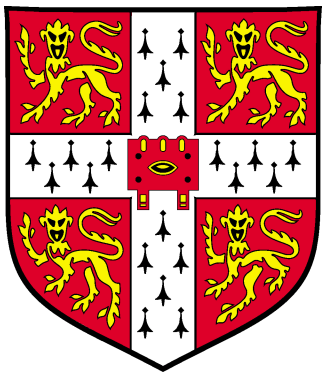

Several applications of a model for dense granular flows

Christopher John Cawthorn

Submitted in total fulfilment of the requirements
of the degree of Doctor of Philosophy

Corpus Christi College
and
Institute of Theoretical Geophysics
Department of Applied Mathematics and Theoretical Physics
University of Cambridge

This document is the result of my own work and includes nothing which is the outcome of work done in collaboration, except where specifically indicated in the text. No part of this dissertation has been submitted for any other qualification.

Christopher John Cawthorn

Several applications of a model for dense granular flows

Christopher John Cawthorn

Abstract

This dissertation describes efforts to evaluate a recently proposed continuum model for the dense flow of dry granular materials (Jop, Forterre & Pouliquen, 2006, *Nature*, **441**, 167-192). The model, based upon a generalisation of Coulomb sliding friction, is known to perform well when modelling certain simple free surface flows.

We extend the application of this model to a wide range of flow configurations, beginning with six simple flows studied in detailed experiments (GDR MiDi, 2004, *Eur. Phys. J. E*, **14**, 341-366). Two-dimensional shearing flows and problems of linear stability are also addressed. These examples are used to underpin a thorough discussion of the strengths and weaknesses of the model.

In order to calculate the behaviour of granular material in more complicated configurations, it is necessary to undertake a numerical solution. We discuss several computational techniques appropriate to the model, with careful attention paid to the evolution of any shear-free regions that may arise. In addition, we develop a numerical scheme, based upon a marker-and-cell method, that is capable of modelling two-dimensional granular flow with a moving free surface. A detailed discussion of our unsuccessful attempt to construct a scheme based upon Lagrangian finite elements is presented in an appendix.

We apply the marker-and-cell code to the key problem of granular slumping (Balmforth & Kerswell, 2005, *J. Fluid Mech.*, **538**, 399-428), which has hitherto resisted explanation by modelling approaches based on various reduced (shallow water) models. With our numerical scheme, we are able to lift the assumptions required for other models, and make predictions in good qualitative agreement with the experimental data.

An additional chapter describes the largely unrelated problem of contact between two objects separated by a viscous fluid. Although classical lubrication theory suggests that two locally smooth objects converging under gravity will make contact only after infinite time, we discuss several physical effects that may promote contact in finite time. Detailed calculations are presented to illustrate how the presence of a sharp asperity can modify the approach to contact.

Acknowledgements

Several people have offered suggestions, comments, and ideas related to the work contained in this thesis. I am particularly indebted to John Hinch for his advice, support, and seemingly infinite patience. I would also like to express my gratitude to Neil Balmforth, who was a wonderful person to work with during my stay in Woods Hole. In addition, I thank my supervisor, Herbert Huppert, for his support and advice throughout my doctoral studies.

My work has been funded by a variety of stipends and grants. I am especially grateful for a doctoral research grant from the Natural Environment Research Council (NERC). In addition, I would like to express gratitude to the organisers of the Geophysical Fluid Dynamics Program at the Woods Hole Oceanographic Institution for allowing me to attend their prestigious program in Summer 2008. I would also like to thank the Oxford Centre for Collaborative Applied Mathematics (OCCAM) for funding a three-month visit in Autumn 2009.

I would like to express my appreciation for those who have helped to keep me sane during my research. I am grateful to the members of the Institute of Theoretical Geophysics for their companionship, support, and stimulating discussions. A special mention should be made of Maurice Blount, Dominic Vella, and Andrew Wells, whose friendship has been an invaluable asset; and of Doris Allen, who kept everything running smoothly. Additional thanks ought to go to those people in the CMS and in Corpus with whom I have played games – it was always nice to ponder over something other than broken code!

Finally, I am eternally grateful to Jenny, who has been an unwavering source of support and encouragement throughout my studies. No matter how bleak things have seemed, she has encouraged me to retain my optimism and press onwards. The very existence of the document before you is proof positive of her success.

Contents

1 Introduction

1.1 Motivation

Granular materials have played an important role in daily life throughout history. For centuries, mankind has used the motion of sand through hourglasses to mark the passage of time. Cereals or rice form a staple part of almost every cultural diet on Earth. In the modern world, vast numbers of people live in dwellings and travel on roads constructed using granular aggregates. Furthermore, few children (at least in developed countries) will not, at some point in their lives, have sculpted castles, faces, or animals out of sand or snow.

There are countless engineering applications involving granular materials, some of which are illustrated by Figure 1.1. The construction industry deals with grains on a daily basis, not only for use as building materials, but also when assessing the quality and safety of ground on which new structures are to be built. Maritime engineers seek to improve and extend harbours by dredging deep channels into the sandy ocean bed. Pharmaceutical manufacturing is dominated by powders and pills, all of which need to be stored, transported, and mixed safely and efficiently. Similar problems exist in the farming of cereal grains, for which Figure 1.1(c) gives an indication that even the simple task of storing grains in a silo is perhaps not sufficiently understood.

Extending our view to include geophysical phenomena reveals many more applications.

Figure 1.1: Some applications involving granular material: (a) Cereal and (b) pharmaceutical materials; (c) Failure and buckling of a grain silo; (d) An hourglass; (e) Sand dunes in the Sahara desert; (f) An avalanche on the Liberty Wall Range, USA; (g) The slumping of a steep hillside following the collapse of a confining wall.

Avalanches and rockslides are a fact of life in mountainous areas. Although many efforts are taken to predict, divert, and arrest these deadly granular flows, millions of dollars' worth of damage are caused each year, and lives are frequently lost amid a tumbling tide of snow. In desert regions, the action of the wind causes the formation of mighty sand dunes, which migrate across the desert. The passage of these dunes can spread sand to new regions, contributing to the inexorable process of desertification in surrounding areas. In river and coastal science, one often observes the motion of grains in the forms of erosion, sedimentation, and scour. Spanning the fields of geophysics and engineering is the problem of soil liquefaction: during earthquakes, violent oscillations of the earth can cause packings of sand and soil alike to lose their solidity, putting at severe risk any buildings they support.

Despite the vast range of applications featuring granular matter, the theoretical understanding of how grains flow is still rather limited. Unlike fluid dynamicists, who have applied the celebrated Navier-Stokes equations to a vast array of complicated problems ranging from weather prediction to the swimming of creatures great and small, students of granular dynamics have yet to agree upon a uniformly valid continuum description of granular flows. Indeed, it may well be the case that such a general mathematical description does not exist at all. What is it that makes the modelling of granular media so complicated?

1.2 A brief review of granular phenomenology

One needs only to read the excellent reviews of Jaeger *et al.* (1996), de Gennes (1999), or Kadanoff (1999), to begin to appreciate the complexities of granular dynamics. Depending on the particular flow conditions, granular material can exhibit very different behaviour. Jaeger *et al.* (1996) separate the different types of granular flow by their analogy with solid, liquide, and gas-like behaviour, as appropriate. For example, a dense packing of grains exhibits the kind of structural rigidity commonly found in solid materials, yet grains poured vertically onto a surface can appear to behave very much like water in a kitchen sink, forming a granular equivalent of the well-studied hydraulic jump (see Boudet *et al.*, 2007, and Figure 1.2b). If a great deal of energy is supplied to a loose packing of grains, it is possible even to observe behaviour like that of a gas, such as in the particle dynamics simulations of Goldhirsch & Zanetti (1993), shown in Figure 1.2(c). This separation into solid-, liquid-,

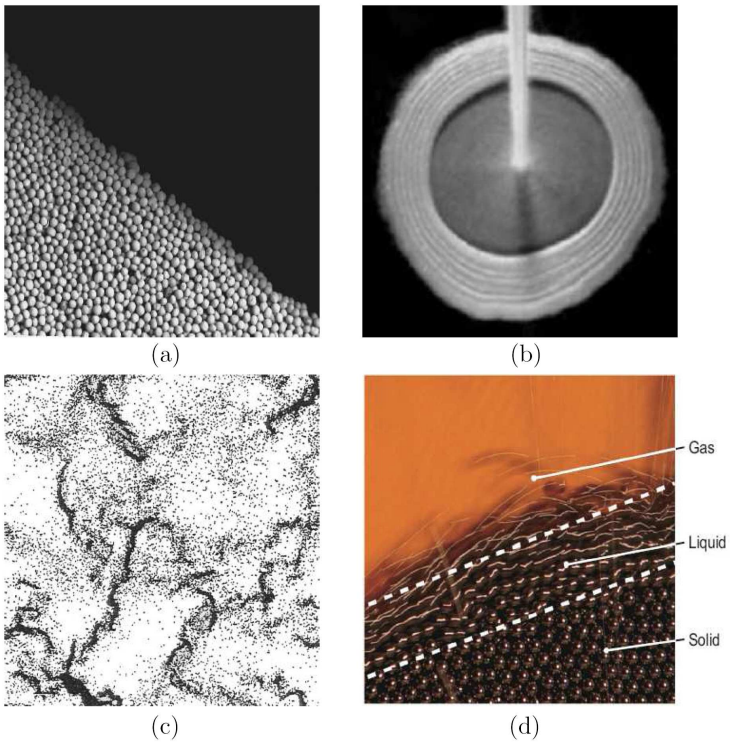

Figure 1.2: Experimental pictures of granular material behaving like (a) a solid, at rest in a pile, from Jaeger *et al.* (1996); (b) a liquid, forming a 'granular jump', (Boudet *et al.*, 2007); (c) a gas, in the particle dynamics simulations of Goldhirsch & Zanetti (1993) (note the clustering of particles, which is caused by the dissipation of kinetic energy due to friction). Panel (d) shows a simple sandpile flow, where all three 'phases' can be observed, from Forterre & Pouliquen (2008).

and gas-like phenomena is a little misleading. Even in fairly simple configurations, such as the flow of grains atop a static pile (Figure 1.2d), granular material can exhibit all three types of behaviour at once. Deep in the pile, the material is essentially solid. Near the surface, however, there is a zone of coherent motion reminiscent of that observed for fluid flow. Finally, a handful of grains bounce along the top of the pile, essentially unaffected by the motion of their nearest neighbours. The dashed lines in Figure 1.2(d) are somewhat arbitrary. In truth, there are no sharp 'phase transitions' between the different types of granular behaviour, as one might hope to observe between ice, water and steam, for example.

With such a wide variety of different behaviour possible, it seems unlikely that a single model will be capable of capturing the entire spectrum of granular phenomenology. However, a good description of gas-like granular flows already exists in modified forms of kinetic theory, such as those described by Pöschel & Brilliantov (2003). Furthermore, the civil engineering community has developed a good understanding of the solid-like behaviour under the umbrella of soil mechanics (see Chen & Baladi, 1985, for example). In this thesis, we focus on the facet of granular flow that is perhaps least well understood at present: that of dense, liquid-like flow, although we will on occasion consider the interaction of liquid-like regions of flow with solid-like static or creeping zones. The presence of moisture in the interstitial spaces between grains, or any cohesion intrinsic to the grains themselves, can greatly modify the observable dynamics. In the interest of pursuing a tractable problem, we avoid such complications by further restricting attention to dry, cohesionless granular materials.

Even with this specialisation to dry, cohesionless, dense flows, many complexities still remain. It is possible for very slow, creeping flows to exist (see Komatsu *et al.*, 2001, for example). In such flows, the dynamics may be very different to the dense, liquid-like flow observed in other experiments, with the creeping dynamics better described by plasticity theory than any form of fluid mechanics. Such differences in behaviour in slow, quasi-static flows are typically due to non-local interactions arising in the form of 'stress chains': sequences of grains through which any applied stress is transmitted throughout the material. As was wonderfully illustrated by the experiments of Howell *et al.* (1999) (Figure 1.3a), these stress chains form a complex network of interactions that are constantly in flux, and result in a rather inhomogeneous stress distribution. Such non-local interactions are also cited as the

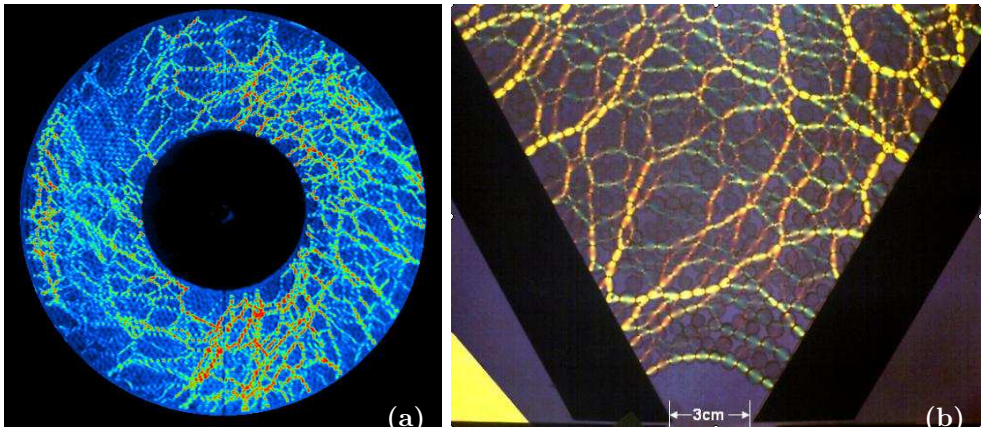

Figure 1.3: Stress chain formation in granular flow, made visible using photoelastic disks. Panel (a) shows the stress chains observed in the annular shear experiments of Howell *et al.* (1999). Panel (b) shows a similar visualisation of stress chains in hopper flow, from Tang *et al.* (2009). Note in particular the 'arch' of grains joining the two walls at the bottom of the hopper. The presence of this arch maintains the jammed state.

chief cause of *jamming* - that phenomenon well-known by industrialists and grain farmers alike, in which granular material suddenly stops flowing and reverts to solid-like behaviour. Jamming poses a significant problem that is still not well explained by existing theoretical models.

Aside from those experiments already mentioned, the last two decades have produced a vast amount of experimental data for granular materials, relevant to many different applications. Some researchers study avalanche dynamics (e.g. Hutter *et al.*, 1995), and the manner in which avalanches may be arrested or diverted (Hákonardóttir & Hogg, 2005; Hákonardóttir *et al.*, 2003); while others might consider the segregation of grains by their size (Möbius *et al.*, 2001), or the efficiency of attempting to mix two different granular species in a rotating drum (Gray, 2001). Other researchers devote their efforts to simpler flows, attempting to gain a detailed picture of the structure and rheology of granular flows. Special mention in this regard should go to the work of the GDR MiDi (2004), who have collated a wealth of experimental data for six benchmark flows. We shall explore their work in detail in Chapter 2.

1.3 Theoretical developments

The first theoretical study of granular materials is attributed to Bagnold (1954), who made observations of the flow properties of sand dunes, specifically involving the transport of grains by an externally applied wind. Following this seminal work, there was little development in the theoretical understanding of granular material until the late 1980s, when the field began to rise in popularity. Since then, however, the field has experienced a surge in interest, with many authors proposing models almost as diverse as the range of experiments performed.

Some models are entirely phenomenological, such as the so-called BCRE model (Bouchaud *et al.*, 1994), which attempts to describe the evolution of a sandpile by way of a two-layer model. One layer represents the static bulk of the pile, whilst the other represents mobile grains rolling and sliding down the surface. The evolution of each layer thickness is governed by an advection-diffusion equation, with a conversion term appropriate to each layer representing the erosion of the static pile or the deposition of rolling grains. This model, though widely regarded, is limited to the studies of sandpile flows, and offers no natural generalisation to more complicated flows. Furthermore, its phenomenological nature offers little insight into the physical effects needed to correctly model granular materials.

Another popular class of theoretical models of granular flows find their roots in shallow water theory. Such models make use of a depth-averaged form of the equations of conservation of mass and momentum, similar to the Saint-Venant equations of shallow fluid dynamics. The first and most popular such model was proposed by Savage & Hutter (1989), who accounted for simple Coulomb sliding friction at the base of the flow, and neglected internal stresses. This simple model, and its later generalisation to two-dimensional flows over complex topography (Gray *et al.*, 1999) have been used with some degree of success to explain experimental observations of avalanches (Hutter *et al.*, 1995). Furthermore, Chen & Lee (2003) have shown good agreement between a similar model and field observations of a landslide on Lantau Island, Hong Kong.

Many other, less renowned models have been proposed over the last few decades. Some are based on physical hydrodynamics, with an order parameter describing the local degree of mobility of the granular medium (Aranson & Tsimring, 2002), whilst other models are based on plasticity theory and allow for the diffusion of small 'spots' of mobility to drive

the flow (Kamrin & Bazant, 2007). Several other authors attempt to derive a continuum theory based on the application of a statistical or averaging procedure to a detailed model of the behaviour of individual grains, or clusters of grains (Jenkins & Savage, 1983; Ball & Blumenfeld, 2002).

In addition to these more traditional continuum models, it is equally possible to perform particle dynamics simulations in order to recreate an experimental flow (see Campbell & Brennen, 1985; Walton & Braun, 1986, for early examples). Detailed simulations following the motion of every particle in a system offer new insights into the internal structure of many experimental flows. Furthermore, it is possible to use the microscopic data (the stresses exerted on individual grains, for example) to inform the development of a continuum model. Two-dimensional approaches (e.g. Campbell & Gong, 1986) have eventually given way to more sophisticated models capable of deriving an approximation to three-dimensional continuum stresses based on the microscopic data (Goldhirsch & Goldenberg, 2002).

1.4 The $\mu(I)$ constitutive law

With the many different models mentioned in §1.3, and the many more fascinating yet unexplained experimental results (only the tip of the iceberg of which are mentioned in §1.2), the field of granular materials is somewhat in a state of flux. Disagreements about the 'correct' model to apply in a given situation are rife, due in part to the highly specialised nature of most models. Many models (such as shallow water models) have strong geometrical constraints, while the more general models are typically rather complicated, and relatively untested against experimental data. Recently, however, a relatively simple model (referred to hereafter as the $\mu(I)$ rheology, or the $\mu(I)$ constitutive law) has been proposed by Jop, Forterre & Pouliquen (2006), which extends a simple one-dimensional friction law to a full three-dimensional model based on the equations of non-Newtonian fluid mechanics. It is the study of this model to which this thesis is dedicated. Before discussing our aims and objectives for this work, we shall take the time to introduce the $\mu(I)$ law, and review the studies that led to its development.

An empirical frictional law

The $\mu(I)$ constitutive law is built largely upon an empirical friction law proposed by Pouliquen & Forterre (2002), which is in turn based upon a simpler law determined by Pouliquen (1999). In order to explain this law, we must describe a certain granular phenomenon. A flow of granular material down a rough inclined plane will come to a halt when the source of material at the top of the incline is removed. At this time, it is observed that the depth of the remaining deposit, referred to as h_{stop}, is a function only of the angle of inclination of the slope, θ, for a given set of material parameters. Following an experimental investigation into this effect, Pouliquen (1999) proposed a relationship between $h_{\text{stop}}(\theta)$ and the Froude number, Fr, of a steady flow at the same inclination angle, which has depth $h > h_{\text{stop}}(\theta)$ and volume flux $\langle V \rangle h$ per unit width:

$$\text{Fr} := \frac{\langle V \rangle}{\sqrt{gh}} = \beta \frac{h}{h_{\text{stop}}(\theta)} \tag{1.1}$$

In this equation, β is a dimensionless constant that depends on the granular material in question. For the glass beads used in his experiments, Pouliquen (1999) determined that $\beta \approx 0.136$.

In addition to the result (1.1), Pouliquen (1999) argued that the the flux in a steady flow must be controlled by the balance between frictional and gravitational forces. Using an analogy with simple Coulomb friction and applying a force balance to the flowing layer of grains, one is forced to conclude that the inclined plane must exert a frictional force, $\rho g h \mu_b$, on the granular medium, where

$$\mu_b \left(h, \langle V \rangle \right) = \tan \theta. \tag{1.2}$$

By inverting the relationship (1.1) to determine θ in terms of the unknown function h_{stop}, one can reduce this basal friction coefficient to a function of a single variable

$$\mu_b(h, \langle V \rangle) = \mu_{\text{stop}} \left(h \frac{\beta \sqrt{gh}}{\langle V \rangle} \right). \tag{1.3}$$

According to Pouliquen (1999), the relationship (1.3) holds provided that the flowing granular layer is deeper than h_{stop} (or, equivalently, that $\text{Fr} > \beta$). The results diverge from this simple form in the case of very slow, creeping flows with smaller Froude numbers. For flows that do satisfy this condition, Pouliquen & Forterre (2002) have empirically determined

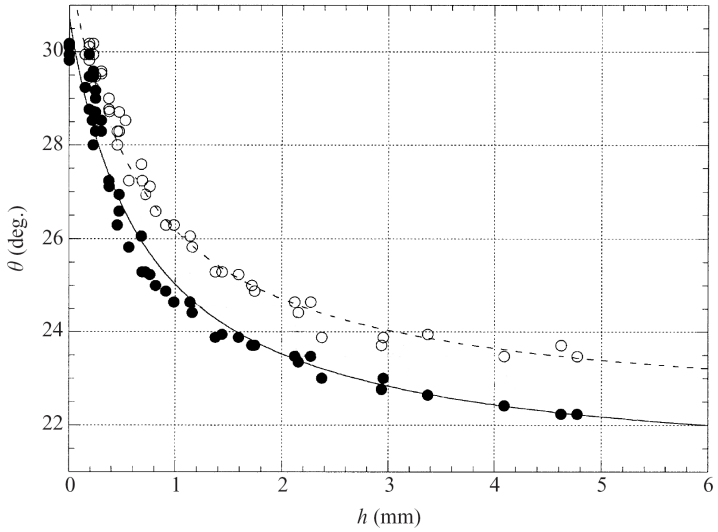

Figure 1.4: Relationship between h_{stop} and θ in experiments on a rough inclined plane (filled circles). Also shown is the threshold h_{start}, at which flow is initiated (open circles). Plot adapted from Pouliquen & Forterre (2002).

the functional dependence of h_{stop} on θ. A plot showing a typical form of this relationship is shown in Figure 1.4, and the key result is that one can use this data to determine the functional form of μ_{stop},

$$\mu_b(h, \langle V \rangle) = \mu_{\mathrm{stop}} \left(h \frac{\beta \sqrt{gh}}{\langle V \rangle} \right) = \mu_1 + \frac{\mu_2 - \mu_1}{1 + \dfrac{\beta h \sqrt{gh}}{\langle V \rangle L_0}}. \tag{1.4}$$

The quantities μ_1, μ_2 and L_0 are all parameters used by Pouliquen & Forterre (2002) to fit the form (1.4) to their experimental data. This friction law, along with some additional features designed to handle slower or shallower flows, was used by Pouliquen & Forterre (2002) in conjunction with a shallow water model in an attempt to predict the spreading of a finite volume of granular material released onto an inclined plane. The outcome was fairly successful, with qualitatively accurate predictions for the evolution of the granular current observed in several experiments.

The inertia number

The next development relevant to the $\mu(I)$ constitutive law came from the GDR MiDi (2004), who carried out a huge array of experiments and particle dynamics simulations in order to gain a thorough understanding of six simple flow configurations. Details specific to each of the six geometries will be discussed in §2.1. For the purpose of deriving the $\mu(I)$ law however, it is necessary only to mention their key result, which was common to all of the experiments.

For a given set of material parameters, the GDR MiDi (2004) observed that the local dynamical properties in each of the experimental flows appear to depend only upon a single dimensionless variable: the *inertia number*[1], which is defined by

$$I = \frac{d|\dot{\gamma}|}{\sqrt{\phi p/\rho}} = \frac{T_p}{T_{\dot{\gamma}}}, \tag{1.5}$$

where

$$T_{\dot{\gamma}} = |\dot{\gamma}|^{-1}, \qquad \text{and} \qquad T_p = d\sqrt{\frac{\rho}{\phi p}}. \tag{1.6}$$

In these definitions, d is a typical grain diameter, ρ the material density, and ϕ the volume fraction of the medium occupied by grains. In addition, p is the local pressure (defined as the magnitude of the isotropic stress in this case), and $|\dot{\gamma}|$ is the modulus of the rate of strain tensor.

The inertia number may be thought of as a ratio of two timescales, which are illustrated by Figure 1.5. The macroscopic timescale, $T_{\dot{\gamma}}$, is defined to be the characteristic time taken for two grains to move past each other because of large-scale shearing motions, set by the local rate of strain, $|\dot{\gamma}|$. By contrast, the microscopic timescale, T_p, is the characteristic time needed for a grain to be forced into the gap between its neighbours via the action of the confining pressure, p. The inertia number therefore measures the local mobility of grains: a large inertia number indicates that the grains are swept along by the mean flow with little deviation, whereas flows at smaller values of I are nearly static, with grains held tightly together by the confining pressure, and the material structure dominated by long-lasting contacts between grains.

[1]Referred to elsewhere in the literature as the Bagnold number, Weissenberg number, Savage number, or Coloumb number. The inertia number appears to be the most commonly applied name in the recent literature for dense granular flows.

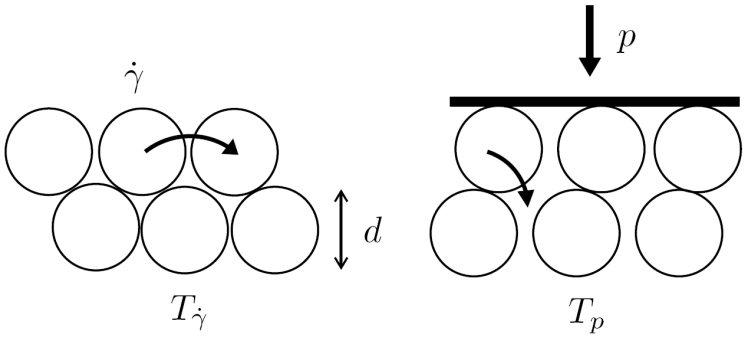

Figure 1.5: Schematic diagrams to illustrate the macroscopic and microscopic timescales ($T_{\dot{\gamma}}$ and T_p, respectively), used in the definition of the inertia number (1.5).

The relevance of the inertia number was revealed by the experiments of the GDR MiDi (2004). By comparing various measures of the local shear stress τ and normal stress p, the GDR MiDi (2004) argue for a local rheological law, in which the shear and normal stresses are related by the equation

$$\left|\frac{\tau}{p}\right| = \mu(I). \tag{1.7}$$

The coefficient of friction $\mu(I)$ is defined by analogy with standard Coulomb sliding friction and is a function only of the local inertia number.

Further interpretation of the experimental results led the GDR MiDi (2004) to conclude also that the local volume fraction, ϕ, is slaved to the local inertia number for a wide range of flows. Moreover, the variation of ϕ with I is essentially linear, but has extremely small gradient over the range of I recorded in the dense flow experiments. One can therefore expect the volume fraction to be essentially constant throughout a given dense flow. Based on this expectation, we shall treat all of the flows studied in this thesis as incompressible, and thus avoid any additional complications introduced by the need to consider dilation or compaction of the granular material.

The scalar $\mu(I)$ law

With the foundations of a local rheological law laid by the GDR MiDi (2004), the next objective was to propose a useful functional form for the friction coefficient $\mu(I)$. This

task was performed by Jop *et al.* (2005), who sought a model capable of reproducing the behaviour observed for granular flow atop a deep sandpile (which we shall discuss in §2.1.5), and found one by building upon the work of Pouliquen & Forterre (2002).

In the case of steady granular flow down an inclined plane, one should expect any predictions of the local $\mu(I)$ rheology to be consistent with the results of Pouliquen & Forterre (2002). More specifically, the $\mu(I)$ rheology must be consistent with the basal friction law (1.4). In such a situation, taking the z-axis perpendicular to the inclined plane (with $z = 0$ being the free surface), a simple force balance argument implies that the pressure distribution should be hydrostatic, and

$$\mu(I(z)) = \tan\theta. \tag{1.8}$$

From this we surmise that I must be a function of θ alone, and hence constant for a given experiment. We then use the definition of I (1.5) in conjunction with a hydrostatic pressure distribution and integrate once to recover the Bagnold velocity profile

$$\frac{V(z)}{\sqrt{gd}} = \frac{2}{3} I(\theta) \sqrt{\phi \cos\theta} \frac{(h^{3/2} - z^{3/2})}{d^{3/2}}, \tag{1.9}$$

It is now a simple matter to integrate (1.9) to find the mean velocity $\langle V \rangle$ and substitute into (1.4) to find that

$$\mu(I) = \mu(I)|_{z=-h} = \mu_b = \mu_1 + \frac{\mu_2 - \mu_1}{1 + \dfrac{5\beta d}{2L_0 I \sqrt{\phi \cos\theta}}}. \tag{1.10}$$

This motivates the definition

$$\mu(I) = \mu_1 + \frac{\mu_2 - \mu_1}{1 + I_0/I} = \mu_1 \left(\frac{I_0 + \lambda I}{I_0 + I} \right), \tag{1.11}$$

where

$$I_0 = \frac{5\beta d}{2L_0 \sqrt{\phi \cos\theta}} \tag{1.12}$$

for flow down an inclined plane. We also define the *friction ratio*,

$$\lambda = \frac{\mu_2}{\mu_1}. \tag{1.13}$$

This definition of λ will be used extensively throughout this thesis.

The parameters μ_1 and μ_2 depend only upon the materials used, and can be easily measured in straightforward experiments for a range of materials (Börzsönyi & Ecke, 2007).

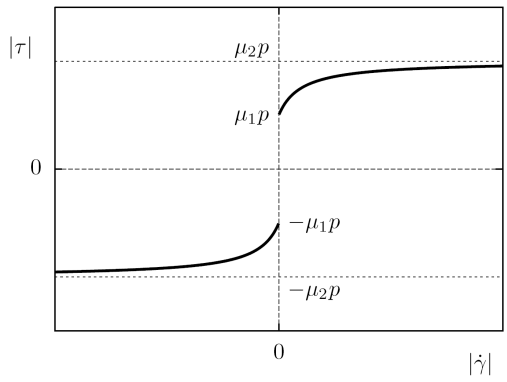

Figure 1.6: Plot of shear stress $|\tau|$ against rate of strain $|\dot{\gamma}|$ for a material obeying the $\mu(I)$ constitutive law (1.11) at constant pressure.

On the other hand, I_0 quite clearly depends upon the flow geometry. Even for a series of experiments on an inclined plane, (1.12) indicates that I_0 varies according to the inclination of the plane. This variation is undesirable in a constitutive model, so most authors (including Jop *et al.*, 2005), take an 'average' value for I_0 from one series of experiments, and attempt to use the same value in other geometries. This approach yields positive results in several cases, but performs poorly in other circumstances (see Chapter 2). We prefer to view I_0 as a fitting parameter that should be determined by the particular flow configuration under examination.

With the form (1.11) for $\mu(I)$ in hand, it is possible to define the *scalar* $\mu(I)$ *rheological law*, which applies to unidirectional flows whose velocity u depends on a single variable, z:

$$\tau = \mu(I) p \ \operatorname{sgn}\left(\frac{\partial u}{\partial z}\right) \quad \text{if} \quad \frac{\partial u}{\partial z} \neq 0. \tag{1.14}$$

This reflects the very sensible physical assertion that friction acts to oppose any shear present in the granular material. Furthermore, the shape of the $\mu(I)$ function (1.11), plotted in Figure 1.6, indicates two more desirable features of the model. First of all, one should observe that the shear stress does not vanish as $|\dot{\gamma}| \to 0$. In other words, there is a *yield stress*, $\mu_1 p$, which must be exceeded in order for the granular material to shear internally. This should hopefully mean that the $\mu(I)$ law is able to capture the transition between regions of solid-like and liquid-like behaviour, albeit in a rather simplistic sense, by predicting which

regions are below and above yield, respectively. In sub-yield regions (i.e. regions with no internal shear, and subject to an external stress less than $\mu_1 p$), the frictional forces will act to oppose any attempt to create shear, represented by the inequality

$$\tau \leq \mu_1 p \quad \text{if} \quad \frac{\partial u}{\partial z} = 0. \tag{1.15}$$

In general, it is not possible to determine the shear stresses inside a sub-yield region locally. Instead, one must consider the stress applied to the entire boundary of the sub-yield region, and use this to govern the motion of the whole region. The need to track the location of yield surfaces for such a calculation adds significant complications in more than one dimension. We shall discuss some strategies for dealing with sub-yield regions numerically in §2.2.

The bounded nature of $\mu(I)$ for large I is also a desirable feature of the rheological law (1.14), because it reflects the notion that the granular friction can only support a finite stress. If the granular medium experiences a stronger driving stress, it will accelerate without limit until the flow becomes sufficiently rapid as to invalidate the $\mu(I)$ model. In fact, it has proved difficult to explore flows at very high inertia number experimentally without crossing into a kinetic regime. The lack of data available in this limit suggests that the particular form of the saturation of $\mu(I)$ for large I is rather arbitrary. Whereas a linear variation of $\mu(I)$ for small I is generic for a wide class of functional forms, the I^{-1} relaxation of $\mu(I)$ towards $\lambda\mu_1$ as $I \to \infty$ is rather specialised, and very different results can be obtained if one were to use a different power law, or an exponential saturation. Thankfully, however, we shall not often need to trouble ourselves with high-I flows in this work.

This simple scalar constitutive law, with $\mu(I)$ defined by (1.11), may be applied to a range of unidirectional flows with varying degrees of success. A detailed account of such applications is given in Chapter 2. For now, we simply note that the use of this law was sufficient for Jop *et al.* (2005) to model their experimental deep pile flow in one dimension, but only if they included an additional term to account for frictional forces at the rough vertical walls containing the flow. For more details of their treatment, see §2.1.5.

Generalisation to three dimensions

In order to treat more complicated flows, it is necessary to generalise the scalar $\mu(I)$ law (1.14) to a form capable of handling two- or three-dimensional flows. Such an extension

was proposed by Jop *et al.* (2006), who attempted to improve upon their previous one-dimensional model in order to better describe their experimental observations for deep pile flow. By analogy with standard approaches in the study of certain non-Newtonian fluids, specifically Herschel-Bulkley and Bingham fluids (see Tanner, 1985, for example), and following similar lines to Goddard (1990), one multiplies the scalar constitutive law by a unit tensor in the direction of the rate of strain tensor to obtain the shear stress tensor, τ_{ij}. The result is the *tensor* $\mu(I)$ *constitutive law*,

$$\sigma_{ij} = -p\,\delta_{ij} + \tau_{ij}, \qquad \text{where} \qquad \tau_{ij} = \mu(I)p\frac{\dot{\gamma}_{ij}}{|\dot{\gamma}|}, \tag{1.16}$$

and $\mu(I)$ is given by (1.11). Throughout our discussion of granular flow, we define the rate of strain tensor to have the form $\dot{\gamma}_{ij} = \partial_j u_i + \partial_i u_j$, and its second moment by $|\dot{\gamma}| = \sqrt{\dot{\gamma}_{ij}\dot{\gamma}_{ij}/2}$. We note also that (1.16) reduces to (1.14) in the special case of unidirectional flow when the velocity depends on only one spatial direction. As with the scalar law, this tensor version also admits a yield stress, below which material does not undergo any shear. The only difference is that one must replace absolute values with second moments in the yield condition, thus

$$|\tau| > \mu_1 p, \qquad \text{where} \qquad |\tau| = \sqrt{\tau_{ij}\tau_{ij}/2}. \tag{1.17}$$

Once again, one ought to worry about sub-yield behaviour. If a material undergoes no local shear, then $|\dot{\gamma}| = 0$, and (1.16) is ill-defined. Our only recourse is to require that the material be sub-yield,

$$|\tau| \leq \mu_1 p, \tag{1.18}$$

and consider the motion of the entire shear-free region as a single entity. We shall elaborate on this issue in §2.2.

Given the form of the tensor $\mu(I)$ law, it is possible to define the *effective (kinematic) viscosity* of a granular medium under this law,

$$\nu_{\text{eff}} = \frac{\mu(I)p}{\rho|\dot{\gamma}|}. \tag{1.19}$$

In effect, the $\mu(I)$ approach models a granular continuum as a non-Newtonian fluid with non-constant, nonlinear viscosity given by (1.19). We note that this hypothetical fluid is shear-thinning (that is, its viscosity decreases as $|\dot{\gamma}|$ increases), but has a hitherto largely

unstudied dependence of viscosity on the local pressure. It is this pressure-dependence in the viscosity that renders difficult the calculations needed to model even rather simple flows. Nevertheless, the pressure dependence is crucial in dictating the frictional dynamics we expect of granular materials.

1.5 Aims of thesis

The $\mu(I)$ constitutive law of Jop *et al.* (2006) is still a relatively new model for granular materials, and only now is it being evaluated by a significant number of authors. Aside from a number of studies involving relatively straightforward flow configurations (Forterre & Pouliquen, 2008), the $\mu(I)$ model remains largely untested. It is the intention of this thesis to apply the $\mu(I)$ law to a wider range of flows, each of which has been investigated experimentally. The objective is not only to evaluate the performance and validity of the $\mu(I)$ law as it stands, but also to use the findings to inform the future development of any refinements to the model.

We begin this task in Chapter 2 with an examination of several unidirectional flows. We initially consider the configurations studied by the GDR MiDi (2004), and present some new calculations alongside the findings of other authors in order to establish a broad overview of the capabilities of the $\mu(I)$ constitutive law. In addition, we introduce some techniques that facilitate the numerical solution of unidirectional flows obeying the $\mu(I)$ rheological law in one and two dimensions.

In order to analyse more complicated flows, it is necessary to solve the equations of motion numerically. The nature of the model (in particular, the pressure-dependent effective viscosity), requires that numerical solutions must be obtained with rather more care than is usually necessary when dealing with Newtonian fluid mechanics. We devote Chapter 3 to the numerical study of the benchmark problem of the collapse of a granular column – a problem that has stubbornly resisted solution by other theoretical approaches in recent years. Before discussing the results and implications for the $\mu(I)$ rheology, a detailed outline of the numerical scheme used is also given in Chapter 3, with a thorough account of (several) unsuccessful attempts to develop a scheme based on a Lagrangian finite element method relegated to the Appendix.

In Chapter 4, we shed further light upon the $\mu(I)$ rheological law by considering problems involving stability and surface waves in granular media. Two quite different problems – the formation of sand dunes and that of longitudinal rolls in flow down an inclined plane – are described in detail, with a discussion of the relevant physical mechanisms leading to instability in each case.

Chapter 5 sits apart from the rest of this thesis, and is essentially independent of the other chapters. Contained within is an account of the largely unrelated problem of the convergence and contact of two objects separated by a viscous fluid. This work was carried out by the author during his Fellowship of the Geophysical Fluid Dynamics Program at the Woods Hole Oceanographic Institution. After this brief interlude, in Chapter 6 we take a synoptic view of the successes and failures of the $\mu(I)$ model. We shall also look to the future, and consider potential extensions, refinements, or novel applications for the $\mu(I)$ rheology.

2

Unidirectional flows

We begin our investigation by considering the application of the $\mu(I)$ rheological law to some simple unidirectional flows. In such flows, the fact that the continuity equation is trivially satisfied offers a significant simplification to the mathematical structure of the governing equations. Typically, steady solutions may be found analytically in closed form or may be obtained by relatively straightforward numerical computations. Where available, we will compare the predictions of the $\mu(I)$ law with data from experiments or particle dynamics simulations. Six simple flow configurations are discussed in §2.1.

In order to compute steady unidirectional flows in which the velocity field varies in two directions, it can be convenient to solve an unsteady problem numerically, allowing material to accelerate from rest and relax to a stable, steady state. In §2.2, we shall discuss some of the difficulties that arise when computing unsteady flows, and suggest methods by which these difficulties can be avoided. With a reliable numerical method in hand, we shall consider the behaviour of some two-dimensional unidirectional flows.

2.1 Steady flows

A great number of granular flow configurations have been studied in the recent literature, using both traditional experimental methods and particle dynamics simulations. In order to

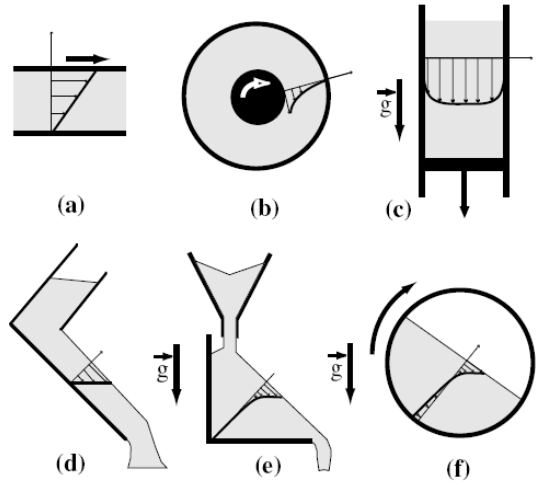

Figure 2.1: The flow geometries to be studied in §2.1, as examined experimentally by the GDR MiDi (2004). The geometries shown are (a) linear (Couette) shear, (b) annular (Taylor-Couette) shear, (c) vertical chute, (d) inclined plane, (e) deep pile, (f) rotating drum.

test the $\mu(I)$ constitutive law, we take advantage of the the work of the GDR MiDi (2004), who have collated a vast amount of data regarding the steady, dense flow of cohesionless grains in the six different configurations shown in Figure 2.1. Though conceptually straightforward, these flows present a range of different boundary conditions, including free surfaces, erodible beds, and frictional rigid walls. Working through each in turn, we shall highlight some of the key successes and failings of the $\mu(I)$ law.

2.1.1 Linear shear flow

Linear shear is the simplest geometry available when considering rheology. A sample of granular material is confined by a pressure p_0 between two rough parallel plates separated by a distance ℓ, and one plate moves at a velocity u_0 relative to the other, as shown in Figure 2.2(a). In the absence of gravity, the stresses must be constant throughout the sample. Under the $\mu(I)$ rheology, we can deduce that both the pressure and inertia number are constant, thus the strain rate must be constant. In order to obey no-slip boundary

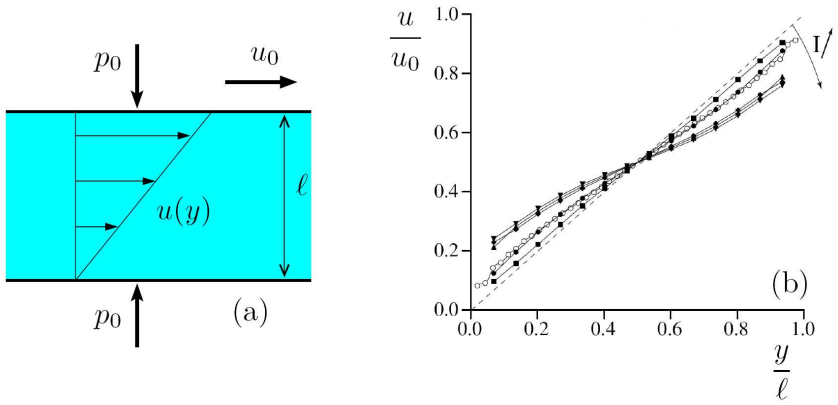

Figure 2.2: Linear shear flow: (a) Diagram and notation, (b) Velocity profiles for values of I between $I = 0.2$ and 0.5, obtained from molecular dynamics computations of da Cruz *et al.* (2005).

conditions on the confining walls, the internal velocity profile must be pure shear,

$$u(y) = u_0 \frac{y}{\ell}. \tag{2.1}$$

Although this gravity-free system is difficult to achieve in experiments, it can be modelled by discrete particle simulations, such as those carried out by da Cruz *et al.* (2005). The velocity profiles obtained from these simulations, normalised by the plate velocity, u_0, are shown in Figure 2.2(b). For small I, defined using the system shear rate $\dot{\gamma} = u_0/\ell$, the simulated velocity profile remains linear, in agreement with (2.1). However, for $I \gtrsim 0.1$, the material begins to slip at the boundaries, and the velocity profile develops a slight S-shape. The $\mu(I)$ rheology is not capable of capturing this non-trivial behaviour observed at larger shear rates. The reason for this failure is not entirely clear, although further probing of the simulations reveals that the volume fraction and velocity fluctuations vary significantly near the rough walls for large I. This slight fluidisation of the granular packing near the walls could account for the apparent slip there, which in turn leads to the S-shaped profiles in Figure 2.2(b). A more physically-realisable form of this system, in which gravity is allowed to play a significant role, will be discussed at length in §2.2.1.

2.1.2 Annular shear flow

A granular medium occupies the annular region $r_0 < r < r_1$ between two concentric cylinders. Both cylinders have rough walls, and impart a no-slip boundary condition to the granular material. The outer cylinder is held fixed while the inner cylinder is forced to rotate with angular velocity ω_0 by the application of a torque, G_0, as illustrated by Figure 2.3(a). Experimental procedures (see Losert *et al.*, 2000, for example) typically control the angular velocity and measure the torque required to maintain this rotation rate. In this analysis, it is more straightforward to solve the inverse problem: we specify the torque driving the inner cylinder, and calculate its rotation rate, along with the velocity profile in the granular medium. For simplicity, we shall neglect the effect of any axial variation in either pressure or velocity. Although not strictly appropriate to the experimental observations, the axially-invariant problem is much more accessible to analysis, and ought to provide a good approximation to the experimental behaviour.

This flow differs from the linear shear of §2.1.1 because the pressure is not constant in the flowing region. The pressure gradient that exists is necessary to balance the centripetal force due to rotation. A simple force balance argument reveals that the torque must remain constant in the radial coordinate, r, while the pressure increases. There will exist, therefore, a radius, r_y, beyond which the azimuthal shear stress, $\tau = G_0/2\pi r_y^2$, is exceeded by the yield stress, $\mu_1 p(r_y)$. In order to obey the no-slip condition at $r = r_1$, the sub-yield material in $r_y < r < r_1$ must remain at rest. Annular shear is thus our first example of a flow containing a static region, whose boundary is not known *a priori*. Our solution must therefore determine the location of this yield surface, $r = r_y$, at which both u and I vanish. It is at this point that our decision to impose the torque, rather than the angular velocity, becomes advantageous. Supposing that the pressure in the sub-yield granular region takes a constant value, p_0, we may determine the yield radius, r_y, from torque balance alone, before solving any differential equations.

Within the shearing region, $r_0 < r < r_y$, the velocity and pressure fields must satisfy equations corresponding to both torque and radial force balance. We choose the nondimensionalisation

$$r = r_0 R, \quad p(r) = p_0 P(R), \quad u(r) = \sqrt{\frac{p_0}{\rho_s \phi}} R\Omega(R)$$

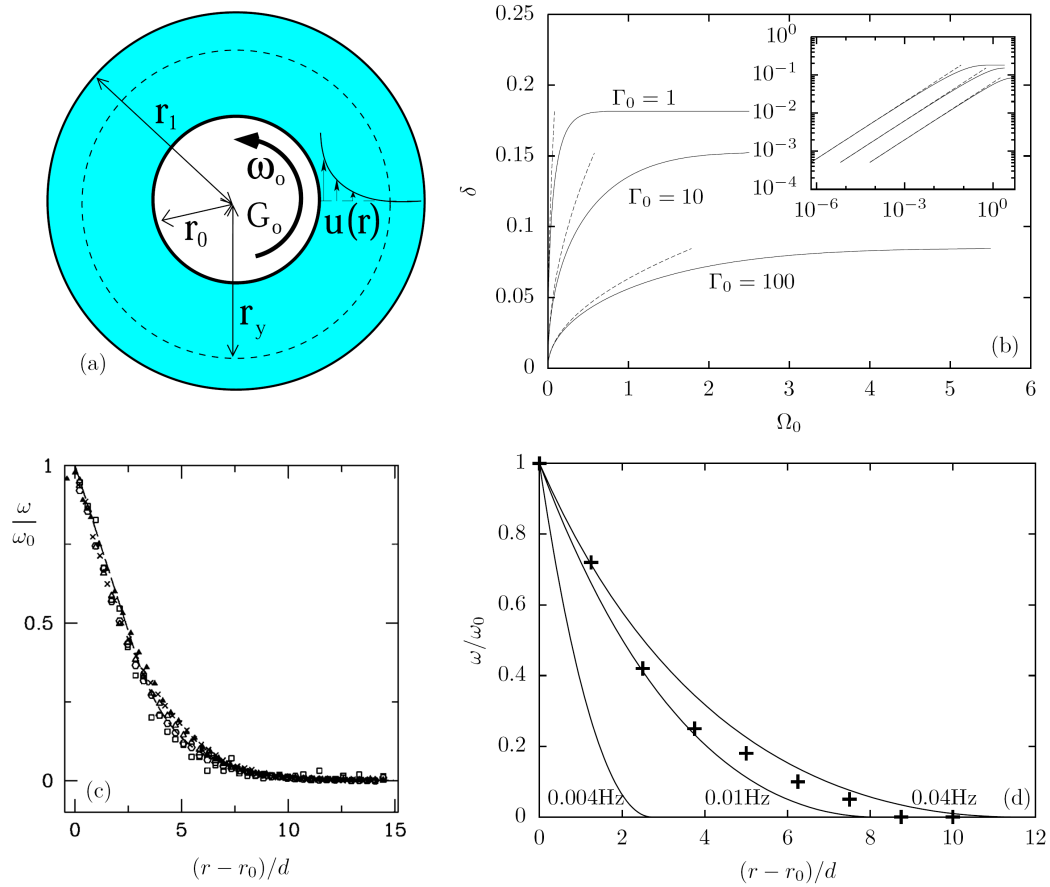

Figure 2.3: Annular shear flow. (a) Diagram and notation. (b) Shear layer thickness $\delta = R_y - 1$ as a function of rotation rate Ω_0 for $\lambda = 1.4$ and Γ_0 as shown. The solid lines show the result calculated numerically, whilst the dashed lines show the asymptotic result (2.11) in each case. The inset shows the same plot on logarithmic axes. (c) The experimental angular velocity profiles (normalised by ω_0) obtained by Losert *et al.* (2000) for a range of rotation rates. (d) Angular velocity profiles calculated using the $\mu(I)$ law, with parameters derived from the experimental parameters of Losert *et al.* (2000) ($\lambda = 1.8, \Gamma_0 = 14.7$) for the range of rotation rates shown. The points shown follow the trend of the data in panel (c).

and apply the $\mu(I)$ constitutive law to cast the dimensionless momentum equations into the form

$$\frac{\mathrm{d}P}{\mathrm{d}R} = R\Omega^2, \tag{2.2}$$

$$\frac{\mathrm{d}\Omega}{\mathrm{d}R} = \Gamma_0 \left(\frac{R_y^2 - R^2 P}{R_y^2 - \lambda R^2 P} \right) \frac{\sqrt{P}}{R}. \tag{2.3}$$

The parameters appearing in (2.3) are defined to be

$$R_y^2 = \frac{G_0}{2\pi\mu_1 p_0 r_0^2} \quad \text{and} \quad \Gamma_0 = I_0 \sqrt{\phi} \frac{r_0}{d}, \tag{2.4}$$

the former being simply a dimensionless form of r_y^2, calculated by balancing the applied torque, G_0, with the frictional torque at the yield point. This is the only control parameter in the problem. The parameter Γ_0 simply reflects the lengthscale imposed by the radius of the inner cylinder, measured relative to a grain diameter. The material parameters, I_0 and $\sqrt{\phi}$, are simply included in the definition of Γ_0 for convenience. The boundary conditions to be applied are

$$\Omega(R_y) = 0 \quad \text{(No-slip)}, \tag{2.5}$$

$$P(R_y) = 1 \quad \text{(Pressure normalisation)}, \tag{2.6}$$

Note that the requirement for I, hence the velocity gradient, to vanish at $R = R_y$ is clearly satisfied by evaluating (2.3) there.

The system (2.2)-(2.6) does not admit an analytical solution. However, it is an initial value problem, and can be readily integrated using an off-the-shelf, non-stiff integration scheme such as Matlab's ode45 routine. We simply integrate from $R = R_y$ to $R = 1$, and read off the angular velocity of the inner cylinder, $\Omega_0 = \Omega(1)$. This straightforward approach is made possible by our choice to impose the torque, rather than the angular velocity of the inner cylinder. In the forward problem, one would have to solve a two-point boundary value problem, as well as determine the free parameter, R_y.

As a check on the validity of the numerical solution, one can obtain an asymptotic solution for the regime in which the torque is only just sufficient to create motion. In this so-called *quasi-static* regime, the shearing layer will be very thin, so we take $R_y = 1 + \delta$, with $\delta \ll 1$. In order to determine the solution to leading order in δ, we make the rescaling

$$R = 1 + \delta\hat{R}, \quad \Omega = \delta^2\hat{\Omega}.$$

After rescaling, equations (2.2)-(2.3) take the form

$$\frac{\mathrm{d}P}{\mathrm{d}\hat{R}} = \mathcal{O}(\delta^5), \tag{2.7}$$

$$\frac{\mathrm{d}\hat{\Omega}}{\mathrm{d}\hat{R}} = \frac{\Gamma_0\sqrt{P}}{\delta}\left(\frac{1-P+2\delta(1-\hat{R}P)+\mathcal{O}(\delta^2)}{1-\lambda P+\mathcal{O}(\delta)}\right), \tag{2.8}$$

thus the pressure remains constant ($P = 1$) to leading order, and the angular velocity obeys the ordinary differential equation

$$\frac{\mathrm{d}\hat{\Omega}}{\mathrm{d}\hat{R}} = \frac{2\Gamma_0}{1-\lambda}(1-\hat{R}) + \mathcal{O}(\delta). \tag{2.9}$$

Together with the no-slip boundary condition (2.5), applied at $\hat{R} = 1$, we obtain the following angular velocity profile, presented in the original dimensionless variables,

$$\Omega(R) = \frac{\Gamma_0}{\lambda-1}(R-1-\delta)^2. \tag{2.10}$$

The angular velocity of the inner cylinder is therefore

$$\Omega_0 = \Omega(1) = \frac{\Gamma_0}{\lambda-1}\delta^2 + \mathcal{O}(\delta^3). \tag{2.11}$$

It is a key observation of many experiments (see GDR MiDi, 2004; Losert *et al.*, 2000, for example) that the thickness of the shearing layer, $\delta = R_y - 1$, is essentially independent of the rotation rate of the inner cylinder. Panel (b) of Figure 2.3 illustrates the predictions made by numerical solution of the $\mu(I)$ model. For high rotation rates, the shear layer thickness is roughly constant for each value of Γ_0, as one might hope. However, the shear layer thickness vanishes quadratically for small rotation rates, in agreement with the asymptotic result (2.11).

