



Kent Academic Repository

Lightwing, Andrew (2011) *Catastrophic disruption of layered ice-silicate bodies.* Doctor of Philosophy (PhD) thesis, University of Kent.

Downloaded from

<https://kar.kent.ac.uk/94483/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.22024/UniKent/01.02.94483>

This document version

UNSPECIFIED

DOI for this version

Licence for this version

CC BY-NC-ND (Attribution-NonCommercial-NoDerivatives)

Additional information

This thesis has been digitised by EThOS, the British Library digitisation service, for purposes of preservation and dissemination. It was uploaded to KAR on 25 April 2022 in order to hold its content and record within University of Kent systems. It is available Open Access using a Creative Commons Attribution, Non-commercial, No Derivatives (<https://creativecommons.org/licenses/by-nc-nd/4.0/>) licence so that the thesis and its author, can benefit from opportunities for increased readership and citation. This was done in line with University of Kent policies (<https://www.kent.ac.uk/is/strategy/docs/Kent%20Open%20Access%20policy.pdf>). If you ...

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal**, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

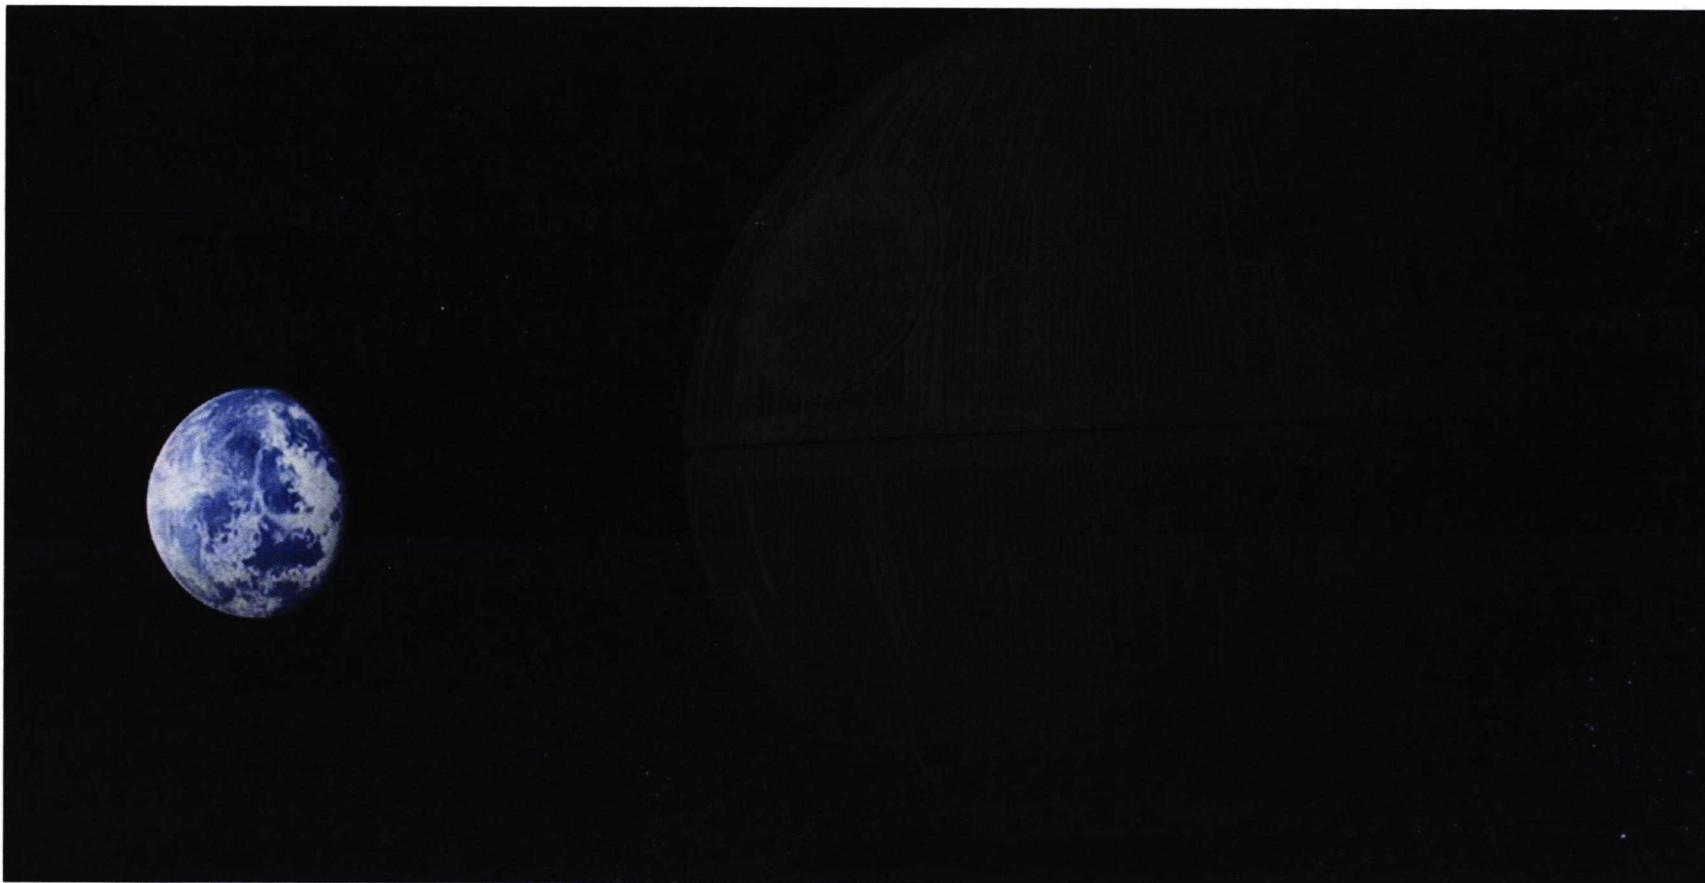
CATASTROPHIC DISRUPTION OF LAYERED ICE-SILICATE BODIES

A THESIS SUBMITTED TO
THE UNIVERSITY OF KENT AT CANTERBURY
IN THE FACULTY OF SCIENCE, TECHNOLOGY AND MEDICAL STUDIES
FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY.

By

Andrew Jonathan Lightwing

April 2011



An impact scientist's ideal piece of lab apparatus.

ACKNOWLEDGEMENTS

My long-suffering supervisor, Professor Mark Burchell. There are many Ph.D.s which bear more than a passing resemblance to a Faustian contract, and so I am grateful to have drawn one of the exceptions to this rule. Thanks for all your advice and understanding over the years; there's not many people that could have sat in that office listening to me talk rubbish and retained enough patience to set me straight afterwards.

Mike Cole, for firing the gun through all 150-odd shots of the experimental programme, for your invaluable help setting up just about every measurement featured in this thesis that required more than a ruler and piece of paper, and for putting up with three years of incessantly stupid questions about where various pieces of lab equipment were located and how they worked.

Jonathan Rowles, for the pool, for the proofreading, and for being used as a general sounding board for ideas over the writing period, even though I'm fairly sure you had no idea what I was talking about most of the time.

Nick Foster, for informing me of the existence of the metal box that was eventually kludged into the improved target chamber.

Mark Price, for the long, drawn-out act of cyber-necromancy that brought my work computer back from the dead, and also for keeping the pair of recalcitrant paper manglers that pass for printers in our office in line.

Jenny Williamson, for proofreading and for the t-shirt. It is an *awesome* t-shirt.

Ben George, Robyn Felix, Sam Lister, Simon Tuohy, Mike Hughes and Xander Meadway, for keeping me sane.

And finally, my Dad. I got there in the end!

ABSTRACT

Large icy Kuiper Belt Objects (KBOs) such as Pluto have observed densities on the order of $1.8 - 2.0 \text{ g cm}^{-3}$. This indicates that they are made up of some combination of ice and silicate material; their size further implies that they will possess a differentiated structure with ice dominating towards the surface and silicate material dominating towards the core. A series of impact experiments using the University of Kent's light gas gun have been carried out in order to determine the impact strength of spherical layered ice-silicate targets that reflect the suspected structure of these bodies. Impacts are also carried out on unlayered ice-silicate targets and pure ice targets for the purposes of comparison. Impact velocities ranged from 1 to 7 km s^{-1} using a range of projectile sizes and materials including stainless steel, titanium, aluminium and copper. Data from previous work undertaken at the University of Kent is incorporated in order to provide a more complete picture of target behaviour.

Basic spherical ice-silicate targets with diameters of 40mm yield critical energy densities (Q^*) of $270 (+2/-26) \text{ J kg}^{-1}$ for 0.8mm stainless steel projectile impacts and $309 (+55/4) \text{ J kg}^{-1}$ for 1mm titanium projectile impacts. This makes them much stronger than large spherical pure ice targets (diameter $0.4 - 1 \text{ m}$) and pure ice cylinders which have critical energy densities of $5.7 (+1.6/-1.2) \text{ J kg}^{-1}$ and $8.1 (+4/-2.7) \text{ J kg}^{-1}$ respectively when impacted with 1mm copper projectiles. 40mm diameter spherical pure ice targets have a higher energy density of $45 \pm 3 \text{ J kg}^{-1}$ when impacted with 1mm aluminium projectiles, indicating some Q^* dependence on target size and projectile size which is borne out by other data sets. Size distribution and energy density plots indicate that targets also exhibit different behaviour when impacted at very high energy densities compared to the behaviour displayed at energy densities around the critical energy density. The addition of an ice layer has little to no effect up to a thickness of five times the projectile diameter, after which target strength is increased dramatically to a Q^* of $912 (+20/-18) \text{ J kg}^{-1}$ for impacts by 1mm titanium projectiles. However this may not be an accurate reflection of the true strength of large KBOs due to the still-basic nature of the layered targets.

Reader Aids – Thesis Overview and Glossary

On the basis that it would be confusing to launch straight into the thesis without any explanation of the overall goals I was working towards during the investigation, I shall first give a very brief explanation of the structure of the thesis and the main points the investigation was attempting to resolve. The primary goal of this investigation was to determine how, if at all, the addition of an ice layer to a body composed of a mixture of ice and silicate would affect the amount of energy required to crater, disrupt or shatter it. The reason for looking at layered bodies is because some of the large bodies in the outer solar system – certain moons of the gas giants such as Callisto and Ganymede and the larger trans-Neptunian objects such as Pluto and Eris, for example – are suspected to possess layered structures due to thermal differentiation, and their composition of a mixture of ice and silicate material implies that the majority of the heavier silicate will have “sunk” into the body, forming a silicate core mixed with some ice that is covered by a layer of pure ice. A better understanding of how these kinds of bodies react to impacts will be useful in understanding the creation and evolution of the areas where these large icy bodies are to be found, such as the Kuiper belt, and we can achieve this understanding by carrying out impacts on small targets which reflect the suspected layered structure of these large trans-Neptunian objects in the laboratory to see how they disrupt.

This thesis describes this investigation and the reasoning behind it in a step-by-step fashion. The first task is to describe in detail the sorts of bodies to be found in the outer solar system – what they are made of, how large they are, and how they are categorised – in order for the reader to build up a good picture of the areas of the solar system being talked about. **Chapter 1 – The Outer Solar System** describes these bodies, and also touches on the reasons why collisions between these bodies have been important in shaping how their general population has evolved over time.

Chapter 2 – Collisions explores the practical side of impacts, explaining how impacts are categorised according to the amount of kinetic energy imparted to a body and the largest fragment remaining after impact, as well as how factors such as the material of a body, its size and its internal structure can affect the amount of energy

required to disrupt it. There is also a brief review of previous work involving laboratory impacts onto icy targets which helps to provide an idea of what this investigation might expect to find.

Chapter 3 – Disrupting Targets in the Laboratory describes how impacts with speeds similar to those encountered by two colliding solar system bodies can be achieved in the laboratory, as well as the methods which were developed to produce, store, impact and measure the many varieties of icy targets – which are required in order to study the target-related factors affecting the impact strength of a target described in Chapter 2 - in detail. **Chapter 4 – Materials** is a preliminary results chapter which describes laboratory measurements providing further background information on the different types of projectiles used to impact targets, the different types of sand used in the targets and the compressive and tensile strengths of pure ice and the various types of ice-silicate mix, which tells us how strong they are in relation to each other and gives some indication of how the different kinds of target will behave under impact.

Chapter 5 – Impact Results and Q^* Calculations is a summary of the main impact results. Q^* is a quantity which represents the impact strength of a target, and Chapter 5 contains Q^* calculations for pure ice cylindrical targets, pure ice spherical targets, ice-silicate cylindrical targets, four different types of ice-silicate spherical targets, targets composed of an ice layer covering an ice-silicate core, and low-temperature targets, as well as discussion on how the impact strength varies from target to target and the reasons why this might be.

Where Chapter 5 uses impact results from many different targets of the same type to build up a single Q^* value for that type of target, **Chapter 6 – Fragment Size Distributions** examines impacts on individual targets to provide a fuller picture of how they are behaving when impacted with varying amounts of kinetic energy. Finally, **Chapter 7 – Discussion** provides a coda to the thesis by discussing in-depth the most interesting phenomena observed during the impact programme and suggests avenues of future work that could be undertaken using the results of this investigation.

Glossary

For the sake of brevity this thesis uses a large number of shorthand terms for certain types of target, target material and some of the quantities used to describe such things as impact energy. These terms are each explained at length in the main body of the text; however I have learned through bitter experience that while they make perfect sense to me, somebody who hasn't spent the last three years dealing with them may not be able to keep track of them quite so easily. This glossary of the most common shorthand terms is included in order to (hopefully) make the reader's job in that regard a little easier.

Sand-ice, ice-silicate - interchangeable terms used to describe water that has been mixed with as much sand as possible (alternatively, sand that has been saturated with water) and then frozen. The percentage ratio of sand to water in this mix is approximately 80:20. The sand used in sand-ice mixes is always kiln-dried sand unless it is specifically stated otherwise.

Target - a body composed of either sand-ice or pure ice intended for use in an impact; usually spherical, although some large cylindrical targets are also used.

Sample - a cylindrical body composed of either sand-ice or pure ice intended for use in tensometer strength testing.

Core - a spherical target about 4 cm in diameter that can be composed of either ice or a sand-ice mix. Used as an inner core covered by an outer ice layer in some shots, hence the name.

M_f/M_o - M_f is the mass of the largest remaining fragment after an impact, and M_o is the original mass of the target. M_f/M_o is therefore the fraction of the original mass of the target that remains intact after an impact.

Cratering - an impact which leaves more than half of the original target mass intact. ($M_f/M_o > 0.5$)

Disruption – an impact which leaves less than half of the original target mass intact. ($M_f/M_o < 0.5$)

Q – the energy density of an impact; obtained by dividing the kinetic energy of the impacting projectile by the total mass of the target, measured in $J\ kg^{-1}$.

Q* - the critical energy density; if a target is struck with this energy density it will leave a largest fragment that has exactly half the mass of the original target. ($M_f/M_o = 0.5$)

Disruption curve – a plot of impact energy densities (Q) against the resulting M_f/M_o values. While this is usually assumed to be a power law relation resulting in a straight line when plotted on a logarithmic scale, the slope of the power law seems to change once specific energy density thresholds are reached and exceeded resulting in a sort of z-shaped plot, hence the term *disruption curve*. Disruption curves are used to calculate Q^* values.

M_{fn}/M_o – The mass of an individual fragment n divided by the original mass of the target; used to plot the size distribution of all fragments recovered after a shot.

LIST OF CONTENTS

<i>Acknowledgements</i>	<i>iii</i>
<i>Abstract</i>	<i>iv</i>
<i>Reader Aids</i>	<i>v</i>
<i>Contents</i>	<i>ix</i>
<i>List of figures</i>	<i>xii</i>
<i>List of tables</i>	<i>xx</i>
1 The Outer Solar System	1
1.0 Introduction	1
1.1 Pluto, Eris, and the dwarf planet debate	1
1.2 Trans-Neptunian Objects	4
1.3 The Kuiper Belt	7
1.4 The Scattered Disc	10
1.5 The Oort Cloud	11
1.6 Centaurs	14
1.7 The Nice Model	14
1.8 Layered Bodies	16
1.9 Collisional Families	18
1.10 Conclusion	20
2 Collisions	21
2.0 Introduction	21
2.1 Examples of collisions in the solar system	21
2.1.1 Shoemaker-Levy 9	21
2.1.2 Tempel 1 and the Deep Impact mission	23
2.2 The critical energy density Q^*	25
2.3 Why collisions are important	29
2.4 Factors that affect the disruption energy	30
2.4.1 Repeated impacts	30
2.4.2 The size of a body	32
2.4.3 Porous bodies	35

2.4.4	Layered bodies	37
2.5	Previous calculations of Q^*	38
2.6	Conclusion	40
3	Disrupting Targets in the Laboratory	41
3.0	Introduction	41
3.1	Light gas guns	41
3.2	University of Kent's Two-Stage Light Gas Gun	43
3.3	Targets	48
3.4	Target holders	53
3.5	Shot process	55
3.6	Tensometer strength tests	56
3.7	Projectiles	58
3.8	Impact Angle	58
3.9	Conclusion	61
4	Materials	62
4.0	Introduction	62
4.1	Projectiles	62
4.1.1	0.8 mm diameter stainless steel projectiles	63
4.1.2	1 mm diameter titanium projectiles	64
4.1.3	1 mm diameter aluminium projectiles	64
4.1.4	1.98 mm diameter soda-lime glass projectiles	65
4.1.5	1mm and 2mm diameter copper projectiles	66
4.1.6	Projectile reliability	68
4.2	Sand types	68
4.3	Tensometer strength tests	73
4.4	Comparison with previous work	81
4.5	Conclusion	82
5	Impact Results and Q^* Calculations	83
5.0	Introduction	83
5.1	Impact outcomes	83
5.1.1	Kiln-sand cores	83

5.1.2	Layered kiln-sand cores	87
5.2	Determining Q^* from impact data	92
5.3	Q^* for homogenous targets	98
5.4	Q^* for layered targets	105
5.5	Varying target temperature and sand grain size	108
5.6	Comparison with previous work	112
5.7	Conclusion	113
6	Fragment Size Distributions	115
6.0	Introduction	115
6.1	Plots of size distribution per target/projectile type	115
6.1.1	Comparison of results	119
6.2	Plots of size distribution per Q/Q^* ratio	128
6.3	Power law fits	132
6.4	Conclusion	145
7	Discussion	147
7.0	Introduction	147
7.1	The shattering regime	147
7.2	Layered targets	148
7.3	Size scaling	162
7.4	Peak shock pressures	166
7.5	The size distribution of the Haumea collisional family	174
7.6	Conclusion	178
	Conclusion	180
	Bibliography	184
	Appendix	193

LIST OF FIGURES

<i>Figure</i>	<i>Page</i>
1.2.1 An image of the dwarf planet Pluto, showing it in true colour and created by tracking brightness changes in Pluto while it was being partially-eclipsed by Charon.	5
1.2.2 The most recent colour-fit of Pluto captured with the Hubble Space Telescope. (Buie <i>et al.</i> 2010)	5
1.2.3 Cartoon illustrating the Kuiper belt in relation to the Oort Cloud.	6
1.3.1 Graph showing the semimajor axis of known TNOs against inclination and eccentricity, with the areas of resonance marked; resonant objects are shown in black while red objects are non-resonant stable objects. (Levison et al. 2007)	7
1.4.1 Hubble Space Telescope image of the dwarf planet Eris, both the largest known TNO and the largest member of the scattered disc, along with its moon Dysnomia.	10
1.5.1 Schematic of the outer solar system.	12
1.9.1 Cartoon illustrating Haumea's elongated shape.	18
2.1.1 Hubble Space Telescope image showing the 21 distinct fragments of comet Shoemaker-Levy 9 spread out over 1.1 million kilometres.	22
2.1.2 A fireball plume rising over the limb of Jupiter	22
2.1.3 A track of dark spots across the Jovian surface where the Shoemaker-Levy fragments have impacted	22
2.1.4 The impactor component of Deep Impact strikes Tempel 1.	24
2.2.1 Graph showing the theoretical divide between the cratering and disruption regimes, and how it is used to calculate Q^* .	27
2.2.2 Graph of a typical Q vs M_i/M_o plot that reflects observed experimental results, including shattering regime	28
2.4.1 Diagram showing the various stages of crack propagation in a homogenous target.	31
2.4.2 Graph showing the two different sets of scaling laws; material strength starts out dominant at small scales but decreases steadily as	33

	the size of the body increases, while the effect of gravity is negligible to start with but quickly becomes dominant at larger size scales. (Durda <i>et al.</i> 1999)	
2.4.3	Graph showing the material and gravity scaling laws combined to form one coherent picture of how the size of a body affects the critical energy density. (Durda <i>et al.</i> 1999)	34
2.4.4	Image of the comet Tempel-1 taken five minutes before the impactor component of the Deep Impact probe struck it.	35
2.4.5	Image of the asteroid 25143 Itokawa taken by the Hayabusa mission in 2005.	36
3.1.1	The typical firing mechanism of a light gas gun (Burchell <i>et al.</i> 1999).	42
3.2.1	The two-stage light gas gun at the University of Kent	44
3.2.2	Examples of the two types of burst disc before and after firing.	45
3.2.3	Diagram illustrating sabot construction (Burchell <i>et al.</i> 1999).	46
3.2.4	Data from sensors within the light gas gun tracing the flight of the projectile through the gun.	47
3.3.1	A spherical ice-silicate target.	50
3.3.2	A sand-ice core and a sand-ice core covered in a layer of ice	51
3.4.1	A cylindrical ice-silicate target sitting in its target holder. A partial ice cap is visible on the rear of the target.	54
3.4.2	Target holder for spherical ice targets with the top removed. A plastic mould has been placed in the polystyrene cup where a target would normally sit.	54
3.6.1	An ice-silicate sample being crushed in the tensometer	57
3.6.2	A cracked ice-silicate sample after being tested in the tensometer, and the plastic mould used to construct it.	57
3.8.1	Diagram showing how a deviation s from a central (normal) impact axis will result in an impact at an angle α .	59
4.1.1	Image of a 0.8mm stainless steel projectile	63
4.1.2	Image of a 1mm titanium projectile	63
4.1.3	Image of a 1mm aluminium projectile	65
4.1.4	Image of a 1.98mm soda-lime glass projectile	65

4.1.5	Image of a 1mm copper projectile	66
4.1.6	Image of a 2mm copper projectile	66
4.2.1	Graph showing how the density of a sand-ice mix varies with the size of the sand grains used in the mix.	70
4.2.2	Microscope image showing fraction E sand grains.	70
4.2.3	Microscope image showing fraction D sand grains.	71
4.2.4	Microscope image showing fraction C sand grains.	71
4.2.5	Microscope image showing kiln-sand sand grains.	72
4.2.6	Microscope image showing fraction B sand grains.	72
4.2.7	Microscope image showing fraction A sand grains.	73
4.3.1	Diagram showing the difference between compressive and tensile strength and where the force is applied to samples to measure each.	74
4.3.2	Graph plotting the average compressive strength of various sand-ice mixes against the mean grain size of the sand type used in the mix.	77
4.3.3	Graph of the tensile strength of various sand-ice mixes plotted against the grain size of the sand type used in the mix.	78
4.3.4	Fraction E samples which have cracked while undergoing tensile strength tests.	79
4.3.5	Fraction B samples which have split while undergoing tensile strength tests.	80
5.1.1	A cratered kiln-sand core.	84
5.1.2	A near disrupted kiln-sand core	84
5.1.3	A disrupted kiln-sand core	85
5.1.4	Diagram showing the differing outcomes of impacts on to spherical cores as they vary with the impact energy density Q relative to Q^*	86
5.1.5	A cratered layered core that has retained its ice layer	88
5.1.6	A cratered layered core that has had its ice layer removed	88
5.1.7	A near-disrupted layered core with the ice layer still attached	88
5.1.8	A near-disrupted layered core with the ice layer removed	88
5.1.9	A disrupted layered core with pieces of the ice layer still attached to some of the fragments.	89
5.1.10	A disrupted layered core where the ice layer has been almost totally removed from the remaining fragments.	89

5.1.11	Diagram showing the differing outcomes of impacts on to layered targets as they vary with the impact energy density Q	90
5.1.12	Stages of disruption observed in a soda-lime glass core/gypsum mantle target (Okomoto <i>et al.</i> 2008)	91
5.2.1	Typical plot of how M_f/M_o varies with Q	93
5.3.1	Results for 2mm copper/1mm copper impacts onto kiln sand cylinders ($Q^* = 150 \text{ J kg}^{-1}$ for 2mm copper impacts) and 0.8mm stainless steel impacts onto kiln sand cores ($Q^* = 278 \text{ J kg}^{-1}$)	99
5.3.2	Results for 1mm copper impacts onto large solid ice spheres (Murray 2004, $Q^* = 5.5 \text{ J kg}^{-1}$), 1mm copper impacts onto solid ice cylinders ($Q^* = 8.1 \text{ J kg}^{-1}$), and 1mm aluminium impacts onto solid ice cores ($Q^* = 45 \text{ J kg}^{-1}$).	100
5.3.3	Results of 1mm aluminium impacts onto solid ice cores ($Q^* = 45 \text{ J kg}^{-1}$), 0.8mm stainless steel impacts onto kiln sand cores ($Q^* = 278 \text{ J kg}^{-1}$), and 1mm titanium impacts onto kiln sand cores ($Q^* = 308 \text{ J kg}^{-1}$).	103
5.4.1	Results of 1mm titanium impacts onto layered kiln-sand cores, with ice layer thickness anywhere between 1 – 7 mm.	106
5.4.2	Results for 1mm titanium impacts onto layered kiln-sand cores, separated into two categories based on ice layer thickness. In red are cores with thin (< 5 mm) ice shells ($Q^* = 360 \text{ J kg}^{-1}$), and in blue are cores with thick (>5 mm) ice shells ($Q^* = 912 \text{ J kg}^{-1}$).	106
5.4.3	Graph showing how the percentage of a layered target's mass made up by the ice layer increases with the layer's thickness	107
5.4.4	Results for 0.8mm stainless steel, 1mm titanium and 1.98mm glass impacts onto thin-shelled cores	109
5.5.1	Comparison of 1mm titanium impacts onto kiln-sand cores frozen at -20°C ($Q^* = 308 \text{ J kg}^{-1}$) and kiln-sand cores frozen at -120°C ($Q^* = 319 \text{ J kg}^{-1}$).	109
5.5.2	Results of thermocouple measurement of target temperature during firing process showing the level of temperature increase between sealing the gun and firing the gun.	111
5.5.3	Results for 0.8mm stainless steel impacts onto sand-ice cores	111

	constructed from fraction B ($Q^* = 43 \text{ J kg}^{-1}$), fraction D (284 J kg^{-1}) and fraction E (293 J kg^{-1}) sand grains.	
6.1.1	Graph showing the theoretical fragment size distribution were the measurements not limited by the sensitivity of the equipment	120
6.1.2	Fragment size distributions for 0.8mm stainless steel shots onto kiln sand cores.	121
6.1.3	Fragment size distributions for 1mm titanium shots onto kiln sand-ice cores.	121
6.1.4	Fragment size distributions for 1mm titanium shots onto low temperature kiln sand cores	122
6.1.5	Fragment size distributions for 1mm aluminium shots on to solid ice cores.	123
6.1.6	Fragment size distributions for 0.8mm stainless steel shots onto fraction E cores (average grain diameter $160 \mu\text{m}$)	123
6.1.7	Fragment size distributions for 0.8mm stainless steel shots onto fraction D cores (average grain diameter $229 \mu\text{m}$)	124
6.1.8	Fragment size distributions for 0.8mm stainless steel shots onto fraction B cores (average grain diameter $1043 \mu\text{m}$)	124
6.1.9	Graph showing how the typical fragment size distribution varies as Q scales with Q^*	126
6.2.1	Fragment distribution plots of various shots falling into the Q/Q^* range of 1.5 – 1.9	129
6.2.2	Graph of various fragment distribution plots which fall into the Q/Q^* range of 2.4 - 2.7	129
6.2.3	Graph of various fragment distribution plots falling into the Q/Q^* range of 3.1 – 3.5	130
6.3.1	Example graph illustrating which data points are excluded when fitting a power law to a size distribution	133
6.3.2	Power law fits for impacts by a 0.8mm stainless steel projectile onto kiln sand cores.	135
6.3.3	Power law fits for impacts by a 1mm titanium projectile onto kiln sand cores.	135
6.3.4	Power law fits for impacts by a 1mm titanium projectile onto low	136

	temperature kiln-sand cores.	
6.3.5	Power law fits for impacts by a 0.8mm stainless steel projectile onto fraction E cores	136
6.3.6	Power law fits for impacts by 0.8mm stainless steel projectiles onto fraction D cores	137
6.3.7	Power law fits for impacts by 0.8mm stainless steel projectiles onto fraction E cores	137
6.3.8	Power law fits for impacts by 1mm aluminium projectiles onto solid ice cores	138
6.3.9	Plots of size distribution slope b against Q/Q^* for 0.8mm stainless steel impacts onto kiln sand cores	140
6.3.10	Plot of size distribution slope b against Q/Q^* for 1mm titanium impacts onto kiln-sand cores	140
6.3.11	Plot of size distribution slope b against Q/Q^* for 1mm titanium impacts onto low temperature kiln-sand cores	141
6.3.12	Plot of size distribution slope b against Q/Q^* for 0.8mm stainless steel impacts onto fraction B cores	141
6.3.13	Plot of size distribution slope b against Q/Q^* for 0.8mm stainless steel impacts onto fraction D cores	142
6.3.14	Plot of size distribution slope b against Q/Q^* ratio for 0.8mm stainless steel impacts onto fraction E cores	142
6.3.15	Plot of size distribution slope b against Q/Q^* for 1mm aluminium impacts onto solid ice cores	143
6.3.16	Hypothetical plot of Q/Q^* vs. b relation, with the different impact regimes marked.	143
7.2.1	Diagram showing one possible mechanism that would concentrate impact energy into the ice shell	149
7.2.2	Diagram showing how the ice shell may absorb the majority of a projectile's impact energy even if it penetrates the shell and strikes the core	150
7.2.3	Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 1 km s^{-1} .	153

7.2.4	Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 1 km s^{-1} , with the impact shockwave visible.	153
7.2.5	Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 1 km s^{-1} , with the impact shockwave visible and spalled fragments included.	154
7.2.6	Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 3 km s^{-1} .	155
7.2.7	Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 3 km s^{-1} , with the impact shockwave visible.	155
7.2.8	Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 3 km s^{-1} , with the impact shockwave visible and spalled fragments included.	156
7.2.9	Separate disruption curves for the ice layer and the sand-ice core in a body where the ice layer has 10% of the mass of the target	157
7.2.10	Separate disruption curves for the ice layer and the sand-ice core in a body where the ice layer has 90% of the mass of the target	158
7.2.11	Separate disruption curves for a body that is 50% ice and 50% sand-ice	159
7.2.12	Plot showing energy bleedover between the two separate types of disruption curve as the ice layer moves into the shattering mode of behaviour	159
7.3.1	Benz & Asphaug's modelling for how Q^* changes for ice targets with the scale of the target.	163
7.3.2	Scaling fit for 0.5 km s^{-1} impacts onto ice adjusted to correlate with observed laboratory results for ice.	164
7.3.3	Scaling fit for 3 km s^{-1} impacts onto ice adjusted to correlate with observed laboratory results for ice.	164
7.4.1	. Plot of projectile velocity versus peak pressures for 1mm aluminium	167

	impacts onto solid ice cores.	
7.4.2	Plot of peak pressure against projectile velocity for 0.8mm steel impacts onto sand-ice cores (approximated by permafrost).	168
7.4.3	Plot of peak pressures versus impact velocity for 1mm aluminium impacts onto ice targets and 0.8mm stainless steel impacts onto permafrost targets using Mizutani's formula.	170
7.4.4	Shockwave propagation through a spherical target, taken from Mizutani (1993). The shockwave radiates outwards from the impact site in all directions and thus dissipates according to an inverse power law.	172
7.4.5	Attenuation of shock pressure for 1mm aluminium impacts onto solid ice cores as the shock travels deeper into a target body.	173
7.4.6	Possible differences in shock wave geometry in a central impact event and an off-centre impact event.	174
7.5.1	Size distribution of the Haumea collisional family.	175
7.5.2	Size distribution of the Haumea collisional family with the Haumea mass point removed.	178

LIST OF TABLES

<i>Table</i>	<i>Page</i>
2.5.1 List of previous Q^* calculations for homogenous and porous icy bodies.	38
4.1.1 Properties of 0.8 mm stainless steel projectiles	63
4.1.2 Properties of 1mm titanium projectiles	63
4.1.3 Properties of 1mm aluminium projectiles	65
4.1.4 Properties of 1.98mm soda-lime glass projectiles	65
4.1.5 Properties of 1mm copper projectiles	66
4.1.6 Properties of 2mm copper projectiles	66
4.1.7 Failure rates of the various projectiles used during the experimental programme.	67
4.1.8 Comparison of measured projectile densities with accepted densities for each material type.	67
4.2.1 Table showing the specified and actual measured properties of the various types of sand grains used in the experimental programme.	68
4.3.1 Table showing the results of compressive strength tests on samples constructed of various different types of sand grains.	76
4.3.2 Table showing the results of tensile strength tests on samples constructed of various different types of sand grains.	76
5.2.1 Table showing fitted Q^* values, associated errors and methods of error calculation	97
6.1.1 Categories of data for plots of size distribution with varying energy density	116
6.3.1 Results from power law fits of fragment size distributions. Fits are of the form $y = ax^b$, where R is the regression coefficient	134
7.4.1 C and S values for ice and aluminium	167
7.4.2 Peak pressures for 1mm aluminium impacts onto solid ice cores (planar approximation)	167
7.4.3 C and S values for steel and permafrost	168
7.4.4 Peak pressures for 0.8mm stainless steel impacts onto sand-ice	168

- cores (approximated by permafrost)
- 7.4.5 Peak pressures for 1mm aluminium impacts onto solid ice cores 170
(Mizutani).
- 7.4.6 Peak pressures for 0.8mm stainless steel impacts onto sand-ice 170
cores (Mizutani, approximated by permafrost).

Chapter 1 – The Outer Solar System

“It’s funny - you live in the universe but you never look at these things until somebody comes to visit.”

- Dr. Zoidberg

1.0 - Introduction

The outer solar system is a region of space that has undergone a dramatic upheaval in our understanding over the last twenty years. Before 1992, the number of known bodies orbiting past Neptune was just two – Pluto and Charon. This is an exceedingly small body of evidence upon which to base hypotheses about the structure of the outer solar system, and so it is no surprise that as more and more objects were discovered from 1992 onwards their physical characteristics and distribution challenged several established theories about how the population of the outer solar system. Today, over a thousand bodies past Neptune’s orbit are known. Some of them, such as Eris, have provoked great public controversy over how we think of the larger members of this population, but even the less well-known new discoveries have greatly enhanced our knowledge of this region of space. In this chapter I shall give an overview of how the outer solar system breaks down into several discrete areas, the type of bodies that can be found there, and the theories of how the different areas interact with each other and how they came to be.

1.1 – Pluto, Eris, the dwarf planet debate and the semantics of Solar System science.

In 1992 the first body orbiting the Sun beyond the orbit of Pluto was discovered (Jewitt & Luu, 1993). This body, termed 1992QB₁, was relatively small at only 160 km in diameter – the typical dimensions of a large asteroid. The existence of a large belt of icy material orbiting the Sun beyond Neptune had been independently posited by the astronomers Kenneth Edgeworth (Edgeworth 1943) and Gerard Kuiper (Kuiper 1951), and the theory persisted over subsequent decades as it provided a plausible source for the constant stream of comets encroaching on the inner Solar System. The discovery of 1992QB₁ appeared to confirm the existence of this belt - popularly

known as the Kuiper belt and less well-known as the Edgeworth-Kuiper belt – and sparked a wave of discoveries of similar bodies orbiting at various distances of the Sun beyond Neptune.

With the number of known Kuiper belt objects rising rapidly these discoveries were potentially ominous for Pluto's status as the Solar System's ninth planet. Pluto had been considered a planet since it was first discovered in 1930; initially thought to be the same size as Earth, its size was continually revised downwards until arriving at the current diameter of 2300 km. Despite being significantly smaller than originally thought it retained its planet status since, while there was no standard definition of what counted as a planet, Pluto certainly seemed to qualify – it was a spherical body far larger than all known asteroids which orbited the Sun, not another planet. As long as no other bodies were found in the same general area of the Solar System which approached it in size Pluto's planet status was secure.

This all changed with the plethora of Kuiper belt objects being found both inside and outside its orbit – and a sizeable proportion even sharing it. At first the objects discovered were small, on the same scale as 1992QB₁ (which was approximately 160 km in diameter) but in 2003 the object Sedna was found, which had a diameter of up to 1600 km (Stansberry *et al.* 2007), proving that objects not much smaller than Pluto existed in the belt. The final nail in the coffin for Pluto was the discovery of Eris (Brown *et al.* 2005) and its companion body Dysnomia in 2005. Eris was calculated to have a diameter of 2400 ± 100 km (Brown *et al.* 2006) making it at least 5% larger than Pluto. Dysnomia, the companion body, was much smaller at only 125 km in diameter, but the fact that it was orbiting Eris in the same manner as Charon (the relatively-recently discovered Nix and Hydra would possibly be better comparisons, being similar in size) only added weight to the case that Pluto was not a unique body that happened to orbit inside this belt of icy material, but rather was part of the belt itself.

The discovery of Eris finally prompted the International Astronomical Union (IAU) to try to come up with a formal definition of a planet so that they could sort out the confusion surrounding what Pluto actually was, but this presented the astronomical

community with a dilemma: they could either accept Eris as a new planet, or downgrade Pluto so that it wasn't. While there was a significant amount of cultural inertia behind the argument to keep Pluto as a planet, any scientific definition that did so would have had to accept Eris – and by extension any subsequent bodies discovered beyond Neptune with sizes on a similar scale – as a planet as well. Faced with a situation in which the number of recognised planets could potentially spiral out of control, the IAU did what any bureaucracy does when faced with a choice between two decisions either of which would displease an awful lot of people: it compromised (IAU resolutions 5A and 6A passed on 24/08/06, which can be found at http://www.iau.org/static/resolutions/Resolution_GA26-5-6.pdf). The formal definition of a planet reached by the IAU was as follows:

(1) A "planet" is a celestial body that: (a) is in orbit around the Sun, (b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape, and (c) has cleared the neighbourhood around its orbit.

This disqualified Pluto as a planet, since while it was orbiting around the Sun and had a spherical shape, the large number of objects discovered inside, outside and sharing its orbit ensured it did not meet the third requirement of having “cleared the neighbourhood around its orbit”. However, as a consolation prize the new category of “dwarf planets” was introduced:

(2) A "dwarf planet" is a celestial body that: (a) is in orbit around the Sun, (b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape, (c) has not cleared the neighbourhood around its orbit, and (d) is not a satellite.

This reclassified Pluto, Eris, the TNOs Haumea and Makemake and the asteroid Ceres as dwarf planets. Everything else found in the Kuiper belt was placed into the following category:

(3) *All other objects, except satellites, orbiting the Sun shall be referred to collectively as "Small Solar System Bodies."*

While the definitions did not pass without criticism, they appeared to do an adequate job of sorting out the mess of objects found in the Kuiper Belt and beyond. However, as we shall see, the category of "Small Solar System Bodies" is a deceptively broad brush with which to paint a wide and fascinating variety of objects.

1.2 - Trans-Neptunian Objects.

"Trans-Neptunian objects" (TNOs) is the general term for bodies that orbit the sun at a greater distance than Neptune. This category of bodies includes the newly-minted dwarf planets Eris and Pluto, which are the largest currently-known TNOs with diameters of 2500 km and 2300 km respectively; most observed TNOs are much smaller with diameters of a few hundred kilometres, if that. These small sizes contribute to the TNO population being incredibly faint, with high apparent magnitudes – Pluto's apparent magnitude is 13.65, but its moons Nix and Hydra have apparent magnitudes of 23. Since they are so faint, it makes ascertaining detailed information about their composition and structure very difficult. Based on the density of the larger TNOs (generally in the range of $1.4 - 2.6 \text{ g cm}^{-3}$; although the density measurements have significant uncertainties associated with them their upper and lower boundaries fall within this range) and the low density of ice (0.997 g cm^{-3}) it is reasonable to assume that TNOs are composed of a mixture of ice and some significant quantity of rocky/silicate material – a solid silicate material such as granite has a density of $2.6-2.7 \text{ g cm}^{-3}$, with significant variation depending on the specific material - the addition of which raises the density to observed levels. Beyond that not much compositional information can be inferred from bands in reflective spectra etc. While light curves can tell us much about the shape of a body and how fast it rotates, attempts to use spectra to gain insight into the composition of a body are hindered by the fact that multiple surface compositions can fit a single observed spectra. Additionally, studying light curves and spectra only tells us information about the *surface* of a TNO, which may have been altered significantly by solar weathering, micrometeorite strikes and so on compared to the internal regions. The difficulty

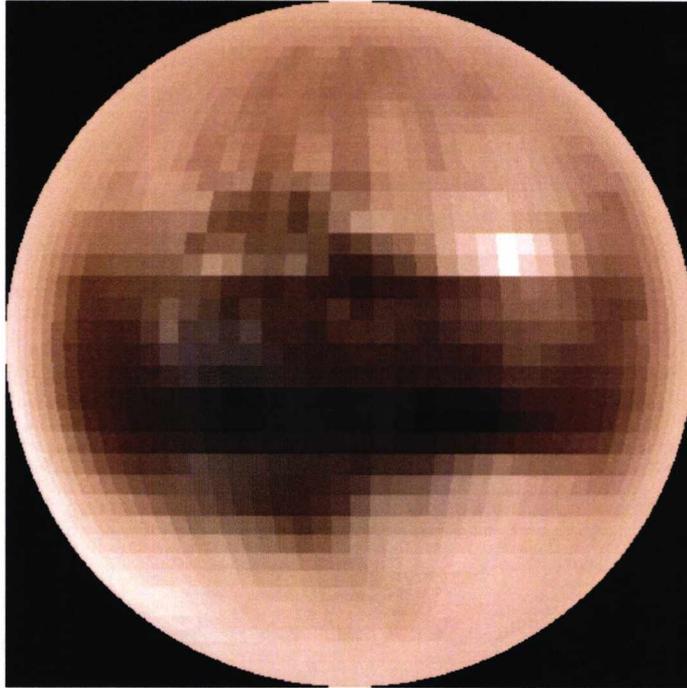


Figure 1.2.1 - An image of the dwarf planet Pluto, showing it in true colour and created by tracking brightness changes in Pluto while it was being partially-eclipsed by Charon. Pluto is the largest member of the Kuiper belt (although not the largest known TNO) with a diameter of 2300 km. Credit Eliot Young (SwRI) *et al.*, NASA. Image source: <http://apod.nasa.gov/apod/ap010319.html>



Figure 1.2.2 – The most recent colour-fit of Pluto captured with the Hubble Space Telescope. (Buie *et al.* 2010)

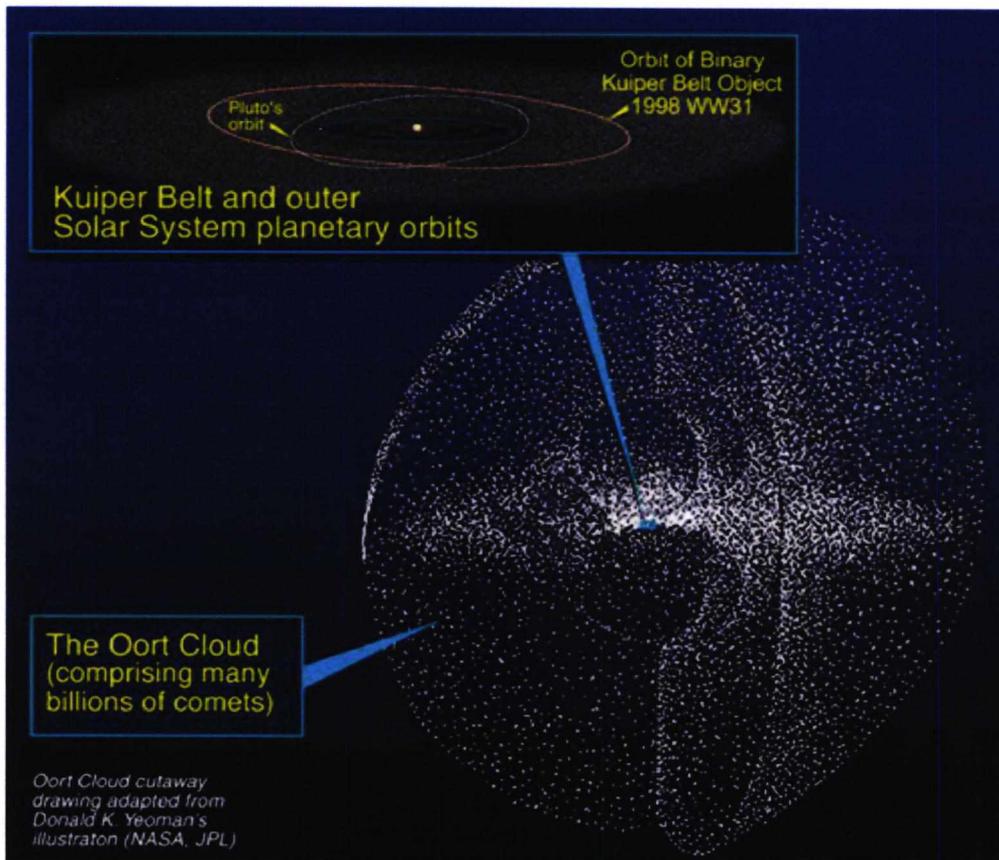


Figure 1.2.3 – Cartoon illustrating the Kuiper belt (and the Scattered Disc, although this is unlabelled; it would be the outer area of the Kuiper belt in the cartoon) in relation to the Oort Cloud. Image source: http://herschel.jpl.nasa.gov/images/kuiper_oort.jpg

involved in deriving the properties of TNOs from observations is illustrated by figures 1.2.1 and 1.2.2; these are the current best images of Pluto available, and Pluto is the second-largest TNO (and closer to us than the largest), but even here the images are pixellated and/or indistinct.

TNOs can be split into several discrete populations based on their orbital characteristics and distance from the Sun. These populations are:

- The Kuiper Belt
- The Scattered Disc
- The Oort Cloud.

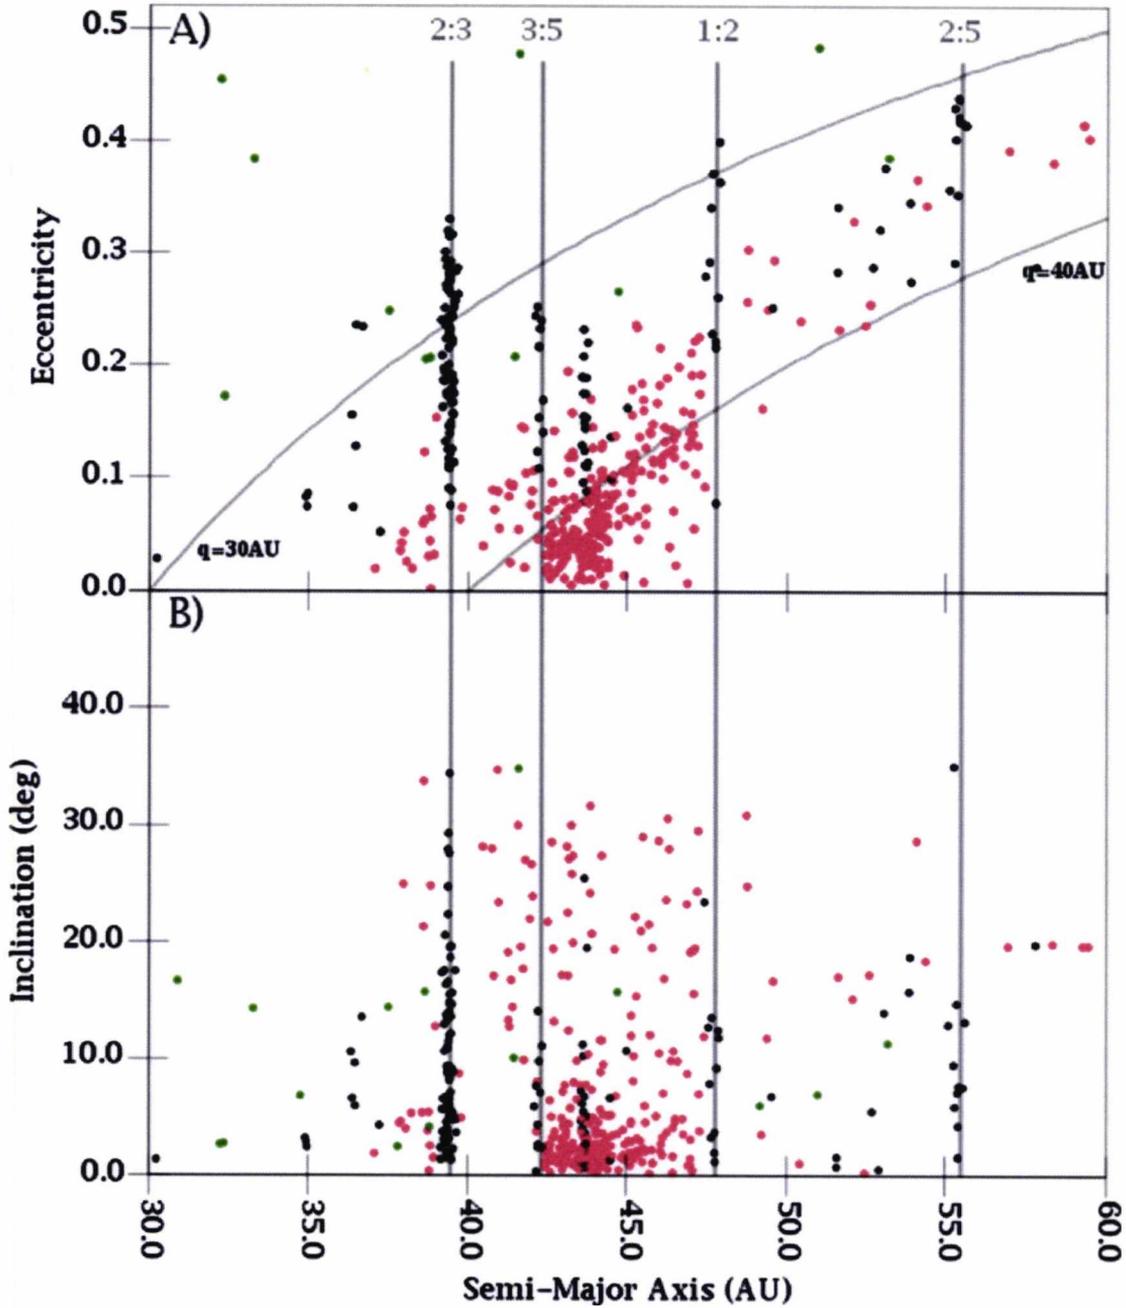


Figure 1.3.1: Graph showing the semimajor axis of known TNOs against inclination and eccentricity, with the areas of resonance marked; resonant objects are shown in black while red objects are non-resonant stable objects. It shows the 2:3 resonance heavily populated by the Plutinos and the area between the 3:5 and 1:2 resonances populated by the mostly low-inclination classical KBOs, as well as the sudden decline in observed objects past the 1:2 resonance. (Levison *et al.* 2007)

1.3 - The Kuiper Belt

The Kuiper belt is a region of space in the Solar System that extends from the orbit of Neptune at 30 AU out to about 55 AU. This region contains the largest number of

TNOs observed, with over a thousand Kuiper Belt Objects (KBOs) currently known. The Kuiper Belt is not tightly bound to the plane of the ecliptic in the same way as the asteroid belt; most member objects orbit at inclinations of anywhere up to 10° , with a minority population possessing even higher inclinations. Figure 1.3.1 illustrates the structure of the Kuiper belt, plotting KBO inclination and eccentricity against orbital semi-major axis.

The vast majority of objects in the Kuiper Belt are found between 39.5 and 47 AU from the Sun. (De Sanctis *et al.* 2001). These two boundaries are defined by two different mean-motion resonances with Neptune; the inner boundary at 39.5 AU by the 2:3 resonance, and the outer boundary at 47 AU by the 1:2 resonance. Resonances with Neptune have a further effect in that there are several discrete populations of Kuiper Belt objects which inhabit orbits corresponding to certain resonances (Levison *et al.* 2007). Neptune has a significant influence on the overall structure of the Kuiper Belt, with gaps similar to the Kirkwood gaps in the asteroid belt in between the populated resonant areas - without the stabilising influence of a resonance, KBOs cannot orbit in these regions without eventually being perturbed inwards by Neptune.

Roughly 20% of the known objects in the Kuiper Belt occupy the 3:2 resonance at 39.5 AU (Kavelaars *et al.* 2008). Since Pluto was the first object found in this resonance, all subsequent objects discovered have been called Plutinos. Past the 1:2 resonance there is then a drop in the number of observed objects until the 3:5 resonance at 42 AU.

The 3:5 resonance contains its own population of KBOs but also marks the inner boundary of a region known as the Classical Kuiper Belt (CKB). This is a region of the belt containing objects that occupy classical orbits that are unperturbed by the gravitational influence of Neptune; they do not need resonances to stabilise their orbits. The first such object discovered was the second KBO discovered after Pluto, 1992 QB₁ (Jewitt & Luu 1993). Subsequent objects in the classical belt are often called cubewanos as a result (“Q-B-one-os”).

While it is not governed by resonances the CKB nevertheless appears to contain two

populations of objects. The first so-called dynamically “cold” population have orbital inclinations of under 4° , while any object in the CKB with an inclination above 4° is classed as part of the dynamically “hot” population (Levison *et al.* 2004). Interestingly this does not appear to be a simple arbitrary delineation based on orbital inclination; the objects in the cold population have a higher albedo than those in the hot population, while there is also a correlation between colour and inclination (Trujillo *et al.* 2002). The cold population are almost entirely red, whereas the hot population display a range of colours running from red to grey – for some reason, there is a deficit of grey objects within the cold populations. These two factors point to the hot and cold populations being two distinctly different classes of objects. A number of explanations have been advanced to explain this difference, such as collisional resurfacing or the two populations having formed in different areas of the solar system, but our current knowledge of the Kuiper belt is insufficient to understand exactly what the cause is.

Within the classical belt there is the 4:7 resonance with Neptune, although this contains far fewer objects than the 2:3 resonance with only twenty bodies in this orbit. Finally the outer edge of the classical belt is marked by the thinly-populated 1:2 resonance at 47.7 AU. Past this point there is a significant drop-off in the number of objects detected that cannot be accounted for by observational bias (Bernstein *et al.* 2004); despite recent surveys having the sensitivity to detect objects larger than 37 km in diameter out to 60 AU very few have been found. Unless there is a sudden increase in the number of objects with a diameter of less than 40 km past the 1:2 resonance then this sudden decline in the number of observed TNOs is real; therefore there must be some reason why most of the known TNOs are clustered into the relatively small space of the Kuiper belt.

The Kuiper belt appears to suffer a considerable mass deficit; while it is estimated to have had 10-30 M_e after its primordial formation, today’s observations and models place its mass at 0.01 – 0.1 M_e (Charnoz & Morbidelli 2006). This is thought to be caused by one of two mechanisms (or even a combination of the two): either collisions between KBOs have continually eroded the smaller members of the belt, reducing the mass to currently observed levels, or there was some event in the Kuiper

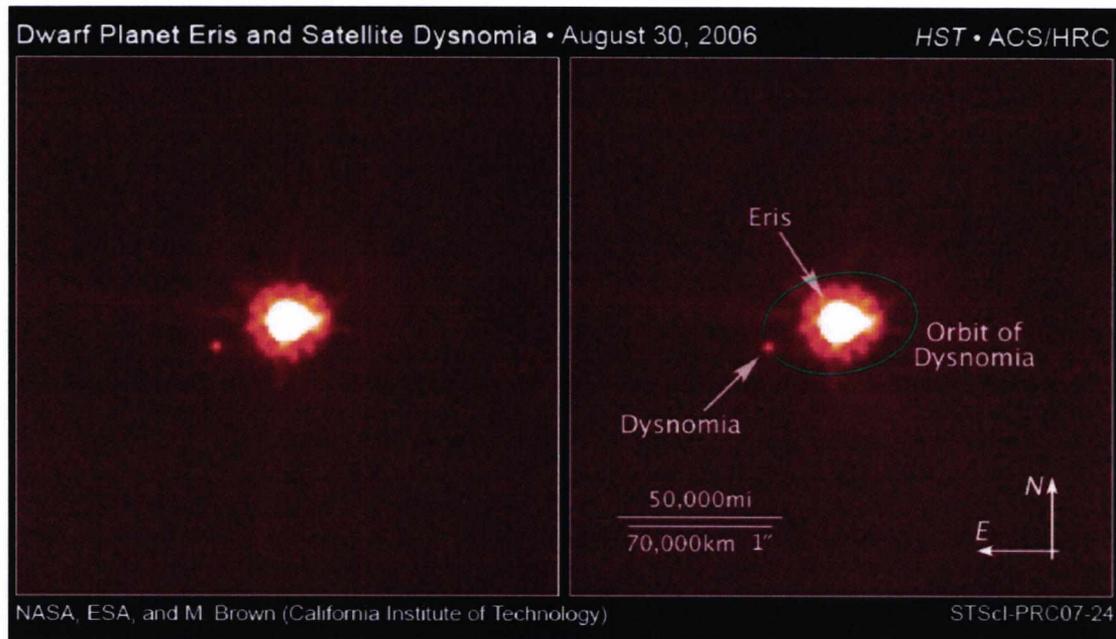


Figure 1.4.1: Hubble Space Telescope image of the dwarf planet Eris, both the largest known TNO and the largest member of the scattered disc, along with its moon Dysnomia. Image credit: NASA, ESA and Mike Brown. Image source: http://www.nasa.gov/mission_pages/hubble/news/eris.html

belt's past that removed this mass dynamically through scattering by the giant planets or orbital resonances.

1.4 – The Scattered Disc

The scattered disc contains a collection of objects that orbit mostly beyond the outer edge of the Kuiper belt, although some of them have perihelia that take them into the

Kuiper belt. Unlike the Kuiper belt, which is dynamically stable with bodies either locked into resonant or classical orbits, the orbits of the TNOs in the scattered disc are highly eccentric and irregular (as can be seen from figure 1.3.1) – with some objects having perihelia of 35 AU or less, but possessing aphelia that extend out past 100 AU from the Sun - and furthermore are subject to perturbations by Neptune. Since it is dynamically unstable it is believed to be the source of short period comets, which are occasionally perturbed inwards towards the rest of the Solar System by the gas giants.

The Scattered disc is very thinly populated compared to the Kuiper belt, with very few objects currently known. This may partly be down to observational bias, since while the objects in the Kuiper belt possess relatively low inclinations (less than 10°) and are therefore located within a thick disc or torus in the plane of the solar system, the bodies in the scattered disc have much higher inclinations (up to 40°) and inhabit a far larger region of space compared to the Kuiper belt. This means that any observations made along the plane of the solar system would find very few TNOs in the scattered disc since they are not clustered in this plane like the Kuiper belt objects. However, even accounting for this the population in this region of space shows a significant drop-off compared to the Kuiper belt.

The best known members are the dwarf planet Eris (figure 1.3.1) and its companion Dysnomia, both of which were discovered by a team of astronomers led by M.E. Brown in 2005 (Brown *et al.* 2005, Brown *et al.* 2005). Eris is the most massive object in either the Kuiper belt or the scattered disc, but the scattered disc contains very few similar bodies.

1.5 – The Oort Cloud

While short-period comets can be accounted for by the scattered disc population, the disc is located too close to the sun to account for comets with periods of many thousands of years – the sheer amount of time they take to orbit means they cover a great deal of distance, and their high eccentricities put their aphelia in a region of space far beyond the outer bounds of the scattered disc. Short period comets tend to orbit the sun in the plane of the ecliptic, but long period comets can have orbits lying at any inclination. Thus the long-period comets are thought to originate from a hypothetical region of the Solar System known as the Oort Cloud. This is a spherical cloud of comets that – unlike the Kuiper belt and the scattered disc, which orbit in the plane of the ecliptic - extends in all directions from the sun out to distances of 50,000 AU.

