Simulation and Analysis of Foam Structure

David Whyte
School of Physics
Trinity College Dublin
The University of Dublin

A thesis submitted for the degree of
Doctor of Philosophy

March 2017

Declaration of Authorship

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university.

Except where otherwise stated, the work described herein has been carried out by the author alone.

I agree to deposit this thesis in the University's open access institutional repository or allow the library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I have read and I understand the plagiarism provisions in the General Regulations of the University Calendar for the current year. I have also completed the online tutorial on avoiding plagiarism, *'Ready, Steady, Write'*.

_____

David Whyte

Acknowledgements

Firstly I must thank my supervisor, Prof. Stefan Hutzler, with whom it has been a total pleasure to work over the last four years. My studies of foam physics have taken countless brief but fascinating detours: whether simulating toy birds or analyzing piano scores, I consider myself lucky for having had a supervisor who encouraged such side projects.

I would like also to thank Prof. Denis Weaire for his numerous insights and wise words, and for never beating around the bush when it came to voicing them.

I must also acknowledge my workmates through the years: Aaron, Steve, Mike, Cathal, Rob, Fritz, Ben and Jens, who have made 3.03 a consistently pleasant work environment.

Thank you to my family for knowing — with a few exceptions — when to mention this thesis and when to steer entirely clear of the topic. Thank you to all of my idiot friends for so consistently distracting me (online and offline) over the last few months.

My four years in the Foams & Complex Systems Group would have been hugely less enjoyable without the good company of my buddy Gav. Thanks for all the coffee and cribbage.

D

Summary

We use the method of bond-orientational order parameter analysis on X-ray tomographic data to investigate the interal structure of an experimental sample of $\sim 25\,000$ microbubbles. By employing appropriate cutoff distances in the parameter space, we see that over the 7 days of the experiment the amount of ordering in the sample increases signficantly. In line with previous experiments and simulations, we see a preference for face-centred cubic (fcc) ordering over hexagonal close-packed (hcp).

We present a simple geometrical argument concerning the ideal shapes of bubbles at an arbitrary liquid fraction between the dry and wet limits. By applying the appropriate transformation to an fcc bubble at a given liquid fraction, we obtain a 'trial' hcp bubble of the same surface area. This surface can be relaxed, proving that the hcp structure has lower energy than the fcc. We perform Surface Evolver simulations of fcc and hcp bubbles over the full range of liquid fractions. The trend observed confirms our proof: the energies are equal at the wet and dry limits, and for intermediate liquid fractions the surface area of the hcp bubbles is very slightly lower.

The Z-cone model is a mathematical formulation which provides analytic approximations to the energy of a bubble as it is deformed. We verify its accuracy for some fundamental test cases: a bubble compressed between parallel plates, a bubble confined to a cube, and a bubble confined to a regular dodecahedron. We see that the energies predicted by the model are accurate.

We apply the model to the case of a bubble in an fcc foam, and see once again that the predicted values of liquid fraction and energy match those obtained from computer simulation. For the fcc case, we obtain from the model an interaction potential similar to that reported in previous simulations. We

derive expressions for the osmotic pressure and, hence, a liquid fraction profile from the cone model expressions which agree with experimental data. Furthermore, we obtain a relationship between liquid fraction and surface liquid fraction which, again, matches experimental data well.

We extend the model to deal with the body-centred cubic structure, resulting in excellent agreement between the model and simulation over the full range of liquid fractions. We investigate the variation of energy with liquid fraction close to the critical liquid fractions at which nearest neighbours and next-to-nearest neighbours are lost. At each point we see logartithmic terms in the variation of energy; however, the forms are different.

We present the results of experiments and simulations concerning the interaction between soap films and fibres: a fibre in the plane of a film, and a fibre in a Plateau border. In each case we see that our simulation predicts the lengths of films and Plateau borders involved. In the latter case we can calculate the force necessary to unpin the fibre from the Plateau border.

Finally, we present simulations concerning pairs of fibres, modelled as infinitely long rigid cylinders, bridged by a liquid drop, for the case of a small contact angle. We see that the drop acts to pull the fibres together, and that, for certain drop volumes, they preferentially orient at an angle which is neither parallel nor perpendicular. We see similar behaviour for a slightly increased contact angle.

List of Publications

1. Stefan Hutzler, Robert Murtagh, David Whyte, Steven Tobin, and Denis Weaire. Z-cone model for the energy of an ordered foam. *Soft Matter*, 10(36):7103–7108, 2014.

2. David Whyte, Robert Murtagh, Denis Weaire, and Stefan Hutzler. Applications and extensions of the Z-cone model for the energy of a foam. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 473:115–122, 2015.

3. Robert Murtagh, David Whyte, Denis Weaire, and Stefan Hutzler. Adaptation of the Z-cone model to the estimation of the energy of a bcc foam. *Philosophical Magazine*, 95(35):4023–4034, 2015.

4. David Whyte, Denis Weaire, Wiebke Drenckhan, and Stefan Hutzler. The relative energy of fcc and hcp foams. *Philosophical Magazine Letters*, 95(6):319–323, 2015.

5. Aaron J Meagher, David Whyte, John Banhart, Stefan Hutzler, Denis Weaire, and Francisco García-Moreno. Slow crystallisation of a monodisperse foam stabilised against coarsening. *Soft Matter*, 11(23):4710–4716, 2015.

6. David Whyte, Nick Didkovsky and Stefan Hutzler. Zero Waste: mapping the evolution of the iterative sight-reading of a piano score. *Music Theory Spectrum* (submitted)

7. David Whyte, Benjamin Haffner, Atsushi Tanaka, Tuomo Hjelt and Stefan Hutzler. The interactions of fibres with soap films. *Colloids and Surfaces A: Physicochemical and Engineering Aspects* (submitted)

List of Presentations

1. Building the pyramids: perfect bubble crystals *(poster)*,
   COST MP1106: Smart and Green Interfaces. Prague, Czechia (2013)

2. Understanding foam–fibre interactions *(poster)*,
   Eufoam. Thessaloniki, Greece (2014)

3. Packings of gas bubbles *(poster)*,
   Shape Up. Berlin, Germany (2015)

4. Extending the cone model to the Kelvin foam *(oral presentation)*,
   Eufoam. Dublin, Ireland (2016)

5. The interaction of foams & fibres *(poster)*,
   Eufoam. Dublin, Ireland (2016)

Image Licenses

- **Fig. 1.1**: licensed under CC BY-NC-ND 2.0[1]
- **Fig. 1.2**: licensed under CC BY-SA 2.0[2]
- **Fig. 1.7**: public domain
- **Fig. 2.1**: reproduced with permission
- **Fig. 4.2**: reproduced with permission
- **Fig. 4.15**: licensed under CC BY-SA 3.0[3]

All other images are my own work.

[1] https://creativecommons.org/licenses/by-nc-nd/2.0/legalcode
[2] https://creativecommons.org/licenses/by-sa/2.0/legalcode
[3] https://creativecommons.org/licenses/by-sa/3.0/legalcode

Contents

List of Figures

Chapter 1

Introduction

1.1 Foams

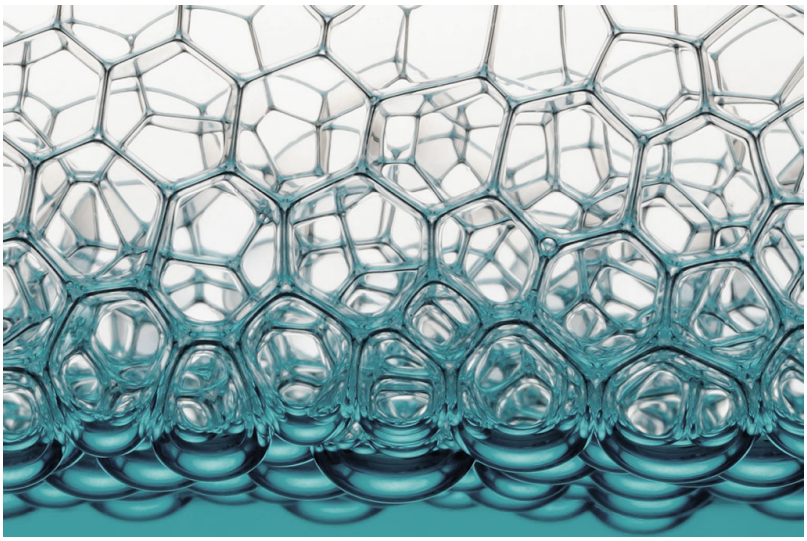

Figure 1.1: A foam. This particularly attractive example is simply made from soap, water and air. *Photo: Long Gao Wu*

A foam is nothing more than a collection of gas bubbles separated by a liquid.

Yet the physics of foams touches on a surprisingly rich variety of topics, including minimal surfaces, sphere packings, and fluid mechanics.

Of course foams are not of interest simply as a topic of academic research. They are ubiquitous in our everyday lives and have applications across fields ranging from firefighting to mining to papermaking.

This is not to mention their aesthetic value — see Fig. 1.1!

1.1.1 Foams research

The modern scientific study of foams began in 1873 when Joseph Plateau published his book *Statique expérimentale et théorique des liquides soumis aux seules forces moléculaires* [1]. C.V. Boys popularized the field with his 1890 book *Soap Bubbles: Their Colours and the Forces Which Mould Them* [2], a compilation of a series of public lectures in which he demonstrated various simple experiments on soap films.

By the 20th century, however, the field had faded into relative obscurity; it was not until 1946 that it resurfaced, thanks to the work of Cyril Stanley Smith and William Lawrence Bragg [3].

Beginning in the 1980s [4], and through the 1990s [5, 6, 7] the field was once again revitalized as computer simulation became a feasible option for modelling the behaviour of 2d and 3d foams.

1.1.2 What is a foam?

Slightly more technically, we may define a foam as a *"two-phase system in which gas cells are enclosed by liquid"* [8]. We note that this definition does seem to exclude *solid foams*, which are the topic of active research. Solid foams, which tend to be formed by solidifying a liquid foam, are very much outside the scope of this thesis. So from here onwards, wherever I write 'foam', I mean 'aqueous foam'.

Being a dispersed two-phase system, a foam is one member of the family of systems called *colloids*. Other members of this family include gels (liquid dispersed in solid), aerosols (liquid dispersed in gas) and emulsions (liquid dispersed in liquid). Of these related systems, emulsions are particularly closely analogous to foams: in fact, many fundamental results pertaining to foams are directly applicable to emulsions.

Bubbles are shaped by surface tension, which we will denote by γ. Under the effects of surface tension, a bubble acts to minimize its surface area. This

Figure 1.2: A soap bubble appears spherical because it wants to minimize its surface area. *Photo: Sara Carter*

is why a free soap bubble attains its familiar spherical shape (as in Fig. 1.2): no body has a smaller surface area than a sphere of the same volume.

When bubbles are packed in a foam, things become more complicated; however much of what goes on in a foam can be understood in terms of the bubbles attempting to minimize their surface area subject to other constraints.

Foams in which all bubbles have the same volume are referred to as *monodisperse*, and ones with a range of bubble volumes are calld *polydisperse*. In practice, precise monodispersity is of course not achievable experimentally: we can define a dispersity:

$$\text{dispersity} = \frac{\sqrt{\langle D^2 \rangle - \langle D \rangle^2}}{\langle D \rangle} \tag{1.1}$$

as the ratio of the standard deviation of bubble diameters D to the mean bubble diameter, and choose an appropriate cutoff value. A dispersity of $\leqslant 5\%$ is often taken as the threshold for monodispersity [9].

1.1.3 Dry foams and wet foams

The two ingredients of a foam are liquid and gas, so it is perhaps natural that the behaviour of a foam depends on how much of each is present. Throughout this thesis we will describe foams in terms of their *liquid fraction*, which we denote ϕ.[1]

$$\phi = \frac{\text{volume of liquid}}{\text{volume of foam}}. \tag{1.2}$$

Foams can broadly be classified as either *dry* or *wet* according to their liquid fraction. The boundary is naturally somewhat fuzzy but it is generally accepted a foam with $\phi \lesssim 0.1$ is dry, and one with $\phi \gtrsim 0.2$ is wet. Liquid fraction is by no means constant throughout a foam; in Fig. 1.1, for example, we see a foam which, as a result of drainage, is dry at the top and wet at the bottom.

The *critical liquid fraction*, ϕ_c, is that liquid fraction at which every bubble in a foam is spherical. When $\phi > \phi_c$, the amount of liquid is such that contact between all bubbles is not maintained. At this point the foam ceases to be a foam and is referred to more properly as a *bubbly liquid* [10]. The value of ϕ_c depends on the structure and polydispersity of the foam. In general, for a monodisperse disordered foam, $\phi_c \approx 0.36$. This corresponds to the void fraction in a random close packing of spheres, and was first computed by Bernal and Mason in 1960 [11]. We refer to $\phi \to 0$ as the *dry limit* and $\phi \to 0.36$ as the *wet limit*.

Fig. 1.3 shows everyday examples of both dry and wet foams. Clearly they look very different. Bubbles in the wet foam are close to spherical and are only just touching. The dry foam adopts a cellular structure, in which liquid channels and thin films separate bubbles which are roughly polyhedral. In this regime, the structure adheres to *Plateau's laws*.

[1] The debate as to whether ϕ should denote liquid fraction or gas fraction rages on. In view of how often I will be referring to liquid fraction, the compromise of ϕ_l, as used in *e.g.* [10] is, in my opinion, too unwieldy to use here. The symbol ϵ is sometimes seen — however we will be using ε to refer to excess energy.

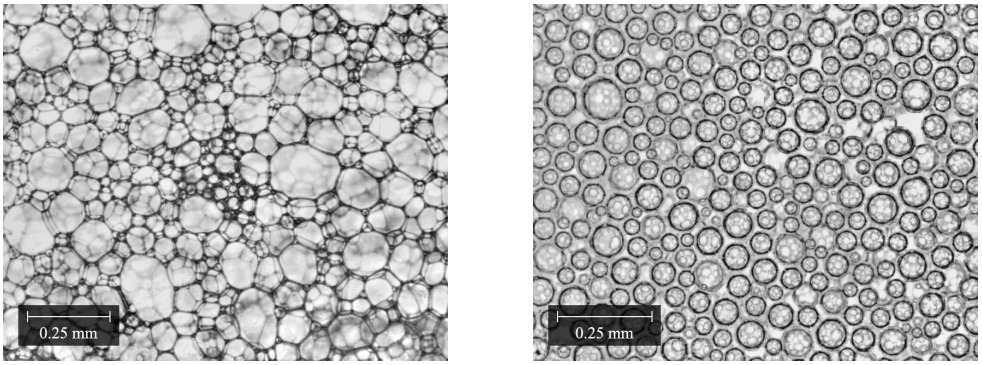

(a) **(b)**

Figure 1.3: Two popular commercially available foams, as viewed under a microscope. (a) Gillette shaving foam: a dry foam, with $\phi < 0.07$ [12]. (b) Guinness beer foam: a wet foam.

1.1.4 Plateau's laws

Joseph Plateau's 1873 work [1] was the first empirical description of the geometry of soap films, applicable to dry foams in equilibrium. It states the following rules:

1. A soap film is a smooth surface of constant mean curvature. The mean curvature is directly proportional to the pressure difference across the film.

2. Soap films meet in threes along an edge, at mutual angles of 120°.

3. Plateau borders meet in fours at nodes, at mutual angles of $\arccos(-1/3) \approx 109°$.

It was not until 1976 that these laws were proven mathematically, by Jean Taylor [13]. The first law in fact predates Plateau, and is called the *Young–Laplace equation*. The edges described in the second law are nowadays referred to as *Plateau borders*. At small liquid fractions the vast majority of the liquid in a foam is contained in its Plateau borders.

The following fourth law is often added [10]:

4. Soap films meet walls at 90°, at *surface Plateau borders*.

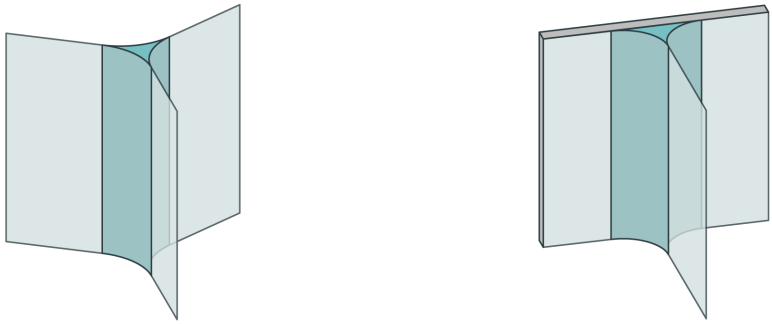

(a) (b)

Figure 1.4: (a) Three soap films meet at a Plateau border at angles of 120°. (b) A soap film meets a solid boundary at a surface Plateau border at an angle of 90°.

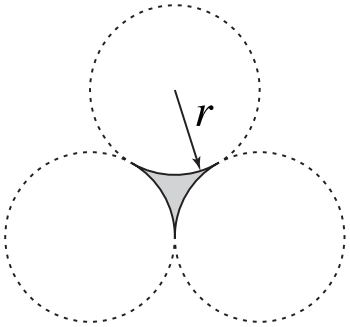

Figure 1.5: Cross section of a Plateau border, as in Fig. 1.4. Its shape consists of three arcs of circle.

This can be understood somewhat intuitively by picturing a wall as a plane of reflection: in order for a soap film to be smooth, as required by Plateau's first law, it must meet the plane at a right angle.

1.1.5 Structure of dry foams

The question of what monodisperse dry foam structure has the lowest energy is an open one. Lord Kelvin conjectured in 1887 [14] that the *body-centred cubic* structure (also now called the *Kelvin foam*), shown in Fig. 1.6 minimizes the surface area per bubble. In this structure, each bubble takes the form of a *truncated octahedron*, with eight hexagonal faces and six square

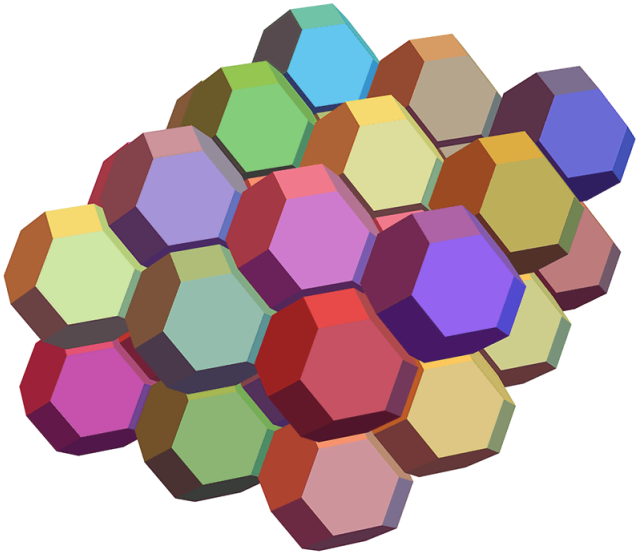

Figure 1.6: Bubbles in a dry body-centred cubic foam or 'Kelvin' foam take the form of truncated octahedra which pack together to fill space.

ones: the hexagonal faces are slightly warped to ensure that Plateau's third law is obeyed. Indeed, the bcc structure is seen to dominate in experimental monodisperse dry foams, created under the correct conditions [9].

In 1993, Weaire and Phelan [6] found a counterexample to Kelvin's conjecture: the so-called *Weaire–Phelan structure*, comprising two different types of bubbles, has a slightly lower surface area per bubble, initially verified by computer simulation, and later proved mathematically [15].

However, unlike the bcc foam, the Weaire–Phelan structure is not seen to spontaneously emerge in monodisperse foams. This is due to the fact that the arrangement of bubbles cannot meet a flat container wall comfortably. It was not until 2012 that a Weaire–Phelan foam was produced experimentally, using a 3d printed container with walls designed to mesh with the structure [16].

1.1.6 Structure of wet foams

In the wet limit, $\phi \approx \phi_c$, bubbles in a foam are approximately spherical. In this regime, bubbles are often approximated as hard spheres acting under

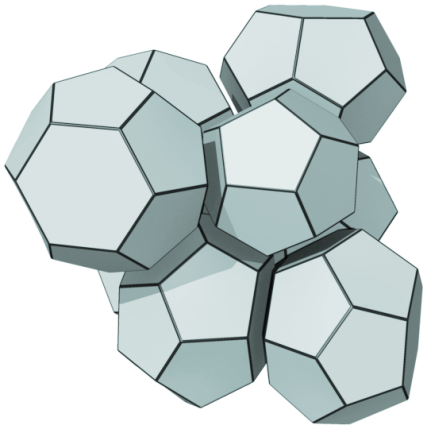

Figure 1.7: The Weaire–Phelan foam is to date the most efficient known partition of space into cells of equal volume.

the influence of buoyancy forces. Hence, when we are dealing with monodisperse foams, finding the minimum energy arrangement of bubbles is equivalent to finding a packing of equal-volume spheres which attains the highest possible density. This is an old problem. Johannes Kepler conjectured in 1611 [17] that the densest such packing is the *face-centred cubic* (fcc) arrangement, whose packing density is $\pi/\sqrt{18} \approx 0.74$; this corresponds to a liquid fraction of $\phi \approx 0.26$. Thomas Hales finally proved this in 2005 via an exhaustive computer-aided minimization process [18].

Nothing beats fcc, but it can be matched. The *hexagonal close-packed* (hcp) arrangement has the same packing fraction. Both strucutres consist of stacked layers of close-packed spheres — they differ in the order in which the layers occur. In hcp, they are arranged in the 'ABA' order, and in fcc, the 'ABC': this difference is perhaps most easily understood visually, as in Fig. 1.8.

It is worth noting here that the so-called *random hexagonally close-packed* (rhcp) arrangement, obtained from arbitrary stacking of A, B and C layers, also achieves the optimum packing fraction of ~ 0.74, but since it can be considered locally either fcc or hcp at any region, it is not considered a separate structure here. Furthermore, the rhcp is very rarely observed experimentally in wet foams [19].

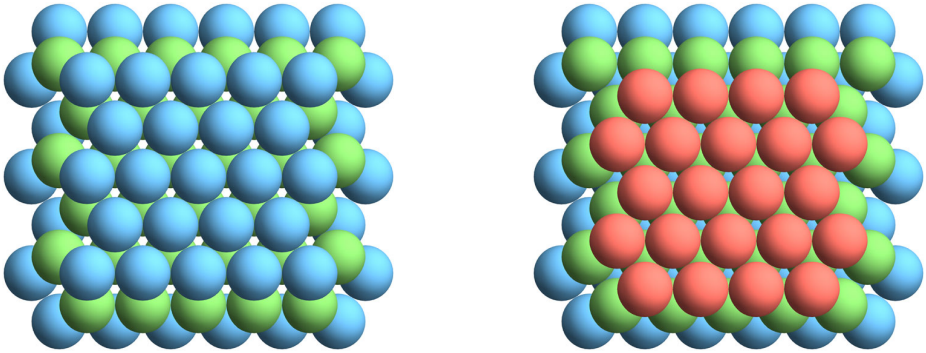

(a) (b)

Figure 1.8: The two densest possible packings of equally-sized spheres. (a) the hexagonal close-packed (hcp) structure. The top layer is aligned with the bottom layer, hence the 'ABA' nomenclature. (b) the face-centred cubic (fcc) packing. The top layer does not line up with either of the other two layers, hence 'ABC'. The full structure is obtained from stacking layers in 'ABABAB...' or 'ABCABC...' order respectively.

1.1.7 Coordination number

The *coordination number* of a bubble, denoted Z, is the number of other bubbles with which it is in contact. Typical values of Z depend on the liquid fraction, polydispersity, and structure of a foam. $\langle Z \rangle$ refers to the mean coordination number taken over all the bubbles in a foam.

For monodisperse foams at the wet limit, we have $\langle Z \rangle = 6$ in the case of a disordered, *i.e.* random, packing of spheres. For an ordered wet foam, $Z = 12$, as we have seen.

At the dry limit, computer simulations on random monodisperse foams give $\langle Z \rangle = 13.7$. For ordered monodisperse dry foams, the bcc has $Z = 14$, and the Weaire–Phelan foam has $\langle Z \rangle = 13$. One can also obtain $\langle Z \rangle = 13.4$ for an idealized dry foam, based on Plateau's laws, by considering the solid angle of the vertex of a tetrahedron [20].

1.1.8 Foam dynamics

We mentioned the Young–Laplace equation in Section 1.1.4. The equation, as it applies to a pair of bubbles separated by a liquid film, may be written as

$$H = \frac{\Delta p}{4\gamma}, \tag{1.3}$$

where H is the mean curvature of the interface between two bubbles, γ is the surface tension of the liquid film, and Δp is the difference in the pressures of the two bubbles.

The mean curvature H can be defined

$$H = \frac{1}{R_1} + \frac{1}{R_2}, \tag{1.4}$$

where R_1 and R_2 are the principal radii of curvature of the interface. H decreases with bubble size, and accordingly (1.3) implies that the pressure of the gas in smaller bubbles is higher than that in larger bubbles. Soap films permit some amount of *diffusion* of gas from bubble to bubble, and so this pressure difference implies that gas tends to flow from small to large bubbles, in a process called *coarsening*. On average, large bubbles grow and small bubbles shrink, and eventually disappear. This means that the average bubble radius increases with time: experimentally one sees $\langle R \rangle \sim \sqrt{t}$ [21]. Hence foams are by their nature dynamical systems, although, as we will see later, the rate of coarsening can be significantly arrested in experiment.

Drainage is another important dynamical effect in foams [22]. A freshly-formed foam with a uniform liquid fraction is not in an equilibrium state. The liquid phase, under the influence of gravity, tends to flow downwards: this occurs, primarily in the Plateau borders of the foam, until an equilibrium state is reached, in which liquid fraction decreases with height. This state represents a balance between forces due to gravity and those due to surface tension [23]. The liquid fraction profile as a reult of drainage can be clearly seen in Fig. 1.1: bubbles are rounder at the liquid interface, where $\phi \approx \phi_c$ is high, and more cellular in the upper region of the foam, where $\phi \to 0$.

Coalescence is a dynamical process closely linked to drainage. The drainage of liquid from films causes them to become thinner and eventually rupture, causing neighbouring bubbles to combine. This occurs with higher frequency in smaller films [24].

In addition to these internal dynamic effects, bubbles in a foam move and rearrange when an external strain is applied. *Foam rheology* is an active area of research focuses on these effects. However, foam dynamics are outside the scope of this thesis. When I discuss foams, it is to be taken that I am referring to a foam in equilibrium.

1.1.9 Capillary length

The *capillary length* l_0 of a fluid–fluid interface (here, between the liquid and gas phases of a foam) is a characteristic length scale which is defined as

$$l_0 = \sqrt{\frac{\gamma}{\rho g}}, \tag{1.5}$$

where γ is the surface tension of the interface (here, the surface tension of the liquid), ρ is the density of the fluid, and g is the acceleration due to gravity. The capillary length of water in air is ~ 2 mm.

Physically, the capillary length is related to the height of a meniscus in a narrow capillary tube, which is given by $\sqrt{2}l_0$. In addition, we may also derive an approximation to the height W_{wet} of a wet foam sitting on a liquid interface to be

$$W_{\text{wet}} \approx \frac{{l_0}^2}{D}, \tag{1.6}$$

where D is the bubble diameter [8]. Equivalently we can say that the number of layers of wet bubbles is approximately $(l_0/D)^2$.

1.1.10 Excess energy

In the wet limit, a bubble is spherical, as this is the shape which minimizes its surface area S. Away from the wet limit, bubbles are not spherical, and

hence have a higher surface area. We describe this increase in surface area in terms of *dimensionless excess energy*, denoted ε, and defined as

$$\varepsilon(\phi) = \frac{S(\phi)}{S_0} - 1, \tag{1.7}$$

where S_0 is the area of a sphere having the same volume as the bubble, *i.e.*

$$S_0 = (36\pi V^2)^{1/3},$$

where V is the bubble's volume.

Naturally, for a spherical bubble, as is the case at $\phi = \phi_c$, this reduces to $\varepsilon(\phi_c) = 0$.

1.1.11 Osmotic pressure

The *osmotic pressure* in a foam is a pressure which arises due to deformation of the bubbles. It was first defined by Henry Princen in 1979 [25] for 2d emulsions, by analogy with a similar quantity in chemistry. In 1986 he offered a much more thorough treatment of the concept, in 3d [26].

The osmotic pressure is denoted Π, and is defined as the force felt per unit area on a semi-permeable membrane (*i.e.*, one which allows liquid to pass through it but not gas bubbles), as it is moved into the foam, as illustrated in Fig. 1.9 (this apparatus is called an *osmometer*). The osmotic pressure of a foam is very sensitive to its liquid fraction: at the wet and dry limits we have

$$\Pi(\phi = \phi_c) = 0, \quad \text{and} \quad \Pi(\phi = 0) \to \infty \tag{1.8}$$

Osmotic pressure ultimately stems from the fact that at a higher liquid fraction, bubbles are closer to spheres, and have lower surface area. Removing liquid from a foam increases the surface area of the bubbles within it, and hence requires an amount of work which is proportional to the incuded increase in surface area.

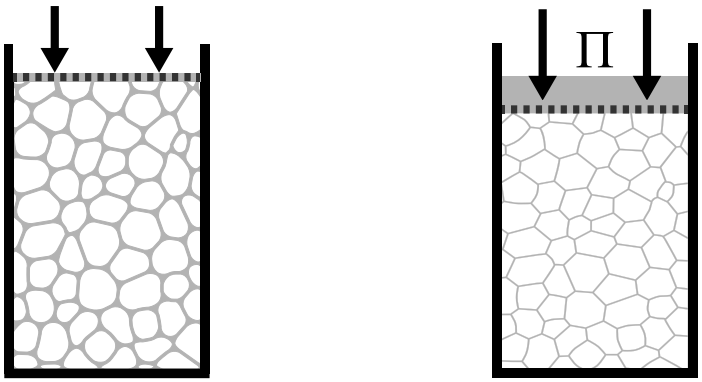

Figure 1.9: An osmometer consists of a semi-permeable membrane which is forced into the foam. The osmotic pressure Π, which is a function of ϕ, is simply the force felt on the membrane per unit area, which is necessary to deform the bubbles.

1.2 Surface Evolver

1.2.1 Overview

Surface Evolver is an interactive program which minimizes the energy of sets of surfaces subject to given constraints [5]. The software was created in 1992 by Kenneth Brakke at the University of Minnesota's Geometry Center. The software is freely available online and is still maintained and updated by Brakke — the most recent version, 2.70, was released in August 2013.

Surface Evolver has been invaluable for the computational study of three-dimensional foams. An early coup for the program was the verification that the Weaire–Phelan structure did indeed beat Kelvin's conjecture [6].

1.2.2 Principle of operation

Surface Evolver represents a surface as a mesh of triangular facets (as shown for the 'double bubble' in Fig. 1.10). The user specifies the positions of a list of vertices, followed by a list of directed edges which link vertex to vertex, and facets which are spanned by a list of edges: the initial setup of vertices typically comprises a crude but topologically faithful approximation to the desired arrangement.

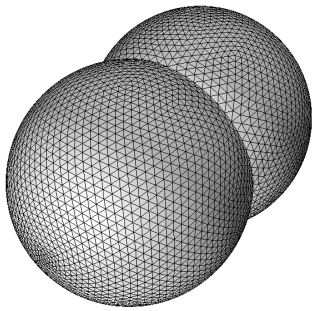

Figure 1.10: Surface Evolver represents surfaces as a mesh of triangular facets. Shown here is the 'double bubble'. The mesh is not particularly fine here.

Also specified are constraints which may apply to any vertices, edges or facets. Where films meet planes, *e.g.* the wall of a container, we may impose a planar constraint, for example. Linear constraints, *e.g.* soap films attaching to a wire in a wireframe, can be implemented as a pair of intersecting planar constraints. Curved wire boundaries may also be implemented by parametric description of the curve in 3d space.

The two fundamental operations of Surface Evolver are *refinement* and *iteration*. In a refinement step, each triangular facet is subdivided into four smaller coplanar facets (see Fig. 1.11). In an iteration step, each vertex is moved according to a gradient descent method, over the total energy function, subject to the appropriate constraints. The result of this is, hopefully, a small reduction in the energy after each iteration step. There are a multitude of other operations at the disposal of the Surface Evolver user; many of these are useful tools which allow systematic 'grooming' of the surface; for example, automatic subdivision of overly long edges, deletion of extremely small facets, *etc*. Nonetheless a Surface Evolver script is at its core a carefully constructed list of refinements and iterations.

The canonical simple Surface Evolver example is perhaps the evolution of a cube into a sphere (a version is included with the software as `cube.fe`).

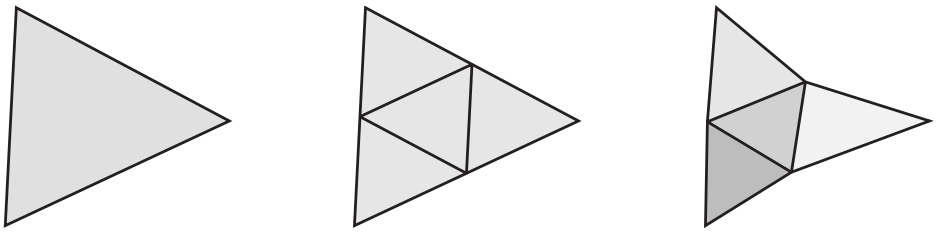

(a) (b) (c)

Figure 1.11: The general principle of operation of Surface Evolver as shown for a single facet. (a) $\rightarrow$ (b) is a refinement step; (b) $\rightarrow$ (c) an evolution step.

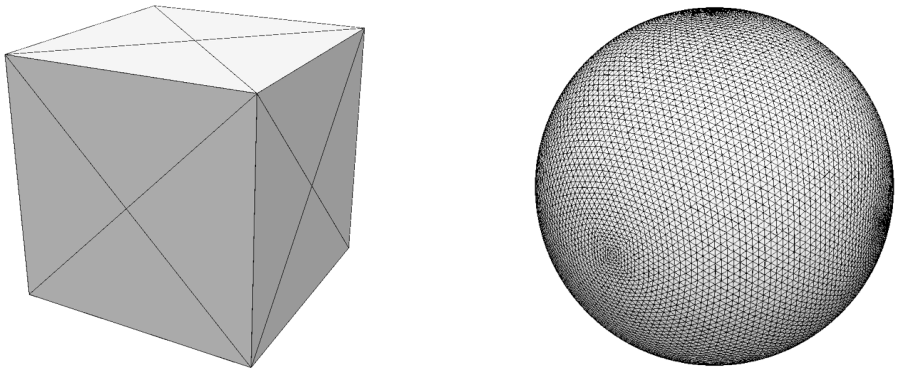

(a) (b)

Figure 1.12: Simulating a free bubble using Surface Evolver. (a) The starting point is a crude approximation to the desired surface. (b) The result, after several iterations and refinements.

The initial arrangement of surfaces, as shown in Fig. 1.12(a), is a cube, which Surface Evolver automatically divides into 16 triangular facets. After 5 refinement steps, with $\sim$ 100 iterations between each, we arrive at the body shown in Fig. 1.12(b): a good approximation to a sphere. If higher accuracy were required, the steps can be repeated as many times as necessary, limited only by memory and performance concerns. After a refinement step, each iteration takes roughly 4 times as long, and the surface requires 4 times as much memory to store.

Surface Evolver also allows the implementation of periodic boundary conditions via the so-called *torus model*: naturally, this makes the file setup and

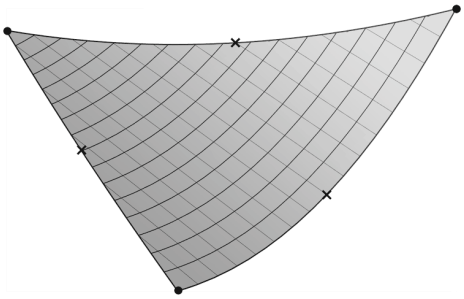

Figure 1.13: Using Surface Evolver's quadratic model, facets take the form of *quadratic patches*: each edge is a quadratic spline which joins two endpoints (•) via a specified midpoint (×). This is a single facet — contour lines are to visualize curvature.

management of edge and facet directions more awkward. For this reason, I tend to avoid using this mode here, and exploit the symmetries in periodic structures to allow them to be represented using standard boundary conditions. This is explained in more detail later.

1.2.3 Quadratic model

For simulations where higher accuracy in edge lengths or facet areas is desirable, at the expense of longer simulation times, Surface Evolver also includes a *quadratic model* to represent facets. The general principle of operation is identical to the linear model described above. Under the quadratic model, each edge is assigned a midpoint vertex in addition to its endpoints. Edges are modelled as quadratic splines which passes smoothly through their three assigned points. Facets are described not by simple triangles but by *quadratic patches*: see Fig. 1.13. Naturally this model slows down computation time — not only are there more vertices whose positions must be calculated for each iteration step; the calculation of surface area is now a quadratic integral in 3d.

1.3 Structure of this thesis

This thesis comprises three main strands of work:

- fcc and hcp foams
- The cone model
- The interaction of foams and fibres

These touch on theoretical, computational and experimental work. Here I will clarify where my own contributions lie in regard to each section.

The chapter concerning fcc and hcp foams begins with analysis of experimental data. The experiment was carried out by Aaron Meagher in Berlin, and the tomographic data was segmented by that group. I performed all of the analysis detailed in the chapter based on 3d bubble position data provided to me. The proof regarding the energies of the structures was a collaborative effort on the part of our group. The various simulations backing up the proof are entirely my work.

The cone model was derived and formulated mathematically by Robert Murtagh. I implemented the model computationally using Wolfram Mathematica. I also derived expressions for, and computed, the physical quantities derived from the model, detailed in the chapter. All of the Surface Evolver simulations in this chapter are my work.

Regarding foams and fibres, the quasi-2d experiment was carried out by our collaborators in VTT, Finland. The frame experiment was carried out by Benjamin Haffner in TCD. Again, all of the simulations in the chapter are my work.

Chapter 2

fcc and hcp Foams

2.1 Introduction

2.1.1 Motivation

It has long been known that small monodisperse bubbles (called *microbubbles*, for diameters < 1 mm), when produced in a pool of surfactant solution, spontaneously form ordered structures at the surface. These assemblages of bubbles were investigated by Bragg and Nye in 1947 as model systems for crystal structures [27], and are now known as *Bragg rafts* in recognition of this work. While Bragg and Nye did not pursue this topic further, Bragg rafts have since been used to model various dynamical processes in crystals and glasses [28, 29]. Specialist equipment is not needed to create a Bragg raft — Fig. 2.1 shows an example of one which formed naturally in the process of washing dishes!

Bragg and Nye's original paper focuses primarily on rafts consisting of a single layer of close-packed bubbles. They briefly discuss rafts of multiple layers, suggesting that both fcc and hcp ordering occurs, but note that observation of 3d structure was difficult, since the effects of refraction mean that each layer of bubbles distorts the view of the layer below it.

More recently, the computer graphics technique of *ray-tracing* has been used to infer the structure up to a depth of four layers of bubbles from optical microscopy [30]. An analysis of a large number of monodisperse foam samples using this ray-tracing method suggested that bubbles in the wet limit

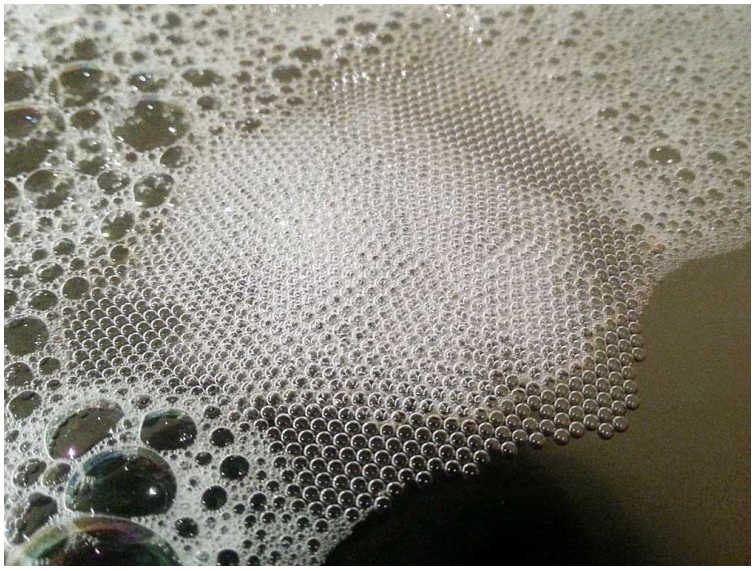

Figure 2.1: A Bragg raft consisting of several close-packed layers of bubbles, spontaneously produced while washing dishes.

tend to order preferentially in the fcc structure over the hcp, at a ratio of approximately 2:1 [31].

