

UNIVERSITY OF SOUTHAMPTON

Numerical issues in the simulation of
orographic gravity waves

Helen Jane Hewson

Thesis submitted for the degree of Master of Philosophy

Faculty of Engineering and Applied Science

Department of Aeronautics and Astronautics

September 2000

© Crown copyright 2000

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND APPLIED SCIENCE
AERONAUTICS AND ASTRONAUTICS

Master of Philosophy

Numerical issues in the simulation of orographic gravity waves

Helen Jane Hewson

Orographic gravity waves have been the subject of active research for several decades. Their effects have been parametrized in global numerical weather prediction models since the 1980s, but our understanding of trapped lee waves remains incomplete. Other phenomena such as rotors are linked to trapped lee waves, and these represent a significant aviation hazard, as well as an aspect of the flow which is not well understood or well represented in numerical models.

Most previous investigations into flow over topography have included the boundary layer over small scale hills, but neglected it for larger scale mountain-induced gravity wave flows. These two situations have traditionally been treated separately, but the boundary layer cannot be ignored in gravity wave cases, because it can significantly change the flow over complex terrain. The simulations presented in this thesis use a high resolution, non-hydrostatic numerical model, which was originally designed to simulate boundary layer flows over small scale hills. The work presented here consists of two-dimensional simulations, with no boundary layer, in order to test the model's ability to predict adequately both upwardly and horizontally propagating waves. The model is used to simulate several two-dimensional idealised cases of orographic gravity waves, without a boundary layer, using a larger mountain and a larger numerical domain than in previous work, and these have shown good agreement with published analytical results. These simulations are carried out with a view to providing a sound base for turbulent boundary layer simulations, including turbulent flow phenomena observed around mountains, such as rotors, in order to better understand the flow patterns involved and the conditions conducive to their formation. This will lead to improvements in parametrizations and NWP performance. In the long term, the work may lead to a more sophisticated local lee wave/rotor forecasting tool for aviation use, to reduce the hazard posed by these phenomena, which are currently very difficult to predict. The work described in this thesis lays further ground work for future studies of this nature.

Various tests are carried out in order to obtain a robust numerical configuration for the model. Grid independence is investigated in both the horizontal and the vertical, using both uniform and stretched grids, and the boundary condition at the top of the model domain is tested. An option was introduced shortly before the start of this work to use inflow/outflow boundary conditions upstream and downstream of the mountain, and the formulation of these boundary conditions, particularly the radiative outflow boundary condition at the downstream edge, is tested. Other sensitivity tests were carried out into the model timestep, the surface boundary condition, the mountain shape, and the initial transience caused by growing the mountain into the model domain at the start of a simulation.

Contents

Title page	i
Abstract	ii
Contents	iii
List of tables and figures	vii
Acknowledgements	xiii
Declaration	xv
List of acronyms, symbols and operators	xvi
1 Introduction	1
1.1 General background	1
1.2 Structure of the atmosphere	3
1.3 Content of this thesis	5
2 The effect of mountains on the atmosphere	6
2.1 Introduction	6

2.2	Orographic gravity waves	8
2.3	Turbulent flow phenomena around mountains	11
2.4	Gravity wave drag parametrizations	18
2.5	Motivation for this work	21
3	Equations of motion	23
3.1	Basic equations	23
3.1.1	Introduction and notation	23
3.1.2	Conservation of mass	24
3.1.3	Conservation of momentum	25
3.1.4	Conservation of energy	27
3.2	Flow parameters	28
3.2.1	Introduction	28
3.2.2	The Reynolds number	29
3.2.3	The Rossby number	30
3.2.4	The Froude number	30
3.3	Turbulent flow	32
3.4	The Boussinesq approximation	36
3.5	Linear mountain waves for a simple idealised case	39
3.6	Momentum fluxes and drags	47
4	The numerical model	51
4.1	Introduction	51
4.2	Model equations and initialisation	52
4.3	Upper and lower boundary conditions	61
4.4	Horizontal boundary conditions	63
4.5	Momentum fluxes and drags	66
4.6	Artificially imposed horizontal diffusion	70

