Very Viscous Flows Driven by Gravity

with particular application to

Slumping of Molten Glass

Yvonne Marie Stokes, B.Sc. (Murdoch), B.Sc. (Hons) (Adelaide)

Thesis submitted for the degree of

Doctor of Philosophy

in

Applied Mathematics

at

The University of Adelaide

(Faculty of Mathematical and Computer Sciences)

Department of Applied Mathematics

July, 1998

Contents

List of Tables

List of Figures

Abstract

This thesis examines the flow of very viscous Newtonian fluids driven by gravity. It is written with concern for specific applications in the optics industry, with emphasis on the slumping of molten glass into a mould, as in the manufacture of optical components, which are in turn used to manufacture ophthalmic lenses. This process is known as *thermal replication*. However, the work has more general applicability, and disc viscometry, used to determine the viscosity of very viscous fluids, is also considered. In addition, one chapter of the thesis is devoted to the flow of dripping honey, as another example of a very viscous flow to which the model can be applied.

The Stokes creeping-flow equations are used to model the very viscous flows of interest. The main solution method is finite elements, and a purpose-written computer program has been developed to solve the creeping-flow equations by this method. The present program is restricted to solving for either two-dimensional or axisymmetric flows but is extendible to three dimensions. In addition, semi-analytic series and asymptotic methods are used for some small portions of the work.

The optical applications of this work demand consideration of the topic of computing surface curvature, and therefore second derivatives, from inexact and discrete numerical and experimental data. For this purpose, fitting of B-splines by a least-squares method to coordinate data defining the surface has been used.

Much of the work assumes isothermal conditions, but in the context of the accuracy required in optical component manufacture it is also possible that non-isothermal effects will be important. Consequently, this restriction is eventually

relaxed and some consideration given to non-isothermal conditions.

In order to validate the creeping-flow model and finite-element program, comparisons of numerical simulations with experimental results are performed. A preliminary assessment of the importance of non-isothermal conditions to the thermal-replication process is also made by comparing isothermal and non-isothermal simulations with experimental results. The isothermal model is found to best match the experimental data.

Signed Statement

This work contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being available for loan and photocopying.

Signed: .....

Date: .....

Acknowledgements

The completion of this thesis after three and a half years of work has not been without greatly appreciated assistance and input from many people.

First and foremost of these is my supervisor Prof. Ernie Tuck, who gave me a splendid start into the subject of very viscous flows, and from whom I have learned much about mathematical modelling in general. I never cease to admire and value his ability to quickly get to the heart of a matter, and give, what has always proved to be, sound advice and direction, even where it may be outside his specific area of expertise.

I am also most grateful to Dr Michael Teubner (Department of Applied Mathematics). His interest in my work and willingness to help, as well as the lengthy discussions we had, provided some much needed stimulus. Dr Peter Gill (Head of the Department of Applied Mathematics), despite an always busy schedule, provided considerable support throughout, and particular assistance and encouragement in some difficult times.

This research was supported by an APA (Industry) scholarship, the industry partner being SOLA International Holdings Ltd. This gave the work some very practical direction, which I believe has been beneficial. I wish to thank all at SOLA who gave time to the project. Of these I remember particularly Mr John Pocket, who was involved with the project at its inception but then moved on to other things. He spent a considerable amount of time with me at the beginning, explaining SOLA's mould manufacturing process. All of the new experimental work done in

the course of this project, was conducted at SOLA's Research Centre at Lonsdale, South Australia, and valuable resources were made available for this. I am especially grateful to Mr Paul Coldrey for all his work in setting up and carrying out all the different experiments I requested.

For provision of data on properties of specific Corning glasses, I am indebted primarily to Dr Paul Danielson, and also to Dr Karen Geisinger, both of Corning Incorporated Sullivan Park Laboratories. Paul first volunteered his services in response to an open request made to the Glass Mailing List — which thus proved itself a useful facility — and then did his very best to answer all of my questions, and fulfil all of my requests. Access to the Corning technical reports and raw experimental data relating to slumping work done at Corning in the mid-1970's was also facilitated by him, but was finally made possible by Dr Suresh Gulati and Mr Tom Loomis. I was lucky enough to be able to visit Corning Incorporated Sullivan Park Laboratories in New York State during May of 1997 and spend some time with Suresh — one of the researchers originally involved in the slumping work — going through his files and identifying information relevant to my work. I am truly very grateful that he made that time and his results available. I also wish to thank Corning Incorporated, and Paul, Suresh, and Tom in particular, for the hospitality shown to me during my visit, which was not only of value to my work, but also greatly enjoyed.

Thanks are also due to Dr Sylvan Elhay (Computer Science) for his direction and suggestions regarding calculation of derivatives from numerical data using B-spline approximations, to Dr A.N. (Nick) Stokes of CSIRO Mathematical and Information Sciences for the help given me while I was coming to grips with the use of *Fastflo*, to Prof. Len Schwartz (Department of Mechanical Engineering, University of Delaware) for his patience and promptness in explaining via email some aspects of the 'dripping-honey' problem, and to Mr Leo Lazauskas for help in chasing literature on the 'thickening-window' myth. Dr Peter Howell (Mathematical Institute, Oxford,

UK) and Prof. Bob Mattheij (Department of Mathematics and Computing Science, Eindhoven Univeristy of Technology) kindly sent me information and papers that were not easily obtainable in Australia.

For assistance with computing matters I owe special thanks to Mr David Beard, systems administrator in the Department of Applied Mathematics, and also to Dr Francis Vaughan (Computer Science) and Dr Jim Denier (Applied Mathematics). Both staff and students within the Faculty of Mathematics and Computer Science at the University of Adelaide provided a friendly and supportive environment in which to work, and I specifically mention Mrs Dianne Parish, the Applied Mathematics secretary, whose always cheerfully given help with administrative matters has been much appreciated.

Chapter 1

Introduction

The viscosity of a fluid is the ratio between shear stress and the resulting time rate of strain or velocity gradient [7, p. 36]. It is a measure of the internal resistance of the fluid to deformation. A very viscous fluid is characterized by a large viscosity, and has large internal resistance to deformation.

Some common household examples of very viscous fluids include honey; syrups such as maple syrup, malt, molasses; motor-oil additives; and cooking mixtures. Such fluids are found to be more difficult to stir or pour than low-viscosity fluids such as water. In nature we see many spiders and insect larvae produce very viscous fluids from which they spin webs and cocoons. Beneath its crust, the earth consists of very viscous molten rock, which is seen only when forced to the surface in volcanic eruptions and larva flows.

Very viscous fluids feature in a number of industrial processes. Polymers, such as nylon, on heating become very viscous fluids, that can be spun into threads for use in textiles, or injection, press or blow moulded to form the plastic items that have become so common in the twentieth century. Molten glass is yet another very viscous fluid which has long been used to produce items of both artistic and practical value — windowpanes, containers of various forms, mirrors, and optical lenses to name just a few.

The primary focus of this thesis is on the simulation of slumping flows of molten glass due to self weight under the influence of gravity, using a creeping-flow model and finite-element methods. Of particular interest is the industrial process of thermal replication [84], which has application to the production of aspheric surfaces for optical uses. Thermal replication involves placing a glass workpiece on a mould, and then heating this combination in an oven so that the glass melts and slumps into the mould. During slumping, the lower surface of the glass workpiece contacts the mould, and hence is affected by the roughness of the mould. However, a very smooth profile is obtained on the upper surface of the glass workpiece. The process is described in more detail in [84] and in Chapter 6 of this thesis. At this stage it is sufficient to note that the idea is to obtain a near replica of the mould surface on the upper surface of the glass workpiece, that is sufficiently smooth for optical applications. The curvature profile of this surface is critical; a model that can accurately predict the final curvature of the upper glass surface from some initial glass workpiece and mould combination, would be of considerable benefit to industry.

A considerable amount of research into the modelling of glass forming has been carried out over the past ten to fifteen years, but this research, with few exceptions, relates to container manufacture by blow moulding of thin glass films. There are two main stages to that process: the forming of an initial shape, known as the parison, usually by pressing a quantity of molten glass in a mould; this still molten parison is then transferred to another mould and air blown into it so that it expands to fill that mould. Most modelling of blow moulding has been focused on the second blowing stage.

In 1984 a paper appeared by Cormeau, Cormeau, and Roose [21] in which a creeping-flow finite-element model for simulating only the blowing stage is presented, along with some initial computational results. Further work of a similar nature followed, by Williams, Owen, and Sa in 1986 [100], and then some years later by Burley and Graham (1991) [14] and Graham, Burley, and Carling (1992)

[38]. Although one flow is driven by gravity and the other by pressure, there is considerable similarity between thermal replication and the blowing of a molten-glass parison out to a mould, and similar modelling techniques to those outlined in these papers on blow moulding have been employed in this thesis. With respect to the first parison-forming stage of blow moulding, just recently (1996) Simons and Mattheij [83] published their work on pressing (or stamping) of parisons, for which both finite-element and boundary-element methods were used.

Methods other than finite elements may also be used to analyse both glass slumping and glass blowing. One used in this thesis is thin-layer approximation, and the equations obtained are comparable with thin-layer equations for a creeping-flow model of pressure driven flows of thin films that have been recently derived by van de Fliert, Howell, and Ockendon (1995) [96], and Howell (1996) [48].

The issue of major concern in container manufacture, and because of which the research referred to above has been carried out, is achieving a uniform wall thickness in the finished product, of a strength that is appropriate to its intended use. Above all, thin weak spots must be avoided, while it is also desirable to keep material costs to a minimum and not have excessive wall thickness either. This issue is quite different to the issue of accurate surface curvature that is of major concern in thermal replication of optical surfaces. Consequently, although a creeping-flow model and finite-element methods similar to those used in some of the above-mentioned work, have been used in the present thesis, it differs from previous work in its concern with calculation of curvature from non-exact discrete numerical data yielded by finite-element simulations. Compared with film thickness, curvature is highly sensitive to small variations in the profile of a surface.

In the area of glass slumping, particularly in relation to the generation of optical surfaces, very little research has been reported in the published literature. A 1976 paper by Gulati, Fontana, and Plummer [42] reports investigations by Corning Glass Works into disc slumping as a method for determining glass viscosities

using a viscoelastic analogy model. The paper [42] also mentions unpublished work, that is nevertheless the subject of an internal Corning Glass Works report [41], on developing of disc sagging schedules for use in the manufacture of lens blanks of a desired surface curvature. Mould contact was not a feature of any of the Corning slumping work. Also, only average curvatures of lens surfaces, computed from the measured sag at the disc centre, assuming lens geometry to be spherical, were obtained, whereas the curvature profile over the whole of a progressive power optical surface is of interest, and is a matter addressed in this thesis.

Although glass slumping and thermal replication are the main focuses of this thesis, the creeping-flow model and finite-element program developed are applicable to a wide range of gravity-driven very viscous flows, both involving contact with a mould, and without mould contact. A simple household example considered here is honey dripping from a spoon. This is a flow that anyone can relate to and experiment with, and for that reason this work at one time aroused some media interest [6]. It is an example of an extensional flow, a class of motions that has received considerable attention in the published literature over many years, for example in relation to fibre spinning [28, 78].

The literature mentioned in this introduction will be further discussed throughout this thesis in the context of specific flow problems. Additional literature, relating to the processes considered and methods used, will also be surveyed, as the context requires and allows. The structure of the thesis is as follows.

First in Chapter 2 a constant-viscosity creeping-flow model applicable for gravity-driven very viscous flows is developed, and its solution by finite-element methods is discussed. A computer simulation program using a finite-element method is also developed, which is referred to throughout this thesis as the "purpose-written" program/code. Where there is no danger of confusion with other finite-element programs it may also be simply referred to as "the finite-element program/code".

In Chapter 3 we consider sagging of a very viscous liquid bridge, which is a two-

dimensional case of slumping without progressive mould contact. Initial motion is obtained by a semi-analytic method, as well as by finite elements using both a commercial CFD package *Fastflo* and the purpose-written code developed in Chapter 2. The purpose-written code is, in this manner, able to be verified, and is then used to solve for the flow at later times. An asymptotic method is used to solve for thin bridges.

The dripping of honey from a spoon is considered in Chapter 4 using both a slender-filament approximation and finite elements. The time and position at which a drop will break is given special attention.

Slumping of glass in the absence of a mould is the subject of Chapter 5, with the sagging work of Gulati et al. considered in some detail, and compared with finite-element simulations. Because the temperature changes in time, it is necessary to compute for a time-varying viscosity, and it is shown how this is possible using the constant-viscosity finite-element model and a time-varying time scale. At this stage it is assumed that, at any point in time, the viscosity is spatially constant throughout the glass. In addition, in this chapter a common myth concerning slow flow of glass windows over centuries is disproved by a thin-layer mathematical analysis not previously employed in relation to this subject.

In Chapter 6 we begin our specific analysis of thermal replication, developing methods for handling mould contact and computing surface-curvature profiles. A viscosity that varies in time but is spatially constant throughout the glass is still assumed. The transfer of mould roughness to the optical surface is also examined, with attention given to the importance of glass thickness relative to the coarseness of the mould surface.

Because even small changes to the shape of an optical surface have a large effect on the curvature profile of that surface, and because the viscosity of a glass, and hence its flow properties, are very sensitive to temperature, there is a real possibility that spatial temperature variations in the glass will significantly affect the final

optical surface. However, measurement of actual temperature throughout a slumping glass disc presents practical difficulties, and developing a full thermo-mechanical model of thermal replication is also very difficult. Hence, in Chapter 7, the finite-element model is extended to handle a spatially-varying viscosity, and a number of simulations run for different prescribed temperature gradients in the glass so as to qualitatively assess the effects of spatial temperature variations on the curvature of the final optical surface. A simple one-dimensional diffusion model is used to examine the magnitude of spatial temperature variations that possibly develop in slumping glass during the thermal-replication process, and give some guidance in the prescription of the temperature gradients.

Then, in Chapter 8, a comparison is made between experimental slumping results and slumping simulations, with and without spatial temperature gradients in the glass, the aim being to find whether a spatially constant-viscosity model is appropriate for thermal replication, or whether more complex heat-flow modelling is warranted.

Some conclusions from this work are given in Chapter 9. In particular it appears that, for the thermal-replication process, curvature profiles computed from finite-element simulations with the assumption of spatially constant glass viscosity correspond well with profiles computed using available experimental data. There are differences between numerical and experimental results in the time over which slumping occurs, which can be explained by incorrect viscosity data and corrected by multiplying viscosity by an appropriate factor.

Chapter 2

A Creeping-Flow Model

2.1 The Equations

As already indicated in the Introduction, there are many fluids that can be described as very viscous. Of these, some may be considered to be Newtonian, with stress being proportional to rate of strain, the viscosity μ being the constant of proportionality. Others are non-Newtonian, requiring more complicated models to describe accurately the relationship between stress and rate of strain (see [10]). Because of the uncertainty associated with non-Newtonian models and their application to specific flow problems, it is common to adopt the Newtonian assumption as a starting point in modelling, even for known non-Newtonian fluids.

A Newtonian fluid is commonly thought of as having constant viscosity both in time and space. However the definition, as given above, does not demand this, and in this thesis a fluid is described as Newtonian provided that there is proportionality between stress and rate of strain, even though the viscosity may vary temporally and/or spatially. In addition, conditions are said to be *isothermal* if the viscosity is spatially constant for all time (i.e. temporal variation is allowed). This terminology is appropriate in circumstances where variations in viscosity are associated with a time-varying temperature.

Throughout this thesis it is assumed that the very viscous fluids with which we shall be concerned are incompressible with constant density ρ, and Newtonian. For the present, it is also assumed that the viscosity is constant in both time and space, although this will be relaxed later. At this stage no attempt is made to justify these assumptions; that is left until we come to deal with specific fluids and flow problems.

Employing tensor notation let us denote the Cartesian coordinate system by x_i, and let x_2 be a vertically upward axis. With gravity g acting as the driving force vertically down in the $-x_2$ direction, and defining u_i as the velocity vector and p as pressure, the flow is given by solving [7, pp. 75,147,175] the Navier-Stokes equations

$$\rho\left(\frac{\partial u_i}{\partial t}+u_j\frac{\partial u_i}{\partial x_j}\right)=-\frac{\partial p}{\partial x_i}+\frac{\partial}{\partial x_j}\left[\mu\left(\frac{\partial u_i}{\partial x_j}+\frac{\partial u_j}{\partial x_i}\right)\right]-\rho g\delta_{2i}, \tag{2.1}$$

and the continuity equation for incompressible flow

$$\frac{\partial u_i}{\partial x_i}=0. \tag{2.2}$$

The use of the Kronecker delta in (2.1) restricts the contribution of the gravitational force (last) term to the second ($i=2$) equation.

Because the fluids to be modelled are very viscous with μ very much larger than $1\mathrm{Pa}\cdot\mathrm{s}$, we can further simplify (2.1). Let us denote characteristic length $\mathcal{L}$, velocity $\mathcal{U}$, time $\mathcal{T}=\mathcal{L}/\mathcal{U}$, and pressure $\mathcal{P}=\rho g\mathcal{L}$, and dimensionless quantities for position x_i^*, velocity u_i^*, time t^*, and pressure p^* such that

$$x_i=\mathcal{L}x_i^*,\quad u_i=\mathcal{U}u_i^*,\quad t=\mathcal{T}t^*,\quad \text{and}\quad p=\mathcal{P}p^*. \tag{2.3}$$

Then equations (2.1) become

$$\frac{\mathcal{U}}{\mathcal{T}g}\left(\frac{\partial u_i^*}{\partial t^*}+u_j^*\frac{\partial u_i^*}{\partial x_j^*}\right)=-\frac{\partial p^*}{\partial x_i^*}+\frac{\partial}{\partial x_j^*}\left[\frac{\mu\mathcal{U}}{\rho g\mathcal{L}^2}\left(\frac{\partial u_i^*}{\partial x_j^*}+\frac{\partial u_j^*}{\partial x_i^*}\right)\right]-\delta_{2i}. \tag{2.4}$$

On defining

$$\mathcal{U}=\frac{\rho g\mathcal{L}^2}{\mu} \tag{2.5}$$

so that all terms on the right-hand side of (2.4) are of order one, we have inertial terms on the left-hand side of these equations of order

$$\frac{\mathcal{U}}{\mathcal{T}g} = \frac{\mathcal{U}^2}{g\mathcal{L}} = \frac{\rho\mathcal{U}\mathcal{L}}{\mu} = \frac{\rho^2 g\mathcal{L}^3}{\mu^2}. \tag{2.6}$$

When this Reynolds number is much less than one, as can be expected for very viscous fluids, the inertial terms may be neglected yielding the (dimensionless) Stokes equations (cf. [7, p. 217])

$$\frac{\partial p^*}{\partial x_i^*} = \frac{\partial}{\partial x_j^*}\left(\frac{\partial u_i^*}{\partial x_j^*} + \frac{\partial u_j^*}{\partial x_i^*}\right) - \delta_{2i}. \tag{2.7}$$

If instead we multiply (2.4) by $\mathcal{T}g/\mathcal{U}$ so that the inertial terms on the left-hand side are of order one, we see that the inertial and viscous terms are of the same order only for a time of size

$$\mathcal{T} = \frac{\rho\mathcal{L}^2}{\mu}. \tag{2.8}$$

For fluids of large viscosity, this is typically an extremely short time period during which the fluid particles accelerate to a speed of order

$$g\mathcal{T} = \frac{\rho g\mathcal{L}^2}{\mu}, \tag{2.9}$$

and move very small distances of order

$$g\mathcal{T}^2 = \frac{\rho^2 g\mathcal{L}^4}{\mu^2}. \tag{2.10}$$

Subsequently, the inertial terms have negligible contribution to the flow and may be dropped from the equations, as is plain from a comparison of (2.5) and (2.9). Because of the smallness of both the time interval in which this happens and the fluid particle displacements, it is reasonable to ignore the inertial terms altogether, and assume instantaneous development of an initial velocity field before any significant movement of the fluid particles.

Thus the flow of very viscous fluids is given by the solution of the Stokes creeping-flow equations (2.7) subject to appropriate initial and boundary conditions, though

we shall need to check the validity of our assumptions in the context of specific flow problems. Deferring a consideration of the initial conditions till a little later, let us first discuss the specification of boundary conditions.

For all problems considered in this work, the boundaries of the flow domain are either in contact with a solid support or are free surfaces. Use will also be made of vertical and horizontal planes of symmetry and anti-symmetry to reduce the amount of computation.

At boundaries that are in contact with a solid support, it is assumed that there is no slip so that

$$u_i = u_i^* = 0. \tag{2.11}$$

At free-surface boundaries, surface stresses are given by [75, p. 14]

$$T_i = -pn_i + \mu n_j \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) = \gamma n_i \left(\frac{1}{R_1} + \frac{1}{R_2} \right) \tag{2.12}$$

where γ is the coefficient of surface tension, n_i is the outward unit normal to the surface, and R_1 and R_2 are the principal radii of curvature. Written in terms of non-dimensional quantities as before, equations (2.12) become

$$-p^* n_i + n_j \left(\frac{\partial u_i^*}{\partial x_j^*} + \frac{\partial u_j^*}{\partial x_i^*} \right) = \frac{\gamma}{\mu \mathcal{U}} n_i \left(\frac{1}{R_1^*} + \frac{1}{R_2^*} \right). \tag{2.13}$$

The dimensionless parameter

$$\mathrm{C}_a = \mu \mathcal{U} / \gamma = \rho g \mathcal{L}^2 / \gamma, \tag{2.14}$$

known as the capillary number, is large in problems of interest here, except in very localized regions of high curvature characterized by a small length scale $\mathcal{L}$. Therefore, at free-surface boundaries, the simplified no-stress free-surface conditions

$$-p^* n_i + n_j \left(\frac{\partial u_i^*}{\partial x_j^*} + \frac{\partial u_j^*}{\partial x_i^*} \right) = 0 \tag{2.15}$$

can be applied. Planes of symmetry and anti-symmetry with normal parallel to axis x_I also require that (2.15) be satisfied, and in addition we must have

$$u_I^* = 0 \tag{2.16}$$

on a plane of symmetry since there can be no flow normal to the plane, and

$$u_i^* = 0, \quad i \neq I \tag{2.17}$$

on a plane of anti-symmetry since there can be no flow tangential to the plane.

We now come to a consideration of the issue of initial conditions. The equations given so far, to describe the very viscous fluid flows with which we shall be concerned, do not involve the time variable. This was removed with the inertial terms in the Navier-Stokes equations, and has not featured since. Thus, initial velocity or pressure fields are not required in order to solve the creeping-flow equations; only the geometry of the flow domain needs to be specified. The solution velocities and pressures so obtained are for a particular instant in time when the geometry is known and specified. In particular, any flow that is apparently started from 'rest' appears to instantaneously establish a flow field appropriate to the initial geometry, as was discussed earlier.

To complete the mathematical formulation and obtain the evolution of geometry and flow in time one further equation must be included, involving the time variable t, and giving the change in the flow domain geometry that results from the flow field that develops. This could be the kinematic condition [75, p. 13] on all free surfaces $\mathcal{F}(x_i, t) = 0$, namely

$$\frac{D\mathcal{F}}{Dt} = \frac{\partial \mathcal{F}}{\partial t} + u_i \frac{\partial \mathcal{F}}{\partial x_i} = 0. \tag{2.18}$$

Alternatively, the Lagrangian equation

$$u_i = dx_i/dt \tag{2.19}$$

can be solved [104, p. 29] to determine the displacement of points x_i, including those that fall on free-surface boundaries. The latter method is preferred for the solution methods used in this thesis, excepting for a few occasions where the kinematic condition (2.18) is more suitable.

There is an interesting consequence deriving from this method of solving the creeping-flow model for the velocity field in a particular geometrical configuration, and then modifying the geometry to reflect the fluid flow. Because there is no need for initial conditions apart from the initial geometry, therefore we do not need to know any of the previous flow history in order to solve for the flow at any point in time; all that is required is the geometrical configuration at that point in time. How that geometrical configuration was arrived at is quite immaterial, so that whether the flow commenced with a specified geometry, or that geometry was attained as the result of some flow history, the future flow behaviour will be identical.

The method used for time advancing the flow geometry will be discussed further, after considering some methods for solving the Stokes equations.

2.2 Methods of Solution

There are a number of methods that might be used to solve Stokes creeping-flow problems, with their suitability depending on the precise nature of the flow problem in hand and the solution sought. Thus in selecting a solution method it is necessary to first consider the important properties of the flow problems to be solved, and the output desired. Already we have the requirement that we be able to handle free surfaces, and accurate tracking of moving free-surface boundaries is of considerable importance to some applications that we shall consider. In addition, for the flows considered in this thesis, it is necessary to be able to model flow in completely general flow domains, and allow possible interaction with a mould. As will be seen, although isothermal conditions have been assumed for the moment, we shall eventually wish to relax this condition and permit non-uniform fluid properties throughout the flow domain.

Of the numerical solution methods applied in the field of computational fluid dynamics (CFD), those most well-known are finite-element, finite-difference, finite-

volume, and boundary-element methods, all of which have been utilized in commercial and purpose-written CFD programs. Of these, the finite-element method has been selected as the most suitable for the type of problem we wish to solve, with reasons as follows.

A well-recognized major drawback to using finite-difference techniques is the requirement to use a structured rectangular mesh, making the boundaries of complex flow domains difficult to represent. Sometimes this problem can be overcome by mapping the physical domain to a rectangular computational domain, which does of course result in an additional level of computational complexity. By contrast the unstructured meshes permissible with finite-element methods can much more easily accommodate complex domain boundaries, and since one of the major aims of this work is to accurately determine the evolution of free-surface boundaries, it makes good sense to use finite elements rather than finite differences. Another known difficulty with finite difference methods is the accurate specification of boundary conditions, and the specification of zero-stress free-surface boundaries, as we require, is extremely difficult, especially by comparison with the finite-element method which handles them quite naturally, as we shall see.

At the commencement of this work, use of the finite-volume commercial package *Phoenics*, developed by CHAM [19], was investigated. Grids used with the finite-volume method must still be reasonably structured arrangements of quadrilateral elements, though variable element sizing does make things easier for non-rectangular domains. There is also a boundary-fitted coordinates facility to assist with mapping from physical to computational domains. However, as with finite differences, the natural boundary conditions of the finite-volume method are Dirichlet and Neumann conditions or a mixed form of these, and, again, specification of the no-stress boundary condition is difficult. Despite communication with the user support facility of CHAM in the United Kingdom, no way was found to specify this using *Phoenics*, and, while it is almost certainly possible to devise a means to specify this

boundary condition by finite-volume methods, this experience along with the desire to represent possibly complex and changing non-rectangular flow domains led to the definite conclusion that finite elements are more appropriate for problems involving free surfaces. Thus, the attempt to use finite-volume methods, and *Phoenics* in particular, was abandoned.

The boundary-element method causes no difficulties when it comes to representing complex domain boundaries, since it is these that are discretized rather than the whole flow domain. In addition this results in a reduction of the order of the problem to be solved. However it does assume that the fluid properties are spatially constant, so that our ultimate desire to permit non-uniform fluid properties (especially viscosity) renders this method inappropriate.

For some simple flow domains and simplifying assumptions, some useful solutions can be readily, and sometimes more efficiently, obtained by methods other than finite elements. As already intimated the boundary-element method could be used so long as we keep the isothermal assumption, though this has not been done because program development for boundary elements is as complex and time consuming as for finite elements, even if the result is computationally more efficient. However use has been made of series methods with collocation to solve for flow in simple rectangular domains, and asymptotic methods for slender and nearly plane geometries. These special solutions yield information that is not easy, or even possible, to obtain by finite elements, and also provide a means of verifying the finite-element solution.

Having settled on finite elements as the generally-preferred numerical solution technique for this work, the question arises as to whether use can or should be made of a commercial CFD finite-element package. At the commencement of this work, it was thought to be a good idea to make use of existing software if applicable, and it was for this reason that the finite volume code *Phoenics* was investigated. Having discarded this for reasons already given, attention was turned to the finite-element code *Fastflo* [22] written by the CSIRO Division of Mathematics and Statistics,

which was accessible during the course of this work. This proved to be much better for handling free surfaces, and was utilized in much early work involving no mould contact. However, updating a boundary segment from free surface to no slip, as is necessary for modelling a free surface contacting a mould, was found to be extremely difficult to achieve within *Fastflo*. Consequently the decision was made to write a special-purpose finite-element program, rather than spend further time trialling other available packages or trying to devise a method that would enable continued use of *Fastflo*.

In the remainder of this chapter, a detailed description of the finite-element method as it has been applied to Stokes creeping flow is given. This description relates not only to the purpose-written program developed in connection with the present work, but also to *Fastflo*, since very similar methods are employed in both programs. It follows, to a considerable extent, that given in the finite-element books by Becker, Carey and Oden, particularly the first volume [8]. Use of series and asymptotic methods for solving very viscous fluid flows will be described in the context of specific problems to which they have been applied.

2.3 The Variational Equations

To implement the finite-element method let us first write the equations in their weak form. This is obtained by multiplying through by an appropriate test function, and then integrating over the flow domain Ω. Dropping asterisks (*) on dimensionless variables, the continuity equation (2.2) becomes

$$\int_\Omega v \frac{\partial u_i}{\partial x_i}\, dV = 0, \tag{2.20}$$

and the Stokes equations (2.7) become

$$\int_\Omega w_i \frac{\partial p}{\partial x_i}\, dV - \int_\Omega w_i \frac{\partial}{\partial x_i}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) dV = -\delta_{2i} \int_\Omega w_i\, dV, \tag{2.21}$$

where w_i is a vector test function and v is a scalar test function.

Next we obtain the variational form suitable for finite-element solution by applying Green's theorem to (2.21) which becomes

$$\begin{aligned} &-\int_\Omega p\frac{\partial w_i}{\partial x_i}\,dV + \int_\Omega \frac{\partial u_i}{\partial x_j}\frac{\partial w_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\frac{\partial w_i}{\partial x_j}\,dV \\ &\quad -\int_{\partial\Omega} w_i\left[-pn_i + n_j\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)\right]dS \\ &\quad = -\delta_{2i}\int_\Omega w_i\,dV \end{aligned} \tag{2.22}$$

with $\partial\Omega$ being the boundary of the flow domain Ω. In this form there are no terms involving the pressure gradient, which have been replaced with terms involving the unknown pressure p, and terms involving second derivatives of velocities have been replaced with terms involving first derivatives. This process has also yielded a boundary-integral term which is readily recognized as containing the expression for the stress on a free surface. Because of its appearance in the variational equations, specification of the surface stress as a boundary condition, as done in (2.15), is known as a natural boundary condition, while the no-slip boundary conditions given in (2.11) are called essential boundary conditions. It is the occurrence of the expression for surface stress in the equations describing the problem, that makes the finite-element method so suitable for modelling the type of free-surface flows that are the subject of this thesis.

At this stage the test functions w_i and v are quite arbitrary, and we now restrict w_i to the space of functions that satisfy the homogeneous essential boundary conditions defined for u_i in (2.11).

Returning again to the boundary integral term in (2.22) we now have $w_i = 0$ on all no-slip boundary segments, while on all free surfaces there is zero surface stress, so that there is no contribution to the boundary integral from these segments of the domain boundary $\partial\Omega$. Similarly, because of the zero stress conditions, the boundary integral evaluates to zero on boundary segments that are planes of symmetry and anti-symmetry. Since these are the only boundary conditions applicable to the problems we shall consider, the boundary integral term evaluates to zero and may

be removed from the variational equation.

2.4 Penalty Formulations

The problem as formulated so far requires the solution of coupled velocity and pressure fields. A common device used with the finite-element method, for decoupling the velocity and pressure solutions, is to introduce an artificial compressibility by writing

$$\frac{\partial u_i}{\partial x_i} = -\frac{p}{\lambda} \tag{2.23}$$

where λ is a large penalty number [18, 34]. For sufficiently large λ this is almost the continuity equation for incompressible flow. Equation (2.23) can be used to eliminate pressure from the Stokes creeping-flow equations, so that our variational problem is, on removing the boundary integral, reduced to solving the equations

$$\int_\Omega \frac{\partial u_i}{\partial x_j}\frac{\partial w_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\frac{\partial w_i}{\partial x_j} + \lambda\frac{\partial u_j}{\partial x_j}\frac{\partial w_i}{\partial x_i}\, dV = -\delta_{2i}\int_\Omega w_i\, dV \tag{2.24}$$

for the velocity field. Once this is known, the pressure field, if required, can be obtained by solving (the weak form of) (2.23).

Because of the introduction of artificial compressibility, there should be, strictly speaking, additional terms on the left-hand side of (2.24), deriving from the momentum equation for compressible flow, all of which are integrals of products of $\partial u_i/\partial x_i$. However, from (2.23), these terms are small compared with the other terms, and hence are neglected.

Equation (2.24) is a penalty form of the equations for incompressible Stokes creeping flow, and has been successfully used in the finite-element modelling of forming processes [70]. However, while we would like to choose λ to be as large as possible, if it is chosen too large, an ill-conditioned matrix will result and the method will not work. Furthermore the maximum permissible size of the penalty number λ is dependent on the mesh size [18].

A variation to the penalty method, known as the augmented Lagrangian method, leads to an iterative scheme for solving for the velocity and pressure fields [105, pp. 357–360]. After modifying (2.23) to

$$p_{n+1} = p_n - \lambda \frac{\partial u_i}{\partial x_i} \tag{2.25}$$

where p_n is the nth pressure iterate, we substitute the expression for p_{n+1} for p in (2.22) leading to (2.24) with an additional term

$$\int_\Omega p_n \frac{\partial w_i}{\partial x_i}\, dV$$

on the right-hand side. We can then start with some initial guess p_0 for the pressure field and iteratively solve this equation for the unknown velocity field followed by the weak form of (2.25) to update the pressure field. When the pressure has converged (i.e. $p_{n+1} = p_n$) then the continuity equation for incompressible flow will be satisfied, so that the solution should be more accurate than obtained by the penalty method. This method has the advantage that λ does not need to be as large as for the penalty method. Rather it is chosen to speed convergence, and hence is called the convergence accelerator.

Both the penalty and augmented-Lagrangian methods were investigated while trialling the CFD package *Fastflo* which was not, at that time, capable of solving the fully-coupled equations. This is reported in Chapter 3, where we also compare with results given by the purpose-written program solving the fully-coupled Stokes creeping-flow problem of Section 2.3.

2.5 The Axisymmetric Equations

So far the equations we have derived are applicable to completely-general three-dimensional flow domains. Three-dimensional problems do however present considerable difficulties when it comes to generating both volume and surface meshes

for complex flow regions, and as a result attention has been restricted to two-dimensional and axisymmetric problems. While the tensor notation of the equations that we have employed is easily applied in two dimensions, the axisymmetric equations do differ slightly requiring the addition of a few extra terms, and these are given here for completeness.

Let u_r, u_θ, and u_z be the radial, angular, and vertical velocity components in cylindrical coordinates (r, z, θ). Then, for axisymmetry, $u_\theta = 0$, and the conversion from cartesian to cylindrical coordinates is given by

$$x_1 = r\cos\theta, \quad x_2 = z, \quad \text{and} \quad x_3 = r\sin\theta, \tag{2.26}$$

$$u_1 = u_r\cos\theta, \quad u_2 = u_z, \quad \text{and} \quad u_3 = u_r\sin\theta. \tag{2.27}$$

Using these, we may convert equations (2.20) and (2.22) to their axisymmetric equivalents. Changing notation, let us now define x_1 and x_2 to be the radial and vertical directions respectively, so that we can express these in tensor notation. Then, we have the axisymmetric variational forms of the continuity equation

$$\int_\Omega v\left(\frac{\partial u_i}{\partial x_i} + \frac{u_1}{x_1}\right) x_1\,dx_1\,dx_2 = 0 \tag{2.28}$$

and the Stokes equations

$$\begin{aligned} &-\int_\Omega p\left(\frac{\partial w_i}{\partial x_i} + \frac{\delta_{1i}w_1}{x_1}\right) x_1\,dx_1\,dx_2 \\ &\quad + \int_\Omega \left(\frac{\partial u_i}{\partial x_j}\frac{\partial w_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\frac{\partial w_i}{\partial x_j} + \frac{2\delta_{1i}u_1w_1}{x_1^2}\right) x_1\,dx_1\,dx_2 \\ &\quad - \int_{\partial\Omega} w_i\left[-pn_i + n_j\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)\right] x_1\,dS \qquad (2.29) \\ &\quad = -\delta_{2i}\int_\Omega w_i\,x_1\,dx_1\,dx_2, \qquad (2.30) \end{aligned}$$

where, again, the boundary integral can be dropped from (2.30) since, as for Cartesian coordinates, it has zero contribution for the types of problem of current interest.

The penalty system of axisymmetric equations is obtained by substituting

$$-p = \lambda\left(\frac{\partial u_j}{\partial x_j} + \frac{u_1}{x_1}\right) \tag{2.31}$$

in (2.30), and as before, the iterative augmented Lagrangian method is obtained from a modification of (2.31).

2.6 The Galerkin Finite Element Method

Having obtained the variational form of either the fully-coupled or penalty equations, we now wish to solve them using finite elements. Before we can proceed further, we need to discretize the flow domain by defining a mesh over it. A good discussion of the issues surrounding mesh generation is contained in [62, Ch. 7]. For two-dimensional and axisymmetric flows, quadrilateral or triangular elements are the usual choices. Historically, prior to the advent of isoparametric mapping, straight-sided triangles were very popular, being much more suitable for approximating irregular boundaries. More recently, with isoparametric formulations, quadrilateral elements have become the more common [32, 69]. However, as is well-known, for any particular problem, convergence of a finite-element method is not assured for all element types (see [62]), and while no rigorous testing has been conducted to determine the reason, experimentation has shown quadrilateral elements to be quite unsuitable for our problems, yielding obviously erroneous results. Structured triangular meshes obtained from quadrilateral meshes by dividing all elements across the same diagonal were also found to be unsuitable, displaying convergence problems as the number of elements was increased. Hence unstructured meshes of isoparametric triangular elements are used, which do yield good results, and in addition are automatically generated by appropriate input to *Fastflo*, a facility that is used extensively in this work.

The discretized problem is to solve the variational equations for the unknowns at the mesh nodes. For stability reasons it is usual to use a mixed method whereby the solution for pressure is obtained using elements of an order lower than those used for solving for velocity [9, 18, 75]. Therefore we mesh with quadratic triangular

elements defined by three corner and three mid-side nodes, noting that the lower order linear elements are defined by just the three corner nodes. Then, we solve for the velocity components at all six nodes, and for pressure at the three corner nodes.

Let N be the total number of nodes in our mesh and M be the total number of corner nodes where $N > M$. We then define the usual set of quadratic polynomial basis functions

$$\{\phi^k(x_i), k = 1, N\} \tag{2.32}$$

that are piecewise continuous between elements, with one per node such that $\phi^k = 1$ at node k and is zero at all other nodes. Each ϕ^k will be non-zero on only those mesh elements that include node k, and over all other elements $\phi^k = 0$. We also define another set of M linear basis functions using only corner nodes,

$$\{\psi^k(x_i), k = 1, M\} \tag{2.33}$$

that are also piecewise continuous between elements and only non-zero on elements containing node k, such that $\psi^k = 1$ at node k and is zero at all other corner nodes. Then we set

$$w_i = \sum_{k=1}^{N} a_i^k \phi^k(x_j), \qquad u_i = \sum_{k=1}^{N} \alpha_i^k \phi^k(x_j), \tag{2.34}$$

and

$$v = \sum_{k=1}^{M} b^k \psi^k(x_i), \qquad p = \sum_{k=1}^{M} \beta^k \psi^k(x_i), \tag{2.35}$$

where the a_i^k and b_k are constant coefficients that we can choose, while the α_i^k and β_i^k are constant coefficients for which we must solve so as to find the velocity and pressure fields throughout the flow domain. Note the use of summation notation in order to clearly show the number of terms involved, while tensor notation is still used in relation to problem dimension.

These expressions for w_i, u_i, v and p are now substituted into the variational equations. This gives for the continuity equation (2.20)

$$-\sum_{k=1}^{M} b^k \sum_{\ell=1}^{N} \int_\Omega \psi^k \frac{\partial \phi^\ell}{\partial x_i} \, dV \, \alpha_i^\ell = 0, \tag{2.36}$$

and for the Stokes equation (2.22)

$$-\sum_{k=1}^{N} a_i^k \sum_{\ell=1}^{M} \int_\Omega \frac{\partial \phi^k}{\partial x_i} \psi^\ell \, dV \, \beta^\ell + \sum_{k=1}^{N} a_i^k \sum_{\ell=1}^{N} \int_\Omega \frac{\partial \phi^k}{\partial x_j} \frac{\partial \phi^\ell}{\partial x_j} \alpha_i^\ell + \frac{\partial \phi^k}{\partial x_j} \frac{\partial \phi^\ell}{\partial x_i} \alpha_j^\ell \, dV$$
$$= -\delta_{2i} \sum_{k=1}^{N} a_i^k \int_\Omega \phi^k \, dV. \tag{2.37}$$

Ignoring for the moment the essential homogeneous boundary conditions, the coefficients a_i^k and b^k can be anything we like to choose, and if we choose successively

$$a_i^k = \delta_{1k}, \delta_{2k}, \ldots, \delta_{Nk},$$
$$b^k = \delta_{1k}, \delta_{2k}, \ldots, \delta_{Mk}, \tag{2.38}$$

for a problem of dimension D we get $M + DN$ equations for the M pressure and DN velocity unknowns, namely

$$-\sum_{\ell=1}^{N} \int_\Omega \psi^k \frac{\partial \phi^\ell}{\partial x_i} \, dV \, \alpha_i^\ell = 0 \qquad \text{for} \quad k = 1, 2, \ldots, M, \tag{2.39}$$

and

$$-\sum_{\ell=1}^{M} \int_\Omega \frac{\partial \phi^k}{\partial x_i} \psi^\ell \, dV \, \beta^\ell + \sum_{\ell=1}^{N} \int_\Omega \frac{\partial \phi^k}{\partial x_j} \frac{\partial \phi^\ell}{\partial x_j} \alpha_i^\ell + \frac{\partial \phi^k}{\partial x_j} \frac{\partial \phi^\ell}{\partial x_i} \alpha_j^\ell \, dV$$
$$= -\delta_{2i} \int_\Omega \phi^k \, dV \qquad \text{for} \quad k = 1, 2, \ldots, N. \tag{2.40}$$

If the penalty equations are used then we have only DN unknown velocity components obtainable by solving the DN equations

$$\sum_{\ell=1}^{N} \int_\Omega \frac{\partial \phi^k}{\partial x_j} \frac{\partial \phi^\ell}{\partial x_j} \alpha_i^\ell + \left(\frac{\partial \phi^k}{\partial x_j} \frac{\partial \phi^\ell}{\partial x_i} + \lambda \frac{\partial \phi^k}{\partial x_i} \frac{\partial \phi^\ell}{\partial x_j} \right) \alpha_j^\ell \, dV$$
$$= -\delta_{2i} \int_\Omega \phi^k \, dV \qquad \text{for} \quad k = 1, 2, \ldots, N. \tag{2.41}$$

The axisymmetric equivalents of (2.39), (2.40) and (2.41) are

$$-\sum_{\ell=1}^{N} \int_\Omega \psi^k \left(\frac{\partial \phi^\ell}{\partial x_i} \alpha_i^\ell + \frac{\phi^\ell}{x_1} \alpha_1^\ell \right) x_1 \, dx_1 \, dx_2$$
$$= 0 \qquad \text{for} \quad k = 1, 2, \ldots, M, \tag{2.42}$$

$$-\sum_{\ell=1}^{M}\int_\Omega\left(\frac{\partial\phi^k}{\partial x_i}+\frac{\phi^k}{x_1}\delta_{1i}\right)\psi^\ell\,x_1\,dx_1\,dx_2\,\beta^\ell$$
$$+\sum_{\ell=1}^{N}\int_\Omega\left(\left[\frac{\partial\phi^k}{\partial x_j}\frac{\partial\phi^\ell}{\partial x_j}+\frac{2\phi^k\phi^\ell}{x_1^2}\delta_{1i}\right]\alpha_i^\ell+\frac{\partial\phi^k}{\partial x_j}\frac{\partial\phi^\ell}{\partial x_i}\alpha_j^\ell\right)x_1\,dx_1\,dx_2$$
$$=-\delta_{2i}\int_\Omega\phi^k\,x_1\,dx_1\,dx_2\qquad\text{for}\quad k=1,2,\ldots,N, \tag{2.43}$$

and

$$\sum_{\ell=1}^{N}\int_\Omega\left(\left[\frac{\partial\phi^k}{\partial x_j}\frac{\partial\phi^\ell}{\partial x_j}+\frac{2\phi^k\phi^\ell}{x_1^2}\delta_{1i}\right]\alpha_i^\ell+\lambda\left[\frac{\partial\phi^k}{\partial x_i}\frac{\phi^\ell}{x_1}+\frac{\phi^k\phi^\ell}{x_1^2}\delta_{1i}\right]\alpha_1^\ell\right.$$
$$\left.+\left[\frac{\partial\phi^k}{\partial x_j}\frac{\partial\phi^\ell}{\partial x_i}+\lambda\left\{\frac{\partial\phi^k}{\partial x_i}\frac{\partial\phi^\ell}{\partial x_j}+\frac{\phi^k}{x_1}\frac{\partial\phi^\ell}{\partial x_j}\delta_{1i}\right\}\right]\alpha_j^\ell\right)x_1\,dx_1\,dx_2$$
$$=-\delta_{2i}\int_\Omega\phi^k\,x_1dx_1\,dx_2\qquad\text{for}\quad k=1,2,\ldots,N. \tag{2.44}$$

In order to satisfy the essential boundary conditions (2.11) we must have $\alpha_i^k=0$ for all nodes k that lie on a no-slip boundary, and in addition at these nodes we must also have $a_i^k=0$, so that our removal of the boundary integral from (2.22) remains valid. Similarly, for nodes k lying on a boundary that is a plane of symmetry with normal parallel to axis x_I, we must have $\alpha_I^k=a_I^k=0$, or if it is a plane of anti-symmetry we must have $\alpha_j^k=a_j^k=0$, $j\neq I$. This removes some of the equations and unknowns leaving a set of equations to be solved which can be represented in matrix form

$$K_{ij}c_j=f_i \tag{2.45}$$

where K_{ij} is the stiffness matrix, c_j is a vector of the unknown α_i^k and β^k, and f_i is the force vector containing the known right-hand side components of the equations.

Because the basis functions ϕ^k and ψ^k are non-zero only on elements which include node k, the stiffness matrix is sparse and may be evaluated mesh element by mesh element. Each non-zero component of the matrix is just the sum of the contributions from each of the mesh elements that have a contribution to make. Thus

$$K_{ij}=\int_\Omega g_{ij}(x_k)\,dV=\sum_{e=1}^{N_e}\int_{\Omega_e}g_{ij}(x_k)\,dV \tag{2.46}$$

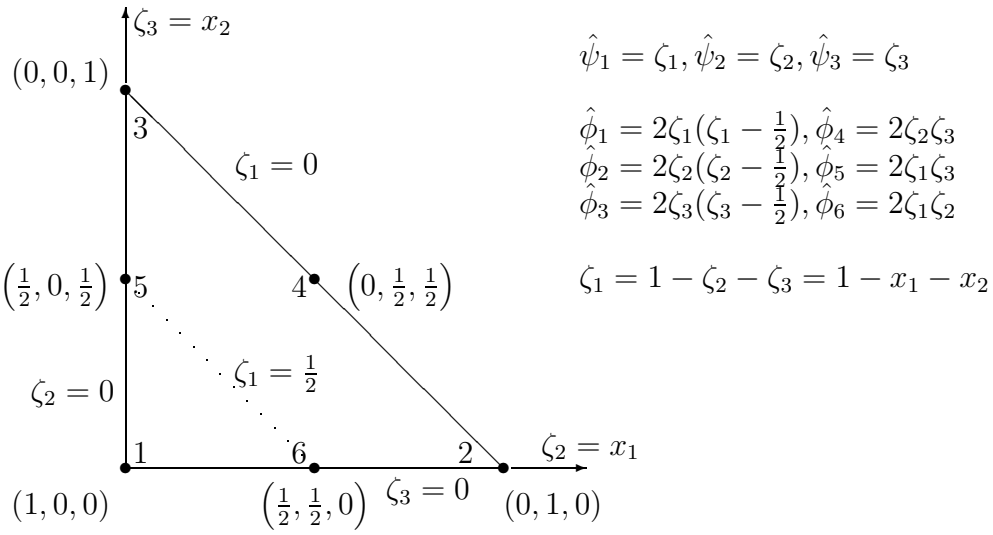

Figure 2.1: Master triangular element $\hat{\Omega}$ for finite-element computations (after Becker et al. [8]).

where there are N_e mesh elements, Ω_e is the region occupied by element e, and the g_{ij} are functions of position representing the integrands. For our problems the stiffness matrix is symmetric and, depending on the node-numbering scheme, may well be banded, though the bandwidth can be large. The force vector is also evaluated by computing the sum of the contributions from each mesh element.

Restricting our attention now to problems in two dimensions or that are axisymmetrical, we follow the usual method as described by Becker et al. [8, Ch. 5–6], and define the master triangular element $\hat{\Omega}$ shown in Figure 2.1. All nodes are labelled by their area coordinates $(\zeta_1, \zeta_2, \zeta_3)$ which are related to cartesian coordinates thus:

$$\zeta_1 = 1 - x_1 - x_2, \quad \zeta_2 = x_1, \quad \text{and} \quad \zeta_3 = x_2. \tag{2.47}$$

This master element has a set of quadratic basis functions which we denote $\hat{\phi}^j, j = 1, \ldots, 6$, such that $\hat{\phi}^j = 1$ at node j and is zero at all other nodes, and a set of linear basis functions which we denote $\hat{\psi}^j, j = 1, \ldots, 3$, such that $\hat{\psi}^j = 1$ at vertex j and is zero at the other two vertices. Now we can compute the components of the stiffness matrix and force vector for each mesh element, by mapping each element to the master element. For a mesh of six-node triangles with nodes $(x_1^j, x_2^j), j = 1, \ldots, 6$,

the transformation is given by

$$x_1 = \sum_{j=1}^{6} x_1^j \hat{\phi}^j, \qquad x_2 = \sum_{j=1}^{6} x_2^j \hat{\phi}^j, \tag{2.48}$$

and for a mesh of three-node triangles the transformation is similar, with the summation being over just the three vertices and the $\hat{\psi}^j$ replacing the $\hat{\phi}^j$. Then,

$$\begin{aligned} \int_{\Omega_e} g(x_1, x_2)\, dx_1\, dx_2 &= \int_{\hat{\Omega}} \hat{g}(\zeta_1, \zeta_2, \zeta_3)\, |J|\, d\zeta_2\, d\zeta_3 \\ &= \int_{\hat{\Omega}} G(\zeta_1, \zeta_2, \zeta_3)\, d\zeta_2\, d\zeta_3 \end{aligned} \tag{2.49}$$

where $|J|$ is the Jacobian determinant of the transformation from Cartesian coordinates (x_1, x_2) to area coordinates $(\zeta_1, \zeta_2, \zeta_3)$ (see Figure 2.1), and $G = \hat{g}|J|$.

To evaluate these integrals we use numerical quadrature rules that evaluate polynomial integrands exactly, and have the form [8]

$$\int_{\hat{\Omega}} G(\zeta_1, \zeta_2, \zeta_3)\, d\zeta_2\, d\zeta_3 = \sum_{j=1}^{n} W_j G(\zeta_{1j}, \zeta_{2j}, \zeta_{3j}). \tag{2.50}$$

Here n is the number of quadrature points, $(\zeta_{1j}, \zeta_{2j}, \zeta_{3j})$ are the area coordinates of quadrature point j, and W_j are the weights. Typical quadrature points and weights, for polynomial integrands of various degrees, are given in numerous finite-element texts, and some that we have used are reproduced (from [55, p. 114], [98, p. 115], and [105, p. 176]) in Table 2.1. The maximum order of polynomial that will be exactly evaluated is given in the table.

The finite-element and true solutions of the variational equations will, of course, differ, due to two sources of error: discretization error arises from representing a fluid domain with a discrete mesh, and from evaluating the integrals using quadrature rules for a discrete number of points; truncation or round-off error is incurred by the representation of numbers with a finite number of digits. In order to keep truncation error to a minimum, all our programming is done using double-precision arithmetic.

With respect to discretization error, let us first assume that all integrals are evaluated exactly and consider the error arising from representing a fluid domain

Table 2.1: Quadrature constants for triangles [55, 98, 105].

Order	*Error*	n	ζ_1	ζ_2	ζ_3	W_i
Linear	$O(\ell^2)$	1	1/3	1/3	1/3	1/2
Quadratic	$O(\ell^3)$	3	1/2	1/2	0	1/6
			0	1/2	1/2	1/6
			1/2	0	1/2	1/6
Cubic	$O(\ell^4)$	4	1/3	1/3	1/3	−27/96
			0.6	0.2	0.2	25/96
			0.2	0.6	0.2	25/96
			0.2	0.2	0.6	25/96

with a discrete mesh. Then, as shown by Zienkiewicz and Taylor [105, pp. 37, 228–230], the finite-element approximation locally, over an element, has error of order $O(\ell^{q+1})$ where ℓ is the mesh element size and q is the order of the polynomial basis functions used in the finite-element expansions. Thus, for the very viscous flow problems with which we are concerned, velocity u_i has error of order $O(\ell^3)$ since quadratic polynomials ϕ are used in the expansions (2.34), and pressure p has error of order $O(\ell^2)$ since linear polynomials ψ are used in the expansions (2.35). Summed over the whole flow domain this error is of order $O(\ell^q)$ [53, Ch. 4].

Discretization error due to evaluation of the integrals in the stiffness matrix and force vector using quadrature rules, varies depending on the integrand and the quadrature rule used. Since any quadrature rule will converge to the exact result as ℓ goes to zero provided that it has error of order $O(\ell)$ or less [105, p. 178], we can just use the simple one-point rule given in Table 2.1, giving a local error of order $O(\ell^2)$. However we can do better than this. For two-dimensional problems using straight-sided triangular elements Ω_e, the Jacobian determinant is twice the element area [55, p. 116] and all terms in the integrand are polynomials — linear for three-node elements and quadratic for six-node elements — so that evaluation using the appropriate order quadrature rule from Table 2.1 is exact and no error arises

from this quarter. Quadrature rules will not be exact, leading to discretization error, when using six-node triangles that are not straight sided so that the integrand is not necessarily a polynomial, and in axisymmetric problems in which one of the terms contains the radius in the denominator which contributes a non-polynomial part to the integrand even with straight-sided mesh elements. In these cases we can at least preserve the order of convergence of the finite-element method by using a quadrature rule that is exact for the degree of polynomial occurring in the integrand [105, p. 178], [53, p. 61]. In two dimensions the degree of the polynomial integrand is $2q - 2$ where q is the degree of the polynomial basis functions, while for an axisymmetric geometry the integrands are of degree $2q - 1$ on account of the multiplication by x_1 (see equations (2.28) and (2.30)). Thus in the two-dimensional quadratic case the three-point quadrature is suitable, and for axisymmetric cases the appropriate choices are the one-point and four-point rules for linear and quadratic mesh elements respectively. In much of the work in this thesis, the four-point rule is used for all integration because it is as good as, or better, than we require. Zienkiewicz and Taylor [105, p. 177] do note that it can be a disadvantage to use higher orders of integration than actually needed, but no problems have arisen on this account.

There is one further point to note in choosing quadrature rules for axisymmetric problems that is briefly discussed in [38]. Evaluation of the term $u_1 w_1 / x_1$ from (2.30), containing the radius in the denominator, will be a problem should any of the quadrature points fall on the axis of symmetry $x_1 = 0$. At these points we also have $u_1 = w_1 = 0$, but it is not obvious that the term will evaluate to zero. Choosing a quadrature rule using points in the interior of the element avoids this problem so that the one-point and four-point rules given in Table 2.1 are suitable while the three-point rule may not be. Graham [38] uses a seven-point quadrature with all quadrature points in the interior of the element, possibly that given in [55], but we cannot see that the extra computation is warranted since the four-point quadrature rule gives an error only as large as the mesh discretization error. For

interest we shall obtain some results using both quadratures with points lying on the element boundaries on the assumption that the difficult term is indeed zero on the symmetry axis, and quadratures with only interior points so that this assumption is not required. A comparison of these may indicate whether the term $u_1 w_1/x_1$ is indeed zero on the axis or not.

Having evaluated the stiffness matrix and force vector, we must solve the system of equations (2.45) for the unknowns c_j. By using a sparse or banded matrix solver we save on memory usage because we do not store all the zero entries in matrix K_{ij}. We can also exploit the symmetry of the matrix. Various library routines, such as NAG and LAPACK, are available. One that has been found to be particularly efficient is DGELB [91]. This routine is applicable to a general banded matrix, but has been slightly modified to make it more efficient for our particular problems. It solves directly by Gauss elimination.

2.7 Advancing in Time

Having solved for the velocity and pressure fields in our flow domain, we shall wish to determine how the fluid region evolves in time due to the motion of the fluid particles. For the unsteady problems, that we consider, a simple method adopted for this work is to solve (2.19) at each mesh node $j = 1, \ldots, N$ having coordinates $[x_i]_j^n$ at the nth time step, for the new node position $[x_i]_j^{n+1}$ after some time step Δt has elapsed, using the velocity $[u_i]_j^n = u_i([x_i]_j^n)$ that has been computed there by the finite-element method. When the mesh nodes have been updated by moving them to their new locations, we have a new flow region on which we can solve for velocity and pressure, and we can continue in this manner to determine the further evolution of the fluid region and the free surfaces in particular. This method is described in [104]. It has the advantage of incorporating automatic remeshing in the time-stepping process. Mesh elements can, however, become quite distorted in

time, when it is necessary to stop the process and define a completely new mesh over the flow region. In general, when this is done, it is necessary to interpolate the velocity and pressure fields onto the new mesh, but because the solution of the Stokes equations requires no information from a previous time step we do not have this complication. As noted in [104], this method also requires that time steps be chosen so that the displacement of any node is not too large, and so that boundary nodes do not penetrate fixed boundaries. We shall discuss these matters further in the context of specific flow problems.

To solve (2.19) we use, depending on the specific problem, either the simple Euler method as given in [104] leading to

$$[x_i]_j^{n+1} = [x_i]_j^n + [u_i]_j^n \Delta t, \tag{2.51}$$

or a fourth order Runge-Kutta method [52]

$$\begin{aligned}
[k_i]_j^1 &= u_i([x_i]_j^n)\,\Delta t,\\
[k_i]_j^2 &= u_i([x_i]_j^n + [k_i]_j^1/2)\,\Delta t,\\
[k_i]_j^3 &= u_i([x_i]_j^n + [k_i]_j^2/2)\,\Delta t,\\
[k_i]_j^4 &= u_i([x_i]_j^n + [k_i]_j^3)\,\Delta t,\\
[x_i]_j^{n+1} &= [x_i]_j^n + ([k_i]_j^1 + 2[k_i]_j^2 + 2[k_i]_j^3 + [k_i]_j^4)/6.
\end{aligned} \tag{2.52}$$

The Euler method is best for contact problems as we shall see, but otherwise the Runge-Kutta method is preferable because of its greater accuracy. The Runge-Kutta method is computationally more expensive than the Euler method, but this is offset by being able to take larger time steps for a comparable or better accuracy, as will be shown.

2.8 What Next?

The Stokes creeping-flow finite-element program that has been developed, as described in this chapter, is applicable to a wide range of very viscous fluid flows, and

throughout the remainder of this thesis we make use of it in modelling just a few of these. As we proceed we shall test and validate the code, introduce further details not yet described, and expand upon some others only briefly touched upon to this point in time.

In addition to this program, the finite-element CFD package *Fastflo*, given appropriate input, solves some creeping-flow problems by a finite-element method very similar to that described in this chapter. We use *Fastflo*, along with the purpose-written program, for the first of the flows that we consider, as well as for generation of meshes that form part of the input to the purpose-written program.

Chapter 3

A Sagging Viscous Bridge

3.1 Introduction

An initially rectangular two-dimensional 'bridge' of width $2w$ and thickness $2h$ spans the gap between two vertical no-slip walls as shown in Figure 3.1. We suppose that at time $t = 0$ the bridge becomes a very viscous fluid such that the assumptions underlying the creeping-flow model of Chapter 2 are justified, and then compute the sag of this bridge under the influence of gravity. This flow is the subject of a paper by Tuck, Stokes, and Schwartz [93].

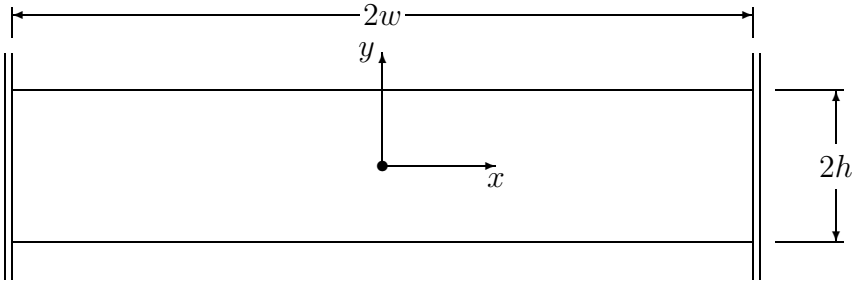

Figure 3.1: An initially rectangular viscous bridge before sagging commences.

Although a very simple idealised flow problem, it does have some 'real-world' applications. One of these is bending-beam viscometry [92] whereby the sag of a horizontal beam is used to measure viscosities of such things as molten glass [44],

butter [60], and ice-cream [80]. A cylindrical beam is common, but a beam of any regular cross section, including rectangular, can be used [44]. This problem is also relevant to disc-bending viscometry [42], being directly comparable to the sagging of an axisymmetrical initially flat disc clamped around its perimeter for which the results are qualitatively similar. We shall return to the subject of disc-bending viscometry in Chapter 5. A final application, that is mentioned here and described in more detail later, is glass slumping/sagging as in the manufacture of windshields and other vehicle components [36, 61], and aspheric optical surfaces [84]. The slumping of optical surfaces is the main focus of this thesis.

In this chapter we analyse this flow using finite-element methods, comparing the solutions obtained using the fully-coupled, penalty and augmented-Lagrangian methods described in Chapter 2. For penalty and augmented-Lagrangian solutions we use *Fastflo* which can readily solve using methods such as these that decouple the pressure and velocity calculations. However the version of *Fastflo* used (2.4, 11 August 1995) cannot simultaneously solve for pressure using linear mesh elements and velocity using quadratic elements, and for this we use the purpose-written finite element code described in Chapter 2. The *Fastflo* solutions provide an independent means of verifying this program code. In using the penalty and augmented-Lagrangian methods we need to consider the size of the penalty number/convergence accelerator, and how this is influenced by mesh size.

We first compute the initial motion of the bridge at time $t = 0^+$. Harking back to the discussion in Chapter 2 justifying the creeping-flow model, we can assume that the initial velocity profile is developed instantaneously, before there is any significant departure of the free surface from its initial rectangular configuration. At this early stage of the flow a semi-analytic solution may be obtained for comparison with the finite-element solutions, providing a check of the finite-element methods used. For the motion and free-surface shape at later times, we must depend upon the finite-element method, excepting that, for nearly plane geometries of small aspect ratio

h/w, asymptotic calculations do provide some useful information.

3.2 A Semi-Analytic Series Solution

Instead of using the primitive variables u, v, and p (or u_i and p) as we have described in Chapter 2, the creeping-flow model may be written in terms of a stream function ψ where the flow velocity is given by $(u, v) = (\psi_y, -\psi_x)$. This stream function automatically satisfies the continuity condition for incompressible flow, while the Stokes equations,

$$p_x = \mu\nabla^2\psi_y, \qquad p_y = -\rho g - \mu\nabla^2\psi_x, \tag{3.1}$$

where subscript variables denote differentiation with respect to those variables, can be reduced to the biharmonic equation

$$\nabla^4\psi = 0. \tag{3.2}$$

On the walls $x = \pm w$ the no-slip conditions become

$$\psi_x = \psi_y = 0, \tag{3.3}$$

and, at $t = 0^+$ when the geometry is still rectangular, the zero normal and tangential stress conditions become

$$-p - 2\mu\psi_{xy} = 0, \qquad \text{and} \tag{3.4}$$

$$\psi_{xx} - \psi_{yy} = 0 \tag{3.5}$$

respectively. The pressure p must be obtained by integrating equations (3.1).

We next write the stream function ψ as a series expansion

$$\begin{aligned}\frac{32h^2\mu}{\rho g}\psi = {} & \frac{1}{5}x^5 + 2x^3y^2 - 3xy^4 + 24h^2xy^2 \\ & + A_0w^4x + B_0w^2(\frac{1}{3}x^3 + xy^2) \\ & + w^5\Re\sum_{j=1}^{\infty} C_j\left[\cos k_j\ \cos(k_jy/h) + \frac{y}{h}\sin k_j\ \sin(k_jy/h)\right]\frac{\sinh(k_jx/h)}{\sinh(k_jw/h)}\end{aligned} \tag{3.6}$$

where A_0 and B_0 are dimensionless real constants, the $C_j = A_j + iB_j$ are dimensionless complex constants, and the eigenvalues k_j are solutions of $\sin 2k = 2k$. The expansion (3.6) satisfies the biharmonic equation, as well as the no-stress boundary conditions on the free surfaces $y = \pm h$. The solution is then obtained by solving for the real coefficients $\{A_j, B_j\}$, $j = 0, 2, \ldots$, such that the wall conditions (3.3) are also satisfied. Since this flow is symmetrical about $x = 0$, and ψ given by the expansion (3.6) satisfies this symmetry requirement, we need, in fact, only satisfy these conditions on one wall, say $x = w$. This is done by truncating the infinite series portion of (3.6) to N terms. We then choose $N + 1$ collocation points on the wall $x = w$ at which we satisfy (3.3), to give $2N + 2$ equations that can be solved for the $2N + 2$ unknowns $\{A_j, B_j\}$, $j = 0, 2, \ldots, N$.

The result should converge to the true solution as the number N of terms in the series increases, although as N becomes large we expect to have problems of numerical precision. Table 3.1 shows the vertical velocity $-v/\mathcal{V}$ at the top centre of a typical bridge of aspect ratio $h/w = 0.2$ for different values of N, computed using a program written by Tuck [95], that utilizes this method. The velocity scale

$$\mathcal{V} = \frac{\rho g w^4}{32\mu h^2} \tag{3.7}$$

is justified below. Although the results for $N = 150$ and $N = 300$ deviate from the general trend for some reason that is not understood, it would appear that we have three figure accuracy, i.e. $-v/\mathcal{V} = 1.60$, with just 50 terms, and four figure accuracy i.e. $-v/\mathcal{V} = 1.599$, with 350 terms. This can be compared with the results of finite-element calculations.

In the thin-layer limit the coefficients C_j in (3.6) tend to zero quite rapidly so that it can be simplified to

$$\frac{32h^2\mu}{\rho g}\psi = \frac{1}{5}x^5 + w^4 x - \frac{2}{3}w^2 x^3. \tag{3.8}$$

Table 3.1: Initial sagging velocity by the series-expansion method.