No confirmed data is known about the Oort cloud; the immense distances and relatively small object sizes involved put the observation of any hypothetical Oort

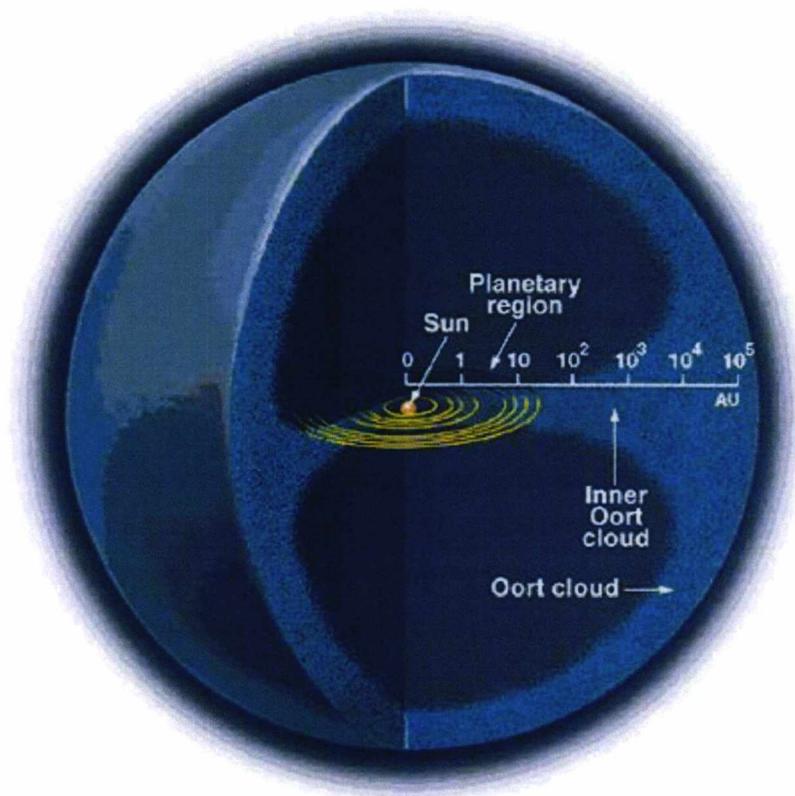


Figure 1.5.1: Schematic of the outer solar system. Note how the shape of the Inner Oort Cloud changes as it becomes less tightly bound to the plane of the ecliptic and eventually expands into the spherical Oort Cloud. Image source: http://abyss.uoregon.edu/~js/images/oort_cloud.gif.

cloud objects beyond the capabilities of today's detection apparatus. However, given that over timescales similar to the lifetime of the Solar System the long-period comets would have eroded away long ago, the fact that we are observing them now implies that there must exist somewhere in the Solar System a reservoir of these comets, and that there is some mechanism that continually replenishes the supply of long-period comets from this reservoir. As with the Kuiper belt this reasoning led the existence of the cloud to be independently hypothesised by both Ernst Öpik (Öpik 1932) and Jan Hendrik Oort (Oort 1950), and the Oort Cloud is also sometimes known as the Öpik - Oort Cloud. Several candidates have been advanced as to what this mechanism (or mechanisms) could be; the great size of the Oort cloud means that the Sun's gravitational influence on bodies contained within it would be very tenuous, and that they could be perturbed from their current orbits by interactions with external forces – for example, an encounter with another star that passes close to the Sun could easily destabilise Oort cloud objects and send them either out of the Solar System or in towards it. Additionally, since the Sun's influence on Oort cloud objects is so tenuous

the normally-negligible galactic tidal forces can have a significant effect on them; not only would the tide distort the overall shape of the cloud away from a perfect sphere, stretching it along an axis pointing towards the galactic centre, the shifting tides could also destabilise individual bodies in the cloud from their orbits and send them inwards or outwards. Finally, while the Oort cloud lies outside the area of Neptune's gravitational influence, the influence of other planets acting in concert with Neptune could be great enough to perturb Oort cloud objects inwards; this is unfortunately hard to prove due to the extreme difficulty of modelling n-body simulations. This is admittedly more of a mechanism affecting the scattered disc than the Oort cloud, but depending on where the inner boundary of the cloud lies it could still be a factor.

Since no Oort cloud objects have been observed in situ - and very few Scattered disc objects have been observed - the precise point at which the Scattered disc gives way to the Oort cloud is not clear. Nevertheless, the constant removal of material from the Oort cloud does allow certain features of its structure to be inferred - if the outer regions of the cloud are as thinly populated as believed then it should have completely dissipated over the course of the lifetime of the solar system. That it has not implies that just as the cloud itself replenishes the supply of long-period comets, the objects in the outer Oort cloud must be replenished from some nearby source.

This source is currently explained by an inner Oort cloud possessing many times more bodies than the outer cloud (Hills, 1981); this inner cloud therefore continually supplies new objects to the outer cloud to replace the ones perturbed inward or outward by external gravitational forces. Computer simulations have shown that this inner cloud extends from 3000 AU out to 20,000 AU from the Sun, and contains roughly five times as many comets as the outer cloud (Duncan *et al.* 1987). It is also more disc-shaped than the outer cloud; since the bodies in the inner cloud occupy a region closer in to the Sun its gravitational influence on them is stronger and so they orbit closer to the plane of the ecliptic.

1.6 - Centaurs

Centaurs are TNO-like objects that have orbital trajectories that cross the orbits of one or more of the giant planets, meaning that they are found in the region between Jupiter and Neptune; the first centaur-like object, 944 Hidalgo, was observed by Walter Baade in 1920, but it was the discovery of 2060 Chiron in 1977 that led to the establishment of the Centaur category. There were approximately 50 known Centaurs in 2005 (Elliot *et al.* 2005); the rate of discovery is slow since the dynamic properties of Centaur orbits make them hard to spot. Since they cross the orbits of the giant planets, the orbits of the Centaurs themselves are unstable on the timescale of several million years; the fact that we are observing them now means that, similarly to the Oort cloud, the centaurs must have a source that continually replenishes the supply over these million-year timescales. This source is believed to be either the Kuiper belt or the scattered disc, which would mean that the Centaurs are a transitional phase between a TNO being perturbed out of the far reaches of the solar system and moving inwards to become a short-period comet. This theory is supported in that the Centaurs Chiron and Echeclus have been observed to display a cometary coma (Delsanti & Jewitt, 2006) while near their perihelion; this combination of cometary properties with the physical characteristics of asteroids has led to Chiron and Echeclus being classed as both, hence the term Centaur being used to describe them.

1.7 – The Nice Model

The Nice Model (Tsiganis *et al.* 2005) is a proposed model for the evolution of the early Solar System that attempts to explain the formation and structure of the Kuiper belt, scattered disc and Oort cloud along with other events in the Solar System's history. The model posits that the large gas giants once inhabited a much more compact area of space than they do now, in near-circular orbits between 5 and 17 AU from the Sun, with Neptune orbiting closer to the Sun than Uranus. Beyond this area containing the gas giants was a large collection of both rocky and icy planetesimals that extended out to roughly 35 AU. This belt was far denser than the Kuiper Belt of today, possessing a combined mass of about 35 Earth masses.

Occasionally these planetesimals were perturbed inwards by gravitational interactions with the outermost gas giant, and as they moved inwards they exchanged angular

momentum with each gas giant that they encountered. Since the belt was more massive than today's Kuiper Belt this happened with great frequency, and while the effect that an individual planetesimal would have on a gas giant's orbit through this process would be negligible, the cumulative effect of many such planetesimals being perturbed inwards over many millions of years would cause the orbits of the gas giants to gradually shift outwards from the Sun.

Eventually, roughly 500 million years after the formation of the Solar System, this slow migration would cause Jupiter and Saturn to cross their mutual 1:2 mean-motion resonance, exacerbating their mutual gravitational effects and moving them into eccentric orbits, which has a profound effect on the rest of the bodies in the Solar System. Jupiter forces Saturn to move outwards, which in turn destabilises Uranus and Neptune outwards into highly eccentric orbits. This propels Neptune past Uranus and into the icy disc of planetesimals, scattering them both inwards and outwards.

The planetesimals scattered inwards result in a dramatic increase in the number of impacts on the inner planets, accounting for the Late Heavy Bombardment impact craters on the Moon. In the gravitational chaos caused by the migrating planets, the planetesimals are free to cross the orbits of Jupiter and Saturn almost unmolested, but once the orbits of the gas giants stabilise any planetesimals caught in this region of space are captured and become the Jupiter and Saturn Trojans (Morbidelli *et al.* 2005). In addition the gas giants also capture many irregular satellites during the migration period, accounting for their high eccentricities and retrograde orbits.

The rest of the planetesimals were mostly scattered outwards into the far reaches of the Solar System, becoming part of either the Scattered Disc or, for those bodies scattered into highly elliptical orbits, the Oort Cloud. This would explain why relatively few large TNOs have been observed past the Kuiper Cliff. Meanwhile, as Neptune moves outwards the shifting areas of gravitational resonance "sweep" up and collect any planetesimals remaining in its path into these resonances, which explains the large number of TNOs observed in resonant orbits with Neptune.

Thus the Nice Model successfully explains how several distinct features of today's Solar System came to be.

- The presence of large numbers of TNOs in resonant orbits with Neptune.
- The sudden drop-off in the number of observed TNOs orbiting past the 1:2 resonance with Neptune.
- The sudden surge in the number of impact craters on geologically dead inner solar system bodies such as Mercury and the Moon during a period roughly 700-900 million years after the formation of the solar system known as the Late Heavy Bombardment.
- The presence of satellites around the gas giants which possess irregular – highly eccentric and often retrograde – orbits.
- The various groups of Trojan asteroids found throughout the solar system.

However, there are elements of the current structure of the outer solar system that cannot be explained by the Nice Model – among them the “cold”, low inclination population of the Kuiper belt.

1.8 - Layered bodies.

As described above, given the dominant primordial elements at that distance from the sun (mainly rocky material and ices, as well as some organic matter) it is expected that the typical TNO would contain some combination of rocky and icy material. Bodies with densities approaching that of ice (0.997 g cm^{-3}) would be dominated by ice, with the rock content increasing with the density. Granite and basalt have densities of 2.7 g cm^{-3} and 3.0 g cm^{-3} respectively, so while it is possible that TNOs with densities approaching these – such as Haumea (Rabinowitz *et al.* 2005) – possess some quantity of a denser material (for example, iron) deep within their core, with a corresponding percentage of ice to compensate, it is reasonable to assume that the majority of their content is composed of rocky material.

However, the majority of the larger TNOs are assumed to have densities of around 2 g cm^{-3} ; furthermore it is assumed that the larger TNOs are non-porous unlike the majority of the cometary material making up the belt, since they possess enough mass to crush out any internal pores and voids through their own self-gravity. The TNO for which the most detailed information is available is Pluto, which has a density of $2.03 \pm 0.06 \text{ g cm}^{-3}$ (Buie *et al.* 2006). This would point to a composition of about 70% rock and 30% ice. This material is not evenly distributed throughout Pluto's structure; the radioactive decay of elements within Pluto would have heated the elements in the interior enough for them to melt. This in turn would allow the denser rock to separate itself from the lighter ice, "sinking" deeper into the interior and forming a dense core. The ice then re-freezes on top of this core. This thermal differentiation therefore produces a body with its component materials separated into a silicate core covered by a thick layer of ice. Thermal differentiation is suspected to have occurred in several medium-sized icy satellites of the gas giants such as Enceladus (Schubert *et al.* 2007) and Callisto (Nagel *et al.* 2004).

The thermal activity inside Pluto may have ceased long ago, in which case it is made up of this simple ice mantle over a rocky core. However, if it is still ongoing then it raises the intriguing possibility of the existence of a subsurface ocean between the icy mantle and the rocky core, where temperatures in the interior are high enough to melt the innermost layers of ice. There is supporting evidence for this general theory; while no probes have visited a TNO, certain of the large icy moons orbiting the gas giants – in particular Europa, Ganymede and Callisto - are good approximations of TNOs in both size and density. The Galileo probe observed perturbations in Jupiter's magnetic field caused by these moons that indicated they were behaving as perfectly conducting spheres; this can be explained if the moons possess a layer of conducting fluid at least several kilometres thick below the outer ice mantle (Carr *et al.* 1998, Zimmer *et al.* 2000); models have shown that the same may be possible for Pluto and Eris (Hussmann *et al.* 2006).

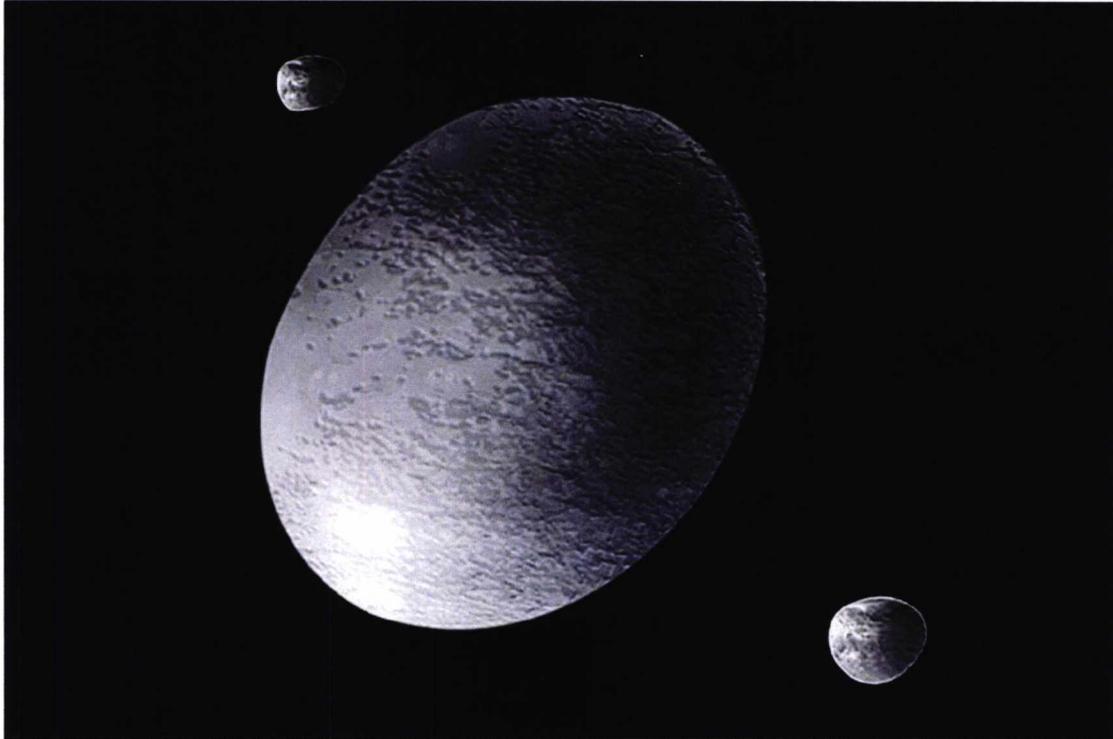


Figure 1.9.1 – Cartoon illustrating Haumea’s elongated shape. Source: http://antwrp.gsfc.nasa.gov/apod/fap/image/0809/haumea_nasa_big.jpg

1.9 - Collisional families.

The dwarf planet Haumea is an unusual example of a TNO. Orbiting at 43 AU and possessing a diameter of 1500 km and a density of $2.6\text{-}3.3\text{ g cm}^{-3}$ (Rabinowitz *et al.* 2005) it is the densest TNO currently known; such a density points to the large majority of its composition being made up of rocky material. It is unlikely that such a body would form in the outer reaches of the solar system without accumulating a substantial amount of ice, as with the other TNOs, and so something must have happened in Haumea’s past that removed all of this icy material.

Observations of Haumea’s light curve indicate that unlike the other dwarf planets Haumea is not spherical, instead being elongated into an ellipsoid that is twice as large along its long axis as its smaller ones (illustrated in figure 1.9.1). Large fluctuations in the light curve that repeat every four hours are explained by a body of Haumea’s unusual shape rotating with an orbital period of that length; this is unusual because it is a faster rotation period than any other large body so far observed in the solar system (Lacerda *et al.* 2008).

Finally, Haumea shares orbital characteristics with several other bodies – besides its moons, five other TNOs with diameters over 500 km and similar colours and IR absorption features consistent with water ice have been detected in orbits that are very close in terms of inclination and distance from the sun (within 5% of the semimajor axis and 1.4° of the inclination). While it is difficult to ascertain any concrete details about these companion objects due to their small sizes they appear to have densities that are substantially less than Haumea. Taken together, these factors point towards Haumea and its associated TNOs being members of the same collisional family (Morbidelli 2007, Brown *et al.* 2007). Haumea is thought to have originally been a typical TNO containing large quantities of both rocky and icy material, with a density of 2 g cm^{-3} , that was impacted by another TNO in its distant past – for the orbital paths of the associated objects to have spread out as much as they have since the impact, this event must have taken place at least a billion years ago. The impact removed most of the icy material in Haumea, accounting for its high density, and set the planet spinning, accounting for its unusually fast rotation. The removed icy material went on to form the bodies that make up Haumea’s current collisional family.

Haumea is a member of the Kuiper belt, and this poses a problem; as we have seen the Kuiper belt objects are all either locked into resonant or on classically stable orbits. Thus while small impacts may occur from time to time, the likelihood of an impact large enough to disrupt a significant proportion of Haumea’s mass is less than 0.1 percent. (Levison *et al.* 2008). While an impact on that scale may have occurred in the chaotic primordial Kuiper belt, the subsequent migration of Neptune that caused the belt to become dynamically stable would have disrupted the orbits of the bodies in Haumea’s collisional family and we would therefore not see them as a discrete group of TNOs today. A better candidate for Haumea’s origin is the scattered disc; here the dynamically unstable orbits of the member objects make the collision that formed Haumea’s family far more likely, as well as providing a mechanism for the family to be moved inwards to its current location in the Kuiper belt.

1.10 - Conclusion.

In this chapter I have described TNOs, the various regions in the Solar System which they inhabit and how these regions differ from one another, and also outlined a theoretical model which explains how these regions were first formed and how TNOs are transferred from one region to another. I have also described the suspected structure of the largest TNOs and evidence indicated by Haumea's family of asteroids that suggests impacts between large TNOs do occur.

Chapter 2 – Collisions

It's five miles wide... it's coming at 30,000 mph... and there's no place to hide!

- Poster tagline for the movie *Meteor*

2.0 - Introduction

Having given an overview of the existence and general properties of the large number of small icy bodies inhabiting the outer Solar System in chapter 1, we now move on to investigating what happens when these bodies collide. Examples of collisions between objects in the solar system both natural and man-made are described, and the basic physics behind these impacts explored: how the impact energy is calculated, how the internal structure of the body being impacted can affect its strength, and how scaling up from laboratory to solar system scales presents difficulties.

2.1 - Examples of collisions in the solar system

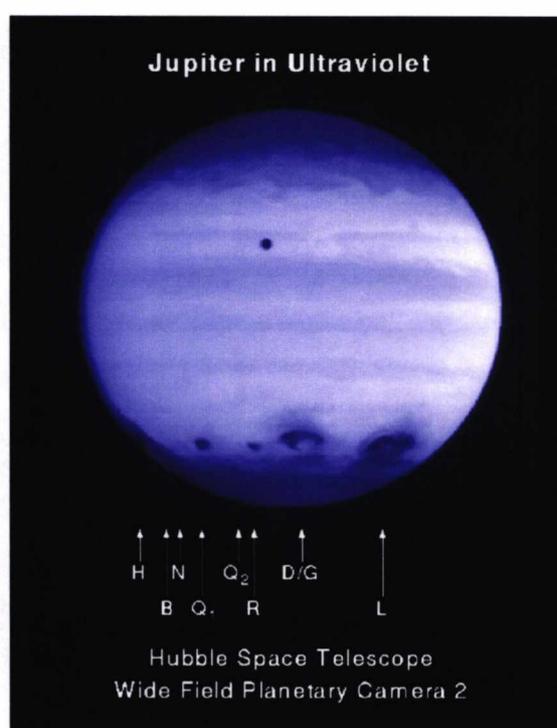
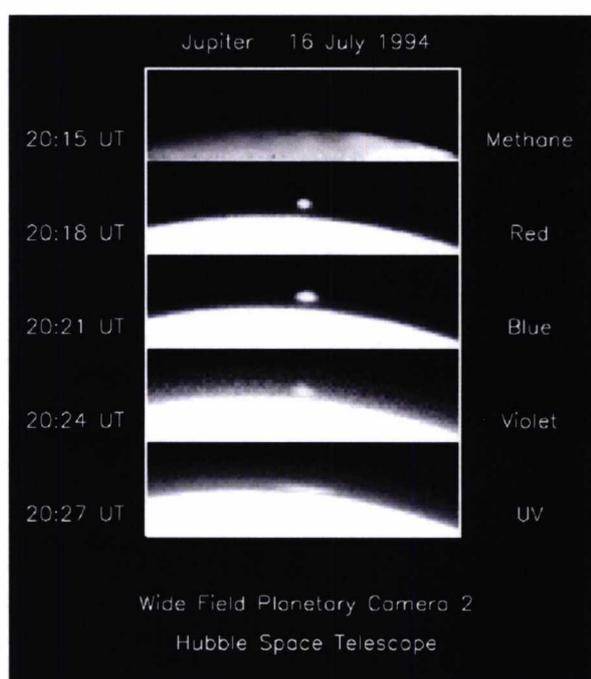
Relative to the lifetime of the solar system, collisions between solar system bodies are fairly frequent events. Relative to the lifespan of the human race however, examples of impacts large enough for us to see are few since it is only recently that we have had the technology to observe impact events on other worlds. Only one naturally-occurring impact event has been observed in detail; the rest have all been the result of a probe or probe component intentionally crashing into a body. Nevertheless, these few examples provide us with valuable data on how impacts in the solar system actually happen.

2.1.1 – Shoemaker-Levy 9

The impact of Shoemaker-Levy 9 in particular revealed much about how destructive impact events could be. Shoemaker-Levy 9 was a Jupiter-orbiting comet, and an orbit in July 1992 brought it inside Jupiter's Roche limit, resulting in the comet being pulled apart by tidal forces into a string of twenty-one separate fragments which



Figure 2.1.1 – Hubble Space Telescope image showing the 21 distinct fragments of comet Shoemaker-Levy 9 spread out over 1.1 million kilometres. Image source: <http://hubblesite.org/newscenter/archive/releases/1994/26/image/c/>



Figures 2.1.2 (left) and 2.1.3 (right) – Images from the Hubble Space Telescope showing the scale of the impacts of the Shoemaker-Levy fragments on Jupiter. 2.1.2 shows a fireball plume rising over the limb of Jupiter, while 2.1.3 shows a track of dark spots across the Jovian surface where the fragments impacted. The D/G and L impact sites are roughly 12,000km in diameter. Image credit Hubble Space Telescope Comet Team and HST Jupiter Imaging Science Team. Image sources: <http://www2.jpl.nasa.gov/s19/image129.html> and <http://hubblesite.org/newscenter/newsdesk/archive/releases/1994/30/image/a>

ranged in size from hundreds of metres up to just over a kilometre in diameter. These fragments impacted Jupiter sequentially between July 16 and July 22, 1994.

Since the comet was only observed after it had been disrupted by tidal forces this made it impossible to accurately determine Shoemaker-Levy's initial size and mass before the breakup, and thus calculate in advance the mass and impact energy that individual fragments would have. Modelling indicated that the parent body could only have a relatively narrow range of densities - between $0.3 - 0.7 \text{ g cm}^{-3}$ - in order for it to produce the observed fragment chain under tidal stresses (Asphaug & Benz, 1994), leading to an estimate of 10^{20} J for the average impact energy. However, the impacts spoke for themselves. The fragments impacted Jupiter at speeds of approximately 60 km s^{-1} ; the first fragment to strike, fragment A, was just 450 m in diameter and yet the impact produced a fireball thousands of kilometres high with a temperature of 24,000 K. The three largest fragments G, K and L were all over 1 km in diameter and contained almost three quarters of the original mass of Shoemaker-Levy (Crawford, 1996), and the results of their impacts were extremely destructive. Fragment G struck with an impact energy of $2.5 \times 10^{19} \text{ J}$ (Carlson *et al.*, 1997) and produced both a fireball plume 3000 km high and a dark spot on the Jovian surface over 12,000 km in diameter (Martin *et al.*, 1995, 1997). The results of the K and L impacts were on a similar scale.

2.1.2 - Tempel 1 and the Deep Impact mission

Deep Impact is a NASA space probe that was sent to rendezvous with the comet Tempel 1 in 2005. The probe's mission was to release a 372 kg impactor component - composed of 49% copper in order to minimise chemical reactions with water in the comet - which would collide with Tempel 1 by maneuvering into the path of the comet and allowing Tempel 1 to hit it (A'Hearn *et al.*, 2005). The purpose of this mission was twofold: the dust plume and outgassing following the collision would yield data on the composition of Tempel 1, and the impact itself was also of interest since it was not known how comets would react to impacts by other solar system bodies; observing an impact first hand was expected to yield useful information on the comet's internal structure.

The probe was launched on January 12 2005 and went into the impact phase of the mission six months later, on June 29. The impactor separated from the main probe and struck Tempel 1 at a relative speed of 10.3 km s^{-1} , delivering 1.96×10^{10} joules of



Figure 2.1.4 – The impactor component of Deep Impact strikes Tempel 1. Image credit NASA Deep Impact team. Image source <http://discovery.nasa.gov/deepimpact.html>

kinetic energy to the target. The impact threw up a large dust cloud which unexpectedly obscured the impact crater; by the time it had cleared it was impossible to pinpoint exactly where the impact site was and thus how large a crater Deep Impact had created, despite stringent searches (Busko *et al.*, 2007); the Deep Impact probe had at this point travelled too far away from Tempel-1 to get decent images of the comet. However an assessment of the expected crater size was still made based on the amount of dust thrown up by the impact. Deep Impact produced a plume containing an estimated $5.6 \times 10^5 - 8.4 \times 10^5$ kg of dust (Sugita *et al.* 2005), leading to a predicted crater diameter of 130 – 220 m (Schultz *et al.* 2007). A crater with a diameter that matches the upper range of this prediction has been observed near the impact zone on the surface of Tempel-1, but variations in possible crater morphology – instead of being a simple bowl shape the crater may instead possess an inverted “hat” shape, featuring a deep central crater surrounded by a much wider and shallower excavation zone – make it impossible to know for sure if this is indeed the Deep Impact crater. In order to locate the true crater the Stardust spacecraft has been redirected to flyby

Tempel-1 in February 2011; this will also provide information on how the topography of a comet nucleus changes after it has made a close approach to the sun.

As mentioned the Deep Impact experiment was also expected to provide clues as to the internal structure of Tempel-1. Comets are thought to be porous bodies with low densities indicating a relatively low internal strength (in comparison to a solid body) so it was presumed that any impact would not be significantly attenuated by the comet's internal structure, but rather that its own self-gravity - while itself extremely tenuous - would be the dominant regime that dictated material loss. The Deep Impact event produced a dust plume that remained attached to Tempel-1 through gravity; while this was interpreted as evidence for the comet being dominated by the gravity regime (A'Hearn et al., 2005, Schultz et al., 2005) it has been shown that in order for the extremely weak gravity of Tempel-1 to win out over its mechanical properties it would have to have a surface strength of essentially zero (Holsapple & Housen, 2006); further modelling has shown that it is possible for a strength dominated comet to produce dust plumes similar to the one observed on Tempel-1, so it remains unclear as to which factor is truly dominant in comets.

2.2 - Critical energy density

In order to quantify how strong a material is, and how much energy is needed to crater or shatter it, the impact energy density Q (measured in J kg^{-1}) is used. This is the kinetic energy of the impactor divided by the amount of mass present in the target, so:

$$Q = \frac{1}{2} \frac{M_p}{M_o} v^2 \quad (\text{Eq. 2.1})$$

where M_o is the mass of the target, M_p is the mass of the projectile striking the target and v is its velocity.

Q is useful for modelling impacts because it is size-invariant; a value of Q obtained by impact experiments on a small centimetre-scale laboratory target should *theoretically* be just as valid for predicting the behaviour of a body many metres in diameter if struck with the same relative amount of energy. In practice, there are other

considerations beyond basic material strength that affect the amount of energy required to disrupt a body which are covered later in the chapter, but Q remains a useful quantity for scaling up laboratory results to solar system scales.

There are in general two main outcomes for a body that has been impacted. Relatively low-energy impacts result in the visible effects of the impact being localised to the immediate area around the impact site; a crater forms, excavating and removing target material to a greater or lesser degree depending on exactly how much energy the impactor delivered, but while the larger part of the target might be cracked and weakened by the impact it would for the most part remain fully intact. This is a *cratering* impact.

The second outcome is an impact which *disrupts* the target. In this scenario the impactor strikes with a relatively high amount of kinetic energy, breaking the whole body of the target into multiple smaller fragments.

While there are impacts that can be seen to clearly disrupt targets and impacts that clearly crater them simply through visual inspection, there is a middle ground where the results of impacts are not so clear. An impact might have a sufficient amount of energy to knock off a significant percentage of the target's mass but leave what remains intact and unbroken – does this count as disruption or cratering? An empirical definition of the two different results is needed in this instance, and this is where the critical energy density Q^* comes in.

The effect of impacting a target with an energy density Q is judged by measuring the mass of the largest remaining fragment as a fraction of the mass of the original target. This yields a value M_f/M_o , where M_f is the mass of the largest fragment and M_o is the mass of the target. A cratering impact that removes 20% of the target's mass would produce an M_f/M_o value of 0.8, while a disruption impact that shatters a target into many small pieces might only have a largest fragment sufficient to give an M_f/M_o value of 0.05.

The critical energy density Q^* is defined as the energy density that will produce an M_f/M_o value of 0.5 (Ryan, 2000). In other words, it is the energy density that will

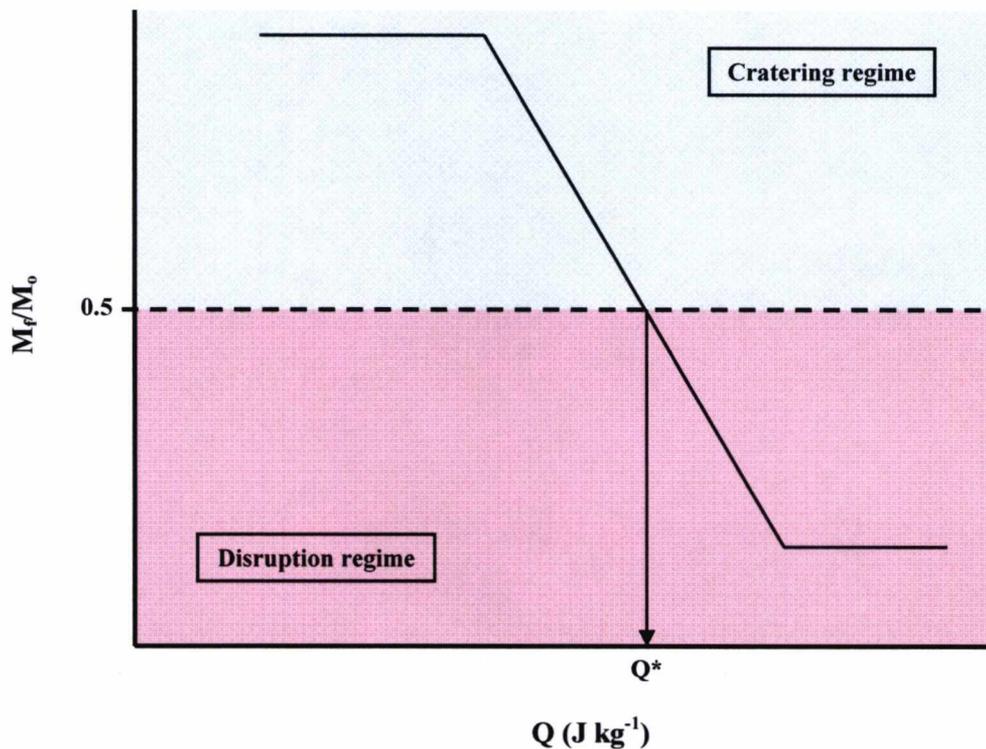


Figure 2.2.1 – Graph showing the theoretical divide between the cratering and disruption regimes, and how it is used to calculate Q^* .

remove exactly half of the mass of the target, and it marks the boundary between the cratering regime and the disruption regime. Any energy density higher than Q^* for a given material will blast away the majority of the original target's mass, and so it is defined as a disruption impact. Conversely any energy density lower than Q^* will leave more than half of the original target intact, and is defined as a cratering impact.

Figure 2.2.1 shows how the theory behind Q^* calculation works; a series of impact experiments at a range of various energy densities Q are carried out on targets composed of that material. The range of Q covers impacts that produce both cratering and disruption results. These Q are then plotted against their associated M_f/M_0 values on a logarithmic scale, which should produce a graph similar to the one shown in Figure 2.2.1. Q^* is then simply calculated as the point on the Q plot that will produce an M_f/M_0 value of 0.5.

In actuality, things are a little more complex. Rather than impacts having simple mechanical outcomes based on which side of the M_f/M_0 line they fall on, the plot of Q vs. M_f/M_0 displays different behaviour depending on the energy range – a shallow

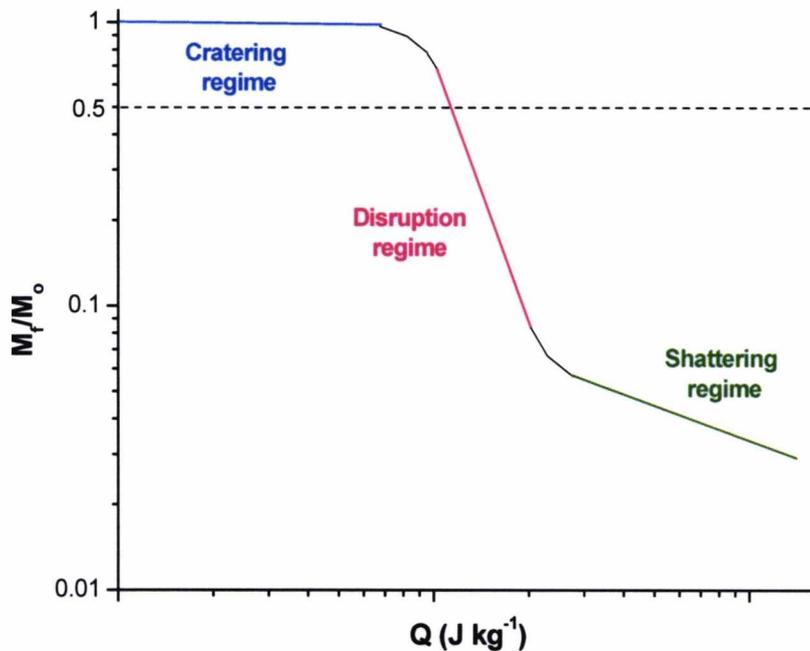


Figure 2.2.2 – Graph of a typical Q vs. M_f/M_o plot that reflects observed experimental results, including the shattering regime. However, Q^* still lies within the disruption regime so the method described above to calculate it is still valid.

decrease in mass with energy as the target is cratered, followed by a very steep decrease as the target starts displaying disruption behaviour, and finally a third regime where the plot shallows out again at high energy densities. This is the shattering regime, illustrated in figure 2.2.2, where targets are shattered into many small fragments instead of several large ones as in the disruption regime, and it will be explored further in chapters 5 and 6.

Fortunately this observed behaviour makes no difference to the method used to calculate Q^* ; semantic discussions of which parts of the plot should be classified in which regime and the existence of a new regime at very high energy densities do not change the physical behaviour of the disruption part of the curve where Q^* is to be found. It simply means that care needs to be taken when attempting to fit the disruption regime to observed data.

2.3 - Why collisions are important

Collisions are an important factor in how the solar system has evolved. It is driven by two distinct mechanisms: accretion, where one body strikes another and “sticks” to it (not quite an accurate description since there might well be some significant loss of material as a result of the collision, but the mutual gravity of the two bodies would be enough to bring the fragments back together after impact). The accretion process is responsible for building up everything in the solar system larger than a grain of dust, but the fact that not all material has been agglomerated into a few large planets and moons is explained by the second mechanism, disruption. Disruption occurs after highly energetic impacts where the two colliding bodies have high relative velocities, shattering them into many smaller fragments that do not reform under gravity. Disruption is responsible for “grinding” collections of bodies down into ever-smaller fragments through successive disruption impacts amongst each other.

The formation of the terrestrial planets is a process that was primarily governed by accretion. The widely-accepted model for their formation is the planetesimal hypothesis (Chambers 2004). This model follows the nebular hypothesis closely; the young Sun is surrounded by a large disk of fine dust grain particles. During the first million years of the solar system’s lifetime these particles stick together, forming numerous small planetesimals ranging from 1-10 km in radius, which then not only sweep up the remaining dust particles but go on to collide and accrete into a few tens of larger proto-planets hundreds of kilometres in radius. These proto-planets are large enough to perturb each other’s orbits drawing each other into chaotic interactions. The proto-planets are either expelled from the solar system by these interactions or collide over a timespan of 10-100 million years to form the terrestrial planets we know today.

Conversely, the asteroid belt is a region of space that has been governed mainly by disruption since very shortly after the birth of the solar system. The asteroid belt is a collection of bodies that never made it past the planetesimal stage; the formation of proto-Jupiter occurred early enough for it to gravitationally perturb the bodies in the belt. This not only ejected much of the mass from the belt but also bumped the velocities of the remaining bodies above the upper limit required for accretion and

into the disruption zone. As a result, very few collisions in the belt result in accretion; typical impact speeds between asteroids are on the order of a few kilometres per second (O'Brien & Greenberg, 2005) which in the majority of cases results in both bodies being disrupted. This means that the majority of the present day population of the asteroid belt are not the original asteroids formed in primordial times, but instead their fragmented remnants.

Similarly the population of the Kuiper Belt has been shaped by both mechanisms during its lifetime. At first the planetesimal theory applies as it did for both the asteroid belt and the terrestrial planets – bodies with radii $\leq 1-10$ km collide and accrete into larger objects. However, as these proto-planets grow so does their gravitational influence and they stir up the orbits of the smaller planetesimals in the belt, increasing their velocity past the disruption limit. The end result is that once a few large bodies exist in the belt they make it impossible for any more to form and facilitate the collisional grinding of the remaining material into smaller and smaller fragments.

2.4 - Factors that affect the disruption energy

It is not simply the strength of the material being impacted that determines the amount of energy required to disrupt it. There are several other factors that affect the critical energy density.

From the perspective of an impact scientist the solar system would be a much simpler place if it consisted solely of bodies that were simply homogenous lumps of rock or ice. Unfortunately we are not so lucky; due to the processes that form small solar system objects and which affect them over their lifetimes they tend to possess relatively complicated internal structures that could alter the amount of energy required to disrupt them to an unknown degree.

2.4.1 – Repeated impacts

Homogenous bodies are the simplest type of body used for modelling and lab

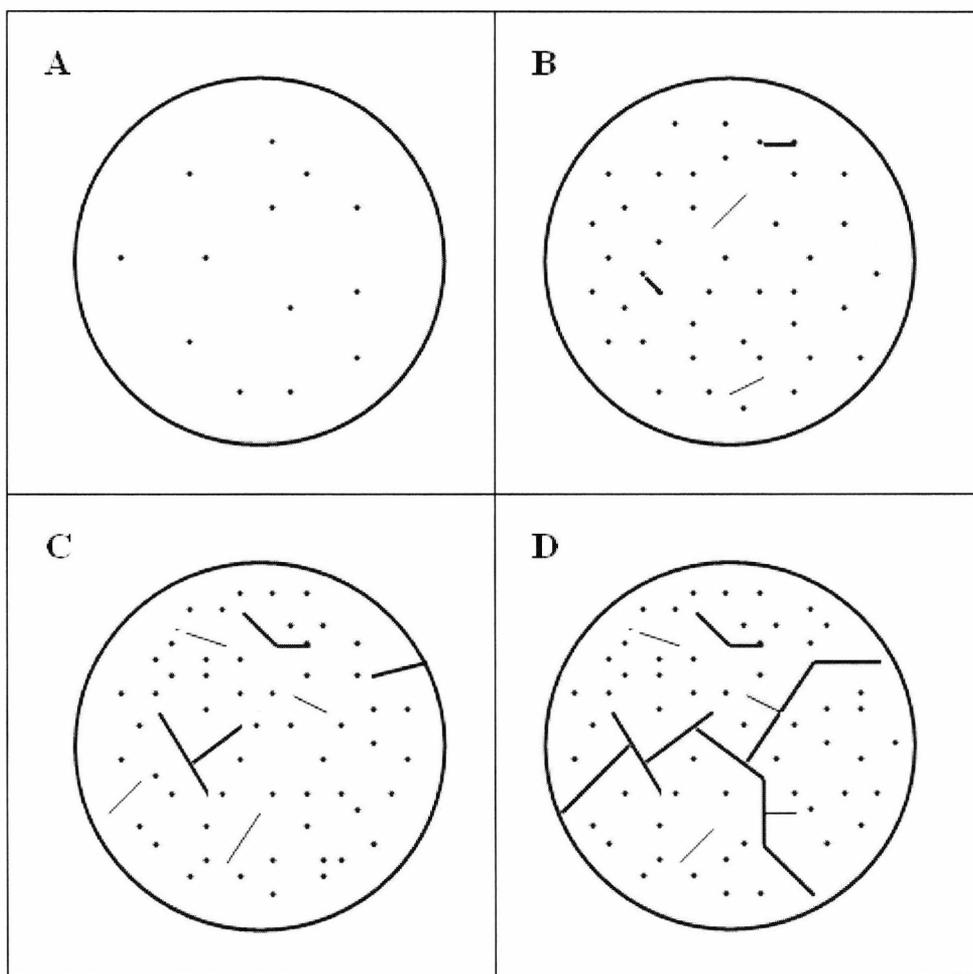


Figure 2.4.1 – Diagram showing the various stages of crack propagation in a homogenous target.

experimentation, and as said they are highly unlikely to occur naturally in the solar system, being used mostly for studies of how a single material or target shape responds to impacts. However, even these simple laboratory examples have their own quirks of behaviour that make predicting disruption outcomes difficult. Successive weak impacts can shock and weaken a target to the point where it can be disrupted by a impactor possessing less energy than Q^* .

While homogenous bodies are initially constructed of a single material type evenly spread throughout the body, relatively weak impacts can shock this material so that tiny flaws form in the body structure. It is through the creation and exploitation of these microflaws that a body is eventually disrupted by successive weak impacts (Bless & Rajendran, 1996).

Figure 2.4.1 provides a rough illustration of how microscopic flaws in a target's structure can combine to disrupt a target after successive sub- Q^* impacts. In part A, the target is unimpacted, and has few flaws in the structure. A sub- Q^* impact, while removing some target mass, will not disrupt the body. However, the impact shocks material throughout the target, and this creates more of the microscopic flaws, shown in Part B. Flaws which form very close together join up into microscopic cracks and fractures; at this stage they are very small, but another sub- Q^* impact will shock the target and increase the number of flaws yet again. In Part C the target is peppered throughout its structure with the flaws, and the fractures formed by nearby flaws joining together are starting to become quite large. A final sub- Q^* impact increases the number of flaws in the target to a critical density; Part D shows that there are now enough flaws present in the target for a single crack to propagate all the way through the target, splitting it into two parts. Despite all three impacts having less than the critical energy density, the target has now been disrupted thanks to repeated impact shock.

In conclusion, flaws in homogenous structures make disrupting them easier with each successive impact. In addition, an overabundance of flaws can lead to a target cracking in the process described above during a single impact – this has particular ramifications as the size of a body is increased, as described below.

2.4.2 - The size of a body.

Extrapolating the results of laboratory impacts on centimetre- or metre-scale targets up to solar system scales is not simply a matter of multiplying the critical energy density by a larger mass. The size of the body being impacted does have an effect on the amount of energy required to disrupt it - as a body is made larger it gains more mass, and bodies larger than a few hundred metres in diameter possess enough mass for the body's self-gravity to have a significant effect on the critical energy density. In order for the body to be disrupted a hypothetical impactor would not only have to deliver enough energy to shatter it, but it would also have to impart sufficient additional energy to accelerate the fragments to speeds greater than the body's escape

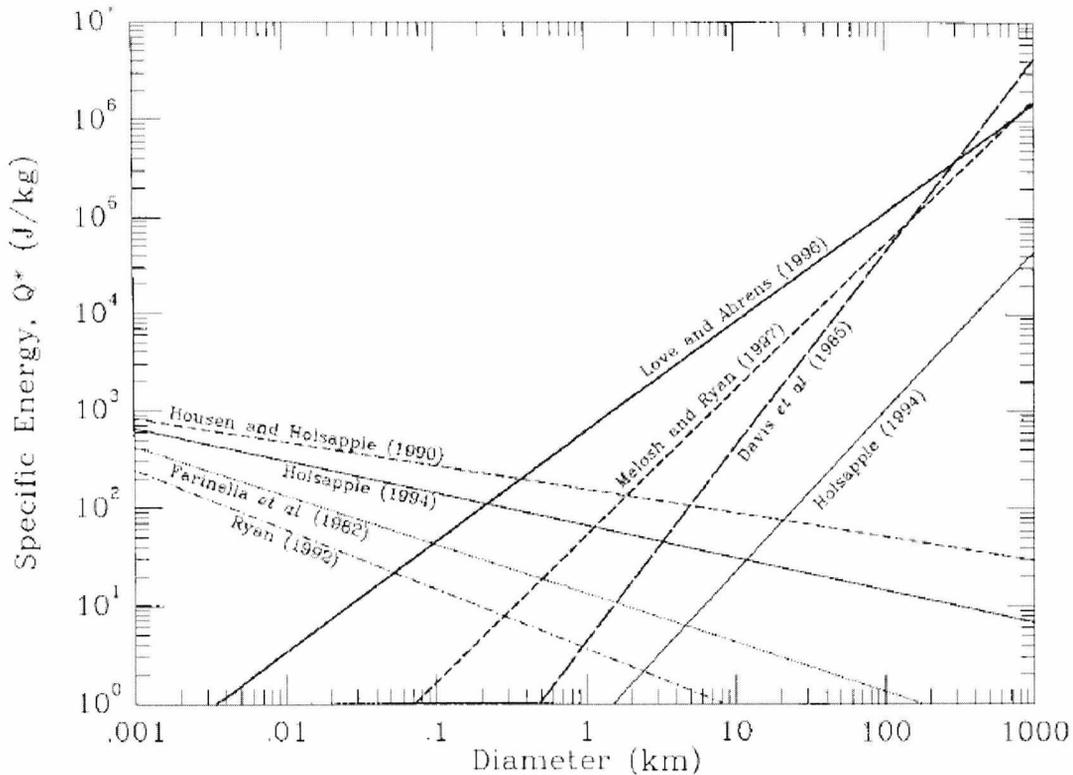


Figure 2.4.2: Graph showing the two different sets of scaling laws; material strength starts out dominant at small scales but decreases steadily as the size of the body increases, while the effect of gravity is negligible to start with but quickly becomes dominant at larger size scales. (Durda *et al.* 1999).

velocity. Any less than this and the fragments would simply fall back towards the centre of mass, eventually reforming the body anew.

Conversely, as the size of a body is increased then flaws within the material composing the body begin to have a cumulative effect. For example, one tiny fracture in a centimetre-scale target would not affect the energy required to disrupt it in any significant way. However, increasing the scale of the target also increases the number of these fractures present inside the target material; where the centimetre body contained only one or a few, a body with a diameter on the scale of tens or hundreds of metres would possess thousands of these flaws within its internal structure. This makes it much easier for cracks to propagate from fracture to fracture, thus decreasing the overall amount of energy required to disrupt the body.

These two different effects are modelled using two different types of scaling law, as shown in figure 2.4.2. At small target diameters the material strength of the target is

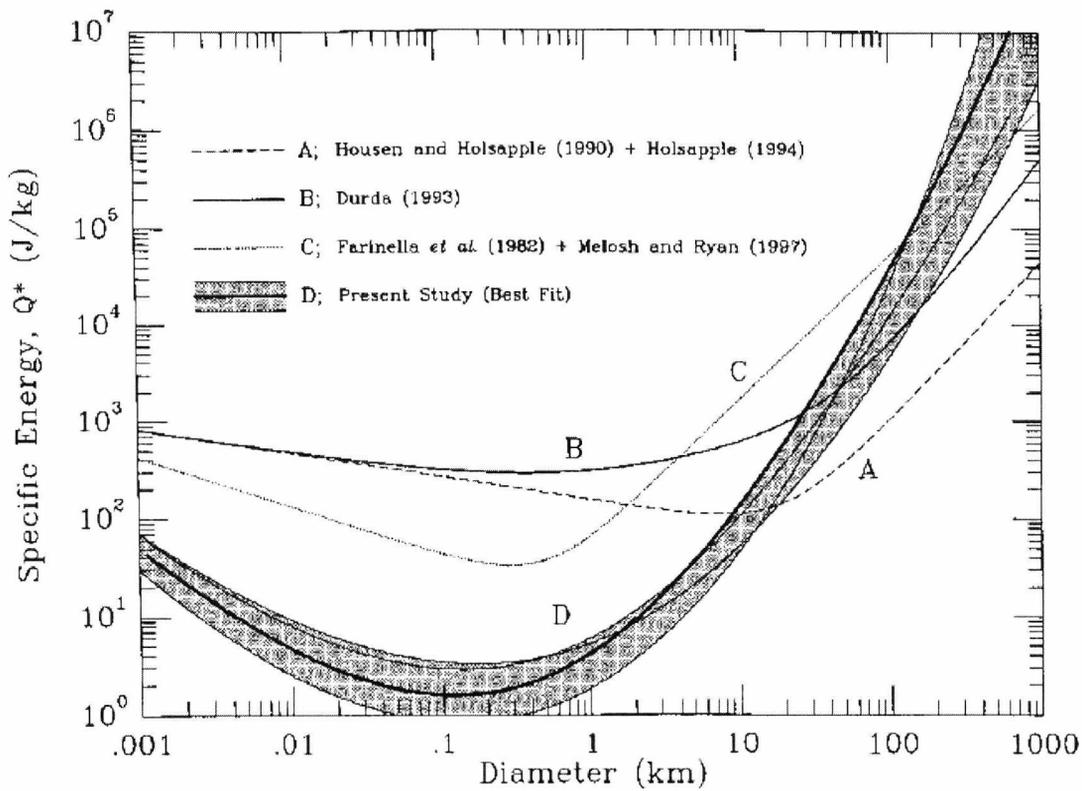


Figure 2.4.3 – Graph showing the material and gravity scaling laws combined to form one coherent picture of how the size of a body affects the critical energy density. (Durda *et al.* 1999)

the dominant factor in determining Q^* , and so the target is said to fall within the strength-scaling regime. For bodies with diameters of tens to hundreds of kilometres the increased energy requirements for disruption caused by a body's self-gravity far outweigh the effects of scaling on material strength, and so the body falls within the gravity-scaling regime.

Where exactly the strength regime gives way to the gravity regime is unclear since several different scaling models exist for each regime – as can be seen in Figure 2.4.3 - creating an intermediate zone for bodies with diameters ranging from 100 metres to several kilometres where it is unknown as to which regime is really dominant. However, the picture these two scaling regimes paint is fairly clear. At first, increasing the size of a body from the laboratory scale will make it easier to disrupt since it falls within the strength regime and gravity effects are negligible. However, once a body reaches the diameters in the intermediate zone the effect of size on material strength begins to be counterbalanced by the body's self gravity, which



Figure 2.4.4 – Image of the comet Tempel-1 taken five minutes before the impactor component of the Deep Impact probe struck it. Image source: http://deepimpact.umd.edu/gallery/jpg/ITS_PR2-PIA02127.jpg

2.4.3 - Porous bodies.

eventually becomes by far the dominant factor at diameters of more than a few kilometres and increases the critical disruption energy far beyond that required to disrupt the base material that comprises the body.

Many of the smaller solar system objects and TNOs are not solid bodies at all. They are instead rough agglomerations of material that has accreted onto the body over time and provided it with an irregular, porous internal structure. The best example of this kind of body are the numerous comets orbiting the sun, which were famously defined as “dirty snowballs” by Fred Whipple in 1950 (Whipple 1950). This is a fairly apt description since comet nuclei are a loosely-packed collection of ices frozen together with dust; voids exist in the internal structure, and comets are not massive enough for them to be squeezed out by self-gravity.



Figure 2.4.5 – Image of the asteroid 25143 Itokawa taken by the Hayabusa mission in 2005. Itokawa is an example of a rubble-pile asteroid, and its nature as a rough agglomeration of rock is clearly visible. Image source: http://www.isas.jaxa.jp/j/snews/2005/1101_hayabusa.shtml

Comets are not the only porous bodies in the solar system; recent observations of the asteroid 25143 Itokawa revealed it to have a similar irregular “rubble pile” makeup (Fujiwara *et al.*, 2006), and many other asteroids – such as 253 Mathilde (Cheng 2003) - are suspected to be porous based on their measured density and their observed composition of rock; Itokawa has a density of 1.96 g cm^{-3} (Abe *et al.* 2006) but the density of rocky material ranges from 2.6 g cm^{-3} to 3.0 g cm^{-3} for heavy silicates such as granite and basalt. These porous asteroids are thought to be formed from the shattered remnants of a previous asteroid that was originally composed of solid rock but was impacted and disrupted by another body; the resulting fragments then reformed into a single body under their own self-gravity (Michel *et al.*, 2001).

Comets and rubble-pile asteroids are examples of two different types of porous body. Comets are more-or-less solid objects frozen together with many internal voids, like a sponge. Rubble-pile asteroids, on the other hand, have no internal tensile strength (except for the internal strength of the individual fragments making up the asteroid); their own self-gravity is the only thing holding them together.

Since both types of porous bodies are riddled with voids this decreases their overall density. A very rough rule for solid bodies would be that the denser they are, the harder they are to disrupt. A low density solid object would require a comparatively low amount of energy to disrupt. However, this is not true for porous bodies; the irregular structure and the presence of the voids actually *increases* the amount of energy required to disrupt it in comparison to a solid body made of the same material. In rubble-pile type bodies the boundaries between the separate fragments composing the body and the voids that are found throughout it inhibit the transmission of impact shockwaves, whereas in sponge-like porous bodies a lot of energy is absorbed in compressing the void spaces (Love *et al.*, 1993). In effect, the energy is transferred from the impacting body to the fragments immediately surrounding the impact site, but the voids and discontinuous boundaries between individual fragments inhibit the transmission of shock energy from those initial fragments to fragments deeper within the targets. Much of the kinetic energy of the target that would otherwise go towards disruption is therefore confined to the first fragments, heating them and causing them to spall. This localised concentration of energy means that while impacts on porous bodies create larger craters and cause more melting around the crater area, outside of the immediate impact zone the effects of the impact are much reduced.

2.4.4 - Layered bodies.

That some objects in the solar system possess a layered structure is not exactly a surprise; the Earth itself is made up of several layers of rocky crust, brittle outer mantle, soft inner mantle, and then a spinning liquid iron core surrounding a solid iron core. It had been suspected for some time that certain of the larger icy bodies in the solar system – the moons of Saturn and Jupiter, and probably the larger TNOs such as Eris and Pluto – would also possess a layered structure due to their size, and the flyby of the Jovian moons by the Galileo probe in 1995-96 provided evidence that supported this hypothesis (Showman *et al.* 1999). The mechanism that provides bodies with a layered structure has already been covered in Chapter 1, but what is not known is how this layered structure might affect Q^* for a body as opposed to a homogenous target made up of the same materials and with the same bulk density.

Table 2.5.1 – List of previous Q^* calculations for homogenous and porous icy bodies.

Study	Target (porosity)	Projectile	Projectile velocities	Q^* ($J\ kg^{-1}$)
Arakawa (1999)	Solid ice	Nylon	2.3 - 4.7 $km\ s^{-1}$	40
Murray (2004)	Solid ice	Copper	1.0 - 5.1 $km\ s^{-1}$	4.77
Fujiwara (1977)	Basalt	Polycarbonate	2.6 $km\ s^{-1}$	700
Ryan (1999)	Porous ice (30-45%)	Fractured ice	73 - 308 $m\ s^{-1}$	50
Giblin (2004)	Porous ice (10%)	Ice	90 - 155 $m\ s^{-1}$	22.6
	Porous ice (37%)	Ice	90 - 155 $m\ s^{-1}$	39
	Porous ice (55%)	Ice	90 - 155 $m\ s^{-1}$	73.4
Arakawa (2002)	Porous ice/pyrophyllite (39%)	Ice	150 - 670 $m\ s^{-1}$	40
	Ice/pyrophyllite (0%)	Ice	150 - 670 $m\ s^{-1}$	100
Arakawa (2004)	Porous ice/serpentine (32%)	Ice	150 - 670 $m\ s^{-1}$	60
	Ice/serpentine (0%)	Ice	150 - 670 $m\ s^{-1}$	121

2.5 - Previous calculations of Q^*

Having established in Chapter 1 that the larger TNOs are composed of an ice-silicate mix and possess probable layered internal structures, it becomes clear that in order for experimental results to be relevant to theoretical models of TNO collisional evolution they have to use targets that incorporate these two features since both could be factors that affect the value of Q^* for that type of body to a greater or lesser degree.

Previous work that directly calculates Q^* has tended to focus on homogenous or porous targets constructed of a single material type – mainly pure ice or rocky material. Arakawa (1999) impacted cubic pure water ice targets up to 100 mm in size with nylon projectiles at velocities of between 2.3 and 4.7 $km\ s^{-1}$. This produced a Q^* for ice of 40 $J\ kg^{-1}$. Previous work undertaken at the University of Kent by Murray (2004) disrupting solid ice spheres with copper projectiles calculated Q^* for solid ice as 4.77 $J\ kg^{-1}$, which is much lower; however it must be taken into account that copper is almost eight times denser than nylon and may prove to be more efficient at disrupting targets. Fujiwara *et al.* (1977) performed impacts on basalt targets using polycarbonate projectiles at speeds of 2.6 $km\ s^{-1}$ and obtained a Q^* of 700 $J\ kg^{-1}$. This indicates that silicate material is at least an order of magnitude stronger than ice under hypervelocity impact.

Ryan *et al.* (1999) performed low speed (73 to 308 m s⁻¹) impacts on porous ice targets with porosities ranging between 30 to 45% using a variety of projectiles (aluminium, solid ice and fractured ice). They observed that the denser aluminium projectiles caused more damage to the targets than the ice/fractured ice indicating a dependence of Q^* on projectile material/size; assuming that impacting bodies were likely to be composed of ice and not a solid mass of aluminium they obtained a value of Q^* for porous ice on porous ice impacts of 50 J kg⁻¹. Giblin *et al.* (2004) similarly carried out a series of low speed impact experiments (projectile velocities ranged from 90 to 155 m s⁻¹) on porous ice targets using solid ice projectiles. The targets used had varying porosities and it was observed that Q^* increased as the porosity increased; at 10% porosity the targets had a Q^* of 22.6 J kg⁻¹ but this increased to 73.4 J kg⁻¹ at 55% porosity. Q^* for a body with 37% porosity was 39 J kg⁻¹ indicating fairly good agreement with Ryan's result. Arakawa *et al.* (2002) performed impacts on mixed ice-silicate porous bodies mixing the ice in a 50:50 ratio with pyrophyllite, a clay mineral with a density of 2.7 g cm⁻³. These mixed porous bodies were impacted at velocities of 150 to 670 m s⁻¹ using ice projectiles, and produced a Q^* of 60 J kg⁻¹. The same process was carried out using serpentine in place of pyrophyllite (Arakawa *et al.* 2004) which has a lower density of 2.6 g cm⁻³ and this time produced a Q^* of 40 J kg⁻¹. Intriguingly, Arakawa finds that the strength of mixed porous bodies decreases as porosity increases, a reverse of the relationship observed in pure ice targets.

These results for porous bodies confirm that, for pure ice porous bodies at least, critical energy density does increase with porosity. A result in the Giblin paper for a body with only 10% porosity put the associated Q^* value at 21 J kg⁻¹; this along with the Arakawa result for solid ice cubes (40 J kg⁻¹) indicates that the critical energy density for ice is on the order of tens of J kg⁻¹. Arakawa also carried out impacts on an ice-silicate body with 0% porosity; these provided Q^* values of 121 J kg⁻¹ (2004) and 100 J kg⁻¹ (2002).