5	Idealised test cases	73
5.1	Introduction	73
5.2	Description of the results	85
5.3	Comparison with published analytical results	91
5.4	Introduction to the numerical issues encountered	97
6	Grid independence	98
6.1	Introduction	98
6.2	Horizontal grid dependence	99
6.2.1	Uniform grid	99
6.2.2	Stretched grid	104
6.3	Vertical grid dependence	116
6.3.1	Vertical resolution	116
6.3.2	Rayleigh damping scheme	121
6.4	Summary	124
7	Inflow/outflow boundary conditions	126
7.1	Introduction	126
7.2	Tests of the radiative outflow boundary condition	127
7.2.1	Periodic boundary conditions	127
7.2.2	Wider horizontal domain	132
7.3	Alternative schemes for the radiative outflow boundary condition	138
7.3.1	Introduction	138
7.3.2	Modifications to the restriction on the magnitude of the outflow phase speed	138
7.3.3	Vertically varying outflow phase speed, with different mass flux corrections	141
7.3.4	Timestep adjusted to avoid clipping the outflow phase speed	143
7.4	Discussion and options for how to proceed	143

8	Other sensitivity testing	148
8.1	Halved timestep	148
8.2	Quasi-free slip surface condition	151
8.3	Tests of different mountain shapes	153
8.4	Initial noise and mountain growth	157
9	Discussion and suggestions for future work	164
9.1	Summary and recommendations	164
9.2	Suggestions for further work	166
	References	171

List of tables and figures

Tables

Table 5.1. The four cases and their physical parameters.	74
--	----

Figures

Figure 2.1. Schematic showing lee wave flow downstream of a mountain ridge, with position of rotor circulation and cloud formations, and different levels of turbulence. After Knable, 1995.	15
Figure 2.2. Postulated mean streamlines in two-dimensional steady flow containing (a) wave breaking; (b) rotor circulation.	16
Figure 3.1. Upwardly propagating waves for the case where $l^2 > k^2$. After Durran (1986).	43
Figure 3.2. Horizontally propagating waves for the case where $l^2 < k^2$. After Durran (1986).	44
Figure 3.3. Potential temperature and wind profiles for the two layer case, showing the change in Scorer parameter with height. After Durran (1986).	45
Figure 3.4. Wave pattern for the two layer case, showing the trapped lee wave at lower levels. After Durran (1986).	46

Figure 4.1(a). The upper diagram shows schematically the vertical grid used in BLASIUS, illustrating how the model fields are held on different levels. The lower diagram illustrates the three-dimensional mesh arrangement. From Wood et al. (1999).	58
Figure 4.1(b). Schematic diagram of the BLASIUS grid at the upstream and downstream edges of the numerical domain, showing how the model fields are held on different points. From Wood et al. (1999).	59
Figure 4.2. Lowest model level surfaces on Z levels (full model levels, shown by a solid line) and ZN levels (model half levels, shown by a dashed line), for a 500 m high mountain.	60
Figure 5.1. Input profiles of wind speed and potential temperature for the four cases.	75
Figure 5.2. Close-up view of input wind speed and potential temperature profiles for the troposphere-stratosphere cases, which were shown in figure 5.1.	76
Figure 5.3. Input profiles of Scorer parameter (compared with the dominant horizontal wavenumber) for the four cases.	78
Figure 5.4. Mountain shapes for the different cases.	79
Figure 5.5. Horizontal gridlengths and expansion ratios for the grid described in the text	80
Figure 5.6. The grid described in the text, shown for a partial domain centred on the mountain.	81
Figure 5.7. Time series of the drag coefficient for all the cases.	82
Figure 5.8. Time series of the drag coefficient for all the cases, with non-dimensional time on the horizontal axis, as defined in the text.	83
Figure 5.9. Vertical velocity (contour interval 0.2 m s^{-1}) and potential temperature (contour interval 2 K) for case 1 after 100 NDTU.	85
Figure 5.10. Vertical velocity (contour interval 0.2 m s^{-1}) and for case 2 after -45 and -25 NDTU (left and right hand plots respectively).	86
Figure 5.11. Vertical velocity (contour interval 0.1 m s^{-1}) and for case 2 after 0 and 35 NDTU (left and right hand plots respectively).	87