N	$-v/\mathcal{V}$	N	$-v/\mathcal{V}$
50	1.601304	400	1.599479
100	1.600067	450	1.599498
150	1.591763	500	1.599540
200	1.599636	600	1.599418
250	1.599706	700	1.599404
300	1.598656	800	1.599414
350	1.599447	900	1.599433

This represents a unidirectional (vertical) flow with a parabolic-squared velocity profile

$$v = -\psi_x = -\frac{\rho g w^4}{32\mu h^2}\left(1 - \frac{x^2}{w^2}\right)^2. \tag{3.9}$$

Setting $x = 0$ in (3.9) gives the thin-layer velocity limit at the centre of the bridge as $v = -\rho g w^4/(32\mu h^2)$, and suggests the velocity scale $\mathcal{V}$ that we use throughout this chapter.

This is a very ad hoc derivation of a thin-bridge result, and is a particular case of a more general result that we shall formally derive using an asymptotic expansion in Section 3.6.

3.3 Initial Motion by Finite Elements

Prior to any movement of the bridge the problem is symmetrical about $x = 0$ and anti-symmetrical about $y = 0$ so that the calculation of initial velocities and pressures may be restricted to one quarter of the total flow domain $0 \le x \le w$, $0 \le y \le h$. We non-dimensionalize choosing the bridge half-width as the characteristic length, i.e. $\mathcal{L} = w$, giving the computational region $0 \le x^* \le 1$, $0 \le y^* \le h/w$ where h/w is the aspect ratio. (As earlier, asterisks denote dimensionless variables.) With the characteristic velocity given by $\mathcal{U} = \rho g w^2/\mu$ (see equation (2.5)), the flow is ob-

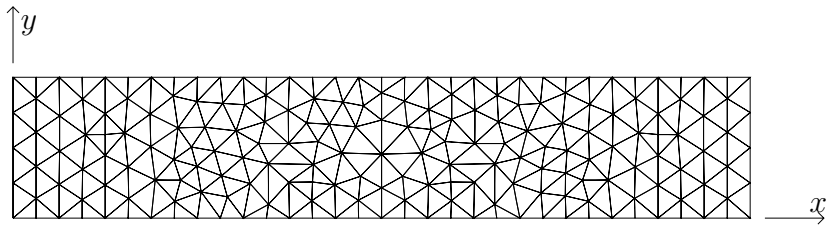

Figure 3.2: A typical uniform finite-element mesh with $N_e = 326$.

tained by solving (2.2) and (2.7) subject to the no-slip boundary conditions (2.11) on $x^* = 1$, the free surface conditions (2.15) on $y^* = h/w$, symmetry conditions (2.15) and (2.16) on $x^* = 0$, and anti-symmetry conditions (2.15) and (2.17) on $y^* = 0$.

For a bridge of some aspect ratio h/w, we define a mesh of N_e uniformly distributed six-node triangles over it. A typical mesh is shown in Figure 3.2. As the number of mesh elements is increased, we expect to see convergence to the true solution, for each of the solution methods used. To show this, we consider a bridge of aspect ratio $h/w = 0.2$ and look at (dimensionless) pressure and vertical velocity at the top centre of the bridge $(x^*, y^*) = (0, h/w)$. We re-scale the vertical velocity using the velocity scale $\mathcal{V} = \rho g w^4/(32\mu h^2)$ derived in the previous section. The vertical velocity should then be comparable with that obtained by the series solution of the previous section. Because the sample point is on the vertical centreline, there is no horizontal velocity component there. The pressure scale used is that defined in Chapter 2, i.e. $\mathcal{P} = \rho g w$.

We first solve this problem using *Fastflo* and the penalty method. In order to obtain a reasonable solution we must use a penalty number λ of suitable size. We require that it be sufficiently large to achieve reasonable mass conservation while we cannot allow it to become so large that the stiffness matrix becomes ill-conditioned. The size of λ is influenced by the size of mesh elements ℓ, and $\lambda = 10^2\ell^{-2}$ is given in the literature [56] as a suitable choice. This means that λ increases as ℓ decreases, and becomes very large as $\ell \to 0$. Now the condition number of the matrix is $\mathrm{O}(\lambda\ell^{-2})$ [56], and, clearly, for small ℓ the condition number is unacceptably large and the

Table 3.2: Initial sagging velocity by the classical penalty method with a uniform mesh.

N_e	λ	$-v/\mathcal{V}$	$p/\mathcal{P}$
158	10000	1.57533	1.23921
326	10000	1.59144	1.24828
518	10000	1.59127	1.24705
718	10000	1.59112	1.24781
918	8000	1.59036	1.24615
1095	8000	1.59631	1.25367
1254	8000	1.59585	1.25046
1454	8000	1.59622	1.25018

penalty method is unstable. Carey and Oden [18, p. 153] discuss the failure of the penalty method as λ becomes very large with reasonably fine meshes, observing that the discrete problem becomes overconstrained yielding a null solution as $\lambda \to \infty$. In order to obtain a non-zero solution the penalty number λ must be reduced.

This behaviour is completely borne out by our results, which are given for the top centre of the bridge in Table 3.2. Taking ℓ^2 as the average element area $(0.2/N_e)$, and $100 \leq N_e \leq 1500$, the literature would suggest a penalty number λ of order $\mathrm{O}(10^5)$ to $\mathrm{O}(10^6)$. However, in solving this problem with *Fastflo*, penalty numbers of these orders are found to be too large, as indicated by clearly erroneous results or by instability with small changes to the penalty number. As shown in Table 3.2 penalty numbers of $\lambda = 10^4$, and even less for very fine meshes, have been used so as to obtain a non-zero and stable solution.

The reduction of the penalty number with mesh element size, while ensuring a solution, results in the continuity condition being less well satisfied so that the convergence seen in Table 3.2 is, in fact, to a flow with small compressibility. Hence we must question how closely this approximates the true incompressible flow solution. Carey and Oden [18, p. 153] refer to experiments that suggest that for rather fine meshes the penalty number has to be chosen so small that the continuity constraint is not satisfactorily approximated. They also discuss reduced integration as a method

that can sometimes give a satisfactory solution with large penalty numbers, but we have not gone down this path. Rather we consider the modified penalty method of Lin et al. [56] which uses solutions obtained with two different penalty numbers of lower order than needed for the classical penalty method. For similar accuracy to the classical penalty method with $\lambda = \mathrm{O}(10^2\ell^{-2})$, penalties of $\lambda_1 = \mathrm{O}(10\ell^{-1})$ and $\lambda_2 = \mathrm{O}(10^2\ell^{-1})$ are used, and higher values will yield greater accuracy. Results from the two solutions are combined to give a more accurate solution using

$$Q_{12} = Q_{\lambda_2} \frac{Q_{\lambda_1} - Q_{\lambda_2}}{\lambda_1 - \lambda_2} \tag{3.10}$$

where Q is replaced by u, v, and p to give each of the velocity components and the pressure. With this method convergence to a result that is closer to the true incompressible solution can be expected, though, even so, there is some error due to small compressibility. Eventually, the penalty numbers still have to be reduced below desirable magnitudes so as to keep the stiffness matrix well conditioned, but this will be required for finer meshes than with the classical penalty method.

The results of the modified penalty method are given in Table 3.3. In addition to the penalties determined from mesh element size as discussed, results are included for $\lambda_1 = \lambda/10$ and $\lambda_2 = \lambda$ where λ is the penalty number used previously with the classical method (see Table 3.2).

By comparing the results of the classical and modified penalty methods in Tables 3.2 and 3.3 respectively, we see that a mesh of around 1500 elements is about as fine as we can get, since for finer meshes ($N_e > 1500$) the penalty numbers required for satisfactory accuracy become too large even for the modified penalty method. In fact, the apparent loss of accuracy for $N_e = 1454$ in the first set of results in Table 3.3 may well be caused by a penalty number that is too large. Apart from this solution, which we ignore, the modified penalty method gives results in which we can have reasonable confidence, and which show more satisfactory convergence (to around four figures) than the classical penalty method. They do however indicate

Table 3.3: Initial sagging velocity by the modified penalty method with a uniform mesh.

N_e	λ_1	λ_2	$-v/\mathcal{V}$	$p/\mathcal{P}$
158	300	3000	1.58363	1.23773
326	400	4000	1.59503	1.24533
518	500	5000	1.59389	1.24488
718	600	6000	1.59457	1.24650
918	700	7000	1.59476	1.24689
1095	750	7500	1.59772	1.25028
1254	800	8000	1.59772	1.24914
1454	850	8500	1.59622	1.24750
158	1000	10000	1.57812	1.23919
326	1000	10000	1.59279	1.24716
518	1000	10000	1.59237	1.24601
718	1000	10000	1.59338	1.24737
918	800	8000	1.59757	1.25053
1095	800	8000	1.59757	1.25043
1254	800	8000	1.59772	1.24914
1454	800	8000	1.59773	1.24865

that the classical penalty method is giving around three figures of accuracy with a mesh of more than 1000 elements, despite the reduced size of the penalty number.

A feature of this flow that compounds convergence difficulties is found by a corner flow analysis as conducted in [93]. To do this we again consider the flow in terms of a stream function ψ as in Section 3.2. Changing to local polar coordinates (r, θ) centred at the intersection of the lower free surface and the left wall $(x, y) = (-w, -h)$, a biharmonic function satisfying the no-stress conditions on the free surface $\theta = 0$ is

$$\psi = r^{\lambda} F(\theta) \tag{3.11}$$

where $\lambda \approx 1.5946$ is the solution of $\sin(\lambda\pi/2) = (\lambda - 1)$, and F is a function of θ [93]. It is easy to show, by change of variables from (x, y) to (r, θ), that both $u = \psi_y$ and $v = -\psi_x$ go to zero as $r \to 0$ like $r^{(\lambda-1)}$, so that the wall boundary conditions are also satisfied. Let us however consider the behaviour of the pressure

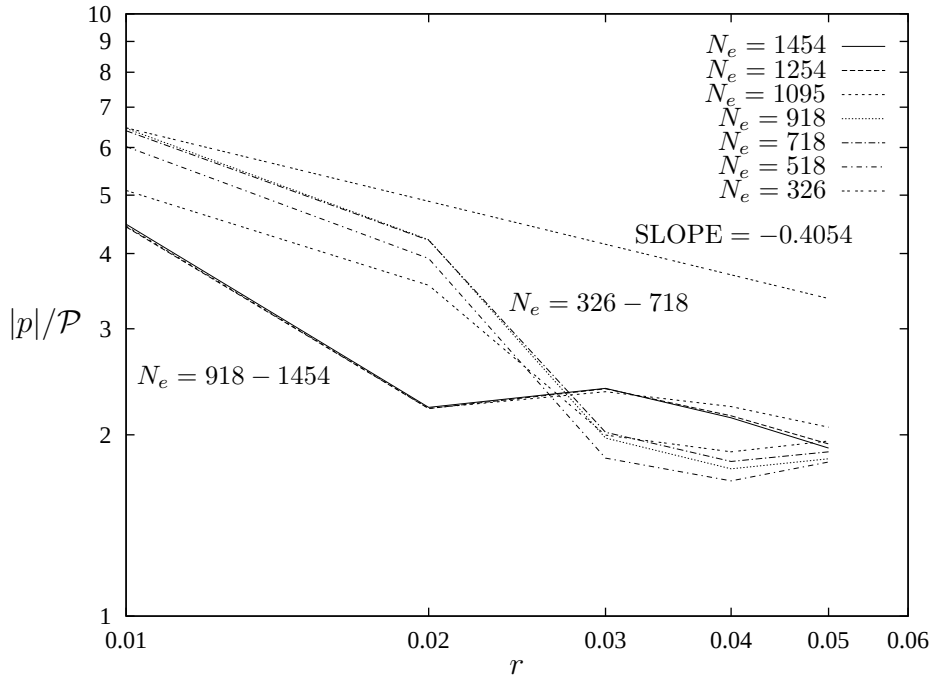

Figure 3.3: Pressure behaviour near a corner singularity for a uniform mesh.

p at this corner. On the free surface $\theta = 0$, pressure is given by (3.4), into which we substitute (3.11). Defining local Cartesian coordinates $\xi = w - x = r\cos\theta$, and $\eta = y + h = r\sin\theta$, we write

$$\begin{aligned} p &= -2\psi_{xy} \\ &= 2\psi_{\xi\eta} \\ &= 2\left(\frac{\psi_{r\theta}}{r} - \frac{\psi_\theta}{r^2}\right) \\ &= 2(\lambda-1)r^{(\lambda-2)}F'(0) \\ &\approx 1.1892\, r^{-0.4054}F'(0) \end{aligned} \tag{3.12}$$

with primes denoting differentiation with respect to θ. Thus, with $p \sim r^{(\lambda-2)} \approx r^{-0.4054}$, there is a mild pressure singularity at the intersection of free surface and wall, and p increases rapidly as $r \to 0$. This behaviour can be expected to slow the convergence of any numerical method, and suggests that the mesh should be constructed so as to concentrate elements about these points.

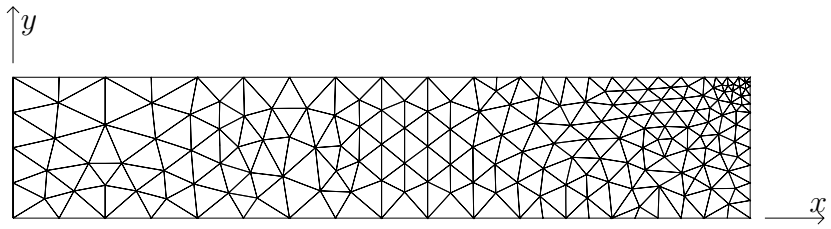

Figure 3.4: A typical finite-element mesh with elements concentrated at the intersection of free surface and wall, and $N_e = 337$.

Writing equation (3.12) as

$$\log|p| = (\lambda - 2)\log r + \log\left[2(\lambda - 1)|F'(0)|\right] \tag{3.13}$$

we see that on the free surface near these corner singularities, $\log|p|$ is a linear function of $\log r$ with slope $(\lambda - 2) \approx -0.4054$. In Figure 3.3 we plot $\log|p|$ versus $\log r$ on the free surface $\theta = 0$ for $0.01 \leq r \leq 0.05$ (i.e. $0.99 \geq x^* \geq 0.95$) using results obtained with the classical penalty method and a uniform mesh. It is clear that for the number of mesh elements used, we are not capturing the behaviour of p at the intersection of wall and free surface particularly well. In fact smaller numbers of mesh elements appear to do better than larger numbers, with $N_e = 326$ giving a slope that best matches our expectation.

In order to better capture the behaviour in the corners we try meshes with elements concentrated around the intersection of wall and free surface. These are generated using the automatic mesh generator within *Fastflo* to which is supplied information specifying the point(s) at which the mesh is to be concentrated. A typical mesh for one quarter of the domain for this problem, with elements concentrated at the corner singularity, is given in Figure 3.4. The vertical velocity and pressure at the top, centre of the bridge, computed using this type of mesh and the classical penalty method, is given for different numbers of mesh elements in Table 3.4.

These results show much faster convergence, to nearly four significant figures on refining the mesh to about 700 elements, than given by a uniform mesh with both the same classical penalty method (see Table 3.2) and the modified penalty method

Table 3.4: Initial sagging velocity by the penalty method with a non-uniform mesh.

N_e	λ	$-v/\mathcal{V}$	$p/\mathcal{P}$
154	10000	1.59341	1.25004
337	10000	1.59771	1.25190
511	10000	1.59722	1.25152
707	10000	1.59839	1.25175
892	8000	1.59896	1.25249
1086	8000	1.59831	1.25034

(see Table 3.3). They are also in good agreement with the results obtained with a uniform mesh having around 1000 elements and the modified penalty method, probably giving increased accuracy in the fourth figure. We do not consider non-uniform meshes of more than about 1000 elements because the small size of the elements around the region of concentration demands that we relax the penalty number even more than for a uniform mesh, giving results that are too suspect.

In Figure 3.5 we plot the relationship between $\log|p|$ and $\log r$ on the free surface near the corner singularity, obtained with the classical penalty method as before, but this time using a concentrated mesh. Now the behaviour of p at the intersection of wall and free surface is much closer to that predicted by our corner flow analysis, with the slope being quite accurate for $N_e = 892$ and 1086.

From this it is clear that a non-uniform mesh with finer resolution about the corner singularity is highly desirable for computational efficiency. In theory, a uniform mesh which matches the resolution of the non-uniform mesh about the corner singularity would also be satisfactory, but the penalty method precludes this because of the large penalty numbers necessary and the consequent ill-conditioning of the stiffness matrix. Thus, in fact, concentrating mesh elements at the corner singularity is essential for this solution method.

We could further improve our solution by using the modified penalty method with a non-uniform mesh, but, rather, we prefer to try a different method that

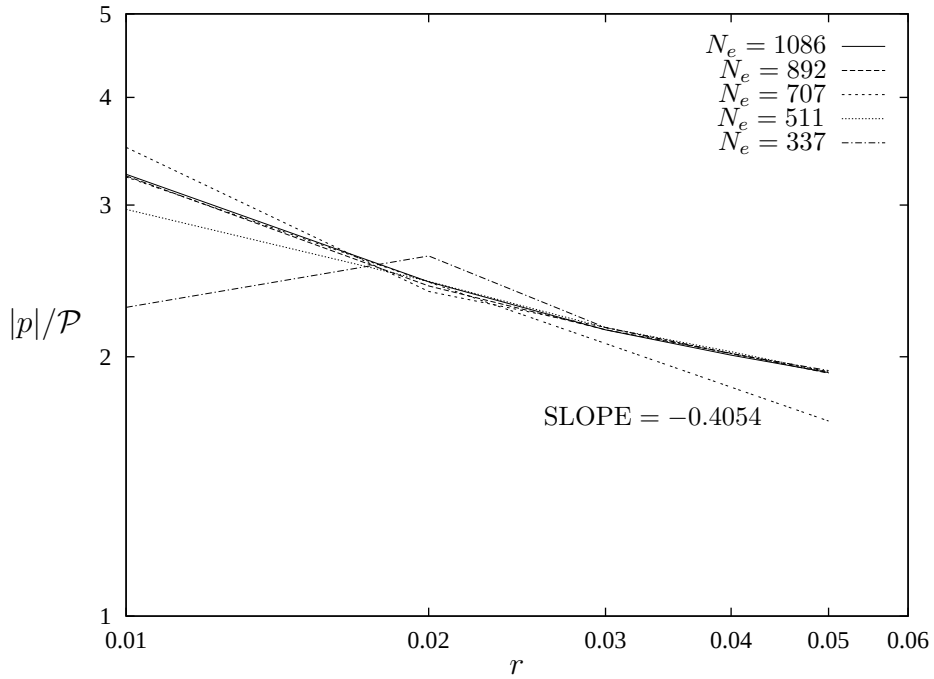

Figure 3.5: Pressure behaviour at a corner singularity with mesh elements concentrated at that corner.

overcomes the need to use large penalty numbers with fine meshes. The augmented Lagrangian method iteratively ensures that the continuity condition is satisfied so that the penalty number acts as a convergence accelerator and can be greatly reduced compared to the penalty method. Thus the augmented Lagrangian method will potentially give better accuracy than penalty methods, and, in addition, should permit finer mesh resolution. Results from using this method for our test problem with non-uniform meshes are given in Table 3.5. Although a larger penalty number will result in faster convergence, a small penalty number has been deliberately chosen so as to avoid ill-conditioning problems for fine meshes.

The results given in Table 3.5 show convergence of velocity to four figures for a mesh of 892 elements, and also show the classical penalty method with a non-uniform mesh to be just a little less accurate for similar numbers of mesh elements and considerably less computation time (cf. Table 3.4). As anticipated, the augmented Lagrangian method does allow us to use much finer meshes with more elements,

Table 3.5: Initial sagging velocity by the augmented Lagrangian method with a non-uniform mesh.

N_e	λ	n	$-v/\mathcal{V}$	$p/\mathcal{P}$	Error
511	100	7	1.59795	1.25138	1.0E-7
707	100	7	1.59835	1.25149	3.2E-7
892	100	8	1.59876	1.25231	3.1E-7
1086	100	8	1.59855	1.25099	9.2E-8
1278	100	7	1.59882	1.25198	1.9E-7
1472	100	7	1.59885	1.25213	6.0E-7
1657	100	7	1.59881	1.25187	3.6E-7
1837	100	7	1.59896	1.25192	7.3E-7

and we can more easily see convergence, and thus have greater confidence in our solution, as a result.

Because the augmented Lagrangian method permits very much finer meshes than the penalty method, it is theoretically possible to achieve the same accuracy with a sufficiently fine uniform mesh matching the resolution of the non-uniform mesh at the singularity corner. However convergence is very slow, and we need many thousands of elements as can be appreciated from a comparison of Figures 3.2 and 3.5. These give some idea of the intense clustering of mesh elements about the corner singularity in a non-uniform grid, and the resolution in the immediate proximity of the corner is extremely fine with the larger numbers of mesh elements that we are using (> 700). The very large computational cost and computer memory requirements associated with the numbers of mesh elements in a corresponding uniform distribution ($\approx$ 7000), make these meshes very inefficient and undesirable, and we do not consider them any further.

By solving the fully-coupled equations for velocity and pressure, we completely do away with the need for penalty numbers. For this a special-purpose finite-element code has been developed, as described in Chapter 2. The results of this method with a non-uniform mesh are given in Table 3.6 and show good convergence that is a little

more rapid than the augmented Lagrangian method.

Table 3.6: Initial sagging velocity from the fully-coupled equations with a non-uniform mesh.

N_e	$-v/\mathcal{V}$	$p/\mathcal{P}$
154	1.597754	1.248916
337	1.598345	1.251754
511	1.598661	1.251161
707	1.598712	1.251163
892	1.598930	1.252068
1086	1.598999	1.251989
1278	1.599011	1.251915
1472	1.599002	1.251776

Allowing for the varying accuracy of the different methods and meshes used, all of the finite-element methods that we have investigated, and both *Fastflo* and the purpose-written program, give a very consistent solution for both vertical velocity and pressure at the top centre of the bridge. With a mesh of uniformly sized and distributed triangles, and using penalty methods with *Fastflo*, three figure accuracy — $v/\mathcal{V} = -1.60$, $p/\mathcal{P} = 1.25$ — was achieved, while with a mesh that clusters elements about the intersection of wall and free surface, and using the more accurate augmented-Lagrangian and fully-coupled methods, with *Fastflo* and the purpose-written program respectively, four figure accuracy — $v/\mathcal{V} = -1.599$, $p/\mathcal{P} = 1.252$ — was achieved. With this consistency alone we have reasonably good verification of the program code that had been developed. But, in addition, we have identical three and four figure results from the series-expansion method of Section 3.2 (cf. Table 3.1), which is very satisfying.

The series-expansion method gives good results with little programming effort compared with finite-element methods, and, hence, is perhaps a preferable technique for computing the initial flow in a rectangular domain. It is also relatively fast in terms of its computer-time expenditure, although, with the computing resources

available today, this is less of a consideration. However we wish to compute the slumping of the bridge at later times when it is no longer rectangular. In addition, in later chapters, we shall consider flow domains that are not necessarily initially rectangular, and compute how the domain shape changes in time. For these, the series-expansion method is no longer convenient, and a finite-element method is preferable.

Of the finite-element methods considered, it finally seems best to use the purpose-written program and solve the fully-coupled equations because this gives the best accuracy without the uncertainties introduced by the penalty number. Accuracy will be of considerable importance to us when we come to model the slumping of optical surfaces. However the classical penalty method does give a reasonably good solution for considerably less computer time, and is therefore an attractive solution method if accuracy is not so important. For this reason it was used in [93] where only graphical accuracy was needed. The modified penalty method greatly improves on the accuracy of the classical penalty method, but it carries the overhead of having to solve the equations twice with two different penalty numbers, and then combine these. Also, because of the iteration necessary, the augmented Lagrangian method is at least as expensive in computer time, if not more so, as solving the fully-coupled equations. Thus there is no benefit in using either the modified-penalty or augmented-Lagrangian methods, and we have chosen to solve the fully-coupled equations from here on.

3.4 Further Initial Velocity Results

Now that we have confidence in the purpose-written finite-element program, and before we consider the slumping of a very viscous bridge over time, we compute the relationship between the aspect ratio h/w and the initial sagging velocity at the top centre of the bridge and at the centroid where the velocity is a maximum.

The results presented here were obtained using the finite-element program, but the semi-analytic solution gives identical results and was in fact used in similar figures in [93].

In Figure 3.6 we plot $-v/\mathcal{V}$ versus aspect ratio h/w at the bridge centroid and at the centre of the top free surface. As previously, the finite-element results have been re-scaled using the thin-layer centre-velocity limit $\mathcal{V}$ of (3.7), given by the semi-analytic method. Velocities were computed for $h/w \geq 0.05$ at increments of 0.025 using meshes of between 337 and 1482 elements, and a spline curve drawn through all points $(h/w, -v/V)$ including the point $(0, 1)$ which we anticipate from the semi-analytic thin-layer solution (3.9). The figure shows that $-v/\mathcal{V} \to 1$ as $h/w \to 0$, confirming $\mathcal{V}$ as the thin-bridge centre-velocity limit. If instead we use the velocity scale

$$\overline{\mathcal{V}} = \frac{\rho g w^2}{2\mu} \tag{3.14}$$

then we have $-v/\overline{\mathcal{V}} \to 1$ as $h/w \to \infty$ as shown in Figure 3.7, showing that $\overline{\mathcal{V}}$ is the thick-bridge centre-velocity limit.

We earlier obtained, from the semi-analytic solution, the thin-limit velocity profile across the width of the bridge (3.9) which has a parabolic-squared shape. In Figure 3.8 we show how the initial velocity $v/\mathcal{V}$ varies across the top free surface of the bridge, for a bridge of aspect ratio $h/w = 0.2$, which also indicates the initial surface deformation. The curve is somewhat bell-shaped, with significantly greater velocities near the centre of the bridge than near the walls (more so than would be the case for a parabolic profile, but less than for a parabolic-squared profile), and with a final rapid decrease toward zero occurring in the last five percent of the distance from the wall. This last effect can be explained using the corner flow analysis discussed earlier (p. 39), which gives that the free-surface vertical velocity v tends to zero like $(1 - |x|/w)^{0.5949}$ as $x \to \pm w$. On taking the x-derivative of this expression we have that the slope of the free-surface vertical velocity becomes infinite as $x \to \pm w$ as we see in Figure 3.8.

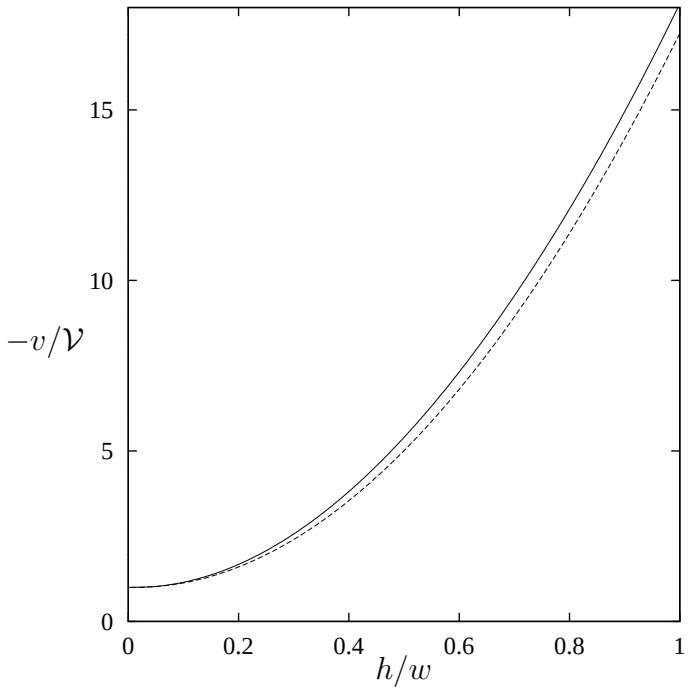

Figure 3.6: Initial sagging velocity $-v$ scaled with respect to the thin-bridge limit $\mathcal{V} = \rho g w^4/(32\mu h^2)$, as a function of aspect ratio h/w. Solid curve is the velocity at the bridge centroid, dashed curve is the velocity at the centre of the top free surface.

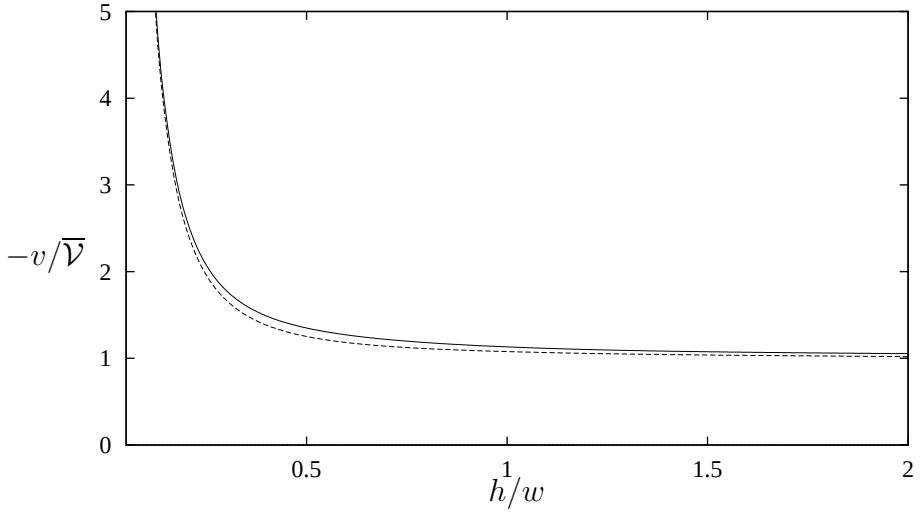

Figure 3.7: Same as Figure 3.6, but with velocities scaled with respect to the thick-bridge limit $\overline{\mathcal{V}} = \rho g w^2/(2\mu)$.

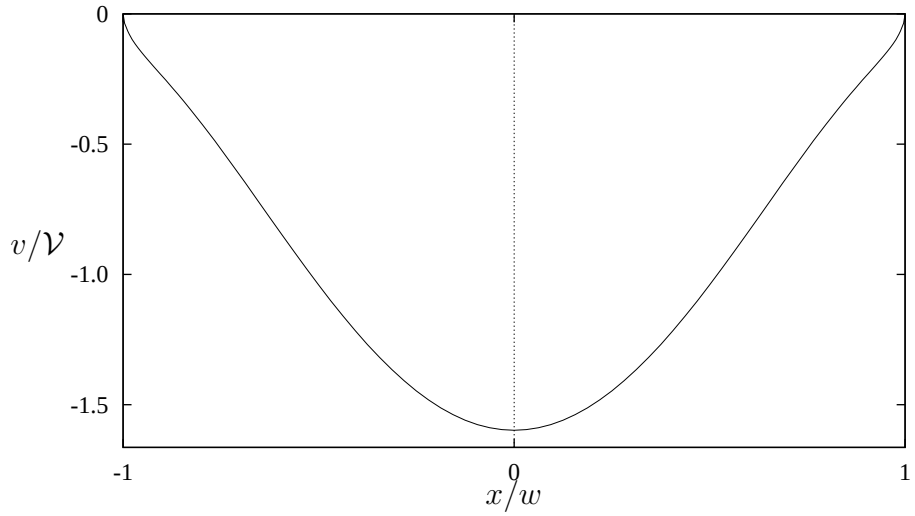

Figure 3.8: Variation across the bridge of the initial free-surface sagging velocity (also proportional to the initial free surface deformation) at $h/w = 0.2$.

In the thick-bridge limit the bridge will fall as if it had no top or bottom free surface, driven solely by the hydrostatic pressure. Hence the flow will approach the parabolic Poiseuille solution for flow down a channel with no dynamic-pressure gradient:

$$v = -\frac{\rho g w^2}{2\mu}(1 - x^2/w^2). \tag{3.15}$$

The thick-limit centre sagging velocity $\overline{\mathcal{V}} = \rho g w^2/(2\mu)$ confirms the numerical result obtained earlier and shown in Figure 3.7. Thus, as the aspect ratio varies over the range $0 \leq h/w < \infty$, the sagging profile across the bridge changes smoothly from parabolic-squared to parabolic.

3.5 Sagging Over Time

Now we move on to consider how the sagging bridge evolves over time. As explained in Section 2.7, the Stokes and continuity equations are first solved for the velocity field at an instant in time by the finite-element method, and this result is then used to solve (2.19) for the node displacements that result over a small time interval.

Table 3.7: Bridge-centroid sag at $t\mathcal{V}/(2h) = 1.953125$.

Δt^*	Euler		Runge-Kutta	
	Sag δ/w	Area	Sag δ/w	Area
0.2	0.806835	0.432237	0.756980	0.400044
0.1	0.781072	0.414370	0.756973	0.400003
0.05	0.768849	0.406762	0.756974	0.400000
0.02	0.761687	0.402605		
0.01	0.759325	0.401286		
0.001	0.757209	0.400127		

After the nodes have been moved to their new locations, this procedure is repeated. We solve the fully-coupled equations using the purpose-written finite-element program, and compare both the Euler and Runge-Kutta methods of solving for the displacements.

Once the bridge moves the anti-symmetry in $y = 0$ that exists initially is lost, and hence we must compute over one half of the flow domain $0 \leq x^* \leq 1$, $-h/w \leq y^* \leq h/w$. The computations are for a bridge of aspect ratio $h/w = 0.2$, and we use a non-uniform mesh of 526 triangular elements which gives the initial velocity and pressure at the top centre of the bridge, to nearly four figures (-1.598287 and 1.251907 respectively, cf. Table 3.6).

We continue to re-scale the velocity output of our computations using the thin-bridge velocity limit $\mathcal{V}$ of equation (3.7). In addition, rather than using the natural time scale of the finite-element method, $w/\mathcal{U}$ where $\mathcal{U} = \rho g w^2/\mu$ (see p. 8), it seems appropriate to introduce a new time scale $2h/\mathcal{V}$, which is the time in which a bridge sagging at a constant velocity of $\mathcal{V}$, will sag a distance equal to its initial thickness $2h$. We shall use this new time scale throughout the remainder of this section.

Table 3.7 shows the sag δ/w of the point initially located at the bridge centroid, $(x, y) = (0, 0)$, at $t\mathcal{V}/(2h) = 1.953125$ ($\equiv t\mathcal{U}/w = 1$) for various values of time step Δt^*, using the Euler and Runge-Kutta methods. (Note that the dimensionless time

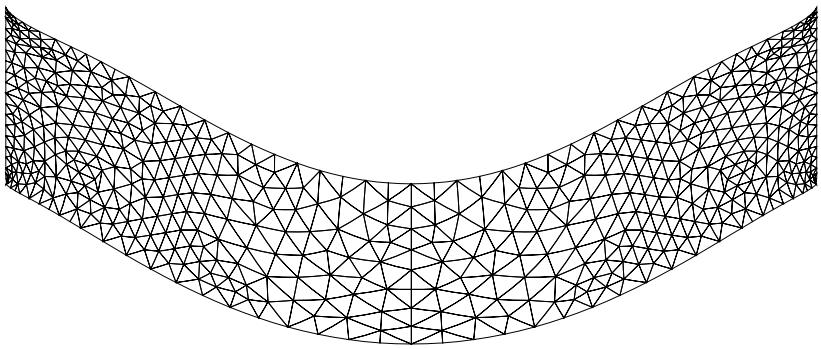

Figure 3.9: Bridge with $h/w = 0.2$ at time $t\mathcal{V}/(2h) = 0.78125$, when it has sagged about one thickness.

step Δt^* is the value actually used in the finite-element program and has not been re-scaled using the new time scale.) In addition, the computed half-area of the bridge is shown, which gives a measure of conservation of mass and should be equal to 0.4 for the bridge under consideration. Mass conservation has been found to be a good indication of accuracy for problems involving a moving mesh. Clearly the Runge-Kutta method is better than the Euler method giving much more accurate results for larger time steps. A very much smaller time step must be used at considerable cost in computational time with the Euler method, to match the results given by the Runge-Kutta method. Even with a time step of 0.001, the Euler result is not as good as that obtained with the Runge-Kutta method and a time step of 0.2. Hence the Runge-Kutta method is here used for time advancing the mesh, choosing a time step of $\Delta t^* = 0.05$.

Figure 3.9 shows the bridge (and the evolving triangular grid with corner concentrations) when it has sagged a distance approximately equal to its initial thickness, at which point the time is given by $t\mathcal{V}/(2h) = 0.78125$. This figure confirms, especially on the upper free surface, the behaviour suggested by the initial velocity profile of Figure 3.8, with a near quadratic-squared shape in the middle parts of the bridge, coupled with a rapid return to the original attachment point near the corner. While probably unimportant in some applications, the multiple changes in curvature

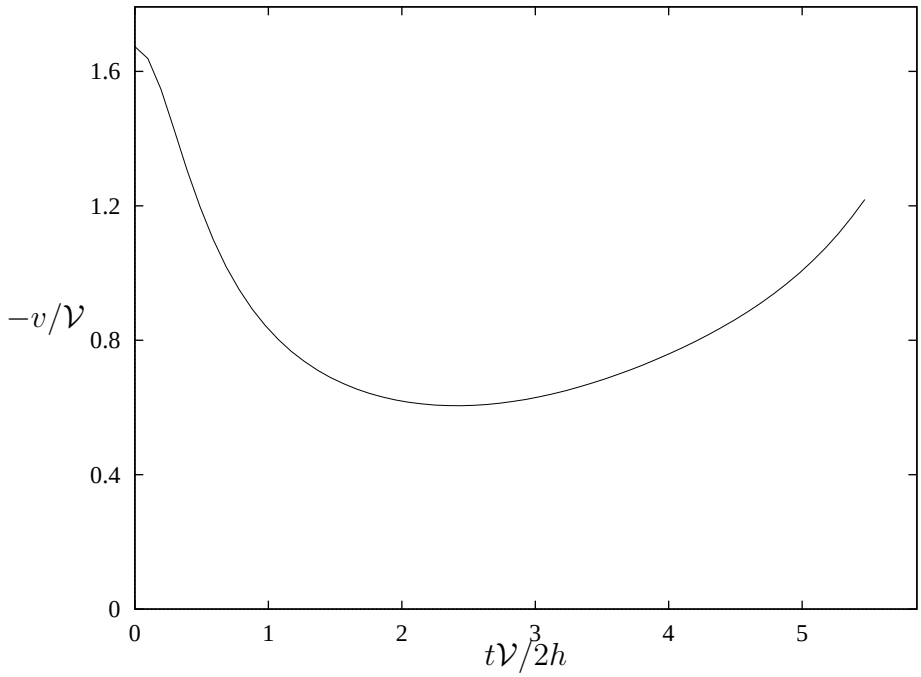

Figure 3.10: Bridge slumping velocity as a function of time for $h/w = 0.2$.

on the free surfaces — fully four inflexion points — would be undesirable if this type of sagging flow were to be used in the manufacturing of optical components.

Figure 3.10 shows the sagging velocity (at the point which starts as the centroid of the rectangle) as a function of time for a bridge of aspect ratio $h/w = 0.2$. At first the sagging velocity decreases rapidly with time from its initial value of $-v/\mathcal{V} = 1.674$, reaching a minimum value of about 37 percent of the initial velocity at time $t\mathcal{V}/(2h) = 2.5$ when the amount of sag equals a little over two and a quarter bridge thicknesses. It then increases again as the central portions of the bridge begin to lose touch with the walls and enter a state of flow which will eventually culminate in free fall. This type of behaviour is qualitatively typical for other aspect ratios too. However, as the aspect ratio falls below 0.2, we see a larger decrease in the sagging velocity over a longer time, with the bridge sagging considerably more that three bridge thicknesses before the velocity begins to increase again. For example, with $h/w = 0.1$ the velocity drops to 13 percent of its initial value by $t\mathcal{V}/2h = 15$,

at which time the bridge has sagged nearly four thicknesses. Conversely, when the aspect ratio increases above 0.2, there is a smaller decrease in velocity over a shorter time period, and the bridge sags a smaller distance in this time. For $h/w = 0.5$ the velocity decreases to 73 percent of its original value by $t\mathcal{V}/(2h) = 0.24$ before increasing again, at which time the bridge has sagged just one thickness.

The final free-fall state will eventually violate the assumption of neglect of inertia, but so long as inertia remains neglected, the velocity (and acceleration) will appear to approach infinity at some finite time $t = t_C = \mu/(\rho g L_e)$. Here L_e is an effective drop length which depends on w and h. This type of behaviour is considered in the context of another creeping flow in Chapter 4 and [89]. Meanwhile, we make no attempt to approach that time in the present computations. All accelerations being computed here are formally very small compared to gravity g, but, in this approach to free fall, they will become comparable to g at times just less than the critical time $t = t_C$, and inertia would have to be re-introduced into the problem in order to complete the solution for times close to and greater than t_C.

Figure 3.11 shows the time taken to sag one bridge thickness $2h$ as a function of aspect ratio. Obviously, we are not able to compute the y-axis intercept using finite-element methods, but can only approach $h/w = 0$. In fact computations were stopped at an aspect ratio of $h/w = 0.05$ since, for smaller aspect ratios, meshing becomes difficult and the finite-element method is not reliable. For this reason the figure given in [93] was left incomplete for very small aspect ratios. We are, however, able to complete the graph via a formal thin-layer asymptotic expansion which is described in the next section (Section 3.6). This gives a time of $t\mathcal{V}/(2h) = 1.487$ (see p. 71) for a bridge of very small aspect ratio $h/w \to 0$ to sag one bridge thickness.

Since the time scale used for Figure 3.11 is $2h/\mathcal{V}$, a unit value would correspond to fall at the constant speed $\mathcal{V}$, which is the initial rate of fall of very thin bridges with small h/w. However, for such thin bridges, even though the initial velocity is close to $\mathcal{V}$, we find a time nearly 50 percent larger than $2h/\mathcal{V}$, since (as in Figure 3.10)

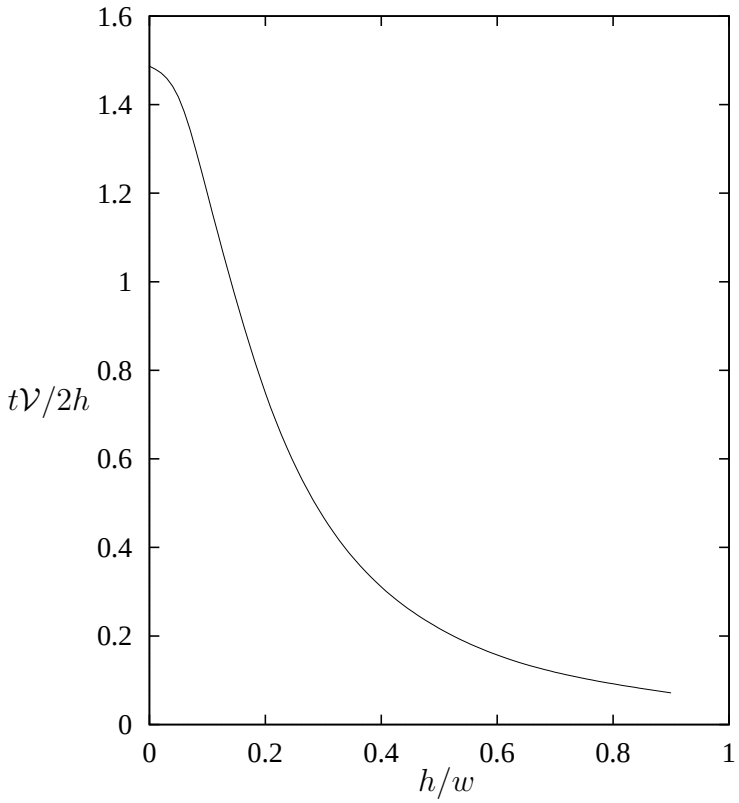

Figure 3.11: Time to sag one bridge thickness as a function of aspect ratio.

there is a significant drop in sagging velocity from its initial value during a sag of one bridge thickness. On the other hand, at larger h/w values, the time to slump one bridge thickness becomes significantly less than $2h/\mathcal{V}$, since the sagging velocity then significantly exceeds $\mathcal{V}$, not just initially (see Figure 3.6), but also at all later times. In fact, for very thick bridges, the time to slump one bridge thickness approaches $2h/\overline{\mathcal{V}}$ as h/w becomes large, corresponding to the flow approaching the Poiseuille thick-bridge limit with centroid slumping velocity $\overline{\mathcal{V}}$.

3.6 Motion of Nearly Plane Bridges

In Section 3.2 a solution for the initial sagging velocity of a thin bridge with $h/w \to 0$ was derived. This is a special case that can also be obtained from a formal asymp-

totic analysis of a two or three-dimensional version of the initial sagging problem, and the fully three-dimensional case is described in the Appendix of [93]. However this analysis can be further generalized for the sagging of a bridge of a general shape, both initially and at later times. This has been considered by Howell [48] and van de Fliert et al. [96] for transverse flow of thin viscous sheets, with both negligible and non-negligible curvature, in the absence of gravity. Here we consider a two-dimensional asymptotic-expansion method for gravity-driven flow of very viscous bridges which are thin and nearly flat (i.e. have small surface slope), sagging distances of up to a few bridge thicknesses in the vertical direction. The sag is small relative to the width of the bridge so that the bridge remains nearly flat. In addition to formally deriving the thin-bridge initial sagging-velocity profile, and perhaps of more interest, this analysis enables us to determine the time for a very thin bridge to sag one bridge thickness, as shown in Figure 3.11. This is a result that cannot be obtained by finite-element methods, or indeed any numerical method. It will also be of interest to see how well this thin-limit model matches the finite-element solution for bridges of finite thickness, and so determine its applicability to real problems.

Figure 3.12 shows the geometry of the problem under consideration. Again we have a very viscous fluid bridge of width $2w$ spanning the gap between two vertical no-slip walls. This time however the initial shape of the bridge is defined more generally by the centreline equation $y = \bar{h}(x,t)$, and the top and bottom free-surface equations $y = \bar{h}(x,t) \pm h(x,t)$. Thus $h(x,t)$ is the half-thickness of the bridge at position x and time t. By definition $h/w = \mathrm{O}(\epsilon)$ where ϵ is a very small positive real number. The only other restriction that we place on the initial geometry, so as to satisfy the requirement of being nearly flat, is that $\bar{h}$ only deviate from the horizontal by distances of order $\mathrm{O}(h)$. For the specific initially rectangular geometry we have been considering, we need the initial conditions at $t = 0$

$$\bar{h}(x,0) = 0 \tag{3.16}$$

$$h(x,0) = \epsilon w \tag{3.17}$$

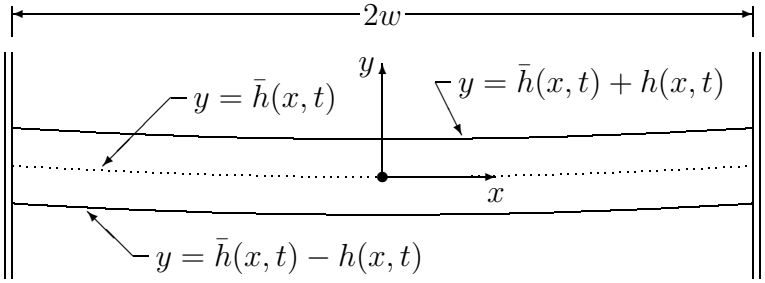

Figure 3.12: A thin and nearly-flat viscous bridge of general shape.

although we shall derive equations for a general initial bridge shape for as long as possible. We consider this bridge slumping distances of order O(h) vertically.

As in [48], we define new dimensionless variables (denoted by asterisks in keeping with our earlier convention) such that

$$\begin{aligned} &x = wx^*, && y = \epsilon w y^*, \\ &u(x,y,t) = \mathcal{V}u^*(x^*,y^*,t^*), && v(x,y,t) = \mathcal{V}v^*(x^*,y^*,t^*), \\ &t = \epsilon w t^*/\mathcal{V}, && p(x,y,t) = \mu\mathcal{V}p^*(x^*,y^*,t^*)/w, \\ &\bar{h}(x,t) = \epsilon w \bar{h}^*(x^*,t^*), && h(x,t) = \epsilon w h^*(x^*,t^*), \end{aligned} \tag{3.18}$$

where, as previously, u is the horizontal velocity component, v is the vertical velocity component, p is the pressure, and $\mathcal{V}$ is a velocity scale to be determined. Then $-1 \le x^* \le 1$ and initial conditions (3.16) and (3.17) become

$$\bar{h}^*(x^*,0) = 0, \tag{3.19}$$

$$h^*(x^*,0) = 1. \tag{3.20}$$

Now, with subscripts denoting differentiation, we express the problem in terms of these dimensionless variables so that the continuity equation (2.2) becomes

$$\epsilon u^*_{x^*} + v^*_{y^*} = 0, \tag{3.21}$$

and the Stokes equations (2.7) become

$$\epsilon^2 p^*_{x^*} = \epsilon^2 u^*_{x^*x^*} + u^*_{y^*y^*}, \tag{3.22}$$

$$\epsilon p^*_{y^*} = \epsilon^2 v^*_{x^*x^*} + v^*_{y^*y^*} - \epsilon^2 \frac{\rho g w^2}{\mu\mathcal{V}}. \tag{3.23}$$

At this stage we need to define the velocity scale $\mathcal{V}$. Already we know, from our series-expansion and finite-element analyses, that the initial central sagging velocity of a thin rectangular bridge, in the limit as $h/w \to 0$, is $\rho g w^4/(32\mu h^2)$ (see Figure 3.6). Since the thin initially-rectangular bridge is a particular case in the class of bridges we are now considering, given by initial conditions (3.16) and (3.17) so that $h/w = \epsilon$, this indicates that a suitable choice for the velocity scale is $\mathcal{V} = \rho g w^4/(\mu h^2) = \rho g w^2/(\epsilon^2 \mu)$. The last term in (3.23) is then ϵ^4.

We also write the boundary conditions in terms of the new variables. On the walls $x^* = \pm 1$ the no-slip boundary conditions (2.12) become

$$u^* = v^* = 0 \tag{3.24}$$

while, on the zero stress free surfaces $y^* = \bar{h}^*(x^*, t^*) \pm h^*(x^*, t^*)$, the no-stress boundary conditions (2.15) become

$$\epsilon^2(p^* - 2u^*_{x^*})(\bar{h}^*_{x^*} \pm h^*_{x^*}) + u^*_{y^*} + \epsilon v^*_{x^*} = 0, \tag{3.25}$$

$$-\epsilon p^* - \epsilon(\bar{h}^*_{x^*} \pm h^*_{x^*})(u^*_{y^*} + \epsilon v^*_{x^*}) + 2v^*_{y^*} = 0. \tag{3.26}$$

In this analysis we do not use the Lagrangian method to compute the evolution of the geometry in time, but rather use the kinematic condition (2.18) which becomes

$$\bar{h}^*_{t^*} \pm h^*_{t^*} + \epsilon u^*(\bar{h}^*_{x^*} \pm h^*_{x^*}) = v^*. \tag{3.27}$$

Next we expand u^*, v^*, p^*, $\bar{h}^*$ and h^* in power series in epsilon

$$\begin{aligned}
u^* &= u^*_0 + \epsilon u^*_1 + \epsilon^2 u^*_2 + \epsilon^3 u^*_3 + \dots, \\
v^* &= v^*_0 + \epsilon v^*_1 + \epsilon^2 v^*_2 + \epsilon^3 v^*_3 + \dots, \\
p^* &= p^*_0 + \epsilon p^*_1 + \epsilon^2 p^*_2 + \epsilon^3 p^*_3 + \dots, \\
\bar{h}^* &= \bar{h}^*_0 + \epsilon \bar{h}^*_1 + \epsilon^2 \bar{h}^*_2 + \epsilon^3 \bar{h}^*_3 + \dots, \\
h^* &= h^*_0 + \epsilon h^*_1 + \epsilon^2 h^*_2 + \epsilon^3 h^*_3 + \dots,
\end{aligned} \tag{3.28}$$

and after substituting into (3.21) to (3.27), equate like powers of epsilon.

Dropping asterisks on dimensionless variables, continuity immediately gives v_0 as a function of x and t only

$$v_0 = v_0(x,t), \tag{3.29}$$

and, in addition, the leading-order kinematic boundary condition on $y = \bar{h} \pm h$ (i.e. $\bar{h}_{0t} \pm h_{0t} = v_0$) immediately gives

$$\bar{h}_{0t} = v_0, \tag{3.30}$$

$$h_{0t} = 0, \tag{3.31}$$

which says that, to leading order, the bridge thickness at any position x remains unchanged in time. Thus, $h_0 = h_0(x)$ for all time, and with initial condition (3.20) we have $h_0(x) = 1$.

Continuity further demands

$$u_{0x} = -v_{1y}, \tag{3.32}$$

$$u_{1x} = -v_{2y}, \tag{3.33}$$

$$u_{2x} = -v_{3y}, \tag{3.34}$$

$$u_{3x} = -v_{4y}, \tag{3.35}$$

and the Stokes equations demand $u_{0yy} = u_{1yy} = 0$ so that

$$u_0 = u_{00} + y u_{01}, \tag{3.36}$$

$$u_1 = u_{10} + y u_{11}, \tag{3.37}$$

where it is to be understood that all coefficients u_{ij} are functions of x and t. From the Stokes equations we also obtain

$$p_{0x} = u_{0xx} + u_{2yy}, \tag{3.38}$$

$$p_{1x} = u_{1xx} + u_{3yy}, \tag{3.39}$$

$$p_{2x} = u_{2xx} + u_{4yy}, \tag{3.40}$$

$$p_{3x} = u_{3xx} + u_{5yy}, \tag{3.41}$$

and

$$p_{0y} = v_{1yy}, \tag{3.42}$$

$$p_{1y} = v_{0xx} + v_{2yy}, \tag{3.43}$$

$$p_{2y} = v_{1xx} + v_{3yy}, \tag{3.44}$$

$$p_{3y} = v_{2xx} + v_{4yy} - 1. \tag{3.45}$$

Note that the last "-1" term in (3.45) is due to gravity, and we must carry the present asymptotic analysis at least as far as this in order to capture the effect of the gravitational driving force.

On the free surfaces $y = \bar{h} \pm h$ we have from (3.25)

$$u_{0y} = 0, \tag{3.46}$$

$$u_{1y} + v_{0x} = 0, \tag{3.47}$$

$$(p_0 - 2u_{0x})(\bar{h}_{0x} \pm h_{0x}) + u_{2y} + v_{1x} = 0, \tag{3.48}$$

$$\begin{aligned}(p_0 - 2u_{0x})(\bar{h}_{1x} \pm h_{1x}) + (p_1 - 2u_{1x})(\bar{h}_{0x} \pm h_{0x}) \\ +u_{3y} + v_{2x} = 0,\end{aligned} \tag{3.49}$$

$$\begin{aligned}(p_0 - 2u_{0x})(\bar{h}_{2x} \pm h_{2x}) + (p_1 - 2u_{1x})(\bar{h}_{1x} \pm h_{1x}) \\ +(p_2 - 2u_{2x})(\bar{h}_{0x} \pm h_{0x}) + u_{4y} + v_{3x} = 0,\end{aligned} \tag{3.50}$$

and from (3.26)

$$-p_0 + 2v_{1y} = 0, \tag{3.51}$$

$$-p_1 + 2v_{2y} = 0, \tag{3.52}$$

$$-p_2 - (u_{2y} + v_{1x})(\bar{h}_{0x} \pm h_{0x}) + 2v_{3y} = 0, \tag{3.53}$$

$$\begin{aligned}-p_3 - (u_{2y} + v_{1x})(\bar{h}_{1x} \pm h_{1x}) \\ +(u_{3y} + v_{2x})(\bar{h}_{0x} \pm h_{0x}) + 2v_{4y} = 0,\end{aligned} \tag{3.54}$$

and from the kinematic condition (3.27)

$$\bar{h}_{1t} \pm h_{1t} + u_0(\bar{h}_{0x} \pm h_{0x}) = v_1, \tag{3.55}$$

$$\bar{h}_{2t} \pm h_{2t} + u_0(\bar{h}_{1x} \pm h_{1x}) + u_1(\bar{h}_{0x} \pm h_{0x}) = v_2. \tag{3.56}$$

Now from (3.36) and (3.46) we must have $u_{01} = 0$ so that

$$u_0 = u_0(x,t). \tag{3.57}$$

This, together with (3.32) leads to

$$v_1 = -yu_{0x} + v_{10}, \tag{3.58}$$

and on substituting for v_1 in (3.51) we obtain

$$p_0 = p_0(x,t) = -2u_{0x} \tag{3.59}$$

which in combination with (3.38) gives

$$u_2 = -\frac{3}{2}y^2u_{0xx} + yu_{21} + u_{20}. \tag{3.60}$$

Equations (3.33) and (3.37) give

$$v_2 = -\frac{1}{2}y^2u_{11x} - yu_{10x} + v_{20}, \tag{3.61}$$

and (3.47) and (3.37) tell us that $u_{11} = -v_{0x}$.

We can now write the two free surface equations (3.48) as

$$-4u_{0x}(\bar{h}_{0x} \pm h_{0x}) - 4u_{0xx}(\bar{h}_0 \pm h_0) + u_{21} + v_{10x} = 0, \tag{3.62}$$

and when one is subtracted from the other we obtain

$$(h_0u_{0x})_x = 0. \tag{3.63}$$

Then

$$h_0u_{0x} = f(t), \tag{3.64}$$

$$u_0 = f(t)\int_{-1}^{x} \frac{1}{h_0(\xi)}\,d\xi, \tag{3.65}$$

where f is some function of t, and we have used the wall boundary condition $u_0(-1,t) = 0$. Since we must also satisfy the other wall boundary condition

$u_0(1,t) = 0$, and $h_0(x) > 0$ for $-1 \leq x \leq 1$, we must have $f(t) = 0$ and hence $u_0(x,t) = 0$. We finally have from (3.62) that

$$u_{21} = -v_{10x} = -v_{1x}. \tag{3.66}$$

Substituting back $u_0 = 0$ into our earlier working, we have in summary

$$\begin{aligned}
&u_0 = 0,\\
&v_0 = v_0(x,t),\\
&h_0 = h_0(x),\\
&\bar{h}_{0t} = v_0,\\
&p_0 = 0,\\
&u_1 = -yv_{0x}(x,t) + u_{10}(x,t),\\
&v_1 = v_1(x,t),\\
&h_1 = h_1(x),\\
&\bar{h}_{1t} = v_1,\\
&u_2 = -yv_{1x}(x,t) + u_{20}(x,t), \quad \text{and}\\
&v_2 = \frac{1}{2}y^2 v_{0xx}(x,t) - yu_{10x}(x,t) + v_{20}(x,t).
\end{aligned} \tag{3.67}$$

It is a simple matter to determine from (3.43) and (3.52) that

$$p_1 = 2yv_{0xx} - 2u_{10x}, \tag{3.68}$$

and from (3.34) that

$$v_3 = \frac{1}{2}y^2 v_{1xx} - yu_{20x} + v_{30}. \tag{3.69}$$

Equations (3.39) and (3.35) give

$$u_3 = \frac{1}{2}y^3 v_{0xxx} - \frac{3}{2}y^2 u_{10xx} + yu_{31} + u_{30}, \tag{3.70}$$

$$v_4 = -\frac{1}{8}y^4 v_{0xxxx} + \frac{1}{2}y^3 u_{10xxx} - \frac{1}{2}y^2 u_{31x} - yu_{30x} + v_{40}. \tag{3.71}$$

Next we can make the appropriate substitutions into (3.49) which must be satisfied on $y = \bar{h} \pm h$ yielding

$$(\bar{h}_0 h_0 v_{0xx} - h_0 u_{10x})_x = 0, \tag{3.72}$$

and

$$u_{31} + v_{20x} = -2(v_{0xx}(\bar{h}_0^2 + h_0^2))_x + 4(\bar{h}_0 u_{10x})_x. \tag{3.73}$$

An equation for p_2 can be obtained from (3.44) and (3.53)

$$p_2 = 2y v_{1xx} - 2u_{20x}, \tag{3.74}$$

and (3.40) then gives

$$u_4 = \frac{1}{2}y^3 v_{1xxx} - \frac{3}{2}y^2 u_{20xx} + y u_{41} + u_{40}. \tag{3.75}$$

We now have sufficient information to solve (3.45) and so include gravity in our solution. Substituting for v_{2xx} and v_{4yy} gives

$$p_{3y} = -y^2 v_{0xxxx} + 2y u_{10xxx} + v_{20xx} - u_{31x} - 1 \tag{3.76}$$

which we integrate yielding

$$p_3 = -\frac{1}{3}y^3 v_{0xxxx} + y^2 u_{10xxx} + y(v_{20xx} - u_{31x} - 1) + p_{30}. \tag{3.77}$$

On the free surfaces we must satisfy (3.54), and, since $u_{2y} + v_{1x} = 0$ (refer (3.68)), this becomes

$$\begin{aligned} p_3 &= (u_{3y} + v_{2x})(\bar{h}_{0x} \pm h_{0x}) + 2v_{4y} \\ &= -(2y^2 v_{0xxx} - 4y u_{10xx} + u_{31} + v_{20x})(\bar{h}_{0x} \pm h_{0x}) \\ &\quad -y^3 v_{0xxxx} + 3y^2 u_{10xxx} - 2y u_{31x} - 2u_{30x}. \end{aligned} \tag{3.78}$$

Substituting $y = \bar{h} \pm h$ into (3.77) and (3.78) and equating gives two expressions for p_{30}

$$\begin{aligned} p_{30} = &-\frac{2}{3}(\bar{h}_0 \pm h_0)^3 v_{0xxxx} + 2(\bar{h}_0 \pm h_0)^2 u_{10xxx} \\ &-(\bar{h}_0 \pm h_0)(v_{20xx} + u_{31x} - 1) - 2u_{30x} - 2(\bar{h}_0 \pm h_0)^2(\bar{h}_{0x} \pm h_{0x}) v_{0xxx} \\ &+4(\bar{h}_0 \pm h_0)(\bar{h}_{0x} \pm h_{0x}) u_{10xx} - (\bar{h}_{0x} \pm h_{0x})(v_{20x} + u_{31}). \end{aligned} \tag{3.79}$$

Subtracting one from the other and using (3.73) gives

$$-h_0 = \frac{4}{3}(h_0^3 v_{0xx})_{xx} + 4(h_0\bar{h}_0\bar{h}_{0x}v_{0xx})_x - 4(h_0\bar{h}_{0x}u_{10x})_x. \tag{3.80}$$

This can be even further simplified by using the relation given in (3.72) which can be integrated to give

$$h_0 u_{10x} = h_0\bar{h}_0 v_{0xx} - f(t) \tag{3.81}$$

for some function f of time t to be determined later. Then (3.80) can be written

$$\begin{aligned} -h_0 &= \frac{4}{3}(h_o^3 v_{0xx})_{xx} + 4\bar{h}_{0xx}(h_0\bar{h}_0 v_{0xx} - h_0 u_{10x}) \\ &= \frac{4}{3}(h_0^3 v_{0xx})_{xx} + 4\bar{h}_{0xx} f(t). \end{aligned} \tag{3.82}$$

If, at this stage, we use the initial conditions (3.19) and (3.20) we obtain at $t = 0$

$$-1 = \frac{4}{3}v_{0xxxx}. \tag{3.83}$$

This compares with the dimensional form of the leading-order equation for the initial vertical velocity of a rectangular bridge obtained by Tuck et al. [93] from a similar asymptotic expansion for the initial sagging problem only. If we remove gravity from the problem by setting $g = 0$ then the left hand side of (3.82) becomes zero, and, after substituting $\bar{h}_{0t}$ for v_0 (see equations (3.68)), we obtain a two-dimensional form of the equation given by Howell [48] for transverse flow of a nearly-flat thin viscous sheet in the absence of gravity. Thus (3.82) contains two special results that have been derived elsewhere.

We can use the wall boundary conditions to determine the function $f(t)$ as follows. Designating $v_0 = V$, $u_{10} = U$ and $\bar{h}_0 = \mathcal{H}$, we rewrite (3.81) as

$$U_x = \mathcal{H}V_{xx} - \frac{f(t)}{h_0}, \tag{3.84}$$

and the equation for u_1 from (3.68) as

$$u_1 = -yV_x(x,t) + U(x,t). \tag{3.85}$$

Since $u_1(\pm 1, y, t) = 0$ for $-1 \leq y \leq 1$ and $t \geq 0$, setting $y = 0$ and then $y \neq 0$ shows that we must have

$$U(\pm 1, t) = 0, \qquad \text{and} \tag{3.86}$$

$$V_x(\pm 1, t) = 0. \tag{3.87}$$

Then

$$U = \int_{-1}^{x} \mathcal{H} V_{xx}\, d\xi - f(t) \int_{-1}^{x} \frac{d\xi}{h_0} \tag{3.88}$$

satisfies $U = 0$ at $x = -1$, and satisfying $U = 0$ at $x = 1$ leads to

$$f(t) = \int_{-1}^{1} \mathcal{H} V_{xx}\, dx \Big/ \int_{-1}^{1} \frac{dx}{h_0}. \tag{3.89}$$

Equations (3.82) and (3.89) give the leading-order flow over time of a thin nearly-flat viscous bridge of arbitrary initial shape given by $y = \bar{h}(x,t) \pm h(x,t)$. To proceed further we need to define this shape via some initial conditions. Since we wish to consider an initially rectangular bridge, the appropriate initial conditions are given in (3.19) and (3.20), and in addition we have established (from (3.31)) that h_0 is constant in time.

Firstly, setting $h_0 = 1$ gives

$$f(t) = \frac{1}{2} \int_{-1}^{1} \mathcal{H} V_{xx}\, dx, \tag{3.90}$$

and then substituting $V = v_0 = \bar{h}_{0t} = \mathcal{H}_t$ from (3.68) gives

$$\boxed{-1 = \frac{4}{3} \mathcal{H}_{txxxx} + 2\mathcal{H}_{xx} \int_{-1}^{1} \mathcal{H} \mathcal{H}_{txx}\, dx.} \tag{3.91}$$

This fourth-order partial differential equation defines the leading-order centreline profile $\mathcal{H}(x,t)$ and sagging velocity $\mathcal{H}_t$ of a thin bridge with initial uniform thickness. Although we obtain only a centreline profile, we can readily determine from this the leading-order time evolution of top and bottom free surfaces using our knowledge that the thickness $h(x,t)$ is constant to leading order for all time (i.e. $h_0(x,t) = 1$). Of course, because of the assumptions we have made en route, (3.91) cannot be

used once the sag exceeds a few bridge thicknesses and $\mathcal{H}_x$ becomes large so that the bridge is no longer nearly flat.