There are few examples of laboratory impact experiments on layered targets, possibly due to the difficulty involved in producing targets with the layered structure. Okamoto *et al.* (2008) impacted layered targets composed of a soda-lime glass or quartz core covered by a porous gypsum mantle with nylon projectiles at velocities of between 1

and 5 km s^{-1} . They observed an increase in target strength as the thickness of the gypsum layer was increased. The base critical energy density of the glass used is stated as 600 J kg^{-1} while the critical energy density of gypsum is 2000 J kg^{-1} , so it is logical that increasing the mass of gypsum present in relation to the weaker core will increase the strength of the target as a whole. However, it was also observed that at certain core/mantle mass ratios (between 80-85% gypsum by mass) the core-mantle target exceeded the impact strength of pure gypsum. This suggests that there are certain layer thicknesses which are extremely effective in increasing a body's total critical energy density.

These results, then, give us a rough indication of what we would expect to see in an experimental programme.

- A critical energy density for solid ice on the order of tens of J kg^{-1} .
- An increased critical energy density for an ice-silicate mixed body on the order of hundreds of J kg^{-1} .
- A critical energy density for layered bodies that exceeds that for a homogenous body made of the layered body's strongest material at certain layer thicknesses.

2.6 – Conclusion

In this chapter I have laid the theoretical groundwork for an exploration of how icy bodies disrupt. Real examples of hypervelocity impacts in the solar system and their destructiveness are discussed, and the relevance of impacts between solar system bodies to the evolution of the Kuiper belt and other collections of rocky/icy material is explained.

Chapter 3 – Disrupting Targets in the Laboratory

“Princess Leia, before your execution, you will join me at a ceremony that will make this battle station operational. No star system will dare oppose the Emperor now.”

“The more you tighten your grip, Tarkin, the more star systems will slip through your fingers.”

“Not after we demonstrate the power of this battle station. In a way, you have determined the choice of the planet that is to be destroyed first. Since you are reluctant to provide us with the location of the Rebel base, I have chosen to test this station's destructive power on your home planet of Alderaan.”

- Grand Moff Tarkin and Princess Leia, *Star Wars*

3.0 - Introduction

This chapter will cover the methods used to create targets that accurately simulate large icy Kuiper Belt objects, and the method used to accelerate projectiles up to speeds of several kilometres per second in order to disrupt them. Various techniques used to investigate the materials used in making the targets are also described, including strength tests and sand grain measurements.

3.1 - Light gas guns

Accelerating projectiles up to impact velocities on the order of several kilometres per second is no easy task. Regular powder-burning guns are limited in that there is a sizable mass of propellant gas that must be accelerated along with the projectile; since the powder burns at an uneven rate over small time scales this creates a large pressure gradient in the column of gas behind the projectile. The front of the gas column accelerates faster than the rear of the gas column, which dissipates the gas kinetic energy and results in relative inefficiency when transferring this energy to the projectile, creating a speed ceiling of approximately 1.2 km s^{-1} .

Two-stage light gas guns have been developed with the specific purpose of getting around these problems and firing projectiles at much faster speeds than the range presented by powder guns (Crozier & Hume, 1957). Rather than accelerating the propellant gas and the projectile at the same time, a light gas gun separates the gas

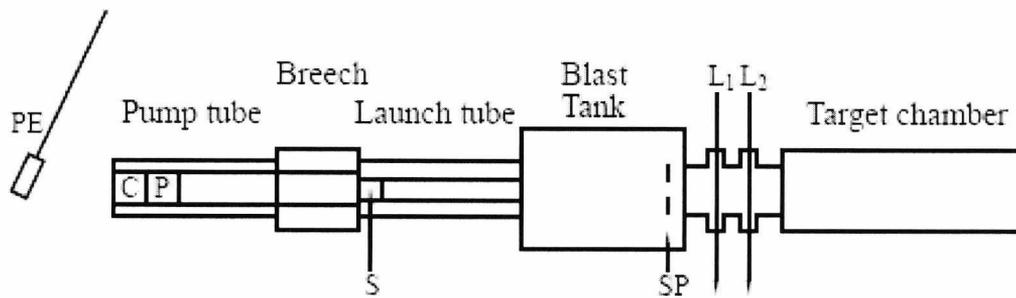


Figure 3.1.1: The typical firing mechanism of a light gas gun (Burchell *et al.* 1999). Legends are as follows: PE represents the pendulum, C is a powder charge (contained within a shotgun shell) which drives P, the piston. The piston compresses a light gas in the pump tube which, when released through the breech, accelerates a projectile contained within a sabot S down the launch tube into the blast tank. The sabot is spun by rifling in the launch tube and flies apart into four separate pieces which travel off-axis and hit the stop plate SP. The projectile carries on through a hole in the centre of the stop plate and through the laser curtains L1 and L2 which measure the projectile speed. Finally the projectile enters the target chamber and strikes the target.

and the projectile into two stages. The propellant gas is compressed in the first stage so that it reaches an optimum pressure before being released. This ensures that when the gas is finally released into the second stage containing the projectile its acceleration down the barrel of the gun is uniform and transfers as much kinetic energy to the projectile as possible.

A typical light gas gun is shown in Figure 3.1.1. Generally, light gas guns work by the following principle: a powder charge C is ignited by a firing pin, accelerating a piston P down the pump tube. The pump tube is sealed at one end by a burst disc and filled with a light gas which the accelerating piston compresses. Once the light gas has been compressed enough to reach a sufficiently high pressure the burst disc ruptures, allowing the gas to expand into the launch tube. It is this rapid expansion along with the still-accelerating piston pushing the gas reservoir into the launch tube that accelerates the projectile contained within the sabot S to hypervelocities.

We can consider a simple model of how the gas expands and transfers energy to the sabot containing the projectile. The average molecular kinetic energy of a gas is given by the equation

$$KE_{avg} = \frac{3}{2}kT \quad (\text{Eq. 3.1})$$

When free to expand, its expansion energy will therefore be roughly equivalent to this molecular kinetic energy.

$$\frac{1}{2}mv^2 = \frac{3}{2}kT \quad (\text{Eq. 3.2})$$

Assuming the temperature remains constant, this means that the expansion velocity of the compressed gas will be inversely proportional to the square root of the mean molecular weight of that gas, as expressed in the following equation.

$$v \propto m^{-\frac{1}{2}} \quad (\text{Eq. 3.3})$$

From this equation it can be seen that a lower molecular weight will result in a higher expansion velocity – in other words, the lighter the gas, the greater the velocity that can be achieved with the gun. It is for this reason that light gases such as hydrogen or helium are used as part of the acceleration mechanism. Using this method light gas guns are capable of accelerating projectiles up to velocities of 8.9 km s^{-1} (Moritoh *et al.* 2001).

3.2 - University of Kent's Two-Stage Light Gas Gun.

The two-stage light gas gun at the University of Kent (Burchell *et al.* 1999) works using the principles outlined above and is capable of accelerating millimetre-sized projectiles to velocities in the range of 1 km s^{-1} to 7 km s^{-1} . The fastest shot recorded on this gun was 8.4 kms^{-1} using a thin slice of plastic as the projectile, but for practical shots that utilise sabots the maximum velocity that can be achieved reliably is 7.5 kms^{-1} .

The cartridges used to provide the initial powder charge are standard 20mm diameter shotgun cartridges, each one of which is loaded with 10g of gunpowder in the laboratory. The pump tube has a bore of 12.7mm diameter and a length of 0.7m, while

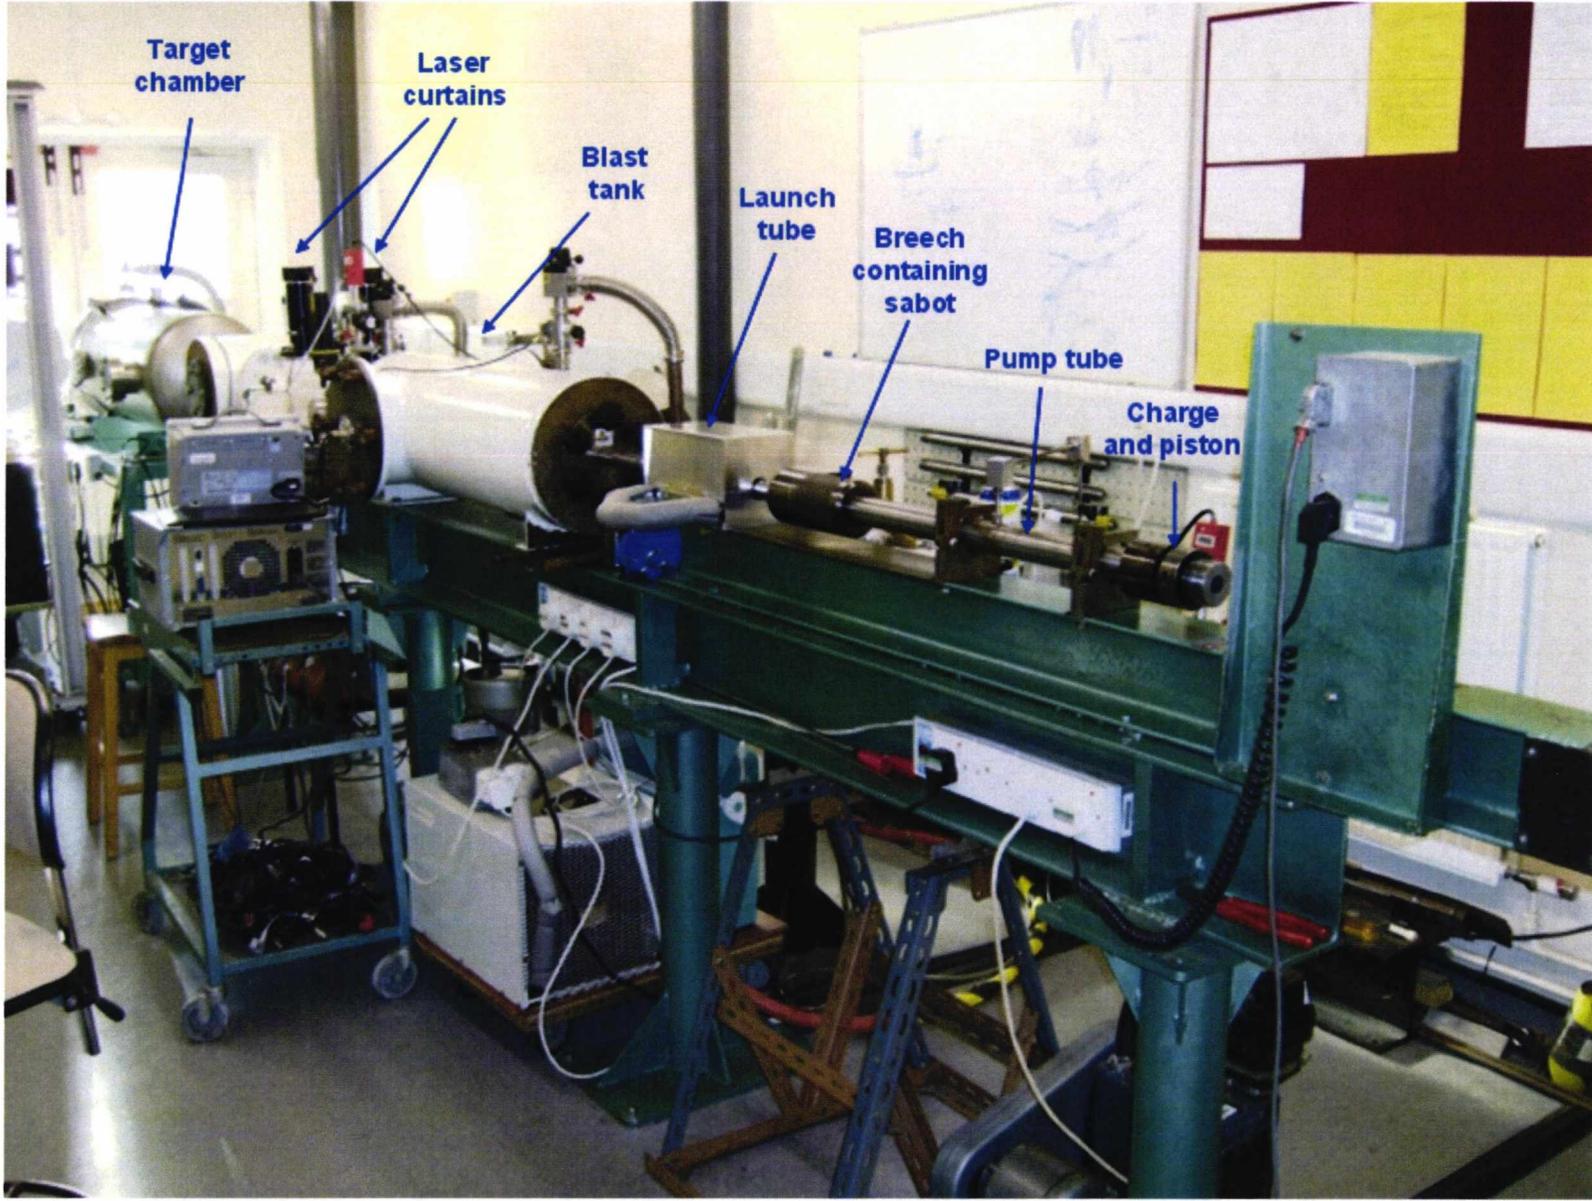


Figure 3.2.1 – The two-stage light gas gun at the University of Kent.

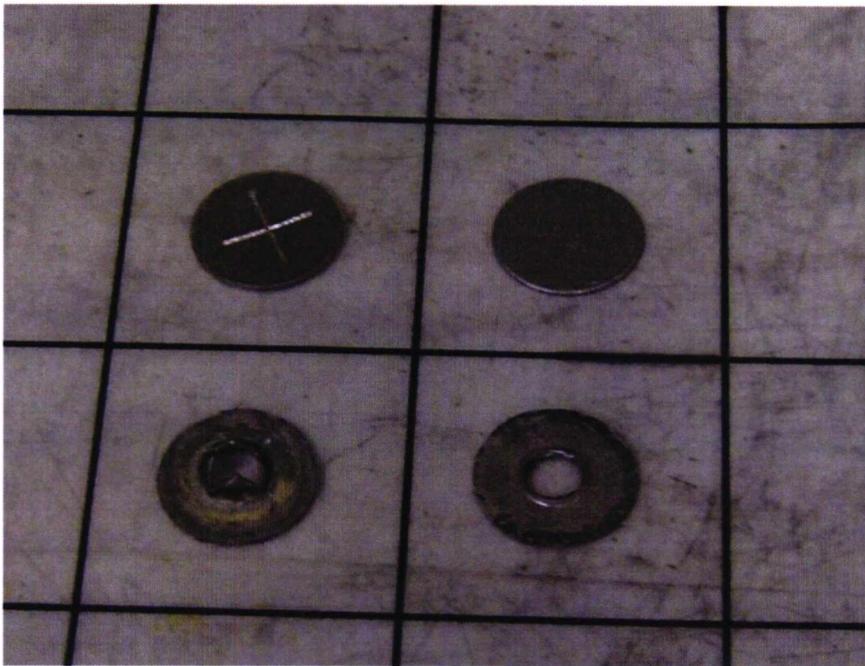


Figure 3.2.2: Examples of the two types of burst disc before and after firing. The burst discs have diameters of 12.7 mm. Top left is a scored burst disc and below it is a scored burst disc after firing with the “petals” visible. Top right is a plain unscored burst disc and below is an unscored burst disc after firing with a hole neatly punched through the middle. Grid squares are 2.55 cm x 2.55 cm.

the piston is made of nylon, has a mass of 12g and a length of 8cm and is also 12.7mm in diameter. While this is a snug fit for the piston it is still not fully airtight, so rubber O-rings are fitted in grooves on the piston at 6 and 6.5cm along from the cartridge end to ensure that no gas can escape.

The pump tube (see figures 3.1.1 and 3.2.1) is separated from the launch tube by a breech into which an aluminium bursting disc has been inserted. The discs are 12.7 mm in diameter and 0.59 mm thick. Two different types of burst disc are used depending on the speed of the shot. For shots up to 5 km s^{-1} , a burst disc made of aluminium 2014 is used. This burst disc is scored in the centre in the shape of a cross, so that when the burst disc ruptures it splits open along the score lines, much like the petals on a flower. For shots above 5 km s^{-1} a plain unscored burst disc of aluminium 7075 is used; in the case of high velocity shots the gas pressure is sufficiently high to punch out the centre of the disc without the need for scoring. The launch tube’s bore is 4.30 mm in diameter and 0.7m long, and possesses internal rifling that completes one turn along its length.

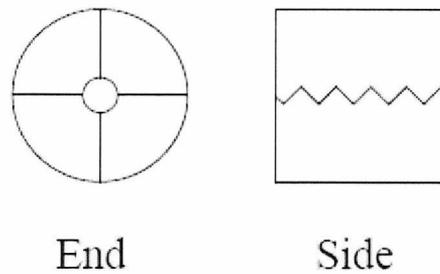


Figure 3.2.3: Diagram illustrating sabot construction (Burchell *et al.* 1999). The projectile is contained in the centre of the sabot pieces (end-on diagram). The sabot pieces are designed to lock together as one while enclosed by the launch tube, but upon exiting the launch tube spin imparted to the sabot by rifling causes the four pieces to separate and fly away from the projectile.

While the burst discs are far too large to be propelled down the launch tube entirely (the neatly punched-out hole in the centre of the bottom right burst disc in figure 3.2.2 is the size of the launch tube diameter), occasionally a fragment of the burst disc (such as one of the “petals” shown in figure 3.2.2) will manage to follow the projectile down the tube. In this case, if it continues to travel on-axis and makes it through the stop plate it will hit the target. While this ruins the shot, it is easy to spot when this has happened – firstly because there will be a second object passing through the light curtains that will show up on the oscilloscope trace used to measure projectile speeds, and secondly targets struck by burst disc fragments are often unexpectedly pulverised thanks to the extra impacting mass. If an impact has been unusually destructive it is likely that a piece of the burst disc is the culprit.

Projectiles are typically placed inside a sabot before firing. The sabots used in the gun are made of nylon and are 4.3mm in diameter and 4.5mm in length. While the sabots can be constructed of a single solid piece of nylon if required, typically they are split sabots cut into four parts with serrated edges that lock together preventing any movement of individual pieces while the sabot is in the launch tube and ensuring that it is accelerated as a single mass. The rifling present in the launch tube makes the sabot spin as it travels down it, and so when the sabot exits the launch tube into the blast tank the individual sabot pieces are no longer constrained by the walls of the launch tube and will travel off-axis at an angle of approximately 6° .

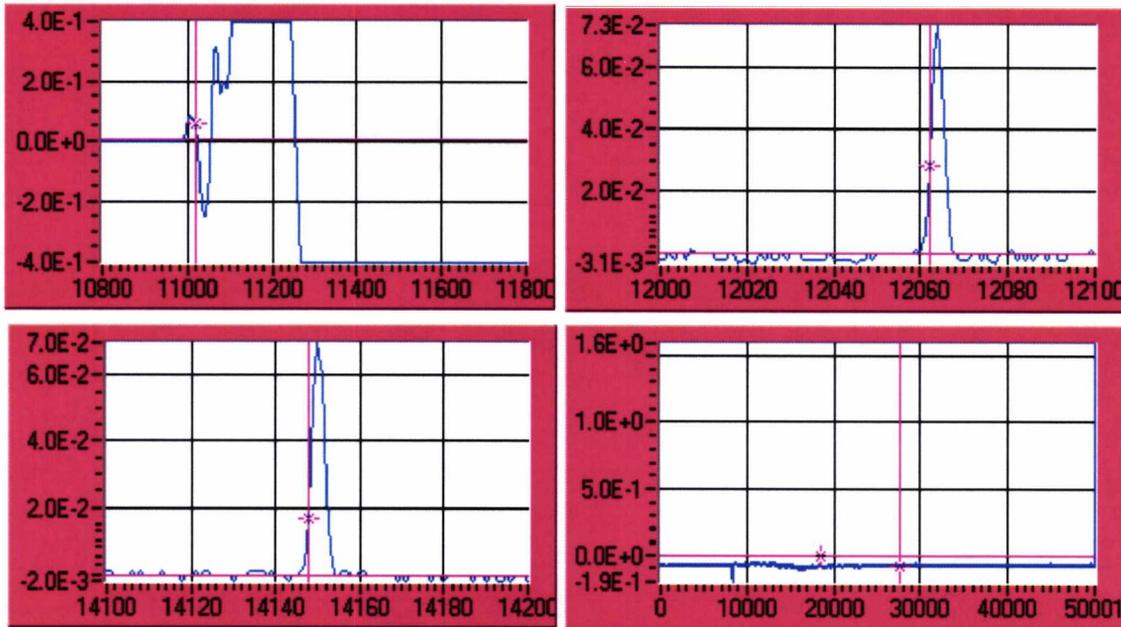


Figure 3.2.4 – Data from sensors within the gun tracing the flight of the projectile through the gun. The top left readout is from the piezoelectric transducer (PZT) attached to the stop plate; this registers the impact of the sabot pieces and burst disc fragments onto the stop plate, giving the time at which the projectile passes through it. The top right readout is from the first laser curtain, and the bottom left readout is from the second laser curtain. The final readout is unused. The x-axis displays the time interval in units of 4×10^{-8} seconds.

The projectile itself continues onwards down the main axis while the sabot pieces fly off to the sides. There is a stop plate placed at the exit from the blast tank which has a 1 cm diameter hole in the centre. The sabot pieces impact this stop plate and do not travel any further down the gun, but the projectile passes through the hole in the centre of the plate and then through two laser curtains spaced 0.499 m apart. The laser curtains are fan-shaped and focused onto a receiving photodiode; this produces a current which is measured by way of a trans-impedance amplifier. When an object breaks the laser beam the level of illumination on the photodiode decreases and the current produced by it increases. This output current is displayed on a Lecroy 9304AM Quad 200 MHz oscilloscope (an equivalent output is shown in figure 3.2.4). By measuring the timing of the two current spikes relative to each other the velocity of the projectile can be calculated since the spacing between the laser curtains is also known. There is also a piezoelectric transducer (PZT) which is used to detect impacts on the stop plate; this PZT can be used as a third sensor for measuring projectile velocity if needed. This velocity measurement is accurate to better than $\pm 1\%$.

After passing through the laser curtains the projectile finally enters the target chamber, where it impacts on the target placed at the end of the gun. The target chamber is 1m in length with an internal diameter of 0.4 m; this gives it a relatively large internal volume with the practical capacity to take targets of up to 30 cm diameter. The end of the target chamber opens to allow targets to be placed inside. Once sealed the target chamber is airtight, and it is evacuated down to a pressure of approximately 0.3 mbar immediately prior to a shot in order to minimise the effects of air drag on the projectile. Electrical feedthroughs are present which allow certain target parameters such as temperature to be measured during the shot process.

3.3 - Targets

For this work a process had to be developed for making targets that roughly resembled the makeup of a large Kuiper Belt Object (KBO), with an ice shell covering a spherical core made of an ice-silicate mix. Previous work (Lightwing *et al.* 2006) had used cylindrical targets made of a plain ice-silicate mix by simply freezing a mixture of sand and water (ratio 4:1 sand to water) inside a cylindrical mould to produce an ice-silicate material with a density of 1.9 g cm^{-3} . Pluto has a density of 2.03 g cm^{-3} (Buie *et al.* 2003) so this material is considered to be a reasonably good match in simulating Pluto and other large icy TNOs as a homogenous body.

Making spherical ice-silicate targets proved to be considerably more complicated however - previous work at the University of Kent had involved the construction of both irregular porous ice spheres and relatively uniform solid ice spheres (Dawe *et al.* 2005). The irregular porous spheres were intended to simulate comets and were made by crushing ice flakes together by hand. The solid ice spheres were made by filling a balloon with water and freezing it, and then cutting off the mould, although due to the shape of the balloons the targets were not exactly spherical. Data from this work is included here for comparison against ice shots made during the current experimental programme, but the methods used to construct these spherical ice targets could not be used to produce spherical sand-ice targets. Trying to mould the sand-ice mixture by hand proved to be unsuccessful since the sand and water mixture is difficult to mould properly due to the effects of gravity – the sand sinks to the bottom leaving a layer of

water on top, resulting in an uneven mix. Furthermore the water in the mix flows just like unmixed water, meaning that any attempt to mould spherical targets without spherical moulds that retained all of the mixture in a fixed, rigid shell resulted in highly misshapen targets unsuitable for experimental use.

The final method fixed upon for this experimental programme involved the use of small spherical plastic moulds with an interior diameter of 39mm – otherwise known as table-tennis balls. These moulds have dimensions that are smaller than what is ideal for a target but have the advantage of being both cheap and disposable. To produce a plain ice-silicate target a hole is cut in the top of the mould and first water and then sand are poured inside. When enough sand had been added for the mould to be completely filled with a mixture that was fairly even (i.e. no loose water layer on top of the mix) the targets were then placed inside a Frigidaire Elite freezer and frozen for a total of four days at a temperature of -20°C .

On the third day the targets were removed and the mould shells cut off from the outside of the now-solid ice-silicate mixture, and the targets are inspected to see if they were suitable for use in a shot – sometimes the sand settled to the bottom of the mixture after the targets had been placed in the freezer, resulting in a sizable water ice cap on the target that made it unsuitable for shooting. Twice the number of required targets were usually made so that any unsuitable ones could be discarded and replaced with backups. All usable targets were then returned to the freezer and frozen for a fourth day.

The end products of this process were small 39mm diameter spherical targets weighing approximately 60g, composed solely of the frozen ice-silicate mix. These targets were not perfectly spherical, since the construction of the moulds meant there was a shallow groove running around the circumference of the target where the two halves of the mould had been joined together. However, this groove was observed to have no effect on the energy required to disrupt the target and so it can be ignored.

These plain spherical ice-silicate targets are useful for comparison with the cylindrical ice-silicate targets, but still do not accurately represent the suspected layered

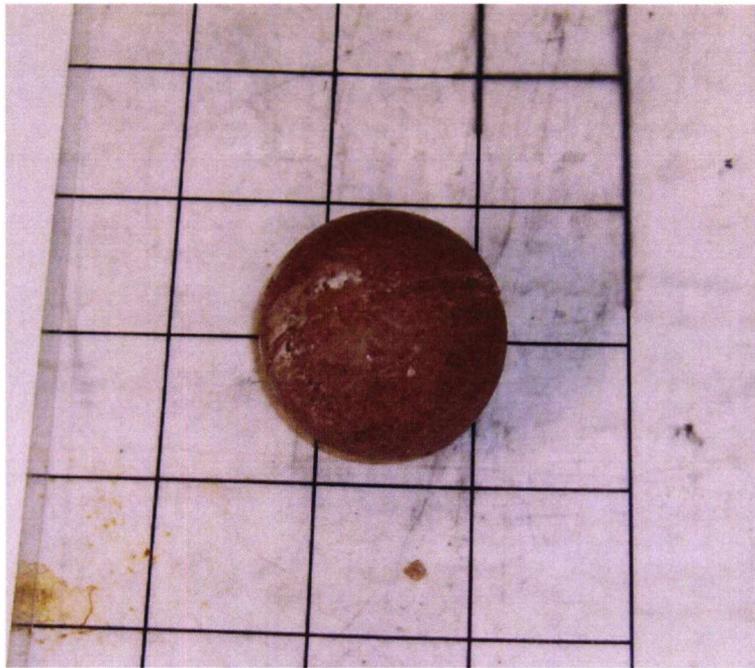


Figure 3.3.1 - A spherical ice-silicate target. The grid squares are 2.55 x 2.55 cm. The groove running around the circumference of all spherical targets caused by the moulds can be seen on the top of the target here.

composition of the large icy bodies in the outer Solar System – they are rough analogues that reflect Pluto-like bodies in density and no more. To produce a more realistic target which does, the spherical targets need a layer of water ice added to the surface of the ice-silicate mix. Several methods were tried and found wanting, due to gravity making freezing an even layer of water on the surface of the target impossible – it would always flow downwards, resulting in lumpy, misshapen targets whenever it was turned to apply the ice layer to a particular spot.

Eventually a method was developed which produced reasonably uniform targets. The first stage of constructing the inner ice-silicate core was as outlined above, except that before the mould was filled with sand and water a length of nylon thread was inserted through the hole at the top of the mould and pushed out of the bottom using a needle. The thread was cut so that there were roughly equal amounts on either side of the mould. This resulted in the mould being strung on the thread much like a pearl on a necklace. The small hole in the bottom of the mould through which the thread has passed was sealed with blu-tac making it watertight, and the rest of the target-making procedure was carried out as normal with one small deviation – the plastic mould



Figure 3.3.2: A sand-ice core, left, and a sand-ice core covered in a layer of ice, right. Grid squares are 2.55×2.55 cm. Note the somewhat-elongated shape of the layered target compared to the core; this is caused by water flowing down towards the bottom of the target during the freezing of the ice layer.

shells were removed the day after freezing rather than three days into the process, to allow the ice layer to be added and then frozen for three days itself.

When the outer plastic shell of the mould was removed it left an ice-silicate target with a length of thread running through it. One end of the thread was then used to dip the target first into a container filled with cold water chilled with ice flakes and then into another container filled with liquid nitrogen in quick succession. Repeating this process slowly built up layers of ice on the surface of the target like an onion. After this process had been carried out about 25 times it produced a 2mm thick shell on the target, but the shell was somewhat thicker towards the bottom of the target due to gravity making the water flow downwards before it is dipped into the liquid nitrogen. To compensate for this effect the target was simply turned over and the other end of the thread was then used to dip it into the water and nitrogen. This resulted in the ice layer being somewhat thicker towards the “poles” of the target where the thread protrudes, but overall the layered targets are reasonably spherical with the ice layer no more than 20% thicker at the poles than at the equator.

After the shell has been added via this process the thread is cut off, leaving an ice-silicate core covered by an ice shell with a length of thread running through it. This method can add layers of ice up to 7mm thick to a target – any thicker than this and the target starts to become seriously misshapen due to gravity effects again. The layered targets are then put back into the freezer for three days before use in a shot. While the shell thickness was varied between 2 and 7 mm in an attempt to investigate how shell thickness affected disruption energy, the median shell thickness used most often was 4mm since this resulted in the most spherical ice layer.

These plain ice-silicate and layered targets were the main types used during the experimental programme, with the process for making them in all cases the same as the one outlined above. However, various properties of the targets were changed in order to investigate whether or not they affected the overall energy required to disrupt a target. These properties included sand grain size, ice layer thickness and target temperature.

A number of targets using differently-sized sand grains were carried out. The default kiln-dried sand grains used had a grain size of 617 μm , but grain sizes of 160-1043 μm were also used in order to investigate what effect, if any, larger or smaller sand grains would have on the results of impact.

Low temperature targets were made as normal using the default kiln-dried sand, but instead of freezing them inside the Frigidaire freezer at $-20\text{ }^{\circ}\text{C}$ (253 K), they were instead placed inside a Heraeus Instruments low-temperature freezer at $-120\text{ }^{\circ}\text{C}$ (153 K) for four days prior to shooting. Other than that, the method used to construct and shoot them was exactly the same as the typical one outlined above.

In addition, both cylindrical and spherical pure ice targets were made so that a comparison could be drawn between the disruption energy of ice and the disruption energy of ice-silicate. The process for making them was somewhat more complicated than simply freezing water from the tap; tap water contains a lot of tiny gas bubbles that result in imperfections in the structure of the ice when it is frozen, at best weakening it and at worst creating cracks through the target rendering it useless.

In an attempt to remove this gas deionised water was poured into a kettle and boiled for half an hour. This purified the water but it must then be cooled down rapidly and frozen so that gas does not seep back into it. To achieve this, the kettle was placed inside a container full of crushed ice which is packed around it. After the purified water had reached a suitably low temperature (not directly measured, but chilled to the point where ice no longer melted in it so close to freezing) it was then poured into the moulds - identical to the moulds used for cylindrical and spherical ice-silicate targets - and frozen in the same way. The ice spheres had a diameter of 39-40 mm as with the sand-ice spheres, while both the sand-ice cylinders and ice cylinders had diameters of 75 mm and heights of 8 – 10 mm.

3.4 - Target holders.

A variety of target holders were required to keep the various different target types in place during a shot. For the cylindrical targets a simple metal V-shape was used, with the target sitting lengthways inside the crevice on top of the holder. The metal holder was then placed on top of an adjustable stand and oriented so that the projectile would strike the flat end of the target. Both the metal holder and the stand were placed inside the same freezer used to chill the targets to -20 °C for at least 24 hours before a shot, ensuring that as little heat as possible was lost from the target while it was placed in the target holder.

The first target holder constructed for the small spherical targets consisted simply of a polystyrene “cup” mounted on top of an adjustable stand upon which the target was placed. The stand was then adjusted to an appropriate height so that the target was placed in the path of the gun projectile. However, a problem which became apparent during the shot process was that due to the small size of the targets and the length of time they had to remain in the gun both before and after a shot, it was very difficult to recover smaller target fragments before they melted – often only the two or three largest fragments survived, making it impossible to plot meaningful size distributions.



Figure 3.4.1 - A cylindrical ice-silicate target sitting in its target holder. A partial ice cap is visible on the rear of the target.

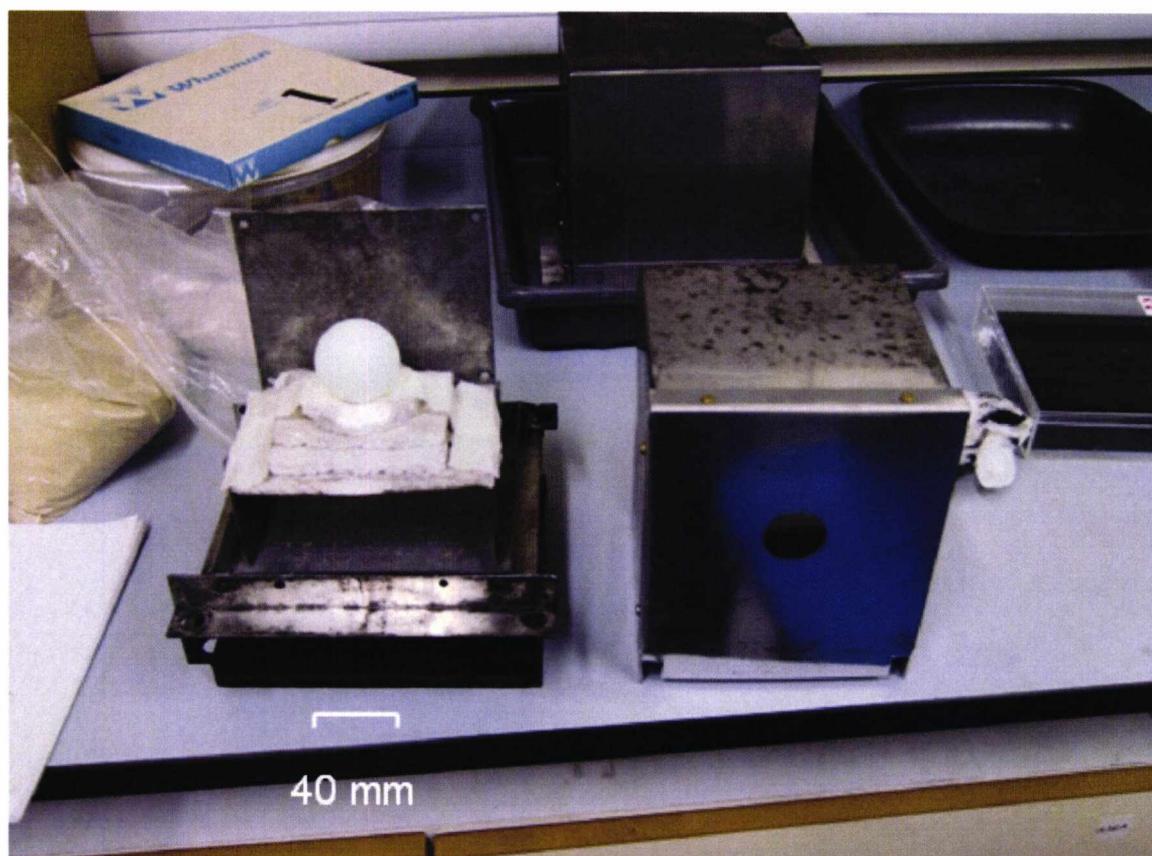


Figure 3.4.2 - Target holder for spherical ice targets with the top removed. A plastic mould has been placed in the polystyrene cup where a target would normally sit.

To solve this problem a new target chamber was constructed. It consists of the same polystyrene cup as before, but this time placed on a stand inside an enclosed metal box with a small hole in the front through which the projectile can pass to strike the target (see figure 3.4.2). The box sits on top of an adjustable stand as before, but prior to a shot it is put into a low temperature freezer and frozen at $-120\text{ }^{\circ}\text{C}$ for 24 hours. This results in both a lower ambient temperature around the target as it sits in the target chamber while the air is evacuated prior to shooting meaning, and also provides a low-temperature surface that catches the smaller fragments after a shot. Both these factors mean that the target as a whole suffers less melting before a shot and that smaller target fragments are much less likely to completely melt before they can be recovered.

3.5 - Shot process

Prior to each shot the targets were briefly removed from the freezer they are stored in order to measure certain properties. The weight of the target was measured using a set of Adam Equipment Co. ADG 6000/L digital scales, accurate to $\pm 0.05\text{g}$, and the diameter of the target was measured using Whitworth digital callipers that are accurate to $\pm 0.005\text{ mm}$. The cylindrical targets also had their height measured using the callipers. These measurements are enough to calculate the density of each target type. The targets were then returned to the freezer until the gun was ready for shooting.

When the gun was ready, the targets were quickly removed from the freezer and placed inside the appropriate target holder. The holder and target were then placed inside the target chamber together and the position of the target was checked by eye using a steel rule. Since the target chamber had a diameter of 0.4 m the centre of the target had to be placed 20 cm away from the bottom and sides of the target chamber to maximise the chances of an impact.

Once the position of the target had been checked and corrected if required, the target chamber was sealed up and evacuated down to a pressure of 0.3 mbar, a process which took about 10 minutes. The firing pin and pendulum were then attached to the

gun, after which the gun room was cleared of all personnel and the gun was fired from outside the room by the engineer.

After firing the oscilloscopes are checked to see if the projectile passed through the laser curtains with no other debris from the sabot or burst disc, and the target chamber is repressurised as quickly as possible. Normally the target chamber would be cleared of harmful byproducts (e.g. noxious vapours) from the shot by pumping it out but since this takes up to twenty minutes it is not practical to do this for ice targets, and so the target chamber was opened up immediately with the experimental personnel wearing gas masks.

The remnants of the target were removed from the gun while still in the holder and any stray fragments that might have escaped were also collected. The mass of as many fragments as possible was then measured. Most important was the mass of the largest fragment remaining after a shot, since this allows the calculation of the critical energy density for a given material. In the case of the small spherical targets the mass was the only fragment property measured, since there were many small fragments which melt quickly making it impossible to measure any more than that (e.g. fragment diameters, which would have been useful for analysing size distributions) in the small timeframe available.

3.6 – Tensometer strength tests

The tensile and compressive strengths of the various material types including samples of all sand-ice types used in the targets was tested by crushing samples of the materials inside a uni-axial Hounsfield tensometer, manufactured by Tensometer Ltd. Materials tested included pure ice and ice-silicate mixes using each of the five different types of sand. The samples were small cylindrical pieces of each type of material, made using much smaller moulds than those used for the large cylindrical ice-silicate targets. These moulds were lengths of plastic tube stoppered at one end by a bung.

To make the samples, the desired material was poured into the stoppered mould and

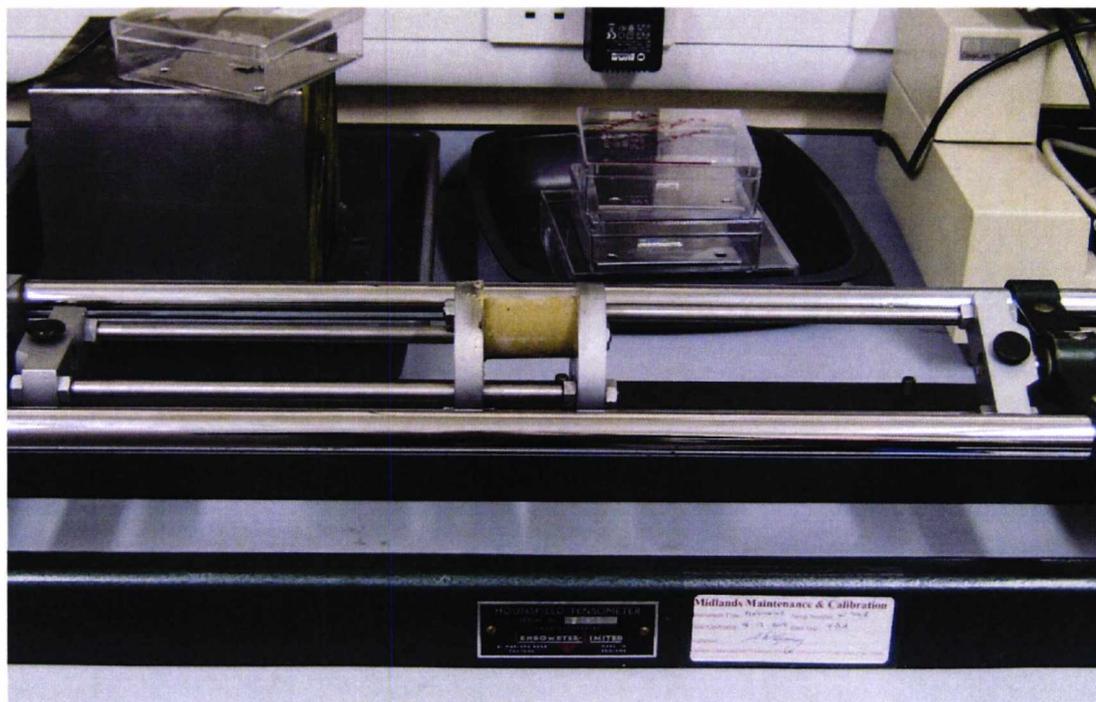


Figure 3.6.1 - An ice-silicate sample being crushed in the tensometer. As the two ends of the tensometer rods are pulled apart the plates in the middle move closer together, compressing the sample until it fractures.



Figure 3.6.2 - A cracked ice-silicate sample after being tested in the tensometer, and the plastic mould used to construct it. In this case the compressive strength has been tested by crushing the sample along its long axis.

the mould was then frozen for the standard four days at -20°C . They were then removed from the freezer and immersed in hot water for approximately twenty seconds, making them easier to remove from the mould. However, both the ice and ice-silicate proved to be very adhesive to the rubber bung, so the bung first had to be “unscrewed” from the mould with the sample still attached to it, and then the sample was snapped away the bung. The expansion of the ice as it froze meant that the top of the sample tended to be rather uneven, so the top was worn down against a piece of warm metal until it was relatively flat. This method produced small samples with an average diameter of 30 mm and heights varying between 5 and 8 cm.

The samples were then placed inside the tensometer and crushed until they cracked, at which point the force being applied to the target was noted. Prior to use the tensometer rods are frozen at -20°C for 24 hours beforehand to minimise the amount of target melting through contact that occurred during the crushing process. The samples were crushed both lengthways and sideways in order to check both the compressive and tensile strength of the materials. The theory behind this method is described in Chapter 4.

3.7 - Projectiles

The light gas gun can fire a range of projectiles from the micron-sized (in the form of buckshot) on up to solid projectiles 2 mm in diameter. Here, since we were interested in disrupting the targets the projectiles used in the shot programme were spheres that ranged in size from 0.8 mm to 2 mm in order to deliver the required range of kinetic energy to the target. A detailed breakdown of the range of different projectiles used and the reasoning behind it can be found in Chapter 4.

3.8 – Impact angle

When impacting a spherical target an impacting projectile needs to strike at a 90° angle in order to impart the full amount of kinetic energy to the target. Since the targets used in the experimental programme are relatively small with radii of

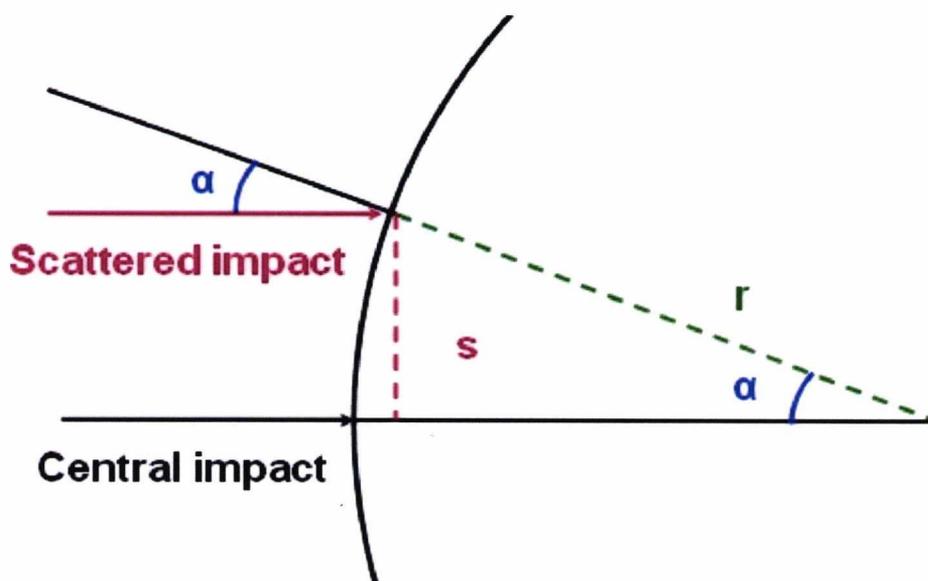


Figure 3.8.1 - Diagram showing how a deviation s from a central (normal) impact axis will result in an impact at an angle α .

approximately 40mm, a deviation of more than a few mm could dramatically change the projectile impact angle. Therefore the accuracy of the gun is important when assessing how much energy has been delivered to the target, especially in the shots described here where it is impossible to determine exactly where a projectile has struck a target after a disruption impact.

Figure 3.8.1 illustrates the problem; when a projectile impacts off-axis by a distance s , it results in an impact that deviates from the normal impact angle of 90° by an angle α . Since α is the same as the angle between the radius r and the central impact axis shown in figure 3.8.1, it is easily calculated by

$$\alpha = \sin^{-1}\left(\frac{s}{r}\right) \quad (\text{Eq. 3.4})$$

Where s is the scatter distance away from the central impact axis and r is the radius of the body being impacted.

Projectiles fired in the light gas gun used to carry out the shot programme detailed here deviate from the central axis by a maximum of 1 cm. Since that is a maximum, it

is reasonable to assume that a projectile will deviate by the full 1 cm only occasionally and that most projectiles will impact within a radius closer to the central axis. A detailed breakdown of what percentage of a given number of projectiles will impact within a certain radius from the central axis would be useful here, but unfortunately it is not available since the accuracy of the gun has never been systematically tested to gather that data. However, it is known that out of a set of twenty shots where the scatter was measured only one had a scatter of more than 7.5 mm.

It is therefore possible to make a rough estimation of the angles projectiles may be impacting at. Using the 1 cm maximum scatter gives an impact angle on a typical 40mm diameter target of 30° ; however we know that - given the sample set earlier - only a small minority of shots will approach the maximum scatter, and the other ~95% will impact at angles up to 22.02° .

These angles seem significant, but the question remains: how will they affect the energy of an impact?

Benz & Asphaug (1999) model how Q^* changes with impact angle as part of their work on size scaling. They find that at laboratory scales Q^* does increase with impact angle, but that this increase only begins to be significant at angles of greater than 30° ; their results are dependent on both target material and projectile speed, but they find that for an impact angle of 30° there is at most a 40% increase in Q^* (since the target is being struck with less effective energy the total energy of an oblique shot needs to be higher than that of a direct shot in order to disrupt a target, hence an increase in Q^*) over an impact angle of 0° for a basalt target. It is reasonable to assume that since most of the shots here are impacting at angles of 22.02° or less, the scale of the increase in Q^* will be even smaller.

Grey *et al.* (2002) investigate how impact craters in ice scale with impact angle. They find that there is a correlation in crater volume and depth with impact angle, pointing to decreased impact efficiency for oblique impacts; again, though, this only becomes significant at highly oblique impact angles, and at the angles likely to be encountered

during a typical shot in the experimental programme described here the change in crater morphology is very small.

These examples indicate that while the scatter in gun accuracy will introduce some scatter in the resulting impact data due to slightly oblique impact angles, this scatter will not be overwhelmingly large and should still allow a reasonable assessment of Q^* for a given target to be made.

3.9 – Conclusion

In this chapter I have described the theory behind the function of light gas guns and how they are used to accelerate projectiles up to hypervelocities. The method used to produce small spherical layered targets that reflect the suspected layered structure of Pluto and other large TNOs has been outlined, and the process of carrying out an impact on these targets using the University of Kent's two-stage light gas gun has been described. The method and apparatus used to measure the fragments after impact and the strength of the different types of material used in the targets are also covered.

Chapter 4 – Materials

“I combined the crystals in a tungsten-titanium matrix at supercooled temperatures, and that’s what did the trick. [...] Its real name has thirty-seven syllables, but I call it “Unobtainium”.”

- Dr. Ed “Braz” Brazzleton, *The Core*

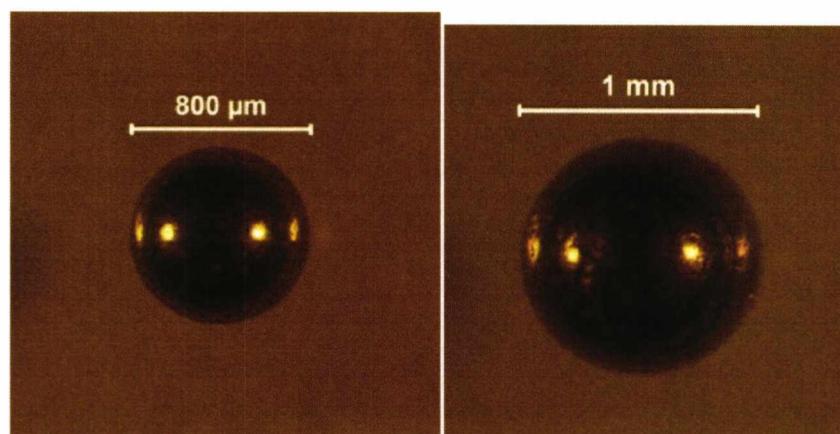
4.0 - Introduction

Before embarking on a detailed description of the experimental results, an overview of the materials used in both the targets and the projectiles is essential since different projectile types are used to impact different target types for different reasons. Detailed knowledge of the physical properties of both the target and the projectile is required for calculation of critical values such as the impact energy density. This chapter explains the rationale behind using each specific material type as well as basic measurements of their physical properties and other relevant qualities.

4.1 - Projectiles

A variety of projectiles were used during the course of the experimental programme. The specific type of projectile was either dictated in response to the type of material being shot at – materials such as ice which are relatively easier to disrupt require lower density projectiles in order to achieve a range of Q values which both crater and shatter the target – or selected in order for a comparison with shots using a different projectile type on the same target material.

In an attempt to determine as accurate a value for the calculation of Q^* as is possible (taking into consideration uncertainties in projectile velocity, impact location etc.), five sample projectiles of each material type and size were imaged and their diameters measured under a Leica MZ16 microscope, and their masses were measured using a Santorius Ultramicro 4504 MPH-1 microbalance, by GMBH Gottingen. The mean masses and diameters were then calculated, and these values were then used to determine the density of the material type.



Figures 4.1.1 (left) and 4.1.2 (right) – Images of a 0.8 mm stainless steel projectile and a 1mm titanium projectile.

Tables 4.1.1 (left) and 4.2.2 (right) – Properties of 0.8 mm stainless steel and 1mm titanium projectiles.

0.8mm Stainless Steel		1mm Titanium	
Mean diameter (mm)	0.79 ± 0.01	Mean diameter (mm)	0.998 ± 0.007
Mean mass (mg)	2.0707 ± 0.004	Mean mass (mg)	2.395 ± 0.004
Density (g cm^{-3})	8.02	Density (g cm^{-3})	4.602

4.1.1 - 0.8 mm diameter stainless steel projectiles

0.8 mm Trafalgar (from Spheric-Trafalgar Ltd., see <http://www.ballbiz.com/>) stainless steel 320 was the default projectile type used at the start of the experimental programme; a 0.8 mm projectile could deliver anywhere up to 900 J kg^{-1} to the standard sand-ice core target depending on the impact speed.. However, as the programme progressed drawbacks to using this projectile type became evident. High velocity shots tended to pulverise the targets beyond the point where any reasonable analysis of fragment size distributions could be performed, and it was thought that this may be an effect of the high density of stainless steel lending the projectile extra penetrating power and increasing its disruptive effect. Using a smaller stainless steel projectile was not practical since 0.8mm was already the smallest size that could be fired in the gun with any reasonable degree of reliability – while smaller stainless steel projectiles were available they were designed to be fired *en masse* in a cloud and not as a single projectile. For these reasons the decision was made to switch away

from stainless steel to a projectile material with a lower density for the latter stages of the experimental programme.

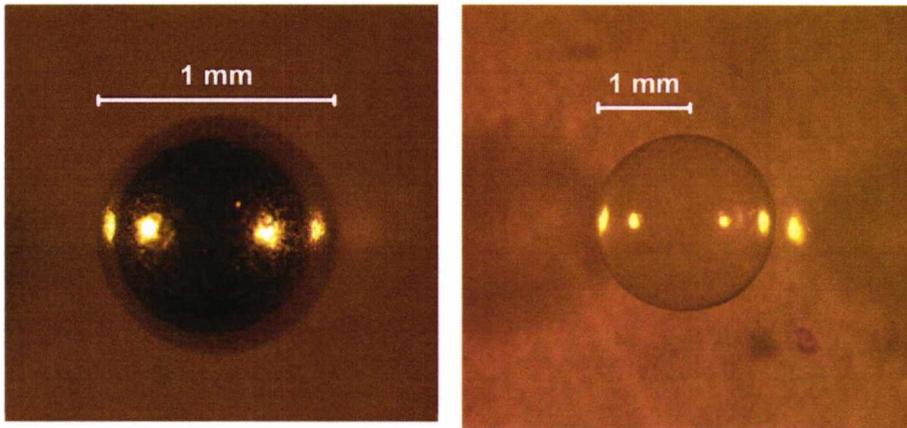
4.1.2 - 1mm diameter titanium projectiles

1 mm Trafalgar titanium projectiles were eventually settled on as the projectile of choice for impacts on both sand-ice cores and layered bodies. With a mass only slightly higher than that of 0.8 mm stainless steel they achieved a similar range of energy densities when fired from the gun – allowing for direct comparisons between titanium shots and the previous stainless steel shots to be made without any difficulty - but the much lower density meant that the targets yielded far more measurable fragments after disruption. In addition the larger diameter of the titanium projectiles ensured a higher degree of reliability than that of 0.8mm stainless steel; 20% of the stainless steel shots made over the course of the experimental programme failed, compared to 8% of the titanium shots.

4.1.3 - 1mm diameter aluminium projectiles

Impacts on small-scale ice cores weighing only 30g posed a problem; not only were they half the mass of the cores constructed from sand-ice, but pure ice has a much lower critical energy – 6 J kg^{-1} (Murray 2005) as opposed to 210 J kg^{-1} (Lightwing 2006) for sand-ice. Therefore a projectile had to be found that would deliver a kinetic energy value more than forty times less than that provided by titanium and aluminium projectiles. With a density of 1.15 g cm^{-3} nylon appeared to meet this requirement, but the low density made it extremely difficult to shoot with success; while shot reliability is dependent on mass rather than density (a heavier projectile is more stable in flight) if a projectile has too low a density then even the largest diameter projectile the gun is capable of firing will not have enough mass to consistently avoid being deflected off-axis.

An abortive shot program in which four out of five 2 mm to 3.175 mm diameter nylon shots failed proved that it was inadequate, and so the next projectile type tried was 1 mm aluminium. With a density of 2.952 g cm^{-3} a 1mm aluminium projectile can



Figures 4.1.3 (left) and 4.1.4 (right) – Images of a 1mm aluminium projectile and a 1.98 mm soda-lime glass projectile.

Tables 4.1.3 (left) and 4.1.4 (right) – Tables displaying the properties of 1mm aluminium and 1.98mm glass projectiles.

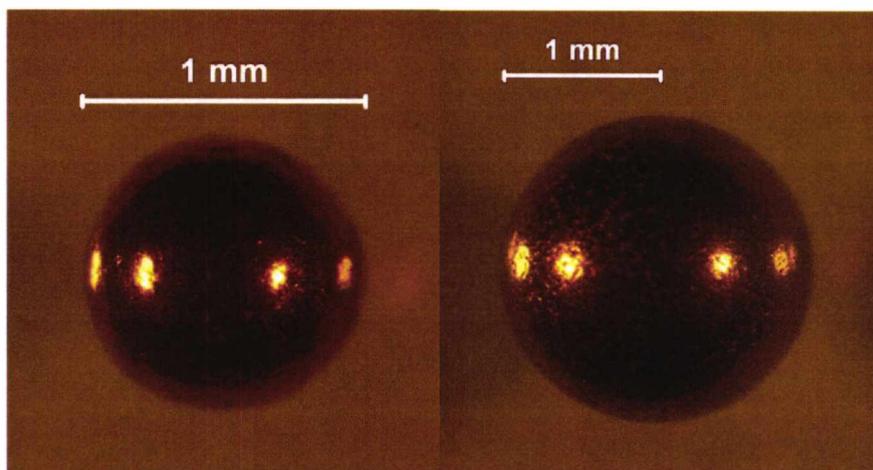
1mm Aluminium		1.98mm Glass	
Mean diameter (mm)	0.992 ± 0.018	Mean diameter (mm)	1.972 ± 0.034
Mean mass (mg)	1.509 ± 0.025	Mean mass (mg)	10.557 ± 0.189
Density (g cm^{-3})	2.952	Density (g cm^{-3})	2.629

deliver a minimum energy density of 25 J kg^{-1} to a 30 g ice target when fired at the slowest velocity possible in the gun ($\sim 1 \text{ km s}^{-1}$). This is several times higher than the previously determined value

of Q^* for pure water ice, but the intent was to cover the upper range of required energy densities for the calculation of Q^* with this larger, more reliable projectile, and then downsize to a smaller diameter for the two or three shots needed to pin down the lower range of required Q . However, the critical energy density for the pure ice cores turned out to be ten times higher than that previously determined for ice (see Chapter 5) and so the 1mm projectiles were sufficient to cover the entire range of Q .

4.1.4 - 1.98mm diameter soda-lime glass projectiles

1.98mm soda-lime glass projectiles manufactured by Potters Ballotini (see <http://www.pottersbeads.com/>, accessed 20/04/10) were used in a sequence of five shots as part of the investigation of how projectile density might affect disruption energy. While a 1mm glass projectile would have been more desirable (in order to



Figures 4.1.5 (left) and 4.1.6 (right) – Images of 1mm and 2mm copper projectiles.

Tables 4.1.5 (left) and 4.1.6 (right) – Tables displaying the properties of 1mm and 2mm copper projectiles.

1mm Copper		2mm Copper	
Mean diameter (mm)	0.962 ± 0.014	Mean diameter (mm)	1.965 ± 0.019
Mean mass (mg)	4.292 ± 0.008	Mean mass (mg)	36.858 ± 0.085
Density (g cm^{-3})	9.208	Density (g cm^{-3})	9.278

roughly match the diameters of other projectiles being used) they turned out to be highly unreliable, possessing too small a mass to be consistently fired in the gun with a reasonable success rate.

4.1.5 - 1mm and 2mm diameter copper projectiles

1mm and 2mm copper projectiles (supplied by Salem Specialty Ball Co., see <http://www.salemball.com/> accessed 20/04/10) were used almost exclusively in a preliminary course of work to determine whether a sand and water mix could be moulded into useable targets and to narrow down exactly where the critical energy density for such a material lay. Since it was unknown what the critical energy density of the sand-ice mix would be and thus what projectile type would be most suitable to use, copper was selected for the sole reason that an accurate comparison could be made with previous work (Murray 2004) that involved impacts on solid ice spheres. These impacts also used copper projectiles so that a comparison could be drawn between those and the recent (at the time) impact of the Deep Impact spacecraft,

Table 4.1.7 – Failure rates of the various projectiles used during the experimental programme.