Figure 5.12. Vertical velocity (contour interval 0.1 m s^{-1}) and potential temperature (contour interval 2 K) for case 2 after 135 NDTU.	88
Figure 5.13. Vertical velocity (contour interval 0.2 m s^{-1}) and for case 3 after 0 and 35 NDTU (left and right hand plots respectively).	89
Figure 5.14. Vertical velocity (contour interval 0.2 m s^{-1}) and potential temperature (contour interval 4 K) for case 3 after 115 NDTU.	90
Figure 5.15. Vertical velocity (contour interval 0.05 m s^{-1}) and potential temperature (contour interval 4 K) for case 4 after 17.5 NDTU.	91
Figure 5.16. (a) Durran (1986) figure 20.2b for $a^{-1}=1$ and 20.2c for $a^{-1}\ll 1$, as described in the text. (b) Potential temperature for the BLASIUS case 1 simulation after 100 NDTU, shown for part of the model domain, with contour interval 3 K.	92
Figure 5.17. Vertical velocity for case 2, the troposphere-only case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 3, which is shown on the left.	94
Figure 5.18. Vertical velocity for case 3, the troposphere-stratosphere case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 5, which is shown on the left.	95
Figure 5.19. Potential temperature for the BLASIUS case 4 simulation after 17.5 NDTU, shown for part of the model domain, with contour interval 1.5 K.	96
Figure 6.1. Vertical velocity after 140 NDTU for a 100 m high Witch of Agnesi mountain of half width 10 km.	99
Figure 6.2. Vertical velocity and potential temperature after 140 NDTU for a 500 m high Witch of Agnesi mountain of half width 10 km.	100
Figure 6.3. Vertical velocity and potential temperature after 250 NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km.	101
Figure 6.4. Horizontal gridpoints at the surface over 500 m high Witch of Agnesi mountains of half width 10 km, 5 km and 2.5 km, using a uniform horizontal grid spacing of 1 km.	102

Figure 6.5. Vertical velocity and potential temperature after 250 (in the left hand plot) and 230 (in the right hand plot) NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km.	103
Figure 6.6. The 'one stretch grid' described in the text, shown here for a partial domain centred on the mountain.	106
Figure 6.7. Vertical velocity plots for case 1 after 140 NDTU, using the uniform (500 m grid spacing) and one stretch grids.	107
Figure 6.8. Horizontal gridlengths and expansion ratios for the one stretch grid (dashed line) compared with those for the additionally smoothed grid (solid line).	108
Figure 6.9. Vertical velocity plot (contour interval 0.2 m s^{-1}) for case 3 after 140 NDTU, using the one stretch grid with halved resolution. The wave pattern was damped as it left the central region; the horizontal gridlengths are shown below for reference.	110
Figure 6.10. Horizontal gridlengths and expansion ratios for the one stretch grid (dashed line) compared with those for the two stretch grid (solid line).	111
Figure 6.11. Vertical velocity plots for case 3 after 140 NDTU, using the one stretch and two stretch grids.	112
Figure 6.12. Vertical velocity plots for case 4 after 10 NDTU, using the one stretch and two stretch grids.	113
Figure 6.13. Time series of the drag coefficient for the results presented in figure 6.11, with non-dimensional time on the horizontal axis, as defined in Chapter 5.	114
Figure 6.14. Horizontal gridlengths and expansion ratios for the one stretch grid (dashed line) and the two stretch grid (dotted line), compared with those for the more smoothly stretched grid (solid line).	115
Figure 6.15. The different grids for the two simulations presented in figure 6.16. The top plot shows the uniform vertical grid, and the lower plot shows the stretched vertical grid.	117
Figure 6.16. Vertical velocity after 75 NDTU for a 100 m high Witch of Agnesi mountain of half width 10 km. The left hand plot has a uniform vertical grid spacing of 350 m; the right hand one has a stretched vertical grid, as described in the text.	118