For an initially rectangular bridge, we must next require that $\mathcal{H} = 0$ at $t = 0$. With this initial condition plus boundary conditions at the walls we can now solve (3.91). By virtue of the derivation of this equation, all solutions $\mathcal{H}(x,t)$ will satisfy the boundary condition of zero leading-order horizontal velocity at the walls at $y = 0$. One solution that also satisfies the boundary condition of zero leading-order vertical velocity at the walls, i.e. $V = 0$ at $x = \pm 1$, is readily obtained by separation of variables:

$$\mathcal{H} = \left(\frac{9t}{32}\right)^{1/3} (x^2 - 1). \tag{3.92}$$

However this does not satisfy the additional boundary condition given by (3.87), so that the horizontal velocity component is zero at the walls only at the centre $y = 0$. This solution is, in fact, for a bending 'beam' pinned at $(x, y) = (\pm 1, 0)$ with $u_1(\pm 1, y) = -u_1(\pm 1, -y)$. Hence we look at obtaining a solution to (3.91) subject to $V = \mathcal{H}_t = 0$ and $V_x = \mathcal{H}_{tx} = 0$ at $x = \pm 1$. Using symmetry about $x = 0$ we can limit our considerations to $0 \le x \le 1$ and satisfy boundary conditions

$$\mathcal{H}_{tx} = \mathcal{H}_{txxx} = 0 \qquad \text{at } x = 0, \tag{3.93}$$

$$\mathcal{H}_t = \mathcal{H}_{tx} = 0 \qquad \text{at } x = 1. \tag{3.94}$$

We proceed as follows, first integrating (3.91) by parts

$$\begin{aligned} -1 &= \frac{4}{3}\mathcal{H}_{txxxx} - 2\mathcal{H}_{xx}\int_{-1}^{1} \mathcal{H}_x \mathcal{H}_{tx}\, dx \\ &= \frac{4}{3}\mathcal{H}_{txxxx} - 2\mathcal{H}_{xx}\int_{-1}^{1} \frac{1}{2}(\mathcal{H}_x^2)_t\, dx \\ &= \frac{4}{3}\mathcal{H}_{txxxx} - 2\dot{\overline{\mathcal{H}_x^2}}\mathcal{H}_{xx} \end{aligned} \tag{3.95}$$

where

$$\dot{\overline{\mathcal{H}_x^2}} = \left(\frac{1}{2}\int_{-1}^{1} \mathcal{H}_x^2\, dx\right)_t = \left(\int_0^1 \mathcal{H}_x^2\, dx\right)_t. \tag{3.96}$$

Because we are only considering bridges that start nearly flat and remain so during the sagging time period, $\mathcal{H}_x$ is small and $\mathcal{H}_x^2$ is even smaller. Thus the rate of strain at time t is given by

$$T(t) = \left(\int_0^1 1 - (1+\mathcal{H}_x^2)^{1/2} dx\right)_t \approx -\left(\int_0^1 \frac{1}{2}\mathcal{H}_x^2 dx\right)_t = -\frac{1}{2}\dot{\overline{\mathcal{H}_x^2}} \tag{3.97}$$

where a Taylor expansion approximation for $(1+\mathcal{H}_x^2)^{1/2}$ has been used. Multiplying the rate of strain by the viscosity yields the tension in the sagging bridge, for which reason we call the function $T(t)$ in (3.97) the "tension" function.

Rearranging equation (3.95) we have

$$\frac{4}{3}\mathcal{H}_{txxxx} = -1 + 2\dot{\overline{\mathcal{H}_x^2}}\mathcal{H}_{xx}, \tag{3.98}$$

and integrating with respect to x gives

$$\frac{4}{3}\mathcal{H}_{txxx} = -x + 2\dot{\overline{\mathcal{H}_x^2}}\mathcal{H}_x + f_1(t). \tag{3.99}$$

Boundary conditions (3.93) dictate that $f_1(t) = 0$. Integrating twice more with respect to x gives

$$\frac{4}{3}\mathcal{H}_{tx} = -\frac{x^3}{6} + 2\dot{\overline{\mathcal{H}_x^2}}\int_0^x \mathcal{H}\, d\xi + x\, f_2(t) + f_3(t), \tag{3.100}$$

and again boundary conditions (3.93) dictate that $f_3(t) = 0$ while boundary conditions (3.94) give

$$f_2(t) = \frac{1}{6} - 2\dot{\overline{\mathcal{H}_x^2}}\int_0^1 \mathcal{H}\, d\xi. \tag{3.101}$$

Integrating once again gives

$$\begin{aligned}\frac{4}{3}\mathcal{H}_t &= -\frac{x^4}{24} + 2\dot{\overline{\mathcal{H}_x^2}}\int_0^x\int_0^\eta \mathcal{H}\, d\xi\, d\eta + \frac{x^2}{2}f_2(t) + f_4(t)\\ &= -\frac{x^4}{24} + 2\dot{\overline{\mathcal{H}_x^2}}\int_0^x (x-\xi)\mathcal{H}\, d\xi + \frac{x^2}{2}f_2(t) + f_4(t).\end{aligned} \tag{3.102}$$

Finally boundary conditions (3.94) give

$$f_4(t) = \frac{1}{24} - \frac{1}{2}f_2(t) - 2\dot{\overline{\mathcal{H}_x^2}}\int_0^1 (x-\xi)\mathcal{H}\, d\xi. \tag{3.103}$$

Defining

$$\overline{\mathcal{H}} = \int_0^1 \mathcal{H}(\xi, t)\, d\xi \tag{3.104}$$

we have, after some manipulation,

$$\begin{aligned} \frac{4}{3} H_t = &-\frac{1}{24}(1-x^2)^2 \\ &-\dot{\overline{\mathcal{H}}_x^2}\left((1-x)^2\,\overline{\mathcal{H}} + 2\int_x^1 (x-\xi)\mathcal{H}(\xi,t)\, d\xi\right). \end{aligned} \tag{3.105}$$

At $t = 0$, $\mathcal{H}(x, 0) = 0$ so that (3.105) becomes

$$V = \mathcal{H}_t = -\frac{1}{32}(1-x^2)^2. \tag{3.106}$$

This is exactly equivalent to the quadratic-squared thin-bridge initial-velocity limit, obtained by the series-expansion technique of Section 3.2, and given in (3.9). Thus our ad hoc derivation of that equation is confirmed.

Equation (3.105) is quite simply solved using finite-difference methods. We define a one dimensional uniform grid on $0 \le x \le 1$ such that

$$x_j = j\,\Delta x, \quad j = 0, 1, \ldots, N_x \tag{3.107}$$

with $\Delta x = 1/N_x$, and also specify a time step Δt so that

$$t_i = i\,\Delta t, \quad i = 0, 1, \ldots\ . \tag{3.108}$$

Then

$$\mathcal{H}_j^i = \mathcal{H}(x_j, t_i). \tag{3.109}$$

Using Euler's method to approximate time derivatives, and central differencing to approximate spatial derivatives gives

$$(\mathcal{H}_t)_j^i = \frac{\mathcal{H}_j^i - \mathcal{H}_j^{i-1}}{\Delta t} \tag{3.110}$$

$$(\mathcal{H}_x)_j^i = \frac{\mathcal{H}_{j+1}^i - \mathcal{H}_{j-1}^i}{2\Delta x}. \tag{3.111}$$

With initial condition given by

$$\mathcal{H}(x,0) = 0, \tag{3.112}$$

and using $\mathcal{H}_{tx} = \mathcal{H}_{xt}$, the boundary conditions (3.93) and (3.94) can be expressed as

$$(\mathcal{H}_x)_0^i = (\mathcal{H}_{xxx})_0^i = 0 \qquad \text{and} \tag{3.113}$$

$$(\mathcal{H})_{N_x}^i = (\mathcal{H}_x)_{N_x}^i = 0 \tag{3.114}$$

for all i.

On using the trapezoidal rule to evaluate integrals, and the boundary conditions (3.113) and (3.114) to eliminate some terms, a discretization of (3.105) is

$$\begin{aligned}\frac{4}{3}(\mathcal{H}_j^i - \mathcal{H}_j^{i-1}) = &-\frac{\Delta t}{24}(1-x_j^2)^2 \\ &-G_j^i(\Delta x)^2 \sum_{k=1}^{N_x-1}\left((\mathcal{H}_x^2)_k^i - (\mathcal{H}_x^2)_k^{i-1}\right)\end{aligned} \tag{3.115}$$

with

$$G_j^i = (1-x_j)^2\left(\frac{\mathcal{H}_0^i}{2} + \sum_{k=1}^{N_x-1}\mathcal{H}_k^i\right) + 2\sum_{k=j+1}^{N_x-1}(x_j - x_k)\mathcal{H}_k^i. \tag{3.116}$$

This system of equations is clearly non-linear and hence we modify it further to render it more easily solvable.

An explicit finite-difference scheme is given by

$$\begin{aligned}\frac{4}{3}(\mathcal{H}_j^i - \mathcal{H}_j^{i-1}) = &-\frac{\Delta t}{24}(1-x_j^2)^2 \\ &-G_j^{i-1}(\Delta x)^2 \sum_{k=1}^{N_x-1}\left((\mathcal{H}_x^2)_k^{i-1} - (\mathcal{H}_x^2)_k^{i-2}\right),\end{aligned} \tag{3.117}$$

and on using central differencing (3.111) to evaluate the x derivatives we finally have

$$\frac{4}{3}\mathcal{H}_j^i = \frac{4}{3}\mathcal{H}_j^{i-1} - \frac{\Delta t}{24}(1-x_j^2)^2 - \frac{1}{4}G_j^{i-1}(F^{i-1} - F^{i-2}) \tag{3.118}$$

where

$$F^i = \sum_{k=1}^{N_x-1}(\mathcal{H}_{k+1}^i - \mathcal{H}_{k-1}^i)^2. \tag{3.119}$$

For better accuracy we can use an iterative explicit finite-difference scheme

$$\frac{4}{3}\left(\left[\mathcal{H}_j^i\right]^n - \mathcal{H}_j^{i-1}\right) = -\frac{\Delta t}{24}(1-x_j^2)^2 - \left[G_j^i\right]^{n-1}(\Delta x)^2 \sum_{k=1}^{N_x-1}\left(\left[(\mathcal{H}_x^2)_k^i\right]^{n-1} - (\mathcal{H}_x^2)_k^{i-1}\right) \tag{3.120}$$

where the superscript n denotes the nth iterate, and

$$\left[\mathcal{H}_j^i\right]^0 = \mathcal{H}_j^{i-1} \qquad \text{and} \tag{3.121}$$

$$\left[(\mathcal{H}_x^2)_j^i\right]^0 = (\mathcal{H}_x^2)_j^{i-1}. \tag{3.122}$$

Iteration ceases when $\left[\mathcal{H}_j^i\right]^n = \left[\mathcal{H}_j^i\right]^{n-1}$ to within some small tolerance for all j. This scheme gives

$$\frac{4}{3}\left[\mathcal{H}_j^i\right]^n = \frac{4}{3}\mathcal{H}_j^{i-1} - \frac{\Delta t}{24}(1-x_j^2)^2 - \frac{1}{4}\left[G_j^i\right]^{n-1}\left(\left[F^i\right]^{n-1} - F^{i-1}\right). \tag{3.123}$$

To compare the results obtained by solving (3.118) and (3.123) let us compute the "tension" function given by (3.97). In discretized form this becomes

$$-T^i = \frac{F^i - F^{i-1}}{8\Delta t \Delta x} \tag{3.124}$$

Let us also compare the sag $\mathcal{H}(0,t)$ computed by the two methods. Some results, at dimensionless time $t = 40$ are given in Table 3.8, and show the explicit method to be reasonably accurate. Three to four figures of accuracy is obtained with a grid spacing of 0.02 and a time step of 0.1, which is satisfactory for our purposes. For better accuracy a very much smaller time step must be used, which increases the computational time considerably.

The explicit method becomes unstable for $t > 89$ when the bridge has sagged just over 1.9 thicknesses, while the iterative explicit method fails to converge after about $t = 66.7$ when the bridge has sagged about 1.59 thickness. Hence we try an iterative implicit method given by substituting

$$\left[F^i\right]^{n-1} \approx \sum_{k=1}^{N_x-1}\left(\left[\mathcal{H}_{k+1}^i\right]^n - \left[\mathcal{H}_{k-1}^i\right]^n\right)\left(\left[\mathcal{H}_{k+1}^i\right]^{n-1} - \left[\mathcal{H}_{k-1}^i\right]^{n-1}\right) \tag{3.125}$$

Table 3.8: Sag of a thin bridge: a comparison of explicit and iterative explicit methods at dimensionless time $t = 40$.

Δx	Δt	Explicit		Iterative Explicit	
		$-\mathcal{H}(0)$	$T \times 10^2$	$-\mathcal{H}(0)$	$T \times 10^2$
0.02	0.1	1.092159	2.896567	1.091500	2.893752
0.01	0.1	1.092069	2.898878	1.091410	2.896063
0.002	0.1	1.092040	2.899618	1.091381	2.896803
0.02	0.01	1.091742	2.895758	1.091654	2.891765
0.01	0.01	1.091652	2.898067	1.091564	2.894054
0.002	0.01	1.091623	2.898807		
0.02	0.001	1.091700	2.895677		
0.01	0.001	1.091610	2.897986		

Table 3.9: Sag of a thin bridge by an iterative implicit method at dimensionless time $t = 40$.

Δx	Δt	$-\mathcal{H}(0)$	$T \times 10^2$
0.02	0.1	1.091500	2.893353
0.01	0.1	1.091410	2.895661
0.02	0.01	1.091632	2.895925

in (3.123) so that we have

$$\begin{aligned}
&\frac{4}{3}\left[\mathcal{H}_j^i\right]^n + \frac{1}{4}\left[G_j^i\right]^n \sum_{k=2}^{N_x-1}\left[\mathcal{H}_k^i\right]^n\left(\left[\mathcal{H}_k^i\right]^{n-1} - \left[\mathcal{H}_{k-2}^i\right]^{n-1}\right) \\
&\quad -\frac{1}{4}\left[G_j^i\right]^n \sum_{k=0}^{N_x-2}\left[\mathcal{H}_k^i\right]^n\left(\left[\mathcal{H}_{k+2}^i\right]^{n-1} - \left[\mathcal{H}_k^i\right]^{n-1}\right) \\
&\quad = \frac{4}{3}\mathcal{H}_j^{i-1} - \frac{\Delta t}{24}(1-x_j^2)^2 + \frac{1}{4}\left[G_j^i\right]^{n-1} F^{i-1}. \qquad (3.126)
\end{aligned}$$

This method is considerably more expensive with regards to computing time, but is more stable and permits computation out to large t. Some results at $t = 40$ are given in Table 3.9 for comparison with those given for the two explicit methods in Table 3.8.

The tension function is plotted against time in Figure 3.13 showing an initial

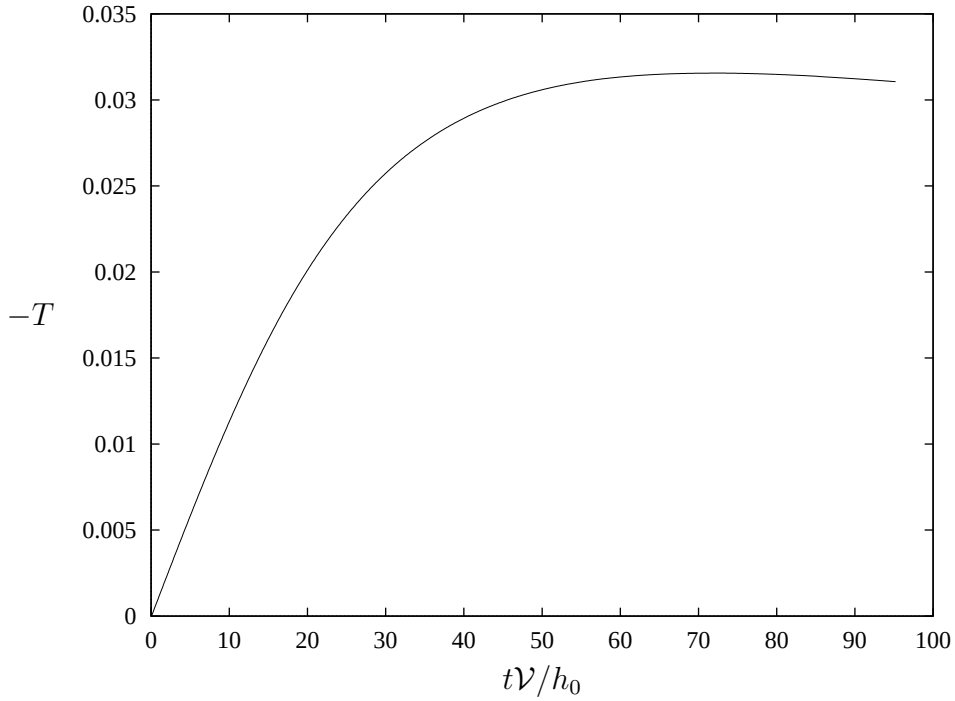

Figure 3.13: Tension versus time for a sagging thin bridge.

rapid increase, with a maximum reached at about $t = 70$ followed by a slow decrease.

With the iterative implicit method we are able to compute the time taken to sag one bridge thickness $2h_0$, not previously able to be determined by finite-element methods, at about $t\mathcal{V}/h_0 = 95.15$ (where t is now a dimensional time, $h_0 = \epsilon w$ is the initial half-thickness of the bridge, and we have returned to denoting dimensionless variables by asterisks). This value was obtained from output of $(t^*,\mathcal{H}^*(0,t^*))$ pairs with a time increment of $\Delta t^* = 0.1$, and using linear interpolation between consecutive time points. For plotting on Figure 3.11 we must re-scale this result using the time scale of the plot. Let us denote the velocity and times scales of these thin-bridge calculations by $\mathcal{V}_1 = \rho g w^4/(\mu h_0^2)$ and $\mathcal{T}_1 = h_0/\mathcal{V}_1$ respectively, while the velocity and time scales appropriate to Figure 3.11 are $\mathcal{V}_2 = \rho g w^4/(32\mu h_0^2)$ and $\mathcal{T}_2 = 2h_0/\mathcal{V}_2$ respectively. Then the y-intercept on Figure 3.11, corresponding to the time taken for an initially rectangular bridge of aspect ratio $h/w \to 0$ to sag one bridge thickness, is given by $95.15\mathcal{T}_1/\mathcal{T}_2 = 95.15/64 = 1.487$. This result is

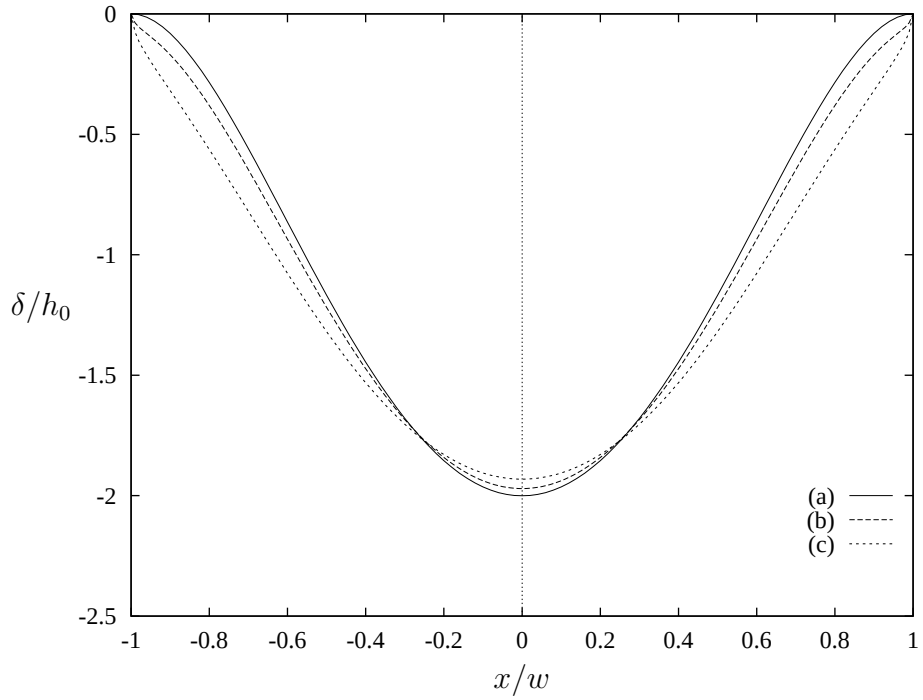

Figure 3.14: The bridge centreline δ/h_0 after sagging one bridge thickness for (a) $h_0/w \to 0$ by asymptotic analysis, (b) $h_0/w = 0.1$ by finite-element methods and (c) $h_0/w = 0.2$ also by finite-element methods.

completely consistent with our finite-element results, as seen in Figure 3.11.

Figure 3.14 shows the sag $\delta(x,t)/h_0$ profile across the width of the bridge as determined to leading order by our thin-limit analysis, i.e. $\delta/h_0 \approx \mathcal{H}^* = \bar{h}_0/h$. The curve given is for time $t\mathcal{V}/h_0 = 95.2$, when the bridge has sagged a distance of just fractionally over one initial bridge thickness $2h_0$. The vertical length scale used for the plot is the initial bridge half-thickness $h_0 = \epsilon w$ so that a sag of one bridge thickness is given by $\delta(0,t)/h_0 = 2$. Also shown, for comparison, are the centreline profiles for bridges of aspect ratios $h_0/w = 0.1$ and $h_0/w = 0.2$ after they have sagged one initial thickness. These were obtained by fitting cubic-spline curves to top and bottom surface nodes given as output from the finite-element program, and then computing the average of the top and bottom free surfaces at positions x/w across the bridge width. The vertical sag data $\hat{\delta} = \delta(x,t)/w$ so obtained has then been re-scaled to correspond to the vertical length scale of the plot, i.e.

$\delta(x,t)/h_0 = \hat{\delta}/(h_0/w)$. From Figure 3.14 we see that the asymptotic method gives a reasonable estimate of central sag even for a bridge of aspect ratio as large as $h_0/w = 0.2$, although the profile towards the edge varies somewhat. As aspect ratio decreases the overall correspondence between the thin-limit and finite-element calculations improves, with $h_0/w = 0.1$ being much closer to the thin limit than $h_0/w = 0.2$.

In addition to giving thin-layer limit results such as we needed to complete Figure 3.11, the asymptotic method is also very suitable for obtaining information on how the bridge changes from a uniform thickness as sagging progresses. This is of considerable importance in optical component manufacture. Already we have (see equation (3.68)) that h_0^* and h_1^* do not vary in time, and the initial condition $h^* = 1$ then dictates that $h_0^* = 1$ and $h_1^* = 0$. Thus to leading order a thin bridge maintains its initial thickness profile as it sags small distances of order $\mathrm{O}(h_0)$. For any change to the thickness of the bridge we must look at the next order term h_2^*. Satisfying the dimensionless kinematic condition (3.56) on the free surfaces leads to

$$\begin{aligned}\bar{h}_{2t} \pm h_{2t} + (-(\mathcal{H} \pm 1)V_x + U)\,\mathcal{H}_x \\ = \frac{1}{2}(\mathcal{H} \pm 1)^2 V_x x - (\mathcal{H} \pm 1)U_x + v_{20}\end{aligned} \tag{3.127}$$

where, as before, asterisks are omitted on dimensionless variables. This then gives

$$\bar{h}_{2t} = \frac{1}{2}(\mathcal{H}^2 V_x + V_x)_x - (\mathcal{H}U)_x + v_{20} \qquad \text{and} \tag{3.128}$$

$$h_{2t} = (\mathcal{H}V_x)_x - U_x. \tag{3.129}$$

After substituting for U and V in terms of $\mathcal{H}$, (3.129) becomes

$$h_{2t} = \frac{1}{2}(\mathcal{H}_x^2)_t - \frac{1}{2}\dot{\overline{\mathcal{H}_x^2}} \tag{3.130}$$

which gives, to a first approximation, the variation in the bridge half-thickness over time. Integrating with respect to t and noting that $h_2 = 0$ at $t = 0$ then gives

$$h_2 = \frac{1}{2}\mathcal{H}_x^2 - \frac{1}{2}\overline{\mathcal{H}_x^2}\,. \tag{3.131}$$

This equation for h_2 is an 'outer' solution that is incorrect in a small boundary layer at the walls, of thickness comparable with the bridge thickness. This becomes apparent when it is observed that $\mathcal{H}_x = 0$ for all time at $x = \pm 1$ (see equation (3.114)) so that (3.131) gives $h_2(\pm 1, t) = -\frac{1}{2}\overline{\mathcal{H}_x^2}$ which is the strain in the bridge and is of necessity non-zero for all $t > 0$. This result at $x = \pm 1$ is inconsistent with a truly non-slip wall boundary which demands that $h_2(\pm 1, t) = 0$ for all time t. In fact, we have not had cause to enforce the zero vertical slip condition at the walls excepting at $y = 0$, and the method is permitting some vertical slip along the wall boundaries with the top and bottom free surface contact points sliding towards $y = 0$. Nevertheless, (3.131) provides useful information over most of the bridge width.

Switching again to dimensional quantities, Figure 3.15 shows the change in the bridge half-thickness $\Delta h/h_0 = (h - h_0)/h_0$ at $t\mathcal{V}/h_0 = 95.2$ when the bridge has sagged just a little more than its initial thickness of $2h_0$ (compare Figure 3.14). In the thin limit this is given to first order by $(h_0/w)^2(h_2/h_0)$. It is clear from this that over time the bridge becomes thinner at its centre and towards its edges, while increasing in thickness between $0.33 < x/w < 0.83$. It attains maximum thickness at about one fifth of its width from the wall. Bearing in mind that the thickness cannot in reality change at the walls, the bridge will be thinnest at its centre.

We can compare this result with our finite-element calculations for bridges of non-zero thickness, and obtain a better idea of thickness changes in close proximity to the walls. Already we have computed the centrelines of bridges of aspect ratios $h_0/w = 0.1$ and $h_0/w = 0.2$ after sagging one bridge thickness, as shown in Figure 3.14. It is then a very simple matter to compute the half-thickness $\hat{h} = h/w$ of these bridges at positions x/w across the bridge width at the time t when the sag is one bridge thickness, and these results can be re-scaled to correspond to the asymptotic output, i.e. $\Delta h/h_0 = \hat{h}/(h_0/w)$. This is shown for $h_0/w = 0.1$ and $h_0/w = 0.2$, together with the thin-limit result, in Figure 3.15. As in Figure 3.14, we see the thin-limit curve correctly predicting the general form of the thickness variation across the

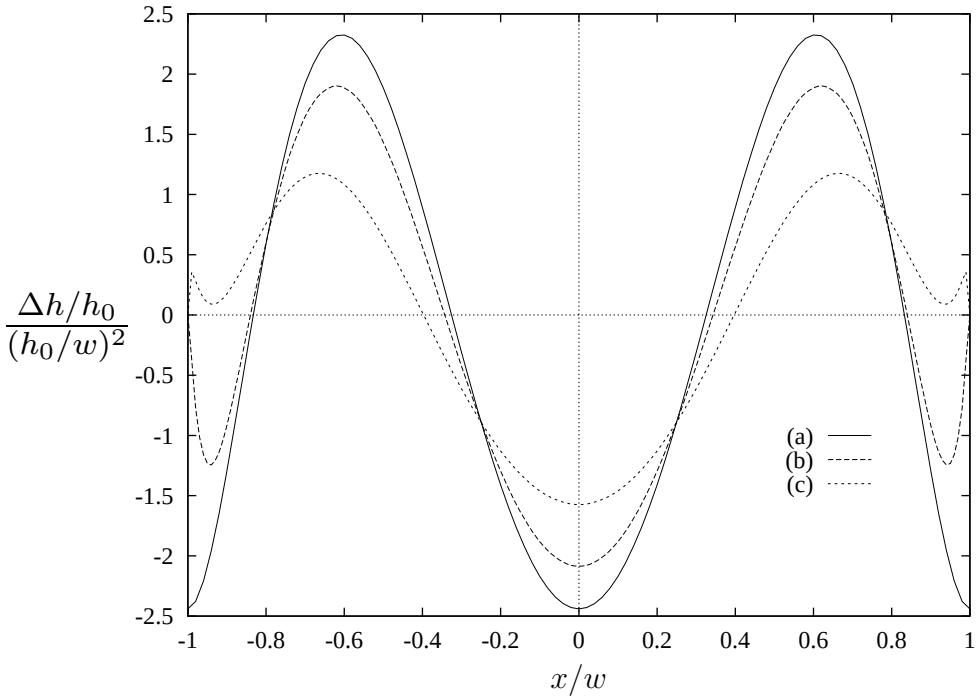

Figure 3.15: Change in bridge half-thickness after sagging one bridge thickness for (a) $h_0/w \to 0$ by asymptotic analysis, (b) $h_0/w = 0.1$ by finite-element methods and (c) $h_0/w = 0.2$ also by finite-element methods.

bridge width, with the agreement between the finite-element and thin-limit results improving as aspect ratio h_0/w decreases. We can also see, from the finite-element result for $h/w = 0.1$, that in the region near the walls, $\Delta h/h_0$ returns rapidly to zero rather than as given by the asymptotic method which allows slip tangential to the wall. However, the boundary region in which this happens is small and the thin-limit result is accurate over most of the bridge width, as we expected.

3.7 Final Remarks

In this chapter we have analysed a simple very viscous flow problem. The initial motion was computed by finite-element methods, using the commercial CFD package *Fastflo* and a purpose-written program, and by a semi-analytic series expansion method. Good agreement between all methods and programs has been demon-

strated. For finite-element analysis the purpose-written code was finally preferred as giving the best accuracy most efficiently, and this was the main tool for computing the motion at later times when the bridge geometry loses its initial rectangular shape and becomes more complex. However for bridges of small aspect ratio h/w a thin-limit asymptotic method is appropriate and yields some information not obtainable by numerical methods. In particular, the thin-limit asymptotic analysis enabled the completion of Figure 3.11 for the time taken to slump one bridge thickness as $h/w \to 0$, with a result that is consistent with, and indirectly validates, the finite-element computations, yet which could not be obtained, or even inferred, from the finite-element results. The thin-limit approximation was also found to be very useful in revealing the magnitude and nature of thickness variation across the width of the bridge, which could then be confirmed from the finite-element computations. This last aspect of the flow is important when applying this type of sagging flow to optical-component manufacture where the finished surface profile is critical.

Our treatment of this sagging flow has had considerable, though not total, emphasis on the validation of the purpose-written computer program. We shall further consider applications of this and similar sagging flows, including the manufacture of optical surfaces which is our main focus, in later chapters of this thesis. However, the finite-element program that has been developed is quite general and can be readily employed in solving some very different very viscous flows. In the next chapter we look at extensional flows such as exhibited by dripping honey.

Chapter 4

Dripping Honey

4.1 Introduction

Honey is a well-known fluid, and its high viscosity relative to that of water is obvious to any who have watched it drip from a spoon onto a slice of bread or into a cup of tea. Initially its motion is slow, but at an ever increasing rate the honey forms itself into a drop which then falls quite quickly under gravity. If the honey is sufficiently smooth and fluid, and perhaps depending on its specific composition, the drop becomes largely detached from the spoon, but remains connected to it by a thin and ever extending thread of honey (see [99, p. 224], [29, p. 988]). Figure 4.1 shows some frames from a simulation of this flow computed using the purpose-written finite-element program; the full MPEG movie is available on the Web, see [94].

The events so described and illustrated occur in a finite time. That time is essentially independent of how far the drop is allowed to fall, the final free-fall time being very short compared to the time during which viscous forces dominate while the drop is being formed. We shall refer to this as the *finite-time phenomenon*, and to the time at which fall occurs as the *crisis time*.

Dripping honey is just one commonly observed example of a class of motions

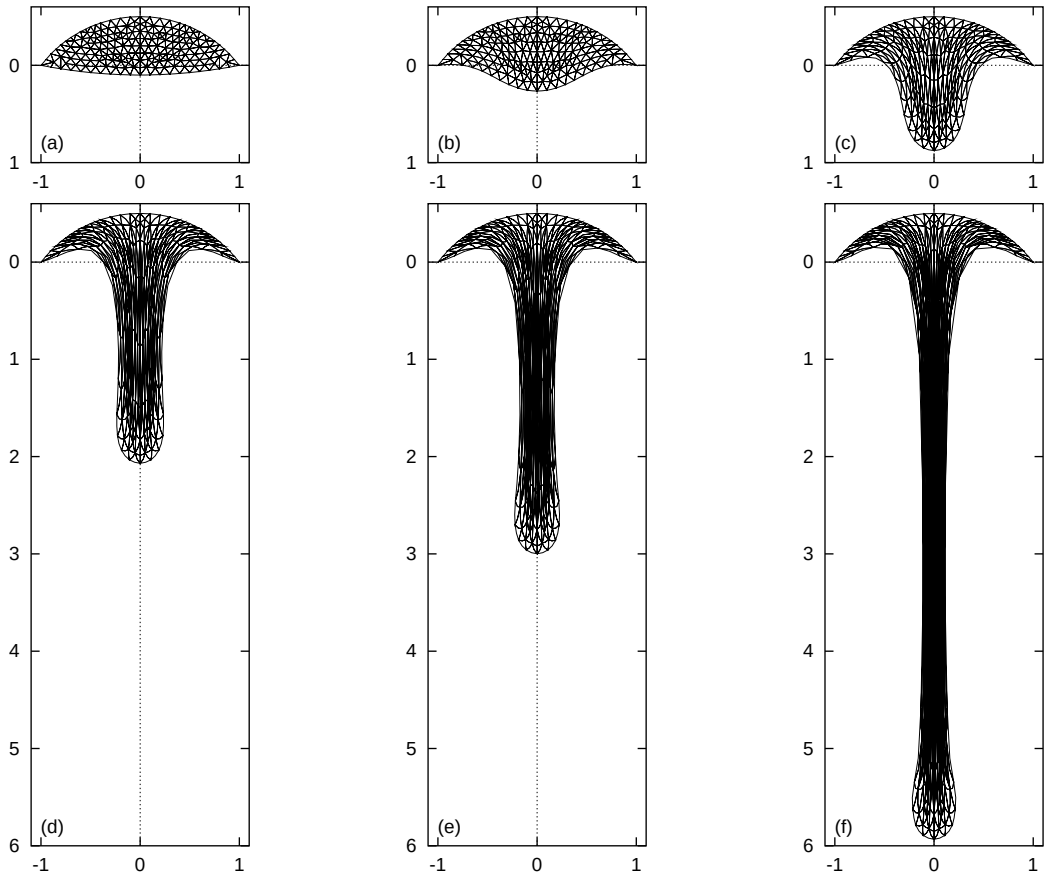

Figure 4.1: Sequence in the fall of an axisymmetric honey drop from a spherical-cap spoon. (a) $t/\mathcal{T} = 0$, (b) $t/\mathcal{T} = 4$, (c) $t/\mathcal{T} = 7$, (d) $t/\mathcal{T} = 8$, (e) $t/\mathcal{T} = 8.2$, (f) $t/\mathcal{T} = 8.4$.

which can be described as extensional flows [10]. Others include spinning and drawing of polymer or glass fibres for use in textiles, glass reinforced plastics, or optical fibres [28, 29, 30], and, on a similar note, web spinning by spiders and insects; continuous drawing of sheet glass [30]; glass-blowing and blow-moulding in the manufacture of containers, light bulbs, and glass tubing [30, 76]; rheological measurement by fibre extension and fibre spinning for polymers and glasses [58, 86]; and in geophysics there are examples in the areas of oil recovery (see [35]) and flows beneath the earth's crust that possibly lead to mountain formation and volcanic activity [16, 20, 47]. Not all exhibit the finite-time phenomenon described for the honey drop, and even in cases where it is very likely to occur, it is not always noted, perhaps because the

flow is modified in the end stages of the processes being considered.

There is a considerable amount of literature concerning the breaking of liquid jets/filaments to form drops. Very viscous through to inviscid liquids have been studied. Most of this work has focused on drop formation when the jet radius goes to zero in a finite time as a result of surface tension. See for example [31, 66, 67]. In these papers, extension under gravity is not an issue, which it certainly is for the case of dripping honey.

The gravity-driven slow dripping of a viscous fluid from a narrow vertical tube has been considered by Wilson [101], both with and without surface tension effects, with particular attention being given to the mass of the drops that form and break away. The flow is modelled by a one-dimensional slender-drop theory utilising conservation of mass, force balance on fluid particles, and the Trouton result [10, 92] relating stress and rate of strain for extensional flows. A finite-time phenomenon and crisis time is identified, similar to that for dripping honey, even with surface tension neglected, when the cross-sectional area of the drop at some point becomes zero so that it breaks. Because of the neglect of inertia, the model also indicates that the drop length becomes infinite at this crisis time. Wilson notes that this physical impossibility can be removed by including inertia in the equations; then the drop length will tend to infinity and the cross-sectional area will tend to zero as the time goes to infinity, with breaking of the drop being caused by mechanisms of instability. These mechanisms, however, become important only at times very close to the crisis time, and hence the model that gives this crisis time is suitable for determining drop volume and the time of rupture.

Gravity-driven extensional flow has also been considered in a geophysical context by Canright [16] and Houseman [47]. Analysis of the stability of a very viscous fluid layer overlying a less dense and much less viscous fluid shows that, when perturbed, vertical 'fingers' of the very viscous fluid grow and extend under gravity into the lower less dense fluid. With the neglect of inertia, these fingers become infinite in

length in a finite time. This type of flow may be a mechanism in the formation of mountains as 'drops' of the earth's lithosphere descend into the asthenosphere [47], and could conceivably cause other catastrophic geological events.

Fibre spinning is well discussed in the literature (e.g [28, 78]). Typically, in the manufacturing process, the molten fibre material leaving the spinnerette rapidly solidifies and is taken up on a winder. Because the fibre is pulled at a constant rate, and because of the solidifying of the fluid, the finite-time phenomenon cannot occur [29]. Schultz and Davis [78] also consider a fibre stretching under its own weight, and, although not discussed in that paper, this is an example where the finite-time phenomenon will occur provided solidification does not happen too quickly. The finite-time phenomenon for a one-dimensional viscous fibre stretching under gravity and/or an applied force is mentioned by Kaye [49].

Drawing of viscous sheets, as in the manufacture of sheet glass [30], is quite similar to fibre spinning, and has been modelled by Howell [48] using a slender approximation to the Stokes equations. Again, a constant drawing velocity and solidification preclude any finite-time phenomenon.

The "liquid bridge" [35] is a simplified extensional-flow model that is widely used for flows such as have been mentioned. A cylindrical column of fluid attached at its ends to two coaxial discs of equal radius is stretched by pulling the discs apart. When the discs are pulled apart with a constant velocity [35], or with a velocity that increases exponentially in time [86], the finite-time phenomenon will not occur, since it will take an infinite time for the liquid bridge to reach infinite length (ignoring other factors such as surface tension that cause it to break) [29]. If, on the other hand, stretching is caused by gravity, say by attaching a falling weight to one disc [58], then the finite-time phenomenon will occur. Much attention has been paid to the mechanisms causing breaking of a liquid bridge and the influence of surface tension, gravity and other factors [35]. As for the slowly-dripping fluid discussed previously, it is likely that, for the types of extensional flow we are considering where

the finite-time phenomenon is a feature, these factors become important mainly at or close to the crisis time.

A different extensional flow to those mentioned so far is pressure-driven flow of high-viscosity thin films such as is used in glass-blowing and blow-moulding [30, 70, 76]. Stretching of a molten viscous material, such as glass or polymer, results from applying a pressure difference across the thickness of the film to cause substantial motion in this transverse direction. Where there is no mould interference, the finite-time phenomenon is a feature of this flow, with the radius of the arc formed by the film becoming infinite in a finite time [96]. The sagging very viscous bridge analysed in Chapter 3, and revisited in an axisymmetric context in Chapter 5, is also a flow of this type, although there we do not consider large deformations such as are typical in blow forming, and which lead to significant extensional flow.

In this chapter, we solve for the time-dependent motion, shape, and length of a finite mass of Newtonian fluid of large constant viscosity that is initially at rest beneath and in contact with a solid boundary. A substantial portion of this is also presented in a paper by Stokes, Tuck, and Schwartz [89].

The most straightforward procedure for solving this problem is to use standard computational fluid-dynamic tools such as the finite-element method to solve the three-dimensional equations of motion, subject to no-slip boundary conditions on the solid boundary and zero-stress conditions on the free surface. In this way we can study drops of arbitrary initial shape. The finite-element method is the preferred method for the numerical modelling of a variety of forming processes [70]. It has been used to simulate liquid bridges [82], although the use there of a mesh with fixed axial length requires that the simulation be restricted to to low-rate and small-deformation experiments. In our dripping-honey example, once the drop begins to move it quickly elongates in the direction of gravity. The Lagrangian method of time advancing with a moving mesh (see Chapter 2) enables us to cope with these large elongations, and, in addition to publications arising from the present

research [93, 88], has been applied to the modelling of blow moulding of containers [14, 38, 100]. In a very similar manner to the present work, a creeping-flow model and finite-element method with a moving mesh has also been used in numerical analyses of the growth of 'fingers' in the earth's lithosphere discussed earlier [47].

However, because of the large elongations seen in the types of flow we are considering, a slenderness assumption is also possible, leading to an approximate one-dimensional theory, cf. [29, 49, 78]. This slenderness assumption may also be valid initially, and it is assumed in constructing the approximation that our interest in the flow begins at a time when this is so. The resulting one-dimensional flow can be described either by Lagrangian or Eulerian methods, and equivalence between these two viewpoints is demonstrated in this chapter. Comparison with the exact finite-element computations helps to clarify the regime where the slender-drop theory is valid.

As discussed in Chapter 2, solutions are obtained with neglect of inertia on the basis that the viscosity is large; the flow is thus creeping or Stokes flow and formally has accelerations that are very small compared to the acceleration of gravity. The most important theoretical conclusion from these solutions is that the drop length eventually increases rapidly, and becomes infinite at a finite crisis time. This was also noted by Wilson [101] and Canright [16], and is an inevitable and explicit property of the model in that work as well as the approximate theory for slender drops employed here. It is also confirmed by the exact finite-element computations. The latter computations must eventually fail when the computational grid becomes unreasonably stretched and the velocities very large. Nevertheless, when applied to drops that are initially sufficiently slender, they predict drop lengths that are increasing extremely rapidly near to finite crisis times and that are in close agreement with the approximate theory.

In the real world, drop lengths cannot become infinite, although the honey example shows that the overall length of the drop plus its connecting filament can indeed

become very large compared to the initial drop length. Meanwhile, however, the main falling head of the drop develops an acceleration comparable to that of gravity, as the remaining ever-thinning filament connection loses its ability to restrain the main drop's transition into free fall. In that phase of the motion, but only at times that are very close to the predicted crisis time, the assumption of neglect of inertia loses its validity. However, as discussed in [101], for fluids of large viscosity, it may be expected that the present results will provide good approximations until quite close to the crisis time, and in particular will provide a useful prediction of the actual value of the crisis or filament-breaking time, and of the break point or fraction of the initial drop that falls.

Although a simulation that applies the creeping-flow model developed in Chapter 2, has already been shown in Figure 4.1, let us pause to validate the assumptions behind the model before we progress any further in applying it. While there are many fluids to which our extensional flow model can be applied, it is the dripping of honey that motivates this chapter, and hence we consider honey as a typical very viscous fluid.

4.2 Properties of Honey

In general, studies have shown (uncrystallized) honeys to be true liquids that are Newtonian, though some notable exceptions include heather honey which is thixotropic (viscosity decreases with shearing e.g. brought about by stirring) and some eucalyptus honeys which are dilatant (viscosity increases with increasing rate of shear) [73]. Thus the Newtonian assumption that we have made in developing our viscous-flow model is generally appropriate for honeys, and is, in any case, a good place to start. Incompressibility is not specifically noted in the literature that has been reviewed, but this assumption seems quite reasonable for the types of flow under consideration since honey primarily consists of sugar with a water content of

some 10 to 20 percent.

Next we have to justify the exclusion of the inertial terms from the Navier-Stokes equation on the basis of the Reynolds number given in (2.6) being sufficiently small. To compute this, let us equate the characteristic length $\mathcal{L}$ with the initial length of the drop; then we require values for the density and viscosity of honeys.

The true specific gravity of honeys at 20°C is tabulated in [99, p. 218], and varies from 1.4457 at 13 percent water content to 1.3950 at 21 percent water content. This variation is quite small in the context of determining the Reynolds number applicable to the flow, and it is sufficient to take an average value of 1.42, i.e. a density of 1420kg/m^3. In accordance with our expectation of incompressibility, and in the absence of any information to the contrary, we shall assume that density does not vary significantly over a reasonable temperature range (say 5 to 50°C).

Viscosity, on the other hand, does vary considerably with both water content and temperature, as well as the content of protein and other constituents that depend on the floral source from which the honey is derived [73]. It is also an important property for honey, influencing its ease of extraction from the honey comb and subsequent processing, and, consequently, its determination has received considerable attention over the years. A thorough review of work carried out to 1953 is found in [73], and this is still considered to be 'state of the art' knowledge since subsequent work has not contributed anything of much significance to the field [99]. Hence the following discussion is largely based on information taken from that review.

In any general treatment of the rheology of honey, it is necessary that attention be limited to those components with which direct correlation can be made, of which water (or moisture) content seems to be the primary one. The other major parameter to which the viscosity of honey is highly sensitive, is temperature. Viscosity data is tabulated in [73, p. 157] for Newtonian honeys of different water content at different temperatures and some of this data is reproduced here in Table 4.1.

Now we require that the Reynolds number given by (2.6) be much less than one.

Table 4.1: Relationship between temperature, water content and viscosity of honey.

14.2% H_2O White Clover		16.1% H_2O Sweet Clover		18.6% H_2O Sage		24.0% H_2O White Clover	
°C	Poise	°C	Poise	°C	Poise	°C	Poise
21.7	417.6	15.7	422.4	6.8	400.8		
29.5	151.2	22.0	156.0	14.0	148.0		
32.0	112.8	24.1	121.2	16.2	105.6	3.4	108.8
49.7	14.7	43.3	16.3	36.2	15.6	18.2	15.6

If we take an initial drop length of the order of one centimetre, and a viscosity of the order of 100 poise = 10 Pa · s then the Reynolds number is about 0.2 which we accept as sufficiently small. From Table 4.1 honeys of 14.2 percent water content up to a temperature of 32°C, or of 16.1 percent water content up to 24°C, or of 18.6 percent water content up to 16°C, can be satisfactorily modelled using our creeping-flow model. Honeys of higher water content are probably not sufficiently viscous in normal temperature ranges.

Note that the Reynolds number increases like the cube of the drop length. Thus as the drop elongates, the Reynolds number increases quite rapidly (along with the flow velocities) until the neglect of inertia is no longer valid. However, as we have already discussed, this occurs only in the final stages of the flow, just prior to breaking.

One further model assumption that needs to be justified is the neglect of surface tension. The only surface tension data that we have come across is for 25 percent honey solutions, and this indicates a surface tension from 0.047 to 0.06 N/m at 20°C, with the value increasing for filtered solutions [99]. Since the value for pure honey will be larger than for honey solutions, it is safe to adopt a surface tension of $\gamma = 0.06$N/m. As before we take a density of 1420kg/m^3 and a drop length of one centimetre. The capillary number given by (2.14) is then about 23 so that the surface tension terms on the right hand side of (2.13) are of size 0.04 compared to

terms of order one on the left hand side. Thus our neglect of surface tension is reasonable.

This brief review of the properties of honeys indicates that the creeping-flow model developed in Chapter 2 is suitable for modelling dripping honey up until times quite close to the crisis time when the drop length increases rapidly and neglect of inertia is no longer valid.

4.3 Mathematical Formulation

Let the x-axis be directed vertically downwards, and suppose that at any time the fluid mass lies in the region $0 < x < L(t), |y| < f(x, z, t)$ and is an incompressible Newtonian viscous fluid of density ρ and viscosity μ, in a gravity field g acting in the x direction. Given suitable initial conditions $L(0)$ and $f(x, z, 0)$ defining the initial flow domain geometry, our task is to solve (2.2) and (2.7) (with the x direction being equivalent to the $-x_2$ direction of Section 2.1), subject to no-slip boundary conditions (2.11) at the wall, which is assumed to lie in or close to the plane $x = 0$, and no-stress free-boundary conditions (2.15) at the side boundary $y = f(x, z, t)$. We must also satisfy a kinematic boundary condition (2.18) on the side boundary, or equivalently the Lagrangian equation (2.19). If the drop has a rectangularly cut-off bottom, i.e. if $f(x, z, t)$ is not zero at $x = L(t)$, similar free-boundary conditions must also be satisfied on the bottom surface. The flow domains that we shall consider are symmetrical about the x-axis, so that we may solve over one half of the flow domain and satisfy (2.15) and (2.16) on this symmetry boundary.

We shall solve for this flow numerically in both two and three dimensions by finite-element methods using the purpose-written computer program, but first let us consider an approximate one-dimensional solution.

4.4 Slender-drop Theory

It is possible to provide an approximate solution of the creeping-flow equations for slender drops, i.e. those drops such that the free-surface slope f_x in the vertical direction is everywhere small. The net effect of such an approximation will be that the flow is determined largely one-dimensionally, the key variables being a measure $u(x,t)$ (averaged over a cross-section at fixed x) of the downward vertical velocity, and the net cross-sectional area

$$A(x,t) = 2\int f(x,z,t)dz \tag{4.1}$$

of the drop at station x.

Derivation of appropriate equations to determine u and A can be done in one of two ways. The simplest way is to convert to a Lagrangian specification $x = X(\xi,t)$, in which $u = X_t$. The label variable ξ is such that $x = \xi$ at $t = 0$. The area A is then to be determined as a function of ξ and t, subject to an initial value $A = A_0(\xi)$ at $t = 0$.

Consider a small piece of the drop between ξ and $\xi+\Delta\xi$, deforming to lie between x and $x+\Delta x$, where $\Delta x = X_\xi \Delta\xi$. Mass conservation implies that $A\Delta x = A_0\Delta\xi$ is time-independent, or $AX_\xi = A_0$. Hence the rate of strain of this piece is

$$E = \frac{1}{\Delta x}\frac{\partial}{\partial t}\Delta x = -\frac{1}{A}\frac{\partial}{\partial t}A. \tag{4.2}$$

The stress acting on this piece due to gravity is $S = W/A$ where

$$W = W_0 = \rho g\int_\xi^{L_0} A_0(\xi_1)d\xi_1 \tag{4.3}$$

is the (time-independent) weight of fluid beneath station ξ. We can now invoke a one-dimensional constitutive equation of the form

$$S = \hat{\mu}E \tag{4.4}$$

where $\hat{\mu}$ is the so-called "Trouton viscosity" ([92], [10, p. 30]), which is related to the actual viscosity μ by $\hat{\mu} = 3\mu$ in three dimensions, and $\hat{\mu} = 4\mu$ in two dimensions. Thus $W_0 = -\hat{\mu} A_t$, which integrates with respect to time to give

$$A(\xi, t) = A_0(\xi) - \frac{\rho g t}{\hat{\mu}} \int_{\xi}^{L_0} A_0(\xi_1) d\xi_1 \,. \tag{4.5}$$

Meanwhile the mass conservation result integrates to give

$$X(\xi, t) = \int_0^{\xi} \frac{A_0(\xi_1)}{A(\xi_1, t)} d\xi_1 \,, \tag{4.6}$$

and a combination of the last two equations yields, via two quadratures, all required information about the flow, given any specification $A_0(\xi)$ of the initial shape. In particular, if required, the velocity u can then be obtained by time differentiation, and the length $L(t)$ of the drop at any time is just given by $L(t) = X(L_0, t)$.

This type of Lagrangian approach is also employed in [49] for similar types of problem. In particular, the extension of a thin fibre of initially constant cross-section A_0 held vertically and fixed at the top, due to gravity and an applied force F_0, is considered. With $F_0 = 0$ the solution obtained [49, p. 67, equation (4.1.17)] is identical to that given by (4.5) and (4.6) above. This particular case will be considered further below.

An alternative derivation of these results proceeds via a formal asymptotic expansion with respect to a small parameter $\epsilon = w_0/L_0$, where $w_0/2$ is the maximum value of $f(x, z, t)$. The details are essentially given later in another context with gravity g acting in the $-x$ direction (see Section 5.3), and hence are omitted here. Eventually the procedure leads to Eulerian equations for the one-dimensional dependent variables $u(x, t)$ and $A(x, t)$, namely

$$\frac{\partial}{\partial x}\left[A \frac{\partial u}{\partial x}\right] + \frac{\rho g}{\hat{\mu}} A = 0 \tag{4.7}$$

involving the Trouton viscosity $\hat{\mu}$ again, and

$$\frac{\partial A}{\partial t} + \frac{\partial}{\partial x}\left[uA\right] = 0, \tag{4.8}$$

cf. equations (5.17) and (5.18). Equation (4.7) has previously been derived (e.g. [29, 48]) in the absence of gravity, when its solution is immediate. With the gravity term present, the best procedure to solve (4.7) coupled with (4.8) is in fact to revert to the Lagrangian representation.

Recognising (4.8) as a one-dimensional continuity equation, it is not difficult to make the appropriate changes of variable to retrieve the Lagrangian-representation solutions above. Integrating (4.7) with respect to x and satisfying the boundary condition $u_x = 0$ at $x = L(t)$, gives $-Au_x$ in terms of the weight W, and this quantity is seen to be equal to the time derivative of A following a fluid particle (i.e. the material time derivative) on writing (4.8)

$$\frac{dA}{dt} = \frac{\partial A}{\partial t} + u\frac{\partial A}{\partial x} = -A\frac{\partial u}{\partial x}\,. \tag{4.9}$$

Hence (4.5) follows by time integration.

The boundary condition $u_x = 0$ at $x = L(t)$ is equivalent to a condition of zero stress as will be seen from the formal asymptotic Eulerian expansion of Chapter 5 (see p. 118). However this boundary condition is implied by (4.9) together with the fact, already known from (4.5), that the cross-sectional area at $x = L(t)$ is constant and possibly non-zero for all time .

The general case of the above approximate solution (4.5) and (4.6) will be discussed later, but for now let us simply observe the simple special case in which $A_0(x)$ =constant. This applies both in two dimensions to an initially rectangular slab, and in three dimensions to an initially cylindrical drop of a general (constant) cross-section. Then we find a length given by

$$\frac{L(t)}{L_0} = -\frac{\hat{t}}{t}\log\left(1 - \frac{t}{\hat{t}}\right) \tag{4.10}$$

where

$$\hat{t} = \frac{\hat{\mu}}{\rho g L_0}. \tag{4.11}$$

Thus it is explicit that the solution "blows up", with $L(t) \to \infty$ at the finite crisis time $t = \hat{t}$. It is also clear by substitution into (4.5) that the area at $\xi = 0$ goes to

zero at this crisis time $\hat{t}$ which can be equated with the fluid breaking away from the wall, cf. [101].

4.5 Finite Element Procedure and Convergence Tests

The problem as formulated in Section 4.3 is suitable for immediate solution using the finite-element method. With the purpose-written computer program we may solve two-dimensional and axisymmetric cases, although in principle the general three-dimensional equations can be solved as easily. The solution procedure does not differ from that already described previously, excepting in two respects. First, for a fixed time step, node displacements become large in the later stages of the flow as velocities increase, and there is a loss of solution accuracy readily identified by non-conservation of flow-domain area (i.e. mass). Because of this the displacement of any node is limited to some maximum value as described below. Second, the large extensions of the flow-domain at times just prior to the crisis time cause considerable distortion of mesh elements resulting in a loss of solution accuracy. When this occurs it is necessary to stop computing and remesh before proceeding further. Eventually a time is reached when there is an extremely rapid distortion of the mesh, a time beyond which computations cannot be continued even with remeshing. This time will be identified with the "crisis time" predicted by the slender-drop theory.

To test convergence of the numerical methods, let us first take an initially rectangular two-dimensional slab with an initial width/length ratio of $w_0/L_0 = 0.2$, attached to a plane no-slip wall $x = 0$. We use a non-dimensional formulation, with a length scale of $\mathcal{L} = L_0$ and a time scale of $\mathcal{T} = \hat{\mu}/(\rho g L_0)$. This means that the slender-drop approximation to the crisis time for this particular initial drop shape is $\hat{t}/\mathcal{T} = 1$. We exploit the lateral symmetry of the problem to reduce computations to half of the rectangular domain, and test convergence with respect to both mesh

Table 4.2: Initial rate of extension of a hanging rectangular slab.

N_e	$L' \cdot (\hat{t}/L_0)$
66	0.029911
146	0.029923
235	0.029926
315	0.029925
409	0.029929
501	0.029929

and time-step sizes.

The local corner flow analysis conducted in Chapter 3 for the sagging viscous bridge is applicable to this problem also, and again there is a mild pressure singularity in the corner where the free surface attaches to the top wall. Consequently non-uniform meshes with elements concentrated at the contact corners give faster convergence than meshes of uniformly distributed elements. They are also preferred because mesh elements are clustered in the area which will see most distortion as time progresses. Convergence with increasing numbers N_e of mesh elements is therefore demonstrated for non-uniform meshes. As a first test, we evaluate the impulsively developed velocities at the initial instant of time $t = 0^+$, and examine convergence of the computed value of the maximum velocity, which occurs at $(x/L_0, y/L_0) = (x^*, y^*) = (0, 1)$ and is the initial rate of extension $L'(0^+)$, the prime denoting time differentiation. The results are given in Table 4.2, from which it can be concluded that a mesh of about 150 elements (in one half of the domain) is suitable for this geometry, and gives about three figures of accuracy.

With respect to convergence with decreasing time step, let us consider both the drop length $L(t)$ and rate of extension $L'(t)$. The Runge-Kutta method has been found to give excellent accuracy at times not too close to the crisis time, even for quite large time steps. Accuracy decreases as the crisis time is approached and mesh movements become large, and is clearly identified by non-conservation of flow-

domain area (equivalent to non-conservation of mass). Because of this, it is a good idea to prescribe a maximum allowable (dimensionless) node displacement $|\Delta\tilde{r}|$ in a single scaled time step Δt^*, and require that the displacements of all nodes i predicted from the known dimensionless velocities $(u^*, v^*)_i$ using the Euler method are no larger than this maximum, i.e.

$$|\Delta\tilde{r}| \geq |(u^*, v^*)_i| \cdot \Delta t^* \quad \text{for all } i \tag{4.12}$$

where the modulus denotes the usual length of a vector. If the maximum displacement over all nodes occurs at node j and exceeds the maximum allowed displacement, then the time step is reduced using

$$\Delta t^* = \frac{|\Delta\tilde{r}|}{|(u^*, v^*)_j|}. \tag{4.13}$$

Because the Euler method is used to predict node displacements and compute the new size of the time-step, while the Runge-Kutta method is used to calculate final node displacements, the actual maximum node displacement may still exceed the nominated allowed value. However the differences are quite small and of no concern. Imposing this restriction on node displacements has the effect of clustering time steps near the crisis time, when there is most action.

Table 4.3 gives some results at $t/\hat{t} = 1$, which is very near the crisis time for our test geometry, for different values of maximum time step $\triangle t^*$ and maximum displacement $|\triangle\tilde{r}|$. The actual number of time steps N_t is also given. A non-uniform mesh of 146 elements was used. In fact, $t/\hat{t} = 1$ is the non-dimensional crisis time in the slender-drop limit, and the effect of thickness is just to increase this time slightly. Considering the extreme proximity of crisis, where numerical inaccuracy is to be expected, these results are extremely good. There is good conservation of mass (dimensionless area) with six figure accuracy for $\triangle t^* = 0.05$ and $|\triangle\tilde{r}| = 0.10$. The corresponding values for L and L' show convergence to seven and five figures respectively.

Table 4.3: Length and rate of extension of an initially rectangular hanging slab at $t/\hat{t} = 1$.

$\triangle t^*$	$\lvert\triangle\tilde{r}\rvert$	N_t	L/L_0	$L' \cdot (\hat{t}/L_0)$	Area
0.25	0.50	5	3.019435	0.772955	0.2001449
0.05	0.50	20	3.017544	0.773823	0.2000013
0.25	0.10	20	3.017533	0.773826	0.2000010
0.05	0.10	27	3.017522	0.773827	0.2000001
0.25	0.05	41	3.017522	0.773829	0.2000000
0.05	0.05	43	3.017522	0.773829	0.2000000

Obviously, we are unable to compute at exactly the crisis time, and can only hope to approach it as closely as possible. Not unexpectedly, as we approach closer to this crisis time, some of the mesh elements become very distorted due to the large elongation and pinching-in near the wall, resulting in numerical error. Remeshing the flow domain before the elements become too distorted and too much numerical error results, enables us to compute closer to the crisis time, until again the mesh elements become excessively distorted. Eventually a point is reached where no further computation is possible, even with remeshing. To obtain an estimate for the crisis time itself, we can plot L^{-2} versus time t, and extrapolate to $L^{-2} = 0$. In Figure 4.2 such a plot is given for the end stages of the flow of a rectangular slab of aspect ratio $w_0/L_0 = 0.2$. Curves obtained with and without remeshing are given to show the loss of accuracy with excessive mesh distortion. All further plots are with remeshing close to the crisis time.

It is instructive to further treat the case of a rectangular initial drop shape, but now with varying aspect ratio (w_0/L_0). As already indicated, there are no extra computational difficulties associated with the axisymmetric equivalent of the two-dimensional problem, namely an initially cylindrical drop, and the results are qualitatively similar. Furthermore, as previously indicated (Section 4.4), in the slender-drop limit the two and three-dimensional crisis times differ only by a factor of three quarters — the ratio of Trouton viscosities. Since the finite-element compu-

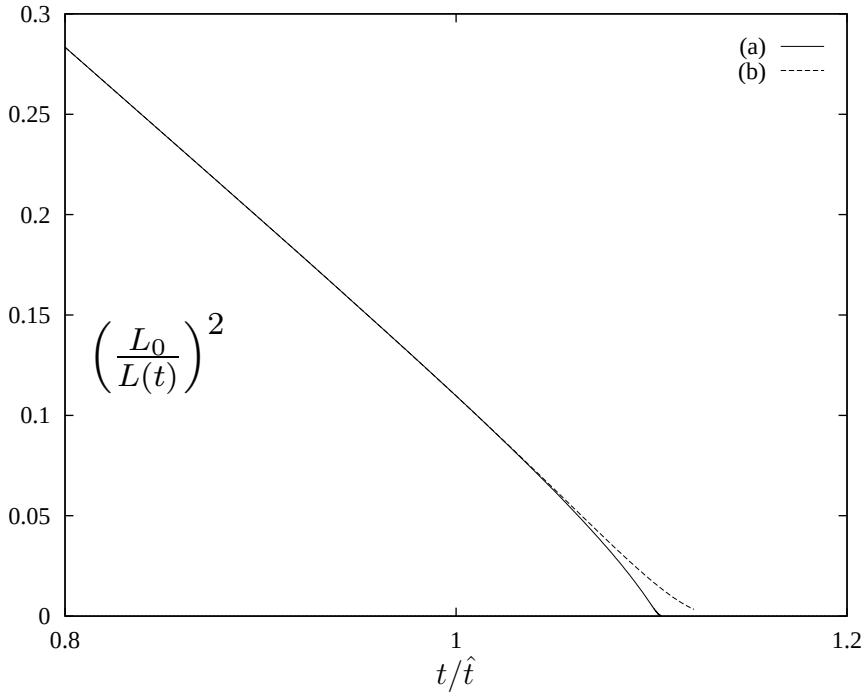

Figure 4.2: Length versus time for a rectangular slab of $w_0/L_0 = 0.2$ (a) with remeshing at $t/\hat{t} = 0.95$, and (b) without remeshing.

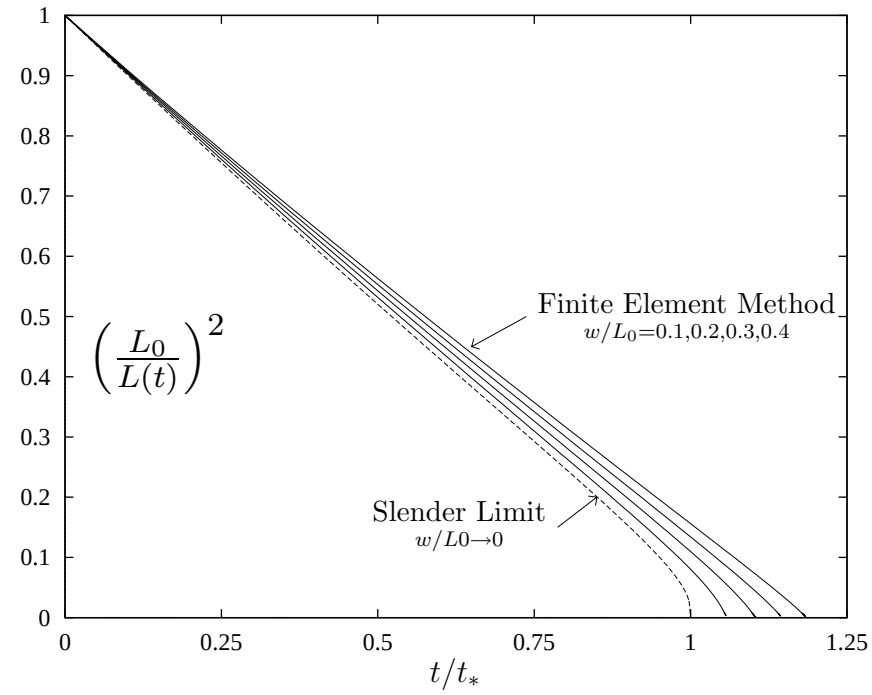

Figure 4.3: Length versus time for rectangular slabs of various aspect ratios.

tations are consistently non-dimensionalized with respect to a time scale involving the Trouton viscosity, the non-dimensional slender-body estimate of the crisis time remains at $\hat{t}/\mathcal{T} = 1$ in both three and two dimensions.

We first consider initially-rectangular two-dimensional slabs with aspect ratios between 0.1 and 0.4, and with meshes of between 155 and 175 elements over the half-domain. Figure 4.3 gives L^{-2} versus t for the different aspect ratios considered. For comparison, the curve obtained from the slender-drop theory is also given. An advantage of plotting L^{-2} is that the small-time behaviour predicted by the slender-drop theory is a straight line passing exactly through the crisis point; the actual curve is then just a small upward deviation from this line. The slender-drop theory is the small-aspect-ratio limit, and the agreement between the slender-drop and finite-element computations is good for the lower aspect ratios, as expected.

Figures 4.4 and 4.5 show in parts (a) the evolution of the shape of the drop of initial aspect ratio 0.2, as given by the finite-element method, at times $t/\hat{t} = 0.75$ and $t/\hat{t} = 0.95$ respectively. The slender-drop approximate profiles are also given for comparison in parts (b) of these figures, the (uncorrected) profiles being the lower of the two curves in each case. There are two types of apparent difference between the exact finite-element computations and the slender-drop theory, namely an upward shift of the profile and a modified profile near the wall. The main effect of finite width is related to what happens at the wall, which then affects the time scale, so that the actual drops fall slightly slower than the slender-drop theory predicts. Otherwise, the exact (finite-difference) and approximate (slender-drop) profile shapes are very close if corrected to compensate for this small time difference, except for the region near the wall, which is now discussed.