One might hope that all of the experimental data originates in the large-Ω_0 regime, but a quantitative comparison with the experiments of the experiments of Losert *et al.* (2000) indicates otherwise. Losert *et al.* (2000) took free-surface measurements of an annular shear flow using glass beads, the properties of which were measured by Börzsönyi & Ecke (2007). Combining the experimental parameters of Losert *et al.* (2000) with the material parameters of Börzsönyi & Ecke (2007), we are able to estimate the parameters $\lambda \approx 1.8$ and $\Gamma_0 \approx 14.7$. The angular velocity profiles obtained by Losert *et al.* (2000) for dimensionless rotation rates in the range $\Omega_0 \approx 10^{-4} - 10^{-1}$, are shown in Figure 2.3(c), with the angular velocity

normalised by its maximum value. Note that the profiles roughly collapse onto a master curve, with a shear layer thickness approximately equal to $10d$. However, calculations carried out using the $\mu(I)$ model, shown in Figure 2.3(d), do not exhibit the a similar collapse over the same range of rotation rates. Although the agreement with the experimental data is reasonable for larger rotation rates (corresponding to $\omega_0 \gtrsim 0.01$Hz), the predictions for slower flows are very poor.

It is possible that uncertainty in the model parameters may explain the discrepancy. In particular, the value of I_0 is estimated for flow on an inclined plane, and even then is only strictly correct for a specific inclination. Daniel *et al.* (2007) offer the possibility that the value of the parameter I_0 determined by Börzsönyi & Ecke (2007) for flow on an inclined plane may not be appropriate to confined shearing flows. This leads to some uncertainty in the value of Γ_0, which can change the range of rotation rates over which the shear layer thickness is constant (as shown by Figure 2.3b). Daniel *et al.* (2007) effectively chose to fit a value of Γ_0 (hence I_0) to the experimental data, and obtained an angular velocity profile in much better agreement with the experiments. Nevertheless, the problem remains that, for a given set of parameters, there exists a range of sufficiently small rotation rates for which the shear layer thickness varies, contradicting the experimental observations.

We conclude our discussion of annular shear with a short summary. The $\mu(I)$ constitutive model is capable of producing shear layers qualitatively similar to those observed in a range of experiments. However, the thickness of these shear layers is consistently underestimated - typically being only one or two grain diameters in thickness. One can fit parameters to match the velocity profiles of particular experiments, but it is not possible to capture more general observations, such as the invariance of shear layer thickness under a changing rotation rate. Of course, one may call into question the validity of a continuum approach when the flowing region is only a handful (6-8 in most experiments) of grain diameters in thickness. Furthermore, such confined shearing experiments are prone the formation of long-range force chains, as shown by the experiments of Howell *et al.* (1999). We can not hope to capture the nonlocal effects introduced by such chains by using a simple, local rheological law like the $\mu(I)$ law. It is the combination of these problems that results in the poor performance of the $\mu(I)$ law when attempting to model confined shear flows.

2.1.3 Flow in a vertical chute

We now consider a simplified version of a problem of great importance to industry, namely the draining of a silo. In the scenario investigated experimentally by the GDR MiDi (2004), granular material occupies the region between two rough, vertical walls separated by a distance, $2L$, as shown in Figure 2.4(a). In the experiments, a piston moves downwards at the base of the granular material, controlling the flow rate. The resulting velocity profiles resemble those shown in Figure 2.4(b). In the particular experiments of Chevoir *et al.* (2001), material in the centre of the chute typically forms a shear-free plug region, which is flanked by a shearing zone necessary to decrease the velocity to zero on the rough walls. The shearing zone is observed to be around 5-10 grain diameters in thickness, independent of all physical parameters except the roughness of the walls. Having recorded the failure of the $\mu(I)$ law in predicting such shear zones for annular shear in §2.1.2, one might expect a similar lack of success here. Nevertheless, in the interest of thoroughness, we shall pursue the calculation.

In an experiment, one controls the flow rate in the chute, and measures the velocity profile, from which one can find the size of the plug. As with annular shear, it is the pressure in the material that controls the size of the shearing region, so solving the inverse problem is the easier task. We would like to assume, therefore, that we know the pressure distribution in the chute. Indeed, the horizontal component of the momentum equation requires that the pressure is constant across the flow, so we need only know about the vertical pressure variation.

In order for the flow to be unidirectional, it is necessary that there must be no vertical variation in the downward velocity, otherwise mass will not be conserved. Clearly, however, there will be be some depth-dependent effects near the surface and base of the granular column. We choose to perform our analysis far from either end, where the velocity and pressure may be assumed to be independent of vertical position. The existence of such a region may be justified by the so-called *Janssen effect* (Janssen, 1895). It is experimentally observed that the pressure in a deep cylinder of granular material does not vary with depth once one is sufficiently far below the free surface. This is unlike the hydrostatic pressure gradient one would expect from a fluid column. This effect occurs because the frictional

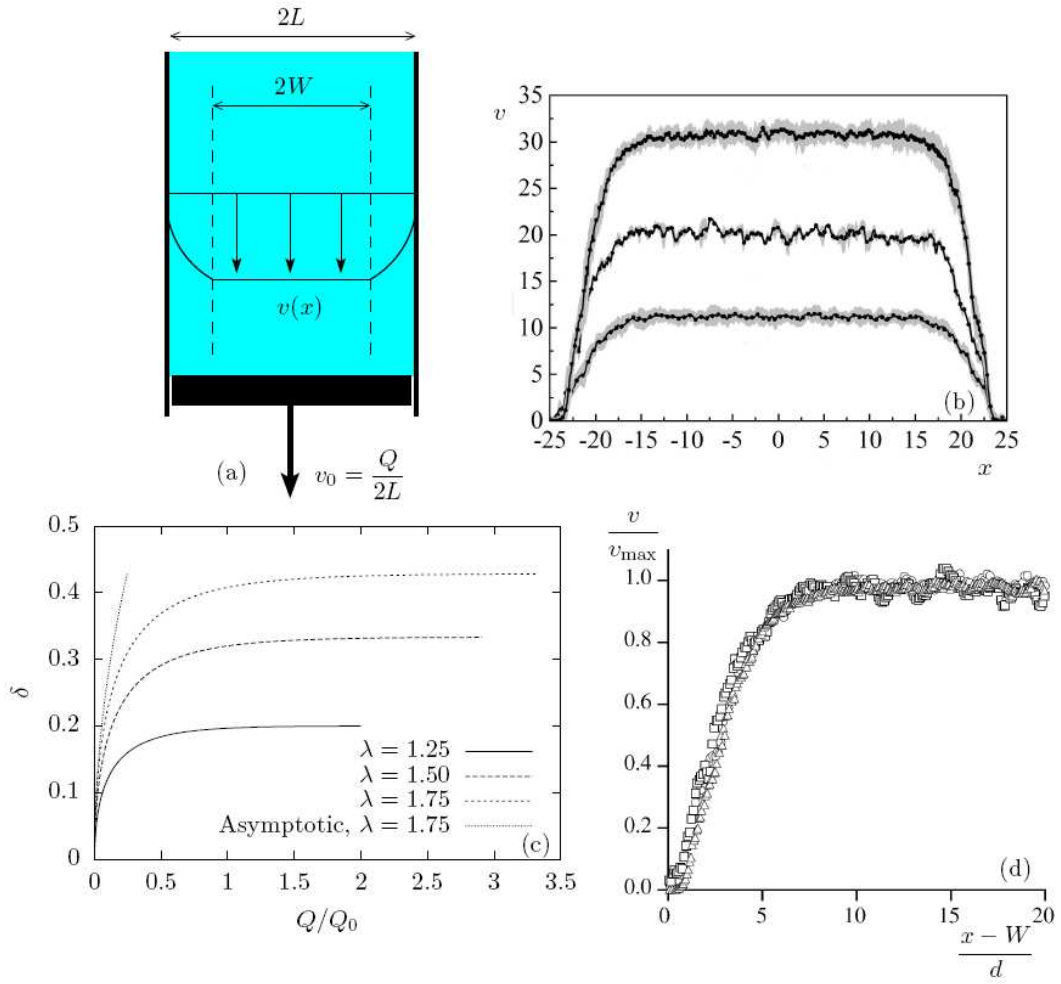

Figure 2.4: Flow in a vertical chute. (a) Diagram and notation; (b) Experimentally-observed velocity profiles for three different flow rates, from Chevoir *et al.* (2001); (c) Graphical representation of the relationship between shear layer thickness and area flux, obtained by integration of (2.16), and by the asymptotic approach leading to (2.22); (d) Experimental profiles of (b), normalised by the plug velocity.

forces on the side walls (which are proportional to the local pressure) support most of the weight of the grains, rather than allowing the weight to transmit throughout the material. To illustrate this point, consider a two-dimensional material slice of thickness δz and width $2L$ held stationary between two rough vertical walls. Balancing the weight of the slice with the difference in pressure across the slice and the (upward) frictional force on the walls, then taking the limit $\delta z \to 0$ gives the relationship

$$\rho g L = \mu_w p - L\frac{\mathrm{d}p}{\mathrm{d}z} \quad \Rightarrow \quad p(z) = \rho g \ell \left(1 - Ae^{-z/\ell}\right), \tag{2.12}$$

where μ_w is the wall friction coefficient, and A is a constant of integration. This exponential distribution is representative of the Janssen effect: for depths greater than the *Janssen length* $\ell = L/\mu_w$, the pressure is essentially constant. For experiments with moving walls (Bertho *et al.*, 2003, for example), the pressure deep in the cylinder remains essentially uniform in depth, but may vary in time according to the flow conditions. It is this depth-invariance that justifies our search for a vertically-invariant, unidirectional solution.

Supposing that the central plug region has width $2W$, we need only solve for the velocity profile in the shearing region, $W < x < L$, and appeal to symmetry to determine the velocity in the identical shearing layer near the wall at $x = -L$. For unidirectional flow, the horizontal component of the momentum equation requires that the pressure maintains a constant value, p_0, across the chute. At the edge of the plug, the inertia number drops to zero, so the vertical shear stress there is $\mu_1 p_0$. This must support the weight of the plug, thus

$$2\rho g W = 2\mu_1 p_0 \quad \Rightarrow \quad W = \frac{\mu_1 p_0}{\rho g}. \tag{2.13}$$

Furthermore, one can integrate the vertical momentum equation,

$$0 = -\rho g + \frac{\mathrm{d}\tau}{\mathrm{d}x}, \tag{2.14}$$

to find that

$$\tau = -\mu(I)p_0 = \rho g(W - x) - \mu_1 p_0 = -\rho g x, \tag{2.15}$$

where we have used the condition $\tau = -\mu_1 p_0 = -\rho g W$ at $x = W$. The negative sign here arises from the fact that velocity is measured vertically downwards. Substituting the form

of the friction coefficient $\mu(I)$ leads to the result

$$\frac{I}{I_0} = \frac{W-x}{x-\lambda W} \quad \Rightarrow \quad \frac{\mathrm{d}v}{\mathrm{d}x} = \sqrt{\frac{gWI_0^2}{\mu_1 d^2}}\frac{W-x}{x-\lambda W}. \tag{2.16}$$

Together with a no-slip condition applied at the rough walls of the chute,

$$v(L) = 0, \tag{2.17}$$

the differential equation (2.16) completely describes the flow in the shearing region near the wall, which has thickness $L - W$. In order to find the plug velocity, one must integrate (2.16), and evaluate the velocity at $x = W$. To find the relationship between the plug size, W, and flow rate, Q, one needs then to integrate the velocity across the chute. Although this calculation may be performed analytically, it is difficult to extract any meaningful interpretation from the resulting expression, so it is omitted. Instead, we present an asymptotic result in the limit of low flow rate. In such a situation, as with the annular shear of §2.1.2, we expect the shearing layer to be very thin, so we take

$$W = L(1-\delta), \quad \text{with} \quad \delta \ll 1. \tag{2.18}$$

After introducing the rescaled coordinate

$$X = \frac{x-L}{\delta L}, \qquad -1 < X < 0, \tag{2.19}$$

one can expand (2.16) in powers of the dimensionless shear layer thickness δ, and integrate the resulting expansion term-by-term. After a little work, we find that

$$v(X) = -I_0\frac{L}{d}\sqrt{\frac{gL}{\mu_1}}\left[\frac{\delta^2}{2(\lambda-1)}(2X+X^2) + O(\delta^3)\right]. \tag{2.20}$$

The plug velocity is then

$$v_{\mathrm{plug}} = v(x=W) = v(X=-1) = I_0\frac{L}{d}\sqrt{\frac{gL}{\mu_1}}\left[\frac{\delta^2}{2(\lambda-1)} + O(\delta^3)\right], \tag{2.21}$$

and the leading-order form of the volume flux is

$$Q = 2\left[Wv_{\mathrm{plug}} + \int_W^L v(x)\,\mathrm{d}x\right] = Q_0\left[\frac{\delta^2}{(\lambda-1)} + O(\delta^3)\right], \tag{2.22}$$

where $Q_0 = I_0 L^2 \sqrt{gL/\mu_1}/d$. As with the case of annular shear, we find that the shear layer thickness scales with the square root of the velocity difference across the layer in the limit of small inertia number. The relationship between the shear layer thickness and the flow rate, obtained via direct integration of (2.16), is shown graphically in Figure 2.4(c), along with the asymptotic result (2.22). As with annular shear, we note that the size of the plug (hence the size of the shearing regions) tends to a constant value for high flow rates.

Panel (d) of Figure 2.4 shows the experimentally observed velocity profiles obtained by Chevoir *et al.* (2001) (shown in Figure 2.4b) normalised by the plug velocity in each case. The striking result is that they coincide, showing that the plug has the same thickness despite the change in flux. Further experiments by Pouliquen & Gutfraind (1996) indicate that the limiting shear zone thickness is also independent of the width of the chute, and depends only upon the material properties of the grains.

Once again, we are reminded of the annular shear, and point out that this observation is not in agreement with the $\mu(I)$ model. Having δ saturate at high flow rates indicates that the limiting shear zone thickness scales with the width of the chute, not the grain diameter. At lower flow rates, which are more typical of those studied in the experiments mentioned, Figure 2.4(c) shows that the shear layer thickness varies with flow rate. Whether in a high or low flow rate regime, the $\mu(I)$ model once again fails to correctly predict the shear layer size in a confined flow.

2.1.4 Flow down an inclined plane

The flow of granular material down an inclined plane is one of the most commonly studied experimental configurations. Given the ubiquity of such a flow in both industrial and geophysical settings, this is hardly surprising. The study of the GDR MiDi (2004) catalogues a number of experimental investigations, yet barely scratches the surface of the work done in the past. Many phenomena have been observed in inclined plane experiments, including roll waves (Forterre & Pouliquen, 2003), levée formation (Deboeuf *et al.*, 2006, for example) and fingering (Pouliquen *et al.*, 1997), to name but a few. In this section, we consider only the simplest possible case of a unidirectional flow of a layer of constant thickness, h, of granular material on a slope of fixed inclination, θ, as shown in Figure 2.5(a).

Even in this very simple configuration, experimental procedures can vary from investi-

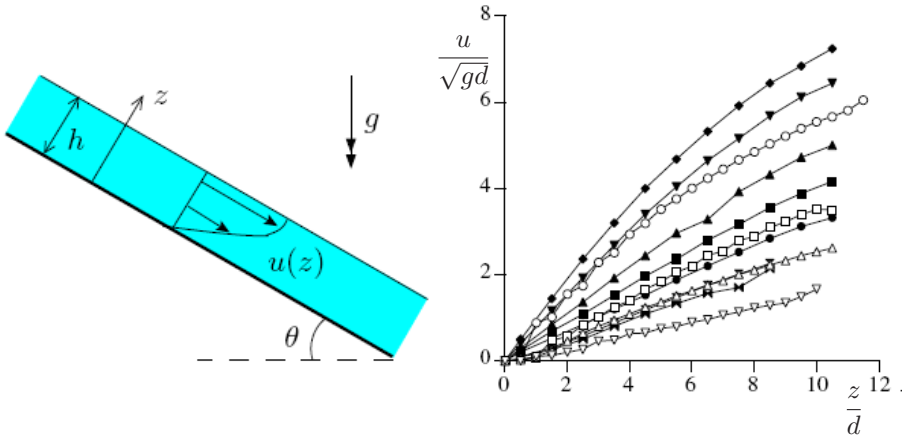

(a) (b)

Figure 2.5: Flow down an inclined plane. (a) Diagram and notation; (b) Velocity profiles for inclination angles $\theta = 21^\circ - 36^\circ$ from a range of experiments and simulations, collated by the GDR MiDi (2004).

gation to investigation. Material properties, basal roughness and the presence and type of any side walls can all make a difference. Of the experiments described by the GDR MiDi (2004), most use smooth side walls and a base made rough by gluing grains to a smooth surface. Typical velocity profiles obtained experimentally for shallow flows are shown in Figure 2.5(b). For deeper flows, the confining side walls can play a role – we will discuss such effects in §2.2.6. The velocity profiles shown all roughly conform to a Bagnold-like profile. That is,

$$\frac{u(y)}{\sqrt{gd}} = A(\theta)\frac{\left[h^{3/2} - (h-z)^{3/2}\right]}{d^{3/2}}, \tag{2.23}$$

where $A(\theta)$ is a dimensionless parameter dependent only upon material properties of the grains and the inclination angle of the plane. More careful investigation reveals that the profile becomes less concave, more linear as the inclination angle decreases toward the critical angle, θ_{stop}, below which no flow occurs. Applying the $\mu(I)$ rheological law and integrating to find the velocity profile yields precisely the form (2.23), with

$$A(\theta) = \frac{2}{3}\mu^{-1}(\tan\theta)\sqrt{\phi\cos\theta}. \tag{2.24}$$

This prefactor agrees very well with the experimentally-observed values of $A(\theta)$ above the flow threshold, but this is no surprise. We must recall that the $\mu(I)$ rheology was originally

derived by considering flow on an inclined plane. In particular, this exact experiment was used to fit the parameters μ_1, λ and I_0, as well as the form of $\mu(I)$ itself. Nevertheless, flow on an inclined plane offers the first flow to which the $\mu(I)$ model may be applied successfully.

Aside from the rather trivial agreement with experimental velocity profiles, the $\mu(I)$ rheology fails to capture some of the more subtle aspects of flow on an inclined plane. In particular, the change in character of the velocity profile from Bagnold to linear near the flow threshold is not predicted by the $\mu(I)$ law. Furthermore, it is well-known that the flow threshold for granular material depends on whether the material is moving. Material at rest will not spontaneously begin to flow until the plane is inclined above an angle, $\theta_{\text{start}}(h)$, but material already in motion will continue to flow steadily until the inclination drops below a smaller angle, $\theta_{\text{stop}}(h)$ (Pouliquen & Forterre, 2002). The $\mu(I)$ rheology cannot capture this hysteretic behaviour: it either predicts Bagnold motion of the form (2.23), or no motion at all, depending on the angle of inclination, θ. A theoretical explanation and description of the subcritical regime, $\theta_{\text{stop}} < \theta < \theta_{\text{start}}$, in which material can be either static or steadily flowing, must await a more sophisticated model, perhaps based on non-continuum or non-local effects.

2.1.5 Flow over a deep pile

The flow of granular material atop a deep pile, as illustrated by Figure 2.6(a), is another flow common in both nature and industry, and as such has been extensively studied in experiments. The results of several such experiments have been reported by the GDR MiDi (2004). Typically, one confines a deep pile of grains between two vertical walls, and introduces new material to the top of the pile at a particular flow rate, Q. In order to allow a steady state to form, material reaching the bottom of the pile is allowed to fall into another container. The surface of the pile will select an inclination angle θ, which is dependent upon the flux.

The presence of the side walls typically causes the observed velocity profiles to vary in both directions perpendicular to motion, although reducing the roughness of the walls can make the velocity less dependent on the cross-stream coordinate. This is a key finding of the work of Jop *et al.* (2005). Typical velocity profiles from these largely two-dimensional experiments are shown in Figure 2.6(b). We note that the velocity is essentially linear near

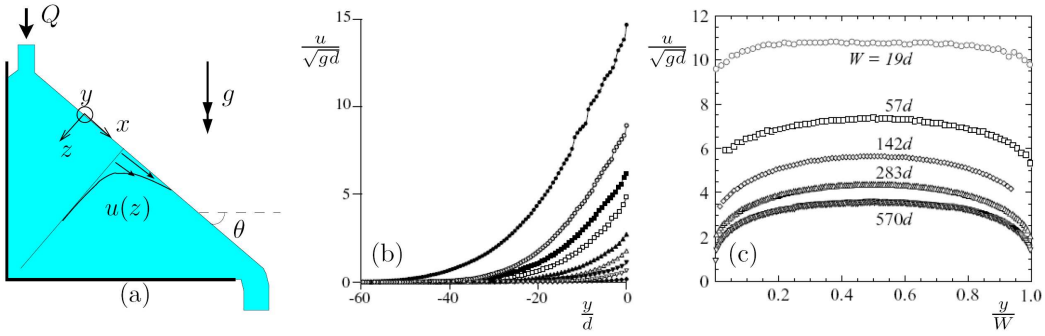

Figure 2.6: Flow over a deep pile: (a) Diagram and notation; (b) experimental velocity profiles in depth for various flow rates reported by the GDR MiDi (2004); (c) Variation of surface velocity across the slope for fixed Q and varying width W, from Jop *et al.* (2005).

the free surface, and decreases exponentially to zero deep in the pile. One might expect there to be an entirely stationary region far below the surface, but careful experiments by Komatsu *et al.* (2001) indicate that a very slow 'creeping' flow exists, so the entire pile is in motion. Such profiles are also observed experimentally for deep flows on an inclined plane. In this case, it is again the influence of the side walls that prevent the formation of a Bagnold profile of the form (2.23).

Given that the side walls are crucial in determining the form of the velocity profile, the problem ought to be solved in two dimensions. The structure of the $\mu(I)$ rheological law means that the most straightforward route to a solution is to perform a numerical calculation, although it is possible to make some analytical progress using a method of characteristics (de Ryck *et al.*, 2008). We shall return to this problem in §2.2.6, after we have developed an appropriate numerical scheme. However, we note that one could take a simplified approach, suggested by Jop *et al.* (2005), by considering the balance of forces on a cuboidal element of depth z spanning the entire channel width, W. We assume that the velocity is independent of the cross-stream coordinate, so that the pressure is hydrostatic, $p(z) = \rho g \cos\theta z$. We then balance the downstream component of the element's weight,

$$F_{\text{weight}} = \rho g \sin\theta W z, \tag{2.25}$$

the internal frictional stress on the base,

$$F_{\text{base}} = p(z)\mu(I)W = \rho g \cos\theta \mu(I) W, \tag{2.26}$$

and frictional stresses from the side walls (which are assumed to have constant friction coefficient μ_w),

$$F_{\text{wall}} = 2\int_0^z \mu_w p(z)\,\mathrm{d}z = \mu_w \rho g \cos\theta z^2, \tag{2.27}$$

to find that

$$\mu(I) = \tan\theta - \mu_w \frac{z}{W}. \tag{2.28}$$

Since $\mu(I)$ is bounded below by μ_1, this simplified model appears to fail below a depth $h = W(\tan\theta - \mu_1)/\mu_w$. Below this depth, we are forced to conclude that that material must not be shearing, hence it must be stationary. This does not, therefore, allow for agreement with the observations of creeping flow by Komatsu *et al.* (2001). Nevertheless, the velocity profiles obtained by integrating (2.28) offer very good agreement with the experimental profiles of Figure 2.6(b). Discussion of the cross-flow variation shown in Figure 2.6(c) will be delayed until we can compare with our numerical calculations in §2.2.6.

Aside from giving good predictions of velocity profiles, the $\mu(I)$ model falls short of capturing more aspects of the phenomenology of pile flows. As with flow on an inclined plane, there is an observable flow threshold that must be exceeded in order to obtain steady flow. Below this threshold, which depends upon the width of the experimental apparatus, one typically observes a series of intermittent avalanches (see Rajchenbach, 1990, for example), causing the angle of inclination of the free surface to oscillate between the two critical angles $\theta_{\text{start}}(W)$ and $\theta_{\text{stop}}(W)$. The $\mu(I)$ rheology is incapable of predicting the hysteresis that occurs for angles between these limits.

2.1.6 Flow in a rotating drum

The final configuration to be investigated experimentally by the GDR MiDi (2004), the rotating drum, is again simple to describe. A cylindrical drum partially filled with grains rotates about its (horizontal) axis with angular velocity Ω, as shown in Figure 2.7(a). This is another flow of great interest to industry. In particular, such flows are typically used to mix two or more granular species, and can be very efficient in doing so (see Gray, 2001, for example).

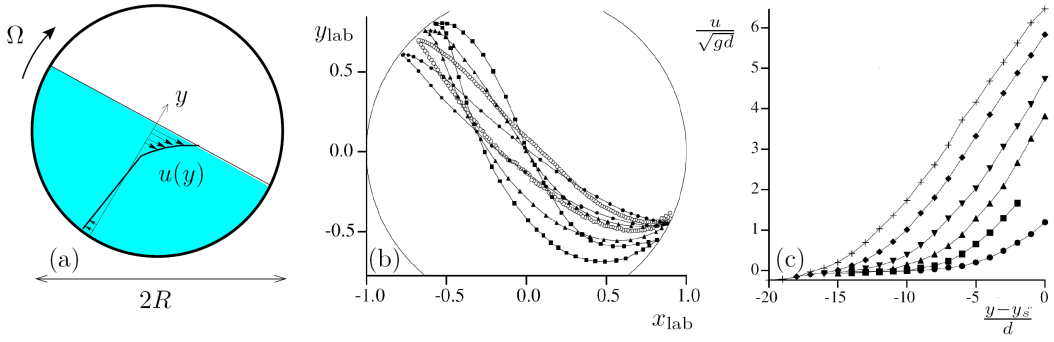

Figure 2.7: Flow in a rotating drum. (a) Diagram and notation; (b) Shape of the shearing layer, described by the free surface and the interface with shear-free material, for a range of rotation rates, from GDR MiDi (2004); (c) Experimental velocity profiles for a range of rotation rates in the centre of the granular region, as indicated by the diagram, from GDR MiDi (2004).

.

Somewhat obviously, rotating drum flows are not unidirectional. The flow exhibits both velocity variations and components in two-dimensions, even without considering the effects of any lateral confining walls. Furthermore, one must consider that the surface is a free boundary. Far from the flat free surface suggested by the diagram in Figure 2.7(a), the free surface evolves towards the S-shape shown in Figure 2.7(b). Any solution for the flow must also predict the shape of this surface. A full solution of the problem will require a sophisticated numerical calculation like the one that will be developed in §3.2. We shall, therefore, not discuss the predictions of the $\mu(I)$ law for this flow, but mention a few experimental observations for completeness.

Despite the complexity of the rotating drum problem, we note that it bears a number of similarities with the flow on a deep pile discussed in §2.1.5. The typical velocity profiles observed in the centre of the drum (Figure 2.7c) are qualitatively similar to those for the deep pile (Figure 2.6b), as long as one remains near to the free surface. Deeper in the granular region, the velocity profile must match onto that of the rigid body rotation of the drum. This explains the small negative velocity of each profile in Figure 2.7(c). Another common feature is the significance of any side walls in the problem. In a three-dimensional

experiment, friction imparted by any confining side walls will introduce axial variation to the velocity profiles, similar to that in Figure 2.6(c). Finally, we note that a steady flow does not exist for sufficiently small rotation rates. Instead, the material displays a series of roughly periodic avalanches (see Rajchenbach, 2000, for example). This once again indicates the existence of a range, $\theta_{\text{stop}} < \theta < \theta_{\text{start}}$, of surface inclinations in which the flow is metastable. Given our previous comments, we should not expect the $\mu(I)$ rheology to be capable of predicting this behaviour. Nevertheless, based on the success in obtaining velocity profiles for the qualitatively-similar deep pile flow, we should expect the $\mu(I)$ rheology to perform well in predicting steady velocity profiles for the rotating drum. Although we do not perform this calculation, the numerical approach used in Chapter 3 would form a good basis for the simulation of a flow in this configuration.

2.2 Unsteady flows

As we have seen in §2.1, it is often difficult to calculate analytical solutions to problems involving the $\mu(I)$ rheology. In one-dimensional unsteady flows, difficulty arises when the pressure is coupled to the velocity, as in the annular flow of §2.1.2. In higher dimensions, where more than one component of the strain rate tensor is non-zero, the presence of $|\dot{\gamma}|$ in the equations is a significant complication, even in cases for which the pressure is known, and even when the flow is steady. In most cases, it is necessary to resort to numerical methods to find solutions. One can allow an arbitrary initial state to evolve and relax onto a steady state, thus allowing the solution of steady (elliptic) boundary value problems to be found by integrating an unsteady (parabolic) initial value problem. Such an approach, of course, requires that we understand how to solve unsteady problems numerically. This section is dedicated to this task. We restrict attention to unidirectional flows in which the pressure is hydrostatic. Solving other problems will typically involve the need to select a pressure distribution to ensure that the continuity equation is satisfied. The complexities introduced by having to calculate pressure warrant a significant amount of additional effort, and will be discussed at length in Chapter 3.

We shall explain some of the difficulties involved with unsteady calculations in one dimension by way of an example, namely the ‘dragged plate’ problem, in §2.2.1. This problem

is largely analogous to the annular shear flow of §2.1.2, yet it admits an analytical solution, which we shall use to validate our numerical approach. With a reliable numerical scheme for calculating one-dimensional unsteady flow in hand, we present its generalisation to two-dimensional unidirectional flows in §2.2.2, and apply it to shear flows with split boundaries in §§2.2.3-2.2.4. Unfortunately, the generalised numerical scheme of §2.2.2 has limited practical use. An alternative approach, based upon a slight modification to the $\mu(I)$ rheology, will be introduced in §2.2.5, and subsequently applied to flow in an inclined channel in §2.2.6.

2.2.1 The 'dragged plate'

The so-called 'dragged plate' problem offers a good example with which to underpin our discussion of unsteady flows. A deep bed of granular material lies beneath a rough, horizontal plate, which exerts a normal stress, $p_0 = \rho g z_0$, and shear stress, τ_0, on the granular medium. One might like to think of z_0 as being the depth of a static bed of grains necessary to provide the pressure, p_0, by its weight alone. This configuration is illustrated in Figure 2.8(a). One should note that this problem bears a significant mathematical similarity with the annular shear of §2.1.2. In both cases, the balance of shear stresses across a region where the pressure increases leads to the appearance of a yield surface, beyond which the material is static. In the case of the dragged plate, the pressure is hydrostatic,

$$p = \rho g(z_0 - z), \tag{2.29}$$

and the decoupled velocity profile may be found analytically, unlike in the case of annular shear.

Steady solution

We begin by calculating the steady velocity profile for the granular flow beneath the dragged plate. In order for a non-trivial steady state to exist, it must be possible to balance the applied stress, τ_0, with internal friction, thus

$$\mu_1 \rho g z_0 < |\tau_0| < \lambda \mu_1 \rho g z_0, \tag{2.30}$$

or, equivalently, the parameter $\theta = \tau_0 / \mu_1 \rho g z_0$ lies in the interval $1 < \theta < \lambda$. If τ_0 is too small, the stress exerted on the material is below the local yield stress, $\mu_1 p_0$, hence the entire

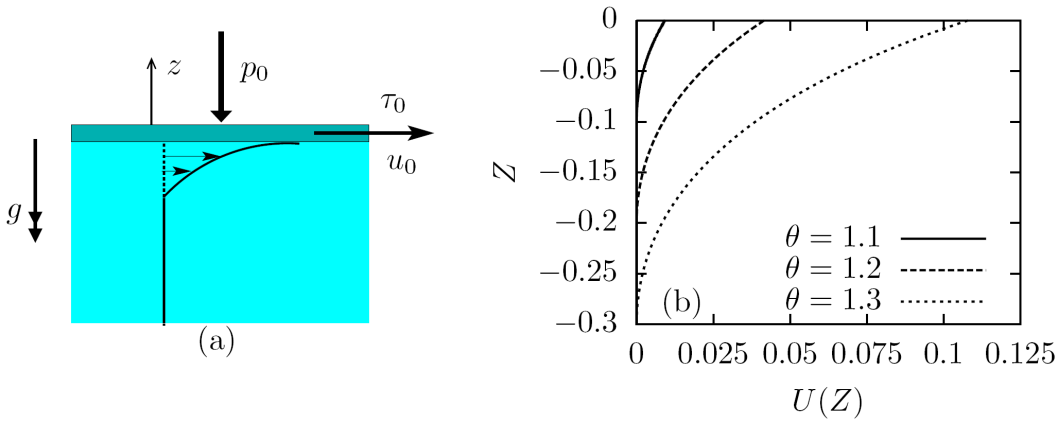

Figure 2.8: The dragged plate problem: (a) Diagram and notation; (b) Steady-state velocity profiles calculated by analytical integration of (2.35) for $\lambda = 1.6$, $\theta = 1.1, 1.2, 1.3$.

granular bed will not move. If τ_0 is too large, then it is impossible (under the $\mu(I)$ model) to generate a large enough internal stress to balance the driving stress, so the plate must accelerate indefinitely. Assuming that a steady state does indeed exist, then balancing the shear stress throughout the material results in the equation

$$\mu(I)p(z) = \tau_0 \quad \Rightarrow \quad \frac{\mathrm{d}u}{\mathrm{d}z} = \frac{I_0\sqrt{g(z_0 - z)}}{d}\left(\frac{\theta - z_0 + z}{\lambda(z_0 - z) - \theta}\right). \tag{2.31}$$

It is now convenient to nondimensionalise the problem, introducing the dimensionless variables

$$z = z_0 Z, \qquad t = \frac{d}{I_0\sqrt{g z_0}}T, \qquad \text{and} \qquad u(z) = \frac{I_0 z_0\sqrt{g z_0}}{d}U(Z). \tag{2.32}$$

We include a scaling for time that will be used when we come to discuss unsteady flow later in this section. In addition, we introduce the (dimensionless) heights

$$Z_1 = 1 - \theta \qquad \text{and} \qquad Z_2 = \left(1 - \frac{\theta}{\lambda}\right). \tag{2.33}$$

At $Z = Z_1$, (2.31) indicates that the velocity gradient drops to zero. In addition, this is the point where the internal friction, $\mu_1 p(z_0 Z_1)$, exactly balances the applied stress, τ_0 – i.e. the yield point. Material in $Z < Z_1$ must be sub-yield, and the no-slip condition applied here gives the boundary condition

$$U(Z_1) = 0. \tag{2.34}$$

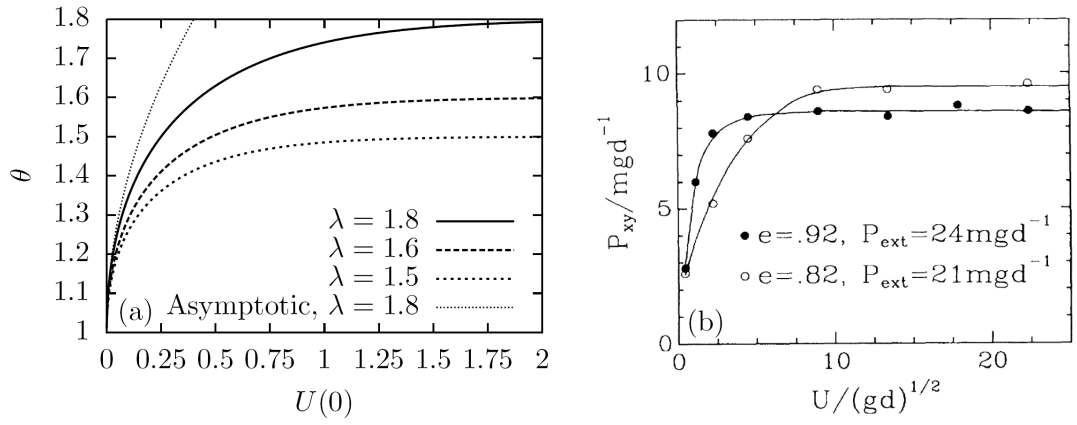

Figure 2.9: Relationship between stress and velocity at $z = 0$ for the dragged plate problem. (a) Sample relationships obtained by direct integration of (2.35) with $\lambda = 1.5, 1.6$ and 1.8, and the asymptotic result (2.38) for $\lambda = 1.8$; (b) Stress-velocity plot from the particle dynamics simulations of Thompson & Grest (1991).

Furthermore, we note that the constraint on the stress for steady flows ($1 < \theta < \lambda$) implies that $0 < |Z_1| < \lambda - 1$. Typical values of λ found in experiments are in the range $1 < \lambda < 2$, so the thickness of the shearing layer will usually be smaller than the equivalent depth of grains needed to supply the pressure p_0. The height $Z_2(> 0)$ is merely a useful notational convenience, but also corresponds to a depth of grains smaller than that required to maintain the confining pressure. In terms of the dimensionless variables, the ordinary differential equation to be solved becomes

$$\frac{\mathrm{d}U}{\mathrm{d}Z} = \frac{1}{\lambda}\left(\frac{Z - Z_1}{Z_2 - Z}\right)\sqrt{1 - Z}, \qquad U(Z_1) = 0. \tag{2.35}$$

Equation (2.35) may be integrated directly, although the form of the resulting velocity profile is sufficiently complicated to prevent it from being particularly illuminating. Instead, we choose to represent the exact solution graphically for a range of θ in Figure 2.8(b).

As we saw on numerous occasions in §2.1, a good experimental outcome one might wish to predict is the dependence of the size of the shearing region on the velocity of a moving wall. The stress or torque on the moving boundary is frequently simply linked to the size of the shearing zone, so one may consider the equivalent problem of predicting the

stress τ_0 on the plate as it moves with velocity u_0. One can obtain this relationship via direct integration of (2.35), giving the results shown in Figure 2.9(a). Alternatively, one can proceed by considering the asymptotic limit in which the imposed shear stress is only slightly above yield, that is,

$$\theta = 1 + \delta, \quad \text{where} \quad \delta \ll 1. \tag{2.36}$$

In this limit, we note that $Z_1 = -\delta$, so it is sensible to introduce a scaled coordinate, $\zeta = Z/\delta$, for which the shearing region is mapped to the interval $-1 < \zeta < 0$. Expanding (2.35) in powers of δ and integrating the resulting expansion yields

$$U(\zeta) = \frac{\delta^2}{2(\lambda - 1)} \left(\zeta^2 + 2\zeta + 1\right) + O(\delta^3). \tag{2.37}$$

Evaluation of the velocity at $\zeta = 0$ gives the desired relationship between the shear layer thickness δ, and the plate velocity

$$U_0 = U(0) = \frac{\delta^2}{2(\lambda - 1)} + O(\delta^3). \tag{2.38}$$

We note that the boundary velocity scales with the square of the shear layer thickness to leading order, as was the case for the annular shear studied in §2.1.2, as well as the chute flow of §2.1.3. This behaviour is generic for flow obeying any $\mu(I)$-like law, provided that the function $\mu(I)$ used is locally linear in I as $I \to 0$. Only the prefactor of this leading order term will change between different geometrical configurations. One can perform a similar asymptotic calculation in the limit of large driving stress, although one must take care not to exceed $\theta = \lambda$, otherwise the internal granular friction will be insufficient to balance the applied stress and a steady solution will not exist. In this limit, the functional form of the relationship between δ and U_0 depends sensitively on the way in which $\mu(I)$ saturates as $I \to \infty$. For the I^{-1} approach to $\mu(I) \to \lambda\mu_1$ studied here, one can show that the plate velocity diverges logarithmically as $\theta \to \lambda$, but one can obtain algebraic behaviour if the limiting form of $\mu(I)$ is appropriately modified. Given the somewhat arbitrary nature of the functional form chosen for $\mu(I)$, we shall not pursue the limit of large I any further here.

Comparing these predictions of the $\mu(I)$ rheology to experimental data is difficult, because a truly one-dimensional flow cannot be realised experimentally. However, Tsai & Gollub (2005) carried out a related experiment in which they sheared the top surface of an

annular channel. In order to avoid any centrifugal effects introduced by the circular geometry, they were restricted to very small plate velocities, giving $U_0 \approx 10^{-8} - 10^{-6}$. Their key observation was that the thickness of the shearing region was independent of the plate velocity. This feature is present in the predictions shown in Figure 2.9(a), but only for $U_0 \gtrsim 1$. For the range of plate velocities examined experimentally, we should expect the asymptotic result (2.38) to be more appropriate. As was the case with the annular shear of §2.1.2, the $\mu(I)$ rheological law struggles to predict the behaviour of slow, quasi-static flows.

Although an experimental realisation of this one-dimensional flow is difficult to construct, the simple geometry makes the dragged plate an ideal candidate for investigation via particle dynamics simulations. Such simulations were carried out by Thompson & Grest (1991), and later by Volfson *et al.* (2003). Interestingly, the relationship between velocity and shear stress obtained by Thompson & Grest (1991) (Figure 2.9b) mirrors that predicted by the $\mu(I)$ rheology. Specifically, they observed that the shear stress, hence the shear layer thickness, does indeed decrease toward zero as the plate velocity tends to zero. Although our inability to calculate the parameters μ_1, I_0 and λ from the microscopic material parameters prevents a quantitative comparison, the qualitative agreement is remarkable.

These two seemingly contradictory observations leave us in doubt as to the validity, or otherwise, of using the $\mu(I)$ law to model such confined shearing flows. They can, however, be reconciled by reconsidering the value of the parameter I_0. Recall that, in this analysis, I_0 features in the velocity scale. For flow on an inclined plane, in which the $\mu(I)$ parameters are typically measured, one tends to find values of I_0 of order unity. Using such a value with the experimental parameters of Tsai & Gollub (2005), one arrives at dimensionless plate velocities of $U_0 \approx 10^{-8} - 10^{-6}$ as previously mentioned. However, a much smaller value of I_0 could give $U_0 = \mathrm{ord}(1)$, placing the experiments in the regime in which the shear layer thickness is essentially invariant with respect to changes in plate velocity, recovering the experimental results of Tsai & Gollub (2005). The same may be said of the annular shear results discussed in §2.1.2, though we have no supporting evidence from particle dynamics in that case.

Placing discussion of the validity of the model to one side, we shall now return to the main theme of §2.2, namely how to compute unsteady flows. Having constructed a steady-state solution for the dragged plate problem analytically, we shall use it to test a numerical

approach for solving unsteady problems.

Discussion of unsteady numerical calculation

In order to investigate the effect of time-dependent forcing on the dragged plate problem, one might naively hope to solve the simple partial differential equation

$$\begin{aligned} \frac{\partial U}{\partial T} =& \frac{\partial}{\partial Z}\left[\mu(I)(1-Z)\operatorname{sgn}\left(\frac{\partial U}{\partial Z}\right)\right], \\ U(Z,0) =& 0, \\ \mu(I)|_{Z=0} =& \mu_1\theta(T), \end{aligned} \tag{2.39}$$

numerically for Z in some finite interval. The most straightforward and obvious approach to take is to use a finite difference scheme to discretise (2.39), and solve using any appropriate timestepping routine. However, a test of this simple approach with $\theta(T) = \text{const.}$ (henceforth referred to as the *start-up problem*), implemented using central differencing and an explicit forward-Euler timestepping routine, does not predict that the system will evolve towards the correct steady state velocity profile calculated previously. The discrepancy is shown in Figure 2.10(a). Though the computed steady profile resembles the exact solution near the moving plate, the solution degrades significantly in depth. In particular, material that should be sub-yield ($Z < -0.2$) and hence stationary has a non-zero velocity. In fact, the flowing region occupies the entire computational domain, however large or small we choose it to be, rather than the fixed interval $Z_1 < Z < 0$ obtained analytically.

The reason for the poor performance of this simple scheme is that we have neglected to include any static frictional stresses that apply when the material is not shearing. Indeed, the differential equation (2.39) is only valid when $\partial_Z U \neq 0$. Where no shear is present, one must instead use the equation

$$\frac{\partial U}{\partial T} = \frac{\partial \tau}{\partial Z}, \quad \text{where} \quad \tau \leq \mu_1(1-Z), \tag{2.40}$$

reflecting the fact that static frictional forces will act to balance any applied stress, provided that it does not exceed the local (dimensionless) yield stress $\mu_1(1-Z)$. The previous, naive approach effectively just sets $\tau = 0$ when $\partial_Z U = 0$, and this leads to the spurious velocities that arise where the material ought to be static. A correct numerical method ought to

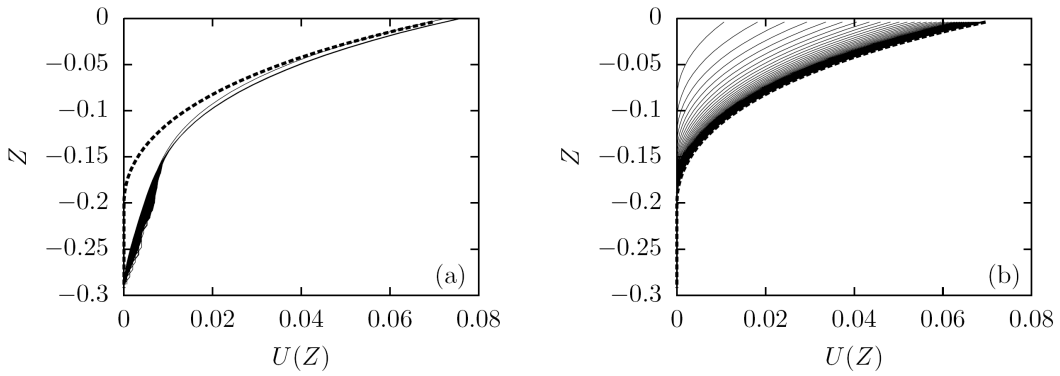

Figure 2.10: Approach to steady state in the dragged plate start-up problem, with $\lambda = 1.4$, $\theta = 1.2$. (a) An incorrect solution, computed using a naive method; (b) The correct solution, obtained after implementing static friction. In each panel, solid lines indicate the numerical solution at every 20,000 timesteps (approximately 0.46 dimensionless time units), whereas the dashed lines show the exact steady solution.

calculate the sub-yield stresses correctly, yet the inequality in (2.40) makes it impossible to calculate the stress locally in a shear-free region. Instead, it is necessary to first find the extent of the shear-free region, and accelerate the entire region as a single entity.

An alternative description of the computational domain may be sought in the form of a stack of rough blocks that slide according to Coulomb friction, but with sliding friction coefficient equal to $\mu(I)$ between any two blocks. If two adjacent blocks are in relative motion, the stress between them will be $\pm\mu(I)p$, where p is the normal stress at their interface. If there is no relative motion, and the driving stress applied to the upper block from above is less than $\mu_1 p$, then the contact between the two blocks is static, and the frictional stress between them is chosen to balance the applied stress. Otherwise, the driving stress overcomes the yield stress, and causes the blocks to begin to slide past one another, with sliding stress $\mu_1 p$. This procedure for calculating static friction suffices for the start-up problem, in which the only shear-free region arises beneath the flowing zone, and remains at rest. In this case, we need only track the base of the flowing zone, $Z = Z_1(T)$, and prohibit any motion beneath this point. This approach recreates the correct steady velocity profile to a high degree of accuracy, as shown by Figure 2.10(b). However, a slightly different approach

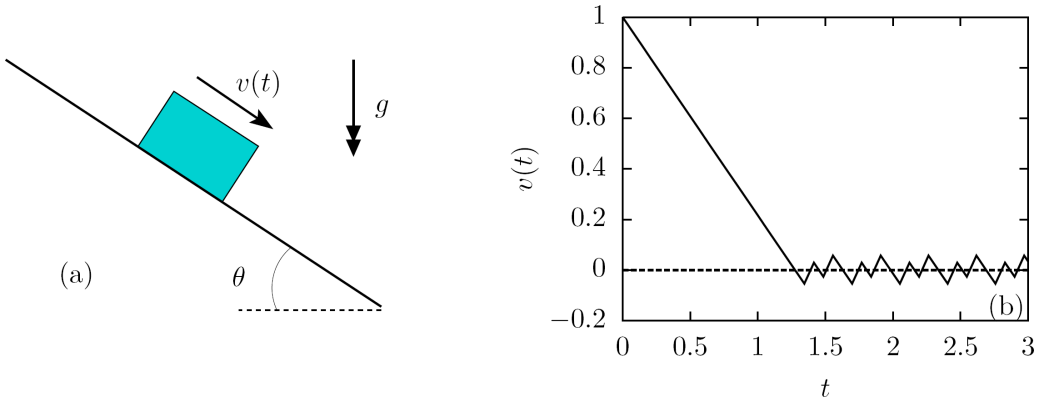

Figure 2.11: An illustration of the problem of computing deceleration and stopping. (a) A block slides down a rough plane inclined at a lower angle than the friction angle between block and plane. (b) A numerical computation of the evolution of the block velocity (in arbitrary units), showing the overshooting described in the main text.

is needed for decelerating flows, in which one can observe the formation of a moving, shear-free region (or 'plug') after the removal of the driving stress provided by the dragged plate. In this case, we follow the same procedure outlined above to determine whether or not each computational block should experience either static or sliding frictional stresses from each of its neighbours. With this task achieved, one can then identify a plug as a group of blocks connected only by static interfaces. This plug then accelerates as a whole in response to the sliding stresses applied from above and below. In this case, there is no need to explicitly calculate the shear stresses within the plug, only those at its (shearing) boundaries, which are found by applying the $\mu(I)$ law as usual.

The inclusion of static friction, as described above, is sufficient to compute time-dependent solutions of the dragged plate problem involving only acceleration. Where deceleration is involved, a further complication arises. Returning to our sliding block analogy, we illustrate the problem via the simple example of a single block sliding down a rough inclined plane with constant friction coefficients. μ_{sliding} and μ_{static}, as shown in Figure 2.11. The equation

of motion of the block in the downhill direction will be

$$\frac{\mathrm{d}v}{\mathrm{d}t} = \begin{cases} g(\sin\theta - \mu_{\text{sliding}}\cos\theta\,\mathrm{sgn}(v)) & \text{if } v \neq 0, \\ g(\sin\theta - \mu_{\text{static}}\cos\theta) & \text{if } v = 0 \text{ and } \mu_{\text{static}} < \tan\theta, \\ 0 & \text{if } v = 0 \text{ and } \mu_{\text{static}} > \tan\theta. \end{cases} \tag{2.41}$$

Suppose that the block is started impulsively with velocity $v = v_0$, and that $\mu_{\text{static}} > \tan\theta$. Under sliding friction, the block velocity should decrease linearly to rest, at which point static friction holds it in place. However, a naive numerical integration of this problem results in the spurious evolution shown in Figure 2.11(b), wherein the block velocity overshoots zero, and proceeds to oscillate about it. This results from the discretisation of time, which makes it very unlikely for the numerical block velocity to exactly reach zero at the end of a timestep. The sign of the sliding frictional stress suddenly reverses, and creates a momentary large downhill acceleration, which gives the block a positive velocity once more, and allowing the process to start again.

In the context of our numerical solution to the unsteady dragged plate problem, this overshooting can come into effect between any adjacent pair of blocks. In practise, it means that a static interface between adjacent blocks, once shearing, will never become static again. This prevents the material from coming precisely to rest, and removes the possibility that any internal plugs can form. We overcome this problem of overshooting by insisting that the shear between any two adjacent blocks may not change sign in a single timestep. Instead, the velocity of each block is set to the same value until the following timestep, when they may shear once again if the applied shear stresses are sufficient to exceed the local yield stress.

Unsteady results

Having remedied the two problems outlined in the previous section, we have constructed a code capable of accurately integrating the initial value problem (2.39) relevant to the dragged plate configuration. In this section, we discuss the behaviour of the system under two impulsive changes in forcing: a sudden stop and a sudden change of direction.

