest that the bulk response is insensitive to morphology, which is consistent with existing literature (Benson and Nellis 1994; Borg and Vogler 2009).

The mesoscale results, however, do appear to be sensitive to morphology. There are some differences in the mesoscale results that can not be easily explained. The increased heating along some shockward-facing grain faces in the 'realistic' grain morphology simulation, for

example. Whilst the bulk response is relatively unaffected by changes to grain morphology, the distribution of peak temperatures within each grain's 'area-of-effect' (AoE) is affected and it is unclear why this is.

It is not known how the AoE is altered by morphology. In order to understand how morphology changes a grain's AoE, additional simulations were performed that examined isolated and individual, grains in a simplified set-up. These considered how the fundamental aspect of a polygonal grain (flat faces) caused variations in a grain's AoE.

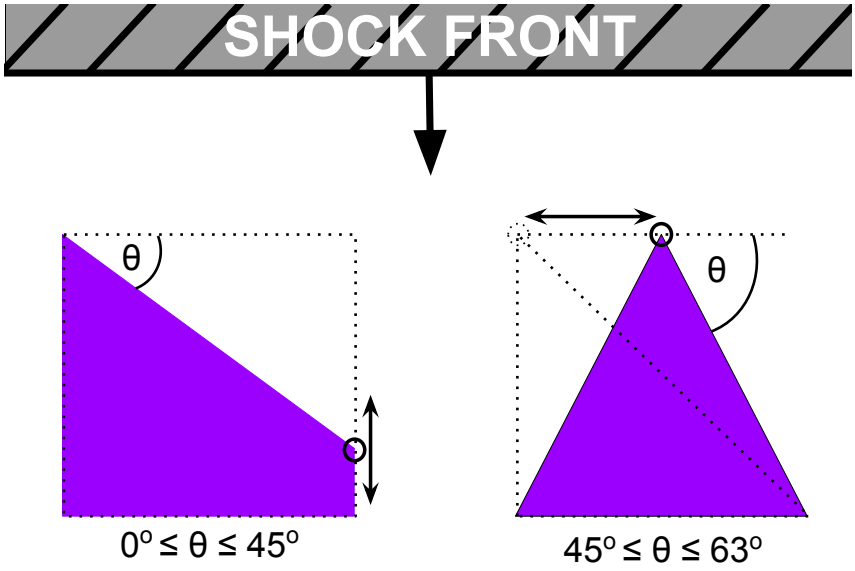


Figure 6.17: Demonstration of how grain shapes were constructed to produce a range of θ. Two shapes were created and the position of one vertex was varied along one side of each.

The fundamental difference between polygonal and circular grain shapes is that polygonal grains (squares and 'realistic' shapes) have flat faces. A flat face interacts with a shock wave in a unique manner depending on the angle between the shock wave and the grain's face. Multiple simple simulations were performed that used two samples consisting of a single grain embedded in matrix ($\phi = 70\%$) colliding at 3 kms^{-1}. Physical dimensions were kept small, such that each grain was of a comparable size to those in the full simulations.

These new simulations were used to investigate how the angle of the grain's face, with the shock wave, affected its AoE's response. Two shapes were considered: a trapezium where the length of one side was varied to achieve multiple angles with the shock front (θ). At one extreme the shape was a square ($\theta = 0°$) and at the other the shape was a right-angled triangle ($\theta = 45°$). For angles greater than 45 degrees, a second shape was used (Figure 6.17). The second set of grain shapes started with the right-angled triangle shape ($\theta = 45°$) and then the tip of the triangle, which pointed toward the shock wave, was moved parallel to the shock front in increments to create higher values of θ (Figure 6.17). The incremental movement of the triangle's vertex continued until the shape became symmetric with respect to the direction of shock propagation, producing a range of angles with the shock of $0 \leq \theta \lesssim 60°$. The grain's area was kept constant across all shapes to avoid any changes in size caused by the altered geometry. When considering the polygons chosen for this study, it can be seen that each grain will actually have two shock-facing sides (although in the trapezium case this will have $\theta = 90°$). In both the trapezium and triangle grain cases, the smallest value of θ out of all the shock-facing faces was taken as the value of θ for that grain because the shock would reach this face first and interfere with the response of any other faces.

The results are presented in Figure 6.18. They suggest that at large angles with the shock wave, the matrix in the grain's AoE experiences less heating along the shock-facing side than at a shallower θ. However, at the shallowest angle (zero) the heating is also smaller than at larger angles. There appears to be an optimal value of θ at which the heating of the matrix in the grain's AoE is greatest.

A similar conclusion can be drawn from the peak pressure results; the peak pressures in the AoE's matrix along the shock-facing side appear to be at a maximum at this mid-range angle. Each grain appears to create a high-pressure region in the matrix that extends away from the shock-facing surface at right angles. This region is approximately

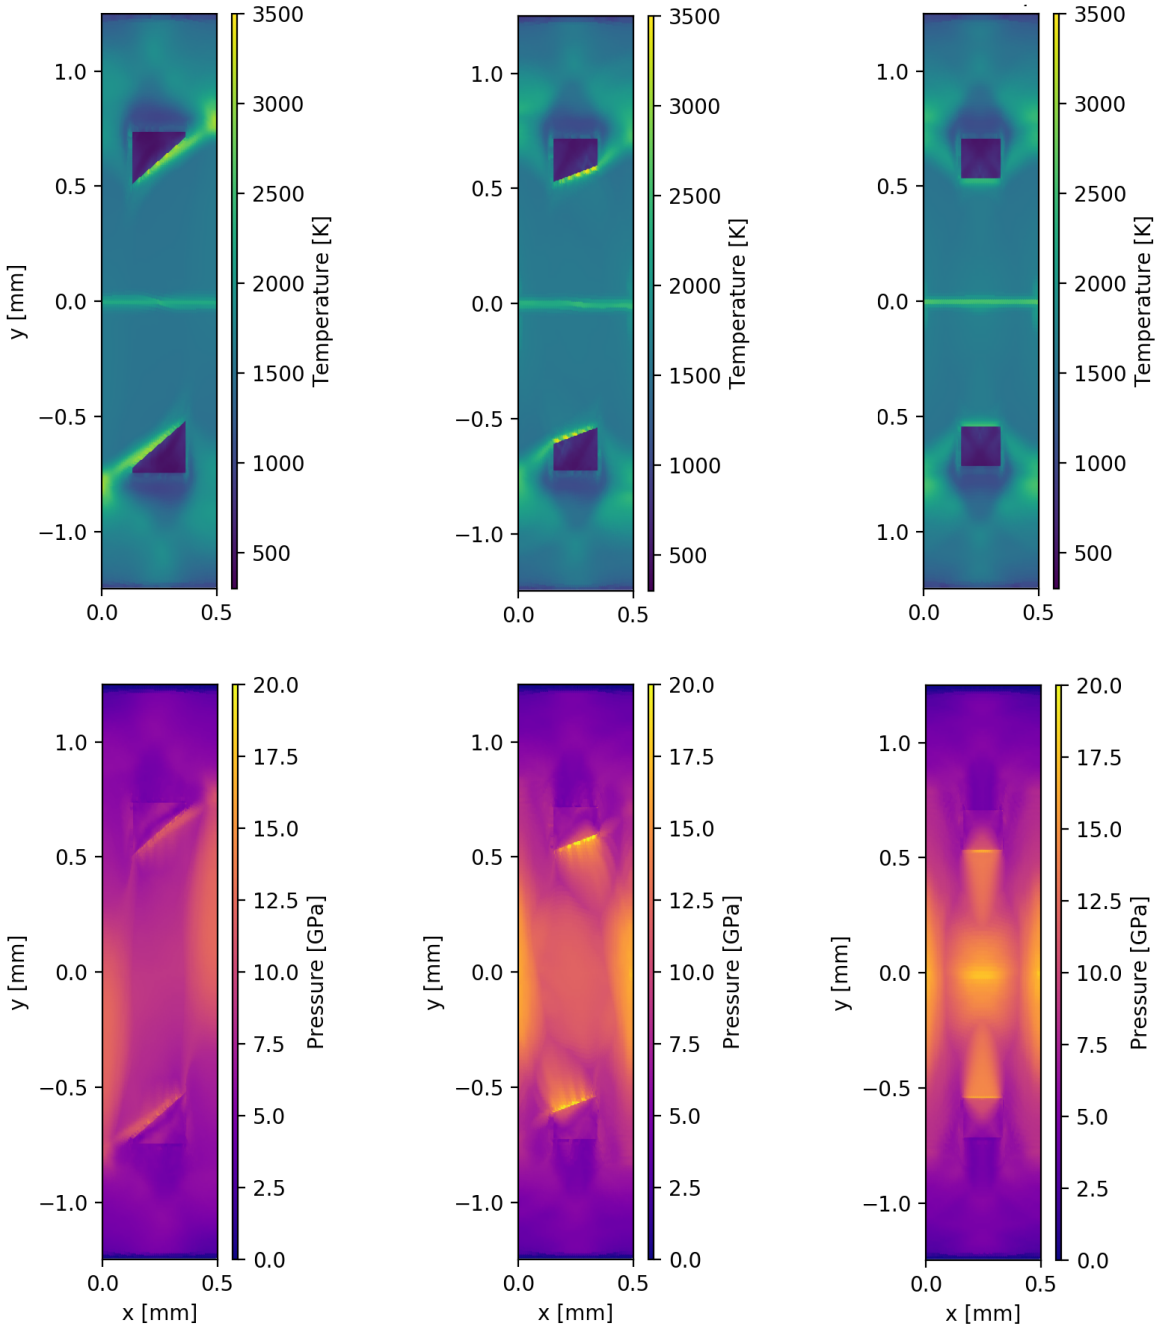


Figure 6.18: Three examples of the results of simulations varying the angle of a grain face with the shock wave. In each case the shock encounters the grains whilst planar. Peak temperatures are on the top row and peak pressures are on the bottom, as in previous figures. The grain area is kept constant as the shape is changed.

the same width as the shock-facing grain side, and tapers away with increasing distance from the face. At sufficiently large θ this region disappears, and a region of lower peak

pressure takes it place. Interestingly, there is also a low pressure region on the leeward grain side and, as θ increases, this appears to increase in magnitude, although, the effect is slight.

This optimal θ effect is notably uncommon in Figure 6.15. This is because of the influence of the heterogeneity of the granular ensemble, which is the cause of most of the localised structure variations. Grain-grain interactions between grains of comparable sizes, will affect the peak shock pressures and temperatures more strongly than the effect of θ through their collisions. Simultaneously, these interactions will distort the shock wave to non-planar shapes, which will also affect the prevalence of the effect. The results in Figure 6.18 suggest that θ does have an effect, but that specific local conditions must be achieved (a lack of heterogeneous effects) for it to be noticeable.