The preference for fcc over hcp observed in this optical experiment is in line with analogous observations in collections of spherical glass particles which, when gently vibrated, are seen to settle into the fcc structure more readily than the hcp [32].

However, all of the ordering observed with the aid of ray-tracing occurs within no more than 4 outer layers of bubbles — refraction once again is the limiting factor. The results, therefore, are perhaps not strictly representative of a bulk foam. To probe the *internal* structure of a foam, more sophisticated techniques are needed.

In this chapter, we focus on a recent experiment which allowed the full *internal* structure of a monodisperse foam to be calculated. We also present a rigorous proof concerning the *energy difference* between the fcc and hcp structures between the wet and dry limits.

2.1.2 Experimental overview

The 'lifetime experiment' [33] was an experiment conducted by Aaron Meagher to investigate the full 3d structure of a foam sample consisting of monodisperse microbubbles, and the evolution of this structure over time. The internal structure was determined using benchtop X-ray computed tomography (CT).

The liquid phase of the sample was a solution of the commercially available Fairy washing-up liquid, at a concentration of 5% by volume. The gas phase was a mixture of nitrogen and perfluorohexane (C_6F_{14}) vapour. C_6F_{14} is not soluble in water and hence does not readily pass from bubble to bubble. This significantly reduces the rate of coarsening, meaning the bubble sizes remain close to constant over the length of the experiment. Bubbles were produced using a flow-focusing device [34], which results in bubbles of a very controllable size. The dispersity (defined as the standard deviation of the distribution of bubble diameters divided by its mean) was 2.4% on day 1, increasing to 3.4% on day 7 due to coarsening: this is below the commonly-used threshold of 5%, allowing us to classify the foam as effectively monodisperse.

Approximately 10 000 bubbles with an average diameter of 0.8 mm were issued into an open cubic container of side length 20 mm. The container was then sealed by sliding a Perspex plate over the open face. Of the $\sim$ 10 000 bubbles in the sample, $\sim$ 2 300 are surface bubbles, in contact with either the container wall or the surface of the liquid. Once per day over a seven-day period, the sample was imaged using a benchtop X-ray computed tomography system. The pixel size of the detector is 50 μm, meaning that the images are sufficiently high quality to allow accurate determination of the centre position of every bubble in the sample.

The 'rule of thumb' described in Section 1.1.9 gives a number of layers which may be considered 'wet', in terms of the capillary length l_0 and the bubble diameter D: $N_{wet} = (l_0/D)^2$. In this case this gives $N_{wet} \approx 6$. Since

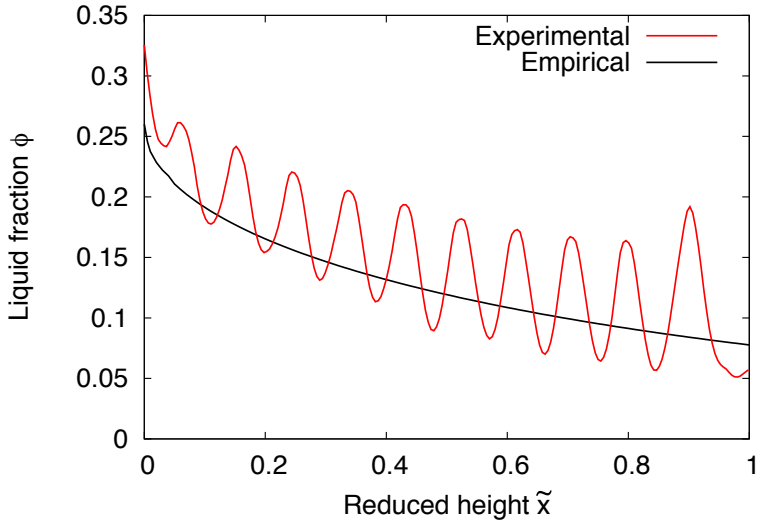

Figure 2.2: The variation of liquid fraction with reduced height as computed from the lifetime experiment data on day 7, as compared to an empirical expression from Höhler *et al.* [9].

the sample consists of 11 layers of bubbles, it is important to note here that the sample cannot be considered to be at the wet limit throughout; however, the average liquid fraction was $\phi \approx 0.2$, as determined from application of the Beer–Lambert law to the X-ray absorption data, and so we are certainly dealing with a wet foam, even if not all the way at ϕ_c.

By representing each bubble as a sphere of appropriate volume, we can compute an approximate local liquid fraction at any given height in the foam. Fig. 2.2 shows the liquid fraction profile, obtained in this way, of our foam on day 7 in terms of the reduced height $\tilde{x} = x\,R/l_0^2$, where R is the bubble radius. Each 'dip' represents a horizontal plane of bubbles. We compare the liquid fraction profile with an empirical expression derived from Höhler *et al.* [9] (discussed in more detail in Section 3.5.4). We see that while our foam is generally wetter than the empirical expression predicts, the liquid fraction generally falls off at approximately the expected rate.

2.2 Bond orientational order parameters

2.2.1 Background

Bond orientational order parameters (BOOPS) are a set of numbers, denoted Q_ℓ, which quantify the local structure of a set of points. Paul Steinhardt *et al.* first employed BOOPS in 1983 in order to study the change in the structure of atoms undergoing the liquid–glass transition [35].

For a given point, its BOOPS are determined based only on the relative positions of its set of nearest neighbours. By computing the BOOPS for a point and comparing the values to a reference set of BOOPS for known lattices, we can determine with good accuracy the local structure around that point. Steinhardt refers to this process as "shape spectroscopy".

In the case of the lifetime experiment, we have accurate position and volume data for every bubble in the foam sample for each day of the experiment. From this we can determine each bubble's nearest neighbours. The method of BOOPS is an ideal tool to help us gauge the local structure throughout the sample in this case.

2.2.2 Computation

We leave the full description of the definition of Q_ℓ to Appendix A, but note that the useful BOOPS are Q_4 and Q_6, and that they depend only on the positions of each bubble's nearest neighbours. The nearest-neighbour positions for fcc and hcp are of course known analytically; hence we can obtain analytic values for their (Q_4, Q_6) signatures. They are as follows:

$$\begin{aligned} \textbf{fcc:} \quad & \left(\sqrt{7/192}, \sqrt{169/512}\right) \approx (0.191, 0.575) \\ \textbf{hcp:} \quad & \left(7/72, \sqrt{29237/124416}\right) \approx (0.097, 0.485) \end{aligned}$$

Plotting these two points in (Q_4, Q_6) space (Fig. 2.3), they appear relatively close together — so much so that we must be cognizant of the risk of 'crosstalk' between these values, *i.e.* bubbles which are fcc being classed as hcp, or *vice versa*. However, as we will see, for our data we can separate the peaks cleanly.

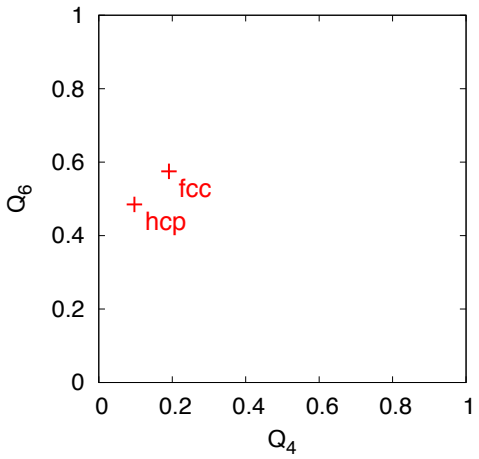

Figure 2.3: The BOOP signatures for the fcc and hcp structures, plotted in (Q_4, Q_6) space. The two points are relatively close together; however, as we will see, the peaks at each point are narrow enough to be distinguishable.

2.3 Lifetime experiment

2.3.1 2d and 3d visualizations

Using image processing software, we can obtain from the CT scans accurate values for the x, y and z coordinates of the centre of each bubble in the sample, for each of the 7 days. A natural first step in probing the ordering of the sample is a direct two-dimensional visualization of the bubble positions.

Fig. 2.4 shows a simple 2d projection of the x and y coordinates of the bubble centres on the first day of the experiment (where y is the direction of gravity). Even from this crude visualization method it is evident that there is some degree of ordering in the sample: we see the bubbles tend to 'line up' parallel to the edges. This is to be expected: the planar walls of the container encourage close packing of the bubbles.

This close packing is evident in simple 3d visualizations too. In Fig. 2.5 we draw a sphere around each computed bubble centre for the first day of the experiment (note that this is *not* a direct visualization of the 3d tomographic data). Immediately we see once again that the bubbles are fairly well-ordered at the walls of the container on the first day.

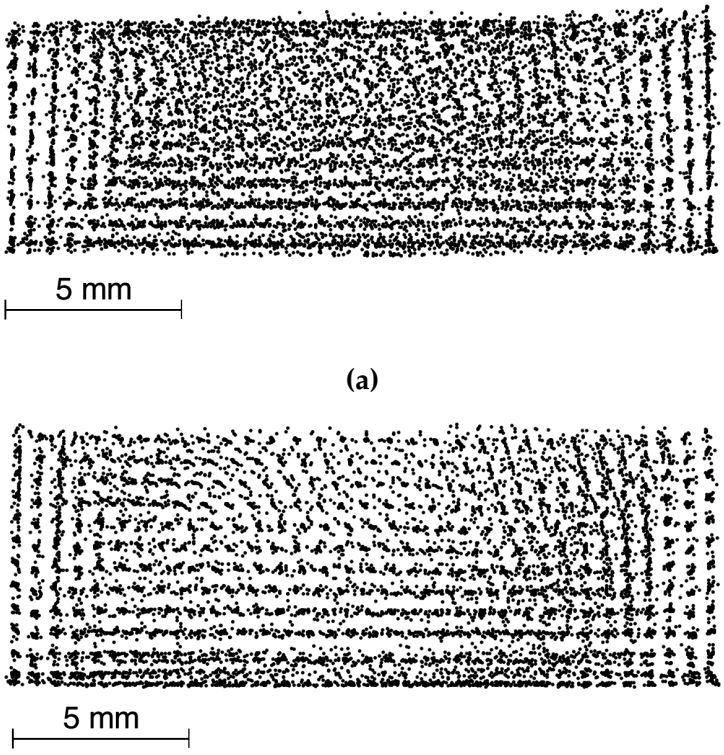

Figure 2.4: x and y positions of bubbles. (a) On day 1 of the experiment, this view of the bubble centres reveals that there is some degree of ordering near the walls. (b) By day 7, the ordering has become clearer: throughout the sample there are horizontal close-packed planes.

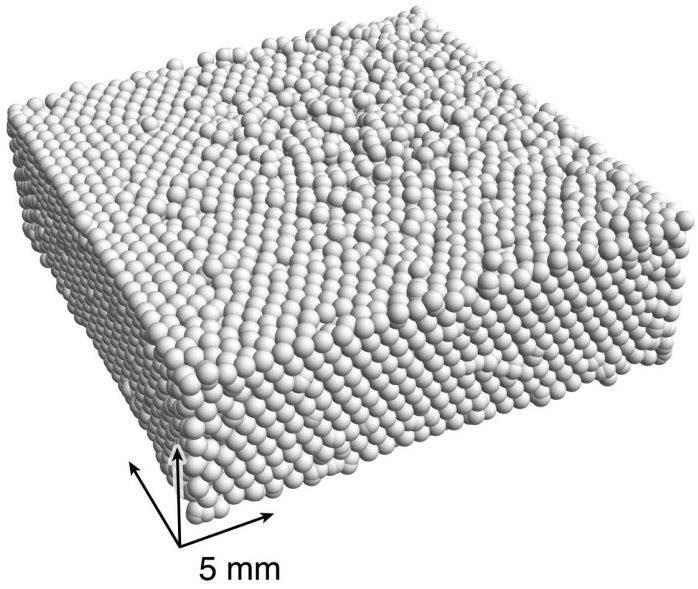

Figure 2.5: 3d visualization of bubble positions on day 1 of the lifetime experiment. It is evident here that the ordering near the walls is due to the bubbles roughly forming close-packed planes. Each of the three arrows represents a length of 5 mm. This image, and all the 3d images in this section, were created using Wolfram Mathematica.

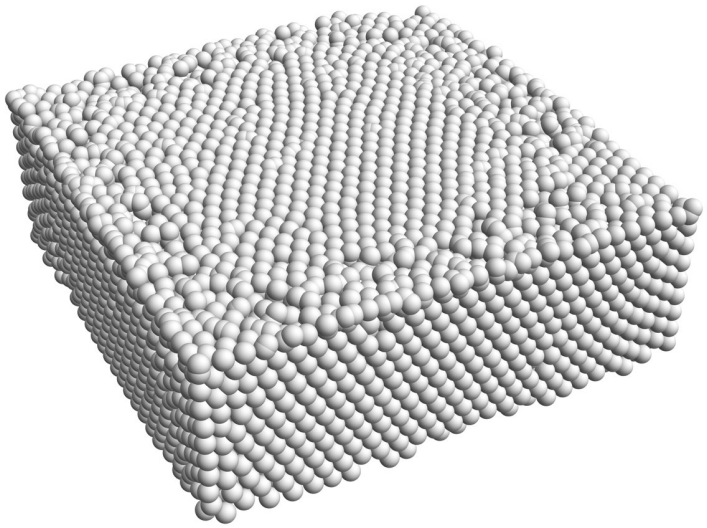

Figure 2.6: 3d visualization of bubble positions on day 7 of the lifetime experiment. Compared to Fig. 2.5 the extent of the ordering has increased: in particular, the close-packing of exterior bubbles is much more pronounced.

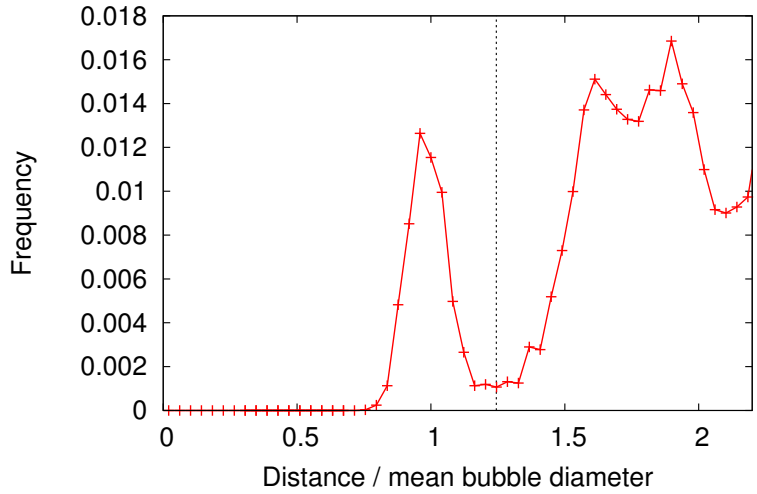

Figure 2.7: A truncated radial distribution function $g(r)$ for the foam on day 1 of the experiment. The position of the minimum between the first two peaks gives us a cutoff distance r_c for bubbles we wish to consider nearest neighbours.

Hence we are compelled to investigate more quantitatively the extent of the ordering in the foam. BOOP analysis is a suitable tool.

It should be noted here that the sample was accidentally dropped on day 3 of the experiment; this significantly disrupted the foam structure. Nonetheless, we can observe separate trends on days $1-2$ and $3-7$.

2.3.2 Coordination number analysis

All of our analysis here hinges on the properties of bubbles' nearest neighbours. To define the set of nearest neighbours for each bubble, we need to choose a cutoff distance r_c for our definition of nearest neighbours.

We do so by computing the radial distribution function, denoted $g(r)$. $g(r)$ is defined as the probability distribution to find a bubble centre at a distance r from the centre of some reference bubble [36]. The first two peaks in $g(r)$ represent the typical separation distances for nearest neighbours and next-to-nearest neighbours respectively. The position of the *minimum* value of $g(r)$ between these two peaks is therefore a natural definition for our near-

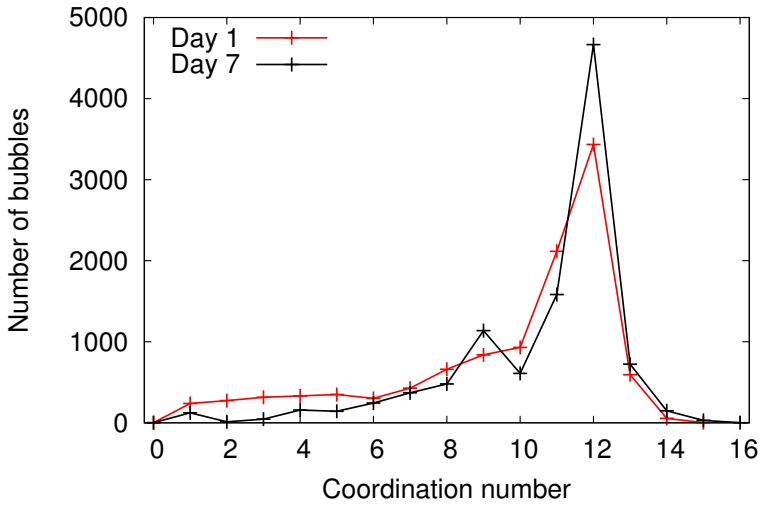

Figure 2.8: Histograms of the coordination numbers Z as obtained from our cutoff radius on (a) day 1 and (b) day 7 of the lifetime experiment. The most common coordination number is $Z = 12$ over the length of the experiment, and the peak grows sharper over time.

est neigbour cutoff distance: Fig. 2.7 shows the distribution on day 1, as well as the computed r_c. For each day we compute a value of r_c by this method.

Having decided on a cutoff distance r_c, we can obtain for each bubble its coordination number Z. The coordination number for a bubble is the number of nearest neighbours it possesses (analogously to the same quantity in crystallography [37]); we interpret this as the number of bubbles whose centres are separated from it by less than r_c. For each day we arrive at a distribution of Z: we plot histograms for days 1 and 7 in Fig. 2.8.

Note that the presence of bubbles with apparent coordination numbers $Z \geqslant 12$ is an artifact of the imperfect nature of the cutoff distance method. For a monodisperse foam a bubble cannot have more than 12 contacts: 12 is the so-called *kissing number* in 3 dimensions, as conjectured by Newton, and proven by Bender in 1874 [38].

We can also chart the evolution of the sample by noting the proportion of bubbles which have $Z = 12$: see Fig. 2.9. This is a simple measure of the amount of ordering in a sample. We note that theoretically, a bubble

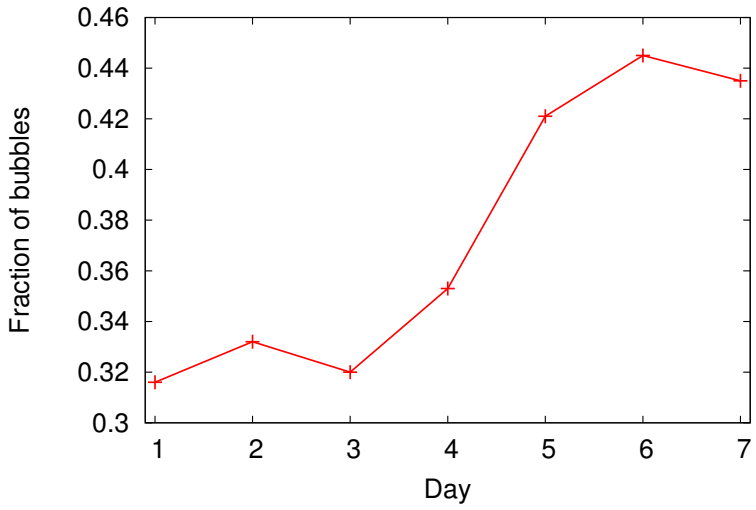

Figure 2.9: The fraction of bubbles with coordination number $Z = 12$ is a measure of order in the sample. Here we see that, with the exception of the dip on day 3 due to mishandling, and a slight decrease on the final day, it increases over the course of the experiment.

with $Z = 12$ must be locally ordered (in either the fcc or hcp structures, or icosahedrally [39]). Due to our choice of r_c not all bubbles with $Z = 12$ will necessarily be perfectly ordered: however it is still a useful and convenient method of obtaining a trend.

2.3.3 BOOP analysis

Now that we have for each bubble a set of nearest neighbours, and hence nearest-neighbour directions, we can compute the BOOPS Q_4 and Q_6 directly, via the formula described in Appendix A. This results in a 'cloud' of points in (Q_4, Q_6) space for each of the 7 days, as shown in Fig. 2.10. It is clear that the large number of points makes it impossible to draw any meaningful conclusion from such plots. For this reason we aim to quantify the shapes of these clouds.

We wish to classify bubbles as 'fcc' or 'hcp' depending on their proximity to the appropriate BOOP signatures, and so we must turn again to cutoff distances: this time, however, in (Q_4, Q_6) space. Unfortunately here we cannot

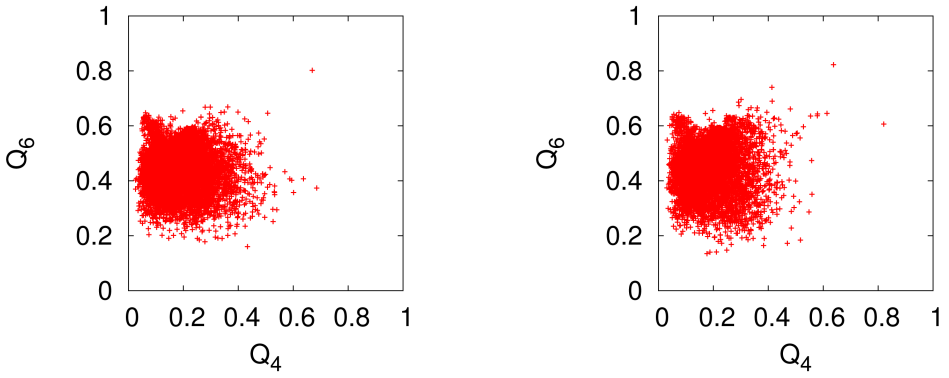

Figure 2.10: The distribution of Q_4 and Q_6 values on (a) day 1 and (b) day 7 of the experiment. The large number of points makes it impossible to draw any meaningful conclusion from such plots.

employ the first-minimum method which we used to define nearest neighbours. The relative proximity of the fcc and hcp BOOP signatures means that we may use a scheme as shown in Fig. 2.11: the threshold radius is determined by the distance between the points, rather than chosen arbitrarily. We obtain a threshold radius $r_{thresh} = 0.065$ by this method.

Luchnikov *et al.* analyzed Q_4 and Q_6 distributions to classify spheres in simulated packings as either fcc or hcp by a similar method [40]: they define elliptical regions around the ideal (Q_4, Q_6) values for fcc with major and minor axes decided "from inspection" — the choice of axes introduces four arbitrary parameters into the classification scheme.

In Fig. 2.12 we plot the fraction of bubbles classified as either fcc or hcp by this method. The trend is clear: as the experiment progresses (with the exception of day 3, when the sample was dropped, as noted) the number of bubbles which are classified as both fcc *and* hcp increases steadily. This strongly suggests that the sample is becoming ordered over time. On day 6, $\sim$ 72% of the bubbles are reported as being either fcc or hcp. Given that $\sim$ 2 500 of the bubbles lie on the walls of the container, and mostly do not bear either the fcc or hcp signatures, this corresponds to an ordering rate of $\gtrsim$ 85% in the bulk.

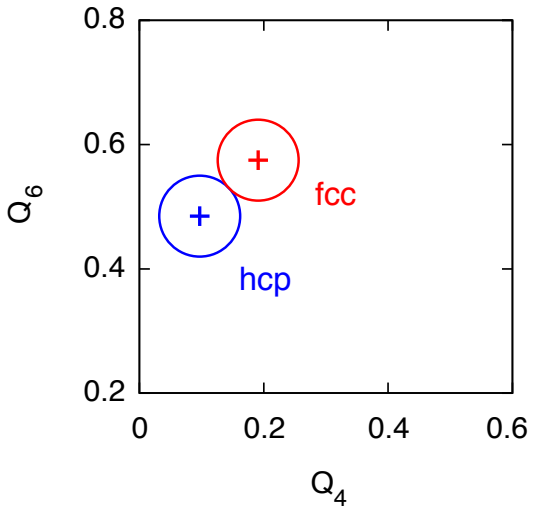

Figure 2.11: We classify a bubble as fcc or hcp if it is sufficiently close to that structure's (Q_4, Q_6) signature. The threshold distance is chosen such that the circles are in contact: $r_{thresh} = 0.065$.

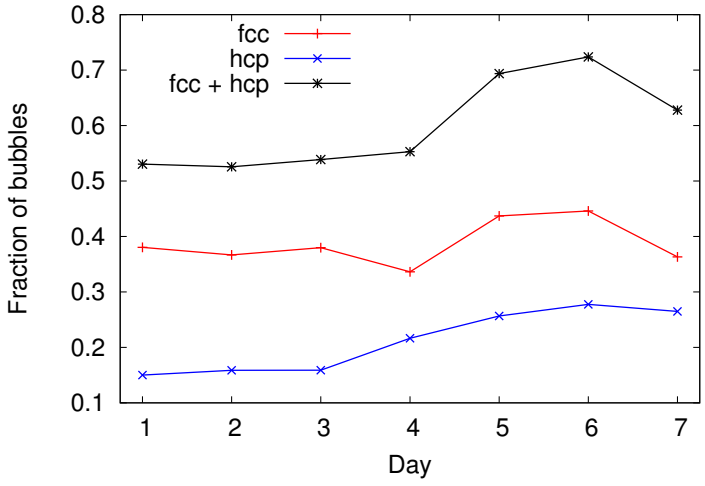

Figure 2.12: The fraction of bubbles classified as fcc or hcp, using a large threshold (Q_4, Q_6) distance as shown in Fig. 2.11.

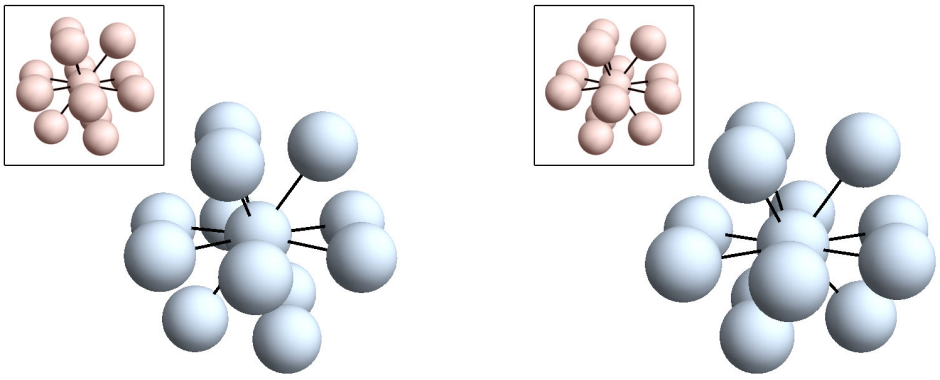

(a) **(b)**

Figure 2.13: Randomly selected examples of bubbles which were counted as being (a) fcc and (b) hcp respectively, but whose distance from the ideal (Q_4, Q_6) signature is near ($> 0.9\times$) the cutoff distance. Inset are visualizations of the theoretical nearest neighbour directions for the relevant structures. Despite the distance between their signatures and the ideal cases, (a) clearly displays ABC ordering in the three layers (of 3, 6 and 3 bubbles respectively), and (b) displays ABA.

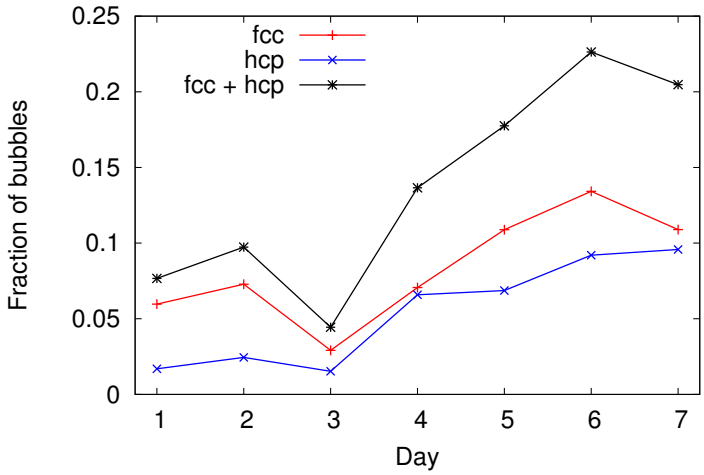

Figure 2.14: The fraction of bubbles classified as fcc or hcp, using a smaller threshold (Q_4, Q_6) distance: namely, 1/5 of that used in Fig. 2.12. The 'dip' on day 3 becomes prominent when using this tighter threshold: this is due to a mishandling of the sample on that day.

It is tempting to believe that this huge amount of ordering may be purely due to the generosity of the threshold distance used in (Q_4, Q_6) space. In Fig. 2.13 we draw the positions of nearest neighbours for two examples of bubbles which are classified as fcc or hcp but whose (Q_4, Q_6) signatures lie very close to the edge of their respective basins. We see that, despite their relatively high distance from the ideal BOOP signatures, they are recognizable as the appropriate ABC or ABA arrangement. For this reason we proceed using this cutoff distance, confident that it does not result in an unduly high rate of false positives.

Nonetheless it is worth repeating this analysis with a stricter threshold: Fig. 2.14 shows the same type of plot as Fig. 2.12, except using a threshold (Q_4, Q_6) distance which is smaller by a factor of 1/5: where before we used $r_{thresh} = 0.065$, here we use $r_{thresh} = 0.013$. Naturally the fraction of bubbles classified as being ordered is lower, but the same trend is clearly visible; in fact, the dip on day 3 is much more pronounced here.

Here it should be pointed out that a potential drawback of the BOOP method

is the possible presence of false positives in the classifications it generates. One technique which has been suggested to ameliorate this drawback is the use of *averaged bond-orientational order parameters* [41]. Using this method, we replace each bubble's Q_4 and Q_6 with the mean value of the quantity, as taken over the bubble and all of its nearest neighbours. In our case, however, this method smears the peaks of the (Q_4, Q_6) distribution, as well as producing a huge number of false *negatives*. Bubbles whose neighbours' positions clearly imply fcc or hcp ordering are dragged away from the correct (Q_4, Q_6) signature by this averaging process. For this reason, we continue to use the 'traditional' BOOP method. Presumably, for less ordered datasets, the number of false negatives introduced is smaller than that of false positives removed, which would explain the advantage of such an averaging process.

Mickel *et al.* [42] highlight how the computed values of Q_4 and Q_6 for a particle can depend strongly on the choice of nearest-neighbour selection algorithm. However here there is less ambiguity than in their examples: looking at Fig. 2.7, the outer edge of the peak representing the first shell is quite well defined. In other words, there is very little overlap between the nearest-neighbour and next-nearest-neighbour peaks.

The sharp decrease on day 3 is explainable by the accidental jolting of the sample, as described earlier. We see a smaller, but noticeable, decrease on day 7: while it was not noted that the sample was knocked as on day 3, it is possible that it was handled less steadily than on previous days. As discussed earlier, a higher degree of ordering corresponds to a lower total energy. Accordingly, a decrease in the ordering implies an increase in the total energy of the system: this cannot occur spontaneously.

We can plot 2d histograms of the joint distribution of (Q_4, Q_6) values to visualize their behaviour as the experiment progresses. Figs. 2.15, 2.16 and 2.17 show such histograms for days 1, 4 and 6 respectively. This method of visualization makes the increase in ordering abundantly clear.

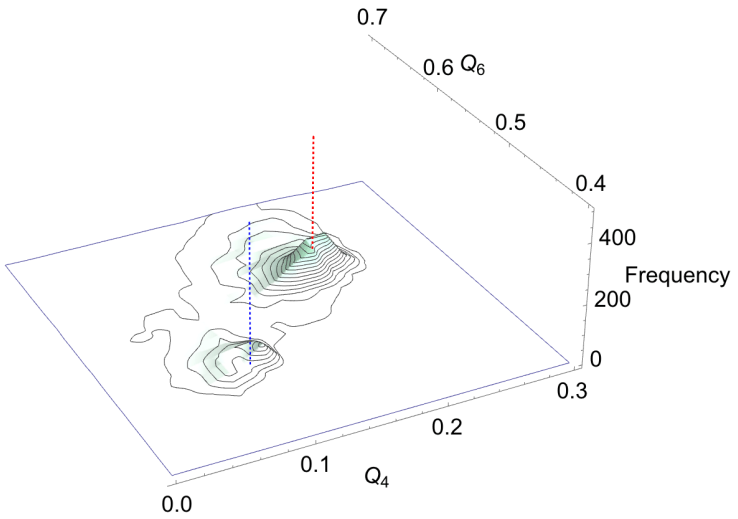

Figure 2.15: Histogram of Q_4 and Q_6 on day 1 of the lifetime experiment. The BOOP signatures for fcc and hcp are marked in red and blue respectively. This is a probability distribution function in two dimensions: the total volume under the surface is 1, and the data is smoothed by a Gaussian kernel of width 0.002.

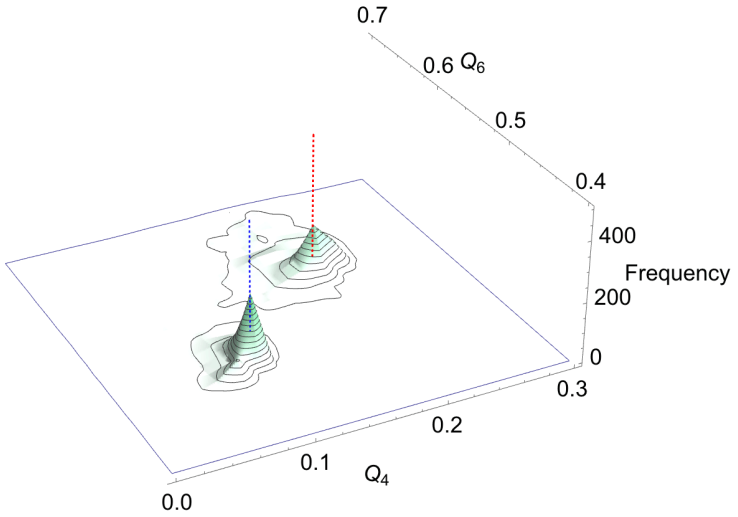

Figure 2.16: Histogram of Q_4 and Q_6 on day 4 of the lifetime experiment.

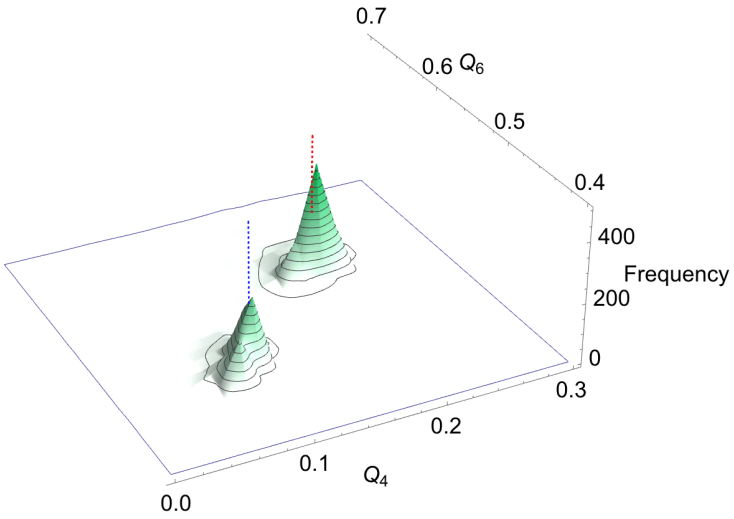

Figure 2.17: Histogram of Q_4 and Q_6 on day 6 of the lifetime experiment. The peaks for fcc and hcp have clearly increased in height and become narrower over the length of the experiment.

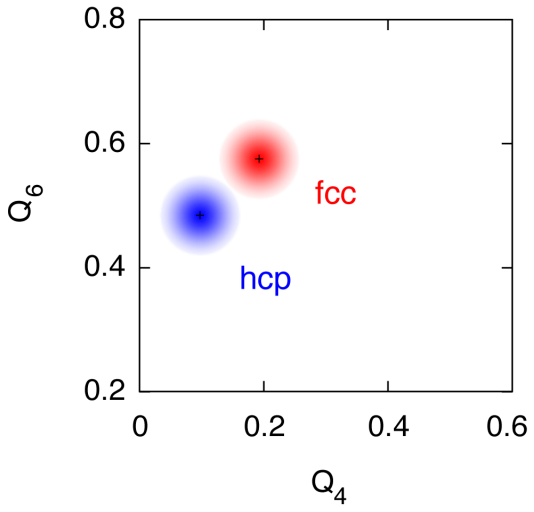

Figure 2.18: Each bubble is coloured according to its position in (Q_4, Q_6) space: the intensity of the colour fades as we move away from the ideal fcc or hcp value.

2.3.4 3d BOOP visualization

We can combine our bubble position data with our BOOP data to create a clearer visualization of the ordering of the sample. As in Fig. 2.5 we place a sphere on every bubble centre, but we now colour the spheres according to their computed BOOP signature: red for fcc, blue for hcp, and white for 'other': the *proximity* of each bubble's (Q_4, Q_6) signature to the ideal values is shown by fading the intensity of the colour with distance from the fcc or hcp value: this scheme is illustrated in Fig. 2.18.

Fig. 2.19 shows the result of such a visualization on days 1 and 7. We see that bubbles at the walls of the container, despite lying in a close-packed plane, are generally not counted as fcc or hcp by our BOOP method.

To account for this effect, we focus now on interior bubbles, by restricting our analysis to a given range of x, y and z near the centre of the sample. By doing so, the process of crystallization becomes much clearer.

Fig. 2.20 shows the evolution of this central section of the foam. The trend is obvious in this view. After the disruption on day 3, ordered layers of fcc and hcp spontaneously form and increase in extent over time. By day 7,

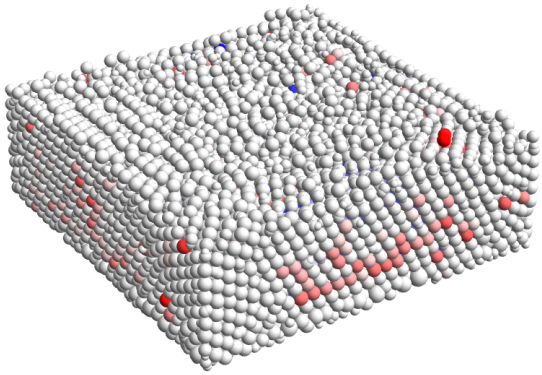

(a)

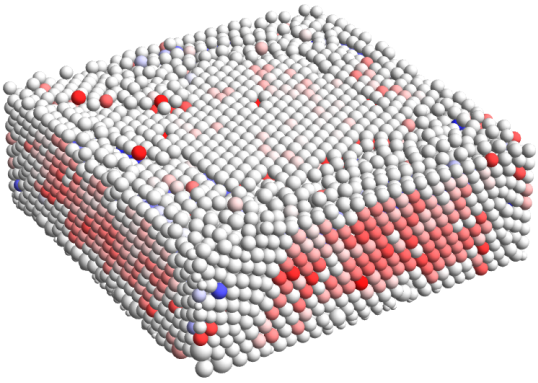

(b)

Figure 2.19: Bubble positions on (a) day 1 and (b) day 7 of the lifetime experiment. Bubbles are coloured according to the scheme in Fig. 2.18. We see that the surface bubbles are generally not classed as either fcc or hcp by our BOOP distance method.