Figure 2.12 illustrates the key features of a decelerating flow that occurs in response to the sudden removal of the driving force. In other words, we apply the dimensionless driving

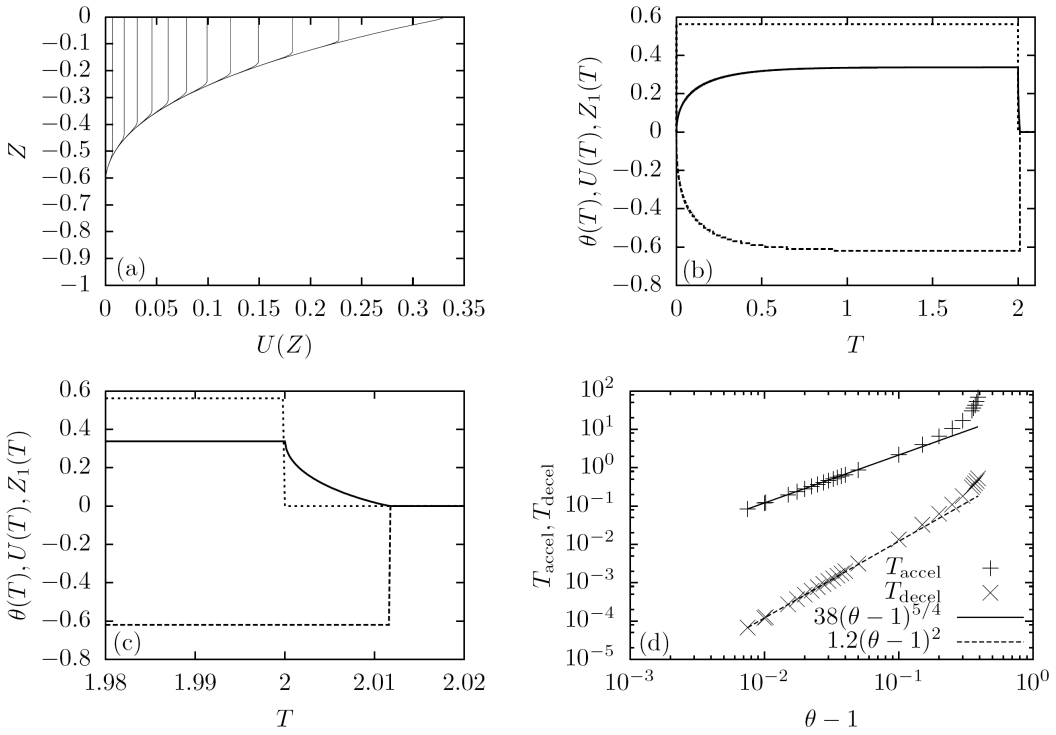

Figure 2.12: Response of the dragged plate to a sudden removal of the driving stress. (a) Velocity profile at regular time intervals after removal of the stress; (b) Time evolution from rest, to steady state, and returning to rest of the plate velocity (solid line) and shear layer depth (dashed line) in response to the driving stress (shown qualitatively as a dotted line); (c) A magnification of (b) over the decelerating period; (d) Log-log plot of the time taken to accelerate to steady state, and to decelerate to rest, with approximate power-law fits.

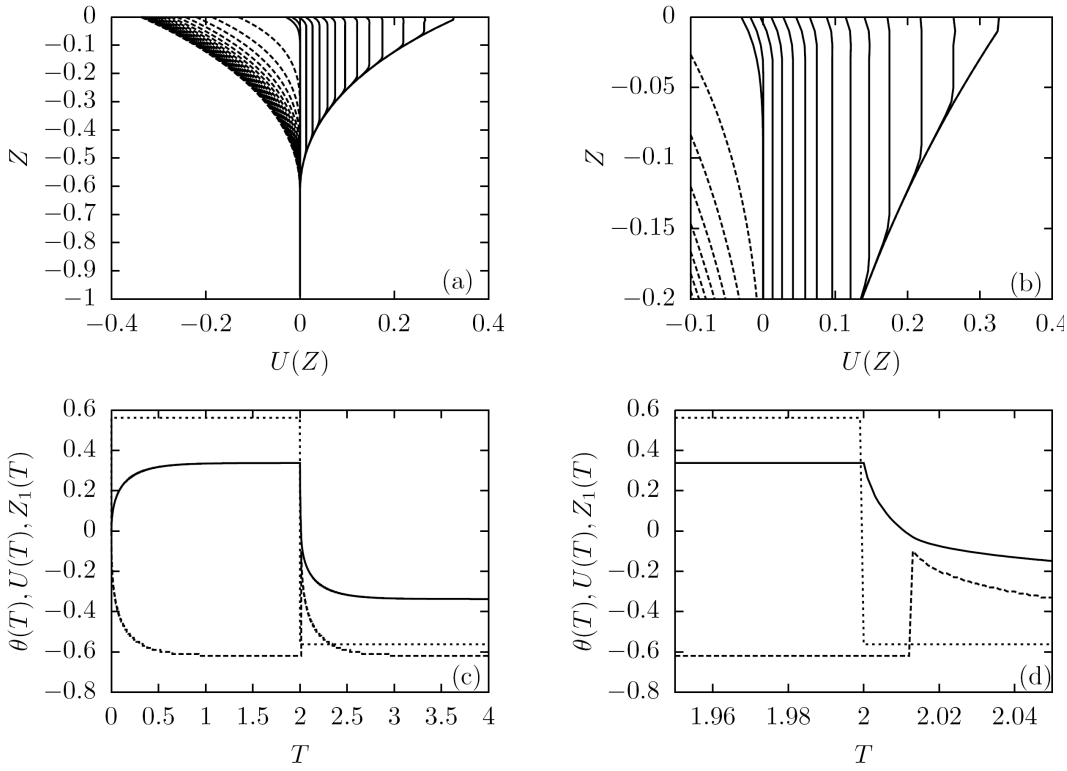

Figure 2.13: Response of the dragged plate to a sudden reversal of the driving stress. (a) Velocity profiles showing deceleration (solid lines) and subsequent acceleration (dashed lines). Decelerating profiles represent timesteps 40 times shorter than accelerating profiles. (b) A magnification of (a), making clear the erosion of the decelerating plug from above. (c) Time from rest of the plate velocity (solid line) and shear layer depth (dashed line) in response to the driving stress (shown qualitatively as a dotted line). (d) A magnification of (c) over the period in which the flow decelerates and reverses.

stress

$$\theta(T) = \begin{cases} \theta_0 & \text{for } 0 < T < T_1, \\ 0 & \text{for } T > T_1, \end{cases} \tag{2.42}$$

where T_1 is chosen to be sufficiently large to ensure that the material reaches steady state before the forcing is removed. The velocity profiles of in Figure 2.12(a) show that a plug region forms at the top of the flowing region, and grows downwards while decelerating. This behaviour is again best understood by using the analogy of a stack of sliding blocks. After the driving stress is removed from the top block, friction causes it to come to rest relative to its neighbour. The newly-formed two-block 'plug', which also lacks a driving stress, must then come to rest relative to the third block, and so on. Eventually, the plug occupies the entire depth of the flowing layer, and the material all comes to rest simultaneously. Figure 2.12(b) shows the evolution of the plate velocity and flow depth in response to the changing force. After a gradual approach to steady state, the material comes to a swift halt after the driving stress is removed. Indeed, the time taken for the material to stop is around 100 times shorter than the time taken to reach steady state originally. We can understand this behaviour in terms of the direction of the frictional force, which always acts against any shear motion. When accelerating the material, friction opposes the motion, and a slow relaxation results. When decelerating, however, the dissipative frictional stresses expedite the deceleration, and the material comes to rest quickly. Further insight into the timescales involved can be gleaned from Figure 2.12(d). For quasi-static flows, where $\theta - 1 \ll 1$, we recall the result (2.38), which tells us that the steady plate velocity scales with $(\theta - 1)^2$. When decelerating, the flow comes to rest in a time $T_{\text{decel}} \propto (\theta - 1)^2$, which is consistent with deceleration from the steady velocity under an ord(1) frictional force. When accelerating, one might expect the time taken to reach steady state, T_{accel} to scale with $(\theta - 1)$, as the net force on the plate will be ord$(\theta - 1)$. However, the numerical data suggests that $T_{\text{accel}} \propto (\theta - 1)^{5/4}$. The discrepancy here is due to the need to accelerate not only the plate, but also the granular material beneath it. Furthermore, the friction coefficient increases with the granular shear, so the net force on the plate gradually decreases in the approach to steady state. It is conceivable that the overall effect of these factors is to produce the $(\theta - 1)^{5/4}$ timescale observed for acceleration in the numerical results.

We now turn attention to the response of the dragged plate to an impulsive reversal in the driving shear stress,

$$\theta(T) = \begin{cases} \theta_0 & \text{for } 0 < T < T_1, \\ -\theta_0 & \text{for } T > T_1, \end{cases} \tag{2.43}$$

as illustrated by Figure 2.13. The key features of the evolution of the velocity profile (Figure 2.13a) can mostly be predicted based on previous observations. After the change in direction of the driving force, a rapid deceleration occurs, almost bringing the flow to rest before beginning a slow approach to the steady velocity profile with opposite sign. Once again, a plug region grows downwards from the plate and decelerates to rest. However, the novel feature in this case is that we observe the plug to be eroded from above due to the action of the reversed driving shear stress. This reversal is made more clear by a magnifying a section near $Z = 0$, as shown by Figure 2.13(b). Examining the evolution of plate velocity and flow depth, we find similar behaviour to the decelerating case discussed above. Deceleration occurs over a much shorter timescale than acceleration in either direction, and the formation of the plug causes much of the flow to come to rest simultaneously. Note, however, that the flow depth does not vanish at any point - a narrow band of reversed shear near the rough surface of the plate exists even when the flowing region is at its thinnest.

2.2.2 Calculating yield surfaces in two dimensions

In the previous section, we developed, tested and used a numerical code to solve problems involving the $\mu(I)$ rheology in one dimension. We now attempt to generalise the block-based scheme used in order to compute unsteady flows in two dimensions.

The sliding block analogy used in §2.2.1 is still useful, though in two dimensions we must think about blocks sliding in a two-dimensional array, allowing for a block to experience frictional stresses from its horizontal neighbours, as well as its vertical neighbours. When blocks are in motion relative to one another, it is a straightforward matter to calculate the sliding frictional stresses furnished by the $\mu(I)$ rheological law. However, complications arise when we consider sub-yield regions of a flow.

In one dimension, a simple force balance can uniquely determine the stress in any sub-yield region, but the same approach leads to a non-unique solution in higher dimensions.

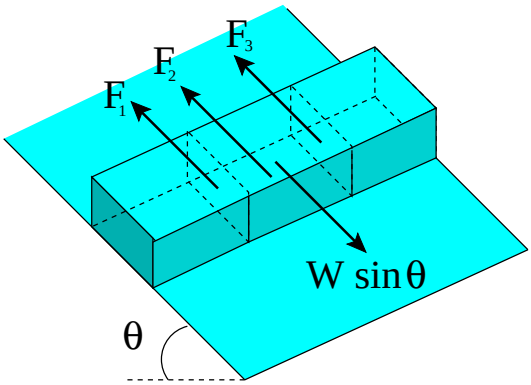

Figure 2.14: Force diagram for three blocks on a rough inclined plane. Used in the main text to explain the non-uniqueness of the static friction problem. Only forces acting downslope are pictured.

This is best explained by way of a simple example. Suppose that three blocks are at rest on an inclined plane, as arranged in Figure 2.14. The central block can experience three frictional forces holding it in place: one from its contact with the base, and one from each of the other two blocks. If we assume for simplicity that the normal force on each contact plane is equal to the component of the block's weight perpendicular to the plane, then the block will remain static provided that

$$\tan\theta < 3\mu_{\text{static}}.$$

Furthermore, the frictional forces must balance the downhill component of the weight, thus

$$F_1 + F_2 + F_3 = W\sin\theta.$$

Unfortunately, we have no further constraints on the three forces, so they may not be uniquely determined. Even after appealing to symmetry to set $F_1 = F_3$, we are left with one degree of freedom in the equations. It is this indeterminacy that makes handling shear-free regions difficult in greater than one dimension.

Despite the indeterminacy of sub-yield stresses, progress is still possible. After all, it is not necessary to know the exact stress distribution within a shear-free plug in order to calculate its motion. One needs only to compute the net force applied to the boundary of the

plug, and accelerate the entire plug as a single entity. This approach was used in the code described in §2.2.1, but was not strictly necessary there, where the frictional stresses could be calculated exactly. In two dimensions, more book-keeping is necessary to keep track of plug regions, but the principle is sound.

The chief difficulty in computing two-dimensional flows with sub-yield regions is the application of the yield condition itself. In one dimension, a yield surface is simply described by a single position. In two dimensions, the yield surface will typically be curved. Depending on the flow and coordinate system, a function describing the yield surface could be multivalued, or the surface could even undergo topological changes as the flow evolves. Such complications require the application of rather sophisticated numerical techniques that we would prefer to avoid. We shall, therefore, restrict attention to simple flows in which the yield surface may be described as a function of one Cartesian coordinate, $z = h_y(y)$. In order to tackle more complicated flows, we shall turn to a different numerical approach, which will be discussed in §2.2.5.

In order to determine the location of a yield surface on a Cartesian (finite difference) computational grid, we use the following procedure. We identify a boundary between two computational cells as *shearing* if those two cells are in motion relative to one another. For such boundaries, one can apply the stresses determined by the $\mu(I)$ law without modification. All other cell boundaries are referred to as *static*. Any cell featuring only shearing boundaries is itself shearing, whilst any cell featuring only static boundaries is static. All velocities associated with a static cell are fixed at zero[1]. The remaining cells, which have a mixture of static and shearing boundaries, are referred to as *incipient*. At each timestep, we consider each incipient cell, and check to see whether it ought to begin moving relative to its static neighbours. We compare the shear stress applied by each shearing boundary to the yield stress, $\mu_1 p$, on each of the static boundaries. If the applied stress exceeds the yield stress on all of the static boundaries, the cell becomes a shearing cell, with the shear stress on its static boundaries chosen to be equal to $\mu_1 p$ on the boundary opposite the largest applied shear. Otherwise, the applied stresses are insufficient to break the incipient cell away from its static neighbours, so it is treated as static. The new static cell then moves at the same

[1]In the more general case of a moving plug, the 'static' cells within are fixed at the plug velocity. If the boundary between a static plug and a rigid wall is static, the plug moves with the velocity of the wall

velocity as the adjacent static cells.

2.2.3 Results: Split-top shear

As a test of the two-dimensional yield condition described in §2.2.2, we now consider a split-top unidirectional shear flow, which is illustrated by Figure 2.15(a). This configuration is very similar to the dragged plate of §2.2.1, except that the moving plate has only finite width in the cross-flow direction. The remainder of the rough, rigid lid is fixed in place. As with the dragged plate problem in one dimension, we specify the shear stress $\tau_0 = \mu_1 \rho g z_0 \theta$ driving the top plate, and calculate the velocity profile, including the velocity of the sliding plate. We retain the nondimensionalisation used in §2.2.1 to allow for straightforward comparison between the two cases.

A contour plot depicting a typical velocity profile is shown in Figure 2.15(b). In the interest of computational economy and figure clarity, we have only computed and shown the solution for half of the domain ($Y > 0$), and note that the solution will be symmetric about $Y = 0$. From this plot, we deduce that the region of grains under shear is confined to the zone beneath the sliding plate ($-1 < Y < 1$ in this example), with no motion at larger values of Y.

Further details of the velocity profile are highlighted by Figure 2.15(c,d), which show, respectively, the velocity variation along horizontal and vertical slices through the material. Figure 2.15(c) makes it clear that there is little horizontal variation in the centre of the shearing zone, and a boundary layer caused by the frictional stresses exerted by the stationary material at each side of the sliding plate. The size of this boundary layer appears to increase with depth. At $Z = 0$, the material is stuck to the sliding plate, so the velocity is uniform. Deeper in the flow, the hydrostatic pressure is higher, and so the width of the region above yield is smaller. The frictional stresses applied to the sides of the shearing region are therefore larger, and apply to a smaller volume of material, thus causing lateral shearing over a greater proportion of the flow.

By considering the variation of the velocity in the vertical direction, as shown in Figure 2.15(d), one can draw similar conclusions. We observe that the flow depth decreases as we move from the centre to the edge of the sliding plate, although it is largely uniform until $Y \approx 0.5$. The form of the vertical variation is reminiscent of that observed in one dimension,

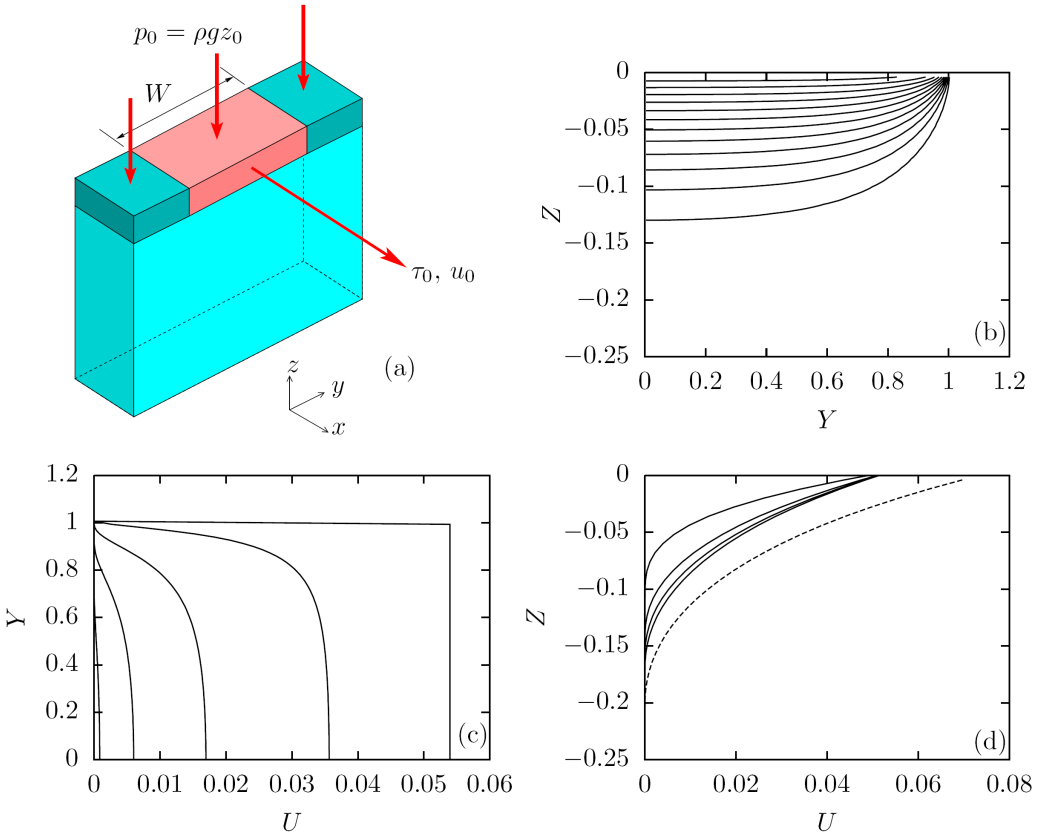

Figure 2.15: Split-top linear shear flow. (a) Diagram and notation; (b) Contour plot of the velocity profile for $(W/z_0, \theta, \lambda) = (2.0, 1.2, 1.4)$ with intervals of 0.004 under the same scalings defined by (2.32); (c) Horizontal variation of velocity at $Z = 0.0, 0.02, 0.06, 0.10, 0.14$ – note that only half of the flow is shown; (d) Vertical variation of velocity at $Y = 0.0, 0.5, 0.7, 0.9$ (solid lines) and the one-dimensional dragged plate result obtained by integrating (2.35) (dashed line).

but there are a few notable differences. For comparison, the one-dimensional case, effectively equivalent to the limiting case $W \to \infty$, is shown as a dashed line in Figure 2.15(d). Because the plate velocity is constant for $0 < Y < 1$, all of the velocity profiles must converge at $Z = 0$. This leads to a slightly more concave profile that the one-dimensional case. However, the velocity of the plate, $U(Y, 0)$, does not match that predicted by the one-dimensional limit – it is significantly smaller. This is no surprise. Consider the balance of the forces on the shearing region as a whole. As in the one-dimensional case, it experiences a driving stress from above and a frictional stress from below, however the two-dimensional shearing region also experiences frictional drag from the stationary material on either side of the sliding plate. This increased drag will inevitably result in a smaller overall steady-state velocity. One would expect the discrepancy between the one- and two-dimensional velocity profiles to decrease as the width of the plate increases, and indeed this is the case. Further numerical experiments with a range of widths (not shown) confirm this limiting behaviour, and serve as a useful validation of our new, two-dimensional code.

2.2.4 Results: Split-bottom shear

In split-bottom shear flow, two L-shaped walls of infinite extent move in opposite directions in response to some applied shear stress (Figure 2.16a). A granular medium occupies a region of depth H between these walls, and is sheared by the motion. While the split-top shear flow serves mainly as a test of our two-dimensional code, split-bottom shear flow offers a genuine opportunity to compare with results available in the literature. Ries *et al.* (2007) carried out contact dynamics simulations for this flow configuration, and reported a number of useful observations, which we shall use to test the $\mu(I)$ rheological law in two dimensions.

Figure 2.16(b) shows the form of the velocity profile in the split-bottom shear apparatus. Near the moving boundaries, the granular material forms a plug, and moves at the same speed as the boundary. In the centre of the channel, the material shears in order to maintain a continuous velocity profile. A more detailed illustration of the velocity profile is provided by Figure 2.16(c), which shows the horizontal variation of velocity at a range of depths in the flow. In the simulations of Ries *et al.* (2007), the authors carefully determined that the behaviour of the free-surface ($z = 0$) velocity profile as $|y| \to \infty$ closely matched an error function fit to the data. The solid line in Figure 2.16(c) is our attempt to approximately fit

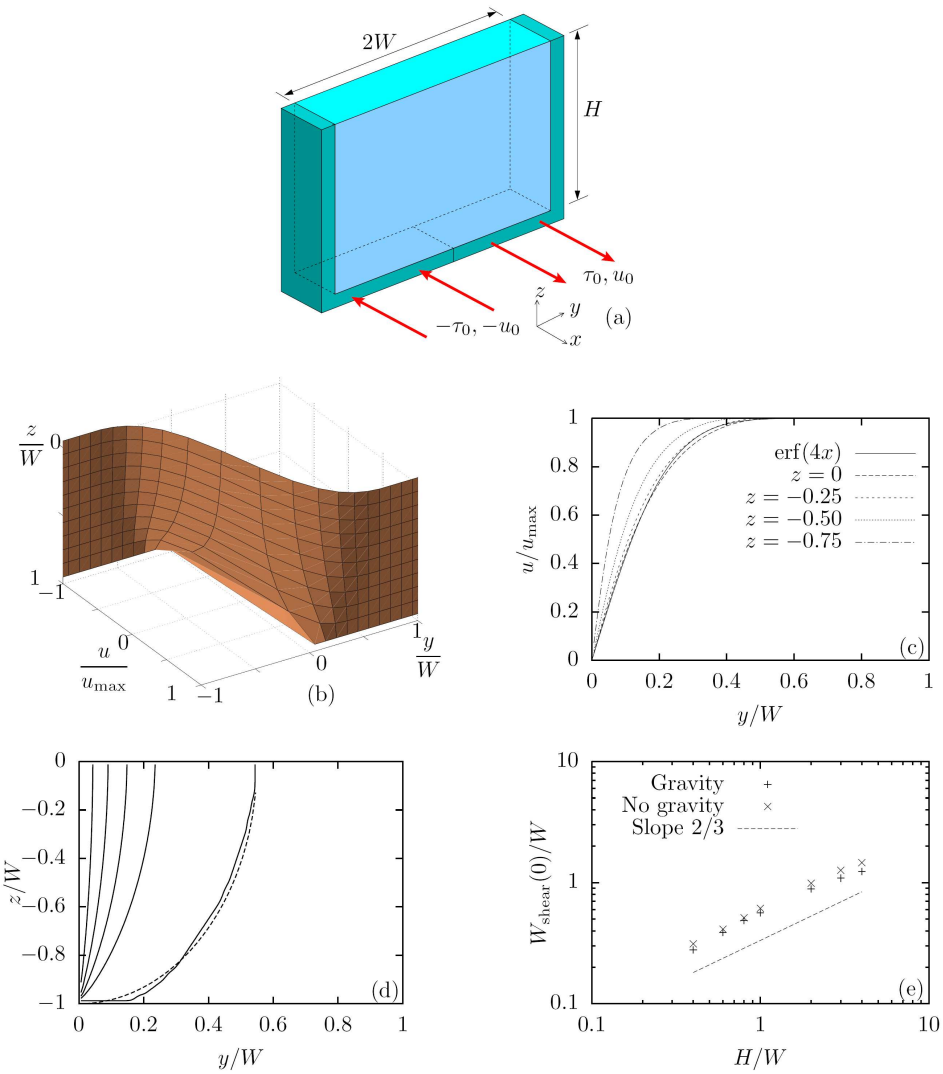

Figure 2.16: Split-bottom shear flow. (a) Diagram and notation; (b) Three-dimensional representation of the velocity profile for $H = W$; (c) Horizontal variation of velocity for $y > 0$ at $z = 0, -0.25, -0.5, -0.75$, with an approximate error function fit to the $z = 0$ profile (solid line); (d) Contour plot of velocity in the right-hand half of the channel, with contours at $u/u_{\max} = 0, 0.2, 0.4, 0.6, 0.8, 1.0$. The dashed line is an approximation of the form (2.44) to the $u = u_{\max}$ contour. (e) Log-log plot of maximum shearing region width as a function of material depth H, for shear with and without gravity. The dashed line has exponent 2/3.

an error function to the numerical solution, but it leaves a lot to be desired - the functional form of the numerical velocity profile does not appear to be in good agreement with that obtained by the simulations.

More promising results may be obtained by considering the structure of the flow within the granular medium. Figure 2.16(d) shows a contour plot of the velocity in the right-hand half of the channel. The outermost solid curve represents the $u = u_{\max}$ contour – that is, the boundary of the shearing region. Our numerical solutions agree well with the qualitative observation that the shearing zone widens as we move from the base to the surface of the granular region. Furthermore, the shape of the region is in good agreement with that observed by Ries *et al.* (2007), who found the shape $W_{\text{shear}}(z)$ of the shearing region to be given by

$$W_{\text{shear}}(z) = W_{\text{shear}}(0)\sqrt{1 - \left(1 - \frac{z}{H}\right)^2}. \tag{2.44}$$

Whilst this functional form, shown as a dashed line in Figure 2.16(d), is not a perfect fit for our numerical solution, we note that both (2.44) and our numerically-determined yield surface meet the bottom boundary tangentially, and the free surface at right angles. The main region of discrepancy between our solution and the elliptical function (2.44) occurs near the base of the flow, where the flow width varies too rapidly for our discrete scheme to capture accurately.

A striking observation of Ries *et al.* (2007) is that the velocity profile obtained under a vertical gravitational force is essentially the same as one obtained under no gravitational force – i.e. with constant pressure, provided that the mean pressure is the same in each case. The only notable change was that the width of the shearing zone, $W_{\text{shear}}(0)$, was larger by a factor of 1.2 ± 0.1 in the absence of gravity. Figure 2.16(e) shows the variation of $W_{\text{shear}}(0)$ with material depth H obtained from our computations with and without gravity. According to these computations, the $\mu(I)$ rheological law is capable of capturing this difference quantitatively. Furthermore, our results are approximately consistent with a power law relationship of exponent $2/3$, which was observed experimentally by Fenistein *et al.* (2004).

Before leaving our discussion of shearing flows, we should draw attention to another type of split-bottom shear to which the $\mu(I)$ rheology has been applied. Jop (2008) has

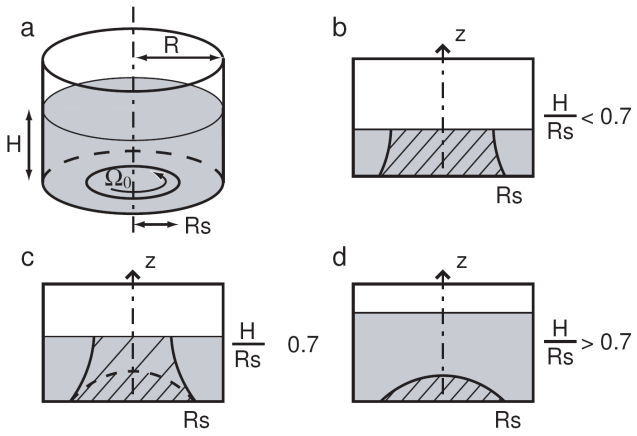

Figure 2.17: Split-bottom Couette shear, copied from Jop (2008). (a) Diagram and notation; (b-d) Schematic of moving region in (b) a shallow cylinder, (c) a moderately deep cylinder, and (d) a deep cylinder.

computed the steady state velocity profile in a split-bottom Couette cell, such as that shown in Figure 2.17(a). Granular material is confined in a cylindrical container, and forced to move by the rotation of a rough disk occupying a portion of the base of the cylinder. In this flow configuration, two distinct types of flow are possible. If the granular layer is shallow (Figure 2.17b), the shearing zone will be similar to that linear split-bottom shear discussed above, although there will be some distortion of the yield surface due to centrifugal effects. For deeper flows (Figure 2.17d), the region of moving material is entirely covered by stationary material, forming a small convex cap that moves with the rotating disk. Jop (2008) has shown that the $\mu(I)$ rheology is capable of producing both types of behaviour, and can recreate the transition between them in a manner that is qualitatively similar to experimental data. However, the width of the radial shear band is not accurately predicted by the $\mu(I)$ model, a finding that resonates with our discussion of shear flows in §2.1.

2.2.5 Regularisation: A numerical simplification

As we have suggested previously, our block-based numerical approach has several limitations, despite its apparent success in calculating unidirectional shear flows. Primarily, our simple

yield condition would need modification in order to account for body forces in the direction of shear. In the examples studied thus far, gravity acts perpendicular to the flow, so the body force only contributes to the determination of the pressure. The modification that must be made to our condition in order to account for general body forces is not obvious, and we have not been able to find a useful approach. Looking at the literature from plasticity theory, we might consider implementing a classical Mohr-Coulomb failure condition (Coulomb, 1776), which predicts the location and orientation of yield surface based upon the principal axes of the stress tensor. However, accounting for the orientation of the yield surface requires significant additional bookkeeping.

Quite apart from the precise form of the yield condition is the difficulty of tracking the evolution of the yield surface. In complicated unsteady flows, the yield surface could vary as much as a free surface obeying the kinematic boundary condition. Topological changes are possible (for example, the appearance of a moving plug during deceleration in §2.2.1), and would require the use of a sophisticated surface-tracking method, which we would rather avoid. Instead, we propose a modification to the $\mu(I)$ constitutive law that will allow us to treat the entire granular material in the same manner, with no need to distinguish between shearing and shear-free zones.

We define a modified deviatoric stress

$$\tilde{\tau}_{ij} = \mu(I)p\frac{\dot{\gamma}_{ij}}{\sqrt{|\dot{\gamma}|^2 + \epsilon^2}}, \tag{2.45}$$

where ϵ is a small positive inverse timescale. This modified stress is illustrated by Figure 2.18(a). For any shearing part of the flow, where $|\dot{\gamma}| \gg \epsilon$, the effect of this change is insignificant. Furthermore, in shear-free regions, the deviatoric stress vanishes, avoiding any dividing-by-zero problems that would arise without the regularisation (2.45). When the rate of strain is very small, i.e. $|\dot{\gamma}| \lesssim \epsilon$, the effective viscosity becomes very large – $O(1/\epsilon)$ – making the material highly dissipative. The idea here is that a static region will not be represented numerically as a true static region, rather as region of very small strain rate and hence very high effective viscosity. This high viscosity quickly eliminates any shear and ensures that 'static' regions remain 'static'.

Before using this regularised expression for the internal stresses, it is important to justify its use. In order to do so, we shall return to the dragged plate problem of §2.2.1. As shown

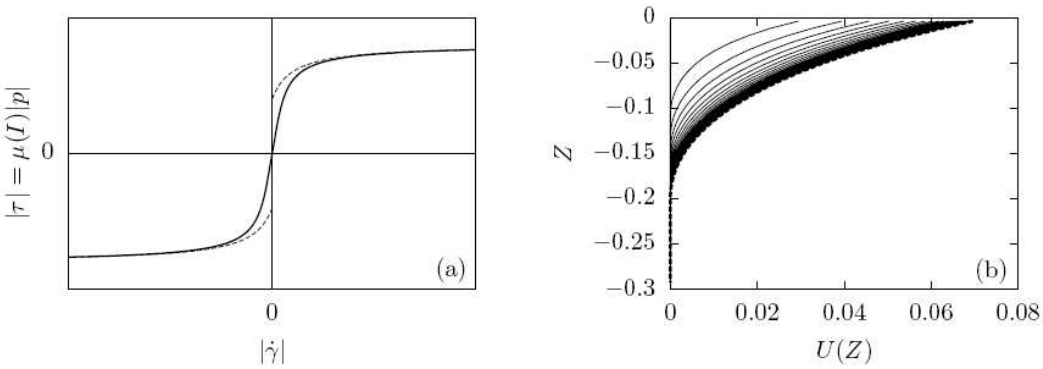

Figure 2.18: Regularisation of the $\mu(I)$ stresses. (a) The modified deviatoric stress (2.45) (solid line), with $\epsilon = 0.5$ to exaggerate the smoothing effect, compared with the unmodified $\mu(I)$ stress (dashed line); (b) Approach to steady state in the dragged plate problem using the regularised $\mu(I)$ rheology.

by Figure 2.18(b), a calculation using the modified rheological law (2.45) evolves toward the correct steady velocity profile, and does so in a very similar way to the careful calculation illustrated by Figure 2.10(b), but does not require the additional effort associated with differentiating static and shearing regions.

One can also test a regularised version of the two-dimensional code. In Figure 2.19, we compare the steady state velocity profile in the split-bottom shear configuration, as computed by the previous block-based method and the new regularised method. The differences are very slight, almost imperceptible to the eye. Nevertheless, the regularised code, which routinely underestimates frictional stresses, leads to a slightly wider shear zone than the block-based code. Furthermore, we reiterate that there is no plug zone in the regularised result – the velocity is strictly less than its maximum value throughout the interior of the flow.

This ease of computation comes at a cost. When using the regularised stresses, the very high effective viscosities that arise necessitate the use of a very small timestep in order to avoid the well-known diffusional instability. Specifically, one must choose the timestep so that

$$\Delta t < A\epsilon\Delta z^2, \tag{2.46}$$

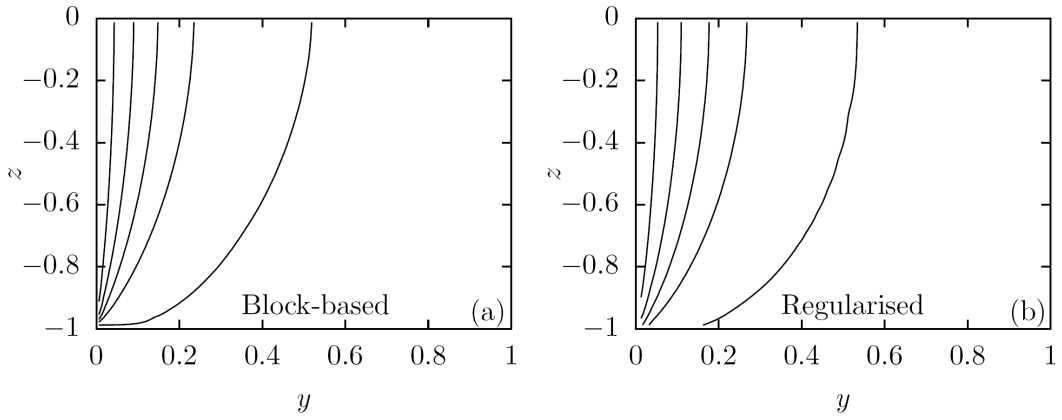

Figure 2.19: Contour plot of velocity in the split-bottom shear of §2.2.4, as computed by (a) the block-based method, with a yield criterion, and (b) the regularised constitutive law (2.45) with $\epsilon = 0.0001$. The contours are drawn at $u/u_{\max} = 0.2, 0.4, 0.6, 0.8$ and 0.999.

where A is a dimensional factor related to the velocity, lengthscale and timescale chosen for the problem. The block-based code discussed in §2.2.1 does not contain such large viscosities, and the timestep may be larger by a factor of $1/\epsilon$ as a result. Based on this argument, we would like to make ϵ as large as possible to save on computational cost, but not so large that it significantly decreases the accuracy of the solution. In practice, we find that $\epsilon = 0.0001$ is a good compromise between these two conflicting concerns.

A further drawback of the regularised rheology becomes apparent when investigating an arresting flow. Unlike the true rheology, which allows for the possibility that the entire granular medium may be sub-yield, and thus not be in motion, a material governed by the modified rheology will always have a large effective viscosity. A slow, creeping flow will frequently result. While not relevant to the flows studied in this chapter, this issue of stopping will play a role in Chapter 3, where we consider a transient dam-break flow.

2.2.6 Results: Flow in an inclined channel

By using the regularised $\mu(I)$ rheological law (2.45), it is now possible to return to our postponed discussion of flow on a deep pile, last seen in §2.1.5. In fact, our numerical approach, which can calculate the velocity profile in a region of rectangular cross-section,

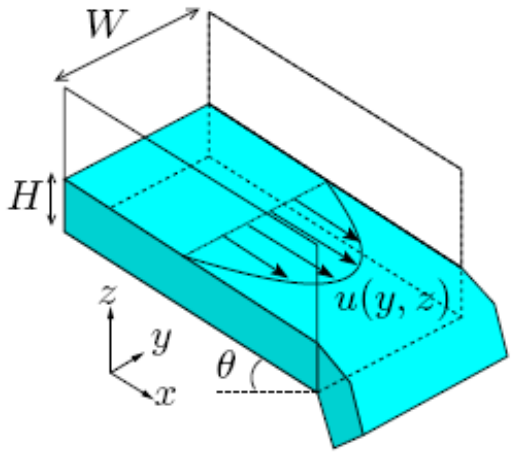

Figure 2.20: Diagram and notation for flow in an inclined channel.

may be applied to the more general problem of flow in an inclined channel, as shown in Figure 2.20. The case of the deep pile is represented by the limiting case in which the depth of the granular material is much larger than the width of the channel, allowing the friction forces from the rough side walls to dominate over the friction exerted at the base of the channel. This limit was studied by Jop *et al.* (2006) as part of their introduction of the $\mu(I)$ rheology, but we improve upon their work by using slightly more careful numerics and by considering the effects of finite depth.

In an experiment, one typically controls the angle of inclination, θ, and width, W, of the channel, as well as the volume flux, Q. As the flow evolves, the depth of granular material, H, will adjust to an appropriate value. However, the problem is more straightforward to solve numerically with fixed geometry, so we scale lengths by W, specify H and θ, and calculate the volume flux from the steady-state velocity profile. Our numerical method uses an explicit, second-order finite difference scheme, with central differencing wherever possible, and second-order extrapolation when applying boundary conditions. This differs from Jop *et al.* (2006)), whose scheme is only first-order accurate at the boundaries. We impose no-slip boundary conditions both at the base and on the vertical side walls of the channel.

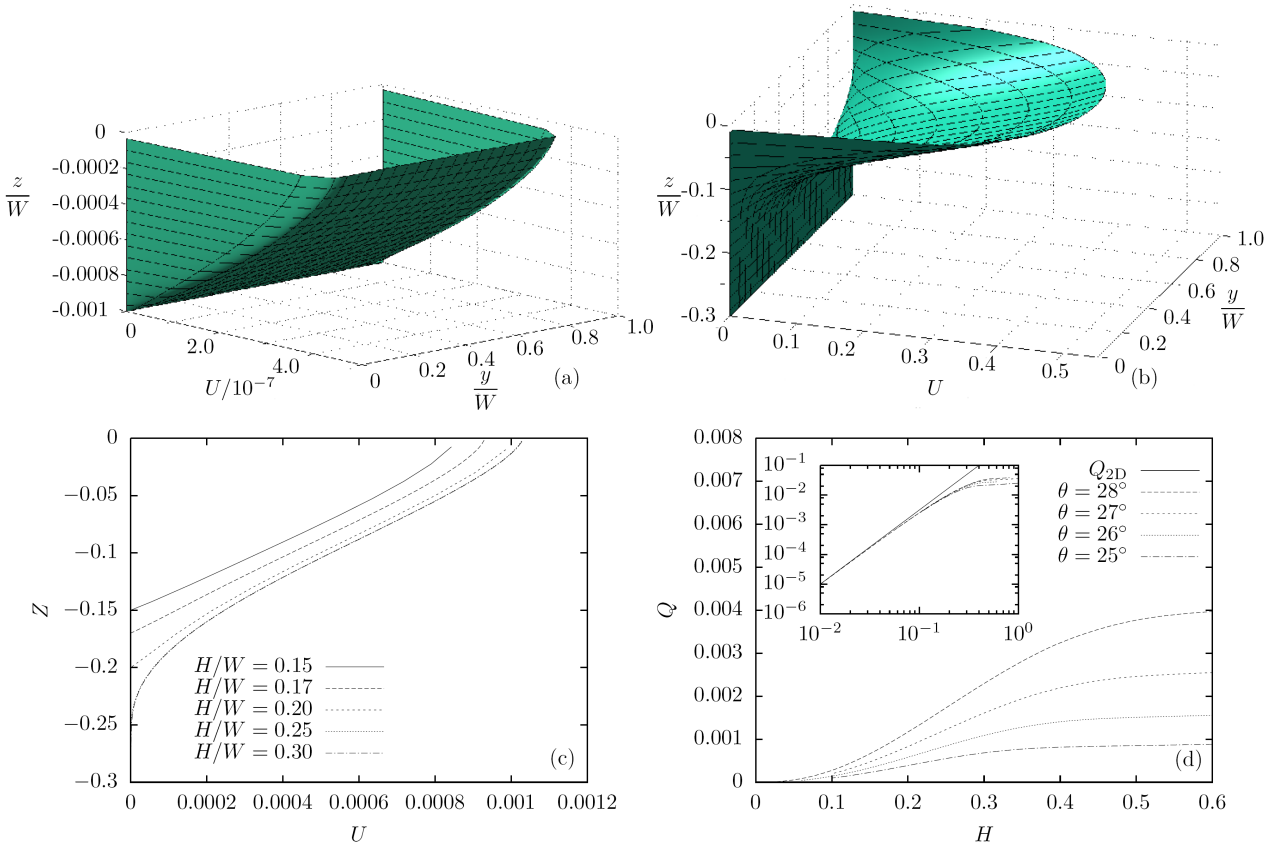

Figure 2.21: Numerical results for flow in an inclined channel: velocity profiles (nondimensionalised by $\sqrt{gW}W/d$) in (a) a shallow channel ($H/W = 0.001$) and (b) a deep channel ($H/W = 0.3$); (c) Velocity profiles at $y = W/2$ for $\theta = 22.6°$ and $H/W = 0.15, 0.17, 0.2, 0.25, 0.3$; (d) Volume flux as a function of flow depth for $\theta = 25°, 26°, 27°, 28°$ – inset: log-log plot of volume flux scaled by $Q_0(\theta)$, with the asymptotic result (2.47) for small H (solid line).

Results from our numerical calculations are shown in Figure 2.21. In a shallow channel, the velocity profile (Figure 2.21a) is largely one-dimensional, with only a thin boundary layer needed to account for the no-slip condition on the side walls. The velocity profile in the central region agrees to $O(\Delta x^2)$ with the Bagnold profile (2.23) for flow on an inclined plane, indicating that our code is performing correctly. In a deep channel (Figure 2.21b), we observe the there is essentially no motion below some finite depth, and the general form of the profile is in good agreement with the calculations and experimental observations of Jop *et al.* (2006).

In a channel whose cross-section has moderate aspect ratio, we can observe some of the characteristics present in both of the extreme cases. As the aspect ratio rises, the width of the cross-stream-invariant region decreases, and the size of the static zones at the base of the channel increases. The gradual change in character of the flow is illustrated by the velocity profiles in the centre of the channel (Figure 2.21c). As the aspect ratio increases, an originally Bagnold-like profile develops an inflection point and an exponential tail at large depths. One could surmise that this captures the creeping flow observed by Komatsu *et al.* (2001), though in truth it is more likely a numerical artefact due to the regularized form of the $\mu(I)$ rheology, which allows for a slow viscous creep at very low rates of strain.

We conclude our discussion of flow in an inclined channel by considering the dependence of the volume flux on the channel depth and inclination angle. For a one-dimensional flow of depth H on a rough inclined plane, the Bagnold profile (2.23) may be integrated to provide the flux per unit width

$$Q_{1\mathrm{D}}(H,\theta) = \frac{2\mu^{-1}(\tan\theta)\sqrt{\phi g\cos\theta}}{5d}H^{5/2} \equiv Q_0(\theta)H^{5/2}. \tag{2.47}$$

In Figure 2.21(d), we have plotted the normalised volume flux, $Q/Q_0(\theta)$, as a function of H for a range of inclination angles. For shallow flows, we see very good agreement with the one-dimensional result (2.47). As the depth of material increases relative to the channel width, the volume flux tends to a finite limit, whose value has a further dependence upon the inclination angle. This agrees qualitatively with the observations of Jop *et al.* (2006). In their experiments, they specified Q, and allowed the free surface of the granular material to find its own inclination angle. The depth, H, plays no role in the deep pile limit, $H/W \to \infty$. In our numerical results, the large H limit produces a one-to-one relationship between the

flux, Q, and inclination angle, θ, which is consistent with the experimental results.

3

Collapse of a granular column

In Chapter 2, we assessed the validity of the $\mu(I)$ rheological law for a wide range of flow configurations. We found that it performs rather well for free surface flows, but struggles when applied to confined shearing flows. As part of our investigation, we developed some simple techniques to facilitate the numerical solution of the governing equations.

In this chapter, we examine the important physical problem of the collapse of a two-dimensional granular column (also referred to as the dam-break problem), which is illustrated by Figure 3.1. A two-dimensional column of granular material of width, W, and height, H, is contained by a vertical wall. At the start of an experiment, the wall is quickly removed, and the granular material spreads horizontally. After a short period of time, the material comes to rest, leaving a deposit with sides inclined at angles less than the angle of repose.

This flow has a good deal of physical relevance. Civil engineers and geophysicists interested in hazard prediction may view it as a laboratory model for several flows involving the slumping of granular material, whether part of a manmade embankment or cutting, or of a naturally-occurring steep slope. Such flows are not restricted to the terrestrial environment – certain formations observed in the Valles-Marineris canyon on Mars are thought to have originated from the collapse of the canyon wall (see Lajeunesse *et al.*, 2006).

We shall discuss the problem of granular column collapse in some detail. A review of the existing experimental and theoretical work is presented in §3.1, where we shall also motivate

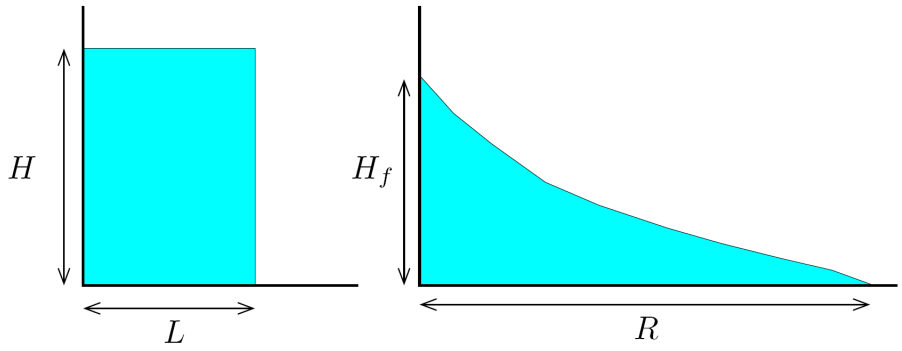

Figure 3.1: Collapse of a granular column: A sketch of the initial column, and the final deposit.

the use of the $\mu(I)$ rheological law for a new theoretical treatment of the problem. In order to solve the appropriate equations, it is necessary to develop a numerical scheme capable of handling the $\mu(I)$ law in a general two-dimensional flow, with a moving free surface. A thorough discussion of this development is given in §3.2, and the predictions generated by the resulting code are discussed in §3.3.

3.1 Previous work

The collapse of a granular column has been the subject of several experimental studies in recent years. Investigations have been carried out in both a quasi-two-dimensional configuration (Balmforth & Kerswell, 2005; Lube *et al.*, 2005; Lajeunesse *et al.*, 2005; Siavoshi & Kudrolli, 2005), and a cylindrical configuration (Lube *et al.*, 2004; Lajeunesse *et al.*, 2004). A range of different granular materials have been used, and the key findings appear to be largely insensitive to the properties of the material used.

We shall discuss the experimental results in much greater detail when we come to compare them with the results of our simulations in §3.3, but some of the key observations are summarised by Figure 3.2. After scaling lengths with the width of the initial column, L, and time with $\sqrt{L/g}$, the flow behaviour appears to depend only upon the initial aspect ratio, $a = H/L$, of the column. The runout distance, R, (the maximum horizontal extent

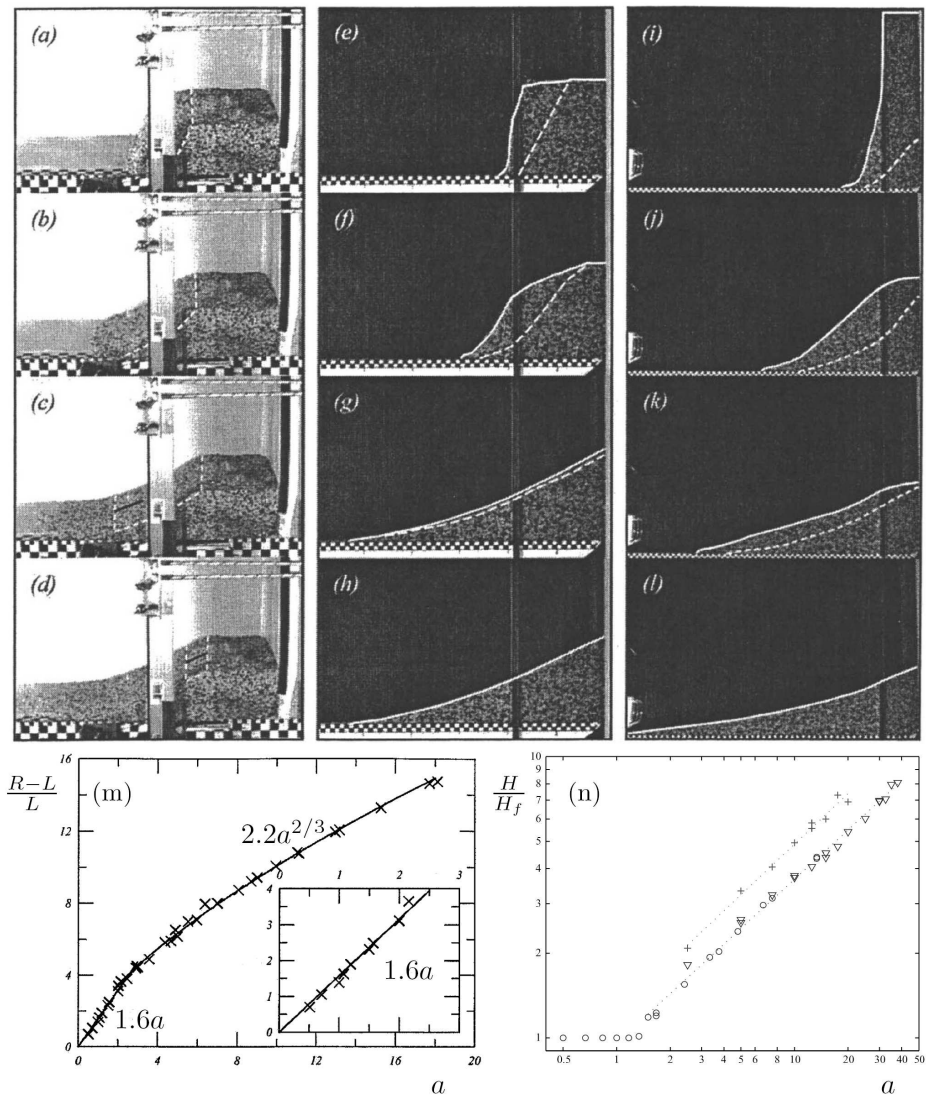

Figure 3.2: Experimental results for the collapse of a two-dimensional column. Panels (a-l) show successive frames of three experiments at aspect ratio $a = 0.5$ (a-d), $a = 1.5$ (e-h), and $a = 7$ (i-l), from Lube *et al.* (2005). Where used, solid lines represent the free surface, and dashed lines highlight the measured interface between static and shearing material. Panel (m) shows the variation of the final runout with aspect ratio from the same paper, with an inset to highlight the linear behavior at small aspect ratio. Panel (n) shows the effect on the final height of varying the aspect ratio, adapted from Balmforth & Kerswell (2005).

to which granular material spreads), an important quantity to understand in the context of hazard prediction, is a prime example of this single parameter dependence, as is shown by Figure 3.2(m). The experiments of Lube *et al.* (2005), which are illustrated by 3.2(a-l), suggest that the runout obeys a power-law

$$R = L\left(1 + \beta a^{\gamma}\right), \tag{3.1}$$

where $\gamma = 1$ for small aspect ratio ($a \lesssim 1.8$), and $\gamma = 2/3$ for large aspect ratio ($a \gtrsim 2.8$), with a smooth transition between the two regimes for moderate aspect ratio. For the collapse of a low aspect ratio column (such as in Figure 3.2a-d), much of the column remains undisturbed, and only the front of the column collapses. If the column has a higher aspect ratio, as in Figure 3.2(i-l), the vast majority of the column falls freely in the initial stages of the collapse, and only a small region at the base of the flow will remain stationary. This effect is perhaps made more clear by Figure 3.2(n), which shows the ratio of the initial to the final height of the column, again as a function of the initial aspect ratio. The height is unchanged for small a, but decreases significantly from its initial value if a is larger. It is thought that this explains the existence of the two different power laws observed experimentally – as the flow is dominated by frictional avalanching (for small a) or the inertia of free fall (for large a).