- Figure 6.17. Vertical velocity after 140 NDTU for a 500 m high cosine-squared mountain of width 15700 m. The left hand plot has the stretched vertical grid described in the text, and the right hand one has the same grid but with approximately double the gridlength. 120
- Figure 6.18. Schematic representation of a wavelength in a discretised numerical model, using five gridpoints (four gridlengths). 120
- Figure 6.19. Vertical velocity after 250 NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km. The Rayleigh damping layer is between 20 km and 35 km in the left hand plot, and between 25 km and 40 km in the right hand one. 122
- Figure 6.20. Vertical velocity after 250 NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km. The Rayleigh damping layer is between 20 km and 35 km in all cases; the damping coefficient is 0.001 s^{-1} in the left hand plot, 0.01 s^{-1} in the centre plot and 0.1 s^{-1} in the right hand plot. 123
- Figure 7.1. Vertical velocity after 80 NDTU (top plots) and 140 NDTU (bottom plots) for case 2. The left hand plots used inflow/outflow boundary conditions, and the right hand plots used periodic boundary conditions. 128
- Figure 7.2. Vertical velocity after 0 NDTU (at the end of mountain growth; top plots) and 100 NDTU (bottom plots) for case 3. The left hand plots used inflow/outflow boundary conditions, and the right hand plots used periodic boundary conditions. 130
- Figure 7.3. Time series of drag coefficient for case 1, using inflow/outflow and periodic boundary conditions, and the original and wider horizontal domain. 133
- Figure 7.4. Vertical velocity after 80 NDTU for case 3. Both these simulations used inflow/outflow boundary conditions; the left hand plot used the original horizontal domain, and the right hand one used the wider one (as described in the text). 135
- Figure 7.5. Time series of drag coefficient for case 3, using inflow/outflow and periodic boundary conditions, and the original and wider horizontal domain. 135
- Figure 7.6. Vertical velocity plots after 7.5 NDTU (top plots) and 10 NDTU (bottom plots) for case 4. The left hand plots used the two stretch horizontal grid and the right hand plots used the wider domain described in the text. 137
- Figure 7.7. Vertical velocity plots (contour interval 0.05 m s^{-1}) for case 4 after 7.5 NDTU, using the original and modified outflow boundary conditions. 139

Figure 7.8. Vertical velocity plots for case 3 after 120 NDTU, using the original and modified outflow boundary conditions.	140
Figure 8.1. Vertical velocity plots for case 3 after 0, 40 and 140 NDTU (top, middle and bottom rows respectively), using the full and halved timestep (left hand and right hand columns respectively).	150
Figure 8.2. Vertical velocity plots for case 4 after 7.5 NDTU, using the no slip and quasi-free slip surface boundary conditions.	153
Figure 8.3. Comparison of Witch of Agnesi and cosine-squared mountains of the same height (500 m in this example) and cross-sectional area.	154
Figure 8.4. Vertical velocity after 140 NDTU for two 100 m high mountains of different shapes but equal cross-sectional areas. The left hand plot is a Witch of Agnesi mountain of half width 10 km; the right hand one is a cosine-squared mountain of width 62.8 km.	155
Figure 8.5. Vertical velocity for case 2, the troposphere-only case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 3, which is shown on the left.	156
Figure 8.6. Vertical velocity for case 3, the troposphere-stratosphere case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 5, which is shown on the left.	157
Figure 8.7. Vertical velocity at the end of mountain growth (0 NDTU). The left hand plot is for the constant U, constant N case, with a 100 m high Witch of Agnesi mountain of half width 10 km; here the mountain grew over 12500 seconds; the Rayleigh damping layer started at 20 km. The right hand plot is for the troposphere-stratosphere case, with a 500 m high cosine-squared mountain of width 15700 m; here the mountain grew over 5000 seconds; the Rayleigh damping layer started at 60 km, so the vertical domain plotted is three times that shown in the left hand plot. There is far more initial noise in the latter case.	158
Figure 8.8. The relationship between the time taken to grow the mountain, the maximum mountain slope and the maximum initial noise level throughout the model domain.	160
Figure 8.9. The relationship between the time taken to grow the mountain, expressed as a percentage of the time it would be expected to take using the equation given in section 8.1, and the maximum mountain slope.	163