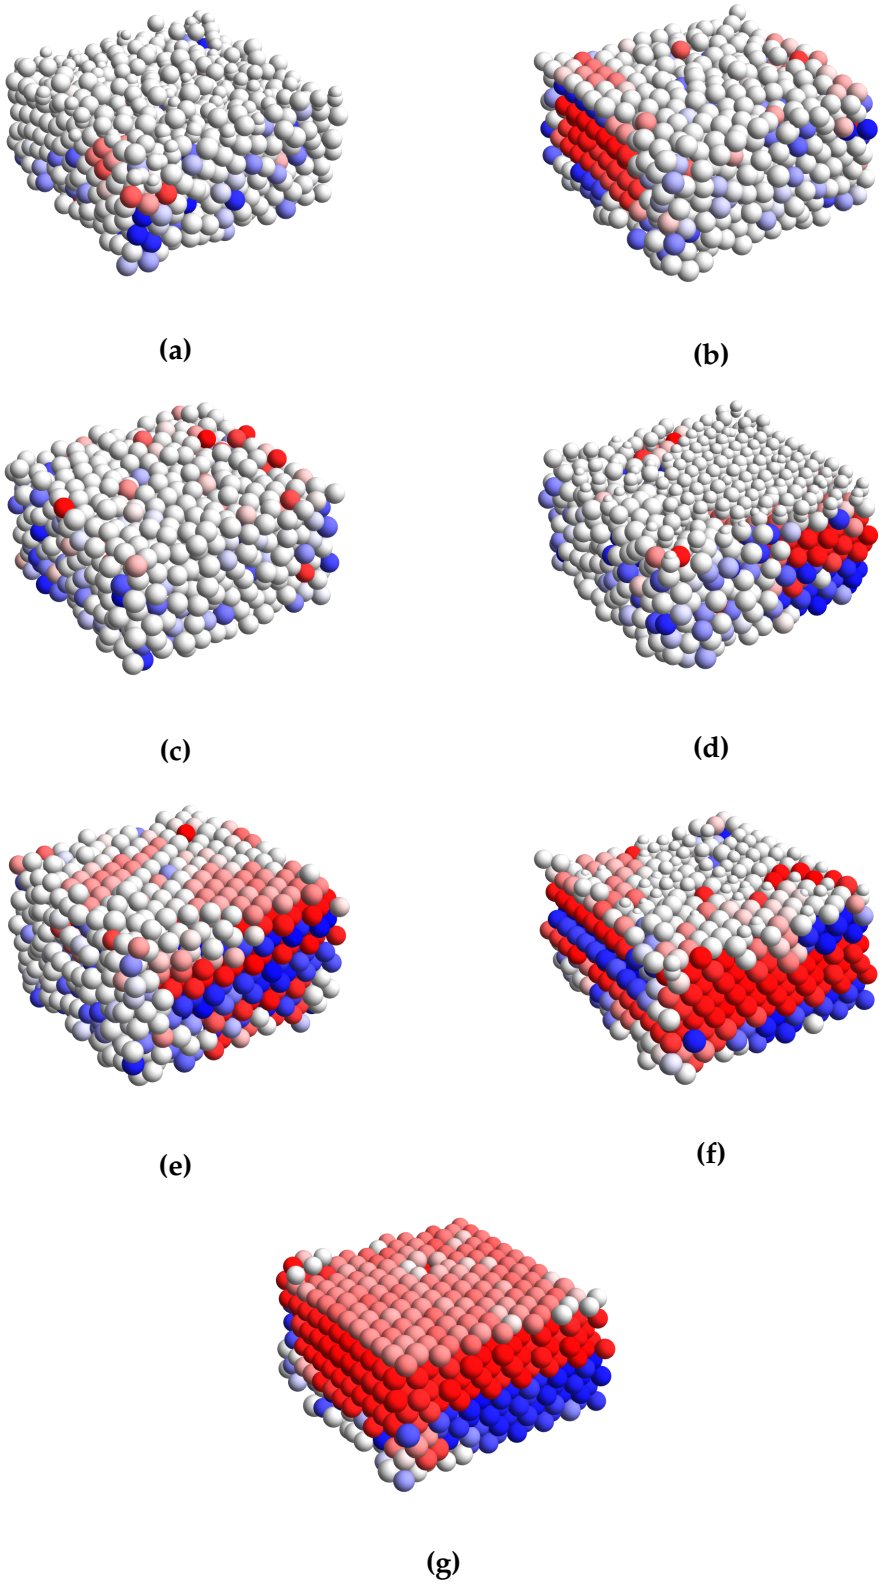

Figure 2.20: Bubble position for an excised central section of the sample on days 1 to 7 of the lifetime experiment. Bubbles are coloured according to the scheme in Fig. 2.18.

clearly divided regions of fcc and hcp are visible. We also see in this region a preference for fcc over hcp.

2.3.5 Relative prevalence of fcc and hcp

The ratio of the number of fcc bubbles to hcp bubbles on day 7, as computed using our BOOP threshold method, depends on the threshold distance r_{thresh} used. For our initial (larger) threshold we have $N_{fcc}/N_{hcp} = 1.38$. For the smaller threshold, $N_{fcc}/N_{hcp} = 1.16$. In either case, this ratio falls well short of the value of ~ 2 obtained from optical observations in [31]. However, as discussed in Section 2.1.1, these measurements were limited to $\leqslant 4$ layers of bubbles, due to the effects of refraction.

Heitkam *et al*. carried out a sophisticated simulation of small gas bubbles rising in a low-viscosity liquid [43]. In their simulation the bubbles were modelled as soft spheres which interact with one another as well as hydrodynamically with the liquid. Periodic boundary conditions were imposed in the horizontal directions. Studying 100 runs of 120 bubbles at a time, they observed N_{fcc}/N_{hcp} ranging between 1 and ~ 1.5 depending on the drainage rate imposed. Our values of N_{fcc}/N_{hcp} are in agreement with these.

2.4 The relative energy of fcc and hcp foams

2.4.1 Introduction

The preference observed in the lifetime experiment for the face-centred cubic structure over the hexagonal close-packed is in line with many other simulations and experiments [9, 31, 43, 44]. This is the case despite the fact that at $\phi = \phi_c \approx 0.26$ all energies — those arising from bubble surface areas, and buoyant energies — are equal. The reason for this preference is unclear: the relative instability of hcp pyramids on the impact of a rising bubble may play a role [43]. Woodcock [45] computes the relative entropy of simulated fcc and hcp hard sphere packings by integrating P-V isotherms around a reversible

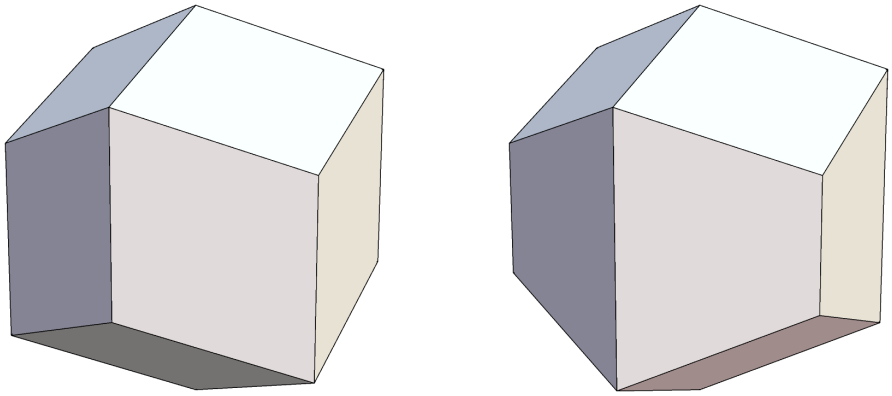

(a) **(b)**

Figure 2.21: The shape of bubbles at $\phi = 0$. (a) fcc: a rhombic dodecahedron, and (b) hcp: a trapezo-rhombic dodecahedron. All faces are flat, and all the angles between adjoining faces are 120°, satisfying Plateau's conditions for equilibrium; however, these structures are unstable equilibria due to the presence of eightfold vertices.

path linking the structures, finding a slightly higher entropy in fcc, hence a theoretical preference for fcc over hcp.

Evidently there is more at play than simply the energies of the two structures. Nevertheless we are drawn to examine the energy of each structure: here, this is equivalent to the surface area of a single bubble.

Here we will prove that it is in fact hcp which has the lower energy for values of liquid fraction ϕ *between* its limiting values: the dry limit of $\phi = 0$ (although the structures are not stable in this limit, as discussed later), and the wet limit $\phi = \phi_c$, with equality of energy at these two limits.

2.4.2 Proof

In the wet limit ($\phi = \phi_c$), bubbles are perfect spheres, so clearly the energies of the two structures are equal. In the dry limit, a bubble in an fcc foam takes the form of a *rhombic dodecahedron*, with twelve identical rhombic faces; in hcp, a *trapezo-rhombic dodecahedron*, with six rhombic and six trapezoidal faces, as in Fig. 2.21. These polyhedra have identical surface areas so the

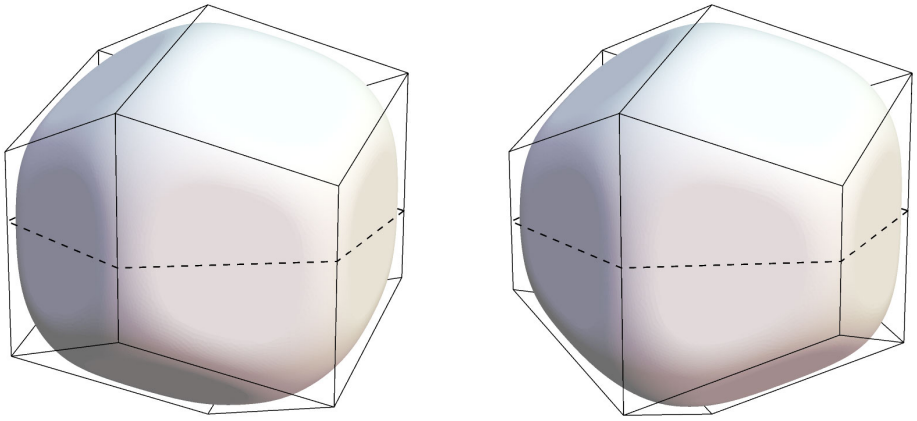

(a) (b)

Figure 2.22: (a) A bubble in an fcc foam at liquid fraction $\phi \approx 0.125$, obtained from Surface Evolver simulation. (b) By reflecting the top half of the fcc bubble in the dividing plane (marked by a dashed line), we obtain a *trial solution* for the shape of a bubble in an hcp foam.

energies are also equal at $\phi = 0$.

For all liquid fractions, we assume a canonical foam, in which the bubble surface assumes a unique shape which minimizes its surface area at constant volume. We assume that in the ordered structures discussed, each bubble possesses all the symmetries of the underlying lattice.

Fig. 2.22(a) shows the form of a bubble in equilibrium for an fcc foam between the wet and dry limits, at $\phi \approx 0.1$. Also indicated is a plane which divides the bubble into two pieces. Since the fcc lattice is centrally symmetric, any plane passing through the centre of the bubble divides its surface into two congruent pieces, equal in area and volume: we use the $(1\,1\,1)$ plane, indicated by the dashed lines in Fig. 2.22. We can reflect one half of the fcc surface (the top half in the figure) through this dividing plane, to obtain the form shown in Fig. 2.22(b). This new surface is everywhere continuous and retains the surface area and volume of fcc, but the positions of its contacts match the hcp structure. It may be regarded as a *trial solution* for hcp, and cannot therefore have a lower energy than the true hcp structure.

Indeed the energy must be *higher*, since this rejoining results in discon-

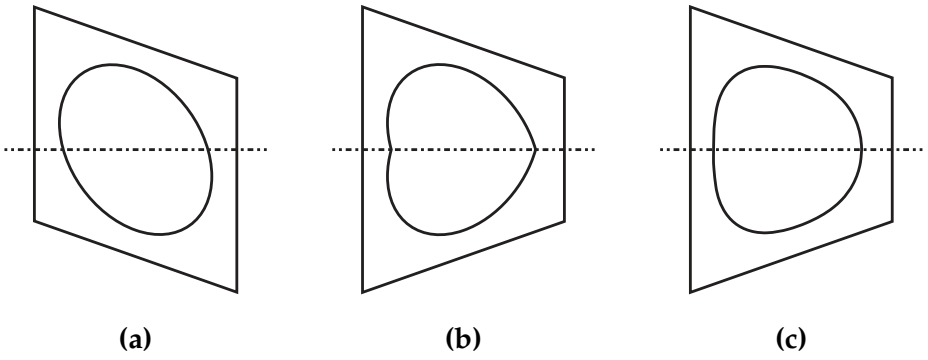

Figure 2.23: The surface obtained by reflection in Fig. 2.22(b) is continuous but not smooth: it can be further relaxed to obtain a surface of lower energy. Here (a) shows the shape of the flat contact for an fcc bubble, (b) for the trial hcp surface and (c) the further relaxed hcp surface. The size of the kink has been exaggerated here for clarity.

tinuities in the surface normal (see exaggerated sketch in Fig. 2.23). Hence the surface can be relaxed, removing these 'kinks' and lowering its energy. This completes the proof. At this point it should be noted that this difference of energies is extremely small. Experimentally it cannot be said to play a role: the effect of any experimental error would far outweigh the size of this energy difference.

It is tempting to consider a similar inverse argument, in which half of an hcp bubble is 'twisted' by 30° and rejoined to yield an fcc bubble of equal surface area. This argument fails: the intersection of the hcp bubble with the dividing plane has only threefold rotational symmetry, as we will see later, and so this process results in a discontinuous surface.

2.4.3 Axial ratio

It should be noted that, unlike the fcc structure, the hcp structure has an extra degree of freedom: its axial ratio, denoted c/a, is not fixed. c/a in this case is equal to twice the ratio of the separation of close-packed planes to the separation of neighbouring bubbles within the planes, as shown in Fig. 2.24. Some metals that crystallize in the hcp structure, such as zinc and cadmium,

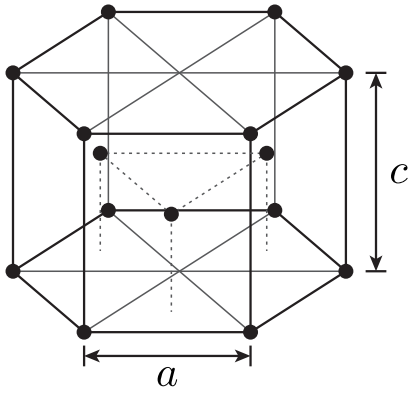

Figure 2.24: The definition of the axial ratio c/a for the hcp structure. For the ideal hcp, $c/a = \sqrt{8/3}$.

have c/a ratios which are very different from the ideal ratio [46].

We do not have this degree of freedom for the fcc structure: the analogous quantity is fixed by symmetries of the lattice.

The ideal value $c/a = \sqrt{8/3}$ is easily calculated at the wet limit by considering the coordinates of sphere centres. In the dry limit, it is easily verified that the surface area of the trapezo-rhombic dodecahedron is indeed minimized at the ideal value of c/a. For some arbitrary intermediate liquid fraction, however, it is possible that varying this value (as illustrated in exaggerated sketches in Fig. 2.25) further reduces the surface area of a bubble in an hcp foam.

Fig. 2.26 shows the excess energy ε, as defined in (1.7), as a function of c/a near its ideal value for three intermediate liquid fractions between 0 and ϕ_c: $\phi \approx 0.06, 0.13, 0.19$. We describe the axial ratio in terms of a 'stretch factor' s:

$$s = \frac{c/a}{(c/a)_{\text{ideal}}}. \tag{2.1}$$

In order for the curves to be viewable on the same plot, we shift the curves by subtracting $\varepsilon_{\text{ideal}}$ (*i.e.* ε at $s = 1$) from each.

The shapes of these plots suggest that the bubble's surface area is indeed minimized at least extremely close to the 'ideal' c/a value. It also shows that the magnitude of any possible reduction is extremely small. We note here that if varying the axial ratio at some liquid fraction $0 < \phi < \phi_c$ does indeed

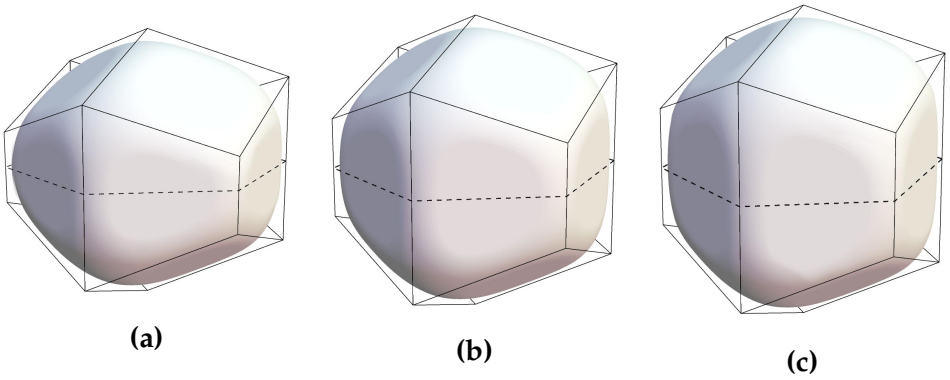

Figure 2.25: The axial ratio c/a can be varied for a bubble in an hcp foam. (a) shows a reduced axial ratio, (b) the ideal, and (c) increased. Note that these are exaggerated illustrations, and *not* properly minimized surfaces. Note also that in creating Fig. 2.26 bubble volumes were kept constant.

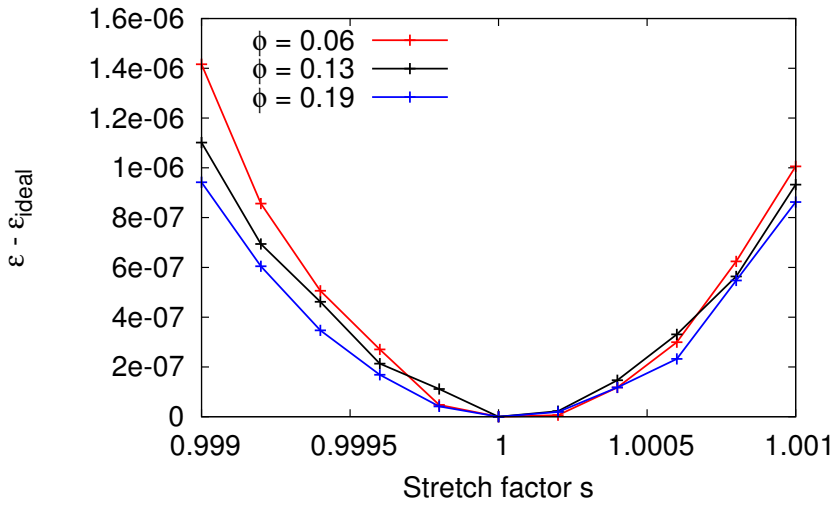

Figure 2.26: Excess energy ε as a function of axial ratio c/a for liquid fractions $\phi \approx 0.06, 0.13, 0.19$. The curves appear seem to achieve a minimum at $s = 1$, *i.e.* the ideal axial ratio. Note that there is a small amount of numerical noise visible in the curves: this is due to the very small changes in ε over the range of stretch factors considered here: we are approaching the limit of accuracy achievable using Surface Evolver.

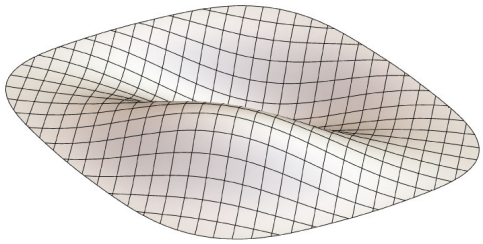

Figure 2.27: Sketch of how an hcp bubble–bubble contact could warp while retaining central symmetry. The size of the warp is greatly exaggerated here for clarity.

result in a small reduction in the surface area of the bubble compared to the ideal ratio, this only strengthens our result.

2.4.4 hcp facets

The underlying symmetries of the fcc lattice mean that every bubble–bubble contact is planar for any given liquid fraction. Bisecting the line connecting the centres of any two bubbles is a plane of reflection of the structure. Since we assume the bubble's symmetries match those of the lattice, we know that the contact is equal to its reflection and must be planar.

The hcp structure does not possess as many symmetries; by the same argument we can only guarantee that 6 of the 12 contacts must be planar (those which are trapezoidal in the dry limit). For the other 6 contacts (those which are rhombi in the dry limit) there is no reflectional symmetry; rather, a point symmetry at the midpoint of the bubble pairs. Hence we know that the contact possesses central symmetry, but is not necessarily planar: it is possible that a 'warping' of the contact reduces the energy of the foam. An exaggerated sketch of this potential warping is shown in Fig. 2.27

This is somewhat analogous to the warping of the hexagonal faces in the Kelvin foam. Plateau's laws dictate that soap films meet in threes at 120°: for the Kelvin structure, this ensures that at $\phi = 0$ the hexagonal faces must be warped. Conversely, for the hcp structure at $\phi = 0$ it ensures that the rhombi must *not* be warped! As one increases the liquid fraction of a Kelvin foam,

the warped faces 'flatten out' fairly quickly. Since the hcp facets begin flat, one may assume they remain flat as ϕ is increased. However we note that it is possible that at some liquid fraction $0 < \phi < \phi_c$ a slight curvature in the bubble–bubble contacts affords a reduction in the bubble surface area. Once again this would only strengthen our result.

We should also note that the fcc and hcp structures are unstable very close to the dry limit, due to the instability of junctions of more than four Plateau borders in the dry limit [47, 48]. Instead, the body-centred Kelvin structure tends to occur in ordered dry foams [9, 49]. This does not affect our calculations — we deal only with a single bubble in a fixed fcc or hcp environment and so no such instability arises. It is also worth noting that experimentally, mixtures of fcc and hcp arrangements are seen when close-packed planes are randomly stacked. However, in this case, each bubble can be considered locally to be either fcc or hcp, depending on the positions of its nearest neighbours.

2.4.5 Evaluation of difference using Surface Evolver

We can evaluate the difference in energies between the two structures using Surface Evolver. Fig. 2.28 is a sketch of our initial simulation method: we compress the bubble between 12 constraint planes, situated in the appropriate nearest-neighbour directions for either structure. Note that since we are using constraint planes here, we are that the rhombic contacts of the hcp bubble are planar. As discussed earlier, if it is the case that they are curved between $\phi = 0$ and $\phi = \phi_c$ this would only *increase* the surface area difference. However since our simulation cannot account for this effect, we consider it preliminary for now.

The energy difference between the two structures is small: so small that it seems to be right at the limit of what our Surface Evolver simulation can resolve. The quadratic model, described in Section 1.2.3, significantly reduces the magnitude of the numerical noise: however some amount does still re-

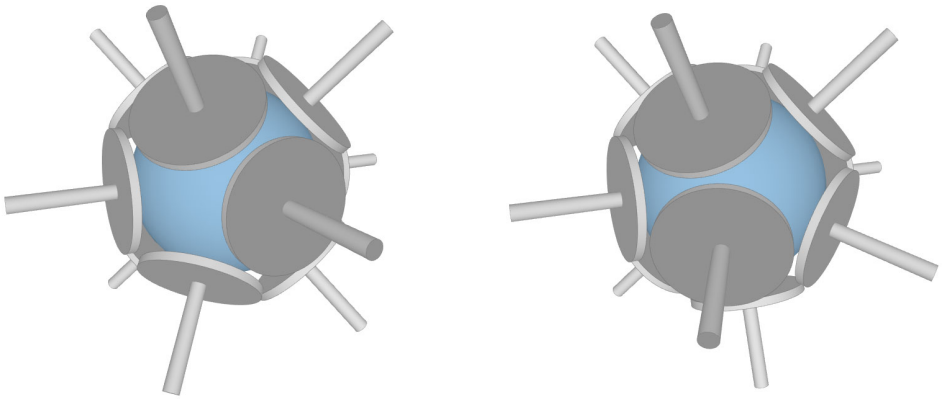

Figure 2.28: We simulate the bubble–bubble contacts in Surface Evolver by compressing the bubble between 12 constraint planes. (a) shows the plate positions for fcc; (b) the plate positions for hcp. The positions of 3 plates (lower right in these pictures) change between (a) and (b).

main. Ideally this could be reduced by using a finer mesh, but we find that even one further refinement leads to an unfeasibly long runtime.

In Fig. 2.29 we plot the difference in surface area between a bubble in an fcc foam and a bubble in an hcp foam as a function of liquid fraction, as obtained from Surface Evolver simulations. Along with our preliminary results we show results obtained by Andy Kraynik, who used a finer mesh as well as a more sophisticated simulation setup — by employing periodic boundary conditions he was able to allow nonplanar rhombic faces in the hcp case.

We see the expected result: $S_{fcc} \geqslant S_{hcp}$, with equality only at the wet and dry limits. The bubbles are of volume 1, meaning their surface area varies from ~ 4.84 in the wet limit to ~ 5.34 in the dry limit. Therefore the maximum surface area difference achieved is of the order of $\sim 10^{-5}$ relative to the total surface area. This maximum difference is achieved at a liquid fraction of $\phi \approx 0.11$.

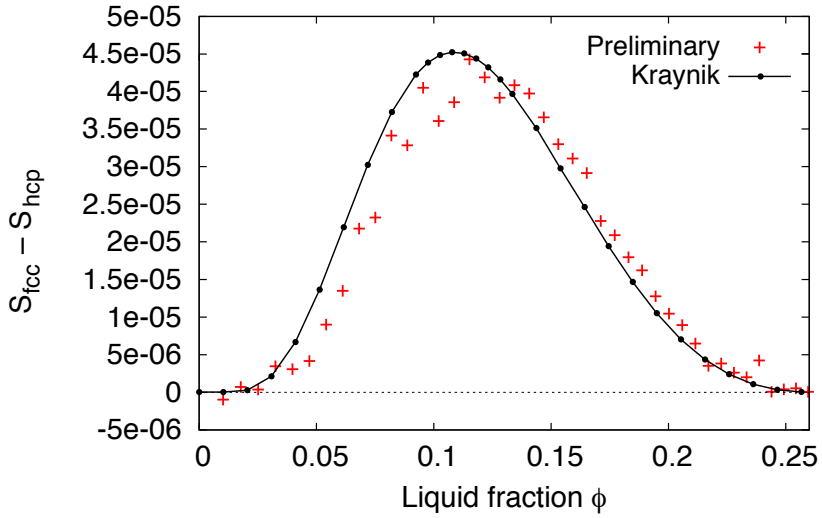

Figure 2.29: The difference in surface area between fcc and hcp bubbles, $S_{fcc} - S_{hcp}$ as a function of liquid fraction ϕ, as obtained from our preliminary simulations, as well as later simulations performed by Andy Kraynik. The difference is very small, and hence due to numerical noise in our preliminary simulations some points near the endpoints lie below 0 here.

2.4.6 Visualization of kinks

It may be easier to picture the relaxation of the kinks, as sketched in Fig. 2.23 in terms of the vertical Plateau borders surrounding the bubble. In Fig. 2.30 we sketch this same relaxation as it appears in the dividing plane.

From our Surface Evolver simulation we can obtain the shape of the intersection of the bubble surface with the dividing plane to further visualize the mechanism by which the relaxing of the kinks in the bubble surface allows our trial hcp surface to be relaxed to obtain a lower energy.

Fig. 2.31 shows such a cross section for fcc and hcp bubbles — we show only part of it, as the rest is equivalent by symmetry. While the shapes of the two figures are clearly very similar, we do indeed see that some relaxation has occurred in the hcp bubble compared with the reflected fcc bubble which we described earlier: namely, three of the vertical Plateau borders have grown in area and three have shrunk. This is manifested as the hcp bubble lying outside the fcc at one vertex and inside it at the other.

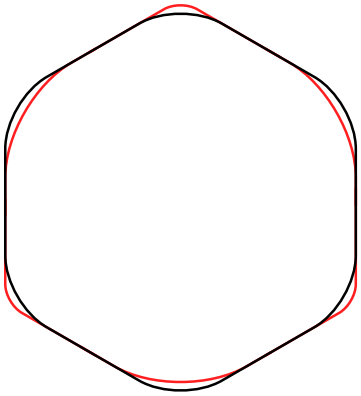

Figure 2.30: An exaggerated sketch of the intersection of the bubble with the dividing plane. The 'trial surface' for hcp, in black, is created from reflecting an fcc bubble. After relaxation of the kinks in the bubble we arrive at the true hcp surface, in red.

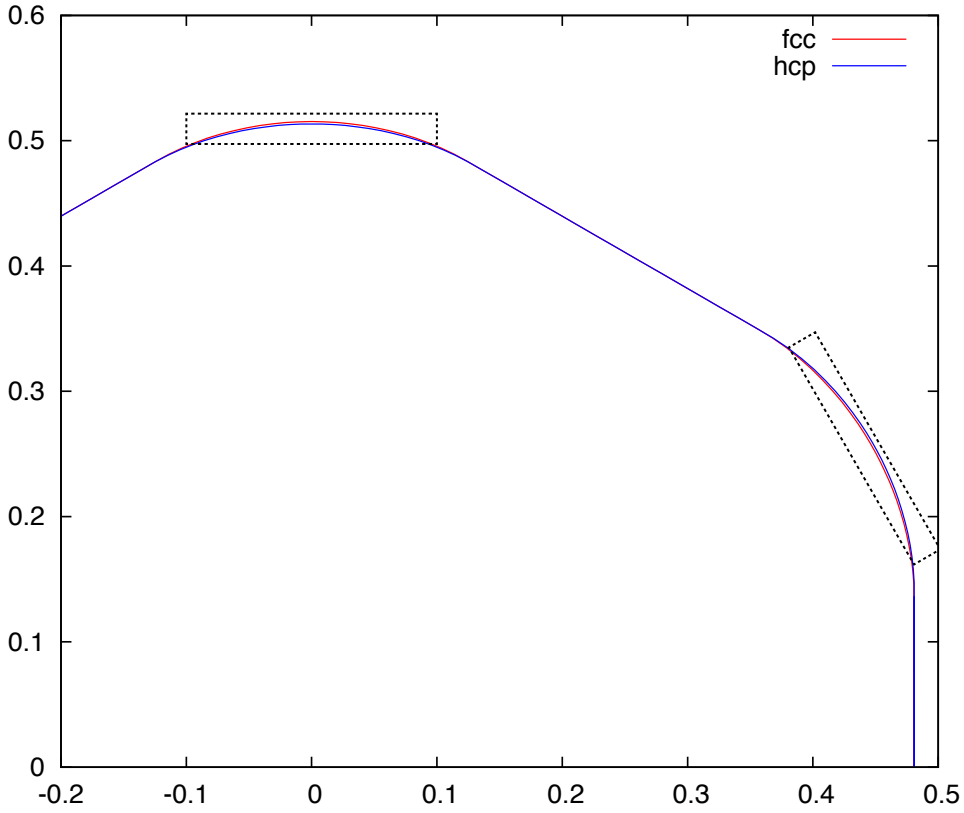

Figure 2.31: The intersection of fcc and hcp bubbles with the dividing plane (as marked in Fig. 2.22) as calculated by Surface Evolver at $\phi = 0.12$. Here the normalization is such that the undeformed bubble had diameter = 1.

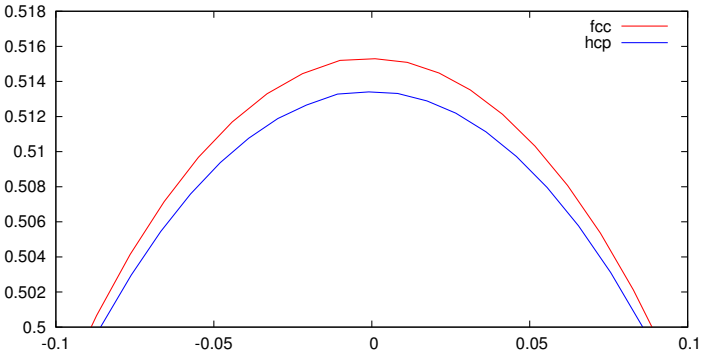

(a)

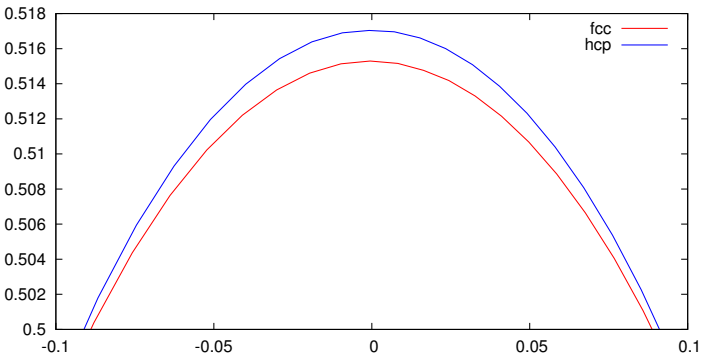

(b)

Figure 2.32: Zoomed in views of the areas marked in Fig. 2.31. In (a) we see that the relaxed hcp surface lies inside the fcc surface, and *vice versa* in (b). We have rotated (b) by $-30°$ for comparison with (a). We may estimate the energy difference between the two bubbles by considering the size of the gaps between the two surfaces. Note that at this zoom level it is visible that the surface is piecewise linear, as represented in Surface Evolver.

In Fig. 2.32 we zoom in on the two areas marked in Fig. 2.31, in which the two surfaces differ most. The behaviour sketched in Fig. 2.30 is clearly visible in these views.

2.5 Conclusion

We have proven that the preference observed in experiments and simulations for fcc ordering over hcp ordering is not due to any energetic advantage. The source of the preference is to some extent still not understood.

The argument presented by Heitkam *et al.* [43] based on the idea that hcp regions are more easily destroyed by the impact of new bubbles is compelling, but as stated seems only to apply to the case where a stream of bubbles is incident on a single point: it does not account for a layer-by-layer deposition process.

Luchnikov *et al*. [40] studied the evolution of an ensemble of hard spheres using molecular dynamics simulations. They found that an initial slight preference for fcc over hcp evolved over time into a 100% dominance of fcc. They refer to it as a *kinetic* state, which is not stable thermodynamically. Perhaps if an experiment similar to the lifetime experiment were allowed to run for a much longer time, such a trend would begin to emerge in the relative abundances of fcc and hcp.

We still cannot account for the *trend* seen in the lifetime experiment, in which the amount of ordering spontaneously increases day by day. The small amount of coarsening which occurred despite the presence of non-diffusing gas may play a role.

Chapter 3

The Cone Model

3.1 Introduction

3.1.1 Motivation

The variation of the energy, *i.e.* surface area, of a bubble as it is deformed in a foam, is vital in understanding the physics of a foam — although, as we saw in the previous chapter, energy alone cannot account for the dominant structure.

In view of this, computational methods, in particular the Surface Evolver [5], are invaluable tools. Given appropriate initial conditions, Surface Evolver can compute a minimal surface area to a high degree of accuracy. Such numerical results, however useful, perhaps fail to provide a more general understanding of the forms of bubble–bubble interactions. For this reason, we seek simpler representations: simplified models, analytic expansions and, as we will see, analytic approximations, which provide more general insights into the form of the interactions between bubbles.

Morse and Witten in 1993 [50] considered the case of a droplet (analogous to a bubble) being weakly compressed between neighbours which are arranged in an ordered manner. They found that the form of the forces involved featured logarithmic terms; however they note that the expansions are only valid for small deformations.

The *soft disk model*, formulated by Durian in 1995 [51], is an example of a model which attempts to simplify the complex interaction between bubbles.

It is a 2d model in which bubbles are represented by disks which are allowed to overlap. The repulsive force between them is harmonic, with spring constant related to the bubble radii. This greatly simplified model allows fairly large ensembles of bubbles to be simulated, and useful quantities such as shear moduli to be calculated.

In 1996 Lacasse *et al.* [52] employed Surface Evolver to obtain curves for various crystalline arrangements of droplets, and computed power-law exponents for the bubble–bubble interactions. They found that the variation of energy with bubble separation can be well approximated by a power law with exponent of 2.2; in other words, it cannot be accurately described as harmonic. In a later paper [53] they also derived an *analytic* form for the shape, and hence surface area, of a droplet in 3d being compressed between two parallel plates. They highlight the presence of logarithmic terms in the energy, as described by Morse and Witten [50], which cannot be accounted for by purely power-law type approximations.

This chapter concerns the cone model, an analytic model which attempts to provide analytic forms for the energy of a bubble as it is deformed by its neighbours. We put its direct predictions to the test by comparison with appropriate Surface Evolver simulations, as well as using the model to calculate various physical quantities, which can then be compared with appropriate experimental data.

3.1.2 Formulation

The *cone model* is a mathematical formulation which provides *analytic approximations* for the energy of a bubble as a function of in a foam as a function of liquid fraction. It was initially inspired by Ziman's 1961 work concerning the electrical properties of metals [54], in particular his treatment of the Fermi surface of copper.

The simplest case of the cone model is for a bubble which has Z planar contacts which are all equivalent [55]. We split the bubble surface into Z equal sections: one for each contact, as shown for the case of an fcc bubble

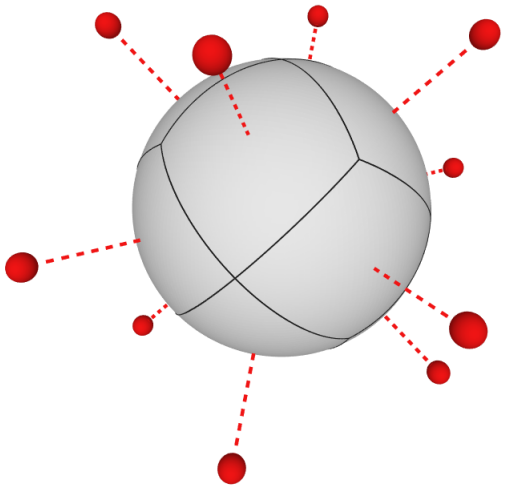

Figure 3.1: A bubble in an fcc-ordered foam. The directions of its neighbours are shown in red. The black lines on the surface divide it into 12 equivalent sections: one for each neighbour.

($Z = 12$) in Fig. 3.1. Since we have imposed that the contacts are all equivalent, minimizing the surface area of the bubble is equivalent to minimizing the surface area of any one of these sections.

We approximate each of these sections as possessing rotational symmetry around the line which connects the bubble centres. This approximation is the essence of the cone model: it reduces the complex multidimensional problem of minimization of bubble surface area to a much simpler minimization problem in one dimension. This allows us to arrive at an analytic approximation for the energy of the bubble.

Fig. 3.2 is a visualization of the approximation of the sections of bubble surface as cones. It is important to note here that if one were to place these circular cones in the positions of the original sections, they would necessarily 'overlap'. This is due to the fact that each cone possesses the same volume and solid angle as the section which it approximates — we want the entire bubble surface to be represented by our cones.

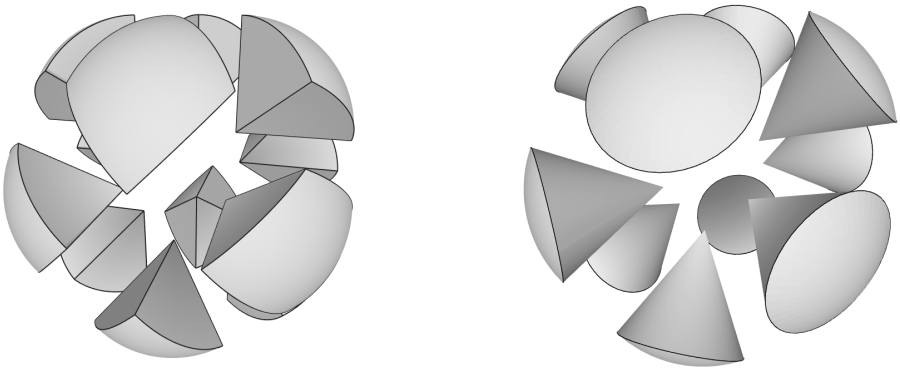

(a) **(b)**

Figure 3.2: (a) We split a bubble in an fcc foam into 12 identical pieces, as in Fig. 3.1. (b) Under the cone model, we approximate each piece as being centrally symmetric.

We leave the derivation of the cone model expressions to Appendix B, but note that we arrive at analytic forms, in terms of elliptic integrals, for the excess energy ε (as defined in (1.7)), in terms of the radius δ of the circular contact, as shown in Fig. 3.3(b).

While δ is a useful parameter in the derivation of the cone model expressions, it is not particularly practical when comparing results to simulation or experiment. Thankfully, the cone model also provides analytic expressions for the *deformation* ξ. Deformation is a dimensionless measure of bubble compression; it is defined as

$$\xi = 1 - \frac{R'}{R} \tag{3.1}$$

where R is the radius of the undeformed cone, and R' is the height of the deformed cone, as shown in Fig. 3.3. For an undeformed bubble, $R' = R$ and we recover $\xi = 0$, as we expect.

The cone model yields expressions for ξ and ε, both in terms of δ. From here we will avoid dealing with δ, and talk in terms of $\varepsilon(\xi)$.

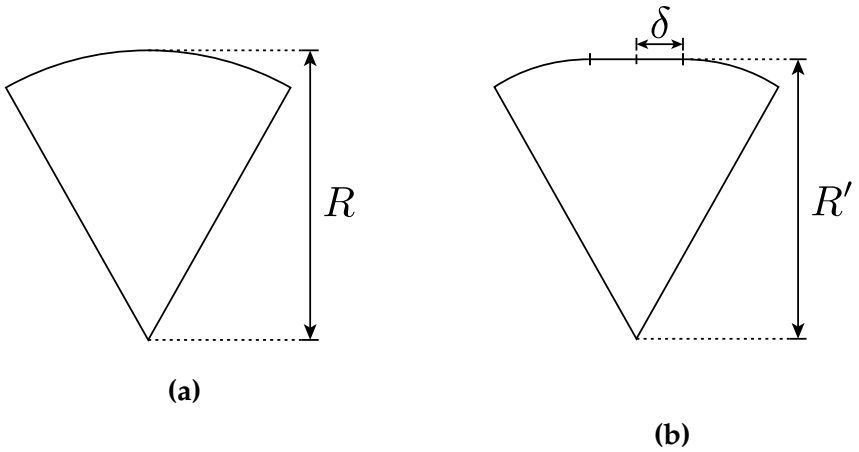

Figure 3.3: Cross sections of undeformed and deformed cones. (a) R is the radius of an undeformed cone. (b) The height R′ of a deformed cone is the distance from its apex to the contact, and δ is the radius of the contact.