Several authors have attempted to model column collapse theoretically. Siavoshi & Kudrolli (2005) drew comparisons with the so-called 'BCRE' model (Bouchaud *et al.*, 1994), which accounts for a small layer of grains rolling across the top of a static region, but does not make any account of frictional stresses. Mangeney-Castelnau *et al.* (2005) attempted to model column collapse using a simple shallow water model, assuming little internal shear, but allowing for a frictional drag at the base of the flow, and achieved reasonable agreement with experimental data for small aspect ratio columns ($a < 1$). Balmforth & Kerswell (2005) presented with their experimental results a shallow water model involving a slightly more sophisticated friction law, similar to that proposed initially by Savage & Hutter (1989). This model was expanded upon and applied to axisymmetric collapse by Kerswell (2005). Unfortunately, the theoretical predictions of both Balmforth & Kerswell (2005) and Kerswell (2005) for the final deposit profile, including the runout, are in poor agreement with the experimental data for the moderate aspect ratio columns considered. At fault is the assump-

tion, intrinsic to the shallow water formulation used, that the free surface height is gently varying in the horizontal direction. While this may be true in the late stages of the flow, it certainly is not in the initial stages, when one face of the column is nearly vertical. In the case of Mangeney-Castelnau *et al.* (2005), one can only assume that any vertical velocity variation occurring in small aspect ratio collapses is not significant enough to affect the evolution of a short pile. In an attempt to account for the non-shallow early stage, Larrieu *et al.* (2006) attempted to model the deep portion of the initial flow as a 'rain' of material into a similar shallow-water model. Using this approach, Larrieu *et al.* (2006) obtain the correct power-law behaviour for the runout in collapses of moderately tall columns ($2 < a < 10$), although their model underestimates the length of time necessary for the flow to come to rest.

The approach of Larrieu *et al.* (2006) was extended by Doyle *et al.* (2007), who were guided by the observation of Lube *et al.* (2005) that the collapsing column may be divided into a thin shearing layer atop a static wedge (the interface between which is illustrated by a dashed line in Figure 3.2a-l). By applying the 'raining' model of Larrieu *et al.* (2006) to the shearing layer only, and using an empirical determination of the location of the static zone, Doyle *et al.* (2007) succeeded in predicting the correct power-law dependence for both the final runout time and distance for high aspect ratio columns ($a > 3$). Although a promising result, the model of Doyle *et al.* (2007) requires an empirical determination of both the shape of the static zone and the sedimentation rate describing the 'rain' of granular material into the flow.

Further progress towards understanding granular slumping can be found in the form of particle simulations. Staron & Hinch (2005) carried out a suite of two-dimensional particle dynamics simulations, and largely reproduced the shape of the final deposit observed in experimental studies. Furthermore, their approach makes it simple to probe the interior of the flow and examine the velocity field. They confirmed the suggestion, made by Lube *et al.* (2005), that the flowing grains only occupy a thin region near the surface of the flow, lending credence to the assumption of a static pile made by Doyle *et al.* (2007).

A more sophisticated set of simulations were performed by Lacaze & Kerswell (2009), who used a commercial particle dynamics code to model the collapse of an axisymmetric column. Using a coarse-graining approach (Goldhirsch & Goldenberg, 2002), they calculated

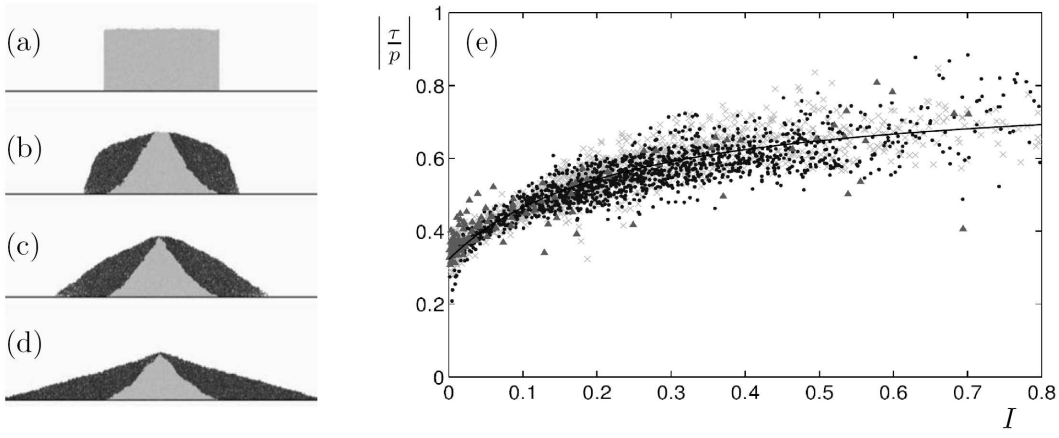

Figure 3.3: Particle simulations of granular column collapse. (a-d) A series of frames from the two-dimensional simulations of Staron & Hinch (2005), showing the static zone. Only particles that move more than a grain diameter from their initial position are coloured black. (b) Coarse-grained approximation to the friction coefficient, $\mu(I) = |\tau|/p$, obtained from a range of axisymmetric simulations by Lacaze & Kerswell (2009), with the form (1.11) shown as a solid line.

an approximation to the components of the stress tensor throughout the granular material. From this data, they calculated the friction coefficient $\mu = |\tau|/|p|$, and plotted its variation with the local inertia number (Figure 3.3b). Interestingly, they found that the relationship between normal and shear stresses agrees quantitatively with the $\mu(I)$ constitutive law (1.14), though there is a large degree of scatter. This observation provides strong motivation for a treatment of the column collapse problem by using the $\mu(I)$ rheological law. It is this treatment to which the remainder of this chapter is devoted.

3.2 Numerical approaches

In §3.1, we briefly discussed the attempts by other authors to model the collapse of a granular column. Each of these attempts was based upon a set of shallow-water equations, which formally require the shape of the flowing domain to vary gradually in the direction of flow. This requirement is clearly not fulfilled in the early stages of flow, and the predictions of

such models are typically inadequate as a result.

We shall lift the key assumption of shallowness and attempt to model column collapse by using a fully two-dimensional version of the $\mu(I)$ rheological law. After scaling lengths with L, velocity with $\sqrt{gL}$, pressure with ρgL, and time with $\sqrt{L/g}$, the equations of conservation of mass and momentum are

$$\frac{\partial \mathbf{U}}{\partial T} + \mathrm{U} \cdot \nabla \mathbf{U} = -\nabla P + \nabla \cdot \left(\frac{\mu(I) P}{\sqrt{|\dot{\gamma}|^2 + \epsilon^2}} \dot{\gamma} \right), \tag{3.2}$$

$$\nabla \cdot \mathbf{U} = 0. \tag{3.3}$$

Note that we use the regularised $\mu(I)$ rheological law described in §2.2.5. To these equations we add boundary conditions corresponding to a no-slip rigid boundary on the base and symmetry at the centre of the pile (located at $X = 0$), thus

$$U(X, 0, T) = 0, \tag{3.4}$$

$$V(X, 0, T) = 0, \tag{3.5}$$

$$U(0, Y, T) = 0, \tag{3.6}$$

$$\frac{\partial U}{\partial X}(0, Y, T) = 0. \tag{3.7}$$

On the free surface (which we describe for now by the curve $Y = \zeta(X, T)$, one must apply the usual kinematic boundary condition, and set the granular pressure to zero, resulting in the conditions

$$\frac{\partial \zeta}{\partial T} + \mathbf{U} \cdot \nabla \zeta = V \quad \text{at } Z = \zeta(X, T), \tag{3.8}$$

$$P = 0 \quad \text{at } Z = \zeta(X, T). \tag{3.9}$$

Note that the condition (3.9) implies that all components of the stress tensor vanish at the free surface.

In the analysis that follows, we shall occasionally refer to simplified versions of the above equations. In particular, when testing our numerical codes, we shall use the *inviscid* form of the equations, which we obtain by setting $\mu_1 = 0$, hence removing the last term in (3.2). The resulting reduction in order of the equations then requires that we omit the no-slip boundary condition (3.4). We shall also make reference to the *constant friction* model, which is produced by setting $\lambda = 1$ (equivalently, $\mu(I) \equiv \mu_1$).

The equations (3.2)-(3.9) must be solved numerically. Unfortunately, the non-standard rheology, coupled with the need to work in a deforming domain, means that we have had to develop our own numerical scheme, as opposed to making simple modifications to off-the-shelf codes. In the remainder of this section, we describe our progress towards the development of a useful scheme.

3.2.1 Lagrangian finite elements

A Lagrangian finite element scheme is arguably the most intuitive method by which to compute the solution to the continuum equations in a deforming domain. One must first make a triangular mesh of the domain, then advect the nodes of this mesh with the local material velocity. The unstructured mesh used should allow the complicated geometry to be handled with ease, whilst the process of advecting nodes updates the position of the free surface without the need for any separate surface-tracking routines. Furthermore, the vanishing of the pressure (hence all components of the stress tensor) at the free surface (3.9) means that no complications are added by the need to accurately compute the free surface stresses based on the interior flow structure.

We have expended a good deal of effort in our attempts to construct a successful Lagrangian finite element scheme, but have ultimately failed. A detailed description of our attempts, ideas, and lessons learned is presented in Appendix A, but we shall summarise the most important points here, in order to put in context some aspects of our successful numerical approach.

The key problem with the Lagrangian finite element approach, and the one that ultimately caused us to abandon this type of scheme in favour of the marker-in-cell method, involved the appearance of spurious pressure modes. In all finite element schemes for fluid flows (involving the primitive variables of velocity and pressure), it is necessary to find the pressure field by solving a linear algebra problem. Depending on the choice of finite elements, this problem may have zero eigenmodes, which manifest as spurious oscillations in the pressure field. By their nature, these eigenmodes have no direct effect on the velocity. However, they have a profound effect on the frictional stresses vital to the $\mu(I)$ rheology. The errors introduced by these spurious eigenmodes eventually accumulate and overwhelm the causal behaviour, invalidating the solution. In order to gain some control over the pres-

sure, we have tried to use a range of different elements, modified the problem solved for pressure in a manner that suppresses the spurious modes, and applied explicit smoothing techniques to the computed pressure. Despite all of these 'improvements', we were unable to construct a scheme that handled the pressure in a satisfactory manner, especially in the later stages of the flow when the mesh is highly distorted.

Deformation of the mesh poses another significant difficulty when attempting to apply the Lagrangian finite element method to the column collapse problem. As observed by Lube *et al.* (2005), the flow is often confined to a thin region just below the surface, while the granular material beneath remains stationary. This leads to large deformations of the mesh elements, which can cause a decline in the quality of the solution. Furthermore, it is possible that the quality of the mesh may be so poor that the iterative method used to calculate the pressure fails to converge, halting the simulation. There are many steps that one may take to improve the quality of the mesh, for example by the introduction or removal of nodes, or retriangulation of the existing set of nodes. However, it is difficult to perform such mesh refinements algorithmically, and the algorithms used may result in very small triangles, necessitating the use of a very small timestep in order to avoid advective or diffusive numerical instabilities. Appendix A.5 details our attempts to implement various refinement techniques. Ultimately, however, we abandoned this approach in favour of different scheme making use of a fixed, finite difference grid.

3.2.2 Finite difference marker-and-cell

When attempting to develop a Lagrangian finite element code, as described in Appendix A, a problem that we frequently encountered was the appearance of spurious pressure modes. Though we were able to eliminate them entirely for particular choices of elements, these elements had other drawbacks that made them unsuitable. Unfortunately, spurious pressure modes are a very common feature of finite element schemes. If one were to use a finite difference scheme on a fixed computational grid, however, then the zero pressure eigenmodes can be eliminated entirely by the use of a staggered grid, with no negative effects on the accuracy of the solution. Of course, one then needs a different method of tracking the free surface. This section is dedicated to the description of a code that does just that.

The marker-and-cell method was first proposed by Harlow & Welch (1965), and was

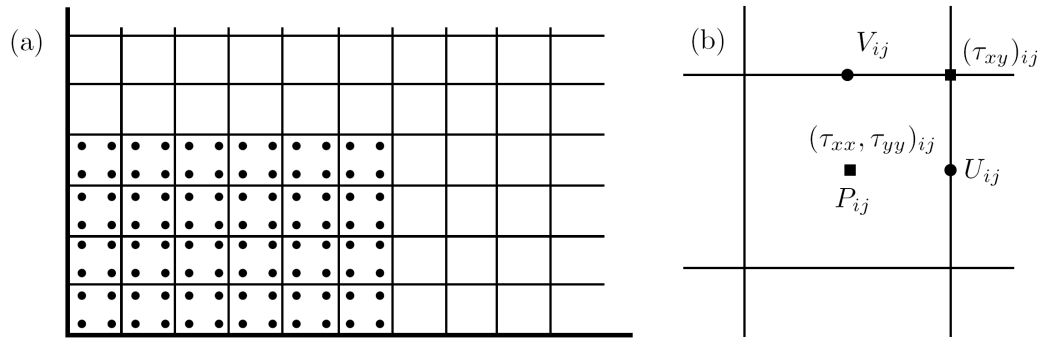

Figure 3.4: (a) A sketch illustrating the finite difference grid and marker particles used in the marker-and-cell method. (b) A computational cell, with the locations at which various quantities are stored. The edges of the cell grid are aligned with the physical boundaries of the system.

originally developed for the simulation of a dam-break problem involving a viscous fluid. One would hope, therefore, that it should be well-suited to compute the same flow with granular material replacing the viscous fluid. The method involves marking the volume occupied by the fluid (or granular material, in our case) by a large number of small particles, which move with the local fluid velocity. A sketch showing schematically the distribution of particles for an undeformed column is shown in Figure 3.4(a). The particles are entirely passive with respect to the equations of motion, except that they determine the region of the computational grid that is occupied by fluid at a given point in time. As time passes, the motion of the particles tracks the motion of the fluid, and the region they occupy mirrors the fluid domain. The difficult part of constructing a marker-and-cell scheme lies in choosing boundary conditions near the free surface, which may have an arbitrary shape.

For the problem of column collapse, one could equally well consider marking only the free surface, and from this work out the volume (or area, in two dimensions) occupied by the fluid. Although this saves a lot of memory and computing time by virtue of not needing to track as many particles, it has the weakness of not being able to handle changes in the flow topology, or any multi-valued (e.g. wave-breaking) phenomena. Neither of these situations should arise in the case of column collapse, yet we insist on marking the volume, rather than

the surface, as it also offers a useful aid to flow visualisation.

Technical details

The marker-and-cell code that we have developed is based on a simple second-order finite difference scheme in space, coupled with forward-Euler timestepping. The incompressibility constraint is enforced by a pressure-update procedure. Suppose that, at time $n\Delta T$, the velocity distribution is $\mathbf{U}^{(n)}$, the pressure distribution is $P^{(n)}$, and the markers are located at positions $(\mathbf{X}_i)^{(n)}$. The procedure to be followed to obtain the state of the system at time $(n+1)\Delta T$ is

1. **Calculate stresses:** Use the current velocity $\mathbf{U}^{(n)}$ and pressure $P^{(n)}$ to calculate the stresses $\tau_{xx}(=-\tau_{yy})$ and τ_{xy} using the (regularised) $\mu(I)$ rheology. We use central differencing wherever possible, and second-order extrapolants otherwise.

2. **Add inertia:** Using the current velocity $\mathbf{U}^{(n)}$, calculate the inertia terms, and apply them to the velocity field. To aid numerical stability, we upwind the velocity gradients, but maintain second-order approximations to all quantities.

3. **Apply gravity:** Increment the vertical velocity of every point by $-\mathrm{d}T$, except for those located at $Y=0$, which are fixed at zero.

4. **Apply stresses:** Apply the stresses calculated in step 1 to the velocity field, obtaining the intermediate velocity field $\mathbf{U}^*$. The velocity boundary conditions at $X=0$ and $Y=0$ are explicitly applied at this step.

5. **Perform pressure update:** We add the contribution from the old pressure $P^{(n)}$ to the velocity field, then find the new pressure $P^{(n+1)}$ by successive over-relaxation of a Gauss-Seidel iterative scheme. The velocity is updated using this new pressure to $\mathbf{U}^{(n+1)}$, which should be divergence-free.

6. **Move markers:** Each marker is moved with the local velocity, obtained by bilinear interpretation and using a forward Euler timestep, obtaining the new marker distribution $(\mathbf{X}_i)^{(n+1)}$.

7. **Update cell status:** Determine which cells contain markers. A cell with no markers is denoted *empty*, a cell containing markers adjacent to an empty cell is a *surface cell*, and all other occupied cells are *bulk cells*. For the purpose of these definitions, the boundaries $X = 0$ and $Y = 0$ are considered to be occupied cells.

8. **Adjust surface velocity:** The velocity at points on the edge of surface cells is only updated during step 3 above. In addition, such velocities are updated during this step in such a manner as to explicitly conserve mass on each of the surface cells.

Although the list above gives a good overview of the numerical scheme used, several further details must be mentioned. In order to avoid the problems associated with spurious pressure modes, it is necessary to use a staggered grid. In our scheme, the pressure, velocity, and stress values are calculated and held at the locations shown by Figure 3.4(b). For ease of applying the no-flux boundary conditions at $X = 0$ and $Y = 0$, the grid is aligned with these boundaries, as sketched in Figure 3.4(a).

The boundary conditions at $X = 0$ are straightforward to apply. In addition to setting $U = 0$, the combination of the no-flux and symmetry conditions (3.6,3.7) require that the off-diagonal component of stress, τ_{xy}, vanishes along $X = 0$. There is therefore no need to extrapolate velocity gradients onto this boundary.

On the rigid base, at $Y = 0$, things are only a little more complicated. The no-slip condition (3.4) is applied to calculate an approximation to the velocity gradient $\partial_Y U$ at $Y = 0$, which is then used to calculate the off-diagonal stress, τ_{xy}, here. Furthermore, the pressure to be used for this calculation is taken by averaging the two nearest pressure values, and adding $\mathrm{d}Y/2$ in order to approximately account for the local pressure gradient. Whilst it might be more appropriate to extrapolate the pressure using more values from the bulk, the first-order error introduced by this approximation is no worse than the error introduced by not accounting for the exact location of the free surface, as we shall discuss next.

The application of stress boundary conditions on the free surface is surprisingly simple. Given that $P = 0$ on the surface, we enforce $P = 0$ on any surface cell. This has the benefit that any stress calculations needed for the bulk cells may be carried out without any need to extrapolate velocity gradients – everything may be done with central differencing. The downside is that the location of the free surface may be in error by as much as the height

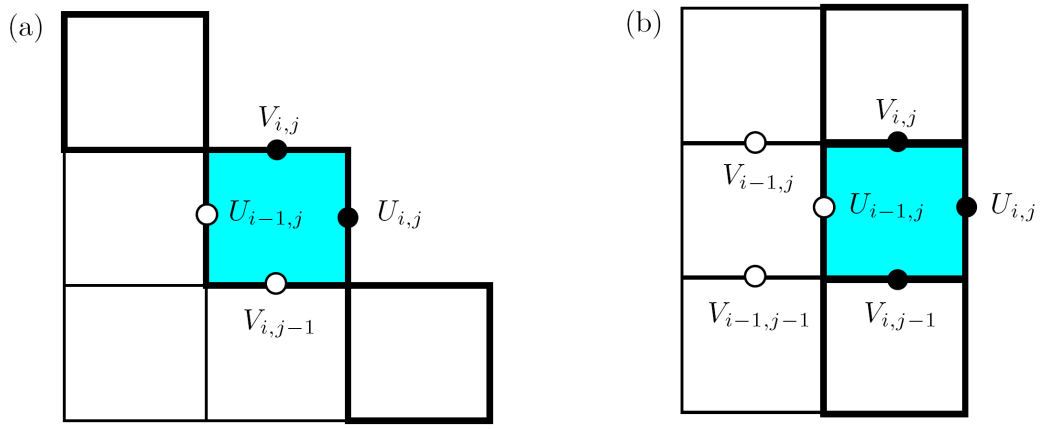

Figure 3.5: Handling of free velocity points located near the free surface. Diagrams illustrate how to assign velocities to the free velocity points (filled circles) for the shaded surface cell, given the bulk velocities (open circles). For clarity, surface cells are denoted by heavier boundary lines than bulk cells. The configurations shown are (a) Two adjacent free velocity points, and (b) Three free velocity points.

of a grid cell, limiting the scheme to have only first-order accuracy. Furthermore, one needs to use a small timestep in order to avoid diffusional instabilities related to the regularised form of the $\mu(I)$ rheology, as discussed in §2.2.5. In practice, we take

$$\delta T = \frac{\epsilon}{5\mu a}\Delta X^2 = \frac{\epsilon}{5\mu a}\Delta Y^2. \tag{3.10}$$

This need for a rather small timestep, taken together with the first-order spatial accuracy of the scheme, means that calculations at a suitably high resolution are quite computationally intensive. A single simulation at large initial aspect ratio ($a = 20$, say) can take as much as two weeks to run from start to finish, although simulations of shorter columns can terminate within two or three days.

An important detail to explain is our treatment of free surface velocity values. If a velocity point is adjacent only to surface and/or empty cells, we call it a *free velocity point*. It is updated as if in free fall, thus it has only the acceleration due to gravity applied to it. Furthermore, the surface cells are constrained to have $P = 0$, so they may not necessarily conserve mass without any further adjustment. In the final part of each timestep, we modify

these free velocities in order to conserve mass on all of the surface cells. If a cell has only one adjacent free velocity point, it is a simple matter to choose this velocity such that it balances the flux into the cell due to the other three velocities. If a cell has two free velocity points on adjacent edges, we force both $\partial_X U$ and $\partial_Y V$ to vanish individually (following Harlow & Welch, 1965), as shown in Figure 3.5(a). A cell containing three free velocity points is a little harder to deal with. Consider the situation illustrated by Figure 3.5(b). In this case, we make a first-order extrapolation of the vertical velocities from those on the left in order to evaluate $V_{i,j}$ and $V_{i,j-1}$. We then fix $U_{i,j}$ by applying conservation of mass to the cell. Other configurations of free surface cells and free velocity points, including the possibility of an isolated cell in free fall, are theoretically possible, but do not appear in practice for the column collapse problem. For this reason, we shall not detail the methods by which they are handled.

Finally, we recall that the modified $\mu(I)$ rheological law does not allow for the formation of static zones, rather it predicts regions of very high viscosity that exhibit a very slow, creeping flow that never comes entirely to rest. For this reason, it is necessary to add a routine that recognises when the flow has 'stopped' – that is, when the strain rate throughout the material is less than some threshold value. Specifically, a little trial and error suggests that one can declare the flow to have stopped when

$$|\dot{\gamma}|_{\max} < \epsilon, \tag{3.11}$$

without any visible change in the long-term behaviour of the solution. When this point is reached, the simulation terminates, and reports the final shape of the granular deposit.

Testing

Before using the marker-and-cell code to make predictions about the collapse of a granular column, it is necessary to validate the code by performing a handful of tests. Given that our principal reason for using a finite difference scheme was to ensure an accurate pressure distribution, checking the pressure seems a sensible place to begin.

Figure 3.6(a) shows the pressure distribution at $T = 0$ calculated using the marker-and-cell code. No trace of any spurious modes is visible, and the agreement with the exact pressure distribution (equation A.9 in the Appendix) shown in Figure 3.6(b) is excellent. A

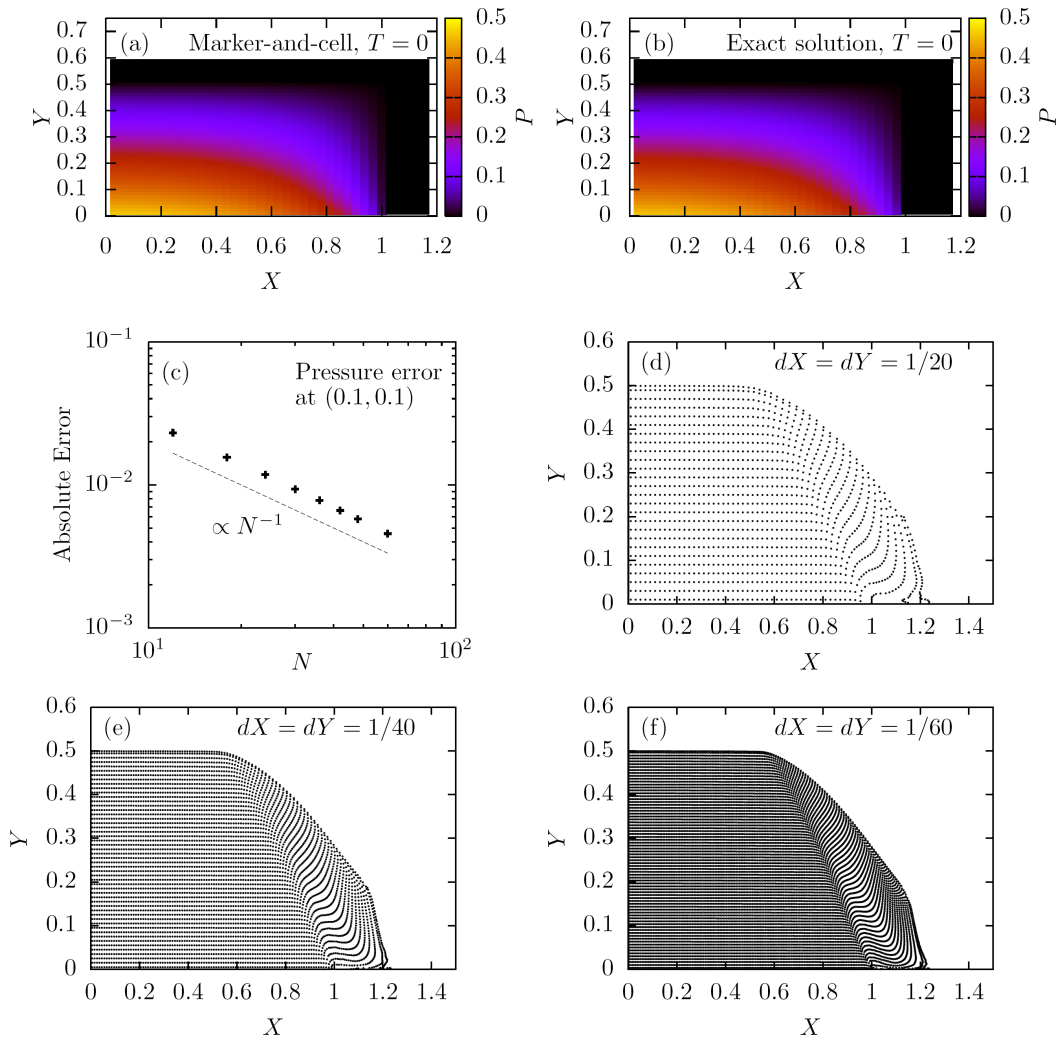

Figure 3.6: Testing the marker-and-cell code for column collapse: (a) Calculated and (b) Exact pressure distributions at $T = 0$ for a column of initial aspect ratio $a = 0.5$; (c) Decay of the absolute error in calculating the pressure at (0.1,0.1) for increasing grid resolution; (d-f) Shape of the granular current and marker particle positions at $T = 0$ for three different grid resolutions, with $a = 0.5$.

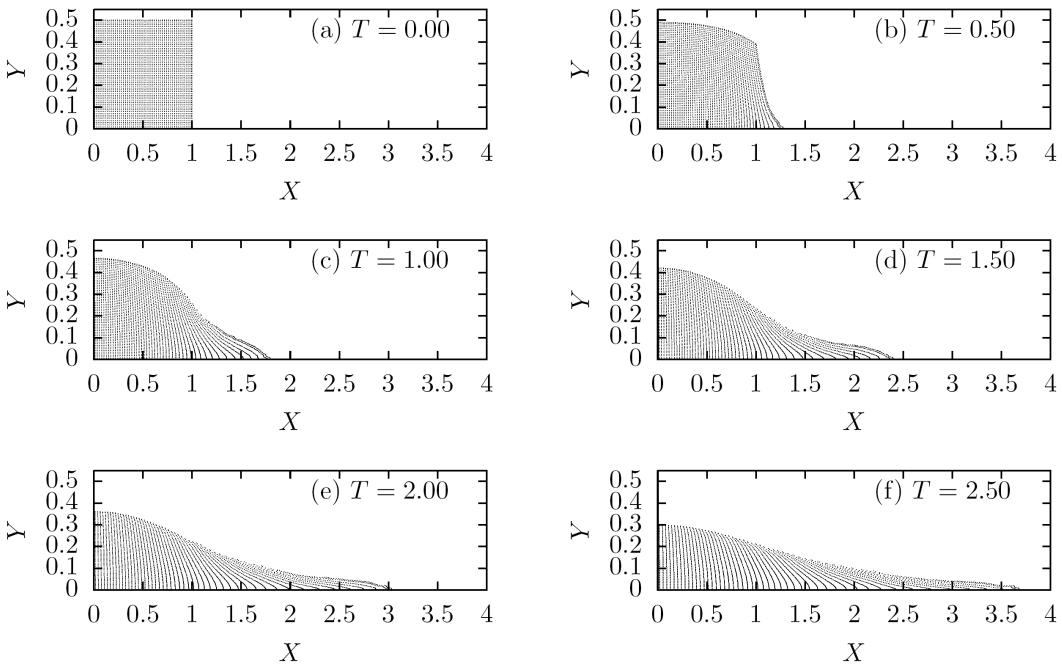

Figure 3.7: Snapshots of marker positions for the inviscid dam-break problem, as calculated by the marker-and-cell code.

formal, quantitative check can be made by considering the dependence the absolute error on the grid resolution. Figure 3.6(c) shows the error in the computed pressure at $(0.1, 0.1)$ for a range of different resolutions. There is clear indication that the errors decrease linearly with the grid spacing, implying that the scheme is first-order accurate. While one might hope for second-order accuracy, consistent with the central differencing used, the first-order-accurate outcome is only to be expected. The fact that our method does not take detailed account of the position of the free surface (as might a volume of fluid method, or a level set method, for example), but rather assumes the free surface to be located along the boundary of the surface cells, we introduce an error proportional to the size of the cells. This uncertainty gives rise to the first-order error observed for the pressure.

Given the lack of any analytical solutions for granular column collapse, it is difficult to make any quantitative statements about the accuracy of the code over the course of the whole collapse. However, we can at least verify that the grid resolution has no significant effect on the dynamics of the column. Figure 3.6(d-f) show the state of the column at $T = 1$ for three different grid spacings. Apart from the differing numbers of particles used (which increases proportionally to the number of cells), the only difference between the three results occurs very near the bottom boundary. On the very bottom row of cells, our treatment of free velocity points can lead to small errors associated with the no-flux boundary condition. We elect to retain the current code on the grounds that such errors are slight, and do not appear to have a great effect on the shape of the final deposit.

As a further test of the time-dependent aspects of the code, we show in Figure 3.7 the calculated evolution of an inviscid fluid column, otherwise referred to as a dam-break. The snapshots shown here agree very well with those presented by Harlow & Welch (1965) in their original work. As they comment, these results are in excellent agreement with the experimental data of Martin & Moyce (1952), but show considerable disagreement with theoretical results obtained by shallow water theory. This would seem to be another vindication for our choice to avoid the shallow water models employed by other authors for the granular column collapse, as discussed in §3.1.

3.3 Results

Having tested the marker-and-cell code and found it to be apparently well-behaved, we are now in a position to compare its predictions to the experimental data available for the collapse of a granular column. Most experimental findings highlight the distinction between collapses at small aspect ratio, in which much of the column remains static, and those at large aspect ratio, which are characterised by having much of the column in free fall during the initial stages of collapse. Accordingly, we shall divide our findings into those relevant to collapses of columns with small (in §3.3.1) and large (in §3.3.2) initial aspect ratio.

For the majority of this section, we shall use only the constant friction ($\lambda = 0$) form of the $\mu(I)$ rheology, because we believe that column collapse, which depends mainly upon inertial effects to spread the granular material, will be relatively insensitive to the precise form of the $\mu(I)$ function used. We shall discuss the differences between constant and non-constant friction in §3.3.3.

3.3.1 Collapse of short columns

When a short granular column collapses under gravity, a wedge of material shears off from the top corner of the column, leaving the remainder of the column relatively undisturbed. The top of the final deposit remaining after such a collapse will be largely horizontal, with a leading edge inclined below the angle of repose (see Figure 3.2d, for example). It is the large region of undisturbed material that characterises the collapse of 'short' columns, rather than any particular range of aspect ratio. Whether a particular column is classified as 'short' depends upon the material parameters involved. Highly frictional materials can produce columns of moderately large aspect ratio that fracture in a manner similar to that described above. In our case, we choose the friction paramters

$$\mu_1 = \tan(36.5^\circ) \quad \text{and} \quad \mu_b = \tan(18.5^\circ), \tag{3.12}$$

in order to draw comparisons with the experiments performed using rough grit by Balmforth & Kerswell (2005). Using these parameters, we find that the columns of aspect ratio $a \lesssim 1$ exhibit the fracturing behaviour that characterises collapses of short columns.

Figure 3.8 shows the evolution of a typical collapse of a short column. In order to make

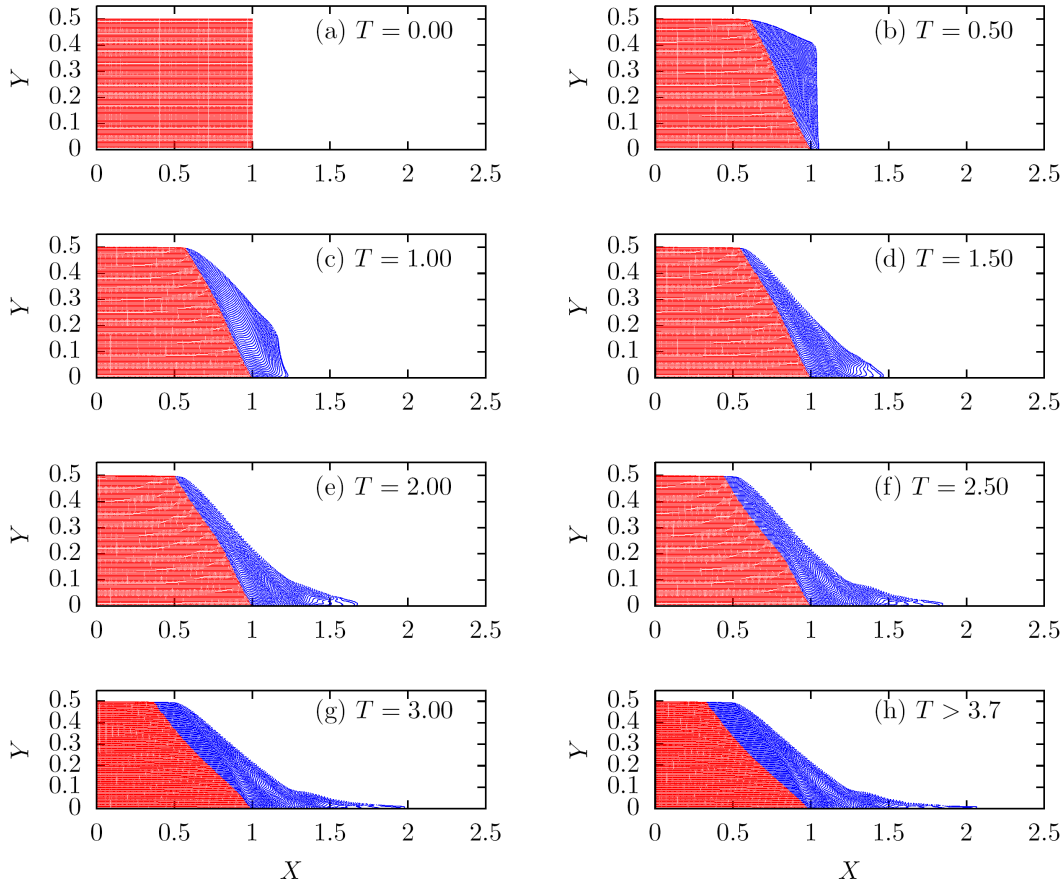

Figure 3.8: Snapshots of marker particle positions for the collapse of a granular column of initial aspect ratio $a = 0.5$. Particles coloured blue have departed from their initial positions by more than a distance, $\frac{1}{2}\Delta X = 1/160$. The remainder are coloured red.

clear the existence of a static region in the centre of the column, only particles that move significantly are coloured blue. The remaining particles, which do not move appreciably throughout the collapse, are coloured red in each frame. As one might expect, following the discussion above, the collapse begins with the failure and slip of the top corner of the column (Figure 3.8b), which then spreads over the horizontal surface. A large proportion of the column remains essentially static throughout the collapse. The final two frames shown suggest that rather more of the column has been in motion than one might expect. This is consequence of the regularised $\mu(I)$ rheology. Although the material in the interior of the column ought to be entirely static, the regularisation allows for a slow viscous creeping flow, which moves most quickly near the region of shearing material.

Figure 3.9(a) shows the final deposits predicted by our numerical simulations, and Figure 3.9(b) illustrates the experimental results of Balmforth & Kerswell (2005) for columns having a similar range of aspect ratios. In the cases shown, there seems to be a generally good agreement between the numerical predictions and the experimental data, in that the shape of each numerical deposit closely resembles that of the experimental deposit for a similar value of the aspect ratio, a. Furthermore, both the numerical and experimental profiles collapse onto a single master profile after scaling heights with the initial column height, H, and lengths with the runout, $R - L$, from $x = L$. Although this agreement holds for most of the deposit profile, one ought to pay close attention to the transition between the horizontal top of the undisturbed central column and the sloping surface to its right. It appears that the transition is much sharper for the experimental data than for the results of the simulation. This is also a numerical artefact, which on this occasion results from our use of the regularised $\mu(I)$ rheology. Instead of the sharp fractures seen in experiments, the regularisation allows for a slow viscous creeping flow throughout the supposedly static region. The effect of this creeping flow is to smooth out the sharp corner, resulting in the transition seen in the simulation results.

The existence of the master deposit profile shown in Figure 3.9 implies that the runout distance, $(R-L)/L$, varies linearly with the initial aspect ratio of the column. This relationship is further illustrated by Figure 3.10(a), in which it is plotted explicitly. Although this linear relationship is observed both in the experiments of Lube *et al.* (2005) and Balmforth & Kerswell (2005), we note that the material-dependent prefactor is overestimated in our

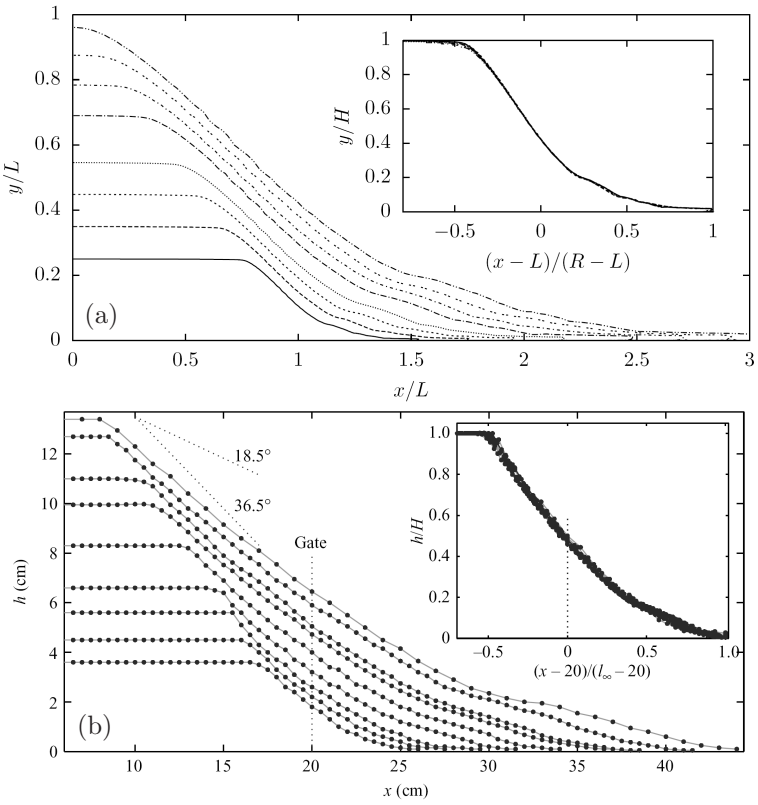

Figure 3.9: The shape of the final deposit for the collapse of short columns. (a) Numerical results for a range of initial aspect ratio ($a = 0.25, 0.35, 0.45, 0.55, 0.70, 0.80, 0.90, 1.00$) using a constant friction model with parameters estimated by Balmforth & Kerswell (2005). The inset shows the same results scaled such that the depths are normalised by the initial height of the column, and horizontal lengths so that the runout from $x = L$ is unity. (b) A similar plot to (a), but showing the experimental results of Balmforth & Kerswell (2005) for a slightly different range of initial aspect ratio. The dotted line marked 'Gate' denotes the initial horizontal extent of the column.

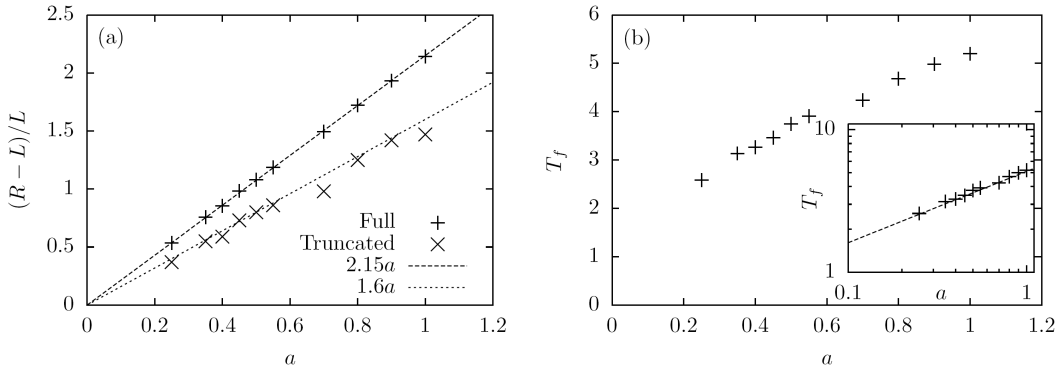

Figure 3.10: Predictions for short columns. (a) The variation of the runout distance, $(R - L)/L$, with initial aspect ratio, a. We include both the 'full' runout (calculated by considering all marker particles) and the 'truncated' runout, described in the main text. Approximate straight line fits are shown for comparsion. (b) The duration of the collapse, T_f, as a funtion of the initial aspect ratio, a. The inset show the same plot on logarithmic axes, with a dashed line proportional to $\sqrt{a}$.

simulations. Balmforth & Kerswell (2005) find that

$$R \approx L(1 + 1.6a), \tag{3.13}$$

for their experiments with coarse grit, whereas our simulations, using the same friction parameters, yields

$$R \approx L(1 + 2.15a). \tag{3.14}$$

There are several possible reasons for this overestimation of the runout. The vertical walls confining the experimental columns exert a frictional stress on the granular material that retards the flow, and results in a shorter runout. Although this effect is certainly present, it is unlikely to be the sole cause of the disagreement. A more likely explanation is based on the fact that the experiments use actual grains, so the measurements are affected by the discrete nature of the material. In the simulation, we treat the granular material as a continuum, and can therefore predict an arbitrarily thin layer of material. In the experiments, the deposit is limited to be at least a grain diameter in depth. One can attempt to truncate the deposit in several different ways. A reasonable method of truncation is to discount any marker particles

in the bottom layer of cells (that is, with $y < \Delta y$) when calculating the runout. The resulting predictions for the runout are shown by diagonal crosses in Figure 3.10(a). Although they predict a runout that corresponds more closely to the experiments of Balmforth & Kerswell (2005), and in particular their approximate expression for the runout (3.13), there is much more scatter in the data, due mainly to the finite spacing between marker particles. Had we chosen a different height below which to ignore the marker particles, the results would be different again. We shall not dwell further on the truncation of the deposit, and instead concentrate on the fact that the simulation predicts the correct power-law dependence of runout on initial aspect ratio for collapses of short columns.

The positive result, predicted by our simulations, that the runout scales linearly with the initial aspect ratio of the column may be deduced via a simple agrument based on dimensional analysis. For the fractured profiles that characterise collapses of short columns, the width of the central, static, flat portion of the deposit is irrelevant to the runout. Two columns of the same height will only spread by the same horizontal distance, as long as they both have small enough aspect ratio to allow for a fractured collapse. Having eliminated the width of the column, L, from the problem, we are left with only two lengthscales: the depth of the column, H, and the grain size, d. Given the wealth of experimental observations (Lube *et al.*, 2004, 2005; Balmforth & Kerswell, 2005) that grain size plays little role in the dynamics, we conclude that the runout, $R-L$, must scale with H, as it is the only remaining lengthscale. Thus

$$R = L + \beta H = L(1 + \beta a), \tag{3.15}$$

where β is a dimensionless constant that will depend upon the frictional properties of the material and of the plane over which it spreads. A similar argument may be made to deduce that the duration of the collapse, T_f, must be proportional to the only timescale available in a short collapse,

$$T_f = \alpha\sqrt{\frac{H}{g}}. \tag{3.16}$$

This dependence is confirmed by the simulation data, as shown in Figure 3.10(b).

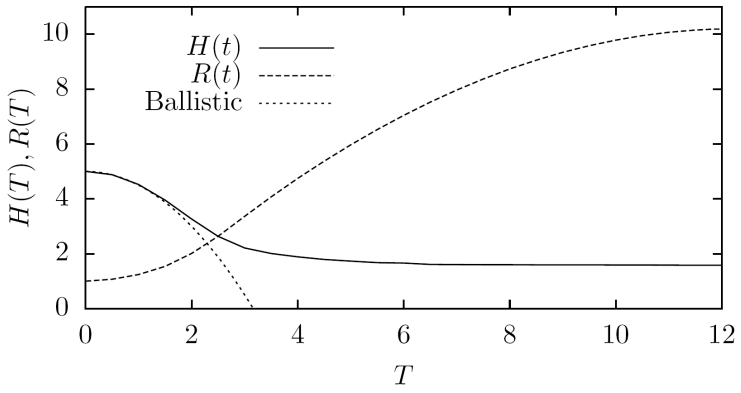

Figure 3.11: Time series showing the evolution of the maximum height, $H(T)$, and runout, $R(T)$, for the collapse of a column with initial aspect ratio $a = 5.0$. The ballistic trajectory, $H(0) - \frac{1}{2}T^2$, is shown for comparison.

3.3.2 Collapse of tall columns

As the intial aspect ratio increases, a point will be reached beyond which the final deposit will no longer feature a horizontal central region near $x = 0$. This may be seen to occur for the tallest deposit shown in Figure 3.9(a), which initially started with aspect ratio $a = 1$. In such cases, the boundary conditions due to symmetry in the vertical axis will affect the solution, and we may no longer neglect the width of the column, as we did in the dimensional analysis leading to linear runout. For the collapse of columns with larger aspect ratio, which we shall refer to as 'tall' columns, there will be an additional dependence of the runout, R, and duration, T_f, on the initial aspect ratio, a, that cannot be deduced by dimensional analysis.

Collapses of tall columns are also characterised by an initial period of free fall immediately following that start of the collapse. With insufficient pressure from below to support its weight, the top of the column accelerates vertically downwards under gravity, forcing material beneath it to spread horizontally. This period of free fall, observed experimentally by Balmforth & Kerswell (2005), is also apparent in our simulations, as shown by Figure 3.11. Once the sides of the falling column have become less steep, the pressure inside becomes sufficient to decelerate the freely-falling material, and the height of the column reaches its

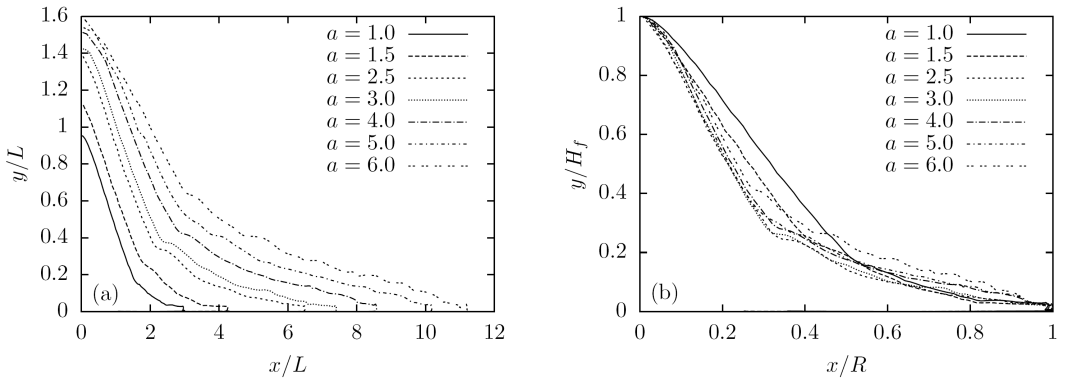

Figure 3.12: The shape of the final deposit for the collapse of taller columns. (a) Numerical results for a range of initial aspect ratio ($a = 1.0, 1.5, 2.5, 3.0, 4.0, 5.0, 6.0$) using a constant friction model with parameters estimated byBalmforth & Kerswell (2005). (b) The same data with both horizontal and vertical lengths scaled by the final runout, R, and deposit height, H_f, respectively.

minimum, final value. However, even when the height of the column has reached its final state, the runout continues to increase. It is the case, both in our simulation and in the experiments of Lube *et al.* (2004), that this continued spreading results from secondary avalanching of the deposit, and ceases when the surface of the entire deposit is inclined at an angle smaller than the angle of repose.

Figure 3.12(a) shows the final deposits predicted by simulations with constant friction for a selection of tall columns. In each case, the deposit is qualitatively similar, with a steep, almost triangular central section giving way to a more shallow, gentle slope towards the edge of the deposit. Unlike the case of shallow columns, we observe no collapse of the deposit profiles here (Figure 3.12b). This result, which is to be expected from the dimensional analysis at the end of §3.3.1, is also in agreement with the experimental findings of Balmforth & Kerswell (2005). However, the experimental results suggests that the free surface inclination of the deposit near $X = 0$ is essentially independent of the initial aspect ratio of the column. This observation is not captured by the numerical results.

The runout of tall column collapses is a matter of some contention in the experimental literature. Both Lube *et al.* (2004) and Balmforth & Kerswell (2005) claim to observe a

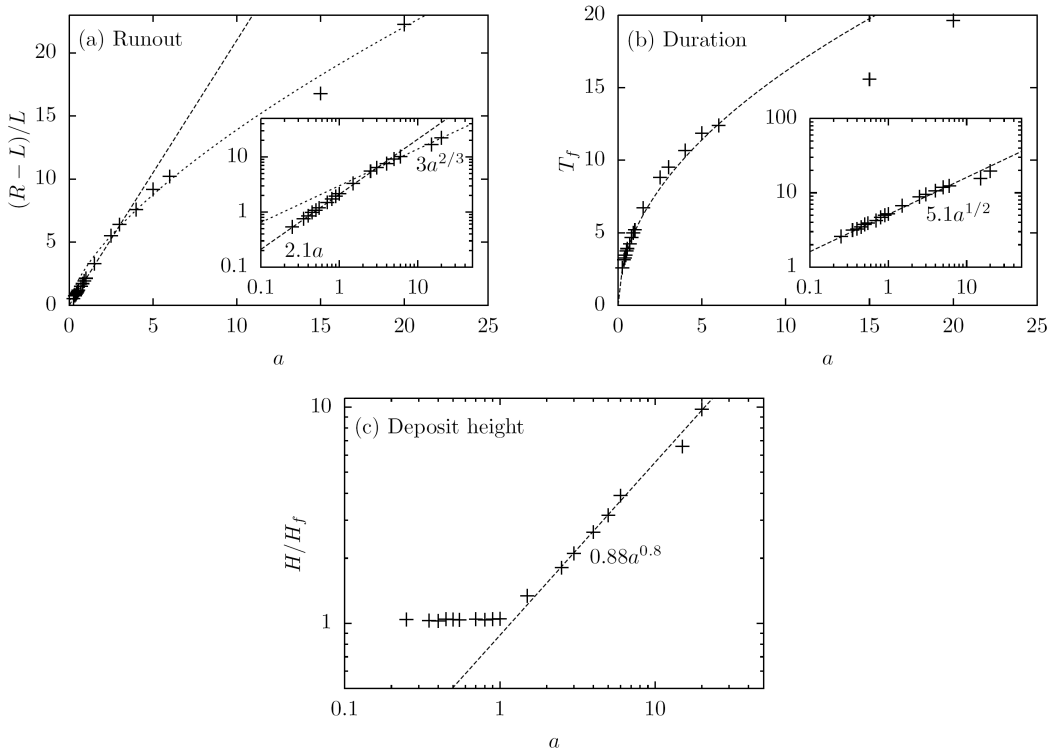

Figure 3.13: Numerical predictions for tall and short columns. (a) The variation of the runout distance, $(R-L)/L$, with initial aspect ratio, a. The straight line fit for short collapses is shown for comparison, as well as a power-law fit more appropriate for taller columns. (b) The duration of the collapse, T_f, as a funtion of the initial aspect ratio, a. The inset show the same plot on logarithmic axes, with a dashed line proportional to $\sqrt{a}$. (c) The ratio of initial to final height, H/H_f, as a function of initial aspect ratio. The dashed line represents an approximate fit to the data for taller columns.

power-law relationship

$$R = L(1 + \beta a^{\gamma}), \tag{3.17}$$

but their experimental data disagree with regard to the exponent, γ. Lube *et al.* (2004) find a good fit of the form (3.17) using an exponent of $\gamma = 2/3$. However, the greater range of data collected by Balmforth & Kerswell (2005) suggest an exponent of $\gamma = 0.9 \pm 0.1$, but their data contains a much greater degree of scatter. In both cases, the power-law behaviour does not become clear until the initial aspect ratio exceeds $a \approx 6$. Unfortunately, our numerical code becomes less stable for collapses of taller columns, mainly due the larger velocities that arise in the free-fall phase of the collapse. To control the instability whilst retaining a good spatial resolution, it is necessary to have a very small timestep, meaning that a single simulation can take several weeks to run to completion on the machines available. Due to time constraints, we are only able to show results for two tall collapses, at $a = 15$ and $a = 20$ in this work, although we hope to present a more thorough investigation in a future journal article.

Taking what data we have available into account, Figure 3.13(a) does not give particularly strong agreement with either set of observations. It definitely seems to be the case that the runout departs from the linear behaviour observed for shorter columns, but the trend for taller columns is not clear. Although we show a power law with exponent 2/3 (as observed by Lube *et al.*, 2004) for comparison with the numerical results, the agreement is far from convincing.