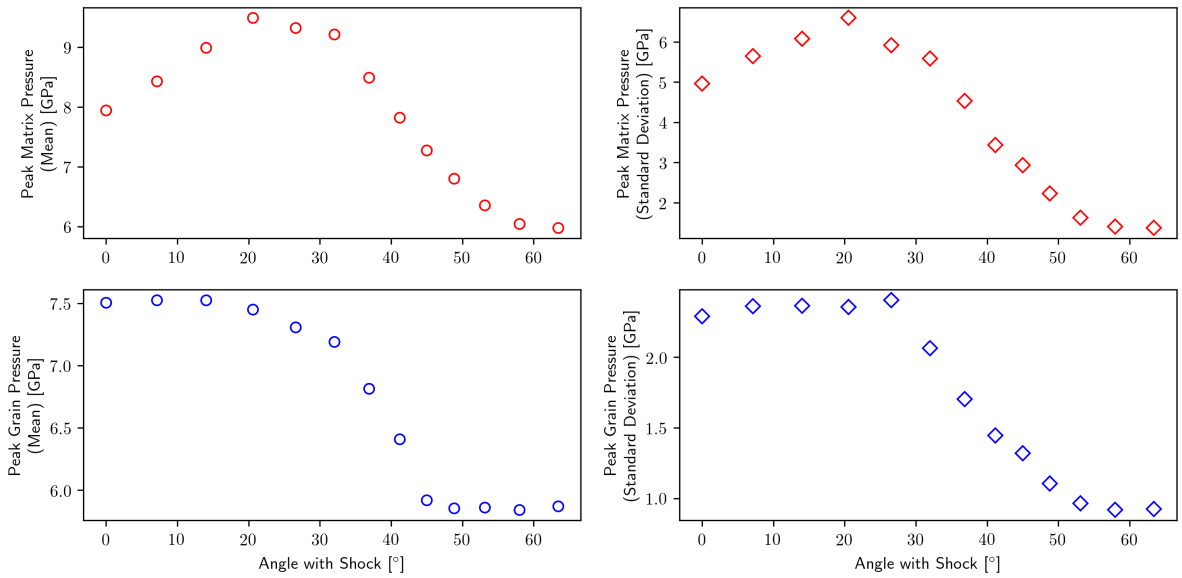


Figure 6.19: Variation of the AoE mean (and standard deviation) peak pressure with the angle of the grain's shock-facing side with the shock wave (θ). The mean (and standard deviation) of the peak pressure of the grain is also included.

To examine the relationship between θ and the temperatures and pressures experienced within the AoE, mean and standard deviations of peak pressure and temperature were extracted from the material around each grain. Any tracers within 10 cells of the grain's

edges (pre-impact) were considered to be within the grain's AoE and used as the 'sample'. Additionally, the results of the top and bottom AoEs were averaged to produce one value per simulation. The results of this analysis are presented in Figures 6.19 and 6.20.

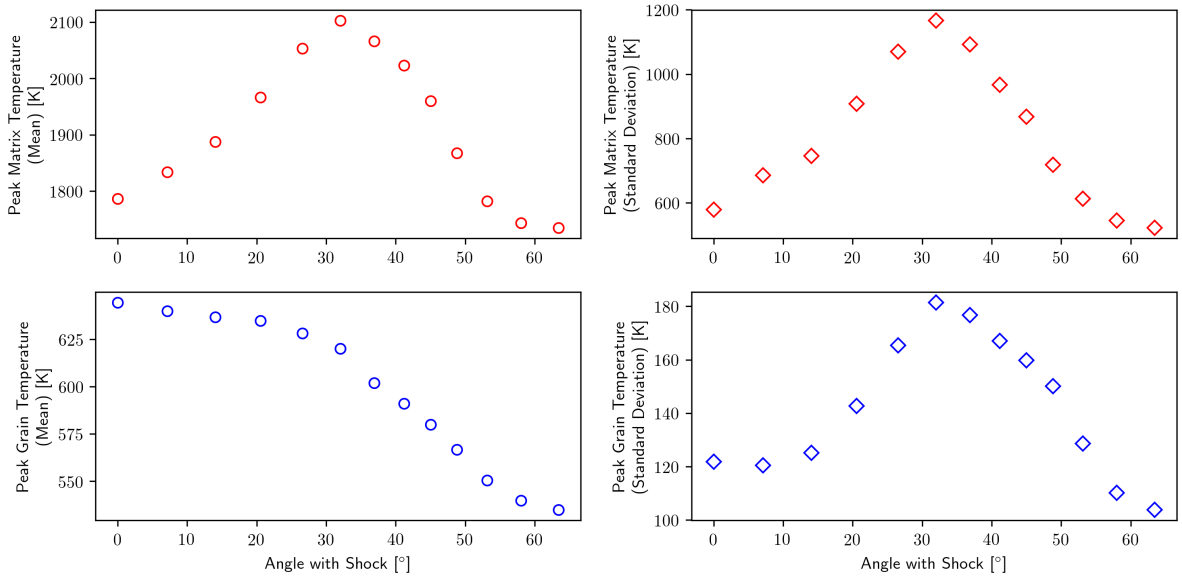


Figure 6.20: Variation of the AoE mean (and standard deviation) peak temperature with the angle of the grain's shock-facing side with the shock wave (θ). The mean (and standard deviation) of the peak temperature of the grain is also included.

The results in Figures 6.19 and 6.20 suggest that there is an optimum angle of θ that maximises the temperatures (and pressures) experienced by the grain's shock-facing edge. The pressures are greatest at $\theta \approx 20°$, and the temperatures are greatest at $\theta \approx 30°$. The standard deviations are also largest in the same places and show a similar dependence on θ to their cognate mean values. This dependency is likely a consequence of the fact that the increase in temperatures and pressures experienced by the matrix in the AoE only occurs in proximity to the shock-facing side. The matrix in the AoE near the other sides is relatively unchanged, as shown in Figure 6.18, which will widen the distribution of field values, and raise a standard deviation on the distribution.

A different pattern is present in the grains themselves. The mean peak pressures and temperatures within the grains are at a maximum when $\theta = 0°$ and drop off as θ in-

creases, with a point of inflection just after the critical angle, to a minimum at $\theta \approx 60°$. The standard deviations of the intra-grain pressures and temperatures, however, exhibit behaviour that is comparable to the standard deviations of the AoE's matrix: they are lower at high and low θ, with a maximum around the critical angle.

The existence of an optimum value of θ that generates the largest pressures and temperatures in the AoE is likely explained by the wave-like nature of the shock wave, combined with the continuous nature of the matrix. When a shock wave in the matrix reaches a grain, the first part of the grain it interacts with will be a vertex. The shock wave will transmit into the grain at this point and begin to travel radially out from this entry point. The shock wave will reach each point along the grain's shock-facing side at different times (depending on the face's value of θ). This radially-travelling shock wave will be bounded by shock-facing sides. Meanwhile, the matrix-shock wave will still be travelling longitudinally past the grain. The grain shock wave will tend to lead the matrix shock wave (owing to its higher impedance), but will also have farther to travel to reach the end of the grain edge facing the shock. There will be an optimum angle θ_c at which the grain shock wave and matrix shock wave propagate along the shockward-facing side at the same speed. If $\theta < \theta_c$ the matrix shock wave propagates along the boundary faster than the grain shock wave and enters the grain; if $\theta = 0°$ the matrix shock wave will reach all points on the face simultaneously. Alternatively, if $\theta > \theta_c$ the grain shock wave will be ahead of the matrix shock wave and the grain shock wave will transmit into the matrix ahead of its main path, distorting the matrix-shock front. At $\theta = \theta_c$, however, the two shock waves will propagate along the boundary at the same speed, and constructively interfere as shocks are reflected/transmitted on both sides simultaneously. This would generate a hot spot along the length of the boundary. This hypothesis predicts that this effect is not related to shear heating, but due to shock speeds and reflections. Additional simulations recreated the results using strengthless materials, which corroborates the hypothesis that this is a relative shock speed phenomenon.

Overall the effect is measurable, but not strong enough to affect the bulk response significantly. The morphology simulations were repeated using a trapezium with a shock-facing side at $\theta \approx \theta_c$ and the bulk results did not appear to change significantly. The grains did appear to exhibit temperature spikes on the shock-facing sides but this was not enough to affect the bulk response significantly.

Angular grains appear to be more likely to experience local temperature spikes along their shock-facing edges. The critical angle, however, depends on the ratio of the shock speed in the grain material to the shock speed in the porous matrix, which makes it difficult to predict what value θ_c will have in a real material—the lunar regolith, for example, which contains many different materials. In addition, as discussed in Chapter 5, the shock wave will not remain planar as it propagates through a multi-grain mixture and this makes predicting macroscale effects difficult because of this. Grains in real powders will also be 3D and, whilst the theory should still be valid, it becomes more challenging to apply it; 3D properties may have unforeseen effects on this optimum angle phenomenon.

6.4.3 Implications for Regolith Studies

Almost all samples of lunar material have come from the Apollo missions, all of which were taken from the regolith to a maximum depth of 3m. This makes the lunar regolith one of humanity's primary sources of material evidence with regards to the Moon and the solar system. Understanding the shock compaction processes experienced by the regolith when the Moon is struck by a meteorite is necessary to interpret the end result correctly. In particular, lunar samples are typically dated using a radiometric method of dating—KAr (potassium-argon) dating—which relies on the presence of ^{40}Ar in a sample (Dalrymple and Lanphere 1971; Turner et al. 1971). KAr dating works on the principle that an isotope of K within a sample decays into ^{40}Ar at a specific rate, which slowly builds up over time. The ratio of un-decayed K to Ar in the sample can be used to infer an age. This

assumes, however, that the sample has always been a closed system and this is typically untrue. The Ar can escape if if the sample melts: this is known as 'degassing'. If a sample has been degassed in the past, this can essentially 'reset' its age. If the complex powder properties of the lunar regolith suggest that it can experience higher temperatures and pressures for lower impact velocities than previously thought, this could have a direct effect on existing interpretations of regolith ages.

There have been few, if any, mesoscale investigations into the compaction of the lunar regolith and the extent of related studies is constrained to those performed by Bland et al. (2014), Davison, Collins, and Bland (2016), and Davison et al. (2017), as discussed in Chapter 5. Their numerical approach presents the potential to analyse the dynamics of lunar regolith impact compaction at an unprecedented level. Most numerical studies of regolith compaction consider it in the broader context of meteorite impacts and parametrise any porous matter as continuous (Yue and Di (2017) or Cintala (1992) for example). The regolith sample responses presented in this chapter exhibit significant heterogeneity within the matrix at the mesoscale. Temperatures and pressures can vary significantly (by as much as 1000 K or 10 GPa) over scales comparable to the DE_{min}. These variations are not reflected in the bulk response of the sample. This suggests that regolith samples used in KAr dating are more likely to be spurious when using smaller grains. Larger grains experience lower temperatures and pressures overall and are thus less likely to have been degassed in the past because of meteorite impacts.