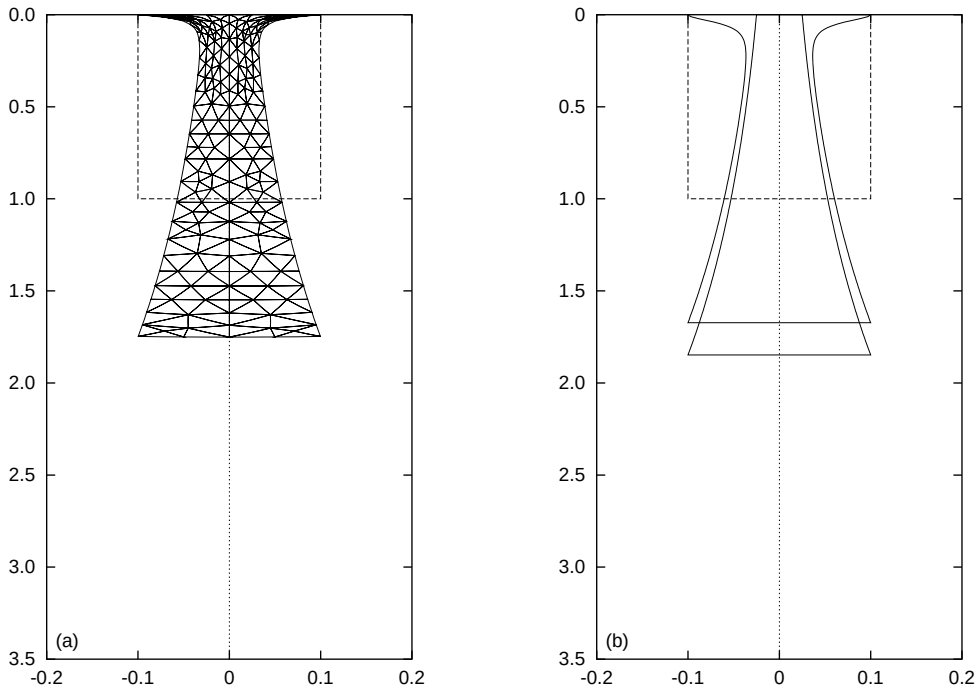

Figure 4.4: Initially-rectangular profile, at $t/\hat{t} = 0.75$. (a) Finite-element method. (b) Slender-drop theory, with and without wall correction.

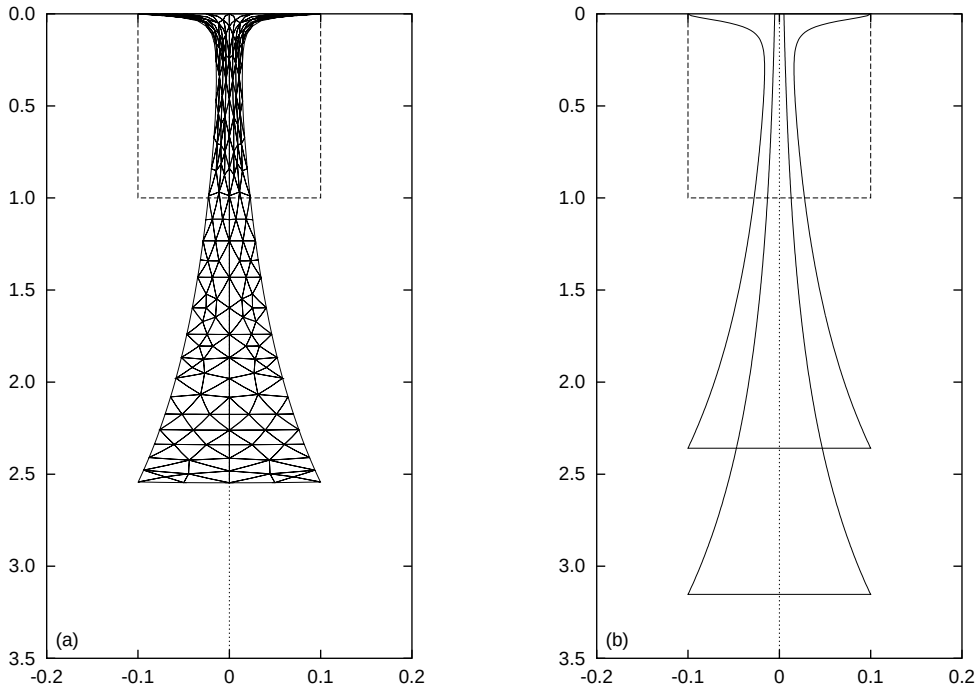

Figure 4.5: Initially-rectangular profile, at $t/\hat{t} = 0.95$. (a) Finite-element method. (b) Slender-drop theory, with and without wall correction.

4.6 Wall Corrections

Because of the second-order character of the partial differential equation (4.7) given by the slender-drop theory only two boundary conditions have been needed for its solution; we required that there be no normal velocity component at the wall, i.e. $u = 0$ at $x = 0$, and a zero velocity gradient at the free end of the drop corresponding to a condition of zero normal stress at the bottom, i.e. $u_x = 0$ at $x = L(t)$. Thus the slender-drop theory permits slip along the wall boundary, which the finite-element calculations do not, as is clear from Figures 4.4 and 4.5. The slender-drop solution is, in fact, an outer solution that is valid to within about one drop width of the ends, where inner solutions are required to give boundary-layer corrections [78]. These inner solutions are given by the higher-order terms in the slender-drop expansion and require extensive numerical analysis (ibid.), so that in one sense the finite-element method is the better way to obtaining the near-wall profile.

It is, however, possible to give an empirical discussion of the inner flow near the wall boundary based on the "squeeze bearing" problem [15, p. 392] of lubrication theory. The details are given in [89], and finally yield a composite empirical expression for the Trouton viscosity, which in two-dimensions is

$$\hat{\mu} = 4\mu \left[1 + \frac{1}{16} \left(\frac{w_0}{x} \right)^2 \right] , \tag{4.14}$$

and for the axisymmetric problem is

$$\hat{\mu} = 3\mu \left[1 + \frac{1}{32} \left(\frac{w_0}{x} \right)^2 \right] . \tag{4.15}$$

These expressions agree with the Trouton result, $\hat{\mu} = 4\mu$ in two dimensions and $\hat{\mu} = 3\mu$ in three dimensions, for large x. They also have the property that $\hat{\mu} \to \infty$ as $x \to 0$, thereby enforcing the pinning of the contact points at $y = \pm w_0/2$ by solidifying the fluid there.

For the two-dimensional initially rectangular drops of length L_0 and width w_0 already considered, the shape evolution $x = X(\xi, t)$ given by the slender-drop theory

employing the above expression (4.14) for the viscosity is the solution of the linear first-order ordinary differential equation

$$\frac{d\xi}{dx} = 1 - \frac{\rho g L_0 t}{4\mu} \quad \frac{1 - \xi/L_0}{1 + w_0^2/(16x^2)} \tag{4.16}$$

which may be integrated in closed form. As an example we show, as the upper curves of parts (b) of Figures 4.4 and 4.5, the corrected shape of the drop with $w_0/L_0 = 0.2$ at times $t/\hat{t} = 0.75$ and $t/\hat{t} = 0.95$. These corrections were obtained by integrating (4.16) and substituting the result for $\xi = \xi(x)$ in (4.5). Although there is some indication that this present theory may slightly over-correct, when compared with the (essentially exact) finite-element solution shown in parts (a) of these figures a significantly improved agreement is clearly achieved by inclusion of this wall correction in the slender-drop theory.

4.7 Break Points

According to the slender-drop theory for a general initial drop shape $A_0(\xi)$, the crisis time occurs when the drop's cross-sectional area $A(\xi, t)$ as given by (4.5) vanishes at some station ξ. Since both $A_0(\xi)$ and its integral

$$V(\xi) = \int_{\xi}^{L_0} A_0(\xi_1) d\xi_1 \tag{4.17}$$

(the volume lying below station ξ) are positive quantities, this must always happen at some sufficiently large time t. From (4.5), it will happen first at the station $\xi = \hat{\xi}$ where $V(\xi)/A_0(\xi)$ takes its maximum value $\hat{L}$, and then the crisis time will be

$$\hat{t} = \frac{\hat{\mu}}{\rho g \hat{L}}. \tag{4.18}$$

The drop will then break at that station $\hat{\xi}$, with the portion $0 < \xi < \hat{\xi}$ of the original drop remaining attached to the wall, and the portion $\hat{\xi} < \xi < L_0$ going into free fall for $t > \hat{t}$. Whenever the crisis time is given by (4.18), it is convenient to refer to the quantity $\hat{L}$ as the "effective length" of the drop.

The break point, effective length and crisis time given in this simple manner by the slender-drop theory, will be correct even though we ignore the wall correction, provided only that the break point does not fall within the near-wall region which has a thickness of the same order as the drop width. Should this occur, then the breaking behaviour will vary somewhat from this prediction.

In many cases, the quantity $V(\xi)/A_0(\xi)$ is monotone decreasing, with its maximum at the wall $\xi = 0$. Hence in such cases, the slender-drop theory predicts breakage at the wall (cf. [49]), with the whole initial mass going into free fall. For example, the family

$$A_0(\xi) = A_0(0)\left[1 - \frac{\xi}{L_0}\right]^n \tag{4.19}$$

for any power $n > -1$, has

$$\frac{V(\xi)}{A_0(\xi)} = \frac{L_0 - \xi}{n+1}. \tag{4.20}$$

Hence $\hat{\xi} = 0$ and the effective length is $\hat{L} = L_0/(n+1)$. Note that $n = 0$ is the rectangular or cylindrical case considered earlier where the effective length is equal to the actual initial length; $n = 1$ for the axisymmetric case is a paraboloidal initial drop of a quite sensible form. Figure 4.6 shows profiles for this paraboloidal initial drop. These profiles are at a scaled time interval of 0.05 where $\mathcal{T} = \hat{\mu}/(\rho g L_0)$, the last profile shown being at $t/\mathcal{T} = 1.95$, compared to a crisis time of $\hat{t}/\mathcal{T} = 2$. Note that, in this case, we no longer have $\mathcal{T} = \hat{t}$ as in the initially rectangular or cylindrical cases because $\hat{L} \neq L_0$.

When $\hat{\xi} = 0$, the break point is formally at the wall itself according to the slender-drop theory, and hence it lies inside the near-wall region. Therefore, the finite-element and slender-drop solutions will differ, as we have already seen for rectangular slabs. Actual breaking occurs at a point which has an initial distance from the wall comparable with the thickness of the near-wall region, with some material in that region being left behind. However, this distance is formally of the same size as the drop width, and hence small relative to the drop length, with

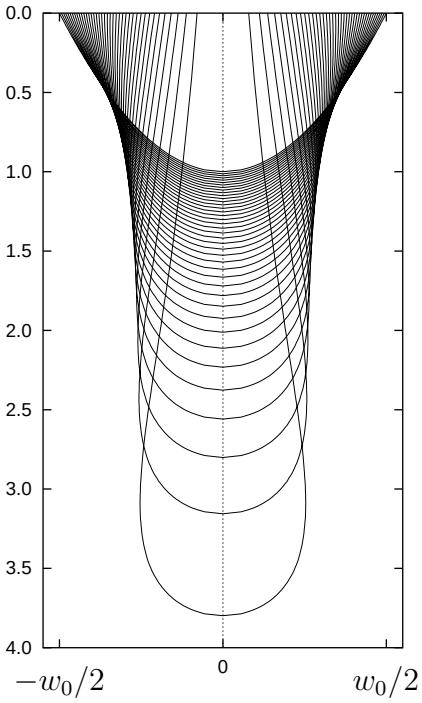

Figure 4.6: Extension of an initially-paraboloidal drop ($\sigma = 0$).

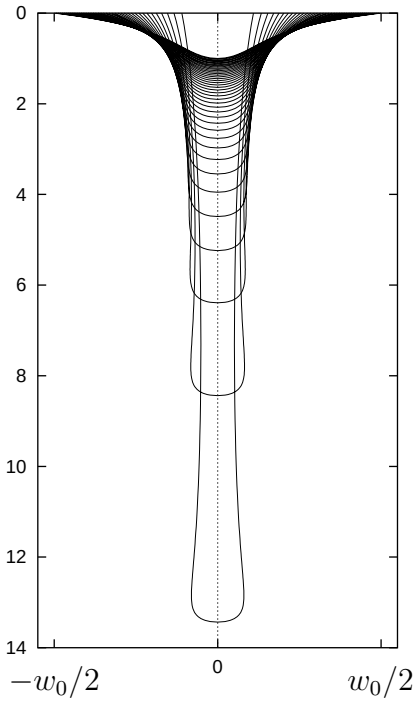

Figure 4.7: Extension of drop with $\sigma = 1.34016$, such that half of the volume falls.

the volume left behind being of the order of ϵ^2, which is small compared to the total order-ϵ volume of the drop. Hence formally we can consider that the drop does indeed break at the wall as indicated by the slender-drop theory, although the actual crisis time is a little longer. The finite-element computations confirm this description.

On the other hand, it is possible for the predicted break point to occur at a non-trivial station $\hat{\xi} > 0$, so long as $V(\xi)/A_0(\xi)$ has an interior maximum. For example, consider an axisymmetric family of drops with initial area profiles given by

$$A_0(\xi) = A_0(0)\left(1 - \frac{\xi}{L_0}\right)\exp(Y^2 - \sigma^2) \tag{4.21}$$

where $Y = \sigma(1 - \xi/L_0)$ and σ is a given parameter. If $\sigma = 0$, we retrieve the paraboloidal drop $n = 1$ discussed above. We find

$$\frac{V(\xi)}{A_0(\xi)} = \left(\frac{L_0}{2\sigma}\right)\frac{1 - e^{-Y^2}}{Y} \tag{4.22}$$

which has a maximum of magnitude $\hat{L} = 0.3191 L_0/\sigma$ at $Y = 1.1209$. Hence if $\sigma > 1.1209$ and the drop is sufficiently slender, breakage occurs at $\xi/L_0 = 1 - 1.1209/\sigma$, with the above effective length $\hat{L}$ and hence a crisis time $\hat{t}$ given by (4.18). For example, if $\sigma = 1.34016$, the break point $\xi/L_0 = 0.1636$ is such that exactly half of the volume of the original drop falls, with $\hat{L} = 0.2381 L_0$. Figure 4.7 shows profiles for this case, at a scaled time interval of 0.1, the last profile shown being at a time $t/\mathcal{T} = 4.1$ close to the crisis time $\hat{t}/\mathcal{T} = 4.200$. These results were computed from the slender-drop theory without any correction for no-slip on the wall.

The same initial drop was also used as input to the finite-element code. A near-crisis profile at $t = 4.0$ is shown in Figure 4.8(a), where a drop initial half-width of $0.05 L_0$ has been chosen such that the break point falls outside of the near-wall region. The resemblance of the near-crisis profile to the slender-drop solution at $t = 4.0$ in Figure 4.7 (the corresponding profile being the "next-to-end" one in that Figure) is very good; note the agreement for the length $L/L_0 \approx 8.3$ which shows that the break point and crisis time are independent of the no-slip wall condition. To emphasize this feature, an artificial "slip" wall condition has also been implemented in the finite-element program, yielding the profile at $t = 4.0$ shown in Figure 4.8(b). The profiles for the slip and no-slip walls are remarkably similar to each other at $t = 4.0$, and this similarity continues for times closer to the crisis time $t \approx 4.2$.

As the drop half-width is increased from $0.05 L_0$, so increasing the vertical extent of the near-wall region, a point is reached where the predicted break point falls within this region. When this happens, the wall boundary conditions do affect the final outcome, and the finite-element solution cannot be expected to compare as well with the slender-drop theory, for either crisis time or break point. The further the break point falls inside the near-wall region, the greater the difference between the two results. For a drop half-width of $0.1 L_0$, the break point is only just within this region, and the solution at $t = 4.0$ is quite similar to those in Figure 4.8. For times closer to crisis, however, a marked difference between a slip and no-slip wall

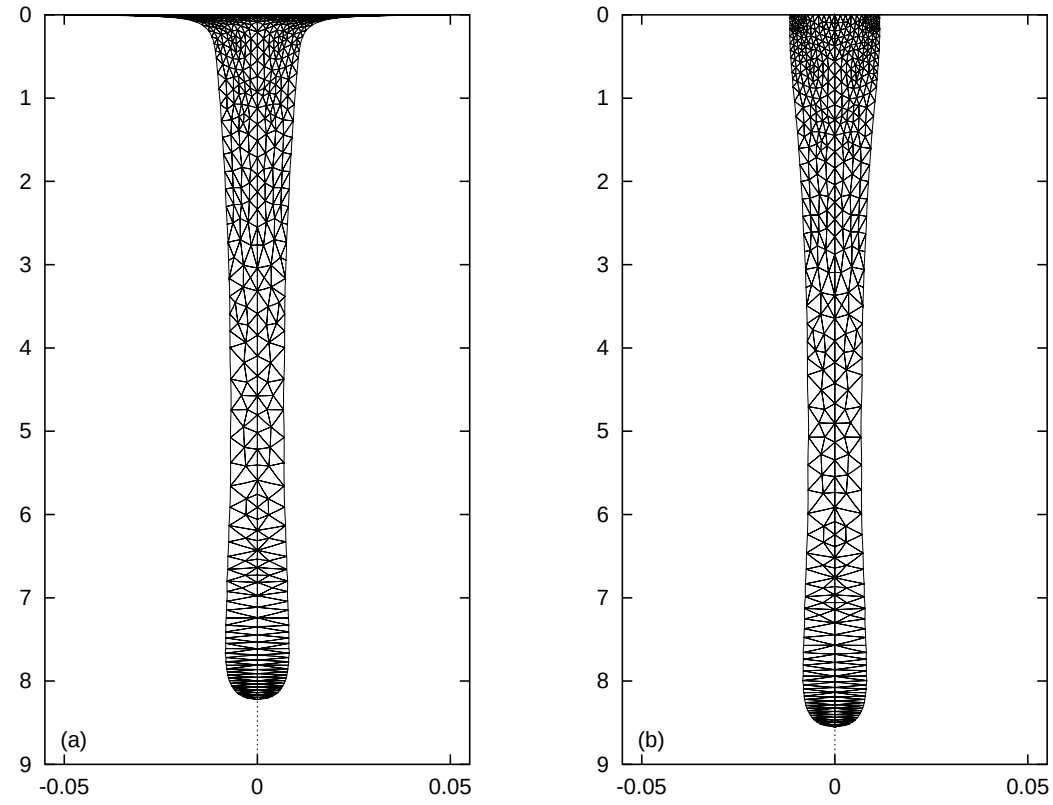

Figure 4.8: Finite-element solution of drop with $\sigma = 1.34016$ at $t = 4.0$. (a) No-slip wall boundary, (b) Slip wall boundary.

boundary develops, showing the importance of the wall boundary condition in this case.

4.8 The Honey Drop Again

We began this chapter with a description of dripping honey, but have digressed from this to consider slender drops which can be modelled using an analytic slender-drop approximation as well as finite-elements. The emphasis has been on the finite-time phenomenon explicitly predicted by the slender-drop theory. This theory is equivalent to one-dimensional theories of both Lagrangian and Eulerian character that have been used extensively in the literature, though the connection is not always easy to see, and enables determination of crisis times, profiles and break-

points in terms of simple quadratures. Because the near-wall region is not modelled by the usual slender-drop approximation, its importance is not always appreciated. A method has been developed for correcting the slender-drop theory for this region, through a modified non-constant expression for the Trouton viscosity.

The results of the slender-drop approximation have been verified by finite-element solutions of the the unapproximated Stokes-flow problem, both in two and three (axisymmetric) dimensions. While the slender-drop approximation is better suited to predicting crisis times and break-points, the finite-element method gives more accurate drop profiles, particularly in the region near the wall. It is also immediately applicable to drops whose initial profile is not slender.

The dripping honey example of Figure 4.1 is, at least in its initial stages, such a non-slender case. The computations were for an axisymmetric geometry, with the "honey" being initially held in a "spoon" consisting of a shallow spherical cap having a depth equal to half the radius R of the rim of the spoon. The free surface of the honey is also assumed to have an initial spherical shape rising a distance $R/10$ above the rim of the spoon. Times (in Figure 4.1) are scaled using the usual $\mathcal{T} = \hat{\mu}/(\rho g L_0)$ with $L_0 = 0.6R$, the initial maximum depth of honey. As for slender geometries, there is an apparent crisis time, which occurs at about $\hat{t}/\mathcal{T} = 8.6$. This corresponds to an effective length of about $\hat{L} = 0.07R$ which is a little less that the maximum initial height of the honey above the rim of the spoon. It is an interesting topic for further research to consider the way in which the effective length $\hat{L}$ depends on actual initial drop size and shape for arbitrary non-slender drops.

Chapter 5

Slumping Glass

5.1 Introduction

Having demonstrated the wide application of the purpose-written finite-element computer program by solving for two quite different very viscous flows in Chapters 3 and 4, we return to a consideration of sagging flows further to the sagging bridge analysed in Chapter 3. Specifically, for the remainder of this thesis we shall be primarily concerned with the slumping of molten glass.

Numerous objects of practical and artistic value are formed from molten glass, which is another very viscous fluid. Glass "slumping", sometimes referred to as sagging or bending, is a forming process whereby a piece of sheet glass suitably supported, usually around its perimeter, is placed into an initially 'cold' oven and heated until it is sufficiently fluid to sag under its own weight. Sometimes the glass sags into a mould, as in the production of some optical components [84] and ornamental items [30, p. 277], and this process will be considered in detail in Chapters 6 to 8. If, however, the item is sufficiently simple, no mould is needed and the glass is allowed to sag until the desired shape is attained. Such slumping is used industrially in the fabrication of windshields and other vehicle components [25, 36, 61], and has been investigated as a means of determining molten glass viscosities [42].

Glass blowing [30] is probably the most well-known glass forming process. The manual procedure involves gathering a quantity of molten glass on a blow pipe through which air is blown. The increased pressure in the interior region results first in the formation of a bubble inside the glass, and subsequently in the expansion and stretching of the glass film similar to the blowing up of a balloon. Throughout this procedure various tools may be used to shape the glass. The similarity between this pressure-driven flow of glass and slumping is easily appreciated, although the deformation resulting from blowing is generally much greater than in slumping with the glass finishing much thinner. Glass blowing is used in the production of a variety of artistic glass items. The process has also been mechanised for the industrial-scale manufacture of glass hollow-ware including containers, tubing [30], and light bulbs [76]. Numerous plastic items are also manufactured by means of blowing [68]. In many cases, the molten glass or plastic is blown into a mould, and the process is then commonly called blow moulding (as in [68]).

Gerhardt [36] comments on the difficulty experienced by the glass industry in manufacturing parts to exacting specifications, particularly in comparison to plastic injection moulding. The problems are more related to the processes than the exact material from which the parts are formed, so that the comment applies also to plastic items manufactured in similar ways. In order to better understand the forming processes and the problems they present, considerable effort has gone into developing models and numerical simulations, much of which is not reported in the published literature because of commercial sensitivity. This thesis is the direct outcome of a desire to better understand slumping as used in the manufacture of optical surfaces which must be accurate to within some reasonably fine tolerances. In blow moulding the major concern seems to be achieving a uniform wall thickness appropriate to the purpose of the item. Excessive thinning in some places with the consequent reduction in strength, or extra thickness with the consequent cost implications are to be avoided (see for example [76]).

Blow moulding of thin viscous films has been the subject of considerable investigation using finite-element methods, with particular application to the manufacture of axisymmetric containers [14, 21, 38, 100]. A slow-flow approximation to the Navier-Stokes equations by dropping the advective terms is used in [14, 38]. Our simplification to the creeping-flow or Stokes equations represents an additional simplification, since, not only the advective terms, but also the partial time derivative have been discarded. This is possible because the time rate of change of velocity as well as the spatial gradient is small. An important consequence of this removal of time from the flow equations, additional to those discussed in Chapter 2, is that temporal changes in viscosity can be accommodated fully within our constant-viscosity model by a change to the time scale. This will be explained in detail later in this chapter (see Section 5.4). A slow-flow approximation, which in neglecting time derivatives is quite similar to that used here, is used in [21, 100], but instead of the Newtonian-fluid assumption inherent in a model based on the Navier-Stokes equations, a viscous Maxwell-fluid model is used. The subject of glass properties will be addressed in the next section, but for now we simply note that in more recent work it is accepted that a Newtonian-fluid assumption is adequate. In all of this previous finite-element modelling, only one or two layers of mesh elements were used across the thickness of the film, which is quite thin. We do not limit our considerations to thin films, and additionally want the greater accuracy given by much finer meshes, especially when modelling the slumping of optical components. As a result the meshes here used are generally similar to that shown in Figure 3.4.

For thin viscous films a thin-layer approximation to the Stokes creeping-flow equations is also valid, and has been used in [48, 96] for analysing pressure driven flows such as blowing, as well as for the stretching of nearly-flat sheets. Indeed, we have already used a similar method for the thin limit of the sagging bridge in Chapter 3. A more intuitive thin-layer analysis based on mechanical principles, and analogous to the Lagrangian derivation of the extensional flow model in Chapter 4, is

given in [68, Ch. 20] for the manufacture of simply shaped items by blow moulding and another process known as thermoforming. In thermoforming a flat sheet is clamped around its edges to a mould, and some combination of vacuum on the mould side, positive air pressure on the opposite side to the mould, and a plunger, forces the sheet into the mould. In the absence of a plunger, the similarity with both blow moulding and slumping is apparent.

In this chapter we shall consider the simple slumping of molten glass supported around its perimeter, in the absence of a mould. A two-dimensional example of this is the sagging bridge of Chapter 3. Here we look at axisymmetric flows in the context of disc viscometry and disc sagging.

To this stage we have been assuming a fluid viscosity that is constant in both time and space. However, as we shall see in the next section, the viscosity of glass is very temperature-sensitive so that, in modelling the slumping of glass as it is heated, we must allow for a changing viscosity. In an ideal slumping environment the temperature will vary in time, but, at any point in time, will be spatially constant. Thus in this chapter it is assumed that conditions are "isothermal" at any point in time, with the viscosity being a function of time only through its dependence on the temperature. The importance of spatial viscosity variations, that are more likely to arise when slumping into a mould, will be investigated in a later chapter. The isothermal assumption permits us to continue using the constant-viscosity model, since no time history is included in the creeping-flow model, which is solved for a point in time when the viscosity is spatially constant. This will be discussed further in Section 5.4, but next we justify the other model approximations for flows of molten glass through a consideration of glass properties.

5.2 Glass Properties

Strictly speaking glasses are any substances that behave as rigid solids, but are like liquids in having no crystalline structure. They are therefore known as both undercooled liquids and amorphous solids [59]. The undercooling of a liquid to form a glass requires that the liquid be very viscous at its freezing point. This high viscosity is a consequence of considerable resistance to atomic rearrangement such as crystallization which normally occurs when a liquid freezes. Most substances are too fluid at their freezing points to prevent crystallization, but a few are sufficiently viscous that the crystallization rate is very slow and cooling can be continued below the freezing point at a sufficiently fast rate that crystallization does not occur. The viscosity continues to rise with decreasing temperature, and at about 10^{12} to $10^{13}\,\mathrm{Pa}\cdot\mathrm{s}$ the substance becomes a glass with mechanical properties similar to an ideal elastic solid [74]. Inorganic mixtures containing a large silica component are notable examples of substances which form glasses, and in this thesis the word "glass" refers to these in accordance with common usage. It is also common to use the word glass to refer to a glass-forming substance in its molten form as well as its 'solid' state, and although this is not in strict accordance with the definition, we have adopted this somewhat loose use of the term in referring to "molten glass".

From this discussion it will already be appreciated that molten glass is a very viscous fluid that might be modelled using the creeping-flow approximation described in Chapter 2. However let us formally validate the model assumptions, and examine the subject of glass viscosity in more detail.

It is generally assumed that molten glass is an incompressible fluid. While this is not strictly true, it is considered to be sufficiently accurate for practical purposes, and in any case is an assumption made in all current methods for determining molten glass viscosities [42]. Furthermore, while a Maxwell-fluid model was used in some early glass modelling work [21, 100], it is now generally accepted that molten glasses, with but few exceptions, are Newtonian fluids [77], and this assumption has been

Table 5.1: ASTM reference viscosities for glass.

Reference	μ (Poise)	Description
Strain point	$10^{14.5}$ (approx.)	Temperature at which the internal stresses in glass are reduced to low values in 4 hours.
Annealing point	10^{13} (approx.)	Temperature at which the internal strains in glass are reduced to an acceptable commercial limit in 15 minutes.
Softening point	$10^{7.6}$ (approx.)	Temperature at which glass will deform under its own weight.
Working point	10^{4} (exact)	
Melting point	10^{2} (exact)	Temperature at which the ingredients of the glass are heated during glass manufacture.

adopted in more recent work [14, 38, 48, 93, 96]. Thus, we shall assume that molten glass is an incompressible Newtonian fluid.

In the following discussion we shall refer to a number of reference viscosities established by the American Society for Testing Materials (ASTM) [4, 44, 77, 81]. These are defined in Table 5.1. Glass viscosities are usually given in "poise" where $1\,\text{poise} = 0.1\,\text{Pa}\cdot\text{s}$.

Glass properties, including viscosity, are very sensitive to chemical composition and temperature. Most commercial glasses fall into one of three main categories as follows [46].

1. *Soda-lime-silica.* The largest group, used for plate glass (as in windows), most bottles and jars, and electric lamp bulbs.

2. *Borosilicate.* Also known as 'pyrex'. Used for domestic ovenware and laboratory glassware.

3. *Lead silicate.* Used for decorative, high quality tableware (crystal), glass to metal seals in electrical components, and radiation shields.

Table 5.2: Temperatures in °C at which reference viscosities are achieved for some typical glasses.

Glass	Strain Point $10^{14.5}$ poise	Annealing Point $10^{13.4}$ poise	Softening Point $10^{7.6}$ poise	Working Point 10^{4} poise
Silica	∼1000	1100	1600	-
Soda-lime-silica	470	510	700	1000
Borosilicate	520	565	820	1220
Lead silicate	390	430	620	970

Each of these categories includes many different glasses having different additives which result in varying properties, and the compositions of a few are given in [46, 59, 74]. Typical reference viscosities for each of these categories are given in [46] some of which are reproduced in Table 5.2. These show both the large change in viscosity for relatively small change in temperature, and also the significant variation in viscosity resulting from different composition. This behaviour is typical of other glass properties also. A comparison of the data for pure-silica glass (see Table 5.2) with that for the other glass types, clearly shows the marked effect that additives have on glass properties such as the temperature and viscosity at which the glass must be worked. No working point is given for silica because increasing the temperature produces other effects so that its viscosity cannot be reduced below about 10^7 poise [46]. Consequently, while silica glass has many excellent qualities, the high temperatures and viscosities at which it must be worked make it impractical for most commercial applications. Soda-lime-silica glasses are used for the applications that we consider.

The relationship between temperature and viscosity must be the most important physical consideration in modelling of glass forming processes, and two complex issues arise in connection with this relationship. We have already touched on one of these, namely the high sensitivity of the viscosity of a molten glass to temperature. In addition, there is the phenomenon of relaxation and equilibrium viscosity.

At large viscosity a change in temperature takes a finite time to be reflected in the corresponding change in viscosity, with this time increasing as viscosity increases (and temperature decreases). Experimental work by Lille published in 1933 is much quoted by way of illustration of this property [77, 59, 74]. It is known as "relaxation", and the final viscosity reached is called the "equilibrium viscosity". With viscosity (μ) measured in decipascal seconds (or poise), Scholze and Kreidl [77] give $L = \log_{10}\mu = 10$ as the point at which viscosity lags by some seconds behind a fast cooling rate, and $13 \leq L \leq 14.5$ as the glass transition in which lag time varies from around 15 minutes to 4 hours and properties are most definitely time dependent. In the numerical modelling of slumping, we are concerned with the change in glass viscosity as it is heated at the commencement of the slumping process, when viscosity is potentially higher than equilibrium. From the author's practical experience (see Chapter 8, Section 8.3) slumping will occur, given sufficient time, at viscosities as high as $L = 11$, which is just on the fringes of the time-dependent property region given in [77]. Practically however, very little slumping would occur at such high viscosities, since oven temperatures are quickly increased to bring the viscosity down into the region $7 \leq L \leq 8$ where slumping occurs at a faster and more acceptable rate and time dependence of the viscosity is not an issue. Furthermore any lag of the viscosity behind the temperature in the time-dependent region will only serve to further prevent slumping at this stage. Thus it is here assumed that relaxation will not significantly affect slumping, and that computation based on equilibrium viscosities will be adequate. In any case, any time lag in the viscosity can be easily corrected by adjusting the time scale as will be shown.

It is generally accepted that the relationship between temperature (T) and viscosity (μ) in a glass is reasonably represented by the Vogel-Fulcher-Tammann (VFT) empirically derived equation (also referred to as the Fulcher equation), dating back to the 1920's [74, 77, 81],

$$\log_{10}\mu = -A + \frac{B}{T - T_0}, \tag{5.1}$$

where base 10 logarithms are used because of the large numbers involved. It involves three constants, A, B and T_0, that are determined for a particular glass from three known temperature-viscosity points spanning a temperature range that is not too extreme. If necessary, a temperature range can be subdivided into smaller intervals and the constants computed for each interval. The temperature range of interest in slumping is sufficiently small to permit a single VFT equation to apply with reasonable accuracy. A two-parameter Walther equation [65] has been found to successfully fit some glass data, and used because it has no singularity at a finite temperature such as in the VFT equation at $T = T_0$ [43]. This singularity does not, however, cause any problem in our work since T_0 is below the temperature range in which slumping can be considered to take place, and hence we stay with the VFT equation.

The density of glass does not vary much with temperature, but can vary substantially with composition. Shand [81] quotes relative densities for silica of 2.203, for a borosilicate glass of 2.13, and for soda-lime-silica glasses of around 2.5, while a high lead glass can have a relative density as great as 6. Morey [59, Table X.3] gives the relative densities of a large number of different glasses ranging from around 1.9 for some borosilicate glasses to around 3 for glasses with a high lead content of around 28 percent. Most glasses, including soda-lime-silica glasses have a relative density of between 2.4 and 2.6.

We have sufficient information at this stage to determine the size of the Reynolds number given by (2.6) for slumping glass flows. For the applications that we consider, data for soda-lime-silica glasses is appropriate. Slumping is usually carried out with the glass at or below its softening point, where the viscosity is of order $10^6\,\mathrm{Pa}\cdot\mathrm{s}$ (10^7 poise). Let us take an average density of $2500\,\mathrm{kg/m^3}$, and a length scale of order 10^{-1} metres. This then gives a Reynolds number of order $\mathrm{O}(10^{-7})$ which is very small and clearly justifies the Stokes creeping-flow approximation.

We must also justify the neglect of surface-tension effects, and hence look at

the capillary number given by (2.14). The surface tension (γ) of glasses decreases with increasing temperature. For a soda-lime-silica glass Shand [81] gives a value of $\gamma = 0.336\,\mathrm{N/m}$ when the viscosity is in the range 10^4 to 10^5 poise. The viscosity, and hence the surface tension, would typically be higher than this for slumping. Data in [59] suggests a surface tension between 0.4 and 0.5 N/m over a temperature range from 1200 to 700°C. Since most slumping occurs at around the softening point, which for soda-lime-silica glasses is about 700°C, a value of $\gamma = 0.5$ is appropriate giving a capillary number of about 490. Then the surface tension terms in (2.13) are of order $\mathrm{O}(10^{-3})$ which is certainly small relative to the other order one terms. Locally at corners, where the length scale is small, the capillary number will be smaller and surface tension will be more important, but this will not affect the overall results.

This completes our consideration of glass properties for the time being, and validates the model assumptions so that we can now proceed to utilize it to consider glass slumping. Before doing this however, we digress to consider a widely circulated, but erroneous, belief that derives from the knowledge that solid glass has properties of a liquid, and seems to have gained currency with many.

5.3 Do Glass Windowpanes Flow?

It is widely reported that apparently solid glass at normal ambient-air temperatures is in fact a liquid, and a common 'proof' advanced for this is that centuries-old glass windowpanes often have greater thickness at the bottom than the top. This is said to be due to the very slow downwards flow of the glass under the influence of gravity, so that over many hundreds of years an initially flat glass windowpane will increase in thickness at the bottom [13, p. 206].

The classification of glass as a solid or a liquid is an issue that has received much discussion over many years, and for which there is no definitive answer [37, 59]. However it is certain that, while glass at normal ambient temperatures has

a molecular structure continuous with its liquid state, it behaves for all practical purposes as a solid and does not flow, so that the variation seen in the thickness of old windowpanes cannot be attributed to gravitational effects over long periods of time. Most arguments that we have encountered to refute this myth, show that thickness variations in windowpanes are a direct result of the processes used to manufacture plate glass prior to the relatively recent introduction of the float glass process in 1959 (e.g. [71, 97]). Of these the "crown" process is most commonly cited as the method of manufacture of most old windowpanes in existence today, and from the description given by Plumb [71] and Morey [59, p. 25] it is easy to appreciate that there would be a general, although small, decrease in thickness from the centre to the edge of the circular "table". It is suggested [71] that, for structural reasons, the panes cut from the "table" were usually mounted with any thicker edge at the bottom. Recently a paper appeared by Zanotto [103] in which relaxation times are shown to be in excess of 10^{32} years for typical medieval glasses, from which it is concluded that "medieval and contemporary window glasses cannot flow at room temperature in human time scales!"

The fallacy of this myth is here shown in a different manner. Supposing that window glass is an extremely viscous liquid that will flow, we use some mathematics to determine the time that must transpire before any appreciable thickening at the base of the windowpane is seen.

The 'flow' of glass in a windowpane may be readily examined by a slender approximation to the creeping-flow model, in a very similar manner to the nearly-plane bridge considered in Section 3.6. The details given here are also directly applicable to the slender drops of Chapter 4, with the simple reversal of the direction in which gravity acts.

Consider a glass window that is thin relative to both its breadth and height. Its large breadth enables us to simplify the problem to two dimensions. Let the x direction be vertically up, with gravity acting vertically down, i.e. the x direction

is equivalent to x_2 of Chapter 2. So as to keep the notation of Section 3.6, let the height of the windowpane be given by $L(t)$, and its faces be given by $y = \pm h(x,t)$, with initial conditions $L(0) = w$ and $h(x,0) = \epsilon w$ where ϵ is a very small number. Since a windowframe does not prevent small movement in a direction parallel to the y-axis (as witnessed by the fact that a window will rattle in a reasonably high wind), the windowframe is assumed to provide a slip boundary support at $x = 0$ that prevents vertical 'flow' of the glass but allows horizontal 'flow'. All other boundaries are considered to be free surfaces.

With the window oriented along the x-axis in this manner, the equations to be solved will be, with just a few exceptions, identical to those given for the nearly-plane sagging bridge in Section 3.6. Dimensionless variables are defined as in (3.18), and $\bar{h} = 0$. However the definition of dimensionless time t^* is more appropriately given by

$$t = wt^*/\mathcal{V}. \tag{5.2}$$

This scale is suggested by the fact that the velocity component in the y direction (v) can be expected to be of order $\epsilon\mathcal{V}$, where $\mathcal{V}$ is a measure of the velocity component in the x direction (u) to be specified later. Then changes in window thickness can be expected to occur in times of order $\epsilon w/(\epsilon\mathcal{V})$, which leads directly to (5.2). Let us also define

$$L^*(t) = L(t)/w \tag{5.3}$$

so that $L^*(0) = 1$. The flow domain of this problem is, at $t^* = 0$, $0 \le x^* \le 1$, $|y^*| \le 1$.

Because gravity is acting in the $-x$ direction (rather than the $-y$ direction as in Section 3.6), this must be shown in the Stokes equations, which now become (on dropping asterisks)

$$\epsilon^2 p_x = \epsilon^2 u_{xx} + u_{yy} - \epsilon^2 \frac{\rho g w^2}{\mu\mathcal{V}} \qquad \text{and} \tag{5.4}$$

$$\epsilon p_y = \epsilon^2 v_{xx} + v_{yy}. \tag{5.5}$$

At this point let us choose the velocity scale $\mathcal{V} = \rho g w^2/\mu$; then the last term on the right of (5.4) becomes $-\epsilon^2$. The continuity equation remains unchanged as

$$\epsilon u_x + v_y = 0, \tag{5.6}$$

but the changed time scale results in the new kinematic condition on $y = \pm h$

$$\pm\epsilon(h_t + u h_x) = v, \tag{5.7}$$

which replaces (3.27). The zero-stress conditions on $y = \pm h$ become, on setting $\bar{h} = 0$,

$$u_y + \epsilon v_x = \mp\epsilon^2(p - 2u_x)h_x, \tag{5.8}$$

$$-\epsilon p + 2v_y = \pm\epsilon(u_y + \epsilon v_x)h_x. \tag{5.9}$$

Finally the wall condition at $x = 0$ is just $u = 0$.

This problem is considerably simpler than the sagging thin bridge of Section 3.6 for two main reasons. First, because of our definition of $\mathcal{V}$, we do not need to use as many of the terms in the power-series expansions in ϵ to solve for the leading-order terms. In addition, symmetry requirements about $y = 0$ in the current problem demand that u and p are even functions of y, while v is odd. Thus, using a Taylor expansion about $y = 0$, we can replace equations (3.28) with

$$\begin{aligned}
u(x,y,t) &= u_0(x,t) &&+ \epsilon^2 y^2 u_2(x,t) &&+ \dots, \\
v(x,y,t) &= \epsilon y v_1(x,t) &&+ \epsilon^3 y^3 v_3(x,t) &&+ \dots, \\
p(x,y,t) &= p_0(x,t) &&+ \epsilon^2 y^2 p_2(x,t) &&+ \dots, \\
h(x,t) &= h_0(x,t) &&+ \epsilon h_1(x,t) &&+ \dots,
\end{aligned} \tag{5.10}$$

where again the asterisks are omitted on dimensionless variables.

Continuity (5.6) immediately requires that

$$v_1 = -u_{0x}, \tag{5.11}$$

and the free-surface condition (5.9) shows that

$$p_0 = 2v_1 = -2u_{0x}. \tag{5.12}$$

From the Stokes equation (5.4) we have

$$p_{0x} = u_{0xx} + 2u_2 - 1, \tag{5.13}$$

from which we can determine that

$$2u_2 = 1 - 3u_{0xx}. \tag{5.14}$$

Satisfying the free-surface condition (5.8) on $y = \pm h$ yields

$$\pm h_0(2u_2 + v_{1x}) = \mp(p_0 - 2u_{0x})h_{0x} \tag{5.15}$$

while the kinematic condition (5.7) yields

$$h_0 v_1 = h_{0t} + u_0 h_{0x}. \tag{5.16}$$

On substituting the above expressions for p_0, v_1 and u_2 and rearranging (5.15) and (5.16) we have finally

$$\boxed{h_0 = 4(h_0 u_{0x})_x\,,} \tag{5.17}$$

$$\boxed{h_{0t} = -(u_0 h_0)_x\,,} \tag{5.18}$$

which compare with the (dimensional) equations (4.7) and (4.8) given for the extensional flow of Chapter 4 with the direction of gravity reversed.

Thus, the 'flow' and changing shape of the window glass, is to leading order obtained by solving the coupled equations (5.17) and (5.18) for $u_0(x,t)$ and $h_0(x,t)$ subject to suitable boundary conditions. One of these is $u_0 = 0$ at $x = 0$. We need one further condition, and this is obtained from the zero stress conditions that must be satisfied on the horizontal boundary at $x = L(t)$:

$$-\epsilon p + 2v_y = 0, \tag{5.19}$$

$$u_y + \epsilon v_x = 0. \tag{5.20}$$

From these we have

$$-p_0 + 2v_1 = 0, \tag{5.21}$$

$$u_{0y} = 0, \tag{5.22}$$

$$y(2u_2 + v_{1x}) = 0. \tag{5.23}$$

With the exception of (5.23), these have been satisfied in our derivation of (5.17) and (5.18), so that the second boundary condition comes from (5.23) which must be true for all $|y| \leq 1$ and hence can be simplified to $2u_2 + v_{1x} = 0$. Substituting for u_2 and v_{1x} and rearranging finally gives

$$u_{0xx} = \frac{1}{4} \tag{5.24}$$

which can in turn be substituted into (5.17) to give

$$h_{0x}u_{0x} = 0 \tag{5.25}$$

on $x = L(t)$. This shows that the zero-vertical-velocity-gradient boundary condition on $x = L(t)$ used in the extensional-flow problem of Chapter 4 (see p. 89) is equivalent to a zero-stress condition as claimed. We are also justified in using this same condition, i.e.

$$u_{0x} = 0, \tag{5.26}$$

for this problem, since there is no mechanism for a change of cross-sectional area at $x = L(t)$. Thus, as previously (p. 89), we may write (5.18) in terms of the material derivative of h_0, which leads directly to (5.26).

With a difference only in the direction of gravity, the problem defined by equations (5.17) and (5.18) with boundary conditions $u_0 = 0$ at $x = 0$ and $u_{0x} = 0$ at $x = L(t)$, and an initial rectangular profile, is identical to that defined for the extending rectangular slab in Chapter 4, and we can immediately adapt (4.5) and

(4.6) to write down the dimensionless solution in terms of the dimensionless label variable ξ where $0 \leq \xi \leq 1$, and $x = \xi$ at $t = 0$:

$$h_0(\xi, t) = 1 + \frac{t}{4}(1 - \xi), \tag{5.27}$$

$$\begin{aligned} x(\xi, t) &= \int_0^{\xi} \frac{1}{h_0(\xi_1, t)} d\xi_1 \\ &= -\frac{4}{t}\left(\log\left[1 + \frac{t}{4}(1 - \xi)\right] - \log\left(1 + \frac{t}{4}\right)\right). \end{aligned} \tag{5.28}$$

From these last two equations we have immediately that the glass thickness at the top of the window is independent of time, i.e. $h_0 = 1$ at $\xi = 1 \equiv x = L(t)$, and that, for all time $t > 0$, the window thickness is a strictly monotone-decreasing function of ξ with maximum thickness at its base $\xi = x = 0$ being given by $2h_0(0, t)$ where

$$h_0(0, t) = 1 + \frac{t}{4}. \tag{5.29}$$

The height of the window $L(t)$ is given by substituting $\xi = 1$ into (5.28):

$$L(t) = x(1, t) = \frac{4}{t}\log\left(1 + \frac{t}{4}\right). \tag{5.30}$$

The result (5.29) can also be obtained in the following simple manner. Since the time scale is very large ($\sim \mu$), the life span of even quite old window glass will be, in terms of dimensionless time, relatively short. Hence, we are justified in considering the solution at (dimensionless) time $t = 0^+$ when for all practical purposes $h_0 = 1$ and $L = 1$. Then integrating (5.24) and using the specified boundary conditions to solve for constants we obtain

$$u_{0xx} = \frac{1}{4}, \tag{5.31}$$

$$u_{0x} = \frac{1}{4}(x - 1), \tag{5.32}$$

$$u_0 = \frac{1}{8}x(x - 2). \tag{5.33}$$

From (5.18) we thus have

$$h_{0t} = -u_{0x} = \frac{1}{4}(1 - x), \tag{5.34}$$

and on integrating with respect to time t and setting $x = 0$ we obtain (5.29).

Thus an increase in thickness of q percent occurs in a time of $t^* = 4q/100$, where we have again introduced asterisks to denote dimensionless variables. In dimensional terms this is a time of

$$t = \frac{w}{\mathcal{V}} t^* = \frac{\mu}{\rho g w} \frac{4q}{100}. \tag{5.35}$$

Now glass becomes substantially rigid at the strain point where the viscosity is of order $10^{14.5}$ poise (see Section 5.2). For soda-lime glasses from which windows are usually made, the strain point would be reached at a temperature of about 500°C, and as the temperature decreases, the viscosity increases enormously. At normal room temperatures, we can expect the viscosity to be around 10^{20} poise, i.e. $\mathrm{O}(10^{19})\,\mathrm{Pa}\cdot\mathrm{s}$, or even higher. With the density of glass at about 2500 $\mathrm{kg/m^3}$, the time taken for a windowpane with an initial height of (say) 250 mm to increase just five percent in thickness is at least

$$\begin{aligned} 4 \times 0.05 \frac{\mu}{\rho g w} &= 0.2 \times \frac{10^{19}}{2500 \times 10 \times 0.25} \\ &= 3.2 \times 10^{14} \text{seconds} \\ &\approx 10^{7} \text{years}. \end{aligned} \tag{5.36}$$

Thus, we must wait some 10 million years (!) and not centuries for a windowpane with initial height and thickness of 250 mm × 5 mm to thicken by just 0.25 mm at the bottom, and even modern windows with heights of a few metres will require around one million years for this magnitude of increase in thickness. Clearly the story of viscous-flow of old glass windowpanes resulting in greater thickness at the bottom than the top, is false.

5.4 Temporal Viscosity Changes

We now return to the subject of glass slumping with which we began this chapter. From the earlier description of slumping, as well as the discussion of glass properties

and the sensitivity of viscosity to temperature in particular, it is apparent that we need to include a time-varying viscosity in the model since some slumping may, and in all probability will, occur during the oven heating and cooling stages at temperatures lower than the maximum reached by oven and glass. Following on from the foregoing analysis of the 'flowing' windowpane we could even consider that slumping commences, albeit very slowly, when the glass is still cold and the oven temperature has only just started to increase. This is of course unnecessary, and we rather consider that slumping occurs over a limited temperature range when the glass is sufficiently fluid for sensible slumping to occur.

It is also possible that spatial temperature and viscosity variations will arise during the heating of the glass due to non-uniformity of oven temperature, supporting of the glass, heat transfer from oven to glass, and other factors. However it seems reasonable for the simple slumping presently being considered, to assume that these will be small, and that at any instant in time conditions are isothermal, i.e. temperature and viscosity throughout the glass are constant. Spatial viscosity variation will be considered later along with slumping into a mould where non-isothermal conditions seem more likely.

As already briefly mentioned in Section 5.1, with the isothermal assumption, temporal viscosity changes can be accommodated within the current constant-viscosity model. This is an important consequence of the neglect of the time derivatives in the field equations (2.7). To see how we can handle temporal viscosity changes let us first consider a viscosity that is everywhere and forever constant, and look again at the dimensionless form of the field equations (2.2) and (2.7). The solution of these equations for flow velocity and pressure is influenced only by the flow-domain geometry at the time of solution, and not by the time itself. The finite-element mesh is then adjusted to reflect the distortion caused by flow over a dimensionless time step Δt^*, and the velocity and pressure fields are recomputed over this new geometry. After a dimensionless time t^* there is only one possible geometrical out-

come from this process (ignoring numerical error caused by discrete approximation). The dimensionless time t^* over which this happens represents a real time given by $t^*\mathcal{T}$, where the time scale $\mathcal{T} = \mathcal{L}/\mathcal{U}$ is dependent on the constant viscosity through $\mathcal{U} = \rho g \mathcal{L}^2/\mu$ as defined in (2.5). Different values of the viscosity μ mean different real slump times, but the same final outcome. Thus, the model tells us that, given identical initial geometries, we can obtain the same item by slumping at a high viscosity for a long time, or a lower viscosity for a shorter time.

Next let viscosity $\mu(t)$ be given as a function of time, and let us define a time scale that is also a function of time

$$\mathcal{T}(t) = \frac{\mathcal{L}}{\mathcal{U}(t)} = \frac{\mu(t)}{\rho g \mathcal{L}}. \tag{5.37}$$

In any very small real-time interval dt the viscosity is constant, and the corresponding dimensionless-time interval is given by $dt^* = dt/\mathcal{T}(t)$. Then a slumping time period of length t_f during which the viscosity is given as a function of time by $\mu(t)$ is obtained in dimensionless terms by summing over all very small intervals dt:

$$t^* = \int_0^{t_f} \mathcal{T}(t)^{-1} dt = \rho g \mathcal{L} \int_0^{t_f} \mu(t)^{-1} dt\,. \tag{5.38}$$

For constant viscosity, (5.38) reduces to $t^* = t_f/\mathcal{T}$, as indeed it should.

Now it is evident from (5.38) that a particular value of t^* can be obtained from a host of different choices of t_f and $\mu(t)$ (keeping ρ, g and $\mathcal{L}$ constant), including constant μ, and the finite-element program will for some initial geometry compute a single slumped shape for this dimensionless time. Another look at (5.38) shows that, provided there is equal area under each curve of inverse viscosity $\mu^{-1}(t)$, from the time at which slumping starts to the time at which it ends (see Figure 5.1), the final product resulting from some initial geometry will be identical. This means that the dimensionless problem (for constant $\mu = 1$) can be solved once to give the shape of the slumping item over dimensionless time. The outcome for different viscosity (i.e. temperature) profiles and real times can then be determined by computing the equivalent dimensionless time as a post-processing operation, and selecting the

slumped shape corresponding to this time. We will make use of this methodology as necessary throughout the remainder of this chapter.

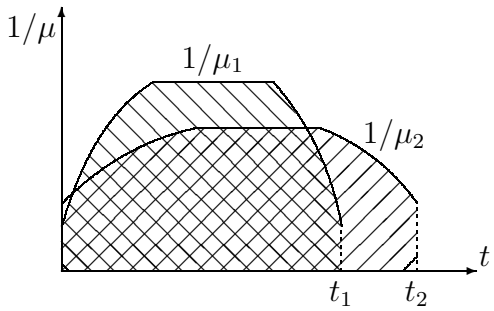

Figure 5.1: Equivalent Viscosity Curves.

5.5 Disc-Bending Viscometry

The rate of flow of a very viscous fluid, such as molten glass, under its own weight or an applied load is obviously highly dependent on the viscosity of the fluid, and this fact is used in a number of methods of viscosity measurement such as fibre-elongation, beam-bending [44], and disc-bending viscometry [42]. Fibre elongation is an extensional flow problem to which the work of Chapter 4 is relevant, while disc bending is a three-dimensional (axisymmetric) flow similar to the two-dimensional sagging problem of Chapter 3.

In 1974 an investigation was made by Corning Glass Works (now Corning Incorporated) into disc bending as a method for determining glass viscosities, as reported in a published paper [42] and an internal Corning report [40]. The detailed results of this work, including the raw experimental data, have been made available to us, and with the aid of these results, we here consider the use of the creeping-flow finite-element program as a viscometry tool. In so doing we compare some finite-element simulations with real and independent experimental data, a necessary step in the development of any numerical tool for simulating industrial processes.

A thin glass disc of radius a and thickness h is supported around its perimeter

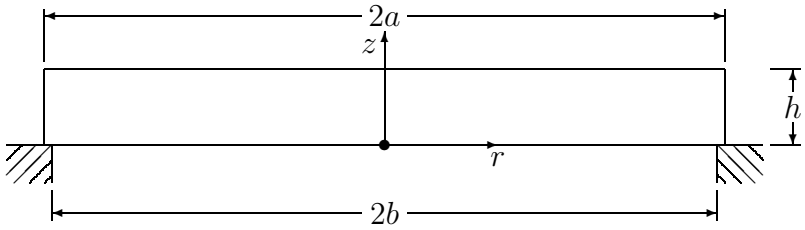

Figure 5.2: Initial setup used in disc bending viscometry.

on a small ledge as in Figure 5.2, and heated in an oven so that it sags under its own weight. Gulati et al. [42, 40] use a viscoelastic analogy to derive equations relating sag δ and rate of sag $\dot{\delta}$ to viscosity μ from the equations for the elastic bending of a disc. The temperature T versus viscosity relationship of a particular glass in the temperature range covering the strain and softening points is then obtained from these equations and experimental measurements of δ and $\dot{\delta}$. The calculated viscosities are shown to be in reasonable agreement with the curve obtained using the generally accepted beam-bending and parallel-plate methods [44].

The gravity-driven creeping-flow model of Chapter 2 can also be applied to disc-bending viscometry. For any particular disc, we are able to compute the relationship between sag and rate of sag. The rate of sag at some point in time is dependent only on the physical geometry (determined by the initial disc dimensions and the subsequent amount of sag) and the viscosity at that time so that we can use the Corning experimental data together with the finite-element computations to derive our own μ versus T relationship which we can compare with that given in [42].

The problem is axisymmetric, and hence the axisymmetric continuity and Stokes equations of Section 2.5 are applicable. Symmetry in $r = 0$ means that, as previously, the computational flow domain may be halved to $0 \leq r \leq a$ by satisfying boundary conditions (2.15) and (2.16) on $r = 0$. On the ledge support, the no-slip conditions (2.11) apply and on all other boundaries the free-surface conditions (2.15) must be satisfied. Let us non-dimensionalize using the radius a as the characteristic length $\mathcal{L}$ so that $0 \leq r^* \leq 1$. Let us also define sag δ to be the magnitude of the vertical

displacement of the point at the centre of the lower surface from its position at time $t = 0$, and rate of sag $\dot{\delta}$ to be the magnitude of the velocity at the same point, with corresponding non-dimensional quantities

$$\delta^* = \delta/a, \qquad \text{and} \qquad \dot{\delta}^* = \dot{\delta}/\mathcal{U}. \tag{5.39}$$

Note that $\mathcal{U}$ is the characteristic velocity that depends on μ and possibly varies in time. However at any given time at which δ and $\dot{\delta}$ are measured, it is a fixed though unknown value. It is only a matter of computation to obtain a plot of $\dot{\delta}^*$ versus δ^* for a disc of aspect ratio h/a from the finite-element program.

Now, given the triplet $(T, \delta, \dot{\delta})$ of experimental data for some disc of radius a and thickness h, we compute the equivalent non-dimensional sag $\delta^* = \delta/a$. For the slumped geometry indicated by this sag there is only one possible value for the rate of sag $\dot{\delta}^*$ which we obtain from the finite-element program output. Then the velocity scale can be calculated from the actual experimental rate of sag $\dot{\delta}$ and this dimensionless value $\dot{\delta}^*$ as $\mathcal{U} = \dot{\delta}/\dot{\delta}^*$, or alternatively, since $\mathcal{U} = \rho g a^2/\mu$, we can use the known values of ρ, g, and a to determine the unknown viscosity μ at temperature T:

$$\mu = \rho g a^2 \frac{\dot{\delta}^*}{\dot{\delta}}\,. \tag{5.40}$$

If the temperature (and hence the viscosity) varies in time as the disc sags, and we are given experimental data triplets $(T, \delta, \dot{\delta})$ over a range of temperatures, we can use this method to calculate the viscosity at each temperature, and plot viscosity versus temperature for a particular glass.

5.6 The Viscoelastic Analogue Equations

The creeping-flow model of disc-bending viscometry can be compared with the Gulati et al. viscoelastic model [42]. This latter model is derived from the elastic

bending of a thin disc with a free but immovable edge support through a viscoelastic analogy whereby the elastic-bending equations are differentiated with respect to time, and the Young's modulus E replaced by $2(1+\sigma)\mu$ where σ is Poisson's ratio for glass. According to [42] incompressibility of glass is assumed, so that $\sigma = 0.5$ and E is replaced by 3μ (the Trouton result [92]). However, the internal Corning report [40] differs from the published paper [42] and indicates that a value of $\sigma = 0.3$ may have been used. This will be commented on further as we proceed, but it need not be of major concern since we can readily compute our own results from the incompressible viscoelastic equations (i.e. with $\sigma = 0.5$) and experimental data, and compare these with the results from the creeping-flow model which definitely assumes incompressibility.

The equations given in the published paper [42] for determining μ from rate of sag $\dot{\delta}$ and cumulative sag δ, assuming glass to be an incompressible fluid, are as follows:

$$\mu = \frac{11}{64}\frac{\rho g a^4}{\dot{\delta} h^2} \qquad \text{for } \delta \leq h/2, \tag{5.41}$$

$$\mu = \frac{11}{64}\frac{\rho g a^4}{[1+7.35(\delta/h)^2]\dot{\delta} h^2} \qquad \text{for } h/2 < \delta \leq h, \tag{5.42}$$

$$\mu = 0.023\frac{\rho g a^4}{\dot{\delta}\delta^2} \qquad \text{for } \delta > h. \tag{5.43}$$

However, on examination of the Corning report [40] it is evident that equation (5.42) has been obtained from a disc-bending equation for a material with $\sigma = 0.3$ by writing the right-hand side only (not the left) in terms of a general σ, time-differentiating and replacing the Young's modulus by $2(1+\sigma)$, and then substituting $\sigma = 0.5$ for an incompressible fluid. According to the Corning report no disc-bending equation for general σ was available at the time of writing, forcing the use of the $\sigma = 0.3$ equation, and in fact this report gives, in place of (5.42), an equation derived from the disc-bending equation with a consistent use of $\sigma = 0.3$. The published paper contradicts the report in giving a disc-bending equation in terms of a general σ, but we are not certain of the correctness of this equation, since setting $\sigma = 0.3$ does not

yield the disc-bending equation quoted in the report, from which (5.42) has been clearly derived. Hence we continue to use (5.42), but noting that it is not strictly valid for either a compressible fluid with $\sigma = 0.3$ or an incompressible fluid with $\sigma = 0.5$. Equations (5.41) and (5.43) are for an incompressible fluid with $\sigma = 0.5$.

We need at this stage to clarify the definition of a. For the creeping-flow model this is the radius of the disc. In (5.41) to (5.43) it is also the disc radius since the equations are for a knife-edge perimeter support. However, in reality, the support was a very small ledge as shown in Figure 5.2, which was equated with a knife-edge support, so that for calculations in [42] the free-span radius of the disc b (i.e. the disc radius minus the support-ledge width) is actually used for a in place of the disc radius. Because the ledge width is very small, this approximation makes very little difference to the results, and it is ignored here with a always taken to be the disc radius.

Equations (5.41) to (5.43) can be written in terms of non-dimensional quantities δ^* and $\dot{\delta}^*$, and aspect ratio $H = h/a$

$$\dot{\delta}^* = \frac{11}{64}\left(\frac{1}{H}\right)^2 \qquad \text{for } \delta^* \leq H/2, \tag{5.44}$$

$$\dot{\delta}^* = \frac{11}{64}\frac{1}{[1 + 7.35(\delta^*/H)^2]H^2} \qquad \text{for } H/2 < \delta^* \leq H, \tag{5.45}$$

$$\dot{\delta}^* = 0.023\frac{1}{(\delta^*)^2}, \qquad \text{for } \delta^* > H. \tag{5.46}$$

For a given aspect ratio h/a, the relationship between $\dot{\delta}^*$ and δ^* so defined, can be readily compared with corresponding output from the creeping-flow model.

5.7 Viscosity Calculations

The disc-bending viscometry experiments conducted by Corning Glass Works were for glass discs of 70 mm diameter and 1.53 mm thickness, i.e. with an aspect ratio $h/a = 0.153/3.5 = 0.0437$. The glass had strain, annealing and softening points of

447, 484, and 645°C respectively, and the temperature range considered was 550–640°C. The dimensions and properties used for computations are given as

$$\begin{aligned} \rho &= 2.356\,\mathrm{g/cm^3}, \\ g &= 980\,\mathrm{cm/s^2}, \\ a &= 3.4925\,\mathrm{cm}, \\ h &= 0.153\,\mathrm{cm}. \end{aligned} \tag{5.47}$$

Note that the value of a above is the free-span radius of the disc, giving a support ledge width of $3.5 - 3.4925 = 0.0075\,\mathrm{cm}$.

The method used to obtain measurements of sag during slumping required that a conducting gold film be painted on the bottom of the disc in a small circular area at the centre, around the circumference of the disc and along a radial line connecting the two [42, 40]. This increased the rigidity of the glass and also seemed to affect heat transfer so that sag was impeded in the painted areas and the sag measurements recorded (for the bottom centre of the disc) were less than would have been without the gold-film paint. Despite this error results were considered to be quite good.

Figure 5.3 shows $\dot{\delta}^*$ as a function of δ^* for this disc as given by the creeping-flow finite-element program, and also the viscoelastic analogue equations (5.44) to (5.46). We see that the viscoelastic analogue equations considerably over-estimate the rate of sag in the early stages of slumping when $\delta < h/2$ (i.e. $\delta^* < H/2$), and under-estimate for larger sag ($\delta^* > H/2$), in comparison with the essentially exact creeping-flow model. Equation (5.40) then indicates that, compared with the creeping-flow model, the viscosity predicted by the viscoelastic model from experimental data $(\delta, \dot{\delta})$ will be too large while $\delta < h/2$ and too small for larger sag.

The raw experimental data for this slumping case, supplied by Corning, shows the temperature T and sag δ at discrete points in time during slumping, with rate of sag $\dot{\delta}$ being computed from sag measurements at two consecutive time points t_1 and t_2 by simple Euler differencing, i.e. $(\delta_2 - \delta_1)/(t_2 - t_1)$. This sag rate is considered to be applicable at a temperature which is the average of the temperatures recorded at t_1

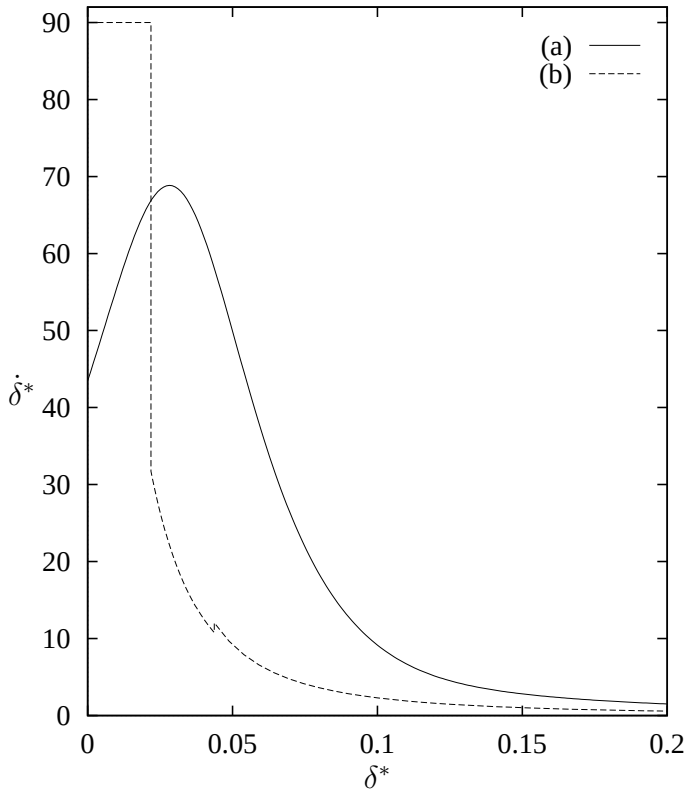

Figure 5.3: Rate of sag versus sag of an initially-flat disc as given by (a) the creeping-flow model, and (b) the viscoelastic analogue equations.

and t_2. Consistent with this, we take the corresponding sag to be the average of the recorded sags at times t_1 and t_2. Non-dimensionalizing the sag data by dividing by the disc radius a, we can then determine the dimensionless sag rate from the finite-element output (see Figure 5.3), and with this and the experimentally measured sag rate calculate the viscosity from (5.40). Some results of this process are shown in Table 5.3.

Also shown in Table 5.3 are viscosities calculated using the viscoelastic analogue equations (5.41) to (5.43). These have been recalculated in preference to using the Corning results since, while the specific equations obtained by substituting all known parameters into (5.41) to (5.43) for the glass disc under consideration are correctly given in the published paper [42] for the incompressible fluid assumption (except of course for (5.42) where both $\sigma = 0.3$ and $\sigma = 0.5$ have been used as already noted),

Table 5.3: Viscosity from sag and rate of sag of a glass disc, using the creeping-flow (CF) and viscoelastic-analogy (VEA) models.