Projectile composition and diameter	Number of shots	Failures	Failure rate
Nylon 2mm	3	3	100%
Nylon 3.175mm	2	1	50%
Stainless Steel 1mm	9	0	0%
Stainless Steel 0.8mm	60	12	20%
Titanium 1mm	36	3	8%
Glass 1mm	3	2	67%
Glass 1.5mm	2	1	50%
Glass 1.98mm	6	1	17%
Copper 1mm	10	2	20%
Copper 2mm	10	0	0%
Aluminium 1mm	6	1	17%

Table 4.1.8 – Comparison of measured projectile densities with accepted densities for each material type. Accepted densities are obtained from Wolfram Alpha (<http://www.wolframalpha.com>, accessed 15/12/10), a website which automatically retrieves requested quantities from a variety of textbook sources; see website for details.

Projectile type	Measured density (g cm ⁻³)	Accepted density (g cm ⁻³)
1mm aluminium	2.952	2.7
1mm copper	9.208	8.92
2mm copper	9.278	8.92
1.98mm glass	2.659	2.52
0.8mm stainless steel	8.02	7.9
1mm titanium	4.602	4.507

which was also made of copper and impacted the comet Tempel 1 in 2005 (A'Hearn *et al.* 2005).

Since the average measured diameter of the projectiles barely deviate from the specified diameter, the specified diameter is used for the purposes of calculating the kinetic energy of a given shot. However, the measured densities *do* differ from specified values by a non-negligible amount (often 0.1 or 0.2 g cm⁻³, which is on the order of a 10% deviation for the least dense projectiles) and so it is these that are used to determine the mass and energy of a projectile.

Table 4.2.1 – Table showing the specified and actual measured properties of the various types of sand grains used in the experimental programme. Errors are represented by the associated standard deviation.

Sand type	Specified grain diameter (μm)	Average grain diameter (μm)	Standard deviation (%)	Density; ice mix (g cm^{-3})	Standard deviation (%)
Fraction A	1180-2360	2007	25.2	1.87	1.1
Fraction B	600-1180	1043	20.2	1.83	4.8
Fraction C	300-600	605	19.3	1.84	2.7
Fraction D	150-300	229	23.1	1.77	1.7
Fraction E	90-150	160	21.3	1.72	7.7
Kiln sand	-	617	22.7	1.90	1.7

Table 4.1.8 compares the measured densities of the projectiles with the accepted bulk densities of the projectile materials. As can be seen the measurements obtained here consistently overstate the density compared to the accepted value, indicating the possibility that some flaw in the measurement process has led to a consistent overestimation of projectile density. If such a flaw exists it is in the microscope measurements of the projectile diameters; projectile masses were measured to a high degree of accuracy.

4.1.6 – Projectile reliability

As has been mentioned, the reliability of firing a given projectile is a major consideration when selecting an appropriate type to be used in a shot. Less-than-ideal projectiles were used in many cases when the projectile type with the best physical qualities turned out to be highly unreliable. Table 4.1.7 displays the failure rate of every type of projectile used for more than one shot in the experimental programme, including those not described in this section due to being discarded as unsuitable for use.

4.2 – Sand types

Three different sand types were used to investigate the effects of impacts on targets composed of differently-sized sand grains. These sand types were graded sand from David Ball Group (see www.davidballgroup.co.uk, accessed 20/04/10) which came

pre-sorted into several specific size ranges. The specific grades used in shots were fractions B, D and E. Two additional sand grades, fractions A and C, were also used along with the above fractions in strength tests to determine how the strength of the sand-ice mix varied with grain size. A sixth sand type was used to construct every other sand-based target used in the experimental programme that was not specifically intended to test the effect of varying sand grain size; this sand type was B&Q kiln-dried sand, used because it was easily available in the large quantities required. The sizes of the various sand types are listed in table 4.2.1; these quantities were found by measuring along the longest and shortest axes that passed through the centre point of a minimum of ten different sand grains under the Leica MZ16 microscope and then taking the average.

Also shown is the standard deviation in the size of each sand type. Interestingly every sand type has a standard deviation of approximately 20-25% including the kiln sand; this indicates that it has been sorted to a similar level of precision as the sand that has been specifically graded into size ranges, and thus is not likely to suffer from any increased scatter in results in comparison. The mean diameter is also constantly towards the high end of the specified range quoted by the supplier, and in two cases even exceeds it. Density measurements were made using the mass and volume measurements of the cylindrical samples used in tensometer testing. While the measurement errors in both the diameter and density measurements are very small (on the order of 0.3% for the diameter and 0.5% for density) the significant variation between separate measurements of each quantity mean that it is reasonable to represent the range of scatter in these quantities as being one standard deviation from the mean.

Figure 4.2.1 shows that there appears to be a very rough relationship between grain size and the density of the sand-ice mix. Certainly larger sand grains produce consistently denser samples than the smaller sand grains; however the different densities obtained for fraction C and kiln sand samples – which have almost identical average grain sizes - indicate that things may not be as simple as this, and the large scatter present in the density measurements for the smallest sand grain type (fraction E) mean that what appears to be a general increase could instead be a flat rate constant

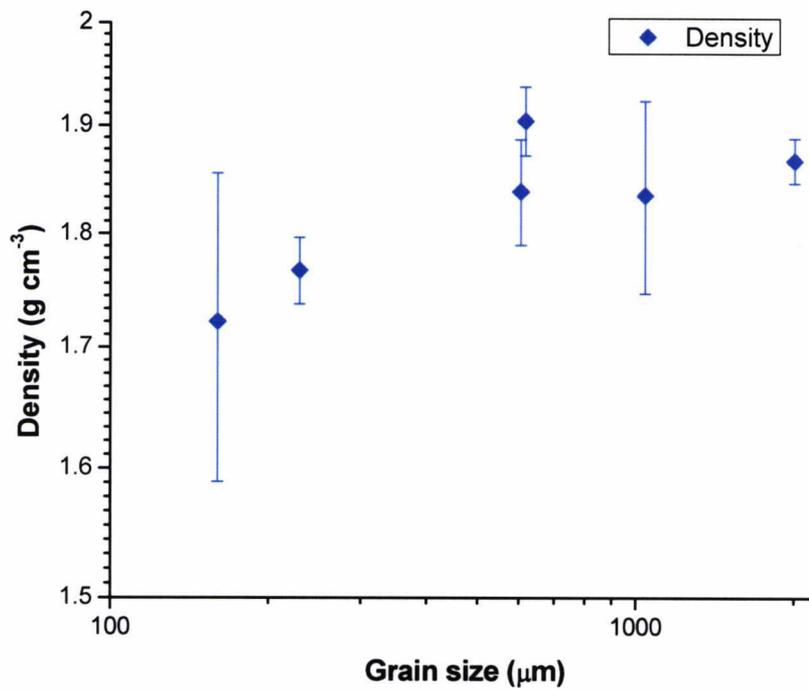


Figure 4.2.1 – Graph showing how the density of a sand-ice mix varies with the size of the sand grains used in the mix. Error bars represent ± 1 standard deviation.

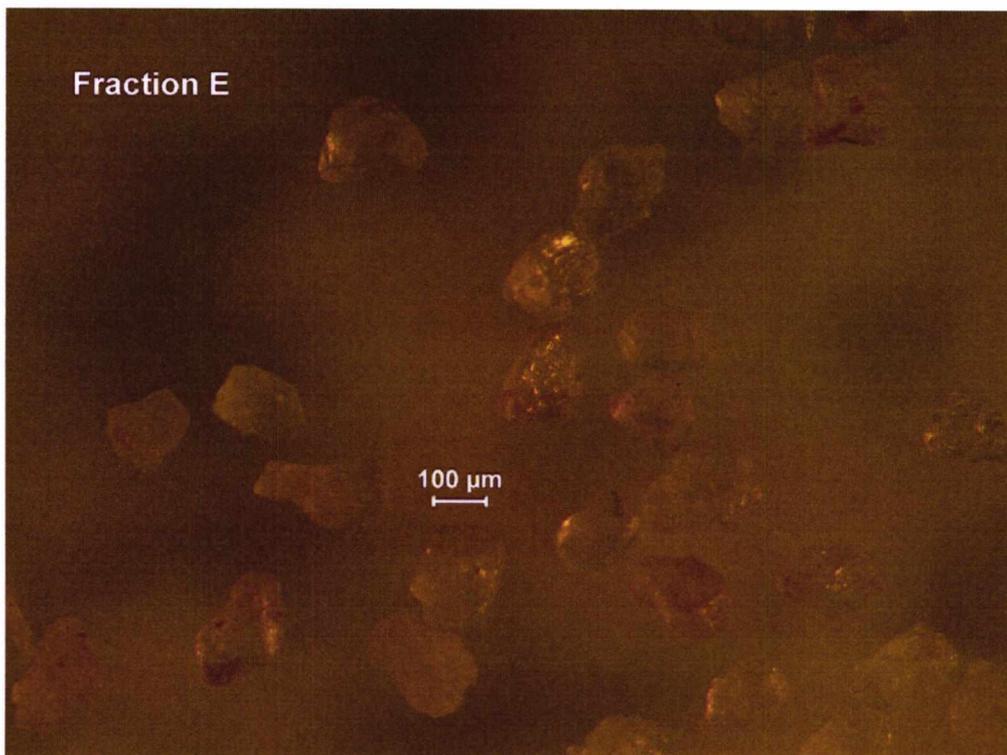


Figure 4.2.2 – Microscope image showing fraction E sand grains.

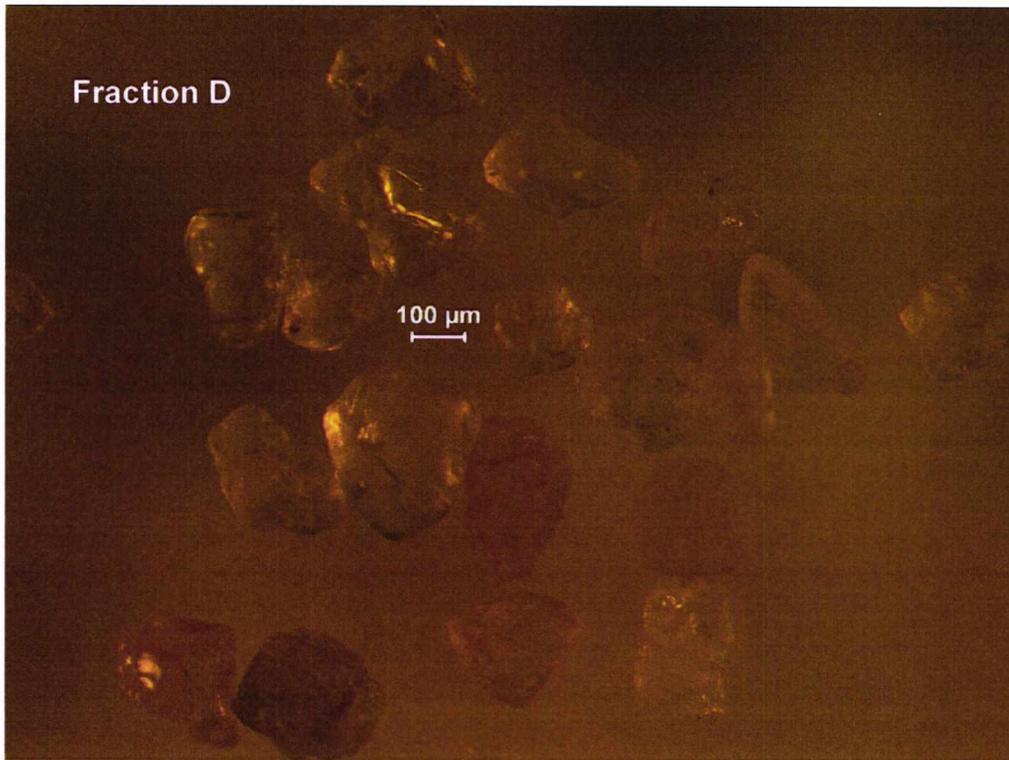


Figure 4.2.3 – Microscope image showing fraction D sand grains.

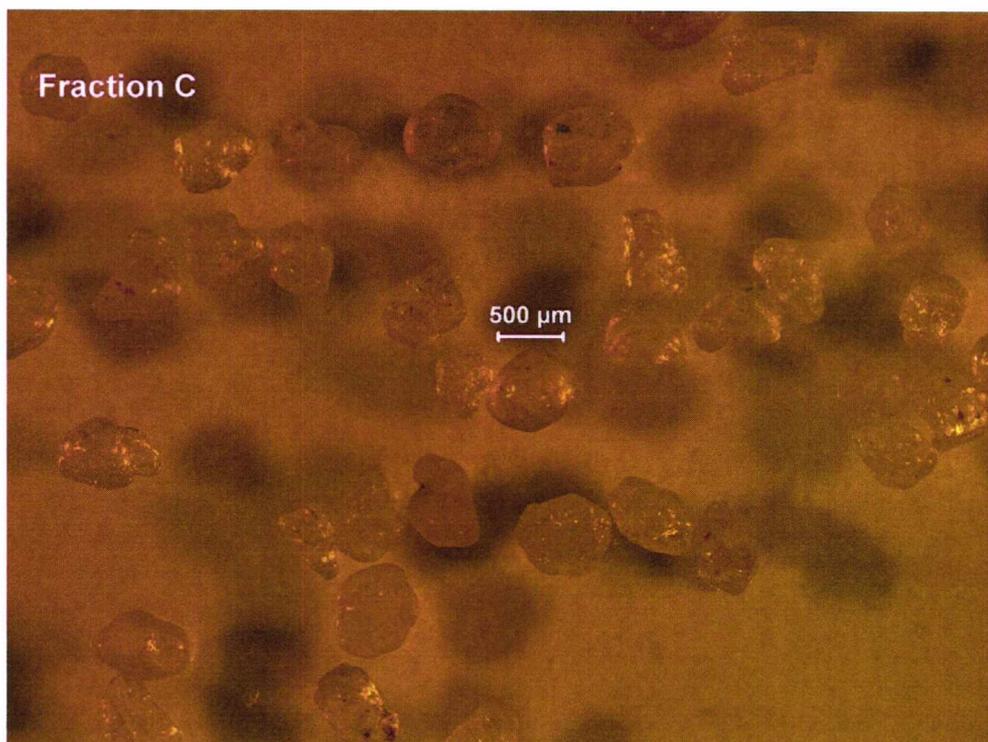


Figure 4.2.4 – Microscope image showing fraction C sand grains.

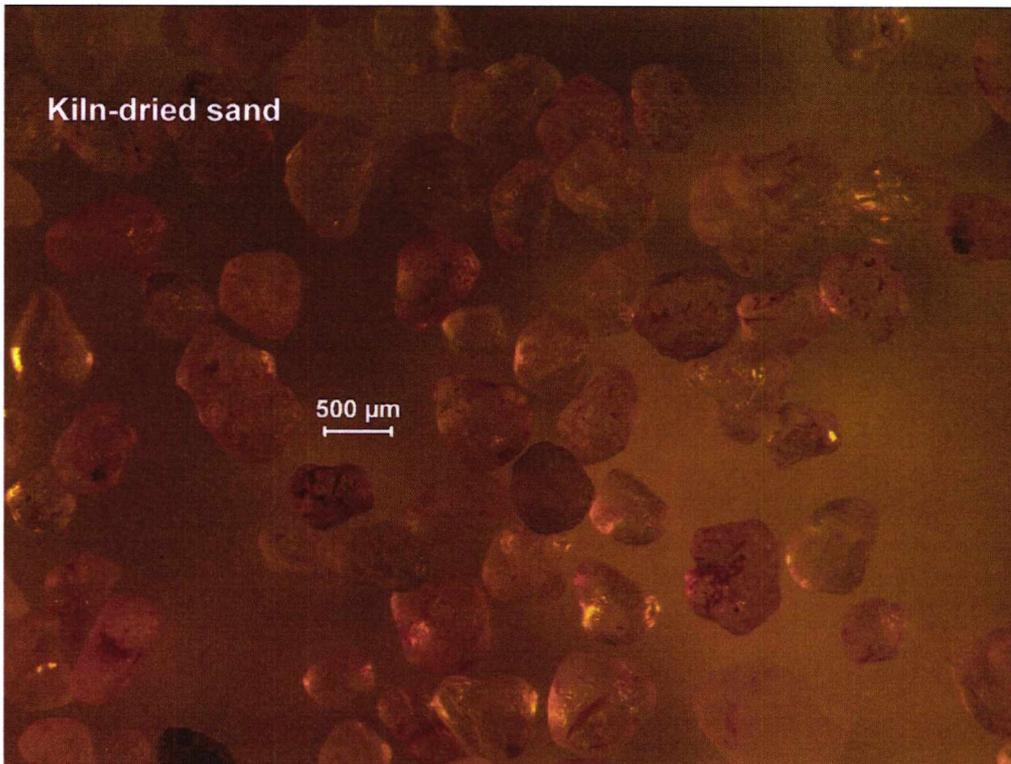


Figure 4.2.5 – Microscope image showing kiln-dried sand grains.

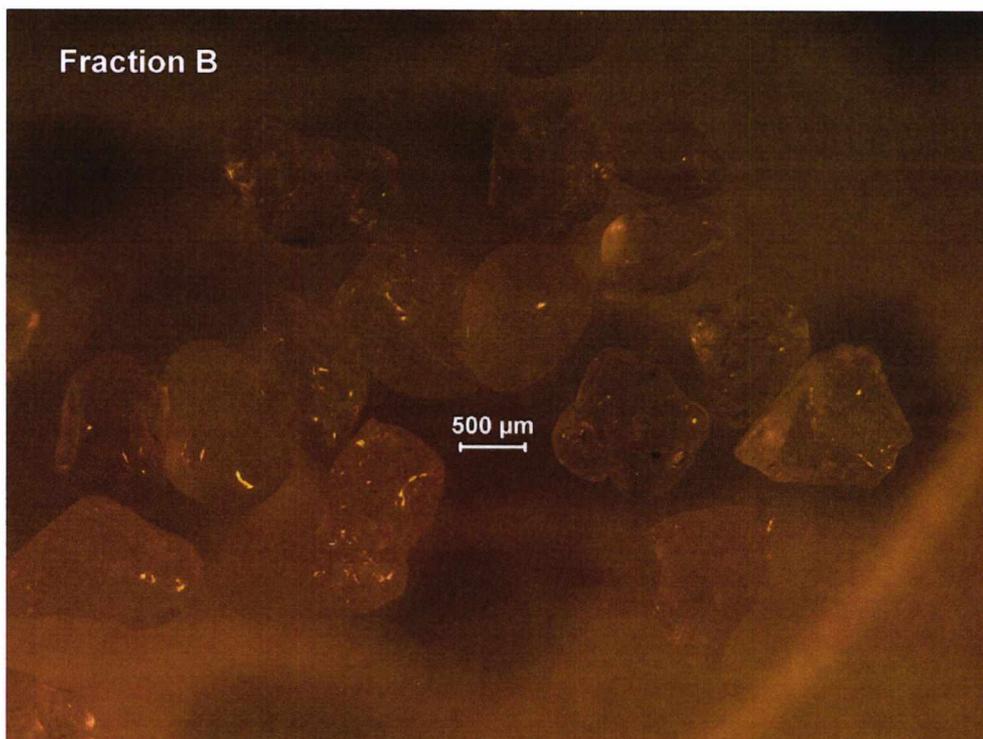


Figure 4.2.6 – Microscope image showing fraction B sand grains.

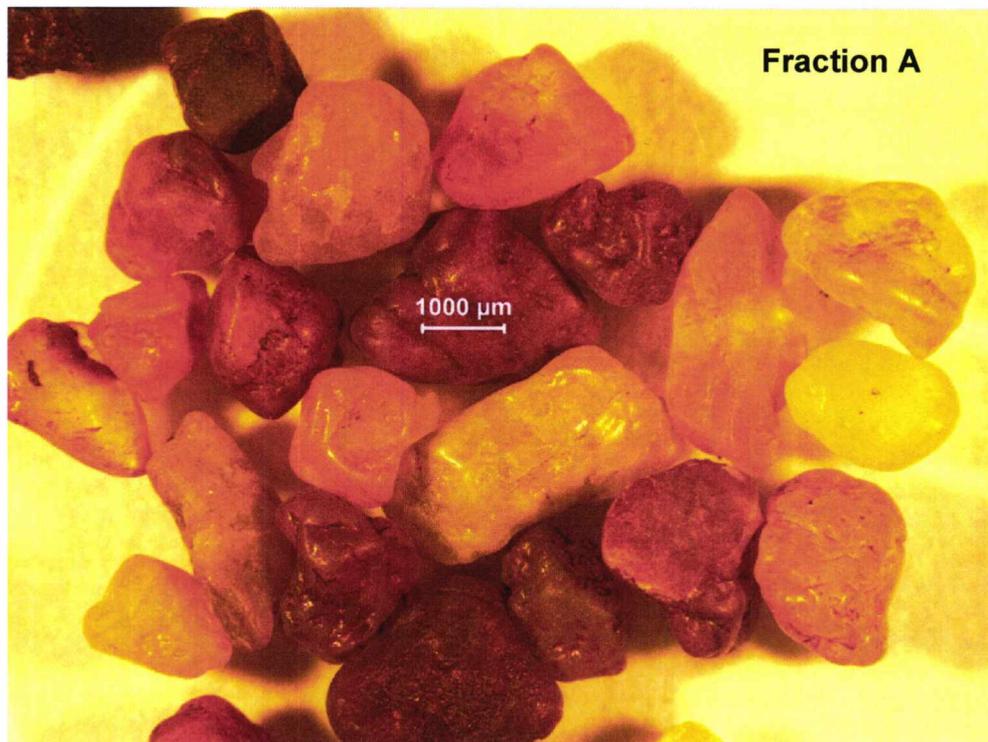


Figure 4.2.7 – Microscope image showing fraction A sand grains.

at around $1.8 - 1.9 \text{ g cm}^{-3}$. Repeat measurements are needed to confirm this relationship.

Figures 4.2.2 through to 4.2.7 show microscope images of the various types of sand grains in ascending order of size. The grains are typically semi-rounded with axis ratios of 6:5.

4.3 – Tensometer strength tests

Multiple samples of sand-ice were made in order to determine sample strength. They were made using each of the sand types described here, along with some identical samples made using pure ice. The samples were cylindrical with a uniform diameter of approximately 32 mm. The samples were then tested in a tensometer to determine their compressive and tensile strengths.

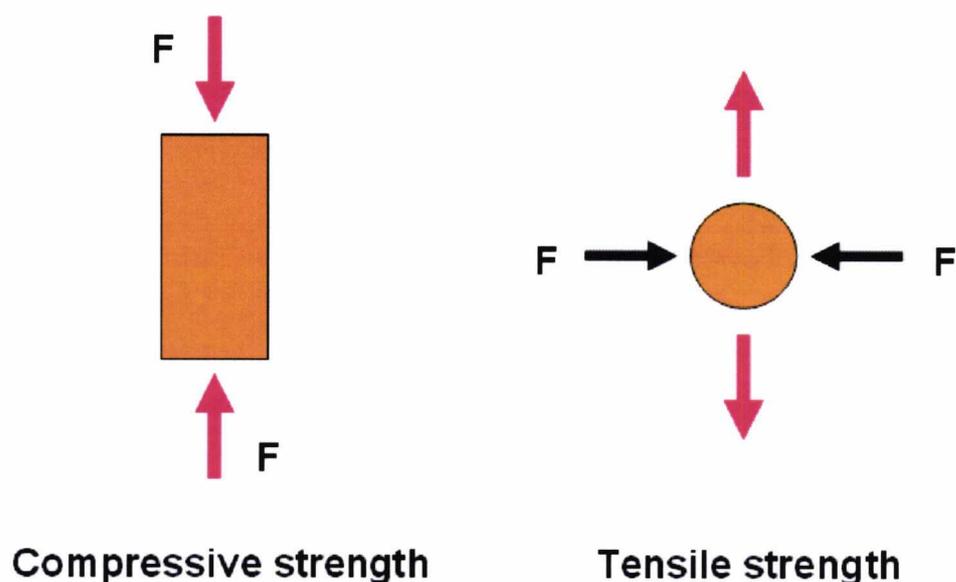


Figure 4.3.1 – Diagram showing the difference between compressive and tensile strength and where the force is applied to samples to measure each.

To determine compressive strength the samples were crushed along their long axes (see Figure 4.3.1), with the force applied to the flat ends of the samples. The force was gradually increased until the samples visibly fractured. While the tensometer used to crush the samples measures kilonewtons of force down to the third decimal place, in practice the fast rate of force increase made it impossible to accurately gauge the precise amount of force being applied past the first decimal place. A slower rate of increase would have been desirable in order to get a better measurement, but a compromise had to be made in order to fracture the sample before it melted – possible future work in this area could involve repeating these measurements inside a cold room in order to slow the rate of sample melt.

To calculate the compressive strength, the breaking force F is divided by the cross-sectional area of the sample it is being applied across, as:

$$\sigma_c = \frac{F}{A} \quad (\text{Eq. 4.1})$$

where σ_c is the compressive strength and A is the cross-sectional area.

Tensile strength measurements are slightly different in that while compressive strength tests simply crush the target, tensile strength tests do the reverse and stretch it. However, applying a stretching force directly to the ends of the icy samples would be complicated and impractical under the best of conditions, and the tensometer used in these tests is incapable of applying such a force regardless. Instead, a Brazilian test is used (Cleasson *et al.* 2002); this is a standard industrial measure (ASTM 1993, Guo *et al.* 1993, Vardar & Finnie 1975) used to calculate tensile strength for isotropic materials such as concrete – and including ice-silicate mixtures (Hiraoka *et al.* 2008) - and should be valid for the reasonably uniform distribution of the sand grains inside the sand-ice samples. The Brazilian test involves applying a crushing force to the sides of a cylindrical sample. This squeezes the sample and produces a strain perpendicular to the direction of application of force, thus producing the stretching effect needed to measure tensile strength.

The equation to calculate Brazilian tensile strength is as follows:

$$\sigma_t = \frac{F}{\pi LR} \quad (\text{Eq. 4.2})$$

where σ_t is the tensile strength, F is the breaking force, L is the height of the sample and R is its radius.

Five to six samples of each type of sand available were crushed in a tensometer to determine the average compressive strength of that sand type; the results are displayed in figure 4.3.2. While the experimental errors involved in the strength calculations are reasonably small, the error bars on the graph are plus/minus one standard deviation from the mean for that set of samples, and thus illustrate the amount of scatter in the strength measurements for that sand type. Also displayed is the average compressive strength of a set of pure ice samples for comparison; this is plotted at a sand grain size of 10000 μm in order to set it apart from the other points on the graph but the samples

Table 4.3.1 – Table showing the results of compressive strength tests on samples constructed of various different types of sand grains. Errors in strength are represented by the associated standard deviation; errors in grain size are given by the standard deviations in figure 4.2.1

Sand type	Grain size (µm)	Number of samples	Compressive strength (MPa)	Standard deviation (%)
Fraction A	2007	4	3.6	48
Fraction B	1043	5	5.11	16
Kiln sand	617	5	5.53	34
Fraction C	605	4	5.7	40
Fraction D	229	4	7.38	44
Fraction E	160	5	7.4	29
Pure ice	None	5	1.64	61

Table 4.3.2 – Table showing the results of tensile strength tests on samples constructed of various different types of sand grains. Errors in strength are represented by the associated standard deviation; errors in grain size are given by the standard deviations in figure 4.2.1

Sand type	Grain size (µm)	Number of samples	Tensile strength (MPa)	Standard deviation (%)
Fraction A	2007	4	1.36	6
Fraction B	1043	5	0.86	36
Kiln sand	617	5	1.21	22
Fraction C	605	4	1.7	6
Fraction D	229	4	1.38	9
Fraction E	160	4	1.56	28
Pure ice	None	4	0.62	29

themselves contained no sand whatsoever. Since adding sand to ice increases the strength of the sample, pure ice samples represent the theoretical lowest limit of strength that an ice-silicate mix should possess. This limit is represented by the red line in figure 4.3.2.

From figure 4.3.2 we can see that adding sand to ice makes the resulting mixture stronger than pure ice by a factor of between two and five depending on the size of the sand grains used. While the amount of scatter in the compressive strength measurements is large and it is conceivable that the compressive strength could be a

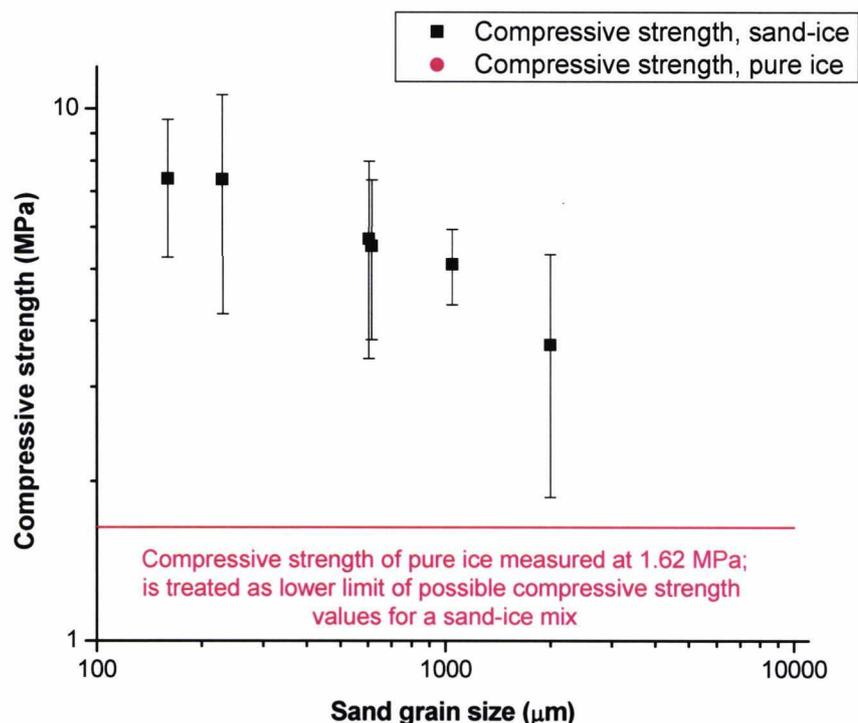


Figure 4.3.2 – Graph plotting the average compressive strength of various sand-ice mixes against the mean grain size of the sand type used in the mix. For comparison, the value for pure water ice is shown in red at the right of the graph.

constant, there does appear to be a slight decrease in strength as the size of the sand grains used in the sand-ice mixes increases. The smallest sand grain sizes (fractions D and E) produced near-identical results and were the strongest samples, with compressive strengths on the order of 7.4 MPa. The kiln sand and fraction C sand grains differ in size by just 12 μm and so it is unsurprising that they too have almost identical compressive strengths of 5.5 and 5.7 MPa respectively. Fraction B samples are only slightly weaker than fraction C/kiln sand at 5.1 MPa, but increasing sand grain size up from 1043 to 2007 μm results in a large decrease in compressive strength for the fraction A samples at 3.6 MPa. Changes in grain size as a *relative* increase seem to result in more or less constant associated increases in compressive strength.

Figure 4.3.3 shows the same sand types tested for tensile strength. Similarly to the compressive strength there is an increase in the strength of sand-ice mixes over pure

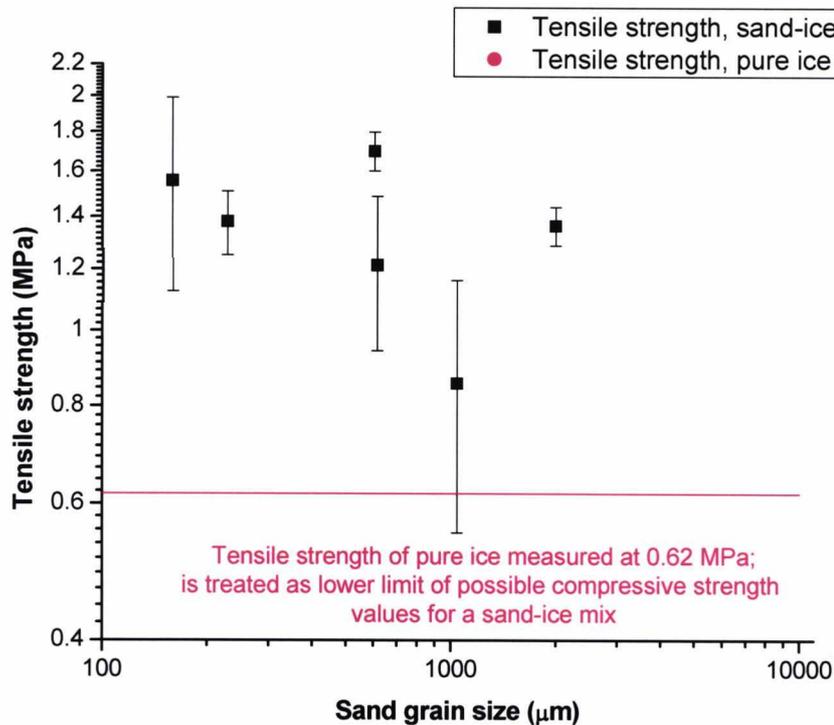


Figure 4.3.3 – Graph of the tensile strength of various sand-ice mixes plotted against the grain size of the sand type used in the mix.

ice alone, although this increase is less pronounced being only two to three times higher than the tensile strength of pure ice. Unfortunately the relationship between grain size and tensile strength is much less clear; indeed, there does not appear to actually *be* one. If the fraction B data point is discounted due to the significant amount of scatter associated with it, then it could be said that there is a near-constant value of tensile strength at around 1.3 – 1.5 MPa that does not vary with sand grain size, but this conclusion is tentative at best.

It is also possible that there is some flaw in the experimental method that led to inconsistent results; one explanation could lie with the way the samples were constructed. Briefly, samples of sand-ice mix that are removed from moulds and left to stand in the freezer for long periods of time (several days) tended to suffer from ice sublimation. Exposed icy surfaces that had previously been a smooth mix of sand and ice now consisted of rough sand half-embedded in an ice layer situated a little way back from the original surface of the sample when removed from the mould.



Figure 4.3.4 – Fraction E samples which have cracked while undergoing tensile strength tests.

While this was a problem that affected samples used for both the compressive and tensile strength tests, the flat top and bottom of the sample cylinders were unaffected thanks to the ice cap protecting the top and the sample being stood on its bottom; these areas were thus unexposed unlike the rest of the surface area of the target. Compressive strength tests use these unexposed areas as contact surfaces to calculate the compressive strength of the sample, and are thus less likely to suffer from this sublimation effect. However, tensile strength tests do use the exposed area in their strength calculations, which may account for the inconsistent nature of the results.

Also of interest is the differing manner in which the samples suffered failure during tensile strength tests. Figures 4.3.4 and 4.3.5 show this, comparing fractured samples made using Fraction E sand, the finest type available, against samples using Fraction B sand, one of the coarsest. Fraction E samples display visible cracks down the axis of the sample when they fail, but increasing the crushing force further does not widen this crack to the point where the sample splits altogether; instead the end result is usually the sample undergoing catastrophic failure, breaking apart into several unequal fragments. In contrast the Fraction B samples almost always split neatly down the long axis upon failure.

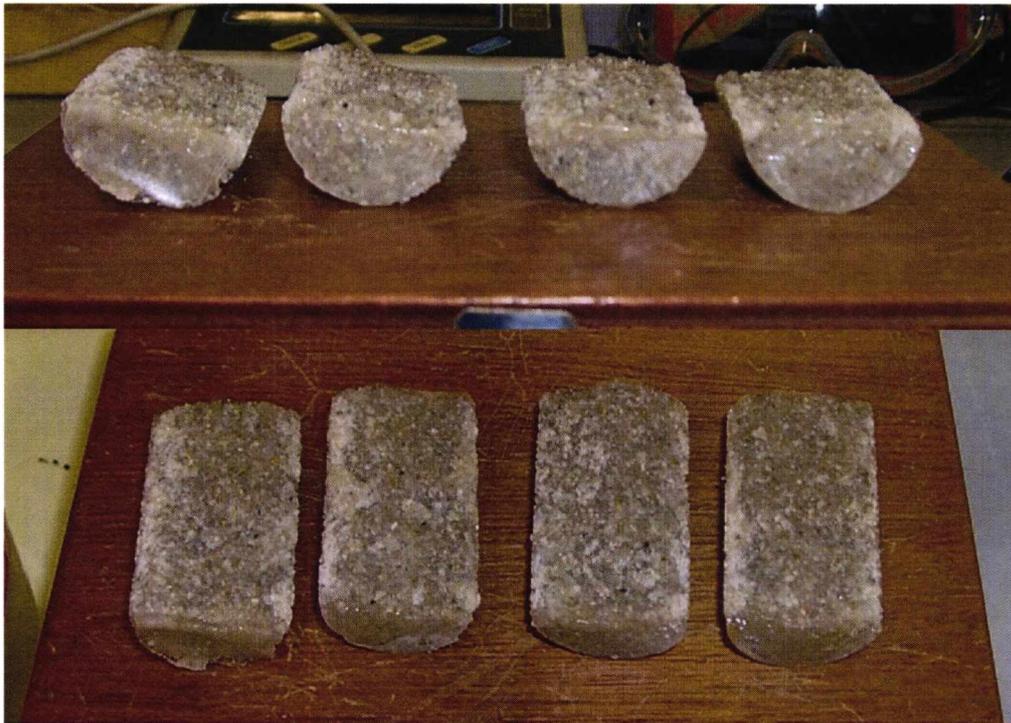


Figure 4.3.5 – Fraction B samples which have split while undergoing tensile strength tests.

Compressive strength tests were a different matter; neat splitting was usually observed in samples made up of the smaller sand grain sizes, but sometimes the samples showed a tendency to deform under pressure rather than cracking or splitting. This made the breaking force of these sample types difficult to determine since there was no obvious point at which the sample could be said to be definitively “broken”. Instead, the breaking force was instead taken to be the point where the measured applied stress on the sample peaked and then suddenly decreased with the rate of strain; this marks either a breaking strength that is not indicated by any large external cracks on the sample (i.e. internal failure) or, failing that, the yield strength for the sample marking the point at which it starts deforming plastically rather than elastically. It is not known why some samples behaved in this manner; plastic properties point to the samples being slightly melted in some way and thus able to deform more freely, but the samples were stored in a freezer until the time came for them to be tested in order to ensure that melting was not possible. Again, repeating these tests in a cold room environment would ensure that any anomalous results caused by melting are eliminated.

4.4 - Comparison with previous work

Lange & Ahrens (1983) measure the tensile strength of ice as 1.6 MPa and frozen silt as 5-6 MPa for similar strain rates to the ones encountered here (on the order of 10^{-3} s^{-1}). Hiraoka *et al.* (2008) used similar methods to investigate how the tensile and compressive strengths of icy samples changed with the percentage rock content present in the sample. While this work investigates strength variance with grain size there are some points of comparison that can be made; in particular the strength values for pure ice can be compared with those obtained here. Hiraoka *et al.* measure the tensile strength of ice as approximately 0.7 MPa, in good agreement with the result of 0.62 MPa measured in this study. However, they give a result for compressive strength calculated in an earlier paper (Arakawa and Tomizuka, 2004) of approximately 5.5 MPa, which is several times higher than the value of 1.64 MPa obtained here. The samples used in the Hiraoka paper are not simple frozen pure water ice; instead they are composed of commercial ice that has been crushed down to small particles on the order of 100 μm , and then mixed with chilled purified water and frozen. This different method may be the reason for the different compressive strength values – indeed, giving their samples an internal structure similar to the sand-ice samples used here (albeit with pure ice particles in place of sand grains) might be consistent with an increased value for compressive strength over simple pure water ice.

However this is purely speculative; Hobbs (1974) provides further strength measurements for commercial pure ice for comparison, and these are on the order of 4-5 MPa for compressive strength and 1.5 MPa for the tensile strength. Each of these measurements exceeds by two to three times the values for compressive and tensile strength obtained here. These values are temperature dependent, however – ice is stronger at low temperatures, and the above values correlate with a temperature of minus 20°C. While this is the temperature the samples used in this work were initially frozen at, strength tests were performed at room temperature mean that the samples may undergo an unknown amount of heating before cracking in the tensometer, which would lead to lower overall strength values. Hobbs gives strength values for ice at 0°C at 2 MPa for compressive strength and 1.4 MPa for tensile strength; this compressive strength value at least is more consistent with the value of 1.62 MPa measured here.

Another reason for the differences may be the quality of the ice used in this work. While the water was purified to a reasonable degree prior to freezing and there were no large-scale cracks present in the samples when they underwent testing, there were tiny flaws in the samples caused by trapped gas bubbles that had not been removed by the purification process. These flaws may have contributed to the lower strength values.

4.5 - Conclusion

In this chapter I have explained the basic properties of the materials used in the main experimental programme. The vital statistics of all major projectile types are given, along with the main reasoning behind their use in impacts on a given set of targets. An explanation of the properties of the four types of sand grains used in impacts is given, as well as the results of strength measurements on cylindrical samples made up of ice mixed with the different grain types. Sand-ice mixes are shown to be much stronger than ice on its own by a factor of between two and five for the compressive strength depending on the size of the sand grains used, and between two and three for tensile strength. Compressive strength decreases as grain size increases, but the results for tensile strength are inconclusive; they appear to indicate a constant value of around 1.3 – 1.5 MPa for all grain sizes but the amount of scatter present in the data makes this conclusion tenuous. Regardless of the overall relation it is clear that adding sand to ice significantly strengthens the resulting mix, and so it is expected that targets composed of sand-ice will require more energy to disrupt via impact than pure ice. There may however be some dependence on grain size depending on whether compressive or tensile strength is most important.

Chapter 5 – Impact Results and Q* Calculations

“You're dead, this is the afterlife... and I'm God.”

“You are not God!”

“Blasphemy! You're lucky I don't cast you out or smite you or something. The bottom line is, your life ended about five minutes ago.”

“No. I am not dead. Because I refuse to believe that the afterlife is run by you. The universe is not so badly designed!”

- Q and Captain Jean-Luc Picard, *Star Trek: The Next Generation*

5.0 – Introduction

Now that the strengths of the different varieties of material used in target construction are known, it is time to investigate how targets composed of these materials actually behave under impact conditions. The main goal of this chapter is to analyse the impact outcomes, and quantify the critical energy density Q^* for each target type and to see how it changes from target to target. This will hopefully provide some insight into how factors such as temperature, grain size and the presence or non-presence of a covering ice layer alter the impact result.

5.1 – Impact outcomes

The most obvious way of analysing an impact outcome is to visually examine the remaining fragments and note any similarities or differences between their general size and makeup from impact to impact and target to target. The two main types of target used in this experimental programme were unlayered kiln-sand cores and layered kiln-sand cores; other target types had too few shots performed for a simple visual inspection to note any common trends in the way those targets disrupted.

5.1.1 - Kiln-sand cores

The impact outcomes of the kiln-sand cores can be divided into several distinct categories. These categories are:

- Cratering (figure 5.1.1), where a core has a very small amount of its mass removed ($M_f/M_o = 0.9$) but otherwise remains mostly intact.



Figure 5.1.1 – A cratered kiln-sand core. Scale: the squares in this and all subsequent images are 2.55 x 2.55 cm.

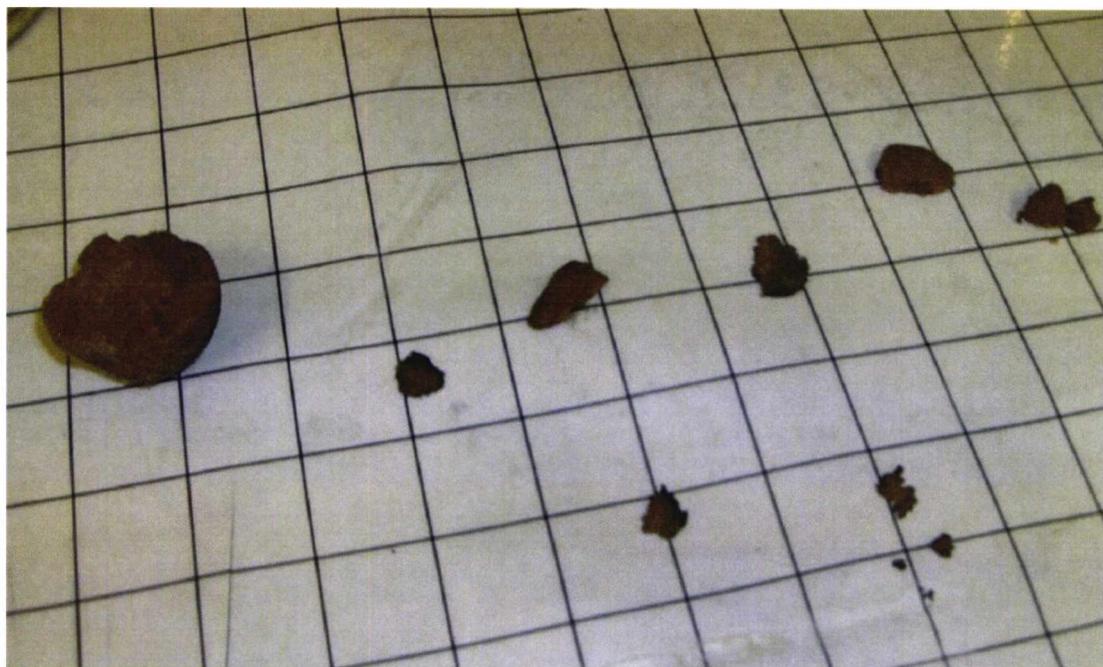


Figure 5.1.2 – A near disrupted kiln-sand core. Fragments have undergone partial melting either due to energy released from the impact itself or else from warming up to room temperature.



Figure 5.1.3 – A disrupted kiln-sand core. Fragments have undergone partial melting since impact as they warm up to room temperature.

- Near disruption/large scale cratering (figure 5.1.2), where a core has a larger amount of its mass removed ($M_f/M_o = 0.5 - 0.9$) but still not enough to disrupt it.
- Disruption (figure 5.1.3), where a target has more than half of its original mass removed but still leaves one sizeable large fragment along with numerous smaller fragments. ($M_f/M_o = 0.1-0.5$)
- Shattering, where a core is completely shattered into many small fragments. Unfortunately the fragments are so small and numerous that after each one has been weighed for size distribution analysis many of them are almost completely melted, making it impossible to get a decent image of a shattered target.

These general categories of impact outcome are solely related to the relative energy density they have been impacted with. While it is difficult to quantify this relationship in exact terms (there is too much variation in the outcome from impact to impact - not to mention too little data overall - to specifically state where one impact outcome

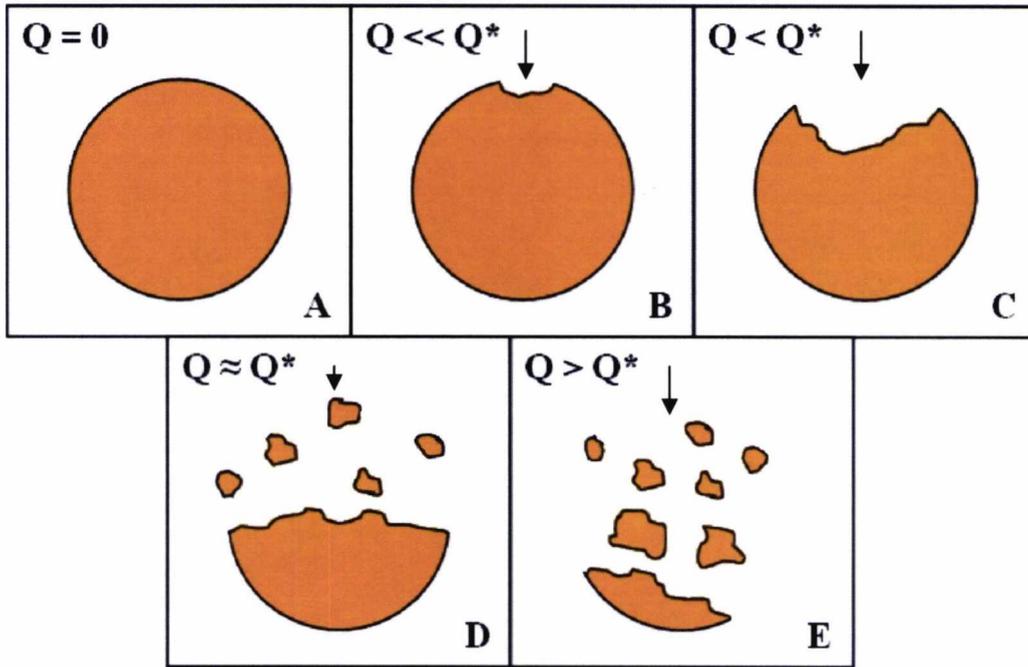


Figure 5.1.4 – Diagram showing the differing outcomes of impacts on to spherical cores as they vary with the impact energy density Q relative to Q^* . Arrows show the orientation of the impacting projectile.

gives way to another) it is at least possible to outline it in a fairly general way, as shown in figure 5.1.4. This scheme can be described in words as:

$Q = 0 \text{ J kg}^{-1}$ - an unimpacted spherical core, shown for reference purposes.

$Q \ll Q^*$ - an impact with an energy density very much smaller than the critical energy density for the spherical core removes only a very small amount of mass from the core, resulting in a small crater. Any fragments produced have either negligible mass (on the order of just a few sand grains) or melt before they can be measured as a result of the room temperature environment.

$Q < Q^*$ - Increasing the energy density so that it is larger but still does not approach the critical energy density simply results in a larger crater. There may be one or two fragments of measurable size produced, but usually the fragments are on the same small scale as in $Q \ll Q^*$; that is, not quite negligible but small enough to melt before they can be measured.

$Q \approx Q^*$ - As the impact energy density approaches the critical energy density Q^* , around half of the original mass of the target is removed; the largest surviving fragment tends to be the rear half of the core relative to where the projectile impacted. Many fragments are produced, several of which are also reasonably sizeable in comparison to the original mass of the target ($M_f/M_o \approx 0.05 - 0.2$).

$Q > Q^*$ - When the impact energy density clearly exceeds the value of Q^* for spherical cores, the target shatters into many fragments. A small portion of the rear of the target may still survive and this is often the largest fragment as for $Q \approx Q^*$, although far smaller at around $M_f/M_o = 0.1$. However, there are invariably now multiple other fragments that approach or sometimes even exceed this in size, as well as numerous other fragments containing significant fractions of the mass of the original target.

Increasing the energy density further ($Q \gg Q^*$) simply grinds the target down into smaller and smaller individual fragments with no significant changes to the disruption mechanism.

The different categories of impact outcome for kiln-sand cores correlate to the relative energy density it has been impacted with; $Q \ll Q^*$ produces a cratering outcome, $Q < Q^*$ near disruption, $Q \approx Q^*$ disruption and $Q > Q^*$ shattering.

5.1.2 - Layered kiln-sand cores

The range of outcomes for layered cores was wider than that for plain cores because of the wide variety of things that could happen to the surface ice layer during impact. It could either be cratered, sheared partially off or sheared entirely off depending on the impact energy, and this sometimes appeared to be independent of what happened to the inner core.

Layered cores appear to follow roughly the same pattern of cratering (figure 5.1.6, 5.1.7), disruption (figure 5.1.9, figure 5.1.10) and shattering as the unlayered cores (figure 5.1.4); while the behaviour of the ice layer varies, it was very



Figures 5.1.5 (left) and 5.1.6 (right) – A cratered layered core that has retained its ice layer (left) and a cratered layered core that has had its ice layer removed (right). The targets were struck with roughly the same energy density (108 and 115 J kg^{-1} respectively) but possessed different ice shell thicknesses (5.98 and 3.02 mm respectively).



Figures 5.1.7 (left) and 5.1.8 (right) – Examples of near-disrupted layered cores, one with the ice layer still attached (left) and one with the ice layer mostly removed (right). This time there is a sizable difference in the energy density each target was impacted with (212 J kg^{-1} and 341 J kg^{-1} respectively) but only a small difference in ice layer thickness (2.23 and 2.73 mm respectively).

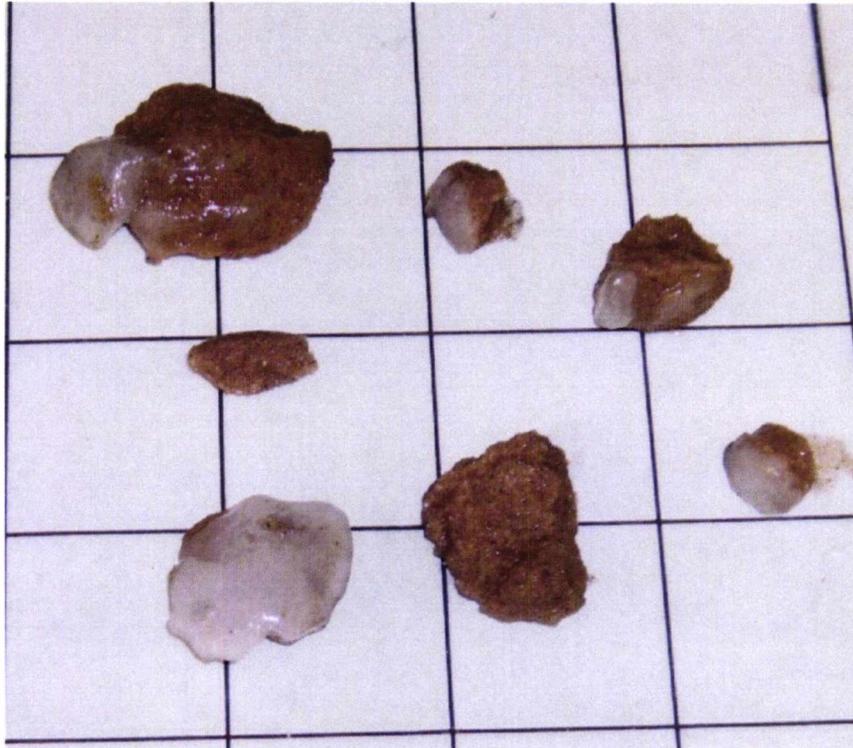


Figure 5.1.9 – A disrupted layered core with pieces of the ice layer still attached to some of the fragments.



Figure 5.1.10 – A disrupted layered core where the ice layer has been almost totally removed from the remaining fragments.

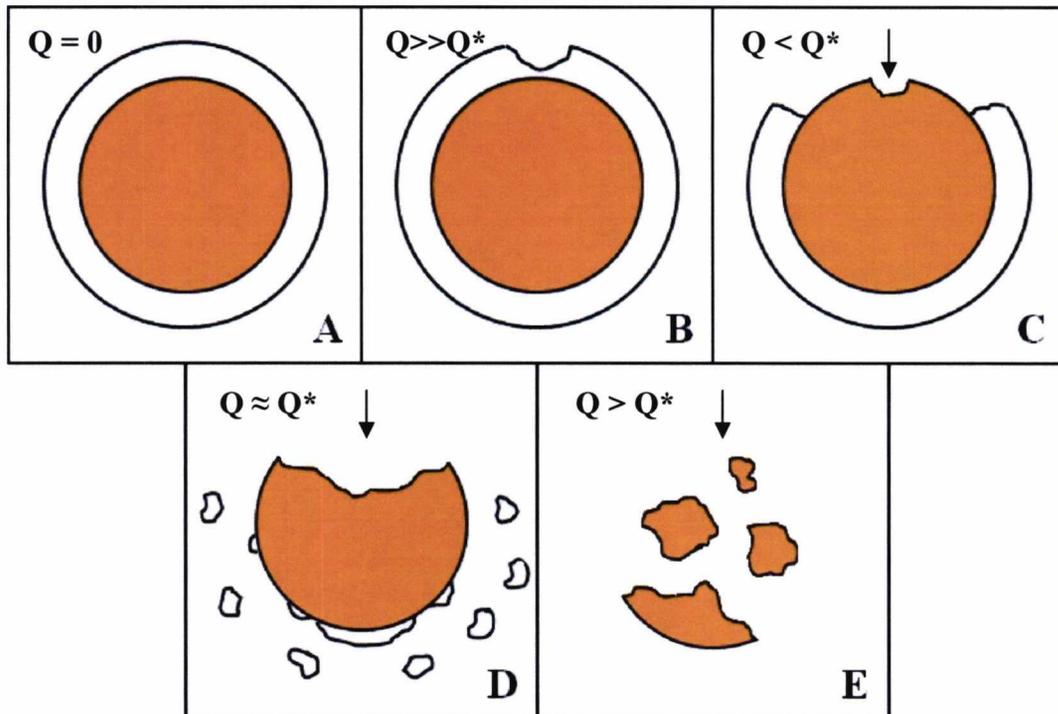


Figure 5.1.11 - Diagram showing the differing outcomes of impacts on to layered targets as they vary with the impact energy density Q . Arrows show the orientation of the impacting projectile.

unusual for it to be sheared off completely at low energy densities, and it was likewise unusual for a fully disrupted target to retain sections of ice layer on the individual fragments. However, the ice layer alters each impact outcome in a subtle way, as seen in figure 5.1.11 and described below.

$Q = 0 \text{ J kg}^{-1}$ – Cross-section of an unimpacted layered target, with a white ice shell covering the orange sand-ice core.

$Q \gg Q^*$ - A very low energy density impact produces a very small crater, similar to the unlayered core seen in figure 5.1.1. However, while roughly the same amount of mass is lost the volume of material removed from the target is greater due to the projectile not possessing enough energy to penetrate down to the sand-ice core and thus excavating material solely from the ice layer; the lower density and disruption energy of ice means that more material is removed relative to an impact with the same energy density onto a sand-ice core alone.

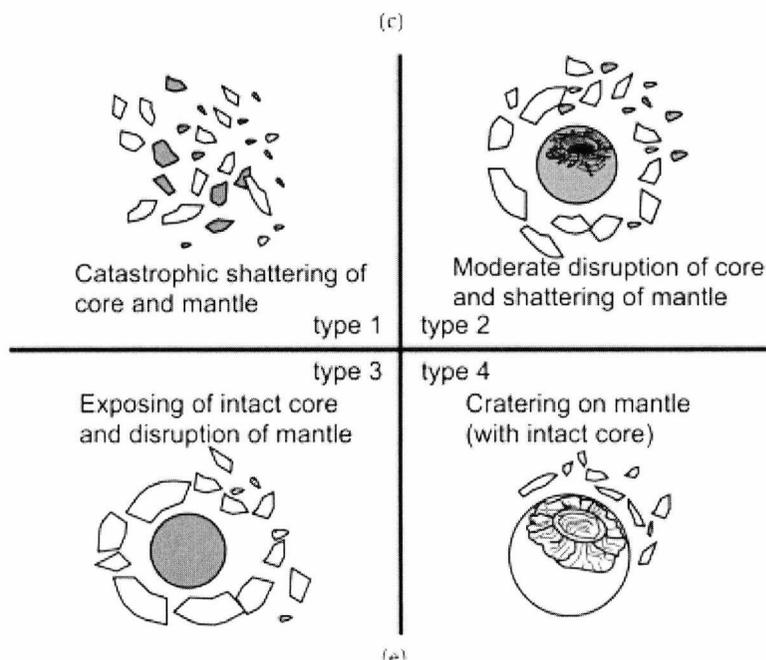


Figure 5.1.12 – Stages of disruption observed in a soda-lime glass core/gypsum mantle target (Okomoto *et al.* 2008)

$Q > Q^*$ - Increasing the energy density gives the projectile enough energy to punch through to the sand ice core. However, much of the projectile's energy is expended removing a large portion of the surrounding ice layer, so the resulting crater in the core is much smaller than it would be on an unlayered core.

$Q \approx Q^*$ - Impacting a layered target with the critical energy density has an interesting effect; much of the energy appears to go into disrupting and removing the entire ice layer, with only a small amount of ice remaining attached to the rear of the target. The amount of mass removed from the core is consequently smaller in comparison to an unlayered core, with much more than 50% of its mass surviving impact. Unlike lower energy density impacts the fragments of the ice layer removed are large enough to allow measurement

$Q > Q^*$ - With impacts of $Q \approx Q^*$ removing the ice layer, increasing the energy density further imparts all additional energy directly into the core. This disrupts it more completely than in the example of $Q \approx Q^*$, producing several fragments. However, this differs from the same relative energy density for unlayered cores; the protection the layer affords the core robs the impact of much of the energy that would

otherwise go into disrupting it, and so the fragments recovered are much larger than those recovered from the shattered unlayered core at similar Q^* values.