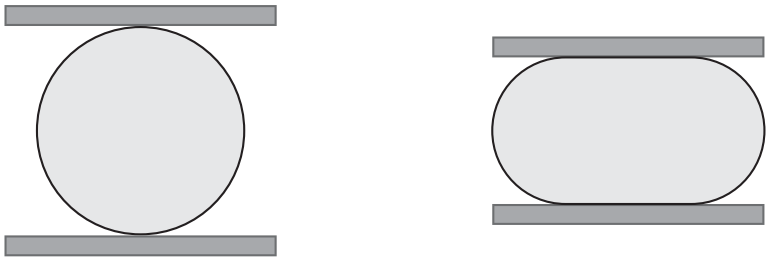

Figure 3.4: The Z = 2 case of the Z-cone model corresponds to a bubble compressed between two parallel plates, shown here in cross section. The system is rotationally symmetric and hence the cone model is exact.

3.2 Simple tests

3.2.1 Parallel plates

For Z = 2, the Z-cone model is not an approximation. A bubble with two contacts — which we visualize as 'plates' in Fig. 3.4 — consists of two equal parts, each of which is rotationally symmetric. Since the only approximation made in the derivation of the Z-cone model is that the cones are rotationally symmetric, our results in this case should be exact. Hence we have an ideal testing ground for the model: we can simulate this setup in Surface Evolver and verify that the results for $\varepsilon(\xi)$ match.

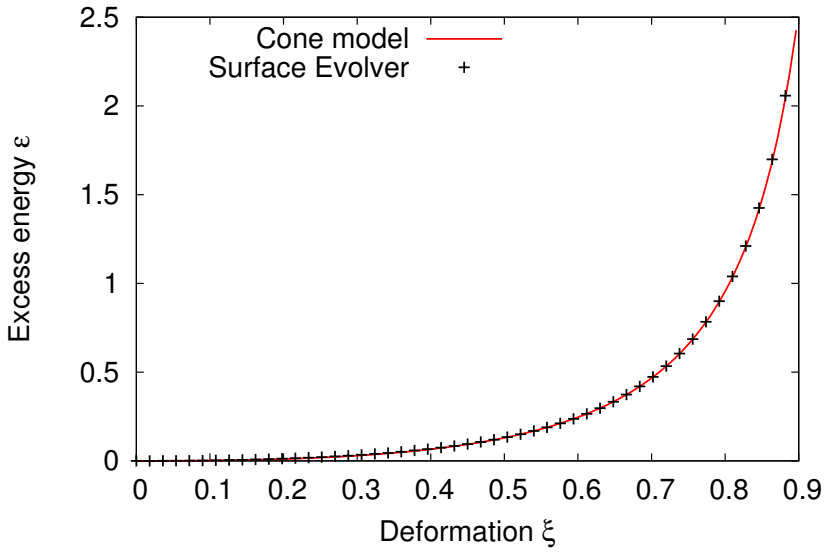

Figure 3.5: Excess energy as a function of deformation for $Z = 2$. We see agreement between Surface Evolver simulations and the Z-cone model.

In practice the 'plates' are implemented in Surface Evolver as constraint planes. To each vertex, edge and facet is added the restriction it must lie between the two planes, and the energy of the bubble is minimized subject to these additional constraints.

Fig. 3.5 shows a plot of the excess energy ε, as obtained from both the Z-cone model expression and Surface Evolver simulation results, for a very wide range of deformations ξ: note that when $\xi = 0.9$ the bubble has been flattened to one-tenth its original height. The difference between the values obtained from Surface Evolver and the model are of the order of 10^{-5}: a difference of this size may be atrributable solely to the error introduced when we approximate the smooth surface as a mesh of triangles, and so we may say that the simulations do indeed agree with the analytic predictions of the Z-cone model, as we expect. We note here that this result is the same as that of Lacasse *et al.* in 1996 [53].

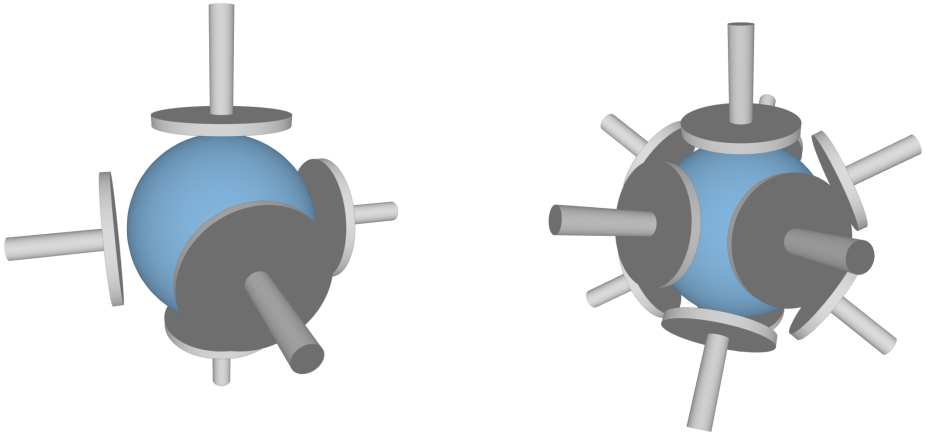

(a) (b)

Figure 3.6: (a) Plates arranged as the faces of a cube for $Z = 6$, and (b) as the faces of a regular dodecahedron for $Z = 12$.

3.2.2 Platonic solids

The Z-cone model applies to bubbles whose contacts are all equivalent. A natural next step in testing its predictions is setting up arrangements of contacts whose positions correspond to the faces of Platonic solids. Here we consider the cube ($Z = 6$) and the regular (pentagonal) dodecahedron ($Z = 12$). Again we can visualize the simulations as consisting of a bubble being compressed by plates, as illustrated in Fig. 3.6.

Figs. 3.7 and 3.8 show the results of Surface Evolver simulations as compared to the Z-cone model's analytic predictions. We see good agreement between theory and simulation over the full range of ξ — it is not exact, as in the $Z = 2$ case, since the contacts are not rotationally symmetric here.

We note that the Z-cone model is more accurate for low deformations ξ. This is to be expected, as at low ξ the contacts are very close to being perfectly circular. It is only as ξ is increased that we can no longer think of the contacts as being isolated from one another: the deformation induced by the plates spreads across the bubble surface, causing an interaction of sorts between the contacts which changes their shape to become less round.

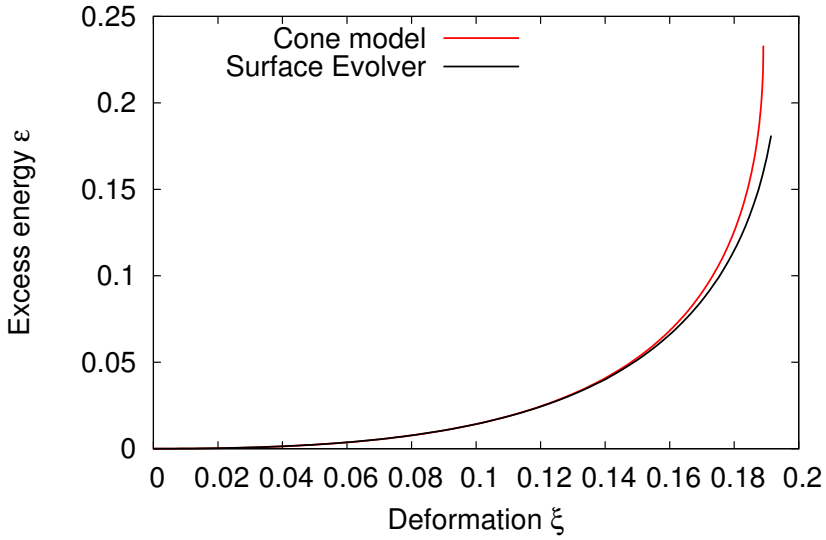

Figure 3.7: Excess energy as a function of deformation for cubic arrangement of contacts. We see that the cone model is less accurate at high deformations ξ.

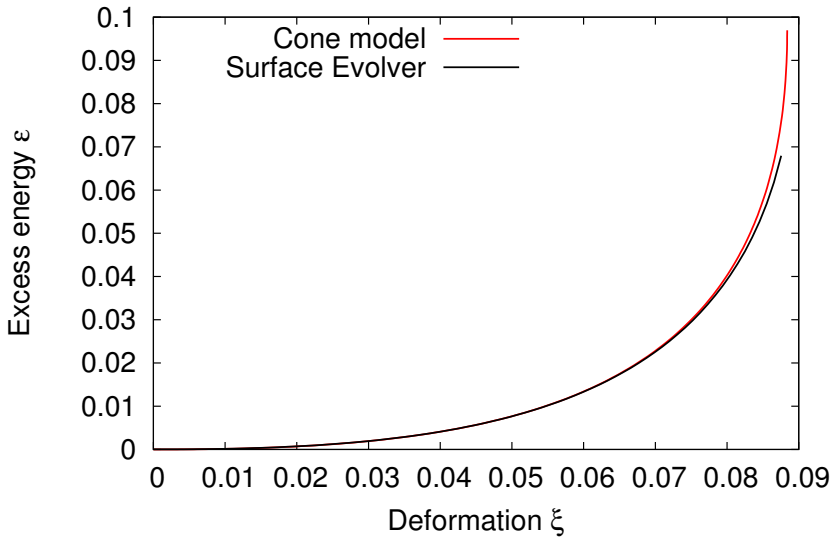

Figure 3.8: Excess energy as a function of deformation for dodecahedral arrangement of contacts. Again, the cone model is less accurate at higher deformations, however the agreement is generally better than the cubic case shown in Fig. 3.7

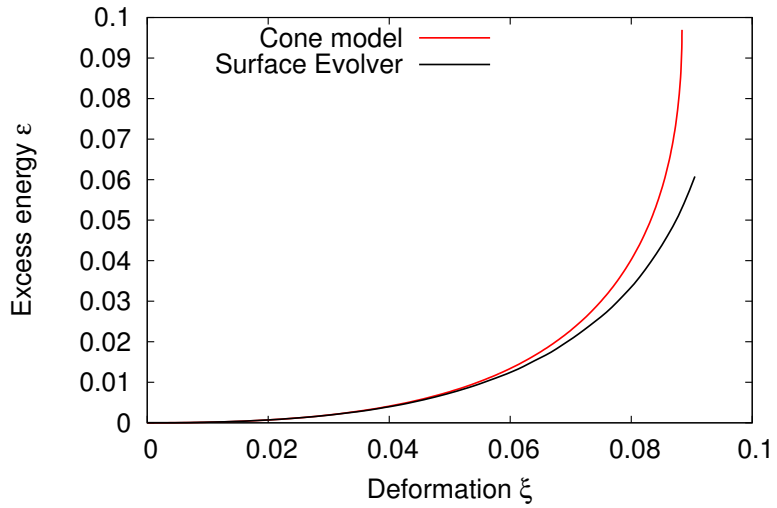

Figure 3.9: Excess energy as a function of deformation for a bubble in an fcc foam.

3.3 Application to fcc foam

3.3.1 Energy

The test cases described above are useful tests of the accuracy of the Z-cone model. However, they are not truly representative of bubbles in foams! The regular dodecahedron does not tile space, and the simple cubic foam is unstable for all liquid fractions [56].

Hence we look to the face-centred cubic structure. A bubble in an fcc-ordered foam has 12 equivalent contacts, so its energy may be estimated using the Z-cone model. Once again we simulate the bubble using constraint planes in Surface Evolver: the plate positions can be seen in Fig. 2.28(a). As we increase the deformation, *i.e.* move the plates inwards, the bubble's shape approaches that of a rhombic dodecahedron, as we saw in the previous chapter.

Fig. 3.9 shows the variation of excess energy with deformation for an fcc bubble. Once again, the Z-cone model provides a good approximation to the results of our simulation. Here the agreement is not as good as the case of the regular dodecahedron. Consider the shape of the contacts at maximum de-

formation — *i.e.*, when the bubbles are completely polyhedral — the regular dodecahedron's faces are regular pentagons, whereas the fcc bubble's faces are rhombi. The approximation of rotational symmetry is more accurate in the case of the pentagon.

3.3.2 Liquid fraction

Now that we are dealing with a bubble in a foam, it is perhaps more useful to speak in terms of liquid fraction rather than deformation. For our simulations this is simple: we consider the ratio of the bubble volume to the cell bounded by the constraint planes:

$$\phi = 1 - \frac{V_{\text{bubble}}}{V_{\text{cell}}}. \tag{3.2}$$

For the cone model, ϕ may be linked to ξ by a geometrical argument [55], yielding the relation

$$\phi = 1 - \frac{1 - \phi_c}{(1 - \xi)^3}. \tag{3.3}$$

ϕ_c is the critical liquid fraction, as before. For the cone model we obtain a value for ϕ_c by considering the volumes of the undeformed spherical cone and the circular cone which contains it. After some simple trigonometry we arrive at the approximation

$$\phi_c = \frac{3 - 4/Z}{Z - 1}. \tag{3.4}$$

Note that this gives $\phi_c = 8/33 \approx 0.24$ for a bubble with 12 contacts: this is close to the true value for an fcc foam; $\phi_c = 1 - \pi/\sqrt{18} \approx 0.26$.

Using (3.3) we can reframe the data shown in Fig. 3.9 in terms of ϕ, giving the plot shown in Fig. 3.10. We see that despite the approximation introduced to convert ξ to ϕ we still see good agreement over the full range.

3.4 Curved interfaces

3.4.1 Curved plates

So far, we have applied the cone model to monodisperse ordered foams. Bubble–bubble contacts in monodisperse foams have zero mean curvature,

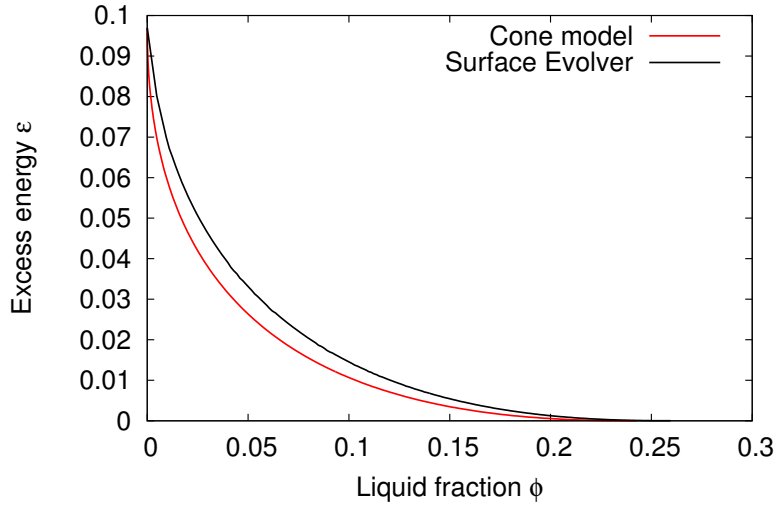

Figure 3.10: Excess energy as a function of liquid fraction for a bubble in an fcc foam.

and so under our approximation of rotational symmetry we have taken them to be planar. In general, the mean curvature of the contact is determined by the difference in pressures between the two bubbles via the Young–Laplace equation, as described in Section 1.1.8. Under the approximation of rotational symmetry, the shape of the interface is a spherical cap, with radius of curvature R_c given by

$$R_c = \frac{4\gamma}{\Delta P}. \tag{3.5}$$

Such curved interfaces are seen in *bidisperse ordered foams*, in which bubbles of two distinct sizes are carefully arranged into an ordered structure. Experimentally bidisperse foams have been produced using pyramidal containers with carefully chosen opening angles [56].

The derivations of the expressions for ε and ξ, in the Z-cone model can be adapted to deal with such a curved interface, and we arrive at analytic forms once again [57]: we leave the details to Appendix B.

We can test the validity of these new forms by simulating the simple case of $Z = 2$ in Surface Evolver. We represent the curved contacts as spherical constraints, as sketched in Fig. 3.11. Note that the setup shown in Fig. 3.11

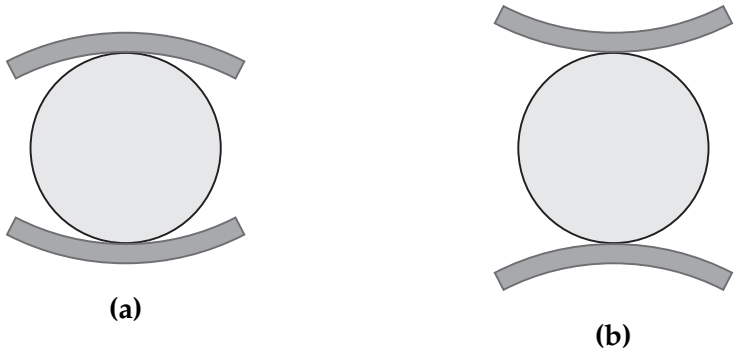

Figure 3.11: Illustration of our simulation setup for $Z = 2$ with curved interfaces: a bubble is compressed between two spherical plates of radius R_c. (a) $R_c > 0$. (b) $R_c < 0$.

(b) is unstable: the bubble can 'slip out' between the plates as they are moved inwards. For this reason we impose symmetry around the vertical axis: we only simulate one quarter of the bubble, which acts to hold the bubble in place, eliminating this possible instability.

In Fig. 3.12 we plot the variation of excess energy with deformation for two cases of curved contacts: $R_c = 3R$ and $R_c = -3R$ — the positive case being contacts which curve 'outwards', the negative 'inwards'. In both cases we see excellent agreement between the cone model theory and the results of our simulations. However we do not see *exact* agreement as we did in the case of $Z = 2$ with flat contacts (Fig. 3.5). This is due to the fact that an approximation is introduced into the definition of ξ to account for the curved contacts. Nonetheless here the difference between theory and simulation is $< 3\%$ for $R_c = 3R$ and $< 8\%$ for $R_c = -3R$ over the entire respective ranges of ξ.

3.4.2 Bidisperse cubic

Fig. 3.13(a) shows an illustration of a simple cubic structure with bidisperse bubbles: note that each small bubble has 6 contacts which are all large. The cone model with curved interfaces applies to these bubbles. Fig 3.13(b) shows a simulation of a single bubble flattened by 6 spherical caps with $R_c = 3R$. In

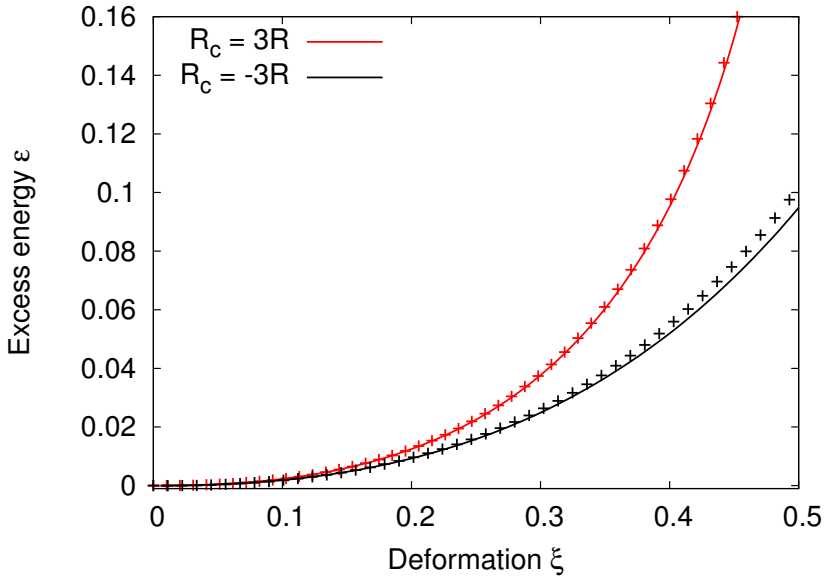

Figure 3.12: Variation of excess energy ε with deformation ξ, for a bubble with two curved contacts, as in Fig. 3.11. Solid lines are cone model expressions; points are results of Surface Evolver simulation. Over the full range of ξ, we see good agreement between theory and simulation.

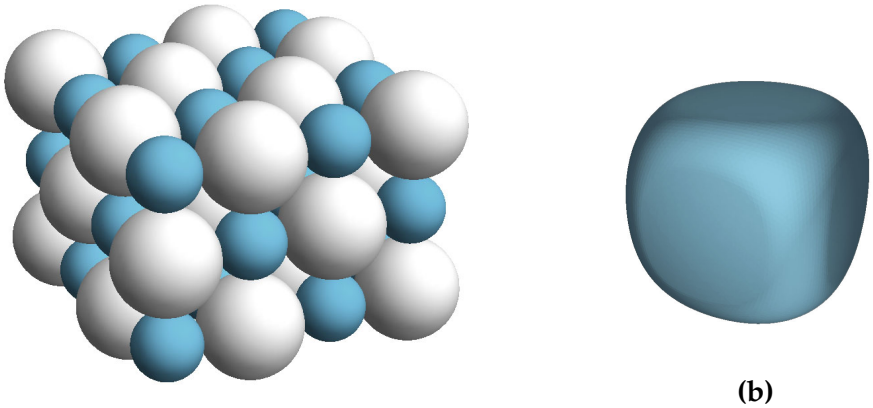

(a) **(b)**

Figure 3.13: (a) A bidisperse simple cubic structure, shown here at $\phi = \phi_c$: each large bubble has 6 small contacts and *vice versa*. (b) Simulation of a single bubble with 6 contacts, curving outwards. This represents a small bubble in (a), for $\phi < \phi_c$.

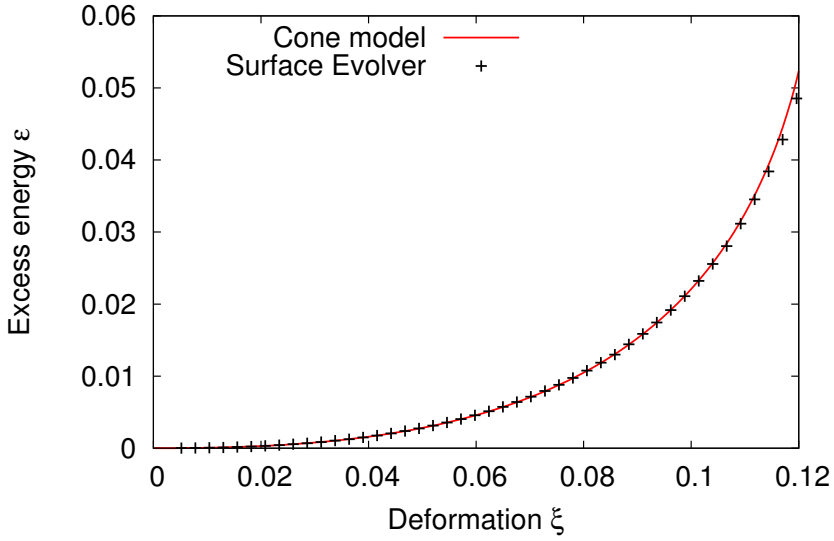

Figure 3.14: Variation of excess energy ε with deformation ξ for a bubble with 6 curved contacts, with radius of curvature $R_c = 3R$. The bubble is pictured at $\xi = 0.1$ in Fig. 3.13(b)

Fig. 3.14 we plot ε as a function of ξ, comparing the results of this simulation to the cone model predictions. We see excellent agreement over the full range of ξ.

Bidisperse foams are frequently employed in two-dimensional rheological experiments and simulations [58, 59]. They are useful in this context since monodisperse 2d foams crystallize spontaneously and hence topological rearrangements are rare. There has been relatively little work focusing on bidisperse foams in three dimensions. Treatment of 3d bidisperse foams has thus far primarily focused on dry disordered foams [60, 61].

It has been shown that it is possible, though difficult, to create ordered polydisperse wet foams using appropriate containers [56], and foams which are initially monodisperse have been seen to spontaneously evolve into a bidisperse regime at $\phi \approx 0.05$ [62]. In addition, a dynamical process has been observed in crystalline foams exposed to air, in which bubbles in the top layer repeatedly shrink and occupy the spaces between the bubbles below, forming a bidisperse top layer, before eventually disappearing completely [63]. This

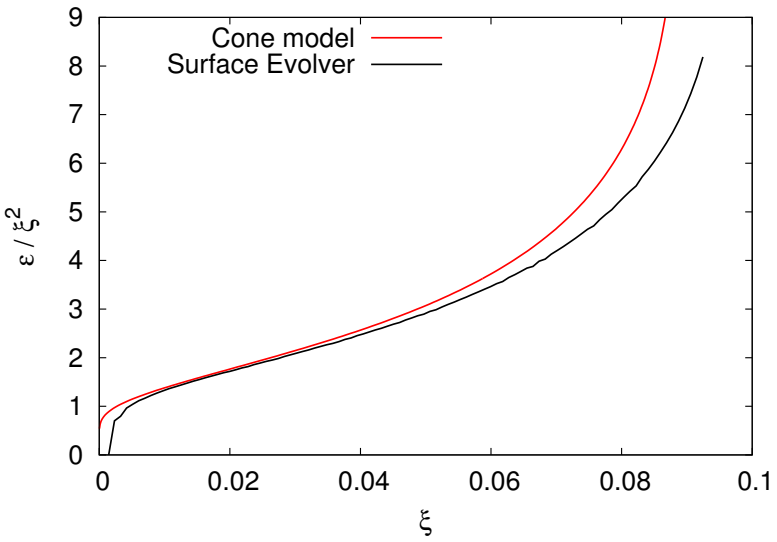

Figure 3.15: ε/ξ^2 as a function of deformation ξ for a bubble in an fcc foam: $Z = 12$. Note that near $\xi = 0$ there is visible noise in the Surface Evolver data. This is due to the fact that the small amount of numerical noise in ε is magnified when divided by very small ξ^2.

process is seen to repeat multiple times.

An analytic approach to the estimation of the energies of such wet bidisperse foams as a function of their liquid fraction may be of use in analysing the stability of such structures.

3.5 Applications of the *Z*-cone model

3.5.1 Interaction potential

The exact form of the interaction between bubbles in a foam is not fully understood. As we mentioned in Section 3.1.1, the interaction has sometimes been approximated as harmonic, *i.e.* Hookean [64]:

$$\varepsilon = k\,\xi^2, \tag{3.6}$$

for some spring constant k which depends somehow on the physical properties of the foam. Indeed, the curve obtained for the fcc case of $Z = 12$, shown in Fig. 3.9, looks to the eye like a quadratic curve.

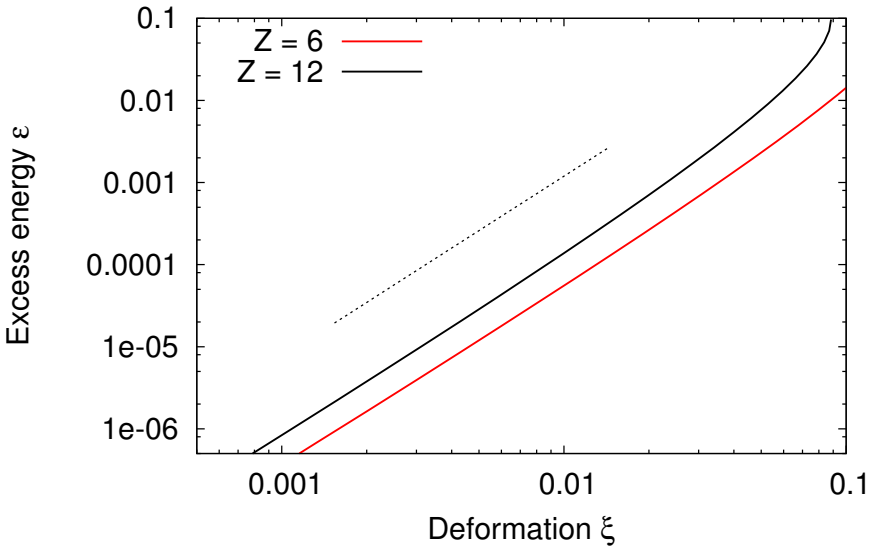

Figure 3.16: log–log plot of ε/ξ^2 as a function of deformation as obtained, as obtained from the Z-cone model for $Z = 6$ and $Z = 12$. A line of slope 2.2 is plotted as a guide to the eye. For intermediate values of Z we see similar slopes; they are not included here for the sake of clarity.

In Fig. 3.15, we plot ε/ξ^2 *vs.* ξ. If the interaction predicted by the model were truly harmonic, ε/ξ^2 would equal k for all values of ξ, *i.e.* a horizontal line on this plot.

Clearly the behaviour of bubbles under the cone model cannot be accurately described as a Hookean interaction.

In Fig. 3.16 we plot ε as a function of ξ, on a log–log plot for a range of values of Z. This confirms that an exponent of 2 is not a satisfactory description of the cone–cone interaction. The exponent which best fits the data is $\alpha \approx 2.2$, in line with results from Lacasse *et al.*, who reported exponents α ranging from 2.1 to 2.5, depending on Z. Interestingly, under the cone model, the exponent does not seem to depend on Z to any significant degree.

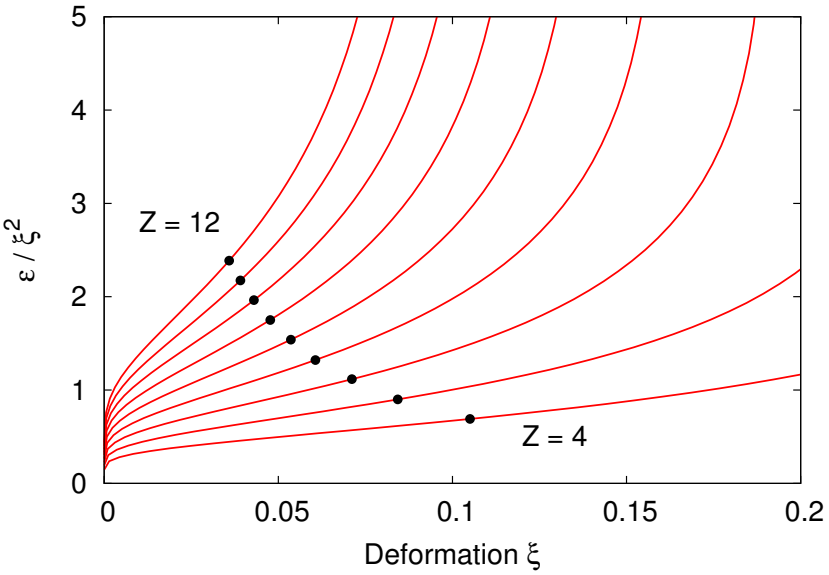

Figure 3.17: The inflection point on a plot of ε/ξ^2 *vs.* ξ may be considered the point at which the graph is flattest: we refer to the value of ε/ξ^2 at this point as an effective spring constant k_{eff}.

3.5.2 Spring constants

While we have shown that a harmonic model does not fully describe the variation of energy with deformation for the Z-cone model, it has been used by others in the past, due at least in part to its computational efficiency. For example, Durian approximates bubble–bubble interactions as harmonic in the bubble model [65]. For this reason we compute an *effective spring constant* k_{eff} for the interaction.

Fig. 3.17 shows ε/ξ^2 as a function of ξ for a range of coordination numbers $Z = 4, 5, \ldots 12$. On each curve we mark the inflection point: that point at which the second derivative is zero, and hence at which the slope is at its minimum. This may be considered the point at which the curve is 'flattest', and hence at which the approximation of harmonicity is best, at least locally.

By evaluating ξ/ε^2 at the inflection point for each value of Z we obtain $k_{eff}(Z)$, which we plot in Fig. 3.18. We see a relationship which is very close

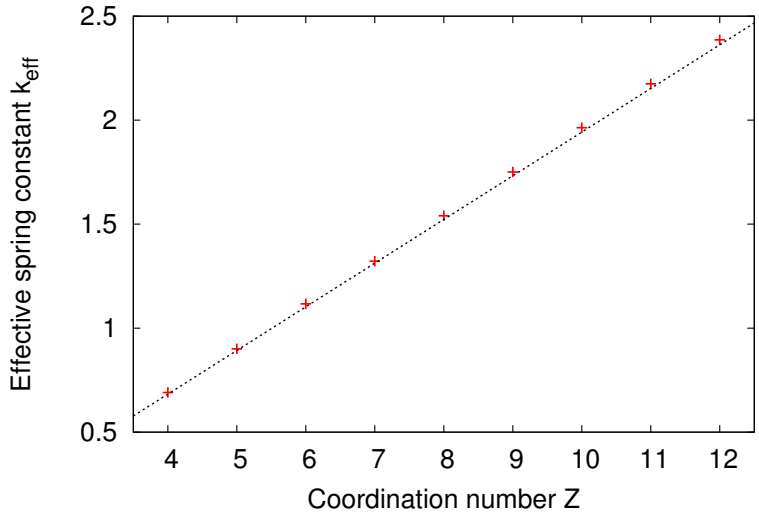

Figure 3.18: The effective spring constants for different coordination numbers, as computed in Fig. 3.17: k_{eff} increases linearly with Z.

to linear. The line of best fit is:

$$k_{eff} = 0.21(Z - 0.75). \tag{3.7}$$

This means that as the coordination number of a bubble is increased, the interaction with its neighbours becomes 'stiffer'.

Lacasse *et al.* [53] suggest that bubble–bubble interactions may be pictured as a harmonic interaction with effective spring constant disappearing as $Z \to 0$. We note that an alternative fit to the data,

$$k_{eff} = 0.15Z^{1.11} \tag{3.8}$$

does vanish at $Z = 0$, and is an equally good fit: the lines are not distinguishable over the range shown in Fig. 3.18.

3.5.3 Osmotic pressure

We now turn to the osmotic pressure Π, as defined in Section 1.1.11. It is possible to measure the osmotic pressure of a foam experimentally, in contrast to the variables which the cone model directly predicts (*e.g.* the surface area of a bubble within a foam, *etc.*).

Dimensional analysis reveals that Π scales as γ/R, where R is the radius of a bubble. Hence we deal instead with the *reduced osmotic pressure*, $\widetilde{\Pi}$, defined as

$$\widetilde{\Pi} = \frac{\Pi}{\gamma/R}. \tag{3.9}$$

The reduced osmotic pressure is a dimensionless quantity which depends not on the physical properties of the foam itself, but more generally on the geometric effect of increasing a bubble's surface area. Höhler *et al.* [9] studied the osmotic pressure of ordered foams in experiments as well as simulations, and arrived at an empirical expression which describes the behaviour of $\widetilde{\Pi}$ over a wide range of liquid fractions:

$$\widetilde{\Pi}(\phi) = k\frac{(\phi - \phi_c)^2}{\sqrt{\phi}}, \tag{3.10}$$

where k is a constant which depends on the structure of the foam. For an ordered foam, $k = 7.3$ and $\phi_c \approx 0.26$.

We can derive an expression for the osmotic pressure as a function from the cone model by considering a convenient definition of osmotic pressure, namely

$$\Pi = -\left(\frac{\partial E}{\partial V}\right)_{V_g}. \tag{3.11}$$

In other words, the osmotic pressure is the partial derivative of the total energy of the foam with respect to its volume, when the volume of gas V_g is kept constant. This can be rewritten in terms of excess energy and liquid fraction as

$$\widetilde{\Pi}(\phi) = -3(1-\phi)^2\frac{\partial \varepsilon}{\partial \phi}. \tag{3.12}$$

In Fig. 3.19 we plot the reduced osmotic pressure $\widetilde{\Pi}$ as obtained from (3.12) using $\varepsilon(\phi)$ from the Z-cone model with $Z = 12$, along with the empirical form from (3.10). As $\phi \to 0$, the form obtained from the cone model diverges quicker than the empirical form, and so the difference between the two curves too diverges. However for $\phi \geqslant 0.01$, the difference between the two values does not exceed 0.12.

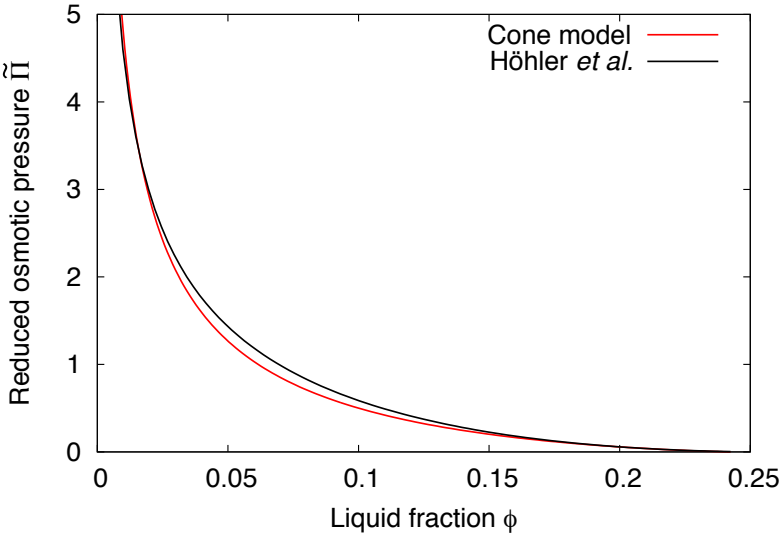

Figure 3.19: The reduced osmotic pressure $\widetilde{\Pi}$, as derived from the cone model expressions for excess energy, compared with the empirical form (3.10) from [9]. We see good agreement over a wide range of liquid fractions.

3.5.4 Liquid fraction profile

We can go one step further with our expression for reduced osmotic pressure, and use it to derive a liquid fraction profile for an ordered foam.

Here we speak in terms of the *reduced height* $\tilde{x}$ in the foam, defined as

$$\tilde{x} = \frac{x}{{l_0}^2/R}, \tag{3.13}$$

where x is the height as measured from the foam–liquid interface (giving $\phi = \phi_c$ at $\tilde{x} = 0$), R is the bubble radius[1] and l_0 is the capillary length of the liquid; ${l_0}^2 = \gamma/(\rho g)$; γ being surface tension, ρ the density of the liquid and g the acceleration due to gravity, as described in Section 1.1.9.

The local liquid fraction at a height $\tilde{x}$ in a foam may be related to the osmotic pressure by

$$d\widetilde{\Pi} = (1 - \phi(\tilde{x}))d\tilde{x}. \tag{3.14}$$

Expanding (3.14) into partial derivatives yields

$$\frac{\partial\phi}{\partial\tilde{x}} = \frac{1 - \phi(\tilde{x})}{\partial\widetilde{\Pi}/\partial\phi}, \tag{3.15}$$

[1]here we are considering a monodisperse foam: in the case of polydisperse foams one uses the *Sauter mean radius* R_{32}, defined as $\langle R^3\rangle/\langle R^2\rangle$.

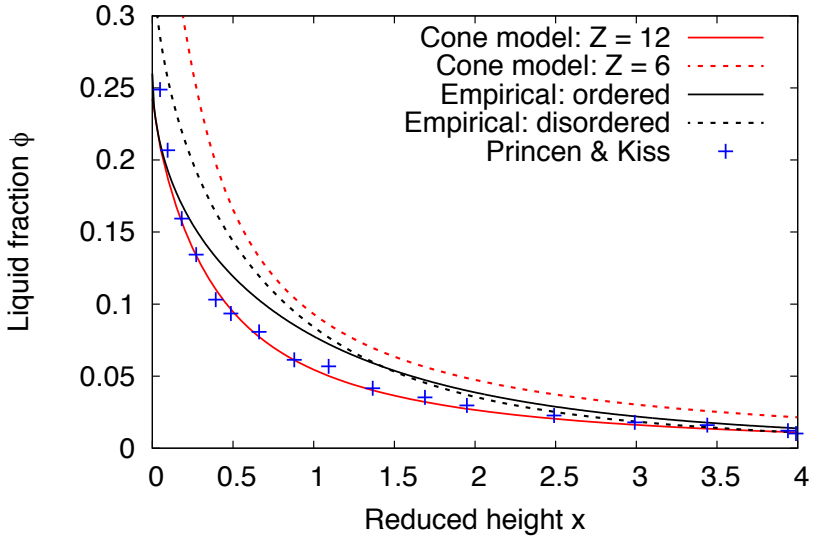

Figure 3.20: The variation of liquid fraction ϕ with reduced height $\tilde{x}$. The curve obtained from the Z-cone model, with $Z = 12$, agrees well with an empirically obtained curve for ordered foams, as well as with experimental data from an emulsion of oil in water.

with the boundary condition that $\phi(0) = \phi_c$.