T	δ	$\dot{\delta}$	δ^*	$\dot{\delta}^*$	$\log_{10}\mu$	
°C	cm	cm/min			CF	VEA
552.5	0.0095	0.002857	2.714×10^{-3}	46.68	10.44	10.73
560.0	0.017	0.004286	4.857×10^{-3}	49.28	10.29	10.55
571.0	0.0375	0.0100	1.071×10^{-2}	56.33	9.98	10.18
580.0	0.070	0.019355	2.000×10^{-2}	64.99	9.76	9.90
592.0	0.155	0.035294	4.429×10^{-2}	58.58	9.45	8.75
601.0	0.250	0.031034	7.143×10^{-2}	24.87	9.13	8.39
609.5	0.310	0.023377	8.857×10^{-2}	13.49	8.99	8.33
614.5	0.340	0.020225	9.714×10^{-2}	9.73	8.91	8.31
622.5	0.385	0.021429	1.100×10^{-1}	6.72	8.73	8.18
629.5	0.435	0.021818	1.243×10^{-1}	4.60	8.55	8.06
635.0	0.480	0.016216	1.371×10^{-1}	3.51	8.57	8.11

the internal Corning report [40] categorically states that Poisson's ratio $\sigma = 0.3$ was used, and certainly derives one of the three viscoelastic analogue equations on this basis, as has already been mentioned. In addition, the final equations given in the report [40], obtained by substituting values for all known parameters, could not be verified with either $\sigma = 0.3$ or $\sigma = 0.5$.

The viscosities calculated from experimental sag and rate-of-sag data using the creeping-flow and viscoelastic-analogy models compare quite well as seen in Table 5.3. Consistent with an earlier comment, the viscoelastic analogy equations give higher viscosities than those based on the creeping-flow finite-element computations while the sag is less than half of the initial disc thickness ($\delta < 0.0765$), after which the situation is reversed. Agreement between the two methods is best for sags of less than half the initial disc thickness.

In Figure 5.4 is shown, as the solid curve, a relationship between temperature and viscosity for the glass used in the sag experiments, taken from [42]. It was calculated from beam-bending and parallel-plate experimental measurements, methods that are

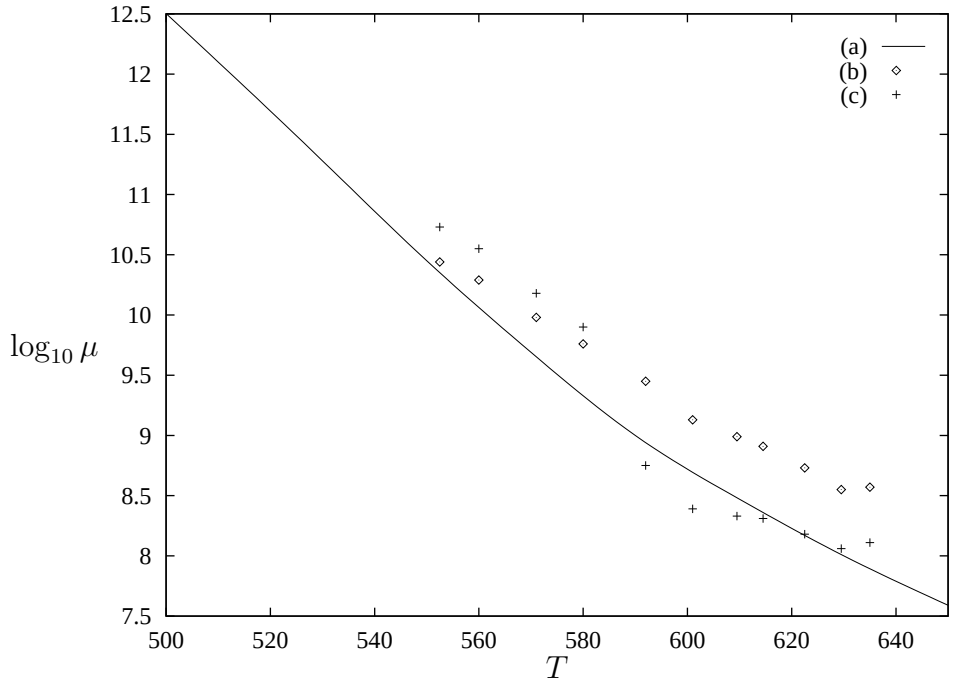

Figure 5.4: The viscosity-temperature relationship given by (a) beam-bending and parallel-plate methods, (b) disc bending with the creeping-flow model, and (c) disc bending with the viscoelastic-analogy model.

generally accepted as reliable within the glass industry. Beam bending is used for high viscosities of 10^8 to 10^{15} poise, while compression between two parallel plates is used at lower viscosities of 10^4 to 10^9 poise [44]. We shall refer to this curve as the "accepted" curve. Also plotted are the results, tabulated in Table 5.3, given by disc bending methods and both the creeping-flow and viscoelastic models.

A comparison of this figure with the similar plot given in [42] shows some difference between the viscosities computed here using the viscoelastic equations and those computed by Gulati et al., particularly for $T < 580°\text{C}$. This may be due to differences in the value of Poisson's ratio. The sudden jump in viscosity given by the viscoelastic equations from $T = 580°\text{C}$ to $T = 590°\text{C}$ is due to the sag exceeding half the initial disc thickness and the consequent change of equation used for computing the viscosity.

This jump is not seen in the viscosities calculated by means of the creeping-flow finite-element computations, which are also more accurate for sags of less than $h/2$. In fact, in the very early stages of slumping at a temperature of about 550°C the viscosity given by the creeping-flow model is very close to the accepted viscosity-temperature curve, but as the temperature increases the difference also increases. This can be understood when it is remembered that the gold-film paint used on the bottom of the disc at the centre resulted in the disc sagging less than it otherwise would have, and the rigidity imparted by the paint can be expected to have a greater effect as sag increases. Thus the rates of sag computed from the sag measurements should be corrected upwards, and the correction should increase with increasing sag. Equation (5.40) shows that this would result in reducing the viscosities calculated, so bringing our results closer to the accepted curve.

From these results, it is concluded that disc-bending experiments in conjunction with creeping-flow finite-element computations can be used to determine the temperature-viscosity relationship of a glass. However a different sag measuring method, perhaps optical, that does not impede sag would be a highly desirable improvement. In fact, the results based on the finite-element computations are essentially more accurate than those given by the viscoelastic analogy equations; the apparent accuracy of the viscoelastic model at temperatures higher than 590°C is probably somewhat fortuitous, and more accurate measurements of sag and rate of sag can be expected to drop these results below the accepted viscosity-temperature curve.

Gulati et al. [42] suggest that the discrepancy between their results and the accepted viscosity-temperature curve is in part due to the assumption of incompressibility, and that using $\sigma = 0.3$ instead of 0.5 may eliminate some of this, particularly when nonlinear disc bending theory is being employed. This does not seem likely for a number of reasons. Firstly it seems highly probable from a reading of the internal Corning report [40] that a value of $\sigma = 0.3$ was used. Secondly, as commented in

[42], the assumption of incompressibility is common to all glass viscometry methods so that the accepted curve is also obtained with this assumption. Finally, the results obtained with the creeping-flow model (with the assumption of incompressibility) and finite-element program are such that we can expect better agreement with the accepted curve should more accurate sag and rate-of-sag measurements be available.

In concluding this consideration of disc-bending viscometry, it is reasonable to say that the comparison of creeping-flow finite-element computations with completely independent experimental data is very good allowing for the known inaccuracies in the data, and gives confidence in our use of the creeping-flow model for analysis of molten glass flows.

5.8 Disc Sagging

In disc-bending viscometry the unknown viscosity-temperature relationship of a very viscous fluid such as molten glass is determined from experimental measurements of temperature and sag. In glass forming by slumping we are, of course, interested in the inverse problem of determining how the disc will sag for a given temperature profile where the viscosity-temperature relationship of the fluid is known.

Following on from their work on disc-bending viscometry Gulati et al., [41] investigated the development of sagging schedules for circular glass lens blanks using viscoelastic analogy equations similar to those for disc-bending viscometry. This work was never published, but has been made available to us by Corning and gives useful experimental results for comparison with our creeping-flow finite-element model.

Circular glass discs of radius 35 mm and thickness 1.5 m were clamped about their perimeter in a mould such that their radial span was $a = 32.77$ mm. The glass was the same as used for the disc-bending experiments with density as given in (5.47), and in this case incompressibility of the molten glass fluid was clearly assumed. The viscosity-temperature relationship for this glass was approximated

by the VFT equation (5.1) with constants

$$A = 1.018, \qquad B = 3303, \qquad \text{and} \quad T_0 = 263. \tag{5.48}$$

The glass discs were heated at oven rate to some temperature which was held for a period of time before the oven was turned off and allowed to cool naturally. The glass was considered to be sufficiently fluid to sag at temperatures in excess of 560°C, but the oven temperature was recorded every half minute while in excess of 540°C. Eight separate experiments were conducted for hold temperatures of 587 to 632°C and hold times of 2 to 30 minutes. It appears that sag measurements are of the top disc surface, rather than the bottom as for the disc-bending viscometry experiments. Note that there was no requirement for any part of the glass to be coated with a gold paint as in the disc viscometry experiments.

Gulati et al. [41] consider the clamped edge support to permit radial slip up to a temperature of 603°C, after which the glass sticks to the mould and the edge becomes immovable. We shall not make this distinction in our model, but assume a no-slip wall condition at $a = 32.77\,\text{mm}$ throughout the whole slumping process. Thus we have the axisymmetric version of the slumping viscous bridge problem of Chapter 3 (see Figure 3.1). As in the computations for disc-bending viscometry, let us non-dimensionalize using the radial span a and make use of symmetry in $r = 0$ so that the dimensionless flow domain is $0 \le r^* \le 1$, $0 \le z^* \le h/a$, where h/a is the aspect ratio and the axes are directed as in Figure 5.2.

This problem is one in which we can make use of the method described in Section 5.4 for temporally varying viscosity. For each case the experimental data consists of a record of temperature at half-minute intervals, and the sag of the top glass surface at the end of the slump. The temperature is considered to be constant in each half-minute interval after which it jumps to the next recorded value. The ith 30-second interval during which the temperature is T_i is then equivalent to a dimensionless time t_i^* that is determined by the glass viscosity μ_i given by the VFT

Table 5.4: Temperature versus time during a slumping experiment of 6 minutes at 585.9°C [41].

t min.	T °C	t min.	T °C	t min.	T °C
0.0	540.1	7.5	585.0	15.0	577.1
0.5	545.7	8.0	585.8	15.5	575.0
1.0	550.5	8.5	585.9	16.0	572.8
1.5	554.9	9.0	585.9	16.5	570.1
2.0	558.7	9.5	585.8	17.0	567.1
2.5	562.0	10.0	585.9	17.5	564.4
3.0	565.2	10.5	586.2	18.0	561.5
3.5	568.0	11.0	586.6	18.5	558.7
4.0	570.8	11.5	586.6	19.0	555.7
4.5	573.2	12.0	586.5	19.5	552.7
5.0	575.6	12.5	585.9	20.0	549.7
5.5	577.8	13.0	584.8	20.5	547.0
6.0	580.0	13.5	583.2	21.0	543.9
6.5	581.9	14.0	581.4	21.5	541.0
7.0	583.7	14.5	579.5	22.0	538.1

equation for that interval, and equation (5.38) giving the dimensionless slump time t^* reduces to the sum of all these t_i^*. Thus

$$t^* = \sum_i t_i^* = \sum_i 30 \frac{\rho g a}{\mu_i}. \tag{5.49}$$

Table 5.4 contains one set of Corning temperature data for a slump in which the temperature was held at around 585.9°C for 6 minutes. The dimensionless slump time given by (5.49) for this case is $t^* = 0.0032188$.

Next the finite-element program is run for a disc of aspect ratio $h/a = 1.5/32.77 = 0.0458$ for a dimensionless time period that covers all the cases that we wish to consider. From this we obtain the dimensionless top-surface sag δ^* at dimensionless times t^* as shown in Figure 5.5, from which we can find the sag $\delta = a\delta^*$ predicted by the creeping-flow model for each of the experimental slump cases described in [41].

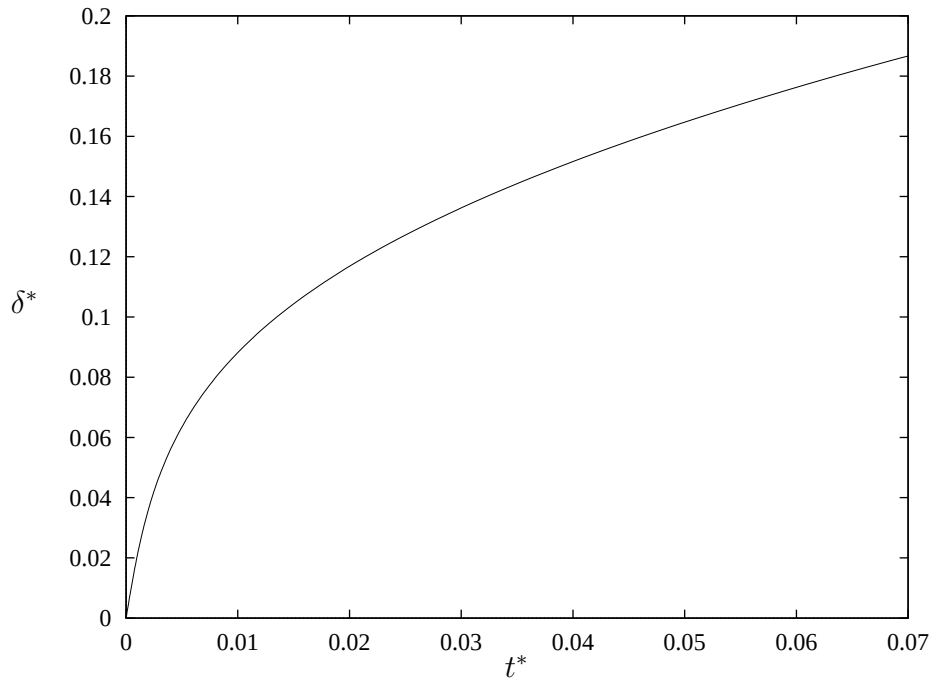

Figure 5.5: Sag versus time for an initially flat disc of aspect ratio $h/a = 0.0458$.

The results are given in Table 5.5 along with the actual measured sag δ_M and the sag δ_G computed by Gulati et al. [41] using viscoelastic analogy equations. The last column contains the differences between the experimental sag and that given by the creeping-flow simulation $\delta_M - \delta$. These show the creeping-flow finite-element simulation to be giving excellent results in all but the first and third cases. There is no obvious explanation for the magnitude of difference seen in these two cases, and it is attributed to fortuitous variations in experimental conditions. It seems unlikely to be a result of slip at the edge support while the temperature was below 603°C, which was not permitted in the creeping-flow simulation, since we would then expect to see similar differences in some of the other cases (e.g. the second to last case). In general, the agreement between the creeping-flow simulation and the experiments is as good as or better than that achieved by Gulati et al., which is very satisfying.

Table 5.5: Top-surface sag of a circular glass disc; computed and experimental values.

Case	δ_M (cm)	t^*	Creeping flow		δ_G (cm)	$\delta_M - \delta$ (cm)
			δ^*	δ (cm)		
16 min. 594.3°C	0.406	9.24897×10^{-3}	0.0856	0.281	0.365	0.125
17 min. 618.7°C	0.617	6.59802×10^{-2}	0.1844	0.604	0.564	0.013
25 min. 602.4°C	0.583	2.66751×10^{-2}	0.1313	0.430	0.412	0.153
30 min. 592.5°C	0.392	1.61873×10^{-2}	0.1082	0.355	0.477	0.037
2 min. 623.1°C	0.471	2.84330×10^{-2}	0.1345	0.441	0.412	0.030
24 min. 604.2°C	0.432	2.75517×10^{-2}	0.1329	0.435	0.408	-0.003
6 min. 585.9°C	0.230	3.12880×10^{-3}	0.0488	0.160	0.185	0.070
7 min. 628.8°C	0.566	5.29428×10^{-2}	0.1697	0.556	0.516	0.010

5.9 Sensitivity of Sag to Viscosity

Figure 5.6 compares the viscosity-temperature relationship for the glass used in the Corning disc-bending and sagging experiments as given by the VFT approximate equation, with constants (5.48), and the generally accepted beam-bending and parallel-plate experimental methods. While the two curves are close in this temperature range of interest, they do differ in some places by 10 to 20 percent. In using a VFT approximation to calculate sag of a disc as above, we raise the question of the sensitivity of slumping to accurate viscosity data and whether the small differences between computational and experimental results could be attributed to differences between the computational and true viscosities.

Accurate determination of the viscosity of molten glass is notably difficult, not least because of its large magnitude and the high temperatures at which measure-

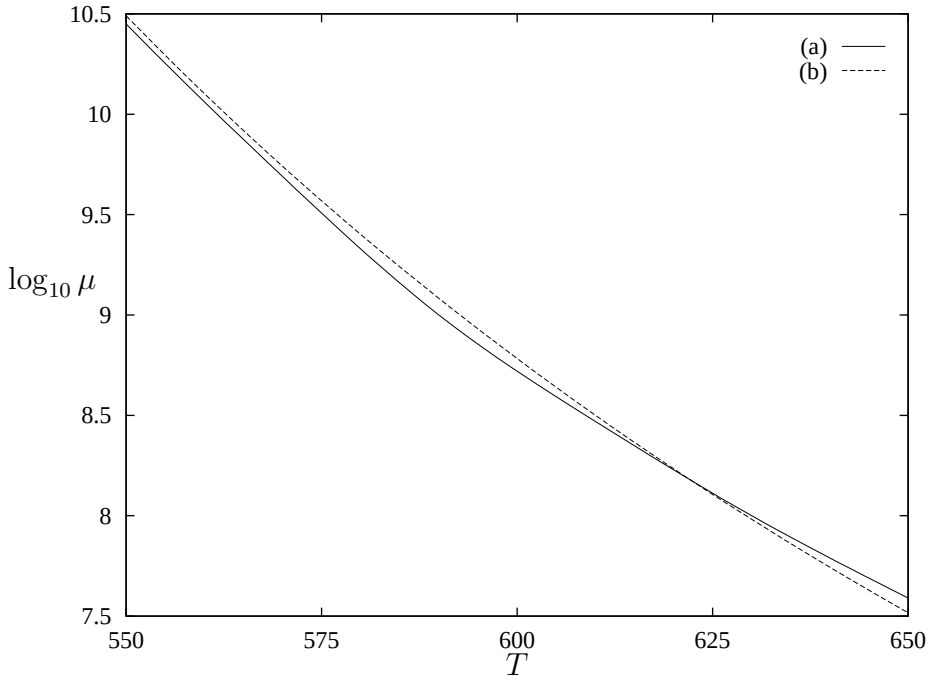

Figure 5.6: Dependence of viscosity on temperature for the glass used in the Corning experiments as given by (a) accepted methods, and (b) the VFT approximate equation with constants (5.48).

ments must be made. Then there are complications caused by time dependence and the consequent need to hold the glass at a temperature for a sufficiently long time period for it to reach equilibrium. Furthermore the large range through which the viscosity varies requires that different methods of measurement must be employed in different viscosity ranges, which do not give identical results in regions of overlap. Glasses with which we have been concerned have viscosities ranging between 10^4 and $10^{14.5}$ poise, i.e. over 10 orders of magnitude, for temperatures between 1000 and 500°C. It is also known that the assumption of incompressibility always used in calculating viscosity from measured data, results in viscosities around 20 percent lower than the true values [42], an error that is considered to be insignificant relative to the large magnitude of the viscosity.

Errors of 20 percent or greater are then quite possible in the viscosity data used for numerical simulation, and, since a primary application of the present work is

to the manufacture of optical surfaces, it is of considerable interest to determine the importance of even small inaccuracy in viscosity data. Also, because glass viscosity is a highly temperature-dependent property, an analysis of the sensitivity of slumping to viscosity will give some idea of the possible effects of inaccurate oven temperature. Practical difficulties in obtaining accurate oven temperatures are common in this type of process. By way of example, a temperature that is 10°C lower than expected can result in an 80 percent higher viscosity at temperatures around the softening point which are usual for slumping. For all of these reasons it is desirable to have some knowledge of the sensitivity of sag to viscosity.

It is clear from the preceding discussion of temporal viscosity changes in Section 5.4 that, provided the assumption of isothermal conditions is reasonable, any error in the viscosity data being used or the oven temperature will modify the time scale of the slumping problem, so that the required product can still be obtained by adjusting the real slump time.

To investigate the sensitivity of sag to viscosity, we again consider the initially flat circular glass disc of aspect ratio $h/a = 0.0458$ with a no-slip-wall edge support of the previous section. From the finite-element simulation of the sagging of this disc we have the (dimensionless) sag δ^* versus time t^* relationship shown in Figure 5.5. Now, suppose that we let a disc slump for a time t_s while the fluid has viscosity μ_1, and another similar disc slump for the same period of time with the fluid viscosity being μ_2. Then, we have non-dimensional slump times of t_1^* and t_2^* respectively, where

$$t_i^* = \frac{\rho g a}{\mu_i} t_s, \tag{5.50}$$

and for each of these times, we may find from the finite-element output the corresponding value of sag δ_i^*. A comparison of these gives the effect of viscosity upon the total sag of the disc for a fixed real slump time of t_s.

This effect is clearly shown by graphing δ^* versus $\log_{10} t^*$ as in Figure 5.7 and

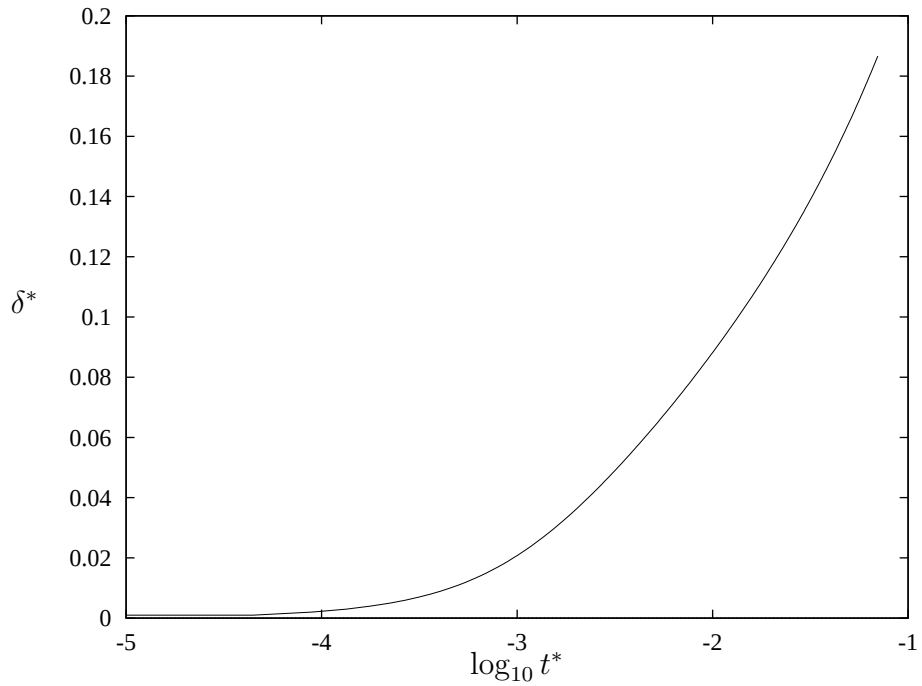

Figure 5.7: Sag versus the logarithm of time for an initially flat disc of aspect ratio $h/a = 0.0458$.

noting that the relation between μ and t^* may be written

$$\log_{10} \mu = \log_{10}(\rho g a t_s) - \log_{10} t^*. \tag{5.51}$$

Thus $\log_{10} t^*$ is in fact a scaled form of $\log_{10} \mu$, with an order of magnitude *decrease* in μ being equivalent to an order of magnitude *increase* in t^*, and Figure 5.7 immediately shows the change in sag resulting from a change in the viscosity. For the disc of aspect ratio $h/a = 0.0458$ under investigation, this is particularly significant in the region $\log_{10} t^* > -3$ which at normal slumping viscosities of 10^8 to 10^9 poise represents real slumping times in excess of 13 seconds to 2.2 minutes.

Now consider viscosity values that are about 20 percent too small compared with the true values, as is anticipated from the assumption of incompressibility. This represents an increase in $\log_{10} \mu$ or a decrease in $\log_{10} t^*$ of 0.08. Then a 2.2 minute slump at an apparent viscosity of 10^8 poise (equivalent to $\log_{10} t^* = -2$) is in reality a 2.2 minute slump at a true viscosity of 1.2×10^8 poise (equivalent to $\log_{10} t^* = -2.08$),

and from Figure 5.5 this represents a decrease in sag of about $0.007 \times 3.277 = 0.023\,\text{cm}$. The order of magnitude of this error is comparable to the differences in computational and experimental sags seen in Table 5.5, from which it is evident that even apparently small and acceptable errors in supplied viscosity data, or small inaccuracies in oven temperatures, can account for at least some of the observed variation between experiment and computational simulations. This variation does not seem to be particularly large, but could still be significant depending on the application of the slumped component, especially when it is appreciated that an error of less than one percent of the disc radius is more than 15 percent of the disc thickness.

Instead of finding the difference in sag for a fixed real slump time with a change in the viscosity, we can determine the required real slump time for fluids of different viscosity in order that some required sag be attained. For this sag δ^* the finite-element program gives the required dimensionless slump time t^* and we then compute the real times t_1 and t_2 taken for fluids of viscosity μ_1 and μ_2:

$$t_i = \frac{\mu_i}{\rho g a} t^*. \tag{5.52}$$

So if the fluid viscosity is in reality 20 percent greater than the supplied viscosity data used in our computations then we must increase our slump times by 20 percent, which is really quite substantial, especially for slump times around one hour in duration. The effect of inaccuracies in oven temperatures can also be ascertained and compensation made by adjusting the slump time.

5.10 Conclusion

In this chapter the creeping-flow finite-element program has been applied to the axisymmetric slumping of molten glass. Comparison of numerical simulations with results of experiments conducted during 1974 and 1975 at Corning Glass Works, in relation to disc-bending viscometry and the development of disc-sagging schedules,

show excellent agreement. In addition we have seen that differences between experiment and computation can probably be largely attributed to known inaccuracy in glass viscosity data, and this knowledge will prove useful later on in Chapter 8 of this thesis.

In the following chapters we continue with our analysis of glass slumping, but consider the process of slumping into a mould. This gives more control of the shape of the item, and enables the forming of more complicated parts that cannot be obtained with the type of free slumping we have so far considered.

5.11 A Note on Quadrature Rules for Axisymmetric Finite-Element Problems

The choice of quadrature rules for numerical evaluation of finite-element integrals was discussed in Chapter 2 (p. 27), with particular attention given to axisymmetric problems. These involve evaluation of a term $u_1 w_1/x_1$ which will cause a divide-by-zero error should any of the quadrature points at which it is evaluated fall on the symmetry axis $r = x_1 = 0$. For this reason a (four-point) quadrature rule with all quadrature points in the interior of the element has been used in all axisymmetric computations in this and following chapters. However, $r = 0$ is a symmetry boundary and all fluid particles lying on it have a zero horizontal velocity component ($u_1 = 0$) so that it is possible, but not obvious, that the terms $u_1 w_1/x_1$ will also be zero. The intriguing question of the value of this term at $r = x_1 = 0$ is mentioned, though not resolved, in [38], but an indication of the answer may be easily obtained by some simple numerical tests. As promised in Chapter 2, the results of these tests are given here as an aside.

The three-point rule given in Table 2.1 has all quadrature points lying on the element boundaries, but, because the integrands are cubic, does not preserve the order of convergence of the finite-element method as does the four-point rule. Hence

Table 5.6: Seven-point quadrature constants for triangular elements.

Order	*Error*	ζ_1	ζ_2	ζ_3	W_i
Cubic	$O(\ell^4)$	1	0	0	1/40
		1/2	1/2	0	1/15
		0	1	0	1/40
		0	1/2	1/2	1/15
		0	0	1	1/40
		1/2	0	1/2	1/15
		1/3	1/3	1/3	9/40
Quintic	$O(\ell^6)$	0.333333	0.333333	0.333333	0.112500
		0.797427	0.101287	0.101287	0.062970
		0.101287	0.797427	0.101287	0.062970
		0.101287	0.101287	0.797427	0.062970
		0.059716	0.470142	0.470142	0.066197
		0.470142	0.059716	0.470142	0.066197
		0.470142	0.470142	0.059716	0.066197

for a proper test we require a quadrature rule having points lying on the element boundaries of an order the same as, or higher than, the four-point rule. One that fits the bill, of the same order as the four-point rule, is a seven-point rule given in [1, p. 893], where there is also a seven-point rule of higher order (quintic) having only interior quadrature points. Since the four-point rule is sufficient to preserve the order of convergence of the finite-element method, this quintic rule should give a similar result despite its higher order. The quadrature points and weights for the two seven-point rules are given in Table 5.6.

To assess the different quadrature rules we consider only the initial vertical velocity and pressure at the centre of the top surface of a disc of aspect ratio $h/a = 0.2$ with a no-slip-wall support around its circumference. The computational results are shown in Table 5.7, where the quadratic three-point rule has also been included to show the loss of accuracy on using a rule of a lower degree than the integrands. Clearly we have the same order of accuracy from the four and seven-point rules

Table 5.7: Effects of different quadratures for axisymmetric problems shown by computing the initial velocity and pressure of a slumping glass disc.

Order	n	$-v/\mathcal{V}$	$p/\mathcal{P}$
Quadratic	3	0.522032	0.921454
Cubic	4	0.521911	0.949024
Cubic	7	0.521909	0.948881
Quintic	7	0.521910	0.948955

(five figures for velocity and three for pressure), including the seven-point cubic rule where the terms $u_1 w_1/x_1$ were assumed to be zero on the axis $r = x_1 = 0$ and hence ignored. By contrast the quadratic three-point rule has reduced accuracy (three figures for velocity and only one for pressure), as expected. Note that the quintic rule gave no better accuracy than the cubic rules, as also expected.

From these computational results we can conclude that the troublesome terms $u_1 w_1/x_1$ in axisymmetric problems do indeed have a zero and not some other finite value on the axis, and this knowledge enables the use of quadrature rules having quadrature points that lie on element boundaries. For the present work, however, the four-point rule is still an obvious choice because it requires less computation and the three-point rule does not preserve the order of convergence of the finite-element method.

Chapter 6

Mould Contact and Replication

6.1 Thermal Replication

Thermal replication is an industrial process used in the manufacture of aspheric optical surfaces. At present accuracy limitations restrict its application to products requiring no more than ophthalmic-optics quality [72] such as mirrors, and ophthalmic lenses and lens moulds [84]. The process is described by Smith et al. [84]. Once the required optical surface has been mathematically defined it is machined into a cast ceramic block which becomes the mould for the thermal-replication process. A glass workpiece, typically but not necessarily flat, having a well polished upper surface is placed on this ceramic mould, and the combination heated in an oven so that the glass softens and slumps into the mould replicating the mould surface on its upper free surface. This part of the process is illustrated in Figure 6.1. The upper surface of the glass that does not contact the mould is the critical optical surface. Unlike the lower surface which is affected by the roughness of the ceramic mould, it is very smooth and requires no polishing — a difficult and costly process for complicated aspheric surfaces. After cooling and annealing, the lower surface of the glass is finished according to its intended use. Sometimes the lower surface must be ground and polished, say if the component is to be used as a lens, but the re-

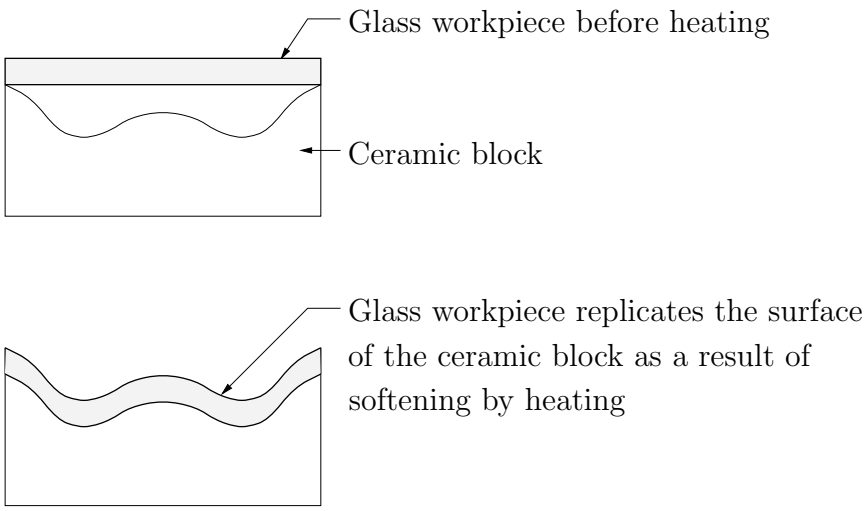

Figure 6.1: Thermal replication (after Smith et al. [84])

quired lower-surface shape will be simple (e.g. spherical or flat) so that polishing to give fine texture is not nearly as difficult or costly as for the complex upper surface.

The idea behind this manufacturing process is to replicate the macro-surface of the ceramic mould on the upper glass surface as nearly as possible, while smoothing out any small-scale imperfections in the mould surface. The upper glass surface does not, however, exactly replicate the mould surface, and consequently the design of the optical surface and its associated ceramic mould, involves an iterative process of slumping, measuring the resultant glass surface, and correcting the mould in order to achieve the desired optical surface at the required accuracy. The specification of the optical surface may also be adjusted based on the experimental results, to render it more easily formable. This iterative process consumes considerable time and resources. The investigations reported in this thesis have been largely directed towards developing a numerical modelling tool for the slumping process that gives a better understanding of the glass flow, and with the aim of greatly reducing, if not eliminating, the experimental iteration currently required during this design stage.

It has already been established in Chapter 5 that molten glass at normal slumping temperatures is a very viscous fluid that can be modelled by the creeping-flow

equations. At this stage we continue to assume isothermal conditions within the glass at any instant in time, which has proved to be reasonable in analyses performed so far. This is, however, now more subject to questioning, since the contact between molten glass and a mould that now occurs leads to an increased likelihood of spatial temperature (and hence viscosity) variation, due to differing thermal properties of glass and ceramic. Certainly for the similar processes of blow moulding [21, 68] and thermoforming [68], as well as the stamping of an initial parison for subsequent use in blow moulding [21, 83], in all of which a hot viscous film is forced against a comparatively cold mould, spatial temperature variations are likely to be quite important, and thermal modelling should be coupled with fluid-flow modelling [14, 21, 83, 100]. In slumping however, both mould and glass are heated together from room temperature, so that spatial temperature variations will be considerably less than for these other processes. Thus the isothermal assumption may well be quite adequate, and its retention is justified, at least for an initial investigation of slumping into a mould. On the other hand, with the accuracy required for optical-quality surfaces such spatial variations as do exist might still be of some importance, and we shall look at relaxing the isothermal assumption in Chapter 7.

In considering the slumping of a glass workpiece into a mould, our particular interest is in the manufacture of glass moulds which are then used in the casting of plastic ophthalmic lenses, specifically lenses of continuously varying focal length, commonly known as "multifocal" or "progressive-power" lenses [84]. A slumped glass mould having a complex surface of varying curvature (or "power") is used as one half of the mould for a plastic lens. Another glass mould of simple shape, made by conventional grind-and-polish processes, is positioned close to the slumped mould so that there is a small gap between them, and the edge is sealed by means of a rubber gasket or similar. Liquid polymer is then injected into this gap and solidifies to form the lens. The slumped half of the mould gives the complex varying curvature (usually) to the outside surface of the lens, while the other lens surface

(usually) nearest to the wearer's face is determined by the simple portion of the mould. The complex surfaces of progressive-power lenses are typically described three dimensionally. However axisymmetric progressive-power lenses are also possible, and axisymmetric modelling is also useful for gaining an understanding of the glass flow in fully three-dimensional situations. Thus, we continue to consider axisymmetric slumping problems. Fully three-dimensional problems are, in principle, handled as easily, once the more difficult issues of of mesh generation for both glass workpiece and ceramic mould have been resolved.

In the initial stages, the slumping flows now under consideration and those considered in Chapter 5 are essentially identical, save only for possible differences in edge support conditions. The major difference arises when the lower surface of the slumping glass contacts the mould, and to model this requires a modification to the finite-element program. The method employed, as explained in the next section, is fundamentally the same as that used in [14, 21, 38, 100] for modelling contact in blow-moulding problems. We have already noted in Chapter 5 the similarity between slumping into a mould (as here in thermal replication) and blow moulding.

In this chapter, in addition to a consideration of this contact problem, interest in the slumping of ophthalmic-lens moulds, for which the curvature of the surface is the critical measure, leads to another interesting computational problem. In order to determine surface curvature, derivatives must be calculated from non-exact numerical data. This is an area of on-going research, e.g. by Anderssen and de Hoog [2] using finite-difference methods. In Section 6.4 we shall consider some other methods more easily applied to non-uniformly spaced data, such as the coordinates of all nodes lying on the top glass surface from which surface curvature must be determined.

A considerable portion of the work detailed in this chapter has been summarized in a publication [88].

6.2 Mould Contact

In order to keep computations relatively simple and fast, we shall follow previous work on finite-element simulation of blow moulding [14, 21, 38, 100] in representing a mould surface quite simply as a set of points connected by straight lines. The accuracy of the mould representation can be improved by increasing the number of points, and by increasing the density of points along segments of rapid change in the mould shape.

When slumping into a mould, the lower free surface of the glass contacts the mould surface as the flow progresses, but cannot, of course, pass through it. Up until the first contact occurs, the time-evolution of the flow geometry may be determined as previously by solving the slumping dynamics equations (2.19) at each node using the known velocity field and the fourth-order Runge-Kutta method (2.53) to give the new node position. The mesh is then adjusted by moving each node to its new position. Eventually, however, the point is reached where one or more nodes will cross the mould boundary. This is illustrated for node n in Figure 6.2, which for the specified time step Δt will move along the vector $\Delta\tilde{x}_n$ and so 'pass through' the mould surface.

The obvious method for preventing nodes from passing through the mould surface, and which is quite typical in modelling a variety of forming processes [14, 21, 38, 100, 104], is to reduce the time step as necessary so that the node(s) that would otherwise first cross the boundary will just reach it. When a node reaches the mould boundary, it is fixed there for all future time. However, use of the Runge-Kutta method for time advancing the mesh, because it requires velocity computations after a node has moved from its current position, prevents exact a priori determination of the new position of a node for any given time step. Thus it is not known for certain whether, or by how much, a node will cross the mould boundary in the next time step Δt, and hence whether, or by how much, to reduce the time step. A similar, though not identical, problem is encountered in modelling of blow

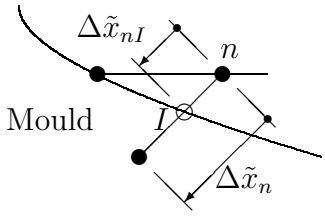

Figure 6.2: Contacting a mould.

moulding by Burley and Graham [14] and Graham et al.[38], in which the partial time-derivative term $(\partial u_i/\partial t)$ from the Navier-Stokes equations is included in the slow-flow equations used, and the chosen method of solution requires that velocities at both the ith and $(i+1)$th time steps be first calculated in order to determine the node positions at the $(i+1)$th time step. The method described in [14] to overcome this difficulty suggests a possible method for handling the current problem, which is also hinted at in [104], a paper describing numerical methods for a general range of forming processes, not specifically aimed at glass.

First, compute the coordinates of the point I (refer Figure 6.2) at which node n will contact the mould if it travels with the known velocity $\tilde{u}_n$, obtained from the finite-element solution of the creeping-flow equations, for a sufficiently long period of time. From this obtain the vector $\Delta\tilde{x}_{nI}$, the maximum permissible displacement of the node. Then estimate the size of the time step Δt_n that will take node n along vector $\Delta\tilde{x}_{nI}$ to the point of intersection with the mould surface I, using the Euler equation (2.51) thus:

$$\Delta t_n = |\Delta\tilde{x}_{nI}|/|\tilde{u}_n|, \tag{6.1}$$

where the modulus denotes the usual length of a vector. Compute this Δt_n for each node $n = 1, 2, \ldots, N$, where N is the total number of nodes not already fixed on the mould, and use the minimum of Δt and all these Δt_n as the time step with which to compute node displacements by the Runge-Kutta method. Return any node that still crosses the mould boundary to the point at which its path intersects that boundary. Once a node reaches the boundary fix it there for all future time.

This method results in some small material loss, caused by returning nodes to the mould boundary. There is also some computational overhead in checking for and computing intersection points of nodes crossing the boundary, not once but twice. However, the major disadvantage in the method is that the value of the Runge-Kutta method in giving excellent accuracy with quite large time steps, as demonstrated earlier in a comparison with the Euler method for the sagging bridge of Chapter 3 (see Table 3.7), is largely lost because of the need to reduce the time step to prevent material loss through the mould. With small time steps and using the Runge-Kutta method, the computational time becomes enormous, and this reduces the intended benefit of the finite-element program as a fast predictive tool for industry.

Because of the large increase in computational time brought about by the introduction of mould contact into the model, the alternative Euler method (2.51) is adopted for sagging problems involving mould contact. In addition, this method is quite typical for problems similar to the present one [21, 100, 104]. Although there is some loss of accuracy in the initial stages of the flow prior to mould contact, by comparison with the Runge-Kutta method, the amount of free sag is usually small, of the order of the glass thickness, so that the error is less than shown for the sagging bridge example in Table 3.7 where the two methods were compared after a much larger sag allowing much more accumulation of error. Of course, a sufficiently small time step must be used so that the flow domain area (i.e. mass) is reasonably-well conserved. Once the lower surface of the glass begins to contact the mould this time step becomes an upper limit on the size of the time step, which is reduced as necessary. Decreasing the time step acts to increase the accuracy of the Euler method. Note that, while not included in the above discussion, the time step is also reduced as necessary to keep the displacement of all nodes at, or below, a prescribed limit, as explained in Section 4.5 of Chapter 4.

Another positive benefit of the Euler method is that (6.1) gives the actual time step required to take node n to the mould boundary, rather than just an estimate,

and hence we no longer have the problem, associated with the Runge-Kutta method, of being unable to predict exact node displacements. Because of this there is no need to check for nodes crossing the mould boundary, or to move them back to the boundary, which further increases computational speed.

Thus the algorithm is as follows . First solve for the velocity and pressure fields in the current geometrical configuration of the flow domain using our finite-element program. Set the time step Δt to the maximum acceptable value, and also specify some maximum allowable displacement $|\Delta\tilde{x}|_{\max}$. Then, for each node $n = 1, 2, \ldots, N$ not already fixed on the mould surface

1. compute the time Δt_n in which node n will move a distance $|\Delta\tilde{x}|_{\max}$ using an equation similar to (6.1);

2. $\Delta t = \min\{\Delta t, \Delta t_n\}$, so restricting the displacement of any node to the maximum allowed;

3. determine if the path of node n will intersect with the mould boundary in time Δt; if no intersection, increment n and go to step 1 otherwise find the intersection point I;

4. compute the distance $\Delta\tilde{x}_{nI}$ (see Figure 6.2) to the mould boundary;

5. compute Δt_n from (6.1);

6. $\Delta t = \min\{\Delta t, \Delta t_n\}$, so reducing the time step if necessary.

After this loop the value of Δt is the maximum time step possible such that no node will cross the mould boundary, but at best just reach it. The displacements of all nodes in this time step is then computed, and their positions adjusted. Any node that contacts the mould is fixed on the mould by changing its boundary specification to that of the mould, i.e. a no-slip wall. Note that nodes can be prevented from crossing symmetry boundaries in a similar manner.

This algorithm gives a fast and effective method for time advancing the flow domain geometry in problems involving contact with a mould. A change to the computationally costly, but probably more accurate, Runge-Kutta method can be easily made in the future should that be desirable, requiring only a minor modification to the program. Albeit, it is worth mentioning that any apparent increase in accuracy is not necessarily very meaningful in the context of other sources of error in the results. We have already considered the effect that small errors in viscosity data relative to the overall magnitude of the viscosity could contribute. Another more important point, that we shall consider in some detail a little later, concerns the measurement of curvature. For lens moulds, it is not the amount of sag that is critical, but the curvature profile of the upper glass surface. Thus, in optical applications, the benefit of the model and simulations is finally determined by the accuracy with which curvature can be computed using the non-exact discrete surface-coordinate data provided by the finite-element program. This is a difficult task which is as yet subject to inaccuracies that far outweigh the error introduced by use of the Euler method in place of the Runge-Kutta method.

6.3 Numerical Slumping: An Example

In this section the method outlined above is illustrated by taking an axisymmetric arrangement of an initially-flat glass disc supported on a concave mould as shown in Figure 6.3. For clarity the vertical scale is twice the horizontal scale. The mould surface may have an aspheric profile, though for present purposes it has been chosen to be spherical in order to more clearly show the differences that slumping produces in the top free surface of the glass compared with the mould surface. This problem is non-dimensionalized using the disc radius a in the usual way. The disc has an aspect ratio of $h/(2a) = 0.0667$ where h is its thickness. The radius of curvature R of the mould surface is $R/a = 5.3293$ giving an initial clearance l between mould

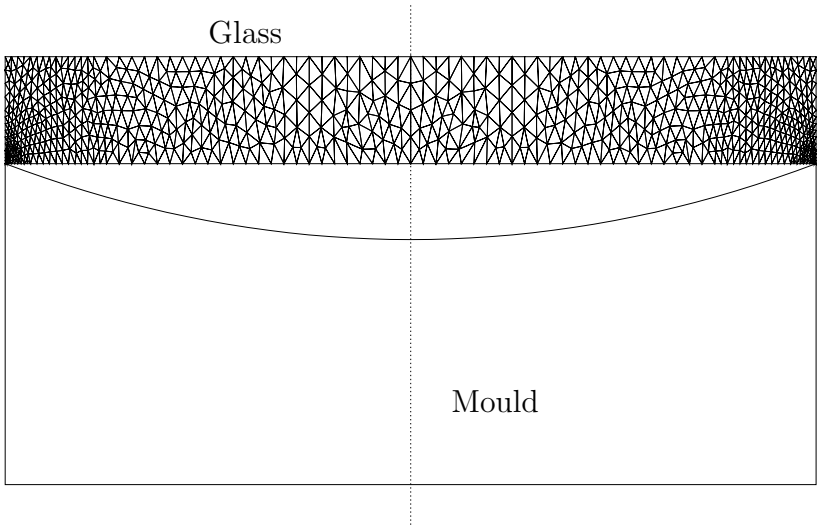

Figure 6.3: Geometry prior to slumping ($t = 0$).

and glass at the centre of $l/a = 0.0947$. The dimensions of this particular problem are based on the slumping of real optical components. Symmetry about the axis $r = 0$ limits the computational domain to a radial cross section of the disc, over which a mesh of 1494 nodes and 695 quadratic triangular elements is defined.

The initial stages of the flow prior to mould contact proceed quite rapidly (though in real terms it may take in excess of one hour), and in order to capture the disc geometry at, or soon after, first mould contact a maximum dimensionless time step of $\Delta t^* = 0.001$ was first used. Then we see contact occurring first at the mould centre at a time of about $t = 0.014\mathcal{T}$, as shown in Figure 6.4. As in previous problems the time scale is $\mathcal{T} = \mu/(\rho g a)$.

Mould contact greatly reduces the rate of flow, and hence to compute the disc geometry at much later times the program was run again using a larger maximum time step of $\Delta t^* = 0.005$, out to a time of $t = 0.15\mathcal{T}$. Although this time represents 30 steps of $\Delta t^* = 0.005$, because of the reduction of the time step for handling mould contact, in actual fact 110 steps were taken of size $\Delta t^* \leq 0.005$. Note also that at a time of $t = 0.015\mathcal{T}$, when the first run was stopped, the disc sags given by both runs agree to three figures, or four decimal places. Full mould contact is achieved,

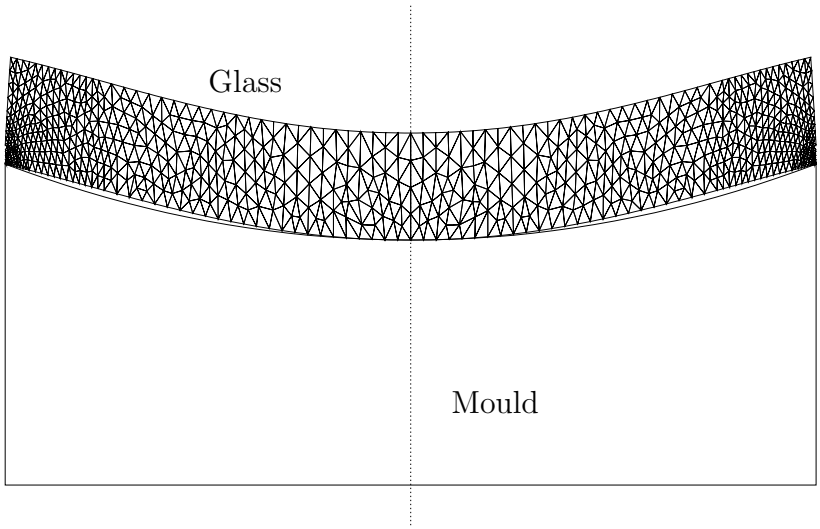

Figure 6.4: Initial mould contact at $t = 0.014\mathcal{T}$.

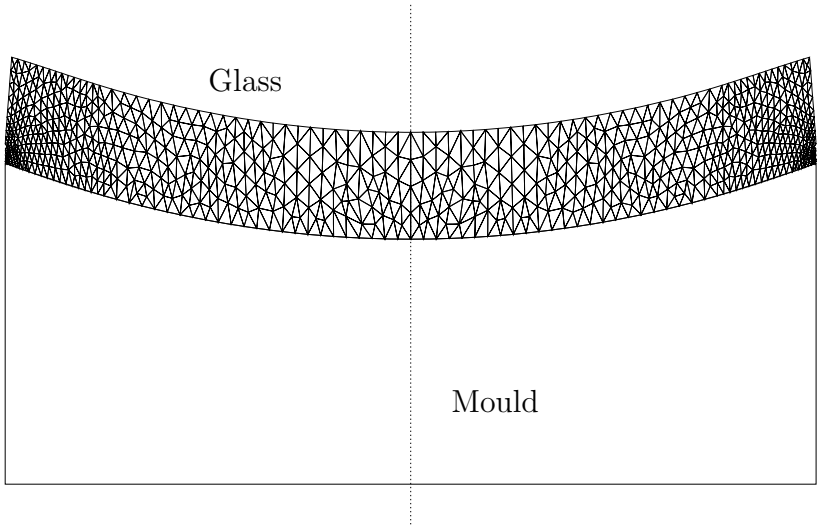

Figure 6.5: Full mould contact at $t = 0.075\mathcal{T}$.

as shown in Figure 6.5, at a time of about $t = 0.075\mathcal{T}$. The first stage of the flow, when there is most slumping action up to initial contact with the mould, takes place in less than 20 percent of the time taken to establish full mould contact. After full mould contact slumping effectively ceases, with further flow taking many hours or even days. Computation time for this run was about 40 minutes using a single 167 MHz processor on a Sun Ultra Sparc 170 computer having 256MB of RAM. Full mould contact was established when the run was about 60 percent complete.

Having computed the slumping of a glass disc into a mould up to full mould contact, we now wish to determine how well, or poorly, the mould surface has been replicated on the top glass surface. It is not possible to determine this at all adequately from Figure 6.5, and hence, for this purpose, we next look at determining the curvature profile of this surface. Surface curvature is a critical measure for optical surfaces, and must meet the specification to within some tolerance which is governed by the product application. In computing curvature, or indeed any other quantity, from numerical data it must be borne in mind that this data cannot be accepted as exact, but is subject to error due to the discretization methods used and a computer's inability to store real numbers to infinite precision.

6.4 Calculating Curvature

The fundamental definition of curvature is [33, p. 553] $\kappa = d\theta/ds$ where s denotes distance along the curve, and θ is the angle made by the tangent to the curve at point s and the x-axis. For the axisymmetrical slumping example just considered, the curve of interest is the top glass surface defined by vertical height $z = f(r)$ where f is an (unknown) function of radial position r. Then the equation for curvature becomes

$$\kappa = \frac{f''}{(1 + f'^2)^{3/2}} \tag{6.2}$$

where primes denote differentiation with respect to r. As output from the finite-element program we obtain (r, z) coordinates of nodes located on, or (allowing for numerical error) close to, this surface, and from these discrete and non-exact data points we need to compute the first and second derivatives of f with respect to r to give curvature, and so obtain a quantitative measure of small-scale changes in surface shape.

The plainly difficult task of accurately differentiating non-exact data, such as is obtained from any discretized computation process, is an area of current research. Finite-difference methods for numerical differentiation have received considerable attention over many years, and this is still an area of active investigation [2, 3]. Such methods, however, generally require that the available data is uniformly spaced, a condition that is typically not satisfied in our problems. As an alternative to finite-difference techniques, splines may be fitted to the (r, z) data, which give the value of the function f as well as its derivatives at positions r [26, 63]. The degree of the splines chosen will determine the extent of continuity and hence the order of derivatives that may be obtained. While interpolating splines may be used, for data that is known to be non-exact it may be preferable to perform a least-squares fit, and so smooth the data. Least-squares B-spline approximations are recommended by de Boor [26, p. 249] as "very effective" for filtering noise from non-exact data.

A very heuristic approach has been taken here to the subject of curvature calculation, with a number of different methods tried to see what 'works' best. Before attempting to compute surface curvature for the slumping example of the previous section, let us consider computing the curvature of a surface such as we might expect from that slump, namely, a spherical cap of (dimensionless) unit radius and with radius of curvature $R = 5$ centred at $(r, z) = (0, R)$. This axisymmetric surface has constant curvature of $\kappa = 0.2$. Let us define this surface by 51 equally spaced points along the radial arc length $[0, 1]$, or 101 equally spaced points across the diameter $[-1, 1]$. Coordinates calculated in double precision are close to exact, and hence are

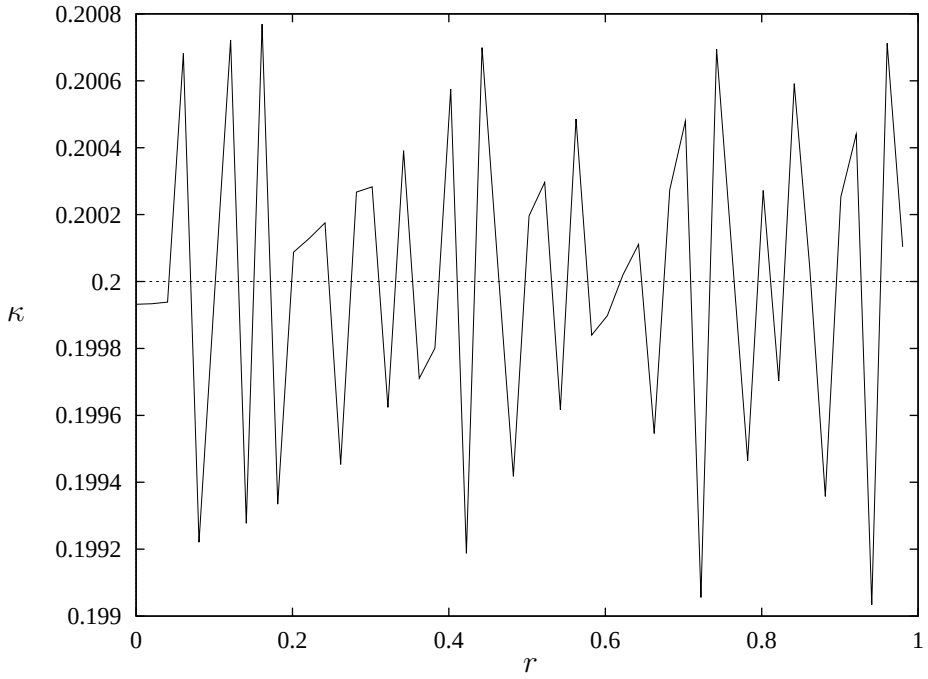

Figure 6.6: Curvature of a spherical cap with $\kappa = 0.2$ found by computing arcs through sets of three consecutive data points, with non-exact single precision data.

refered to as *exact* data points, although this is not strictly accurate terminology. Coordinates calculated in single precision are then *non-exact*, relative to the double-precision data. Note that the data is not uniformly spaced with respect to the r coordinate.

Probably the most intuitive method of calculating curvature, especially for this type of test data, is to find the radius R of the arc that passes through three consecutive data points, and then let $\kappa = 1/R$ at the central point of the three. The curvature profile for the spherical cap obtained using this method with exact (i.e. double precision) coordinate data is very accurate, but with non-exact (i.e. single precision) coordinate data the result is quite unstable as shown in Figure 6.6. From this it is clear that the method is quite sensitive to small errors in the data, as is to be expected because of the extremely localized nature of the data used in the computations.

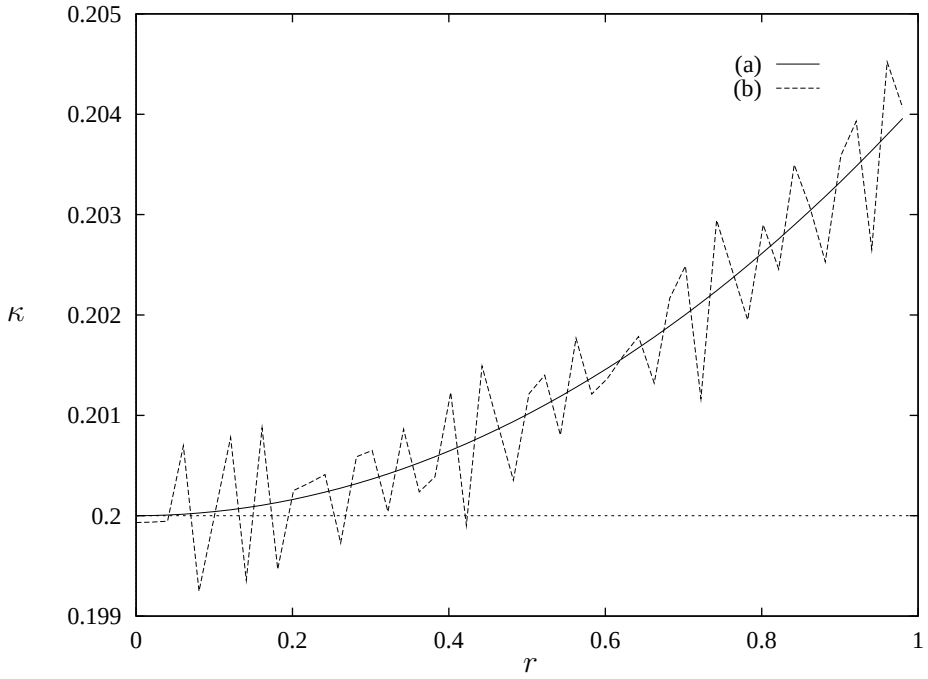

Figure 6.7: Curvature of the spherical cap computed using a finite-difference method and sets of three consecutive data points, with (a) exact and (b) non-exact data.

Next we consider a simple method based on finite differences, again using three consecutive data points. Because the data is not uniformly spaced in r the usual centred difference formulae cannot be used. Instead let us derive some formulae as follows. Let the three data points be denoted (r_i, z_i), $i = 1, 2, 3$, and let them be ordered such that $r_i < r_{i+1}$. Let us define $\Delta r_1 = r_2 - r_1$ and $\Delta r_2 = r_3 - r_2$. Then Taylor expansions about r_2 give two equations in the two unknowns z_2' and z_2'', where primes denote differentiation with respect to r,

$$z_1 = z_2 - \Delta r_1 z_2' + \frac{\Delta r_1^2}{2} z_2'' - \mathrm{O}(\Delta r_1^3) \qquad \text{and} \tag{6.3}$$

$$z_3 = z_2 + \Delta r_2 z_2' + \frac{\Delta r_2^2}{2} z_2'' + \mathrm{O}(\Delta r_2^3). \tag{6.4}$$

Combining (6.3) and (6.4) to eliminate z_2' gives

$$z_2'' \approx \frac{2}{\Delta r_1 + \Delta r_2} \left(\frac{z_3 - z_2}{\Delta r_2} - \frac{z_2 - z_1}{\Delta r_1} \right), \tag{6.5}$$

which substituted into (6.3) gives

$$z_2' \approx \frac{1}{\Delta r_1 + \Delta r_2}\left(\frac{\Delta r_2}{\Delta r_1}(z_2 - z_1) + \frac{\Delta r_1}{\Delta r_2}(z_3 - z_2)\right). \tag{6.6}$$

Now with these expressions for the first and second derivatives at r_2 we can compute curvature at r_2 from equation (6.2). A plot of the result for the spherical cap with exact and non-exact coordinate data is given in Figure 6.7. Again the method is sensitive to small errors in the initial coordinate data. It is surmised that the curve for the exact data bends upwards due to increasing inaccuracy in calculating the second derivative ($f''(r)$) as it increases more rapidly towards the edge of the cap, indicating that this method is not particularly appropriate for this type of data, no matter how close to exact the data is.

Aside from the obvious problem we see on calculating derivatives with this finite-difference method and our particular data, it is a known fact that standard centred difference formulae are generally too unstable for computing second and higher order derivatives of non-exact data, even when that data is quite precise [3]. Thus the unstable oscillatory behaviour that we see in the curvature (which is basically the second derivative) computed using single precision data, is to be expected. Because of this instability, Anderssen, de Hoog and Hegland [3] have developed a finite-difference method for higher-order differentiation that effectively gives derivatives of averaged data. This is, however, for data that is uniformly spaced (in r), and cannot be directly used with non-uniform data. Some work has been undertaken by Hegland [45] to adapt the method for our type of data, but for this thesis a different method, namely least-squares fitting of B-splines [26, Ch. 14] to the data, has been used. By using a least-squares approximation instead of interpolation, non-exact data is also averaged.

The construction of B-splines for a given data set is described in [26], and only a brief summary is given here of the method employed. Suppose we have a curve (or axisymmetric surface) defined by N discrete data points (r_i, z_i), $i = 1, 2, \ldots, N$ that we wish to represent by a degree-k B-spline. We define a set of points, or "knots",

$\{t_i,\, i = 1, 2, \ldots, M\}$ on the interval $[r_1, r_N]$ such that

$$r_1 = t_1 = \ldots = t_{k+1} < \ldots < t_j < t_{j+1} < \ldots < t_{M-k} = \ldots = t_M = r_N.$$

The degree-k B-spline is defined on these knots by the recurrence relation

$$\mathrm{B}_i^k(r) = \left(\frac{r - t_i}{t_{i+k} - t_i}\right) \mathrm{B}_i^{k-1}(r) + \left(\frac{t_{i+k+1} - r}{t_{i+k+1} - t_{i+1}}\right) \mathrm{B}_{i+1}^{k-1}(r) \tag{6.7}$$

for $i = 1, 2, \ldots, (M - k - 1)$ and $k \geq 1$, with

$$\mathrm{B}_i^0 = \begin{cases} 1 & t_i \leq r < t_{i+1} \\ 0 & \text{otherwise} \end{cases}. \tag{6.8}$$

Then, the curve is approximated by

$$z = \mathrm{B}^k(r) = \sum_{i=1}^{D_f} \alpha_i \mathrm{B}_i^k(r), \tag{6.9}$$

where the α_i are coefficients to be determined and $D_f = M - k - 1$ is the number of degrees of freedom, i.e. the number of coefficients to be determined. Substituting the N discrete data points (r_i, z_i) in (6.9) gives N equations in the D_f unknown α_i, and, assuming for the moment that these N equations are linearly independent, they can be solved if $N \geq D_f$. If we have $N = D_f$ then there is a unique solution that interpolates the data. If we have $N > D_f$ then we solve using the method of least squares to get the spline of best fit. Choosing the (internal) knots such that $t_i \leq r_i < t_{i+k}$ for all $i = 1, 2, \ldots, N$ ensures linear independence of the N equations [26, pp. 200, 254], although this is allowed to 'come out in the wash' and knots are not specifically selected to satisfy this condition. Note that B-spline fitting to data is very similar to the finite-element method in one-dimensional space.

From this description it is apparent that, even for interpolation, there is not a unique B-spline for a particular set of (r, z) data, since the knots t_i can be chosen in many different ways, although a B-spline with M knots does not change much with small changes in the knot placements [26, p. 126]. It is common to choose the $M - 2k - 2$ interior knots t_i, $i = k + 2, \ldots, M - k - 1$, to coincide with r-coordinates

of data points, as does the NAG library routine (E02BEF) which computes a cubic B-spline for a supplied set of data points. Even then, for $k > 1$ there is some choice concerning which r_i should be coincident with knots, since the number of interior data points $N-2$ always exceeds the number of interior knots $M-2k-2 \leq N-k-1$, the maximum number $(N-k-1)$ of knots being required by an interpolating spline. For example, a cubic ($k = 3$) interpolating spline requires $N - 4$ of the $N - 2$ interior data-point r-coordinates, so that two of these must be omitted from the knot sequence. The NAG routine omits the first and last of the interior r_i.

Now, the position of knots affects the continuity of the B-spline approximation. Specifically, if there are n knots at coordinate r, then there are $k+1-n$ conditions of continuity there, for a spline of degree-k. That is, the function value and its first $k - n$ derivatives exist and are continuous [26, p. 114][1]. At the end points of the curve we have $n = k + 1$ knots, and hence no conditions of continuity, but this is of no consequence provided that we compute over the complete curve and make no use of symmetries. However, in the interior of the curve we require that the curvature be continuous, that is a continuous second derivative. Therefore in this region we require a minimum of three conditions of continuity, i.e. that the function value $z = \mathrm{B}^k(r)$ and its first two derivatives be continuous. Preferably, the curvature should be smooth and differentiable, in which case we must also require that the third derivative be continuous. Since, for the B-spline method employed, multiple knots are not permitted in the interior of the curve, $n = 0$ or 1 on the interval (r_1, r_N) and there will be adequate continuity for all splines of degree $k \geq 3$. For a cubic spline (with $k = 3$) the second derivative will be continuous but not differentiable at knot locations, so that, while we can compute curvature, the profile obtained will not be smooth, i.e. differentiable, at knots. Using splines of higher degree (> 3) ensures a smooth and differentiable curvature profile. Splines of degree 2 or less cannot be used.