While the behaviour of the runout changes as we move from short to tall columns, the duration of the collapse continues to vary like $a^{1/2}$ until $a \approx 10$, as shown by Figure 3.13(b). Although the dimensional analysis argument used in §3.3.1 no longer holds, the result appears to be unchanged. This is consistent with the experimental observations of Lube *et al.* (2004), and suggests that the time-dependent behaviour is dominated by free-fall of grains at the leading edge of the flow.

Another quantity that can be easily compared with the experimental data is the maximum height of the final deposit. Unlike short columns, for which the top of the column remains essentially undisturbed, the large degree of free fall involved for higher columns causes the height of the deposit to be significantly lower than the height of the initial col-

umn. In Figure 3.13(c), we show the ratio of the initial to final column heights as a function of the initial aspect ratio. In the experiments of both Lube *et al.* (2004) and Balmforth & Kerswell (2005), one finds that

$$H_f \propto L\, a^{0.4}, \quad \text{or} \quad \frac{H}{H_f} \propto a^{0.6}, \tag{3.18}$$

for tall columns. The data for our simulations exhibits similar power-law behaviour of the form

$$\frac{H}{H_f} \propto a^{0.8}, \tag{3.19}$$

even for the collapse of very tall columns. The reason for the discrepancy in the exponent between experimental and simulation data is unclear. It is possible, though unlikely, that the additional viscous creep present in the simulated collapse is responsible for this disagreement.

3.3.3 Discussion

Having simulated column collapse using a constant-μ version of the $\mu(I)$ rheological law, we can draw several conclusions. The model clearly captures correctly many details observed for experimental collapses. In particular, the simulations for short columns predict the correct power-law dependence of the duration and runout of the collapse on the initial aspect ratio for $a < 1$. For taller columns, the comparison with experimental data is less favourable. Our results share aspects in common with both Lube *et al.* (2004) and Balmforth & Kerswell (2005), although a lack of knowledge about the material parameters μ_1, μ_2 and I_0 prohibit a quantitative comparison. For collapses of very tall columns, our simulation data are rather sparse, but indicate weak agreement with the observations of Lube *et al.* (2004).

The differences between the two sets of experimental data and our simulations could arise from several possible effects. In the experiments, the vertical side walls used to confine the column could impart frictional stresses to the column, limiting its spread. Our simulations, on the other hand, are entirely two-dimensional: no account of any friction from the side walls is made. This effect could certainly be responsible for the disagreement between experimental and simulated results for the final runout and column height. However, the majority of the experiments by Balmforth & Kerswell (2005) were carried out in a wide channel, so the influence of lateral walls ought to be negligable. Balmforth & Kerswell

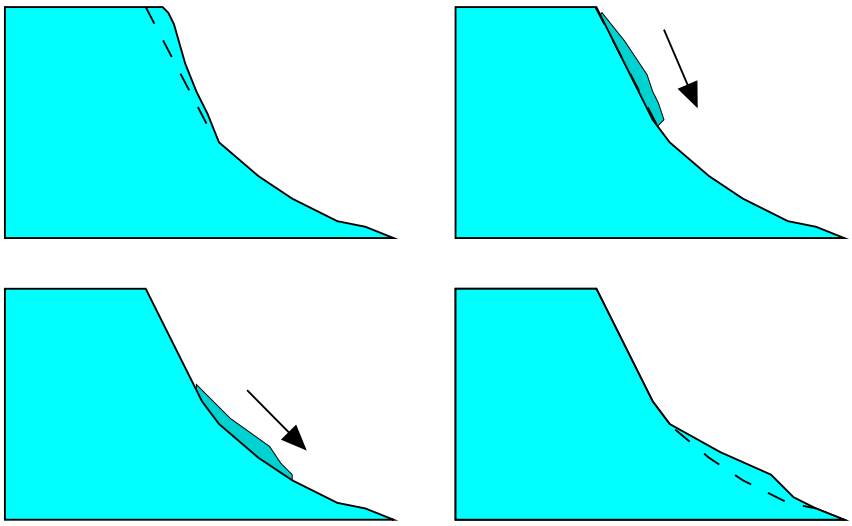

Figure 3.14: Illustration of secondary avalanching. Material whose surface is inclined at too great an angle will shear off (Panel a) and begin to slide down the surface below (Panels b,c). Eventually, the surface inclination decreases to a point where motion can no longer be sustained, leaving a 'kink' in the deposit profile (Panel d).

(2005) also report a slight dilation of the granular material over the course of the collapse. We make no account of this, treating the granular medium as strictly incompressible.

Putting discussion of possible sources of error aside for a moment, we mention that a notable feature of the deposit profiles obtained both experimentally (by Balmforth & Kerswell, 2005) and numerically is the existence of a slight 'kink' in the free surface near the front of the deposit. One might speculate that such a feature might originate as a product of the secondary avalanching that occurs in the latter stages of a collapse, as shown by Figure 3.14. After the initial period of free fall, it may be that a portion of the free surface of the deposit is still inclined at an angle greater than the angle of repose. In this case, a small avalanche is initiated, and material begins to slide down the surface of the pile. This current will eventually reach a region of free surface inclined at an angle shallow enough that frictional drag dominates over the gravitational force acting downhill. This causes the avalanching current to decelerate and eventually stop. It seems reasonable to suggest that the 'kink'

observed in many of the deposit profiles arises as one of these arrested avalanches. Despite the qualitative similarity between the experimental and numerical results, the quantitative agreement between the two is relatively poor. One typically observes the numerical 'kink' to be located higher up the surface of the pile, and to be larger than its experimental counterpart. The reason for this discrepancy is not clear, although it may well be linked to the rather coarse description of the free surface afforded to us by the marker-and-cell method. A higher spatial resolution may help to eliminate this problem, but it would require a much smaller timestep, hence a much longer simulation, in order to avoid numerical instabilities. It is likely that the use of a semi-implicit timestep or a more sophisticated means of upwinding the advection term, would improve the numerical stability, and allow for solution at higher spatial resolution in a reasonable length of time.

It is entirely possible that the less than perfect agreement between the simulated and experimental results is due to some inadequacy of the rheological model used. In light of the particle dynamics simulations carried out by Lacaze & Kerswell (2009), and mentioned previously in §3.1, one would hope that this is not the case. Perhaps, then, it is the parameters used in our model that are at fault. Unfortunately, the inefficiency of the numerical scheme used means that a single simulation can take a very long time to run; as much as two weeks on a reasonably fast computer for the tallest columns. This makes a thorough exploration of the parameter space a costly exercise, and one that our time constraints prohibit us from undertaking in the present work. Instead, we hope to report the results of such an investigation in a future journal article. Before concluding our discussion of column collapse, however, we ought to briefly examine a simulation carried out using the full $\mu(I)$ rheological law, rather than the constant friction approximation. The use of a variable friction coefficient requires the estimation of the two additional parameters: λ and I_0. As neither Lube *et al.* (2004) nor Balmforth & Kerswell (2005) offer any data in this regard, one cannot expect to obtain a significant quantitative improvement in the simulated solution. Instead, we offer only a brief comparison between the collapse of a column of initial aspect ratio $a = 0.5$ with constant friction (i.e. $\lambda = 0$), and that of an identical column with $\lambda = 1.2$. Some images of the collapse of each column at certain times are shown in Figure 3.15. Although the column obeying the $\mu(I)$ law comes to rest more quickly and spreads over a shorter horizontal distance than the column with constant friction, the differences between the two

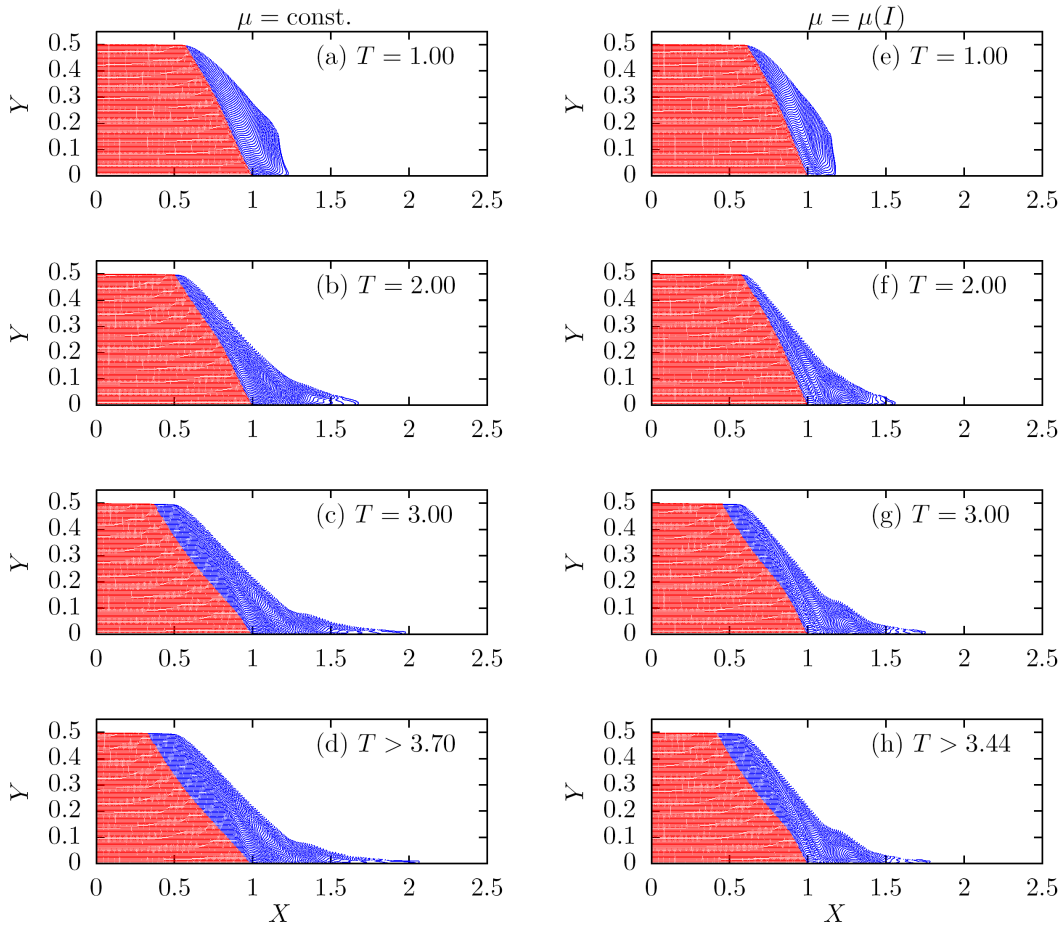

Figure 3.15: Several snapshots of marker particle positions for the simulated collapse of a column with initial aspect ratio $a = 0.5$, using (a-d) constant friction coefficients: $\mu_1 = 0.75$, $\mu_b = 0.33$; (e-f) the $\mu(I)$ friction coefficient, with $\lambda = 1.2$, and $I_0 = 0.0005$. In both sets of plots, blue particles have moved more than a distance, $\frac{1}{2}\Delta X$, from their initial positions. All other particles are coloured red.

are relatively slight. Furthermore, they may be attributed to the fact that the $\mu(I)$ medium is more dissipative than the constant μ material, because $\mu(I) \geq \mu_1$ throughout the column. This enhanced dissipation could explain not only the shorter and more controlled collapse, but also the more prominent 'kink' near the front of the deposit in Figure 3.15(h). In this case, an avalanching current (as discussed previously) is brought to rest sooner than that apparent in Figure 3.15(d). It may well be that the enhanced dissipation afforded by the variable friction coefficient, $\mu(I)$, can reconcile the simulated predictions for the runout with the experimental results. However, time constraints require that a detailed comparison will have to await a future investigation.

4

Stability

With the numerical scheme developed in Chapter 3, we have enabled the application of the $\mu(I)$ constitutive law to more-or-less arbitrary two-dimensional flows (although, of course, some modifications will be necessary depending on the specific flow configuration under study). In the context of evaluating the performance of the $\mu(I)$ law, this numerical approach has a significant drawback, in that any discrepancies between simulation and experiment may be due either to the model or the numerical scheme used, and classifying errors according to their cause is not necessarily straightforward. In this chapter, we adopt an analytical approach, in the hope of testing the predictions of the $\mu(I)$ rheology against another kind of experimental result. However, as we have mentioned previously, the structure of the governing equations often prohibits analytical solutions to even relatively simple problems (such as the annular shear of §2.1.2). Seeking tractable problems, we turn to problems of stability, which still exhibit the non-trivial, multi-directional shear necessary to test the $\mu(I)$ properly in a three-dimensional setting, but allow one to linearise the equations about a known basic state, and calculate asymptotic solutions in physically-relevant regimes.

The first application of the $\mu(I)$ constitutive law to a stability problem was made by Forterre (2006), who considered the formation of long surface waves (otherwise known as roll waves, or the Kapitza instability) in the flow of a relatively shallow layer of granular material on a rough inclined plane. This instability gives rise to wave structures running

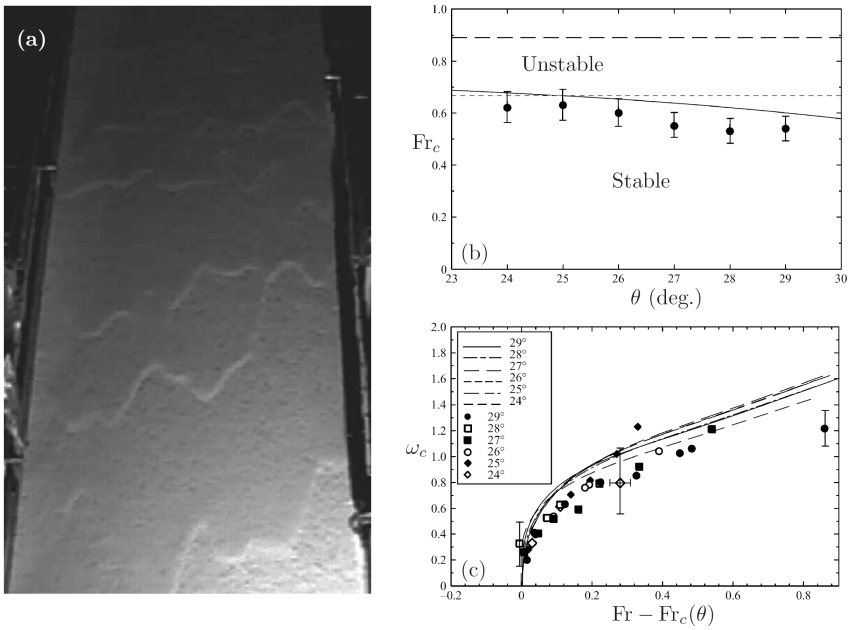

Figure 4.1: Roll waves in a shallow granular inclined plane flow: (a) A photograph of the experiments by Forterre & Pouliquen (2003); (b) The dependence of the threshold Froude number, Fr_c, on the inclination angle, θ, from experiments (circles), the $\mu(I)$ model (solid line), and two Saint-Venant models (dashed lines); (c) Cutoff wavenumber ω_c as a function of Froude number Fr from experiments (symbols) and the $\mu(I)$ model (lines) for a range of inclination angles. Panels (b) and (c) are adapted from Forterre (2006).

perpendicular to the direction of flow, as shown in Figure 4.1(a). The cross-slope variation in Figure 4.1(a) is thought to be caused by a secondary instability. This roll wave instability was studied experimentally by Forterre & Pouliquen (2003), who attempted to model the instability by using a simple Saint-Venant (shallow water) model. Although they found that this model made reasonable predictions for the stability threshold, this simple model did not capture the short-wave cutoff observed in the experiments. By contrast, a linear stability analysis using the $\mu(I)$ constitutive law (Forterre, 2006) not only offers an improved prediction for the threshold (Figure 4.1b), but also correctly predicts the existence of a cutoff wavelength, and predicts its magnitude to a reasonable degree of accuracy when compared

with experiments (see Figure 4.1c).

We aim to build upon the success of Forterre (2006) by applying the $\mu(I)$ constitutive law to two different stability problems. In §4.1, we attempt to explain the formation of longitudinal wavelike structures observed for flow on an inclined plane, studied experimentally by Forterre & Pouliquen (2001), and more recently by Börzsönyi *et al.* (2009). Afterwards, we consider (in §4.2) the well-studied problem of sand dune formation, breaking new ground by using a full granular constitutive law to describe the transport of granular material driven rom above by fluid stresses.

4.1 Longitudinal vortices in inclined plane flow

The roll wave instability studied by Forterre (2006) occurs in fluid as well as granular dynamics, and as a result its mechanism was well-understood before the experimental study of Forterre & Pouliquen (2003). By way of contrast, the instability we consider next has no direct counterpart in fluid mechanics. When granular material flows down an inclined plane, as illustrated by Figure 4.2(a), it is possible that, instead of roll wave formation, one may observe the formation of longitudinal vortices (i.e. with axes pointing downslope), as shown in Figure 4.2(b). A surface waveform appears that is aligned with the longitudinal rolls (Figure 4.2c).

This instability was first recorded by Forterre & Pouliquen (2001), who conducted experiments on a relatively steep incline. In these experiments, the flowing granular material was quite loosely packed ($\phi \approx 0.3$), so Forterre & Pouliquen (2001) turned to the kinetic theory of dissipative granular gases to model the instability. Based on their observation that the material beneath the peaks was less densely packed than beneath the troughs, Forterre & Pouliquen (2001) likened the instability to the celebrated Rayleigh-Bénard instability experienced by fluids heated from below. By replacing the fluid temperature with the *granular temperature* (a measure of the fluctuations in the velocities of individual grains about the mean flow), they were able to use the equations of granular kinetic theory to demonstrate that the analogy was valid. In this case, the strong shear provided by the roughness of the inclined plane is the source of agitation ('heat') needed to drive the convection rolls. This model also gave a plausible explanation for the inability of Forterre & Pouliquen (2001) to

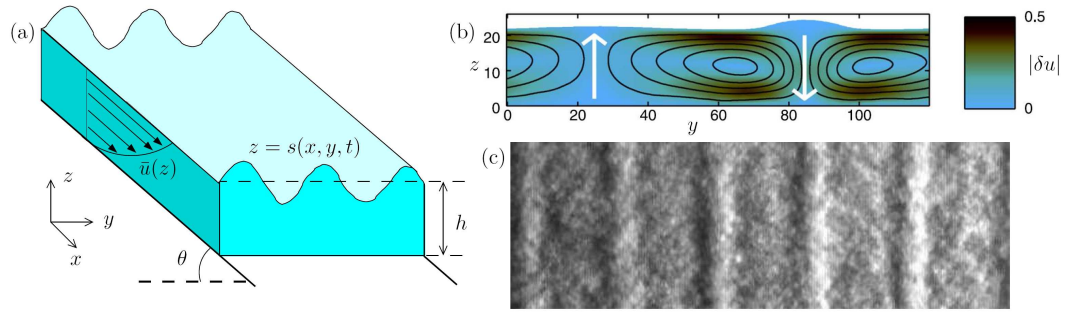

Figure 4.2: Longitudinal roll formation in inclined plane flow: (a) Diagram and notation; (b) Roll structure and variation in perturbation speed $|\delta u|$ obtained via DEM by Börzsönyi *et al.* (2009); (c) Photograph of the free-surface deformation over several wavelengths, also from Börzsönyi *et al.* (2009).

observe the instability for lower flow rates (using either a less steep incline, or a shallower granular layer).

Later, Börzsönyi *et al.* (2009) carried out a similar series of experiments involving inclined plane flow, but observed the formation of longitudinal vortices even in relatively dense flow (with $\phi \approx 0.55$). Further investigation using particle image velocimetry and particle dynamics simulations revealed that there is no sharp transition between vortex formation in dense and dilute flow regimes, but that the vortices in dense flow have a different character to those occurring in the dilute flow studied by Forterre & Pouliquen (2001). In dense flow, the flow beneath the surface peaks is faster and more dense than elsewhere, and convection rolls observed are typically more horizontally elongated when compared to Rayleigh-Bénard-type rolls observed elsewhere. For the rapid flow studied by Forterre & Pouliquen (2001), however, the peaks were found to be more dilute and to move more slowly than the troughs. This evidence seems to suggest that the 'thermal convection' mechanism, proposed by Forterre & Pouliquen (2001) to explain the dilute flow instability, may not be appropriate to dense flows, posing the question of exactly what provides the mechanism for the dense flow instability.

In order to investigate the formation of longitudinal vortices in dense flows, we shall perform a linear stability analysis of inclined plane flow using the $\mu(I)$ rheology. Although it is clear from both the experiments and simulations of Börzsönyi *et al.* (2009) (and through

personal communication with the authors) that variation in the packing fraction of grains is present wherever the longitudinal vortices occur, it is not clear that it is necessary for the instability to occur. Furthermore, preliminary results from the simulations (Jim McElwaine, personal communication) suggest that the $\mu(I)$ rheological law seems to hold over a significant proportion of the flows studied numerically. We shall therefore continue to use the incompressible $\mu(I)$ model employed throughout this thesis, and consider whether it alone is capable of providing a mechanism by which the longitudinal vortices can arise.

4.1.1 Governing equations

Consider the inclined plane flow shown in Figure 4.2. We recall (from §2.1.4) that the steady velocity profile in a granular layer of depth h on a plane inclined at angle θ to the horizontal is

$$\bar{u}^*(z^*) = \frac{2\sqrt{g\phi\cos\theta}}{3d}\mu^{-1}(\tan\theta)\left[h^{3/2} - (h-z^*)^{3/2}\right], \tag{4.1}$$

and the steady pressure is

$$\bar{p}^*(z^*) = (h-z^*)\rho g\cos\theta. \tag{4.2}$$

In this section, an asterisk denotes a dimensional variable. Dimensional parameters are left unstarred, and will soon be scaled out. Based on the steady-state velocity and pressure profiles, we rescale the variables according to

$$\begin{aligned}
\text{Length} \qquad & (x^*, y^*, z^*) = h(x, y, z), \\
\text{Velocity} \qquad & (u^*, v^*, w^*) = \mu^{-1}(\tan\theta)\frac{h\sqrt{\phi g h\cos\theta}}{d}(u, v, w), \\
\text{Time} \qquad & t^* = \mu^{-1}(\tan\theta)\sqrt{\phi g h\cos\theta}\, t, \\
\text{Pressure} \qquad & p^* = \rho g h\cos\theta\, p.
\end{aligned} \tag{4.3}$$

Under these scalings, the dimensionless equations of motion are

$$\mathrm{Fr}^2\left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u}\right) = -\nabla p + \begin{pmatrix} \tan\theta \\ 0 \\ -1 \end{pmatrix} + \nabla\cdot\left(\mu(I)p\frac{\dot{\gamma}}{|\dot{\gamma}|}\right), \tag{4.4}$$

$$\nabla\cdot\mathbf{u} = 0, \tag{4.5}$$

where

$$\mathrm{Fr}^2 = \phi\mu^{-1}(\tan\theta)\frac{h}{d} \tag{4.6}$$

is a Froude number that is in this case related to the thickness of the flowing layer measured relative to a grain diameter.

The rough inclined plane exerts a no-slip condition at the base of the flow,

$$\mathbf{u} = 0 \quad \text{at} \quad z = 0. \tag{4.7}$$

On the free surface, which is located at $z = 1 + s(x, y, t)$, we ought to apply the kinematic boundary condition and a stress-free condition

$$\frac{\partial s}{\partial t} + u\frac{\partial s}{\partial x} + v\frac{\partial s}{\partial y} = w \quad \text{at} \quad z = 1 + s(x, y, t), \tag{4.8}$$

$$\sigma \cdot \mathbf{n} \to 0 \quad \text{as} \quad z \to 1 + s(x, y, t), \tag{4.9}$$

where $\mathbf{n} = \nabla(s - z)$ is the normal to the surface. The limit in (4.9) may be replaced by an equality at $z = s(x, y, t)$, although we retain the current form because the vanishing of the steady pressure, $\bar{p}(z)$, and velocity gradient, $\bar{u}'(z)$, at $z = 1$ results in a (regular) singularity in the linearised equations, which is better handled by the form (4.9). The singularity has little effect on the asymptotic solution of the linearised equations, but will be discussed further in the context of our numerical approach in §4.1.4.

4.1.2 Linearisation

The equations (4.4)-(4.9) completely describe three-dimensional flow down an inclined plane. With the nondimensionalisation chosen, the steady-state velocity and pressure distributions are given by

$$\bar{u}(z) = \frac{2}{3}\left[1 - (1 - z)^{3/2}\right], \qquad \bar{p}(z) = 1 - z. \tag{4.10}$$

In order to investigate the formation of longitudinal vortices, we linearise equations (4.4)-(4.9) about the steady state (4.10). We follow the usual procedure of linear stability analysis, making a perturbation to the free surface of the form

$$s = s_0 e^{ik(y+ct)}, \tag{4.11}$$

and defining the perturbed velocity and pressure

$$\begin{aligned} u &= \bar{u}(z) + U(z)e^{ik(y+ct)} \qquad & v &= V(z)e^{ik(y+ct)}, \\ w &= W(z)e^{ik(y+ct)} \qquad & p &= \bar{p}(z) + Pe^{ik(y+ct)}, \end{aligned} \tag{4.12}$$

with any downslope variation explicitly omitted. This simplification is not merely a convenience. In fact, it is designed to suppress the Kapitza (roll wave) instability studied by Forterre (2006) from appearing in our calculations. By eliminating downstream variation, the mechanism leading to roll waves is made unavailable, and allows us to focus on the stability of inclined plane flow to transverse perturbations, such as may give rise to the longitudinal vortices studied by Börzsönyi *et al.* (2009).

Substitution of the perturbed dependent variables (4.12) into the momentum equation (4.4), followed by elimination of the cross-slope velocity, $V = iW/k$, via the continuity equation (4.5), yields the linear problem

$$\mathrm{Fr}^2\left[ikcU + \sqrt{1-z}W\right] = -k^2\tan\theta\sqrt{1-z}U + \frac{\mathrm{d}}{\mathrm{d}z}\left[e_0\sqrt{1-z}\frac{\mathrm{d}U}{\mathrm{d}z} + (\tan\theta - \tfrac{1}{2}e_0)P\right], \tag{4.13}$$

$$-ikc\mathrm{Fr}^2\frac{\mathrm{d}W}{\mathrm{d}z} = k^p + 2k^2\tan\theta\sqrt{1-z}\frac{\mathrm{d}W}{\mathrm{d}z} - \frac{\mathrm{d}}{\mathrm{d}z}\left[\tan\theta\sqrt{1-z}\left(\frac{\mathrm{d}^2W}{\mathrm{d}z^2} + k^2W\right)\right], \tag{4.14}$$

$$ikc\mathrm{Fr}^2W = -\frac{\mathrm{d}P}{\mathrm{d}z} - \tan\theta\sqrt{1-z}\left(\frac{\mathrm{d}^2W}{\mathrm{d}z^2} + k^2W\right) + \frac{\mathrm{d}}{\mathrm{d}z}\left[2\tan\theta\sqrt{1-z}\frac{\mathrm{d}W}{\mathrm{d}z}\right], \tag{4.15}$$

where

$$e_0 = \frac{(\mu_2 - \tan\theta)(\tan\theta - \mu_1)}{\mu_2 - \mu_1} \tag{4.16}$$

is the only parameter that depends upon the precise functional form of the $\mu(I)$ law.

The no-slip boundary condition at $z = 0$ becomes

$$U(0) = W(0) = \frac{\mathrm{d}W}{\mathrm{d}z}(0) = 0, \tag{4.17}$$

and the three linearised stress-free boundary conditions are

$$\tan\theta W - ikc\left(e_0\sqrt{1-z}\frac{\mathrm{d}U}{\mathrm{d}z} + (\tan\theta - \tfrac{1}{2}e_0)P\right) \to 0, \tag{4.18}$$

$$\sqrt{1-z}\left(\frac{\mathrm{d}^2W}{\mathrm{d}z^2} + k^2W\right) \to 0, \tag{4.19}$$

$$W - ikc\left(P - 2\tan\theta\sqrt{1-z}\frac{\mathrm{d}W}{\mathrm{d}z}\right) \to 0, \tag{4.20}$$

as $z \to 1$. In order to derive the stress boundary conditions in the above form, we make use of the linearised kinematic boundary condition,

$$ikcs_0 = W(1), \tag{4.21}$$

to eliminate the amplitude of the free surface perturbation, s_0.

It is important to note that the downslope velocity perturbation, U, is coupled to the other variables only in the downslope component of the momentum equation, and in the corresponding component of the stress condition of the free surface. Instead of solving the full sixth-order problem defined by (4.13)-(4.20), it is possible to solve a fourth-order subsystem (for P, W and the complex wave speed, c), obtained by omitting equations (4.13) and (4.18). If needed, U may then be calculated by substituting the solution for P, W and c into the omitted equations.

4.1.3 Long-wave asymptotic result

Although even the fourth-order subsystem is sufficiently complicated to require numerical solution for arbitrary wavenumber, it is possible to make analytical progress by performing an asymptotic expansion of the solution for small k. Following a procedure similar to that of Yih (1967), we expand the pressure, velocity and complex wave velocity in the form

$$\begin{aligned} P &= P_0(z) + (ik)P_1(z) + (ik)^2P_2(z) + \ldots, \\ W &= W_0(z) + (ik)W_1(z) + (ik)^2W_2(z) + \ldots, \\ c &= c_0 + (ik)c_1 + (ik)^2c_2 + \ldots . \end{aligned} \tag{4.22}$$

To leading order in k, (4.14), (4.17) and (4.19) combine to give $W_0 \equiv 0$. Equation (4.15) then requires that $P_0 = \text{const}$. No conclusions may yet be drawn regarding c_0, which does not appear at the leading order.

At $O(k)$, we similarly conclude that $W_1 \equiv 0$ and that $P_1 = \text{const}$. However, we now receive useful information from (4.20), which requires that $c_0 P_0 = 0$. Without loss of generality, we assume that $P_0 \neq 0$, and conclude that $c_0 = 0$.

Expanding (4.14) to $O(k^2)$, and applying boundary conditions (4.17) and (4.19) yields

$$W_2(z) = \frac{2P_0}{15\tan\theta}\left[2(1-z)^{5/2} + 5z - 2\right]. \tag{4.23}$$

Substitution of this result into (4.20) then furnishes an expression for the leading-order wave velocity,

$$c_1 = \frac{2}{5\tan\theta}. \tag{4.24}$$

One can conclude, therefore, that the growth rate of a long-wave transverse instability is

$$\sigma = -\mathrm{Im}(c)k = -\frac{2k^2}{5\tan\theta} + O(k^3), \tag{4.25}$$

which is strictly negative for flow at all inclination angles, and independent of the details of the rheology. An incompressible $\mu(I)$ rheology is therefore unable to predict the longitudinal vortex instability observed by Börzsönyi *et al.* (2009) for dense granular flows, at least for long wavelengths.

It is worth mentioning at this point that the asymptotic result of a negative growth rate with quadratic dependence on small wavenumber is also predicted by the Navier-Stokes equations for a viscous fluid, though of course the prefactor differs between the two cases.

4.1.4 Numerical approach

In the experiments of Börzsönyi *et al.* (2009), it was observed that the wavelength of the free-surface disturbances was similar to the depth of the flowing granular layer. Given that we scale lengths with the layer depth, h, this corresponds to a dimensionless wavenumber of order unity. It is entirely possible that the $\mu(I)$ model may lead to instability for such moderate wavenumbers, but not for wavenumbers in the range of validity of the asymptotic result in §4.1.3. We shall therefore solve the linear problem numerically to obtain the full dispersion relation, $c(k)$, which links the complex wave velocity to the wavenumber.

The solution may be found using a standard numerical solver for boundary value problems[1], but the singularity at $z = 1$ requires some careful treatment. In particular, the

[1] We use the bvp4c routine from Matlab.

boundary condition (4.19) is difficult to implement in a form appropriate to the numerical routine. However, one can consider a slightly different problem, which has the same solution, but avoids any singularities in the momentum equations.

Suppose that, instead of setting $\bar{p}^* = 0$ on the free surface, we equate the pressure to $\bar{p}^* = \rho g d$, approximately corresponding to the weight of a monolayer of grains on the surface of the granular medium. The dimensionless steady pressure and velocity distributions would then become

$$\bar{p}(z) = 1 + \delta - z, \tag{4.26}$$

$$\bar{u}(z) = \frac{2}{3}\left[(1+\delta)^{3/2} - (1+\delta-z)^{3/2}\right], \tag{4.27}$$

where $\delta = d/h$. The resulting linearised equations will then resemble (4.13)-(4.14), except with the factors of $\sqrt{1-z}$ replaced by $\sqrt{1+\delta-z}$. Crucially, if $\delta > 0$, the modified equations are regular for $0 \leq z \leq 1$. Furthermore, the linearised stress-free boundary conditions are modified to

$$ikc\delta W - \sqrt{\delta}\tan\theta\left(\frac{\mathrm{d}^2 W}{\mathrm{d}z^2} + k^2 W\right) = 0, \tag{4.28}$$

$$W - ikc\left(P - 2\sqrt{\delta}\tan\theta\frac{\mathrm{d}W}{\mathrm{d}z}\right) = 0, \tag{4.29}$$

on $z = 1$.[2] The boundary conditions (4.28) and (4.29) are much easier to implement numerically, as are the (regular) modified differential equations.

In order to connect this modified problem to the true problem (with $\bar{p} = 0$ on $z = 1$), one needs only to take the limit $\delta \to 0$. It can be readily demonstrated that this limit is regular by repeating the asymptotic calculation of §4.1.3 for the modified problem, in which case one finds the growth rate to be

$$\sigma = -\frac{4}{15\tan\theta}\left[\delta^{3/2}(1-5\delta) + \tfrac{5}{2}(1+\delta)^{3/2} - \tfrac{15}{2}\delta\sqrt{1+\delta} + 5\delta(1+\delta)^{3/2} - (1+\delta)^{5/2}\right]k^2 + O(k^3). \tag{4.30}$$

This result reduces to (4.25) in the limit as $\delta \to 0$.

A typical dispersion relation calculated numerically from the modified equations is shown in Figure 4.3. The numerics agree excellently with the asymptotic result as $k \to 0$, vindicating our decision to use the modified pressure boundary condition for the numerical approach.

[2] We omit the condition analogous to (4.18), as it has no bearing on the stability analysis.

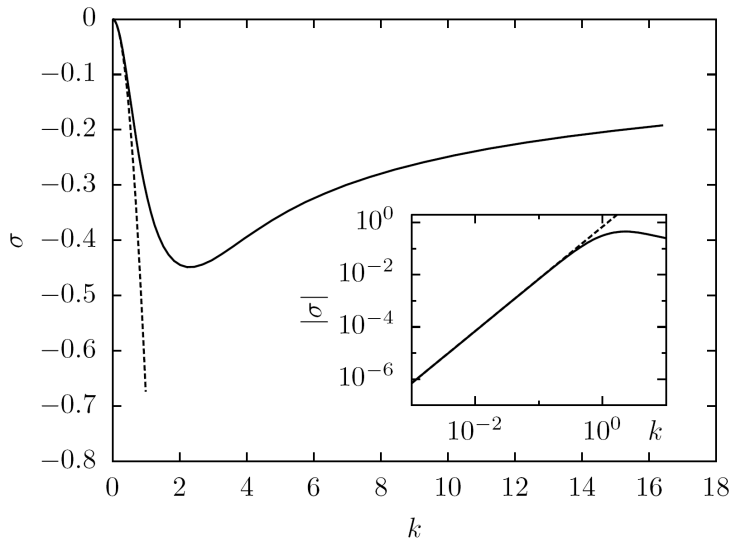

Figure 4.3: Growth rate of transverse perturbation of wavenumber k in for flow on an inclined plane. A typical numerical result ($\mathrm{Fr}^2 = 0.1$, $\theta = \pi/6$) is shown as a solid line, whilst the asymptotic result (4.25) shown as a dashed line. The inset shows a the same plot on a logarithmic scale, to make clear the long-wave behaviour. The wave speed, $-\mathrm{Re}(c)$, is not shown, but is zero for all k.

As k increases, we observe that the growth rate reaches a minimum, whose location depends upon Fr, but is typically near $k = 2$. After this minimum, the growth rate increases toward zero for very short waves. At no point does the growth rate ever become positive, thus we are forced to conclude that the $\mu(I)$ model is insufficient to describe the formation of longitudinal vortices.

4.1.5 Discussion

It is clear from the results in §§4.1.3-4.1.4 that the instability giving rise to longitudinal vortices is not captured by the $\mu(I)$ rheological law studied in this thesis. In order to explain this phenomenon, it will be necessary to introduce new physical effects into the model.

In their numerical study, Börzsönyi *et al.* (2009) note that the packing fraction in the

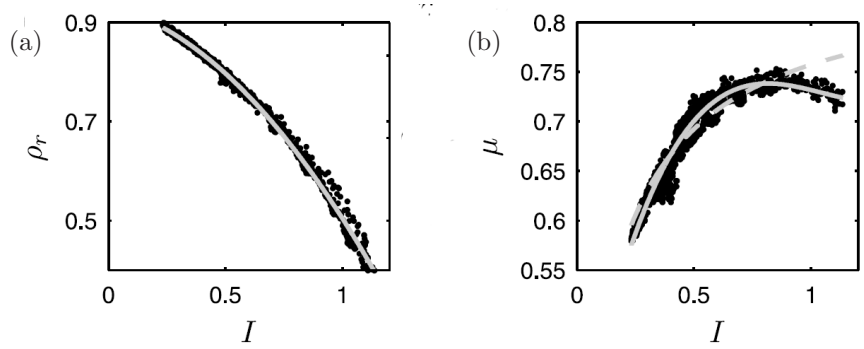

Figure 4.4: Dilation in the longitudinal vortex problem. The simulations of Börzsönyi *et al.* (2009) provide relationships between the local inertia number and: (a) the relative density, ρ_f, which is linearly related to the packing fraction; (b) the effective friction coefficient. The solid lines in (a) and (b) show the best quadratic and cubic fits respectively, whilst the dashed line in (b) shows the best fit of the form (1.11) for $\mu(I)$.

granular layer varies significantly after the formation of the longitudinal vortices. Typically, they observe that a relatively sparsely-packed region forms at the base of the flow beneath the peaks, with a densely-packed region forming directly above it. One might argue that the density difference between these two regions could drive an upwelling flow toward the peaks, but this seems to contradict the velocity distribution shown in Figure 4.2(b), which suggests that material flows in the opposite vertical direction. Nevertheless, the variation in packing fraction is likely to have an noticeable effect on the dynamics of the longitudinal vortex flow, so it would be interesting to repeat the stability analysis using a compressible analogue of the $\mu(I)$ rheology. As with compressible fluid flow, it would be necessary to include an equation of state relating the local density and pressure. To obtain such a relationship, we recall the observation of the GDR MiDi (2004), who noted that the volume fraction is a function only of the local inertia number. Using either the linear dependence determined by the GDR MiDi (2004), or the quadratic fit obtained from the simulations by Börzsönyi *et al.* (2009) (and reproduced in Figure 4.4a), it would be a relatively simple matter to derive an appropriate equation of state.

Despite having an equation of state in hand, further difficulty would arise when con-

sidering the rheological law. Just as a fluid's viscosity may vary depending on the local density, so too might the granular shear stresses depend upon the local packing fraction. To our knowledge, no detailed model for the rheology of compressible granular flows has yet been proposed. Although one might hope that the $\mu(I)$ relationship may be used without modification, coarse-graining of the DEM simulation by Börzsönyi *et al.* (2009) shows that the form of the function (shown in Figure 4.4b) is quite different to that used throughout this thesis. In particular, Börzsönyi *et al.* (2009) found that the effective friction coefficient for longitudinal vortices is not a monotonic increasing function of I, and actually decreases for $I \gtrsim 0.7$. Börzsönyi *et al.* (2009) argue that this turnover is crucial to the instability. Where the flow is deepest (that is, beneath the peaks), the inertia number may become sufficiently large to cause the effective friction to decrease, leading to stronger fluidisation near the rough plane under the peaks. Near the free surface, the friction coefficient remains relatively high, and a plug region develops. This plug can move more quickly over the fluidised region beneath than the shearing material in the neighbourhood of the troughs, further enhancing the disparity in inertia number, thus leading to instability.

Although Börzsönyi *et al.* (2009) consider the use of a $\mu(I)$ type rheology to model the stress-strain relationship in their DEM simulations, they note that some of the key assumptions used in its derivation are not met by the flow. Aside from dilation, they note that the deviatoric stress tensor is not always aligned with the rate of strain tensor. They conclude, as shall we, that it may be necessary to use a more complicated rheology to fully understand the mechanism giving rise to longitudinal vortices in dense granular flow.

4.2 Growth of sand dunes

The study of sand dunes and related phenomena has commanded the attention of a great many researchers over the years. Indeed, it was a set of field observations of desert sand dune migration that inspired Bagnold (1954) to complete his seminal work, which in turn set into motion much of the modern physical study of flowing granular materials. Since Bagnold, many attempts have been made to model the growth and evolution of dunes (e.g. Kennedy, 1969; Hoyle & Woods, 1997; Prigozhin, 1999), but few models have come close to being able to reproduce the myriad landforms observed in a natural setting. Nevertheless,

efforts appear to be converging. Kroy *et al.* (2002) have proposed a 'minimal model' for dune migration: one they hope contains sufficient physical detail to describe properly the behaviour of several different types of sand dunes.

A common feature of previous models for dune dynamics is that they do not use a full rheological model for the granular material. Instead, they use an empirical transport law to govern the flux of grains depending on the stresses imposed by the fluid flow above the dune. One might speculate that the inclusion of a proper constitutive model for the grain flux (or *bedload transport*, as it is referred to in much of the literature) might improve the quality of the various models available. In this section, we shall apply the $\mu(I)$ rheology as part of a two-layer model for sand dunes, in the hope of correctly modelling the instability that gives rise to dune formation.

The problem of dune formation is simple to describe. A viscous, often turbulent fluid overlies a granular bed of arbitrary depth[3]. If the mean fluid flow is sufficiently large, the surface of the granular bed will be unstable to long-wave disturbances. It is this instability that leads ultimately to the formation of large sand dunes. The mechanism driving the instability was first studied by Kennedy (1969), who considered only an inviscid fluid, and used a heuristic rule linking the slip velocity on the granular interface to the grain flux. He found that instability could occur, but only if the fluid velocity disturbance was allowed to lag behind the peaks of the disturbed interface. Although Kennedy (1969) chose to apply an artificial phase shift to the velocity disturbance in his inviscid calculation, a more recent work by Charru & Hinch (2000) showed that in fact a lag arises naturally between the disturbed interface and the fluid shear stress exerted there in the case of two viscous fluid layers. Advection of the vorticity disturbance creates a small out-of-phase vorticity component, which in turn gives rise to an out-of-phase shear stress, the maximum of which occurs slightly upstream of the crests. In a later work, Charru & Hinch (2006*a*) tackle the case of a granular lower layer by proposing a more sophisticated model for the erosion and deposition of grains due to the fluid shear stress. Although this model captures many of the qualitative features of dune formation seen in experiments, it relies on the selection of some

[3]Certain types of dunes, particularly the parabolic barkhan dunes, are observed only when the layer of sand above the bedrock is relatively thin. We shall not consider the case of a finite granular bed, but expect that the linear stability analysis would remain essentially unchanged.

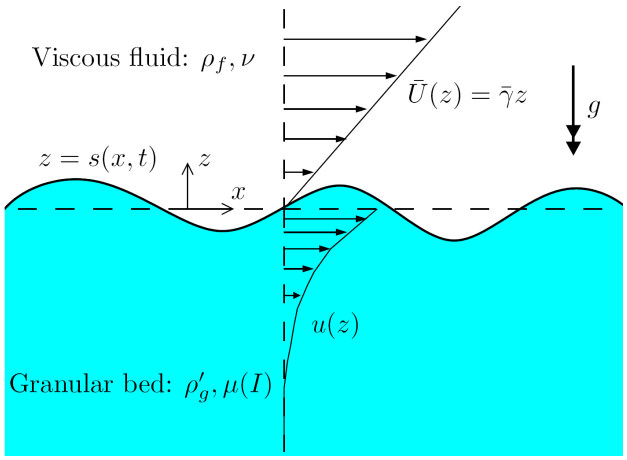

Figure 4.5: Diagram and notation for the two-layer dune model. The granular velocity $u(z)$ has been greatly exaggerated for clarity – the velocity should in fact be continuous across the interface.

free dimensionless parameters, limiting the power of any quantitative predictions. Instead of the transport model of Charru & Hinch (2006*a*), we shall use the $\mu(I)$ constitutive law to describe the motion of the granular bed, and develop a linear stability analysis to investigate whether its predictions are also compatible with the experimental data.

4.2.1 Two-layer model

In order to investigate the predictive power of the $\mu(I)$ rheology for the problem of dune formation, we carry out a linear stability analysis of a granular bed driven from above by fluid shear. We describe the problem using a two-layer model, perform a long-wave stability analysis asymptotically, and solve numerically for the full dispersion relation, before moving to discuss the results in §4.2.5. Throughout this section, an asterisk denotes a dimensional variable.

Consider the situation sketched in Figure 4.5. A granular bed of arbitrary depth is overlain by a deep viscous fluid of density ρ_f and kinematic viscosity ν. The interface between the fluid and granular media is positioned at $z^* = s^*(x^*, t^*)$. We suppose that the fluid flow takes the form of the linear shear $\bar{\gamma} z^*$ near the interface, and dispense with any

details of the outer flow, even though the outer flow may be complicated, even turbulent in the case of atmospheric flow over a desert sand dune (see Hunt *et al.*, 1988, for example). In this case, one might like to think of the linear velocity as being the first term in a power series expansion of the full velocity about the undisturbed surface, $z^* = 0$. The granular medium has reduced density[4] ρ_g' and is assumed to obey the $\mu(I)$ constitutive law. We nondimensionalise the variables according to

$$\begin{aligned}
\text{Length} \qquad & (x^*, z^*) = d(x, z), \\
\text{Fluid velocity} \qquad & (U^*, W^*) = d\bar{\gamma}(U, W), \\
\text{Granular velocity} \qquad & (u^*, w^*) = \sqrt{dg}(u, w), \\
\text{Time} \qquad & t^* = \sqrt{d/g}\, t, \\
\text{Fluid pressure} \qquad & P^* = \rho_f g d\, P, \\
\text{Granular pressure} \qquad & p^* = \rho_g' g d\, p.
\end{aligned} \tag{4.31}$$

Under these scalings, the dimensionless Navier-Stokes equations for the fluid layer take the form

$$\epsilon \partial_t U + U \partial_x U + W \partial_z U = -\epsilon^2 \partial_x P + \frac{1}{\mathrm{Re}_d} \nabla^2 U, \tag{4.32}$$

$$\epsilon \partial_t W + U \partial_x W + W \partial_z W = -\epsilon^2 \partial_z P - \epsilon^2 + \frac{1}{\mathrm{Re}_d} \nabla^2 W, \tag{4.33}$$

$$\partial_x U + \partial_z W = 0. \tag{4.34}$$

The boundary conditions to be applied on the granular interface are the familiar kinematic and no-slip conditions

$$\epsilon \partial_t s + \mathbf{U} \cdot \nabla s = \mathbf{U} \cdot \hat{\mathbf{z}} \qquad \text{on } z = s(x, t), \tag{4.35}$$

$$\mathbf{U} = \epsilon \mathbf{u} \qquad \text{on } z = s(x, t), \tag{4.36}$$

The remaining fluid boundary condition then depends on the flow configuration. If the fluid is unbounded, we assume that the linear shear is a local approximation to a more complicated flow near $z = 0$, and employ a matching condition,

$$\mathbf{U} \to z\hat{\mathbf{x}} \quad \text{as } z \to \infty. \tag{4.37}$$

[4] We use the reduced density, $\rho_g' = \rho_g - \rho_f$, in order to account for any buoyancy afforded to the grains by the fluid invading the porous granular bed.

If the fluid is instead driven by a rigid plate located at $z = h$, as in the experiments of Charru & Mouilleron-Arnould (2002) and the theoretical model of Charru & Hinch (2006*a*), then we must apply a no-slip condition there, thus

$$\mathbf{U} = z\hat{\mathbf{x}} \quad \text{at } z = h. \tag{4.38}$$

The dimensionless groups arising in the fluid equations (4.32)-(4.38) are the familiar fluid Reynolds number and a quantity reminiscent of a reciprocal Froude number, both based on the grain lengthscale:

$$\mathrm{Re}_d = \frac{\bar{\gamma} d^2}{\nu} \qquad \text{and} \qquad \epsilon = \frac{1}{\bar{\gamma}}\sqrt{\frac{g}{d}} = \left(\frac{d\bar{\gamma}}{\sqrt{gd}}\right)^{-1}. \tag{4.39}$$

The parameter ϵ may be usefully thought of as a ratio between the timescales associated with the fluid and granular velocities. The various experimental observations all concur that bed deformation occurs on a much slower timescale than the fluid flow, so we may assume that ϵ is very small. In fact, we take the formal limit $\epsilon \to 0$. In this case, we may regard the granular bed as a fixed, rigid boundary when calculating the fluid flow, with any disturbance to the fluid flow caused only by the wavy shape of the disturbed interface, $z = s(x,t)$. In terms of the governing equations, we omit the time derivatives in (4.32) and (4.33), and replace the velocity boundary conditions with the no-slip condition

$$\mathbf{U} = 0 \qquad \text{on } z = s(x,t). \tag{4.40}$$

The continuum equations for the granular bed, using the $\mu(I)$ constitutive law, are

$$\partial_t u + u\partial_x u + w\partial_z u = -\partial_x p + \frac{\partial}{\partial x}\left(\frac{\mu(I)p}{|\dot{\gamma}|}\dot{\gamma}_{xx}\right) + \frac{\partial}{\partial z}\left(\frac{\mu(I)p}{|\dot{\gamma}|}\dot{\gamma}_{xz}\right), \tag{4.41}$$

$$\partial_t w + u\partial_x w + w\partial_z w = -\partial_z p - 1 + \frac{\partial}{\partial x}\left(\frac{\mu(I)p}{|\dot{\gamma}|}\dot{\gamma}_{zx}\right) + \frac{\partial}{\partial z}\left(\frac{\mu(I)p}{|\dot{\gamma}|}\dot{\gamma}_{zz}\right), \tag{4.42}$$

$$\partial_x u + \partial_z w = 0. \tag{4.43}$$

The granular deformation is driven by shear stresses imposed by the fluid flow on the surface of the bed. On the interface, in addition to the usual kinematic boundary condition, we must ensure that the shear stress is continuous between the two media, and apply a boundary

condition on the granular pressure there,

$$\partial_t s + \mathbf{u}\cdot\nabla s = w \qquad \text{on} \quad z = s(x,t), \tag{4.44}$$

$$\mathbf{n}\wedge(\sigma_{\mathbf{g}}\cdot\mathbf{n}) = \frac{\rho_f}{\rho_g'}\frac{1}{\epsilon^2\mathrm{Re}_d}\mathbf{n}\wedge(\sigma_{\mathbf{f}}\cdot\mathbf{n}) \qquad \text{on} \quad z = s(x,t), \tag{4.45}$$

$$p = 1 \qquad \text{on} \quad z = s(x,t). \tag{4.46}$$

The pressure condition (4.46) used here represents the weight of a monolayer of grains at located at the interface. Using the more natural condition, $p = 0$, forces the granular stresses to vanish on the interface, meaning that nothing may balance the fluid stress on the right-hand-side of (4.45), prohibiting a steady solution. In order to ensure that a steady solution exists, we adopt (4.46) as a reasonable alternative.

Ordinarily, we would use the friction rule

$$\mu(I) = \mu_1\left(1 + \frac{\lambda - 1}{1 + I_0\sqrt{p}/|\dot{\gamma}|}\right). \tag{4.47}$$

However, this rule, combined with the pressure and shear stress applied from above, ought to result in a steady velocity profile similar to that calculated previously for the dragged plate problem in §2.2.1, which we recall had a rather complicated analytical expression. Using such a velocity profile renders even the linear stability problem difficult to solve analytically. Instead, we opt to use a linearised form of the $\mu(I)$ law,

$$\mu(I) = \mu_1\left(1 + b\frac{|\dot{\gamma}|}{\sqrt{p}}\right) \tag{4.48}$$

Using this form will modify the granular mean flow, but otherwise should not affect the stability analysis, which will require the $\mu(I)$ function to be linearised about the steady state anyway. Furthermore, it is a reasonable approximation to the 'exact' form (4.47) provided that the stresses exerted on the grains are not too large.

As with the dragged plate problem of §2.2.1, the presence of a yield stress in the rheology means that the granular material will only shear in a layer of finite depth, $-z_b < z < 0$. The depth, z_b, is unknown *a priori*, and must be determined as part of the solution (cf. §2.2.1). This is achieved by applying the boundary conditions

$$\mathbf{u} = 0 \quad \text{and} \quad |\dot{\gamma}| = 0 \quad \text{at } z = -z_b. \tag{4.49}$$

4.2.2 Linearisation

The field equations (4.32)-(4.34) and (4.41)-(4.43), together with the boundary conditions (4.37) (or 4.38), (4.40), (4.44)-(4.46) and (4.49) completely define the two-layer model used to study dune formation. Seeking a unidirectional steady solution with a flat interface, $s(x,t) = 0$, we find the basic velocity and pressure profiles in each layer

$$\bar{U}(z) = z, \tag{4.50}$$

$$\bar{P}(z) = z_0 - z, \tag{4.51}$$

$$\bar{u}(z) = \frac{2}{3b}\left(2\theta^{3/2} - 3\theta\sqrt{1-z} + (1-z)^{3/2}\right), \tag{4.52}$$

$$\bar{p}(z) = 1 - z, \tag{4.53}$$

where z_0 is an arbitrary fluid reference pressure that we shall henceforth take to be zero without any loss of generality. In the expression (4.52) for $\bar{u}(z)$, the quantity

$$\theta = \frac{1}{\mu_1}\frac{\rho_f}{\rho'_g}\frac{1}{\epsilon^2 \mathrm{Re}_d} \tag{4.54}$$

is the ratio of the applied shear stress (from 4.45) to the yield stress, $\mu_1 p$, at the interface. This definition of θ, except for the appearance of μ_1, coincides with that of the Shields number (Raudkivi, 1998), which is a ratio between the hydrodynamic force acting on a particle, $\rho_f \nu \bar{\gamma} d^2$, and its effective weight, which is of order $\rho'_g g d^3$. In previous studies, including Charru & Hinch (2006*a*), it is asserted that particles may only be set in motion if the Shields number exceeds a certain threshold value. Our analysis suggests that motion is possible only if the stress exerted on the granular bed is above yield – that is, $\theta > 1$. The friction coefficient, μ_1, is therefore equivalent to the threshold Shields number in this case. Furthermore, application of the boundary condition (4.49) implies that the depth of the flowing granular layer is $z_b = 1 - \theta$.