As before, we can use the Z-cone model to obtain an expression for $\widetilde{\Pi}(\phi)$ via (3.12), and hence an expression for $\partial\widetilde{\Pi}/\partial\phi$, as required in the denominator of (3.15). Therefore we can integrate this up from $\tilde{x} = 0$ to yield a full liquid fraction profile of the foam.

As a means of comparison we turn again to the empirical form of Höhler *et al.*. One can derive from (3.10) a somewhat unwieldy expression for the reduced height $\tilde{x}$ in terms of the liquid fraction ϕ:

$$\begin{aligned}\tilde{x} = k\Bigg[&(\sqrt{\phi_c} - \sqrt{\phi})\left(3 + \frac{\sqrt{\phi_c^3}}{\sqrt{\phi}}\right) + \\ &+\frac{1}{2}(3 - 2\phi_c - \phi_c^2)\log\left(\frac{(\sqrt{\phi}+1)(\sqrt{\phi_c}-1)}{(\sqrt{\phi}-1)(\sqrt{\phi_c}+1)}\right)\Bigg],\end{aligned} \tag{3.16}$$

where again we use the parameters $k = 7.3$ and $\phi_c = 0.26$ for an ordered foam.

In Fig. 3.20 we plot ϕ as a function of $\tilde{x}$ as obtained from the Z-cone model,

compared with the empirically derived expression of Eq. (3.16). We use the parameters previously mentioned for an ordered foam, as well as those obtained by Maestro *et al.* for polydisperse disordered foams [66]. In addition, we plot results obtined for polydisperse emulsions of paraffin oil in water, as measured directly by Princen and Kiss [67].

We keep Z constant for the sake of simplicity and to obtain analytic expressions; we note that in a real foam $\langle Z \rangle$ decreases steadily as we increase ϕ, as shown by Jorjadze *et al.* [64]. Despite this we see good agreement for $Z = 12$. For comparison we show the curve obtained by setting $Z = 6$ in the Z-cone model, as is the case for a disordered monodisperse foam at the wet limit. As one would expect this overestimates the empirical and experimental values.

3.5.5 Surface liquid fraction

It is simple to measure the average liquid fraction of a foam sample experimentally. If a known volume of liquid is used in generating the foam, then the volume of the foam can be read off *e.g.* a graduated cylinder, immediately yielding an average value of ϕ.

There exist several methods for measuring the *local* liquid fraction at a certain height. These include measurement of electrical conductivity across the foam using electrodes [68], measurement of optical scattering [69], and even measurement of X-ray transmission [33, 70]. These methods are much more awkward than the straightforward determination of average liquid fraction.

The *surface liquid fraction*, denoted ϕ_s, of a foam in a container is the fraction of the walls of the container which is coated by liquid. This quantity can be measured experimentally by imaging the surface of a container of foam using a telecentric lens and a 45° prism [71]. This method produces clear images of very high contrast from which ϕ_c is readily obtained through simple image binarization, with errors of $< 5\%$.

Since surface liquid fraction is easier to measure than local liquid fraction, an expression relating ϕ with ϕ_s is of great use experimentally. Via consid-

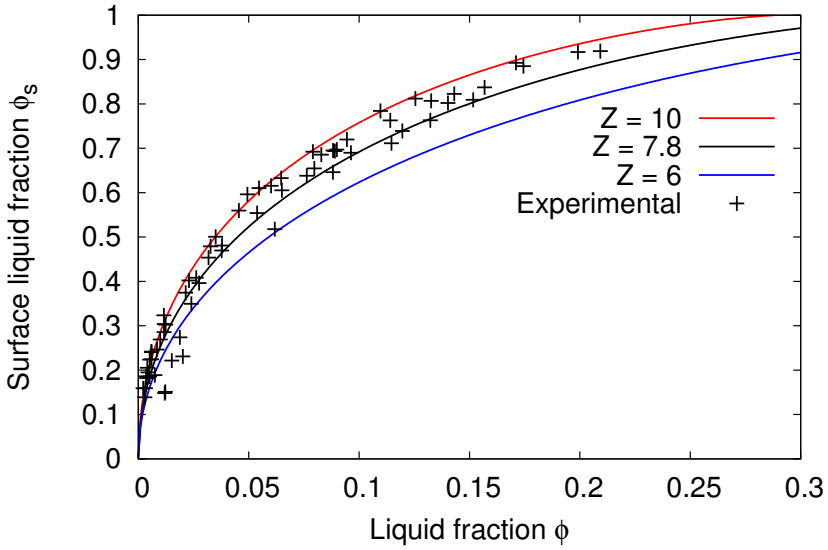

Figure 3.21: Variation of surface liquid fraction ϕ_s with liquid fraction ϕ, as derived from the Z-cone model for various values of Z, compared with experimental data from Drenckhan *et al.*

eration of osmotic and capillary pressures in the foam, one arrives at such an expression [72]:

$$\frac{1-\phi}{1-\phi_s} = 1 + \frac{2}{3}\frac{E}{\Pi V}. \tag{3.17}$$

As noted in Section 1.1.11, $\Pi \to 0$ in the wet limit $\phi \to \phi_c$. Hence in the wet limit, $\phi_s \to 1$: the surface is completely covered by liquid. This is intuitively understood when one notes that bubbles are spheres at ϕ_c, and hence any contact between a bubble and the surface consists of a single point, yielding a total area of 0.

(3.17) can be nondimensionalized to give

$$\frac{1-\phi}{1-\phi_s(\phi)} = 1 + \frac{2(1+\varepsilon(\phi))}{\widetilde{\Pi}(\phi)}. \tag{3.18}$$

We see that the relationship between ϕ and ϕ_c depends only on $\varepsilon(\phi)$ and $\widetilde{\Pi}(\phi)$ — these are both quantities which can be provided by the cone model, and so we can obtain a curve relating ϕ_c to ϕ analytically.

The question remains as to what value of Z is appropriate. The contact number Z of a bubble in a foam depends on the liquid fraction and the poly-

dispersity, as well as the degree of ordering in the foam, and can vary from 6, for a random wet foam, to 13.4 for an idealized dry foam [20].

In Fig. 3.21 we plot ϕ_s as a function of ϕ as obtained from the Z-cone model for various values of Z, along with experimental data from Drenckhan *et al.* [72]. Least-squares fitting yields $Z = 7.8$ as the best fit to the experimental data: this makes some sense intuitively, as it lies between these upper and lower limits. It is worth noting that the experimental data covers a range of dispersities, so it is perhaps surprising that the Z-cone curve is satisfactory to describe the data. This may be explained by the fact that for an ordered foam, the form of Π is relatively independent of what type of structure is present [19]. It should be noted, however, that ordered and disordered foams have distinct osmotic pressures at higher liquid fractions since they must approach 0 at different values of ϕ.

3.6 Kelvin foam

3.6.1 Motivation

As discussed in Section 1.1.5, ordered foams below a liquid fraction $\phi \lesssim 0.1$ do not tend to have bubbles arranged in the face-centred cubic arrangement. In the dry regime, it is the *body-centred cubic* (bcc) arrangement, also called the *Kelvin foam*, which is energetically favourable. The Z-cone model, as described above, is *not* directly applicable to the bcc structure, for two reasons.

Firsly, we note that in a bcc foam at $\phi = 0$, each bubble has 14 contacts: 8 nearest neighbours, in the $\langle 1\,1\,1 \rangle$ directions[2], and 6 next-to-nearest neighbours, in the $\langle 1\,0\,0 \rangle$ directions. The distance to next-to-nearest neighbours is $2/\sqrt{3} \approx 1.15\times$ the distance to nearest neighbours. These directions and distances are shown in Fig. 3.22(a). Fig. 3.22(b) shows the shape of a bubble at $\phi = 0$. Clearly the contacts are not all equivalent, and so already we have violated one of the requirements for validity of the Z-cone model.

[2]Here we are using the convention that $\langle h\,k\,\ell \rangle$ refers to the direction (h, k, ℓ) in terms of the *direct* unit vectors, as well as all directions which are equivalent by symmetry [73].

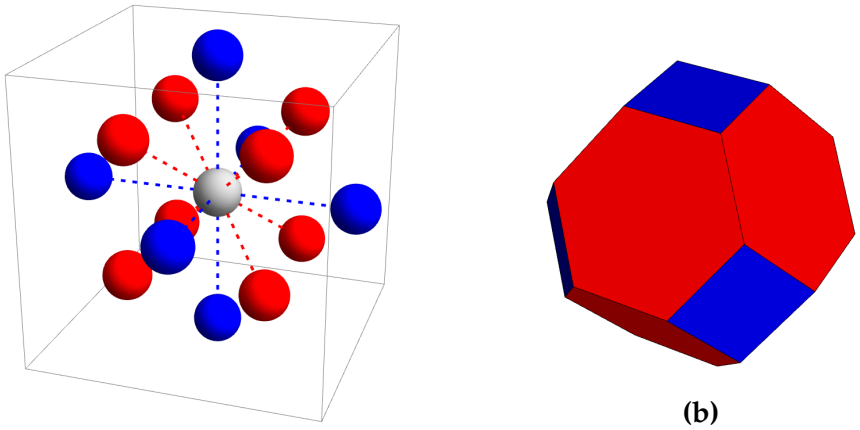

(a) (b)

Figure 3.22: (a) A bubble in a Kelvin foam has 8 nearest neighbours in the $\langle 1\,1\,1 \rangle$ directions (red) and 6 next-to-nearest neighbours, in the $\langle 1\,0\,0 \rangle$ directions (blue). (b) In the dry limit, the nearest neighbours correspond to the hexagonal faces and the next-to-nearest neighbours to the square faces.

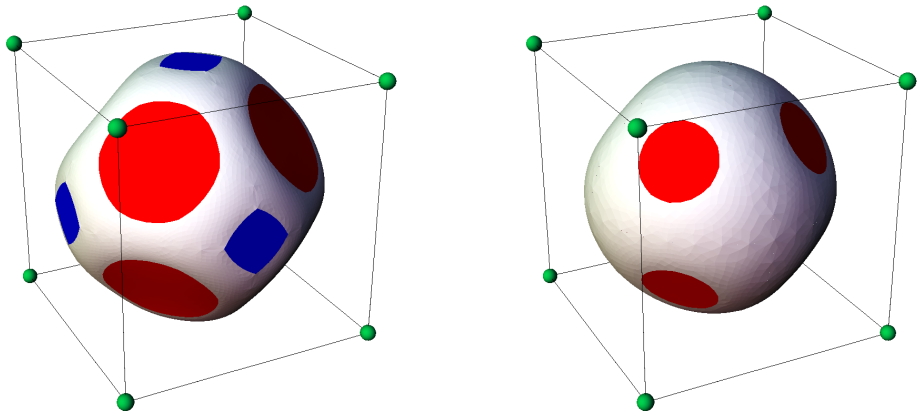

(a) (b)

Figure 3.23: Surface Evolver simulations of equilibrium bubble shapes in a wet bcc foam; centres of neighbouring bubbles are marked. (a) For liquid fraction $\phi < \phi^*$ ($\phi = 0.05$ here) there are two sets of contacts, corresponding to the (dry) hexagonal $\langle 1\,1\,1 \rangle$ and square $\langle 1\,0\,0 \rangle$ faces. (b) When ϕ exceeds ϕ^* ($\phi = 0.15$ here) the square contacts are lost.

Secondly, we note that when the liquid fraction is increased above a certain point, the $\langle 1\,0\,0 \rangle$ contacts are lost, leaving the bubble with only 8 of its 14 contacts. We will denote this intermediate critical liquid fraction by ϕ^* here. Fig. 3.23 shows examples, created using Surface Evolver, of a bubble in a bcc foam for the two cases of $\phi < \phi^*$ and $\phi > \phi^*$. So not only are the contacts not all equivalent for a bcc bubble, we find that Z is not constant! This loss of contacts is itself a motivation for attempting to treat the bcc foam analytically: it is a well-defined structure that can be used for the study of a general feature of foams: namely, the gain or loss of a face at some critical liquid fraction.

From Surface Evolver we find $\phi^* \approx 0.107$. This corresponds roughly to the point at which the bcc structure is no longer stable for an ordered foam: Weaire *et al.* first suggested in 1993 [47] that this loss of next-to-nearest-neighbour contacts may be the mechanism by which the bcc foam becomes unstable. Here we discuss the energy of a bcc-ordered foam over the full range of liquid fractions $0 \leqslant \phi \leqslant \phi_c$, neglecting this instability, just as we did for fcc.

Part of our motivation concerns the details of the variation of energy close to the liquid fraction at which contact is lost with the six neighbours in the $\langle 1\,0\,0 \rangle$ directions. Such subtle questions are difficult to pursue with Surface Evolver simulations, due to the high level of precision required, and it was expected that the extended cone model would shed some light on the matter. Among other things, it should bear on the precise position and nature of the instability associated with the loss of $\langle 1\,0\,0 \rangle$ contacts.

3.6.2 Direct application of *Z*-cone model

Despite these difficulties, it is worth investigating how well the Z-cone model can approximate the energy of a bcc foam, if at all.

At $\phi < \phi^*$ a bubble in a bcc foam has 14 contacts; at $\phi > \phi^*$ it has 8. Hence one might assume that using $Z = 14$ or $Z = 8$ in the Z-cone model may approximate the energy of such a bubble over some range. In Fig. 3.24 we plot

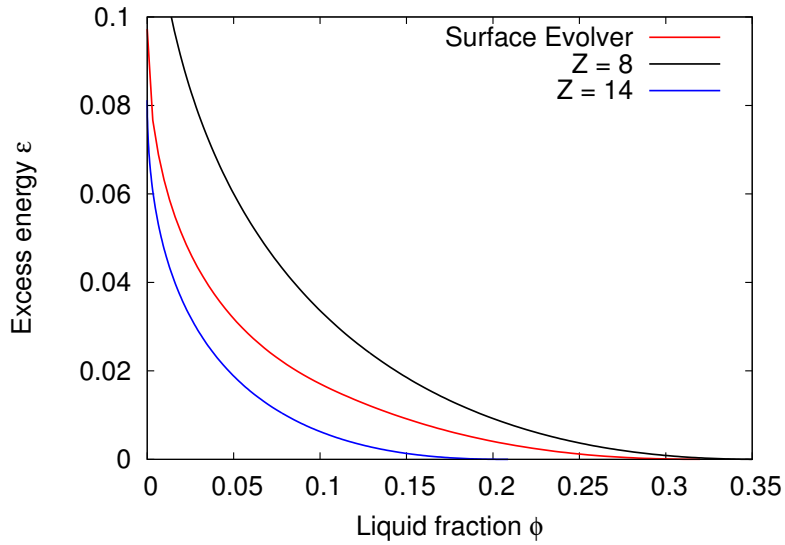

Figure 3.24: Applying the Z-cone model to a bubble in a bcc foam. Evidently, neither $Z = 8$ or $Z = 14$ captures the variation of energy with liquid fraction adequately over any appreciable range. The energy does however lie closer to the $Z = 14$ curve in the dry limit and the $Z = 8$ curve in the wet limit, as one might expect.

ε *vs.* ϕ as obtained from a Surface Evolver simulation, as well as the curves obtained from the Z-cone model for $Z = 8$ and $Z = 14$. Evidently, neither value of Z adequately describes the form of the energy, although broadly one could say that the true curve lies closer to the $Z = 14$ curve in the dry limit and closer to the $Z = 8$ curve in the wet limit, in line with what one would expect to observe.

In fact, the value of Z which best approximates $\varepsilon(\phi)$ for bcc is $Z = 10.5$, as shown in Fig. 3.25. In some sense this is not physical, as Z should refer to the integer number of contacts per bubble. However we know that the $\varepsilon(\phi)$ curves vary smoothly with Z, so we can read this as representing some sort of an 'average' coordination number over the full range of liquid fractions.

3.6.3 Extension of Z-cone model

Fig. 3.26 shows a sketch of how we adapt the Z-cone model to account for the behaviour of a bcc bubble. In short, the bubble surface is split into 14 pieces

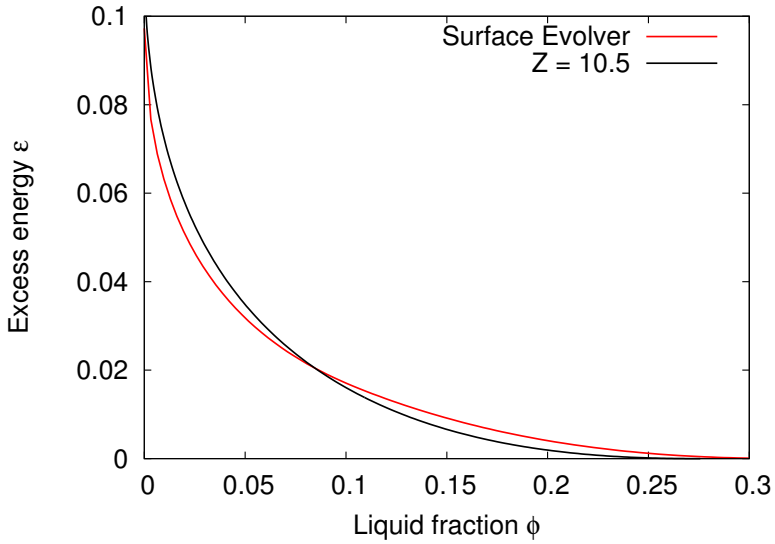

Figure 3.25: The value of Z which results in the best Z-cone approximation to the energy of a bubble in a bcc foam obtainable is $Z = 10.5$.

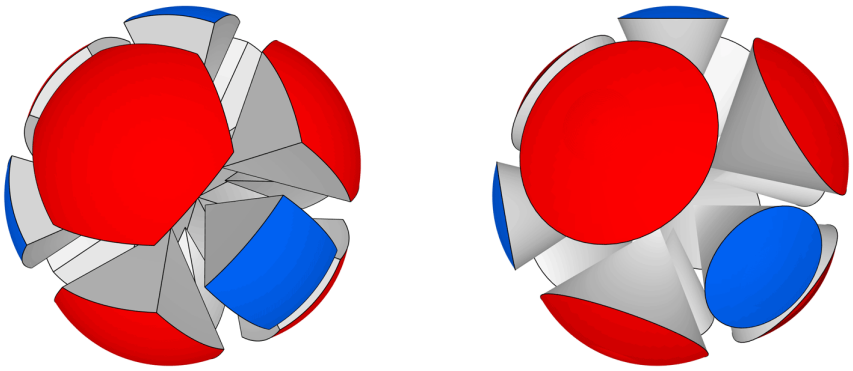

(a) **(b)**

Figure 3.26: (a) The surface of a bubble in a bcc foam, shown here undeformed, *i.e.* at $\phi = \phi_c$, can be split up into 14 parts: 8 bounded by hexagons, corresponding to the $\langle 1\,1\,1 \rangle$ contacts, and 6 bounded by squares, corresponding to the $\langle 1\,0\,0 \rangle$ contacts. (b) In the cone model, each of these is approximated as a *circular* cone. The total surface area of the caps of the cones is minimized subject to appropriate constraints.

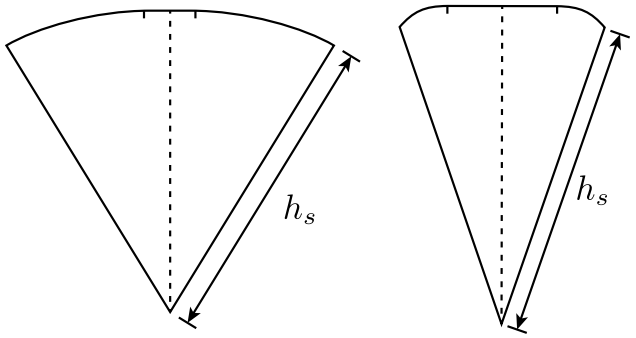

Figure 3.27: To model the Kelvin foam, the cone model requires two different types of cone, as sketched here. The cones have different opening angles and deformations, but have the same slant height h_s.

of two distinct types: 8 for the $\langle 1\,1\,1 \rangle$ contacts, and 6 for the $\langle 1\,0\,0 \rangle$ contacts. As before, we approximate these as being circular cones, but now the cones are *not* all identical.

We obtain a value for the opening angle θ of each type of cone by considering the solid angle subtended by the corresponding face in the dry structure, *i.e.* the truncated octahedron. The deformation ξ now differs for the two sets of cones, and the volume of each type of cone can vary as ϕ is varied, but we require that the slant heights h_s of each type of cone match, in order for the surface to be smooth: this is sketched in Fig. 3.27.

Once again, we leave the full details of the mathematics of the model to Appendix B, but note that as before, the model yields analytic expressions for ε and ϕ.

3.6.4 Results

As before, we test the accuracy of the model by comparing the values of ε predicted by the model with those obtained from Surface Evolver simulations, over the full range of ϕ.

Fig. 3.28 shows the variation of the dimensionless excess energy $\varepsilon(\phi)$ with liquid fraction, obtained from both Surface Evolver and the cone model — a surprisingly good estimation of the excess energy over the entire range of ϕ, with the difference not exceeding one percent of $\varepsilon_0 = \varepsilon(0)$, *i.e.* the value of ε

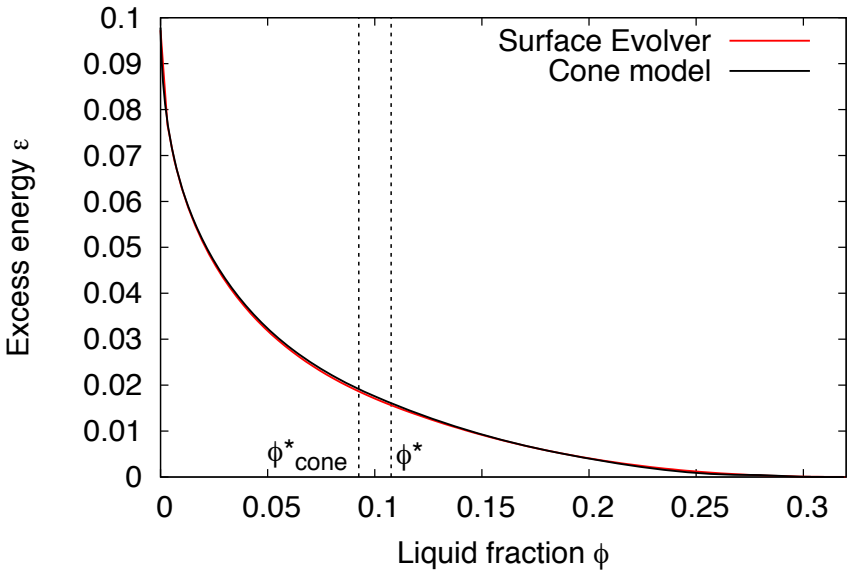

Figure 3.28: Variation of dimensionless excess energy ε with liquid fraction ϕ for a bubble in a bcc foam, obtained from Surface Evolver calculations, and its approximation using the generalized cone model. Increasing ϕ leads to the loss of the six square faces. This takes place at $\phi^* = 0.108$ for Surface Evolver simulations, and at $\phi^*_{\text{cone}} = 0.092$ in the cone model: see dashed vertical lines on plot.

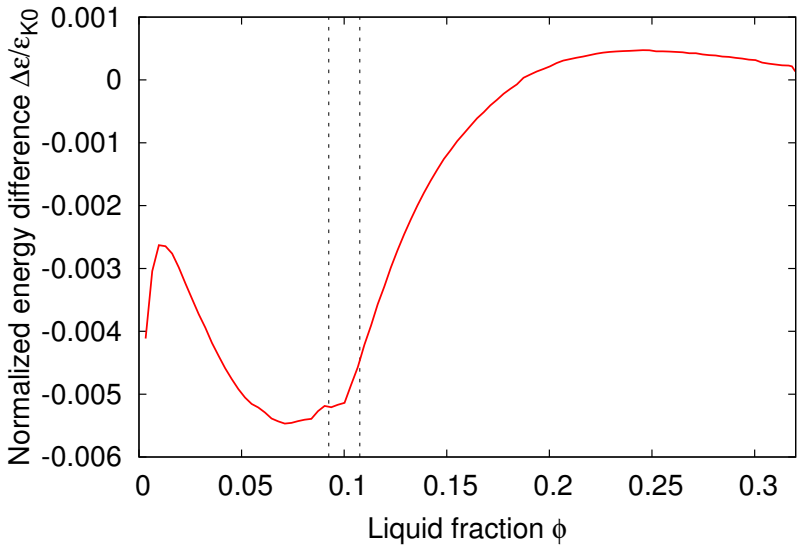

Figure 3.29: The difference in excess energy $\Delta\varepsilon$ between the cone model and Surface Evolver for bcc, normalized by the dry energy ε_0. The maximum difference is less than 1%.

in the dry limit: this normalized difference is plotted in Fig. 3.29. It is worth noting that we see much better agreement here than in the fcc case of the original Z-cone model — this is perhaps surprising in view of the fact that extra approximations need to be made (regarding the heights of the cones, which are no longer all equal, as described in Appendix B) in order to extend the model to deal with bcc. However this improved accuracy is explainable by the fact that we approximate the contacts as circles: the square and hexagonal contacts in a bcc foam are more rotationally symmetric than the rhombi of an fcc foam. Similarly, for $Z = 12$ in the original Z-cone model we saw much better agreement with simulation for a pentagonal dodecahedron than the rhombic dodecahedron.

The excess energy in the dry limit $\varepsilon(\phi = 0)$ is found to be $\varepsilon_0 = 0.0970$ from the Surface Evolver and $\varepsilon_{0,C} = 0.0980$ from the cone model. The value of ϕ at which the square $\langle 1\,0\,0\rangle$ contacts vanish is given as $\phi^* = 0.108$ by the Surface Evolver, and $\phi^*_{\text{cone}} = 0.092$ by the cone model.

It is worth noting that Weaire *et al.* [47] arrived at a remarkably accurate

early estimate of ϕ^*. They approximate all Plateau borders as having a uniform cross-sectional area, and hence the liquid fraction going as the square of the transverse dimension of the Plateau borders. If one thinks of the loss of faces as arising from the Plateau borders spanning the corresponding face, the ratio of ϕ_c/ϕ^* is then given by the square of the ratio of the widths w of the hexagonal and square faces. $w_{hex}/w_{square} = \sqrt{3}$ by simple geometry, and hence $\phi^* \approx \phi_c/3 = 0.108$.

The critical liquid fraction for the wet limit is $\phi_c = 0.320$ for the Kelvin foam; the cone model arrives at an extremely good approximation $\phi_{c,cone} = 0.319$.

Fig. 3.30 shows the variation in area of both square and hexagonal faces with ϕ: recall that ϕ^* is defined as the point at which the area of the hexagonal faces disappears. When ϕ is very slightly less than ϕ^* we encounter problems in accurately modelling the surface using the Surface Evolver, due to difficulties in allowing the area of facets to go to zero.

3.6.5 Logarithimic terms in energy

We now turn our attention to the variation of energy with liquid fraction at each of the two critical points: ϕ^*_{cone} and $\phi_{c,cone}$. In doing so, results are clearer when viewed in terms of derivatives. We show in Fig. 3.31 the variation of the derivative $d\varepsilon/d\phi$ with liquid fraction as obtained from the cone model. The asymptotic behaviour of $d\varepsilon/d\phi$ near the wet limit $\phi = \phi_{c,cone}$, as obtained from differentiating the expression from the original Z-cone model, and keeping the highest-order term:

$$\frac{d\varepsilon}{d\phi} \sim a\frac{\phi_{c,cone} - \phi}{\log(\phi_{c,cone} - \phi)}, \tag{3.19}$$

where a is a constant. This describes the present data well, as one might expect: see Fig. 3.32(a).

For $d\varepsilon/d\phi$ at $\phi = \phi^*_{cone}$ we did not succeed in finding an analytical expression from our new cone model, on account of the numerical procedures

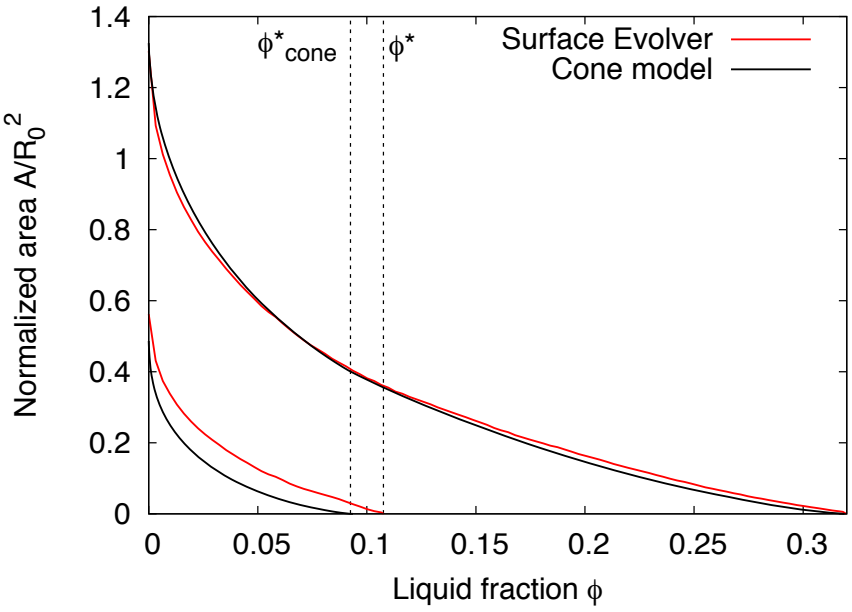

Figure 3.30: Variation of the areas of hexagonal and square contacts with liquid fraction, obtained from the Surface Evolver and the cone model. The areas are normalized by R_0^2, where the bubble volume is $4/3\pi {R_0}^3$. The upper curves refer to the hexagonal faces; the lower curves to the square faces.

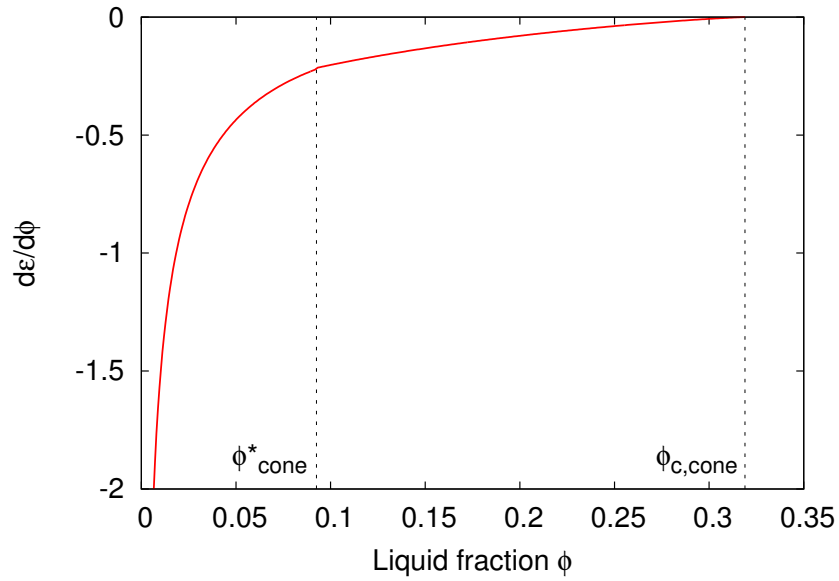

Figure 3.31: The derivative of excess energy with respect to liquid fraction, $d\varepsilon/d\phi$, computed numerically from the cone model. There is a discontinuity of slope at $\phi = \phi^*_{\text{cone}} \approx 0.092$, *i.e.* the point at which the square faces are lost.

involved. By trial and error we arrived at the following empirical expression, which is a good description of the data near $\phi = \phi^*_{\text{cone}}$ (see Fig. 3.32(b)):

$$\frac{d\varepsilon}{d\phi} \sim b_1 + \frac{b_2}{(\log(\phi^*_{\text{cone}} - \phi))^2}, \tag{3.20}$$

with two parameters $b_1 < 0$ and $b_2 > 0$.

There is a discontinuity of the slope of $d\varepsilon/d\phi$ at $\phi = \phi^*_{\text{cone}}$, which is clearly visible in Fig. 3.32(b).

Of note is the presence of logarithmic terms in both expressions, a feature known from various studies of bubble–bubble interactions [53, 50]. The discrepancy between the two forms (3.19) and (3.20) suggests that results which describe simple bubble–bubble interactions are not directly applicable to contact losses at intermediate critical liquid fractions, away from the wet limit.

It has been argued that the limit of mechanical stability of the Kelvin structure is directly attributable to the loss of the square faces [74]. A bcc crystal of interacting points is well known to require second-nearest-neighbour interactions to stabilize it when simple pairwise potentials are applied [75, 76].

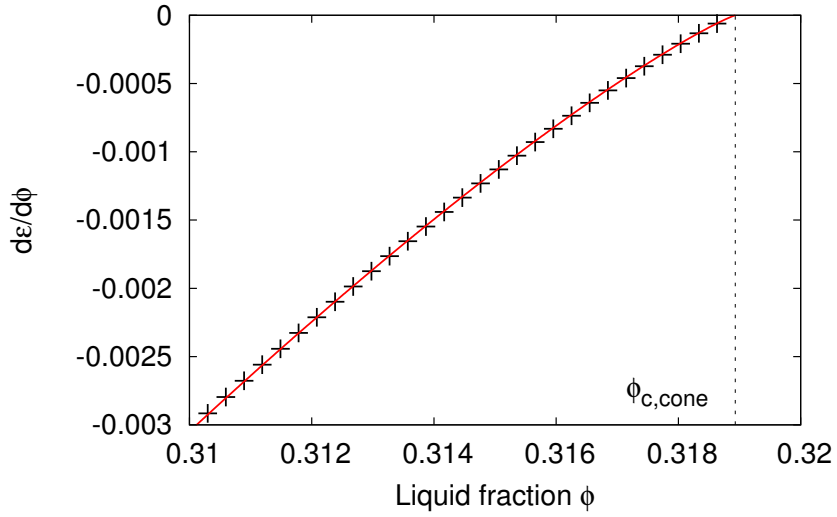

(a)

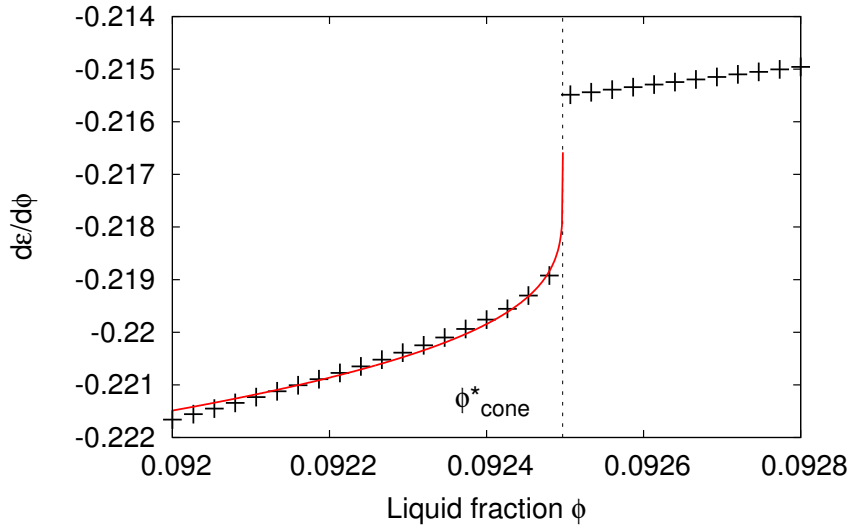

(b)

Figure 3.32: We zoom in on the two regions of interest of Fig. 3.31. (a) Near $\phi = \phi_{c,cone}$ (the wet limit), the variation of the derivative $d\varepsilon/d\phi$ (points) is well approximated by the form in Eq. (3.19) (continuous line), obtained from the simple Z-cone model. (b) Near $\phi = \phi^*_{cone}$ (loss of square faces), the variation is quite different, and is reasonably well approximated by the proposed empirical form of Eq. (3.20) (continuous line).

This appeared supported by Phelan *et al.* [7], who found a negative elastic constant at values of $\phi > 0.11$, *i.e.* very close to the value of $\phi = 0.11 \pm 0.005$ that these authors identified for the face loss.

3.7 Conclusion

We have shown that the Z-cone model achieves its goal of estimating the energy of a bubble whose contacts are all flat and equivalent, over a fairly wide range of liquid fractions. We see the same when we extend it to deal with curved contacts, *i.e.* spherical caps here. We have verified this by comparison with Surface Evolver simulations.

The cone model as extended to deal with the bcc foam is remarkably accurate over the full range of liquid fractions. We see logarithmic terms in the variation of energy near $\phi = \phi_c$. At $\phi = \phi^*$, an empirical expression of a similar form is a good fit to the data. It is possible to obtain these curves at arbitrarily high resolution close to critical points.

To this end, the cone model perhaps presents an *advantage* over simulation methods. Numerical noise hinders the determination of such sensitive derivatives using the Surface Evolver. Understanding the variation of the energy of a bcc foam with liquid fraction at our intermediate critical value of ϕ^* may be applicable to the more general phenomenon of the gain and loss of contacts at intermediate liquid fractions.

Chapter 4

The Interaction of Fibres and Foams

4.1 Introduction

In the paper making industry, *foam forming* is a relatively new manufacturing technique in which the suspension of fibres in water is foamed, before drying it to form paper [77]. The bubbles act to space out the fibres more evenly, improving the homogeneity and the strength of the final product. For this reason, foam forming is of great interest to the papermaking industry. The method is also of interest as a method of producing fibrous material for thermal insulation. The technique has been known since 1974 [77] but is of renewed interest recently.

Adding fibres to a liquid foam alters its physical properties (*e.g.* drainage rate, bubble size distribution, *etc.*) significantly. Recent work [78, 79, 80] concerning the interaction of foams and fibres concerned mainly *wet* foams, *i.e.* those with liquid volume fraction $\phi \gtrsim 0.2$, and focuses on the effect of the presence of fibres on global properties of the foam.

In this chapter we consider the *local* effects of fibres on configurations of soap films. We present two experiments investigating the interaction of a single fibre with the simplest possible model systems for dry foams: one essentially 2d in character, the other fully 3d. We compare data obtained from these experiments to that obtained from simulations using Surface Evolver. We consider the role that this fibre–film interaction may play in the observed

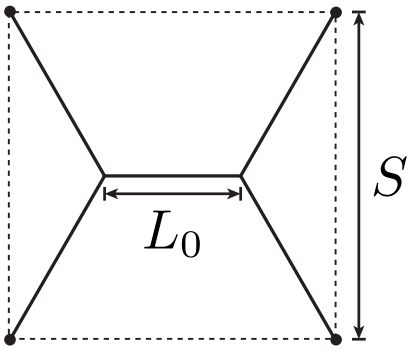

Figure 4.1: The shortest possible network of lines connecting four points arranged in a square has lines meeting in threes at 120°.

reduction of coarsening rates in fibre-laden foams.

We also present some simulations of configurations of cylinders bridged by liquid drops which wet the cylinder surfaces. Such fibre–liquid–fibre bridges are present in the microstructure of paper *after* the drying process, and act to hold the mesh together. We examine the energy as a function of fibre separation and angle, as well as considering evaporation rates via computation of the exposed surface area.

4.2 The interaction of fibres and soap films

4.2.1 The effect of fibres in a quasi-2d arrangement of soap films

Finding the shortest possible length of lines linking a set of 2d points is called the *Steiner problem*. Solutions to the Steiner problem consist of sets of straight lines meeting in threes at angles of 120°, analagously to Plateau's second law [81]. The solution for the vertices of a square is shown in Fig. 4.1. Simple geometry gives the length L_0 of the central line

$$\frac{L_0}{S} = 1 - \frac{1}{\sqrt{3}} \approx 0.42, \tag{4.1}$$

where S is the side length of the square.

Fig. 4.2 shows a simple apparatus consisting of two parallel Perspex plates bridged by four graphite pins arranged in a square. When it is dipped in a

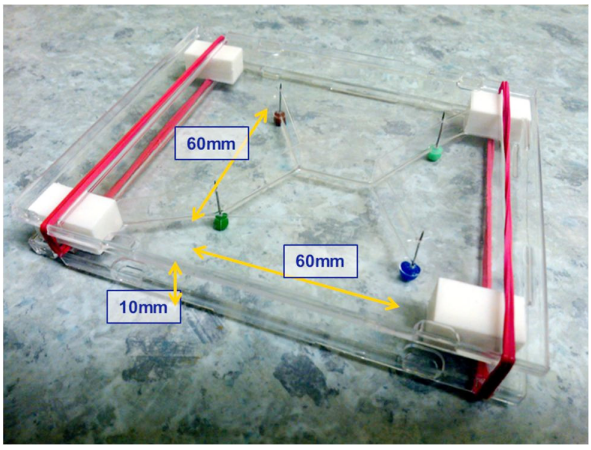

Figure 4.2: An experimental analog to Fig. 4.1: two Perspex plates bridged by graphite pins after being dipped in a soap solution yields the corresponding arrangement of soap films.