Acknowledgements

I'd really like to thank absolutely everyone I can think of, but somehow I doubt that it would look very good if the acknowledgements were longer than the thesis itself, so I'll settle for just picking out the biggest and brightest stars.

Firstly, a big thank you to both my supervisors for their time, effort, support, encouragement and inspiration over the last few years: Adrian Broad at The Met. Office for his constructive criticism, helpful suggestions and welcome jokes; Ian Castro at the University of Southampton (and formerly at the University of Surrey) for helping me to see things from different angles, for his diligent proof-reading, and for always, always believing in me. I am, of course, very grateful to The Met. Office for making this possible financially, and for all the other training that they gave me before I embarked on this. In addition, I'd like to thank Nigel Wood, who has been an excellent mentor and tutor to me, and many other colleagues currently or formerly working in or with the Atmospheric Processes Research branch for their advice and for many useful discussions: they are too numerous to name here, but they know who they are. To Derrick Ryall and Alistair Manning, thank you for giving me a target that I really wanted to meet, rather than leaving me with a deadline that I couldn't afford to miss.

This list would be incomplete without a big thank you to my family and friends for their moral support, and for putting up with me during the more difficult times. I would particularly like to thank my father for asking the most thought-provoking questions ever, and for stretching me beyond what I naively thought were my limits, as he has been doing throughout my entire life. Heartfelt thanks to my mother and my sister for always being there when I needed to talk through how I was feeling about my work, and for their unfailing faith in me. Words cannot express my gratitude to my wonderful fiancé, Andy, who frequently knows me better than I know myself, and who has provided a fresh pair of eyes to proof-read every word you see before you: thank you, my love. To all the other friends who have repeatedly told me that I can do it, and listened to endless complaints when things weren't going well, a big thank you. I can't write this without thanking Zoe for not throwing me out of the car half way down the motorway after some of the more stressful days, AJ for always being prepared to spend his lunch break just listening, and Susan for understanding perfectly how utterly horrible and absolutely wonderful the whole thing has been.

Unconventionally, perhaps, I'd also like to thank myself. I've had two Helens inside my head for the past six months: one saying that I've had enough and it's too stressful and I can't possibly finish it and it's not worth it, and the other standing silently and serenely in front of me every single day just holding up a big list entitled "Reasons Why You Are *Choosing* To Do This". I've never once felt, even on the hardest days, that I was doing this because I had to; I've done it because I wanted to do it and because I've enjoyed it. Thank you to everyone who gave their support to the me with the list, and quietly kicked the other me in the teeth while no one else was looking.

Helen Hewson, Bracknell, September 2000.

The influence of fine scenery, the presence of mountains, appeases our irritations and elevates our friendships.

Ralph Waldo Emerson (1803-1882), *Culture, The Conduct of Life* (1860)

There's only us
There's only this
Forget regret
Or life is yours to miss
No other road
No other way
No day but today

Jonathan Larson (1960-1996), *Rent* (1996)

List of acronyms, symbols and operators

Acronyms

ABL	Atmospheric boundary layer
BLASIUS	Boundary Layer Above Stationary, Inhomogeneous Uneven Surfaces
CFL	Courant-Friedrichs-Lewy (refers to the CFL number or CFL criterion)
CSO	Centred second order
GCM	General circulation model
GWD	Gravity wave drag
LES	Large eddy simulation
NDTU	Non-dimensional time units
NWP	Numerical weather prediction
SGSO	Sub-grid scale orography
TVD	Total variance diminishing
UQ	Ultimate quickest