It is worth noting that this disruption sequence roughly matches that observed in Okomoto's work on soda-lime glass core/gypsum mantle targets (2008), with a lightly cratered mantle suddenly shearing off completely leaving an undamaged core, moving on to a sheared-off mantle and damaged core, and finally giving way to total disruption of the target

The interesting thing is that since each impact outcome is connected to the relative energy density (essentially the ratio Q/Q^*) that produced it, the changes in the disruption mechanism caused by the addition of an ice layer are theoretically independent of any increase in strength an ice layer might impart, because that would be reflected in a higher value of Q^* . The logical next step here would be to conduct an analysis of the size distribution of the recovered fragments in an attempt to investigate this difference in disruption mechanism in more detail (see Chapter 6 for size distribution analysis of other target types). Sadly, all of the impacts on layered targets took place before the introduction of the improved target chamber described in Chapter 3, and so most of the smallest fragments produced after impact – especially those of the ice layer – melted before they could be measured. Such an investigation will therefore have to be left for future work.

5.2 – Determining Q^* from impact data

Having taken a general look at impact outcomes, we now move on to calculating specific values of Q^* for each target type.

As stated in Chapter 2, plots of Q vs M_f/M_0 tend to follow the trend shown in figure 5.2.1. The trend shows three distinct types of behaviour depending on the magnitude of Q , which can be split into three separate regimes according to the categories identified in section 5.1. These are:

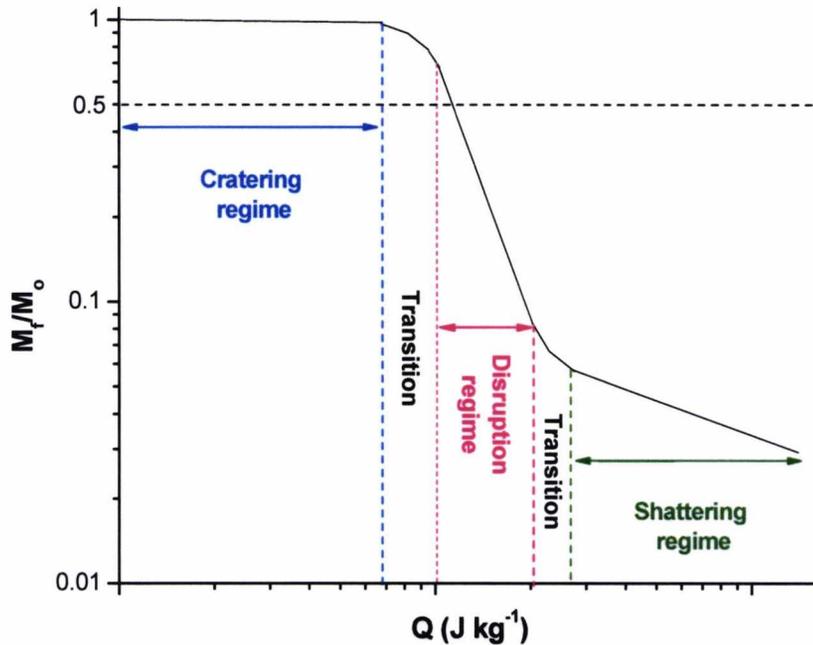


Figure 5.2.1 - Typical plot of how M_f/M_0 varies with Q (henceforth referred to as a *disruption curve*). The width of certain features (e.g. transition zones) have been exaggerated to make them clearer.

- The *cratering* regime, involving a shallow decrease in M_f/M_0 as Q increases; the impacts are not yet powerful enough to remove more than small amounts of mass from the targets.
- The *disruption* regime, where the amount of mass being removed increases dramatically with Q , resulting in a very steep decrease in M_f/M_0 as Q increases. This can be thought of as grinding one large fragment down into smaller and smaller pieces; since it is the size of the large fragment that dictates M_f/M_0 the fact that much of the target mass survives in other fragments is irrelevant. Q^* always falls within the disruption regime.
- The *shattering* regime; at this point the target has been shattered into many small fragments comparable to each other in terms of size. Further increases in Q must reduce each of these small fragments down further (producing yet more small fragments that increase in number exponentially) in order to

reduce M_f/M_0 (as opposed to reducing the mass of one large fragment), dissipating the effect of energy increase and resulting in a shallow decrease in M_f/M_0 with increase in Q similar to the cratering regime.

The transition from one regime to another is not quite instant, and so separating each regime are transition zones in which the trend of one regime makes a gradual transition to the trend of the next regime. Exactly how broad these transition zones are in terms of Q is not known, but for the purposes of illustration it is shown to be quite large in figure 5.2.1.

The addition of the shattering regime to the disruption curve means that further constraints are placed on which data points can be used to calculate Q^* . Impact results from the shattering regime must be excluded from any such calculation; including them – or even worse, relying on data taken from the shattering regime alone – will result in a systematic underestimation of Q^* thanks to the shallower slope of the curve.

To find Q^* , then, the data are examined and the points corresponding to the disruption regime are singled out. Since it is not known *a priori* which factors affect Q^* - such as projectile density, projectile size, target size, target shape etc. – the impacts here are first separated into specific projectile/target combinations to ensure that the same impact conditions apply across a data set. The data gathered from impacts mostly covers the cratering and disruption regimes, with few shots possessing a high enough energy density to put them in the shattering regime. Since by definition it is during the disruption regime that the plot crosses the $M_f/M_0 = 0.5$ boundary, it is this portion of the data that we attempt to fit. Unfortunately since Q^* and the shape of the disruption curve change from projectile/target combination to projectile/target combination, it is not possible to formulate a specific set of criteria that can be applied across all cases to select which data points should be excluded from the Q^* calculation.

Therefore any fits to the data are done on an ad hoc basis that varies from data set to data set. Since we are only interested in the part of the curve that describes the disruption regime, any points that are obviously part of the cratering regime are

excluded. Points that may be part of the shattering regime are also excluded, although these are far less in number due to the high energy densities required to place a point in this data range. Certain anomalous data points that greatly distort the fit away from the majority of points are excluded as well; these are easy to spot since they will have an abnormally large or small M_f/M_o value for their energy density compared to the rest of the data.

The error on an individual measurement of Q is very small. The projectile velocity is known to within 1% and target mass to within less than 1%; the error in projectile mass is unknown, but judging from the data used to calculate the averages described in Chapter 4 it too is less than 1%; taken together the statistical error in a given measurement of Q is no more than $\pm 2\%$.

M_f/M_o is a different matter in that the scales used to weigh fragment masses measure them to a constant $\pm 0.05\text{g}$, resulting in a percentage error that increases as the size of the fragments decreases. The smallest value of M_f/M_o recorded during the experimental program is 0.01; this has an associated error of 8%, which therefore is the upper limit on errors in M_f/M_o .

All Q^* values presented in this chapter were calculated by imposing the exclusion criteria described above on the data (as a general rule, anything in the $M_f/M_o = 0.9$ - 1.0 range was excluded, while anything within $M_f/M_o = 0.8$ - 0.9 may also have been excluded depending on how anomalous it appeared to be compared to the rest of the data set; this should have had the result of excluding any data points that might be part of the cratering regime) and then using Origin to fit the remaining data points using an allometric fit.

An allometric fit produces a fit according to the following equation:

$$\frac{M_f}{M_o} = aQ^b, \quad (\text{Eq. 5.1})$$

which can be rearranged to

$$Q = \left[\frac{\left(\frac{M_f}{M_o} \right)}{a} \right]^{\frac{1}{b}}, \quad (\text{Eq. 5.2})$$

Since Q^* is the value of Q corresponding to $M_f/M_o = 0.5$, substituting this into the equation gives the final form

$$Q^* = \left[\frac{0.5}{a} \right]^{\frac{1}{b}} \quad (\text{Eq. 5.3})$$

Calculating Q^* is then a simple matter of substituting the values for a and b from the specific power law for a given target type into equation 5.1.

This method allows very precise calculation of Q^* ; however calculation of the associated errors is not so easy. Origin provides errors in both the a and b terms of the above equations, but there is no way to disassociate these errors from each other. As a general rule we are only interested in the b term of the equation since it is this which determines the slope of the fit, with a being a simple scaling factor. Since it is not possible to separate the a and b errors to see how the errors in b only affect the resultant value of Q^* , it is only possible to use the Origin-derived errors to provide an overall Q^* error when the error in a is reasonably small.

Faced with this limitation, two different methods have been used to calculate error. For cases where the error in a is small, error in Q^* is calculated by using the error in b to calculate maximum and minimum Q^* values; the difference between these maximum and minimum values and the actual value of Q^* is then treated as the error. For cases where the error in a is large a less empirical method must be used: data points are systematically excluded from the set one by one starting with the most extreme values (i.e. those furthest away from $M_f/M_o = 0.5$) and a range of Q^* values are generated which reflect the sensitivity of the fit to the data. If removing a single data point produces a large change in the determined value of Q^* then that data set is

Table 5.1 – Table showing fitted Q^* values, associated errors and methods of error calculation.

Projectile	Target	a	a error	b	b error	R^2	Q^* ($J\ kg^{-1}$)	Q^* error + ($J\ kg^{-1}$)	Q^* error – ($J\ kg^{-1}$)	Error method
1mm titanium		1083	856	-1.34	0.14	0.99	309	55	4	Exclusion
0.8mm stainless steel	Kiln-sand core	3641	7836	-1.58	0.31	0.64	278	7	13	Exclusion
1mm aluminium	Solid ice core	291.1	0.0	-1.67	0.01	0.98	45	3	3	b error
1mm titanium	Thick shelled layered kiln sand core	1.4	0.1	-0.15	0.02	0.86	913	20	18	b error
1mm titanium	Thin shelled	3412	5414	-1.50	0.28	0.85	360	1	12	Exclusion
0.8mm stainless steel	layered kiln sand core	61.6	143.8	-0.86	0.42	0.48	270	2	26	Exclusion
1.98mm glass		217	492	-1.24	0.50	0.79	134	110	11	Exclusion
1mm titanium	Low temperature kiln sand core	757	1504	-1.27	0.14	0.73	319	22	85	Exclusion
0.8mm stainless steel	Fraction B	36.9	8.7	-1.14	0.07	0.99	44	13	11	b error
0.8mm stainless steel	Fraction D	1707	1597	-1.44	0.18	0.99	284	4	100	Exclusion
0.8mm stainless steel	Fraction E	1506	313	-1.41	0.03	0.99	293	26	9	Exclusion
2mm copper	Kiln-sand cylinder	91.7	114.4	-1.04	0.25	0.84	150	18	1	Exclusion
1mm copper	Solid ice cylinder	4.6	1.4	-1.06	0.18	0.99	8.1	4	2.7	b error
1mm copper	Murray sphere	3.8	0.7	-1.17	0.12	0.96	5.7	1.6	1.2	b error

presenting a very uncertain picture of what Q^* for that projectile/target type combination actually is, and this is reflected in the error, which is calculated by taking the maximum and minimum Q^* values from the range of fits generated and treating the difference between them and the actual value of Q^* as the error. This is not an ideal method of error determination, but it at least provides some idea of how accurate a given fit of Q^* is.

Q^* values for all data sets are presented in table 5.2.1 along with associated errors and R^2 values. R^2 is another measure of goodness of fit calculated by comparing a least-squares fit of the data to a simple average; the better the fit is the closer R^2 is to 1. Most data sets seen here are within the 0.8-1.0 range, indicating a reasonably good fit. The major exceptions are for 0.8mm stainless steel projectile impacts onto kiln-sand cores and thin-layered kiln-sand cores; the low R^2 values associated with these fits indicate that the Q^* values derived from them must be treated with a certain degree of suspicion.

5.3 – Q^* for homogenous targets

Before looking at what effect – if any – adding an ice layer to a body has on the critical energy density for that body, we must first establish what would happen to that body were it unlayered – what its base value of Q^* is, and which other factors might affect that value of Q^* besides the addition of an ice layer. Only after all these factors have been accounted for can a proper comparison between layered and unlayered bodies be made.

First, as an illustration of the baseline Q vs M_f/M_0 relationship for the kiln-sand material on its own, we have the results for 0.8mm stainless steel impacts onto kiln-sand cores compared with results for 2mm copper onto kiln-sand cylinders (figure 5.3.1). This shows a markedly lower Q^* value for kiln-sand cylinders ($150 \pm 18/-1 \text{ J kg}^{-1}$) as opposed to that for kiln-sand cores ($278 \pm 7/-13 \text{ J kg}^{-1}$). This could be caused by two factors aside from the difference in the shape of the target; firstly, this Q^* value is calculated using the results for 2mm copper projectiles (the 1mm copper impact results were all firmly in the cratering regime and thus not used for the

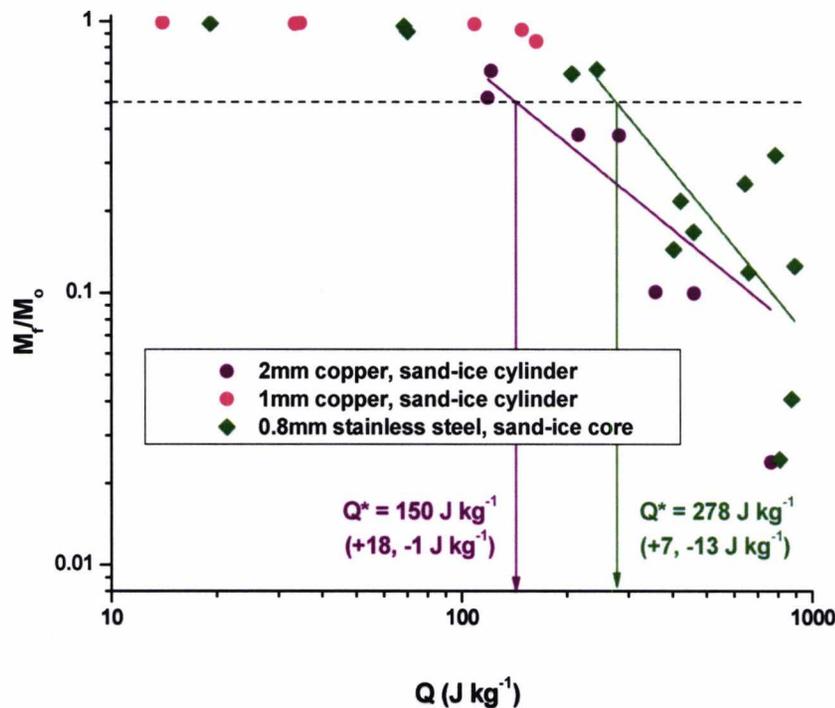


Figure 5.3.1 – Results for 2mm copper/1mm copper impacts onto kiln sand cylinders ($Q^* = 150 \text{ J kg}^{-1}$ for 2mm copper impacts) and 0.8mm stainless steel impacts onto kiln sand cores ($Q^* = 278 \text{ J kg}^{-1}$)

purposes of calculating Q^*). While copper has a similar density to the stainless steel projectiles used to impact the kiln-sand cores (8.96 g cm^{-3} and 8.00 g cm^{-3} respectively), the copper projectile has a diameter over twice that of the stainless steel. Moreover, 1mm copper impacts at the top end of the range of energy densities the light gas gun was capable of producing merely cratered the target. However, the upper range of Q for 1mm copper projectiles overlaps with the lower range of Q for 2mm copper projectiles; 2mm copper projectiles that impacted with similar energy densities to the 1mm copper projectiles did not crater the target, but instead removed 40-50% of the total target mass, placing them in the disruption regime. Since the projectile material, target type and impact energy density Q are all the same, the conclusion to draw from this is that it is the doubling of the projectile diameter that has resulted in this more destructive impact result. A larger projectile does therefore appear to disrupt a target at a lower value of Q^* than would be obtained by impacts using a smaller projectile, and so the difference in Q^* between the kiln-sand cores and the kiln-sand cylinders is at least partially caused by the larger copper projectile.

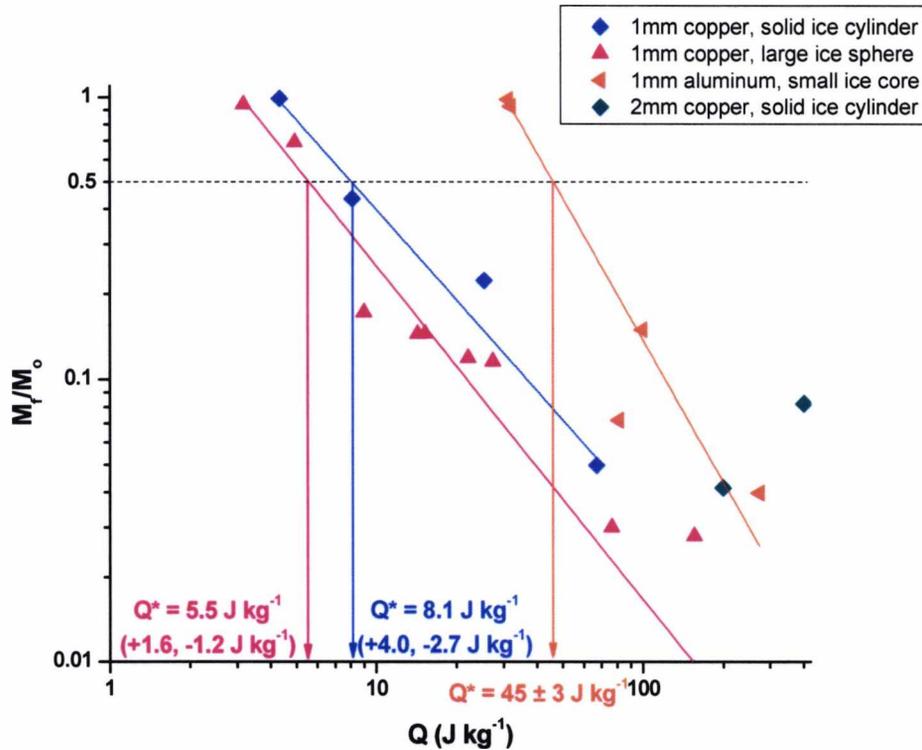


Figure 5.3.2 – Results for 1mm copper impacts onto large solid ice spheres (Murray 2004, $Q^* = 5.5 \text{ J kg}^{-1}$), 1mm copper impacts onto solid ice cylinders ($Q^* = 8.1 \text{ J kg}^{-1}$), and 1mm aluminium impacts onto solid ice cores ($Q^* = 45 \text{ J kg}^{-1}$).

Secondly, the kiln-sand cylinders had many times the mass of the kiln-sand cores; the approximate mass of a core is 60g, whereas the average mass of the cylinders was 747g, over twelve times that of the cores. While a direct size comparison is not necessarily correct due to the differing target shapes, since both target types are constructed of the same kiln-sand material, mass is equivalent to volume; the cylinders are therefore twelve times as large as the cores. As stated in chapter 2 the size of the body being impacted affects the amount of energy required to disrupt it; the critical energy density decreases as target size increases due to easier crack propagation in larger bodies, which continues until a body becomes large enough for its own self-gravity to counteract this effect. While this effect was modelled for bodies on the scale of a metre diameter and above, it is not inconceivable that a difference of an order of magnitude in target volume would be enough to have a small effect on the value of Q^* at centimetre scales. Benz and Asphaug (1999) model the decrease in strength with size for both ice and basalt bodies. Over an order of

magnitude increase in the volume of their body (3 cm to 6 cm radius increase) the critical energy density decreases by 25% for ice and 30% for basalt.

It therefore becomes difficult to determine which of these factors is responsible for the different Q^* values since it could be any combination of projectile size, target size, and target shape. Further investigation of the size and shape factors is possible, however; a series of impacts were carried out on large solid-ice spheres at Kent (Murray 2004). The data from these impacts can be combined with the results for impacts on solid ice cylinders and solid ice cores carried out during the shot programme for this work to produce the comparison seen in figure 5.2.5

The large solid ice spheres used in Murray (2004) had diameters of 8 – 12 cm and masses of 0.4 – 1 kg. The solid ice cylinders had masses of approximately 330g, while the solid ice cores had diameters of 4 cm and masses of 30g. Once again the cylinders are eleven times larger than the cores, while the large ice spheres are anywhere from thirteen to thirty-three times as large. All projectiles used to impact these targets were

1mm in diameter, although the projectiles used to impact the ice cores were aluminium and thus only a third as dense as the copper projectiles used for the cylinders and large spheres (2.952 g cm^{-3} as opposed to 8.9 g cm^{-3}).

The resulting values of Q^* shown in figure 5.3.2 are interesting. The values calculated for the large ice spheres and the ice cylinders are very close together at $5.5 (+1.6/-1.2) \text{ J kg}^{-1}$ for the ice spheres and $8.1 (+4.0/-2.7) \text{ J kg}^{-1}$ for the ice cylinders. This tallies well with what we know of the targets and projectiles; the same projectiles were used to impact each target type, and the mass of the ice cylinders approaches the lower limit of the mass range of the large ice spheres. Since some of the ice spheres are much larger than the cylinders however (up to three times for the largest spheres) we would expect the results for these larger spheres to lower the Q^* value from where it might be had a set of constant mass ice spheres been used. This would account for much of the difference in the Q^* value between large ice spheres and ice cylinders, and indicates that the shape of the target is not a significant factor in determining Q^* .

However, the result for the small solid ice cores is markedly different. At $45 \pm 3 \text{ J kg}^{-1}$ it is significantly larger than the Q^* values obtained for the cylinders and the large spheres. Since we have determined that target shape does not significantly affect Q^* , explaining this discrepancy must rely on the differing factors of target size and projectile material. Unfortunately, as with sand-ice cores and sand-ice cylinders it is impossible to disentangle these factors; we now know that target size has some effect on Q^* , but we cannot say how much of the Q^* increase is due to the smaller target size and how much is due to the lower density projectile.

For the sort of volume increase seen here (solid ice cylinders are ten times larger than the small ice spheres, while the large ice spheres are thirty times larger), Benz and Asphaug (1999) model decreases in Q^* of 25% for ice and 30% for basalt for a volume increase of a factor of ten (3 cm to 6 cm radius increase), and 45% for ice and 50% for basalt for a volume increase of a factor of thirty (3 cm to 9 cm radius increase). The result for the small ice spheres exceeds this modelled decreased value in both cases; there is a decrease in Q^* of over 80% from the small ice spheres to the ice cylinders, and the decrease approaches 90% when compared to the result for large ice spheres. Therefore it is fair to say that size alone cannot account for this dramatic drop in Q^* , but isolation and investigation of how projectile density affects Q^* is necessary to take this further.

An attempt was made to test how solely varying projectile density varies Q^* by impacting 1mm diameter titanium projectiles into kiln-sand cores and comparing the results with impacts by 0.8mm stainless steel projectiles into the same target type; there is only a small decrease in diameter between these two projectile types, and stainless steel has a density of 7.7 g cm^{-3} while titanium has a density of 4.5 g cm^{-3} . If projectile density had any effect on Q^* we would expect it to be reflected in a difference in Q^* between these two data sets. Unfortunately comparing the two leaves us no better off: titanium projectiles produce a Q^* of $308 (+55/-4) \text{ J kg}^{-1}$ while stainless steel projectiles produce a Q^* of $278 (+7/-13) \text{ J kg}^{-1}$. While there is a difference of 30 J kg^{-1} between Q^* for the two data sets, they are close enough together that it is impossible to separate the effects of the small decrease in projectile

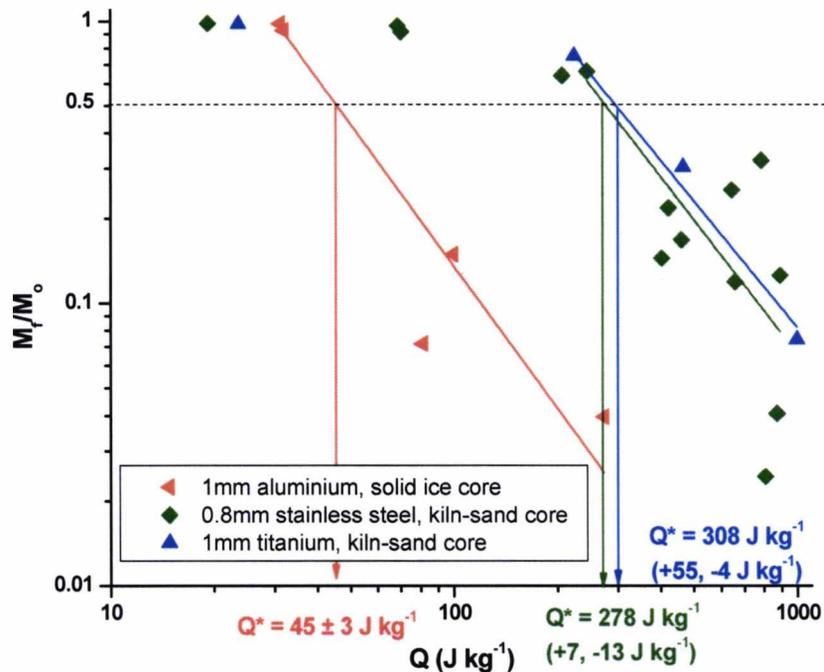


Figure 5.3.3 – Results of 1mm aluminium impacts onto solid ice cores ($Q^* = 45 \text{ J kg}^{-1}$), 0.8mm stainless steel impacts onto kiln sand cores ($Q^* = 278 \text{ J kg}^{-1}$), and 1mm titanium impacts onto kiln sand cores ($Q^* = 308 \text{ J kg}^{-1}$).

diameter and the increase in projectile density, making it very unclear as to whether there is any increase in Q^* due to projectile density at all.

We have so far managed to isolate two factors of an impact that can produce changes in the resulting value of Q^* :

- Projectile size
- Target size

But how do these factors relate to each other? How much will increasing the projectile size change Q^* in comparison with increasing the size of the target? Even under ideal conditions this would be a very tricky thing to investigate properly; since light gas guns are only capable of a limited range of impact velocities, it places constraints on how much one factor can be varied while keeping the others constant. For example, an increase in projectile size only will produce a commensurate increase in the values

of Q produced by that projectile. Assuming that target mass and projectile density remain constant, the initial radius of a projectile r_i may only be increased to $3.66r_i$; increasing the radius further will cause the larger projectile to reach values of Q when fired at even the slowest velocity of 1 km s^{-1} that cannot be matched by the smaller projectile fired at the fastest velocity of 7 km s^{-1} .

Target mass gives more leeway; this can be increased by up to forty-nine times theoretically allowing any practical range of target masses to be used. However, this theoretical maximum increases ignore two large problems. One is the need for a specific range of Q covering Q^* to be achieved with a specific target-projectile combination, not just one static value that may well fall in the shattering regime for that particular material. For example, a 1mm copper projectile fired at 7 km s^{-1} will match the Q of a 3.66 mm copper projectile fired at 1 km s^{-1} , but Q^* might fall at 5 km s^{-1} for the 1mm copper projectile, meaning that - even taking into account the fact that changing the projectile size will change Q^* - the range of Q produced by the 3.66 mm projectile will fall well outside the range of Q necessary to calculate Q^* .

The other problem is the laboratory issues involved with changing these attributes. First, it is not possible to select any projectile diameter between, say, 1 mm and 3.66 mm, since the specific projectile diameters available are limited by what can be sourced from a manufacturer, and cannot in any case exceed 3 mm since this is the maximum diameter of projectile that can be fired in the light gas gun. Similarly it is not easy to simply double the mass of a spherical ice-silicate target; the range of target sizes possible is limited by the size of moulds available. While it is possible to have moulds custom-made, this is both expensive and time-consuming; at the very least an investigation would be limited to only a few shots per size range rather than the dozens featured here. In addition to this the target size is also constrained by practical considerations imposed by the gun: the lower limit is on the order of a few centimetres since this is the smallest size of target that can be reliably hit centre-on, while the maximum diameter of target that can be fitted inside the target chamber is approximately 30 cm. For these reasons any future work that did attempt to study these Q^* -altering factors in isolation would be severely constrained in scale due to practical limitations.

Here, however, we are limited to these data sets that have two or more factors changing from set to set making quantification of their individual effect impossible. While we can make a crude approximation of how significant each factor might be based on the magnitude of the Q^* shift, a future, more specific investigation will then be required in order to attach a more precise value to them.

5.4 – Q^* for layered targets

Having looked at how basic factors such as target size and projectile type affect Q^* , we now move on to examining the significance of adding an ice layer to a target.

As explained in chapters 3 and 4, the majority of impacts on layered cores were carried out using 1mm titanium projectiles. The layered cores had ice shells that varied in thickness anywhere from 1mm to 7mm thick. Plotting the results of these impacts en masse results in figure 5.4.1. At first glance, the plot does seem to show the typical Q^* vs M_f/M_0 relationship, albeit with a lot of scatter. However, segregating the targets by layer thickness reveals the reason for this scatter, as shown in figure 5.4.2. There are in fact *two* different M_f/M_0 plots depending on the thickness of the ice layer. The targets are separated into two categories: *thin* layered targets, where the ice layer is less than 5mm thick, and *thick* layered targets, where the ice layer is more than 5mm thick. Delineating the boundary between thick and thin layered targets at 5mm may seem somewhat arbitrary, but it is the separation that best fit the available data. This points to this boundary reflecting some property of the layers that maximises their effectiveness, such as the ratio of projectile diameter to ice layer thickness or the mass of the ice layer as a percentage of the total mass of the target.

From figure 5.4.2 it is apparent that the critical energy density Q^* changes dramatically depending on how thick the ice layer is. Targets with ice layers over 5mm thick appear to possess a shallow cratering regime that extends almost down to the disruption limit. From section 5.1 we have seen that layered targets tend to disrupt in such a way that the entire ice layer is stripped away before any significant disruption of the core occurs; therefore it follows that this initial shallow trend

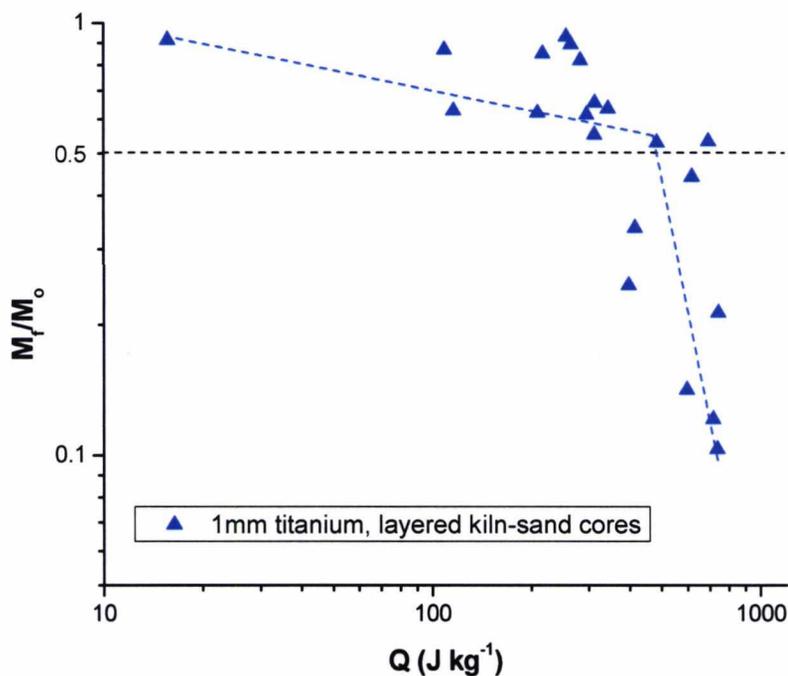


Figure 5.4.1 – Results of 1mm titanium impacts onto layered kiln-sand cores, with ice layer thickness anywhere between 1 – 7 mm.

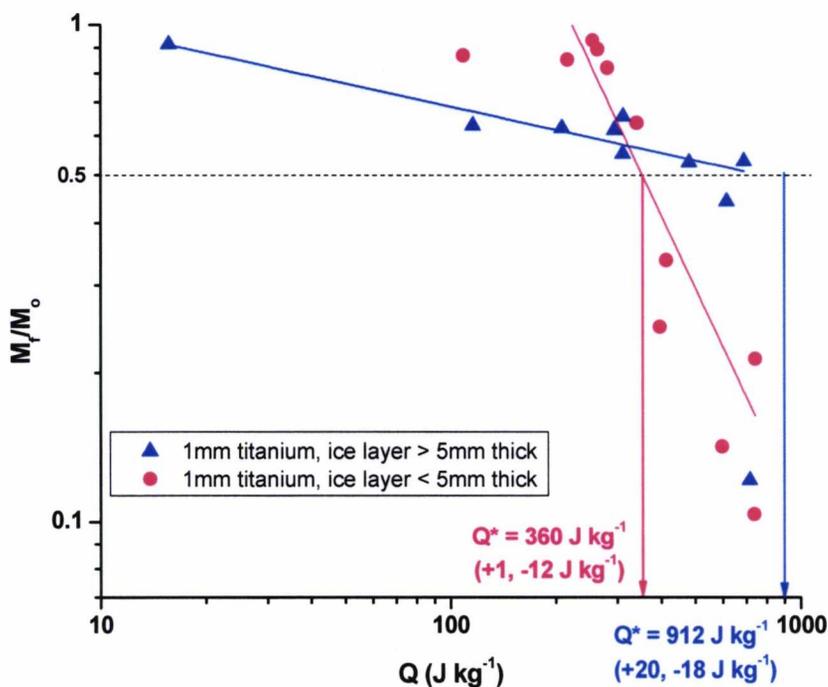


Figure 5.4.2 – Results for 1mm titanium impacts onto layered kiln-sand cores, separated into two categories based on ice layer thickness. In red are cores with thin (< 5 mm) ice shells ($Q^* = 360\ J\ kg^{-1}$), and in blue are cores with thick (>5 mm) ice shells ($Q^* = 912\ J\ kg^{-1}$).

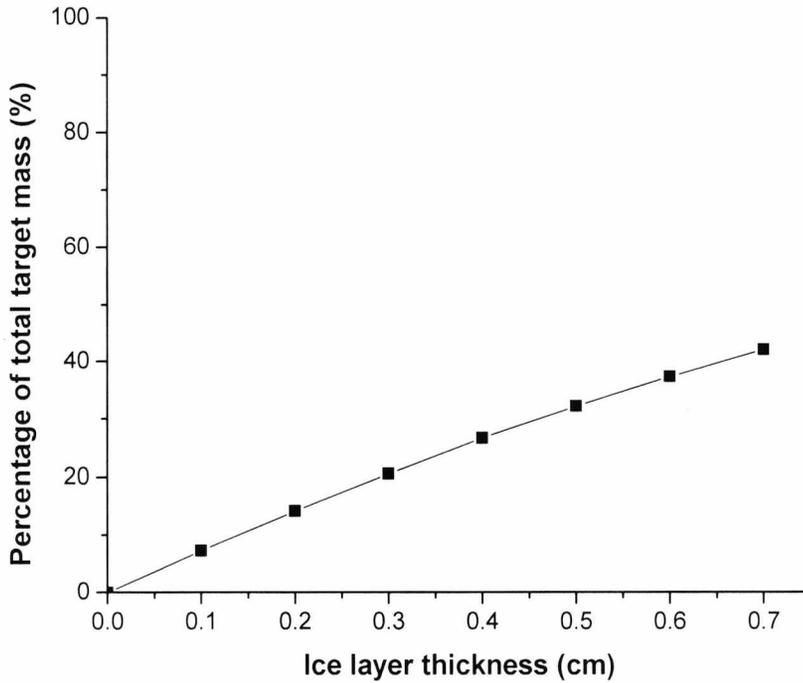


Figure 5.4.3 – Graph showing how the percentage of a layered target’s mass made up by the ice layer increases with the layer’s thickness, assuming even ice thickness over the entire surface area of the target and typical densities and masses for the ice layer and ice-silicate core respectively.

represents the ice shell being disrupted with very little energy transmitted to the denser core.

One reason for the difference between the two data sets in figure 5.4.2 lies in the rising proportion of the target’s mass provided by the ice layer as its thickness increases. In the “thick” size range of 5mm – 7mm, the ice layer accounts for anywhere from 30 – 40 % of the total target mass, as seen in figure 5.4.3, while the majority of ice layers in the “thin” size range were only 10-20% of the mass.

Since it should require far less energy to crater and disrupt the ice layer (based on the lower critical energy densities obtained for solid ice targets) than it does to reach and disrupt the kiln-sand core, this results in a shallower trend as the 40% of target mass provided by the ice layer is removed at lower energy densities while the inner core, where the majority of the target’s mass is locked up, remains largely untouched. As shown in section 5.1, once the ice layer has been entirely disrupted, increasing the

impact energy further causes the inner core to begin to disrupt in the manner of an unlayered core, which would produce a relatively steep trend at that particular range of energy densities. This is why it is expected that the shallow trend would give way to a much steeper trend at around the $M_f/M_o = 0.5$ point, represented by the dashed blue line. Unfortunately there is little data available for thick-shelled targets at high energy densities, so this remains to be demonstrated for sure.

Results for kiln-sand cores with thin layers are remarkably similar to the results for unlayered cores shown in figure 5.4.4. The 0.8mm stainless steel projectiles produce a Q^* value of 269 J kg^{-1} , only slightly less than the 279 J kg^{-1} produced by the same projectile type onto unlayered cores. In comparison, 1mm titanium projectiles produce a Q^* value of 360 J kg^{-1} as seen in figure 5.3.4; this is a larger difference from the unlayered result of 308 J kg^{-1} shown in figure 5.2.6 but not a significant one. Finally, 1.98mm glass projectiles produce a Q^* value of $133 (+110/-11) \text{ J kg}^{-1}$ in figure 5.3.4. While glass is a low density material at 2.68 g cm^3 , the near doubling of projectile size over the standard 0.8mm/1mm projectiles used in the majority of impacts is enough to account for the lower Q^* value. Unfortunately there are no glass impacts onto unlayered cores that can be used for comparison.

It is clearly the case that adding a thin ice layer to a sand-ice core makes no appreciable difference to the critical energy density needed to disrupt it. At the same time, a reasonably thick ice layer appears to confer a disproportionate increase in Q^* compared to the increase in the mass of ice present. Why does an ice layer with 30-40% of the overall mass of the target provide far better shielding for the core than an ice layer with 10-20% of the overall target mass? Possible reasons for this effect are discussed in Chapter 7.

5.5 – Varying target temperature and sand grain size

From figure 5.5.1 it is apparent that lowering the temperature of a kiln-sand core by 100°C produces very little difference in the critical energy density required to disrupt it. This is in stark contrast to previous experimental work that has shown low

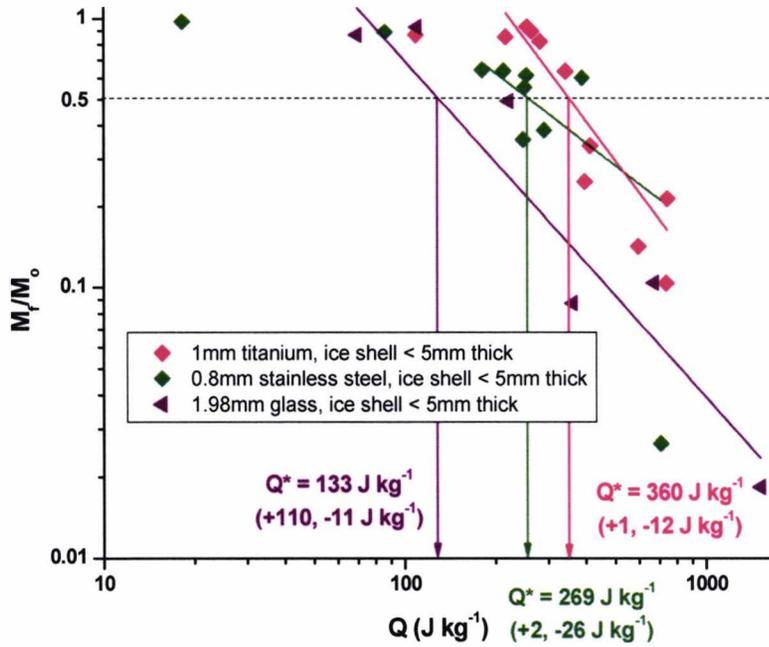


Figure 5.4.4 – Results for 0.8mm stainless steel, 1mm titanium and 1.98mm glass impacts onto thin-shelled cores.

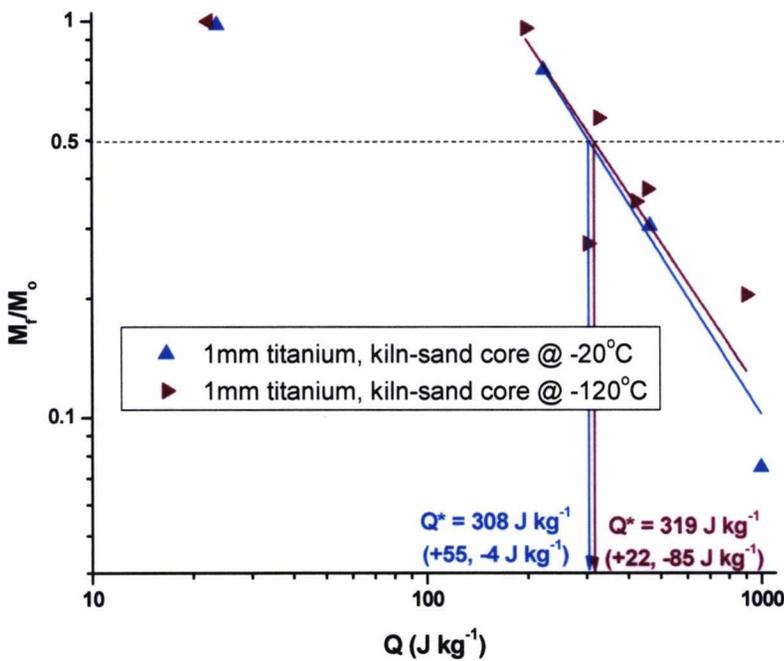


Figure 5.5.1 – Comparison of 1mm titanium impacts onto kiln-sand cores frozen at $-20^\circ C$ ($Q^* = 308\ J\ kg^{-1}$) and kiln-sand cores frozen at $-120^\circ C$ ($Q^* = 319\ J\ kg^{-1}$).

temperature ice cratering to be dependent on ice temperature (Grey & Burchell, 2003), where ice targets frozen at 100 K used here displayed much less depth penetration when impacted compared to targets frozen at 253 K; extrapolating from this the low temperature targets used here should be harder to penetrate and disrupt. However the stated temperatures are only the temperatures the targets were frozen at, and therefore the temperatures of the targets when they were removed from their respective freezers. The approximately 10-minute delay between sealing a target inside the light gas gun and the target being impacted will result in some amount of heating.

One attempt was made to measure the heating of a sand-ice core using a PT100 thermocouple with the sensor embedded into the core during target construction. Figure 5.5.2 shows the results of the thermocouple measurement. While the target was chilled to -120°C in the freezer, by the time the thermocouple sensor had been attached to a voltage meter and placed in the gun the target had already warmed to -85°C . Additionally the presence of the thermocouple cabling meant that the top enclosing part of the target chamber could not be used and therefore the measurements obtained will probably overestimate the amount of heating due to both the target not being enclosed in a chilled environment during the impact process and the extended amount of time the target was out of the freezer before being impacted.

Even despite these factors, however, the amount of heating the target undergoes is significant. There is a temperature increase of sixty degrees centigrade between the gun being sealed and the target being impacted; in this case the supposed low-temperature target was only slightly colder than a target frozen in a regular freezer at the time of impact. Again, this measurement does overstate the amount of heating involved, but it is reasonable to assume based on this that a target placed straight from the low-temperature freezer into the gun inside the target chamber would still undergo a significant amount of heating. Targets frozen at -20°C would also experience some amount of heating, but to a lesser degree since they are closer to thermal equilibrium with their surrounding environment. The conclusion to draw from this is that while an apparent reduction in target temperature appeared to make little or no difference to the critical disruption energy of a target at the moment of impact, the targets themselves

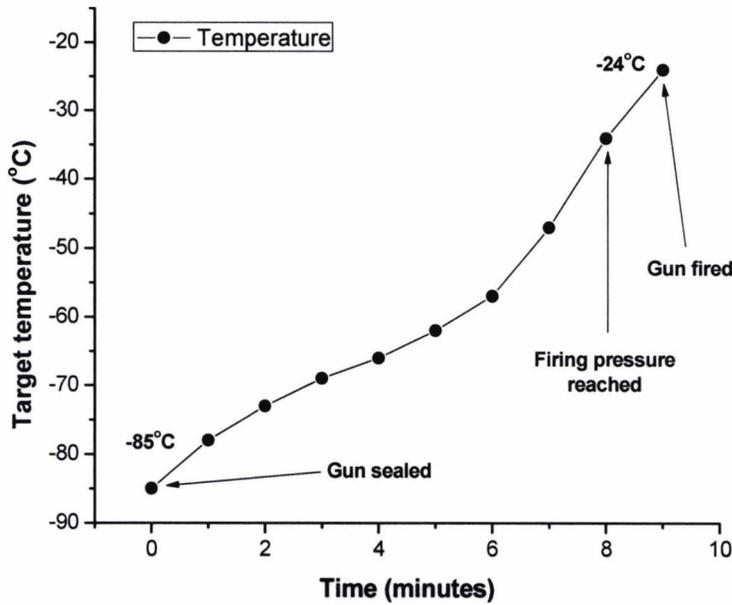


Figure 5.5.2 – Results of thermocouple measurement of target temperature during firing process showing the level of temperature increase between sealing the gun and firing the gun.

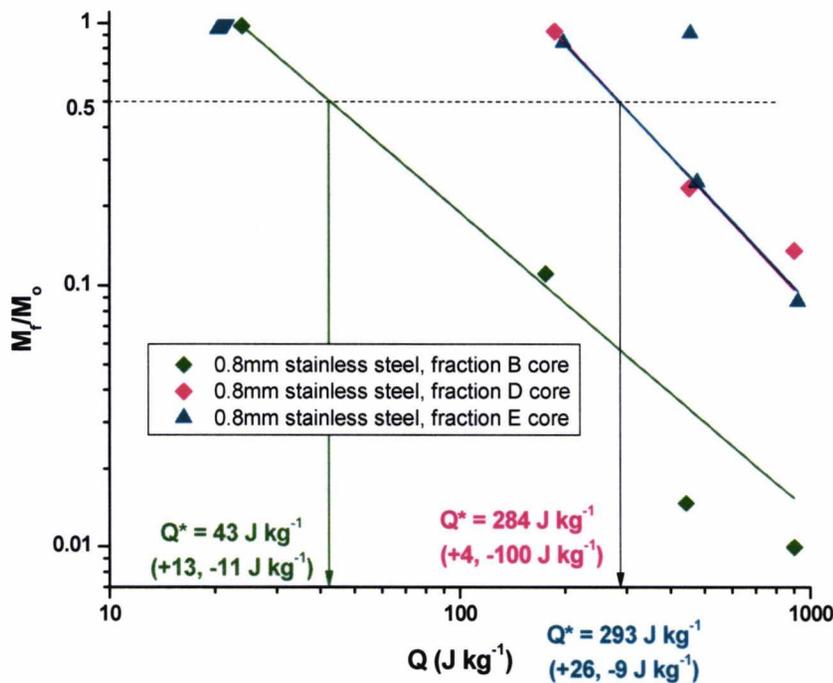


Figure 5.5.3 – Results for 0.8mm stainless steel impacts onto sand-ice cores constructed from fraction B ($Q^* = 43 \text{ J kg}^{-1}$), fraction D (284 J kg^{-1}) and fraction E (293 J kg^{-1}) sand grains.

may be much closer together in terms of temperature than was initially believed, and so the small shift in Q^* may reflect a real change in Q^* due to a colder target. It simply does not arise from a 100°C temperature change as originally thought.

Figure 5.5.3 shows the results of the impacts onto targets constructed from differently-sized sand grains. In Chapter 4 we saw that sand-ice samples composed fraction D and E grains had roughly similar compressive and tensile strengths while fraction B sand-ice samples were significantly weaker. These strengths are borne out by the impact results; fraction B sand-ice targets display a critical energy density of 43 (+13/-11) J kg⁻¹, the same as that measured for pure ice targets, while fraction D and fraction E targets have critical energy densities on par with regular kiln-sand targets (themselves roughly analogous to fraction C targets) at 284 and 293 J kg⁻¹.

5.6 - Comparison with previous work

Many of the Q^* values obtained here are in good agreement with the previous work researching Q^* values described in chapter 2. The Q^* for small ice spheres impacted by 1 mm aluminium projectiles of 45 J kg⁻¹ matches well with the Q^* value for ice cubes impacted by nylon projectiles obtained by Arakawa (1999) of 40 J kg⁻¹, while impacting a larger ice cylinder with denser copper projectiles produced a lower Q^* very similar to that of Murray (2004) – 8.1 J kg⁻¹ for the cylinders as opposed to 4.77 J kg⁻¹ for the large ice spheres used in Murray's work. This would appear to corroborate Q^* dependence on projectile density, with the lower density nylon and aluminium projectiles requiring much higher energy densities to disrupt ice targets than high-density copper.

Moving onto ice-silicate homogenous mixes, the nearest approximations in previous research are the 0% porosity ice/pyrophyllite and ice/serpentine mixtures described in Arakawa *et al.* (2002) and Arakawa *et al.* (2004). It should be noted that the method of target construction and the proportion of silicate content differed greatly: Arakawa *et al.* used targets that were mixed in a 50:50 ratio of ice and silicate and, rather than being composed of silicate material mixed with water and frozen, the 0% porosity targets are instead composed of powdered ice and powdered silicate compacted down

until all the pore space had been compressed out. This explains why the 0% porosity targets are somewhat less robust than the ice-silicate mixtures featured here with Q^* values of 100 J kg^{-1} for ice impacts onto ice/pyrophyllite and 121 J kg^{-1} for ice impacts onto ice/serpentine. This compares with values of 278 J kg^{-1} for impacts on small ice-silicate cores using stainless steel and 308 J kg^{-1} for impacts on the same targets using 1mm titanium. These Q^* values are at least on the same order of magnitude, indicating that Q^* for a homogenous ice-silicate mix will be on the order of hundreds of J kg^{-1} depending on the method used to construct the target and the type and quantity of silicate used in the mix.

Finally, the effect the addition of a layer has on impact strength mostly matches that noted in Okamoto *et al.* (2008). There are several key differences in target construction and material type (in particular Okamoto uses a denser material to cover a less-dense one, and layer mass goes up to 90% of the total target mass instead of 40% here) but the same increase in impact strength over that of the strongest material used in the target was observed at specific layer masses. The scale is different, however; here, the critical energy more than doubles from 360 J kg^{-1} for a plain homogenous ice-silicate core impacted by a titanium projectile to 912 J kg^{-1} for a core covered by an ice layer composing up to 40% of the mass of the target. Okamoto *et al.* observe the impact strength of their core mantle body rising above the impact strength of gypsum, the layer material, at 80-85% mass gypsum, but not to the extent that the impact strength is doubled. Possible reasons for this are discussed in Chapter 7.

5.7 - Conclusion

Analysis of the Q^* results for the various projectile/target combinations has yielded a lot of useful information about the factors affecting Q^* . A key discovery has been that targets do not undergo a steady disruption process as the impact energy is increased; instead there are obvious changes in the behaviour of the target under impact as the energy increases through the cratering, disruption and shattering regimes that are reflected in the slope of the disruption curve.

Also of note is that the layered targets used here do not disrupt as a single mass. Rather, the ice layer is disrupted first, and only when it has been mostly removed does

any significant disruption of the sand-ice core occur. This can lead to significant increases in Q^* above and beyond what would be expected depending on the overall thickness of the ice layer added.

Finally, shifts in Q^* due to changes in projectile size as well as target size are identified. Increases in projectile size result in lower Q^* values, while an increase in target size results in a higher Q^* value. Any calculation of Q^* must therefore be considered under the proviso that it only applies for that specific target/projectile combination.

Chapter 6 – Fragment Size Distributions

Definition of statistics: The science of producing unreliable facts from reliable figures.

-Evan Esar

6.0 – Introduction

While calculating the critical energy density from Q^* using the mass of the largest fragment is the main goal of the impact programme carried out during this work, additional information can be gleaned from looking at the other fragments formed as a result of the impact; how they are distributed in terms of mass is not apparent from a simple study of Q^* . For example, two impacts might produce similarly sized largest fragments but have wildly different fragment distributions, meaning that they are not being disrupted in similar ways. Fragment size distributions can therefore tell us important things about the character of a specific target type and how it is disrupted. This chapter attempts a detailed exploration of the size distribution data produced during the experimental programme.

6.1 – Plots of size distribution per target/projectile type

It is important to recognise that studying size distributions was something of a secondary goal during the experimental programme. All prospective energy densities throughout it were pitched at values designed to investigate Q^* , not size distribution. This makes a proper study of size distributions difficult; several size distributions have two or three sets of data while others that would be useful for comparison have none. The paucity of good data is exacerbated by the fact that, prior to the introduction of the improved target chamber described in Chapter 3, measuring the masses of all but the four or five largest fragments was impossible due to melting, and this is nowhere near enough data for a detailed picture of the size distribution to be built up. The improved target chamber was only used for the last fifty shots out of a hundred and fifty total, meaning that detailed data only exists for a third of the shots carried out during the experimental programme.

Table 6.1.1 – Categories of data for plots of size distribution with varying energy density.

Projectile type	Target type
1mm titanium	Sand-ice core
	Low temperature sand-ice core
1mm aluminum	Solid ice core
0.8mm stainless steel	Sand-ice core
	Fraction B core
	Fraction D core
	Fraction E core

The following discussion of the size distributions that were obtained is a thus little less comprehensive than that in the previous chapter due to gaps in the data. However, there is enough to allow us to draw some useful conclusions.

The first step towards analysing the size distributions is to separate them into data sets with common properties. Since it is not known which factors might affect size distribution – projectile type, target material and so on – it is a good idea to keep these constant while looking at the effects of a factor that *is* known to affect size distributions: projectile impact energy. Once the effect of the impact energy has been isolated and understood, it will be possible to move on to comparing impacts on different target types that take place at different impact energies.

If the available data is divided off into projectile/target type combinations, it produces the categories shown in table 6.1.1.

Since the experimental programme was designed to investigate Q^* , different target/projectile combinations in general undergo at least four impacts. These impacts are pitched at a range of energy densities designed to bracket the suspected Q^* value for that target type, with two on either side of the Q^* line. This means that, typically, impact energy densities fall into one of the following four categories:

- *Low energy cratering impacts* ($Q \ll Q^*$) remove a very small fraction of the original mass of the target, leaving the target almost entirely intact.

- *High energy cratering impacts* ($Q < Q^*$) remove a much larger fraction of the original mass of the target, leaving a single very large fragment plus several smaller but still sizeable fragments.
- *Disruption impacts* ($Q = Q^*$, $Q > Q^*$) smash the target into several large fragments and many small fragments.
- *Shattering impacts* ($Q \gg Q^*$) vaporise much of the original target mass leaving a large number of small fragments behind but almost no large fragments.

For each target type there is at least one example of each of these impact outcomes. In this analysis the fragments are sorted into descending order of mass and then plotted as relative mass (M_{fn}/M_o) against the cumulative fragment number (with the heaviest fragment being number 1, the next heaviest being number 2 and so on).

The error in relative mass $\Delta M_{fn}/M_o$ is calculated by the equation

$$\left(\frac{\Delta A}{A}\right)^2 + \left(\frac{\Delta B}{B}\right)^2 = \left(\frac{\Delta C}{C}\right)^2 \quad \text{(Eq. 6.1)}$$

where A is M_o , B is M_{fn} , C is M_{fn}/M_o , ΔA and ΔB are both 0.05g (the absolute error in any measurement taken with the set of scales used), and ΔC is $\Delta M_{fn}/M_o$, the desired quantity. This equation can therefore be rearranged to the form

$$\left(\frac{0.05}{M_o}\right)^2 + \left(\frac{0.05}{M_o}\right)^2 \left(\frac{M_{fn}}{M_o}\right)^2 = \left(\Delta \frac{M_{fn}}{M_o}\right)^2 \quad \text{(Eq. 6.2)}$$

to calculate the absolute error (where M_{fn} and M_o are measured in grams).

For one size distribution plot, M_o is constant. Furthermore, since the targets used were almost all cores constructed of sand-ice, M_o was nearly always $60\text{g} \pm 2\text{g}$. Since M_o is constant for one specific size distribution plot and practically constant between

different size distribution plots, the first term and the first bracketed part of the second term in equation 6.1.2 are dependent solely on M_o and can therefore be treated as constants. As a result, the error in relative mass is almost entirely dependent on the second bracketed part of the second term $(M_{fn}/M_o)^2$; if M_f is small in comparison to M_o then this term is negligible in comparison to the first $(0.05/M_o)^2$ and the error is a constant value.

In practice, only the largest fragments possess a large enough mass to change the error from this constant even by a small amount, and these fragments are so large that the error is negligible in comparison to them anyway. Errors for fragments below the mass limit $M_{fn}/M_o = 0.1$ (where the vast majority of fragments in the size distribution plots lie) vary by less than 0.5% from $\Delta(M_{fn}/M_o) = 0.00085$, so the error in relative mass is treated as having this constant value unless otherwise stated. For this reason, error bars are only included on the first graph for illustration purposes. Subsequent graphs have no error bars since they obscure the data points at low fragment masses.

Fragments were inspected visually after impact and weighed in apparent order of descending mass. They were handled using insulated gloves in order to reduce any melting from body heat. While the fragments remained in the target chamber they remained reasonably intact since the target chamber surfaces were still very cold; however upon being transferred to the unfrozen metal surface of the scales they underwent immediate melting. This was most pronounced when dealing with the smallest fragments which melted almost completely during the mass measurement. The scales were wiped clean after every measurement to ensure that the weighing of subsequent fragments was not distorted by leftover melted mass from previous fragments.

The laboratory equipment used to measure fragment mass had a sensitivity limit of 0.1g; this was the smallest mass it was capable of measuring. This means that the measurement of fragments with masses approaching 0.1g become increasingly imprecise, as reflected by the increasing error (see Figure 6.1.2). In addition, while this sets the mass of the smallest fragments plotted on the graphs below at different points depending on the mass of the original target, the smallest fragments measured

are by no means the smallest fragments produced by an impact and subsequently collected. There were many fragments which, when measurement was attempted, did not register on the scales used (indicating masses of less than 0.05g).

In theory the number of fragments in each mass bin may increase further as the fragment mass decreases, continuing the upwards trend of the plot (as read from right to left) until finally reaching the mass of individual sand grains. An approximation of this mass can be made by assuming the sand grain is a sphere and then calculating it using the typical diameter and density of a grain. This yields an approximate mass of 3×10^{-4} g, which combined with the typical core mass of 60g gives the relative mass of a sand grain as 5×10^{-6} . At this point it is expected that there would be a huge spike in the number of fragments recovered due to the large number of sand grains produced after an impact, after which the number of fragments recovered drops off to nothing since the mass of a sand grain is the smallest mass unit making up a sand-ice

target; producing smaller masses would require an impact to physically break up the sand grains themselves. This is possible, but given the makeup of the target (sand embedded into a matrix of ice) it would require highly energetic impacts to occur on anything but the smallest of scales. While it would be more than a little impractical to weigh or even count every individual sand grain produced after an impact, a theoretical fragment size distribution where every single fragment has been recovered and measured is shown in figure 6.1.1.