[1]Note that k in [26] is equivalent to $k + 1$ in our terminology

Let us now compute the curvature of the spherical cap used earlier, using B-spline fitting. We first try the NAG library routine (E02BEF), already mentioned in the above discussion, which computes a cubic B-spline for an arbitrary set of data points. In line with the above comments on continuity, we must fit our B-spline over the full diameter $[-1, 1]$ of the cap, defined by 101 data points spaced so that there is an equal arc length between any two consecutive points. It is possible to assign different weights to the data points, but all are given an equal unit weight. The value of a positive parameter s, known as the smoothing factor, controls the smoothness of the spline and its closeness of fit to the data, with a zero value specifying that the spline interpolate the data. The knots of the spline are located automatically, but are always some subset of the set of r-coordinates of the data points. The smoothing factor s also influences the number of degrees of freedom D_f in the B-spline approximation, with $D_f = 4$ being the minimum possible for a cubic spline corresponding to no interior knots, and $D_f > 4$ implying interior knots. When $s = 0$ then $D_f = N$ and the data is interpolated, and as s increases D_f decreases until $D_f = 4$ and there are knots at the end points only. In this case, the spline is simply a cubic polynomial, with four continuity conditions everywhere in the interior of the cap so that the curvature profile is smooth. For the spherical cap, it was found necessary to choose $s < 10^{-6}$ for $D_f > 4$. Having solved for the coefficients α_i, another NAG routine (E02BCF) evaluates the spline and its first three derivatives at given points r along the curve, from which results curvature can be computed. The result for the spherical cap given in this manner is shown in Figure 6.8 for different values of s (and numbers of degrees of freedom D_f). Exact (double precision) data was used for these computations, and we see the curvature profile approaching the known straight line solution $\kappa = 0.2$ as D_f becomes larger and the interpolation B-spline result approached. Note the lack of differentiability at knot points. For non-exact data, the curvature profiles obtained with small D_f (≤ 12) are virtually identical to those ((a) to (d) in Figure 6.8) for the exact data. However, when D_f

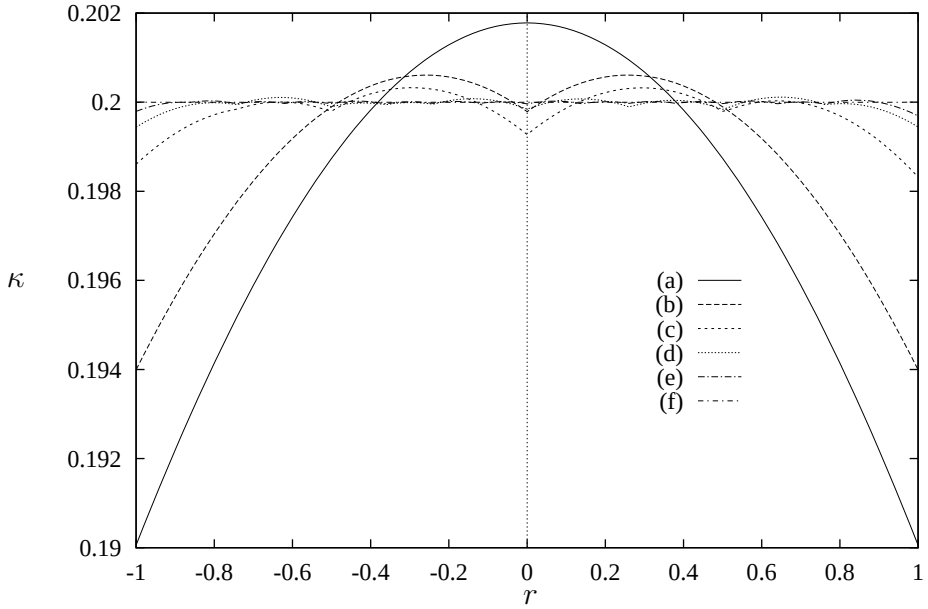

Figure 6.8: Curvature of the spherical cap found using NAG routines. Exact coordinate data was used with (a) $s \leq 10^{-6}$, $D_f = 4$, (b) $s = 10^{-7}$, $D_f = 5$, (c) $s = 10^{-9}$, $D_f = 8$, (d) $s = 10^{-11}$, $D_f = 12$, (e) $s = 10^{-13}$, $D_f = 18$, and (f) $s \geq 10^{-16}$, interpolation ($D_f = 101$).

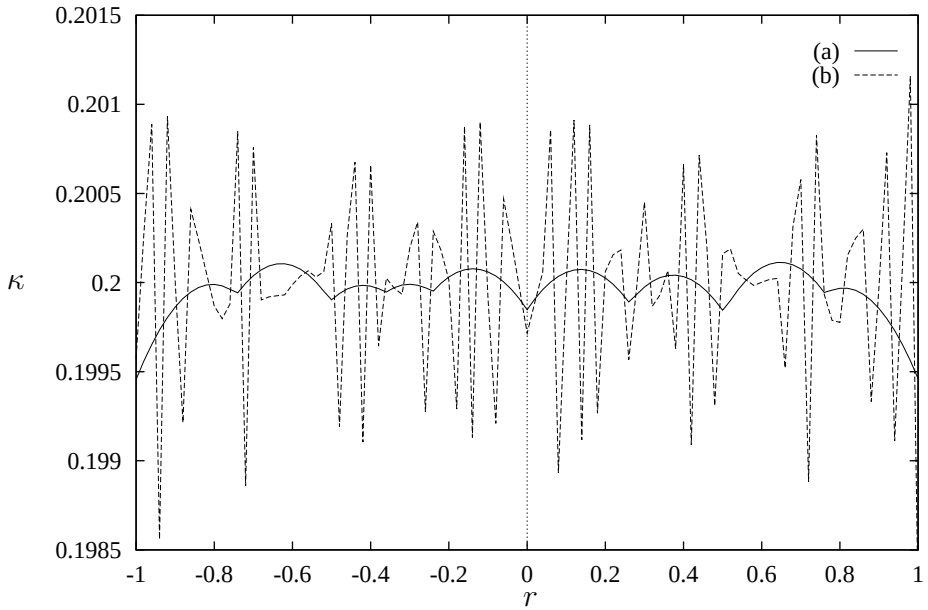

Figure 6.9: Curvature of the spherical cap found using NAG routines. Non-exact data was used with (a) $s = 10^{-13}$, $D_f = 12$ and (b) $s = 10^{-13}$, $D_f = 76$.

becomes large, approaching interpolation, the curvature profile for the non-exact data becomes substantially affected by the error in the data and fluctuates greatly, as shown in Figure 6.9. From this we can expect there to be some optimum number of degrees of freedom, between 12 and 76, that gives enough averaging of non-exact data for a satisfactory determination of curvature. With less than this optimum, there would be too much averaging and the loss of surface features. With more, there would be insufficient averaging with the curvature computations being too much influenced by inaccuracies in the data.

In order to consider higher degree splines, and have more control over the number of degrees of freedom in the spline approximation, a B-spline routine was written. For the spherical cap it seems reasonable to use uniformly spaced interior knots in the interval (-1,1) over which the curve is defined. D_f is varied to obtain splines with more or less averaging of the data, including interpolating splines with no averaging. By increasing the D_f and comparing the curvature profiles for exact and non-exact data the value of D_f is found above which error in the data starts to substantially affect our computations. Let us consider this value to be the optimum number of degrees of freedom. A similar methodology for determining an appropriate level of averaging of the data, using many fewer degrees of freedom than there are data points, is described by de Boor [26, Ch. 14].

Figure 6.10 shows this process for cubic (degree-3) B-splines. For clarity the plot is limited to the interval $[0, 1]$. Because of symmetry in $r = 0$, the interval $[-1, 0]$ is virtually a mirror image of this, as is seen in earlier plots. The curves for exact and non-exact data are virtually identical for $D_f = 14$, vary a little for $D_f = 19$, and are quite different for $D_f = 24$ with the error in the data beginning to have a substantial affect. From this it may be considered that around 19 degrees of freedom is optimum.

We now consider higher degree quartic ($k = 4$) and quintic ($k = 5$) B-splines. These show the same type of behaviour as the cubic splines, but the curves are

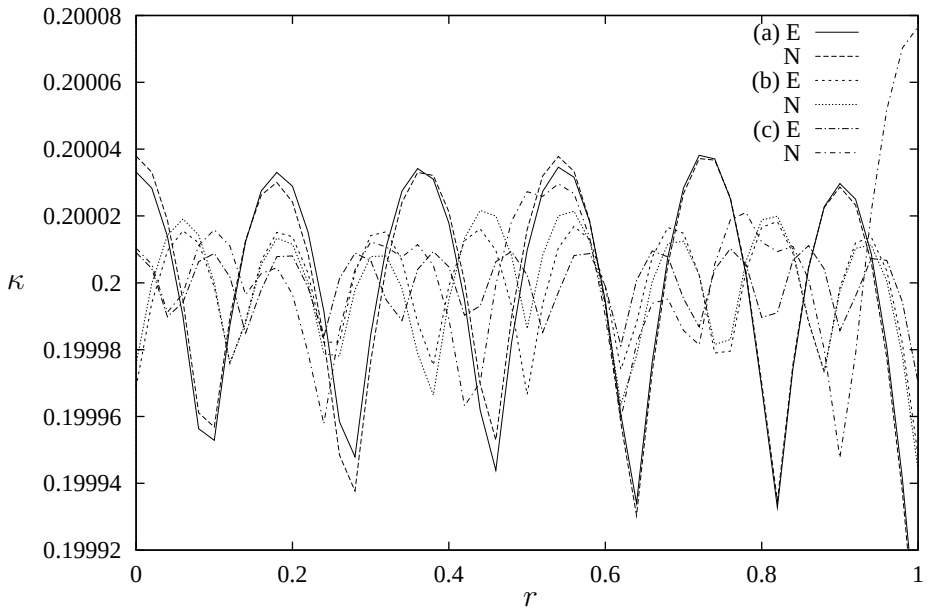

Figure 6.10: Curvature of the spherical cap found using a cubic B-spline fit with (a) 14, (b) 19, and (c) 24 degrees of freedom. In each case the curve denoted 'E' is for exact data and that denoted 'N' is for non-exact data.

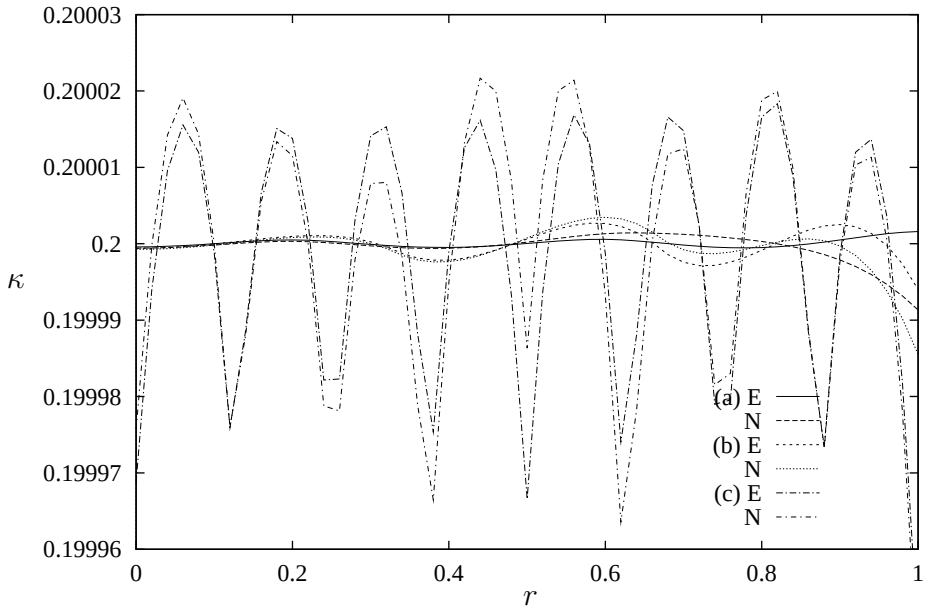

Figure 6.11: Curvature of the spherical cap found using B-spline fits of degree (a) 5, (b) 4, and (c) 3, with optimum degrees of freedom. In each case the curve denoted 'E' is for exact data and that denoted 'N' is for non-exact data.

smoother due to the fact that the second derivative is not only continuous but also differentiable. On repeating the above described process for determining an appropriate number of degrees of freedom, for quartic and quintic B-spline approximations, 8–9 degrees of freedom is found to be around the optimum. Figure 6.11 shows a comparison of the curvature profiles obtained using cubic, quartic and quintic B-splines, with 19, 9, and 8 degrees of freedom respectively. Approximating the data by a quartic or quintic B-spline clearly gives a better result than a cubic spline. The quintic spline appears to be a little better than the quartic spline, but in view of the stretched vertical scale used to show the variations in these curvature profiles, there is really very little difference between them and both are very acceptable.

Thus, with appropriate choice of parameters, we have been able to quite accurately compute the constant curvature of a spherical surface from least-squares B-spline fitting to the surface coordinate data, even when that data is non-exact. We next proceed to use this method to find the curvature of the slumped surface from the finite-element computations of the previous section.

6.5 Numerical Slumping: Surface Curvature

In order to compute the curvature of the upper surface of the numerically-slumped glass disc of Section 6.3 at some point in time, we take the node coordinate data of the slumped glass disc at that time as given by the finite-element program and extract from this the coordinates of nodes located on the upper surface. To this discrete and non-exact data we fit a least-squares B-spline of at least degree 4, to ensure adequate continuity of the second derivative, as described in the previous section. If we keep the proportion of degrees of freedom to data points about the same as in those computations, then we might expect to need as few degrees of freedom as one tenth of the number of data points. However, this can also be expected to depend on the shape of the surface, which is not necessarily spherical,

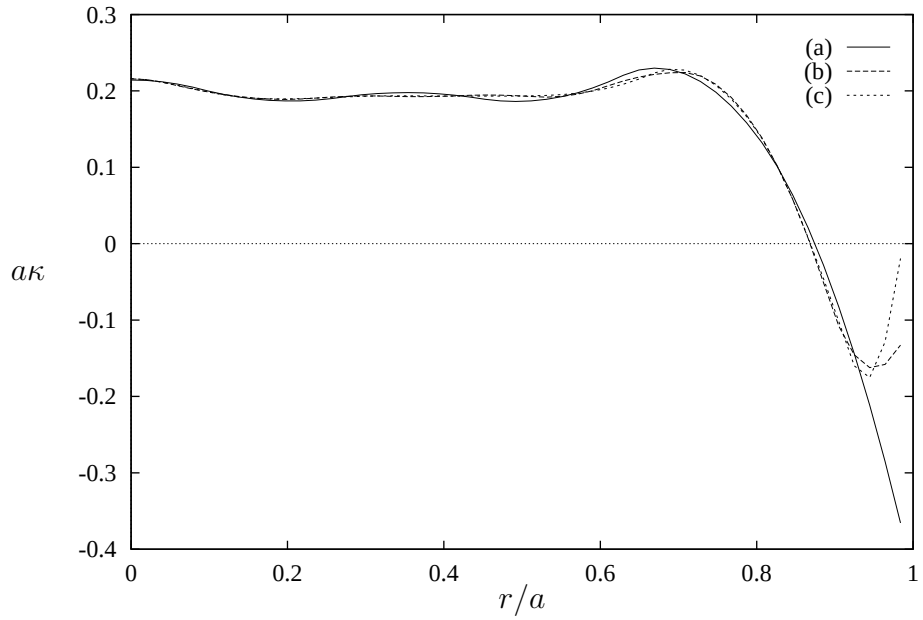

Figure 6.12: Determining surface curvature of the slumped disc of Figure 6.5 using quartic B-spline approximations. (a) $D_f = 15$, (b) $D_f = 25$, and (c) $35 \leq D_f \leq 65$.

so that, as in our earlier curvature computations, D_f needs to be determined by experimenting with this particular coordinate-data set. As a first try we space internal knots uniformly across the diameter of the component.

Firstly let us compute the curvature profile of the slumped disc at time $t = 0.075\mathcal{T}$, after it has attained full contact with the mould as shown in Figure 6.5. This glass disc has 157 nodes on its upper surface, and hence we begin with $D_f = 14$ and increase this number watching for instability in the curvature profile due to error in the coordinate data. The process is illustrated for quartic splines in Figure 6.12. For $35 \leq D_f \leq 65$ the curvature profile barely changes, and further this same profile is also obtained with B-spline approximations of degrees 5 and 6. The plot of curvature shown in Figure 6.13, and at a larger scale in Figure 6.14, may therefore be accepted as correct for this data set. On increasing the number of degrees of freedom above 65, there are increasing fluctuations in the curvature profile due to error in the data. Although the vertical scale of the plot is considerably smaller than those used for plots of curvature of the spherical cap in the previous section,

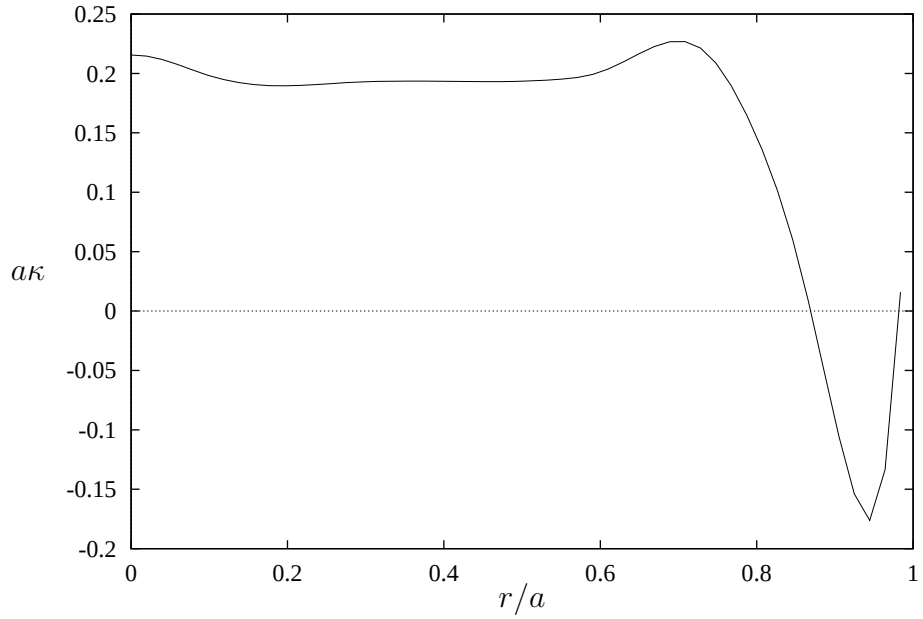

Figure 6.13: Surface curvature of the slumped disc Figure 6.5 computed using a quintic B-spline approximation.

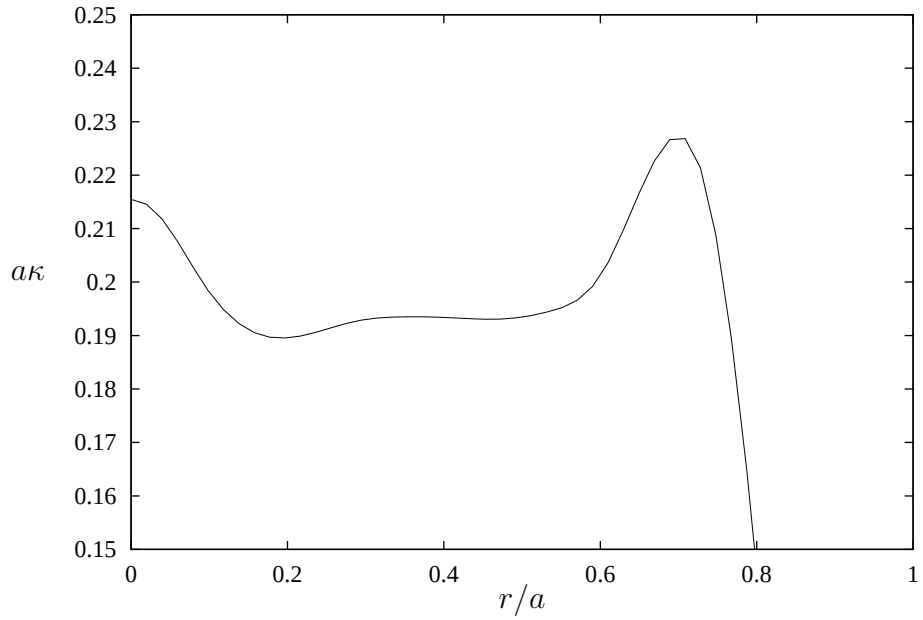

Figure 6.14: As above but at a larger scale.

instability due to error in the data does not become obvious, even at much larger scales, until many degrees of freedom (> 65 compared with ~ 10 for the spherical cap) are used, leading to the conclusion that the numerical data is smoother than the single-precision discretization of the spherical cap. This does not however prove the result to be correct for this slumping case, since the curvature profile is dependent on the space and time discretizations used in the finite-element simulation.

Having computed the curvature profile Figure 6.13 for the slumped glass disc, we can now see that it is considerably different from the constant mould curvature $a\kappa = 0.1876$ that we are trying to replicate. It also shows regions along the radial span of the disc of rapid change in curvature. It is possible that we might compute curvature more efficiently with a non-uniform distribution of the internal knots such that there is a greater knot density in regions of rapid change in curvature, but we can be satisfied with the accuracy of the results with uniformly spaced knots since they remain basically unchanged on increasing the number of degrees of freedom of the B-spline approximations and hence the density of knots. Furthermore the curvature profile of Figure 6.13 is also obtained on using a quartic B-spline approximation with 30 degrees of freedom and Chebyshev knot spacing with the knot density increasing towards the edge of the disc where curvature changes most rapidly. However, in general Chebyshev knot spacing is not appropriate for this data, leading to problems with matrix invertibility. By contrast uniform spacing is very reliable in giving a solvable problem, and the price of any lack in computational efficiency is small, especially when compared with the computational time that can be spent in searching for some better knot distribution.

To give some idea of the dependence of curvature profile on the discretization used in the finite-element computations, and hence determine how accurately it represents the true curvature for this slumping case, two further simulations were run. For one of these a coarser mesh was used over the glass disc of 329 elements instead of 695, giving a top surface defined by 105 data points instead of 157. For

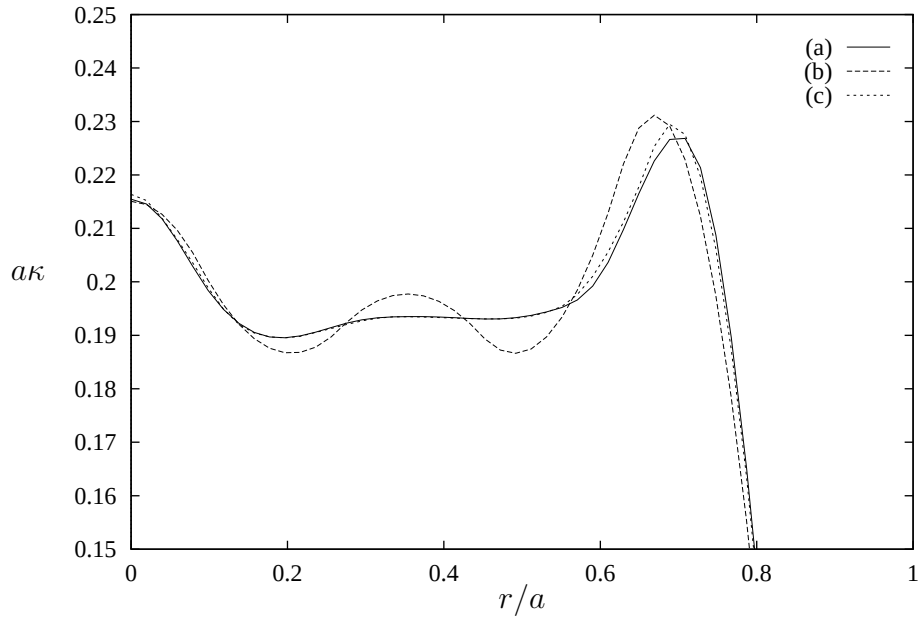

Figure 6.15: Curvature computed from numerical simulations using meshes of (a) 695, (b) 329, and (c) 880 triangular elements.

the other a finer mesh was used of 880 elements with elements concentrated at approximate locations of change in curvature, $r/a = 0.1$ and $r/a = 0.7$ on the top surface. This last simulation gave a top surface defined by 229 data points. Figure 6.15 compares the curvature profile at $t = 0.075\mathcal{T}$ for all of the simulations. The coarse mesh differs a little from the previously computed curve, which is not unexpected, but the finest mesh agrees with it very closely and indicates that the first discretization is sufficient to give a quite accurate result.

Next we see in Figure 6.16 the progression of the curvature profile as slumping proceeds up to full mould contact at $t = 0.075\mathcal{T}$. It is evident from this plot and Figures 6.3 to 6.5 that once the glass contacts the mould the curvature reaches a stable value and subsequently changes very little. Thus the changes to the curvature profile are seen more and more toward the edge of the disc as time proceeds. After full mould contact, the curvature at the disc edge continues to change for a while as shown in Figure 6.17, and it appears that a 'limit state' will be attained after a sufficiently long period of time. This is not strictly correct, but slumping after

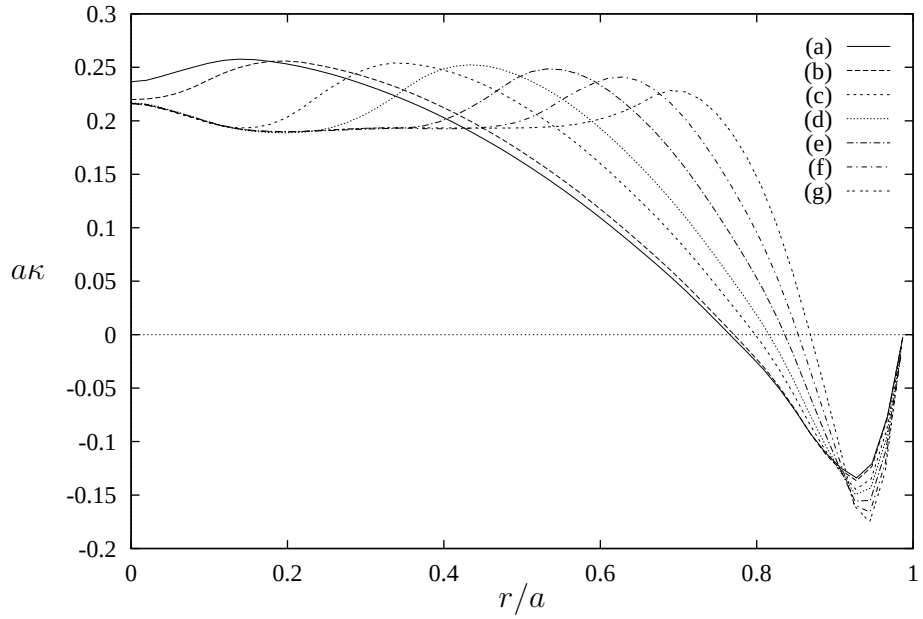

Figure 6.16: Curvature of the slumping disc at (a) $t = 0.014\mathcal{T}$ (initial mould contact), (b) $t = 0.015\mathcal{T}$, (c) $t = 0.020\mathcal{T}$, (d) $t = 0.025\mathcal{T}$, (e) $t = 0.035\mathcal{T}$, (f) $t = 0.050\mathcal{T}$, and (g) $t = 0.075\mathcal{T}$ (full mould contact).

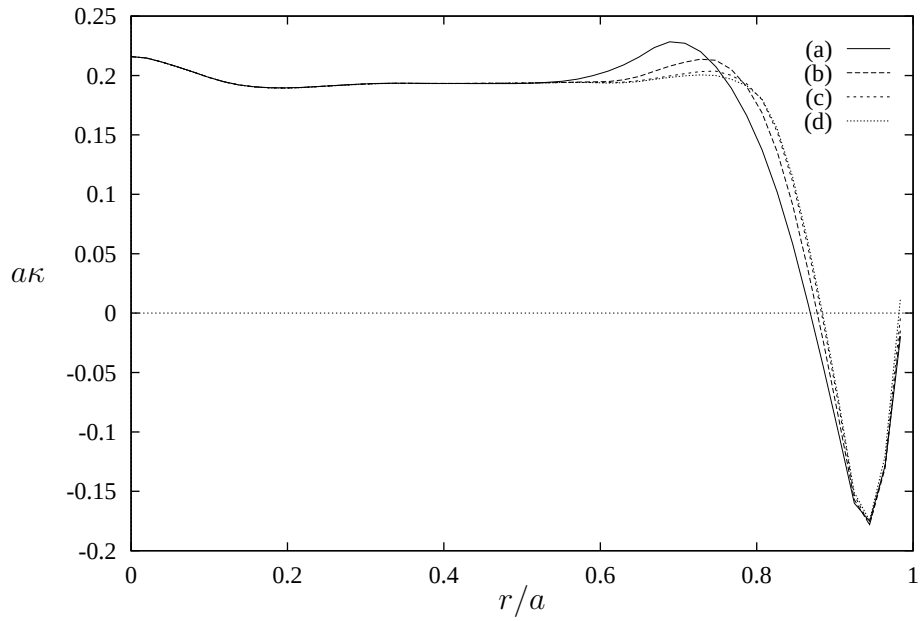

Figure 6.17: Curvature of the slumping disc after full mould contact. (a) $t = 0.075\mathcal{T}$ (full mould contact), (b) $t = 0.100\mathcal{T}$, (c) $t = 0.125\mathcal{T}$, and (d) $t = 0.150\mathcal{T}$.

this settling stage proceeds so very slowly, that it is a good approximation to the truth. Even, the settling stage of the slumping process, from full mould contact to the limit state, takes as much time again as slumping from commencement up to full mould contact, so that it may be desirable to stop the slump before arriving at the limit state.

6.6 Effect of Initial Geometry on Surface Replication

Pearson [68, p. 244] in discussing the extrusion of polymers comments on the importance of geometry above all other factors in determining flow fields and process behaviour. In slumping of optical surfaces, geometry is also very important, and in this section we look at the slumping of a glass disc into a spherical mould and how initial geometry affects final free-surface curvature.

It is obvious that the replication of the mould surface on the upper surface of the slumping glass disc will be greatly affected by the initial thickness h of the glass disc relative to the size of the cavity into which it slumps. Possibly the initial shape of the disc, e.g. flat or spherical, will also be of significance. Initially the radius a of the mould cavity, which is also the radius of the glass disc, will be the only cavity-size parameter, but as the flow progresses and the lower surface of the disc contacts the mould, the depth δ of the cavity can also be expected to have an influence. Now, for a disc of very small aspect ratio $h/(2a) \to 0$, the mould curvature κ_M must be exactly replicated on the upper free-surface of the glass regardless of initial geometry, i.e. $\kappa \to \kappa_M$, where κ is the curvature of the free surface. However, for the thick-disc limit $h/(2a) \to \infty$ the upper disc surface will not see the mould at all and will retain its initial shape, i.e. $\kappa \to \kappa_0$ where κ_0 is the initial upper-surface curvature at time $t = 0$. For $0 < h/(2a) < \infty$ the final upper-surface curvature of the disc will vary, in some manner depending on the disc and mould geometries,

from $\kappa = \kappa_M$ to $\kappa = \kappa_0$. It is desirable to obtain a better understanding of how this surface varies with geometry so that in specific industrial applications of the thermal-replication process, appropriate choices to yield the required product can be more easily identified.

We first consider flat glass discs of varying aspect ratio $h/(2a)$ slumping into a spherical mould with a maximum cavity depth of $\delta = 0.1a$. For each aspect ratio computed, a finite-element numerical simulation is run until full contact between the mould and the lower disc surface is established, and we are close to the limiting curvature profile (see previous section) for the upper free surface. Then as $h/(2a) \to 0$ we have $\kappa \to 1/(5.05a) \approx 0.1980/a$, while as $h/(2a) \to \infty$ we have $\kappa \to 0$. Figure 6.18 shows results for a number of different disc aspect ratios. The vertical axis gives κ/κ_M so that a unit value represents a free-surface curvature equal to the mould curvature. As expected the curvature profile for $h/(2a) = 0.025$ is close to the constant mould curvature ($\kappa/\kappa_M = 1$) everywhere, excepting at the disc edge where 'edge effects' can be anticipated. These edge effects have an increasing radial extent as the disc aspect ratio increases until there is a definite change in the nature of the curvature profile as seen in the curve for $h/(2a) = 0.2$.

In the range of disc aspect ratios $h/(2a) \leq 0.2$ the free-surface curvature is, if we ignore the edge effects, larger than the mould surface curvature, i.e. $\kappa > \kappa_M$, and κ increases with aspect ratio. This suggests that, to a first approximation, the glass forms a layer of uniform thickness over the mould, so that if $R_M = 1/\kappa_M$ is the radius of curvature of the mould then the curvature of the free surface is approximated by $\kappa = 1/(R_M - h)$, which exceeds that of the mould in keeping with the computations. Of course, the computations show curvature to be non-constant, and in addition, at the very centre of the discs, the curvature is significantly larger than this explanation alone would permit, clearly indicating that this approximation is still significantly lacking in accuracy.

Above $h/(2a) = 0.2$ there exists an aspect ratio for which the curvature at the

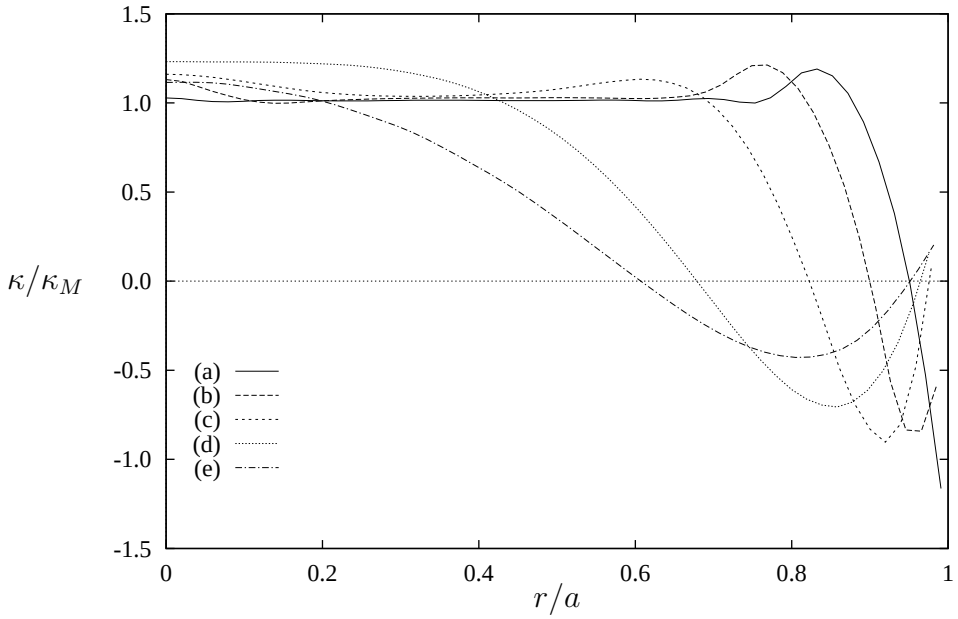

Figure 6.18: Curvature after full mould contact of an initially flat slumping disc of aspect ratio (a) $h/(2a) = 0.025$, (b) $h/(2a) = 0.05$, (c) $h/(2a) = 0.1$, (d) $h/(2a) = 0.2$, and (e) $h/(2a) = 0.3$. The mould cavity is spherical with $\kappa_M = 0.1980/a$.

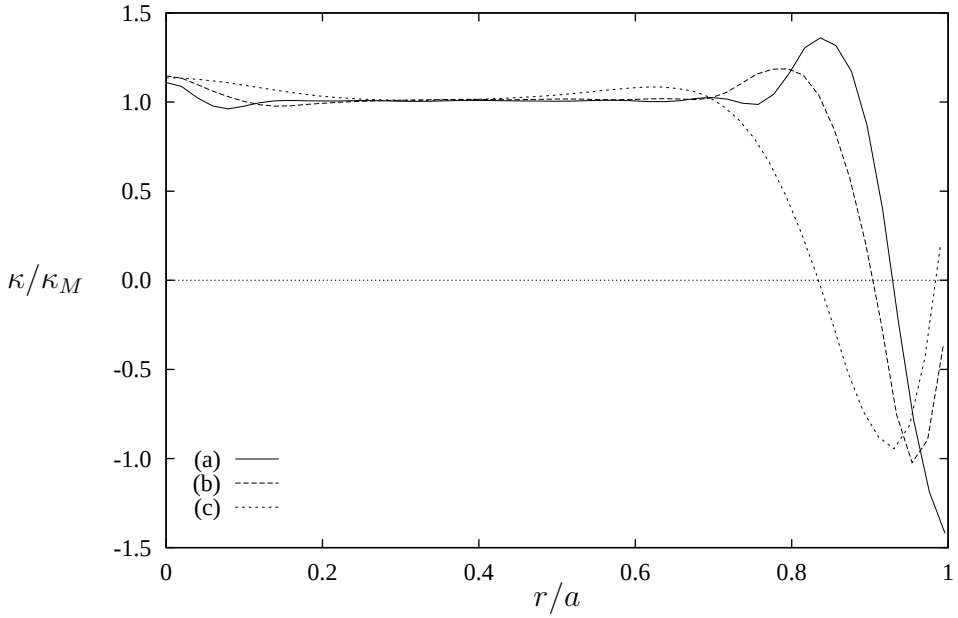

Figure 6.19: Curvature after full mould contact of an initially flat slumping disc of aspect ratio (a) $h/(2a) = 0.025$, (b) $h/(2a) = 0.05$, and (c) $h/(2a) = 0.1$. The mould cavity is spherical with $\kappa_M = 0.0998/a$.

disc centre reaches a maximum, and on further increasing the aspect ratio above this value, κ decreases with increasing $h/(2a)$ and the curvature profile becomes flatter as it approaches the thick-disc limit $\kappa = 0$. In Figure 6.18 this behaviour is evident at $h/(2a) = 0.3$.

The type of behaviour illustrated in Figure 6.18 is typical for spherical mould cavities of other depths δ, although the precise curvature profiles do change a little. On halving the cavity depth to $\delta = 0.05a$, i.e. $\kappa_M = 0.0998/a$ results are obtained as shown in Figure 6.19. A comparison of Figures 6.18 and 6.19 shows less variation in the scaled-curvature profile as aspect ratio increases. For example the curves for $h/(2a) = 0.1$ are quite similar, and this similarity can be expected to continue for larger aspect ratio as the effect of the mould is lost on the free surface and $\kappa \to 0$.

So far we have considered only the slumping of initially flat discs. Spherical discs of some initial curvature κ_0 can also be readily obtained, and it is of interest to consider how using such discs might influence the final free-surface curvature. We consider discs of horizontal radius a having the same initial curvature on both top and bottom surfaces. Thickness h is measured vertically. Note that for regions having curved surfaces, as for the present problem, curvature calculation is sensitive to the accuracy of boundary representation, meshes generated by *Fastflo* must be processed to improve this accuracy. A more detailed description of this fact and the method of mesh improvement is given further on in Chapter 8.

Figure 6.20 shows some results for discs of initial curvature $\kappa_0 = 0.0998/a$ slumping into a mould with $\kappa_M = 0.1980/a$. Then, while the cavity depth is $0.1a$, the distance between the lower surface of the disc and the mould is half of this, i.e. $0.05a$. For ease of comparison, the vertical scale is as before in Figures 6.18 and 6.19. Clearly, while the general behaviour with increasing aspect ratio is similar to that seen with flat discs, the final free-surface curvature is considerably closer to the mould curvature that we are trying to replicate. It is perhaps almost obvious that an initially spherical disc would replicate a spherical mould better than an initially

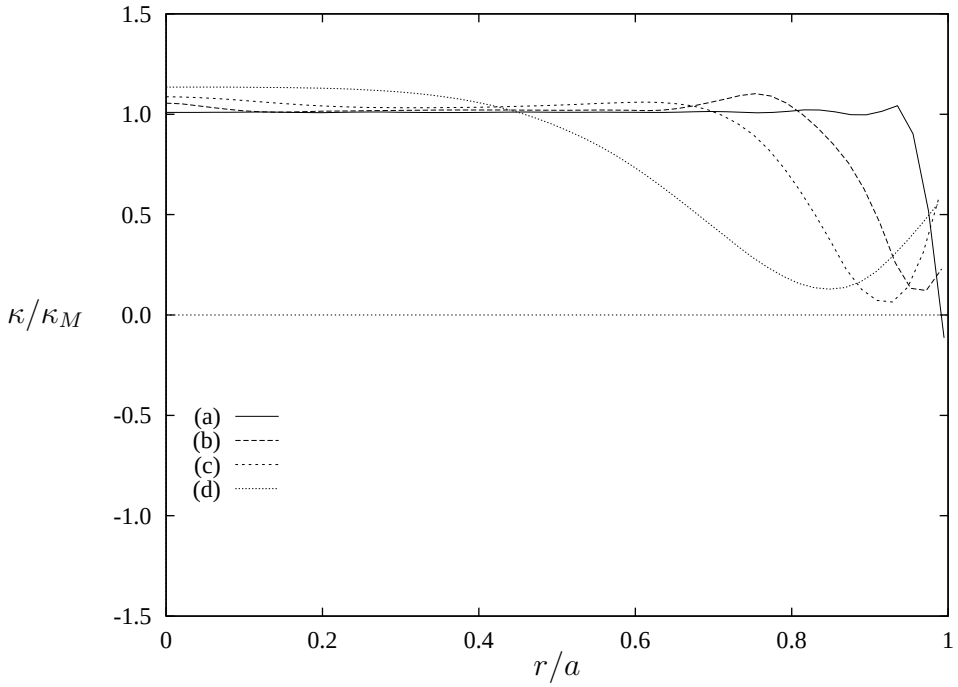

Figure 6.20: Curvature after full mould contact of an initially spherical ($\kappa_0 = 0.0998/a$) slumping disc of aspect ratio (a) $h/(2a) = 0.025$, (b) $h/(2a) = 0.05$, (c) $h/(2a) = 0.1$, and (d) $h/(2a) = 0.2$. The mould cavity is spherical with $\kappa_M = 0.1980/a$.

flat disc, but this test case does illustrate how that, in addition to thickness, the initial shape of the glass disc influences the final product.

For practical applications, where the mould and desired final surface are typically of general non-spherical shape, numerical computations such as these, could assist considerably in determining the best initial shape and thickness of the glass disc, and the shape of the mould cavity.

A feature of all of the slumping problems that we have considered that warrants comment is the increasing time taken to achieve full contact between the lower disc surface and the mould with increasing aspect ratio. We take by way of example the first slump case considered in this section, namely a flat glass disc slumping into a spherical mould of curvature $\kappa_M = 0.1980$. For small aspect ratios slumping occurs relatively quickly, with a disc of $h/(2a) = 0.025$ achieving full mould contact in a time of about $0.01\mathcal{T}$, where $\mathcal{T} = \mu/(\rho g a)$. For a disc of $h/(2a) = 0.3$ the total

slumping time increases more than 20 fold to $0.22\mathcal{T}$. This can be explained by considering the way in which the lower disc surface contacts the mould, and with reference to results for the sagging viscous bridge of Chapter 3.

For small aspect ratio ($h/(2a) \leq 0.025$) the disc slumps rapidly, contacting the mould first at the centre, after which the slump rate slows considerably. For larger aspect ratio ($h/(2a) = 0.2$) there is first contact in a band at the edge of the mould cavity, then at the centre, and lastly along the intermediate radial region. For sufficiently large aspect ratio ($h/(2a) = 0.3$) contact begins at the cavity edge and proceeds radially inwards with contact at the centre coming last.

Consider for a moment Figure 3.7 in Chapter 3 showing the relationship between aspect ratio and initial centroid velocity scaled with respect to $\overline{\mathcal{V}} = \rho g w^2/(2\mu)$ for a sagging viscous bridge of width $2w$ and thickness $2h$. Ignoring constant multipliers, the velocity scale of this figure and that used in our present problem $\mathcal{V} = \rho g a^2/\mu$ are equivalent. Furthermore, as $h/(2a) \to 0$ the present flow becomes equivalent to an axisymmetric form of the sagging viscous bridge with no-slip wall supports, while as $h/(2a) \to \infty$ the present flow must approach a uniform free stream. Thus the relationship between initial velocity at the centre of a slumping disc and aspect ratio must be generally similar to that shown for the sagging viscous bridge in Figure 3.7, with velocity becoming infinitely large in the thin limit and approaching a constant value in the thick limit. Now, for time $t = t^*\mathcal{T}$ with $\mathcal{T} = \mu/(\rho g a)$ and vertical slumping velocity $v = v^*\mathcal{V}$ with $\mathcal{V} = \rho g a^2/\mu$, as in the present problem, the time taken to slump a fixed distance ϵa is given by $t^* = \epsilon/v^*$, assuming that the slumping velocity v^* remains constant over the time interval taken to slump this distance, which for small epsilon is a reasonable approximation.

Let us first consider the thick disc limit $h/(2a) \to \infty$. In this case mould contact proceeds radially inwards from the disc edge, so that the radius of the unfilled mould cavity, which we denote $r_c(t)$ with $r_c(0) = a$, decreases over time. At any time t, this problem is approximately equivalent to the initial problem at $t = 0$, with

the slumping velocity $v^* \to c$, for some constant c, and a velocity scale $\rho g r_c^2/\mu = \mathcal{V}(r_c/a)^2$, $\mathcal{V}$ being the initial velocity scale when $r_c = a$. Thus, for a constant velocity scale $\mathcal{V}$, as in present computations, we have a slumping rate that decreases like $(r_c/a)^2$ as slumping proceeds. Next let us assume that for an unfilled mould cavity of radius r_c the disc slumps a distance ϵr_c before further mould contact results in a decrease in r_c. The time taken to slump this distance is $\Delta t^* = \epsilon/c$ multiplied by a time scale $\mu/(\rho g r_c) = \mathcal{T}(a/r_c)$, $\mathcal{T}$ being the initial time scale. Thus as slumping proceeds and $r_c \to 0$ we have the slumping time $t \to \infty$, and in fact, although we get very close, we never actually get to full mould contact. If instead we consider a constant slumping increment of ϵa (instead of ϵr_c), then the slumping time increment is $\Delta t = (\epsilon/c)\mathcal{T}(a/r_c)^2$, and, while with this formulation of the problem the slump is completed, again the total slump time is very large.

Now we consider the thin disc limit $h/(2a) \to 0$. In this case there is a period of time during which the disc slumps freely, until contact with the mould at the centre. For this first stage of the slumping process, the velocity is very large, with $v^* \to \infty$ initially, and the velocity scale remains constant at $\mathcal{V}$. Even allowing for a decrease in the velocity during this period (cf. Figure 3.10) it is readily appreciated that this first stage of the slump up to initial contact at the centre of the mould is of short duration. After contacting the mould at the centre, the slumping of the disc between its centre and edge is approximately the same as the initial problem, but the size of the mould cavity is much reduced with an effective radius of $a/4$ so that the time scale increases to $4\mathcal{T}$ and the velocity scale reduces to $\mathcal{V}/16$, where again $\mathcal{T}$ and $\mathcal{V}$ are the initial time and velocity scales respectively. However, this next stage of the slumping process appears to proceed in a manner similar to that described above for a thick disc, i.e. with mould contact proceeding radially inwards. Thus, a slowing of the slumping rate is to be expected, although the large slumping velocity applicable to a very thin disc still results is a faster process than for a thick disc.

For discs with some thickness, between the thin and thick limits, the first stage

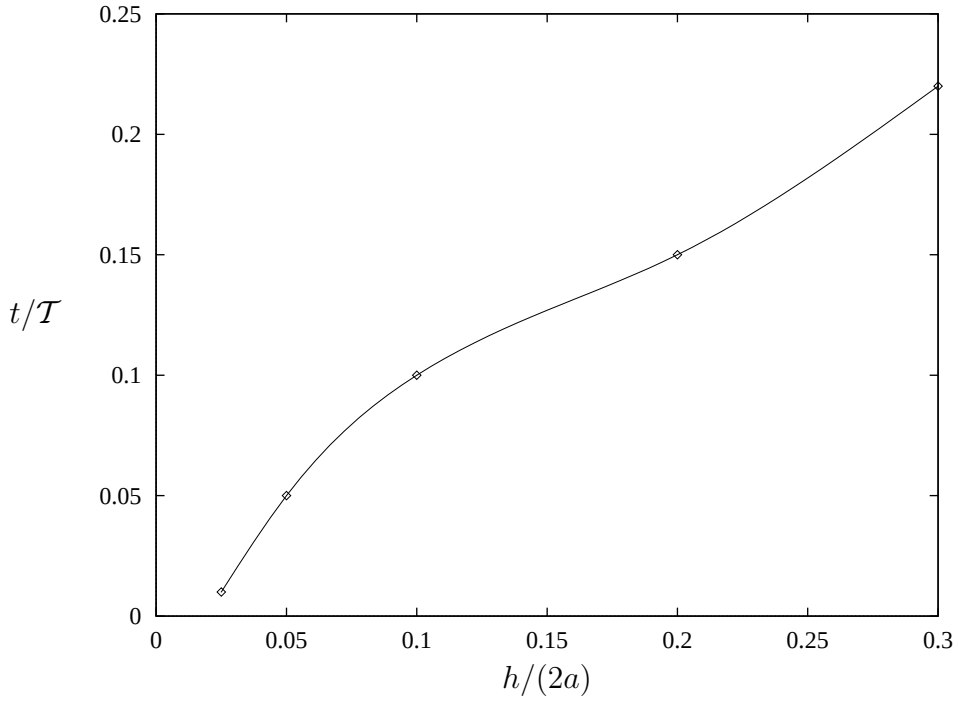

Figure 6.21: Slumping time for full mould contact.

of the slumping process (up to contact with the mould at the centre) occurs over a longer time period than for very thin discs due to the lower slumping velocities (cf. Figure 3.7). The second stage of the process will also take a longer time than for very thin discs, as the aspect ratio h/r_c, r_c being the effective cavity radius, increases and towards the end of the slump approaches the thick-disc limit.

Thus, although this argument is very approximate, the increasing total slump time with increasing disc thickness is understandable. Figure 6.21 shows the relationship between disc aspect ratio and the time taken for full mould contact. There is no obvious tendency to approach a limiting slump time as aspect ratio increases, although this is possible for aspect ratios exceeding the range of our computations. However, even if such a limit exists, in terms of real times it will exceed practical slump times, and hence has no real relevance.

One conclusion that we can draw from our considerations so far is that, if thermal replication is to be practical as a manufacturing process, then thin discs must be

used, both to keep manufacturing times to a minimum, and so that the mould shape can be effectively transferred to the upper free surface of the disc. Because of this, thick-limit calculations may appear to be irrelevant. However, such calculations may be of benefit in determining just how thin a disc must be for a reasonable slumping time and sufficient accuracy of replication, and, furthermore, thick-limit calculations are relevant to a consideration of the transfer of mould roughness to the upper free surface.

6.7 Print-Through of Surface Roughness

In describing the thermal-replication process at the beginning of this chapter, mention was made of the reason for using such a process being to obtain a replica of the macro-surface of a mould while not transferring the small-scale imperfections in the mould. The ceramic mould is machined in some type of milling machine, which gives it certain roughness characteristics. Obviously the surface of the product that comes into direct contact with the mould will be directly affected by the mould roughness, and some work on small scale surface changes resulting in such circumstances has been done by others [85]. This type of roughness transfer is not, however, of importance in considering replication of the mould surface on the upper surface that makes no contact with the mould. Rather we are concerned, in this section, with "print-through" of the mould surface roughness to the upper free surface. It is of interest to know how coarse the mould can be and still yield a very smooth surface, such as is necessary for optical applications, since the finer the required milling, the greater the cost of the mould, and it is obviously important on an industrial scale where these moulds are mass produced, to keep costs to a minimum. From this study we also obtain some further insight into the size of the surface features in the mould that can be transferred to the upper surface of the glass in the thermal-replication process. This is also important, since it strongly influences the design of

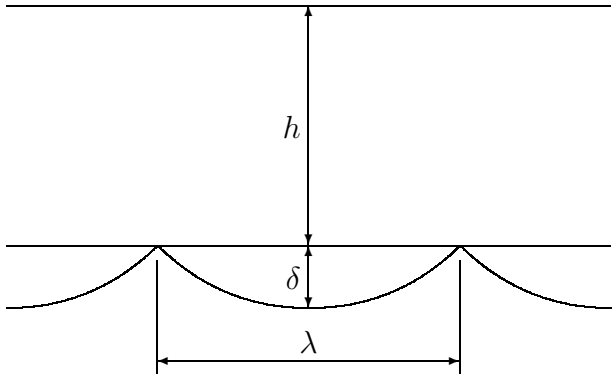

Figure 6.22: Enlargement of fluid-mould interface showing mould roughness.

products that can be manufactured with this process.

The milling machines which have come to the author's attention during the present work, cut the required mould surface shape by making a series of closely spaced parallel runs which are long compared to the width of the cut. Thus, at the scale at which the roughness is seen, the problem is two dimensional. At this scale also and for the gravity slumping process with which we are primarily involved, it is reasonable to take the viscous fluid layer to be of uniform thickness lying in the horizontal plane. This suggests the initial geometry shown in Figure 6.22. Here roughness is represented as a periodic "wave", with wavelength λ and depth from crest to trough δ. The wave shape from crest to crest has been shown as an arc because something like this is expected from a milling machine, but this is not critical to the final formulation of the problem. A layer of viscous fluid (e.g. molten glass) of height h is initially supported on the mould crests, and typically $h \gg \delta$. This fluid layer sags under gravity g. Interest is in the shape of the upper free surface that results from the fluid layer sinking down into the troughs in the mould surface.

The expected periodicity of the mould roughness, enables the restriction of the computational flow domain to a half wavelength $w = \lambda/2$, with symmetry conditions on the side boundaries. Let us also simplify the mould representation to a series of short, horizontal supports (at the "wave" crests) with infinite depth holes in between. This is reasonable while the only contact of the lower free surface with

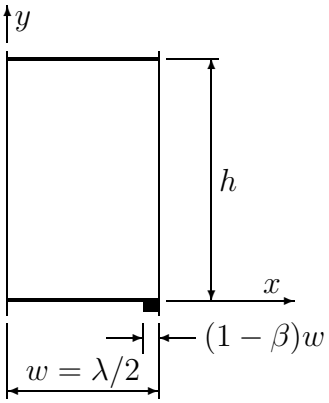

Figure 6.23: A simplified mould roughness element.

the mould is at the crests, which is certainly true for $t = 0^+$. Furthermore, the time scale of this problem $\mathcal{T} = \mu/\rho g w$ (using the half wavelength w as the length scale) is very large for the large viscosity μ and small length scale w involved, so that even typical real slumping times of one to two hours, are very small relative to this scale, and little further mould contact can be expected. In any case, this is a worse-case scenario. Thus, the final geometry is that shown in Figure 6.23.

We first solve for the flow field in this computational domain, and then determine how this affects the upper free surface of the glass. Since in the industrial application the slumping time t is very small, the deformed shape of the initially horizontal upper free surface given by $y = h$, will be, to a good approximation, given by

$$y = h + tv(x, h). \tag{6.10}$$

In particular, the difference in sag of the top surface from the centre to the edge of the roughness element, given by

$$\Delta y = t[v(w, h) - v(0, h)] = t\Delta v, \tag{6.11}$$

gives a measure of the print-through of mould roughness to the upper surface.

This problem is quite similar to the slumping liquid bridge analysed in Chapter 3; only the boundary conditions must be modified. It is most readily analysed using a semi-analytic series method to solve for the stream function ψ, in a similar manner to

that outlined in Section 3.2. The basic slow-flow equations and boundary conditions are the same as in Section 3.2, but are repeated here for convenience. Thus, we must solve the Stokes equations

$$p_x = \mu\nabla^2\psi_y \qquad \text{and} \tag{6.12}$$

$$p_y = -\rho g - \mu\nabla^2\psi_x, \tag{6.13}$$

which for an incompressible fluid combine to give the biharmonic equation

$$\nabla^4\psi = 0, \tag{6.14}$$

with boundary conditions as given below.

As always throughout this thesis, the mould surface is assumed to be a no-slip wall so that on $y = 0$, $\beta w \leq x \leq w$ we must satisfy

$$\psi_x = \psi_y = 0. \tag{6.15}$$

The zero-stress free-surface conditions at $t = 0^+$,

$$-p - 2\mu\psi_{xy} = 0 \qquad \text{for no normal stress and} \tag{6.16}$$

$$\psi_{xx} - \psi_{yy} = 0 \qquad \text{for no tangential stress,} \tag{6.17}$$

apply on $y = 0$, $0 \leq x < \beta w$ and $y = h$, $0 \leq x \leq w$. Differentiating (6.16) with respect to x and substituting into (6.12) yields

$$\psi_{yyy} + 3\psi_{xxy} = 0 \tag{6.18}$$

which replaces (6.16) and temporarily eliminates pressure from the equations [93]. However, (6.16) must be used at some stage of the solution process to re-introduce pressure back into the problem. Symmetry requires that there be no horizontal velocity component and no horizontal vertical-velocity gradient on the left and right side boundaries so that

$$\psi = 0 \qquad \text{on}\, x = 0, \tag{6.19}$$

$$\psi = Vw \qquad \text{for constant } V \text{ on}\, x = w, \text{ and} \tag{6.20}$$

$$\psi_{xx} = 0 \qquad \text{on both}\, x = 0 \text{ and } x = w. \tag{6.21}$$

Choosing the characteristic length $\mathcal{L} = w/\pi$, we next define non-dimensional variables

$$x^* = \pi x/w \tag{6.22}$$

$$y^* = \pi y/w \tag{6.23}$$

so that the computational flow domain at $t = 0^+$ is the rectangle $0 \le x^* \le \pi$, $0 \le y^* \le h^*$ where $h^* = \pi h/w$, and the mould surface is the line segment $\beta\pi \le x^* \le \pi$, $y^* = 0$. Then

$$\psi = \frac{Vw}{\pi}\left(x^* + \sum_{n=1}^{N} \sin(nx^*)[e^{-ny^*}(a_n + b_n y^*) + e^{ny^*}(c_n + d_n y^*)]\right) \tag{6.24}$$

satisfies the biharmonic equation (6.14) and the side-boundary conditions (6.19) to (6.21).

On substituting (6.24) into equations (6.15), (6.16), and (6.18), we obtain

$$\sum_{n=1}^{N} n\cos(nx^*)(a_n + c_n) = -1 \qquad \text{and} \tag{6.25}$$

$$\sum_{n=1}^{N} \sin(nx^*)(-na_n + b_n + nc_n + d_n) = 0 \tag{6.26}$$

on $\beta\pi \le x^* \le \pi$, $y^* = 0$,

$$\sum_{n=1}^{N} n\sin(nx^*)(na_n - b_n + nc_n + d_n) = 0 \qquad \text{and} \tag{6.27}$$

$$\sum_{n=1}^{N} n^3\sin(nx^*)(a_n - c_n) = 0 \tag{6.28}$$

on $0 \le x^* < \beta\pi$, $y^* = 0$, and

$$\sum_{n=1}^{N} n\sin(nx^*)[e^{-nh^*}(na_n + nb_n h^* - b_n) + e^{nh^*}(nc_n + nd_n h^* + d_n)] = 0 \qquad \text{and} \tag{6.29}$$

$$\sum_{n=1}^{N} n^3\sin(nx^*)[e^{-nh^*}(a_n + b_n h^*) - e^{nh}(c_n + d_n h^*)] = 0 \tag{6.30}$$

on $0 \le x^* < \pi$, $y^* = h^*$.

Choosing M collocation points on $y^* = 0$ and another M collocation points on $y^* = h^*$, with $M > N$, we have a linear least-squares problem of $4M$ equations in $4N$ unknowns.

To completely define ψ, we lastly need to determine the constant V. To do this we substitute (6.24) into (6.12), then integrate with respect to x and differentiate with respect to y. Comparing this result with equation (6.13) yields an expression for pressure,

$$p = p_0 - \frac{\rho g w}{\pi} y^* - \frac{\mu \pi V}{w} \sum_{n=1}^{N} 2n \cos(nx^*)(e^{-ny^*} b_n + e^{ny^*} d_n) \tag{6.31}$$

where p_0 is a constant. This is then substituted into boundary condition (6.16) which, when satisfied on $0 \leq x^* < \beta\pi$, $y^* = 0$, gives an expression for p_0 that is in turn used in satisfying (6.16) on $y^* = h^*$ to give an expression for V

$$\begin{aligned} V^{-1} = \frac{2\mu\pi^2}{\rho g w^2 h^*} \sum_{n=1}^{N} n^2 (a_n - c_n \\ -e^{-nh^*}(a_n + b_n h^*) + e^{nh^*}(c_n + d_n h^*)). \end{aligned} \tag{6.32}$$

The equations that have been derived may be solved for a fluid layer of any thickness h. In the thick-layer limit $h \to \infty$ we must have $\psi \to Vx$, requiring $c_n = d_n = 0$ for all n. Then equations (6.29) and (6.30) are automatically satisfied for large h, and they may be dropped from the system of equations leaving a least-squares problem of $2M$ equations in $2N$ unknowns.

For the present calculations let us fix some of the problem parameters. Let the mould support to be one tenth of the roughness 'wavelength' λ, i.e. $\beta = 0.1$ cf. Figure 6.23. In general let us also set $M = 2N$, which has been found by experimentation to be about the right proportion of terms in the series expansions (N) to collocation points (M). A larger proportion of terms leads to numerical difficulties. Sometimes, when solving the full set of equations for fairly large aspect ratios $h/(2w)$ it is necessary to use less terms to avoid computational problems with numerical overflow caused by the terms involving $\exp(nh^*)$. Aspect ratio $h/(2w)$

Table 6.1: Velocities on the top free surface of a roughness element of aspect ratio $h/(2w) = 0.5$.

M	$-v(0,h)/\mathcal{V}$	$-v(w,h)/\mathcal{V}$	$\Delta v/\mathcal{V}$
40	0.71283	0.50706	0.20578
60	0.73915	0.50003	0.23912
80	0.73776	0.49876	0.23900
100	0.73900	0.49875	0.24025
120	0.73999	0.49926	0.24073
140	0.74009	0.49934	0.24075
160	0.74071	0.49970	0.24101
180	0.74070	0.49967	0.24103
200	0.74054	0.49955	0.24099
220	0.74055	0.49955	0.24099
240	0.74054	0.49955	0.24099

and the number M of collocation points on each of $y^* = 0$ and $y^* = h^*$ are left as adjustable parameters. With respect to spacing of the collocation points, both uniform spacing and Chebyshev spacing about $x = (1-\beta)w$ have been tried with the same point spacing on both $y^* = 0$ and $y^* = h^*$. Not surprisingly Chebyshev spacing, which clusters points about the singularity where the boundary changes from no-slip to free-surface, was found to be best. Collocation points were located near to, but not at, this point, and $0.1M$ points were placed on the wall segment according to the size of β.

Having solved for ψ it is easy to obtain expressions for initial horizontal and vertical velocity components $(u, v) = (\psi_y, -\psi_x)$ at some vertical position y, and in particular find these at $y = h$. Table 6.1 shows the top surface velocities at the centre and edge of a roughness element of aspect ratio $h/(2w) = 0.5$, computed using the full system of equations. Velocities are scaled with the usual $\mathcal{V} = \rho g w^2/\mu$. From this we see that $M = 80$ gives two figures of accuracy and $M = 120$ gives almost three figures of accuracy, which is near the best the method can do. Increasing M to 200 with $N = 100$ gives slightly better accuracy, but at considerable computational cost.

Table 6.2: Centre to edge velocity difference from the full and thick-limit series equations for roughness elements of large aspect ratio.

$h/(2w)$	Full System		Thick Limit	
	M, N	$\Delta v/\mathcal{V}$	M, N	$\Delta v/\mathcal{V}$
0.5	120, 60	0.2407	1500, 750	0.1651
1.0	100, 50	0.0343	1500, 750	0.0174
1.5	100, 30	0.0032	1500, 750	0.0016
2.0	100, 20	0.0003	1500, 750	0.0001

Numerical overflow occurs with $N > 100$ so that the results shown for $M > 200$ were computed with $N = 100$ rather than $N = M/2$. Consequently, the consistency of the results for $N \geq 100$ cannot be taken as definite convergence, and the results can only be quoted to three significant figures.

Table 6.2 gives the difference in the velocities at the centre and edge of the roughness element, computed using the full and the thick-limit systems of equations for a range of large aspect ratios. As the aspect ratio increases so do numerical problems when solving the full system of equations, and the number N of terms computed must be significantly reduced. Although we would like to see a better correspondence between the two methods, it is necessary, and probably adequate, to use the thick-limit equations for $h/(2w) > 1$.

Now we try to ascertain the minimum aspect ratio at which the top free surface can be considered unaffected by the mould. As stated at the beginning of this section, the difference in height Δy from the centre to the edge of the roughness element gives a measure of the mould effect. This is given by equation (6.11), and is simply the product of the velocity difference $\Delta v = v(w, h) - v(0, h)$ between the two points and a time period t, so that we may also use Δv to measure mould effects. The top-surface velocity profile is much like a sine curve, and, when scaled so that

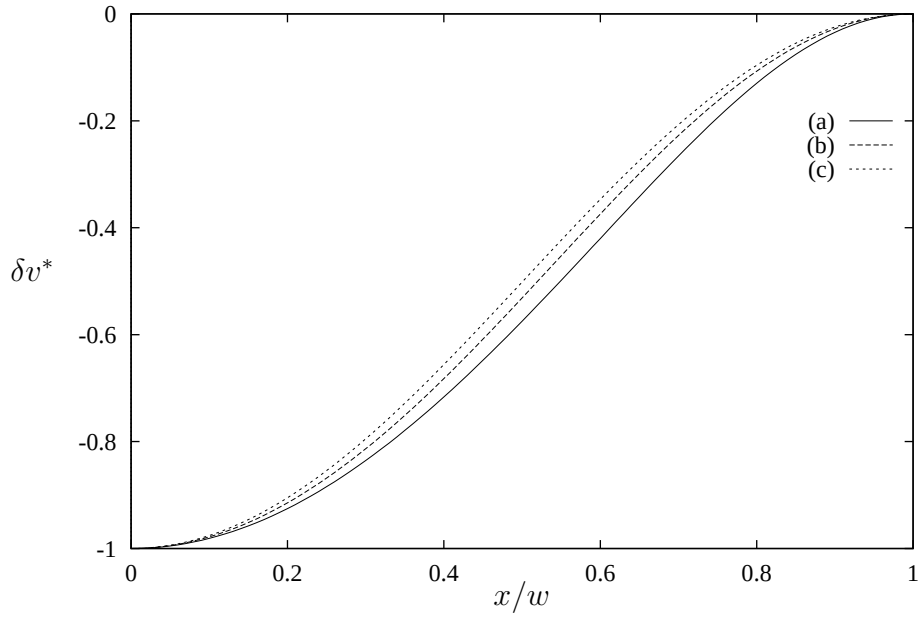

Figure 6.24: Roughness-element top surface velocities $\delta v^*(x/w)$ for aspect ratios of (a) $h/(2w) = 0.25$, (b) $h/(2w) = 0.5$, and (c) $h/(2w) \geq 1$.

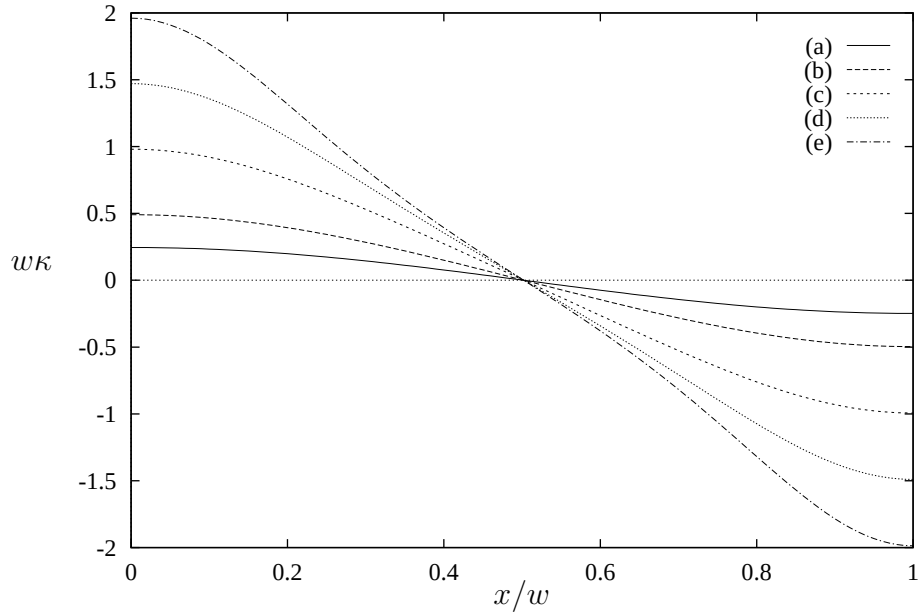

Figure 6.25: Roughness-element top surface curvature profiles computed for $h/(2w) = 1$ when (a) $\Delta y/w = 0.05$, (b) $\Delta y/w = 0.1$, (c) $\Delta y/w = 0.2$, (d) $\Delta y/w = 0.3$, and (e) $\Delta y/w = 0.4$.