Whilst this work further highlights the applicability of this technique to complex granular structures, it does present several issues that must be considered in any study. The exact porosity of the matrix has a significant effect on the bulk temperatures, and the explicit resolution of pores has a lesser, but still measurable, effect on the bulk temperatures and pressures. Altering grain morphology does not have a significant effect on the bulk response of regolith matter, which suggests that when considering the bulk response ex-

clusively, the circular grain approximation is reasonable. Such an assumption does have a measurable effect on the local material, however. The existence of an optimal angle of a grain face with the shock front to maximise nearby matrix temperatures and pressures has the potential to be exploited within existing meteorites, if combined with the porosity vector from the previous chapter. Lunar meteorites, like NWA 5000, contain angular grains; these are lumps of material that were ejected from the moon in other impact events. By determining the shock direction in a lunar meteorite (using the porosity vector) it might then be possible to determine the critical angle for grains in the meteorite and calibrate the strength of the shock it was exposed to. This could additionally be used to assist dating of lunar meteorites, by highlighting material that should have last been degassed in the meteorite's formation (along the hot-spots). Dating of samples taken from these hot-spot regions could indicate when the meteorite was shocked.

6.5 Conclusions

The simulation technique of resolving a granular material semi-explicitly, introduced by Bland et al. (2014), presents a possible way of examining the impact compaction of complex powders, such as the lunar regolith. The method's efficacy, when applied to a complex powder (with the lunar regolith as a case study), was tested in three ways. First, the effects of modelling a continuous particle size distribution were tested by varying the particle size below which grains were parametrised and comparing the corresponding response of grains in different grain size categories (GSCs) across simulations. Second, density variations in the parametrised material were tested by initialising a base simulation and introducing an increasing amount of explicit pore space, whilst keeping the matrix, and bulk, porosity constant (by introducing additional grains). Third, the sensitivity of the powder's shock response to changes in grain morphology was investigated by examining the effects of three grain shapes: circular, square, and 'realistic'.

The results obtained when varying the size of the smallest, explicitly resolved, grain (DE_{min}) suggest that resolving a granular material semi-explicitly accurately reproduces the bulk response in terms of temperature and pressure. The response of grains in each GSCs, however, were not similarly consistent across simulations. In particular, as the amount of material that is being parametrised (the matrix of fines) increases and its porosity decreases, the explicit grains exhibit lower peak temperatures and pressures. For low matrix volume (high matrix porosities) the matrix temperature appears to be spuriously high compared with the fully explicit simulations.

The temperatures and pressures experienced by GSCs of different size ranges demonstrated a small dependency on grain size. GSCs containing smaller grains exhibited higher temperatures and pressures than GSCs containing larger grains. Whilst the temperatures and pressures of equivalent GSCs weren't necessarily consistent across simulations, the relative variation between GSCs within a simulation did appear consistent. This is evidence that the temperature and pressure results observed when simulating matrix-chondrule mixture compaction (considered in Chapter 5 and Bland et al. (2014), Davison, Collins, and Bland (2016), and Davison et al. (2017)) by resolving the powder semi-explicitly are physical effects.

Varying the density distribution of the matrix, by introducing explicit pores and grains to maintain constant porosity, demonstrated that increasing the matrix's non-uniformity increased the pressures and temperatures experienced by grains in the GSCs of the sample. The mass-weighted bulk temperatures and pressures, however, were consistent across all simulations (with the same V_I), although there did appear to be a slight increase in pressure, and decrease in temperature, with increasing non-uniformity. The mass-weighted bulk porosity also decreased significantly as non-uniformity increased.

The disagreement in shock response between grains in the same GSC in simulations with different DE_{min} was alleviated by introducing matrix non-uniformity. The high matrix

temperatures experienced by samples using a high-porosity matrix were mitigated by introducing explicit pores and decreasing the initial matrix porosity accordingly. This also lowered the final matrix porosity and was able to bring the results of the $DE_{min} = 20$ cppd in line with the explicitly resolved regolith simulation's bulk porosity. The closest overall agreement was found when introducing a sufficient number of explicit pores that the initial matrix porosity was equal to the bulk porosity of the regolith as a whole (50%). The agreement between the semi-explicitly resolved regolith simulation (including pores) and the explicitly resolved regolith simulation suggests that the semi-explicit method is suitable for investigating the shock response of other complex powders, providing the parametrised porosity is engineered to be equal to the bulk porosity.

Grain morphology was not found to exhibit a significant effect on either the bulk shock response, or the response of any grain size categories. Differences, however, were observed at the mesoscale in the matrix surrounding individual grains. These were likely caused by changes in the size of the region affected by the presence of a grain; shock interaction with a square grain will not necessarily extend over the same region of material as a shock interacting with a circular grain of the same size in the same position. These differences between areas-of-effect were investigated further by conducting a detailed mesoscale analysis of the response of angular grains in the semi-explicit method. The angle of the grain face to the shock front was found to have a direct relationship with the temperatures and pressures experienced on its shock-facing side. There is an optimum angle at which the grain generates the highest temperatures and pressures in its surrounding material. This phenomenon does not have enough of an effect to influence the bulk response. The evidence suggests that the 'circular grains' assumption is reasonable in semi-explicit set-ups, even when considering grains as complex, and irregular, as those in the lunar regolith. In addition, the existence of the critical angle may present a method of extending the porosity vector (described in Chapter 5) to meteorites with non-spherical grains.

Chapter 7

Conclusion

7.1 Summary

Among the many aspects of granular materials that can affect their response to shock compaction, mesostructure is one that has rarely been considered or quantified in the shock physics literature. This thesis aimed to document the effects of mesostructure, and related aspects, on the shock response of granular materials at both the macro and mesoscales.

This thesis also attempts to build confidence in the use of mesoscale modelling to simulate complex granular properties that have otherwise been considered too challenging in the past. In particular, the efficacy of semi-explicitly resolving granular matter in simulations was considered. The conclusion that semi-explicitly resolving a granular material can accurately replicate the response of explicitly resolving the granular material establishes confidence in the efficacy of the technique. This would enable future studies to use this technique to study powders that would otherwise have been numerically challenging to simulate, like the lunar regolith, precursor meteorite matter, or terrestrial soil.

The existing, and extensive, work of Borg & Vogler (Borg and Vogler 2008; Borg and Vogler 2009; Borg and Vogler 2013) on granular systems of tungsten carbide was used as the basis to consider the most pure form of mesostructure: the arrangement of granular material. It was found that the average number of grain-grain contacts per grain (the coordination number Z) had a distinct effect on the shock response. In particular, increasing Z artificially stiffened the response of the bulk material. Z is a post-set-up parameter that is calculated after a bed a has been created. This is a unique feature of this work because all previous attempts to quantify arrangement used pre-set-up parameters, which parametrised arrangement by the relative application of their creation algorithms. Borg and Vogler 2009, for example, varied the length of application of their randomisation algorithm, as opposed to considering the arrangements generated post-creation.

Mesostructure has rarely been adequately accounted for in previous investigations; its influence, however, on the shock response of granular matter may be able to explain discrepancies in experimental results presented in previous studies. For example, different values of Z could be used to explain the experimental variations in shock speed seen by Borg and Vogler (2008). The variations observed by Borg and Vogler (2008) are on the same order of magnitude as the variations in U_s observed in this study ($\sim 5 - 10\%$) and significantly greater than their reported uncertainty on U_s ($\sim 0.5\%$). This work suggests that if mesostructure can not be quantified, multiple experiments/simulations should be conducted at a single impact velocity, and their result averaged, to minimise any mesostructure's effect on the shock response.

The response of a granular system with a bimodal particle distribution is a topic that has ramifications in several fields, such as shock chemistry (Guo et al. 2017), composite material structure (Olney et al. 2011), and meteorite formation (Bland et al. 2014). Chapter 5 presents a joint numerical-experimental examination of the shock compaction of an example bimodal system at the mesoscale. The work of Bland et al. (2014) and

Davison, Collins, and Bland (2016) focused on the compaction of primitive meteoritic material and was used as a basis for this study. The present investigation used materials analog to the precursor meteorite matter that consisted of large spherical solidified melt droplets (chondrules) surrounded by a highly porous, extremely fine, powder (matrix). The experiments were performed at the ESRF and generated real-time, X-ray images of the shock compaction. iSALE simulations were performed to replicate the experiments and extract additional information that would be otherwise inaccessible. It was found in both the experiments and simulations that the larger chondrules 'shielded' the matrix in their vicinity, producing a dichotomy in matrix porosity on opposite sides of a chondrule. The matrix shockward of the chondrule was seen to be more compacted than the material in its lee. This effect was identified, and quantified, whilst considering a system of only one chondrule, surrounded by matrix. Additional simulations that used multi-chondrule mixtures also observed the compaction dichotomy, though to a lesser extent, around the majority of grains. The implications of this finding are that it (in conjunction with existing studies of direction dependence in meteorite matrix Forman et al. (2017)) presents a potential new method of decoding the direction, and strength, of low-intensity shocks experienced by a meteorite during its formation.

The method of resolving a granular material semi-explicitly in a simulation was introduced by Bland et al. (2014), and used in Chapter 5. The efficacy of resolving more complex granular materials semi-explicitly that possess a continuous, rather than bimodal, particle size distribution was tested in Chapter 6. The lunar regolith was a motivation for this study. It was found that changing the smallest grain size to be explicitly resolved did not have a significant effect on the bulk shock response of the powder. Alongside the bulk response, grains within the powder were divided into four explicitly resolved grain size categories (GSC) and one category for fines implicitly resolved in the matrix. Each GSC of a simulation was compared with the equivalent GSC in the other simulations. For example, a GSC might consist of grains in the size range 20 cppd to 40 cppd; these were

implicitly resolved in one simulation (part of the matrix GSC), but explicitly resolved in another. The response of grains in individual GSCs within each simulation, however, did not agree unless the parametrised matrix material was engineered to possess the same porosity as the bulk porosity, by introducing explicitly resolved pore space within it.

The effect of varying the morphology of grains when using the semi-explicit method was also investigated. Three shapes were examined: circular grains, square grains and 'realistic' grains that were based on known lunar regolith shape parameters. Changing the morphology of a simulation was not sufficient to affect the bulk, or GSC, responses appreciably. The mesoscale response, however, demonstrated unusual local responses around grains. Additional simulations revealed that there is an optimum angle for a grain's face to have in orientation to an oncoming shock front, which produces temperature and pressure spikes in the matrix material along that face. This result could potentially be combined with the porosity vector introduced in Chapter 5 to extend the porosity vector to the non-spherical grains that are common in lunar meteorites like NWA 5000.