6.1.1 – Comparison of results

Size distributions are generated by arranging the recovered fragments from an impact in order of size, with the largest fragment labelled as number 1, second largest fragment as number 2 and so on. The fragment masses M_{fn} are normalised to the mass of the original target M_0 by dividing one by the other in a similar manner to M_f/M_0 ; these M_{fn}/M_0 values are then plotted against the size-order number in order to get some idea of the overall distribution of the fragment masses. Only solid fragments are measured; while it would have been useful to measure all the remaining target mass present in the gun after impact in order to make an estimate of the total mass loss due

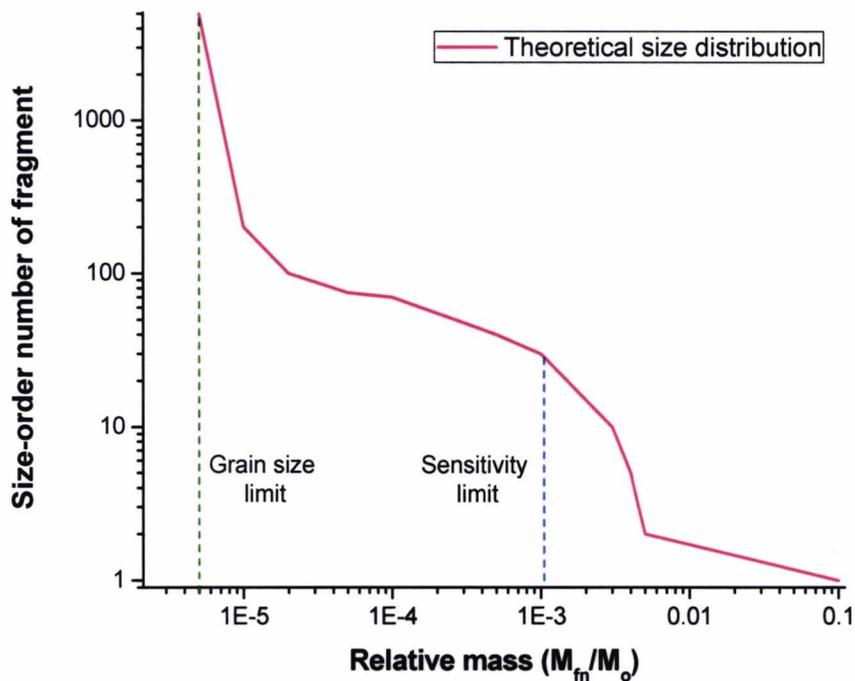


Figure 6.1.1 – Graph showing the theoretical fragment size distribution were the measurements not limited by the sensitivity of the equipment.

to vapourisation, recovering all remaining target mass would have required sweeping out the gun and the remaining mass would therefore become contaminated with byproducts from the firing process that were not part of the original target.

Firstly, some comparisons can be drawn between the results of the impacts of 1mm titanium onto kiln sand and low temperature kiln sand core plots (Figures 6.1.3 and 6.1.4); using the same projectile type means that the only varying factors between these two data sets are the temperature of the target and the range of energy densities used in impacts. Two things immediately stand out: one is that all the size distributions for the low temperature target are following roughly the same trend, merely being shifted along the x-axis depending on how massive the largest fragment is. The other is that compared to the normal temperature cores the low temperature size distributions are reversed with respect to energy density (from right to left they are red blue green as opposed to green blue red for normal cores). Unlike the normal cores low temperature targets consistently produce more and heavier fragments after

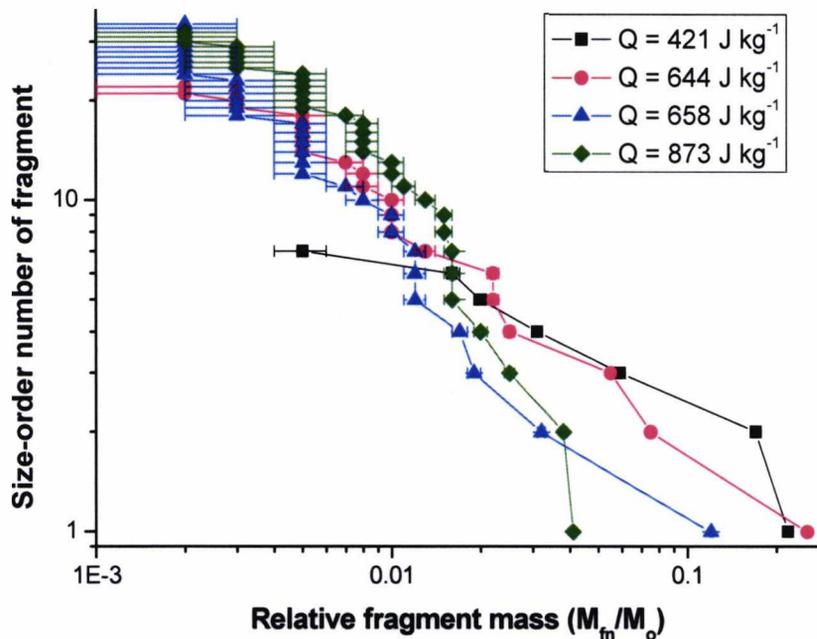


Figure 6.1.2 – Fragment size distributions for 0.8mm stainless steel shots onto kiln sand cores.

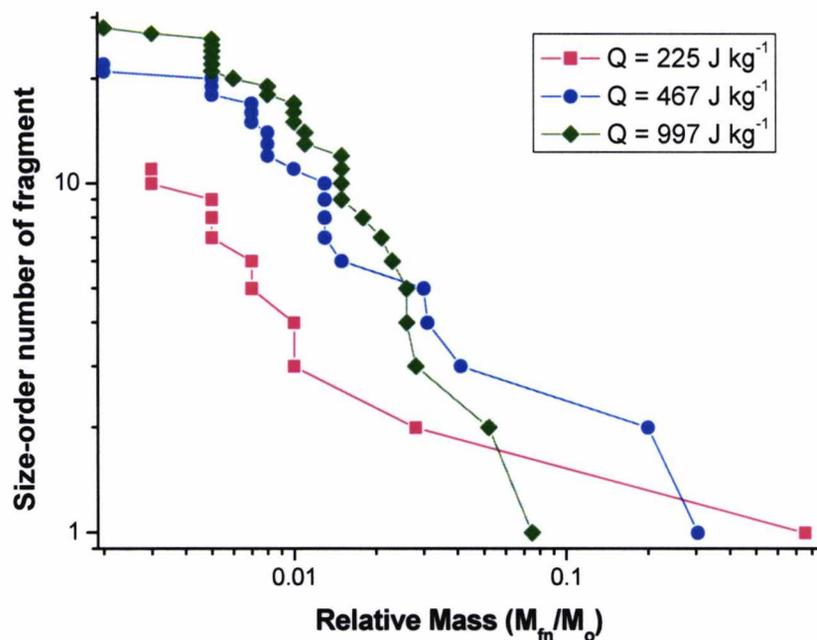


Figure 6.1.3 – Fragment size distributions for 1mm titanium shots onto kiln sand-ice cores.

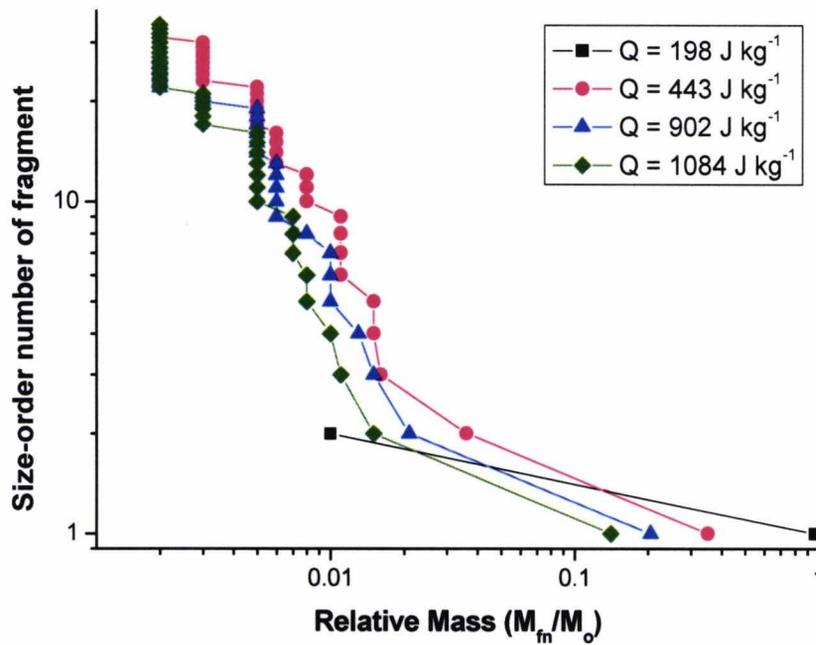


Figure 6.1.4 – Fragment size distributions for 1mm titanium shots onto low temperature kiln sand cores.

disruption shots ($Q > Q^*$), even down to very small fragment sizes. Typically a shattering shot ($Q \gg Q^*$) produces less large fragments and more small fragments than a disruption shot as seen in figure 6.1.3. The increased number of fragments is partially a direct result of the low temperature, but not through any change in the shattering mechanism in the target. Rather, lower temperature fragments are more likely to survive the time period separating the shot and fragment measurement without melting, resulting in more of them overall.

Next I compare impacts onto cores composed of differently-sized sand grains. The impacts on cores composed of fraction D and fraction E sand grains (figures 6.1.6 and 6.1.7) produce almost identical size distribution plots at similar energy densities (low red plot at $\sim 180 \text{ J kg}^{-1}$ and overlapping blue and green plots at $\sim 450 \text{ J kg}^{-1}$ and 900 J kg^{-1}), suggesting that these two target types are being disrupted in the same way. Fraction D and fraction E sand grains have both very similar grain sizes ($229 \mu\text{m}$ and $160 \mu\text{m}$ respectively) relative to the size of the projectile, so this may be a reason for the similar size distributions.

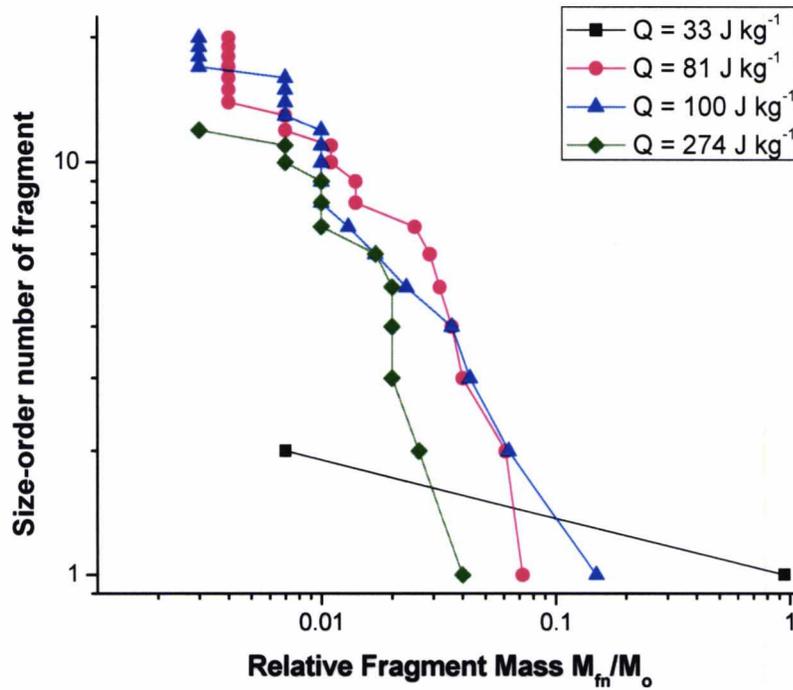


Figure 6.1.5 – Fragment size distributions for 1mm aluminium shots on to solid ice cores.

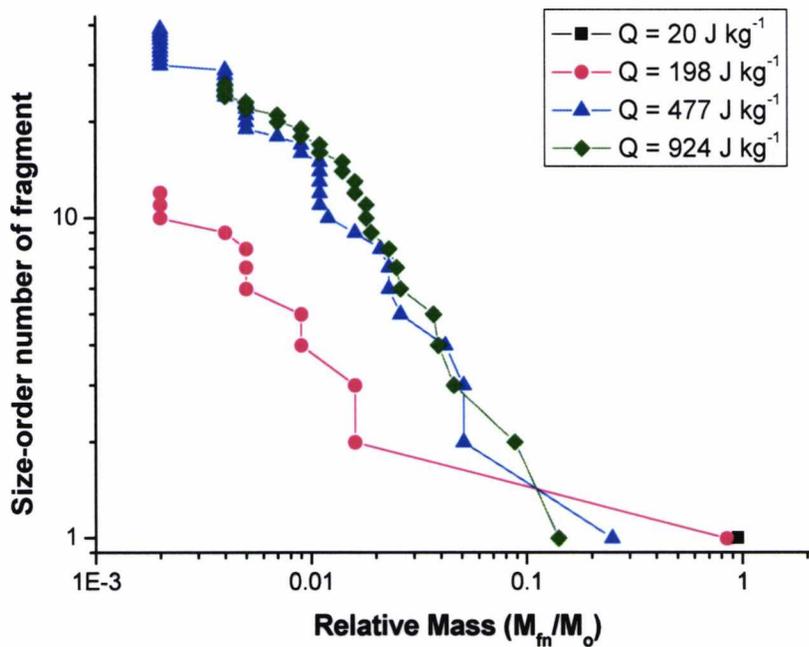


Figure 6.1.6 – Fragment size distributions for 0.8mm stainless steel shots onto fraction E cores (average grain diameter $160 \mu\text{m}$).

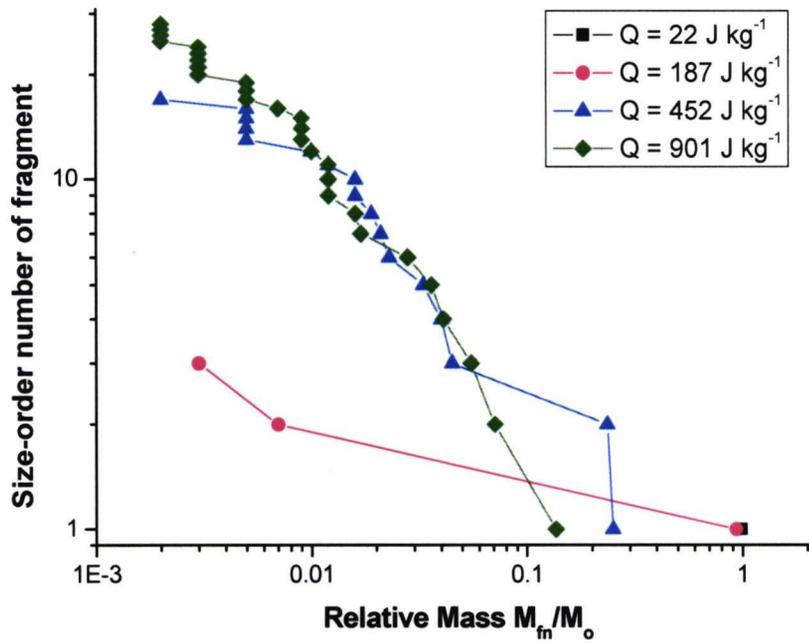


Figure 6.1.7 – Fragment size distributions for 0.8mm stainless steel shots onto fraction D cores (average grain diameter 229 μm).

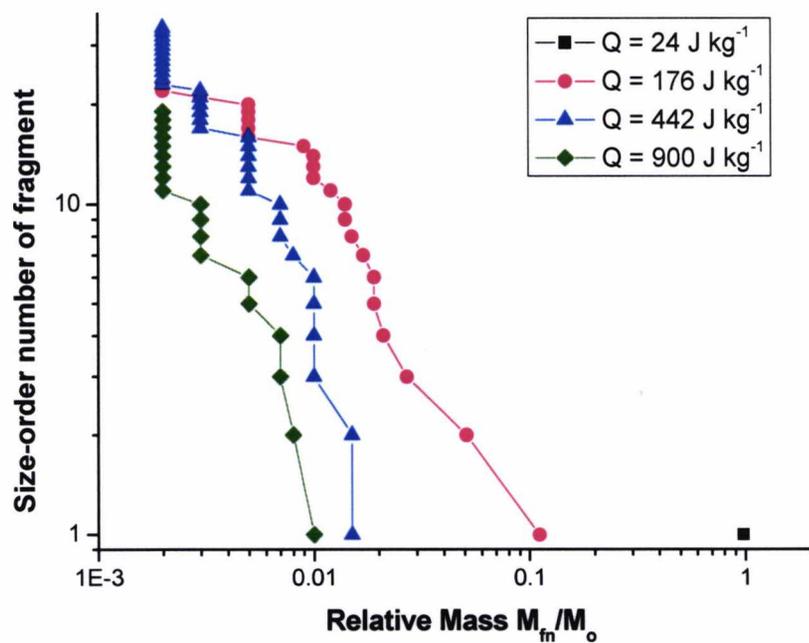


Figure 6.1.8 – Fragment size distributions for 0.8mm stainless steel shots onto fraction B cores (average grain diameter 1043 μm).

The impacts on fraction B cores (figure 6.1.8) produce markedly different plots to those on fraction D and E, which again is to be expected – Q^* for fraction B samples is much lower than that for fraction E/D (90 J kg^{-1}) and the impacts are taking place at the same energy densities as those on fraction E/D samples, so the impact outcomes are going to be shifted towards the shattering end of the spectrum. This is an example of the difficulty of comparing even impacts with identical projectile type and energy density if they use different targets; the difference in Q^* makes a direct comparison impossible.

However, by looking at these size distributions some general observations about their properties relative to Q^* can be made. Firstly, cratering impacts ($Q < Q^*$, black square symbols in figures 6.1.3 – 6.1.8) produce too few fragments to allow detailed study of the size distribution; due to the extremely low mass loss during cratering shots they invariably result in an almost-intact target with only two or three measurable fragments besides that. Cratering impacts must therefore be removed from any study of fragment size distributions.

Near-disruption and disruption impacts ($Q = Q^*$, red circle symbols in figures 6.1.2 to 6.1.8) are more useful. These produce size distributions where the first fragment has a relatively high M_f/M_0 value (between 0.2 – 0.5). Since most of the surviving mass is locked up in this first fragment, there are less surviving fragments overall and they tend to be much smaller than the first fragment. This produces a very flat and shallow trendline between the relatively massive first fragment and the small second fragment. Past this point the fragments form what could be referred to as a “typical” size distribution; relative mass steadily decreases as the cumulative number of fragments increases, with the trendline starting to level out as the fragment size approaches the sensitivity limit of the laboratory instruments used to measure mass.

As Q is increased ($Q > Q^*$, blue triangle symbols in figure 6.1.2 – 6.1.8) the observed size distribution increasingly resembles the distribution commonly seen in the literature for disrupted bodies (e.g. Arakawa 1999, Ryan *et al.* 1999); the largest surviving fragment decreases as Q increases smoothing out the initial shallow kink in the trendline. Smaller largest fragments mean there is more mass available to form

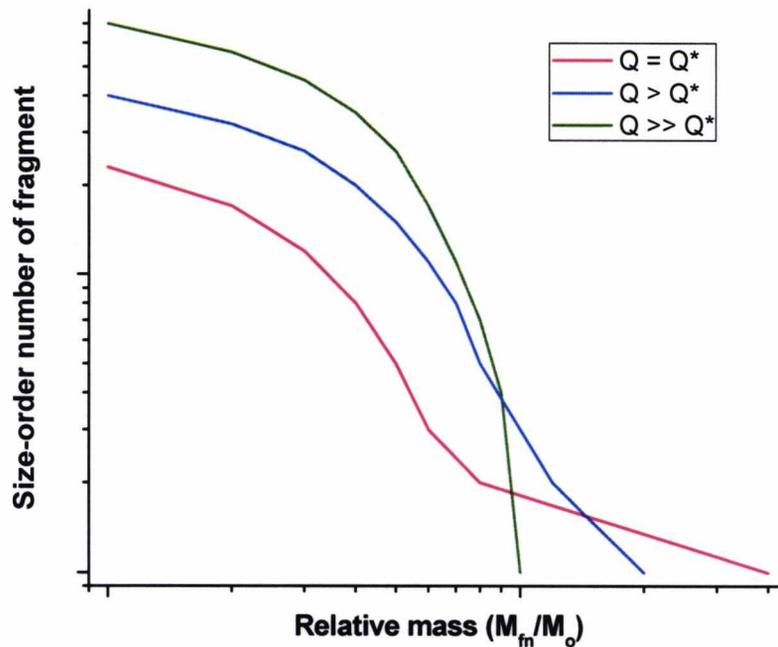


Figure 6.1.9 – Graph showing how the typical fragment size distribution varies as Q scales with Q^* .

smaller fragments. This produces two effects; one is that the second, third etc. largest surviving fragments are larger in comparison to the largest fragment rather than being a fraction of the size. The other is that there are a lot more surviving fragments overall.

Finally for shattering impacts at very large values of Q ($Q \gg Q^*$, green diamond symbols in figures 6.1.2 – 6.1.8) the size distribution almost entirely follows that of the typical distribution. The largest fragment is only fractionally larger than the second and third fragments eliminating the initial kink entirely and a high number of tiny fragments are produced.

These general trends are displayed in an example graph in figure 6.1.9. Three different trendlines are plotted showing the observed size distributions for the impacts $Q = Q^*$, $Q > Q^*$, and $Q \gg Q^*$.

It is interesting to note that these typical trends are similar to both the observed size distributions of asteroid families (Tanga *et al.* 1999, Leliwa-Kopystynski *et al.* 2009) and models that simulate how these families might have formed after impacts of varying energy (Durda *et al.* 2007). Asteroid families tend to have a size distribution which follows either a “convex” trend corresponding to that observed here for impacts with $Q = Q^*$ (red plot in figure 6.1.9) or a “concave” trend that matches impacts with $Q \gg Q^*$ (green plot in figure 6.1.9). The models in Tanga *et al.* support the reliance of the overall size distribution on the size of the largest fragment, while Durda *et al.* show a similar reliance of the size distribution on impact energy; cratering impacts produce convex size distributions, while a supercatastrophic impact (highly energetic and broadly similar to the shattering $Q \gg Q^*$ impacts carried out here) produces a concave one. Thus these observed trends seem to broadly reflect what happens to asteroids when they undergo impact.

The next step is to begin drawing comparisons between size distributions for the different projectile/target data sets. However, this poses a problem; as described above, the fragment size distribution follows a trend that depends on the relation of the Q the target was impacted with the Q^* for that particular target type. Since Q^* changes from target type to target type, a comparison between a solid ice core impacted at a Q of 400 J kg^{-1} and a kiln sand-ice core impacted at a Q of 400 J kg^{-1} will produce wildly different size distributions simply through dint of this Q being $\gg Q^*$ for the ice target, and merely $> Q^*$ for the kiln sand-ice target.

One solution to this problem is not to make straight comparisons between different Q , but instead to compensate for the variance in Q^* for each target/projectile combination by making a comparison between Q/Q^* ratios for each type of target – a solid ice target impacted with twice its disruption energy might be expected to display an approximately similar fragment size distribution to a kiln sand-ice target impacted with twice its disruption energy, since the size distributions should in theory both follow the general $Q > Q^*$ trend. If they are dissimilar, then using ratios goes some way to eliminating the Q^* variance and allowing analysis of the other factors affecting the size distribution such as difference in target material.

6.2 – Plots of size distribution per Q/Q* ratio

Exactly which data plots fall into a specific Q/Q* range was not the result of a preplanned programme of work; since Q* for a target type was not known before the shots were carried out, it is impossible to tailor shots to produce one particular Q/Q* ratio. Now that the shots have been performed and the Q* values *are* known it may be useful for future work to investigate the Q/Q* relationship more thoroughly; nevertheless we can still draw some initial conclusions from the data gathered here.

It is not really meaningful to compare just two shots with similar Q/Q* ratios since this will only tell us how those two shots compared to each other rather than allowing the identification of any overall trend. While several of the shots with fragment

distribution data shared a Q/Q* ratio with another shot, none of them shared Q/Q* ratios with multiple shots.

The best that can be done is to see which shots fall into a range of Q/Q* that is small enough to allow a reasonable comparison to be made between them, but large enough to contain several shots. Three Q/Q* ranges were selected, each of which contained four shots. These Q/Q* ranges are 1.5 – 1.9, 2.3 – 2.6 and 3.0 – 3.5.

Excluding the first two data points from each plot (which are the most variable from shot to shot), the plots for impacts on kiln sand, fraction D and fraction E cores within the Q/Q* range of 1.5-1.9 (figure 6.2.1) all follow the same trend fairly closely; there is some divergence but this is to be expected through natural scatter. The plot for solid ice has the same basic shape as the other plots, but there are fewer fragments and it is shifted further to the left; this shows that, generally, there are fewer fragments recovered and those that are are much less massive than those recovered from the sand-ice targets. This is almost certainly caused by the difference in target material, but the exact mechanism is unknown. The effects of melting can be excluded for the fragment sizes seen here; as previously stated mass loss due to melting was kept to a

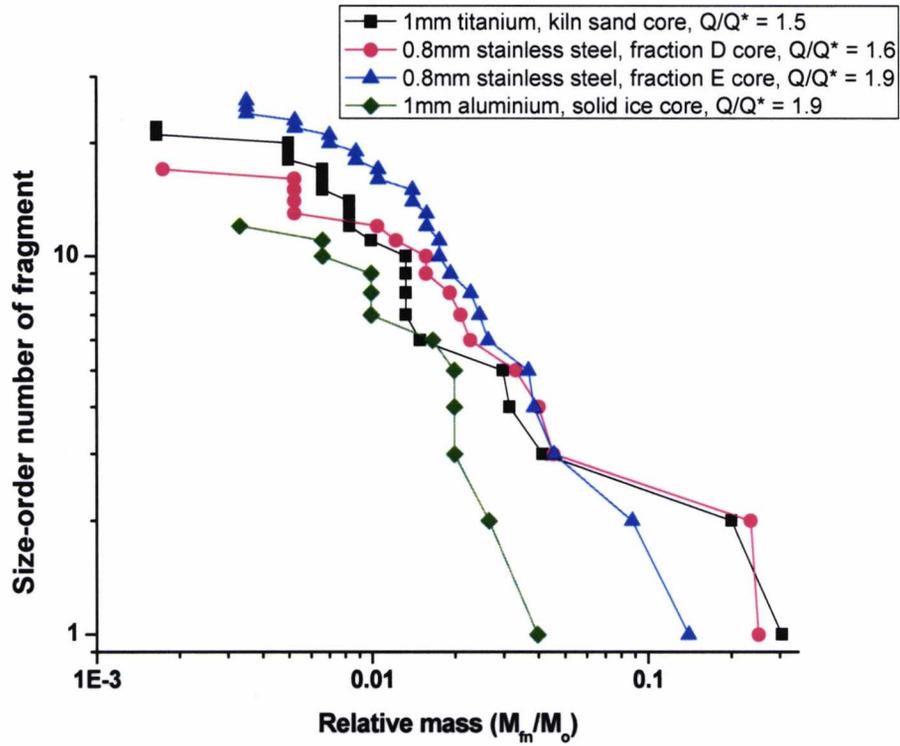


Figure 6.2.1 – Fragment distribution plots of various shots falling into the Q/Q^* range of 1.5 – 1.9.

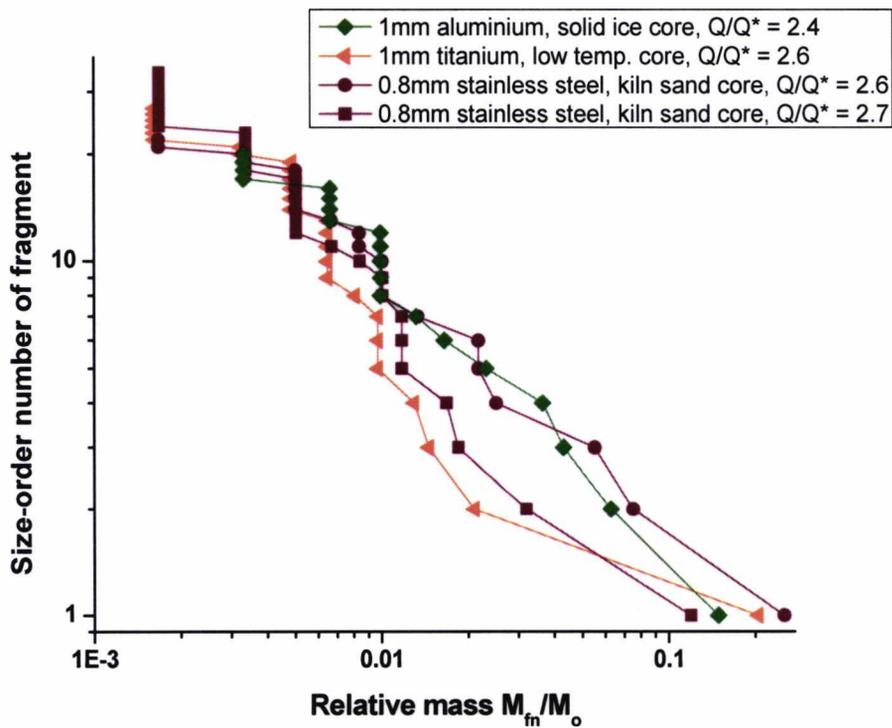


Figure 6.2.2 – Graph of various fragment distribution plots which fall into the Q/Q^* range of 2.4 - 2.7.

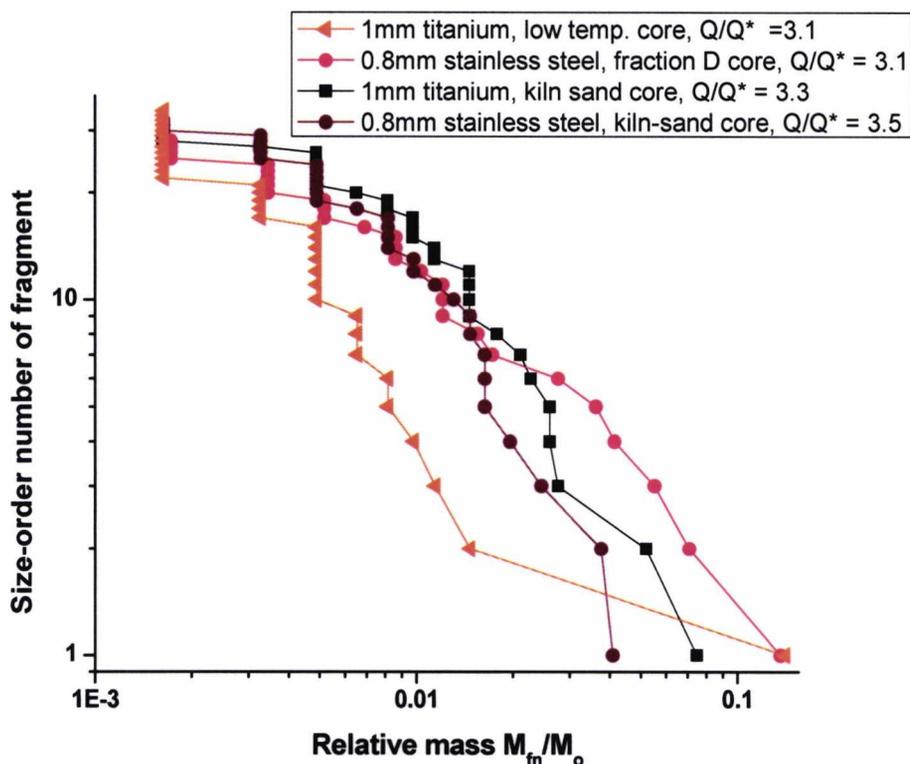


Figure 6.2.3 – Graph of various fragment distribution plots falling into the Q/Q^* range of 3.1 – 3.5.

minimum until the fragments were weighed – thus, the effect seen here is real and not a systematic flaw in the measuring process. It is possible that pure ice vapourises more easily under impact than sand-ice, which would account for the decrease in total mass recovered; however this is conjecture and further work in this area is required to investigate the causes behind the smaller fragments.

The plots for $Q/Q^* = 2.4 - 2.7$ displayed in figure 6.2.2 again have the same general shape. However, this shape is different from the one shown in figure 6.2.1; it is a shallower, steadier decrease in fragment mass with fragment number. All the plots here tell us something of interest; the two stainless steel impacts used the same projectile on the same target type at almost exactly the same Q/Q^* ratio, yet there is a moderate amount of divergence in the resulting size distribution plots. This is a good indication of how much natural scatter is present in the data. Next we note that the low temperature core shot produces a consistently shallower size distribution than any of the other shots. Since the increased Q^* value of the low temperature cores should

have been compensated for by using a Q/Q^* comparison, a putative reason for this may be that low temperature cores are somehow inherently stronger than regular cores – that is, instead of requiring a simple increase in energy to disrupt them after which they fracture in much the same way as the normal cores, they have a wholly different fracturing mechanism resulting in the shallower plot. While ice strength does increase at lower temperatures (Hobbs 1974), Chapter 5 showed that the low temperature targets used here were not so far away from the normal targets in terms of temperature as previously thought; however there may still be a large enough difference to alter the plot.

Finally the solid ice core size distribution in figure 6.2.2 – which was somewhat anomalous on the $Q/Q^* = 1.5 - 1.9$ plot – is here almost exactly in line with the plot for stainless steel onto kiln sand cores.

Figure 6.2.3 shows that at relatively high energy densities of greater than three times the critical energy density, there is still a rough similarity between the various types of sand-ice targets regardless of projectile. The exception, again, is the low temperature core; this has a size distribution with the largest fragment – albeit only just – for this Q/Q^* range, but the remainder of the fragments are smaller than those recovered for the other sand-ice targets. This is consistent with the plot for low temperature cores in the $Q/Q^* = 2.3 - 2.6$ range, except much more pronounced; it confirms that low temperature cores do not fragment in the same way as regular sand-ice cores.

From these three Q^* -independent plots, the following conclusions can be drawn.

- Small changes in target material type do not affect the fragment size distribution. Fraction D, fraction E and kiln sand cores do not produce markedly different size distributions at similar Q/Q^* values despite the difference in the size of the various types of sand grains used to make these cores being sizable. However, no comparisons between these sand types and size distributions for fraction B cores could be made, so it can only be said that this holds true for projectiles made of sand grains with diameters smaller than the size of the projectile used to impact them

- The projectile material does not affect the fragment size distribution; 1mm titanium and 0.8mm stainless steel projectiles consistently produce overlapping size distributions, and in one case a 1mm aluminium projectile also produced an overlapping size distribution.
- Dramatic changes in target material type *do* affect the fragment size distribution, at least at low Q/Q^* ratios. A solid ice target impacted at a Q/Q^* ratio of 1.9 produced a smaller size distribution when compared to sand-ice targets impacted at similar ratios. The mechanism behind this is unknown.
- Changing the target material temperature also affects the fragment size distribution, as borne out by the two differing plots for low temperature cores. While this is to be expected when dealing with icy targets, the mechanism behind this is also unknown.

One additional observation that can be made is that the general trend of each of the three graphs follows one of the typical plots shown in figure 6.1.9. $Q/Q^* = 1.5 - 1.9$ follows the $Q \approx Q^*$ plot, $Q/Q^* = 2.3-2.6$ follows the $Q > Q^*$ plot, and $Q/Q^* = 3.1 - 3.5$ follows the $Q \gg Q^*$ plot.

6.3 – Power law fits.

It is useful to quantify how similar/dissimilar these fragment size distributions are by fitting them with a power law. However, this is not quite as simple as it first appears; due to the nature of the size distributions and the sensitivity limit of the instruments used to measure fragment masses, certain parts of the distributions will distort the overall trend.

Firstly, fragment size distributions that are the result of shots with a low Q/Q^* ratio will have comparatively heavy largest fragments, so much so that they often outweigh the rest of the fragments combined. This produces the “kink” described for $Q \approx Q^*$ distributions in section 6.1 which heavily distorts the first part of the trend. Since the

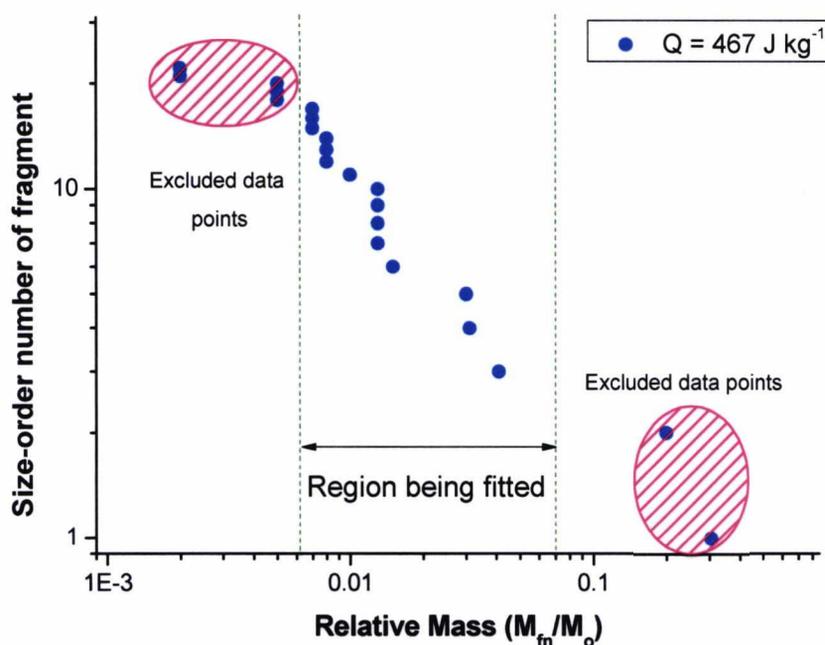


Figure 6.3.1 – Example graph illustrating which data points are excluded when fitting a power law to a size distribution.

largest fragments are not characteristic of the rest of the distribution this effect must be compensated for by removing them from the fit.

Conversely, as the distribution approaches the size of the smallest measurable fragments the trend becomes distorted again, this time through measurement bias. When the fragments approach this scale two effects start to dominate: fewer fragments overall are measured due to melting destroying them and/or the difficulty of selecting, by eye, a fragment that weighs enough to register on the scales used, while the precision of weighing the fragments that *are* measured becomes increasingly uncertain as the distribution moves closer to the sensitivity limit. These small fragments too need to be isolated and removed from any fit made.

The data is therefore made suitable for a power law fits in the following way: the two largest fragments from a size distribution are removed from the plot, as are all fragments with a relative mass of 0.005 or below (although the 0.005 points are left in if this method removes so many data points from a size distribution that it is no longer

Table 6.3.1 – Results from power law fits of fragment size distributions. Fits are of the form $y = ax^b$, where R is the regression coefficient.

Projectile	Target	Q	Q*	Q/Q*	a	b	± b	R ²
1mm titanium	Kiln sand core	224	306.4	0.7	0.063	-0.9	0.29	0.79
		465		1.5	0.158	-0.92	0.07	0.95
		997		3.3	0.061	-1.2	0.07	0.92
0.8mm stainless steel	Kiln sand core	644	247.3	2.6	0.260	-0.78	0.06	0.95
		658		2.7	0.092	-1.0	0.10	0.91
		873		3.5	0.046	-1.21	0.09	0.94
1 mm titanium	Low temperature kiln sand core	423	353.3	1.2	0.032	-1.2	0.11	0.92
		902		2.6	0.018	-1.3	0.18	0.87
		1084		3.1	0.000	-2.2	0.41	0.87
0.8mm stainless steel	Fraction B core	176	30.9	5.7	0.034	-1.3	0.09	0.96
		443		14.3	0.008	-1.4	0.17	0.88
		900		29.1	0.001	-1.6	0.44	0.80
0.8mm stainless steel	Fraction D core	452	288.7	1.6	0.318	-0.8	0.07	0.95
		901		3.1	0.280	-0.82	0.06	0.96
0.8mm stainless steel	Fraction E core	198	486	0.4	0.139	-0.7	0.18	0.82
		477		1.0	0.219	-0.9	0.08	0.94
		924		1.9	0.173	-1.0	0.05	0.94
1mm aluminium	Solid ice core	81	42	1.9	0.626	-0.61	0.06	0.93
		100		2.4	0.130	-0.9	0.11	0.92
		274		6.5	0.163	-0.8	0.12	0.89

practical to fit a power law to it). Figure 6.3.1 illustrates which data points are typically excluded. This process leaves a single more-or-less consistent data set to which a reasonable fit can be made.

Following are the power law fits for various data sets.

One interesting thing to note from figure 6.3.4 is that despite the general order of the plots being reversed compared to that seen in Figure 6.3.3 for low temperature cores (red blue green instead of green blue red) the fits of these size distributions follow the same trend as those for regular titanium cores (and other targets) – red and blue are shallower and fairly similar, whereas green is much steeper.

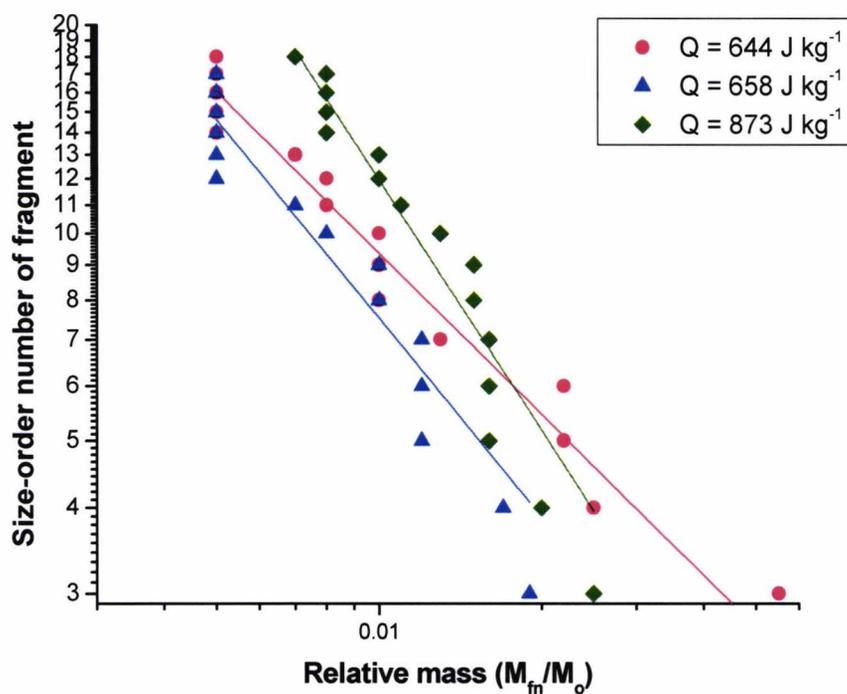


Figure 6.3.2 – Power law fits for impacts by a 0.8mm stainless steel projectile onto kiln sand cores. Fit coefficients are shown in Table 6.3.1.

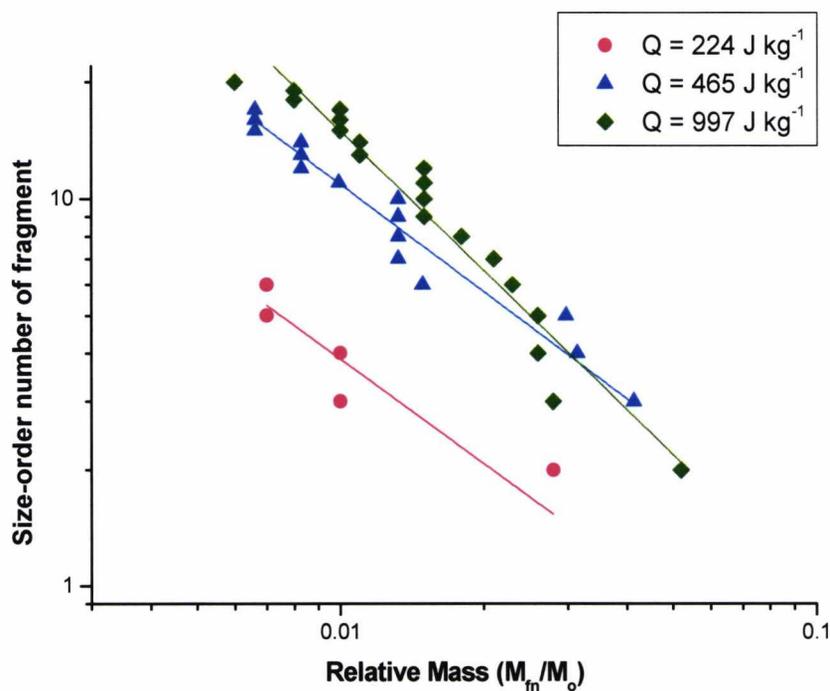


Figure 6.3.3 – Power law fits for impacts by a 1mm titanium projectile onto kiln sand cores. Fit coefficients are shown in Table 6.3.1

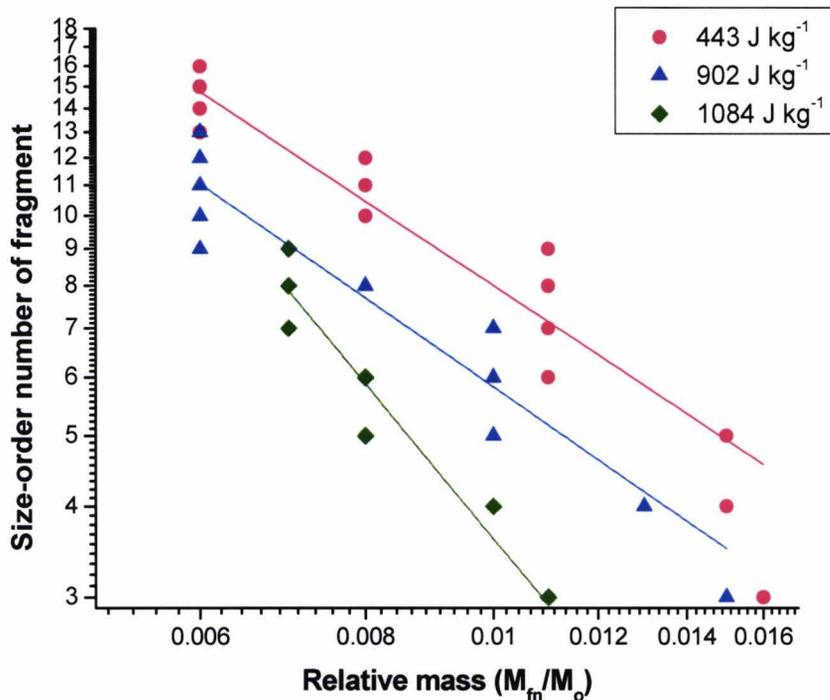


Figure 6.3.4 – Power law fits for impacts by a 1mm titanium projectile onto low temperature kiln-sand cores. Fit coefficients are shown in Table 6.3.1

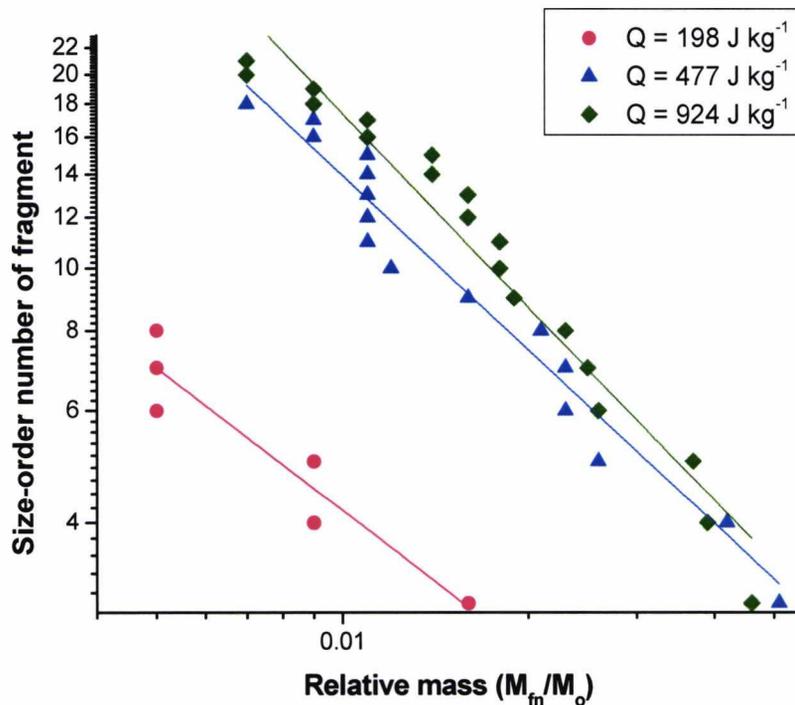


Figure 6.3.5 – Power law fits for impacts by a 0.8mm stainless steel projectile onto fraction E cores. Fit coefficients are shown in Table 6.3.1

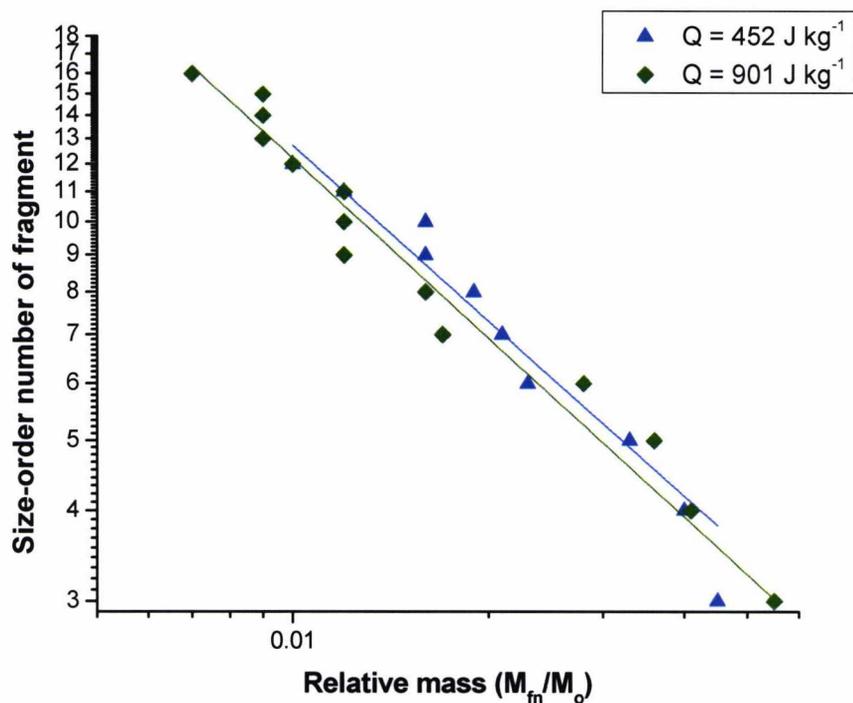


Figure 6.3.6 – Power law fits for impacts by 0.8mm stainless steel projectiles onto fraction D cores. Fit coefficients are shown in Table 6.3.1

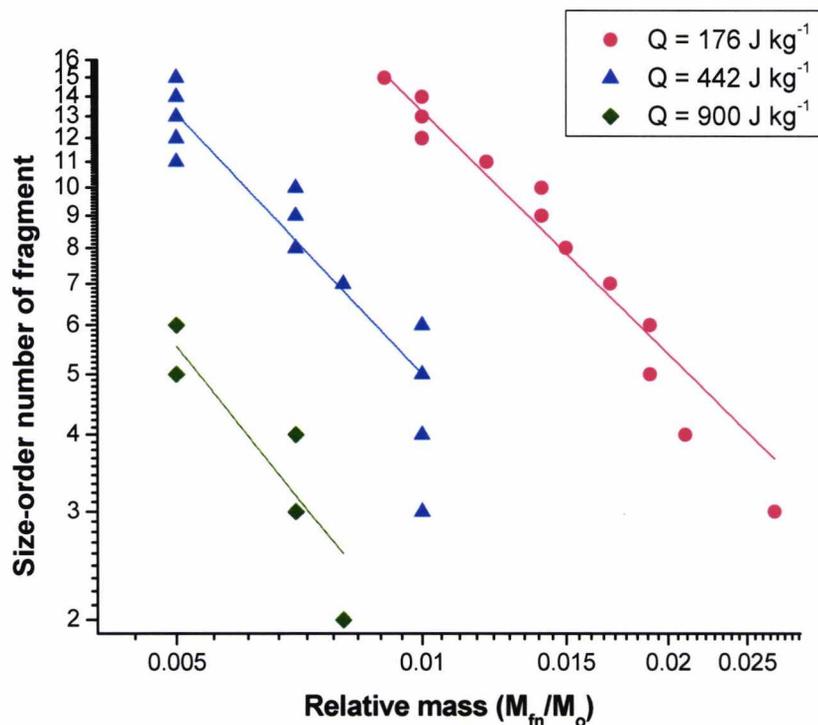


Figure 6.3.7 – Power law fits for impacts by 0.8mm stainless steel projectiles onto fraction E cores. Fit coefficients are shown in Table 6.3.1

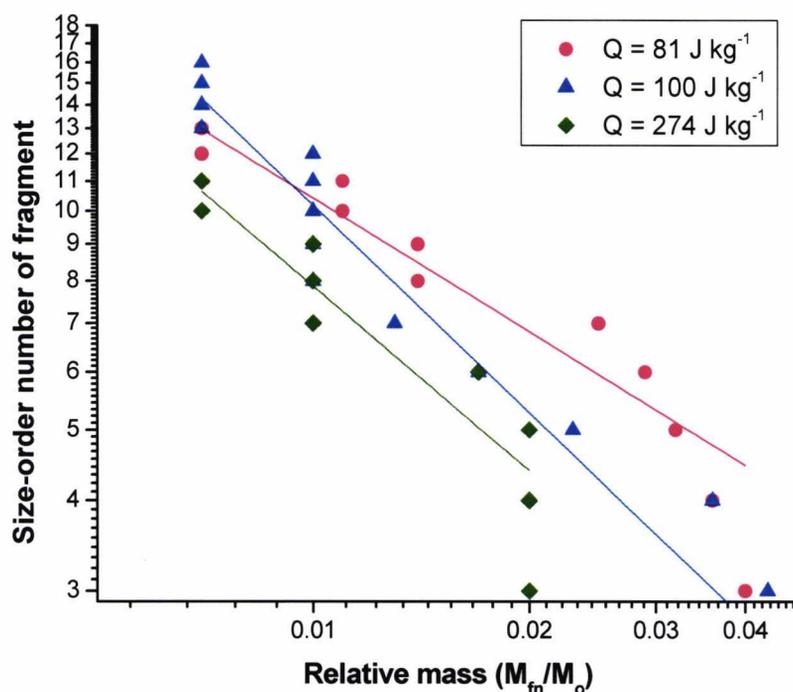


Figure 6.3.8 – Power law fits for impacts by 1mm aluminium projectiles onto solid ice cores. Fit coefficients are shown in Table 6.3.1

The coefficients derived from these fits are shown in Table 6.3.1. The square of the regression coefficient R^2 is an indication of how good the associated power law fit is; it essentially works by doing a least squares fit to compare the positions of the data points on the graph to the positions the generated power law says they should occupy. An R^2 value of 1 would be a perfect match while 0 indicates no correlation of the data with the fit function; all but one value obtained here are 0.80 or above, indicating the power laws obtained are reasonably good fits.

The power laws fitted to the size distribution plots have the form:

$$y = ax^b$$

Since the size distributions are plotted on log scales, this equation becomes:

$$\log(y) = \log(ax^b) \quad (\text{Eq. 6.3})$$

which simplifies to:

$$\log(y) = \log(a) + b \log(x) \quad (\text{Eq. 6.4})$$

This resulting equation after the logs are taken is of the form:

$$y = a + bx \quad (\text{Eq. 6.5})$$

which is the equation for a straight line on a linear scale. The a term therefore controls the overall scale of the fit, whereas b represents the slope. Since it is the slope of the fit we are interested in it is this term which is focused on in subsequent discussion. Considering the power b only, then, we can plot how this varies with Q/Q^* ratio for each projectile/target combination to see what trends, if any, are produced.

The 0.8mm stainless steel onto kiln-sand cores plot shown in figure 6.3.9 features two shots with almost identical Q/Q^* ratios, but with b values that are outside each other's error range. This points to a heavy amount of natural scatter present in all data that could have been caused by any one or a combination of factors – minor variations from target to target and projectile to projectile, where the projectile impacted the target, how many fragments were recovered and so on. These results should be considered with this caveat in mind. However, some general conclusions can be drawn from the data available.

Six of the seven plots of b against Q/Q^* (figures 6.3.9 through 6.3.14) display rough increases in b with Q/Q^* . While b values cannot be directly compared between different projectile/target combinations due to the shift in Q^* and relative target strength with projectile type, an approximate trend can be observed; Q/Q^* ratios of less than 1 produce shallow size distributions, but the slope of the distribution increases sharply as Q/Q^* approaches 1 (see fraction E cores, figure 6.3.14). This is to be expected since the amount of mass being removed from the target and ground down into smaller fragments increases exponentially as Q approaches the disruption limit – it represents the shift from the cratering regime to the disruption regime.

Between Q/Q^* ratios of 1 and 3 the change in b is much less pronounced (figures 6.3.10, 6.3.11, 6.3.13). In this Q/Q^* range the impact outcome remains in the

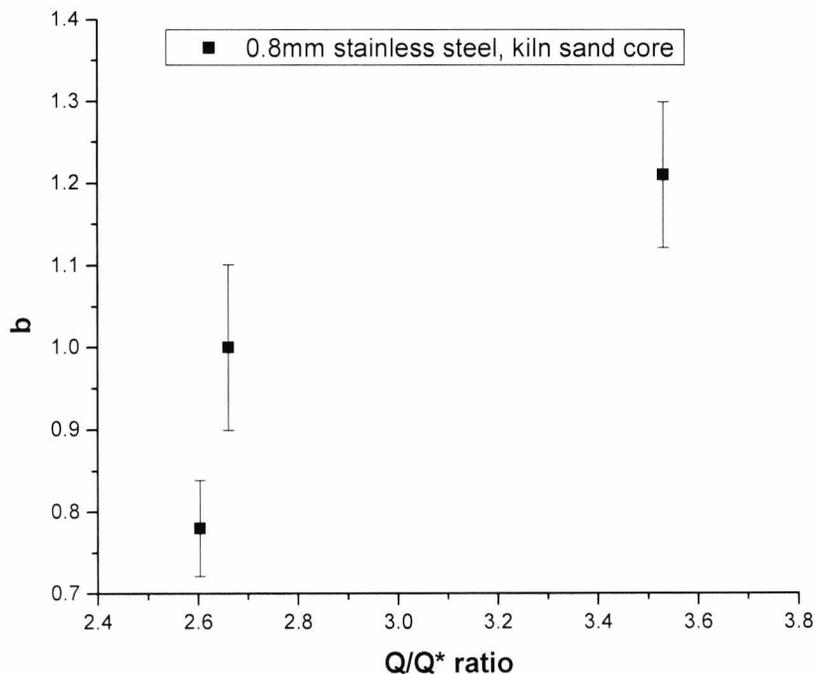


Figure 6.3.9 – Plots of size distribution slope b against Q/Q^* for 0.8mm stainless steel impacts onto kiln sand cores. Data taken from Table 6.3.1.