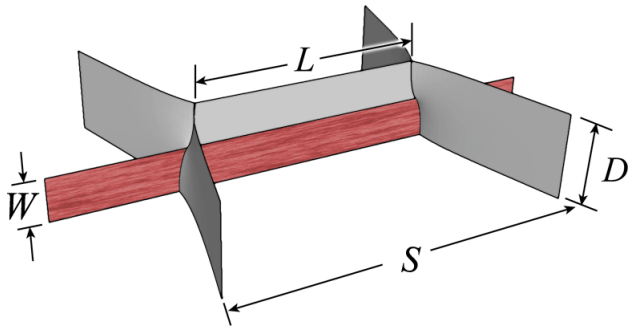

Figure 4.3: A fibre of width W is inserted into the plane of the central film. The length L of the film is measured on the *top* plate, as shown here. The plates are separated by a distance D, with D $<$ S in our quasi-2d setup.

surfactant solution, the network of films formed between the pins is a Steiner tree of the same form as Fig. 4.1. Since the arrangement of films mimics the 2d Steiner tree, we refer to this a *quasi-2d* apparatus. We will use this simple setup as a first step to investigate the interaction of soap films and fibres.

We place a narrow fibre into the plane of the central film, and measure the variation in the length of the central film as we vary the thickness of the fibre. Fig. 4.3 shows the definitions of the relevant lengths (*i.e.* the pin separation S, the plate separation D, the fibre width W and the length of the central film L) for our setup. As a consequence of Plateau's fourth rule, the soap films

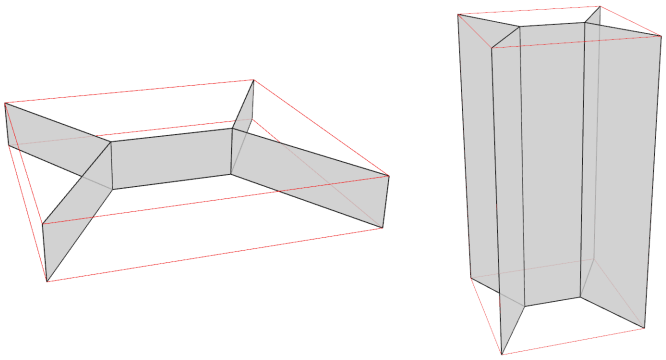

Figure 4.4: We can vary the aspect ratio D/S of the apparatus. We show here two extreme cases without a fibre: $D/S = 0.17$ and $D/S = 2$ respectively.

must meet the fibre at right angles, resulting in an increase in the length L as compared to L_0, its length in the absence of a fibre. It is important to note here that L is measured along the *top* plate here.

Before measuring the variation of film length L with fibre thickness W we note the following facts from the outset. When $W = 0$ there is no fibre, and so we have, from earlier, $L = L_0 \approx 0.42S$. At $W = D$, the fibre spans the space between the two plates. We know in this case that the films will meet the fibre at right angles, giving $L = S$. Hence, regardless of the aspect ratio of the setup, we have the theoretical endpoints: $L/S = 0.42$ at $W/D = 0$, and $L/S = 1$ at $W/D = 1$.

Experimentally, we introduce wood fibres of various widths into the plane of the central film of the apparatus shown in Fig. 4.2 to obtain values for L/S between these endpoints. We vary the plate separation D to obtain data for three different aspect ratios D/S: we illustrate two extreme aspect ratios in Fig. 4.4.

We will compare the experimental data to that obtained from simulation using Surface Evolver. The film length L can be output for a range of fibre widths and aspect ratios. Full details of the simulation are left to Appendix C.

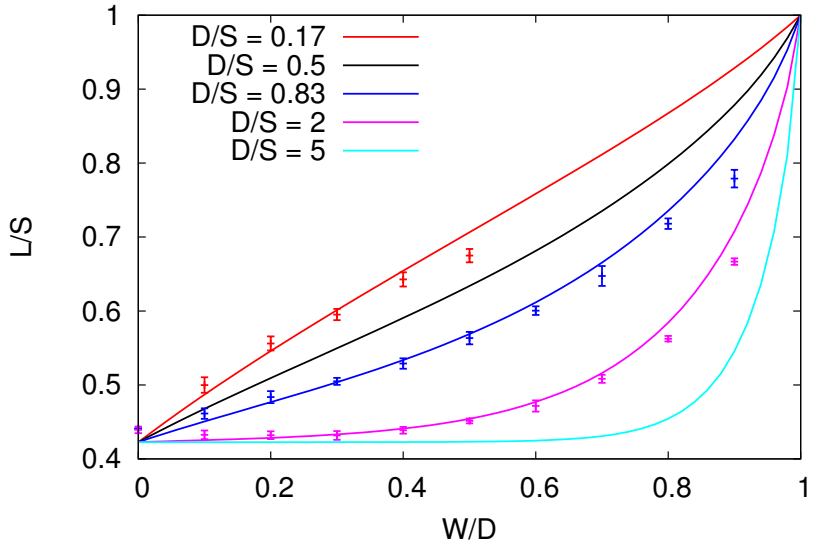

Figure 4.5: As we increase the fibre thickness W from 0 to its maximum value of D, the length L of the film, as measured on the top plate, increases from $\sim 0.42S$ to S. The shape of the curve between these endpoints depends strongly on the aspect ratio D/S.

Fig. 4.5 shows the variation of L/S with W/D, as obtained from our Surface Evolver simulations, along with experimental data from various fibre widths and aspect ratios. Indeed, the theoretical endpoints are recovered. The form of the variation of L/S with W/D between the endpoints depends greatly on the aspect ratio D/S; however, for all three aspect ratios used experimentally, our simulation accurately predicts the variation of film length with fibre thickness.

We look to link the various curves obtained for differing aspect ratios. It turns out that the curves obtained from simulation in Fig. 4.5 can be fairly well approximated by empirical equations of the form

$$\frac{L}{S} = k_1 + k_2 \frac{\exp(\beta W/D) - 1}{\exp(\beta) - 1}, \tag{4.2}$$

with one free parameter β, and constants $k_1 = 1 - 1/\sqrt{3}$ and $k_2 = 1/\sqrt{3}$. The theoretical endpoints are fixed by this form, and in the limit $\beta \to 0$ we recover a straight line between these endpoints. In Fig. 4.6 we show such fits

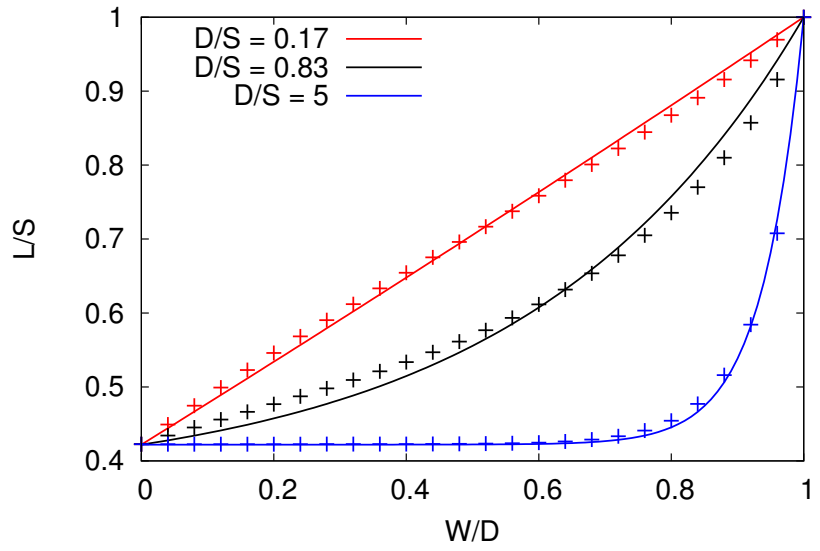

Figure 4.6: For various aspect ratios D/S, the variation of the film length with fibre thickness as obtained from simulations (points) can be approximated by the empirical form (4.2) (solid lines), with only one free parameter β. These curves represent, $\beta \approx 0.3, 2.5, 16$ respectively.

for three values of D/S. These fits are satiscfactory, especially considering we only have one free parameter.

Fig. 4.7 shows the variation of the parameter β with aspect ratio D/S. We see a roughly linear relationship over the range of aspect ratios we consider. We note that the limit $\beta \to 0$ represents a straight line, supporting our observation that smaller aspect ratios produce a more linear form.

Despite the initially large number of variables involved, we see that the response of the central film to the presence of fibres can be fairly well described in terms of a single parameter. However we note that the setup is somewhat artificial, and perhaps not truly representative of the geometry of films in a foam.

4.2.2 The effect of fibres on a single Plateau border

For a more realistic model system for soap films in a dry foam, we must move to three dimensions. Plateau's third rule states that in a dry foam, Plateau

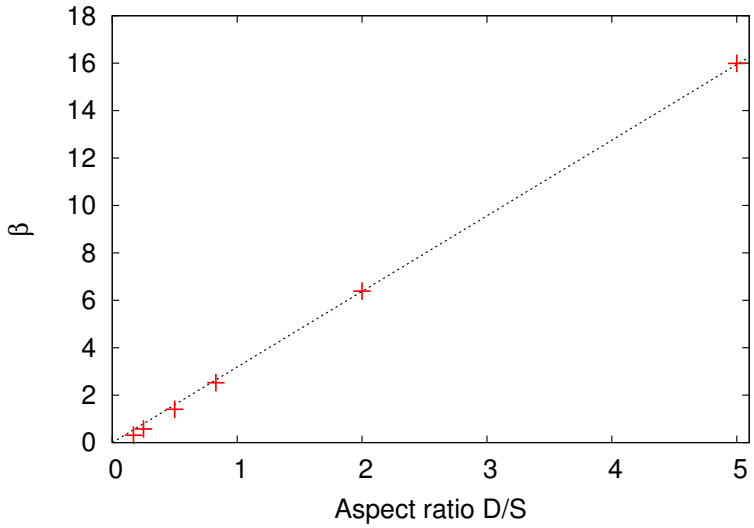

Figure 4.7: The free parameter β which characterizes the variation of film length with fibre thickness depends on the ratio of plate separation to pin separation, D/S. Over the range of aspect ratios we consider here, β increases approximately linearly with D/S.

borders meet in fours at vertices. We want to look at a single Plateau border as it appears in a foam, *i.e.* running from a fourfold node to a fourfold node.

The simplest experimental model system to set up for such a Plateau border is based on a wire frame in the shape of an equilateral triangular prism: a familiar model system [82]. Such wire frames have long been used to study dry foams, dating back to Plateau's work of 1873. More recently, a cubic wire frame was used to investigate the stability of eightfold vertices in microgravity [83].

When the wire frame is dipped into a surfactant solution, a single central Plateau border is formed, connecting two fourfold Plateau border junctions, similarly to Plateau borders in a bulk foam. The wire frame is fed from the top with a burette containing the surfactant solution to prevent drainage of the film and hence its breakup. This arrangement is shown in Fig. 4.8(a), with the relevant lengths (the triangle side length C, the height A and the Plateau border length L_0) labelled.

From here on we will normalize lengths by the side length C of the tri-

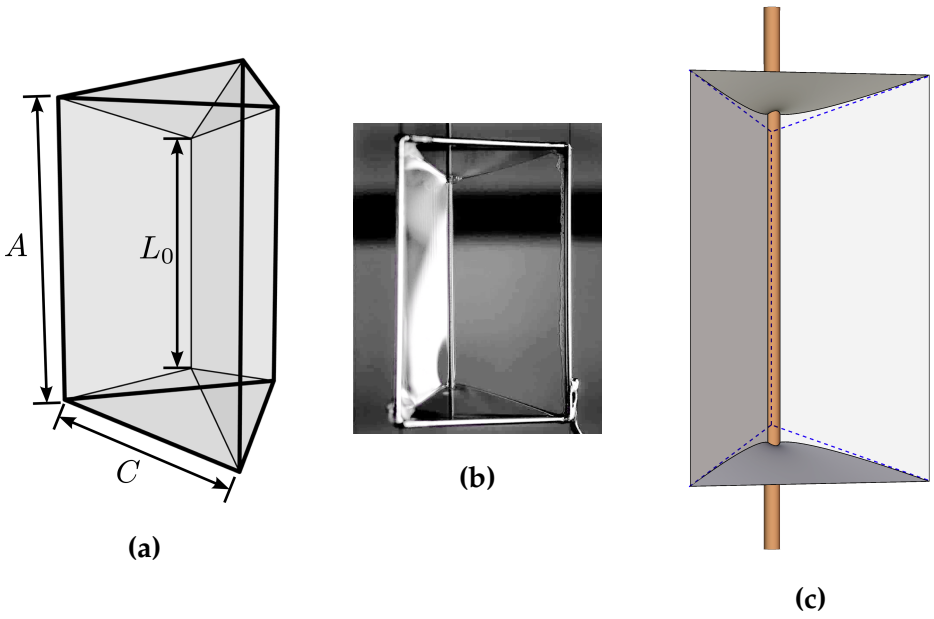

Figure 4.8: (a) Our wire frame apparatus for creating a single isolated Plateau border. The central Plateau border, of length L_0, joins two vertices. (b) Photograph of our experimental setup; a cylindrical fibre has been introduced into the central Plateau border. (c) Simulation of experimental setup. Here one can see that the surface Plateau borders on the fibre are longer than the original Plateau border length (dashed lines).

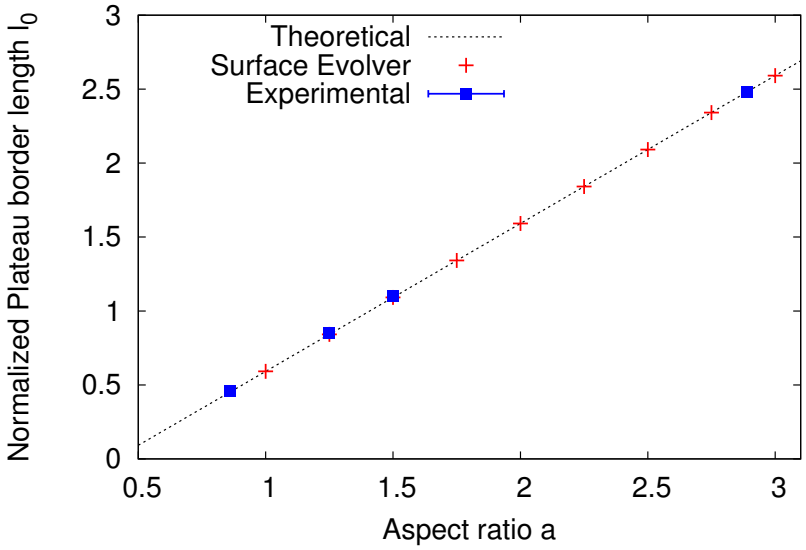

Figure 4.9: The normalized Plateau border length l_0 as a function of the frame aspect ratio a, as obtained experimentally, from simulation, and the theoretical relationship (4.3). The error bars for the experimental data are too small to be visible here.

angle, defining our aspect ratio $a = A/C$ and a normalized Plateau border length $l_0 = L_0/C$. According to Plateau's third rule, the Plateau borders meet at angles of $\arccos(-1/3) \approx 109°$. Hence via simple geometry we can determine the normalized length l_0 of the central Plateau border as a function of the aspect ratio a,

$$l_0 = a - \frac{1}{\sqrt{6}} \approx a - 0.41. \tag{4.3}$$

In Fig. 4.9 we plot l_0 *vs.* a as obtained experimentally for several frames, as obtained from Surface Evolver simulation of the same geometry, along with the linear relationship (4.3). The Surface Evolver values are all accurate to the theoretical values to within 10^{-5}, and the experimental data lie within experimental error of the theoretical values.

4.2.3 Variation of surface Plateau border length with fibre diameter

When a cylindrical fibre is introduced into the central Plateau border, ash shown in Fig. 4.8(c) three *surface Plateau borders* form where the three films meet the fibre. Surface Plateau borders meet Plateau borders at angles of 90°, rather than 109°. Hence the length of these surface Plateau borders (once again, normalized by the side length C of the triangle), which we denote l, is greater than the length l_0 of the original central Plateau border. This can be seen in Fig. 4.8(c).

We find that the surface Plateau border length l depends on the normalized diameter d of the fibre which is introduced into the central Plateau border, as well as on the aspect ratio a of the frame. However, we can eliminate the effects of varying aspect ratios by considering the quantity $\Delta l = l - l_0$, *i.e.* the increase in the length of the Plateau border compared to the length of the Plateau border in the absence of a fibre.

As before, we compare our experimental data with results from the Surface Evolver program. Details of the simulation are given in Appendix C.2.

In Fig. 4.10 we plot Δl as a function of d, as obtained from experiments using a variety of aspect ratios and fibre diameters, along with results from Surface Evolver simulations (these collapse onto a single line in the plot). Within experimental error, our simulations correctly predict the length of the surface Plateau borders.

4.2.4 Variation of Plateau border length with fibre position

We now wish to investigate the behaviour of the surface Plateau borders as the fibre is perturbed from this central position. To avoid having to deal with an overwhelming number of variables, from here on we use a frame with $C = 46\,\text{mm}$ and $A = 69\,\text{mm}$, and use a fibre of diameter 1 mm, giving $a = 1.5$, and $d \approx 0.022$. We move the fibre along the axis of symmetry of

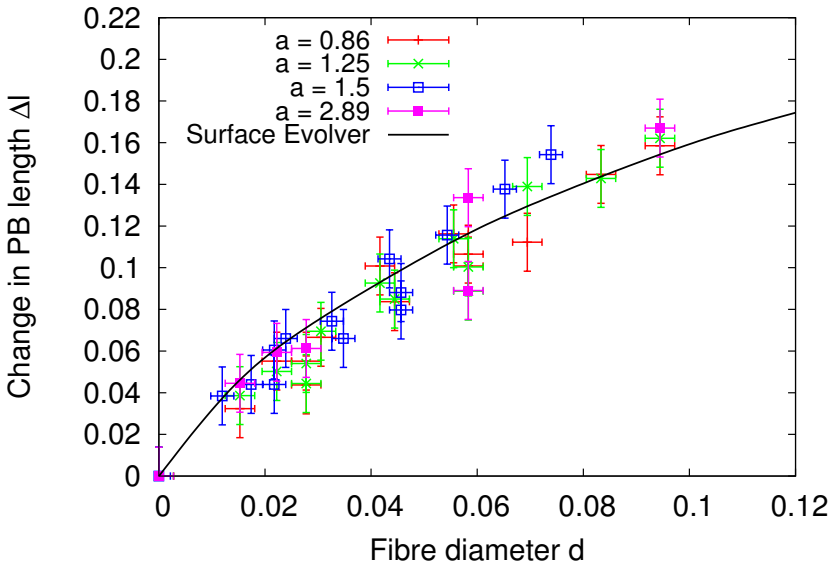

Figure 4.10: As fibres with diameter d are introduced into the central Plateau border, its length increases by an amount Δl. For various aspect ratios, Surface Evolver simulations accurately predict the variation of Δl with d.

the equilateral triangle by a distance x (again normalized by C), as shown in cross section in Fig. 4.11.

Moving the fibre breaks the threefold rotational symmetry, and so the surface Plateau borders are not of equal length. We denote the two different surface Plateau border lengths by l_1 and l_2. Similarly to before, we define $\Delta l_1 = l_1 - l_0$ and $\Delta l_2 = l_2 - l_0$ to denote their respective increases relative to the length of the Plateau border when no fibre is present.

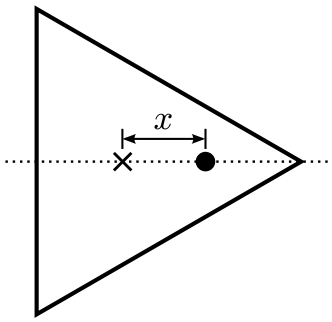

Figure 4.11: We move the fibre (•) a distance x from the centre of the triangle (×), normalized by the triangle side length C.

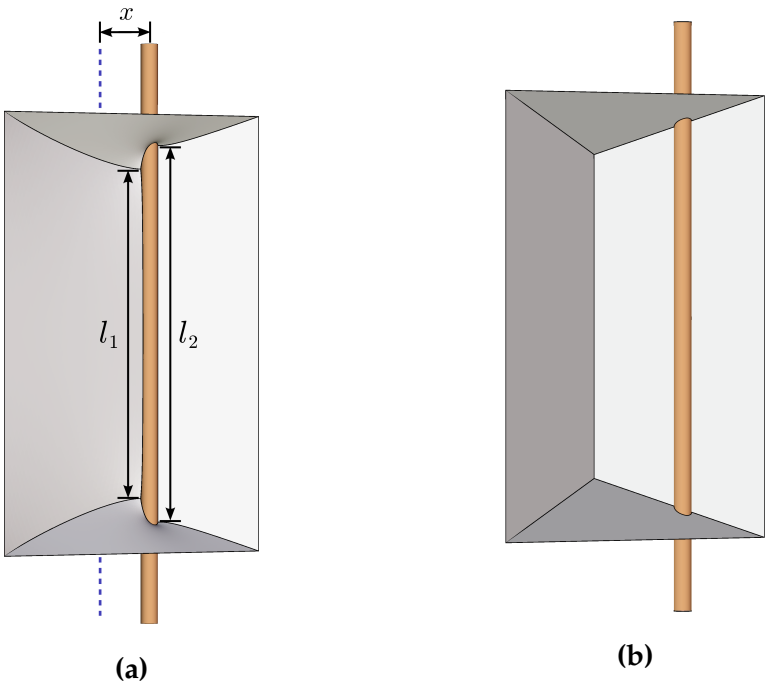

Figure 4.12: (a) Moving the fibre from the centre of the triangle (dashed line), to $x = 0.2$ here, changes the lengths of the surface Plateau borders: l_1 decreases and l_2 increases. (b) When the fibre is moved beyond a point $x > x_b$ the arrangement of films becomes unstable and the fibre no longer lies in the Plateau border.

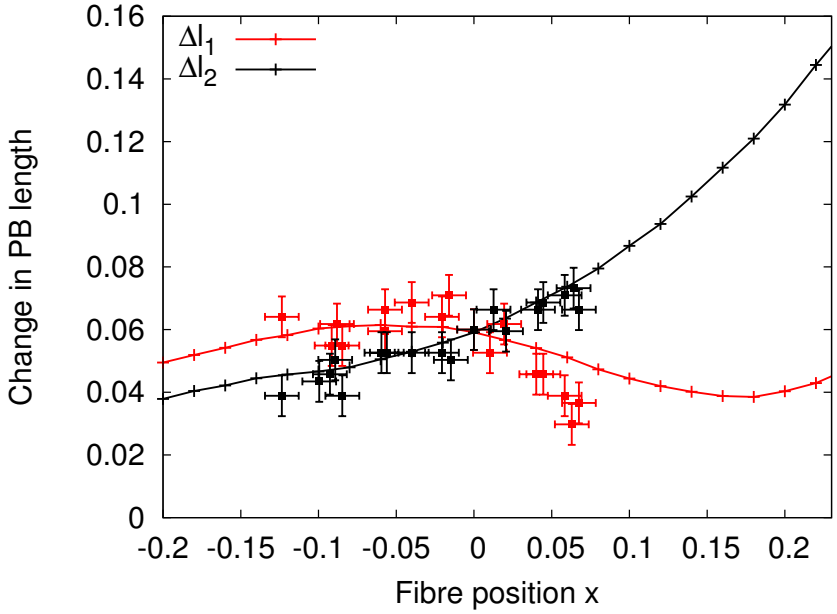

Figure 4.13: Variation of the normalized change in surface Plateau border lengths Δl_1 and Δl_2 as the fibre is moved along the central axis of the triangle. The datapoints represent experimental data; the solid lines are results from Surface Evolver. Note that when $x = 0$, $\Delta l_1 = \Delta l_2$ due to symmetry.

Fig. 4.12(a) shows, for our simulation, the effect of moving the fibre in the positive x direction, away from its initial central position: the Plateau border on the leading edge lengthens while those on the trailing edge shorten. In simulation, when the displacement x exceeds some 'breaking point', $x_b \approx 0.23$, the arrangement of films becomes unstable, prompting a topological change, as shown in Fig. 4.12(b). In our experiments this instability is reached earlier: the surface Plateau borders begin to 'pinch off' preferentially at one end, implying that this preempting of the transition is due to the fact that the fibre is not truly perfectly aligned with the frame. In the negative x direction, the arrangement of films remains stable even as the fibre reaches the edge of the frame.

In Fig. 4.13 we plot the variation of Δl_1 and Δl_2 as the fibre is moved in the x direction, as obtained from Surface Evolver simulations as described above, and from experiments using the setup shown in Fig. 4.8b. The experimental

results are the combination of two runs using the same frame and fibre as before: triangle side $C = 46$ mm, height $A = 69$ mm, fibre diameter 1 mm. The fibre is held in place by a clamp which allows the position of the fibre to be varied from the central position by ± 30 mm.

l_2, the length of the 'leading edge' as in Fig. 4.12(a), increases monotonically with x: the simulation describes this variation to within experimental error over the full range of motion of the fibre. Generally, l_1 tends to decrease as the fibre is moved away from its central position. Over most of the range of x we see a good approximation of l_1; however, it is systematically overestimated for large x. In practical terms, this represents the films peeling off the fibre earlier for the experiment than for the simulation due to asymmetry, as discussed above. The simulation data also shows an 'uptick' for very large x, near x_b. We have not been able to see this effect in experiment, since we have not been able to reach such a large value of x without the films breaking off.

4.2.5 Variation of energy with fibre position

Having seen that our Surface Evolver simulations have been successful in modelling the shape of the films as they interact with the fibre in the cases we examined, we now use the same simulations to compute the area of the soap films in the presence of fibres: the curvature of the films means that this quantity is not directly measurable from experiment. We define a dimensionless energy E as the total surface area of free films, as reported by Surface Evolver for each fibre position, normalized by C^2, where C is the triangle side length, as before.

In Fig. 4.14 we plot the variation of normalized energy E with fibre position x. The minimum value of E is achieved at $x = 0$, meaning that force acts to pull the fibre back towards its central position when it is perturbed in the x direction.

The variation of E with x near $x = 0$ is quadratic: for small x we have $E(x) \approx 3.27 + 1.96x^2$. In other words, for small displacements the action of the films is comparable to a Hookean spring.

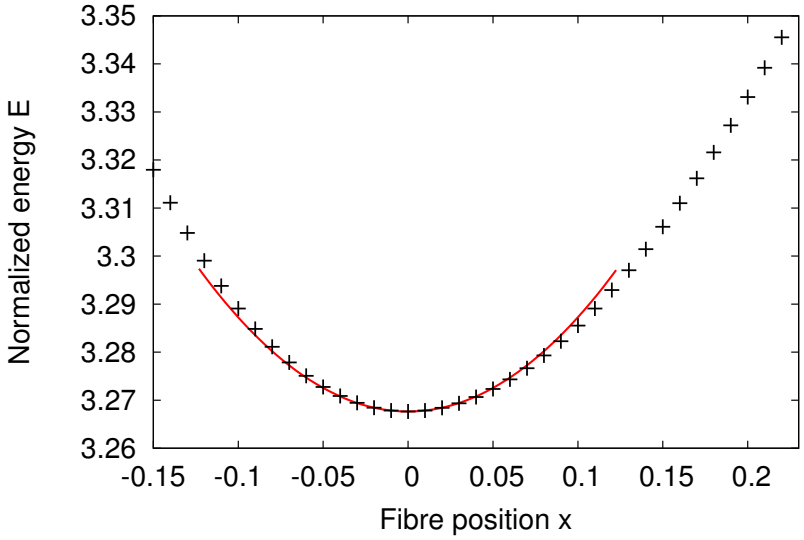

Figure 4.14: Variation of the total normalized surface area E of the system with fibre position. The points are data obtained from Surface Evolver; the line is a quadratic fit for small x. The energy takes a minimum at $x = 0$ meaning that the system is in equilibrium when the fibre is in its central position.

We now turn our attention to $x \to x_b$: the breaking point, at which the films peel off the fibre (as shown in Fig. 4.12(b)). From our simulation, we have $x_b \approx 0.23$. We can take a numerical derivative of our dimensionless $E(x)$ curve to compute the force necessary to induce this transition: we have $E'(x_b) \approx 0.67$, which corresponds to a force of $0.67C\gamma$ ($C = 46$ mm here).

We can also compute the total *energy* necessary to move the fibre from its equilibrium position to x_b. $E(x_b) - E(0) \approx 0.084$, corresponding to an energy requirement of $0.084C^2\gamma$. Note that experimentally the films break off before the theoretical breaking point is reached, and so the force and energy required will be somewhat lower than these values.

We may compare the magnitude of the force necessary to detach the fibre from the films to the forces experienced as a result of coarsening. We know that the Laplace pressure across the surface of a bubble is

$$\Delta p = \frac{4\gamma}{R} \tag{4.4}$$

where R is the bubble radius. If we approximate the area of the contact as 1/13 the surface area of the bubble (using here the often-cited 'ideal' number of contacts $\langle Z \rangle \approx 13.4$, as described in Section 1.1.7), we obtain a corresponding force

$$F = \frac{4\gamma}{R} \cdot \frac{4\pi R^2}{13} \approx 4\gamma R. \tag{4.5}$$

Our simulations gave a value for the force of $0.67C\gamma$ for the force necessary for detachment of the fibre. We may assume that C (the side length of the triangle in the experiment) is of the same order of magnitude as R.

This suggests that the local force due to coarsening is greater than, but of a similar order of magnitude to, the force necessary to detach the fibre from a Plateau border. This would suggest that there is some degree of 'pinning' of Plateau borders to fibres in a fibre-laden foam, to the extent that the rate of coarsening is impeded, but not prevented entirely. This is in line with what has been observed in experiments.

4.3 The interaction of fibres and liquid drops

4.3.1 Motivation

Fig. 4.15 shows the structure of paper on the microscopic level. We see that the paper consists of long narrow criss-crossing cellulose fibres surrounded by a perhaps surprising amount of empty space. In the manufacturing process, these fibres are dispersed in a liquid mixture which is then dried by heating. The resultant network of fibres is stabilized by the small amount of liquid which was not evaporated in this process. This liquid is distributed throughout the paper in the form of small drops situated at the fibre–fibre contacts. These drops 'bridge' the gaps between pairs of fibres, and the effect of surface tension acts to hold the fibres together.

Using Surface Evolver, we can simulate a model system which is an approximation to these fibre–liquid–fibre bridges: namely, a pair of infinitely long cylinders bridged by a liquid drop. We can vary a multitude of parameters in this system and probe the resultant energies, forces, *etc.* involved.

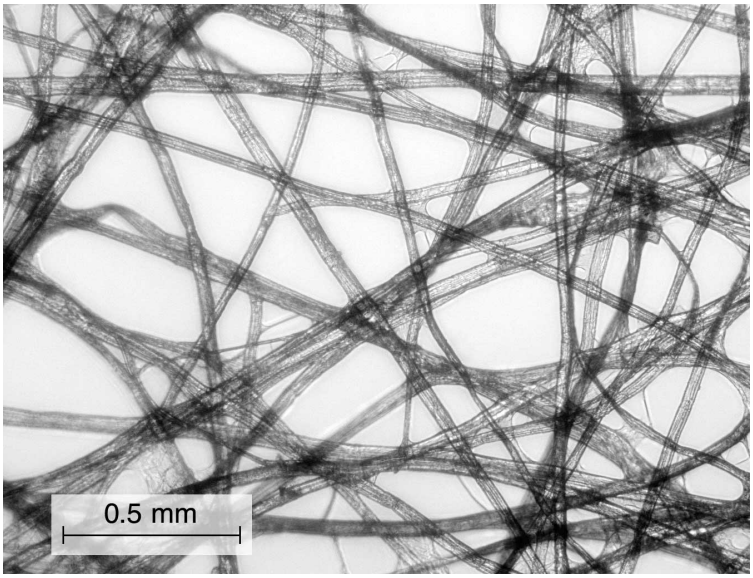

Figure 4.15: Bright-field micrograph of paper. At this scale we see that paper consists of a network of thin fibres surrounded by mostly empty space. *Image: Richard Wheeler*

4.3.2 Contact angle

The shape of the liquid drop will depend strongly on the *contact angle* θ_c between the liquid and the fibre. Different forms of cellulose have different contact angles with water [84]: the kinds used for the manufacture of paper tend to have $15° \lesssim \theta_c \lesssim 30°$ [85]. Here, we will use $\theta_c = 30°$ for the sake of improved stability of simulations.

In our simulations, we impose the contact angle by introducing a solid–

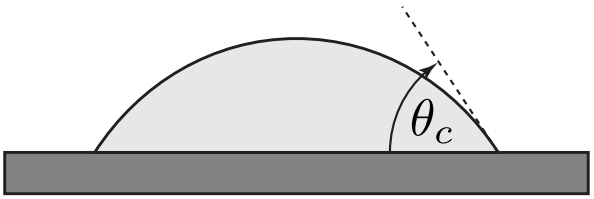

Figure 4.16: The *contact angle* θ_c is measured at the line where a liquid–vapour interface (here, water and air) is in contact with a solid (here, cellulose). Here we show a low contact angle: the liquid is said to *wet* the surface in this case.

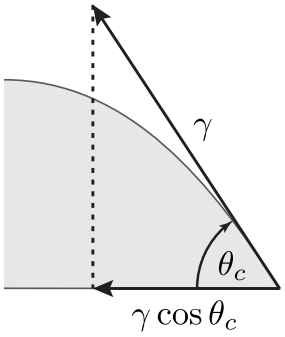

Figure 4.17: We impose the contact angle θ_c in Surface Evolver by introducing a solid–liquid surface tension γ_{sl} to facets which lie on the cylinders. By balancing forces here we see that $\gamma_{sl} = -\gamma \cos \theta_c$. Since this tension is negative, we can picture the lower film as 'pushing' the contact line outwards.

liquid surface tension to those facets which lie *on* the cylinders. This surface tension is given by

$$\gamma_{sl} = -\gamma \cos \theta_c, \tag{4.6}$$

as shown in Fig. 4.17. Since we use a value of 1 for our liquid–air surface tension γ, the tension we assign to these facets is $-\cos(30^\circ) \approx -0.866$. This is negative, which means that the liquid preferentially spreads out on the cylinders.

We define the total energy E as

$$E = \sum_{\text{facets}} \gamma A, \tag{4.7}$$

where A is the area of a facet, and γ its surface tension. Note that by this definition, E can be negative.

4.3.3 Variation of energy with fibre separation

Fig. 4.18 shows a schematic of our simulation setup. We model the fibres as two cylinders of radius r with their central axes separated by a distance $s+2r$, giving a separation s between the cylinder surfaces. From here on, we will fix the cylinder radius $r = 1$ — one can think of all distances here as being normalized by r, energies by r^2 and volumes by r^3.

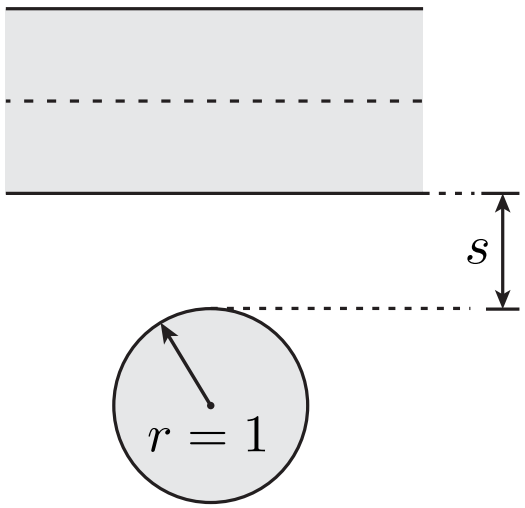

Figure 4.18: We model cellulose fibres as cylinders. Initially we orient a pair of cylinders perpendicularly, and vary their separation s. We set the radius r to 1.

For now, we orient the cylinders at right angles to one another. We add a liquid drop of volume V, initially in the shape of a cube, with opposite faces constrained to lie on each cylinder.

In Fig. 4.19 we plot the variation of total energy E with cylinder separation for a range of drop volumes. We see that the energy increases steadily as we increase the separation between the fibres, gradually 'levelling off' with increasing separation. At low separations, the forces involved (*i.e.* the derivative of the energy with respect to separation) are roughly equal: at $s = 0$, $dE/ds \approx 7$ for all three drop volumes.

4.3.4 Variation of energy with angle

The variation of energy with separation seems to be relatively stable with varying volumes. Hence, for simplicity, from here we will set $V = 4$. We now vary the *angle* φ between the cylinders, as shown in Fig. 4.21, between 0° to 90°, representing the full range of angles by symmetry.

In Fig. 4.22 we plot the variation of energy E with angle φ for a range of separations s. We see that for high separations, the angle has very little effect on the total energy: this is intuitively understood by inspection of *e.g.*

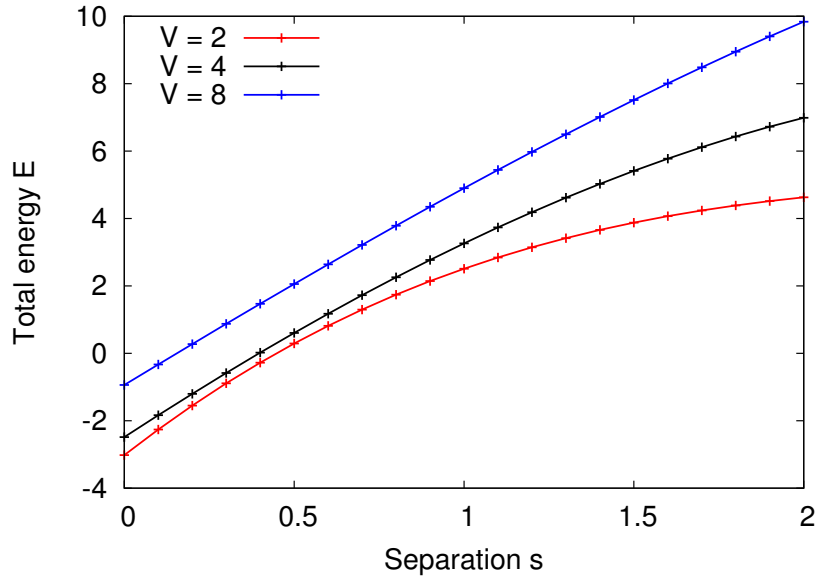

Figure 4.19: Total energy E as a function of separation s for perpendicular fibres, for various drop volumes V. The equilibrium position is $s = 0$, *i.e.* fibres in contact, for all volumes. The forces dE/ds slowly decrease as separation increases.

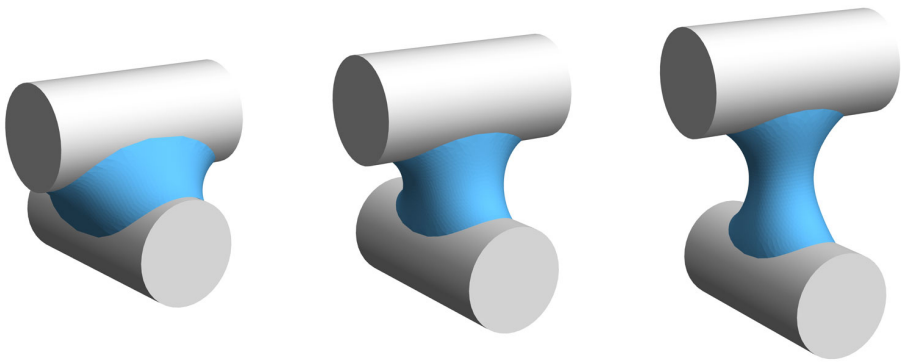

(a) $s = 0$ **(b)** $s = 1$ **(c)** $s = 2$

Figure 4.20: A drop of volume $V = 4$ bridging two fibres oriented in perpendicular directions ($\varphi = 90°$). As we increase the separation s from its minimum value of 0 the energy increases steadily.

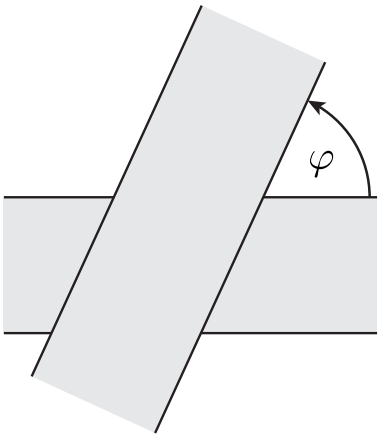

Figure 4.21: We vary the angle φ between the axes of the cylinders, for various separations s.