Having established the steady state (4.50)-(4.53), we consider the effect of making a small sinusoidal perturbation to the interface, thus

$$s = s_0 e^{ik(x+ct)}. \tag{4.55}$$

We linearise the governing equations about the steady state, making use of perturbation

streamfunctions, Ψ and ψ, such that

$$U = \bar{U}(z) - \Psi'(z)e^{ik(x+ct)}, \qquad W = ik\Psi(z)e^{ik(x+ct)}, \tag{4.56}$$

$$u = \bar{u}(z) - \psi'(z)e^{ik(x+ct)}, \qquad w = ik\psi(z)e^{ik(x+ct)}, \tag{4.57}$$

(note that $' \equiv \mathrm{d}/\mathrm{d}z$) in order to automatically satisfy the continuity equations (4.34) and (4.43). We also define pressure perturbations, Φ and ϕ, such that

$$P = \bar{P}(z) + \Phi(z)e^{ik(x+ct)} \quad \text{and} \quad p = \bar{p}(z) + \phi(z)e^{ik(x+ct)}. \tag{4.58}$$

In order to obtain the linearised fluid equations, one can take the curl of the momentum equations (4.32)-(4.33), with $\epsilon = 0$, to arrive at an equation of Orr-Sommerfeld type

$$ikz\left(\Psi'' - k^2\Psi\right) = \frac{1}{\mathrm{Re}_d}\left(\Psi'''' - 2k^2\Psi'' + k^4\Psi\right), \tag{4.59}$$

with linearised boundary conditions

$$\Psi(x,0) = 0, \tag{4.60}$$

$$\Psi'(x,0) = s_0, \tag{4.61}$$

$$\Psi, \Psi' \to 0 \quad \text{as } z \to \infty, \tag{4.62}$$

or

$$\Psi, \Psi' = 0 \quad \text{at } z = h, \tag{4.63}$$

depending on whether we apply condition (4.37) or (4.38) at the top of the fluid region.

The equations (4.59)-(4.62) describe the perturbation to a uniform shear flow due to the presence of a (fixed) wavy bottom boundary, as calculated by Benjamin (1959). The solution, or rather some useful asymptotic results based upon it, will be discussed shortly. Now, however, we press on to present the linearised granular momentum equations, which do not lend themselves to such an elegant expression as (4.59). After some effort, the equations may be reduced to

$$ik[(c+\bar{u})\psi' - \bar{u}'\psi] = ik\phi - 2\mu_1\theta k^2\frac{\psi'}{\bar{u}'} + \mu_1 b\left[\sqrt{1-z}\left(\psi'' + k^2\psi - \bar{u}''\phi\right)\right]', \tag{4.64}$$

$$k^2(c+\bar{u})\psi = \phi' + ik\mu_1 b\left[\sqrt{1-z}\left(\psi'' + k^2\psi - \bar{u}''\phi\right)\right] - 2ik\mu_1\theta\left[\frac{\psi'}{\bar{u}'}\right]', \tag{4.65}$$

with linearised boundary conditions

$$b(\psi'' + k^2\psi) - \tfrac{1}{2}(\theta+1)\,\phi = \theta\Psi'' \quad \text{at } z = 0, \tag{4.66}$$

$$\phi(0) = s_0, \tag{4.67}$$

$$[c + \bar{u}(0)]\phi(0) - \psi(0) = 0, \tag{4.68}$$

$$\psi(1-\theta) = 0, \tag{4.69}$$

$$\psi'(1-\theta) = 0. \tag{4.70}$$

4.2.3 Asymptotic solutions in the long wave limit

Equations (4.59)-(4.70) together form the full linear problem to be solved in order to calculate the dispersion relation, $c(k)$. While an exact solution is elusive, it is possible to make progress by considering the asymptotic limit of small wavenumber (hence large wavelength). This limit is, conveniently, the limit appropriate to dune formation: desert dunes typically have horizontal lengthscales of order 100m, whilst grains of sand are around a millimetre in size. The asymptotic approach begins with the solution of the fluid problem (4.59)-(4.63). Because of the $\epsilon \to 0$ limit taken earlier, the fluid equations may be solved independently of the motion of the granular bed. The calculation, however, depends upon the boundary condition applied at the top of the fluid region. If the fluid is unbounded, then a thin boundary layer dictates the shear stress exerted on the granular bed. If instead the fluid is driven by a rigid lid, then one must consider the 'shallow viscous' limit of Charru & Hinch (2000).

The boundary layer regime

In the case of an unbounded fluid, with $k \ll \mathrm{Re}_d$, it is possible to solve the Orr-Sommerfeld problem asymptotically via the method of matched asymptotic expansions (see Drazin & Reid, 1981, for example – the details are not presented here). In order to solve the granular problem, we need know only the shear stress resulting from the fluid velocity perturbation, which is

$$\Psi''(0) \sim -\frac{3^{1/3}}{\Gamma(2/3)}(ik)^{1/3}\mathrm{Re}_d^{1/3}s_0 \equiv -\mathcal{M}s_0(ik)^{1/3} \quad \text{as} \quad \frac{\mathrm{Re}_d}{k} \to \infty. \tag{4.71}$$

This expression also arises as the 'boundary layer' limit of the two-fluid calculation by Charru & Hinch (2000).

In order to find the asymptotic solution for the granular layer in this long wave limit, we adopt a procedure similar to that of Yih (1967). We posit the expansions

$$\phi = \phi_0 + (ik)^{1/3}\phi_1 + (ik)^{2/3}\phi_2 + ..., \tag{4.72}$$

$$\psi = \psi_0 + (ik)^{1/3}\psi_1 + (ik)^{2/3}\psi_2 + ..., \tag{4.73}$$

$$c = c_0 + (ik)^{1/3}c_1 + (ik)^{2/3}c_2 + ..., \tag{4.74}$$

substitute into (4.64)-(4.70) and solve at successive orders in $k^{1/3}$. The solution at $O(k^0)$ is

$$\phi_0 = s_0, \tag{4.75}$$

$$\psi_0 = s_0\bar{u}(z), \tag{4.76}$$

$$c_0 = 0. \tag{4.77}$$

At $O(k^{1/3})$, we obtain

$$\phi_1 = 0, \tag{4.78}$$

$$\psi_1 = -\frac{2\theta\mathcal{M}s_0}{3b}\left[2(1-z)^{3/2} - 3\theta^{1/2}(1-z) + \theta^{3/2}\right], \tag{4.79}$$

$$c_1 = -\frac{2\theta\mathcal{M}}{3b}\left[\theta^{3/2} - 3\theta^{1/2} + 2\right], \tag{4.80}$$

and so the growth rate is

$$\sigma = -\mathrm{Im}(c)k = \frac{\theta\mathcal{M}}{3b}\left[\theta^{3/2} - 3\theta^{1/2} + 2\right]k^{4/3} + O(k^2). \tag{4.81}$$

The corresponding phase speed of these long surface waves is

$$\mathrm{Re}(c) = -\frac{\theta\mathcal{M}}{b\sqrt{3}}\left[\theta^{3/2} - 3\theta^{1/2} + 2\right]k^{1/3} + O(k). \tag{4.82}$$

The observant reader will have noted the omission of the ord($k^{5/3}$) term in (4.81) and the ord($k^{2/3}$) term in (4.82). This is intended – continuing the asymptotic analysis to the next few orders reveals that no additional contribution appears until $O(k^1)$, where more terms of (4.64) and (4.65) are able to play a role.

By inspection of the leading order expression for the growth rate (4.81), we see that the growth rate is positive for $\theta > 1$, and zero (to $O(k^2)$) at $\theta = 1$. Recalling that $\theta > 1$ is

precisely the condition needed in order to the granular bed to be above yield, we conclude that the granular bed is unstable to long wavelength disturbances as long as there is some non-zero bedload transport.

The shallow viscous limit

The 'boundary layer' asymptotic regime described above is appropriate to desert sand dunes. The experiments of Charru & Mouilleron-Arnould (2002), however, are carried out in an annular channel, in which the fluid shear is driven by the motion of a rigid lid. In this case, the boundary layer limit (4.71) is no longer appropriate. Instead, the correct limit to use is the 'shallow viscous' limit considered by Charru & Hinch (2006*a*), which gives

$$\Psi''(0) \sim -\frac{4s_0}{h} - \frac{h^2 \mathrm{Re}_d z_0}{30}(ik) \quad \text{as} \quad kh \to 0, \tag{4.83}$$

where h is the undisturbed depth of the fluid layer. With this new limiting form for the shear stress, it is still possible to perform an asymptotic analysis of the granular problem. However, on this occasion it is necessary to expand the pressure, streamfunction, and complex wave speed in powers of (ik), rather than $(ik)^{1/3}$. The resulting expression for the phase speed is

$$\mathrm{Re}(c) = -\frac{8\theta}{3bh}\left(2 - 3\sqrt{\theta} + \theta^{3/2}\right) + O(k^2), \tag{4.84}$$

and the growth rate is

$$\sigma = A(\mu_1, h, b, \theta)k^2 + O(k^3), \tag{4.85}$$

where $A(\mu_1, h, b, \theta)$ is a rather complicated nonlinear function in each of its arguments. The precise form of the function $A(\mu_1, h, b, \theta)$ is not particularly illuminating, and as such is not given explicitly. It is important to note, however, that $\sigma > 0$ whenever $\theta > 1$, just as in the boundary layer case. Once again, the $\mu(I)$ rheology predicts that any flow involving motion of the granular bed is unstable to a long-wave dune forming instability.

4.2.4 Numerical solution

In addition to solving the linearised equations asymptotically for small k, it is possible to solve the equations numerically for arbitrary k. As in our study of longitudinal vortex formation, the inhomogeneous pressure boundary condition on the free surface eliminates

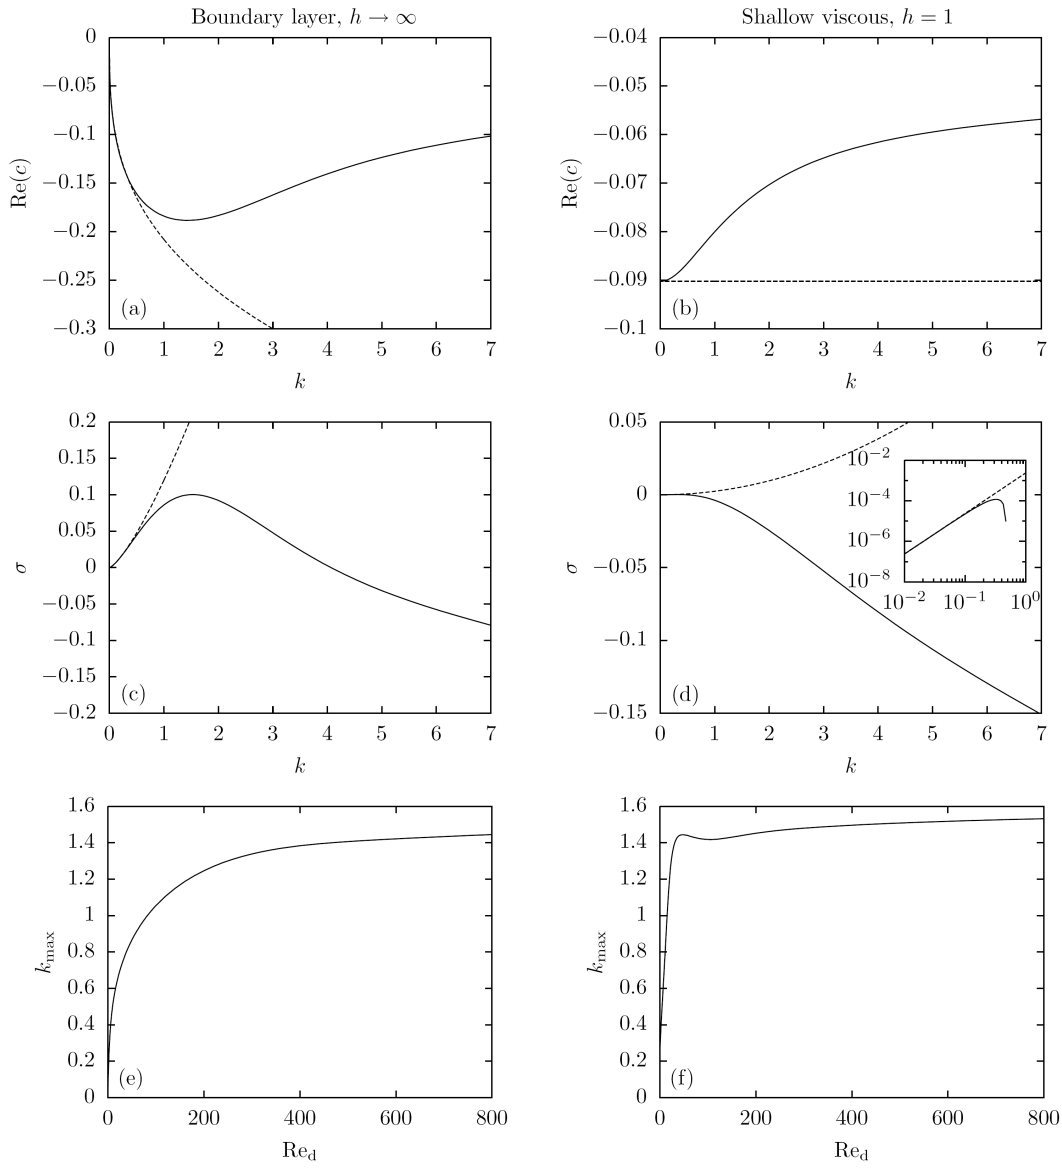

Figure 4.6: Numerical results for dune formation in the boundary layer (panels a, c, and e) and shallow viscous (panels b, d and f) limits. In panels (a-d), the numerical solution (solid curves) is compared with the appropriate asymptotic limit from §4.2.3. Panels (a,b) show the phase speed, whilst panels (c,d) show the growth rate. The inset in panel (d) shows the same data as the main plot on logarithmic axes, to highlight the long-wave limiting behaviour. Panels (e) and (f) show the variation of the least stable wavenumber with the fluid Reynolds number in each limiting case.

any potential singularity at the free surface, so it is possible to solve the linearised system using an off-the-shelf routine, such as the bvp4c routine in Matlab. In order to explore the parameter space more efficiently, we make use of a crude continuation approach, whereby the solution at one wavenumber is used as an initial guess from which to begin the iteration that solves the stability problem at a slightly different wavenumber. A similar approach may be used to investigate the effect of varying the material parameters, if desired.

One important point to consider when performing the numerical calculation is the handling of the limit $z \to \infty$ in the boundary layer regime. Although it is not straightforward to solve on an infinite interval, the result should be the same (asymptotically, at least) as long as h is sufficiently large. A cursory inspection of the Orr-Sommerfeld equation (4.59), allows one to conclude that the perturbation streamfunction must decay like e^{-kz} as $z \to \infty$. In order to correctly capture this exponential decay, we must therefore choose $h \gg k^{-1}$. Otherwise, the presence of the rigid lid could affect the solution. In practice, we choose $h = 100k^{-1}$.

The key results of the numerical calculations are shown in Figure 4.6. We show the complex wave velocity calculated in both the boundary layer and shallow viscous limits, with the appropriate asymptotic results from §4.2.3 shown for comparison. In all cases, the strong agreement between the asymptotic and numerical results for small wavenumber serves to validate the code. For large wavenumbers, and in both regimes, the growth rate reaches a maximum at $k = k_{\max}$, before decreasing and ultimately becoming negative for all sufficiently large k. This indicates that, although there are always some long-wave disturbances that will grow, short-wave perturbations will always decay.

The least stable wavenumber, $k_{\max}$, is an important quantity to calculate, as it offers a prediction for the lengthscale of structures occurring in natural dunes. In order to calculate $k_{\max}$, we use a standard extremum-finding routine (fminbnd in Matlab), coupled with a crude continuation scheme to explore a range of fluid Reynolds numbers. To begin the continuation, we offer an initial estimate of a specific Reynolds number based on the full dispersion relations calculated previously.

4.2.5 Discussion

Considering both the asymptotic results of §4.2.3 and the numerical dispersion relations determined in §4.2.4, we can conclude that the $\mu(I)$ constitutive law provides a description of the bed motion that is compatible with dune formation. Regardless of the details of the fluid flow, we find that the granular bed is susceptible to a long-wave instability, provided that that stress imparted to the grains in steady state is above yield (i.e. $\theta > 1$). Unlike previous results, which rely either on an algebraic law linking the bed flux to the fluid shear stress (e.g. Kroy *et al.*, 2002) or a separate model for the moving grains (e.g. Charru & Hinch, 2006*a*), this calculation uses the full $\mu(I)$ constitutive law to describe the grain motion, which can be verified in several flow configurations (see Chapter 2, for example).

A common failing of dune formation models is that they tend to underestimate the wavelength of naturally-occurring dunes by several orders of magnitude. The two-layer $\mu(I)$ model used here is no exception. Typically, one might think that dunes should form with approximately the least stable wavelength predicted by the linear stability analysis. However, the numerical results shown in Figure 4.6(e,f) indicate that the least stable wavelength, $k_{\max}$, is of order unity for the moderate to large Reynolds numbers expected in natural dune formation. Restoring dimensions, these results suggest that the least stable wavelength is of a size with individual grain diameters, in contrast to the size of natural dunes, which can be hundreds of metres in length. It is of course entirely possible that the various linear models do not tell the whole story, and that some nonlinear process acts to suppress instability at these short wavelengths, allowing longer structures to develop.

Although the $\mu(I)$ model can capture the dune-forming instability, the predictions do not entirely agree with the experimental observations. In experiments conducted in an annular channel, Charru & Mouilleron-Arnould (2002) found that stable granular motion can exist for a range of Shields numbers. A similar result may be found in the experimental work of Ouriemi (2007), who studied dune formation in a circular pipe. Furthermore, Mouilleron (2002) also observed that the range of stable Shields numbers decreases if the bed is subjected to an oscillating fluid flow of the same magnitude. In the present calculation, the granular bed is either stationary or unstable. There can, therefore, be no reduction of the stable range of Shields numbers upon consideration of oscillating flow using the $\mu(I)$ model. By

contrast, the erosion-deposition model of Charru & Hinch (2006*a*) is capable of predicting both the stable regime in steady flow and the decrease in stability in oscillating flow (Charru & Hinch, 2006*b*).

One can speculate that these failures of the $\mu(I)$ model may occur for several reasons. We note that the experiments of Charru & Mouilleron-Arnould (2002) and Ouriemi (2007) were both conducted underwater. In such a case, one might need to consider the additional effects of particle settling and resuspension in a viscous fluid. Indeed, this is the basis of the erosion-deposition model used by Charru & Hinch (2006*a*), although they too fail to correctly predict the order of magnitude of the least stable wavelengths. It is also true that the layer of granular material seen to be moving in experiments is very thin – only a few grain diameters across. It therefore may not be valid to apply a continuum model to the saltating grains atop a sand dune. This resonates with our findings for unidirectional shear flows in §§2.1.2-2.1.3: when the thickness of the shearing region is small and the flow is very slow, the $\mu(I)$ rheology tends to provide poor predictions for the relationship between the granular motion and its driving stress.

5

Contact in a viscous fluid

This chapter, which is largely unrelated to the rest of this thesis, describes work carried out mainly at the Geophysical Fluid Dynamics Summer Program at the Woods Hole Oceanographic Institution, USA. This work was done under the supervision of Neil Balmforth of the University of British Columbia, but the calculations described are my own. This work has also been presented in the Journal of Fluid Mechanics (Cawthorn & Balmforth, 2010)

The $\mu(I)$ rheological law used throughout the preceding chapters is based upon Coulomb friction, which is a widely-used classical model for the frictional stresses between two solid objects. When considering granular materials, the $\mu(I)$ law extends this frictional stress, in an averaged sense, to granular packings. However, as the grains are in motion, are we correct to assume that they make contact for a sufficiently long time in order to impart frictional stresses of this kind to one another? For very rapid flows, in which particle contacts are extremely short-lived, the $\mu(I)$ continuum model is not appropriate, and we must turn instead to kinetic theories (e.g. Pöschel & Brilliantov, 2003). However, even for relatively slow flows, there is cause to question the assumption that the internal stresses in granular materials arise due to frictional contacts between particles.

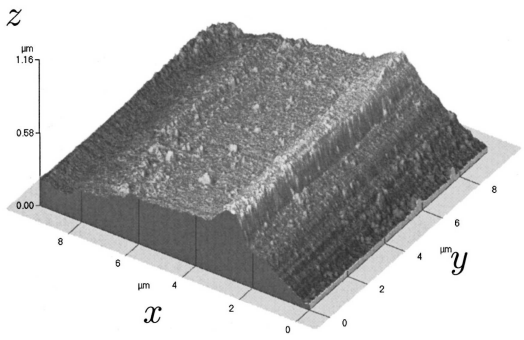

Figure 5.1: Anisotropic roughness of a machined surface, viewed using an atomic force microscope, adapted from Plouraboué & Boehm (1999).

Classical results in fluid mechanics suggest that a single grain falling under gravity through a viscous fluid should make contact with an underlying surface, or another grain, only after an infinite time (Brenner, 1961). This slow approach to contact results from the divergence of the lubrication pressure force stemming from the squeeze flow in the closing gap between the two rigid objects. This result, although simple to derive, is a little unsatisfying in view of the intuition built from everyday experiences, which suggests that objects can touch in finite time, whether or not they are immersed in a fluid.

This problem of contact has much farther-reaching applications than a pedantic justification for one aspect of the $\mu(I)$ rheological model, of course. Finite-time contact has also recently been invoked to explain a number of related fluid mechanical phenomena, such as the fore-aft asymmetry of the paths of particles colliding in viscous fluid (Davis *et al.*, 2003), the terminal velocities reached by particles rolling down inclined surfaces (Smart *et al.*, 1993), and the steady rolling of viscously lubricated, nested cylinders (Balmforth *et al.*, 2007). In all of these examples, it is postulated that asperities on the surfaces are required to make contact and hold the colliding objects apart by a roughly fixed distance, allowing fluid flow through the residual gap to provide drag. Such asperities are present even on seemingly smooth, machined surfaces, such as the metal sheet shown in Figure 5.1, or the glass beads typically used for experiments involving granular materials. Although it has never been explicitly shown that true contact is necessary to promote frictional drag

forces, such behaviour is typically asserted.

Nevertheless, there are no obvious flaws in the theoretical prediction that contact occurs after infinite time. The simplest route by which the result may be established is via the well-known lubrication theory of Reynolds (1886), which applies to relatively smooth, slowly moving and nearly touching objects. A more recent formal analysis by Gérard-Varet & Hillairet (2010), that does not adopt the usual approximations of lubrication theory (the Stokes approximation in conjunction with small aspect ratio), also demonstrates that objects with smooth surfaces may not make contact in finite time.

In this chapter, we address the question of where contact can be allowed to occur in finite time when one includes some additional (and plausibly relevant) physical effects. We shall give a detailed account in §5.1 of how the approach to contact is modified in the presence of rough asperities, illustrated by the model problem of a sharp wedge settling under gravity onto a horizontal plane. A brief discussion of effects arising from the compressibility of the fluid, the elastic deformation of the solid surfaces, and intermolecular forces will be made in §5.2.

5.1 The effect of asperities

The idea that asperities could affect contact is suggested by lubrication theory, which predicts that if the approaching surfaces are not locally smooth and flat, but sharper, then fluid flows more easily out of the intervening gap, weakening the divergence of the lubrication forces. Contact can then be approached in finite time. Unfortunately, the geometry of the sharpened problem contradicts one of the underlying assumptions of lubrication theory (that the characteristic horizontal scales must be much less than the vertical ones), calling into question the validity of that result.

In this section, we consider the model problem of the descent of a sharp wedge and solve the appropriate two-dimensional Stokes problem numerically. An alternative simple geometry would be an axisymmetrical cone, which arguably has greater physical realism. However, the wedge offers a more accessible geometry for the method with which we solve the flow problem (and, in particular, for our choice of boundary conditions). We shall find that the upwards resistive force exerted on the descending wedge by the fluid scales logarithmically

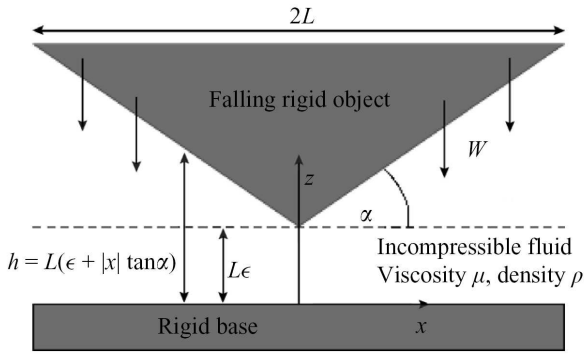

Figure 5.2: Diagram and notation for a vertically-falling wedge.

with the minimum gap thickness, as predicted by lubrication theory. Thus, the prediction that contact can be made in finite time is vindicated (at least for two-dimensional geometry; the corresponding result for a cone remains unproven, though plausible). This section is a mildly edited version of an article published in the Journal of Fluid Mechanics (Cawthorn & Balmforth, 2010), with additional clarification of certain points.

5.1.1 The Stokes problem for a falling wedge

In order to investigate the effect of a sharp asperity on contact, we consider the model Stokes problem of a falling wedge illustrated in Figure 5.2. This model problem encapsulates how a corner-shaped asperity modifies the squeeze flow beneath an object sedimenting towards a flat wall. It is, of course, possible to consider a different geometry more relevant to a physical problem (two approaching spheres, for example), but the falling wedge offers the simplest possible geometrical setting for studying the effect of the asperity. However, using a wedge introduces a complication in that the forces on the upper surface are not finite, if the object is taken to flare outwards to infinity. On the other hand, however, were we to demand that the wedge end at a finite distance from its apex, we would need to add the detailed boundary conditions there. We avoid these issues by solving for the Stokes flow beneath an infinite wedge, and then computing the upward resistive force on wedge due to the flow beneath a pre-defined central section.

The wedge is symmetric about $x = 0$ and is pitched at angle α to the horizontal; as shown

in Figure 5.2, we consider a central section of width $2L$. The wedge apex lies a distance $L\epsilon$ above the plane, and the wedge falls vertically. Motivated by gravitational settling, we cast the governing Stokes' equations into a dimensionless form by scaling lengths by L, velocities by $\sqrt{Lg}$, time with $\sqrt{L/g}$, and pressure with $\mu\sqrt{g/L}$, where g is the gravitational acceleration and μ the fluid viscosity. The falling speed can then be written as $W = -\sqrt{Lg}\dot{\epsilon}$, and we arrive at the biharmonic problem for the (dimensionless) streamfunction, $\psi(x, z)$:

$$\nabla^4\psi = 0, \tag{5.1}$$

$$\psi(x, 0) = \psi_z(x, 0) = 0, \tag{5.2}$$

$$\psi_z(x, h(x)) = 0, \quad \psi_x(x, h(x)) = -\dot{\epsilon}, \tag{5.3}$$

$$\psi(0, z) = \psi_{xx}(0, z) = 0, \tag{5.4}$$

which incorporates the no-slip conditions on the lower and upper surfaces, and the symmetry condition at $x = 0$ (the streamfunction is antisymmetric in x). The problem also demands far-field conditions to be placed on ψ for $|x| \to \infty$; these will be described later, after we derive an approximate solution. The velocity field, $(u, w) \equiv (\psi_z, -\psi_x)$, and the local thickness of the fluid layer is given by

$$h(x) = \epsilon + |x| \tan\alpha. \tag{5.5}$$

Using the solution to this problem, we will compute the upward force per unit width on the wedge from the section, $-1 \le x \le 1$:

$$F_z = -\int_{-1}^{1} \hat{\mathbf{z}} \cdot \left(\begin{array}{cc} 2u_x - p & u_z + w_x \\ u_z + w_x & 2w_z - p \end{array} \right)\Bigg|_{z=h(x)} \cdot \left(\begin{array}{c} -\tan\alpha \\ 1 \end{array} \right) \mathrm{d}x \tag{5.6}$$

$$= 2\int_0^1 \left[p - 2\frac{\partial w}{\partial z} + \left(\frac{\partial u}{\partial z} + \frac{\partial w}{\partial x} \right) \tan\alpha \right]_{z=h(x)} \mathrm{d}x, \tag{5.7}$$

where p is the (dimensionless) fluid pressure (which converges to an ambient pressure level far from the apex of the wedge that can be set to zero for our current problem).

Lubrication solution

As we have already asserted, the neighbourhood of the sharp corner of the wedge constitutes a region in which the characteristic lengthscale tends to zero, so it is not formally appropriate

to use lubrication theory to calculate an approximate solution to the Stokes problem (5.1)-(5.4). Nevertheless, we present the calculation here both to put into context our valid approach, and to show later that lubrication theory qualitatively captures the dynamics despite the violation of its assumptions.

In lubrication theory, we reduce the dimensionless governing fluid equations for the velocity field (u, w) and pressure p to

$$p_x = u_{zz}, \qquad p_z = 0, \qquad u_x + w_z = 0. \tag{5.8}$$

Given that $u = 0$ on the two surfaces, we arrive at the Poisseuille velocity profile,

$$u(x, z) = -\frac{1}{2} z(h - z) p_x. \tag{5.9}$$

Integrating the continuity equation across the gap, and imposing the conditions, $w(x, 0, t) = 0$ and $w(x, h, t) = -1$, furnishes

$$-\dot{\epsilon} = \frac{\partial}{\partial x} \int_0^h u \, \mathrm{d}z = -\frac{1}{12} (h^3 p_x)_x, \tag{5.10}$$

from which it follows that

$$p = -\frac{6\dot{\epsilon}}{\tan^2 \alpha} \left(\frac{2}{h} - \frac{\epsilon}{h^2} \right), \tag{5.11}$$

satisfying $p \to 0$ as $x \to \infty$. Finally, the upward resistive force may be found by integrating this pressure distribution over the interval $-1 \le x \le 1$. Note that the other terms in (5.7) constitute a second order correction to the force exerted by pressure alone under the assumptions of lubrication theory. After performing the integration, we calculate the resistive force to be

$$F_z \approx F_L = \int_{-1}^{1} p \, \mathrm{d}x = -\frac{24\dot{\epsilon}}{\tan^3 \alpha} \left[\log \left(1 + \frac{\tan \alpha}{\epsilon} \right) - \frac{\tan \alpha}{2(\epsilon + \tan \alpha)} \right]. \tag{5.12}$$

Had the gap thickness been locally smooth, then $h(x) \approx \epsilon + x^2/2$. In this situation, the analogous result to (5.11) is $p = -6\dot{\epsilon}/h^2$, and the lubrication pressure force on the entire object (which is now integrable) is

$$\int_{-\infty}^{\infty} p \, \mathrm{d}x = -\frac{3\pi\sqrt{2}\dot{\epsilon}}{\epsilon^{3/2}}. \tag{5.13}$$

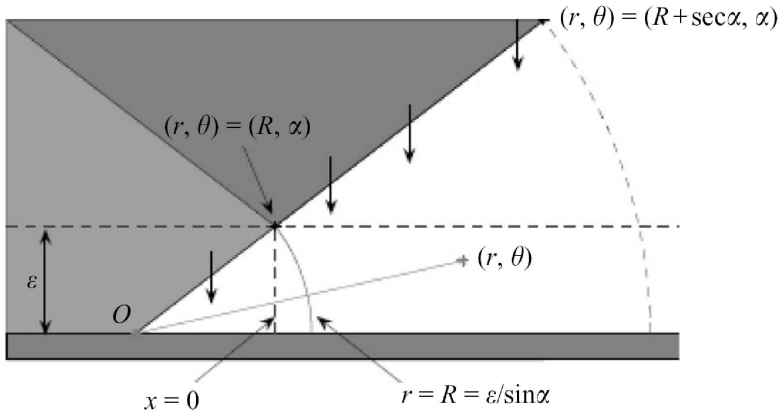

Figure 5.3: Diagram and notation for flow in a corner. The darker shaded region illustrates the wedge in Figure 5.2, with symmetry axis $x = 0$.

Asymptotic or outer solution

If the wedge is relatively close to the lower plane (and so $\epsilon \ll 1$), an approximate solution follows on considering the corner flow illustrated in Figure 5.3. The velocity boundary conditions on the solid surfaces correspond to those for the falling wedge. A simple analytical solution can be constructed for this configuration using polar coordinates (r, θ) centred at the corner itself (labelled O in the figure). In these coordinates, the streamfunction, $\psi(r, \theta)$, satisfies

$$\nabla^4 \psi = 0 \quad \text{in } 0 < \theta < \alpha, \tag{5.14}$$

$$\psi(r, 0) = \frac{\partial \psi}{\partial \theta}(r, 0) = 0, \tag{5.15}$$

$$\frac{\partial \psi}{\partial r}(r, \alpha) = -\dot{\epsilon} \cos \alpha, \quad \frac{\partial \psi}{\partial \theta}(r, \alpha) = r\dot{\epsilon} \sin \alpha, \tag{5.16}$$

with the radial and angular velocity components given by $U = r^{-1}\partial_\theta \psi$ and $V = -\partial_r \psi$, respectively.

A solution which satisfies these relations and corresponds to a bounded velocity field for $r \to \infty$, is given by

$$\psi(r, \theta) = \dot{\epsilon}[rg(\theta) + f(\theta)], \tag{5.17}$$

where

$$g(\theta) = -\frac{(\alpha + \sin\alpha\cos\alpha)(\theta\cos\theta - \sin\theta) + \theta\sin^2\alpha\sin\theta}{\alpha^2 - \sin^2\alpha} \tag{5.18}$$

and

$$f(\theta) = \frac{\epsilon\cos\alpha[(\cos 2\alpha - 1)(\cos 2\theta - 1) + \sin 2\alpha(\sin 2\theta - 2\theta)]}{4\sin^2\alpha(\sin\alpha - \alpha\cos\alpha)}. \tag{5.19}$$

Although (5.17) produces no net mass flux through $x = 0$ (or, equivalently, $r = R = \epsilon/\sin\alpha$), it does not satisfy the full symmetry conditions. [1]

The upward resistive force on the section of the upper surface, $R \le r \le R + \sec\alpha$ and $\theta = \alpha$, calculated using the corner flow solution (5.17)-(5.19), offers an approximation of $\frac{1}{2}F_z$, thus (using polar vector notation)

$$F_z \approx -2\int_R^{R+\sec\alpha} (\sin\alpha, \cos\alpha)\cdot\begin{pmatrix} 2U_r - p & V_r + \frac{U_\theta - V}{r} \\ V_r + \frac{U_\theta - V}{r} & 2\frac{V_\theta + U}{r} - p \end{pmatrix}\cdot\begin{pmatrix} 0 \\ 1 \end{pmatrix} \mathrm{d}r, \tag{5.20}$$

$$= 2\cos\alpha\int_R^{R+\sec\alpha}\left[p + \frac{2}{r}\psi_{r\theta} - \frac{2}{r^2}\psi_\theta - \tan\alpha\left(\frac{\psi_{\theta\theta} + r\psi_r}{r^2} - \psi_{rr}\right)\right]\mathrm{d}r. \tag{5.21}$$

To calculate p:

$$\begin{aligned} p_r &= U_{rr} + \frac{1}{r}U_r + \frac{1}{r^2}U_{\theta\theta} - \frac{1}{r^2}U - \frac{2}{r^2}V_\theta \\ p_\theta &= rV_{rr} + V_r + \frac{1}{r}V_{\theta\theta} + \frac{2}{r}U_\theta - \frac{1}{r}V, \end{aligned} \tag{5.22}$$

[1] The approximation (5.17) also corresponds to the first two terms of a series solution with the separable form,

$$\psi = \dot{\epsilon}\left[f(\theta) + rg(\theta) + \ \mathrm{Re}\left\{\sum_{k=1}^{\infty} r^{2-\lambda_k}\left[A_k f_k(\theta) + B_k g_k(\theta)\right]\right\}\right], \tag{†}$$

where

$$f_k(\theta) = \frac{\cos[\lambda_k(\theta - \alpha/2)]}{\cos[\lambda_k\alpha/2]} - \frac{\cos[(\lambda_k - 2)(\theta - \alpha/2)]}{\cos[(\lambda_k - 2)\alpha/2]}, \qquad \sin(\lambda_k - 1)\alpha = -(\lambda_k - 1)\sin\alpha,$$

and

$$g_k(\theta) = \frac{\sin[\mu_k(\theta - \alpha/2)]}{\sin[\mu_k\alpha/2]} - \frac{\sin[(\mu_k - 2)(\theta - \alpha/2)]}{\sin[(\mu_k - 2)\alpha/2]}, \qquad \sin(\mu_k - 1)\alpha = (\mu_k - 1)\sin\alpha$$

(*cf. Moffatt, 1964; Liu & Joseph, 1977*). The terms in the series expansion in (†) decay rapidly with r, implying that (5.17) provides a useful far-field approximation to the true streamfunction. Although the summed terms could also lead to a pattern of alternating eddies (Moffatt, 1964), they are not apparent in any of the solutions we present later because they are dominated by the leading two terms, $f(\theta) + rg(\theta)$, which exhibit no such behaviour. In other words, we observe no infinite sequence of "Moffatt" eddies here, unlike in several other problems involving Stokes flow in and around wedges.

which lead to

$$p = -\frac{f'''}{2r^2} - \frac{(g''' + g')}{r}. \tag{5.23}$$

Hence,

$$\begin{aligned} F_z &\approx -2\cos\alpha \int_R^{R+\sec\alpha} \left[\frac{f''' + 4f'}{2r^2} + \frac{g''' + g'}{r} + \tan\alpha \left(\frac{f''}{r^2} + \frac{g'' + g}{r} \right) \right] \mathrm{d}r \\ &= -2\cos\alpha [g''' + g' + (g'' + g)\tan\alpha] \log\left(1 + \frac{\tan\alpha}{\epsilon}\right) \\ &\quad - \frac{\sin^2\alpha (f''' + 4f' + 2f''\tan\alpha)}{\epsilon(\tan\alpha + \epsilon)}, \end{aligned} \tag{5.24}$$

where the various derivatives of f and g are evaluated at $\theta = \alpha$. A little more effort yields the result

$$F_z \approx F_O = -\dot{\epsilon} \left[\frac{2(2\alpha + \sin 2\alpha)}{(\alpha^2 - \sin^2\alpha)} \log\left(1 + \frac{\tan\alpha}{\epsilon}\right) - \frac{4\tan\alpha}{(\tan\alpha + \epsilon)(\tan\alpha - \alpha)} \right]. \tag{5.25}$$

Note that (5.25) reduces to (5.12) in the limit of a shallow wedge, $\alpha \ll 1$.

Numerical solution

The preceding arguments suggest that the approximate solution (5.17)-(5.19) provides the limiting far-field form of the full wedge solution. Thus, one can simplify numerical computations of that solution by truncating the domain at $x = \pm 1$ (which can be considered to lie in the far field when $\epsilon \ll 1$), and matching the streamfunction there to (5.17). Given this truncation, we numerically solve equations (5.1)-(5.5) by first mapping the region onto the unit square, via the transformation,

$$\xi = x, \qquad \zeta = z/h(x). \tag{5.26}$$

It is then a simple matter to discretise the equations on a grid on the $(\xi, \zeta)-$plane by using a second-order finite difference scheme, and perform the linear algebra necessary to solve the discrete problem.

In order to calculate accurately the solution in the vicinity of the vertex of the wedge, one would like to ensure that the grid there is well-resolved. The nature of the transformation (5.26) is sufficient to ensure a good vertical resolution with no further effort required, however to improve the horizontal resolution we use a simple multigridding method. Initially, the

problem is solved for $0 < x < 1$, using the outer solution at $x = 1$ as described above. The calculation is then repeated with a finer spatial resolution on the interval $0 < x < 0.1$. In this case, we evaluate the streamfunction obtained from the previous calculation at $x = 0.1$ in order to supply the boundary condition there. This process may then be repeated further to achieve as accurate a solution as is desired. The upward resistive force is dominated by the high stresses near the vertex of the wedge, and the benefit of this scheme is that we can accurately calculate the force without wasting computational effort on the region further from the vertex.

Two sample numerical solutions with different wedge angles are shown in Figure 5.4. For smaller angles ($\alpha = \pi/12$ in Figure 5.4), both the lubrication and outer solutions provide a good approximation for the horizontal velocity profile at $x = 0.1$ (Figure 5.4d). However, the parabolic velocity profile predicted by lubrication theory does not match the computed solution for a sharper wedge ($\alpha = \pi/3$, Figure 5.4c), because the aspect ratio is no longer small. Instead, the horizontal flux is concentrated towards the lower boundary, as predicted by the outer solution.

Plots of the vertical force as a function of ϵ are shown in Figure 5.4(e,f). The force is a linear function of $\log \epsilon$ for small ϵ, with a pre-factor that is given by the approximate solution. At larger wedge angles, lubrication theory inaccurately estimates that pre-factor, leading to a force law with the wrong slope in Figure 5.4(e).

5.1.2 Smoothing out the corner at a smaller scale

The argument that all physical surfaces are rough on some small lengthscale can readily be countered by an assertion that even a sharp asperity must be rounded on a yet smaller lengthscale. To illustrate the consequence of smoothing out the corner at a shorter scale, we carried out some further numerical computations for a rounded wedge, in which the fluid gap is given by

$$h(x) = \begin{cases} \epsilon + \frac{1}{2}\sigma \tan^2 \alpha + \frac{1}{2}\sigma^{-1} x^2, & |x| < \sigma \tan \alpha, \\ \epsilon + |x| \tan \alpha, & |x| > \sigma \tan \alpha. \end{cases} \tag{5.27}$$

Sample numerical solutions, computed using the same scheme as described in §5.1.1, are presented in Figure 5.5. Away from the smoothed corner, the streamfunctions largely re-

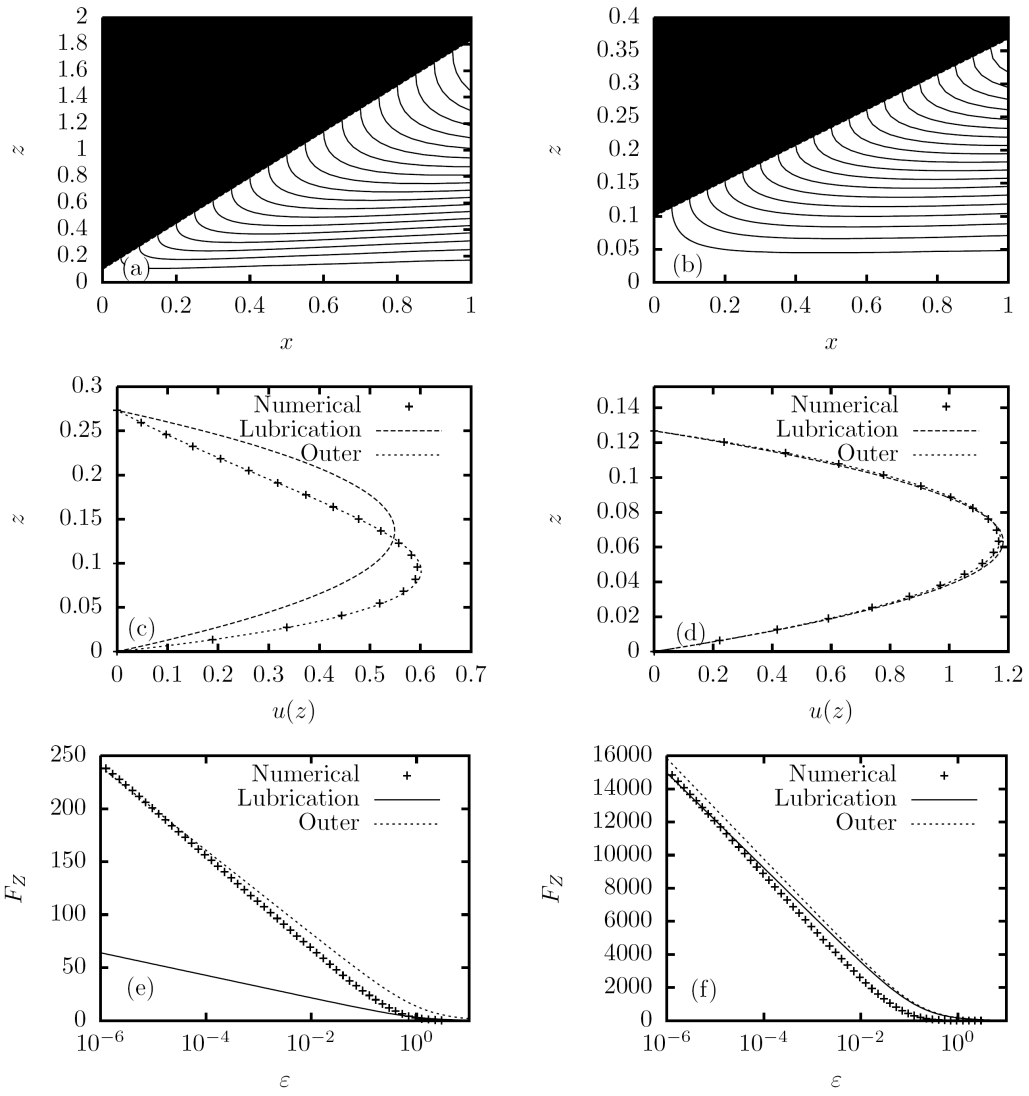

Figure 5.4: Numerical Stokes solutions with $\dot{\epsilon} = -1$ for wedges with $\epsilon = 0.1$ and $\alpha = \pi/3$ (panels (a), (c) and (e)) or $\alpha = \pi/12$ (panels (b), (d) and (f)). Shown in panels (a) and (b) are contour plots of the streamfunction on the $(x, z)-$plane. Panels (c) and (d) show the horizontal velocity profile along the line $x = 0.1$, with predictions from the lubrication and outer solutions. Panels (e) and (f) are plots of the vertical force F_z, with $\dot{\epsilon} = -1$, as determined by lubrication theory (5.12), the approximate Stokes solution (5.25), and the numerical solution, via (5.7).

produce those computed for the full wedge; only for $x \sim \sigma \tan\alpha$ is there an appreciable difference.

The force on the rounded wedges is shown in Figure 5.5(e,f), together with the approximation for the full wedge (5.25) and the leading-order lubrication result as the gap closes (*cf.* equation (5.13)),

$$F_{RL} = -3\sqrt{2}\pi\dot{\epsilon}\left(\frac{\sigma}{\delta}\right)^{3/2}, \tag{5.28}$$

where $\delta = \epsilon + \frac{1}{2}\sigma\tan^2\alpha$ is the actual minimum gap thickness. The full numerical solution converges to the lubrication prediction at the smallest separations; for larger gap thicknesses, the numerical result falls closer to the approximate solution for the full wedge. In other words, at the smallest separations, the lubrication pressure force reflects the smoothed corner, but at large separations the smoothing of the wedge is not felt and the descending object appears to have a real corner. The cross over between the two regimes occurs for $\delta \sim \sigma\tan^2\alpha$ (separations of order the radius of curvature of the rounded tip).

5.1.3 Consequences for sedimentation and contact

Having established the form of the upwards resistive force on a wedge converging with a plane, we now consider the dynamics of a wedge-shaped object falling under gravity. As mentioned earlier, however, the gravitational and hydrodynamic forces diverge for an infinitely wide wedge. We therefore shift focus, and consider a finite wedge with our previous calculations describing the flow beneath its vertex, and neglect any effect of the flow around the finite edges of the wedge.

Vertical settling

We first consider purely vertical settling, with the left-right symmetry of the wedge ensuring that it does not tilt during the descent. The limiting form of the resistive force from the squeeze flow beneath the corner of the wedge indicates that, near contact, the equation of motion of the falling object reduces to

$$\mathcal{M}\frac{d^2\epsilon}{dt^2} \sim -1 - \nu(\alpha)\frac{d\epsilon}{dt}\log\frac{1}{\epsilon}, \tag{5.29}$$

where the dimensions of time have been chosen to scale gravitational accelerations to unity, $\nu(\alpha)$ is a constant, but angle-dependent, drag coefficient, and $\mathcal{M}$ measures the inertia of the

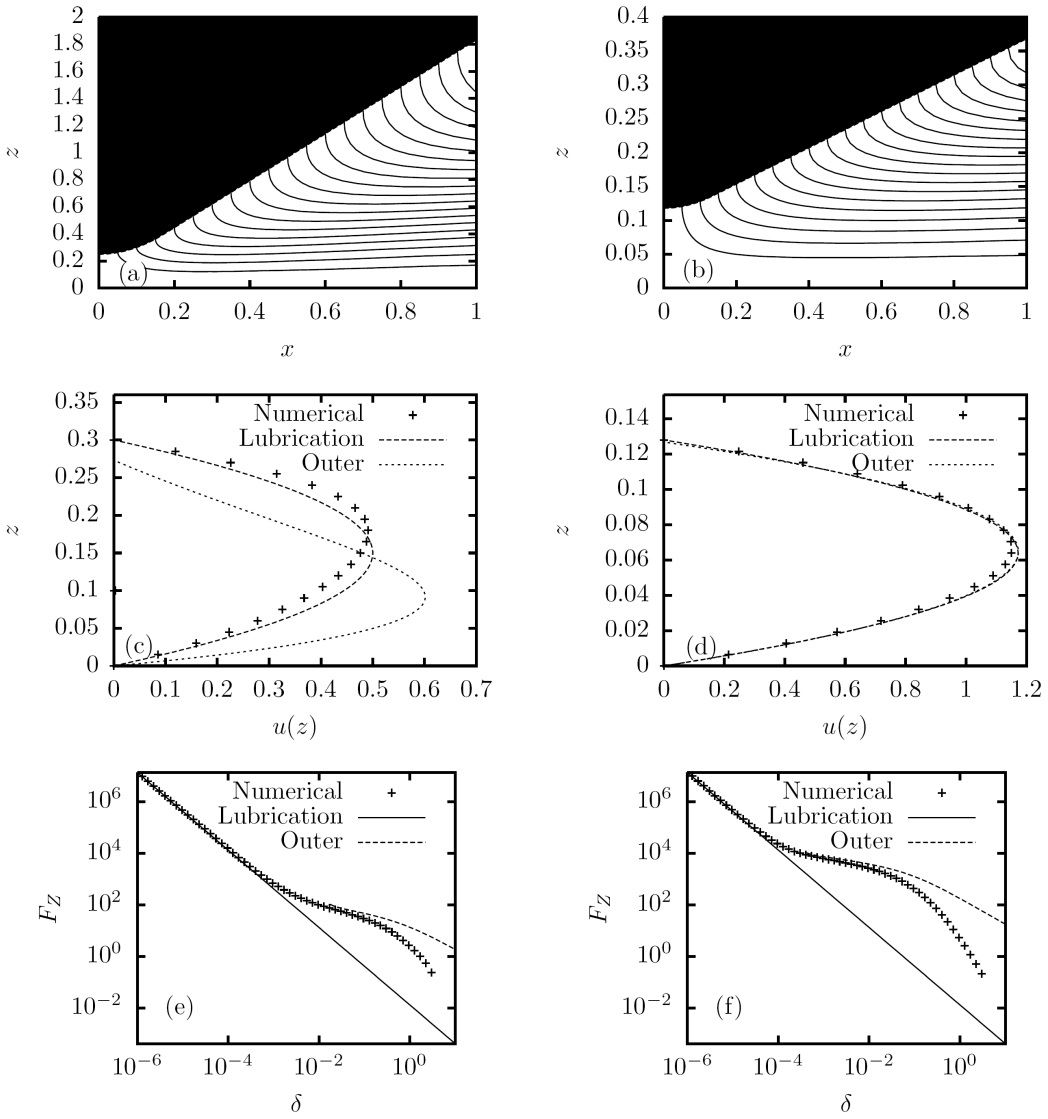

Figure 5.5: Stokes solutions with $\dot{\epsilon} = 1$ for rounded wedges with $\epsilon = 0.1$ and $(\alpha, \sigma) = (\pi/3, 0.1)$ (panels (a), (c) and (e)) or $(\alpha, \sigma) = (\pi/12, 0.5)$ (panels (b), (d) and (f)). Shown in panels (a) and (b) are contour plots of the streamfunction on the $(x, z)-$plane. Panels (c) and (d) show the horizontal velocity profile along the line $x = 0.1$, with predictions from the lubrication and outer solutions. Panels (e) and (f) feature plots of the vertical force F_z, with $\dot{\epsilon} = -1$, as determined by lubrication theory (5.28), the approximate Stokes solution (5.25), and the numerical solution, via (5.7).