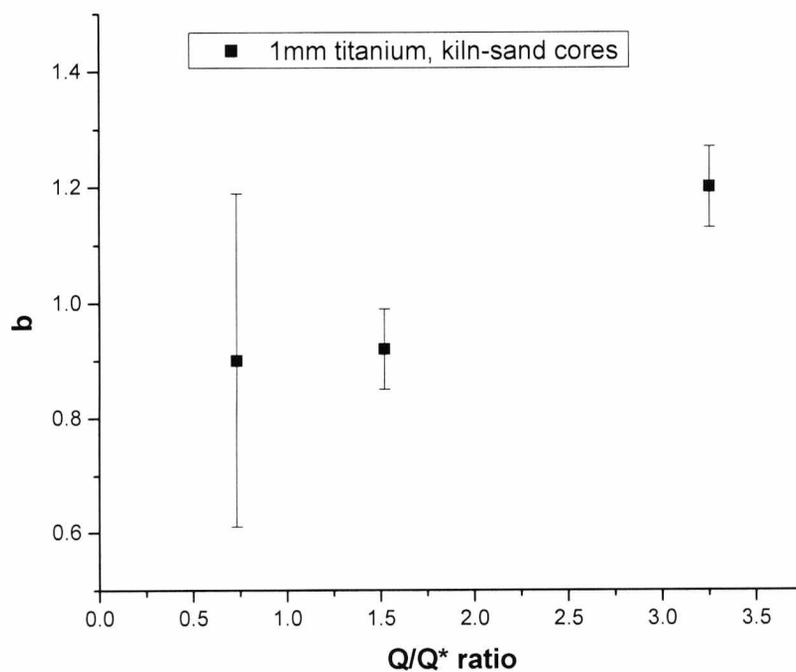
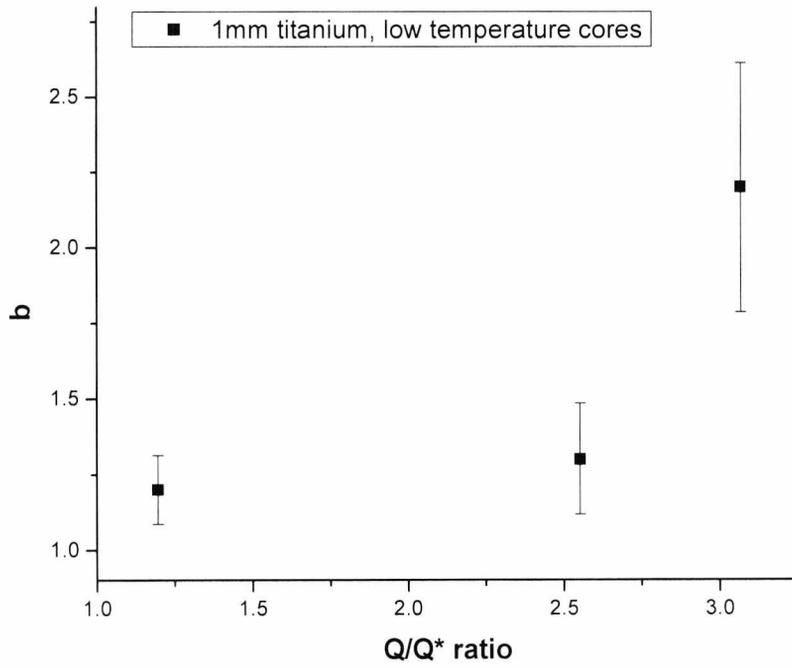
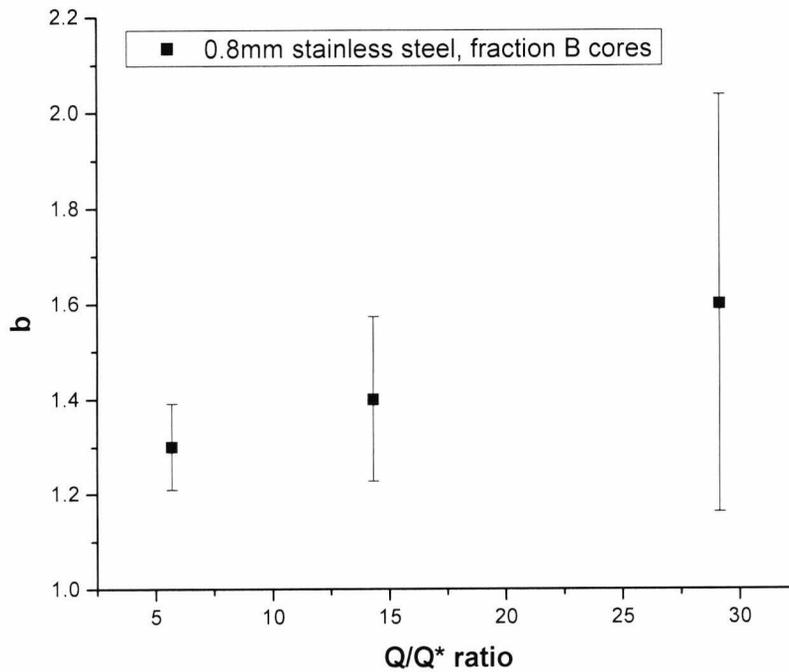


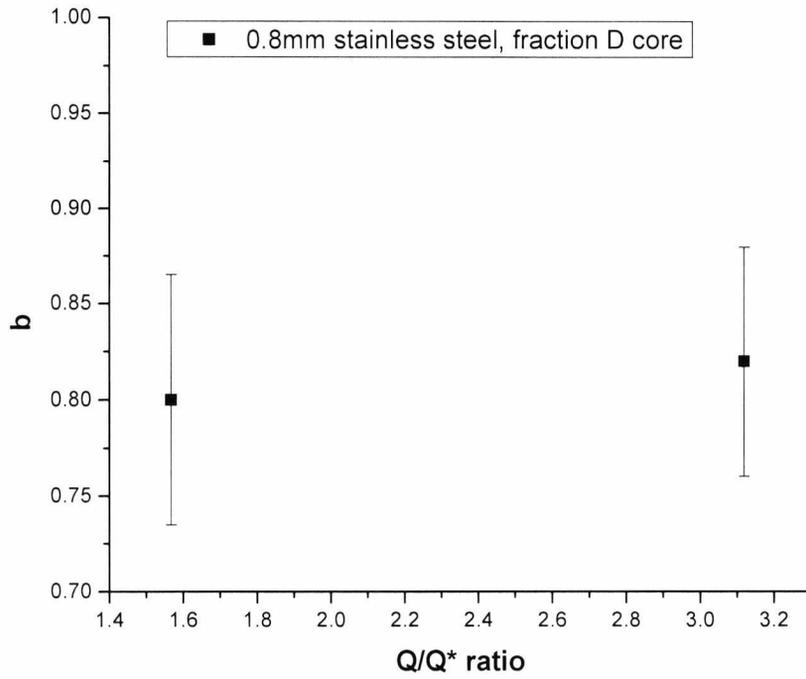
Figure 6.3.10 – Plot of size distribution slope b against Q/Q^* for 1mm titanium impacts onto kiln-sand cores. Data taken from Table 6.3.1.



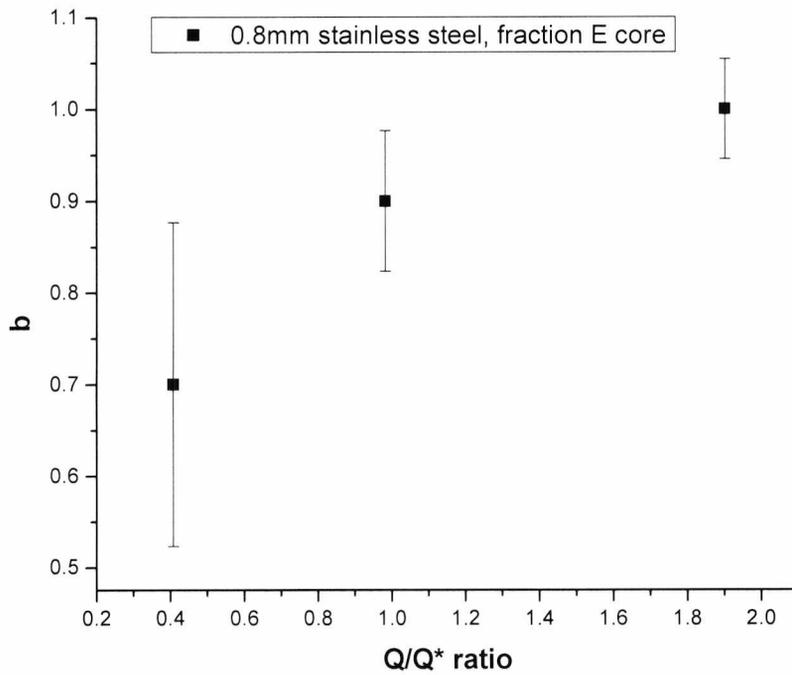
Graph 6.3.11 – Plot of size distribution slope b against Q/Q^* for 1mm titanium impacts onto low temperature kiln-sand cores. Data taken from Table 6.3.1.



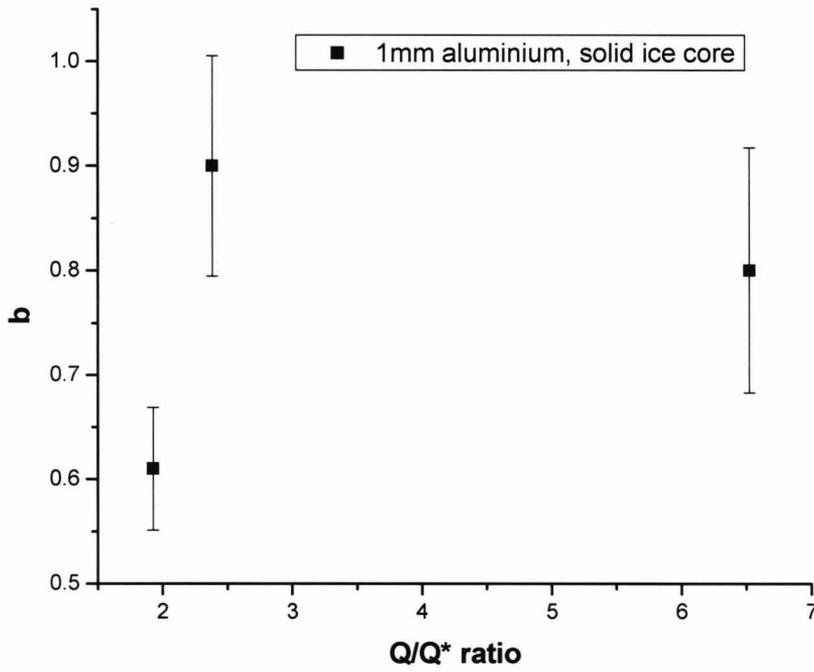
Graph 6.3.12 – Plot of size distribution slope b against Q/Q^* for 0.8mm stainless steel impacts onto fraction B cores. Data taken from Table 6.3.1.



Graph 6.3.13 – Plot of size distribution slope b against Q/Q^* for 0.8mm stainless steel impacts onto fraction D cores. Data taken from Table 6.3.1.



Graph 6.3.14 – Plot of size distribution slope b against Q/Q^* ratio for 0.8mm stainless steel impacts onto fraction E cores. Data taken from Table 6.3.1.



Graph 6.3.15 – Plot of size distribution slope b against Q/Q^* for 1mm aluminium impacts onto solid ice cores. Data taken from Table 6.3.1.

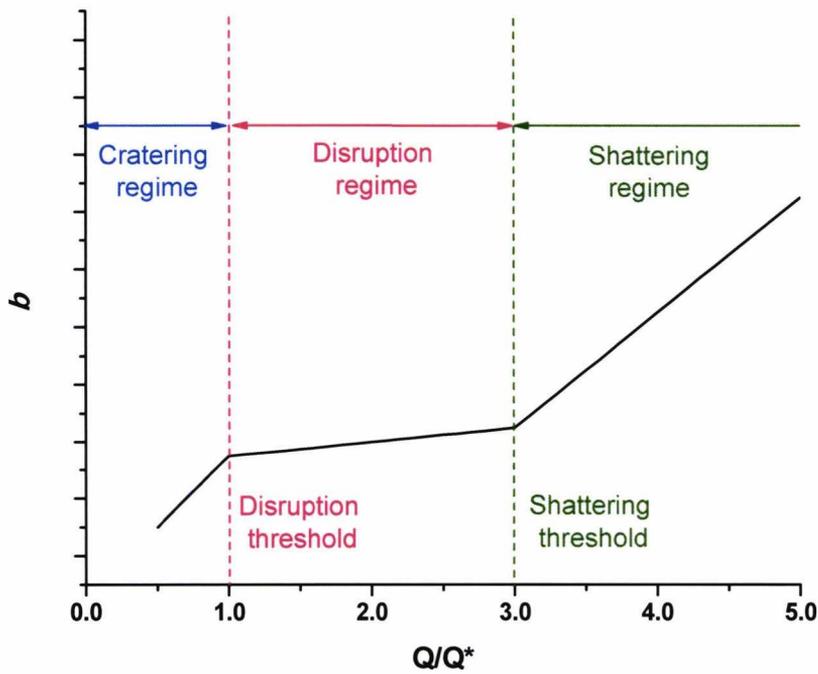


Figure 6.3.16 – Hypothetical plot of Q/Q^* vs. b relation, with the different impact regimes marked.

disruption regime - single large fragments being produced along with several smaller fragments, although the number of smaller fragments will still increase with Q – and since impacts of this type produce similar fragment spreads there is only a small change in b .

Finally, once Q/Q^* moves past 3 (6.3.9, 6.3.10, 6.3.11) b once again starts to increase at a larger rate, producing very steep slopes. This rate of increase seems to remain constant past this point, as evidenced by the graph for fraction B targets (figure 6.3.12). This can be explained by the energy density becoming large enough for the impact outcomes to enter the shattering regime; there is no longer a largest fragment which is comparable in mass to the rest of the fragments put together. Instead the target is totally shattered into many tiny fragments, the largest five or six of which

(for example) might all be comparable in size to each other. The large number of fragments and the very similar fragment masses produce very steep slopes, and hence large values of b .

A theoretical plot of this relation is shown in figure 6.3.16. From this, we can identify the various impact regimes as follows depending on how b behaves within them:

- **Cratering regime:** $Q/Q^* = 0.5 - 0.8$, steep increase in b .
- *Disruption threshold:* $Q/Q^* = 0.8 - 1.2$; theoretically this should be at $Q/Q^* = 1.0$ but judging from the experimental data the change in rate of increase of b takes place in this range of Q/Q^* .
- **Disruption regime:** $Q/Q^* = 1.2 - 2.8$, shallow increase in b .
- *Shattering threshold:* $Q/Q^* = 2.8 - 3.2$; again, there are examples of b that increase sharply before this point and examples that increase after it, making an exact threshold hard to determine.
- **Shattering regime:** $Q/Q^* > 3.2$, steep increase in b .

The Q/Q^* values for these regimes mesh well with the values for each regime determined in section 6.1.

Again, this analysis needs to be considered with the scatter on the data and the large error bars in mind, and there is one example of a set of size distributions that goes against the general relationship described here (figure 6.3.14). However, the shifting value of b at the $Q/Q^* = 2.5-3.0$ mark is undeniable; something is definitely happening to drastically increase the slopes of the size distributions produced at that point for sand-ice cores. Further work in this area would be extremely useful in determining if b values do indeed follow this pattern – twenty or thirty different size distributions for a single target/projectile combination spanning a range of Q/Q^* from 0.5 to 6.0 will definitively prove or disprove it.

6.4 - Conclusion

In this chapter we have seen a detailed analysis of the fragment size distribution data produced by the experimental programme. The factors that affect size distribution are investigated and identified: the material of the projectile used in an impact does not alter the resulting size distribution in a significant way, but large changes in the material of the target (i.e. sand-ice to ice) do. Lowering the target temperature also changes the resulting size distributions. A rough classification of the various types of size distribution produced according to impact outcome is suggested and an attempt to quantify this by fitting power laws to the size distribution data is made. Impacts appear to fall into three categories - cratering impacts, disruption impacts, and shattering impacts - with distinctly different types of size distribution produced for each one. These categories appear to be borne out by the power law fits to the cumulative mass distributions, which behave differently according to which impact regime a particular impact is in.

In essence:

- Cratering impacts occur in an energy density range of $Q/Q^* = 0.5 - 1.0$; the rate of change in the slope of the fragment size distributions with impact energy density over this range is constant and rapid. This reflects the initial shift from an intact target to one which is having up to half of its mass converted into fragments; the more mass disrupted, the more fragments present.
- Disruption impacts occur in an energy density range of $Q/Q^* = 1.0 - 3.0$. The rate of change in the slope of the fragment size distributions with impact energy density is shallower in comparison to cratering impacts. This may be because in this phase the bulk of the impact energy is going into reducing the size of the largest fragment rather than splitting the smaller fragments down (which would result in an exponentially increasing number of fragments and a steep increase in b similar to the cratering regime).
- Shattering impacts occur at energy densities greater than $Q/Q^* = 3.0$. The rate of change in the slope of the fragment size distributions with impact energy density becomes steeper again, reflecting a shift away from the reduction in size of the largest fragment and back to many smaller fragments being broken down.

Chapter 7 - Discussion

7.0 - Introduction

In Chapter 5 we have seen that adding an ice layer to an ice-silicate core has some curious effects on the strength of the body as a whole. Thin layers only a few millimetres thick have little to no effect on the critical energy density Q^* , while thick layers 5mm thick or greater increase the disruption energy by a factor of three. This chapter discusses reasons why this might be, as well as exploring other observed behaviour such as the shattering regime and the relevance of results for Q^* obtained in the laboratory for bodies on the scales encountered in the TNO region.

7.1 – The shattering regime

Disruption results from pure ice targets as well as fragment size distributions from all target types point towards targets undergoing a third mode of behaviour at a range of energy densities above those of the cratering and disruption regimes. This has been termed the shattering regime, since the fragment size distribution data shows the targets breaking into a large number of small, equally-sized fragments rather than the more graded fragmentation of the target that occurs in the disruption regime.

Why this third regime exists is not known. Judging from the progression in fragment size distributions from cratering to disruption to shattering, it appears that as the target fragments become smaller they become harder to disrupt further, as indicated by the evening-out in the range of sizes observed in the largest fragments. This points to an even distribution of energy over these largest fragments; I suggested in chapters 5 and 6 that the difficulty in breaking down a target at high energy densities was somehow caused by this partitioning of the impact energy over a large number of fragments. This assumes, however, that small target fragments on the scales observed in the shattered targets are intrinsically harder to fragment than a single larger fragment. Another possible culprit is the conversion of impact shock energy to fragment kinetic energy – in other words, the higher energy impacts not only break down a target into

many smaller fragments but they also cause these fragments to fly away from the impact site with a great deal of energy. Any impact energy converted to kinetic energy in this way should be effectively lost for the purposes of fragmentation and disruption, and would explain the increased difficulty in grinding down a target at high energy densities.

One way of checking this would be to capture images of an impact using a high-speed camera. Fragment velocities could then be measured during an impact event to see if there is any unusual relationship between impact energy and fragment velocity.

7.2 - Layered targets

That thin-shelled targets produce only a small change in Q^* while thick-shelled targets amplify Q^* considerably indicates that the thickness of the ice shell is the determining factor in causing this strength increase. However, whether this is a simple reliance on shell thickness (possibly as a ratio to the size of the projectile) or a function of the overall mass of the ice shell remains to be seen. The two are related in any case - a thicker ice shell results in more mass and vice versa, so the question remains the same: why do thick ice shells increase the strength of the target where thinner ice shells do not? It certainly has little to do with the simple act of adding more and more ice to the target; since ice has a lower disruption energy compared to the ice-silicate core increasing the fraction of the target that is composed of ice should decrease the required disruption energy for the target as a whole, not raise it.

As stated in chapter 5, the ice layer showed a tendency to disrupt completely before any significant damage is done to the core. This suggests that the impact energy is somehow being concentrated and contained in the ice shell, and that only impacts sufficiently energetic to shatter the shell completely succeed in transferring significant energy to the core. Moreover, the ice shell is absorbing an amount of energy disproportionate to its mass. Contributing factors that may produce effects are as follows.

First, the method used to produce the layered targets creates a physical discontinuity

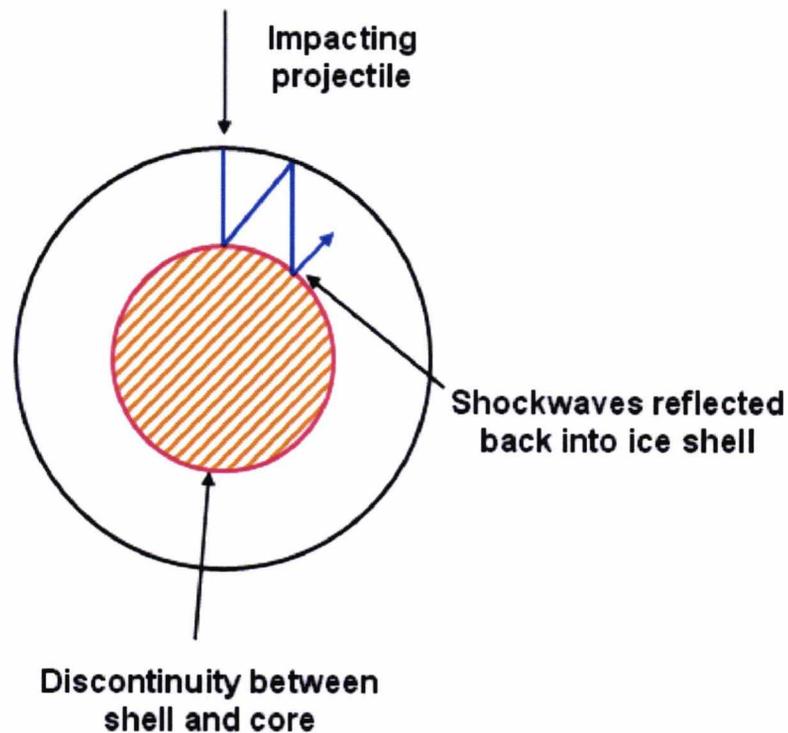


Figure 7.2.1 – Diagram showing one possible mechanism that would concentrate impact energy into the ice shell. The discontinuity between the shell and the core reflects the impact shockwaves back into the ice shell, permitting only a small amount of “contact” energy to be transmitted to the core. This results in minimally damaged cores and totally destroyed shells, as observed.

between the ice shell and the ice-silicate core. There is a definite point where one material stops and another begins, and this could have the effect of altering the behaviour of the shockwave; instead of being transmitted throughout the target as would happen in a single material homogenous body, the shift in refractive index between the two materials may refract or even reflect the shockwave so that the majority of the impact energy is contained within the shell.

The addition of the shell on top of the core may also be a factor. Ideally these targets would consist of one solid piece of ice with sand added to the centre of the target to make up the core. As the shell instead consists of layers of water consecutively frozen on to the core it follows that, since it was added later instead of the entire target being frozen at once, the target consists of two ice elements frozen together instead of a single contiguous mass of ice. The bond between the shell and the core will therefore be much weaker – in addition to adding to the effect of the physical discontinuity

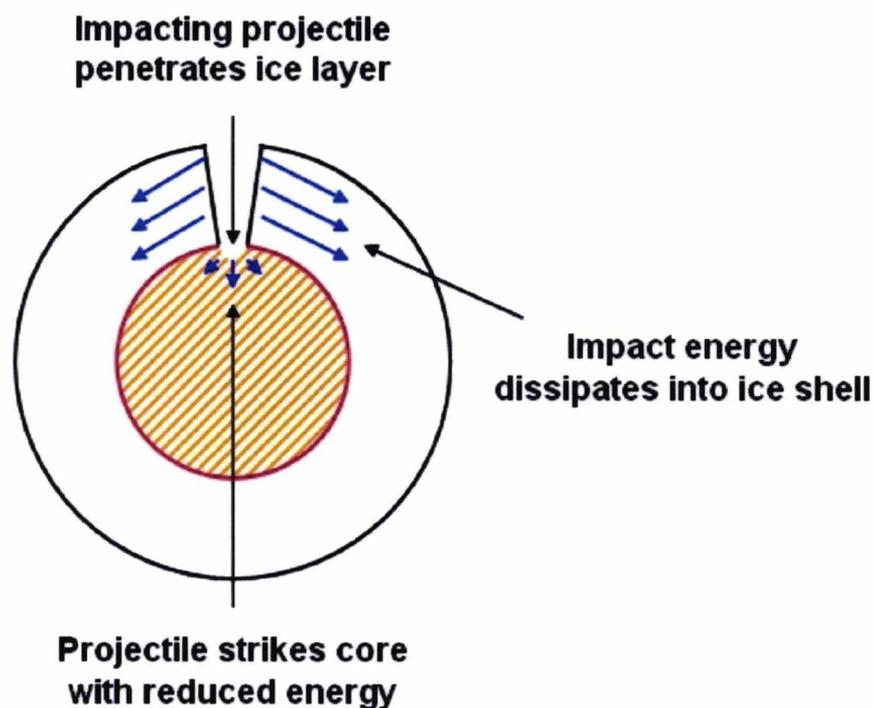


Figure 7.2.2 – Diagram showing how the ice shell may absorb the majority of a projectile’s impact energy even if it penetrates the shell and strikes the core.

between the materials mentioned above, if it is relatively easy to separate the shell from the core then any energy still contained within the shell fragments can no longer be transmitted to the core. This is supported by the behaviour observed in the core-mantle targets impacted by Okamoto *et al.* (2008); there, as here, the targets were created by adding a mantle onto an already existing core. The gypsum mantles used in Okamoto *et al.* are only weakly coupled to the glass core through baking, and similar removal of the mantles is observed before any significant damage is done to the core.

This brings me to the third factor, which is that a significant amount of impact energy may not go into breaking the target as usual, but rather into imbuing the shell fragments with kinetic energy causing them to fly away from the core as described in section 7.1. Any energy spent in this way is effectively lost for the purposes of disruption. Assuming that once total removal of the shell occurs all shell fragments are given some amount of kinetic energy, this would greatly increase the amount of energy needed to disrupt the core as well.

In order to test some of these suppositions and model how the shock wave is transmitted through a target – in particular, how an impact shock wave might interact with the interface between the ice shell and the core in a layered target – hydrocode modelling was carried out by M. Price at the University of Kent. Hydrocodes are a type of computer program designed to model how materials behave under a variety of impact conditions. They achieve this by breaking down the material in question into a matrix of cells (which can be two- or three-dimensional) and modelling fluid flow between these cells using equations of state. A detailed explanation of the mechanics of hydrocode modelling can be found in Zukas (2004).

There are two main types of hydrocodes used to model solid materials: Euclidean, and Lagrangian. Euclidean hydrocodes model the cells as fixed structures with fixed shapes and volumes through which the material flows according to the equations of state. Lagrangian hydrocodes fix the amount of mass in a cell, and instead deform the cells themselves in order to model an impact. There is a third type of hydrocode called SPH (Smooth Particle Hydrocode) which breaks down a material into a series of point-like particles with no connecting bonds; this makes it useful for modelling the behaviour of gases, but it is of little use when modelling solid structures.

Hydrocodes are complex computer programs, but a very basic description of how they work would be to say that they operate in a series of iterative loops using equations of state to calculate how various physical quantities – such as pressure, volume and temperature – change from loop to loop. The initial state of every cell is fed into the equations of state, which then output how that cell will change over a small period of time. That output state is then used as the input state for the next loop.

Calculating a single loop for a single cell would be relatively easy, but to model a solid structure accurately requires calculations for thousands of cells through thousands of loops in order to make the time step between loops as small as possible. This is why hydrocodes require extremely powerful computers to model, with the amount of processor power available being the main constraining factor on the complexity and accuracy of any given hydrocode.

The modelling described here used a program called Autodyn (version 12.1, from ANSYS) to simulate a spherical layered body under impact. The model for the impacted body was fully Lagrangian, with 282,000 cells. The projectile was composed of stainless steel 304 (simulated using a model for that material from the in-built Autodyn library of materials) and modelled using SPH; this is to ensure that the projectile dissipates upon impact since in this case we are only interested in what happens to the target. The ice layer was modelled using a five-phase equation of state devised by Senft & Stewart (2008), while the best approximation that could be found for the ice-silicate core was concrete, modelled using the in-built Autodyn library. The target was composed of a 40mm diameter concrete core surrounded by a 5mm thick ice layer, while the projectile was 1mm in diameter.

Two impacts were successfully simulated, one at 1 km s^{-1} and another at 3 km s^{-1} . An attempt was made to simulate a 6 km s^{-1} impact, but the model ran into problems with glitched cells and had to be halted.

Figure 7.2.3 shows the 1 km s^{-1} target 130000 cycles – or 2.794×10^{-2} seconds – after impact. This visualisation is intended to show clearly the division between the ice layer and the concrete core, with the shock wave absent. The projectile has penetrated the ice layer and impacted the concrete core, creating a small crater – smaller than would be expected from an ice-silicate core given the increased material strength of concrete. The ice layer immediately surrounding the impact site is in the process of spalling away from the core.

Figure 7.2.4 shows the target at the same point in time as figure 7.2.3, but this time with the shock wave visible. The shock pressure immediately below the impact site is 3.8 MPa; with the shock wave propagating inward and upward relative to the orientation of the picture. Interestingly, the interface between the ice layer and the core does seem to inhibit and reflect the transmission of the shockwave somewhat. There is a shock front being transmitted from the core into the ice layer, but this is not transmitted further into the ice layer; instead, it is causing the ice layer to spall and

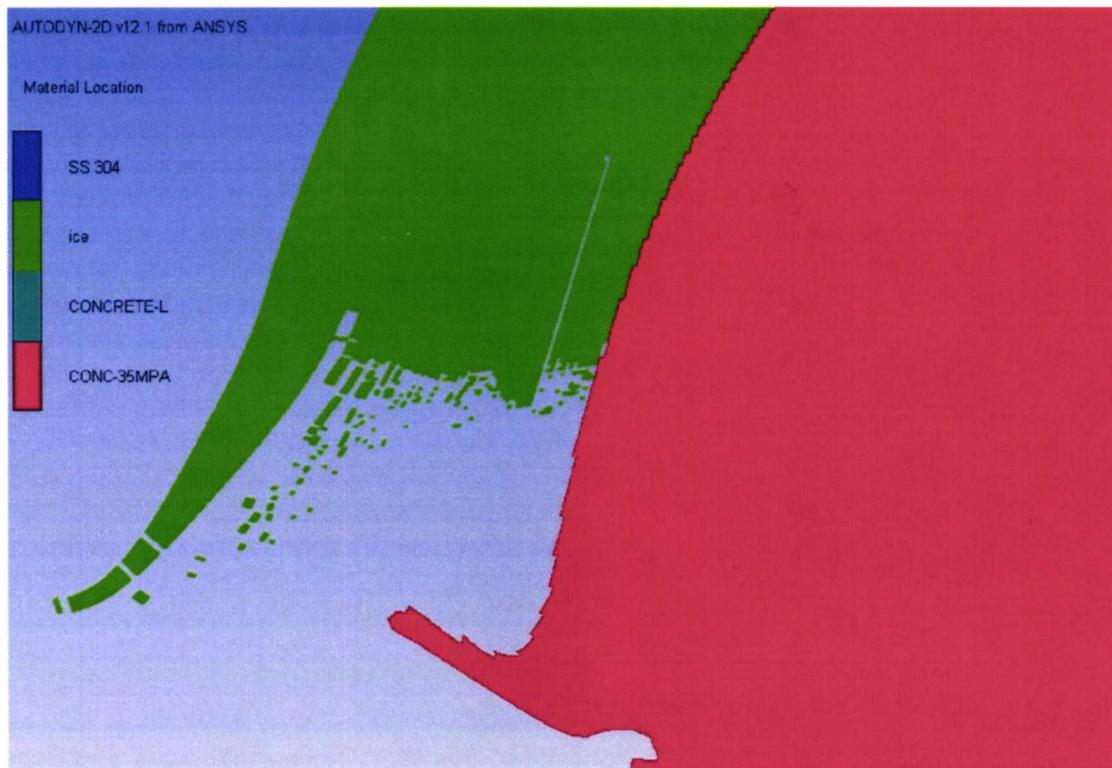


Figure 7.2.3 – Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 1 km s^{-1} .

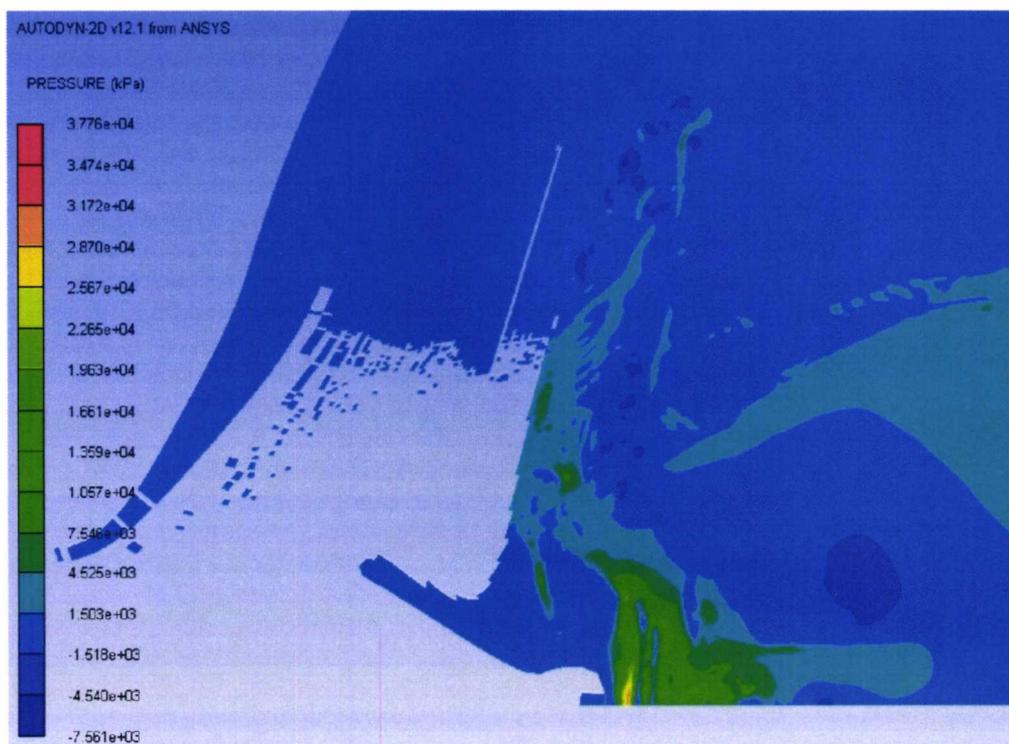


Figure 7.2.4 - Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 1 km s^{-1} , with the impact shockwave visible.

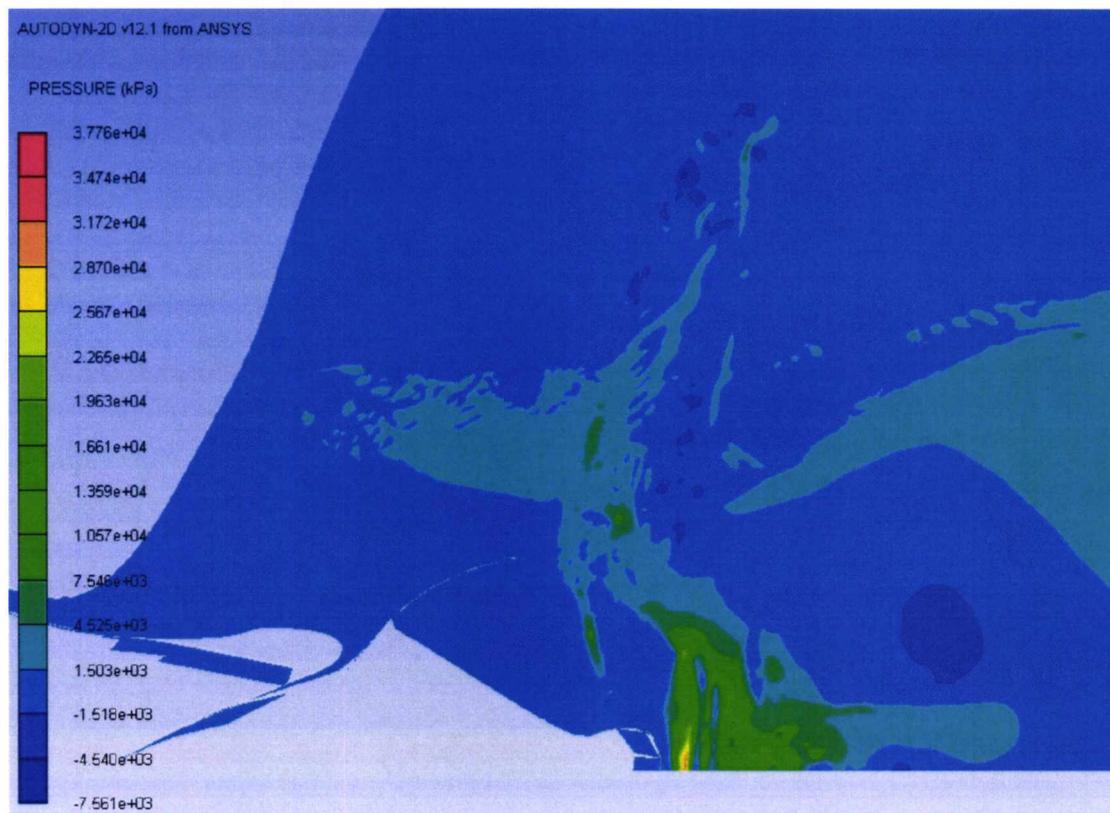


Figure 7.2.5 - Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 1 km s^{-1} , with the impact shockwave visible and spalled fragments included.

separate away from the core. Figure 7.2.5 shows the simulation with these spalled fragments removed, and here the reflection of the shock wave from the interface is much clearer.

Similar behaviour is observed in the 3 km s^{-1} model (observed earlier in the simulation at 53000 cycles, or 9.768×10^{-3} seconds after impact; figures 7.2.6, 7.2.7, 7.2.8) with more ice spalling from the ice layer and a larger crater in the core thanks to the more energetic impact, which also produces a slightly higher shock pressure of 4 MPa.

This modelling is very interesting, since it partly confirms some of the theories put forward in section 7.2, yet not quite in an expected way. Reflection of the shockwave from the interface between the ice layer and the core was predicted, but it appears that there is very little shock energy transmitted directly into the ice layer from the

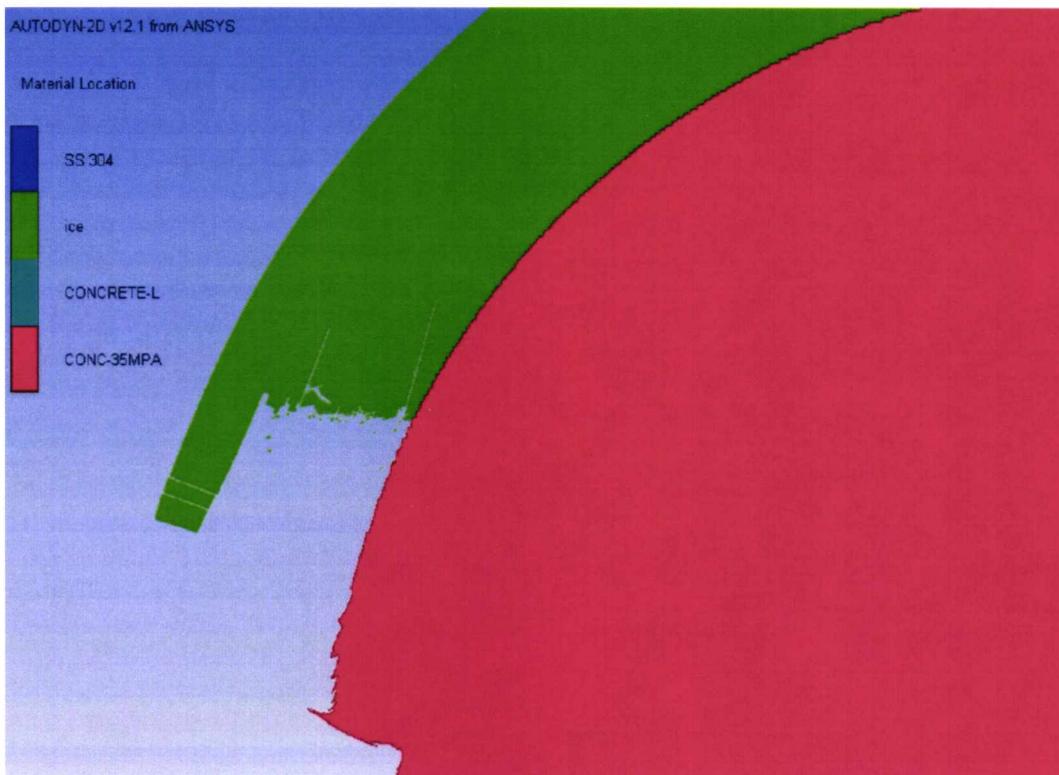


Figure 7.2.6 - Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 3 km s^{-1} .



Figure 7.2.7 - Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 3 km s^{-1} , with the impact shockwave visible.

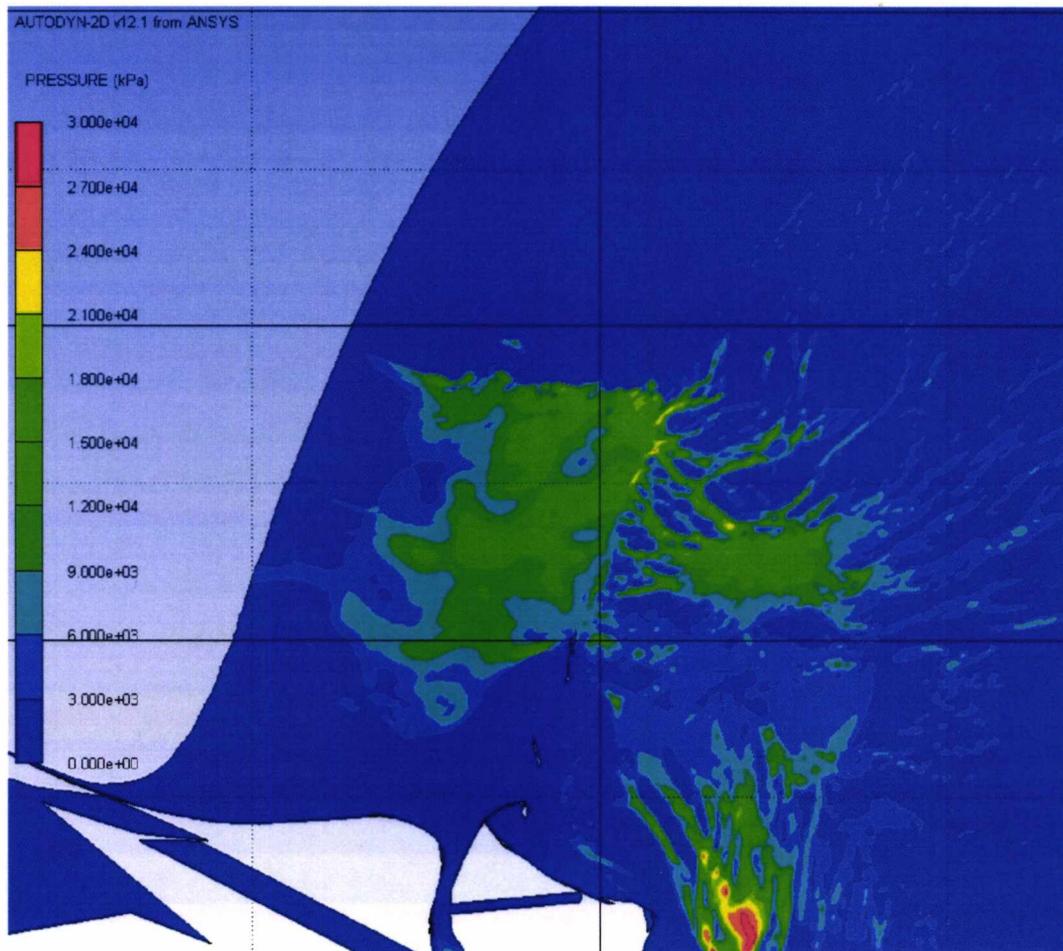


Figure 7.2.8 - Hydrocode simulation of a 1mm diameter stainless steel projectile impacting a spherical target with an ice layer resting on top of a concrete core at 3 km s^{-1} , with the impact shockwave visible and spalled fragments included.

projectile during impact. The ice immediately surrounding the impact site melts and spalls, but this inhibits transmission of a shockwave through the ice itself. Instead, shock energy is delivered to the core, where the shock wave travels out and up (relative to the figures). When the shockwave meets the interface between the core and the ice layer some of the shock energy bleeds over into the ice, causing it to spall away from the core, while the rest of the shockwave is reflected back into the core.

The modelling has one significant limitation beyond the material ones already described; it does not simulate a bond between the ice layer and the core. The ice layer is simply “resting” on top of the core, and thus the only thing keeping it attached to the core during impact is its own material strength forcing it to retain a spherical shape. Fortunately this is not such a bad approximation of the layered targets used in the experimental programme, since in those cases the ice layers were frozen onto the

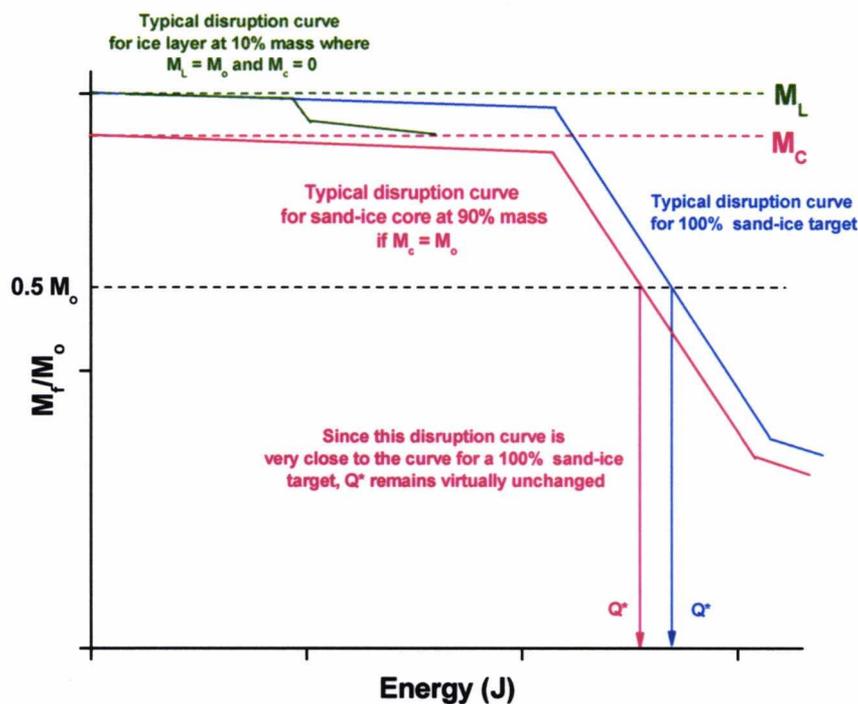


Figure 7.2.9 – Separate disruption curves for the ice layer and the sand-ice core in a body where the ice layer has 10% of the mass of the target. In this case the energy density required to disrupt the ice layer is so low that it falls within the cratering regime for the sand-ice core, rendering the addition of the shell irrelevant once energy densities enter the disruption regime which defines Q^* . Thin-layered targets seen in chapter 5 fall into this category, and so this helps to explain why a thin layer has little to no effect on the critical energy density Q^* .

core after the core had been manufactured rather than the whole body being created at once. This results in the bond between ice layer and sand-ice core being a simple attachment between two pieces of ice that is relatively easy to sever under impact.

Regardless, these assimilations broadly agree with the observed laboratory results, with wide-ranging disruption of the ice layer coupled with a relatively small crater in the concrete core, and thus provide a little more insight on how a shell protects a core during an impact event – according to the evidence seen above, the ice layer provides a sort of “cushion” that absorbs shock energy and then spalls away from the core, effectively insulating the core from the effects of impact.

Thinner shells are less effective at mitigating impacts. The opposite is true for thick shells; a shell with 40% of the total mass of the target and 5-7 times the thickness of

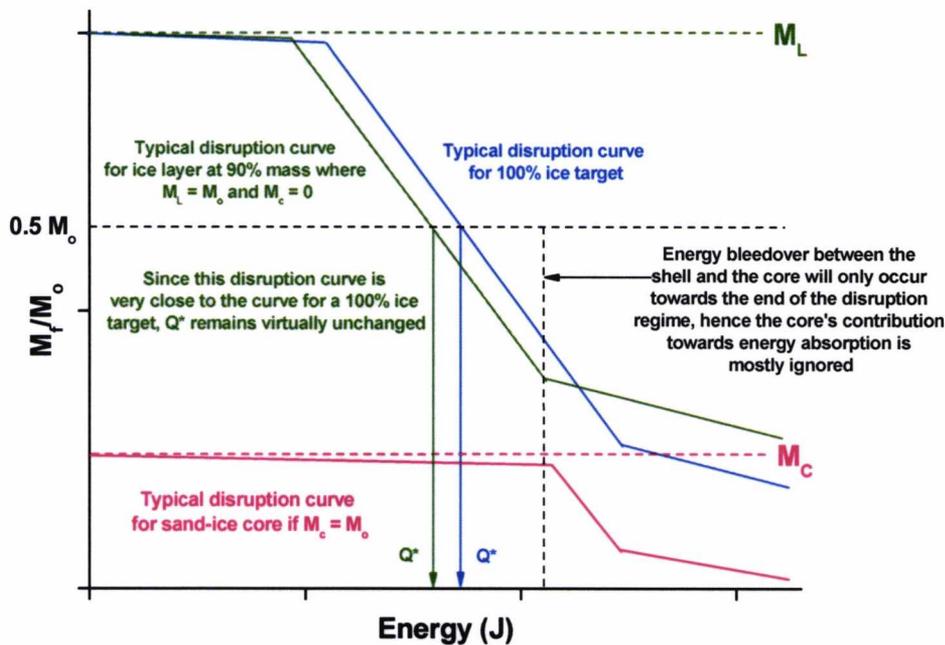


Figure 7.2.10 – Separate disruption curves for the ice layer and the sand-ice core in a body where the ice layer has 90% of the mass of the target. In this case there is so much ice present in the target that disrupting 50% of the target only requires disruption of 56% of the ice layer, which ensures that the critical energy density of this body will be almost the same as the critical energy density of a body composed of 100% ice; the sand-ice core distorts the result slightly, but not significantly.

the projectile striking it increases the amount of energy required to disrupt a target by nearly 300%. However, it does not automatically follow that increasing the shell thickness still further will improve the energy absorption qualities of the shell.

Consider the definition of Q^* used in this work. If Q^* is the impact energy density that will leave a largest fragment that is half the mass of the original target, then increasing the mass of the ice shell while keeping the mass of the core constant will

have diminishing returns in terms of preserving a largest fragment if the shell approaches a dominant percentage of the mass of the whole target. Ice has a much lower critical energy density than the sand-ice core, and the shell in isolation disrupts much in the same way as a body constructed solely of pure ice would disrupt. While the core is insulated from the effects of impact except at very high energy densities, if

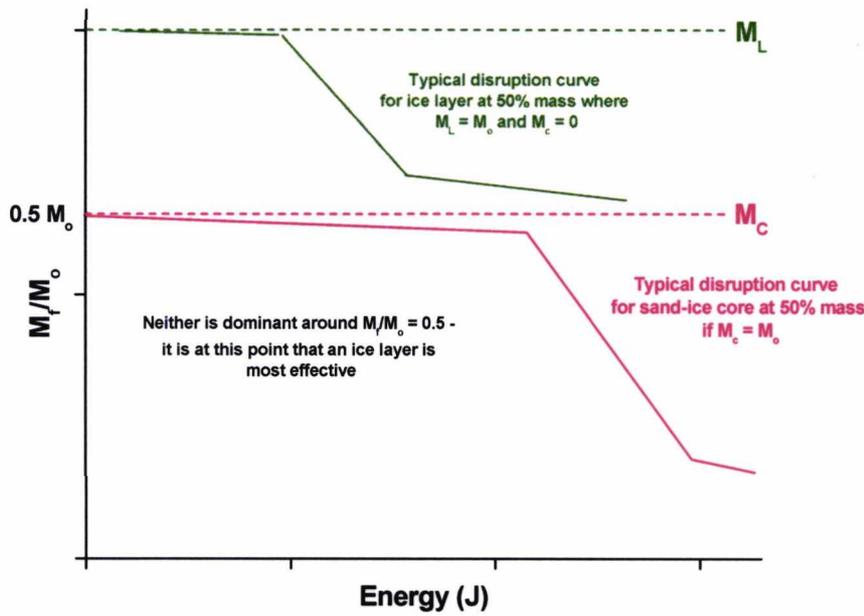


Figure 7.2.11 – Separate disruption curves for a body that is 50% ice and 50% sand-ice. Neither material is dominant around the $M_f/M_o = 0.5$ region that defines Q^* . As a result the shallower curve of the ice layer’s shattering regime has the greatest effect on Q^* , distorting it to a much greater degree than it would were the body dominated by a single material.

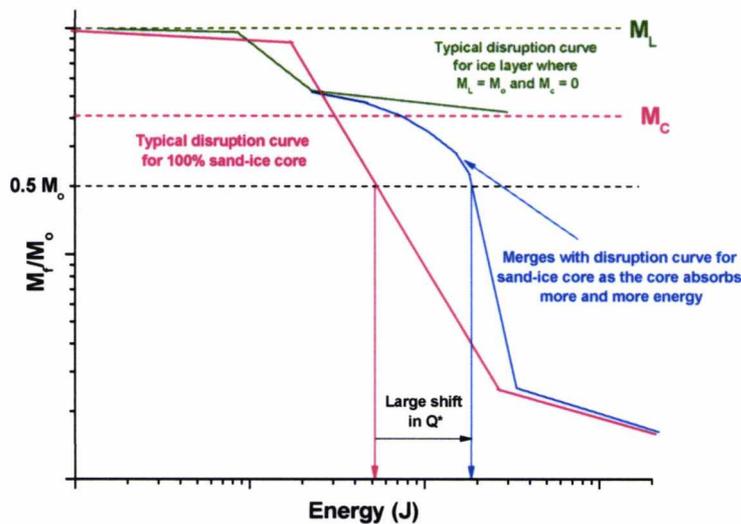


Figure 7.2.12 – Plot showing energy bleedover between the two separate types of disruption curve as the ice layer moves into the shattering mode of behaviour; increasing amounts of energy are absorbed by the core rather than the ice layer as the impact energy is increased eventually merging with what would be expected of a 100% sand-ice target. However, at the M_f/M_o region there is still a significant difference between the 100% sand-ice plot and the merged plot, explaining the large increase in Q^* observed in targets with ice shells comprising 30-40% of the target mass.

the mass of the core makes up less than half of the mass of the target then the critical energy density of the target as a whole starts to approach the critical energy density of ice.

This can be illustrated by modelling the layered body as two separate bodies: an ice body and a sand-ice body, each with the same proportion of mass present in the layered target. It is required that the ice body be disrupted before any disruption of the sand-ice body can occur. This means that the disruption curve for both bodies can be modelled as a combination of the disruption curve for each body separately. The mass of the target M_0 is the sum of the mass of the ice layer M_1 plus the mass of the sand-ice core M_c . A body with a thin ice layer of 10% of the mass of the target would display the behaviour shown in figure 7.2.3. As can be seen, the disruption curve for the sand-ice core is easily dominant in the $M_f/M_0 = 0.5$ region, and thus the value of Q^* for a thin-layered body is close to that of a body with no ice layer. On the other hand if the mass proportions are switched the reverse is true. For a body where the ice shell makes up 90% of the mass of the target, as in figure 7.2.4, the $M_f/M_0 = 0.5$ region is squarely in the domain of the disruption curve for ice.

From this, I draw the conclusion that a body with only a 10% mass ice layer will only alter the top of the disruption curve away from what would be expected of a sand-ice target, while a body with a 10% mass sand-ice core will only alter the bottom of the disruption curve away from what would be expected of a pure ice target. Since it is not the top or bottom of the curve that dictates Q^* but rather the central $M_f/M_0 = 0.5$ region, Q^* remains unchanged for a body dominated by a single material. This explains why the thin ice layers had little effect on Q^* ; there simply wasn't a large enough mass of ice present to significantly shift the disruption curve away from what would be expected of a pure sand-ice target.

Now, for a body which is 50% ice and 50% sand-ice things are a little more complex. The $M_f/M_0 = 0.5$ region for the body as a whole is passes right along the boundary between M_1 and M_c . In this theoretical model where the shell must be completely removed before the core can be disrupted, all that has to happen to this body *is* for the

shell to be completely removed; since the core makes up 50% of the target mass it would count as disruption.

There is a catch, however. Since we are not talking about disrupting the shell (which would require the simple removal of half its mass) but rather shattering it completely, this will require energy densities that put the shell firmly into the shattering mode of behaviour. As seen in chapter 5, the shattering regime has a shallower Q^* vs M_f/M_0 plot, indicating that once the transition has been made from disruption to shattering it becomes more difficult to reduce the mass of the largest fragment. As a consequence of the region where the shell gives way to the core lying within this regime, the energy density required to disrupt the body by removing the shell will be increased significantly.

Of course the separation between the shell and the core is not absolute. Small amounts of mass are removed from the core even when the shell has not been removed, and increasing amounts are lost as the ice layer moves further into the shattering regime. The behaviour of the body as a whole under disruptive impact gradually switches from being mostly dominated by the ice layer component to being wholly dominated by the sand-ice core component. Figure 7.2.6 shows one way in which this transition may work: the curve for the combined body follows that of the ice layer until it moves out of the disruption regime. Once the layer reaches the shattering regime, it is assumed that the impacting projectile is now smashing through the ice layer with enough energy to cause increasingly significant amounts of mass loss to the sand-ice core, and the combined curve for the body will start to trend increasingly towards the disruption curve for a pure sand-ice body.

At what mass percentage is the ice shell most effective? Experimental evidence confirms that ice shells with 30-40% of the mass of the target more than double the critical energy density needed to disrupt the target, but Q^* must drop again at some point as the percentage of ice approaches 100%. My opinion is that increasing the mass of the ice shell much past 40% will result in lower critical energy densities. The simple model above indicates that maximum effectiveness should be achieved at 50% ice mass, but it does not take into account core mass loss while the ice shell is still

present. A largest fragment of $M_f/M_o = 0.5$ could therefore be possible without total disruption of the ice shell being required; an impact which disrupts only 80% of the ice shell (which would avoid the effects of the energy-sapping shattering regime) and 20% of the core would count as being disrupted. An ice mass of 40% compensates for the small amount of core mass loss that occurs, and so I think the point of maximum effectiveness at increasing Q^* occurs at or near 40%.

It is worth noting that, if the above model is correct, the increase in critical energy density experienced by targets with ice layers making up approximately 30-40% of the mass of the target may be a simple reflection of the way Q^* is defined. Q^* is the energy density at which half of the mass of the original target is disrupted. It therefore follows that ice shells making up just under half of the mass of the target will be most effective at increasing Q^* .

Consider the hypothetical scenario of “cratering” being defined as a result in which more than 90% of the target is left intact. Using this definition, the “thin” ice shells with 10-20% ice mass which had almost no effect on disruption would instead have a significant effect on this new cratering phenomena.

The above discussion should therefore be read with this in mind.

7.3 – Size scaling

Benz & Asphaug (1999) have performed modelling which allows a crude extrapolation of laboratory scale experimental data to solar system scale for rocky and icy bodies.

Benz & Asphaug’s method involved modelling the effect of impact on ice and basalt targets of differing sizes using a smooth particle hydrocode, and then fitting the results by eye to describe the changing critical energy density required to disrupt these targets as the size of the body increased. Their fit is described by the equation

$$Q^*_{D} = Q_o \left(\frac{R_{pb}}{1cm} \right)^a + B\rho \left(\frac{R_{pb}}{1cm} \right)^b \quad (\text{Eq. 7.1})$$

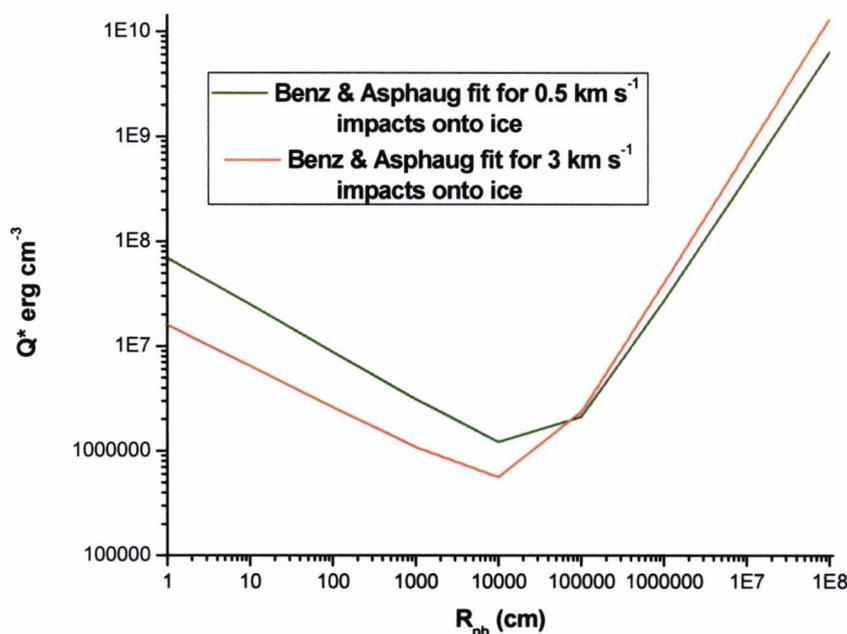


Figure 7.3.1 – Benz & Asphaug’s modelling for how Q^* changes for ice targets with the scale of the target.

Where Q^*_D is the modelled critical disruption energy, R_{pb} is the radius of the body being disrupted and ρ is the density of the body being disrupted. Q_0 , B , a and b are all constants that vary from impact to impact; these make it a little tricky to apply the modelling to my own laboratory data, but since this model involved impacts on ice and basalt it is at least possible to draw some conclusions about how the ice disruption results described in this thesis might extrapolate to larger size scales.