Fig. 4.20(d). The profile of the drop is such that there is little interaction between the angles of the two cylinders: the 'neck' is close to circular in cross section.

For $s \geqslant 0.25$ the energy increases monotonically with angle, such that 0° is the lowest energy configuration: in other words, torque acts to *align* the fibres. This is not the case for $s = 0$, *i.e.* fibres in contact, in which the equilibrium angle is ~ 24° (or, equivalently, 156°).

The fact that the energy takes a minimum at $\varphi = 0$, *i.e.* parallel fibres, for most separations may be understood intuitively: we have a positive contact angle, so the liquid preferentially *wets* the fibres, as we discussed earlier. When they are oriented in a parallel direction, the liquid can coat fairly long sections of each fibre without requiring a significant increase in the surface area of the drop.

We note that this is not the case for $s = 0$: two fibres in contact. In this case, the energy is minimized at $\varphi \approx 24°$. The reason for this may be best understood by inspection of Fig. 4.23(a). When parallel fibres are in contact, the liquid drop is 'pinched' at its edges: the 30° contact angle means that there is a high degree of curvature near these points, resulting in an increased surface area. As we increase φ, as in Fig. 4.23(b), this deformation of the surface can relax, lowering the energy. When φ exceeds 24°, the energy begins to

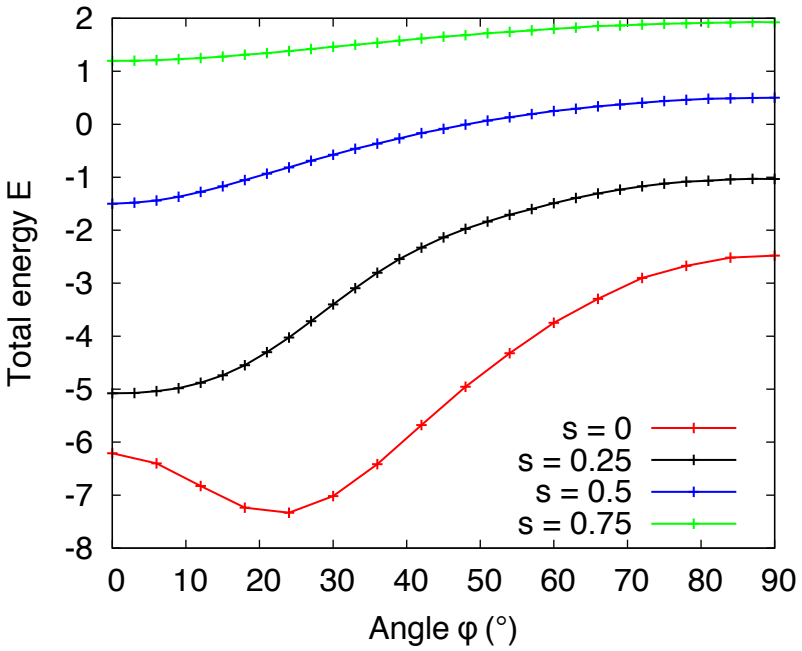

Figure 4.22: The variation of total energy E with angle φ for various separations s. Here we set the drop volume $V = 4$. This range of φ captures the full behaviour: by symmetry α is equivalent to $-\alpha$, and to $180° - \alpha$, for any α.

increase once again.

4.3.5 Parallel and perpendicular fibres

Fig. 4.22 shows that the total energy of the fibre–liquid–fibre system varies quite strongly with angle at low separations. For higher separations we do not see much of a change in energy over the full range of φ, raising the question as to at what separation does the angle between the fibres become a significant factor?

To investigate this, we compare the variation of energy with separation for pairs of fibres which are oriented in parallel and perpendicular directions respectively. We plot this, once again for $V = 4$, in Fig. 4.24.

As we already saw in Fig. 4.19, the curve for $\varphi = 90°$ increases monotonically with s. In contrast, however, when the fibres are parallel, the energy takes a minimum at $s \approx 0.1$. As the separation increases, the difference between the two curves becomes much smaller, as we noted earlier.

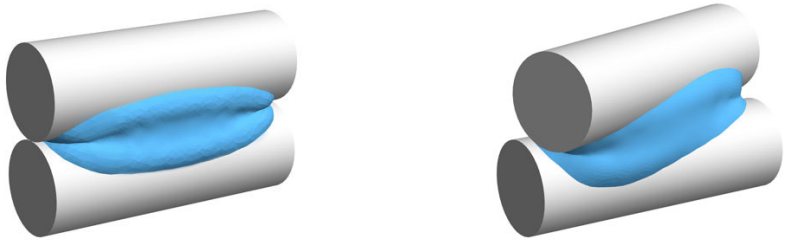

(a) $\varphi = 0°$ **(b)** $\varphi = 24°$

Figure 4.23: For fibres which are parallel in contact, the liquid drop is 'pinched' at that part of its surface near the axes of the cylinders. As φ is increased, the surface can relax, lowering its energy.

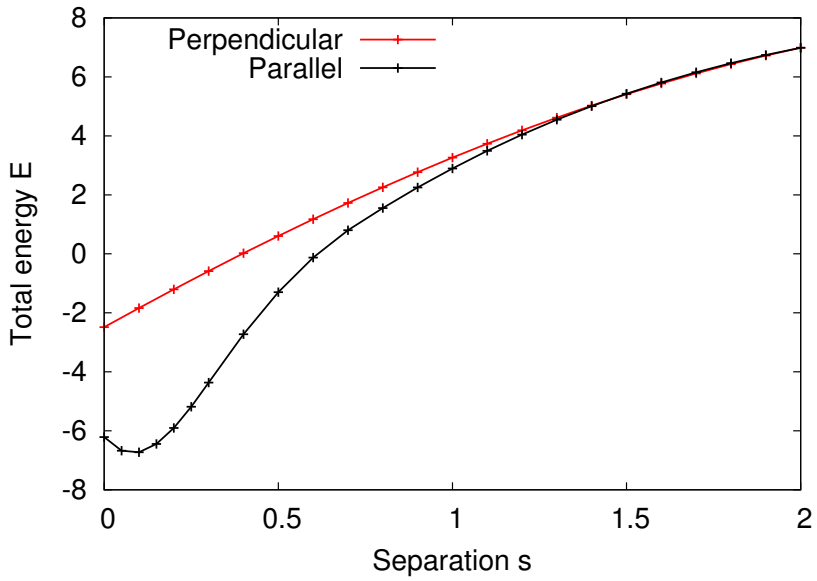

Figure 4.24: The variation of total energy E with separation s takes different forms for parallel ($\varphi = 0°$) and perpendicular ($\varphi = 90°$) fibres, again for $V = 4$.

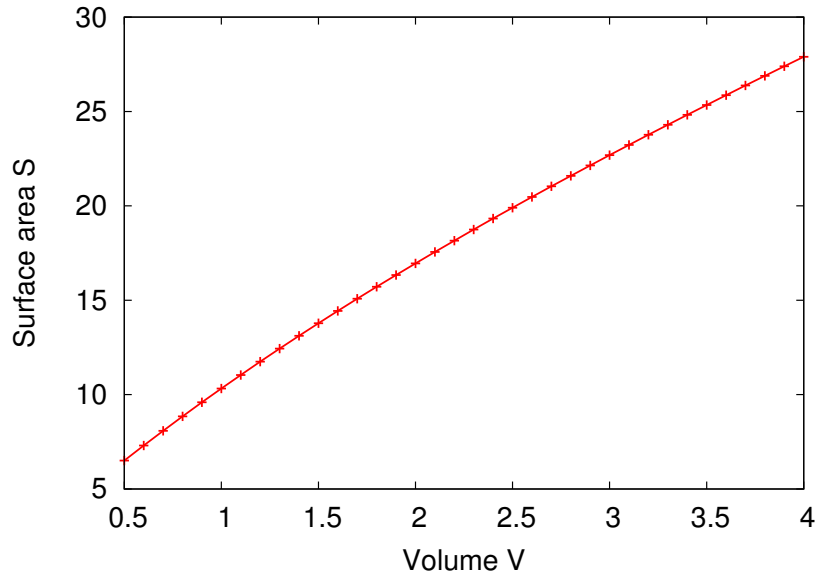

Figure 4.25: For a drop between perpendicular fibres in contact, the surface area S increases steadily with volume V.

This again due to the 'pinching' effect, which comes into play at very low separations. As the distance between the cylinders is slightly increased, this effect is lessened, and so we arrive at a minimum of energy at $s = 0.1$ here.

4.3.6 Surface area and evaporation

We can compute the surface area S of the drop, *i.e.* the area of the facets which are exposed to the air, not counting those on the cylinders, as a function of the drop volume. We plot S as a function of V in Fig. 4.25 for a range of drop volumes V for perpendicular fibres in contact: $\varphi = 90°$; $s = 0$. As we may expect, the surface area increases monotonically with the volume. For $V \lesssim 0.5$ the simulations become unstable as we are dealing with a very narrow sliver of liquid.

As we discussed earlier, liquid evaporates as paper is heated during the paper manufacturing process. As a first approximation, we may take the rate of evaporation as being proportional to the surface area exposed to air, *i.e.*

$$\frac{dV}{dt} \propto -S(t) \tag{4.8}$$

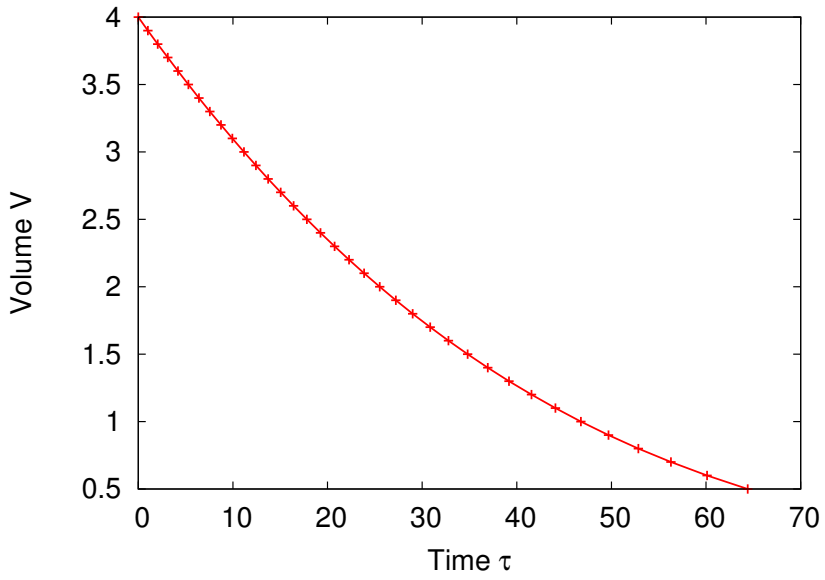

Figure 4.26: Evaporation curve for a drop between two perpendicular fibres in contact. We assume that the rate of change of volume is proportional to the surface area of the drop.

Hence from our $S(V)$ data we can approximate the decrease in volume as a function of time, for some rescaled time τ, where we set $V(\tau = 0) = 4$. The scale of τ will depend on the temperature and humidity of the air.

We plot V as a function of τ in Fig. 4.26. We see that the rate of evaporation slows down gradually over time. This is to be expected, since the surface area S decreases with volume, as above.

4.3.7 Varying contact angle

As we mentioned, the contact angle between water and cellulose can vary depending on the particular makeup of the cellulose used. Bartell and Ray [84] determined the advancing and receding contact angles for various cellulose derivatives and water. Using a formula from Tadmor [86] we can derive the static contact angle θ_c from these. Their measurements yield a value $\theta_c \approx 60°$.

In Fig. 4.27 we plot the total energy as a function of separation using both values of θ_c. For the sake of clarity, we shift the $\theta_c = 60°$ curves vertically so

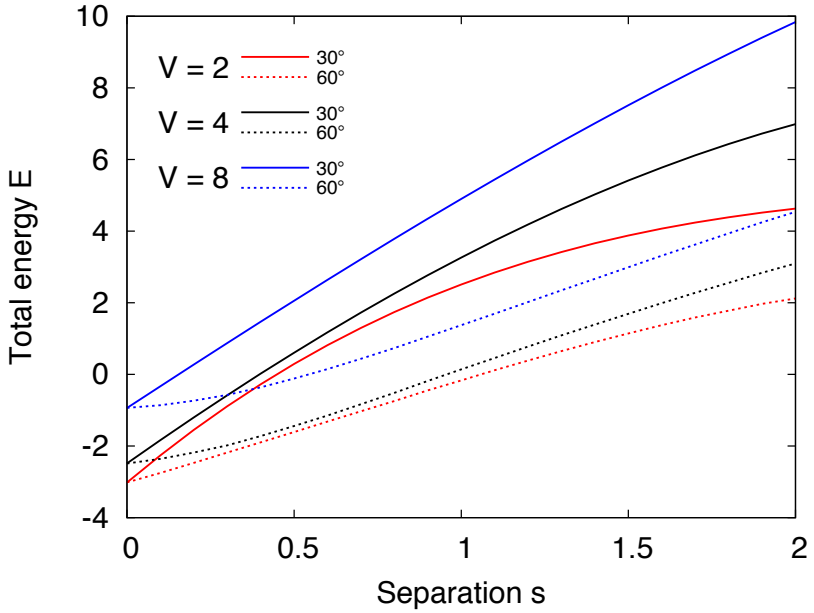

Figure 4.27: Total energy E as a function of separation s for various drop volumes V. Solid lines are for a contact angle $\theta_c = 30°$; dotted lines are for $\theta_c = 60°$.

that they match their $\theta_c = 30°$ counterparts at $s = 0$.

In all cases, we see that the rate of increase of energy E with separation s is significantly lower when a 60° contact angle is used: in other words, the more hydrophilic the fibres are, the greater the forces are which hold them together.

4.4 Conclusion

4.4.1 Fibres and soap films

We see that the introduction of a fibre into an arrangement of soap films alters its geometry from the equilibrium described by Plateau's laws. The response of the soap films to the presence of a fibre in both a film and a Plateau border can be accurately modeled using the Surface Evolver, at least in these simple model systems.

When a fibre introduced into a Plateau border is moved, the effect is that the Plateau border tends to move with it; in other words, the Plateau border is

pinned to the fibre until it pinches off at a large enough displacement. Surface Evolver simulations recreate the response of the surface Plateau borders to the displacement of the fibre, and hence we can use the same simulations to calculate the forces involved.

This effect is relevant to the phenomenon of coarsening in a fibre-laden foam. This pinning of Plateau borders by fibres will act to reduce the rate of coarsening: for bubbles to change in size, motion of films and hence Plateau borders is necessary. This pinning effect has previously been seen to significantly reduce coarsening in experiments on 2d foams [87], in which fixed pins were used. Recent experiments on 3d foams [88] similarly show a reduction in the rate of coarsening in fibre-laden foams. Pinning of fibres to Plateau borders may represent a considerable contribution to this observation.

4.4.2 Fibres and liquid drops

Our simulations, which are a very simpilified representation of fibre–liquid–fibre bridges present in paper, broadly confirm that the drops act to pull the fibres together. This is somewhat unsurprising in view of the fact that the contact angle used, $\theta_c = 30°$, means the fibres are quite hydrophilic. Using a larger contact angle we see the same effect, but weaker.

The variation of energy with angle φ is more subtle: generally once the fibres have been drawn sufficiently close together they preferentially achieve a mutual angle of $\sim 24°$. This fact may help to explain the structure of paper: if the fibres preferentially aligned at $\varphi = 0°$, *i.e.* parallel, perhaps a criss-crossing network as seen in Fig. 4.15 would not be achievable.

4.5 Conclusion

Using foam forming in papermaking has been seen to produce lighter and stronger yields, meaning it is of interest not just commercially but also environmentally. However, the technique has thus far only been explored in

smaller 'pilot' paper plants. In order to fully understand this complex process one must consider the behaviour of foam–fibre mixtures at both the wet and dry limits. The work described in this chapter concerns very idealized model systems, and represents a first step towards understanding this dry limit.

Chapter 5

Conclusion

5.1 Summary

In this thesis we have analyzed various aspects of foam structure. We have interpreted experimental data, considered the shape of bubbles in an ideal theoretical ordered foam, and run countless Surface Evolver simulations, with a view to understanding the shapes and energies of soap films, bubbles, and foams.

In this section I will summarize our most significant findings.

We have successfully employed the method of bond-orientational order parameter analysis to quantify the structure of an experimentally produced foam. We observed a preference for fcc over hcp throughout the foam, in line with previous simulations and optical experiments. Intriguingly, the number of fcc bubbles *and* the number of hcp bubbles increased over the course of the experiment.

Using a simple geometric argument, we have proven that for an ideal foam between the wet and dry limits, the hcp structure has a *lower* energy than the fcc structure. Accordingly, the preference for fcc over hcp must ultimately derive from some other source than energetic considerations. We have confirmed this energy difference using Surface Evolver simulations.

Through various simulations, we have tested the accuracy of the analytic predictions of the cone model in its various forms. In the initial case of equivalent flat contacts, its extension to curved contacts, and the adapted model which deals with the bcc foam, our simulations confirm that its expressions for energy are accurate over wide ranges of liquid fraction.

Furthermore, we have used the analytic predictions of the cone model to derive expressions for measurable properties of foams, and found remarkably good agreement with experiments — even, in some cases, for disordered or polydisperse foams.

~

We have presented simulations and experiments on simple model systems which represent the interaction of fibres with dry foams. We see that our simulations reliably predict film and Plateau border lengths, and hence we can compute the force necessary to 'unpin' a fibre from a Plateau border. We see that this force is comparable in size to those experienced in coarsening.

Our preliminary simulations of cellulose fibres bridged by liquid drops suggest that the drops act to hold the fibres together and, interestingly, ideally hold them at an angle which is neither parallel nor perpendicular.

~

In the following section I will address questions which arose in the course of this work which remain unanswered. I offer my own suggestions as to the direction of possible future work: starting points which I believe represent logical continuations of the research discussed in this thesis.

5.2 Outlook

5.2.1 Prevalence of fcc and hcp

The amount of ordering in the lifetime experiment was seen to increase steadily over the seven days. The natural question is for how long and how far this trend would continue. A similar experiment simply allowed to run for a

longer period would surely shed some light on this question. While the perfluorohexane vapour used in the lifetime experiment reduced the coarsening rate significantly, there was some small change in bubble volumes over the length of the experiment. The gas phase would need to be chosen carefully in a longer experiment to ensure that the bubble sizes remained close to constant.

Another question which may be addressed by a longer experiment is whether the ratio N_{fcc}/N_{hcp} evolves significantly, in analogy to the molecular dynamics simulations of Luchnikov *et al.* [40]. Having more than 7 data points (of which 6 or, perhaps, 5 are reliable) would allow for a more extensive treatment of the evolution of this ratio.

5.2.2 Relative energy of fcc and hcp

The fact that fcc and hcp have different energies is an important consideration if one wants to build an expansion of the energy of a foam of the form

$$E = E_{dry} + E_{PBs} + E_{nodes} + \dots, \tag{5.1}$$

as described in *e.g.* [8], where E_{dry} is the energy of the foam in the dry limit, E_{PBs} represents the contribution of Plateau borders and goes as $\sqrt{\phi}$, and E_{nodes} accounts for Plateau border junctions and goes as ϕ.

The fcc and hcp structures have the same energy in the dry limit, the same total line length per dry bubble, and the same number of nodes. Furthermore they have the same numbers of *types* of nodes (fourfold and eightfold) in the dry limit.

Yet, they have different energies for $\phi > 0$. Hence, if such an expansion were to fully describe the variation of the energy of a foam with liquid fraction, a further term is necessary which reflects the *interaction* of the nodes. Fig. 5.1 shows the shape of fcc and hcp cells in the dry limit: we note that while the total length of edges (Plateau borders) is the same for the two polyhedra, the individual lengths differ. For $\phi > 0$ the shorter Plateau borders

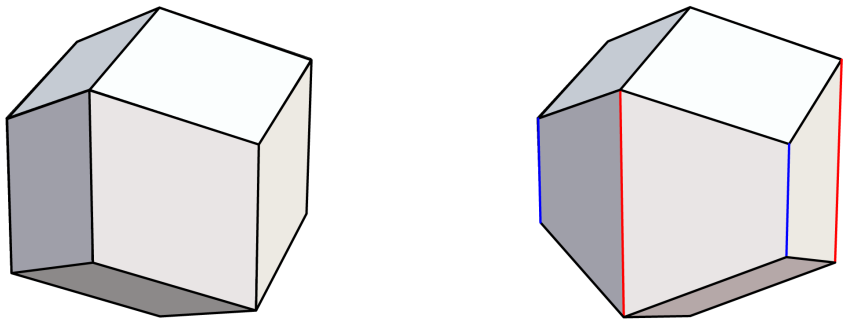

(a) (b)

Figure 5.1: Plateau border lengths for fcc and hcp in the dry limit. (a) the rhombic dodecahedron has 24 edges, all of equal length. (b) the equivalent trapezo-rhombic dodecahedron retains the same length for 16 of the edges. 3 (red) are lengthened; 3 (blue) are shortened.

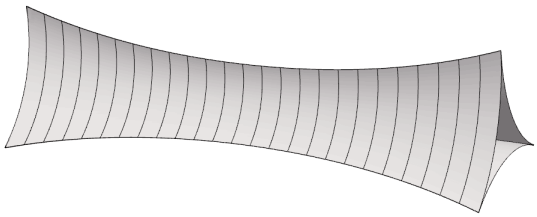

Figure 5.2: The cross section of a Plateau border varies along its length. Its radius of curvature r_c goes approximately as a hyperbolic cosine..

are thicker, *i.e.* have a greater cross-sectional area, and a greater radius of curvature, than the longer ones.

In order to derive a term which accounts for the interaction between nodes, one must consider the *profile* of a Plateau border: *i.e.* the variation of its radius of curvature r_c along its length: see Fig. 5.2. For a Plateau border which possesses reflectional symmetry, as is the case for the 'long' and 'short' Plateau borders in the hcp structure, the property of constant mean curvature (Plateau's first law) implies that the radius of curvature varies approximately as $r_c \propto \cosh(kx)$, for some constant k. These profiles can be obtained with high accuracy from Surface Evolver simulations for each of the three types of Plateau border.

We note that the practical usefulness of such a term as a contribution to

the total energy is limited: the difference between the energies of the fcc and hcp structures is very small, as we have seen, and when comparing structures which *do* differ in edge lengths, node types, *etc.*, the 'interaction term' will be dwarfed by the contributions of the leading terms.

5.2.3 The cone model

The cone model as it stands, strictly speaking, applies to ordered foams. Initially we dealt with the case where all contacts were equivalent (such as an fcc foam).

The model was quite successfully extended to capture the more complex properties of the Kelvin foam, in which contacts differ in size and shape, and are gained and lost at $\phi = \phi^*$. Such an extension required careful consideration of the shape of the dry cell and is not directly applicable to the more general case of a bubble with contacts of differing size. However the implications of its predictions of the behaviour of $d\varepsilon/d\phi$, the derivative of excess energy with liquid fraction, at $\phi = \phi^*$ may be worth pursuing. We see logarithmic behaviour, but not of the same form seen at the wet limit $\phi = \phi_c$. The precise nature of the link between the loss of square contacts and the instability of the Kelvin foam is not fully understood: the Surface Evolver cannot provide curves of high enough resolution for such derivatives to be closely examined, and so an analytic treatment such as the cone model may be of use.

The eventual goal of the cone model was for it to deal with bubbles with arbitrary contacts, to model *e.g.* a disordered foam. Whether this is achievable is uncertain. However, as saw in Section 3.5.5, using the Z-cone model for some carefully chosen value of Z can produce results which are remarkably close to experimental data for disordered foams, despite the derivation of the model hinging on all contacts being equivalent. In view of this, there may be many other physical properties of foams which can be computed analytically using the Z-cone model, even in the case of disordered foams.

5.2.4 Fibres and soap films

From our wire frame experiment, we see that Surface Evolver simulations predict the length of the surface Plateau borders on a fibre fairly well as it is moved from its equilibrium position. For that reason we are confident that it faithfully represents the geometry of the surfaces involved, and hence the total energy of the configuration.

However it must be noted that for simplicity we restricted our analysis to motion in a single direction, in both experiments and simulations. Perhaps more notably, we held the fibre parallel to the axis of the frame and to the Plateau border in all cases. As we noted, misalignment of the fibre and the Plateau border leads to an earlier breaking off of the films. A more advanced set of simulations encompassing various directions of motion and fibre orientations would build a fuller picture of the nature of this topological transition, and the energies involved.

Our comparison of the forces required for unpinning of the Plateau border from the fibres to those involved in coarsening is of course very approximate. A more quantitative understanding could be obtained from simulations which, in place of bare soap films on a wire frame, implemented complete bubbles. The bubble volumes could be varied as well as the fibre position, and minimum-energy configurations found. This would still of course be a highly idealized representation of a real foam.

5.2.5 Fibres and liquid drops

These simulations of drops on fibres are very much a first step towards a fuller understanding of the dynamics of fibres held together by liquid drops.

Modelling fibres as infinitely long rigid cylinders is a natural starting point for any treatment of such systems, and it is convenient computationally. Future work must consider the fact that such fibres can bend, and the energies associated with bending. Duprat and Protiere [89] present a simple experiment in which two parallel silicone fibres are clamped at either end

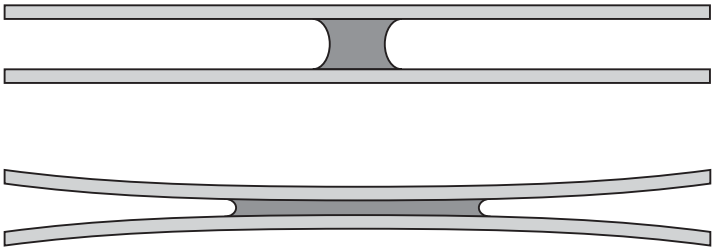

Figure 5.3: Depending on the bending energies involved, the presence of a liquid drop on two parallel fibres may cause the fibres to buckle inwards, as illustrated here.

and bridged by a drop of mineral oil, causing the fibres to buckle inwards to facilitate wetting, as sketched in Fig. 5.3.

Such a setup would be a very sensible first step to understanding the interaction of bendable fibres with liquid drops; however, implementing such a setup in Surface Evolver would require a much more sophisticated definition of the fibres than the simple cylindrical constraints used in Section 4.3.

The fibres could themselves be implemented as full free surfaces, with custom energy integrals defined over their surfaces to account for bending energies.

Alternatively one could attempt a 'trial-and-error' type method, in which the cylinder axis is specified in terms of some free parameters. Running such a simulation for various combinations of parameters would, somewhat tediously, eventually yield ones which minimized the total energy of the setup. However this approach would fail to replicate the strongly hysteretic curves in [89].

Our modelling of the drop shrinking upon evaporation, with the contact line moving freely, is a very idealized approximation. Experimental results from Birdi *et al.* [90] concerning the evaporation of water drops on glass suggest that pinning of the contact line plays an important role, and that the evaporation rate is hence not exactly proportional to surface area.

Appendix A

Bond-orientational order parameters

A.1 Definition

For a bubble with Z nearest neighbours, we first compute the mean value of the spherical harmonic $Y_{\ell m}$ as evaluated in the direction of each of the bubble's nearest neighbour bonds.

$$\overline{Y_{\ell m}} = \frac{1}{Z} \sum_{i \in \text{nns}} Y_{\ell m}(\theta_i, \varphi_i). \tag{A.1}$$

The BOOP Q_ℓ for any ℓ is defined as

$$Q_\ell = \sqrt{\frac{4\pi}{2\ell+1} \sum_{m=-\ell}^{\ell} \left| \overline{Y_{\ell m}} \right|^2}. \tag{A.2}$$

This combination over m gives us rotational invariance. This is extremely important — we want, for example, an fcc lattice to maintain its BOOP signature if it is rotated! The factor $4\pi/(2\ell+1)$ arises from the spherical harmonic addition identity (sometimes called *Unsöld's theorem*)

$$\sum_{m=-\ell}^{\ell} \left| Y_{\ell m}(\theta, \varphi) \right|^2 = \frac{2\ell+1}{4\pi}, \tag{A.3}$$

and ensures that Q_ℓ is normalized such that $0 \leqslant Q_\ell \leqslant 1$ for any set of neighbours.

A.2 Choice of ℓ

This definition of Q_ℓ works for any integer $\ell \geqslant 0$. Not all of these are useful quantities however. Under the central symmetry transformation $\mathbf{r} \to -\mathbf{r}$ the spherical harmonics transform as $Y_{\ell m} \to (-1)^\ell Y_{\ell m}$. This means that for structures which possess central symmetry, $Q_\ell = 0$ for any odd ℓ. For this reason, only even ℓ tend to be considered.

Furthermore, we find that $Q_0 = 0$ for all sets of bonds, and that $Q_2 = 0$ for any set of bonds that possesses a degree of reflectional symmetry. This leaves us with $Q_4, Q_6, Q_8, \ldots$ as the only 'useful' BOOPS.

Of these, one only tends to see Q_4 and Q_6. One reason for this may be to do with the oscillatory nature of the spherical harmonics. Mapping the angles where $\mathrm{Re}(Y_{\ell m}) = 0$ yields *nodal lines*, as shown in Fig. A.1 for Y_{63}. In general these take the form of $2|m|$ lines of longitude and $\ell - |m|$ lines of latitude. This means that for large ℓ, $Y_{\ell m}$ varies greatly for small changes in angle — this effect is not desirable when dealing with experimental data, in which bond angles will naturally not be perfect.

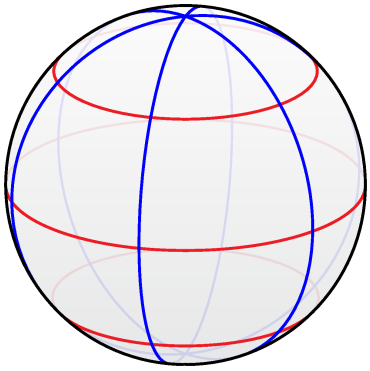

Figure A.1: Nodal lines for Y_{63}, *i.e.* the angles for which $\mathrm{Re}(Y_{63}) = 0$. These take the form of 6 lines of longitude (blue) and 3 lines of latitude (red).

It is worth noting that Steinhardt also describes an *averaged* BOOP, $\overline{Q_\ell}$, which is similar to Q_ℓ except that it takes into account *all* bonds in the system, rather than computing on a bubble-by-bubble basis. This is not useful for our case, since the sample contains a mixture of different structures.

A.3 Selection of nearest neighbours

A bubble's BOOPS depend only on the positions of its nearest neighbours. Unfortunately, there is no one universally accepted defition of "nearest neighbour"; however there are several which are commonly used.

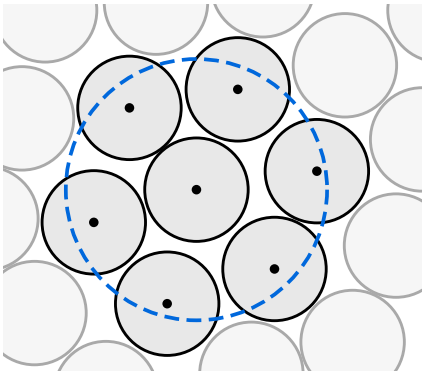

Figure A.2: 2D illustration of the cutoff distance method. Any two bubbles whose centres are separated by less than the cutoff distance are considered to be nearest neighbours.

The simplest way is to define a cutoff distance. We consider two particles as nearest neighbours if their centres are separated by a distance less than this cutoff distance, as illustrated in Fig. A.2. Steinhardt *et al.* use a cutoff distance of $1.2r_0$, where r_0 was the minimum of the Lennard-Jones potential. In the case of experimental data, setting r_0 as the position of the first minimum of the radial distribution function ensures that the first shell is captured as much as possible, while minimizing false positives in the form of next-to-nearest neighbours.

The other commonly used method is to construct the Voronoi cell around each particle and consider particles as nearest neighbours if their Voronoi cells are in contact. This method is slower, and in the case of relatively ordered structures (as we are dealing with) will give very similar results.

More recently, there have been several attempts at constructing more sophisticated and efficient schemes [91, 92]. However these are too computa-

tionally intensive to consider here, especially in view of how cleanly separated the first two peaks in $g(r)$ are, as described in Section 2.3.2.

Appendix B

Mathematics of the cone model

B.1 A note

These expressions were derived by Robert Murtagh. I include them here for the sake of completeness.

B.2 *Z*-cone model

The total surface area per contact, A, can be written as

$$A = A_f + 2\pi \int_0^h r(z) \sqrt{1 + \left(\frac{dr(z)}{dz}\right)^2}\, dz, \tag{B.1}$$

where A_f is the area of the bubble-bubble contact; $A_f = \pi\delta^2$ here.The second term in this equation is the general expression for the area of the surface of revolution of $r(z)$. The volume under this curve is given by

$$V = \pi \int_0^h r(z)^2 dz + \frac{\pi r(0)^3 \cot\theta}{3}. \tag{B.2}$$

To minimize A under the constraint of constant V requires the integrated form of the Euler–Lagrange equation:

$$\frac{dr(z)}{dz} \frac{dL}{d\left(\frac{dr(z)}{dz}\right)} - L = C, \tag{B.3}$$

with the Lagrangian function

$$L\left(r(z), \frac{dr(z)}{dz}\right) = 2r(z)\sqrt{1 + \left(\frac{dr(z)}{dz}\right)^2} - \lambda r(z)^2, \tag{B.4}$$

and its derivative to obtain

$$-\frac{2r(z)}{\sqrt{1+\left(\frac{dr(z)}{dz}\right)^2}}+\lambda r(z)^2=C. \tag{B.5}$$

We determine the unknown constants λ and C are from the boundary conditions

$$\left.\frac{dr(z)}{dz}\right|_{z=h} = \infty \tag{B.6}$$

$$\left.\frac{dr(z)}{dz}\right|_{z=0} = \cot\theta, \tag{B.7}$$

The first of these ensures that the bubble surface meets the flat contact smoothly while the second ensures that the bubble surface meets the cone at a right angle: θ is the opening angle of the cone, given by

$$\theta=\arccos\left(1-\frac{2}{Z}\right). \tag{B.8}$$

After imposing these boundary conditions we have

$$\frac{r(z)}{\sqrt{1+\left(\frac{dr(z)}{dz}\right)^2}}=\frac{r(0)\left(r(z)^2-\delta^2\right)}{\left(r(0)^2-\delta^2\right)\sqrt{1+\frac{(Z-2)^2}{4(Z-1)}}}. \tag{B.9}$$

Rescaling this equation in terms of the dimensionless quantities $\rho(z)=r(z)/r(0)$ and $\rho_\delta=\delta/r(0)$ yields

$$\sqrt{1+\left(\frac{dr(z)}{dz}\right)^2}=\frac{\rho(z)\frac{Z}{2\sqrt{Z-1}}(1-\rho_\delta^2)}{\rho(z)^2-\rho_\delta^2}. \tag{B.10}$$

This is a dimensionless first-order differential equation which can be solved by integration between the limits of ρ_δ and $\rho(z)$. Rearranging for dz and noting that $dr(z)=r(0)\,d\rho(z)$, we find that

$$\int_{-h}^{z} dz = z+h = r(0)I(\rho(z),\rho_\delta,Z) \tag{B.11}$$

and so

$$z=-h+r(0)I(\rho(z),\rho_\delta,Z). \tag{B.12}$$

where $I(\rho(z),\rho_\delta,Z)$ is a definite elliptic integral defined below.

By considering $\rho(z = 0) = 1$ in equation B.12, we obtain the important identity

$$r(0) = \frac{h}{I(\rho_\delta, Z)}, \tag{B.13}$$

allowing us to express the bubble profile as

$$z(\rho(z), \rho_\delta, Z) = h \left[\frac{I(\rho(z), \rho_\delta, Z)}{I(\rho_\delta, Z)} - 1 \right]. \tag{B.14}$$

The elliptic integrals $I(\rho(z), \rho_\delta, Z)$ and $I(\rho_\delta, Z)$ are given by

$$I(\rho(z), \rho_\delta, Z) = \int_{\rho_\delta}^{\rho(z)} (x^2 - \rho_\delta^2) f(x, \rho_\delta, Z)\, dx \tag{B.15}$$

and

$$I(\rho_\delta, Z) = \int_{\rho_\delta}^{1} (x^2 - \rho_\delta^2) f(x, \rho_\delta, Z)\, dx, \tag{B.16}$$

with

$$f(x, \rho_\delta, Z) = \left[\frac{Z^2}{4(Z-1)} x^2 (1 - \rho_\delta^2)^2 - \left(x^2 - \rho_\delta^2\right)^2 \right]^{-\frac{1}{2}}. \tag{B.17}$$

The volume V of our single cone is equal to $1/Z$ of the volume of a spherical bubble, giving

$$V = \frac{4\pi R_0^3}{3Z}. \tag{B.18}$$

Inserting this expression into equation B.2 and solving for $r(0)$ yields

$$r(0) = \frac{h}{I(\rho_\delta, Z)} = R_0 \left[\frac{\left(\frac{4}{Z}\right)}{3J(\rho_\delta, Z) + \frac{Z-2}{2\sqrt{Z-1}}} \right]^{\frac{1}{3}} \tag{B.19}$$

where $J(\rho_\delta, Z)$ is another elliptic integral given by

$$J(\rho_\delta, Z) = \int_{\rho_\delta}^{1} x^2 \left(x^2 - \rho_\delta^2\right) f(x, \rho_\delta, Z)\, dx. \tag{B.20}$$

Making use of equations B.10, B.11 and B.19, we can re-express the surface area per contact A as

$$A(\rho_\delta, Z) = \pi R_0^2 \left(\frac{\left(\frac{4}{Z}\right)}{3J(\rho_\delta, Z) + \frac{Z-2}{2\sqrt{Z-1}}} \right)^{\frac{2}{3}} \left[\rho_\delta^2 + \frac{Z}{\sqrt{Z-1}} \left(1 - \rho_\delta^2\right) K(\rho_\delta, Z) \right] \tag{B.21}$$

where $K(\rho_\delta, Z)$ is further elliptic integral given by

$$K(\rho_\delta, Z) = \int_{\rho_\delta}^{1} x^2 f(x, \rho_\delta, Z)\, dx. \tag{B.22}$$

The dimensionless excess surface energy is defined as

$$\varepsilon(\rho_\delta, Z) = \frac{A(\rho_\delta, Z)}{A_0(Z)} - 1 \tag{B.23}$$

where $A_0(Z)$ is the surface area of the top of a spherical sector corresponding to our undeformed cone. From simple geometry, this is

$$A_0(Z) = 2\pi R_0^2(1 - \cos\theta) = \frac{4\pi R_0^2}{Z}. \tag{B.24}$$

Combining Z of these spherical sectors recovers the total surface area of a spherical bubble of $4\pi R_0^2$, as expected.

Therefore, the dimensionless excess energy is

$$\varepsilon(\rho_\delta, Z) = \frac{\rho_\delta^2 + \frac{Z}{\sqrt{Z-1}}\left(1 - \rho_\delta^2\right) K(\rho_\delta, Z)}{Z^{-\frac{1}{3}}\left(6J(\rho_\delta, Z) + \frac{Z-2}{\sqrt{Z-1.}}\right)^{\frac{2}{3}}} - 1 \tag{B.25}$$

The dimensionless deformation is defined, to the middle of the flat contact, as

$$\xi = 1 - \frac{h + h_c}{R_0} \tag{B.26}$$

where the height of a cone h_c is given by

$$h_c = r(0)\frac{Z-2}{2\sqrt{Z-1}}. \tag{B.27}$$

Using equation B.19, this dimensionless deformation is

$$\xi(\rho_\delta, Z) = 1 - \left(\frac{\left(\frac{4}{Z}\right)}{3J(\rho_\delta, Z) + \frac{Z-2}{2\sqrt{Z-1}}}\right)^{\frac{1}{3}}\left[\frac{Z-2}{2\sqrt{Z-1}} + I(\rho_\delta, Z)\right]. \tag{B.28}$$

The plots of dimensionless excess energy as a function of dimensionless deformation, *e.g.* Fig. 3.5, *etc.*, are produced by plotting equations B.25 and B.28 parametrically as a function of ρ_δ.

B.3 Curved interfaces

We can extend this model to deal with curved interfaces, as in the case of bubbles of unequal sizes.