Symbols

a	Half width of a Witch of Agnesi mountain
a_n	Height of the vertical model levels
b''	Buoyancy

C_{FLOW}	} Lower and upper bounds on the CFL number in the model
C_{FLUPP}	
C_D	Drag coefficient
c	Phase speed
\bar{c}	Vertically constant outflow phase speed in the model
\underline{c}_g	Group velocity
c_p	Specific heat capacity for air at constant pressure
c_{prop}	Propagation speed of the wave head during model spin-up
F	Wave drag
F_D	Drag force over the full domain
F_h	Froude number (used to refer to the vertical Froude number)
F_i	Body forces
F_x	Surface pressure force per unit length in the model
F_α	Horizontal Froude number
f	Coriolis parameter
g	Gravitational acceleration (also appears as $\underline{g} = -g\underline{k}$)
H	Vertical length scale
H_D	Depth of the model domain
H_j	Turbulent heat flux
H_ρ	Density scale height
h	Mountain height
h_m	Maximum mountain height
I	Represents a model gridpoint in the x -direction
IIP	Number of model gridpoints in the x -direction
I_{max}	Maximum initial noise over the whole model domain
\underline{i}	Unit vector in the x -direction
J	Represents a model gridpoint in the y -direction
JJP	Number of model gridpoints in the y -direction
\underline{j}	Unit vector in the y -direction
K	Represents a model gridpoint in the z -direction

KKP	Number of model gridpoints in the z -direction
K_h	Eddy conductivity (also called heat diffusivity)
K_m	Eddy viscosity (also called momentum diffusivity)
K_ϕ	Represents either the eddy conductivity or the eddy viscosity
\underline{k}	Unit vector in the z -direction
k	Dominant horizontal wavenumber
L	Horizontal length scale; width of a cosine-squared mountain
L_x	Domain width
l	Scorer parameter
M	Horizontally averaged vertical flux of horizontal momentum
m	Vertical wavenumber
m_*	Related to the vertical wavenumber when m is imaginary: $m_* = im$
N	Brunt-Väisälä frequency (also called buoyancy frequency or static stability)
n_g	Number of timesteps used to grow the mountain into the model domain
P	Pressure (discretised in the model)
p	Pressure (in the continuous equations)
Q	Discretised moisture variable in the model
R	Universal gas constant for dry air; Rayleigh damping coefficient
RHO	Discretised density variable in the model
Re	Reynolds number
Re_g	Grid Reynolds number
$Re_{g_{crit}}$	Critical value of the grid Reynolds number for the artificial diffusion scheme
Ro	Rossby number
R_1	Expansion ratio for the vertical model grid
r_ϕ	Non-dimensionalised outflow phase speed in the model
S	Entropy
S_{max}	Maximum mountain slope
T	Temperature, normalised model temperature, or period
$TAUI3$	Discretised model variable representing τ_{13}
t	Time
t_e	Expected time taken to grow the mountain into the model domain

t_g	Actual time taken to grow the mountain into the model domain
t_{nd}	Non-dimensional time
t_T	Total simulation time
U	Horizontal velocity scale; velocity in the x -direction (discretised in the model)
$u \equiv u_1$	Velocity in the x -direction (in the continuous equations); $\underline{u} = (u, v, w)$
V	Velocity in the y -direction (discretised in the model)
$v \equiv u_2$	Velocity in the y -direction (in the continuous equations)
W	Vertical velocity scale; velocity in the z -direction (discretised in the model)
$w \equiv u_3$	Velocity in the z -direction (in the continuous equations)
X, XN	Represent the staggered model grid levels in the x -direction
X_d	Downstream fetch
$x \equiv x_1$	Horizontal Cartesian co-ordinate in the standard rotating frame of reference
x_-, x_+	Ends of the horizontal model domain
Y, YN	Represent the staggered model grid levels in the y -direction
$y \equiv x_2$	Horizontal Cartesian co-ordinate in the standard rotating frame of reference
Z, ZN	Represent the staggered model grid levels in the z -direction
Z_s	Model surface height
Z_1	Height of the lowest internal vertical model level
$z \equiv x_3$	Vertical Cartesian co-ordinate in the standard rotating frame of reference; transformed co-ordinate in the equation for the model grid
\tilde{z}	Cartesian co-ordinate in the equation for the transformed model grid
z_b	Height of the base of the Rayleigh damping layer
z_t	Height of the top of the Rayleigh damping layer
α	CFL number
Δt	Model timestep
ΔX	Model gridlength in the x -direction
ΔY	Model gridlength in the y -direction
ΔZ	Model gridlength in the z -direction