$\Delta v^* = 1$, is quite similar for all aspect ratios. Let us define

$$\delta v^*(x^*) = \frac{v(x,h) - v(w,h)}{\Delta v} \tag{6.33}$$

so that $-\delta v^*(0) = \Delta v^* = 1$. Figure 6.24 shows curves of δv^* for several aspect ratios. At this graphical scale the curves for $h/(2w) \geq 1$ appear identical. Now the dimensionless time period

$$t^* = t\frac{\Delta v}{w} \tag{6.34}$$

is equivalent to the measure $\Delta y/w$ of roughness effect. Multiplying δv^* by t^* gives a curve of the same shape as the top free surface of the fluid, from which the curvature profile $w\kappa$ associated with the roughness measure $\Delta y/w$ can be computed. Curvature is computed using the B-spline method described earlier in, and used extensively throughout, this chapter. Because aspect ratio does not greatly affect the curve δv^*, especially when large, it is reasonable, at least for the present, to do all computations for one aspect ratio, and take them as representative for all aspect ratios. Figure 6.25 shows curvature profiles for a number of different values of $\Delta y/w$, computed using $h/(2w) = 1$.

In the present context of determining the extent of print-through of mould surface roughness to the upper free surface in thermal replication, it is the maximum variation of the top-surface curvature from a value of zero that is of most interest. This can be obtained from curvature profiles such as in Figure 6.25, but there is a simpler way. From Figure 6.25 we see that the maximum variation occurs at $x/w = 1$. Now, at this position symmetry requires that the curve of δv^* versus x/w have zero slope, as we see in Figure 6.24, so that the curvature at $x/w = 1$ is given by (refer equation (6.2))

$$w\kappa = t^* \times (\delta v^*)'', \tag{6.35}$$

where primes indicate differentiation with respect to $x^* = x/w$. Thus, we can compute the curvature profile for $t^* = \Delta y/w = 1$, find the maximum variation from

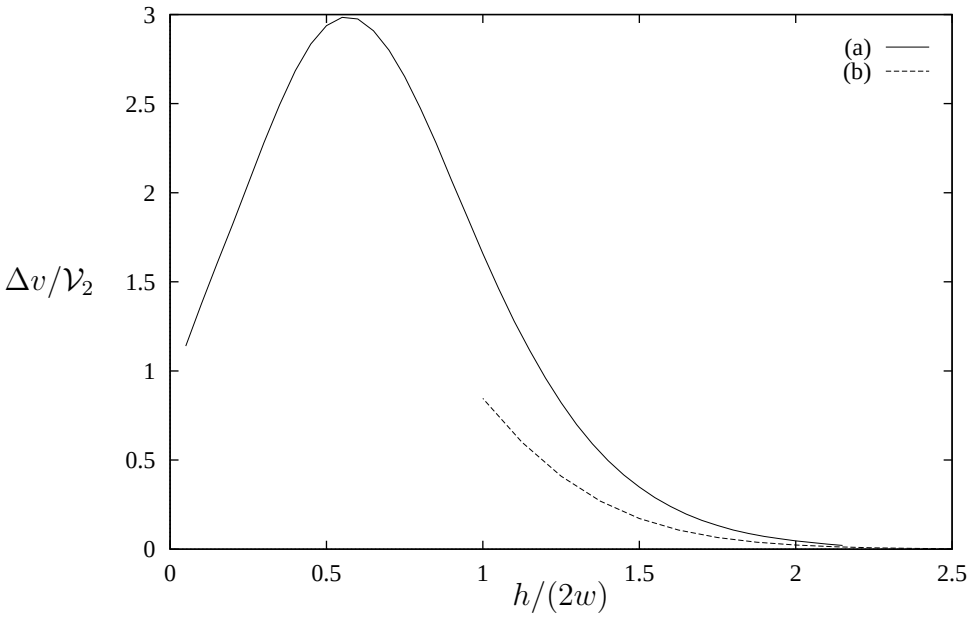

Figure 6.26: Relationship between aspect ratio and roughness print-through as given by the difference in the centre and edge velocities on the top free surface $y = h$. Solution by (a) full system of equations and (b) thick limit.

zero, which is $|w\kappa| = 4.966739$ at $x/w = 1$ and then just use

$$|w\kappa|_{\max} \approx 5.0t^* \tag{6.36}$$

to find the maximum variation for other values of t^*.

Lastly the plot in Figure 6.26 shows the relationship between aspect ratio $h/(2w)$ and Δv. This plot has been scaled using $\mathcal{V}_2 = \rho g a^4/(32\mu\bar{h}^2)$, where $a = (1-\beta)w$ is the roughness cavity half-width and $\bar{h} = h/2$ is the half-thickness. This scale is the thin-bridge centroid velocity limit of the sagging viscous bridge of Chapter 3 (see equation (3.7)). Since the thin limit of the present problem is identical to that sagging viscous-bridge problem, we must have $v(0,h)/\mathcal{V} \to 1$ and $v(w,h)/\mathcal{V} \to 0$, and therefore $\Delta v/\mathcal{V} \to 1$ as $h/(2w) \to 0$. This is seen in Figure 6.26. Note that for optical, and perhaps other, applications that have quite fine tolerances on the specified surface curvatures, aspect ratios exceeding $h/(2w) = 2.5$, which is the maximum shown in Figure 6.26, must be considered.

Now, if for a particular application of the thermal-replication process we know the variation in curvature that can be tolerated with the top surface still considered unaffected by the mould, we can determine from (6.36) the corresponding measure $\Delta y/w = t^*$. From Figure 6.26 we can find the value of Δv for different aspect ratios $h/(2w)$, and using (6.34) we can then find possible combinations of glass thickness h, roughness element half-width w and slumping time t such that the tolerable curvature variation is not exceeded. If, as is likely, t and w are known and fixed by the process and milling machine, then the value of Δv is specified by (6.34) and we must find from Figure 6.26 the aspect ratio that corresponds to this value, to give the minimum glass thickness such that the top surface is unaffected by the mould roughness.

To illustrate this procedure, let us take a roughness element with $w = 0.5\,\text{mm}$, i.e. a very rough mould, and assume a slumping time of one hour. Suppose that the permissible tolerance on curvature is $\pm 0.2 \times 10^{-5}\,\text{mm}^{-1}$. Then the allowable variation in dimensionless curvature $w\kappa$ is $\pm 10^{-6}$. Equation (6.36) then gives that $t^* = \Delta y/w \leq 2 \times 10^{-7}$, and from (6.34) we have $\Delta v \leq 10^{-7}\,\text{mm/hr}$. Now, while Figure 6.26 is not of a suitable range or scale to show this, when $h = 2.5\,\text{mm} \equiv h/(2w) = 2.5$ then $\Delta v/\mathcal{V}_2 = 2.383 \times 10^{-3}$ and $\Delta v = 1.724 \times 10^{-7}\,\text{mm/hr}$, and when $h = 3.0\,\text{mm}$ then $\Delta v/\mathcal{V}_2 = 2.114 \times 10^{-4}$ and $\Delta v = 1.062 \times 10^{-8}\,\text{mm/hr}$, where a density for glass of $\rho = 2500\,\text{kg/m}^3$ and a viscosity of $10^6\,\text{Pa}\cdot\text{s}$ have been used. Thus, a glass component of just 2.5 mm thickness is almost adequate, and a component of 3 mm thickness is more than adequate. Note that glass of about 6 mm in thickness is quite typical for slumping ophthalmic lens moulds, and the ceramic moulds used generally have less roughness than here assumed, so that these calculations indicate that cost savings may be possible with both more coarsely milled moulds and/or thinner glass components.

In this section we have concentrated on the unwanted transfer of mould roughness to the upper free surface of the fluid layer. The minimum aspect ratio for eliminating

this transfer can also be considered, for a larger scale problem, as the maximum aspect ratio at which the mould will have some effect on the upper surface. Thus, as mentioned earlier, these calculations can be used to give some idea of the size of features that can be replicated from a mould for different glass thicknesses and slumping times.

6.8 In Summary

In this chapter we have considered thermal replication by slumping into a mould, developing a method for handling mould contact, and, more importantly, a method for determining surface curvature when slumping is complete. This latter method appears to give quite good results, although we have not attempted any formal quantifying of error which calls for further work.

With these tools we are able to simulate and compare different slumping cases, and we considered the slumping of initially flat and spherical glass discs of varying thickness. In addition we considered the transfer of mould roughness to the critical upper surface of the glass, using a series method to compute velocities and a B-spline method to compute surface curvature profiles.

In the next chapter we shall relax the isothermal assumption, and consider spatially varying viscosity and how this might affect the final product.

Chapter 7

Non-Isothermal Effects in Replication

7.1 The Issue

To this point in this thesis, we have, with some justification, assumed viscosity to be spatially constant throughout the fluid, i.e. that conditions are *isothermal* at any point in time. However, in modelling molten glass flows, and thermal replication of optical surfaces requiring accurate surface curvatures in particular, further investigation of this assumption is warranted.

As was seen in our considerations of glass properties in Chapter 5, the viscosities of molten glasses are highly sensitive to temperature. Soda-lime-silica glasses, which are by far the most commonly used, including for the types of optical components that we are considering, typically have viscosities that increase from the softening point of $10^{7.6}$ poise at around 700°C to the annealing point of $10^{13.4}$ poise at around 500°C (see Table 5.2), i.e. by nearly six orders of magnitude for a temperature drop of about 200°C. Thermal replication involves heating the glass and mould from room temperature over this temperature and viscosity range to the softening point, after which the temperature is held constant for a period of time. Cooling

then takes the glass through another period of rapid viscosity change. The viscosity changes with temperature most rapidly at low temperatures, the rate decreasing with increasing temperature, but even at the upper end of the slumping-temperature range the viscosity changes by about an order of magnitude in 50°C, as is indicated by the VFT equation (5.1) with constants ($A = 2.0886$, $B = 4915.8820$, $T_0 = 192.6136$) determined using the annealing, softening and working point data given in Table 5.2 for a typical soda-lime-silica glass. With the viscosity being so sensitive to temperature there is clearly potential for even small spatial temperature variations to significantly affect the glass flow, all the more so when the quality of the final product is measured in terms of second derivatives, i.e. curvature.

As already discussed in Chapter 6, the importance of temperature has been recognized in the modelling of other glass forming processes, specifically blow moulding and parison forming [14, 21, 83, 100]. Cormeau et al. state [21, p. 221]

> It is important to realize that a full thermal analysis must be conducted in parallel with the mechanical simulation; indeed, the flow of molten glass is strongly dependent on the glass viscosity which varies rapidly with temperature in the vitreous transition temperature interval in which one operates ...

However, as also noted in Chapter 6, in these particular processes a hot viscous film is forced against a comparatively cold mould, so that significant spatial temperature variations are most likely. This is not the case in thermal replication where the mould is heated with the glass, so that the hot glass contacts a mould of a similar, if not the same, temperature, and the isothermal assumption is more justifiable. On the other hand, small surface variations resulting from small spatial temperature variations, are considerably magnified when calculating second derivatives, and may therefore significantly affect the critical surface curvature of the optical components being formed.

There are a number of possible causes of non-isothermal conditions in the thermal-replication process. First there are factors relating to specific oven designs such as non-uniform oven temperatures. Then, the finite time taken for heat transfer into glass, as well as geometrical arrangement which may allow better heat transfer in some areas of the glass compared to others, could result in non-uniform glass temperatures. Lastly, although the mould is heated with the glass, the probably differing thermal properties of the glass and ceramic (or other material) from which the mould is made may also lead to small spatial temperature variations as the lower free surface of the glass progressively contacts the mould.

It is well recognized that full thermo-mechanical modelling of an industrial process involving complex heat transfer mechanisms of conduction, convection and radiation, and equipment specific factors, is a very difficult task. Where ceramics, including glass, are involved, there is additional complication resulting from the dependence of thermal properties on composition and temperature in a complex, and incompletely understood, manner [50]. The level of difficulty presented by thermo-mechanical modelling may, in fact, be a good reason to first get an initial understanding of a process by mechanics-only modelling, as done so far in this thesis. When thermal modelling is included numerous simplifying assumptions are necessarily made. In relation to glass forming processes, such work as is reported in the published literature mostly concerns the blowing phase in blow moulding after the initial parison has been formed; in [14] some assumptions are made about the temperature or viscosity distribution in the glass before and after mould contact, avoiding solution of the heat flow equation; a coupled heat and fluid-flow model is used in [21, 100] with an initial temperature field established by heat-flow-only computations on an already formed parison prior to commencement of blowing. A recent paper by Simons and Mattheij [83] reports a more comprehensive thermo-mechanical analysis of parison forming which is, however, not continued through the blowing phase.

In considering the development of a full thermo-mechanical model of thermal replication, we are immediately confronted by very complex issues of heat transfer in both glass and ceramic, including composition and temperature dependent thermal properties, and possible heat transfer by radiation between the glass and ceramic. Then, at the end of the day, such a model may not even prove worthwhile since some, or even all, of the factors identified as possibly contributing to non-isothermal conditions may not cause significant temperature variation, and there may be others that do. Furthermore, while the isothermal assumption should be questioned and investigated, it has not yet been established that temperature variations such as might arise, will have an important effect on the slumping process. It must also be borne in mind when contemplating a significant increase to the complexity of a model, as we are now doing, that a simple model that provides fast answers may be of much more value to industry than a complicated and slow model that gives only a little more accuracy. As Yuen [102, p. 30] states

> In general, it is advisable to adopt the simplest model possible for the purpose intended. Overly complicated models are often a cause for the failure of their applications.

In the light of the foregoing discussion, it is apparent that, before attempting to couple heat flow to the fluid-flow model, there is considerable value in conducting some numerical simulations of thermal replication under both isothermal and some arbitrary non-isothermal conditions. A comparison of the isothermal and non-isothermal results will indicate how temperature variations can influence the final outcome. Then, by comparing the results of isothermal and non-isothermal simulations with experimental data we might determine whether, in fact, non-isothermal conditions of real practical importance arise in typical industrial manufacturing circumstances. This methodology is due to the fact that temperatures within a slumping glass disc are not at all easy to measure. Our results should indicate whether a proper thermo-mechanical analysis is warranted.

We look at a comparison of experimental and numerical results in Chapter 8. In this chapter we compare numerical simulations of thermal replication under isothermal and non-isothermal conditions. Different non-isothermal conditions are specified by imposing different temperature distributions in the glass which can be converted to viscosity distributions via the VFT equation. This approach is similar to that of [14], and requires only a simple modification to the finite-element program to allow an arbitrary non-uniform viscosity distribution in the fluid.

In comparing different numerical simulations, and also numerical and experimental results, we need to think about the type and magnitude of differences that should be considered important. This is not easy to decide since it will be greatly determined by specific industry needs, but it seems that, at the least, major differences in the top free-surface curvature profiles would be significant, and hence we shall focus on this in our analyses. Even in this, what constitutes a *major* difference needs to be defined by industry, and we can only indicate the scale of variation that the numerical modelling predicts.

7.2 A Fluid Flow Model with Non-Constant Viscosity

As the first step towards simulating glass slumping under non-isothermal conditions, we next consider modifying the finite-element fluid-flow model to cater for a fluid viscosity that may vary in both time and space, i.e $\mu = \mu(x_i, t)$, employing tensor notation. Let us define a characteristic constant viscosity μ_0 and a dimensionless viscosity $\mu^*(x_i, t)$ such that

$$\mu(x_i, t) = \mu_0 \mu^*(x_i, t). \tag{7.1}$$

With reference to equations (2.4) and (2.5) we now define

$$\mathcal{U} = \frac{\rho g \mathcal{L}^2}{\mu_0}. \tag{7.2}$$

For very viscous fluids such as molten glass, μ_0 is very large ($\mathrm{O}(10^6)\,\mathrm{Pa}\cdot\mathrm{s}$) so that the inertial terms on the left of (2.4) may still be neglected, and, in place of (2.7) we obtain the Stokes creeping-flow equations for a fluid of non-constant viscosity

$$\frac{\partial p^*}{\partial x_i^*} = \frac{\partial}{\partial x_j^*}\left[\mu^*\left(\frac{\partial u_i^*}{\partial x_j^*} + \frac{\partial u_j^*}{\partial x_i^*}\right)\right] - \delta_{2i}. \tag{7.3}$$

This differs from (2.7) only in the μ^* multiplier, and on setting $\mu^* = 1$ for a fluid of constant viscosity we recover that equation. Similarly the no-stress free-surface conditions (2.15) become

$$-p^* n_i + \mu^* n_j \left(\frac{\partial u_i^*}{\partial x_j^*} + \frac{\partial u_j^*}{\partial x_i^*}\right) = 0. \tag{7.4}$$

Following the same route as taken in Chapter 2 and using Green's theorem to obtain the variational form of (7.3), an equation identical to (2.22) is obtained, excepting that each term inside the integrals involving a component of the velocity u_i must be multiplied by μ^*. As then, the boundary integral term is removed from the equation by virtue of (7.4). Likewise, the axisymmetric Stokes equations (2.30) must just have each term inside the integrals involving a component of u_i multiplied by μ^*. This carries through the Galerkin finite-element formulation of the equations, so that at the end of the route μ^* is a factor in some of the integrals to be evaluated using a numerical quadrature rule. To compute these integrals by means of the quadrature rule, the value of μ^* must be known at each quadrature point.

Now, it is most convenient to store the value of μ^* at each mesh node, and this is done. However, a transformation similar to those in (2.48) for mapping a mesh element onto the master triangular element, maps the viscosity field applicable over a mesh element onto the master element, namely

$$\mu^*(\zeta_i) = \sum_{j=1}^{6} \mu^{*j}\hat{\phi}^j(\zeta_i), \tag{7.5}$$

where the μ^{*j} are the (dimensionless) viscosity values at the six mesh nodes $j = 1, 2, \ldots, 6$, the $\hat{\phi}^j$ are the quadratic basis functions defined for the master element (see Figure 2.1), and ζ_i, $i = 1, 2, 3$ are the area coordinates of a point within the

master element. On substituting the coordinates of each quadrature point, the corresponding values of μ^* are obtained for use in evaluating integrals with the chosen quadrature rule.

The creeping-flow equations can now be solved with spatially-varying viscosity specified at the mesh nodes. In order to run a slumping simulation with a spatially-varying viscosity field that also changes in time with temperature the following procedure has been used. First note that below some known temperature T_s the glass can be considered a solid that will not flow, so that time $t = 0$ may be defined as that point in time when the maximum temperature in the glass is T_s. Then, considering the temperature at a node n to be a function of time, i.e. $T^n = T^n(t)$, the initial temperature distribution within the glass is given by specifying the temperature $T^n(0)$ at each mesh node such that $\max\{T^n(0)\} = T_s$. The maximum temperature $T_{\max}$ to which the glass will be heated is also specified, as well as some heating-rate rule that defines how the temperature at each node n will vary from its initial value $T^n(0)$ to $T_{\max}$. On reaching $T_{\max}$ it is assumed that the temperature remains at that value. Next, it is assumed that over a specified time interval $\Delta t_T = t_2 - t_1$, the temperature at any node remains fixed at $T^n(t_1 + \Delta t_T/2)$. Then the algorithm is as follows:

1. set $t = 0$;

2. compute $T^n(t + \Delta t_T/2)$ using the known heating-rate rule, for all nodes n;

3. compute $\mu^* = \mu/\mu_0$ at each node by substituting the known nodal temperatures into the VFT equation (5.1) with constants appropriate to the particular glass under consideration;

4. compute the slumping of the glass over a time interval Δt_T using the algorithm given on page 152, but now with the Stokes equations for a non-constant viscosity field;

5. set $t = t + \Delta t_T$;

6. if $t = t_{\max}$ then stop (or use some other stopping criterion), else go to step 2.

The finite-element method is very convenient in removing the need to compute the gradients of μ^* that arise in (7.3), so that the problem with spatially varying viscosity is not much harder than the constant-viscosity problem, basically only requiring a little extra in computer resources and book-keeping to store the temperature and viscosity at each mesh node, and then use these in computations.

7.3 Isothermal Slumping - A Check

In Chapter 5 a method was described for computing slumping in circumstances where the viscosity is spatially constant but varies in time due to a time-varying temperature. We saw that a time-varying viscosity is equivalent to a time-varying time scale, and that slumping with a time-varying viscosity can be computed by post-processing of results from a constant-viscosity simulation having the same initial geometry. A spatially constant time-varying slumping problem may also be computed with the non-constant-viscosity model described in the previous section, by specifying $T^n(0) = T_s$ for all nodes n and using a heating-rate rule that changes the temperature by the same amount at every node. Comparing such a simulation with the results obtained from post-processing of a constant-viscosity simulation provides a check of the non-constant-viscosity model.

Let us again consider the initially-flat glass disc of radius a slumping into a mould as used to demonstrate numerical slumping in Chapter 6 (Section 6.3). Suppose that the disc is a typical soda-lime-silica glass of density $\rho = 2500\,\mathrm{kg/m^3}$, and with reference viscosities as given in Table 5.2 yielding VFT equation constants $A = 2.0886$, $B = 4915.8820$, and $T_0 = 192.6136$. It is reasonable in a slumping context to consider glass to be solid at its annealing point, and hence define $T_s = 510°\mathrm{C}$. We consider the slumping of the disc into the mould as the temperature is increased at

5°C per minute from 510°C at $t = 0$ to a maximum temperature of $T_{\max} = 700$°C, after which the temperature is then held constant at 700°C.

Following the algorithm given in the previous section, a uniform initial temperature distribution is first specified by setting $T^n(0) = 510$°C for all nodes n. The heating-rate rule is, of course, just

$$T^n(t) = \begin{cases} 510 + 5t, & 0 \le t \le 38, \\ 700, & t \ge 38, \end{cases} \tag{7.6}$$

where t is the elapsed time in minutes, and $T^n(t)$ is the temperature at node n and time t in °C. Next, the temperature is assumed to remain constant during time intervals of $\Delta t_T = 1$ minute. Then the temperature in the ith interval at each node n is given by

$$T_i^n = \min(510 + 5(i - 1/2), 700) \tag{7.7}$$

Choosing $\mu_0 = 10^6\,\mathrm{Pa}\cdot\mathrm{s}$, the dimensionless viscosity μ_i^* can be computed at each mesh node during the ith one minute constant-temperature time interval. The time scale of these computations is constant at $\overline{\mathcal{T}} = \mu_0/(\rho g a) \approx 15.1172$ minutes, so that in dimensionless terms the constant-temperature time intervals are of length $\Delta t_T^* = 0.06615$. Within each of these intervals the slumping of the disc is computed according to the algorithm on page 152, using a maximum dimensionless time step of $\Delta t^* = 0.01$ which is reduced as necessary to prevent nodes from passing through the mould boundary.

Because we wish to compare the results of this numerical slumping simulation with the constant-viscosity simulation conducted in Chapter 6, it should be stopped at the point corresponding to Figure 6.5, i.e. at $t = 0.075\mathcal{T}$, where $\mathcal{T} = \mu(t)/(\rho g a)$ is the time scale of the constant-viscosity simulation. Note that the time scale $\mathcal{T}$ for the constant-viscosity simulation is given as a function of time t, as a reminder that, although referred to as a 'constant-viscosity' simulation, a time-varying viscosity, as in the present problem, is readily accommodated in a time-varying time scale as

Table 7.1: Dimensionless time t^* after the ith one minute interval at temperature $T_i\,^\circ$C and viscosity $\mu_i^* \times 10^6\,\mathrm{Pa}\cdot\mathrm{s}$, as given by equation (7.8) for a constant-viscosity slumping simulation.

i	T_i	μ_i^*	t^*
29	652.5	3.987847E+01	6.614433E-03
30	657.5	3.060357E+01	8.775945E-03
31	662.5	2.361843E+01	1.157672E-02
32	667.5	1.832739E+01	1.518608E-02
33	672.5	1.429700E+01	1.981292E-02
34	677.5	1.121019E+01	2.571380E-02
35	682.5	8.833594E+00	3.320226E-02
36	687.5	6.994445E+00	4.265977E-02
37	692.5	5.564134E+00	5.454841E-02
38	697.5	4.446408E+00	6.942559E-02

described in detail in Chapter 5. We need then to compute the value of $t = 0.075\mathcal{T}$ so as to find the dimensionless time in terms of the time scale $\overline{\mathcal{T}}$ of the present simulation, i.e. $t^* = t/\overline{\mathcal{T}}$, at which to stop this simulation.

For a constant-viscosity slumping simulation, the relationship between true and dimensionless time through a time-varying time scale is given by equation (5.38). For the present case, this becomes

$$t^* = \frac{\rho g a}{\mu_0}\left(\sum_{i=0}^{n-1}\frac{\Delta t_T}{\mu_i^*} + \frac{t-(n-1)\Delta t_T}{\mu_n^*}\right), \tag{7.8}$$

where $(n-1)\Delta t_T < t \le n\Delta t_T$ is the true time, and t^* is the corresponding dimensionless time. It is then a simple matter to compute the dimensionless time after each one minute constant-temperature time interval over the 38 minute time period during which the temperature increases from 510 to 700°C. This is shown, for the last part of the heating period, in Table 7.1 from which we see that the maximum temperature is reached at $t^* = 0.069426$. From $t^* = 0.069426$ to the stopping point at $t^* = 0.075$ the temperature is held constant at 700°C, so that $\mu^* = 3.9814$. This adds a further 20.13 seconds to the slump time, giving a total

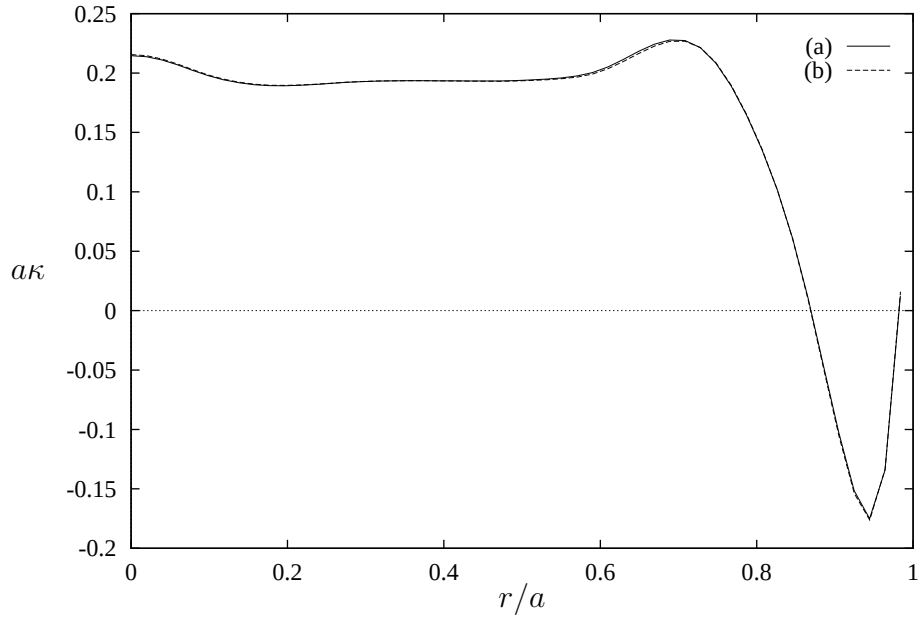

Figure 7.1: Surface curvature of a slumped disc computed using a quintic B-spline approximation. (a) Non-constant-viscosity simulation and (b) constant-viscosity simulation.

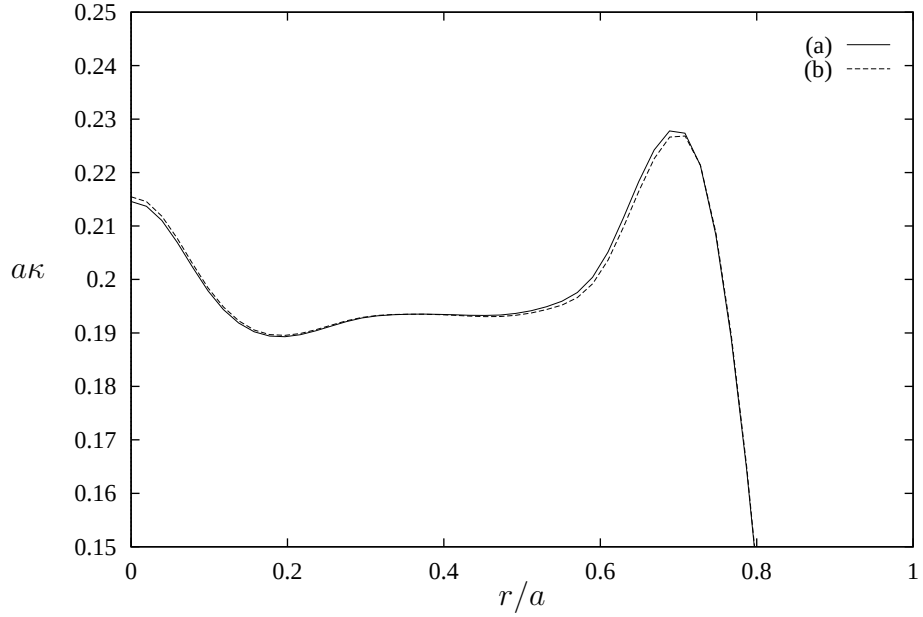

Figure 7.2: As above but at a larger scale.

time of 38.3355 minutes. For the non-constant viscosity model having a constant time scale of $\overline{\mathcal{T}} = 15.1172$ minutes, this represents a dimensionless slumping time of $t^* = 2.5359$, at which the simulation was stopped.

Figures 7.1 and 7.2 compare the top free-surface curvature of the slumped disc after full mould contact as obtained from the constant-viscosity simulation of Chapter 6 and the present non-constant-viscosity simulation. Considering the different numerical slumping methods the two curves are remarkably similar, and we can certainly be satisfied with the accuracy of the non-constant-viscosity finite-element program.

7.4 Temperature Variation in Slumping Glass

We could now proceed directly to compare slumping simulations with identical initial geometry but a variety of arbitrary temperature distributions in the glass, and so examine the sensitivity of thermal replication to non-isothermal conditions. However, before we do this we briefly consider the magnitude of temperature variations that are likely to arise in thermal replication.

In this connection, we first note that, as part of their disc sagging experimental work, Gulati et al. [41, 42] made some measurements of temperature uniformity within the region occupied by the glass disc during a sagging experiment. In [41] they report a radial temperature increase of about 2.6°C over about 30 mm, and a vertical temperature increase of less than 0.5°C over 10 mm. In [42] they report a similar radial temperature increase of less than 3°C over 30 mm, but about 4°C over 10 mm vertically. This last result is somewhat surprising, and is ignored because an error is suspected in the presentation. Otherwise, the measurements indicate a temperature gradient of up to 1°C in 10 mm. These measurements were made with a small oven in carefully controlled laboratory conditions, and greater oven-temperature variations can be anticipated in typical industrial manufacturing conditions. Oven

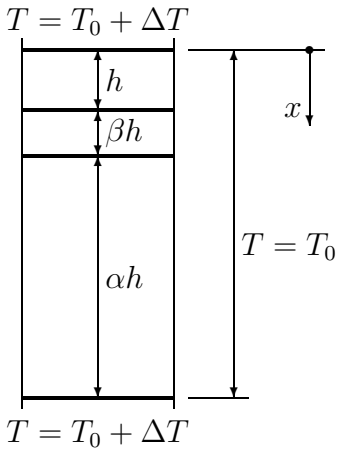

Figure 7.3: A simple one-dimensional heat-conduction model at $t = 0$.

air temperature is not, of course, necessarily the same as glass temperature, and hence it can only be used as indicative of glass temperature in the absence of better data.

At the beginning of this chapter mention was made of the difficulties presented by thermal modelling of industrial processes, because of which it was decided to investigate the effects of arbitrary spatial temperature variations on thermal replication as a precursor to possible full thermo-mechanical modelling of the process. However, in order to get some very rough feel for the magnitude of spatial temperature variation that might arise due to differential heat transfer in the thermal-replication process, the very simple one-dimensional problem shown in Figure 7.3 has been considered. Here we have a 'sandwich' consisting of an upper layer of glass of thickness h, a middle layer of air of thickness βh, and a bottom layer of ceramic of thickness αh being heated in an oven. Initially the temperature in the glass, air, and ceramic is given by $T = T_0$. Then, at time $t = 0$ the oven temperature, which is assumed to be uniform, is increased to $T = T_0 + \Delta T$, and it is supposed that the temperature of the surfaces in direct contact with the oven environment at $x = 0$ and $x = (1 + \alpha + \beta)h$, also instantaneously reach this value. The temperature profile in the glass, air and ceramic is then computed after a time period Δt, assuming heat transfer by conduction only. At the end of this period the oven temperature, and

hence the temperature at $x = 0$ and $x = (1 + \alpha + \beta)h$, is again increased by an amount ΔT, and the resulting temperature profile computed at the end of another time period Δt, and so on.

The one-dimensional geometry is justified for a glass disc supported on a ceramic mould as in thermal replication, where the disc radius is large compared to its thickness h, and where the radial change in the thicknesses of the air gap and ceramic mould is slow. Then the computed temperature profile is locally relevant at a radial position appropriate to the choice of α and β, and an idea of the radial temperature distribution is given by solving the one-dimensional model with different values of α and β corresponding to different radial positions in the glass-mould combination. Only in a narrow band around the perimeter of the disc will the additional edge surface area of both disc and mould contribute anything to the heating process, and this is ignored in the computations.

In assuming heat transfer by conduction alone, any natural convection in the air gap is neglected, which seems reasonable because of the small thickness of this gap. By contrast, neglect of radiative heat transfer between the glass and mould across the air gap has no physical justification. However, since at this stage we only want to get an idea of the magnitude of possible temperature variation in slumping glass, this neglect is justified by the resulting computational simplification.

This model has been used to compute the temperature differences that might arise in a glass disc and ceramic mould combination as it is heated from normal ambient temperatures to a temperature in excess of the annealing point where the glass will begin to sag. For the isothermal slumping case just computed in the previous section, significant sagging began at a temperature of about 600°C, so that, at a heating rate of 5°C per minute there was a heating period of some 115 minutes (from 25°C) during which non-isothermal conditions could arise to affect the glass flow in the following time period of a little over 20 minutes during which the glass slumped into the mould.

The one-dimensional equation of heat conduction in a solid of density ρ, thermal conductivity k, and specific heat c, where there are no heat sources, is [75]

$$\rho c \frac{\partial T}{\partial t} - \frac{\partial}{\partial x}\left(k \frac{\partial T}{\partial x}\right) = 0, \tag{7.9}$$

where it is allowed that all thermal properties, and k in particular, may vary with temperature and hence in both time and space. If k is independent of x at any point in time, then (7.9) may be written in terms of the diffusion coefficient $\alpha = k/(c\rho)$

$$\frac{\partial T}{\partial t} - \alpha \frac{\partial^2 T}{\partial x^2} = 0. \tag{7.10}$$

For boundary conditions, we have

$$T = T_0 + i\Delta T, \qquad \text{for} \quad i\Delta t \leq t < (i+1)\Delta t \tag{7.11}$$

at $x = 0$ and $x = (1 + \alpha + \beta)h$, while at the interface of two different materials, with thermal conductivities k_1 and k_2, and temperature profiles $T_1(x)$ and $T_2(x)$, we must satisfy [17, p. 13]

$$k_1 \frac{\partial T_1}{\partial x} = k_2 \frac{\partial T_2}{\partial x} \qquad \text{and} \tag{7.12}$$

$$T_1 = T_2. \tag{7.13}$$

Before (7.9) or (7.10) can be solved subject to the boundary conditions (7.11) to (7.13) data is required for ρ, c, and k for each of the three different materials involved in the problem (i.e. glass, air, and ceramic), and for how these vary with temperature in the range $25 \leq T \leq 600°\text{C}$.

Already we know from earlier discussion of glass properties in Chapter 5 that the density of a glass changes only very little with temperature, and that a constant density of $\rho = 2500\,\text{kg/m}^3$ is typical for soda-lime-silica glasses. A plot given in [81, p. 25] shows the specific heat of a soda-lime-silica glass increasing from 0.18 to $0.31\,\text{cal/(g°C)}$ over the temperature range of present interest (see Table 7.2), and this data is reasonably well confirmed by a formula given in [5, p. 266] relating temperature and true specific heat in the temperature range $500°\text{C} \leq T \leq 1400°\text{C}$. The

Table 7.2: Typical thermal properties for soda-lime-silica glasses.

T °C	k cal/(s · cm · °C)	c cal/(g°C)	α cm^2/s
25	0.003	0.18	0.006
100	0.003	0.22	0.005
200	0.003	0.25	0.005
400	0.0035	0.29	0.005
600	0.004	0.31	0.005

conductivity of a semi-transparent solid such as glass is a more complex subject, being influenced by the thickness of the glass sample, and involving "radiation conductivity" due to internal radiation within the glass, as discussed in [5, 11, 50, 81]. From the data given in these sources, it seems reasonable for soda-lime-silica glass discs of around 5 mm in thickness, as typically used in thermal replication, to assume that the thermal conductivity is constant at around $0.003\,\mathrm{cal/(s \cdot cm \cdot °C)}$ to a temperature of 200°C, and that it then increases linearly to $0.004\,\mathrm{cal/(s \cdot cm \cdot °C)}$ at 600°C. With such a small increase in the conductivity for a fairly large increase in temperature of 400°C, and since we would not expect a temperature difference of this magnitude across the thickness of the glass, it is reasonable to take k to be independent of x and work with (7.10) and the diffusion coefficient α. As we see from Table 7.2 it is quite justifiable to use a constant diffusion coefficient of $\alpha = 0.005\,\mathrm{cm^2/s}$ for glass.

There are a number of different classes of engineering ceramics having widely varying properties, and even within the same class properties can vary significantly (see [79]). As well as chemical composition, temperature and the method of manufacture which determines the porosity of the material, strongly influence thermal properties [39, Ch. 14]. Ceramics suitable for the moulds used in thermal replication must be thermally stable and must not adhere to glass [84]. They must also be machinable, and from observation they can be quite porous. With respect to chem-

Table 7.3: Typical thermal properties for alumina ceramic.

T °C	k cal/(s · cm · °C)	α cm^2/s
25	0.093	0.112
100	0.066	0.079
400	0.032	0.038
800	0.015	0.018

ical composition of ceramics used for thermal replication, no specific information could be found. However, the metal oxide ceramics alumina and stabilized zirconia are suitable for the high temperature corrosive environments of glass melting furnaces [12, p. 71], and it seems likely that these would also be used for molten glass forming in thermal replication. Alumina is a much better heat conductor than zirconia [79, p. 249] but otherwise their properties are expected to be similar. Assuming alumina for our current problem, we have from [12, p. 377] a density of $3.96\,\mathrm{g/cm^3}$, a specific heat at 25°C of $0.21\,\mathrm{cal/(g°C)}$, which is assumed to be constant over the temperature range of present interest, and a conductivity varying with temperature as shown in Table 7.3. Although the variation of conductivity with temperature is larger than for glass, it is still sufficiently small to consider k to be independent of x for the magnitude of temperature variation that can be expected in the ceramic. Thus, values of the diffusion coefficient α have been computed as shown in Table 7.3. Values of k and α at other temperatures are computed using

$$k = -0.051 \log_{10} T + 0.164, \qquad \text{and} \tag{7.14}$$

$$\alpha = -0.061 \log_{10} T + 0.197, \tag{7.15}$$

where T is in °C, k is in cal/(s · cm · °C), and α is in cm^2/s.

Air behaves very nearly as an ideal gas, and hence obeys Charles's law which states that at constant pressure the volume occupied by a unit mass of air is directly proportional to the absolute temperature [87, p. 29-4]. Then, using data given in

Table 7.4: Thermal properties of air.

T °C	ρ g/cm^3	k cal/(s · cm · °C)	c cal/(g°C)	α cm^2/s	m, b $\alpha = mT + b$
4	1.274×10^{-3}	0.579×10^{-4}	0.239	0.190	1.551×10^{-3}, 0.1838
60	1.060	0.703	0.240	0.276	
171	0.795	0.868	0.243	0.449	
282	0.636	1.074	0.249	0.678	2.108×10^{-3}, 0.0885
393	0.530	1.240	0.255	0.917	
504	0.454	1.446	0.261	1.220	2.703×10^{-3}, -0.1452
615	0.398	1.612	0.267	1.517	

[7, p. 594] we have

$$\rho = \frac{0.353}{273 + T} \tag{7.16}$$

where ρ is in g/cm^3 and T is in °C. Data for specific heat and conductivity of air at various temperatures, as obtained from [87, p. 33-12], is given in Table 7.4. Compared to both the glass and ceramic, air is seen to be a poor heat conductor and can be expected to act as an insulator between these two layers. Again it seems reasonable to consider k as independent of x, since we can expect temperature differences of much less than 100°C across the air gap. Table 7.4 gives the diffusion coefficient α at various temperatures, calculated using the data from [87] and equation (7.16) for density. A plot of α versus T indicates that the relationship can be approximated by piecewise continuous line segments $\alpha = mT + b$ over the intervals $4 - 171$°C, $171 - 393$°C, and $393 - 615$°C, as given in the last column of Table 7.4. For computational purposes, a relationship is also needed between k and T for air. This is close to linear and can be approximated by

$$k = (1.638 \times 10^{-3}T + 0.605) \times 10^{-4} \tag{7.17}$$

with k in units of cal/(s · cm · °C) and T in °C.

Sufficient information has now been given to solve the simple heat conduction problem. Because thermal conductivity k is assumed to be independent of x in each

of the three layers of the 'sandwich', the appropriate equation is (7.10), subject to the boundary conditions (7.11) to (7.13). Let us denote the diffusion coefficient for the glass by α_g. Then, non-dimensionalizing using a length scale $\mathcal{L} = h$ and a time scale $\mathcal{T} = h^2/\alpha_g$, and using asterisks to denote dimensionless variables, gives

$$\frac{\partial T}{\partial t^*} - \frac{\alpha}{\alpha_g}\frac{\partial^2 T}{\partial x^{*2}} = 0 \tag{7.18}$$

in place of (7.10), and

$$k_1\frac{\partial T_1}{\partial x^*} = k_2\frac{\partial T_2}{\partial x^*} \tag{7.19}$$

in place of (7.12).

To solve (7.18) in each of the glass, air, and ceramic layers, the Crank-Nicholson finite-difference method [64, Ch. 3, p. 31] is used. Defining a one-dimensional uniform grid on $0 \leq x^* \leq (1 + \alpha + \beta)$ such that

$$x_j^* = j\Delta x^*, \qquad j = 0, 1, \ldots, N_x \tag{7.20}$$

with $\Delta x^* = (1 + \alpha + \beta)/N_x$ chosen so that a grid point falls at every material interface, and specifying a time step Δt^* with

$$t_i^* = i\Delta t^*, \tag{7.21}$$

equation (7.18) becomes

$$-sT_{j-1}^{n+1} + 2(1+s)T_j^{n+1} - sT_{j+1}^{n+1} = s(T_{j-1}^n + T_{j+1}^n) + 2(1-s)T_j^n, \tag{7.22}$$

where $s = (\alpha/\alpha_g)\Delta t^*/(\Delta x^*)^2$ and $T_j^n = T(x_j^*, t_n^*)$ is the temperature at the jth grid point and the nth time step. At a material interface, simple forward and backward differencing is used to compute the temperature gradients, so that (7.19) becomes

$$-k_1T_{j-1}^{n+1} + (k_1 + k_2)T_j^{n+1} - k_2T_{j+1}^{n+1} = 0. \tag{7.23}$$

Equations (7.22) and (7.23) lead to a tri-diagonal system of equations that must be solved for the unknowns T_j^{n+1}. The initial conditions at time $t = 0$ are

$$T_j^0 = T_0, \qquad j = 0, \ldots, N_x, \tag{7.24}$$

and at the $(n+1)$th time step the boundary conditions are given by

$$T_j^{n+1} = T_j^n + \Delta T, \qquad j = 0, N_x, \tag{7.25}$$

where ΔT is the increase in the temperature in one time step.

Now, the glass, air, and ceramic have thermal properties that may change in time with temperature, although they are assumed to be constant in any one time step. The values of the material properties to be used in the $(n+1)$th time step, are just computed in the manners discussed previously using the average of the temperatures on the exterior boundaries at the n and $(n+1)$th time steps,

$$\overline{T} = (T_0^n + T_0^{n+1})/2. \tag{7.26}$$

Let us compute the temperature distribution that develops in the flat glass disc considered in the previous section and in Chapter 6. Let the glass thickness be $h = 0.6\,\text{cm}$ corresponding to a radius of $a = 4.5\,\text{cm}$. Then the time scale is $\mathcal{T} = h^2/\alpha_g = 0.36/0.005 = 72\,\text{seconds} = 1.2\,\text{minutes}$, a time step size of $\Delta t^* = 1$ is, for a heating rate of 5°C/min., equivalent to temperature increments at each step of $\Delta T = 6°\text{C}$, and 96 such steps takes the temperature from 25 to 601°C in a time of 115.2 minutes. For this particular slumping case, the air gap increases from nothing at $r = a$ to $0.0947a \approx 0.7h$ at $r = 0$, so that we need to consider $0 \le \beta \le 0.7$. Let the mould thickness be $5h = 3\,\text{cm}$ at the edge, so that $\alpha = 5 - \beta$, and let $\Delta x^* = 0.02$.

Figure 7.4 shows temperature profiles across the glass-air-ceramic 'sandwich' that have developed while heating from 25 to 313°C, i.e. half way through the heating period being considered. Figure 7.5 shows similar profiles at the end of this period when the oven temperature has reached 601°C. Profiles are shown close to the centre of the disc ($r/a = 0.12$, $\beta = 0.7$), near the edge ($r/a = 0.9277$, $\beta = 0.1$), and at the edge ($r/a = 1$, $\beta = 0$). In Figure 7.6 we have the relationship between radial position and temperature on the lower face of the glass, after the oven temperature has reached 601°C. These figures show the air gap acting as an insulating layer (as anticipated), slowing heat conduction from the glass into the mould. However, in

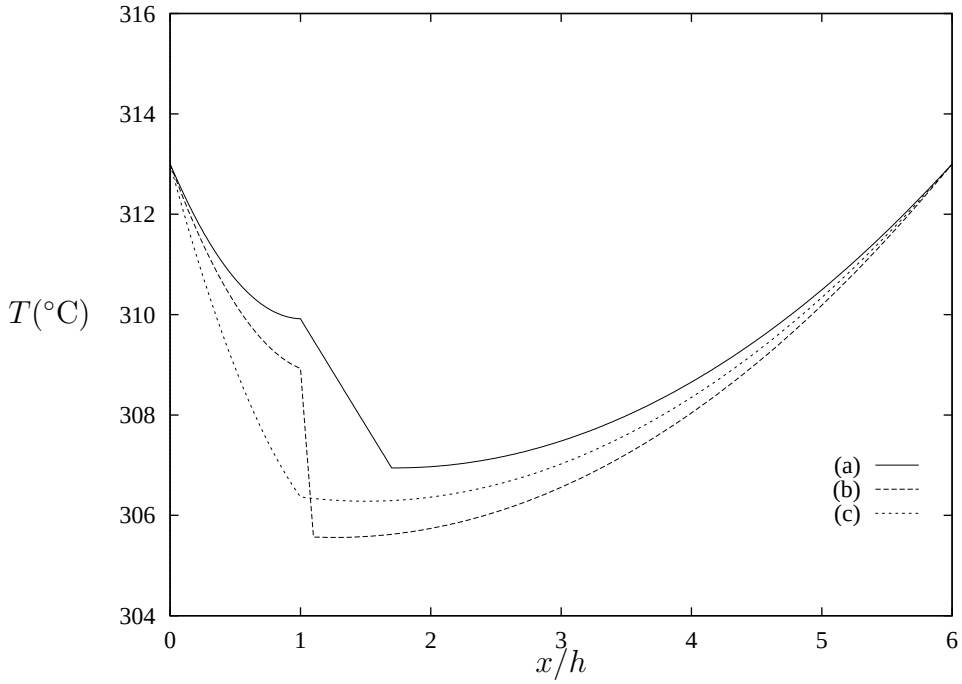

Figure 7.4: Temperature variation across a glass-air-ceramic 'sandwich' at $t^* = 48$, with (a) $\alpha = 4.3$, $\beta = 0.7$, (b) $\alpha = 4.9$, $\beta = 0.1$, and (c) $\alpha = 5$, $\beta = 0$.

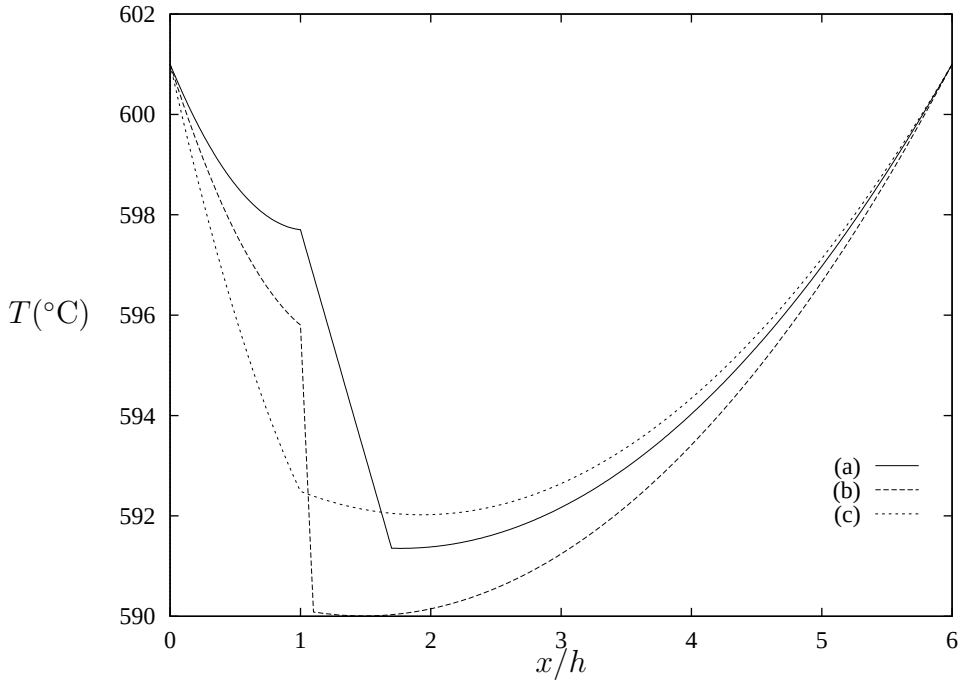

Figure 7.5: As above but at $t^* = 96$.

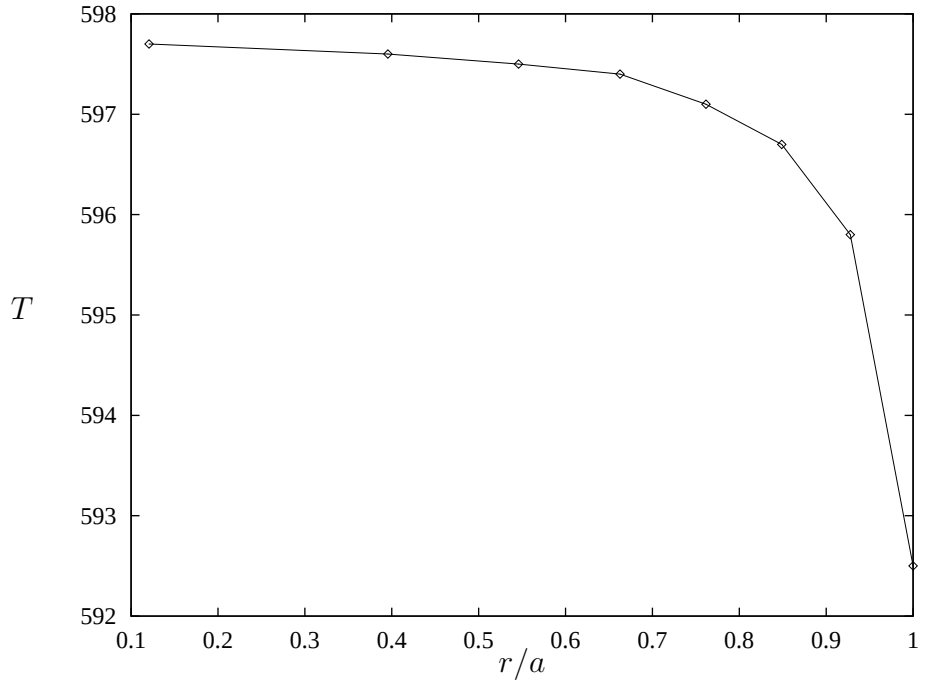

Figure 7.6: Relationship between radial position and temperature at the bottom surface of the glass for the disc-mould combination of Figure 6.3.

practice it is likely that the decrease in glass temperature towards the edge will be less rapid than indicated, because of the additional surface area of both glass and mould at the edge which is ignored in the model. Thus we might expect a radial decrease in temperature of $1-2$°C, and, in view of the shape of the curve plotted in Figure 7.6, it is reasonable to assume that the temperature varies linearly with disc radius. These computations also show an approximately linear temperature drop of $3-5$°C across the thickness of the glass.

The temperature gradients in the glass change little on increasing the temperature from 313 to 601°C (cf. Figures 7.4 and 7.5), and since the radial gradient caused by the changing thickness of the air gap (see Figure 7.6) is considerably smaller than the vertical gradient, it is reasonable to suppose that these gradients would remain throughout the slumping of the glass as the air gap decreases in size. Contact of some portion of the lower glass surface with the ceramic mould might result in a fairly sudden temperature drop of perhaps 4 or 5°C in that portion of the

glass, but this is unlikely to significantly affect glass flow, since, as we have already seen, mould contact itself, in the absence of a change in temperature, greatly arrests glass flow.

These results, and their interpretation, must be treated with a considerable degree of caution, considering both the assumptions underlying the simple heat-conduction model used, and also the lack of specific thermal-properties data for the glass and ceramic. Nevertheless, they are used in the following section as a rough guide in deciding the temperature distributions in the glass to be considered.

7.5 Non-Isothermal Replication

By far the steepest temperature gradient suggested by our considerations so far, is 5°C in 0.6 cm through the thickness of the glass disc given by the heat-flow computations. (The gradient suggested by the oven temperature measurements of Gulati et al. is considerably less than this.) Because of this, let us first consider the influence of linear temperature variations across the thickness of the glass on the top-surface curvature. Although the heat-conduction model used resulted in the glass temperature varying from hottest at the top surface to coldest at the bottom surface, results are also given for the reverse case where the bottom surface is hottest, just to see how this compares. For this work we take a glass disc of 4.5 cm radius and 0.6 cm thickness, as used for the heat-flow computations.

The computational method is the same as used in Section 7.3, excepting for the definition of the initial temperature distribution. Again, let $t = 0$ be the time at which the maximum temperature in the glass reaches the annealing point of 510°C, and let the simulation start at this point in time. However, instead of setting the initial temperature at all mesh nodes at 510°C, it is assumed that during the heating period prior to $t = 0$ a straight-line temperature gradient develops across the thickness of the glass disc, so that the temperature at each node is set according

to its z-coordinate. The simulation then proceeds exactly as before. Thus, at any node a constant temperature gradient is assumed until the temperature at that node reaches the maximum $T_{\max}$ to which the glass is to be heated, which seems reasonable based on the earlier heat-conduction computations.

Let T_d be the difference between the upper and lower surfaces of the glass, with a positive value denoting a temperature increase in the direction of the z-axis. Hence, a positive value of T_d indicates an initial linear temperature distribution varying from $510 - T_d\,^\circ$C on the lower surface $z = 0$ to 510°C on the upper surface $z = h$, i.e.

$$T^n(0) = 510 - T_d(1 - z/h), \tag{7.27}$$

where $T^n(0)$ is the temperature at node n having vertical coordinate z, at $t = 0$. A negative value of T_d indicates the reverse distribution, i.e.

$$T^n(0) = 510 - T_d z/h. \tag{7.28}$$

Figures 7.7 and 7.8 show the top-surface curvature of the glass disc after full mould contact has been established for different values of T_d, with $T_d = 0$°C being just the isothermal case computed in Section 7.3. Note that the time taken to achieve full mould contact increases slightly as $|T_d|$ is increased. These figures shows that even a probably excessively steep temperature gradient of ±10°C across the 0.6 cm thickness of the disc has only a small effect on the finished top-surface curvature. Relative to the isothermal curve, at the centre of the disc a small decrease in curvature results from a temperature increasing in the z-direction, while a small increase in curvature results from a temperature decreasing in the z-direction. In this region the maximum curvature variation from the isothermal curve is about 1.5 percent for $T_d = 10$°C and less than one percent for $T_d = 5$°C. Much of the variation between the curves for the different slumping conditions seen towards the edge ($r/a > 0.6$), is quite likely attributable to the fact that we are not comparing at exactly equivalent instants in the slumping process. Overall, considering that the

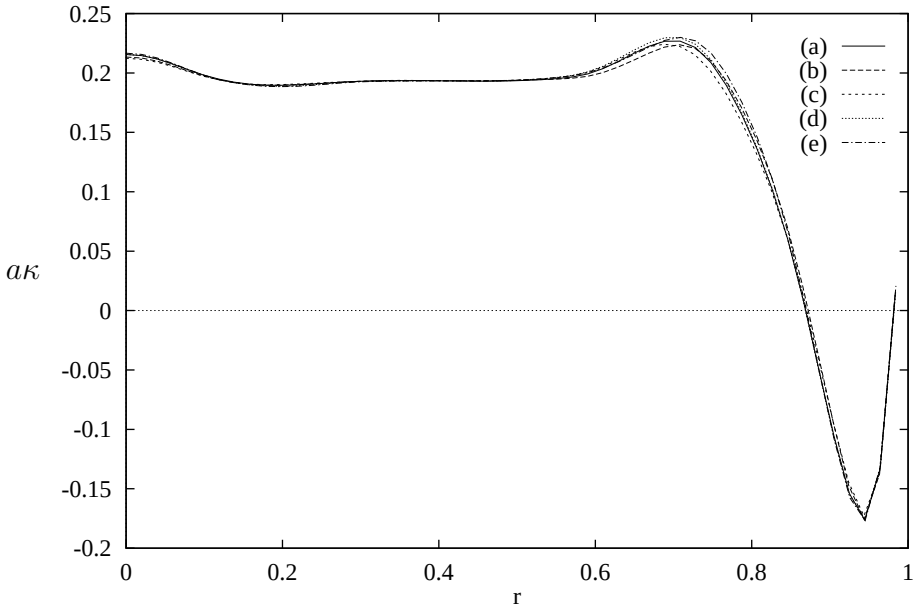

Figure 7.7: Top-surface curvature resulting from a temperature drop T_d from the top surface to the bottom surface of the glass disc of aspect ratio $h/(2a) = 0.0667$. (a) $T_d = 0°\mathrm{C},\ t = 38.3355\,\mathrm{min.}$, (b) $T_d = 5°\mathrm{C},\ t = 39\,\mathrm{min.}$, (c) $T_d = 10°\mathrm{C},\ t = 39\,\mathrm{min.}$, (d) $T_d = -5°\mathrm{C},\ t = 39\,\mathrm{min.}$, and (e) $T_d = -10°\mathrm{C},\ t = 40\,\mathrm{min.}$

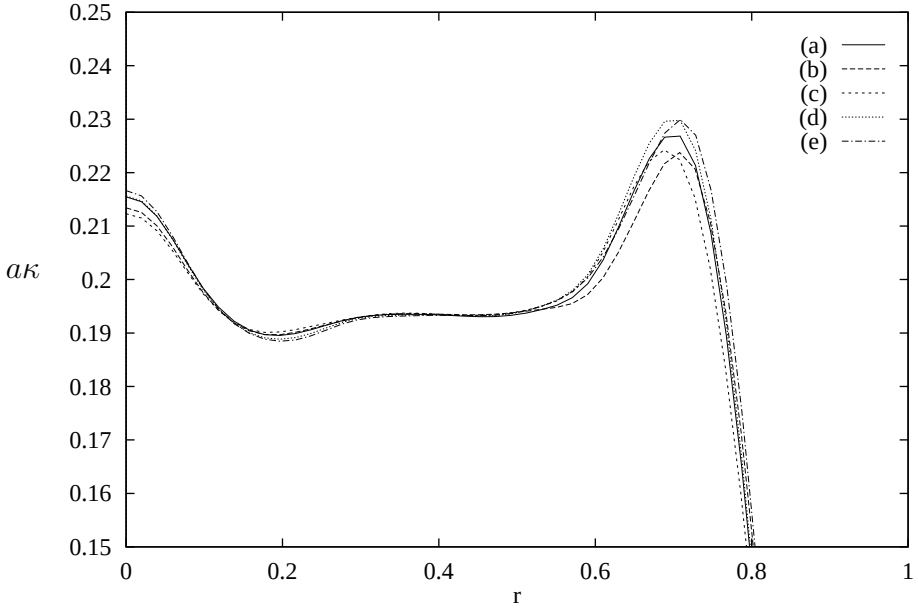

Figure 7.8: As above but at a larger scale.

method for calculating curvature is subject to some degree of uncertainty and cannot be considered to give an exact profile, it is reasonable to interpret these results as suggesting that under these non-isothermal conditions, the disc basically slumps as though conditions were isothermal but at a lower effective temperature and hence higher effective viscosity, increasing the slumping time but not much affecting the final product. From this we can draw the important conclusion that slumping is not particularly sensitive to transverse temperature gradients, and hence that transverse temperature gradients in the glass can be ignored at this stage in this work.

Next we consider the effect of linear radial temperature gradients in the glass. The Gulati et al. measurements of air temperatures in their experimental oven suggest a possible gradient of 1°C in 0.1 cm or more, i.e. 4.5°C or more across the radius of the slumping disc under present consideration, while the heat-flow computations of Section 7.4 suggested a radial temperature decrease of $1 - 2$°C. We consider the effect of a temperature difference of 5°C, which seems a possibility, and also a larger temperature difference of 10°C, which is probably quite excessive. This time, let T_d be the temperature increase in the direction of the r-axis, from the centre to the edge of the disc. Thus, a positive value of T_d indicates an initial linear temperature distribution varying from $510 - T_d$ °C at $r = 0$ to 510°C at $r = a$, i.e. the temperature at node n with radial coordinate r, at $t = 0$, is given by

$$T^n(0) = 510 - T_d(1 - r/a). \tag{7.29}$$

A negative value of T_d indicates the reverse initial temperature distribution, i.e.

$$T^n(0) = 510 - T_d r/a. \tag{7.30}$$

Figures 7.9 and 7.10 show that a radial temperature gradient has an effect on surface curvature at the centre of the disc which is considerably greater than seen with a transverse temperature gradient, even though the gradient of up to ±10°C in 4.5 cm is small relative to the transverse temperature gradient of up to ±10°C in 0.6 cm. A comparison of the variation from the isothermal result for transverse

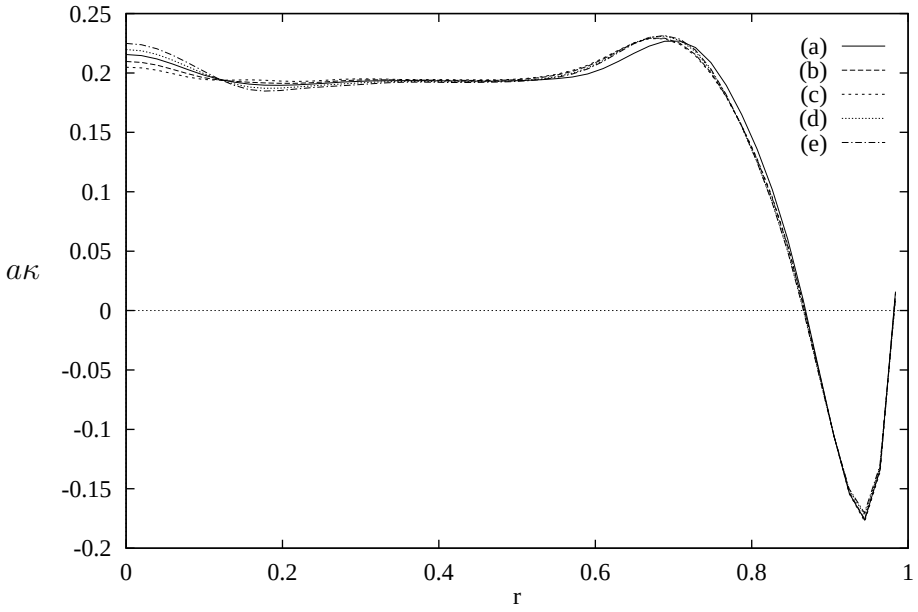

Figure 7.9: Top-surface curvature resulting from a temperature drop T_d from the centre to the edge of the glass disc of aspect ratio $h/(2a) = 0.0667$. (a) $T_d = 0°\mathrm{C},\ t = 38.3355\,\mathrm{min.}$, (b) $T_d = 5°\mathrm{C},\ t = 38\,\mathrm{min.}$, (c) $T_d = 10°\mathrm{C},\ t = 38\,\mathrm{min.}$, (d) $T_d = -5°\mathrm{C},\ t = 39\,\mathrm{min.}$, and (e) $T_d = -10°\mathrm{C},\ t = 40\,\mathrm{min.}$

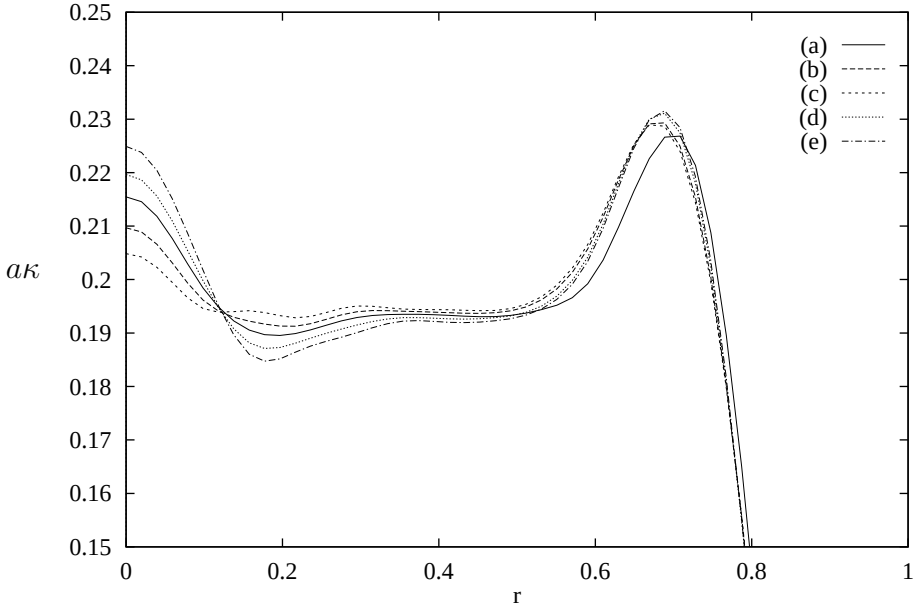

Figure 7.10: As above but at a larger scale.