Multiple aspects of the work presented in this thesis has been published. The work presented in Chapter 4 (including the benchmark in Chapter 3) was published in Derrick et al. (2018b). Some aspects of the work presented in Chapter 5 has been published in several different papers including Davison et al. (2017), Rutherford et al. (2017), and Derrick et al. (2018a). Other parts of the work presented in this chapter are also in the process of being peer-reviewed for additional publications.

7.2 Future Work

This work has revealed that the coordination number, Z, has a distinct effect on the shock response of a granular material and future research into its effect would ideally focus on extending Z's effects to 3D. Z can possess a much greater range of values in 3D than 2D and a 3D investigation would enable Z to be used to predict the result in an experimental

study, where Z is calculated pre-shot. Success would validate Z as a means of predicting the shock response of granular matter and even enable it to be added to existing porosity models. In addition, as the arrangements become more complex, a more complex metric for structure may be better suited to capturing the extent of the mesostructural variation: the full fabric tensor (or variation on it), for example.

The existence of the porosity vector C has important implications for the analysis of shock compacted bimodal granular mixtures. In particular, it presents a potential method of decoding the shock direction and magnitude from shocked samples. Future work might involve calculating the porosity vector using analytical data from real meteorite samples and investigating if there is a similar directional trend to that shown here. Once this is complete, these porosity vectors can be compared with features known to be diagnostic of, or influenced by, shock direction, such as the crystallographic preferred orientation of the meteorite's matrix (as found in Forman et al. (2017)). Agreement would suggest they are both diagnostic of the same event and could be used to calibrate the strength of the shock the meteorite experienced.

Chapter 6 demonstrated that resolving granular materials semi-explicitly (as opposed to wholly explicitly or wholly implicitly) is an accurate method for modelling more complex powders than those it has been used for so far. In addition, it demonstrated that there are significant effects local to grains that can be attributed to grain morphology. These results suggest that the effect could be used in lunar meteorites to calibrate the porosity vector. Before that can be done, however, work needs to be performed that investigates the effect of morphology on the shock response when matrix non-uniformity is present (explicit pores). Additionally, this effect has been studied purely in 2D so far and an extension to 3D would assist in making meaningful comparisons with meteorite samples.

Bibliography

Amsden, A. A., H. M. Ruppel, and C. W. Hirt (1980). *SALE: A Simplified ALE computer program for fluid flow at all speeds*. Vol. LA-8095. June, 101p. ISBN: LA-8095.

Anderson, Charles E. (1987). "An overview of the theory of hydrocodes". In: *International Journal of Impact Engineering* 5.1. Hypervelocity Impact Proceedings of the 1986 Symposium, pp. 33 –59. ISSN: 0734-743X. DOI: https://doi.org/10.1016/0734-743X(87)90029-7.

Beitz, E. et al. (2013). "Experiments on the consolidation of chondrites and the formation of dense rims around chondrules". In: *Icarus* 225.1, pp. 558–569. ISSN: 00191035. DOI: 10.1016/j.icarus.2013.04.028.

Benson, D J et al. (1994). "An analysis by direct numerical simulation of the effects of particle morphology on the shock compaction of copper powder". In: *Modelling and Simulation in Materials Science and Engineering* 2.3A, pp. 535–550. ISSN: 0965-0393. DOI: 10.1088/0965-0393/2/3A/008.

Benson, David J. and W. J. Nellis (1994). "Dynamic compaction of copper powder: Computation and experiment". In: *Applied Physics Letters* 65.4, p. 418. ISSN: 00036951. DOI: 10.1063/1.112319.

Bland, P A et al. (2014). "Pressure-temperature evolution of primordial solar system solids during impact-induced compaction." In: *Nat Commun* 5, p. 5451. DOI: 10.1038/ncomms6451.

Blott, Simon J. and Kenneth Pye (2012). "Particle size scales and classification of sediment types based on particle size distributions: Review and recommended procedures". In: *Sedimentology* 59.7, pp. 2071–2096. ISSN: 00370746. DOI: 10.1111/j.1365-3091.2012.01335.x.

Borg, John P and Tracy J Vogler (2007). "Mesoscale Calculations of Shock Loaded Granular Ceramics". In: *AIP Conference Proceedings* 955.1, pp. 227–+. ISSN: 0094-243X.

— (2008). "Mesoscale calculations of the dynamic behavior of a granular ceramic". In: *International Journal of Solids and Structures* 45.6, pp. 1676–1696. ISSN: 0020-7683. DOI: http://dx.doi.org/10.1016/j.ijsolstr.2007.10.027.

— (2009). "Aspects of simulating the dynamic compaction of a granular ceramic". In: *Modelling and Simulation in Materials Science and Engineering* 17.4, p. 45003. DOI: 10.1088/0965-0393/17/4/045003.

— (2013). "Rapid compaction of granular material: characterizing two- and three-dimensional mesoscale simulations". In: *Shock Waves* 23.2, pp. 153–176. ISSN: 0938-1287. DOI: 10.1007/s00193-012-0423-6.

Borg, John P et al. (2006). "Computational simulations of the dynamic compaction of porous media". In: *International Journal of Impact Engineering* 33.112, pp. 109–118. ISSN: 0734-743X. DOI: http://dx.doi.org/10.1016/j.ijimpeng.2006.09.045.

Borg, John P. et al. (2009). "a Review of Mesoscale Simulations of Granular Materials". In: 1331.2009, pp. 1331–1336. DOI: 10.1063/1.3295053.

Brace, W. F. (1981). "The effect of size on mechanical properties of rock". In: *Geophysical Research Letters* 8.7, pp. 651–652. DOI: 10.1029/GL008i007p00651.

Chapman, David J. (2010). "Shock-compression of porous materials and diagnostic development". PhD thesis. University of Cambridge.

Cheng, A.F. et al. (2016). "Asteroid Impact & Deflection Assessment mission: Kinetic impactor". In: *Planetary and Space Science* 121, pp. 27–35. ISSN: 00320633. DOI: 10.1016/j.pss.2015.12.004.

Cintala, Mark J. (1992). "Impactinduced thermal effects in the lunar and Mercurian regoliths". In: *Journal of Geophysical Research: Planets* 97.E1, pp. 947–973. DOI: 10.1029/91JE02207.

Collins, G. S. et al. (2011). "Improvements to the epsilon-alpha porous compaction model for simulating impacts into high-porosity solar system objects". In: *International Journal of Impact Engineering* 38.6, pp. 434–439. ISSN: 0734743X. DOI: 10.1016/j.ijimpeng.2010.10.013.

Collins, Gareth S., H. Jay Melosh, and Boris A. Ivanov (2004). "Modeling damage and deformation in impact simulations". In: *Meteoritics & Planetary Science* 39.2, pp. 217–231. ISSN: 1945-5100. DOI: 10.1111/j.1945-5100.2004.tb00337.x.

Collins, Gareth S. et al. (2016). *iSALE-Dellen manual*. DOI: 10.6084/m9.figshare.3473690.v2.

Dalrymple, G. Brent and Marvin A. Lanphere (1971). "40Ar/39Ar technique of KAr dating: a comparison with the conventional technique". In: *Earth and Planetary Science Letters* 12.3, pp. 300 –308. ISSN: 0012-821X. DOI: https://doi.org/10.1016/0012-821X(71)90214-7.

Dandekar, Dattatraya P. and Dennis E. Grady (2002). "Shock Equation of State and Dynamic Strength of Tungsten Carbide". In: *AIP Conference Proceedings* 620.1, pp. 783–786. DOI: 10.1063/1.1483654. eprint: http://aip.scitation.org/doi/pdf/10.1063/1.1483654.

Davison, Thomas M., Gareth S. Collins, and Philip A. Bland (2016). "Mesoscale Modeling of Impact Compaction of Primitive Solar System Solids". In: *The Astrophysical Journal* 821.1, p. 68. ISSN: 1538-4357. DOI: 10.3847/0004-637X/821/1/68.

Davison, Thomas M. et al. (2017). "Impact-induced compaction of primitive solar system solids: The need for mesoscale modelling and experiments". In: *Procedia Engineering* 204.Supplement C. 14th Hypervelocity Impact Symposium 2017, HVIS2017, 24-28 April 2017, Canterbury, Kent, UK, pp. 405 –412. ISSN: 1877-7058. DOI: https://doi.org/10.1016/j.proeng.2017.09.801.

De Cola, F. et al. (2018). "Effect of Particle Morphology, Compaction, and Confinement on the High Strain Rate Behavior of Sand". In: *Experimental Mechanics* 58.2, pp. 223–242. ISSN: 1741-2765. DOI: 10.1007/s11340-017-0331-x.

Derrick, James G. et al. (2018a). "Interrogating heterogeneous compaction of analogue materials at the mesoscale through numerical modeling and experiments". In: *AIP Conference Proceedings* 1979.1, p. 110004. DOI: 10.1063/1.5044923.

Derrick, James Gordon et al. (2018b). "Mesoscale simulations of shock compaction of a granular ceramic: effects of mesostructure and mixed-cell strength treatment". In: *Modelling and Simulation in Materials Science and Engineering*.

Di, Kaichang et al. (2016). "Lunar regolith thickness determination from 3D morphology of small fresh craters". In: *Icarus* 267, pp. 12 –23. ISSN: 0019-1035. DOI: https://doi.org/10.1016/j.icarus.2015.12.013.

Dodge, Yadolah (2008). *The Concise Encyclopedia of Statistics*. New York, NY: Springer New York, pp. 502–505. ISBN: 978-0-387-32833-1. DOI: 10.1007/978-0-387-32833-1_379.

Drucker, D. C. and W. Prager (1952). "Soil Mechanics and Plastic Analysis or Limit Design". In: *Quarterly of Applied Mathematics* 10, pp. 157–165. DOI: 10.2307/43633942.

Dwivedi, S. K., L. Pei, and R. Teeter (2015). "Two-dimensional mesoscale simulations of shock response of dry sand". In: *Journal of Applied Physics* 117.8, p. 085902. ISSN: 0021-8979. DOI: 10.1063/1.4913479.

Eakins, D E and N N Thadhani (2008). "Mesoscale simulation of the configuration-dependent shock-compression response of Ni + Al powder mixtures". In: *Acta Materialia* 56.7, pp. 1496–1510. ISSN: 1359-6454. DOI: http://dx.doi.org/10.1016/j.actamat.2007.12.009.

Fang, Zhigang, Burton R. Patterson, and Malcolm E. Turner (1993). "Modeling particle size distributions by the Weibull distribution function". In: *Materials Characterization* 31.3, pp. 177 –182. ISSN: 1044-5803. DOI: https://doi.org/10.1016/1044-5803(93)90058-4.

Fenton, Gregg, Blaine Asay, and Devon Dalton (2015). "Impact Compaction of a Granular Material". In: *Procedia Engineering* 103, pp. 121–128. ISSN: 1877-7058. DOI: 10.1016/j.proeng.2015.04.017.