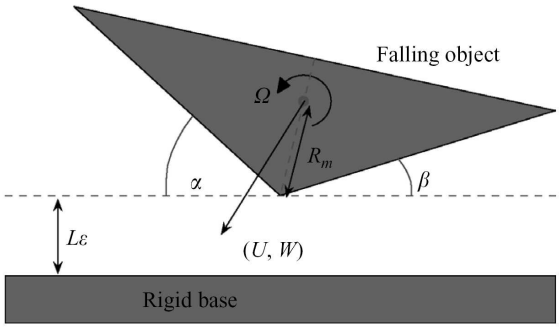

Figure 5.6: Diagram and notation for a falling, rotating, and translating wedge.

object. As we approach contact, the inertial term may be discarded, and we arrive at the key result,

$$\epsilon \log \frac{1}{\epsilon} \sim \frac{(t_c - t)}{\nu(\alpha)}, \tag{5.30}$$

where t_c is a finite contact time.

For a wedge that is rounded at a small length scale, the logarithmic divergence of the drag force applies over an intermediate range of gap thickness, but ultimately switches back to the algebraic divergence expected for smooth surfaces. The switch occurs for separations that are of the same order as the smoothing scale. Thus, although contact no longer occurs in finite time, the sharpness of the asperity accelerates the initial approach to contact.

Settling with translation and rotation

We now allow the wedge full freedom of movement so that it may tilt as it descends as a result of some initial asymmetric orientation, in order to determine whether this influences the approach to contact. One danger is that the wedge may tip over onto one side, falling with one face horizontal, significantly increasing the lubrication force and delaying contact. Alternatively, horizontal motion of a tilted wedge can generate lift, as in the classical Reynolds bearing. Thus, we need to check that the case of purely vertical settling is not overly idealized before we can conclude that contact occurs in finite time in generic situations. Note that the wedge will also necessarily tilt due to the torque induced by any left-right asymmetry.

The geometry of the problem is sketched in Figure 5.6. The position of the wedge's apex is $(x_0(t), \epsilon(t))$, which can be related geometrically to the position of the centre of mass using the distance between those points, R_m, the opening angle of the wedge, $\pi - \alpha(t) - \beta(t)$, and the tipping angle, $\alpha(t) - \beta(t)$, as is evident in Figure 5.6. The centre of mass of the wedge has velocity, $(U(t), W(t))$, and rotation rate, $\Omega(t)$; these can again be connected geometrically to dx_0/dt and $d\epsilon/dt$. The problem now reduces to formulating the equations of motion of the wedge, which amount to ordinary differential equations for $(\epsilon(t), x_0(0), \alpha(t))$, once the forces and torque on the wedge are computed.

To ensure the analysis remains tractable, we restrict attention to a wedge with a shallow opening angle, so that angles (α, β) remain small. We can then exploit the lubrication approximation, which, as demonstrated above, yields an adequate expression for the vertical hydrodynamic force. With this approximation, we can analytically compute the force on the wedge, (F_x, F_z), and the torque, G, about its apex. To $O(\alpha^2, \beta^2)$, these are given (dimensionlessly) by

$$F_x = \int_{-1}^{1} \left(-p \frac{\partial h}{\partial x} - \frac{\partial u}{\partial z} \right) \mathrm{d}x, \tag{5.31}$$

$$F_z = \int_{-1}^{1} p \, \mathrm{d}x, \tag{5.32}$$

$$G = \int_{-1}^{1} |x| p \, \mathrm{d}x, \tag{5.33}$$

where h is the thickness of the fluid-filled gap beneath the wedge,

$$h(x) \approx \begin{cases} \epsilon + (x - x_0)\alpha & \text{for } x > x_0, \\ \epsilon - (x - x_0)\beta & \text{for } x < x_0. \end{cases} \tag{5.34}$$

The pressure follows from integrating the lubrication equations along the lines outlined between equations (5.8-5.11), albeit with the more general velocity boundary conditions for the translating and rotating wedge, and imposing $p = 0$ at the wedge's edges. Substitution of that result into the forces and torque leads to the relation

$$(F_x, F_z, G)^T = \mathbf{A}(U, W, \Omega)^T, \tag{5.35}$$

where coefficients of the "resistance matrix", $\mathbf{A}$, are given by

$$\begin{aligned} A_{11} &= 3(I_{02}^2 - \tfrac{4}{3} I_{01} I_{03})/I_{03}, \\ A_{12} = A_{21} &= 6(I_{03} I_{12} - I_{13} I_{02})/I_{03}, \\ A_{13} = A_{31} &= 3(I_{03} I_{22} - I_{23} I_{02})/I_{03}, \\ A_{22} &= 12(I_{13}^2 - I_{03} I_{23})/I_{03}, \\ A_{23} = A_{32} &= 6(I_{13} I_{23} - I_{03} I_{33})/I_{03}, \\ A_{33} &= (I_{23}^2 - I_{03} I_{43})/I_{03}. \end{aligned} \tag{5.36}$$

In the above formulae, the integrals I_{mn} are defined by

$$I_{mn} = \int_{-L}^{L} \frac{x^m \, \mathrm{d}x}{[h(x)]^n}. \tag{5.37}$$

Finally, for the eventual approach to contact, we neglect the wedge's inertia to arrive at the equations of motion of that object, which reflect force and torque balance:

$$F_x = 0, \tag{5.38}$$

$$F_z = 1, \tag{5.39}$$

$$G = R_M(\alpha - \beta)/2. \tag{5.40}$$

(Note that the horizontal displacement of the centre of mass from the apex of the wedge leads to a gravitational torque.)

To calculate the motion of the wedge, we first find and invert the matrix, $\mathbf{A}$, given the current geometry. Equation (5.35) then determines the velocity and angular velocity of the centre of mass of the wedge. Geometrical relations connect these quantities to $(\dot{\epsilon}, \dot{x}_0, \dot{\alpha})$, providing the governing differential equations that we integrate numerically. Numerical computations were carried out using a standard explicit non-stiff integration routine[2], with a variable timestep designed to accurately identify any instance of contact in the simulation[3].

Figure 5.7 shows a sample, typical solution for a wedge beginning with a slight tilt, but no initial horizontal translation speed or rotation (indeed, without inertia, the initial velocity

[2]We used the ode45 routine featured in Matlab.

[3]This is achievable using the standard options of the ode45 routine. One needs only specify an 'event' ($\epsilon = 0$ in our case), and the timestep is dynamically controlled in order to accurately calculate the approach to this event

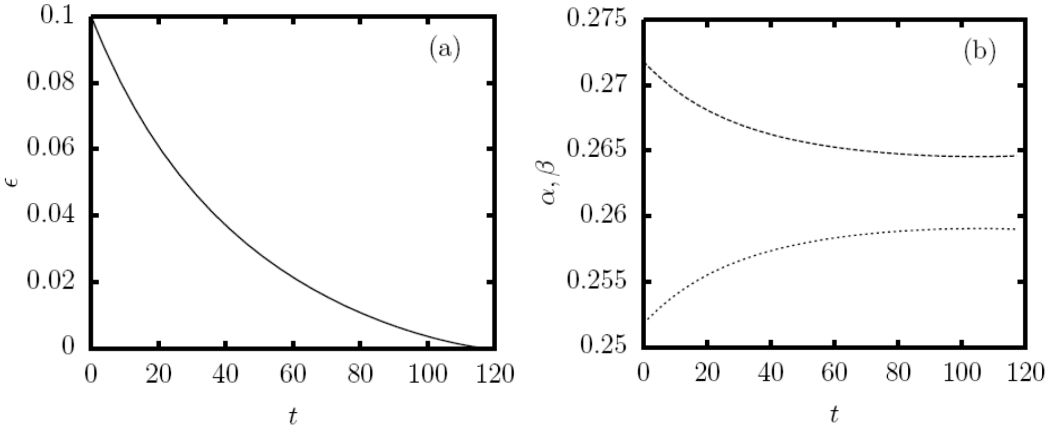

Figure 5.7: Numerical evolution of a sedimenting, tilting wedge, computed using lubrication theory. Panel (a) shows the evolution of height ϵ, and panel (b) shows the change in the angles (α, β). The computation begins from the initial conditions, $\epsilon(0) = 0.1$, $x_0(0) = 0$, $\alpha(0) = \pi/12 - 0.1$ and $\beta(0) = \pi/12 + 0.1$. Contact occurs in finite time near $t = 117$.

and rotation cannot be independently prescribed). A key feature is that the wedge still makes contact with the horizontal plane in finite time. Although the wedge does not righten itself completely by the time of contact, the wedge tends to rotate in the direction that restores symmetry because the side of the wedge above the narrower fluid gap experiences a greater lubrication pressure than the opposite side. This pressure imbalance exerts a restoring torque which overcomes the gravitational torque, and indicates that vertical descent is stable even towards relatively large perturbations (such motion can be shown to be stable towards infinitesimal perturbations using linear stability theory). Hence, we conclude that a sharp wedge settling under gravity will, unless released with one face very nearly horizontal, make contact with a horizontal plane in finite time.

5.2 Other physical effects

Although any the presence of any rough asperities will promote contact in finite time, it is by no means the only mechanism by which finite-time contact may occur. In this short section, we give a brief overview of several other physical effects that could conceivably affect

contact. Details of the effects of fluid compressibility and solid elasticity are presented in a published article (Balmforth *et al.*, 2010), but are not included here. Instead, we aim to give an overview of the processes involved for completeness.

Fluid compressibility

If we account for fluid compressibility, contact in finite time may occur via two possible mechanisms. Firstly, any inertia the object may have can trigger a series of oscillations during its descent. If the ambient pressure if low enough, these oscillations can become sufficiently severe to force contact between the two surfaces at the minimum point of one such oscillation. If the ambient pressure is higher, however, the oscillations decay and cease to offer a route to contact in finite time.

The second pathway to contact relies upon the dependence of the fluid viscosity on the local pressure. If the viscosity weakens sufficiently rapidly with increasing pressure, the lubrication force is unable to diverge rapidly enough to slow the descent of the falling object. Notably, density changes play little role directly in either of these mechanisms, rather it is the indirect effect on the pressure distribution or fluid viscosity that allows contact to occur.

Solid elasticity

If one or both of the solid surfaces are in fact elastic, then they may deform in response to the pressure exerted by the fluid. The details vary depending on the geometry of the undeformed objects and the form of elasticity assumed (e.g. an elastic foundation, half-space, membrane or beam), and are thoroughly discussed by Balmforth *et al.* (2010). In the current work, we merely note that a typical feature of such problems is the formation of a pressurised reservoir (or 'dimple') of fluid separated from the ambient fluid by a very narrow gap. The thickness of this gap controls the flux out of the dimple, hence the approach to contact. The dynamics of the dimple bear similarities with those studied for impacting viscous drops (Gopinath & Koch, 2002; Griggs *et al.*, 2008). The presence of the pressurised dimple greatly retards the approach to contact beneath the centre of the falling object, but it is interesting to note that, for certain combinations of geometry and elasticity, contact can occur in finite time at the boundary of the dimple. In such cases, the narrow gap closes

in finite time, leaving a disconnected, permanently pressurised dimple supporting much of the weight of the object.

van der Waals forces

During the approach to contact between two objects, the gap separating them may become small enough for intermolecular forces to have a significant effect (*cf.* Hocking, 1973; Gopinath *et al.*, 1997; Serayssol & Davis, 1986). The additional attractive force provided by van der Waals forces, for example, can accelerate contact. The van der Waals forces are represented by a disjoining pressure, αh^{-3}, where α is a Hamaker constant appropriate to the materials used. As the gap between the two objects becomes very small, the attractive forces stemming from this disjoining pressure dominate the gravitational force, so the motion of the object is determined by the balance between the resistive lubrication force and the attractive intermolecular forces only. For two smooth objects with minimum separation ϵ, one can perform a straightforward lubrication calculation (similar to that in §5.1.1) to show that the motion is dictated by

$$2\epsilon\dot{\epsilon} \sim -3\alpha. \tag{5.41}$$

It follows that contact can occur at a finite time, $t = t_c$, with

$$\epsilon \sim \sqrt{3\alpha(t_c - t)}. \tag{5.42}$$

Microscopic slip

It is often asserted that the no-slip boundary condition is inadequate to describe flows in microfluidic channels, and that real boundaries exhibit a degree of slip on a scale usually too small to be relevant (see Lauga *et al.*, 2007, for example). For objects approaching contact that are separated by a very small distance, however, it is necessary to allow for slip on the surfaces of the converging objects. This slip can enhance the fluid transport out of the thinning gap, weaken the resistive pressure force, and thus accelerate contact. One can repeat the lubrication analysis for the convergence of two smooth objects, using Navier slip conditions on the rigid boundaries,

$$u_z(z = 0) = Su(z = 0), \qquad u_z(z = h) = Su(z = h), \tag{5.43}$$

where S is the inverse of a characteristic slip length. The result of the lubrication calculation is that resistive lubrication force between two smooth objects is weakened, meaning that the motion is governed by

$$\frac{\dot{\epsilon}}{\sqrt{\epsilon}} \sim -\frac{3\sqrt{2}}{\pi S} \tag{5.44}$$

(*c.f.* 5.13), and contact can occur in finite time, with

$$\epsilon \sim \frac{1}{2}\left[3(t_c - t)/\pi S\right]^2 . \tag{5.45}$$

6

Conclusions and extensions

Throughout this thesis, we have described the application of the $\mu(I)$ rheology of Jop, Forterre & Pouliquen (2006) to a wide range of different granular flow configurations. Although some of the problems considered were sufficiently simple to be described by a unidirectional velocity field varying in only one dimension, we have developed numerical techniques that allow for the calculation of more complicated flows, in which all components of velocity, as well as pressure, vary in two spatial dimensions. Furthermore, we have applied the $\mu(I)$ rheological law to problems of stability, where a correct description of the interaction between different velocity components is vital to the understanding of any mechanisms leading to instability. Our aim was to assess the robustness of the $\mu(I)$ law over a range of configurations, and in particular to test its validity for more complicated, multidimensional flows. In this final chapter, we collate our findings and suggest how the model may be improved upon or extended in the future.

6.1 Successes and limitations of the $\mu(I)$ rheological law

Our investigations have shown the $\mu(I)$ rheology to be a good model for several flow configurations, particularly those involving a free surface. The $\mu(I)$ law provides an excellent description of flow on an inclined plane by design, and can be applied successfully to flow

in a rectangular channel, or atop a deep pile of grains. In the latter case, it is possible to obtain good agreement with experimental results for both the velocity profile and the depth of flowing material. Similar success can be achieved in split-bottom shear cells, whether in the linear case studied in §2.2.4, or the cylindrical case studied by Jop (2008). For problems involving sediment transport, the $\mu(I)$ law is consistent with the notion that motion occurs only for sufficiently large values of the Shields number, which relates the applied hydrodynamical stress to the weight of individual grains. One can therefore make use of the $\mu(I)$ rheology to model bedload transport in a manner that correctly predicts the existence of a dune-forming instability when a bed of grains is subjected to a fluid shear.

Promising results were also obtained for the important problem of the collapse of a granular column. The construction of a numerical code based on the marker-and-cell method has enabled us to perform the first simulations of granular slumping using not only the $\mu(I)$ rheological law, but also the first simulations using any fully two-dimensional continuum model. These simulations indicate that a constant-friction version of the $\mu(I)$ law is capable of capturing all of the key features of the collapse of a short column, as well as several details relating to the collapse of taller columns. Although we fail to quantitatively predict the correct dependence of the runout on the initial aspect ratio of the column under the constant friction model, it is hoped that the enhanced dissipation introduced to shearing regions by employing the $\mu(I)$ model will improve our predictions, although a thorough investigation in this regard will be reported in a future article.

One should not believe, however, that the $\mu(I)$ law is without its limitations. Indeed, it fails to even qualitatively predict some of the key observations in certain simple flow geometries. In particular, the $\mu(I)$ law fails to predict the existence of S-shaped velocity profiles for linear shear flow, and often underestimates the thickness of shear bands in several confined shear flows. In experiments, such as those performed in an annular shear cell, one observes that the size and location of shear bands are essentially independent of the external flow conditions. Although the $\mu(I)$ rheological model reproduces this behaviour for relatively fast flows, it fails for slower, quasi-static flows, by predicting that the shear layer vanishes as the driving stress decreases. This last observation exemplifies the key failing of the $\mu(I)$ law. Its simple, local nature carries with it the impossibility of capturing the multitude of complicated effects that can occur when granular material moves very slowly.

Experiments show that the slow, quasi-static flows for which the $\mu(I)$ law performs poorly are characterised by a very small inertia number, I. In such flows, the macroscopic shear is sufficiently weak to allow the formation of force chains, which can transmit information over large distances. In such situations, we are forced not only to abandon attempts to use a local rheological law, but also to question the validity of a continuum approach, given that the position of individual grains can have a profound effect on the global force chain structure. It seems reasonable to believe that this omission of the non-local, quasi-static physical effects is also responsible for the inability of the $\mu(I)$ constitutive law to reproduce phenomena associated with the arrest of granular flows. The existence of a minimum thickness, $h_{\mathrm{stop}}(\theta)$, and the formation of levées for flow on an inclined plane are both good examples of such effects, as is the widely-studied effect of jamming in more confined flows. Furthermore, the inability of the $\mu(I)$ law to correctly predict a range of stable, slow, laminar flows in the dune problem may also result from the poor behaviour of the $\mu(I)$ model in the quasi-static regime.

Despite the apparent failure of the $\mu(I)$ law for slow flows, it remains possible that the model may still be able to explain some experimental observations. We note that, in the functional form of $\mu(I)$ (1.11), I is scaled by the parameter I_0. In addition, we know that the $\mu(I)$ law performs rather well for flows at moderate inertia number. One can speculate, therefore, that predictions may be improved by decreasing the value of I_0 used in the model. After all, the very definition of I_0 (1.12), calculated by considering flow on an inclined plane, depends upon the inclination angle. Why, then, should a different value of I_0 not be used for the very different case of confined shearing flows? By way of justification, we note that using a much smaller value of I_0 in the case of the annular shear flow studied in §2.1.2 shifts the experimental data entirely into the range over which the width of the shearing layer does not vary with the speed of the moving boundary, as observed in the experiments.

6.2 Lessons for numerical solutions

A key aspect of the work described in this thesis is the development of various numerical techniques for calculating granular flows according to the $\mu(I)$ rheology. After all, solution of any of the problems in more than one dimension would have been impossible if one could

not resort to numerical computations.

The treatment of any sub-yield regions, and their interaction with shearing regions, is an important factor to consider when working with this granular rheology, or indeed with any rheology featuring a non-zero yield stress. Although we began by constructing a careful treatment of the yield criterion and any sub-yield stresses, we later showed that it is possible to regularise the equations by replacing sub-yield regions with regions of very high effective viscosity. The loss of accuracy associated with this regularisation is minimal, although the relative simplicity of computing regularised flows is offset by the need to use a much smaller timestep in order to avoid the diffusional instability.

With a valid regularisation for sub-yield regions in hand, we were able to turn attention to the development of a code capable of solving for general two-dimensional flows using the $\mu(I)$ rheology in more-or-less arbitrary geometry, using the problem of granular column collapse as a test case. Initial attempts to construct a Lagrangian finite element scheme were unsuccessful, primarily because of the difficulty in maintaining a smooth pressure distribution for use when calculating the deviatoric stresses. Although various smoothing and stabilisation techniques can be employed to effect several improvements, we were ultimately unsuccessful in obtaining a satisfactory result. We believe that a working Lagrangian finite element scheme could be constructed, but the large deformations of the computational domain apparent in the inertia-dominated column collapse problem pose a significant further complication to overcome, even if the spurious pressure behaviour can be controlled.

An alternative, simpler numerical approach was found in the form of a marker-and-cell scheme applied on a finite difference grid. This scheme has yielded good results for the test problem of two-dimensional granular column collapse, and could be applied to other flow configurations with minimal modification. The rotating drum would be an excellent candidate for future study using this code. With a little modification to the code, it would possible to compute any steady state velocity and surface profiles, or to investigate whether the $\mu(I)$ model can capture the time-dependent avalanching behaviour seen in slowly-rotating drum experiments. The marker particles used to denote the granular region could also be used to investigate mixing in the rotating drum, as in the experiments of Gray (2001).

To conclude our discussion of numerical techniques, one ought to mention that the marker-and-cell method used is a little crude, particularly as regards accurately tracking

the location of free surfaces. It might be beneficial to consider developing a more careful scheme with a more sophisticated method of tracking the free surface (e.g. the volume of fluid method, level-set methods, etc.) before attempting to model flows with particularly complicated free surfaces.

6.3 Extensions and outlook

In this thesis, we have concentrated our efforts on a thorough evaluation of the $\mu(I)$ rheological law for dense granular flows. Having determined its limitations, it is left to future researchers to improve upon this relatively simple model by including more physical effects intended to increase its range of validity. In this final section, we offer some speculation regarding the form that such extensions might take.

Even without looking beyond the boundaries of relatively dense flows of dry granular material, significant challenges are yet to be overcome. It should be important to address the inability of the $\mu(I)$ rheology to correctly describe flows at very low inertia number. Modelling such flows, however, will be far from straightforward. We have seen that these slow, quasi-static flows are controlled by highly inhomogeneous, anisotropic stress distributions in the form of long force chains connecting many grains. It seems reasonable that any successful model for these flows must somehow take account of this nonlocal force structure, and consider its effect on the dynamics. One might worry that the discrete nature of the system will become important for quasi-static flow. Nevertheless, progress in continuum modelling of such flows has been made by several authors. Furthermore, the most successful approach (Kamrin & Bazant, 2007) also seems to be one of the most straightforward. Using an approach involving the stochastic motion of small patches of mobility in an otherwise static medium, Kamrin & Bazant (2007) are able to correctly predict the velocity profiles observed for the same shear flows in which the $\mu(I)$ model is typically found to fail. Although both models are successful in their respective regimes, it is not clear whether they can be made compatible when one considers the transition between dense, liquid like flow and quasi-static flow.

Several authors who have performed particle dynamics simulations (including Börzsönyi *et al.*, 2009, for example) claim to have observed that the deviatoric stress tensor, calculated

by applying an averaging process to the inter-particle stresses in their simulation, appears to be asymmetric. Other authors, including Lacaze & Kerswell (2009), report a symmetric stress tensor. If the latter observations reflect the physical truth, then no additional problems are posed for the $\mu(I)$ approach. However, any need to account for an asymmetric stress tensor would add significant complication to the continuum modelling. In such a case, one would need to construct a model for granular material based upon a Cosserat continuum (see Mohan *et al.*, 1999, for example), in which it is necessary to take account of internal degrees of freedom related to the angular momentum of individual grains.

At the opposite end of the spectrum of granular flows, one might also consider flow at very high inertia number. Although the functional form of $\mu(I)$ proposed initially by Jop *et al.* (2006) has $\mu(I)$ saturating like I^{-1}, the experimental data is insufficient to support this, or any other, asymptotic behaviour. Other algebraic or even exponential relaxations are possible, and can significantly change the predictions of the model for flow at large I. It may be possible to resolve this ambiguity by making a careful comparison between the kinetic theories appropriate to granular gases and the large-I limit of a continuum obeying the $\mu(I)$ rheological law for various asymptotic forms of the function $\mu(I)$.

Although attempts to link the $\mu(I)$ model to other models appropriate to flow at very small or very large inertia number will probably be rather difficult, a much more fruitful avenue of research would be to attempt to generalise the $\mu(I)$ framework to other types of granular flow. Throughout this thesis, we have dealt exclusively with incompressible granular flows. This approach seems to be justified by the experiments of the GDR MiDi (2004), who noted that the packing fraction of grains varied only by 1-2% in each of the simple flows studied. However, we later saw in §4.1 that certain flow instabilities depend upon dilation and compaction of the granular material. Extension of the $\mu(I)$ model to encompass compressible flow would therefore be a significant positive development. The foundations of such an extension are already available in the literature. The simple dependence of packing fraction, $\phi(I)$, on the local inertia number (GDR MiDi, 2004) ought to offer an adequate equation of state, although some modification to the constitutive law may also be necessary to correctly account for the stresses associated with material dilation (Börzsönyi *et al.*, 2009).

Further generalisations of the $\mu(I)$ rheology appropriate to different types of granular flows are already under development. For example, Lemaître *et al.* (2009) have proposed a

modification to the $\mu(I)$ rheological law to create a new model appropriate to the study of dense, non-Brownian suspensions. This model is based upon the observation by Courrech du Pont *et al.* (2003) that the dynamics of a granular flow depend on whether the motion of individual grains is governed by the confining granular pressure, as for dry grains, or by viscous stresses due to the interstitial fluid. This motivates an alternative definition of the inertia number as the ratio of the inertial to the viscous timescales. Remarkably, using this new 'viscous inertia number' without otherwise modifying the existing $\mu(I)$ model yields positive results when compared to experimental data for submarine granular avalanches (Cassar *et al.*, 2005; Pailha & Pouliquen, 2009).

In addition, Rognon *et al.* (2006) have begun to consider the inclusion of cohesive effects into the $\mu(I)$ framework. Cohesion can arise from several physical effects: direct interactions (such as van der Waals forces); capillary forces from any fluid present between the grains; or solid bridges (e.g. ice forming between grains of powder snow). Regardless of the source, the approach of Rognon *et al.* (2006) is to parameterise cohesion by the local *cohesion number*: a ratio between the maximum possible attractive force and the average normal force on a material surface. The generalisation of the $\mu(I)$ rheological law then follows by allowing the parameters μ_1, μ_2 and I_0 to depend upon the local cohesion number. This generalisation shows significant promise when compared with the results of particle dynamics simulations, but would benefit from a more thorough evaluation, including testing against a range of data from physical experiments.

Although the extensions mentioned above may all be considered to be viable next steps in the development of our understanding of granular material, they barely scratch the surface of the myriad possible developments that could be made based on the $\mu(I)$ law. In this thesis, we have shown that this model, although it has its limitations, is generally suitable for a wide range of uses. Furthermore, we have made progress in the creation of a robust numerical scheme for simulating granular flows in an arbitrary two-dimensional geometry. It falls to future investigators to pursue further development of the $\mu(I)$ model, and the construction of numerical schemes capable of simulating more complicated granular flows. Given the wealth of fascinating experiments reported across the literature, such a model and numerical scheme will certainly have no shortage of possible applications.

APPENDIX

A

Lagrangian finite element approaches for simulating the collapse of a granular column

In this appendix, we provide a detailed account of the development of a Lagrangian finite element code to compute numerical solutions to the problem of granular column collapse using the $\mu(I)$ rheological law. This approach was eventually abandoned in favour of the finite difference marker-and-cell scheme described in §3.2.2. Nevertheless, the great deal of effort expended in attempting to construct the Lagrangian finite element scheme has not been entirely in vain. Several useful lessons were learned during the course of the development. It is hoped that our progress, charted in this appendix, may be useful to any researchers planning to develop a similar scheme in the future, if only to aid them in avoiding potential pitfalls.

A.1 A simple beginning

The initial version of the column collapse code is very straightforward. It uses triangular elements with piecewise linear velocity and piecewise constant pressure. We use a slightly modified version of a standard pressure-update Lagrangian finite element scheme (see Chung, 2002, for example), which works as follows. Suppose that, at time $n\Delta T$, the mesh configu-

ration is $(\mathbf{X}_i)^{(n)}$, the velocity is $\mathbf{U}^{(n)}$, and the pressure is $P^{(n)}$, where the index i runs over all vertices. The procedure to be followed in order to update the system for time $(n+1)\Delta T$ is

1. **Calculate geometric quantities:** Based on the current mesh configuration, $(\mathbf{X}_i)^{(n)}$, find the normals and area of each triangle. This geometrical information is used throughout the following steps.

2. **Calculate deviatoric stresses:** Using the current velocity, $\mathbf{U}^{(n)}$, and *current* pressure, $P^{(n)}$, together with the geometric quantities calculated in step 1, calculate the deviatoric stresses according to the regularised $\mu(I)$ rheology.

3. **Apply deviatoric stresses and gravity:** Update the velocity field to an intermediate state $\mathbf{U}^*$ by using the previously-calculated stresses. The boundary conditions (3.5)-(3.7) are applied this stage.

4. **Apply gravity:** Increment the vertical velocity of every vertex by $-\mathrm{d}T$, except for those located at $Y = 0$, which are fixed at zero.

5. **Perform pressure update:** We must now add the $-\nabla P$ term, with the new pressure distribution, $P^{(n+1)}$, chosen to ensure that $\nabla \cdot \mathbf{U}^{(n+1)} = 0$. This is achieved via successive over-relaxation (SOR) of a simple Gauss-Seidel method. As part of this iterative method, we automatically update the velocity from $\mathbf{U}^*$ to $\mathbf{U}^{(n+1)}$.

6. **Update mesh:** Advect the mesh vertices with the local material velocity $\mathbf{U}^{(n+1)}$ to obtain the new mesh configuration $(\mathbf{X}_i)^{(n+1)}$. We use a simple forward Euler update for the mesh position. Elements are unchanged except for the motion of their vertices.

Several points should be emphasised here. First of all, our method differs from the standard finite element approach, in that we use *mass lumping*. That is, we choose not to invert the mass matrix exactly in favour of approximating the mass matrix by a diagonal form, and inverting that instead. To be precise, suppose that the velocity field is defined by the sum

$$\mathbf{U} = \sum_{i=1}^{n_{\mathrm{vel}}} \mathbf{U}_i \phi_i(\mathbf{X}), \tag{A.1}$$

where n_{vel} is the total number of velocity nodes, which is dependent upon the elements used, and $\phi_i(\mathbf{X})$ is the basis function associated with node i. In the simplest version of our code, which has piecewise linear velocity, n_{vel} is equal to the number of vertices in the mesh. The standard mass matrix has components defined by

$$M_{i,j} = \langle \phi_i, \phi_j \rangle = \int\int \phi_i(\mathbf{X})\phi_j(\mathbf{X})\,\mathrm{d}^2X,$$

whereas we define the lumped mass matrix to be diagonal, with components

$$M_{i,j} = \begin{cases} 2\langle \phi_i, \phi_i \rangle & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

In the context of piecewise-linear velocity, the component $M_{i,i} =: M_i$, otherwise referred to as the *mass of vertex* i, corresponds to the sum

$$M_{i,i} = \frac{1}{3}\sum_k A_k, \tag{A.2}$$

where the sum runs over all triangles that include vertex i, and A_k is the area of triangle k. Using the lumped mass matrix in this form removes the need for a computationally-intensive matrix inversion, and does not significantly affect the accuracy of the results. For more details and discussion of mass lumping, see Hinton *et al.* (1975). We have implemented versions of the Lagrangian finite element code using the exact mass matrix, and observe no significant differences in their output from that of their mass-lumped counterparts.

A second important feature of our scheme is the way in which we handle pressure. In a Newtonian fluid, the pressure only appears in the isotropic term ∇P. As such, when we update the pressure in step 5 above, we use only the new pressure, $P^{(n+1)}$, to calculate the new velocity. For granular material, however, the $\mu(I)$ law requires the pressure also to appear in the deviatoric stress. Rather than build the complicated $\mu(I)$ term into a pressure update routine, we elect to use the old pressure, $P^{(n)}$, to calculate the stresses, but use the new pressure, $P^{(n+1)}$, when updating the velocity in step 5. This ensures that mass is conserved between timesteps, and should cause no loss of accuracy provided that the pressure varies smoothly in time. Indeed, if the change in pressure from one timestep to the next is $O(\Delta\mathrm{T})$, the error introduced by this approach will be no worse than that introduced by the first-order forward Euler timestepping scheme we use to update the vertex positions.

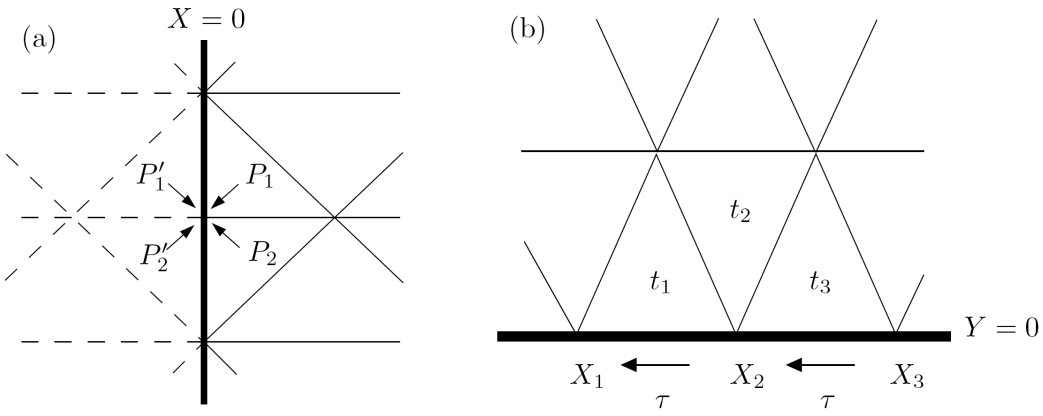

Figure A.1: Diagrams illustrating the treatment of boundary conditions in the finite element scheme. (a) The symmetry axis, $X = 0$, showing image points and forces. (b) The rigid base, $Y = 0$, with some notation to facilitate our description of the basal friction (A.4).

We turn now to the boundary conditions used in our scheme. At the free surface, things are rather simple. As the pressure vanishes on the free surface, so too do all of the components of stress. The kinematic boundary condition is dealt with implicitly when the mesh position is updated in step 6. On the symmetry axis, $X = 0$, we must apply a no flux condition ($U = 0$), but also we recognise that vertices located on this axis should be considered to have twice the mass indicated by (A.2), owing to the implied presence of more granular material located at negative X positions. One could consider the existence of image points mirroring the mesh in $X > 0$, as illustrated by Figure A.1. Combining the pressure forces with their images tells us further that, whilst the horizontal components cancel out (as one expects for a no-flux boundary), the vertical components act in the same direction. We must remember, therefore, to apply twice the contribution from the pressure term to the vertical velocity in step 5.

On the horizontal surface, we again apply a no-flux condition ($V = 0$). The other condition to be applied is dependent upon the nature of the surface. One might like to apply a no-slip condition, as one would do with a viscous fluid. However, this can lead to 'locking' behaviour, where the conservation of mass applied to each triangle ultimately results in the entire domain having zero velocity. More convenient for the numerics, and much more

appropriate to the real granular physics, would be to apply an additional frictional stress to velocity nodes on the bottom boundary. We choose to mimic the form of the (regularised) $\mu(I)$ rheology as closely as possible, and apply a frictional stress,

$$\tau_b = -\mu_b P \frac{U}{\sqrt{U^2 + \epsilon^2}}, \tag{A.3}$$

to all vertices on the boundary $Y = 0$. The friction coefficient μ_b need not be related to $\mu(I)$. For example, Balmforth & Kerswell (2005) used grit flowing over a perspex base in their experiments, and measured the internal friction coefficient of the grit to be $\mu_1 \approx 0.75$, while the basal friction coefficient was measured to be only $\mu_b \approx 0.33$. In the context of our finite difference scheme, we approximate this condition by the update

$$U(X_2, 0) \leftarrow U(X_2, 0) - \Delta T \left(\frac{\mu_b P(t_1)}{M_{X_2}} |X_2 - X_1| + \frac{\mu_b P(t_3)}{M_{X_2}} |X_3 - X_2| \right) \frac{U(X_2)}{\sqrt{U^2(X_2) + \epsilon^2}}, \tag{A.4}$$

where the quantities referred to are indicated by Figure A.1(b).

A.2 Key problems

The numerical scheme described in §A.1 appears to contain all of the details necessary to solve the problem of column collapse, as defined by (3.2)-(3.9), at least from a physical viewpoint. We have included the (regularised) granular rheology as well as the physically-relevant boundary conditions. However, code run using this basic scheme rapidly develops an instability in pressure, and causes the Gauss-Seidel iteration to diverge, effectively halting the code. A closer look at the state of the system shortly before this divergence, shown by Figure A.2, reveals the extent of the instability. The resolution in Figure A.2 is, obviously, rather coarse. Nevertheless, such plots serve as adequate examples to illustrate the key issues, and using a finer mesh does little to improve matters.

Figure A.2(b) makes it clear where the root of the problem lies. Despite the mesh having a reasonable form (as in Figure A.2a), the pressure distribution is extremely rough, with values that can vary by at much as 10^6 in dimensionless units between one triangle and the next. This numerical pressure instability rapidly leads to unstable motion of the mesh itself, via the granular stress terms. The situation only five timesteps later is as shown in Figure

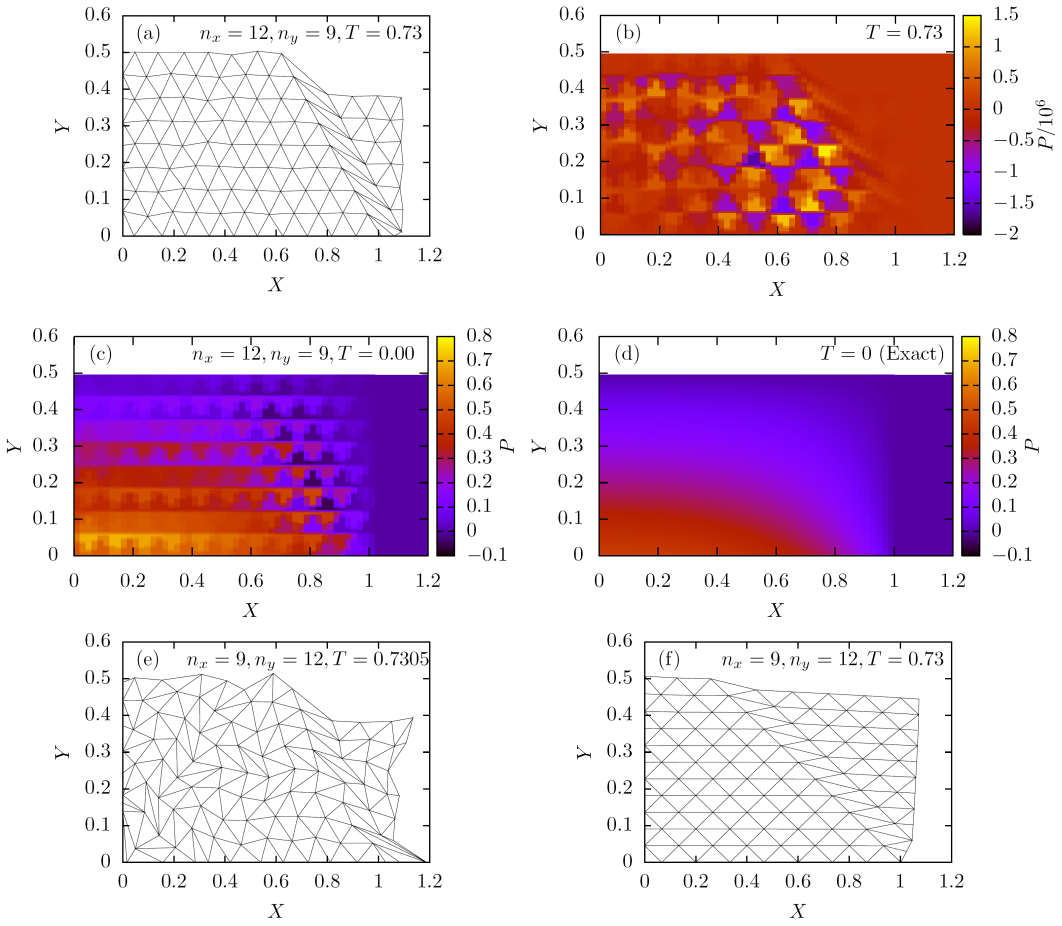

Figure A.2: Some problems associated with the finite element scheme for collapse of a granular column. All panels use a column of initial aspect ratio $a = 0.5$ and constant friction $\mu_1 = \mu_b = 0.4$. (a) Mesh position at $T = 0.73$, shortly before divergence. (b) Pressure distribution at $T = 0.73$; note that the colour scale represents pressure divided by 10^6. (c) Computed pressure distribution at $T = 0$. (d) Exact pressure at $T = 0$, given by (A.9). (e) The onset of divergence of the mesh, five timesteps after the situation in panel (a). (f) Mesh position at $T = 0.73$, with elements of different aspect ratio.

A.2(e). Clearly, this instability is able to grow to the point at which it has a significant destructive effect on the solution.

Further investigation reveals that a version of the pressure instability is present from the very beginning of the simulation. Figure A.2(c) shows the pressure distribution in the column at $T = 0$, where the only non-zero velocity is due to the application of gravity in the first timestep. In this special case, one can calculate the pressure field analytically, by solving the Laplace problem

$$\nabla^2 P = 0 \quad \text{in } 0 < X < 1, 0 < Y < a, \tag{A.5}$$

$$P(1, Y) = P(X, a) = 0, \tag{A.6}$$

$$\frac{\partial P}{\partial X}(0, Y) = 0, \tag{A.7}$$

$$\frac{\partial P}{\partial Y}(X, 0) = -1. \tag{A.8}$$

Solving by Fourier series yields the solution,

$$P(X, Y) = \frac{2}{\pi^2} \sum_{n=1}^{\infty} (-1)^{n+1} \frac{\sinh\left[\left(n + \frac{1}{2}\right)\pi(Y - a)\right] \cos\left[\left(n + \frac{1}{2}\right)\pi X\right]}{\left(k + \frac{1}{2}\right)^2 \cosh\left[\left(n + \frac{1}{2}\right)\pi a\right]}, \tag{A.9}$$

which is plotted in Figure A.2(d) for comparison with the numerical output. Whilst the numerically-determined pressure vaguely resembles the exact solution both qualitatively and quantitatively, some discrepancies are clear, and not solely due to the coarse numerical mesh. Particularly clear along the line $Y = 0.3$ in Figure A.2(c) is an oscillation of approximate amplitude 0.1 in the pressure. Essentially, the causal solution is present, but has added to it a spurious pressure mode. The presence of this spurious mode, small at present, will cause small deviations in velocity via the granular stress term, which in turn influence the pressure at the next timestep. Over many timesteps, this positive feedback eventually leads to the runaway instability shown in Figure A.2(b).

Attempts to control the instability by using a finer spatial resolution or a smaller timestep will not succeed in this case. Indeed, one would observe the wavelength of the spurious pressure modes to be similar to the lengthscale of the mesh elements, in some cases causing a more rapid degradation of the solution. This problem is very well known, and well documented in the finite element literature (see Chung, 2002, for an overview), and results from the following. When performing the pressure update, one attempts to adjust the pressure

in each element so that the net (Lagrangian) mass flux into every element is zero. The intermediate step involved is to modify the velocity at each vertex of the element according to the change in pressure. For a simple explicit timestepping scheme, such as the one used here, the relationship between the pressure change and the velocity change is linear. One could, therefore, describe the velocity update by the matrix equation

$$\begin{pmatrix} \delta U \\ \delta V \end{pmatrix} = \begin{pmatrix} \mathrm{K} \\ \mathrm{L} \end{pmatrix} \delta P, \tag{A.10}$$

where K and L are matrices of dimension $n_{\mathrm{pres}} \times n_{\mathrm{vel}}$. For our piecewise-linear velocity and piecewise constant pressure, velocity is held on the vertices, so $n_{\mathrm{vel}} = n_{\mathrm{vert}}$, and pressure is held on the triangles, so $n_{\mathrm{pres}} = n_{\mathrm{tri}}$. Unfortunately, for all useful computational meshes, $n_{\mathrm{tri}} > n_{\mathrm{vert}}$. This inevitably means that there will exist certain pressure modes that have no effect whatsoever on the velocity field. In other words, the matrices K and L have nontrivial null spaces. When attempting to update the pressure, zero eigenvectors belonging to the null space of pressure distributions will not affect the conservation of mass, so are not controlled during the pressure update. In a fluid dynamics code, this is often not a great drawback, as pressure is usually calculated only in order to determine its effect on velocity. The presence of null pressure eigenmodes does not, therefore, affect the dynamics. However, using the granular $\mu(I)$ rheology necessitates the accurate computation of pressure, so we must not allow such eigenmodes in the pressure computed by our finite element code.

One might hope to eliminate the null pressure modes by explicitly calculating all eigenvectors of the matrices K and L, and projecting out the null eigenvectors from the solution for pressure. Unfortunately, the intrinsic approximations made by the finite-precision code result in a continuous spectrum of eigenvalues, from which it is difficult to identify those whose corresponding eigenvectors should be retained and which should be discarded. It is also worth mentioning at this stage that it is not our iterative Gauss-Seidel scheme for the pressure update, nor the mass-lumping, that is at fault. An exact solution using the routines available in the LAPACK package[1] produces the same problematic results. In reality, two main approaches exist to overcome the problem of null pressure eigenmodes. Crucially, the null eigenmodes mainly have a wavelength of the order of the mesh spacing. This means

[1]Specifically, the DGETRF and DGRTRS routines, which solve by first computing a QR factorisation.

that various methods exist by which the pressure, once calculated, may be smoothed locally to largely eliminate the spurious modes. Several attempts at using these methods are documented in §A.3. An alternative, more mathematically-inclined approach is to review our choice of elements in order to reduce the number of degrees of freedom in pressure relative to those in the velocity. In this case, we aim to reduce (ideally to zero) the number of null eigenmodes. To this end, several different elements and schemes associated with them are considered in §A.4.

Aside from any concerns regarding the pressure, we draw attention to another potential difficulty that must be overcome in order to construct an accurate code. Under the Lagrangian framework, it is possible that individual elements may undergo significant deformations as their vertices move with the flow. Consider the mesh in Figure A.2(a). Even at this early stage of the flow, the triangular elements located along the yield surface have become elongated. Were the simulation to continue (without halting due to the pressure instability discussed above), one would expect these elements to stretch even further. This distortion of the elements can lead to several problems. Indeed, it may well be the long thin triangles near $(1.0, 0.1)$ in Figure A.2(e) that ultimately cause the divergence of the pressure iteration in the following few timesteps. One can be a little more specific regarding the effects of such large deformations. For example, the three vertices of a stretched element could be located at a significant distance from each other, relative to the vertices of unstretched elements. This can result in a strange sort of nonlocality, with information being moved around the mesh faster than one might expect given the average element size. A more worrying problem occurs when two very long, thin elements lie next to each other. In such a situation, it is possible for small numerical errors to allow a vertex to cross an element boundary, and cause two (or more) elements to overlap. This obviously violates the conservation of mass, and in practice leads rapidly to further numerical instabilities. In order to avoid issues such as these, it is necessary to dynamically improve the mesh according to some definition of the mesh quality. A discussion of the mesh refinement techniques used is presented in §A.5.

We conclude this section by identifying one final problem. We observe that the mesh shown in Figure A.2(a) actually appears to be deforming in agreement with our intuition. In this early stage of the collapse, the top corner of the column seems to be shearing along

a particular plane, leaving a region of (almost) plug motion on either side of the shearing zone. However, given the relatively small friction coefficient, $\mu_1 = 0.4$, we should expect the yield surface to be inclined at a much more shallow angle than that shown in Figure A.2(a). In this case, the shear plane is suspiciously closely aligned with the diagonal lines present in the mesh. Repeating the simulation using a mesh with elements of different aspect ratio (Figure A.2f) indicates that, indeed, the orientation of the shear plane is dependent upon the mesh used. One should expect to remedy this problem by using an unstructured mesh, instead of the regular mesh seen in Figure A.2. Until we discuss issues relating to the quality and type of mesh further in §A.5, however, we shall retain the structured initial mesh.

A.3 Pressure smoothing

As we demonstrated in §A.2, a naive finite element scheme using simple elements is susceptible to the influence of spurious pressure modes. In this section, we shall retain the simple elements and attempt to control the fluctuations in pressure by applying two different methods of smoothing.

The first method, which we shall refer to as *explicit smoothing*, employs a local averaging method by which the pressure on each triangle is replaced by a weighted average of the pressure on this and all adjacent triangles. Other than requiring additional logic to determine which triangles share an edge, this is simple to implement, and the weightings used in the average may, in principle, be tuned to improve the effect. The exact details of the weighting used do not have a significant effect on the quality of the smoothing achieved, so we elect to keep things simple and use the averaging update

$$P(\text{Triangle } i) \leftarrow \frac{3P(\text{Triangle } i) + \sum_j P(\text{Triangle } j)}{6}, \tag{A.11}$$

where the index j runs over the three triangles adjacent to triangle i. This averaging step may be repeated several times in each timestep if desired.

A more standard approach to pressure smoothing is to modify the conservation of mass equation in such a way as to suppress the spurious eigenmodes. A good review of several popular methods for 'pressure stabilisation' is presented by Tezduyar (1991). In this work, we consider only one of the many possible stabilisation strategies, which we shall refer to as

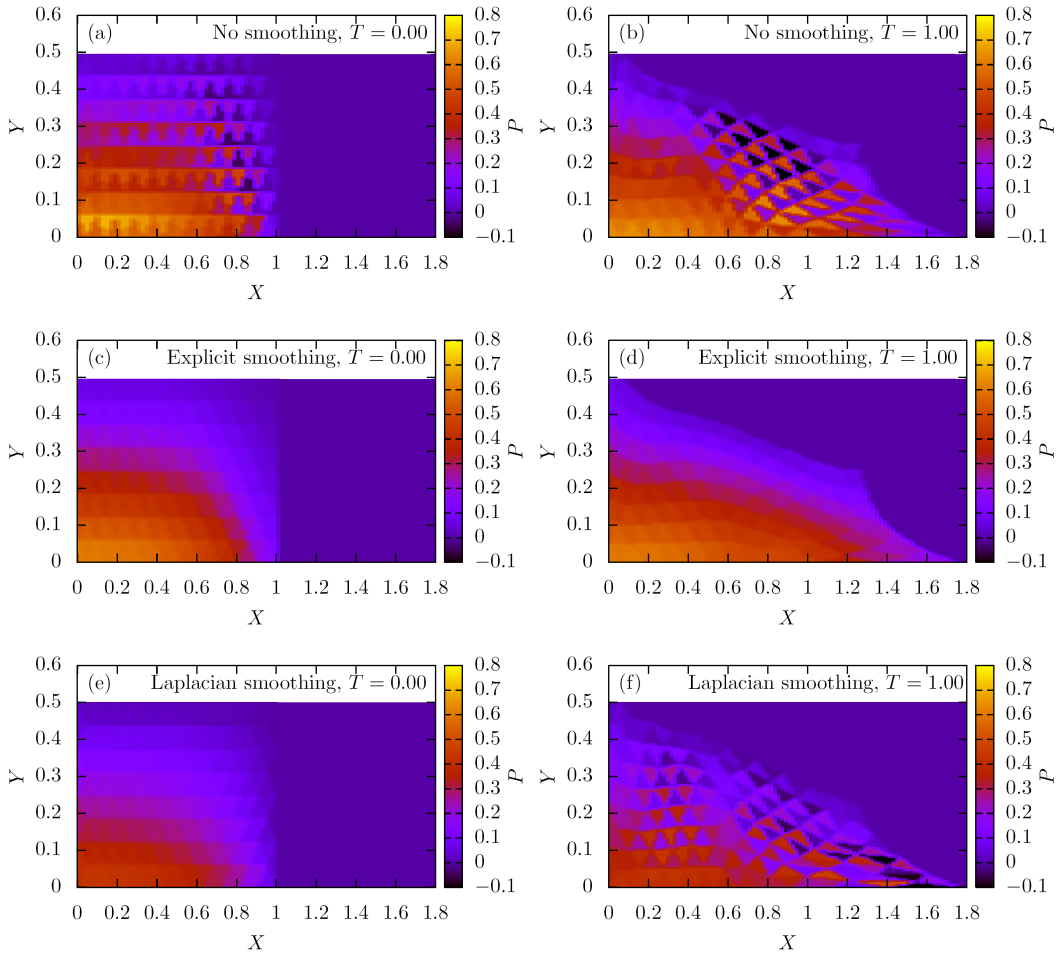

Figure A.3: Effects of pressure smoothing on Lagrangian finite element code. Colourmaps show the pressure distribution and have the same scales. Panels (a)-(h) show inviscid calculations under various conditions: (a) No smoothing, $T = 0$; (b) No smoothing, $T = 1$; (c) Explicit smoothing, $T = 0$; (d) Explicit smoothing, $T = 1$; (e) Laplacian smoothing, $T = 0$; (f) Laplacian smoothing, $T = 1$; (g) Both smoothing techniques, $T = 0$; (h) Both smoothing techniques, $T = 1$.

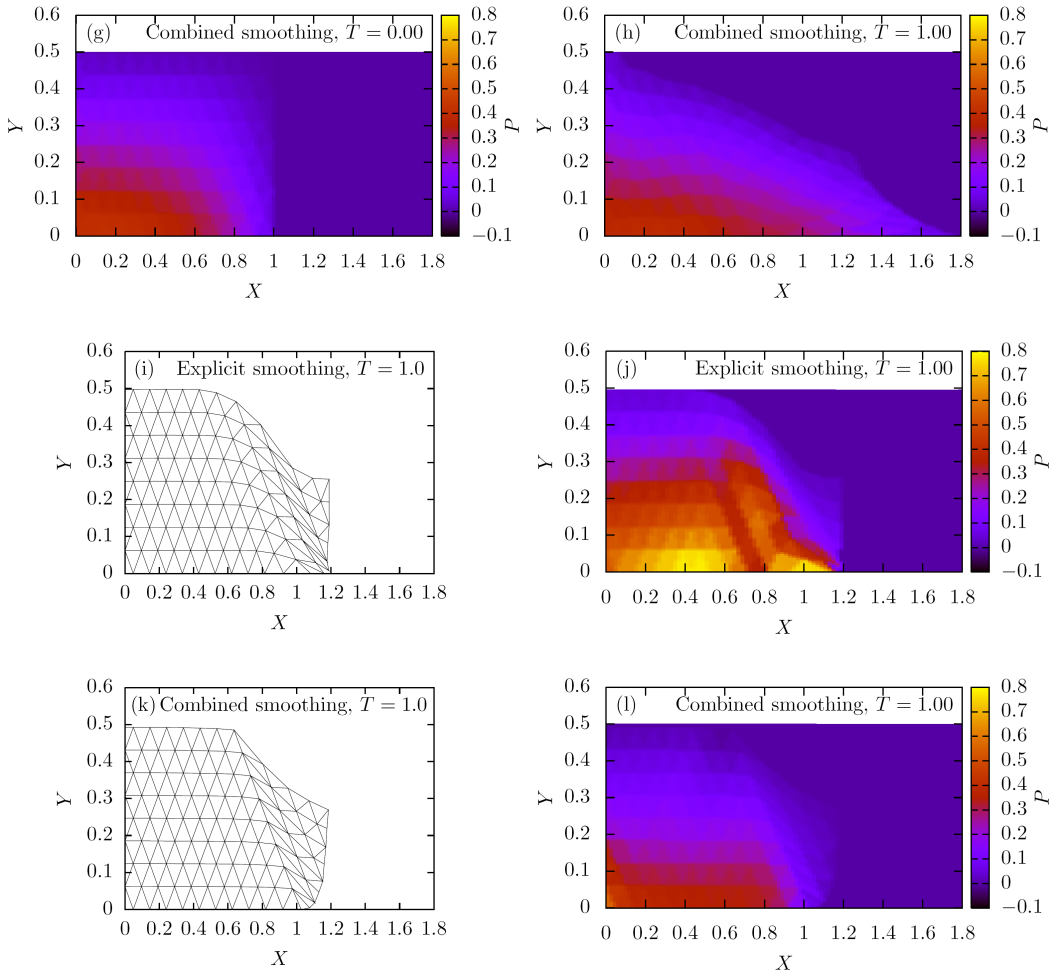

Figure A.3: (continued) Effects of pressure smoothing on Lagrangian finite element code. Panels (i)-(l) show calculations using the granular rheology. Panels (i) and (k) show the position of the mesh at $T = 1$, while panels (j) and (l) show the pressure distribution at $T = 1$ using (i,j) Explicit smoothing; (k,l) Both smoothing techniques. Computations using just the Laplacian smoothing technique fail quickly, and are not shown.