Table 7.5: Maximum variation in curvature from the isothermal result for transverse and radial linear temperature gradients.

T_d °C	Percent Variation	
	Transverse	Radial
−10	+0.54	+4.38
−5	+0.08	+1.93
+5	−0.96	−2.69
+10	−1.47	−4.93

and radial temperature gradients is given in Table 7.5. A clear pattern is evident in the figures with the central curvature profile becoming flatter for a temperature that increases in the r-direction, and showing much more change with a temperature that decreases in the r-direction. Note that a temperature that decreases with r can be expected from the heat-flow computations. Again the non-isothermal conditions cause a small increase in the time taken for full mould contact, and the variations of curvature toward the edge of the disc are attributable to disc slumping having progressed to slightly different points at the times of the curvature profile 'snap-shots'. These results lead to the conclusion that any radial temperature variations of around 5°C over 4.5 cm or more that occur, should be included in numerical simulation of thermal replication.

7.6 Concluding Remarks

The investigations conducted in this chapter into the effects of non-isothermal conditions in slumping glass on the final surface curvature of a thermally-replicated component, have shown that even quite large and probably excessive transverse temperature variations alter the curvature profile by only quite small amounts. By contrast, relatively smaller radial temperature gradients in the glass have a significantly greater affect on surface curvature of the final product. Thus, while the

simple heat-flow computations conducted show that reasonably large temperature differences ($T_d \sim 5°\text{C}$) may exist through the thickness of the glass, these can probably be ignored, and the temperature through the thickness of the glass at some radial position r be assumed to be a constant average value. Then we have the glass temperature being independent of vertical position z, and a function of r and time t only, i.e. $T = T(r, t)$. Further, if radial temperature gradients are only 1 or 2°C over 4.5 cm, as also suggested by the heat-flow computations, then these too will have little effect on surface curvature, leading to the conclusion that, unless significant radial oven air temperature variations exist throughout the thermal replication process, an isothermal simulation is likely to be reasonably accurate. However, as intimated at the beginning of this chapter, a proper judgement can only be made in the context of specific industry needs, which may modify this assessment.

Although we have only considered linear temperature gradients in the glass, the finite-element program developed for non-isothermal conditions is not limited to these, and simulations could be run for any other temperature distributions. However, in the absence of clear evidence to the contrary, approximate linear gradients seem most likely, and are supported by the rough heat-flow modelling done. Further, it has been established that isothermal modelling may, in fact, be quite adequate for thermal replication. Thus, it is now appropriate to proceed to a comparison of numerical and experimental results, which might supply better evidence as to whether radial temperature variations, of sufficient magnitude to affect the final surface curvature, do, in practice, arise in the glass.

Chapter 8

Thermal Replication: A Comparison of Numerical and Experimental Results

8.1 Introduction

The mathematical model and finite-element simulation of thermal replication that has been developed and used for the work reported in this thesis, has been necessarily based on numerous simplifying assumptions. These render the problem more tractable, and indeed solvable, using current knowledge, and allow a solution to be obtained sufficiently quickly to be of value in a real industrial manufacturing context. Much useful insight into and understanding of the process has been gained from a simple constant-viscosity creeping-flow model that only involves what appear to be, its most important aspects. A simple extension to this model, enabling slumping simulations with a spatially varying viscosity field, has permitted an assessment of the possible effects of a non-uniform temperature distribution in the glass. It must now be decided whether the additional complexity of this extended model is required, or indeed whether the even more difficult task of developing a

full thermo-mechanical model that couples heat transfer and fluid flow should be addressed, for the simulation to be of real benefit to industry in respect to both the predictions of product shape yielded, and the time taken to obtain these.

Unfortunately, determining of actual temperature variations in slumping molten glass, by either experimental or computational methods, is extremely difficult. Hence, in order to find whether these do arise in the thermal-replication process and significantly affect the final product, requiring their inclusion in the numerical model, we now compare the results of numerical simulations, using both the simple and extended models, with experimental data. Clearly industry input is highly desirable in making this comparison, to indicate whether differences between numerical and experimental results are at acceptable levels, and also to determine acceptable trade-offs between complexity and accuracy on the one hand, and computational time on the other. In the absence of this input, the aim is to give some comparative results and an assessment of them, that can be used by industry to assess the present models, and determine whether something more is needed.

This type of model-development procedure, involving numerical simulation followed by validation against experimental data to ensure that the model gives results that are consistent with reality, which may then lead to modifications to the model and further experimental validation, is typical for, and a necessity in, the modelling of many industrial processes [27, 102].

All experimental work reported here was conducted at SOLA International Holdings Ltd Research Centre, Lonsdale, South Australia. The work presented in this chapter has also been reported in a paper by this author [90].

8.2 The Experiments

The particular test geometry used for current validation purposes consisted of a glass disc with small initial spherical curvature supported on a circular concave

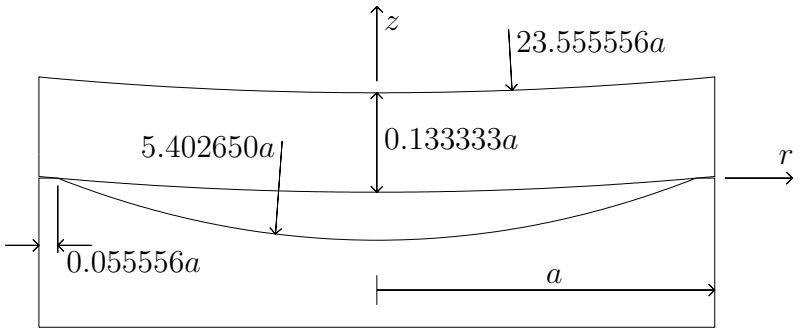

Figure 8.1: Experimental test geometry.

mould with larger spherical curvature, excepting for a small flat ledge around its perimeter. This arrangement is shown in Figure 8.1, with all dimensions scaled by the mould radius a which in practice was 45 mm. The glass had a density of 2620 kg/m^3, and strain, annealing, and softening points of 505 ± 3°C, 545 ± 3°C, and 735 ± 3°C respectively.

The glass disc and mould combination was placed in an oven and heated from room temperature at 4°C/minute to a temperature of 700°C. The temperature was then held at this value for a period of time known as the "soak", before the oven was turned off and the system allowed to cool naturally.

In order to find the point in time at which full contact between the lower glass surface and the mould is achieved, a number of experiments with different soak times were conducted. Full contact was deemed to have occurred if the slumped glass component showed marking from the mould over the whole of its lower surface. This is not easy to see, and it is possible that full contact occurs at an earlier time than identified.

Figure 8.2 shows the curvature (κ) profile (multiplied by the mould radius a) on the upper glass surface for 30, 40 and 50 minute soaks. After 30 minutes of soak the glass did not appear to have fully contacted the mould. After 40 minutes of soak it was still not certain that full contact between glass and mould had been established, but after 50 minutes of soak this appeared to be fairly certain. Curvature is calculated in the manner described in Chapter 6, using quintic B-splines with 26

degrees of freedom, from vertical sag (δ) measurements made across a diameter of the glass component using a profilometer. Two data sets of 81 (r, δ) coordinates at 1 mm intervals across two perpendicular diameters (denoted "0°" and "90°") intersecting at $r = 0$ were obtained in the region $-40 \leq r \leq 40$, to give two curvature profiles for each of the two experiments, as shown.

The curves shown in Figure 8.2 bear a strong resemblance to those calculated in previous chapters for numerical slumping simulations. However, since the slumping geometry is axisymmetric (or at least nearly so), we expect all of these curves to be symmetrical about the vertical axis and the two curves for a particular soak time to be identical. Neither of these expectations is realised in Figure 8.2. Rather, the central peak in the curves is offset to the right for the 0° profiles and to the left for the 90° profiles, and all curves (with the possible exception of the 90° curve for the 30 minute soak) have oscillations on the left-hand side that are much larger than those on the right. The magnitude of these differences cannot be explained by error introduced through the method used to compute curvature. A number of possible causes are, however, suggested by the nature of the differences. Also, it must be borne in mind that quite significant changes in curvature can arise from even small variations in the surface profile of the glass.

The large oscillations in the left portion of the curves compared with the right portion, and the fact that this behaviour is consistent in all curves, strongly indicates imperfections in the experimental components and/or setup resulting in irregular slumping. Since the mould is the only component common to both experiments, small imperfections in its shape, possibly resulting in non-continuous support of the glass disc in the initial stages of slumping, are strongly suspected. The way in which the central peaks of the curves are offset either side of the vertical axis for the 0° and 90° diameters, initially suggests that the glass components were not quite centrally positioned on the profilometer, so that the intersection of these 'diameters' identified with the disc centre is in fact offset from the true centre, to the positive

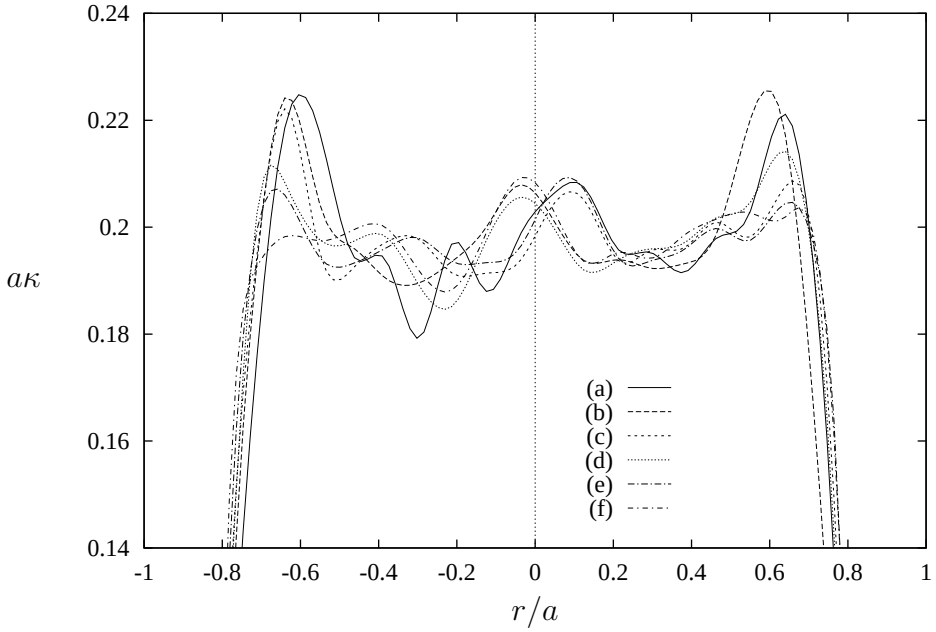

Figure 8.2: Top-surface curvature for (a) 0° diameter, 30 min. soak, (b) 90° diameter, 30 min. soak, (c) 0° diameter, 40 min. soak, (d) 90° diameter, 40 min. soak, (e) 0° diameter, 50 min. soak, and (f) 90° diameter, 50 min. soak.

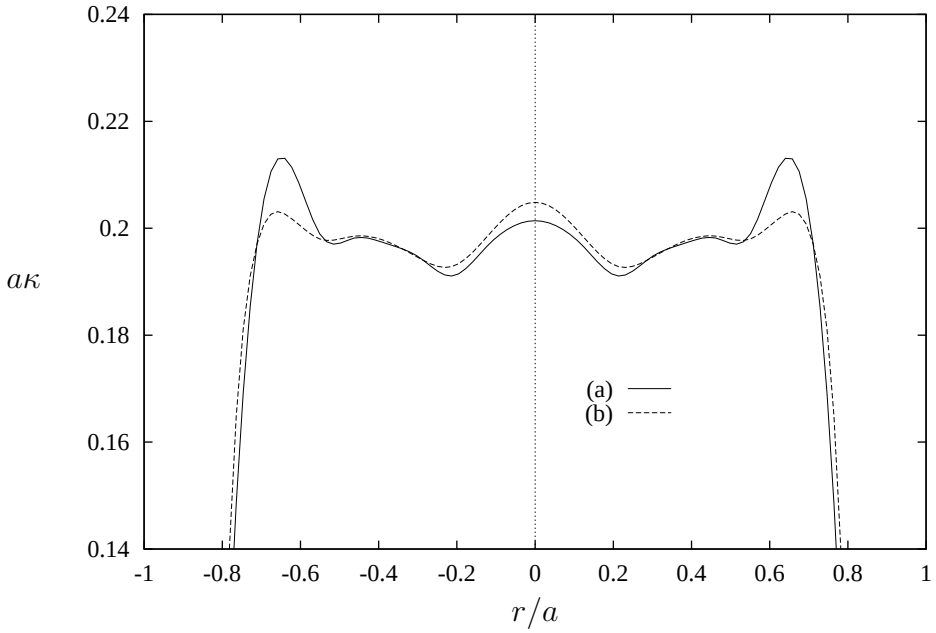

Figure 8.3: Average top-surface curvature for (a) 40 min. soak, and (b) 50 min. soak.

Table 8.1: Raw profilometer data.

Soak (min.)	r (mm)	δ (mm) 0° diameter	90° diameter
40	-3.000	54.449925	54.449205
	-2.000	54.460255	54.459945
	-1.000	54.466335	54.465965
	0.000	54.468050	54.467390
	1.000	54.465375	54.464285
	2.000	54.458140	54.456825
	3.000	54.446655	54.444765
50	-3.000	54.395180	54.395825
	-2.000	54.406280	54.407110
	-1.000	54.413225	54.413970
	0.000	54.415680	54.415885
	1.000	54.413590	54.413325
	2.000	54.406980	54.406070
	3.000	54.395785	54.394285

side of the 0° diameter and to the negative side of the 90° diameter. However, fitting a parabola to the central three points of the raw profilometer sag data, some of which is shown in Table 8.1, indicates an offset in both horizontal directions of less than 0.2 mm, and it is certainly much less than the $3 - 5\,\text{mm}$ seen from Figure 8.2 to be necessary to account for this problem. Thus, it too is probably mostly due to imperfections in the experimental components and/or setup. Another source of error is small inaccuracies in the profilometer measurements themselves, which is clearly evident in the data for each soak time at $r = 0$. The two sag measurements for the 0° and 90° diameters should be identical since they are supposedly taken at the same point on the glass component. In fact they differ a little, which must be due to profilometer error in exactly locating the same point $r = 0$ for each of the 0° and 90° measurements, and/or error in measuring sag. Profilometer measurement error should, however, have a uniform effect over all data sets and cannot explain the magnitude of asymmetry that we are seeing. Note that direct comparison of raw sag

data for one soak time with that of another is not meaningful since the profilometer was reset for each set (consisting of both 0° and 90° diameters) of measurements.

Rather than using the curves shown in Figure 8.2 for comparison with numerical slumping results, average symmetric curves are obtained for each of the 40 and 50 minutes soak times, for which there is most certainty of full mould contact, in the following way. First the sign of the r-coordinate of each data point (r, δ) is ignored and only its distance $|r|$ from the disc centre $r = 0$ considered. The 0° and 90° data sets for a particular soak time are then combined, so that there are four sag measurements for each value of $|r|$, excepting for $|r| = 0$ which has two. The curvature profile is calculated from this new set of $(|r|, \delta)$ points. This yields indicative (axisymmetric) profiles for 40 and 50 minute soak times as shown in Figure 8.3, which we shall use in comparisons with numerical simulations. Note however, that this averaging causes a reduction in the height of the central peak, as well as a significant change to its shape, as seen in Figures 8.4 and 8.5. Some of the peaks towards the edge of the discs are also flattened.

Numerical simulations of slumping with a slightly different geometry than the present test case, showed the top-surface curvature profile in the central portion of the disc remaining largely unchanged once contact between the lower surface and the mould had been established (see Figures 6.16 and 6.17). By contrast, Figures 8.2 and 8.3 show a considerable variation in the curvature in the central region of the disc on increasing the soak time from 40 to 50 minutes. However, considering the magnitude of difference between 0° and 90° curvature profiles for a single soak time, the central portions of the 40 and 50-minute-soak curves are really quite comparable. Furthermore, there is much similarity between all three curves in the set of 0° curves, and those in the set of 90° curves, so that our observation from numerical simulations that mould contact strongly impedes further flow of the glass to the point where flow can be considered to have ceased, is supported by these experimental results.

From the foregoing discussion, it can be appreciated that the experimental data

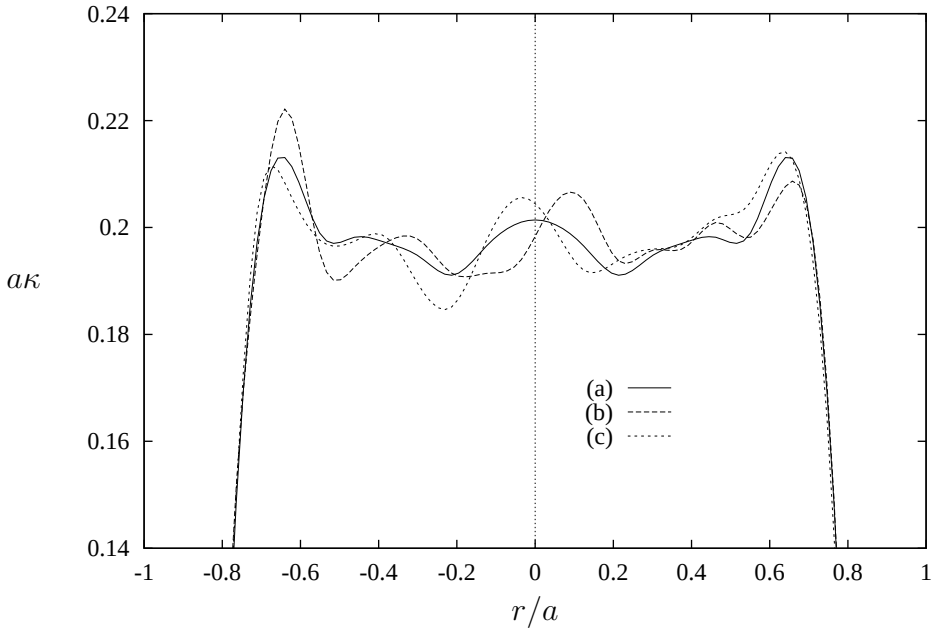

Figure 8.4: Top-surface curvature profiles for a 40 min. soak. (a) Average, (b) 0° diameter, and (c) 90° diameter.

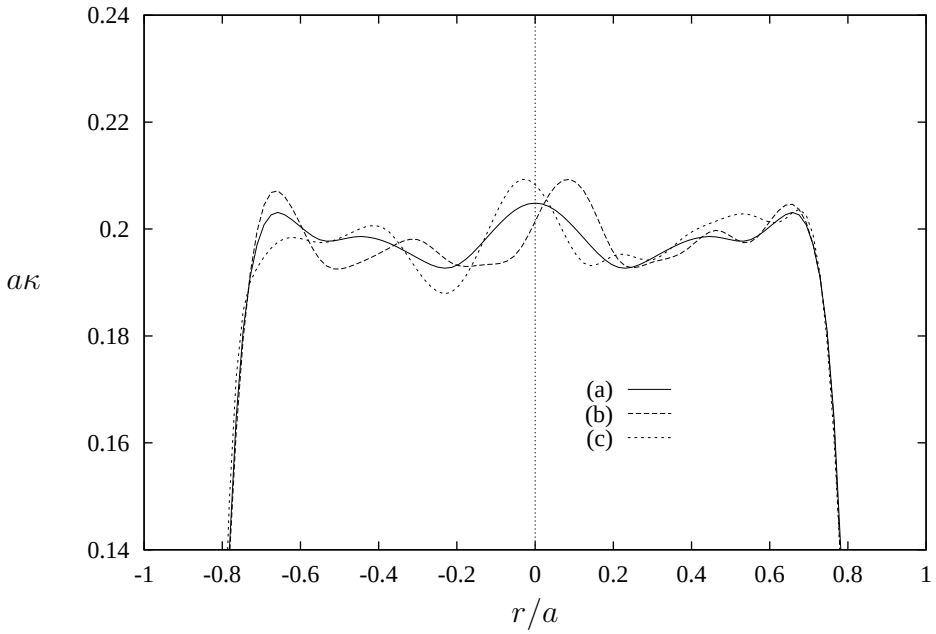

Figure 8.5: As above but for a 50 min. soak.

currently available for comparison with numerical simulations, are not as consistent as we would like, despite the considerable efforts and care taken during its collection. Further work in refining experimental methods and techniques is certainly warranted. In the meantime, however, let us compare what we have with some finite-element simulations.

8.3 When Does Slumping Begin?

As in previous numerical work, computations could begin when the glass reaches its annealing point, since flow can be considered to begin at some point after this. However, as part of this work, some tests were conducted for the glass used in the experimental work described above, to find the temperature at which the glass will begin to flow. From the disc sagging work of Gulati et al. [42], we know that a glass having an annealing point of 484°C and a softening point of 645°C slumped to a measurable extent at temperatures higher than 550°C, i.e. the minimum temperature for glass flow was 66°C higher than the annealing point and 95°C below the softening point. This leads to the expectation that slumping of the glass used in these new experiments, with annealing and softening points of 545 and 735°C respectively, will begin at a temperature of 600°C or greater.

To find the temperature T_s below which the glass is considered to be solid, a number of tests were conducted as follows.

1. Two glass discs of the geometry described in Section 8.2 were supported on moulds with radius of curvature 88.33 mm = $1.96a$.

2. The disc-mould combinations were placed in a rectangular slumping oven, one at a location in the centre of the oven and the other in a corner of the oven.

3. The oven was then heated to a particular temperature T as quickly as possible, and held for one hour at this temperature before being turned off and allowed

to cool naturally.

4. Finally, the curvature of the slumped glass components was measured to determine whether there had been any sag.

Tests were done for temperatures T of 450, 500, 550, 600 and 650°C. Note that the curvature of the mould was chosen to be large relative to that of the glass disc, so that no contact between the lower disc surface and the mould occurred.

A dioptre-scope was used to measure curvature of the upper surface of the glass disc. This instrument gives lens power P in dioptres (D) which, for a thin glass 'lens' in air, as the disc can be considered, is related to the radius of curvature R in metres by $P = (n-1)/R$ where n is the refractive index of the glass [57, pp. 85–87], [51, p. 51]. For the glass used $n = 1.53$ so that $P = 0.53/R$, and curvature in m^{-1} is given by $\kappa = 1/R = P/0.53$. Since the accuracy of the dioptre-scope was ± 0.1 D, and the initial power of the glass discs was 0.5 D, only power measurements greater than 0.51 D were considered to indicate actual sag. Allowing for this measurement error, it was found that slumping occurred for $T \geq 600^{o}C$, but not at the lower temperatures. Hence we have $T_s = 600$°C, at which temperature the glass viscosity is around 10^{11} poise. No effort was made to find T_s to a better accuracy, since a long soak time of one hour was required to detect slumping at this temperature, whereas in thermal replication the oven is quite quickly heated above this to a temperature near to the softening point of the glass (~ 700°C).

8.4 Isothermal Numerical Simulation

The first numerical simulation we consider is for isothermal conditions, that is, the temperature throughout the glass is uniform at any point in time, although that temperature varies with time. The slumping of the glass could be computed at a temperature that is constant in time, and temporal temperature and viscosity changes included by post-processing using a time-varying time scale, as described

in Chapter 5. However, because we also wish to consider non-isothermal slumping, the time scale is kept constant at $\mathcal{T} = \mu_0/(\rho g R)$, with computation proceeding in the same way as for the isothermal-slumping check of Chapter 7. The glass density is $2620\,\mathrm{kg/m^3}$, and a characteristic viscosity of $\mu_0 = 10^7\,\mathrm{Pa}\cdot\mathrm{s}$ is used, so that $\mathcal{T} = 144.2477\,\mathrm{minutes}$.

Whereas in all previous numerical slumping simulations considered in this thesis the glass disc has been initially flat, in the present case, shown in Figure 8.1, the glass disc has some small initial spherical curvature. This makes no difference whatsoever to the computations, but some extra care is needed in mesh generation. For this *Fastflo* is still used, which with appropriate input will generate meshes over domains having boundaries that are segments of circles. For the current problem a mesh of 889 6-node triangles is generated over the glass component. However, *Fastflo* only generates straight-sided triangular elements (even for 6-node triangles) so that curved boundaries are approximated by straight-line segments between element corner nodes that lie on the boundaries. Therefore, in all probability, mid-side nodes between these corner nodes do not lie on the boundaries. Further error arises in the representation of curved boundaries when *Fastflo* writes the mesh specification to file, because all node coordinates are output to only four decimal places. Although the total error in the coordinates of nodes lying on curved boundaries is small, calculation of surface curvature becomes increasingly sensitive to this as the number of degrees of freedom of the B-spline approximation is increased. For a better estimation of curvature it is preferable to use more degrees of freedom rather than less, and consequently, before running a numerical simulation involving curved boundaries, as in the present case, the mesh file generated by *Fastflo* is processed to move all element nodes — both corner and mid-side nodes — that lie on the curved boundaries back to those boundaries, with all node coordinates output in double precision (i.e. to sixteen significant figures). Apart from the increased accuracy of double precision, the effect of this is that a curved boundary is now approximated

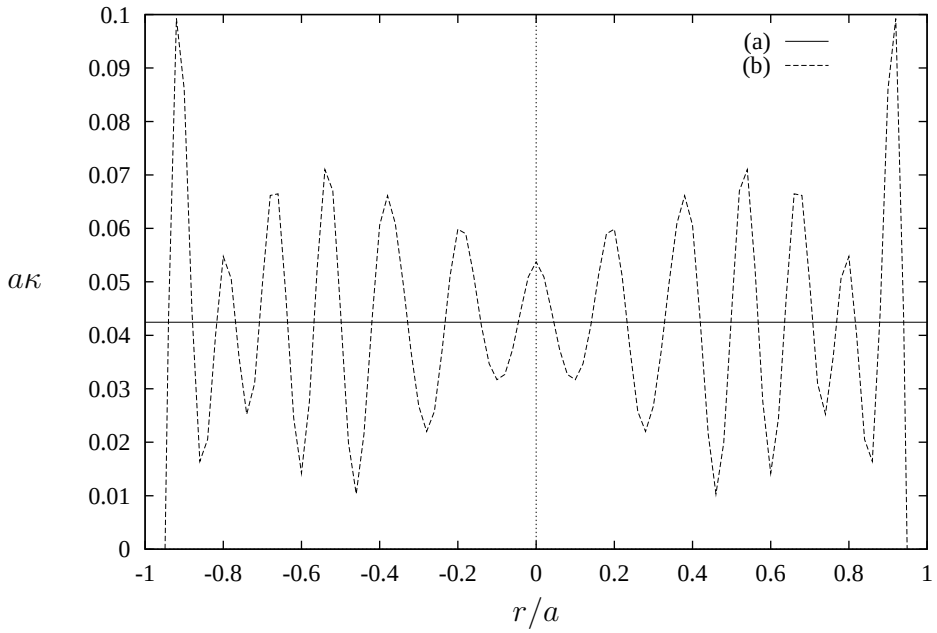

Figure 8.6: Comparison of error in curvature calculated using (a) a modified mesh giving the correct constant curvature (horizontal line) of $a\kappa = 0.04245$, and (b) an unmodified mesh causing considerable instability.

by parabolic segments defined by one mid-side and two corner nodes, all of which lie on the boundary. The greatly increased accuracy of this approximation, over the straight-line approximation, is seen in much smoother surface-curvature profiles obtained from quintic B-spline approximations having more than 10 degrees of freedom. A comparison of the initial upper-surface curvature of the glass disc prior to slumping, computed from the *Fastflo* mesh for the present problem before and after modification, using a quintic B-spline approximation with 35 degrees of freedom, is given in Figure 8.6. The curvature should be constant at $a\kappa = 0.04245$ as is given by computations using the modified mesh. By contrast, the unmodified mesh causes significant instability in the curvature profile due to numerical error. With a modified mesh there is no such instability, even for a B-spline approximation having in excess of 100 degrees of freedom relative to 165 data points on the upper disc surface.

With reference back to Sections 7.2 and 7.3 of Chapter 7, let $\Delta t_T^* = 0.01$, and

$\Delta t^* = 0.005$. Then, for a heating rate of 4°C/minute, the temperature increment after a constant-temperature time interval Δt_T^* is $\Delta T = 4\,\Delta t_T^*\,\mathcal{T} \approx 5.770$°C, and, having decided, based on the tests described in the previous section, to begin computations at $T_s = 600$°C, the temperature at node n during constant-temperature time interval i is given by

$$T_i^n = \min(600 + 4\,\Delta t_T^*\,\mathcal{T}(i - 1/2), 700). \tag{8.1}$$

To determine viscosity from temperature using the VFT equation (5.1), constants appropriate to the glass used in the experiments are required. Substituting the strain, annealing and softening point data for this glass into (5.1) and solving for the three constants gives one possibility, but the relationship between viscosity and temperature is probably more accurately given by data supplied by Corning [23, 24] for another very similar glass. This glass compares with that used in the experiments described earlier in this chapter in having an annealing point ($\mu = 10^{13}$ poise) of 544°C instead of 545°C, and a softening point ($\mu = 10^{7.6}$ poise) of around 730°C instead of 735°C. For a viscosity expressed in poise, VFT-equation constants of $A = -1.6837$, $B = 1908.5$ and $T_0 = 407.689$ have been obtained using viscosity measurements at 10°C intervals over the temperature range $650 \leq T \leq 800$°C. The viscosity in SI units of Pa · s is obtained using $A = -0.6837$.

Slumping is now computed to a time of $0.3\mathcal{T}$ in 138 steps of size 0.005 or less. From this simulation, the point of full contact between the mould and lower disc surface is identified for comparison with the experimental results presented above. Thus it is effectively assumed that there is no glass flow once the oven has been turned off. This is not too unreasonable given that flow is very slow once full mould contact has been established and that the experimental results with which we are comparing are for full mould contact.

The numerical simulation gives full mould contact occurring at about $0.25\mathcal{T}$ to $0.26\mathcal{T}$, which in real terms is a total slumping time of 36 to 38 minutes. Since this includes the time for the temperature to rise from 600 to 700°C at a rate of

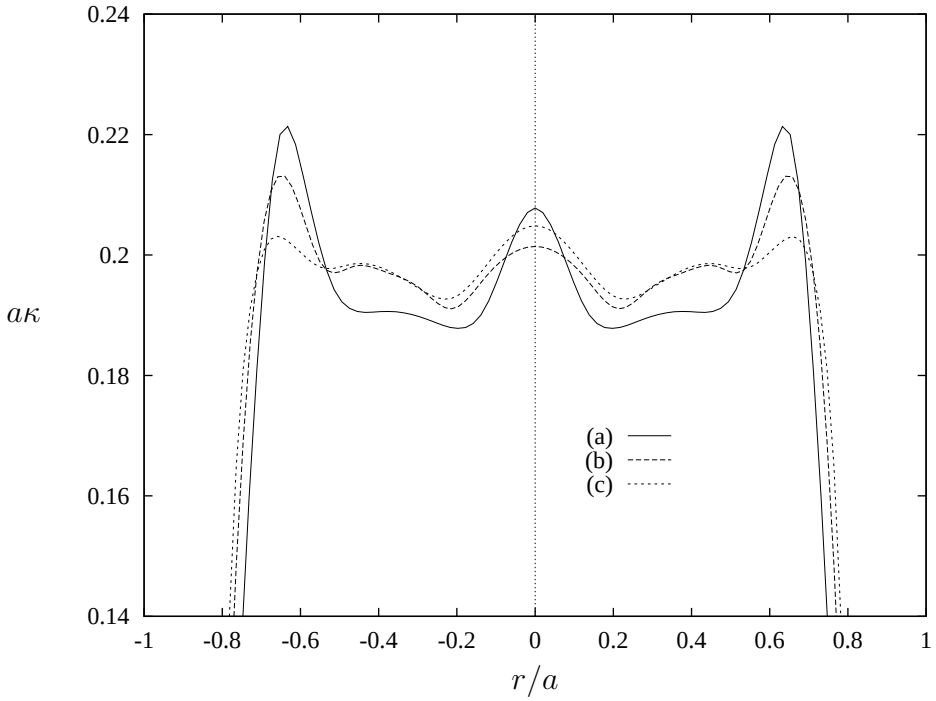

Figure 8.7: A comparison of upper-surface curvature given by (a) an isothermal numerical simulation, (b) experiment with a 40 min. soak, and (c) experiment with a 50 min. soak.

4°C/min., the soak time is only 11 to 13 minutes. This is considerably different from the 40 to 50 minute soak determined experimentally, and we will return to this later in this chapter.

Figure 8.7 shows the upper-surface curvature profile obtained from the numerical data, compared with those obtained from the experimental data. A quintic B-spline with 35 degrees of freedom was used to compute curvature from the numerical data. The general behaviour of the numerical solution is in good agreement with the experiments, especially considering the nature of the raw experimental data and the processing it has undergone. In the region $|r/a| < 0.6$ where, at this stage in the slump, the curvature profile is quite stable and changes only very slowly, the maximum difference between numerical and experimental curvatures is about $|a\kappa| = 0.008$. Considering that we see this magnitude of difference between experimental data sets that should be identical (see Figure 8.2), this is really quite good. It is more

difficult to compare experimental and numerical results in the region $0.6 < |r/a| < 0.8$. Recalling the difficulties experienced in experimentally identifying the point of full mould contact, it is possible that this occurs after a soak of between 30 and 40 minutes. If so, then our numerical result corresponds to an earlier stage in slumping than the 40 and 50 minute soak results with which we are comparing, and, since there is some flattening of these edge peaks in the last stage of slumping after full mould contact (see Figure 6.17), we can expect that the numerically obtained prediction of curvature in this region will vary from that actually obtained experimentally in the manner seen in Figure 8.7. It would even be quite easy to "prove" this as the reason for the observed variation, by computing a curvature profile using numerical data at a later stage in the slump with edge peaks that correspond with either the 40 or 50 minute experimental curves. However, there are other possible reasons for the discrepancy in the numerical and experimental results in this region. Averaging of data sets, and the consequent reduction in some peak heights, is a likely cause of variation, as indicated by the fact that a large peak approaching the height of that seen in the numerical result is obtained from one of the data sets for the 40 minute soak (see Figure 8.4), which is lost when the 0° and 90° curves are averaged. Differences in this region could also be due to using too coarse a grid when measuring sag with the profilometer, leading to an inability to properly capture what is happening in this region of rapid curvature change. Compared with the profilometer measurements, the numerical simulation gives about double the number points on the upper glass surface in this region, from which we calculate curvature.

8.5 Non-isothermal Numerical Simulation

Next, thermal replication under non-isothermal conditions is considered, and curvature profiles computed from simulations after full mould contact for comparison

with the curves obtained from experimental data. Having learned in Chapter 7 that transverse temperature gradients in the glass are of little consequence, we consider only radial temperature gradients. Again following the work of Chapter 7, we look at just two arbitrary temperature distributions.

Firstly, we consider that the temperature in the glass at time $t = 0$ is a linear function of radius only, decreasing from 600°C in the centre to 590°C at the edge (a drop of 10°C), i.e. the initial temperature distribution (in °C) is given by

$$T^n(0) = 600 - 10(1 - r/a), \tag{8.2}$$

where r is the radial coordinate of node n at time $t = 0$. We then suppose that for $t > 0$ the temperature at any node is given by $T^n(0) + 4°\text{C/min.} \times t$, where t is given in minutes. When the temperature at any node reaches 700°C it is not further increased. Thus, the temperature at node n during constant-temperature time interval i is given by

$$T_i^n = \min(T^n(0) + 4\,\Delta t_T^*\,\mathcal{T}(i - 1/2), 700). \tag{8.3}$$

At any time t, the viscosity distribution is obtained from the temperature distribution via the VFT equation.

The second example is similar to the first, excepting that at $t = 0$ the temperature increases from 590°C at the centre to 600°C at the edge, i.e. the initial temperature distribution (in °C) is given by

$$T^n(0) = 590 + 10\,r/a. \tag{8.4}$$

The rough heat-flow computations of Chapter 7 indicate that a radial temperature gradient of 10°C in 45 mm, as in both of these non-isothermal cases, is probably extreme. Nevertheless, they have been chosen in order to be able to clearly see the nature of the changes to the curvature profile resulting from these types of non-isothermal conditions.

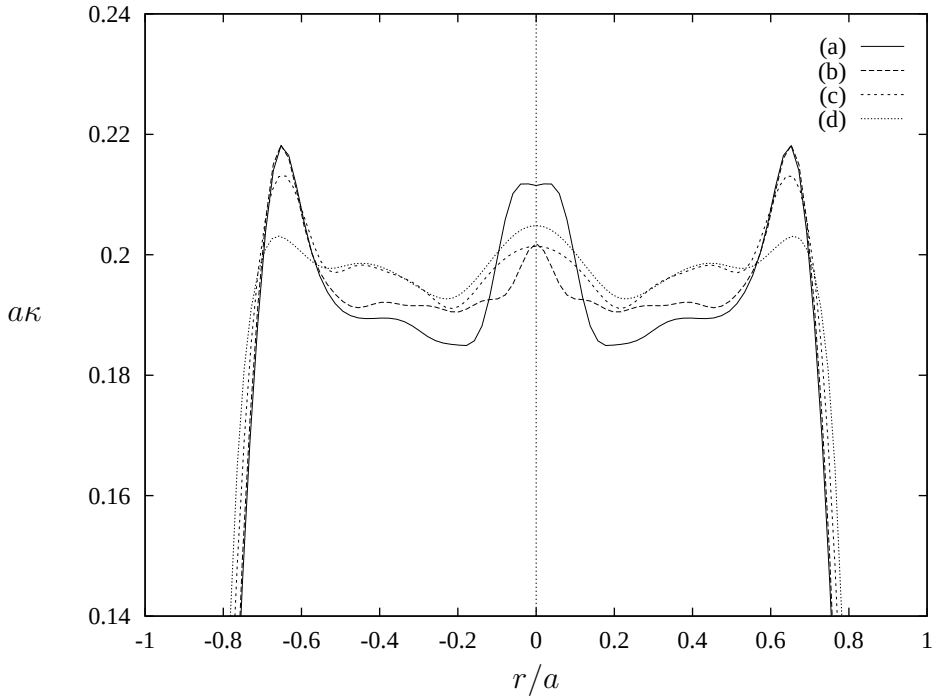

Figure 8.8: A comparison of upper-surface curvature given by (a) a numerical simulation with temperature decreasing linearly from the centre to the edge, (b) a numerical simulation with temperature increasing linearly from the centre to the edge (c) experiment with a 40 min. soak, and (d) experiment with a 50 min. soak.

Figure 8.8 compares the upper glass-surface curvature profiles after full mould contact for both of these non-isothermal cases, with the experimental results obtained with 40 and 50 minute soak times. We restrict our attention to the central region of stable curvature, $|r/a| < 0.6$. The outcome with temperature decreasing from the disc centre to the edge, which the heat-flow computations of Chapter 7 suggests as more likely than the reverse, has a general shape that does not compare with the experiments nearly as well as did the isothermal result. The differences between experiment and the second case with temperature increasing from the disc centre to the edge, are about the same magnitude as with the isothermal simulation, but, again, the overall shape of the curve in the central region is less in keeping with the experimental results. Thus it appears that the isothermal simulation gives the most satisfying comparison.

The non-isothermal simulations show full mould contact occurring in only slightly longer times than the isothermal simulation ($0.28\mathcal{T}$ to $0.29\mathcal{T}$). Thus the time taken to achieve full mould contact, as determined by numerical simulation, is not much affected by a 10°C radial temperature difference, and is significantly less than the time found by experiment to be required.

8.6 Discussion and Conclusion

We have compared isothermal and non-isothermal slumping simulations with experimental slumping data. The best agreement is obtained with an isothermal model, so that, at least with the present quality of experimental data available, it seems reasonable to suppose that temperature variations of sufficient magnitude to substantially change the outcome do not arise in the thermal-replication process. Rather the experimental results tend to support the assessment given in Chapter 7 that temperature gradients are sufficiently small to ignore, and that an isothermal simulation is reasonably accurate. Thus it is not necessary, at this stage, to be concerned with coupling heat transfer to the fluid flow model. However, there is a clear need to refine experimental methods and techniques, which might lead to an improved agreement between numerical prediction and experimental outcome. Alternatively we might then be able to more easily assess whether non-isothermal conditions can account for small variations in curvature profiles. Once we have more accurate experimental, and possibly numerical, results, some consideration can be given to meeting industry tolerances on surface curvature. However, at the present stage, the experimental results are too uncertain to be able to meaningfully address this subject.

The most major difference between all the numerical simulations and the experimental work, that has not yet been explained, is in the time taken for the lower surface of the glass disc to achieve full contact with the mould. The experiments

showed a soak time of 40 to 50 minutes to be necessary before full mould contact was achieved, while the numerical simulations give soak times of around 10 minutes, and even quite large temperature gradients in the glass do not appreciably increase the slumping time to the extent necessary for agreement with experiment. The most likely major cause of this time difference is that the viscosities computed from the known temperature using a VFT equation, and which vary over several factors of 10, are in error. If we increase them by a factor of two, then we double the time scale so that the numerical soak time is of the order of 45 minutes, which is in accord with the experiments. On a logarithm scale, such as is typical for quoting glass viscosities, doubling the viscosity is an increase of only $\log_{10} 2 \approx 0.3$, which is of small magnitude relative to the viscosity range of a molten glass and represents an error that is quite a possibility. Already we know that the VFT equation used is for a different, though similar, glass to that used in the experiments. Then, even if the VFT equation for the specific glass was used, it is still an approximation with an associated error. Finally, the methods of determining glass viscosity in different temperature ranges are known to be subject to error, with different methods not generally giving exactly the same result in regions of overlap, as is plain from the discussion by Hagy [44] of the various methods of measurement.

In summary, allowing for all the known and possible causes of error in both the experimental and numerical results, the agreement between an isothermal finite-element simulation based on a creeping-flow model and some actual experimental data is very good.

Chapter 9

In Conclusion

A creeping-flow finite-element program has been developed that is capable of solving a variety of very viscous fluid-flow problems that are driven by gravity. In this thesis it has been used to solve for flows which are not impeded by a mould or other object, such as the dripping of honey from a spoon and the slumping of glass in disc viscometry. In addition, the slumping of glass in thermal replication of optical quality surfaces, where the flow is modified by progressive contact with a mould, has been given special attention. At this stage the program is restricted to geometries that are two-dimensional or axisymmetric.

The finite-element program is able to compute for fluids of constant viscosity, or fluids having a viscosity that varies with temperature in time and/or space in a way that is known and can hence be prescribed. In the isothermal case where the fluid viscosity is a function of time only, and hence is constant throughout the fluid at any instant in time, the flow differs from that for a constant-viscosity fluid only via a time scale that varies with the viscosity. Hence a constant-viscosity model can always be solved first, and the results then re-processed if necessary to include effects of time-varying viscosity. In the non-isothermal case where the viscosity may also vary throughout the fluid at any instant in time, computations are a little more complicated since the viscosity distribution in the fluid at every time step must be

known and specified at the time of finite-element solution.

Because of specific application of this research to manufacture of optical surfaces for which curvature is the critical measure, a method for computing the curvature profiles of surfaces from coordinate data was required. For the types of surfaces that arise in this thesis, which can be considered as functions of a single variable, the method adopted was to fit B-splines by a least-squares method to surface coordinate data, from which first and second derivatives can be obtained in order to compute curvature. Least-squares fitting of B-splines is recommended by de Boor [26, p. 249] as an effective method of filtering noise from non-exact data such as can be expected from both numerical simulations and experiments. High degree ($k \geq 4$) B-splines were used to provide sufficient continuity so that second derivatives were continuous and differentiable, and hence a smooth curvature profile could be computed.

Some simulations conducted using the creeping-flow finite-element program have been compared with experimental results. Disc-sagging experiments in the absence of a mould done at Corning Glass Works in the mid 1970's provided some experimental data. Other new experiments, involving slumping simple glass discs into moulds of spherical shape, were conducted in conjunction with the work reported in this thesis, providing data of direct relevance to the thermal replication process. The numerical and experimental results were in good agreement in all cases, and likely explanations exist for the differences seen. The viscosity-temperature relation obtained using creeping-flow finite-element computations and Corning disc-viscometry experimental data differed from the relation obtained by other accepted methods, but this can be attributed to the method used to measure sag, which modified the flow of the glass. Comparisons of finite-element simulations with other Corning disc-sagging experiments in which the flow was not affected by sag-measuring techniques, as well as comparisons with the new thermal-replication experiments, suggest inaccuracy in the viscosity data and/or the approximated viscosity-temperature relationship used in the simulations as a likely cause of differences between numerical

and experimental results.

Multiplying viscosities by a 'fudge factor' to improve the agreement of numerical simulations with experimental results is not as dubious as it might first appear. The corrections are small relative to the magnitude of the viscosity, and small error can easily arise from one or more of (i) inaccuracy in available viscosity data, (ii) inaccuracy due to a VFT approximation to that data, and (iii) glass temperatures differing a little from nominal oven temperatures. Other factors that have not been examined in this thesis might also be responsible for differences between numerical and experimental results. For example, surface skins can form on glasses, which have higher viscosity than the bulk of the glass and act to increase the effective viscosity [54].

Thermal replication has been the major focus of this thesis, and for this reason it is appropriate to give some specific conclusions in relation to this process. Preliminary comparisons of curvature profiles at full mould contact from numerical simulations and experiments indicated that an isothermal model most closely matches what is seen experimentally. The differences, although quite considerable in an optical context, were no larger than differences between what should have been identical curves computed from experimental data. At least there is a need for more consistent experimental data against which to assess numerical simulations before very much effort is put into full thermo-mechanical modelling of the thermal replication process. Considering the sensitivity of surface curvature to small changes in surface shape, obtaining such consistent experimental data is quite a challenge for the future. In the meantime, the creeping-flow finite-element simulation program may still prove a useful tool to industry, not least for the good qualitative understanding of thermal replication that may be, and indeed has already been, gained from it. It has shown the progress of glass slumping, something that was not physically observable with the equipment available for experiments and which may even be difficult to see with sophisticated and expensive equipment. Specifically it

has shown, for the thermal-replication cases considered, that mould contact occurs first at the centre of a disc. It has shown that the curvature of the upper surface of the glass disc quickly stabilizes once the opposite lower surface is in contact with a mould. From this it is obvious that continued soaking after full mould contact has been achieved is virtually wasted time, something that was unknown prior to this investigation. Furthermore, if the outer annulus of the slumped disc is subsequently discarded, then soaking can cease just as soon as the desired central portion of the disc has attained contact with the mould and its curvature stabilized.

If thermal replication is an example of a process for which, as Pearson [68, p. 244] found in relation to polymer processing,

> ... the geometry of flow boundaries is often far more important than all other factors in determining flow fields and hence process behaviour ...

so that small temperature variations in the glass are not important and an isothermal model is adequate, then an important implication is that the slumping temperature history has no effect on the final product, provided only that the area under the curve of inverse viscosity is constant. Thus, heating rates and soak times can be adjusted so as to optimize heating costs and slumping times, and/or match throughput to other inter-dependent processes.

Since thermal replication is most commonly used for generating quite general three-dimensional surfaces, the most serious limitation of the current creeping-flow finite-element program is probably the fact that only two-dimensional and axisymmetric flows can be simulated. Full three-dimensional geometries will significantly affect final surface curvature profiles, as indicated by the above quotation from Pearson and the experiments conducted as part of this work. Extending the current two-dimensional and axisymmetric program to handle flows in three dimensions is reasonably simple; the difficulties lie in the generation of three-dimensional volume and surface meshes. These are not automatically generated by *Fastflo*, the CFD code used so far to generate two-dimensional meshes, but other commercial mesh

generation packages are no doubt available that can be used. In addition a three-dimensional code will require considerably more in computer resources, with computational time also increasing significantly. A comment by de Boor [26] indicates that it may be possible to extend the least-squares B-spline method for calculating surface curvature profiles to three dimensions.

While, however, a full three-dimensional creeping-flow model is highly desirable, it is also best to sort out some, if not all, of the current problems with simpler two-dimensional or axisymmetric geometries, in the same spirit in which Pearson [68, p. 245] also stated

> ... only after ruthless simplification can reliable quantitative analyses be undertaken.

To this end there is, at this stage of the work, a need for substantial industry input, to assess the importance of the differences between numerical and experimental results in an industrial context, to decide the extent to which these can or should be addressed, and to assist with the development of better experimental techniques and methods of measurement. Further progress can then be made towards a (possibly three-dimensional) program that can assist designers by reducing or eliminating experimental iteration to determine the mould shape required to give a particular optical surface. And this may also result in better techniques and controls for the manufacturing process itself.

Bibliography

[1] M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functons with Formulas, Graphs and Mathematical Tables*. US Dept. of Commerce, 1965.

[2] R.S. Anderssen and F.R. de Hoog, Finite difference methods for the numerical differentiation of non-exact data. *Computing* 33 (1984) 259–267.

[3] R.S. Anderssen, F.R. de Hoog, and M. Hegland, *A stable finite difference ansatz for higher order differentiation of non-exact data.* Mathematics Research Report MRR96-023, Centre for Mathematics and its Applications, Australian National University. Available at `http://discus.anu.edu.au:80/~marcus/litheg.html`

[4] American Society for Testing Materials, *Annual Book of ASTM Standards*. 15.02 Glass; Ceramic Whitewares. 1994.

[5] C.L. Babcock, *Silicate Glass Technology Methods*. John Wiley and Sons, 1977.

[6] Stephanie Balogh, Money in dripping honey. *The Australian*, Wednesday January 29, 1997, p. 23.

[7] G.K. Batchelor, *An Introduction to Fluid Dynamics*. Cambridge University Press, 1967.

[8] Eric B. Becker, Graham F. Carey and J. Tinsley Oden, *Finite elements: An Introduction*, The Texas Finite Element Series v.1. Prentice-Hall, New Jersey, 1981.

[9] Eric B. Becker, Graham F. Carey and J. Tinsley Oden, *Finite elements: Fluid Mechanics*, The Texas Finite Element Series v.6. Prentice-Hall, New Jersey, 1986.

[10] R.B. Bird, C.F. Curtiss, R.C. Armstrong and O. Hassager, *Dynamics of Polymeric Liquids* v.1 Fluid Mechanics. Wiley, New York, 1977.

[11] A. Blazek, *Review of Thermal Conductivity Data in Glass, Part I Thermal Conductivity at Low and Moderated Temperatures*. International Commission on Glass, 1983.

[12] C. Bliem, D.J. Landini, J.F. Whitbeck, R. Kochan, J.C. Mittl, R. Piscitella, J. Schafer, A. Snyder, D.J. Wiggins, J.M. Zabriskie, B.A. Barna, S.P. Henslee, P.V. Kelsey, J.I. Federer, and E.S. Bomar, *Ceramic Heat Exchanger Concepts and Materials Technology*. Noyes Publications, New Jersey, 1985.

[13] Kenneth G. Budinski, *Engineering Materials: Properties and Selection*, 5th ed. Prentice-Hall, New Jersey, 1996.

[14] D.M. Burley and S.J. Graham, The blowing of thin films into moulds with applications in the glass container industry. In *The Mathematics of Finite Elements and Applications VII*. Academic Press Ltd, 1991, 279–286.

[15] A. Cameron, *The Principles of Lubrication*. Longmans, London, 1966.

[16] D. Canright and S. Morris, Buoyant instability of a viscous film over a passive fluid. *J. Fluid Mech.* 255 (1993) 349–372.

[17] H.S. Carslaw, *Introduction to the Mathematical Theory of the Conduction of Heat in Solids*, 2nd ed. Macmillan, London, 1921.

[18] Graham F. Carey and J. Tinsley Oden, *Finite elements: A Second Course*, The Texas Finite Element Series v.2. Prentice-Hall, New Jersey, 1983.

[19] Concentration, Heat & Momentum Ltd (CHAM), Bakery House, 40 High Street, Wimbledon Village, London, UK, `http://www.cham.co.uk/`

[20] J. Chery, A. Bonneville, J.P. Vilotte, and D. Yuen, Numerical modelling of caldera dynamical behaviour, *Geophys. J. Int.* **105** (1991) 365–379.

[21] A. Cormeau, I. Cormeau and J. Roose, Numerical simulation of glass-blowing. In J.F.T. Pittman, O.C. Zienkiewicz, R.D. Wood and J.M. Alexander (eds), *Numerical Analysis of Forming Processes*. John Wiley and Sons, 1984, 219–237.

[22] CSIRO Mathematical and Information Sciences, Macquarie University Campus, Herring Road, North Ryde, NSW, Australia. Computational Fluid Dynamics, *Fastflo*, `http://www.cmis.csiro.au/fastflo/`

[23] P. Danielson, Corning Incorporated, Corning, New York. Personal communication by email, 24 July, 1995.

[24] P. Danielson, Corning Incorporated, Corning, New York. Personal communication by email, 22 March, 1996.

[25] J. Daoudi, E. Irving and N. Pons, A multivariable predictive control with internal model of a laminated windshield bending furnace. *IEEE Conference on Control Applications*, Proceedings v.3 (1994) 1897–1902.

[26] Carl de Boor, *A Practical Guide to Splines*, Applied Mathematical Sciences Vol. 27. Springer-Verlag, 1978.

[27] F.R. de Hoog and N.I. Robinson, Some Case Studies in Industrial Mathematics. In R.L. May and A.K. Easton (eds), *Computational Techniques and Applications: CTAC95*. World Scientific, Singapore, 1996, 1–12.

[28] Morton M. Denn, Continuous drawing of liquids to form fibres. *Ann. Rev. Fluid Mech.* 12 (1980) 365–387.

[29] J. De Wynne, J.R. Ockendon and P. Wilmott, On a mathematical model for fiber tapering, *SIAM J. Appl. Math.* 49 (1989) 983–990.

[30] P.J. Doyle, *Glass Making Today*. R.A.N. Publishers, Ohio, 1994.

[31] Jens Eggers, Theory of drop formation. *Phys. Fluids* 7 (1995) 941–953.

[32] C.A.J. Fletcher, *Computational Galerkin Methods*. Springer Verlag, New York, 1984.

[33] J.B. Fraleigh, *Calculus with Analytic Geometry*. 2nd ed. Addison-Wesley, 1985.

[34] R.H. Gallagher, J.T. Oden, O.C. Zienkiewicz, T. Kawai and M. Kawahara, *Finite Elements in Fluids*, v.5. John Wiley and Sons, 1984.

[35] S. Gaudet, G.H. McKinley and H.A. Stone, Extensional deformation of Newtonian liquid bridges. *Phys. Fluids* 8 (1996) 2568–2579.

[36] S.P. Gerhardt, What every modular framer should know about glass. *S.A.E. Transactions* 100 (1991) 689–694.

[37] Philip Gibbs, Is glass liquid or solid? *Physics FAQ*, 10 October 1996, `http://www.weburbia.com/physics/glass.html`

[38] S.J. Graham, D.M. Burley and J.C. Carling, Fluid flow in thin films using finite elements. *Math. Engng Ind.* 3 (1992) 229–246.

[39] R.W. Grimshaw, *The Physics and Chemistry of Clays and Allied Ceramic Materials*, 4th ed. Ernest Benn Ltd, London, 1971.

[40] S.T. Gulati, E.H. Fontana and W.A. Plummer, *Disc bending viscometry*. Report No. R-5702, Corning Glass Works, Corning, New York, 10 March 1975.

[41] S.T. Gulati, W.A. Plummer, and E.H. Fontana *Sagging of ATON Circular Blanks with Clamped Edge*. Report No. R-5962, Corning Glass Works, Corning, New York, 21 October 1975.

[42] S.T. Gulati, E.H. Fontana and W.A. Plummer, Disc bending viscometry. *Phys. and Chem. of Glasses* 17 (1976) 114–119.

[43] G.K. Gupta, W.W. Schultz, E.M. Arruda and X. Lu, Nonisothermal model of glass fiber drawing stability. *Rheologica Acta* 35 (1996) 584–596.

[44] H.E. Hagy, Rheological Behaviour of Glass. In L.D. Pye, H.J. Stevens and W.C. LaCourse (eds), *Introduction to Glass Science*. Plenum, 1972, 343–371.

[45] M. Hegland, Computer Sciences Laboratory, Research School of Information Sciences and Engineering, Australian National University, Canberra, ACT. Personal communication.

[46] D.G. Holloway, *The Physical Properties of Glass*. Wykeham Publications, London, 1973.

[47] Gregory A. Houseman and Peter Molnar, Gravitational (Rayleigh-Taylor) instability of a layer with non-linear viscosity and convective thinning of continental lithosphere. *Geophys. J. Int.* 128 (1997) 125–150.

[48] P.D. Howell, Models for thin viscous sheets. *Euro. Jnl of Applied Math.* 7 (1996) 321–343.

[49] A. Kaye, Convected coordinates and elongational flow. *J. Non-Newtonian Fluid Mech.* 40 (1991) 55–77.

[50] W.D. Kingery, *Introduction to Ceramics*. Wiley, 1967.

[51] R. Kingslake, *Lens Design Fundamentals*. Academic Press, New York, 1978.

[52] Erwin Kreysig, *Advanced Engineering Mathematics*, 7th ed. Wiley, New York, 1993.

[53] M. Křížek and P. Neittaanmäki, *Finite Element Approximation of Variational Problems and Applications*. Longman, Essex, 1990.

[54] W. LaCourse, Center for Glass Research, Alfred University, Alfred, New York. Personal communication.

[55] Leon Lapidus and George F. Pinder, *Numerical Solution of Partial Differential Equations in Science and Engineering*. Wiley, 1982.

[56] San-Yih Lin, Yan-Shin Chin and Tsuen-Muh Wu, A modified penalty method for Stokes equations and its application to Navier-Stokes equations. *SIAM J. Sci. Comput.* 16 (1995) 1–19.

[57] D. Malacara and Z. Malacara, *Handbook of Lens Design*. Marcel Dekker Inc., New York, 1994.

[58] Joseph E. Matta and Raymond P. Tytus, Liquid stretching using a falling cylinder. *J. Non-Newtonian Fluid Mech.* 35 (1990) 215–229.

[59] G.W. Morey, *The Properties of Glass*. Reinhold Publishing Corp., 1938.

[60] H. Mulder, The consistency of butter. In G.W. Scott Blair (ed), *Foodstuffs: Their Plasticity, Fluidity and Consistency*. North-Holland Publishing Co., Amsterdam, 1953, 91–123.

[61] O. Narayanaswamy, Detroit Motor Company, Dearborn, Michigan. Personal communication.

[62] National Agency for Finite Element Methods and Standards, *A Finite Element Primer*. Department of Trade and Industry, National Engineering Laboratory, East Kilbride, Glasgow, UK, 1987.

[63] A.H. Norton, The use of convolution splines in the NQS Einstein evolution code. In B.J. Noye, M.D. Teubner and A.W. Gill (eds), *Computational Techniques and Applications: CTAC97*. World Scientific, Singapore, 1998, 473–480.

[64] B.J. Noye, *Finite Difference Methods for Partial Differential Equations*, Lecture Notes. Department of Applied Mathematics, University of Adelaide, 1993.

[65] Robert J. O'Donnell and Jack A. Zakarian, Survey of two-parameter equations for the viscosity-temperature behaviour of lubricating oils. *Ind. Eng. Chem. Process Des. Dev.* 23 (1984) 491–495.

[66] D.T. Papageorgiou, On the breakup of viscous liquid threads. *Phys. Fluids* 7 (1995) 1529–1544.

[67] D.T. Papageorgiou, Analytical description of the breakup of liquid jets. *J. Fluid Mech.* 301 (1995) 109–132.

[68] J.R.A. Pearson, *Mechanics of Polymer Processing*. Elsevier, London, 1985.

[69] Darrel W. Pepper and Juan C. Heinrich, *The Finite Element Method: Basic Concepts and Applications*. Hemisphere Publishing Corporation, Washington, 1992.

[70] J.T. Pittman, O.C. Zienkiewicz, R.D. Wood and J.M. Alexander, *Numerical Analysis of Forming Processes*. Wiley, New York, 1984.

[71] Robert C. Plumb, Antique windowpanes and the flow of supercooled liquids. *J. Chem. Educ.* 66 (1989) 994–996. Also available at `http://www.ualberta.ca/~bderksen/windowpane.html`

[72] Harvey M. Pollicove, Survey of present lens molding techniques. In M.J. Riedl (ed), *Replication and Molding of Optical Components*, Vol. 896, Proceedings of The Society of Photo-Optical Instrumentation Engineers. Washington; Bellingham, 1988, 158–159.

[73] J. Pryce-Jones, The rheology of honey. In G.W. Scott Blair (ed), *Foodstuffs: Their Plasticity, Fluidity and Consistency*. North-Holland Publishing Co., Amsterdam, 1953, 148–176.

[74] Harold Rawson, *Properties and Applications of Glass*. Glass Science and Technology Vol. 3. Elsevier Science Publishers, 1980.

[75] J.N. Reddy and D.K. Gartling, *The Finite Element Method in Heat Transfer and Fluid Dynamics*. CRC Press, Boca Raton, Florida, 1994.

[76] Simon Rekhson, Zhong-Hao Lu and Claud Day, Computer Modeling of Glass Processing. *Coll. Papers from the 52nd Conf. on Glass Problems* Nov 12 & 13 (1991) 65–77.

[77] Horst Scholze and N.J. Kreidl, Technological aspects of viscosity. In D.R. Uhlmann and N.J. Kreidl (eds), *Glass Science and Technology*, Vol. 3 Viscosity and Relaxation. Academic Press, London, 1986, 233–273.

[78] W.W. Schultz and S.H. Davis, One-dimensional liquid fibres. *J. Rheol.* 26 (1982) 331–345.

[79] M.M. Schwartz (ed), *Engineering Applications of Ceramic Materials*, American Society for Metals, Ohio, 1985.

[80] G.W. Scott Blair, Rheology of milk, cream, ice-cream and similar products. In G.W. Scott Blair (ed), *Foodstuffs: Their Plasticity, Fluidity and Consistency*. North-Holland Publishing Co., Amsterdam, 1953, 80–90.

[81] E.B. Shand, *Glass Engineering Handbook*, 2nd ed. McGraw Hill Book Company, 1958.

[82] Robert W.G. Shipman, Morton M. Denn, and Roland Keunings, Mechanics of the "falling plate" extensional rheometer. *J. Non-Newtonian Fluid Mech.* 40 (1991) 281–288.

[83] P.J.P.M. Simons and R.M.M. Mattheij, The cooling of molten glass in a mould. In R. Ertekin et al. (eds) *Boundary Element Technology XI*. Computational Mechanics Publications, 1996, 3–12.

[84] L. Smith, R.J. Tillen and J. Winthrop, New directions in aspherics : glass and plastic. In M.J. Riedl (ed), *Replication and Molding of Optical Components*, Vol. 896, Proceedings of The Society of Photo-Optical Instrumentation Engineers. Washington; Bellingham, 1988, 160–166.

[85] S.L. de Snoo, R.M.M. Mattheij, and G.A.L. van de Vorst, Modelling of glass moulding, in particular small scale surface changes. In E.A. Lipitakis (ed), *Hellenic European Research on Mathematics and Informatics: HERMIS 96*. LEA, Athens, 1997, 162–183.

[86] T. Sridhar, V. Tirtaatmadja, D.A. Nguyen and R.K. Gupta, Measurement of extensional viscosity of polymer solutions. *J. Non-Newtonian Fluid Mech.* 40 (1991) 271–280.

[87] W. Staniar (ed), *Plant Engineering Handbook*, 2nd ed. McGraw Hill, New York, 1959.

[88] Y.M. Stokes, Creeping-flow computational modelling of optical quality free surfaces formed by slumping of molten glass. In B.J. Noye, M.D. Teubner and A.W. Gill (eds), *Computational Techniques and Applications: CTAC97*. World Scientific, Singapore, 1998, 671–678.

[89] Y.M. Stokes, E.O. Tuck and L.W. Schwartz, Extensional fall of a very viscous fluid drop. Department of Applied Mathematics, University of Adelaide, 1997.

[90] Y.M. Stokes, Thermal replication: a comparison of numerical and experimental results. In E.O. Tuck and J.A.K. Stott (eds), *The Proceedings of the 3rd Bi-*

ennial Engineering Mathematics and Applications Conference: EMAC98. The Institution of Engineers, Australia, 1998, 471–474.

[91] M. Teubner, Department of Applied Mathematics, University of Adelaide, South Australia. Personal communication.

[92] Fred T. Trouton, On the coefficient of viscous traction and its relation to that of viscosity. *Proc. Roy. Soc. A* 77 (1906) 426–440.

[93] E.O. Tuck, Y.M. Stokes and L.W. Schwartz, Slow slumping of a very viscous liquid bridge. *J. Eng. Math.* 32 (1997) 31–44.

[94] E.O.Tuck, Mathematics of honey on toast. In R.R. Moore and A.J. van der Poorten (eds), *National Symposium on the Mathematical Sciences: Multiplying Australia's Potential*, Australian Academy of Science, 1997. See also: `http://www.maths.adelaide.edu.au/Applied/staff/etuck/honey_on_toast/Tuck.html`

[95] E.O. Tuck, Department of Applied Mathematics, University of Adelaide, South Australia. Personal communication.

[96] B.W. van der Fliert, P.D. Howell and J.R. Ockendon, Pressure-driven flow of a thin viscous sheet. *J. Fluid Mech.* 292 (1995) 359–376.

[97] Andrew Warinner, 7 April 1997, `http://www.urbanlegends.com/science/glass.flow/glassmaking.html`

[98] W. Weaver, Jr. and P.R. Johnson, *Finite Elements for Structural Analysis*. Prentice-Hall, New Jersey (1984).

[99] Jonathan W. White, Jr., Physical Characteristics of Honey. In Eva Crane (ed), *Honey: A Comprehensive Survey*, Heinemann, London, 1975, 207–239.

[100] J.H. Williams, D.R.J. Owen and J.M.A.C. Sa, The numerical modelling of glass forming processes. *Collected papers, XIV Intl. Congr. on Glass* (1986) 138–145.

[101] S.D.R. Wilson, The slow dripping of a viscous fluid. *J. Fluid Mech.* 190 (1988) 561–570.

[102] W.Y.D. Yuen, Mathematical modelling for industrial applications. In B.J. Noye, M.D. Teubner and A.W. Gill (eds), *Computational Techniques and Applications: CTAC97*. World Scientific, Singapore, 1998, 27–50.

[103] Edgar Dutra Zanotta, Do cathedral glasses flow? *Am. J. Phys.* 66 (1998) 392–395.

[104] O.C Zienkiewicz, Flow Formulation for Numerical Solution of Forming Processes. In J.F.T. Pittman, O.C. Zienkiewicz, R.D. Wood and J.M. Alexander (eds), *Numerical Analysis of Forming Processes*. John Wiley and Sons, 1984, 1–44.

[105] O.C. Zienkiewicz and R.L. Taylor, *The Finite Element Method*, 4th ed. v.1. McGraw Hill Book Company, 1989.