Forman, L.V. et al. (2017). "Defining the mechanism for compaction of the CV chondrite parent body". In: *Geology* 45.6, p. 559. DOI: 10.1130/G38864.1.

Fredenburg, D. A. and N. N. Thadhani (2013). "Predicting the shock compression response of heterogeneous powder mixtures". In: *Journal of Applied Physics* 113.22, p. 223513. DOI: 10.1063/1.4810929. eprint: http://dx.doi.org/10.1063/1.4810929.

Fredenburg, D. A. et al. (2014). "The influence of morphology on the low- and high-strain-rate compaction response of CeO2 powders". In: *Journal of Applied Physics* 115.12, p. 123511. ISSN: 10897550. DOI: 10.1063/1.4868356.

Fredenburg, D.A., D. Dennis-Koller, and D. Dattelbaum (2012). "Configurational affects on the compaction response of CeO_2 powders". In: *EPJ Web of Conferences* 26, p. 02011. ISSN: 2100-014X. DOI: 10.1051/epjconf/20122602011.

Freim, John, J. McKittrick, and W. J. Nellis (1995). "Densification behavior of dynamically shock compacted AI2O3/ZrO2powders synthesized through rapid solidification". In: *Metallurgical and Materials Transactions A* 26.10, pp. 2503–2509. ISSN: 10735623. DOI: 10.1007/BF02669408.

Güldemeister, Nicole et al. (2013). "Propagation of impact-induced shock waves in porous sandstone using mesoscale modeling". In: *Meteoritics & Planetary Science* 48.1, pp. 115–133. ISSN: 1945-5100. DOI: 10.1111/j.1945-5100.2012.01430.x.

Guo, Jun et al. (2017). "Reaction Behavior of Polytetrafluoroethylene/Al Granular Composites Subjected to Planar Shock Wave". In: *Propellants, Explosives, Pyrotechnics* 42.3, pp. 300–307. ISSN: 15214087. DOI: 10.1002/prep.201600115.

H., Hewins Roger and Zanda Brigitte (2012). "Chondrules: Precursors and interactions with the nebular gas". In: *Meteoritics & Planetary Science* 47.7, pp. 1120–1138. DOI: 10.1111/j.1945-5100.2012.01376.x.

H. Tillotson, J (1962). "Metallic Equations of State for Hypervelocity Impact". In: GA-3216, p. 141.

Han, Shin-Chan et al. (2014). "Global characteristics of porosity and density stratification within the lunar crust from GRAIL gravity and Lunar Orbiter Laser Altimeter topography data". In: *Geophysical Research Letters* 41.6, pp. 1882–1889. ISSN: 1944-8007. DOI: 10.1002/2014GL059378.

Heiken, G., D. Vaniman, and B.M. French (1991). *Lunar Sourcebook: A User's Guide to the Moon*. LPI contribution. Cambridge University Press.

Heiken, Geant H, David S McKay, and R.W Brown (1974). "Lunar deposits of possible pyroclastic origin". In: *Geochimica et Cosmochimica Acta* 38.11, pp. 1703 –1718. ISSN: 0016-7037. DOI: https://doi.org/10.1016/0016-7037(74)90187-2.

Herbold, E. B. et al. (2008). "Particle size effect on strength, failure, and shock behavior in polytetrafluoroethylene-Al-W granular composite materials". In: *Journal of Applied Physics* 104.10, p. 103903. ISSN: 00218979. DOI: 10.1063/1.3000631. arXiv: 0806.1775.

Ivanov, B. A. (2005). "Numerical Modeling of the Largest Terrestrial Meteorite Craters". In: *Solar System Research* 39.5, pp. 381–409. ISSN: 1608-3423. DOI: 10.1007/s11208-005-0051-0.

Ivanov, B. A., D. Deniem, and G. Neukum (1997). "Implementation of dynamic strength models into 2D hydrocodes: Applications for atmospheric breakup and impact cratering". In: *International Journal of Impact Engineering* 20.1. Hypervelocity Impact Proceedings of the 1996 Symposium, pp. 411 –430. ISSN: 0734-743X. DOI: https://doi.org/10.1016/S0734-743X(97)87511-2.

Ivanov, B.A., H.J. Melosh, and E. Pierazzo (2010). *Basin-forming impacts: Reconnaissance modeling*. DOI: 10.1130/SPE465.

Johnson, Gordon R and William H Cook (1985). "Fracture characteristics of three metals subjected to various strains, strain rates, temperatures and pressures". In: *Engineering Fracture Mechanics* 21.1, pp. 31–48. ISSN: 0013-7944. DOI: http://dx.doi.org/10.1016/0013-7944(85)90052-9.

Kerley, Gerald I (2013). *The Linear US-uP Relation in Shock-Wave Physics*. Tech. rep. arXiv:1306.6916. KTS06-1.

Kowitz, A et al. (2013). "Diaplectic quartz glass and SiO2 melt experimentally generated at only 5 GPa shock pressure in porous sandstone: Laboratory observations and mesoscale numerical modeling". In: *Earth and Planetary Science Letters* 384.0, pp. 17–26. ISSN: 0012-821X. DOI: http://dx.doi.org/10.1016/j.epsl.2013.09.021.

Kruyt, N.P. (2012). "Micromechanical study of fabric evolution in quasi-static deformation of granular materials". In: *Mechanics of Materials* 44, pp. 120–129. ISSN: 01676636. DOI: 10.1016/j.mechmat.2011.07.008.

Lammi, Christopher J. and Tracy J. Vogler (2012). "Mesoscale simulations of granular materials with peridynamics". In: *AIP Conference Proceedings* 1426.1, pp. 1467–1470. ISSN: 0094243X. DOI: 10.1063/1.3686559.

Li, X, H S Yu, and X S Li (2009). "Macromicro relations in granular mechanics". In: *International Journal of Solids and Structures* 46.2526, pp. 4331–4341. ISSN: 0020-7683. DOI: http://dx.doi.org/10.1016/j.ijsolstr.2009.08.018.

Li, Xia and Xiang-Song Li (2009). "Micro-Macro Quantification of the Internal Structure of Granular Materials". In: *Journal of Engineering Mechanics* 135.7, pp. 641–656. ISSN: 0733-9399. DOI: 10.1061/(ASCE)0733-9399(2009)135:7(641).

Lomov, Ilya. et al. (2009). "Mesoscale Simulations of Powder Compaction". In: *AIP Conference Proceedings* 1195.1, pp. 1327–1330. DOI: http://dx.doi.org/10.1063/1.3295052.

Lyzenga, Gregory A., Thomas J. Ahrens, and Arthur C. Mitchell (1983). "Shock temperatures of SiO2 and their geophysical implications". In: *Journal of Geophysical Research: Solid Earth* 88.B3, pp. 2431–2444. DOI: 10.1029/JB088iB03p02431. eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/JB088iB03p02431.

Marchi, S et al. (2014). "Widespread mixing and burial of Earth's Hadean crust by asteroid impacts". In: *Nature* 511, pp. 578–582.

Marsh, Stanley P. (1980). *LASL shock Hugoniot data / Stanley P. Marsh, editor*. English. University of California Press Berkeley [etc.] ; London, xix, 658p : ISBN: 0520040082.

McCoy, C. A. et al. (2016). "Shock-wave equation-of-state measurements in fused silica up to 1600GPa". In: *Journal of Applied Physics* 119.21, p. 215901. DOI: 10.1063/1.4952975.

McGlaun, J M, S L Thompson, and M G Elrick (1990). "CTH: A three-dimensional shock wave physics code". In: *International Journal of Impact Engineering* 10.14, pp. 351–360. ISSN: 0734-743X. DOI: http://dx.doi.org/10.1016/0734-743X(90)90071-3.

Melosh, H J (1989). *Impact cratering: A geologic process*. Vol. 49. 14-15. ISBN: 0195042840. DOI: 10.1103/PhysRevB.46.10411.

Melosh, H. J., E. V. Ryan, and E. Asphaug (1992). "Dynamic fragmentation in impacts: Hydrocode simulation of laboratory impacts". In: *Journal of Geophysical Research: Planets* 97.E9, pp. 14735–14759. ISSN: 2156-2202. DOI: 10.1029/92JE01632.

Mises, Richard von (1913). "Mechanik der festen Korper im Plastisch deformabllem Zustand". In: *Nachr Ges Wissenschaft Göttingen* 1913.4, pp. 582–592.

Nakamura, Tomoki, Kazushige Tomeoka, and Hiroshi Takeda (1992). "Shock effects of the Leoville CV carbonaceous chondrite: a transmission electron microscope study". In: *Earth and Planetary Science Letters* 114.1, pp. 159 –170. ISSN: 0012-821X. DOI: https://doi.org/10.1016/0012-821X(92)90158-R.

Neal, W D, D J Chapman, and W G Proud (2012). "Shock-wave stability in quasi-monodisperse granular materials". In: *European Physical Journal-Applied Physics* 57 (3), p. 31001. DOI: epjap/2011110054.

Neal, William David (2012). "Department of Physics The Role of Particle Size in the Shock Compaction of Brittle Granular Materials". PhD thesis.

Nesterenko, Vitali F. (2001). *Dynamics of Heterogeneous Materials*. New York, NY: Springer New York, p. 510. ISBN: 0387952667. DOI: 10.1007/978-1-4757-3524-6.

Ohnaka, Mitiyasu (1995). "A shear failure strength law of rock in the brittle-plastic transition regime". In: *Geophysical Research Letters* 22.1, pp. 25–28. DOI: 10.1029/94GL02791.

Olney, K. L. et al. (2011). "Role of material properties and mesostructure on dynamic deformation and shear instability in Al-W granular composites". In: *Journal of Applied Physics* 110.11, p. 114908. ISSN: 00218979. DOI: 10.1063/1.3665644.

Ormel, C. W., J. N. Cuzzi, and A. G. G. M. Tielens (2008). "CoAccretion of Chondrules and Dust in the Solar Nebula". In: *The Astrophysical Journal* 679.2, pp. 1588–1610. ISSN: 0004-637X. DOI: 10.1086/587836.

P. de Gennes (1999). "Granular Matter: A Tentative View". In: *More Things in Heaven and Earth: A Celebration of Physics at the Millennium*. Ed. by Benjamin Bederson. New York, NY: Springer New York, pp. 629–643. ISBN: 978-1-4612-1512-7. DOI: 10.1007/978-1-4612-1512-7_41.

Perry, J. I. et al. (2015). "Behaviour of moist and saturated sand during shock and release". In: *Applied Physics Letters* 107.17, p. 174102. ISSN: 0003-6951. DOI: 10.1063/1.4934689.