Beginning with the smaller of the contacting bubbles, the radius of our neighbouring bubbles R_n can be written as $R_n = aR$ with $a > 1$. Our cone has a smaller radius than its neighbour, meaning a higher Laplace pressure and, hence, the interfaces between the bubbles are curved "outwards". The radius of curvature R_c of the interfaces between the bubbles is obtained from the difference in Laplace pressures of the undeformed bubbles. The introduction of curved interfaces introduces two new angles θ_{min} and α into our model. The angle θ_{min} corresponds to the angle made between the axis of symmetry and a line from the edge of the curved contact to the apex of the cone. The angle α corresponds to the angle made at the edge of the curved contact between the contact and a plane perpendicular to the axis of symmetry.

For each Z and a, there are unique values of θ_{min} and α obtained by numerically solving

$$\widehat{I}_\delta(\rho_\delta, \theta_{min}, Z) = \frac{\rho_\delta}{\tan\theta_{min}} - \frac{Z-2}{2\sqrt{Z-1}} \tag{B.29}$$

and

$$\sin\alpha = \rho_\delta \left(\frac{a-1}{a}\right) \left[\frac{\left(\frac{4}{Z}\right) - \left(\frac{a}{a-1}\right)^3 \left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\widehat{J}_\delta(\rho_\delta, \theta_{min}, Z) + \frac{(Z-2)}{2\sqrt{Z-1}}}\right]^{\frac{1}{3}}. \tag{B.30}$$

The dimensionless excess energy for the small bubble $\varepsilon(\rho_\delta, \theta_{min}, \alpha, Z, a)$ is written as

$$\varepsilon(\rho_\delta, \theta_{min}, \alpha, Z, a) = \frac{Z}{4}\left(\frac{\left(\frac{4}{Z}\right) - \left(\frac{a}{a-1}\right)^3 \left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\widehat{J}_\delta(\rho_\delta, \theta_{min}, Z) + \frac{(Z-2)}{2\sqrt{Z-1}}}\right)^{\frac{2}{3}} \times$$
$$\left[\rho_\delta^2 + \left(\frac{a}{a-1}\right)^2 (1-\cos\alpha)^2 \left(\frac{\left(\frac{4}{Z}\right) - \left(\frac{a}{a-1}\right)^3 \left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\widehat{J}_\delta(\rho_\delta, \theta_{min}, Z) + \frac{(Z-2)}{2\sqrt{Z-1}}}\right)^{-\frac{2}{3}} + \right.$$
$$\left. \frac{Z}{\sqrt{Z-1}}(1-\rho_\delta^2)\widehat{K}_\delta(\rho_\delta, \theta_{min}, Z)\right] - 1. \tag{B.31}$$

The dimensionless deformation for the small bubble $\xi(\rho_\delta, \theta_{min}, \alpha, Z, a)$, measured to the middle of the curved contact, is expressed as

$$\xi(\rho_\delta, \theta_{min}, \alpha, Z, a) = 1 - \left(\frac{\left(\frac{4}{Z}\right) - \left(\frac{a}{a-1}\right)^3 \left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\widehat{J}_\delta(\rho_\delta, \theta_{min}, Z) + \frac{(Z-2)}{2\sqrt{Z-1}}} \right)^{\frac{1}{3}} \times \left[\frac{Z-2}{2\sqrt{Z-1}} + \widehat{I}_\delta(\rho_\delta, \theta_{min}, Z) \right] - \left(\frac{a}{a-1} \right) (1 - \cos\alpha) . \tag{B.32}$$

The definite elliptic integrals $\widehat{I}_\delta(\rho_\delta, \theta_{min}, Z)$, $\widehat{J}_\delta(\rho_\delta, \theta_{min}, Z)$ and $\widehat{K}_\delta(\rho_\delta, \theta_{min}, Z)$ are given by:

$$\widehat{I}_\delta(\rho_\delta, \theta_{min}, Z) = \int_{\rho_\delta}^{1} \left[\left(x^2 - \rho_\delta^2\right) - \rho_\delta \sin\theta_{min} \frac{Z}{2\sqrt{Z-1}} \left(x^2 - 1\right) \right] \widehat{f}(x, \rho_\delta, \theta_{min}, Z)\, dx, \tag{B.33}$$

$$\widehat{J}_\delta(\rho_\delta, \theta_{min}, Z) = \int_{\rho_\delta}^{1} x^2 \left[\left(x^2 - \rho_\delta^2\right) - \rho_\delta \sin\theta_{min} \frac{Z}{2\sqrt{Z-1}} \left(x^2 - 1\right) \right] \widehat{f}(x, \rho_\delta, \theta_{min}, Z)\, dx, \tag{B.34}$$

$$\widehat{K}_\delta(\rho_\delta, \theta_{min}, Z) = \int_{\rho_\delta}^{1} x^2 \widehat{f}(x, \rho_\delta, \theta_{min}, Z)\, dx, \tag{B.35}$$

with

$$\widehat{f}(x, \rho_\delta, \theta_{min}, Z) = \left[\frac{Z^2}{4(Z-1)} x^2 (1 - \rho_\delta^2)^2 - \left[\left(x^2 - \rho_\delta^2\right) - \rho_\delta \sin\theta_{min} \frac{Z}{2\sqrt{Z-1}} \left(x^2 - 1\right) \right]^2 \right]^{-\frac{1}{2}}. \tag{B.36}$$

For the large bubble we need to determine the dimensionless ratios ρ_{min} and ρ_{max} by numerically solving

$$\rho_{min} = (a-1)^{-1} \sin\alpha \left(\frac{\left(\frac{4}{Z}\right) + (a-1)^{-3} \left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\left[\widetilde{J}_\delta(\rho_{max}, Z) + \breve{J}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z)\right] + \frac{Z-2}{2\sqrt{Z-1}}} \right)^{-\frac{1}{3}} \tag{B.37}$$

for the paired values of ρ_{min} and ρ_{max} which give the smallest excess energy provided $\rho_{min} < \rho_\delta$ and $\rho_{min} < \rho_{max}$.

The dimensionless excess energy for the large bubble $\varepsilon(\rho_{min}, \rho_{max}, \theta_{min}, \alpha, Z, a)$ is written as

$$\varepsilon(\rho_{min}, \rho_{max}, \theta_{min}, \alpha, Z, a) = \frac{Z}{4}\left(\frac{\left(\frac{4}{Z}\right) + (a-1)^{-3}\left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\left[\widetilde{J}_\delta(\rho_{max}, Z) + \breve{J}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z)\right] + \frac{Z-2}{2\sqrt{Z-1}}}\right)^{\frac{2}{3}}\Bigg[\rho_{min}^2 + \\ + (a-1)^{-2}\left(\frac{\left(\frac{4}{Z}\right) + (a-1)^{-3}\left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\left[\widetilde{J}_\delta(\rho_{max}, Z) + \breve{J}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z)\right] + \frac{Z-2}{2\sqrt{Z-1}}}\right)^{-\frac{2}{3}}(1-\cos\alpha)^2 + \\ +2(\rho_{max}^2 - \rho_{min}^2)\breve{K}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z) + \frac{Z}{\sqrt{Z-1}}(1-\rho_{max}^2)\widetilde{K}_\delta(\rho_{max}, Z)\Bigg] - 1. \tag{B.38}$$

The dimensionless deformation for the small bubble $\xi(\rho_{min}, \rho_{max}, \theta_{min}, \alpha, Z, a)$, again measured to the middle of the curved contact, is expressed as

$$\xi(\rho_{min}, \rho_{max}, \theta_{min}, \alpha, Z, a) = 1 + \frac{1-\cos\alpha}{a-1} - \\ - \left(\frac{\left(\frac{4}{Z}\right) + (a-1)^{-3}\left(2 - 3\cos\alpha + \cos^3\alpha\right)}{3\left[\widetilde{J}_\delta(\rho_{max}, Z) + \breve{J}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z)\right] + \frac{Z-2}{2\sqrt{Z-1}}}\right)^{\frac{1}{3}} \times \\ \times \left[\frac{Z-2}{2\sqrt{Z-1}} + \breve{I}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z) + \widetilde{I}_\delta(\rho_{max}, Z)\right]. \tag{B.39}$$

The definite elliptic integrals defined for the large bubble are given by

$$\breve{I}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z) = \int_{\rho_{min}}^{\rho_{max}} \sin\theta_{min}\rho_{min}\left(x^2 - \rho_{max}^2\right)\breve{f}(x, \rho_{min}, \rho_{max}, \theta_{min}, Z)\,dx, \tag{B.40}$$

$$\breve{J}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z) = \int_{\rho_{min}}^{\rho_{max}} \sin\theta_{min}\rho_{min}x^2\left(x^2 - \rho_{max}^2\right)\breve{f}(x, \rho_{min}, \rho_{max}, \theta_{min}, Z)\,dx, \tag{B.41}$$

$$\breve{K}_\delta(\rho_{min}, \rho_{max}, \theta_{min}, Z) = \int_{\rho_{min}}^{\rho_{max}} x^2\breve{f}(x, \rho_{min}, \rho_{max}, \theta_{min}, Z)\,dx, \tag{B.42}$$

$$\widetilde{I}_\delta(\rho_{max}, Z) = \int_{\rho_{max}}^{1} \left(x^2 - \rho_{max}^2\right)\widetilde{f}(x, \rho_{max}, Z)\,dx, \tag{B.43}$$

$$\widetilde{J}_\delta(\rho_{max}, Z) = \int_{\rho_{max}}^{1} x^2(x^2 - \rho_{max}^2)\widetilde{f}(x, \rho_{max}, Z)\, dx, \tag{B.44}$$

$$\widetilde{K}_\delta(\rho_{max}, Z) = \int_{\rho_{max}}^{1} x^2\widetilde{f}(x, \rho_{max}, Z)\, dx, \tag{B.45}$$

with

$$\check{f}(x, \rho_{min}, \rho_{max}, \theta_{min}, Z) = \left[x^2(\rho_{min}^2 - \rho_{max}^2)^2 - \sin^2\theta_{min}\rho_{min}^2\left(x^2 - \rho_{max}^2\right)^2\right]^{-\frac{1}{2}}, \tag{B.46}$$

and

$$\widetilde{f}(x, \rho_{max}, Z) = \left[\frac{Z^2}{4(Z-1)}x^2(1 - \rho_{max}^2)^2 - \left(x^2 - \rho_{max}^2\right)^2\right]^{-\frac{1}{2}}. \tag{B.47}$$

B.4 The bcc foam

While the total volume of the bubble, V_0, is constant, the volume of each of the cones is no longer required to be constant. The constraint on the individual cone volumes is now given by

$$8V_h + 6V_s = V_0,$$

where V_s and V_h denote the volumes of the cones associated with square and hexagonal bubble contact areas, respectively.

The second additional complication is the determination of the opening angles θ_h and θ_s of each type of cone (see Fig. B.1). We choose to retain the values of the solid angles subtended by each type of face in the “dry" Kelvin structure. This ensures that the sum of the solid angles subtended by the eight hexagonal and six square faces is equal to the 4π steradian solid angle of our bubble.

Two cones which meet each other are required to have a common slant height r_s (see Fig. B.1) so that their curved caps match. In the original Z-cone model (with identical cones) we required each of the curved caps to meet

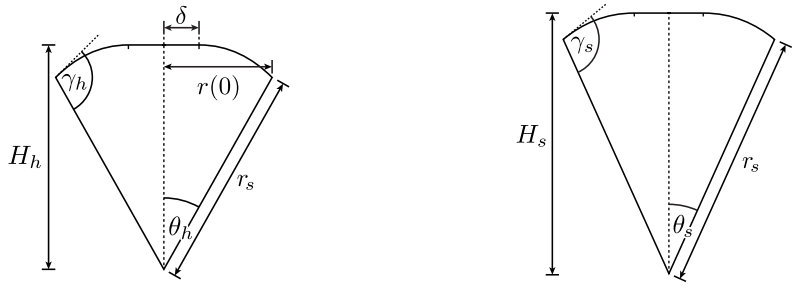

Figure B.1: In the extended cone model we deal with *two* types of cones: (a) for the hexagonal $\langle 1\,1\,1\rangle$ contacts and (b) for the square $\langle 1\,0\,0\rangle$. They share a common slant height r_s. The ratio $\delta/r(0) = \rho_\delta$ (where δ is the radius of the contact, and $r(0)$ the maximum width of the cone, as shown) features in the derivation of the cone model expressions.

their respective cones at right angles. In the case of the bcc bubble, this is more subtle. As can be seen in Fig. 3.24(a), a square cone joins only with four hexagonal cones (with corresponding angle γ_s), whereas a hexagonal cone joins with three square cones and three hexagonal (with corresponding angles γ_{hs} and γ_{hh} respectively). Smoothness requires $\gamma_{hh} = \pi/2$, and that

$$\gamma_s + \gamma_{hs} = \pi \tag{B.48}$$

The cone model requires a single angle, γ_h, which we define as an average $\gamma_h = (\gamma_{hh} + \gamma_{hs})/2$. Hence we can rewrite (B.48) as

$$2\gamma_h + \gamma_s = \frac{3\pi}{2}.$$

Similarly to the generalized volume condition above, the angles γ_h and γ_s are no longer fixed, as in the Z-cone model. We can make use of the following constraints to uniquely determine V_h, θ_h and γ_h for any given liquid fraction ϕ.

The first constraint we impose is that the ratio ν of cone heights is constant with respect to liquid fraction. Adopting the solid angles from the dry Kelvin structure, and requiring that cone slant heights match, fixes this ratio:

$$\nu = \frac{H_h}{H_s} = \frac{\cos\theta_h}{\cos\theta_s} = 0.8644. \tag{B.49}$$

Note that for a real Kelvin foam the corresponding ratio is $\sqrt{3}/2 \approx 0.866$.

The final constraint on our problem is that the internal pressure p in each of the neighbouring cones should be equal. This is simply the statement that pressure does not depend on the position in the bubble. The internal pressure of a bubble is responsible for the curvature of its surface and, by considering the work done to increase the volume of each cone by a small amount ΔV_i, while keeping the size of the contact constant (*i.e.* blowing it up slightly), we arrive at the following expression for the internal pressure of a cone,

$$p_i = \frac{\Delta E_i}{\Delta V_i^*} - 2\pi r_{s_i} \cos\gamma_i \cos\theta_i \frac{\Delta r_{s_i}}{\Delta V_i^*}, \tag{B.50}$$

where Δr_{s_i} is the slant height change of a cone, ΔE_i is the surface energy change and ΔV_i^* the change in the volume associated with the curved surface of the bubble. The first term represents the work necessary to increase the free surface of the cone by an amount ΔV_i^*, while the second is the work done by the surface tension σ in changing the surface energy of the cone to account for the increase in slant height. It can be thought of in terms of reduction of curvature:

$$W_{\sigma_i} = 2\pi r_{s_i} \Delta r_{s_i} \cos\theta_i \cos\gamma_i.$$

The above constraints are sufficient to determine all of the variables in our problem and to write the excess surface energy ε for the Kelvin cone model as

$$\varepsilon(\rho_{\delta_h}, \rho_{\delta_s}, \theta_h, \theta_s, \Gamma_h, \Gamma_s, q_h, q_s) = \frac{8A_h(\rho_{\delta_h}, \theta_h, \Gamma_h, q_h) + 6A_s(\rho_{\delta_s}, \theta_s, \Gamma_s, q_s)}{4\pi} - 1.$$

Here, the area of each face is given by

$$A_i(\rho_{\delta_i}, \theta_i, \Gamma_i, q_i) = \pi R_0^2 \left(\frac{4q_i}{3J(\rho_{\delta_i}, \Gamma_i) + \cot\theta_i}\right)^{\frac{2}{3}} \left[\rho_{\delta_i}^2 + 2(1-\rho_{\delta_i}^2)K(\rho_{\delta_i}, \Gamma_i)\right],$$

and the angles Γ_h and Γ_s are related to γ_h and γ_s via

$$\Gamma_i = \gamma_i + \theta_i - \frac{\pi}{2}.$$

The quantities q_s and q_h are fractions of the total volume V_0 taken up by any one of the square or hexagonal cones:

$$q_s = \frac{V_s}{V_0}; \quad q_h = \frac{V_h}{V_0},$$

and ρ_δ is given by the ratio of lengths $\delta/r(0)$ as shown in Fig. B.1(a).

The liquid fraction can be similarly expressed in these terms as

$$\phi(\rho_{\delta_h}, \theta_h, \theta_s, \Gamma_h, q_h) = 1 - \frac{3J(\rho_{\delta_h}, \Gamma_h) + \cot\theta_h}{2q_h\left[I(\rho_{\delta_h}, \Gamma_h) + \cot\theta_h\right]^3\left(4\tan^2\theta_h\left(\frac{3}{v^3}\right) + \tan^2\theta_s\right)}$$

The elliptic integrals are given explicitly by:

$$\begin{aligned}
I(\rho_\delta, \Gamma_i) &= \int_{\rho_{\delta_i}}^{1} \sin\Gamma_i\left(x^2 - \rho_{\delta_i}^2\right) f(x, \rho_{\delta_i}, \Gamma_i)\, dx, \\
J(\rho_\delta, \Gamma_i) &= \int_{\rho_{\delta_i}}^{1} \sin\Gamma_i x^2\left(x^2 - \rho_{\delta_i}^2\right) f(x, \rho_{\delta_i}, \Gamma_i)\, dx, \text{ and} \\
K(\rho_\delta, \Gamma_i) &= \int_{\rho_{\delta_i}}^{1} x^2 f(x, \rho_{\delta_i}, \Gamma_i)\, dx,
\end{aligned}$$

with

$$f(x, \rho_{\delta_i}, \Gamma_i) = \left[x^2(1-\rho_{\delta_i}^2)^2 - \sin\Gamma_i^2\left(x^2 - \rho_{\delta_i}^2\right)^2\right]^{-\frac{1}{2}}.$$

Appendix C

Simulation details

C.1 fcc and hcp bubbles

For simplicity, and ease of visualization, we do not use full periodic boundary conditions to simulate fcc and hcp bubbles. Rather, we take a single bubble and impose one-sided constraint planes on every vertex, edge and facet of the bubble. These are perhaps best visualized as plates which compress the bubble in its nearest-neighbour directions.

The nearest-neighbour directions for an fcc bubble are $(\pm 1, \pm 1, 0)$ plus permutations, to give 12 directions in total. We use these directions to construct the constraints. For the plate in the $(1, 1, 0)$ direction, for example we write:

```
constraint 1 nonpositive
formula: x + y = limit
```

where `limit` is defined as $1/\sqrt{2}$ times the bubble–bubble distance. Hence, for a bubble of diameter 1, we begin with a value of `limit = 1/sqrt(2)` to represent the wet limit, recovering an undeformed sphere as expected. We decrease the value of `limit` to deform the bubble.

For hcp, 9 out of the 12 directions are the same. For the three which differ, we replace $(1, 1, 0)$ with $(1/3, 1/3, 4/3)$, $(1, 0, 1)$ with $(1/3, 4/3, 1/3)$, and $(0, 1, 1)$ with $(4/3, 1/3, 1/3)$. Otherwise the process is identical.

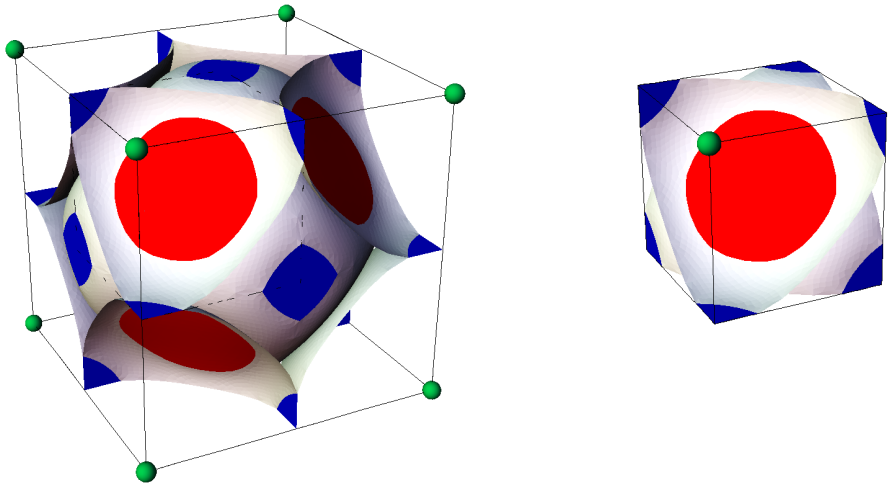

(a) (b)

Figure C.1: (a) The equilibrium structure for a Kelvin foam, including all surfaces within the conventional cell. The $\langle 1\,1\,1 \rangle$ contact faces are shown in red, and the $\langle 1\,0\,0 \rangle$ faces in blue. (b) We exploit reflectional symmetries to obtain a representative cell one eighth the size of (a). The full foam can be built from reflected and translated copies of the representative cell.

C.2 Kelvin foam

A method using constraint planes, as above, does not suffice for the bcc foam. We know that at $\phi = 0$ the hexaonal $\langle 1\,1\,1 \rangle$ faces cannot be flat, as this would violate Plateau's third and fourth laws. By contrast, we know that the square $\langle 1\,0\,0 \rangle$ faces *must* be flat, by the symmetries of the lattice.

Fig. C.1(a) shows the conventional bcc cell for a foam. The Kelvin foam consists of repeated translated copies of this cell. However, for simplicity we can exploit some of the symmetries of the conventional cell: namely, reflectional symmetry in the x, y and z directions (Brakke and Sullivan [93] exploit even more symmetries to yield a minimal representation of the fully dry Kelvin foam). Hence we arrive at a reduced cell (Fig. C.1(b)), which has one eighth of the volume of the conventional cell, and is composed of a cube containing one eighth of a bubble at each of two opposite corners. This in-

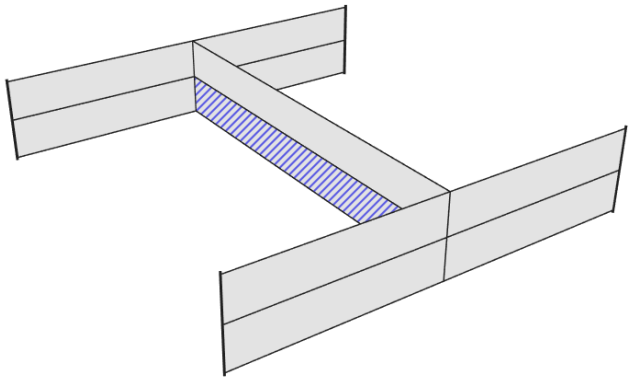

Figure C.2: Initial configuration of films for the quasi-2d simulation. The shaded film has a surface tension of 0. This is iteratively refined and relaxed to arrive at an arragement of minimal surface area.

creases the speed of computation considerably.

We begin with a very roughly triangulated approximation of the configuration in Fig. C.1(b), with appropriate film edges constrained to lie within the faces of the cube, *i.e.* planes of reflection. Minimisation of energy results in films meeting the faces of the cube at 90°, which ensures that the resulting foam structure is smooth. We note that in order to faithfully represent the full foam, films which lie *within* these planes (in this case, the blue $\langle 1\,0\,0 \rangle$ faces) are given half of their 'real' surface tension. Hence we give the red $\langle 1\,1\,1 \rangle$ contact face a tension of 2 and all other facets a tension of 1. Iterated mesh refinements and gradient-descent minimisations yield the configuration shown in Fig. C.1(b): the same surface is visualized as a single bubble in Fig. 3.23(a).

C.3 Quasi-2d fibre setup

Fig. C.2 shows the initial configuration of facets input into Surface Evolver for the quasi-2d setup. Edges which lie on the top plate are constrained to remain within that plane, and similarly for the bottom plate. The edges located at the pins are fixed in place. The shaded facet is assigned a surface tension of 0 to

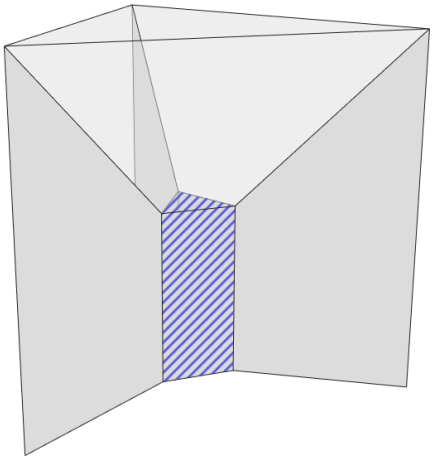

Figure C.3: Initial configuration of films for the Plateau border simulation. The shaded films have a surface tension of 0. The upper films have been rendered semi-transparent for ease of visualization.

represent the absence of a soap film in the region between the plates spanned by the fibre; all other surfaces have tension 1. The top edge of the shaded facet is also constrained to remain at its assigned height.

C.4 Single Plateau border

For this simulation we exploit a symmetry of the experiment. The setup has reflectional symmetry in the direction of the axis of the prism, and so we only need to simulate half of the films.

Fig. C.3 shows the initial configuration of facets we used in this case. The fibre is represented initially by three facets of surface tension 0 which are constrained to lie within a cylindrical shell, whose radius is specified to match the radius of the fibre in the experiment, and whose position can be varied in the x direction, in line with the experiment. The edges which represent the wires of the wireframe are fixed in place, and those edges which lie in the plane of symmetry (the bottom, as viewed here) we constrain to remain in that plane.

Here we use the *quadratic model*, in which edges are represented by quadratic splines. We found that the default linear model resulted in much noisier $l_1(x)$

and $l_2(x)$ curves as the fibre was moved, despite attaining energies very close to the quadratic case.

References

[1] J Plateau. *Statique expérimentale et théorique des liquides soumis aux seules forces moléculaires*. Gauthier-Villars, 1873.

[2] CV Boys. *Soap-bubbles, and the forces which mould them*. Cornell University Library, 1890.

[3] D Weaire. A commentary on the historical background to the modern study of foam structures. *Philosophical Magazine*, 93(31-33):4132–4137, 2013.

[4] D Weaire and JP Kermode. Computer simulation of a two-dimensional soap froth: I. method and motivation. *Philosophical Magazine B*, 48(3):245–259, 1983.

[5] KA Brakke. The Surface Evolver. *Experimental Mathematics*, 1(2):141–165, 1992.

[6] D Weaire and R Phelan. A counter-example to Kelvin's conjecture on minimal surfaces. *Philosophical Magazine Letters*, 69(2):107–110, 1994.

[7] R Phelan, D Weaire, and KA Brakke. Computation of equilibrium foam structures using the Surface Evolver. *Experimental Mathematics*, 4(3):181–192, 1995.

[8] D Weaire and S Hutzler. *The Physics of Foams*. Oxford University Press, 2001.

[9] R Höhler, Y Yip Cheung Sang, E Lorenceau, and S Cohen-Addad. Osmotic pressure and structures of monodisperse ordered foam. *Langmuir*, 24(2):418–425, 2008.

[10] I Cantat, S Cohen-Addad, F Elias, F Graner, R Höhler, R Flatman, O Pitois, F Rouyer, and A Saint-Jalmes. *Foams: Structure and Dynamics*. Oxford University Press, 2013.

[11] JD Bernal and J Mason. Packing of spheres: co-ordination of randomly packed spheres. *Nature*, 188:910–911.

[12] J Pierre, V Leroy, A Saint-Jalmes, B Dollet, I Ben Salem, J Crassous, R-M Guillermic, W Drenckhan, and F Elias. Shaving foam: A complex system for acoustic wave propagation. *The Journal of the Acoustical Society of America*, 133(5), 2013.

[13] JE Taylor. The structure of singularities in soap-bubble-like and soap-film-like minimal surfaces. *Annals of Mathematics*, 103(3):489–539, 1976.

[14] Sir W Thomson. On the division of space with minimum partitional area. *Acta Mathematica*, 11(1-4):121–134, 1887.

[15] R Kusner and JM Sullivan. Comparing the Weaire–Phelan equal-volume foam to Kelvin's foam. *Forma*, 11(3):233–242, 1996.

[16] R Gabbrielli, AJ Meagher, D Weaire, KA Brakke, and S Hutzler. An experimental realization of the Weaire–Phelan structure in monodisperse liquid foam. *Philosophical Magazine Letters*, 92(1):1–6, 2012.

[17] J Kepler. Strena, seu de nive sexangula. 1611.

[18] TC Hales. A proof of the Kepler conjecture. *Annals of Mathematics*, 162(3):1065–1185, 2005.

[19] W Drenckhan and S Hutzler. Structure and energy of liquid foams. *Advances in colloid and interface science*, 224:1–16, 2015.

[20] JM Sullivan. The geometry of bubbles and foams. In *Foams and emulsions*, pages 379–402. Springer, 1999.

[21] DJ Durian, DA Weitz, and DJ Pine. Scaling behavior in shaving cream. *Physical Review A*, 44(12):R7902, 1991.

[22] D Weaire, S Hutzler, G Verbist, and EAJF Peters. A review of foam drainage. *Advances in Chemical Physics, Volume 102*, pages 315–374, 1997.

[23] SA Koehler, HA Stone, MP Brenner, and J Eggers. Dynamics of foam drainage. *Physical Review E*, 58(2):2097, 1998.

[24] A Bhakta and E Ruckenstein. Decay of standing foams: drainage, coalescence and collapse. *Advances in Colloid and Interface Science*, 70:1–124, 1997.

[25] HM Princen. Highly concentrated emulsions. I. cylindrical systems. *Journal of Colloid and Interface Science*, 71(1):55–66, 1979.

[26] HM Princen. Osmotic pressure of foams and highly concentrated emulsions. i. theoretical considerations. *Langmuir*, 2(4):519–524, 1986.

[27] L Bragg and JF Nye. A dynamical model of a crystal structure. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 190(1023):474–481, 1947.

[28] A Gouldstone, KJ Van Vliet, and S Suresh. Nanoindentation: Simulation of defect nucleation in a crystal. *Nature*, 411(6838):656–656, 2001.

[29] AS Argon and HY Kuo. Plastic flow in a disordered bubble raft (an analog of a metallic glass). *Materials science and Engineering*, 39(1):101–109, 1979.

[30] A van der Net, L Blondel, A Saugey, and W Drenckhan. Simulating and interpretating images of foams with computational ray-tracing techniques. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 309(1):159–176, 2007.

[31] A van der Net, GW Delaney, W Drenckhan, D Weaire, and S Hutzler. Crystalline arrangements of microbubbles in monodisperse foams. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 309(1):117–124, 2007.

[32] AB Yu, XZ An, RP Zou, RY Yang, and K Kendall. Self-assembly of particles for densest packing by mechanical vibration. *Physical Review Letters*, 97(26):265501, 2006.

[33] AJ Meagher, D Whyte, J Banhart, S Hutzler, D Weaire, and F García-Moreno. Slow crystallisation of a monodisperse foam stabilised against coarsening. *Soft Matter*, 11(23):4710–4716, 2015.

[34] P Garstecki, I Gitlin, W DiLuzio, GM Whitesides, E Kumacheva, and HA Stone. Formation of monodisperse bubbles in a microfluidic flow-focusing device. *Applied Physics Letters*, 85(13):2649–2651, 2004.

[35] PJ Steinhardt, DR Nelson, and M Ronchetti. Bond-orientational order in liquids and glasses. *Physical Review B*, 28(2):784, 1983.

[36] T Aste, M Saadatfar, and TJ Senden. Geometrical structure of disordered sphere packings. *Physical Review E*, 71(6):061302, 2005.

[37] AD McNaught, A Wilkinson, et al. Compendium of chemical terminology, IUPAC recommendations. 1997.

[38] C Bender. Bestimmung der grössten anzahl gleich grosser kugeln, welche sich auf eine kugel von demselben radius, wie die übrigen, auflegen lassen. *Archiv Math. Physik*, 56:302–306, 1874.

[39] VN Manoharan and DJ Pine. Building materials by packing spheres. *Mrs Bulletin*, 29(02):91–95, 2004.

[40] V Luchnikov, A Gervois, P Richard, L Oger, and JP Troadec. Crystallization of dense hard sphere packings: competition of hcp and fcc close order. *Journal of Molecular Liquids*, 96:185–194, 2002.

[41] W Lechner and C Dellago. Accurate determination of crystal structures based on averaged local bond order parameters. *The Journal of Chemical Physics*, 129(11):114707, 2008.

[42] W Mickel, SC Kapfer, GE Schröder-Turk, and W Mecke. Shortcomings of the bond orientational order parameters for the analysis of disordered particulate matter. *The Journal of Chemical Physics*, 138(4):044501, 2013.

[43] S Heitkam, W Drenckhan, and J Fröhlich. Packing spheres tightly: influence of mechanical stability on close-packed sphere structures. *Physical Review Letters*, 108(14):148302, 2012.

[44] A Van der Net, W Drenckhan, D Weaire, and S Hutzler. The crystal structure of bubbles in the wet foam limit. *Soft Matter*, 2(2):129–134, 2006.

[45] LV Woodcock. Entropy difference between the face-centred cubic and hexagonal close-packed crystal structures. *Nature*, 385:141–143, 1997.

[46] NG Szwacki and T Szwacka. *Basic elements of crystallography*. CRC Press, 2016.

[47] D Weaire, N Pittet, S Hutzler, and D Pardal. Steady-state drainage of an aqueous foam. *Physical Review Letters*, 71(16):2670, 1993.

[48] KA Brakke. Instability of the wet cube cone soap film. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 263(1):4–10, 2005.

[49] D Weaire. *The Kelvin Problem*. CRC Press, 1997.

[50] DC Morse and TA Witten. Droplet elasticity in weakly compressed emulsions. *EPL (Europhysics Letters)*, 22(7):549, 1993.

[51] DJ Durian. Foam mechanics at the bubble scale. *Physical Review Letters*, 75(26):4780, 1995.

[52] M-D Lacasse, GS Grest, D Levine, TG Mason, and DA Weitz. Model for the elasticity of compressed emulsions. *Physical Review Letters*, 76(18):3448, 1996.

[53] M-D Lacasse, GS Grest, and D Levine. Deformation of small compressed droplets. *Physical Review E*, 54(5):5436, 1996.

[54] JM Ziman. A theory of the electrical properties of liquid metals. I: the monovalent metals. *Philosophical Magazine*, 6(68):1013–1034, 1961.

[55] S Hutzler, R Murtagh, D Whyte, ST Tobin, and D Weaire. Z-cone model for the energy of an ordered foam. *Soft Matter*, 10(36):7103–7108, 2014.

[56] AJ Meagher, D McAteer, S Hutzler, and D Weaire. Building the pyramids: perfect bubble crystals. *Philosophical Magazine*, 93(31-33):4138–4150, 2013.

[57] D Whyte, R Murtagh, D Weaire, and S Hutzler. Applications and extensions of the Z-cone model for the energy of a foam. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 473:115–122, 2015.

[58] I Cantat, C Poloni, and R Delannay. Experimental evidence of flow destabilization in a two-dimensional bidisperse foam. *Physical Review E*, 73(1):011505, 2006.

[59] G Katgert, A Latka, ME Möbius, and M van Hecke. Flow in linearly sheared two-dimensional foams: From bubble to bulk scale. *Physical Review E*, 79(6):066318, 2009.

[60] EB Matzke and J Nestler. Volume-shape relationships in variant foams. a further study of the role of surface forces in three-dimensional cell shape determination. *American Journal of Botany*, pages 130–144, 1946.

[61] AM Kraynik, DA Reinelt, and F van Swol. Structure of random bidisperse foam. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 263(1):11–17, 2005.

[62] K Feitosa, Olivia L Halt, Randall D Kamien, and DJ Durian. Bubble kinetics in a steady-state column of aqueous foam. *EPL (Europhysics Letters)*, 76(4):683, 2006.

[63] A van der Net, D Weaire, and S Hutzler. Rearrangement and elimination of ordered surface layers of crystalline bubble structures due to gas diffusion. *Soft Matter*, 5(2):318–324, 2009.

[64] I Jorjadze, L-L Pontani, and J Brujic. Microscopic approach to the nonlinear elasticity of compressed emulsions. *Physical Review Letters*, 110(4):048302, 2013.

[65] DJ Durian. Bubble-scale model of foam mechanics: mmelting, nonlinear behavior, and avalanches. *Physical Review E*, 55(2):1739, 1997.

[66] A Maestro, W Drenckhan, E Rio, and R Höhler. Liquid dispersions under gravity: volume fraction profile and osmotic pressure. *Soft Matter*, 9(8):2531–2540, 2013.

[67] HM Princen and AD Kiss. Osmotic pressure of foams and highly concentrated emulsions. 2. determination from the variation in volume fraction with height in an equilibrated column. *Langmuir*, 3(1):36–41, 1987.

[68] K Feitosa, S Marze, A Saint-Jalmes, and DJ Durian. Electrical conductivity of dispersions: from dry foams to dilute suspensions. *Journal of Physics: Condensed Matter*, 17(41):6301, 2005.

[69] MU Vera, A Saint-Jalmes, and DJ Durian. Scattering optics of foam. *Applied Optics*, 40(24):4210–4214, 2001.

[70] E Solórzano, S Pardo-Alonso, JA de Saja, and MA Rodríguez-Pérez. Study of aqueous foams evolution by means of x-ray radioscopy. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 438:159–166, 2013.

[71] S Mukherjee and H Wiedersich. Morphological and viscoelastic properties of dense foams generated from skin cleansing bars. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 95(2):159–172, 1995.

[72] E Forel, E Rio, M Schneider, S Beguin, D Weaire, S Hutzler, and W Drenckhan. The surface tells it all: Relationship between bulk and surface fractions of liquid dispersions. *Soft Matter*, 2016.

[73] NW Ashcroft and ND Mermin. Solid State Physics (Holt, Rinehart and Winston, New York, 1976). 2005.

[74] D Weaire. Structural transformations in foam. *Philosophical Magazine Letters*, 69(2):99–105, 1994.

[75] C Zener. *Elasticity and anelasticity of metals*. University of Chicago press, 1948.

[76] B Fultz. *Phase Transitions in Materials*. Cambridge University Press, 2014.

[77] MK Smith, VW Punton, and AG Rixson. The structure and properties of paper formed by a foaming process. *Tappi*, 57(1):107–111, 1974.

[78] AM Al-Qararah, T Hjelt, A Koponen, A Harlin, and JA Ketoja. Bubble size and air content of wet fibre foams in axial mixing with macro-instabilities. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 436:1130–1139, 2013.

[79] AM Al-Qararah, T Hjelt, A Koponen, A Harlin, and JA Ketoja. Response of wet foam to fibre mixing. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 467:97–106, 2015.

[80] AM Al-Qararah, A Ekman, T Hjelt, JA Ketoja, H Kiiskinen, A Koponen, and J Timonen. A unique microstructure of the fiber networks deposited from foam–fiber suspensions. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 482:544–553, 2015.

[81] M Hazewinkel. *Encyclopaedia of Mathematics*. Springer, 1987.

[82] S Hutzler, D Weaire, SJ Cox, A Van der Net, and Eric Janiaud. Pre-empting Plateau: The nature of topological transitions in foam. *EPL (Europhysics Letters)*, 77(2):28002, 2007.

[83] DGT Barrett, S Kelly, EJ Daly, MJ Dolan, W Drenckhan, D Weaire, and S Hutzler. Taking Plateau into microgravity: The formation of an eight-fold vertex in a system of soap films. *Microgravity-Science and Technology*, 20(1):17–22, 2008.

[84] FE Bartell and B Roger Ray. Wetting characteristics of cellulose derivatives. I. Contact angles formed by water and by organic liquids. *Journal of the American Chemical Society*, 74(3):778–783, 1952.

[85] A Koponen. Private communication, 2016.

[86] R Tadmor. Line energy and the relation between advancing, receding, and Young contact angles. *Langmuir*, 20(18):7659–7664, 2004.

[87] O Krichevsky and J Stavans. Coarsening of two-dimensional soap froths in the presence of pinning centers. *Physical Review B*, 46(17):10579, 1992.

[88] B Haffner, FF Dunne, S Burke, and S Hutzler. Ageing of fibre-laden aqueous foams. *Cellulose*, 2016 (submitted).

[89] C Duprat and S Protiere. Capillary stretching of fibers. *EPL (Europhysics Letters)*, 111(5):56006, 2015.

[90] KS Birdi, DT Vu, and A Winter. A study of the evaporation rates of small water drops placed on a solid surface. *The Journal of Physical Chemistry*, 93(9):3702–3703, 1989.

[91] JA van Meel, L Filion, C Valeriani, and D Frenkel. A parameter-free, solid-angle based, nearest-neighbor algorithm. *The Journal of Chemical Physics*, 136(23):234107, 2012.

[92] PW Jones, A Osipov, and V Rokhlin. Randomized approximate nearest neighbors algorithm. *Proceedings of the National Academy of Sciences*, 108(38):15679–15686, 2011.

[93] KA Brakke and JM Sullivan. Using symmetry features of the Surface Evolver to study foams. In H-C Hege and K Polthier, editors, *Visualization and Mathematics: Experiments, Simulations and Environments*, pages 95–117. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.