Laplacian smoothing. Given that the spurious modes have a typical wavelength that is proportional to the mesh lengthscale, h, one can replace the usual incompressibility constraint with

$$\nabla \cdot \mathbf{U} - \beta h^2 \nabla^2 P = 0, \tag{A.12}$$

where β is a constant whose optimal value is believed to be 0.025 for the Navier-Stokes equations. Any causal variation in the pressure should vary over an lengthscale of order unity, thus the errors introduced by the slightly incompressible condition (A.12) are $O(h^2)$, and negligible in our first-order scheme. The spurious pressure modes vary over a lengthscale of order h, and the $\nabla^2 P$ term acts to suppress this variation when we come to find the pressure needed to enforce incompressibility. For a piecewise constant pressure, the weak form of (A.12) used in the finite element method would require us to evaluate

$$\langle \nabla^2 \psi_i, \psi_j \rangle = -\langle \nabla \psi_i, \nabla \psi_j \rangle, \tag{A.13}$$

where ψ_i is the piecewise constant basis function equal to one on triangle i and zero otherwise. The discontinuity in pressure prevents the evaluation of such terms exactly, so we instead choose to approximate the Laplacian term by

$$(\nabla^2 P)(\text{Triangle } i) \approx \frac{3P(\text{Triangle } i) - \sum_j P(\text{Triangle } j)}{h^2}, \tag{A.14}$$

where the index j runs over the three triangles adjacent to triangle i, as before. This modification is also straightforward to implement, although it requires some calibration in order to determine a good value for β. A little experimentation, largely by trial and error, with the granular code suggests that one should take $\beta \approx 0.022$ for best results.

Figure A.3 presents the results of several versions of the finite element code, and attempts to illustrate the varying degree of success that can be achieved with these smoothing methods. First of all, we consider the pressure distribution in an inviscid flow calculation. In this case, any spurious pressure modes that may arise do not affect the dynamics of the mesh, but may grow from one timestep to the next. For reference, the pressure distribution calculated without any smoothing is shown in Figure A.3(a,b), at $T = 0$ and $T = 1$ respectively. The spurious eigenmodes are clearly visible in each case, and have grown to an amplitude of roughly unity by $T = 1$. If we apply three iterations of the explicit averaging update (A.11) at each timestep, we obtain the pressure distributions shown in Figure

A.3(c,d). The fluctuations in pressure have been nearly entirely eliminated at $T = 0$, and are only barely visible near the spreading front at $T = 1$. Unfortunately, however, a comparison with the exact solution (shown previously in Figure A.2d) reveals that the pressure at $T = 0$ is now slightly overestimated, probably due to an amplification of the fluctuation near the rigid boundary being swept into the mesh by the iterative averaging routine. Applying the slightly-compressible Laplacian smoothing, we obtain the pressure distributions shown in A.3(e,f). On this occasion, the pressure at $T = 0$ agrees very well with the exact solution, and is essentially free of spurious pressure oscillations. However, by $T = 1$ the oscillations have grown to a visible level, and are of similar magnitude to those occurring in the absence of any smoothing. One might hope that applying both of the smoothing techniques would result in a method that combines the best features of both approaches, and this indeed seems to be the case if we consider the results, shown in Figure A.3(g,h). In this experiment, the pressure appears to remain smooth throughout, and agrees well with the exact solution at $T = 0$.

As has been mentioned previously, the inviscid calculation has the advantage that any spurious pressure modes that may arise have no effect on the dynamics. By comparing the outline of the mesh at $T = 1$ for each of the four approaches in Figure A.3, one could be convinced that this is indeed the case. When working with the granular $\mu(I)$ rheology, however, we have already seen that any pressure fluctuations can have a profound destabilizing effect of the dynamics of the computational column. Indeed, an attempt to use only the Laplacian smoothing (which still allows spurious modes to appear in the inviscid case) with the granular code rapidly failed in a similar manner to that described in §A.2. No results of this simulation are shown here. We consider instead explicit averaging, both alone and in combination with the Laplacian smoothing. The pressure at $T = 0$ is not shown, because it will be identical to that calculated in the inviscid case (if $\mathbf{U} = 0$ everywhere, as is the case for a column initially at rest, the modified granular stresses vanish everywhere, and the pressure problem reduces to that solved in the inviscid calculation).

Using the explicit smoothing alone, it is clear from the pressure distribution at $T = 1$ (Figure A.3j) that something is not correct. In this case, spurious pressure behaviour arises near the spreading front of the granular current, and causes significant disruption to the mesh there, as shown in Figure A.3(i). This distortion of the mesh, as well as the asso-

ciated velocity field, feeds back on the pressure in successive timesteps, and produces the strange pressure distribution shown in Figure A.3(j). By contrast, a simulation performed with both the explicit and Laplacian smoothing techniques behaves rather well, with the pressure distribution (Figure A.3l) remaining both smooth and intuitively correct throughout. Furthermore, the mesh deforms in a reasonable manner, with no ill behaviour for these early times under consideration. We clearly observe a static core (with an essentially hydrostatic pressure distribution) from which the top corner shears off, and begins to settle on the horizontal surface. Combining the explicit and Laplacian smoothing in this manner is, therefore, a useful method by which one may suppress the spurious pressure modes for the problem of granular column collapse.

A.4 Different elements

Parallel to the development of the pressure smoothing techniques used in §A.3, the possibility of using different elements was considered. Recall that the undesirable pressure modes are zero eigenvectors of the coupling matrices K and L defined by (A.10), which describe the effect of changes in the pressure on the velocity. It is possible to select elements for which these matrices have no zero eigenvectors. In fact, one can derive a criterion known as the Babuška-Brezzi condition, or 'lim-sup' condition, that will specify whether or not the coupling matrices will have any zero eigenvectors (see Fortin & Brezzi, 1991). The details are not relevant – the useful outcome is that we can select elements proven to be 'Babuška-Brezzi stable' (or LBB stable, recognising the additional contribution of Ladyzhenskaya, 1969), and expect the resulting scheme to exhibit no spurious pressure behaviour.

Figure A.4 shows a schematic representation of several types of element, indexed by the differentiability of the velocity and pressure fields. Using notation common in the literature (see Gresho & Sani, 1998, for example), a PmPn element has a velocity field that whose $(m-1)^{\text{th}}$ derivative is continuous, and a pressure field whose $(n-1)^{\text{th}}$ derivative is continuous. Having m or n equal to zero means that we take the corresponding field to be piecewise constant, and hence discontinuous across element boundaries. In Figure A.4, we show the P1P0 elements used in the preceding sections along with three Babuška-Brezzi stable elements: P1P1, P1P0(macro) and P1P1(macro). In the latter two cases, the word 'macro'

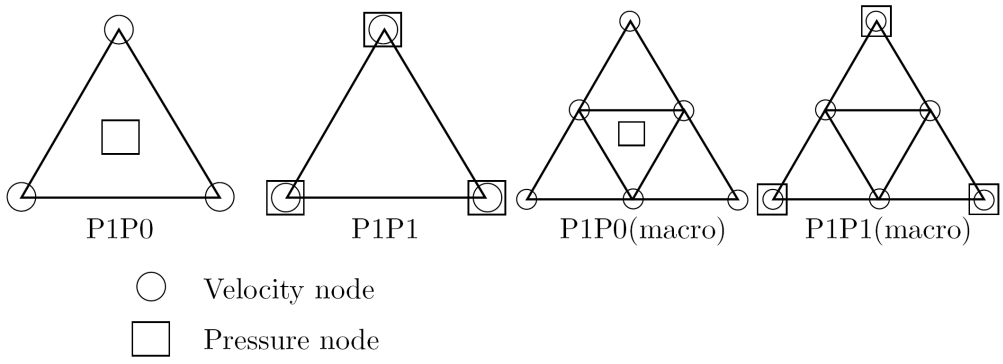

Figure A.4: Four examples of finite elements, showing the locations where velocity and pressure values are held. Interpolations are piecewise linear between values held on vertices, and piecewise constant if held on the centre of a triangle. Apart from the P1P0 element used previously, all examples are Babuška-Brezzi stable.

denotes that the velocity is a piecewise linear function over each triangle in the mesh, while the pressure is piecewise constant or linear over a patch of four similar triangles.

Many more element schemes exist that are Babuška-Brezzi stable. Gresho & Sani (1998) provide an extensive list of schemes along with their benefits and drawbacks, including some useful schemes that are not LBB-stable. However, the schemes not already discussed involve a higher order approximation to the velocity field. For example, the P2P1 (or Taylor-Hood) elements are Babuška-Brezzi stable elements that use piecewise quadratic velocity and piecewise linear pressure. The use of such elements, while requiring slightly more complicated code, should improve the accuracy from first to second order in the mesh resolution, and should result in a higher quality solution. Indeed, such elements are highly recommended for the simulation of viscous fluids. However, having velocity interpolated by a nonlinear function means that the strain rate $|\dot{\gamma}|$ will be non-constant on each triangle, and will almost always require the integrals associated with the granular frictional stress terms (which are proportional to $|\dot{\gamma}|^{-1}$) to be evaluated by quadrature. The benefit of a low-order scheme is that these integrals may be performed analytically. The numerical code necessary to include quadrature constitutes an additional complication that we would rather not pursue before exhausting other options. Accordingly, we shall not discuss these higher

order elements in any further detail.

We have implemented schemes for the column collapse problem using the three Babuška-Brezzi stable elements shown in Figure A.4. In order to study their stability properties and practicality for the problem at hand, it suffices to consider an inviscid fluid calculation. Some results of the inviscid simulations using these three schemes are shown in Figure A.5.

The P1P1 scheme achieves a smooth, accurate pressure distribution at $T = 0$ (Figure A.5a), with no trace of spurious eigenmodes. Unfortunately, the fact that pressure is held on the mesh vertices means that, essentially, conservation of mass is enforced on a vertex-by-vertex basis, rather than being explicitly enforced on each triangular element. Although this ought to maintain the global conservation of mass, it is possible that individual elements may grow or shrink during the simulation. An example of this behaviour may be seen in Figure A.5(b), particularly for elements near $(1.2, 0.1)$. The inset shows two triangles that began the simulation with the same area, but clearly are not the same size at $T = 1$. As the simulation progresses further, some triangles become extremely small, allowing diffusive instabilities to arise and destroy the solution.

A scheme based on the P1P0(macro) should be Babuška-Brezzi stable, and furthermore should explicitly conserve the area of each element. Indeed, the pressure distribution at $T = 0$ (Figure A.5c) is smooth, and the elements do not appreciably change in area during a simulation. However, the pressures predicted at $T = 0$ are noticeably too small. The exact solution has a maximum pressure (at the origin) of almost 0.5, whereas the pressure in Figure A.5(c) peaks at around 0.3, though it has the correct distribution. The reason for this underestimate is made more clear by the velocity distribution, as shown in Figure A.5(d). The horizontal velocity component exhibits an oscillation, aligned with the mesh, whose wavelength changes with the mesh resolution. The problem lies with our choice of a piecewise constant pressure. Consider the forces acting on the midpoint, $\mathbf{X}_{m3}$, of the element t_1 sketched in Figure A.6(a). The weak form of momentum equations, using a piecewise constant basis function for pressure, implies that the pressure force in this element can only exert a force normal to its edge upon the midpoint $\mathbf{X}_{m3}$. This means that the velocity at any mid-point may only comprise a fixed vertical component from gravity, and an additional component of variable magnitude in the direction normal to the edge. In our example, the midpoint $\mathbf{X}_{m3}$, whose edge is horizontal, is constrained to move only in the

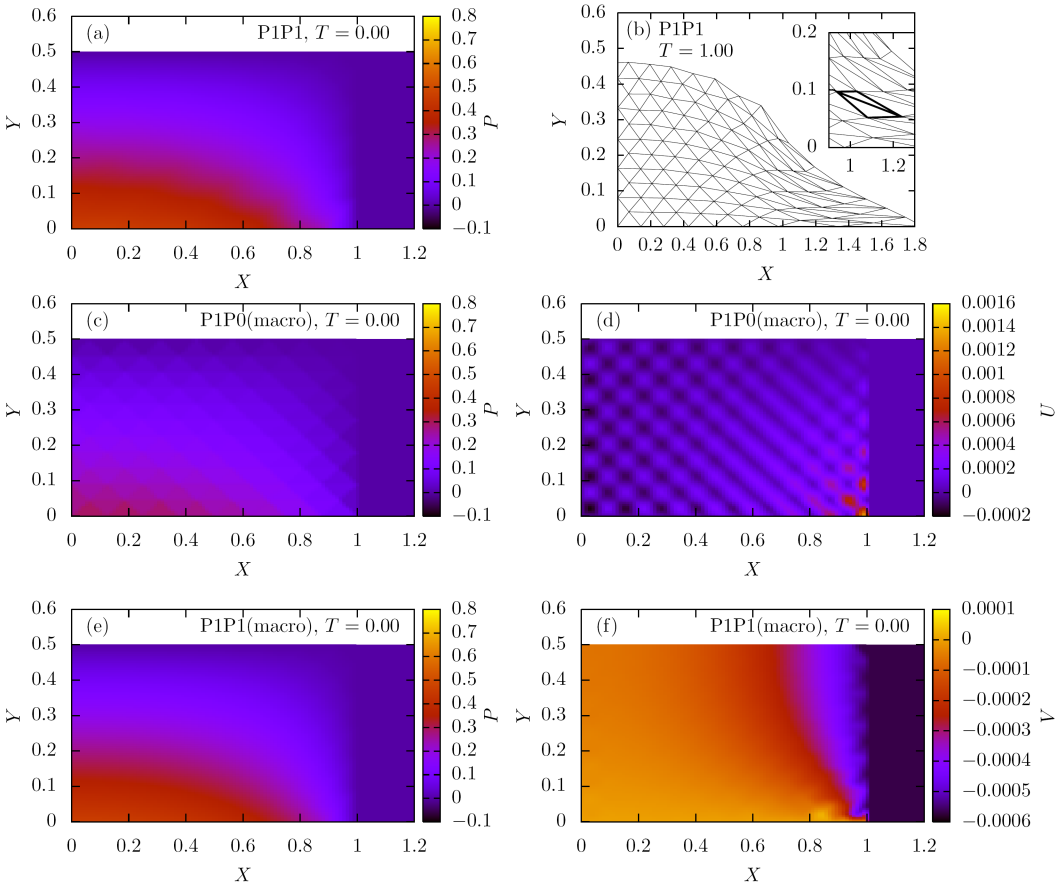

Figure A.5: Effects of using different elements for the Lagrangian finite element scheme: (a) Pressure at $T = 0$ using P1P1 elements; (b) Mesh position at $T = 1$ using P1P1 elements, inset shows a magnification of two adjacent triangles, which started the simulation with equal areas; (c) Pressure at $T = 0$ using P1P0(macro) elements; (d) Horizontal velocity at $T = 0$ using P1P0(macro) elements; (e) Pressure at $T = 0$ using P1P1(macro) elements; (f) Vertical velocity at $T = 0$ using P1P1(macro) elements.

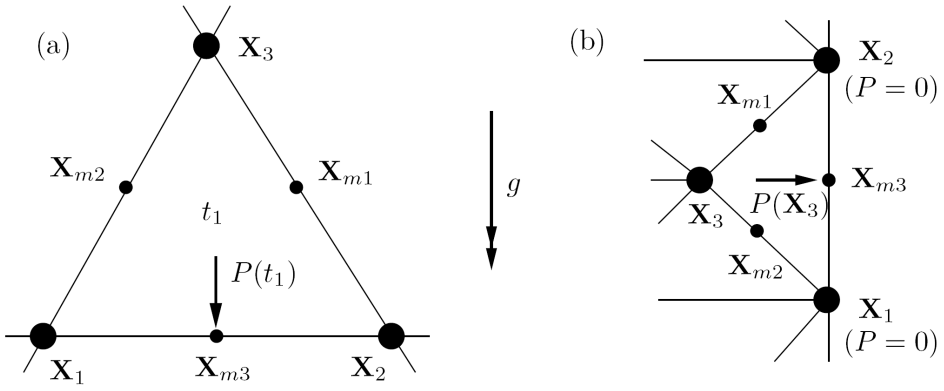

Figure A.6: Diagrams illustrating the problem of constrained velocities when using (a) P1P0(macro) elements and (b) P1P1(macro) elements. The triangles shown represent a patch of four similar triangles - that is, the triangles used to define the pressure basis functions.

vertical direction. It is this restriction on the mid-point velocities that causes the strange oscillation in velocity, and ultimately the under-estimation of the pressure.

Finally, we consider a scheme involving P1P1(macro) elements. As with the P1P1 elements, the pressure distribution at $T = 0$ (Figure A.5e) closely matches the analytical solution (A.9). On this occasion, however, one can observe spurious behaviour in the vertical velocity even at $T = 0$, as shown in Figure A.5(f). Although the velocity is essentially correct in the bulk of the flow, some strange fluctuations can be seen on the vertical free surface. These fluctuations depend upon the resolution of the mesh, but are present regardless of quality of the mesh. Once again, the problem is rooted in the way in which pressure influences velocity. Consider the free surface element shown in Figure A.6(b). The two vertices on the free surface are fixed at zero by the stress-free boundary condition, so only the pressure at vertex $\mathbf{X}_3$ may affect the velocity at midpoint $\mathbf{X}_{m3}$. Moreover, the linear basis function used for pressure means that the pressure force from $\mathbf{X}_3$ may only act in the direction normal to the edge containing $\mathbf{X}_{m3}$. In a similar manner to the case of P1P0(macro) elements, this restricts the form of the velocity on these free surface midpoints, and causes the spurious velocity fluctuations. Over time, these errors in the free

surface velocity have a significant impact upon the behaviour of the spreading front, thus rendering the P1P1(macro) scheme unusable for the column collapse problem.

A.5 Mesh quality and refinement

Having considered several of the Babuška-Brezzi stable elements with linear velocity, and found that none of them can be made to work with the column collapse geometry, we were forced to consider different options. Still hoping to avoid the need to include any quadrature routines when handling the granular stresses, we opt to build upon a simple P1P0 scheme, using the explicit averaging and Laplacian smoothing schemes described in §A.3.

As we touched upon at the end of §A.2, the form of the mesh may have a significant impact upon the accuracy of the solution. For the structured initial meshes used, we saw that the location of the yield surface, for example, was dependent upon the aspect ratio of the triangular elements. Furthermore, we noted that the Lagrangian approach of moving vertices means that triangles can become extremely stretched, especially in regions of large shear. These stretched triangles can cause strange, almost non-local dependencies in the pressure and velocity fields, and ultimately can result in the divergence of the iterative scheme used to find the pressure. In this section, we give a brief discussion of our efforts to eliminate these problems by using various routines designed to maintain a good quality mesh.

We must first define what we mean by a good quality mesh. For a finite element code, it is generally accepted that the bad behaviour described above can be minimised by maximising the smallest angle in any triangle. Given a set of vertices satisfying a mild set of conditions, this maximum may be attained by constructing a Delaunay triangularisation of the vertices (Delaunay, 1934). This triangularisation has the property that interior of the circumcircle of each triangle contains no vertices belonging to other triangles.

In the context of the Lagrangian finite element scheme, the fact that the vertices are constantly in motion means that the Delaunay triangulation of the granular domain may change frequently. One could perform a global Delaunay remeshing at every timestep, but this is very time-consuming. Instead, we use a set of small local updates at each timestep (see Edelsbrunner, 2000, for example), designed to either improve or remove the 'worst'

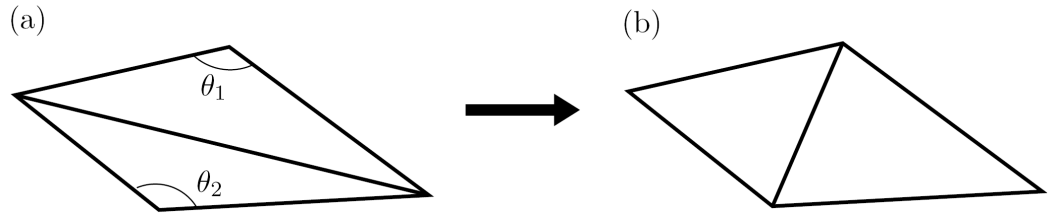

Figure A.7: Diagram to illustrate the 'flips' used to maintain a Delaunay triangulation. (a) A pair of 'bad', non-Delaunay triangles, and (b) the 'good' triangles that result from a flipped diagonal.

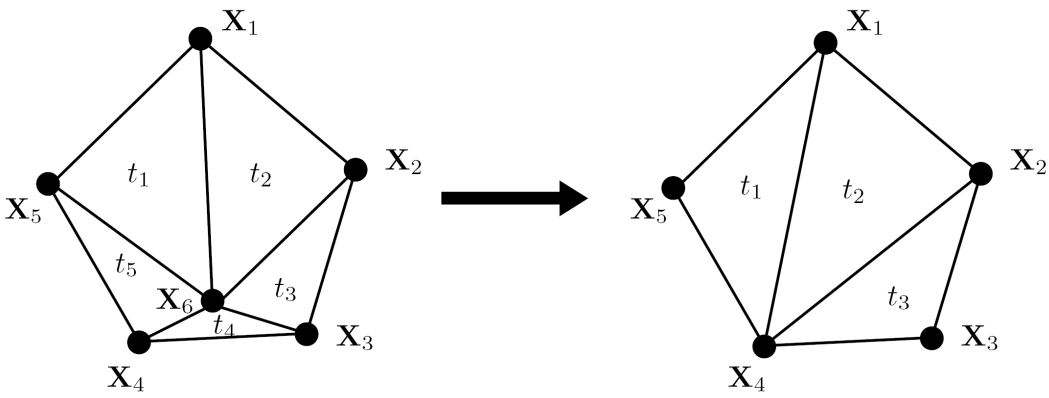

Figure A.8: Diagram illustrating the removal of a vertex. In order to remove vertex $\mathbf{X}_6$, one can imagine merging it with vertex $\mathbf{X}_4$. This action destroys triangles t_4 and t_5.

triangles and maintain an approximately Delaunay mesh. These updates are discussed only briefly here. For a more thorough discussion of these strategies, as well as their extension to three dimensions, one might like to consult the Ph.D thesis of Morrison (2008), who implemented many such local mesh refinements in both two and three dimensions. Consider the two triangles shown in Figure A.7(a). We declare them to be of poor quality, and thus candidates for adjustment, if

$$\theta_1 + \theta_2 > \pi. \tag{A.15}$$

If this condition is met, then we flip the diagonal so that we are left with the configuration shown in Figure A.7(b), in which both triangles now have a more favourable aspect ratio. Unfortunately, such flips can, in rare situations, result in the creation of very small triangles, which in turn make the scheme more susceptible to numerical (typically diffusional) instabilities. One can attempt to eliminate a very small triangle by removing one of its vertices from the mesh, and adjusting the neighbouring triangles appropriately, as shown by Figure A.8.

Aside from our desire to improve the quality of the mesh, one may like to increase or decrease the spatial resolution in certain areas of the mesh. For example, along the yield surface, where the shear is greatest, it would be beneficial to have a finer mesh resolution. By contrast, in the static portion of the column, where there is no motion, such increased resolution would be wasteful. With this in mind, we have implemented a simple routine that adds vertices to the centre of triangles with $|\dot{\gamma}|$ greater than some threshold value $|\dot{\gamma}|_{\text{add}}$, while a second routine attempts to remove any triangles with strain rate smaller than a lower threshold value $|\dot{\gamma}|_{\text{remove}}$. Because of the changes to triangle areas made by all of the routines described so far, it is necessary to adapt the timestep in order to avoid any diffusional instabilities that may arise if the timestep is too large compared to the mesh lengthscales. Inspired by the commonly-used condition

$$\Delta T < \frac{1}{4}\text{Re}\Delta X^2 \tag{A.16}$$

for a finite difference scheme for fluid flow at Reynolds number Re, we choose to set

$$\Delta T = \frac{a\epsilon}{10\mu_1} \min_i A_i, \tag{A.17}$$

where A_i is the area of triangle i.

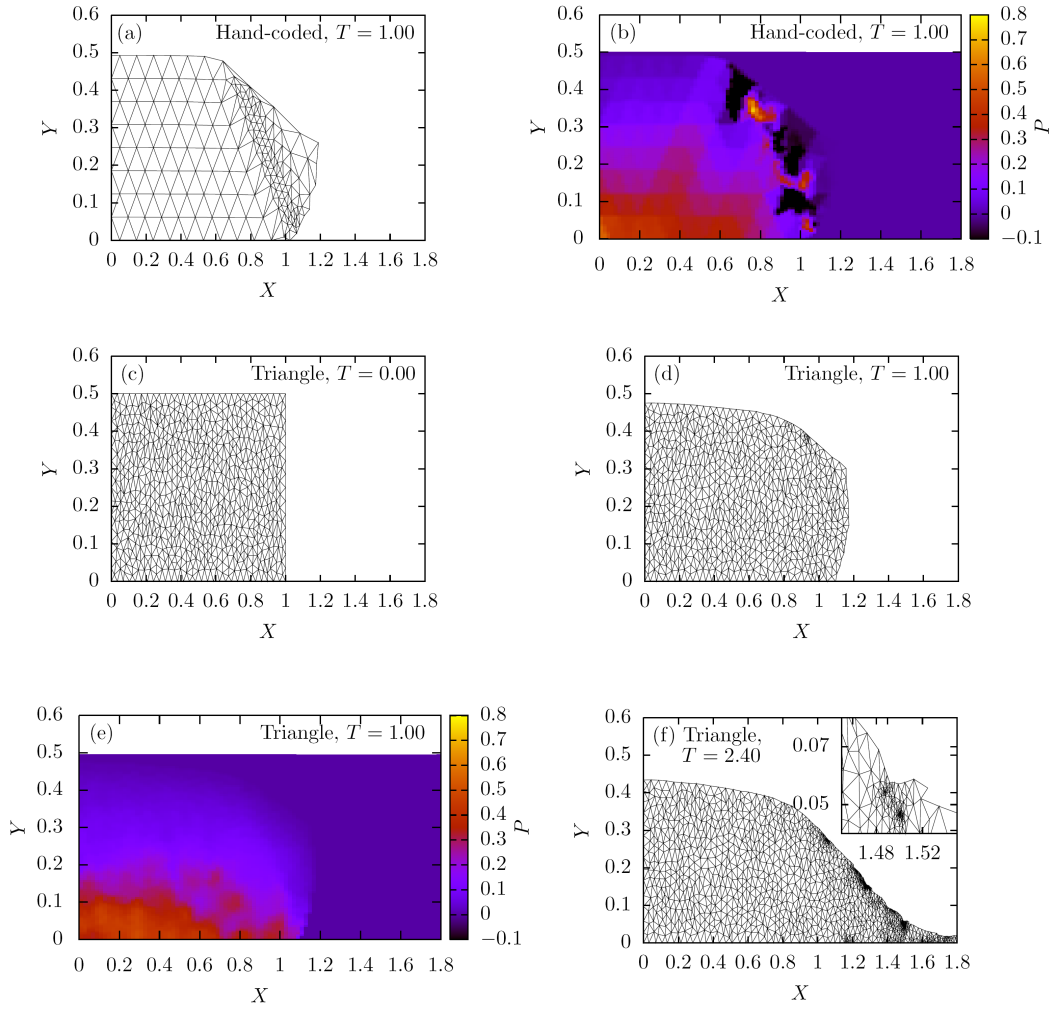

Figure A.9: The effect of mesh refinement on the Lagrangian finite element scheme. (a-b) Mesh and pressure distribution at $T = 1$ from a code using our hand-coded mesh generation routines. (c-f) Output from a code linked to the Triangle mesh generation software: (c) Mesh at $T = 0$; (d) Mesh at $T = 1$; (e) Pressure at $T = 1$; (f) Mesh at $T = 2.4$, with an inset magnifying the rough free surface and accompanying very fine resolution.

After having written and tested local update routines for Delaunay flips, small triangle removal, and increasing or decreasing the resolution, it was possible to apply the mesh refinement process to the column collapse code. The initial results, shown in Figure A.9(a), look promising. There are no long, thin triangles in the mesh, and the resolution has been made more fine along the shearing/yield surface. One might like to compare this mesh to that shown in Figure A.3(k), which is generated by code identical save for the lack of mesh refinement routines. Unfortunately, making such changes to the mesh appears to have a detrimental effect on the smoothness of the pressure distribution, which is shown in Figure A.9(b). Despite extensive testing and modification, a successful adaptive mesh code that maintains a smooth pressure distribution was not found.

This was perhaps, in retrospect, a little ambitious. After all, entire theses have been devoted to algorithms for mesh generation and refinement, and the problem of generating a quality mesh for a wide range of applications is still very much open. Eventually, the home-grown mesh refinement code was abandoned in favour of using well-developed, freely-available mesh generation software. In this case, we use Triangle, by Shewchuk (1996). This program is capable of generating a good mesh for a given geometry, subject to various constraints on the mesh quality (e.g. minimum angles, maximum area). After gaining an understanding of Triangle, and its input and output data structures, it is possible to implement a new mesh refinement strategy. At regular intervals, we output the vertex data in a format readable by Triangle, which then creates a new mesh using the existing vertices (possibly introducing new vertices). We then read in this new mesh, and continue with the finite element simulation. In addition, Triangle may be used to generate an unstructured initial grid for the simulation, in the hope of removing the structure-dependent features mentioned at the end of §A.2. An example of such a grid is shown in Figure A.9(c). After allowing the simulation to progress, we find that by $T = 1$ the mesh takes the plausible shape shown in Figure A.9(d). At this point, the pressure distribution (Figure A.9e) is not entirely smooth, and exhibits some minor fluctuations near the rigid base. Nevertheless, the general trend of the pressure is in good agreement with the distributions observed earlier in §A.3, and there seems to be no adverse effect on the dynamics. Unfortunately, as the simulation progresses yet further, the free surface becomes noticeably rough, with several small concavities arising in otherwise convex sections of the the surface. The constraints

adhered to by Triangle then require that we dramatically increase the resolution in the neighbourhood of these concavities. Ultimately, the situation becomes similar to that shown by Figure A.9(f). Several patches of very high resolution arise in areas that do not have a complicated flow structure, and some strange overhanging surface effects can develop, as seen in the inset. Given the 'black box' nature of the Triangle program, it is difficult to suppress the emergence of such phenomena. Indeed, they are present in essentially every simulation attempted, and eventually cause the code to fail by introducing smaller and smaller triangles, continuing to the point where the timestep needed to avoid diffusional instability becomes impractically small.

A.6 Epilogue

Despite the great amount of effort expended in attempting to develop a numerical scheme suitable for the problem of granular column collapse, we have ultimately been unsuccessful. Nevertheless, we have made a good deal of progress towards understanding how such a scheme could be constructed. Our best attempt as a working code involves simple P1P0 elements, with both Laplacian and explicit smoothing to suppress any spurious pressure modes. Promising results can be obtained by using a freely-available mesh generation program to maintain a good quality mesh. With this program, it is possible to maintain control over the pressure, whilst retaining a high quality mesh. However, the repeated process of mesh regeneration seems to cause a lack of smoothness in the free surface. While working on a solution to this problem, the author was made aware of an alternative numerical free-surface approach that might be appropriate for solving a problem like that of the collapse of a granular column. This alternative, and ultimately successful, approach, was the marker-and-cell method of Harlow & Welch (1965), and is discussed in the main text in §3.2.2.

Bibliography

Aranson, I. S. & Tsimring, L. S. 2002 Continuum theory of partially fluidized granular flows. *Phys. Rev. E* **65**, 061303.

Bagnold, R. A. 1954 Experiments on a gravity-free dispersion of large solid spheres in a newtonian fluid. *Proc. Roy. Soc. Lond. A* **225**, 49.

Ball, R. C. & Blumenfeld, R. 2002 Stress field in granular systems: Loop forces and potential formulation. *Phys. Rev. Lett.* **88** (11), 115505.

Balmforth, N. J., Bush, J. W. M., Vener, D. & Young, W. R. 2007 Dissipative descent: rocking and rolling down an incline. *Journal of Fluid Mechanics* **590**, 295–318.

Balmforth, N. J., Cawthorn, C. J. & Craster, R. V. 2010 Contact in a viscous fluid. Part 2. A compressible fluid and an elastic solid. *J. Fluid Mech.* **646**, 339–361.

Balmforth, N. J. & Kerswell, R. R. 2005 Granular collapse in two dimensions. *J. Fluid Mech.* **538**, 399–428.

Benjamin, T. B. 1959 Shearing flow over a wavy boundary. *Journal of Fluid Mechanics* **6**, 161–205.

Bertho, Y., Gioriutti-Dauphiné, F. & Hulin, J.-P. 2003 Dynamical Janssen effect on granular packing with moving walls. *Phys. Rev. Lett.* **90**, 144301.

Börzsönyi, T. & Ecke, R. E. 2007 Flow rule of dense granular flows down a rough incline. *Phys. Rev. E* **76**, 031301.

Börzsönyi, T., Ecke, R. E. & McElwaine, J. N. 2009 Patterns in flowing sand: Understanding the physics of granular flow. *Phys. Rev. Lett.* **103** (17), 178302.

Bouchaud, J. P., Cates, M. E., Prakash, J. R. & Edwards, S. F. 1994 A model for the dynamics of sandpile surfaces. *J. Phys. I France* **4**, 1383–1410.

Boudet, J. F., Amarouchene, Y., Bonnier, B. & Kellay, H. 2007 The granular jump. *J. Fluid Mech.* **572**, 413–431.

Brenner, H. 1961 The slow motion of a sphere through a viscous fluid towards a plane surface. *Chemical Engineering Science* **16**, 242 – 251.

Campbell, C. S. & Brennen, C. E. 1985 Chute flows of granular material: Some computer simulations. *Journal of Applied Mechanics* **52**, 172–178.

Campbell, C. S. & Gong, A. 1986 The stress tensor in a two-dimensional granular shear flow. *Journal of Fluid Mechanics* **164**, 107–125.

Cassar, C., Nicolas, M. & Pouliquen, O. 2005 Submarine granular flows down inclined planes. *Phys. Fluids* **17**, 103301.

Cawthorn, C. J. & Balmforth, N. J. 2010 Contact in a viscous fluid. Part 1. A falling wedge. *J. Fluid Mech.* **646**, 327–338.

Charru, F. & Hinch, E. J. 2000 'Phase diagram' of interfacial instabilities in a two-layer Couette flow and mechanism of the long-wave instability. *Journal of Fluid Mechanics* **414**, 195–223.

Charru, F. & Hinch, E. J. 2006*a* Ripple formation on a particle bed sheared by a viscous liquid. Part 1. Steady flow. *Journal of Fluid Mechanics* **550**, 111–121.

Charru, F. & Hinch, E. J. 2006*b* Ripple formation on a particle bed sheared by a viscous liquid. Part 2. Oscillating flow. *Journal of Fluid Mechanics* **550**, 123–137.

Charru, F. & Mouilleron-Arnould, H. 2002 Instability of a bed of particles sheared by a viscous flow. *Journal of Fluid Mechanics* **452**, 303–323.

Chen, H. & Lee, C. F. 2003 A dynamic model for rainfall-induced landslides on natural slopes. *Geomorphology* **51**, 269–288.

Chen, W. F. & Baladi, G. Y. 1985 *Soil plasticity: Theory and implementation*. Elsevier.

Chevoir, F., Prochnow, M., Moucheront, P., da Cruz, F., Bertrand, F., Guilbaud, J.-P., Coussot, P. & Roux, J.-N. 2001 Dense granular flows in a vertical chute. In *Powders and Grains* (ed. Kishino), pp. 399–402. Swets & Zeitlinger, Lisse.

Chung, T. J. 2002 *Computational Fluid Dynamics*. Cambridge University Press.

Coulomb, C. A. 1776 Essai sur une application des régles des maximis et minimis a quelquels problémes de statique relatifs, a la architecture. *Mem. Acad. Roy. Div. Sav.* **7**, 343–378.

Courrech du Pont, S., Gondret, P., Perrin, B. & Rabaud, M. 2003 Granular avalanches in fluids. *Phys. Rev. Lett.* **90**, 044301.

da Cruz, F., Emam, S., Prochnow, M., Roux, J.-N. & Chevoir, F. 2005 Rheophysics of dense granular materials: Discrete simulation of plane shear flows. *Phys. Rev. E* **72**, 021309.

Daniel, R. C., Poloski, A. P. & Sáez, A. E. 2007 A continuum constitutive model for cohesionless granular flows. *Chem. Eng. Sci.* **62**, 1343–1350.

Davis, R. H., Zhao, Y., Galvin, K. P. & Wilson, H. J. 2003 Solid-solid contacts due to surface roughness and their effects on suspension behaviour. *Royal Society of London Philosophical Transactions Series A* **361**, 871–894.

de Ryck, A., Ansart, R. & Dodds, J. A. 2008 Granular flows down inclined channels with a strain-rate dependent friction coefficient. Part I: Non-cohesive materials. *Granular Matter* **10**, 353–360.

Deboeuf, S., Lajeunesse, E., Dauchot, O. & Andreotti, B. 2006 Flow rule, self-channelization, and levees in unconfined granular flows. *Phys. Rev. Lett.* **97**, 158303.

Delaunay, B. 1934 Sur la sphére vide. *Izv. Akad. Nauk. SSSR, Otdelenie Matematicheskikh i Yestestvennykh Nauk* **7**, 793–800.

Doyle, E. E., Huppert, H. E., Lube, G., Mader, H. M. & Sparks, R. S. J. 2007 Static and flowing regions in granular collapses down channels: Insights from a sedimenting shallow water model. *Physics of Fluids* **19**, 106601.

Drazin, P. G. & Reid, W. H. 1981 *Hydrodynamic Stability*. Cambridge University Press.

Edelsbrunner, H. 2000 Triangulations and meshes in computational geometry. *Acta Numerica* **9**, 133–213.

Fenistein, D., van de Meent, J. W. & van Hecke, M. 2004 Universal and wide shear zones in granular bulk flow. *Phys. Rev. Lett.* **92**, 094301.

Forterre, Y. 2006 Kapiza waves as a test for three-dimensional granular flow rheology. *J. Fluid Mech.* **563**, 123–132.

Forterre, Y. & Pouliquen, O. 2001 Longitudinal vortices in granular flows. *Phys. Rev. Lett.* **86**, 5886–5889.

Forterre, Y. & Pouliquen, O. 2003 Long-surface-wave instability in dense granular flows. *J. Fluid Mech.* **486**, 21–50.

Forterre, Y. & Pouliquen, O. 2008 Flows of dense granular media. *Ann. Rev. Fluid Mech.* **40**, 1–24.

Fortin, M. & Brezzi, F. 1991 *Mixed and hybrid finite element methods*. Springer.

GDR MiDi 2004 On dense granular flows. *Eur. Phys. J. E* **14**, 341–366.

de Gennes, P. G. 1999 Granular matter: A tentative view. *Rev. Mod. Phys* **71**, S374–S382.

Gérard-Varet, D. & Hillairet, M. 2010 Regularity issues in the problem of fluid structure interaction. *Arch. Ration. Mech. An.* **195**, 375–407.

Goddard, J. D. 1990 Nonlinear elasticity and pressure-dependent wave speeds in granular media. *Proc. Roy. Soc. Lond. A* **430**, 105–131.

Goldhirsch, I. & Goldenberg, C. 2002 On the microscopic foundations of elasticity. *Eur. Phys. J. E* **9**, 245–251.

Goldhirsch, I. & Zanetti, G. 1993 Clustering instability in dissipative gases. *Phys. Rev. Lett.* **70**, 1619–1622.

Gopinath, A., Chen, S. B. & Koch, D. L. 1997 Lubrication flows between spherical particles colliding in a compressible non-continuum gas. *J. Fluid Mech.* **344**, 245–269.

Gopinath, A. & Koch, D. L. 2002 Collision and rebound of small droplets in an incompressible continuum gas. *J. Fluid Mech.* **454**, 145–201.

Gray, J. M. N. T. 2001 Granular flow in partially filled slowly rotating drums. *J. Fluid. Mech.* **441**, 1–29.

Gray, J. M. N. T., Wieland, M. & Hutter, K. 1999 Gravity-driven free surface flow of granular avalanches over complex basal topography. *Proc. Roy. Soc. Lond. A* **455**, 1841–1874.

Gresho, P. M. & Sani, R. L. 1998 *Incompressible flow and the finite element method*. Wiley.

Griggs, A. J., Zinchenko, A. Z. & Davis, R. H. 2008 Gravity-driven motion of a drop or bubble near an inclined plane at low reynolds number. *Int. J. Multiphase Flow* **34**, 408–418.

Hákonardóttir, K. M. & Hogg, A. J. 2005 Oblique shocks in rapid granular flows. *Phys. Fluids* **17**, 077101.

Hákonardóttir, K. M., Hogg, A. J., Batey, J. & Woods, A. W. 2003 Flying avalanches. *Geophys. Res. Lett.* **30** (23), 2191.

Harlow, F. H. & Welch, J. E. 1965 Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *Phys. Fluids* **8**, 2182–2189.

Hinton, E., Rock, T. & Zienkiewicz, O. C. 1975 A note on mass lumping and related processes in the finite element method. *Earthq. Eng. Struct. D.* **4**, 245–249.

Hocking, L. M. 1973 The effect of slip on the motion of a sphere close to a wall and of two adjacent spheres. *J. Engng. Math.* **7**, 207–221.

Howell, D., Behringer, R. P. & Veje, C. 1999 Stress fluctuations in a 2d granular couette experiment: A continuous transition. *Phys. Rev. Lett.* **82**, 5241–5244.

Hoyle, R. B. & Woods, A. W. 1997 Analytical model of propagating sand ripples. *Phys. Rev. E* **56**, 6861–6868.

Hunt, J. C. R., Leibovich, S. & Richards, K. J. 1988 Turbulent shear flows over low hills. *Q. J. Roy. Meteor. Soc.* **114**, 1435–1470.

Hutter, K., Koch, T., Pluüss, C. & Savage, S. B. 1995 The dynamics of avalanches of granular materials from initiation to runout. part ii. experiments. *Acta Mechanica* **109**, 127–165.

Jaeger, H. M., Nagel, S. R. & Behringer, R. P. 1996 Granular solids, liquids and gases. *Rev. Mod. Phys.* **68**, 1259.

Janssen, H. A. 1895 Versuche uber getreidedruck in silozellen. *Zeitschr. d. Vereines deutscher Ingenieure* **39**, 1045–1049.

Jenkins, J. T. & Savage, S. B. 1983 A theory for the rapid flow of identical, smooth, nearly elastic, spherical particles. *Journal of Fluid Mechanics* **130**, 187–202.

Jop, P. 2008 Hydrodynamic modeling of granular flows in a modified Couette cell. *Phys. Rev. E* **77**, 032301.

Jop, P., Forterre, Y. & Pouliquen, O. 2005 Crucial role of sidewalls in granular surface flows: Consequences for the rheology. *J. Fluid Mech.* **541**, 167–192.

Jop, P., Forterre, Y. & Pouliquen, O. 2006 A constitutive law for dense granular flows. *Nature* **441**, 727–731.

Kadanoff, L. P. 1999 Built upon sand: Theoretical ideas inspired by granular flows. *Rev. Mod. Phys.* **71**, 435–444.

Kamrin, K. & Bazant, M. Z. 2007 Stochastic flow rule for granular materials. *Phys. Rev. E* **75**, 041301.

Kennedy, J. F. 1969 The formation of sediment ripples, dunes, and antidunes. *Ann. Rev. Fluid Mech.* **1**, 147–168.

Kerswell, R. R. 2005 Dam break with coulomb friction: A model for granular slumping? *Phys. Fluids* **17**, 057101.

Komatsu, T. S., Inagaki, S., Nakagawa, N. & Nasuno, S. 2001 Creep motion in a granular pile exhibiting steady surface flow. *Phys. Rev. Lett.* **86**, 1757–1760.

Kroy, K., Sauermann, G. & Herrmann, H. J. 2002 Minimal model for sand dunes. *Phys. Rev. Lett.* **88**, 054301.

Lacaze, L. & Kerswell, R. R. 2009 Axisymmetric granular collapse: A transient 3D flow test of viscoplasticity. *Phys. Rev. Lett.* **102**, 108305.

Ladyzhenskaya, O. A. 1969 *The mathematical theory of viscous incompressible flow*, 2nd edn. Gordon and Breach, New York.

Lajeunesse, E., Mangeney-Castelnau, A. & Vilotte, J. P. 2004 Spreading of a granular mass on a horizontal plane. *Phys. Fluids* **16**, 2371–2381.

Lajeunesse, E., Monnier, J. B. & Homsy, G. M. 2005 Granular slumping on a horizontal surface. *Phys. Fluids* **17**, 103302.

Lajeunesse, E., Quantin, C., Allemand, P. & Delacourt, C. 2006 New insights on the runout of large landslides in the valles-marineris canyons, mars. *Geophys. Res. Lett.* **33**, L00403.

Larrieu, E., Staron, L. & Hinch, E. J. 2006 Raining into shallow water as a description of the collapse of a column of grains. *J. Fluid Mech.* **554**, 259–270.

Lauga, E., Brenner, M. P. & Stone, H. A. 2007 Microfluidics: The no-slip boundary condition. In *Handbook of Experimental Fluid Dynamics* (ed. C. Tropea, A. Yarin & J. F. Foss). Springer.

Lemaître, A., Roux, J.-N. & Chevoir, F. 2009 What do dry granular flows tell us about dense non-brownian suspension rheology? *Rheologica Acta* **48**, 925–942.

Liu, C. H. & Joseph, D. D. 1977 Stokes flow in wedge-shaped trenches. *J. Fluid Mech.* **80**, 443–463.

Losert, W., Bocquet, L., Lubensky, T. C. & Gollub, J. P. 2000 Particle dynamics in sheared granular media. *Phys. Rev. Lett.* **85**, 1428–1431.

Lube, G., Huppert, H. E., Sparks, R. S. J. & Freundt, A. 2005 Collapses of two-dimensional granular columns. *Phys. Rev. E* **72**, 041301.

Lube, G., Huppert, H. E., Sparks, R. S. J. & Hallworth, M. A. 2004 Axisymmetric collapse of granular columns. *J. Fluid Mech.* **508**, 175–199.

Mangeney-Castelnau, A., Bouchut, F., Vilotte, J. P., Lajeunesse, E., Aubertin, A. & Pirulli, M. 2005 On the use of Saint Venant equations to simulate the spreading of a granular mass. *J. Geophys. Res.* **110**, B09103.

Martin, J. C. & Moyce, W. J. 1952 Part IV. An experimental study of the collapse of liquid columns on a rigid horizontal plane. *Philos. T. Roy. Soc. A* **244**, 312–324.

Möbius, M. E., Lauderdale, B. E., Nagel, S. R. & Jaeger, H. M. 2001 Brazil-nut effect: Size separation of granular particles. *Nature* **414**, 270.

Moffatt, H. K. 1964 Viscous and resistive eddies near a sharp corner. *J. Fluid Mech.* **18**, 1–18.

Mohan, L. S., Nott, P. R. & Rao, K. K. 1999 A frictional cosserat model for the flow of granular materials through a vertical channel. *Acta Mechanica* **138**, 75–96.

Morrison, N. F. 2008 Computations of flows of elastic liquids. PhD thesis, University of Cambridge.

Mouilleron, H. 2002 Instabilités d'un milieu granulaire cisaillé par un fluide. PhD thesis, Université Paul Sabatier, Toulouse, France.

Ouriemi, M. 2007 Erosion, transport et instibilités d'un lit de particules dans un tube. PhD thesis, Université de Provence.

Pailha, M. & Pouliquen, O. 2009 A two-phase flow description of the initiation of underwater granular avalanches. *J. Fluid Mech.* **633**, 115–135.

Plouraboué, F. & Boehm, M. 1999 Multi-scale roughness transfer in cold metal rolling. *Tribol. Int.* **32**, 45–57.

Pöschel, T. & Brilliantov, N. 2003 *Granular gas dynamics*. Springer.

Pouliquen, O. 1999 Scaling laws in granular flows down rough inclined planes. *Phys. Fluids* **11**, 542–548.

Pouliquen, O., Delour, J. & Savage, S. B. 1997 Fingering in granular flows. *Nature* **386**, 816–817.

Pouliquen, O. & Forterre, Y. 2002 Friction law for dense granular flows: Application to the motion of mass down a rough inclined plane. *J. Fluid Mech.* **453**, 133–151.

Pouliquen, O. & Gutfraind, R. 1996 Stress fluctuations and shear zones in quasistatic granular chute flows. *Phys. Rev. E* **53**, 552–561.

Prigozhin, L. 1999 Nonlinear dynamics of Aeolian sand ripples. *Phys. Rev. E* **60**, 729–733.

Rajchenbach, J. 1990 Flow in powders: From discrete avalanches to continuous regime. *Phys. Rev. Lett.* **65**, 2221–2224.

Rajchenbach, J. 2000 Granular flows. *Adv. Phys.* **49**, 229–256.

Raudkivi, A. J. 1998 *Loose boundary hydraulics*. A. A. Balkema, Rotterdam.

Reynolds, O. 1886 On the theory of lubrication and its application to Mr. Beauchamp Tower's experiments, including an experimental determination of the viscosity of olive oil. *Philos. T. Roy. Soc. I* **177**, 157–234.

Ries, A., Wolf, D. E. & Unger, T. 2007 Shear zones in granular media: 3D contact dynamics simulation. *Phys. Rev. E* **76**, 051301.

Rognon, P. G., Roux, J.-N., Wolf, D., Naaïm, M. & Chevoir, F. 2006 Rheophysics of cohesive granular materials. *Europhys. Lett.* **74**, 644–650.

Savage, S. B. & Hutter, K. 1989 The motion of a finite mass of granular material down a rough incline. *J. Fluid Mech.* **199**, 177–215.

Serayssol, J.-M. & Davis, R. H. 1986 The influence of surface interactions on the elastohydrodynamic collision of two spheres. *J. Colloid Interface Sci.* **114**, 54–66.

Shewchuk, J. R. 1996 Triangle: Engineering a 2D quality mesh generator and delaunay triangulator. In *Applied Computational Geometry: Towards Geometric Engineering* (ed. Ming C. Lin & Dinesh Manocha), *Lecture Notes in Computer Science*, vol. 1148,

pp. 203–222. Springer-Verlag, from the First ACM Workshop on Applied Computational Geometry.

Siavoshi, S. & Kudrolli, A. 2005 Failure of a granular step. *Phys. Rev. E* **71**, 051302.

Smart, J. R., Beimfohr, S. & Leighton, Jr., D. T. 1993 Measurement of the translational and rotational velocities of a noncolloidal sphere rolling down a smooth inclined plane at low Reynolds number. *Phys. Fluids* **5**, 13–24.

Staron, L. & Hinch, E. J. 2005 Study of the collapse of granular columns using two-dimensional discrete-grain simulation. *J. Fluid Mech.* **545**, 1–27.

Tang, J., Sagdiphour, S. & Behringer, R. P. 2009 Jamming and flow in 2D hoppers. In *American Institute of Physics Conference Series* (ed. M. Nakagawa & S. Luding), *American Institute of Physics Conference Series*, vol. 1145, pp. 515–518.

Tanner, R. I. 1985 *Engineering Rheology*. Clarendon.

Tezduyar, T. E. 1991 Stabilized finite element formulations for incompressible flow computations. *Adv. Appl. Mech.* **28**, 1–44.

Thompson, P. A. & Grest, G. S. 1991 Granular flow: friction and the dilatancy transition. *Phys. Rev. Lett.* **67**, 1751–1754.

Tsai, J.-C. & Gollub, J. P. 2005 Granular packings sheared in an annular channel: Flow localization and grain size dependence. *Phys. Rev. E* **72**, 051304.

Volfson, D., Tsimring, L. S. & I, S. Aranson 2003 Partially fluidized shear granular flow: Continuum theory and molecular dynamics simulations. *Phys. Rev. E* **68**, 021301.

Walton, O. R. & Braun, R. L. 1986 Viscosity, granular-temperature, and stress calculations for shearing assemblies of inelastic, frictional disks. *J. Rheol.* **30**, 949–980.

Yih, C.-S. 1967 Instability due to viscosity stratification. *J. Fluid Mech.* **27**, 337–352.