Perry, J. I. et al. (2016). "The significance of grain morphology, moisture, and strain rate on the rapid compaction of silica sands". In: *Applied Physics Letters* 109.17, p. 174103. DOI: 10.1063/1.4966178.

Petel, Oren E. and Francois X. Jetté (2010). "Comparison of methods for calculating the shock hugoniot of mixtures". In: *Shock Waves* 20.1, pp. 73–83. ISSN: 1432-2153. DOI: 10.1007/s00193-009-0230-x.

Rutherford, Michael E. (2017). "Studies of shock compression at the mesoscale with synchrotron X-ray phase-contrast radiography". PhD thesis. Imperial College London.

Rutherford, Michael E et al. (2017). "Probing the early stages of shock-induced chondritic meteorite formation at the mesoscale". In: *Scientific Reports* 7.November 2016, p. 45206. ISSN: 2045-2322. DOI: 10.1038/srep45206.

S, S. P. A. V. and S. Basu (2015). "Numerical modelling of shock-induced chemical reactions (SICR) in reactive powder mixtures using smoothed particle hydrodynamics (SPH)". In: *Modelling Simul. Mater. Sci. Eng.* 23.7, 075005, p. 075005. DOI: 10.1088/0965-0393/23/7/075005.

Schulte, Peter et al. (2010). "The Chicxulub Asteroid Impact and Mass Extinction at the Cretaceous-Paleogene Boundary". In: *Science* 327.5970, pp. 1214–1218. ISSN: 0036-8075. DOI: 10.1126/science.1177265.

Simon, Franz and Gunther Glatzel (1929). "Bemerkungen zur Schmelzdruckkurve". In: *Zeitschrift fr anorganische und allgemeine Chemie* 178.1, pp. 309–316. DOI: 10.1002/zaac.19291780123.

Spearman, C. (1904). "The Proof and Measurement of Association Between Two Things". In: *American Journal of Psychology* 15, pp. 88–103.

Stöffler, Dieter, Klaus Keil, and Scott Edward R.D (1991). "Shock metamorphism of ordinary chondrites". In: *Geochimica et Cosmochimica Acta* 55.12, pp. 3845–3867. ISSN: 00167037. DOI: 10.1016/0016-7037(91)90078-J.

Summers, Randall M et al. (1997). "Recent progress in ALEGRA development and applications to ballistic impacts". In: *International Journal of Impact Engineering* 20.6-10, pp. 779–788.

Sun, Qicheng, Guangqian Wang, and Kaiheng Hu (2009). "Some open problems in granular matter mechanics". In: *Progress in Natural Science* 19.5, pp. 523–529. ISSN: 1002-0071. DOI: http://dx.doi.org/10.1016/j.pnsc.2008.06.023.

Thompson, S.L. and H.S. Lauson (1974). *Improvements in the CHART D radiation-hydrodynamic code III: revised analytic equations of state*. Tech. rep.

Turner, G. et al. (1971). "40Ar-39Ar ages and cosmic ray exposure ages of Apollo 14 samples". In: *Earth and Planetary Science Letters* 12.1, pp. 19 –35. ISSN: 0012-821X. DOI: https://doi.org/10.1016/0012-821X(71)90051-3.

Vogler, T. J. and J. D. Clayton (2008). "Heterogeneous deformation and spall of an extruded tungsten alloy: plate impact experiments and crystal plasticity modeling". In: *Journal of the Mechanics and Physics of Solids* 56.2, pp. 297–335. ISSN: 00225096. DOI: 10.1016/j.jmps.2007.06.013.

Vogler, T J, M Y Lee, and D E Grady (2007). "Static and dynamic compaction of ceramic powders". In: *International Journal of Solids and Structures* 44.2, pp. 636–658. ISSN: 0020-7683. DOI: http://dx.doi.org/10.1016/j.ijsolstr.2006.05.001.

Vogler, Tracy J. (2015). "Shock Wave Perturbation Decay in Granular Materials". In: *Journal of Dynamic Behavior of Materials* 1.4, pp. 370–387. ISSN: 2199-7454. DOI: 10.1007/s40870-015-0033-3.

VonNeumann, J. and R. D. Richtmyer (1950). "A Method for the Numerical Calculation of Hydrodynamic Shocks". In: *Journal of Applied Physics* 21.3, pp. 232–237. DOI: 10.1063/1.1699639.

Wackerle, Jerry (1962). "Shock-Wave Compression of Quartz". In: *Journal of Applied Physics* 33.3, p. 922. ISSN: 00218979. DOI: 10.1063/1.1777192.

Wieczorek, Mark A. et al. (2013). "The Crust of the Moon as Seen by GRAIL". In: *Science* 339.6120, pp. 671–675. ISSN: 0036-8075. DOI: 10.1126/science.1231530.

Williamson, R. L. (1990). "Parametric studies of dynamic powder consolidation using a particle-level numerical model". In: *Journal of Applied Physics* 68.3, pp. 1287–1296. ISSN: 00218979. DOI: 10.1063/1.346730.

Wünnemann, K, G S Collins, and H J Melosh (2006). "A strain-based porosity model for use in hydrocode simulations of impacts and implications for transient crater growth in porous targets". In: *Icarus* 180.2, pp. 514–527. ISSN: 0019-1035. DOI: http://dx.doi.org/10.1016/j.icarus.2005.10.013.

Yue, Zongyu and Kaichang Di (2017). "Hydrocode simulation of the impact melt layer distribution underneath Xiuyan Crater, China". In: *Journal of Earth Science* 28.1, pp. 180–186. ISSN: 1867-111X. DOI: 10.1007/s12583-017-0741-9.

Zheng, Xiaojing and Dengming Wang (2010). "Multiscale mechanical behaviors in discrete materials: a review". In: *Acta Mechanica Solida Sinica* 23.6, pp. 579 –591. ISSN: 0894-9166. DOI: https://doi.org/10.1016/S0894-9166(11)60005-0.

Zhu, Huaxiang, Franois Nicot, and Flix Darve (2016). "Meso-structure organization in two-dimensional granular materials along biaxial loading path". In: *International Journal of Solids and Structures* 96, pp. 25 –37. ISSN: 0020-7683. DOI: https://doi.org/10.1016/j.ijsolstr.2016.06.025.

Appendix A

Chapter 4: additional material

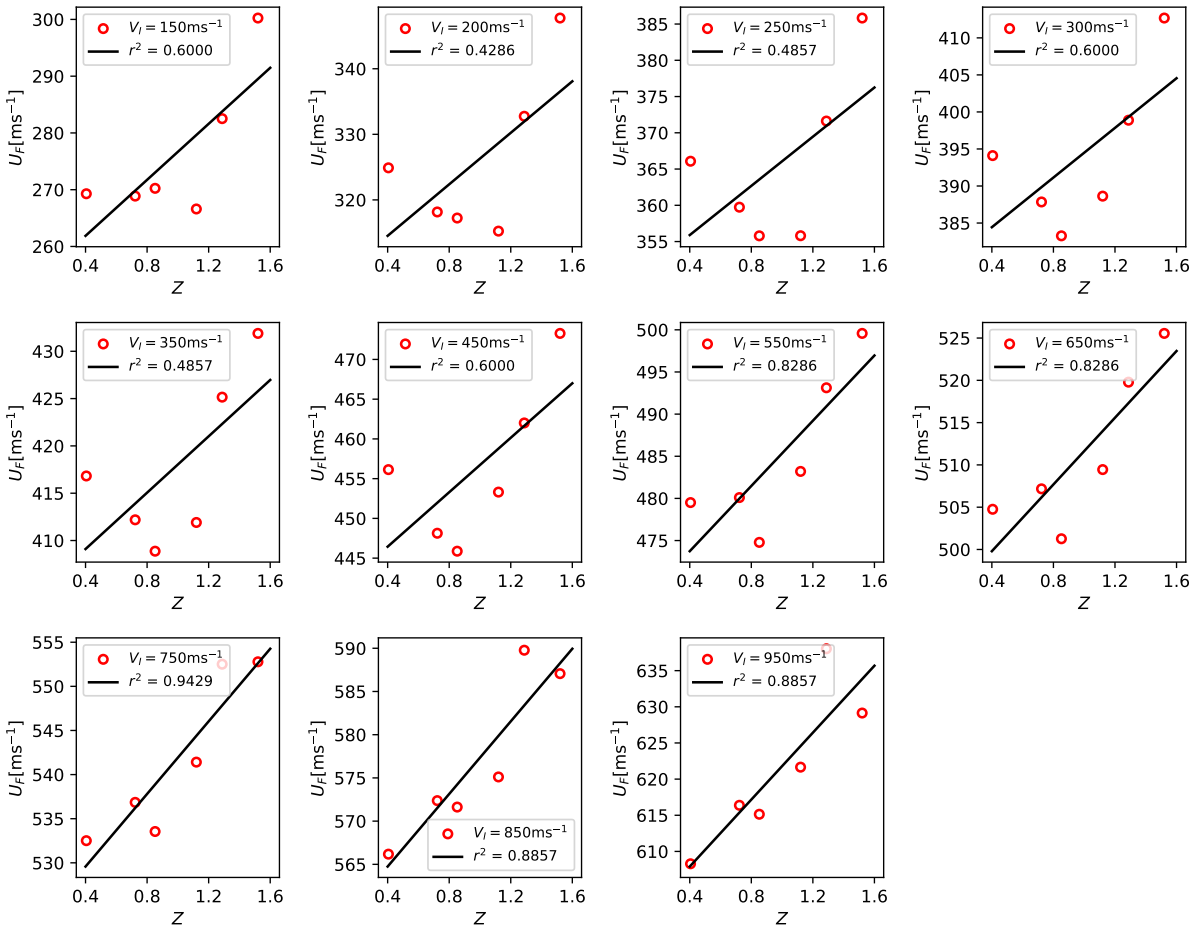


Figure A.1: Correlation of $U_F - Z$ for all impact velocities from Chapter 4 (summarised in Table 4.2)