Equation 7.1 models the strength and gravity regimes as two separate, non-dependent halves of the equation. This is useful since the gravity regime is entirely unrelated to the material strength of an object. At extremely large size scales it does not matter if a body is a solid homogenous object or a rubble pile with no shear strength – the total strength of the body will be the same either way since the gravity regime will be overwhelmingly dominant. The consequence of this is that while the a term of equation governing the strength regime will need to be adjusted to reflect the laboratory data collected in this work, the b term governing the strength regime must

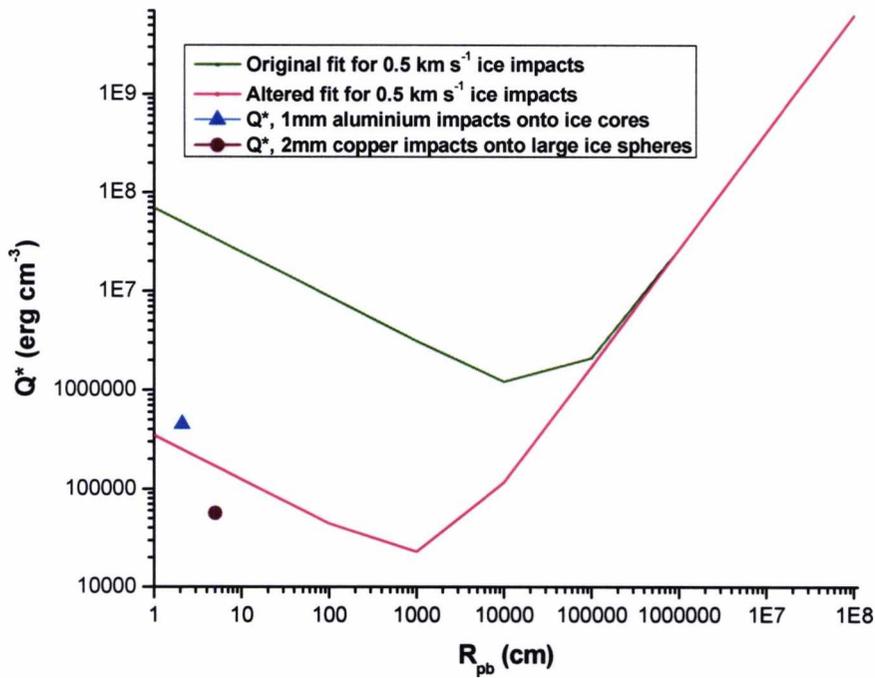


Figure 7.3.2 – Scaling fit for 0.5 km s⁻¹ impacts onto ice adjusted to correlate with observed laboratory results for ice.

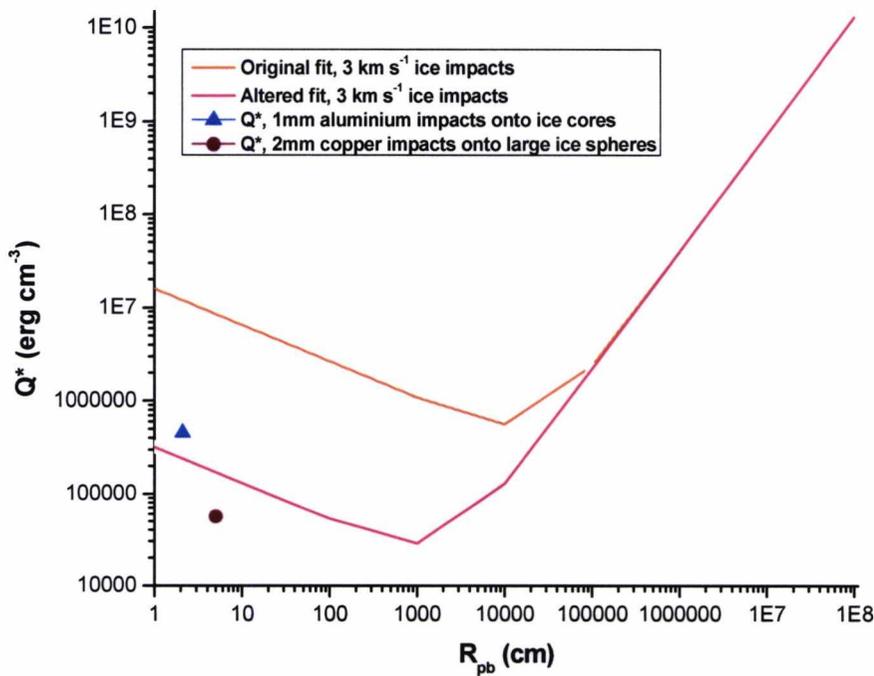


Figure 7.3.3 – Scaling fit for 3 km s⁻¹ impacts onto ice adjusted to correlate with observed laboratory results for ice.

remain as-is in order to reflect the non-variance of gravitational strength on material strength.

Benz & Asphaug's modelling for ice impacts at 0.5 and 3 km s⁻¹ is shown in figure 7.3.1. Q^* values are – for the sake of adherence to the model – expressed here in erg g⁻¹. These two fits predict a value for Q^* at laboratory size scales that is two or more orders of magnitude above that which was actually observed in laboratory targets composed of solid ice. In order to adjust these fits to make them compatible with the laboratory data, the strength regime for each will have to be scaled downwards.

Figure 7.3.2 shows the adjusted model for 0.5 km s⁻¹ impacts, with the fit now passing through a zone straddled by the laboratory results. The fits are very easy to change; all that is needed is to divide the Q^* values output by equation 7.1 by a scaling constant in order to adjust it up or down. The fit for 0.5 km s⁻¹ ice impacts has been divided by 200. Figure 7.3.3 shows 3 km s⁻¹ impacts having undergone the same adjustment, with the original fit divided by 50. As can be seen from the figures, the laboratory results reduce the fitted values of Q^* by approximately two orders in magnitude for bodies of up to 10 m radius. Between 10 m and 1 km radius the adjusted fit begins to converge with the original fit thanks to the increasing dominance of the gravity regime over material strength, with bodies larger than 1 km radius following the same fit line as the original fit.

These fits demonstrate several things. First, the laboratory results for Q^* for ice targets indicate that the strength of a body will reach a minimum at approximately 10 metres radius. This is the point where the material strength is weakest due to the increased size of the body, yet it is not yet large enough for the body's gravitational strength to have a significant effect on the disruption energy. Secondly, for bodies on the size scales commonly observed in large KBOs (on the order of tens to hundreds of kilometres) the material strength of the body is effectively irrelevant since it is dwarfed by the gravitational strength. This implies that, as far as scaling goes, laboratory impact experiments are only useful for fixing the lower end of the scale for bodies smaller than ~100 m radius. While this is helpful for modelling the critical

disruption energy of the majority of smaller KBOs, it is of little importance when determining Q^* for large icy KBOs such as Plutos.

Nevertheless, that laboratory experiments are not relevant when modelling Q^* for large bodies does not mean they have no utility at all in telling us *how* such bodies might disrupt. For example, laboratory impacts described in this thesis determined the disruption mechanism of layered targets where the ice layer spalls away from the core before any significant disruption of the core takes place.

There is a problem inherent in the scaling of laboratory impact results to solar system scales. Most models for predicting shock wave behaviour under impact assume an infinite target with no boundaries, through which an impact shock wave can travel unimpeded. In the case of bodies on the scale of large TNOs, this is largely correct; the body is so big that it is effectively infinite for the purposes of shock wave propagation unless the impacting body is also very large. The small scale of laboratory targets, however, do not enjoy such an easy approximation; the finite limits of the spherical targets used in the experimental programme mean that the impact shock front does not have far to travel before it reaches the boundary of the target, at which point things get complicated since this reflects the shock wave back into the target again. Therefore a small laboratory target will not have just one shock wave that spreads out and dissipates upon impact, but will instead focus the shockwave back into itself thanks to its finite size. Mizutani (1993) finds that this effect goes some way towards counterbalancing the assumed decrease in material strength with size; larger targets might have more flaws for a shockwave to exploit, but they also have more space for the impact shockwave to spread out and dissipate.

7.4 - Peak shock pressures

The methods used to calculate peak shock pressures during impact are described in detail in Melosh (1989). Melosh describes a planar impact approximation which assumes a projectile and a target which are both effectively infinite. Peak shock pressure calculations are reliant on knowing the projectile density, the target density, the projectile impact speed and the quantities C (dimension of velocity) and S

Tables 7.4.1 (left) and 7.4.2 (right) – Table 7.4.1 shows the C and S values used to calculate peak pressure, while table 7.4.2 shows the results of the peak pressure calculations themselves.

Material	Density (g cm ⁻³)	C (km s ⁻¹)	S	Projectile velocity (km s ⁻¹)	Peak pressure (Gpa)
Ice	2.952	1.28	1.58	1.00	2.26
Aluminium	0.997	5.3	1.37	2.00	6.29
				3.00	11.75
				4.00	18.54
				5.00	26.58
				6.00	35.86
				7.00	46.36

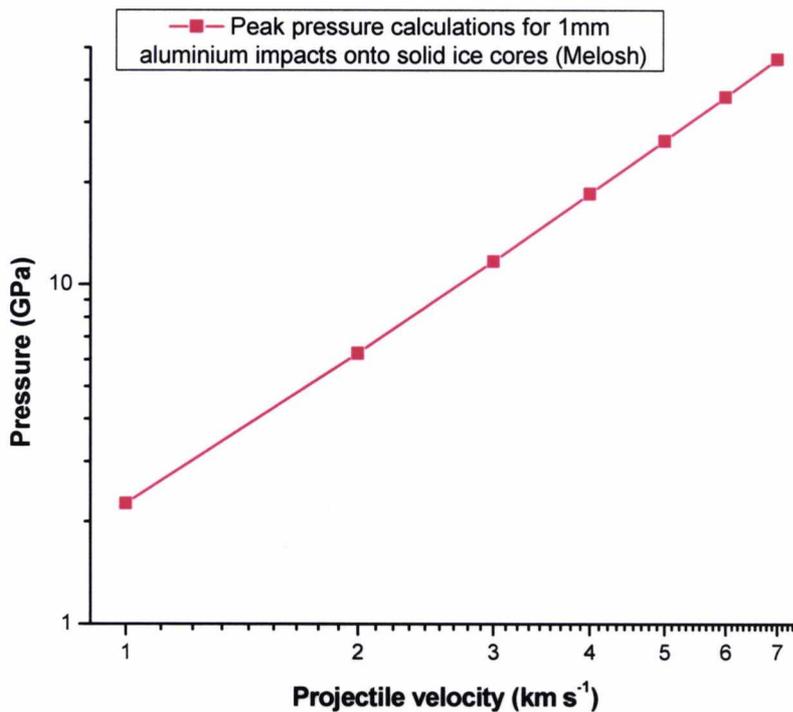


Figure 7.4.1 – Plot of projectile velocity versus peak pressures for 1mm aluminium impacts onto solid ice cores.

(dimensionless), which are predetermined parameters that describe the target and projectile materials. For the planar impact approximation the derivations in Melosh are followed exactly, so I shall not waste time deriving the method used to calculate peak shock pressures here. Projectile density, target density and impact speed are all known from laboratory measures while Melosh provides a limited number of C and S

Tables 7.4.3 (left) and 7.4.4 (right) – Table 7.4.1 shows the C and S values used to calculate peak pressure, while table 7.4.2 shows the results of the peak pressure calculations themselves.

Material	Density (g cm ⁻³)	C (km s ⁻¹)	S	Projectile velocity (km s ⁻¹)	Peak pressure (Gpa)
Permafrost	2	2.51	1.29	1.00	5.91
Steel (304)	8.02	4.56	1.5	2.00	14.72
				3.00	26.14
				4.00	40.07
				5.00	56.43
				6.00	75.19
				7.00	96.35

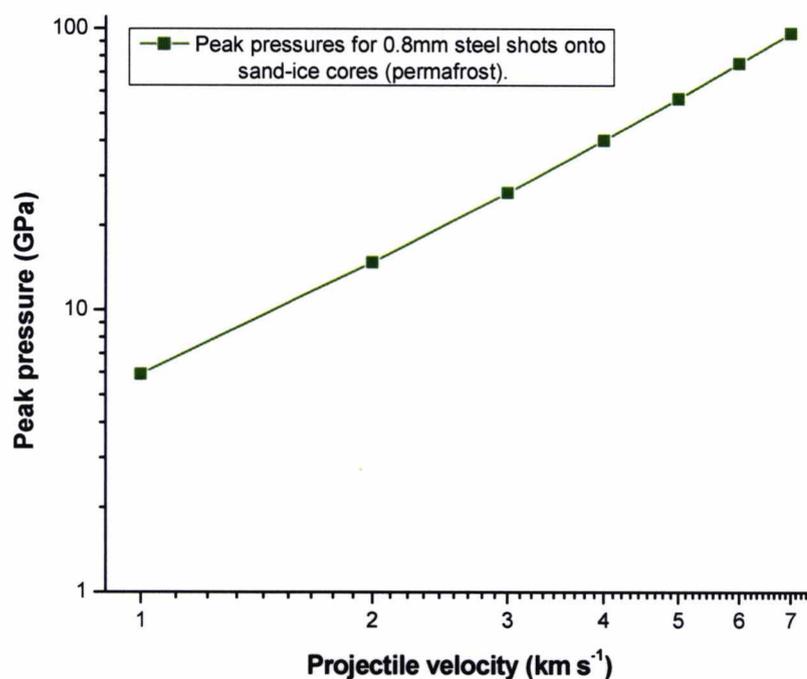


Figure 7.4.2 – Plot of peak pressure against projectile velocity for 0.8mm steel impacts onto sand-ice cores (approximated by permafrost).

values for various material types. Unfortunately no C/S values exist for the ice-silicate mix used in most of the experimental impacts, but sample peak shock pressures can be calculated for 1mm aluminium impacts onto ice cores.

Tables 7.4.1 and 7.4.2 present the appropriate C and S values and the results of the calculations using them, while figure 7.4.1 plots the peak pressures against projectile velocity. As can be seen from the graph it is a power law relationship. Projectile velocities in the laboratory impacts of 1mm aluminium projectiles onto ice cores ranged from $\sim 1 \text{ km s}^{-1}$ to 3 km s^{-1} , producing outcomes of very light cratering and total disruption respectively. This may seem odd in light of the fact that the measured compressive strength of ice was 1.6 MPa whereas here we are seeing theoretical peak shock pressures on the GPa scale, but the peak pressures are the shock pressures immediately after impact. As the shockwave travels through the target the shock pressure decreases according to an inverse power law (proportional to d^{-3} , where d is the distance travelled) so the pressure diminishes very rapidly; indeed, high initial peak pressures are required to make any significant impact on the target at all.

While no C/S values exist for an ice-silicate mix, I shall attempt a very crude approximation using Melosh's provided values for permafrost, which is the most analogous material from the examples available. The density of permafrost is highly variable but is approximately 2 g cm^{-3} . Tables 7.4.3 and 7.4.4 display the C/S values and the peak pressure values, while figure 7.4.2 plots projectile velocity against peak pressure. Laboratory impacts spanned the full range of velocities from 1 km s^{-1} to 7 km s^{-1} , and peak pressures therefore ranged from 5.91 GPa all the way up to 96 GPa.

Now, as has previously been stated a model which assumes the impact shock is propagating through an infinite body is not a very good approximation for small, laboratory scale targets. Mizutani (1993) attempts a model for such targets using the equation

$$P_o = \frac{1}{2} Z \rho_t C_t^2 \left(1 + \frac{1}{2} S_t Z \left(\frac{v}{C_t} \right) \right) \left(\frac{v}{C_t} \right) \quad (\text{Eq. 7.2})$$

Where P_o is the peak shock pressure, Z is the acoustic impedance ratio of the material, ρ_t , C_t and S_t are the density, C and S values for the target, and v is the projectile impact velocity. Z is calculated by

Tables 7.4.5 (left) and 7.4.6 (right) – Peak pressure calculations for 1mm aluminium impacts onto ice targets (left) and 0.8mm stainless steel impacts onto permafrost targets (right) using Mizutani's formula.

Projectile velocity (km s ⁻¹)	Peak pressure (Gpa)	Projectile velocity (km s ⁻¹)	Peak pressure (Gpa)
1.00	2.53	1.00	6.41
2.00	7.75	2.00	16.81
3.00	15.66	3.00	31.20
4.00	26.27	4.00	49.57
5.00	39.57	5.00	71.94
6.00	55.56	6.00	98.30
7.00	74.24	7.00	128.64

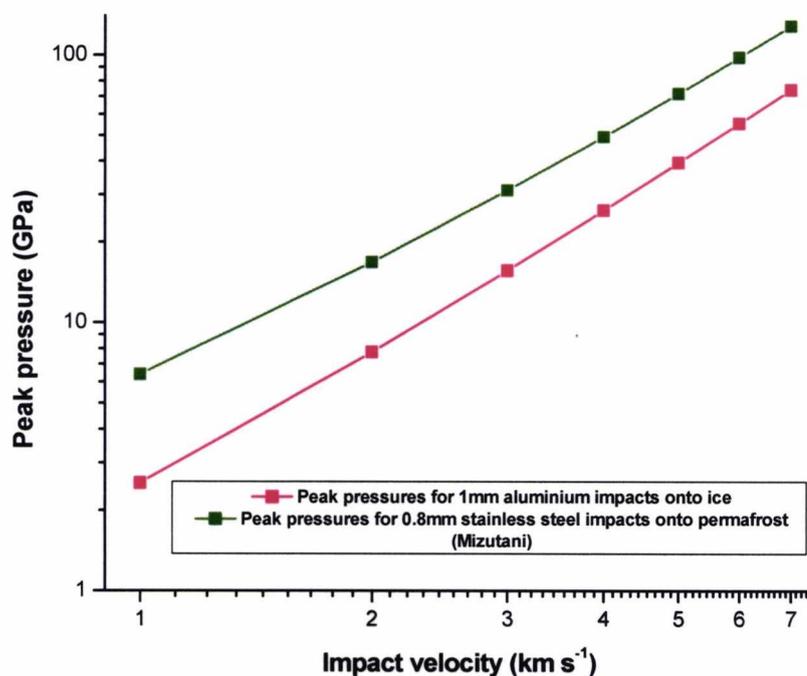


Figure 7.4.3 – Plot of peak pressures versus impact velocity for 1mm aluminium impacts onto ice targets and 0.8mm stainless steel impacts onto permafrost targets using Mizutani's formula.

$$Z = \frac{2}{\left(1 + \frac{\rho_i C_i}{\rho_p C_p}\right)} \quad (\text{Eq. 7.3})$$

Where ρ_p and C_p are the density and C values for the projectile. Using these two equations along with the C and S values used in the planar impact approximation method, Mizutani's model produces the following peak shock pressures for ice targets and permafrost targets.

Tables 7.4.5 and 7.4.6 show that Mizutani's formula (equation 7.2) for peak pressures produces peak pressures that are about 30% higher than those calculated using the planar impact approximation formula. 1mm aluminium impacts in the laboratory produced peak pressures ranging from 2.5 to 15.7 GPa, while 0.8mm stainless steel impacts ranged further from 6.4 to 128.6 GPa.

These peak pressures tell us about the shock pressure immediately after impact, and they tell us that it is very high – approaching or exceeding the pressure at which ice starts to vapourise under impact (5 GPa) in all the 1mm aluminium impacts, for example. If taken at face value these peak pressures mean that even relatively low velocity 1 km s^{-1} impacts should melt a large portion of the target mass. However, the fact that 1 km s^{-1} impact merely produced very small craters reflects the very fast rate of attenuation of the impact shock as it spreads out and dissipates into the target. The area immediately around the impact site would be vapourised or melted in accordance with the peak pressure, but the shock then attenuates at a rate determined by the diameter of the impacting projectile; a 1mm projectile's shock pressure will have dissipated to 10% of its original value after travelling just 1mm through a spherical target.

Mizutani (1993) models this shockwave attenuation by first assuming that the shockwave propagates through the impacted body in the manner shown in figure 7.4.4. The black area of the diagram is the area where peak shock pressure is unattenuated, with the shockwave then propagating through the target in all directions, creating an expanding sphere of shock. As the “surface” of the shock front

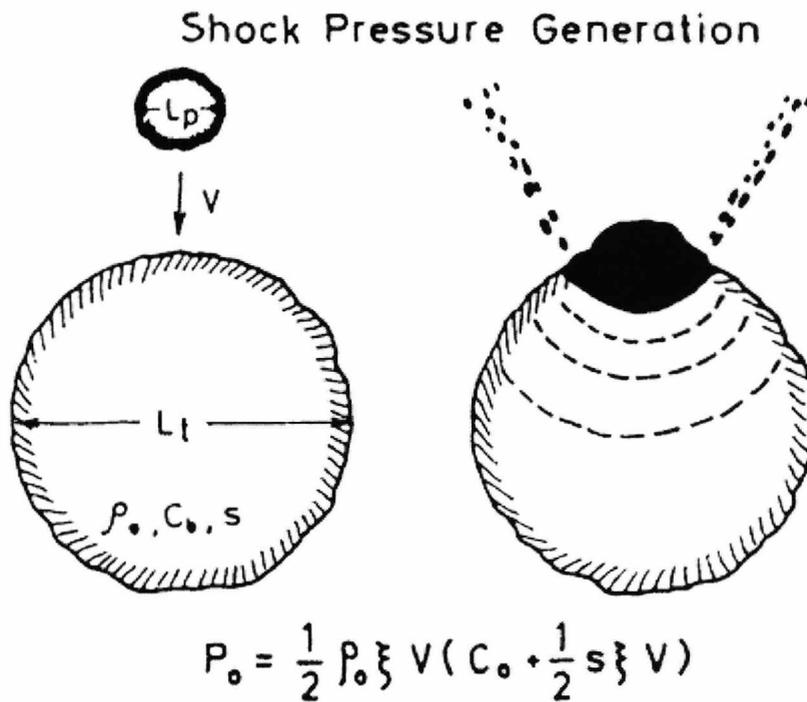


Figure 7.4.4 – Shockwave propagation through a spherical target, taken from Mizutani (1993). The shockwave radiates outwards from the impact site in all directions and thus dissipates according to an inverse power law.

increases, the shock pressure will decrease according to the inverse power law

$$P(r) = P_o \left(\frac{r}{L_p} \right)^{-b} \tag{Eq. 7.4}$$

Where $P(r)$ is the pressure of the shock after travelling a distance r through the target, P_o is the initial peak pressure, L_p is the diameter of the impacting projectile and b is a constant describing the rate of the decay of the shock pressure; this is stated to be ~ 3 for high shock pressures above 1 MPa, and ~ 2 below that. This equation indicates a dependence of the rate of shock attenuation on the size of the impacting projectile; a larger projectile will create a shock front that can travel further through the target before suffering significant attenuation.

Using equation 7.x, the rate of shock attenuation for 1 km s^{-1} , 2 km s^{-1} and 3 km s^{-1} impacts for a 1mm aluminium projectile onto ice has been calculated as the shock front propagates deeper into the target, and the results are shown in figure 7.4.5. In all cases the shock pressures dissipate extremely rapidly, reducing to 1% of its original

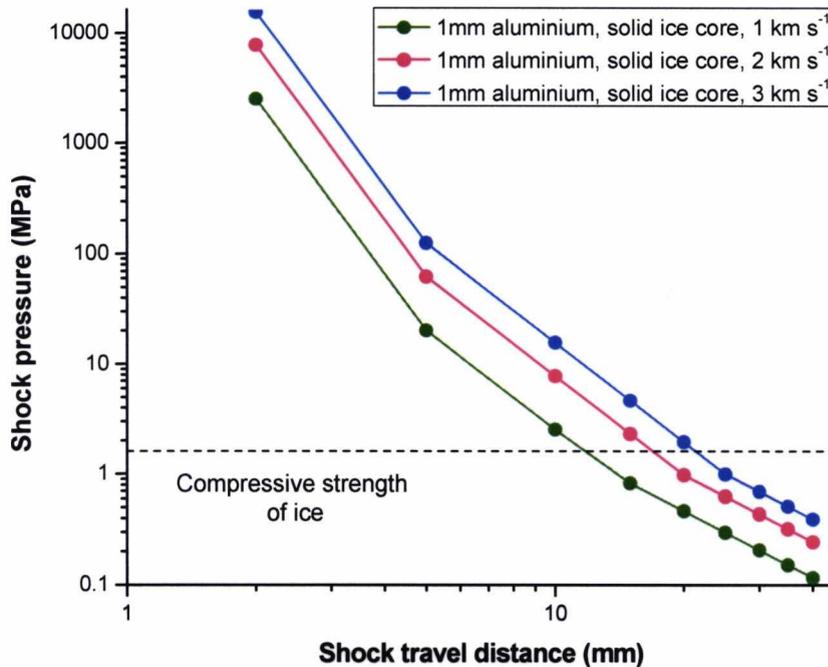


Figure 7.4.5 – Attenuation of shock pressure for 1mm aluminium impacts onto solid ice cores as the shock travels deeper into a target body.

value 10mm from the impact site. This is relevant when considering how an ice target fragments, since this will occur when the shock pressure exceeds the compressive strength of the material; a rarefaction wave follows behind the impact shockwave which further tests the tensile strength of the material. In chapter 4 the compressive strength of pure ice was found to be 1.62 MPa. Figure 7.4.5 tells us that the shock pressure will exceed this up to approximately 11mm for a 1 km s⁻¹ impact and 21mm for a 3 km s⁻¹ impact. However this cannot completely account for the observed experimental results; 1 km s⁻¹ impacts in the lab removed up to 8% of the target mass, which is somewhat consistent with this attenuation model, but 3 km s⁻¹ impacts almost completely pulverised targets leaving only 3-8% of the initial mass intact. The rarefaction wave may account for the extra mass removed due to the lower value of the tensile strength for ice (0.6 MPa), but it is not possible to adequately model the strength of the rarefaction wave so this remains speculative.

The spherical shape of propagation of the shock front through a spherical target has implications for the geometry of the shock waves in off-centre impacts. In chapter 3 it

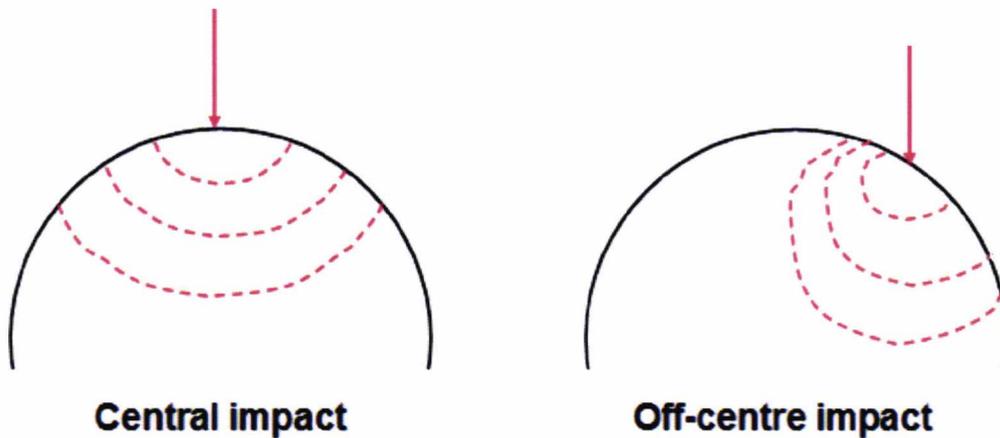


Figure 7.4.6 – Possible differences in shock wave geometry in a central impact event (left) and an off-centre impact event (right).

was stated that impacts could vary by up to ± 1 cm off-centre; while this was stated to have little overall effect on the mass of the largest fragment (and thus Q^*) an impact 1 cm off centre on a 4 cm diameter target will significantly affect how the shock front propagates, and could thus alter the mechanism of disruption of a spherical body, even if the mass of the largest fragment is largely unaffected. This is illustrated in figure 7.4.6; the shock front in the off-centre impact is propagating asymmetrically compared to the central impact, a phenomenon that will be exacerbated whenever the shock front encounters the physical boundary of the material.

7.5 – The size distribution of the Haumea collisional family

While no size distribution data is available for the layered bodies impacted in the experimental programme, it is still useful to compare the observed size distributions for the homogenous bodies with the size distribution of the Haumea collisional family in order to get a general picture of the impact that created it. Unfortunately size distribution data is not directly available for the majority of objects in the Haumea collisional family; while the masses of Haumea and its moons Hi'iaka and Na'maka have been accurately measured (Ragozzine & Brown, 2009) all that is available for the rest of the family is a collection of absolute magnitudes (JPL Small Body database browser). It is therefore necessary to make some assumptions in order to derive approximate mass values from the absolute magnitudes: first, that the bodies in the

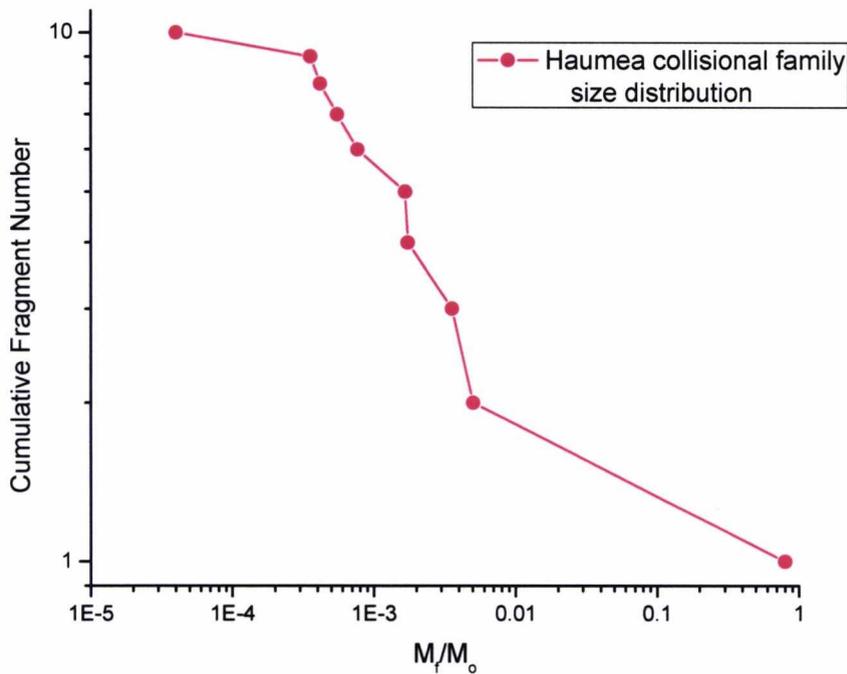


Figure 7.5.1 – Size distribution of the Haumea collisional family.

Haumea collisional family will have albedos close to that of Haumea itself (0.7), and secondly that they are composed almost entirely of pure ice with a density of 0.997 g cm⁻³. This first allows calculation of the diameter of the bodies from

$$D = \frac{1329 \times 10^{-0.2H}}{\sqrt{A}} \quad (\text{Eq. 7.5})$$

Where D is the diameter of the body, H is the absolute magnitude and A is the albedo.

Then it is simple to determine the volume of the body and thus the mass using the density for pure ice. This is the same method used to determine the size distribution of the Haumea collisional family in most literature e.g. Leinhardt *et al.* 2010. The Haumea parent body is thought to have lost roughly 20% of its mass in the collision that created the Haumea family (Brown *et al.* 2007); therefore the mass used for the parent body M_0 in determining M_f/M_0 is 1.25 times the mass of Haumea.

The majority of the Haumea collisional fragments are - predictably - several orders of magnitude smaller than Haumea itself; comparing the shape of this distribution to the ones observed in Chapter 6, the massive largest fragment (Haumea) is followed by a collection of similarly-sized tiny fragments, giving the distribution an extremely concave shape. This is consistent with the distribution curves for low velocity, low energy impacts that do not exceed the critical energy density required to disrupt a body, and is therefore consistent with the general theory of how Haumea was formed – as the result of an impact that was energetic enough to remove the surface layer of the parent body (and a significant percentage of mass), but which left the majority of the mass locked up in the higher-density core material relatively intact.

More generally, the density difference between Haumea and its collisional fragments mirrors what was observed during the disruption of layered targets in the laboratory – the ice shell fractured and spalled away almost entirely before any significant damage was done to the core, proving that it is entirely possible for an impact to remove low-density material covering the surface of an object while leaving higher density core material relatively intact.

While the above analysis is valid for a collisional family created by a single impact event, there is reason to believe that the impact event that created Haumea is not quite as simple as it first appears. The results seen for the layered targets indicates that it is perfectly possible for an impact to strip away a surface layer of ice while leaving a denser core mostly intact, but the velocity dispersions of both the resulting fragments and Haumea indicate that the actual event was a little more complex than this (Schlichting & Sari, 2009). To explain the current observed orbital path of Haumea requires a dispersion velocity of 400 m s^{-1} . The observed dispersion of velocity of the fragments making up Haumea's collisional family is 160 m s^{-1} , which is consistent with their current observed positions. A dispersion velocity of 400 m s^{-1} would put them well beyond those positions and is thus incompatible with the collisional family, so something must have happened to either Haumea or the fragments after the original impact stripping the ice layer in order to explain the discrepancies in the velocities.

Theories to explain this involve either an event which changes Haumea's velocity dispersion – a subsequent non-disruptive impact providing a velocity change, or else gravitational interactions with the outer gas giants somehow boosting Haumea's velocity – or an event which changes the velocity dispersion of the fragments. To change the velocity dispersion of all the fragments at once requires some creative thinking, but it is possible; to do so would require the fragments to form themselves into a single moon of Haumea after impact. This moon is then struck by a second impact which creates the fragment family with the current observed dispersion velocity. This second possibility has a small problem in that Haumea's escape velocity is 900 m s^{-1} ; the observed fragment dispersion velocity of 160 m s^{-1} is too low for fragments created from a moon which remained *in situ* where it formed since those fragments would not then be able to escape Haumea's gravity. Therefore some way of gradually moving this ice moon away from Haumea is required so that it is impacted at a sufficient distance that the escape velocity is much lower – an unstable or resonant orbit would fit the bill, since this could cause the orbit of the moon to slowly tidally evolve outwards.

Either scenario involves a fairly complex sequence of events to explain the observed properties of the Haumea collisional family, and simulations show that there is a 47% chance of a Haumea scenario happening over the lifetime of the solar system if the collision happened in the shattered disc (Levison *et al.* 2008). If the collisions are restricted to the KBO itself then the increased stability of the orbits involved drops to just 0.1%. This makes it likely that Haumea is the only example of a collisional TNO family of its type we are likely to find in the TNO region.

If Haumea was indeed created by a two-step impact process, then the size distribution of the fragments will be unrelated to the original impact which separated them from Haumea, and should be examined separately to Haumea itself. Unfortunately this is impossible to do in a qualitative fashion without knowledge of the likely size of the ice moon that formed the fragments to derive accurate M_f/M_o values. Nevertheless, M_f/M_o is ultimately just a scaling quantity to quantify fragment mass measurements in terms of the original body's mass, and the overall spread of the fragment masses

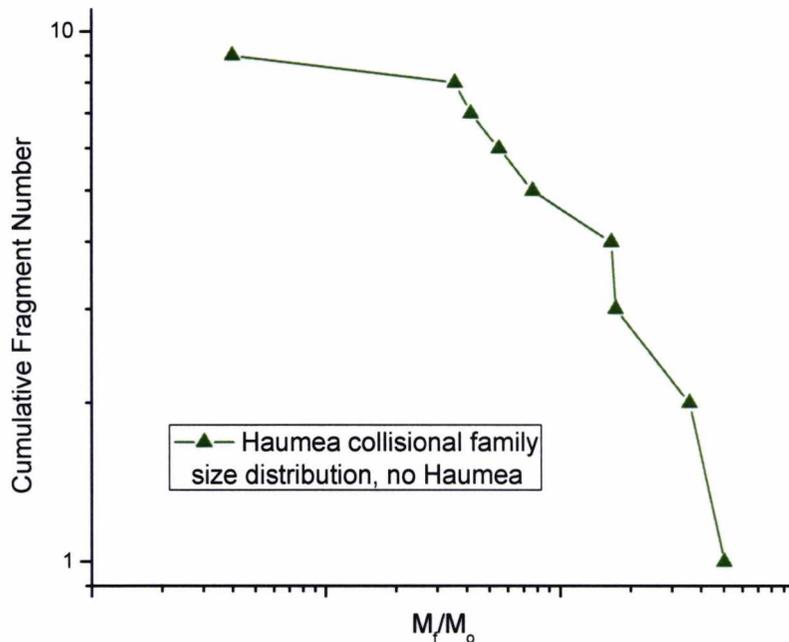


Figure 7.5.2 – Size distribution of the Haumea collisional family with the Haumea mass point removed. This attempts to model the size distribution of the fragments from the ice moon which is thought to have been their source. The M_f/M_0 axis is scaleless because the original mass of the moon is not known.

should remain unchanged no matter how large the original body which formed them was.

With the Haumea mass point removed from the plot the general shape of the size distribution of the collisional family changes as seen in figure 7.5.2; it is no longer concave but is instead now slightly convex, pointing to a disruptive impact event with a Q that was somewhat larger than Q^* ($Q > Q^*$). This changed plot agrees with the changed theory on how the fragments were formed.

7.5 – Conclusion

In this chapter I have discussed possible reasons that may explain both the different mode of disruption observed in ice and some other targets at high energy densities, and the dramatic increase in critical energy density when a core is covered with a shell

composing about 40% of its mass. It is very likely that the latter effect is dependent on the specific way the target is constructed; in the targets used here the ice layer is frozen onto the core, so it is only weakly coupled to the rest of the target. It is therefore relatively easy to remove the shell from the core, causing energy to be focused in the shell rather than the core in a number of ways. As long as the core still makes up more than 50% of the total target mass the layered target will not be disrupted according to the traditional definition of disruption, even though the shell may be shattered and pulverised. However, it is unlikely that the ice layers of the big TNOs are weakly coupled to their parent bodies in the same way as the shells here; the thermal differentiation process would result in a steadily increasing percentage of silicate content encased in solid, monolithic ice rather than a discrete two-step body. Thus the effect seen here, while interesting, may not reflect the actual behaviour of layered solar system bodies when impacted.

Additionally, modelling work has been undertaken both to test the propagation of impact shock through a spherical layered target as well as sample peak shock pressures for solid ice and sand-ice cores, and scale up laboratory results for Q^* to solar system scales, as well as comparing the actual size distribution of the Haumea collisional family with the general size distribution trends observed in Chapter 6. Laboratory results determining a value of Q^* that is dominated by a body's material strength do not have an effect on scaling bodies up past 1 km radius, while the size distribution trends of the Haumea family are consistent with the impact events thought to have formed it.

Conclusion

In this work I have investigated how adding a layer of ice to a body affects the amount of energy required to disrupt it. In order to do this several experimental hurdles have been overcome: I have developed methods for constructing spherical sand-ice targets that reflect the general composition of some of the larger icy bodies found in the solar system, for adding relatively uniform solid ice layers to these spherical targets, and for subjecting both the layered and unlayered targets to hypervelocity impacts at velocities of between 1 and 7 km s⁻¹ while maximising the number of fragments recovered.

Secondary goals of this work were to investigate other factors that might affect the critical energy density required to disrupt a body Q^* , and to examine the size distribution of target fragments after impact to see if they varied at all with various combinations of impact energy, target type and projectile type. To this end several other target types – such as solid ice spheres and sand-ice cylinders - were manufactured using variations on existing methods and all target types were then impacted with a number of different projectile types that varied in both composition and size. The largest remaining fragments were collected and measured along with other fragments if they were still relatively intact, and Q^* values have been calculated for all major target types.

The results of this work lead to three main conclusions. The first is that the critical energy density Q^* is not solely a function of impact energy. Instead, it shifts up or down depending on several external factors: the ones positively identified here are the size of the impacting body and the size of the target. While target size has long been known to affect the required critical energy density (Durda 1999) the properties of the impacting body are not often taken into account in previous studies of Q^* . Therefore it is fair to say that these previous studies are not providing an accurate assessment of what the disruption strength of a particular material actually is.

Second, the results for both the sand-ice spherical targets and the layered targets point to some interesting effects on Q^* . Most experimental and modelling work assumes

that the icy solar system bodies being emulated are made up of ice only, or ice with a small fraction of silicate content (Arakawa 2004). However, since the observed densities of these bodies are significantly higher than they would be if composed of pure ice, it follows that they possess a sizeable amount of silicate content – for example, Pluto’s observed density of 2.03 g cm^{-3} (Buie *et al.* 2006) corresponds to a silicate content of up to 80%. The ice-silicate targets used in this work have a similar silicate content and density, and display critical energy densities of $270\text{-}309 \text{ J kg}^{-1}$ depending on the impacting body. Pure ice targets of the same size display a critical energy density of $45 \pm 3 \text{ J kg}^{-1}$. Therefore ice-silicate targets require approximately 5-6 times more energy to disrupt than pure ice targets, and possibly even more when the lower density of the projectiles used to impact the ice targets is taken into account.

Adding a layer of ice to the sand-ice target has an effect on Q^* that varies depending on the thickness of the ice layer. It is unclear whether this variation is dependent to a ratio of the layer thickness to the diameter of the impacting projectile or a ratio of the mass of the ice layer to the mass of the impacting projectile, but in order to quantify the effect it is described in terms of percentage of target mass. Thin ice layers with masses of up to 20-25% of the total target mass produce little or no increase in the critical energy density, with the critical energy densities of a set of targets with layer thicknesses in this range producing Q^* values in a range of $270\text{-}360 \text{ J kg}^{-1}$. However, once an ice layer passes about 30-35% of the total target mass the energy density required to disrupt it jumps significantly to $912 (+20/-19) \text{ J kg}^{-1}$.

These results indicate that icy solar system bodies consisting of an ice-silicate core covered by a layer of pure ice may be far harder to disrupt than previously thought. Of course while the targets used here have attempted to model the suspected makeup of these bodies more accurately they are by no means perfect. In particular the layered targets exhibit behaviour whereby the weak bond between the ice layer and the sand-ice core results in very easy fragmentation and spalling of the ice layer away from the core, which is shown in hydrocode simulations to absorb a lot of shock energy in these spalled fragments. In reality a differentiated layered body would not be a simple two material body as used here, but would rather be composed of pure ice at the surface with a steadily increasing percentage of silicate content deeper into the body.

While previous attempts to determine Q^* for icy bodies almost certainly underestimate the true value when applied to large icy TNOs, the value calculated here likely overestimates it by some unknown degree.

Still, a direct comparison between previous critical energy densities calculated for ice (on the order of $4 - 40 \text{ J kg}^{-1}$) (Murray 2004, Arakawa 1999) and the critical energy densities for the ice-silicate mix calculated here indicates that - whatever the structure of the target - modelling outer solar system bodies as being made of pure ice only is not providing an accurate value of Q^* for the larger, denser ones.

Finally, both the size distribution plots and the energy density plots indicate the existence of three separate impact regimes reflecting three different types of impact behaviour - cratering, disruption and shattering. In light of this, previous calculations of Q^* (Arakawa 1999, Ryan *et al.* 1999, Giblin *et al.* 2004) are too simplistic since they only assume a disruption regime; while the existence of the cratering regime is implicit in the way the disruption regime is modelled, the shattering regime is not, only revealing itself if a body is impacted at both medium and high energy densities. Much experimental work involves impacts over a small range of energy densities for a certain material and then assuming that this range falls within the disruption regime. This work shows that it could instead fall in the transition zone or within the shattering regime, which would make any value of Q^* derived from the data inaccurate to some degree.

Future work modelling the behaviour of large icy solar system bodies must take the three main conclusions described here into account: first, that Q^* is not a fixed quantity for a given material but is instead dependent on the impacting body as well; second, that ice-silicate bodies are harder to disrupt than pure ice bodies; and third, that calculations of Q^* that only use data gathered at high values of Q - in other words, within the shattering regime - are suspect since they may not accurately reflect the disruptive behaviour of a body.

There is much scope to improve upon this work. First and foremost would be a more detailed investigation of the effect adding an ice layer has on Q^* . It would be easy to

determine if the increase in Q^* was caused by layer thickness or layer mass by using differently-sized cores; a 5mm ice layer covering a 40mm diameter core impacted by a 1mm projectile would have the same layer thickness/projectile diameter ratio as a 5mm ice layer covering a 60mm core impacted by a 1mm projectile, but the mass of the ice layer would be greater. An improved attempt to isolate the factors upon which Q^* is dependent could also be made – if the size of the targets or the size of the projectiles could be varied at will it would become possible to fine-tune a projectile/target combination so that it tested one factor and one factor only rather than a change in projectile size forcing a change in projectile density as well to achieve the same range of Q values. Finally, it would be useful to image an impact on a layered body using an ultra-high-speed camera in order to see what happens to the layer during disruption, and the manner in which it shears away from the target completely at particular energy densities.

Bibliography

- A'Hearn, M.F. and 34 others, (2005). Deep Impact: Excavating comet Tempel 1. *Science*, **310**, 258-264.
- Abe, S., and 15 others, (2006). Mass and local topography measurements of Itokawa by Hayabusa. *Science*, **312**, 1334-1347.
- Arakawa, M., (1999). Collisional Disruption of Ice by High-Velocity Impact. *Icarus*, **142**, 34-45.
- Arakawa, M., Leliwa-Kopystynski, J., Maeno, N., (2002). Impact Experiments on Porous Icy-Silicate Cylindrical Blocks and the Implication for Disruption and Accumulation of Small Icy Bodies. *Icarus*, **158**, 516-531.
- Arakawa, M., Tomizuka, D., (2004). Ice-silicate fractionation among icy bodies due to the difference of impact strength between ice and ice-silicate mixture. *Icarus*, **170**, 193-201.
- Asphaug, E., Benz, W., (1994). Density of comet Shoemaker-Levy 9 deduced by modelling breakup of the parent 'rubble pile'. *Nature*, **370**, 120-124.
- ASTM (1993). Standard test method for splitting tensile strength of cylindrical concrete specimens. (C496-90) 1993 *Annual book of ASTM standards*, Sect. **4**, Vol. **04.02**, 269-272.
- Bernstein, G.M., Trilling, D.E., Allen, R.L., Brown, M.E., Holman, M., Malhotra, R., (2004). The size distribution of trans-Neptunian bodies. *Astronomical Journal*, **128**, 1364-1390.
- Bless, S.J., Rajendran, A.M., (1996). Initiation and Propagation of Damage Caused by Impact. *High Pressure Shock Compression of Solids II*, (Eds. Davison L., Grady D., Shahinpoor M.), Springer-Verlag New York, pp. 194-218.

Brown, M.E., (2001). The inclination distribution of the Kuiper belt. *Astronomical Journal*, **121**, 2804-2814.

Brown, M.E., (2005). S/2005 (2003 UB313) 1. *Central Bureau Electronic Telegrams*. **240**, 1.

Brown, M.E., Trujillo, C.A., Rabinowitz, D.L., (2005). Discovery of a planetary-sized object in the scattered Kuiper Belt. *Astrophysical Journal*. **635**, 97-100.

Brown, M.E., Schaller, E.L., Roe, H.G., Rabinowitz, D.L., Trujillo, C.A., (2006). Direct measurement of the size of 2003 UB313 from the Hubble Space Telescope. *Astrophysical Journal*, **643**, 61-63.

Brown, M.E., Barkume, K.M, Ragozzine, D., Schaller, E.L. (2007). A collisional family of icy objects in the Kuiper belt. *Nature*, **446**, 294-296.

Buie, M.W., Grundy, W.M., Young, E.F, Young, L.A, Stern, S.A. (2006). Orbits and photometry of Pluto's satellites: Charon, S/2005 P1, and S/2005 P2. *Astronomical Journal*, **132**, 290-298.

Buie, M.W., Grundy, W.M., Young, E.F., Young, L.A., Stern, S.A., (2010). Pluto and Charon with the Hubble Space Telescope. II. Resolving changes on Pluto's surface and a map for Charon. *Astronomical Journal*, **139**, 1128-1143

Burchell, M.J., Cole, M.J., McDonnell, J.A.M, Zarnecki, J.C. (1999). Hypervelocity impact studies using the 2 MV Van de Graff accelerator and the two-stage light gas gun of the University of Kent at Canterbury. *Measurement Science and Technology*, **10**, 41-50.

Carlson, R.W., Drossart, P., Encrenaz, T., Weissman, P.R., Hui, J., Segura, M., (1997). Temperature, Size, and Energy of the Shoemaker-Levy 9 G-Impact Fireball. *Icarus*, **128**, 251-274.

Carr, M.H. and 21 others (1998). Evidence for a subsurface ocean on Europa. *Nature*, **391**, 363-365.

Chambers, J.E., (2004). Planetary accretion in the inner Solar System. *Earth and Planetary Science Letters*, **223**, 241-252.

Charnoz, S., Morbidelli, A. (2007). Coupling dynamical and collisional evolution of small bodies II. Forming the Kuiper belt, the Scattered Disk and the Oort Cloud. *Icarus*, **188**, 468-480.

Cheng, A.F., (2004). Implications of the NEAR mission for internal structure of Mathilde and Eros. *Advances in Space Research*, **33**, 1558-1563.

Claesson, J., Bohlooli, B., (2002). Brazilian test: stress field and tensile strength of anisotropic rocks using an analytical solution. *International Journal of Rock Mechanics and Mining Sciences*, **39**, 991-1004.

Crawford, D.A, Remo, J.L., (1997). Near-Earth Objects – United Nations International Conference. *Annals of the New York Academy of Sciences*, **822**, 155-173.

Crozier, W.D., Hume, W., (1957), High-velocity, light gas gun. *Journal of Applied Physics*, **28**, 892-894.

Dawe, W., Murray, M., Burchell, M.J., (2005). Catastrophic disruption of porous and solid ice bodies. Lunar and Planetary Science Conference 36 abstract 1096, available at <http://www.lpi.usra.edu/meetings/lpsc2005/pdf/1096.pdf>

De Sanctis, M.C., Capria, M.T., Coradini, A. (2001). Thermal evolution and differentiation of Edgeworth-Kuiper belt objects. *Astronomical Journal*, **121**, 2792-2799.

Duncan, M., Quinn, T., Tremaine, S., (1987). The formation and extent of the solar system comet cloud. *Astronomical Journal*, **94**, 1330-1338

Durda, D.D., Greenberg, R., Jedicke, R., (1999). Collisional models and scaling laws: A new interpretation of the shape of the main-belt asteroid size distribution. *Icarus*, **135**, 431-440.

Durda, D.D., Bottke, W.F., Nesvorny, D., Enke, B.L., Merline, W.J., Asphaug, E., Richardson, D.C., (2007). Size-frequency distributions of fragments from SPH/N-body simulations of asteroid impacts: Comparison with observed asteroid families. *Icarus*, **186**, 498-516.

Delsanti, A., Jewitt, D. (2006) The Solar System Beyond The Planets. *Solar System Update*, Pub. Springer-Verlag, 267-293.

Edgeworth, K.E., (1943). The Evolution of Our Planetary System. *Journal of the British Astronomical Association*, **53**, 181-188.

Elliot, J.L., Kern, S.D., Clancy, K.B., Gulbis, A.A.S., Millis, R.L., Buie, M.W., Wasserman, L.H., Chiang, E.I., Jordan, A.B., Trilling, D.E, Meech, K.J., (2005). The Deep Ecliptic Survey: A search for Kuiper Belt Objects and Centaurs. II. Dynamical classification, the Kuiper Belt plane, and the core population. *Astronomical Journal*, **129**, 1117-1162.

Giblin, I., Davis, D.R., Ryan, E.V., (2004). On the collisional disruption of porous icy targets simulating Kuiper belt objects. *Icarus*, **171**, 487-505.

Grey, I.D.S., Burchell, M.J., (2003). Hypervelocity impact cratering on water ice targets at temperatures ranging from 100 K to 253 K. *Journal of Geophysical Research – Planets*, **108**, 5019.

Guo H. (1993) Rock Fracture-Toughness Determination by the Brazilian Test. *Engineering Geology*, **33**, 177

Fujiwara, A., and 21 others, (2006). The Rubble-Pile Asteroid Itokawa as Observed by Hayabusa. *Science*, **312**, 1330-1334.

Hills, J.G. (1981). Comet showers and the steady-state infall of comets from the Oort cloud. *Astronomical Journal*, **86**, 1730-1741.

Hiraoka, K., Arakawa, M., Setoh, M., Nakamura, A.M., (2008). Measurements of target compressive and tensile strength for application to impact cratering on ice-silicate mixtures. *Journal of Geophysical Research*, **113**,

Hobbs, P.V., (1974). *Ice Physics*. Clarendon Press, Oxford, pp. 330-333.

Husmann, H., Sohl, F., Spohn, T. (2006). Subsurface oceans and deep interiors of medium-sized outer-planet satellites and large trans-Neptunian objects. *Icarus*, **185**, 258-273.

Jewitt, D.C., Luu, J.X., (1993). Discovery of the Candidate Kuiper Belt Object 1992QB1. *Nature*. **362**, 61-63.

Kavelaars, J.J., Jones, L., Gladman, B., Parker, J.M., Petit, J-M., (2008). The Orbital and Spatial Distribution of the Kuiper Belt. *The Solar System beyond Neptune*, U. Arizona Press.

Kuiper, G.P., (1951). On the origin of the Solar System. *Astrophysics*, McGraw-Hill, New York. 357-424.

Lacerda, P., Jewitt, D., Peixinho, N., (2008). High-precision photometry of extreme KBO 2003 EL₆₁. *Astronomical Journal*, **135**, 1749-1756.

Lange, M.A., Ahrens, T.J., (1983). The dynamic tensile strength of ice and ice-silicate mixtures. *Journal of Geophysical Research*, **88**, 1197-1208.

Leliwa-Kopystynski, J., Burchell, M.J., Wlodarczyk, I., (2009), The Impact origin of Eunomia and Themis families. *Meteoritics & Planetary Science*, **44**, preprint.

Levison, H.F., Morbidelli, A., (2005). The formation of the Kuiper belt by the outward transport of bodies during Neptune's migration. *Nature*, **426**, 219-221.

Levison, H.F., Morbidelli, A., VanLaerhoven, C., Gomes, R., Tsiganis, K., (2008). Origin of the structure of the Kuiper belt during a dynamical instability in the orbits of Uranus and Neptune. *Icarus*, **196**, 258-273.

Levison, H.F., Morbidelli, A., Vokrouhlicky, D., Bottke, W.F., (2008). On a scattered-disk origin for the 2003 EL₆₁ collisional family – an example of the importance of collisions on the dynamics of small bodies. *Astrophysical Journal*, **136**, 1079-1088.

Lightwing, A., Burchell, M.J., (2006). Catastrophic disruption of mixed ice:sand bodies. Lunar and Planetary Science Conference 37 abstract 1565, available at <http://www.lpi.usra.edu/meetings/lpsc2006/pdf/1565.pdf>

Love, S.G., Horz, F., Brownlee, D.E., (1993). Target porosity effects in impact cratering and collisional disruption. *Icarus*, **105**, 216-224.

Meech, K.J, and 208 others, (2005). Deep Impact: Observations from a worldwide Earth-based campaign. *Science*, **310**, 265-269.

Martin, T.Z., Orton, G.S, Travis, L.D, Tamppari, L.K., Claypool, I., (1995). Observation of Shoemaker-Levy impacts by the Galileo Photopolarimeter Radiometer. *Science*, **268**, 1875-1879.

Martin, T.Z., Orton, G.S., (1997)., Shoemaker-Levy 9: distribution of radiant energy. *Planetary and Space Science*, **45**, 1287-1298.

- Melosh J. (1989). Impact Cratering: A Geologic Process. Oxford Monographs on Geology and Geophysics #11, pp. 37-57, 232-233
- Michel, P., Benz, W., Tanga, P., Richardson, D.C., (2001). Collisions and Gravitational Reaccumulation : Forming Asteroid Families and Satellites. *Science*, **294**, 1696-1700.
- Mizutani H (1993). Scaling law on impact phenomena. *Primitive Solar Nebula and Origin of Planets* pp. 297-317.
- Morbidelli, A., Levison, H.F., Tsiganis, K., Gomes, R. (2005). Chaotic capture of Jupiter's Trojan asteroids in the early Solar System. *Nature*, **435**, 462-465.
- Morbidelli, A., (2007). Portrait of a suburban family. *Nature*, **446**, 273-274.
- Moritoh, S., Kawai, N., Nakamura, K.G., Kondo, .K., (2001). Optimization of a compact two-stage light-gas gun aiming at a velocity of 9 km/s. *Review of Scientific Instruments*, **72**, 4270-4272
- Murray, M., (2006). Catastrophic disruption of solid ice bodies and Pluto by impactors at normal and oblique incidence. Undergraduate dissertation, School of Physical Sciences, University of Kent.
- Nagel, K., Breuer, D., Spohn, T., (2004). A model for the interior structure, evolution and differentiation of Callisto. *Icarus*, **169**, 402-412.
- O'Brien, D.P., Greenberg, R., (2005). The collisional and dynamical evolution of the main-belt and NEA size distributions. *Icarus*, **178**, 179-212.
- Okamoto, C., Arakawa, M., (2008). Experimental study on the impact fragmentation of core-mantle bodies: Implications for collisional disruption of rocky planetesimals with sintered core covered with porous mantle. *Icarus*, **197**, 627-637.

Oort, J. H., (1950). The structure of the cloud of comets surrounding the Solar System and a hypothesis concerning its origin. *Bulletins of the Astronomical Institute of the Netherlands*, **11**, 91-110

Öpik, E., (1932). Note on Stellar Perturbations of Nearby Parabolic Orbits. *Proceedings of the American Academy of Arts and Sciences*, **67**, 169-182

Rabinowitz, D.L., Barkume, K., Brown, M.E., Roe, H., Schwartz, M., Toutellotte, S., Trujillo, C. (2006). Photometric observations constraining the size, shape and albedo of 2003 EL61, a rapidly rotating, Pluto-sized object in the Kuiper belt. *Astrophysical Journal*, **639**, 1238-1251.

Ragozzine D., Brown M.E. (2009) Orbits And Masses Of The Satellites Of The Dwarf Planet Haumea = 2003 EL61. *Astronomical Journal*, **137**, 4766-4776.

Ryan, E.V., Davis, D.R., GIBLIN, I. (1999). A Laboratory Impact Study of Simulated Edgeworth-Kuiper Belt Objects. *Icarus*, **142**, 56-62.

Ryan, E.V., (2000). Asteroid Fragmentation and Evolution of Asteroids. *Annual Review of Earth and Planetary Sciences*, **28**, 367-89.

Schubert, G., Anderson, J.D., Travis, B.J., Palguta, J., (2007). Enceladaus: Present internal structure and differentiation by early and long-term radiogenic heating. *Icarus*, **188**, 345-355

Schultz, P.H., Ernst, C.M., Anderson, J.L.B., (2005). Expectations for crater size and photometric evolution from the deep impact collision. *Space Science Review*, **117**, 137-160

Showman, A.P., Malhotra, R., (1999). The Galilean Satellites. *Science*, **286**, 77-84.

Stansberry, J., Grundy, W., Brown, M., Cruikshank, D., Spencer, J., Trilling, D., Margot, J-L., (2008). Physical Properties of Kuiper Belt and Centaur Objects: Constraints from Spitzer Space Telescope. *The Solar System beyond Neptune*, U. Arizona Press.

Sugita, S., and 22 others, (2005). Subaru telescope observations of Deep Impact. *Science*, **310**, 274-278.

Tanga, P., Cellino, A., Michel, P., Zappala, V., Paolicchi, P., Dell'Oro, A., (1999). On the Size Distribution of Asteroid Families : The Role of Geometry. *Icarus*, **141**, 65-78.

Trujillo, C.A., Brown, M.E. (2002). A correlation between inclination and color in the classical Kuiper belt. *Astrophysical Journal*, **566**, 125-128.

Trujillo, C.A., Brown, M.E. (2001). The radial distribution of the Kuiper belt. *Astrophysical Journal*, **554**, 95-98.

Tsiganis, K., Gomes, R., Morbidelli, A., Levison, H.F. (2005). Origin of the orbital architecture of the giant planets of the Solar System. *Nature*, **435**, 459-461

Vardar O., Finnie I. (1975). Analysis of Brazilian Disk Fracture Test using Weibull Probabilistic Treatment of Brittle Strength. *International Journal of Fracture*, **11**, 495-508.

Whipple, F.J., (1950). A comet model. I. The acceleration of comet Enke. *Astrophysical Journal*, **111**, 375-394.

Zimmer, C., Khurana, K.K, Kivelson, M.G., (2000). Subsurface Oceans on Europa and Callisto: Constraints from Galileo Magnetometer Observations. *Icarus*, **147**, 329-347.

Zukas J. (2004). Introduction to Hydrocodes. *Studies in Applied Mechanics.*, Elsevier

Appendix