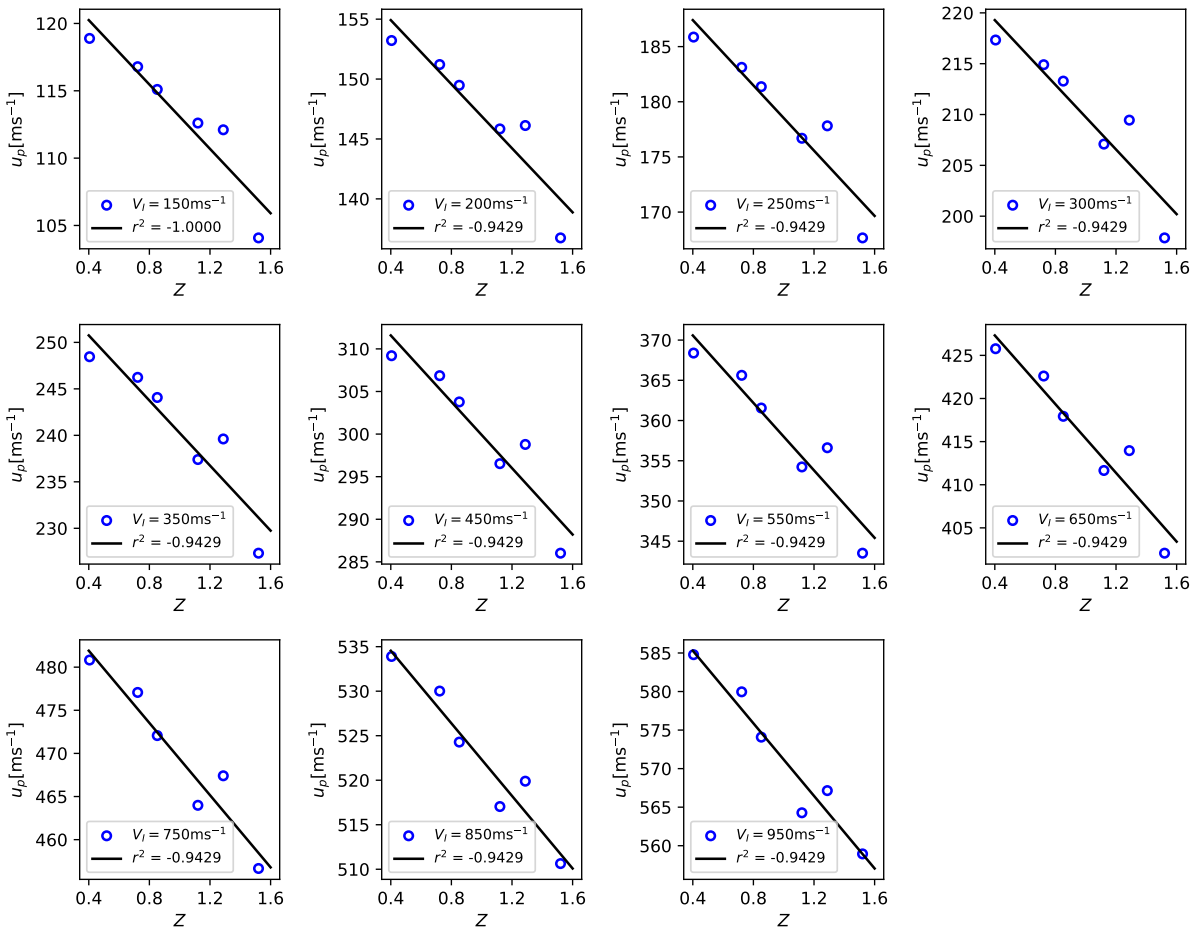


Figure A.2: Correlation of $u_p - Z$ for all impact velocities from Chapter 4 (summarised in Table 4.2)

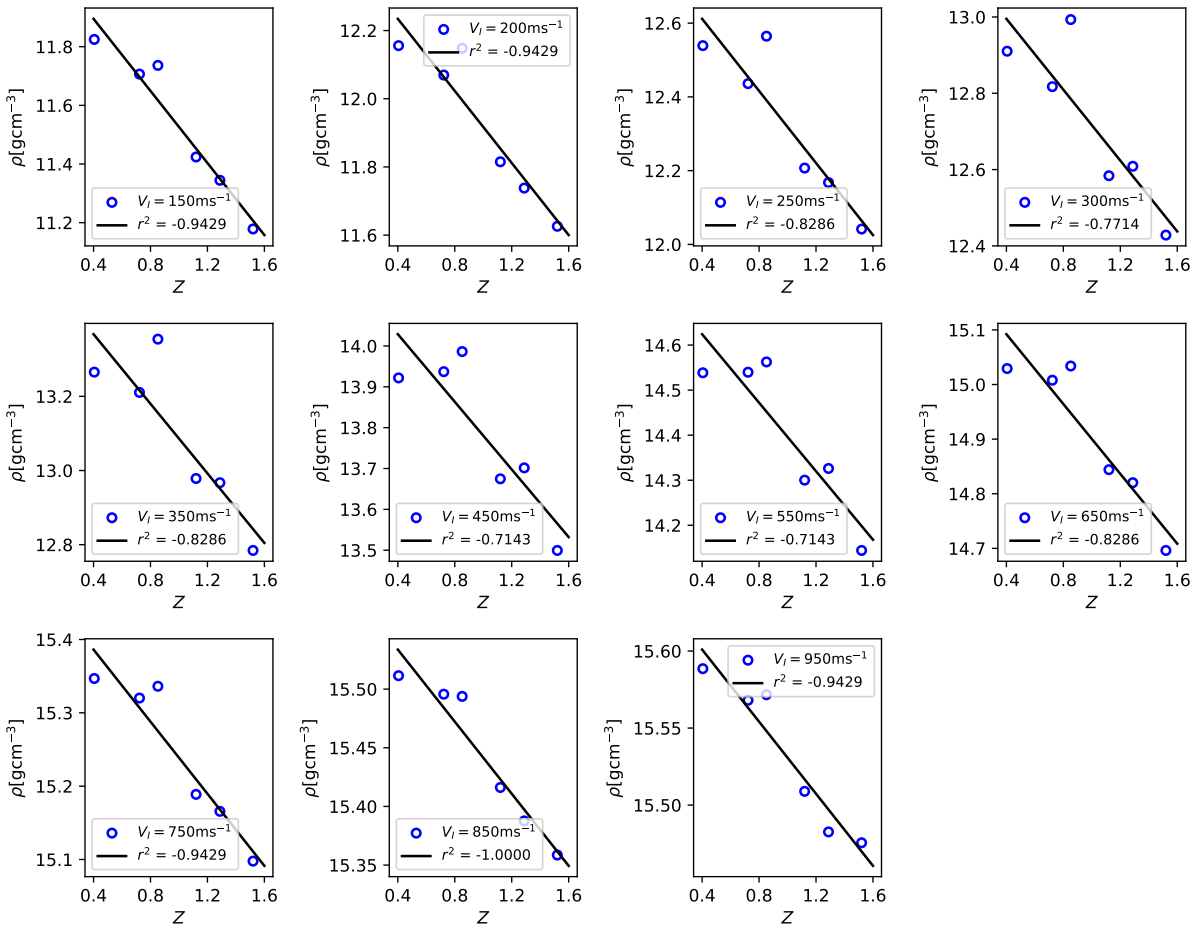


Figure A.3: Correlation of $\rho - Z$ for all impact velocities from Chapter 4 (summarised in Table 4.2)

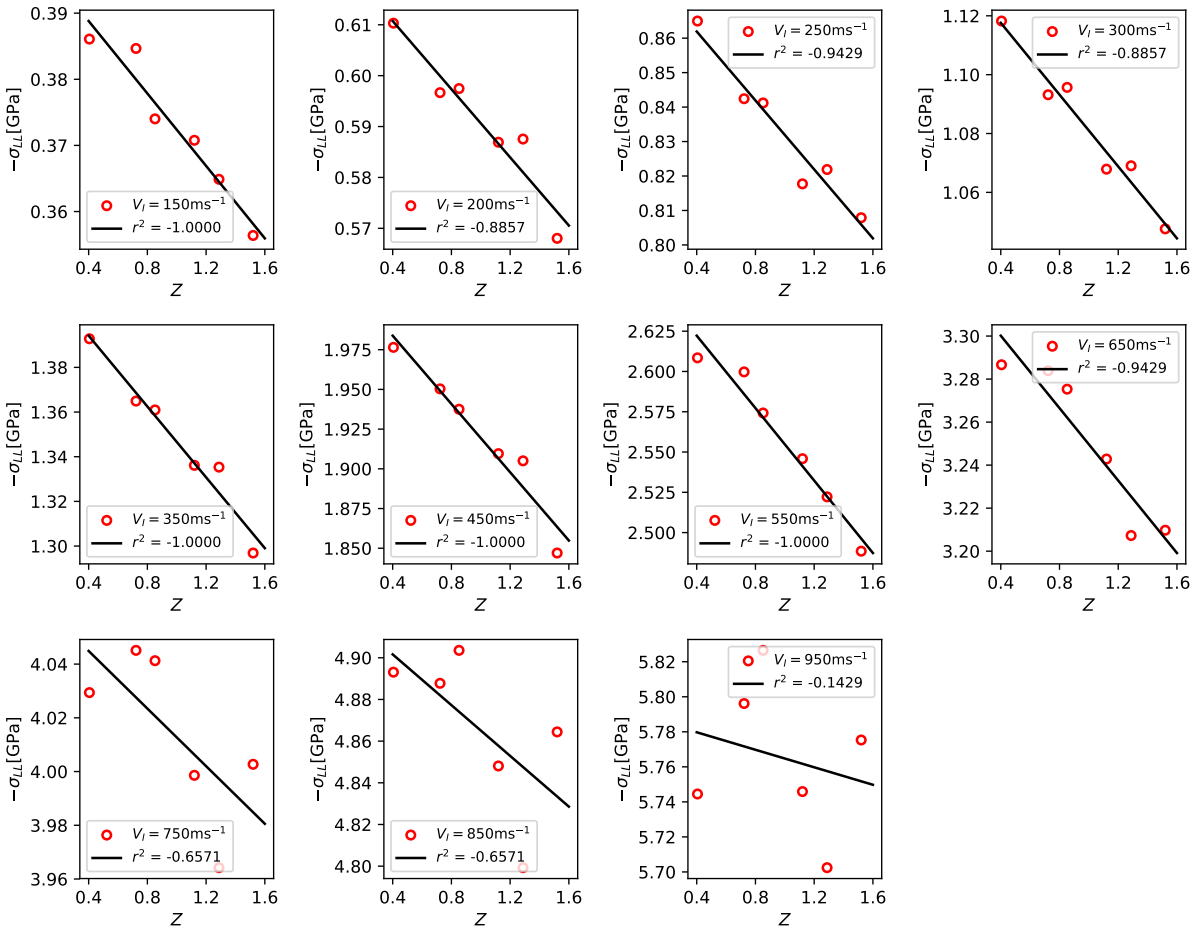


Figure A.4: Correlation of $-\sigma_{LL} - Z$ for all impact velocities from Chapter 4 (summarised in Table 4.2)

Appendix B

Chapter 5: additional material

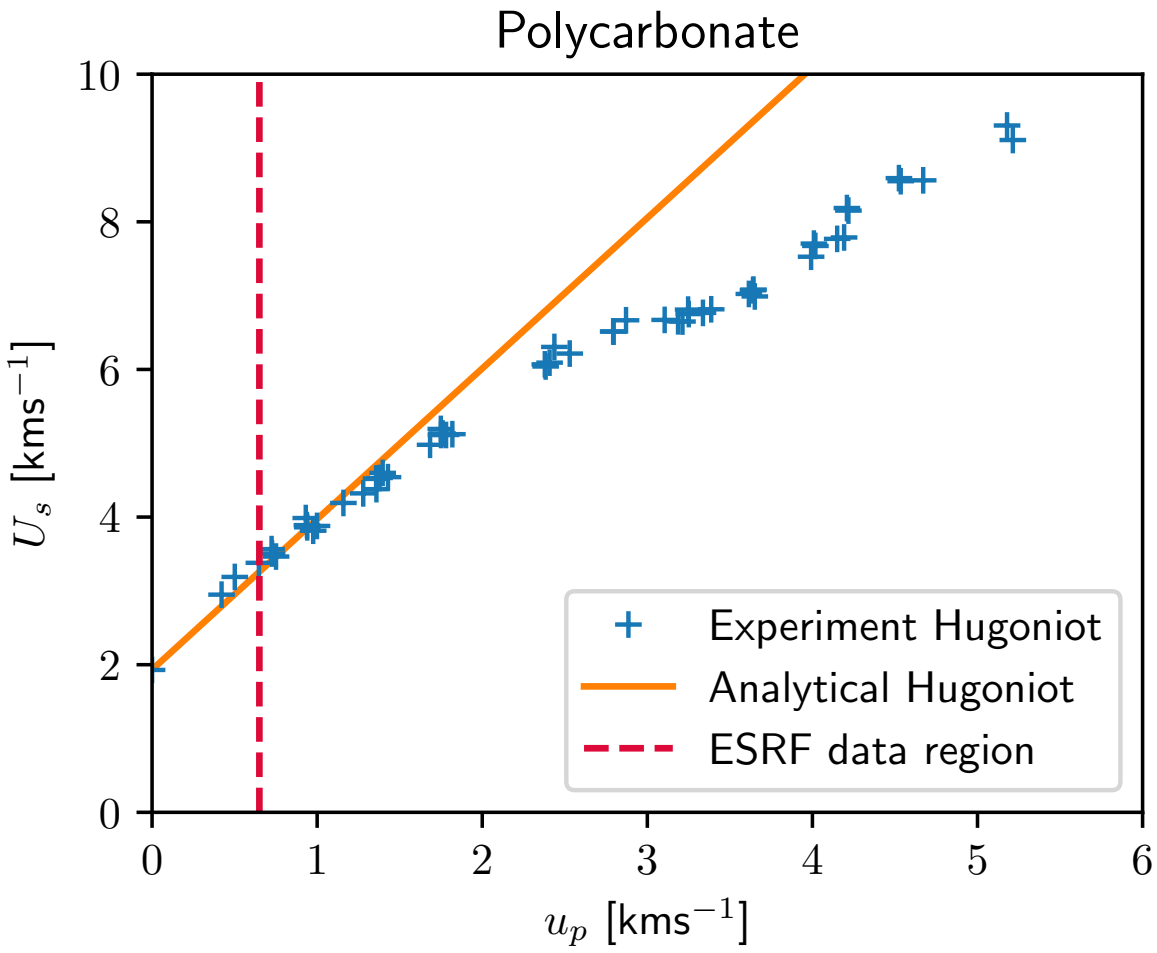


Figure B.1: Direct comparison of the modelled polycarbonate Hugoniot with experimental data. The 'Ballpark' region at which the ESRF experiments occur is illustrated by a dashed line. The experimental data is taken from (Marsh 1980)

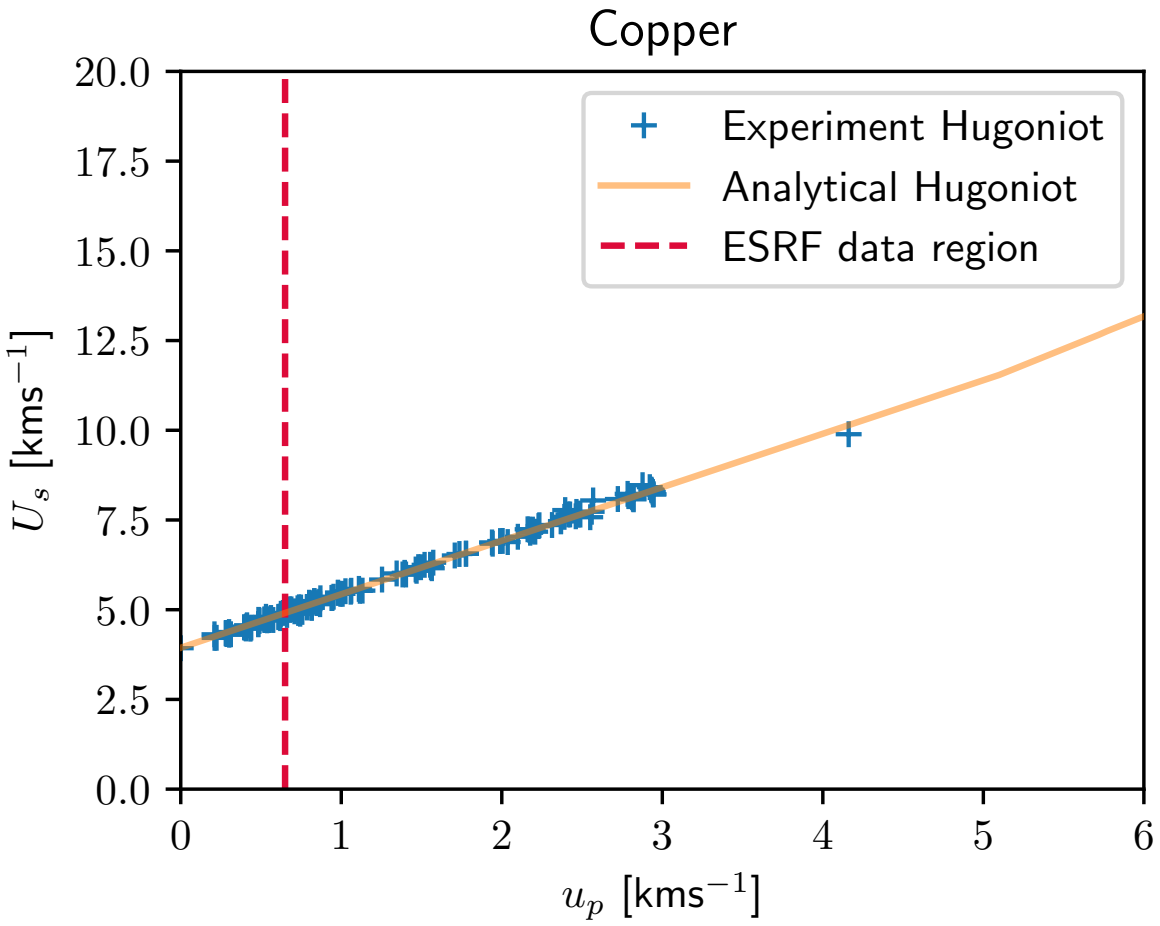


Figure B.2: Direct comparison of the modelled copper Hugoniot with experimental data. The 'Ballpark' region at which the ESRF experiments occur is illustrated by a dashed line. The experimental data is taken from Marsh (1980)

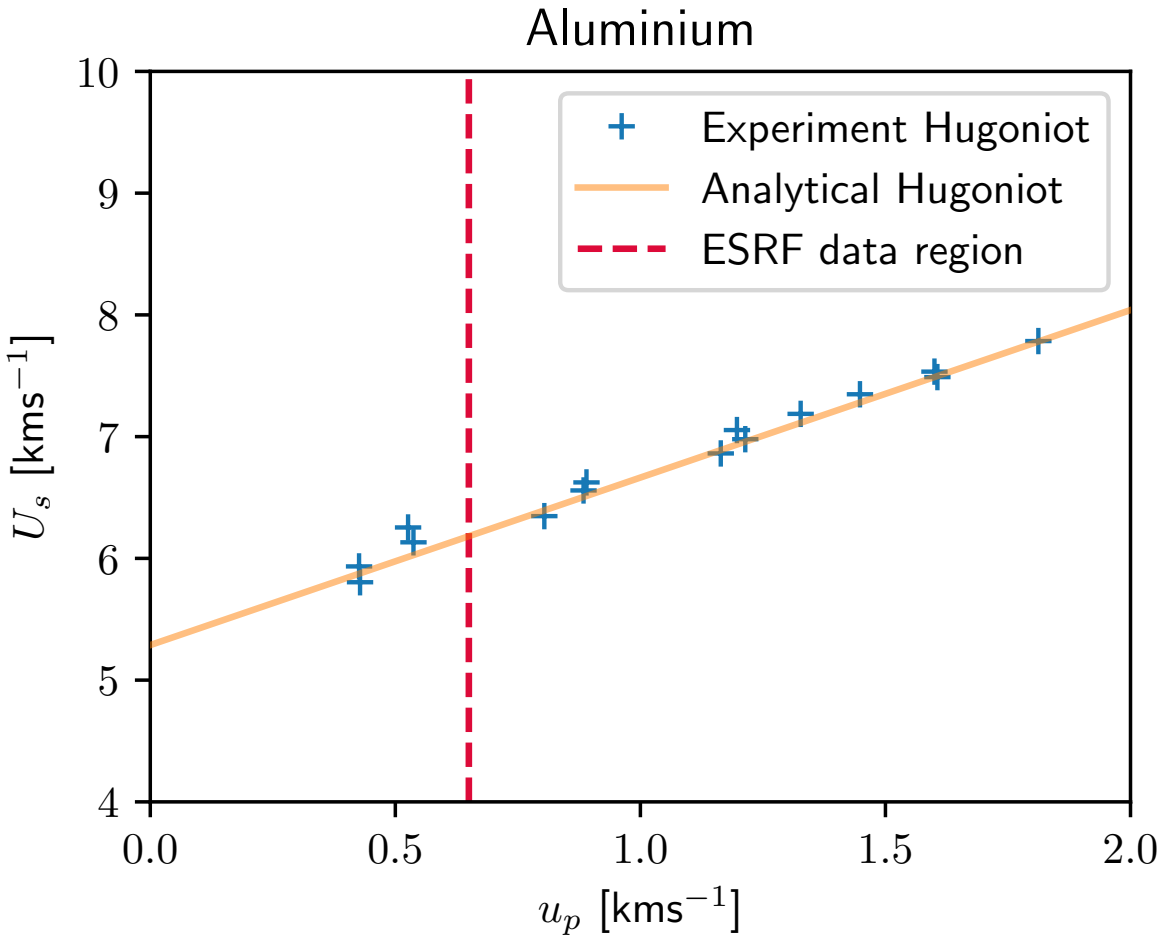


Figure B.3: Direct comparison of the modelled aluminium Hugoniot with experimental data. The 'Ballpark' region at which the ESRF experiments occur is illustrated by a dashed line. The experimental data is taken from Marsh (1980)