δ_{ij}	Kronecker delta (equal to unity if $i = j$ and zero otherwise)
θ	Potential temperature
θ_{init}	Initial potential temperature profile in the numerical model
κ	Defined as $\frac{R}{c_p}$, where here R is the universal gas constant for dry air
λ_h	Horizontal wavelength
μ	Dynamic viscosity
ν	Kinematic molecular viscosity
ρ	Density
σ_{ij}	Stress tensor
τ_{ij}	Reynolds stress tensor
ϕ	Normally used to represent any field in a general equation or expression; also used to represent latitude in the discussion of the Coriolis force
$\underline{\Omega}$	Angular velocity of the earth, with magnitude $\Omega = \underline{\Omega} $
ω	Angular frequency
$\hat{\omega}$	Intrinsic frequency

Operators

$\hat{\phi}$	Non-dimensional variable (except $\hat{\omega}$, the intrinsic frequency)
$\bar{\phi}$	Mean state (may be a time, space or ensemble average)
ϕ_0	Background profile, varies with height only
ϕ_r	Reference state
ϕ'	Turbulent fluctuation from the mean state
ϕ''	Perturbation from the mean state or background profile (non-turbulent)
$\frac{\partial \phi}{\partial t}$	Eulerian derivative (the local rate of change, at a fixed point)

$$\nabla\phi \quad \text{Gradient: } \nabla\phi \equiv \left(\frac{\partial\phi}{\partial x}, \frac{\partial\phi}{\partial y}, \frac{\partial\phi}{\partial z} \right)$$

$$\nabla \cdot \underline{\phi} \quad \text{Divergence: } \nabla \cdot \underline{\phi} \equiv \frac{\partial\phi_1}{\partial x} + \frac{\partial\phi_2}{\partial y} + \frac{\partial\phi_3}{\partial z}$$

$$\nabla^2\phi \quad \text{Laplacian (also referred to as del-squared): } \nabla^2\phi \equiv \frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial y^2} + \frac{\partial^2\phi}{\partial z^2}$$

$$\underline{u} \cdot \nabla\phi \quad \text{Advection term: } \underline{u} \cdot \nabla\phi \equiv u \frac{\partial\phi}{\partial x} + v \frac{\partial\phi}{\partial y} + w \frac{\partial\phi}{\partial z}$$

$$\frac{D\phi}{Dt} \quad \text{Lagrangian, material, total or substantive derivative (the rate of change following the motion of an air parcel): } \frac{D\phi}{Dt} \equiv \frac{\partial\phi}{\partial t} + \underline{u} \cdot \nabla\phi$$

$$\overline{\frac{D}{Dt}}(\bar{\phi}) \quad \text{Lagrangian derivative with advection by the mean flow:}$$

$$\overline{\frac{D}{Dt}}(\bar{\phi}) \equiv \frac{\partial\bar{\phi}}{\partial t} + \bar{\underline{u}} \cdot \nabla\bar{\phi}$$

1 Introduction

1.1 General background

Meteorology is concerned with understanding and predicting weather patterns and their temporal and spatial evolution. Observations are made by various means in order to learn about the state of the atmosphere, but due to the high costs involved, they cannot provide detailed high resolution information on all the atmospheric variables over a wide area, so other methods are employed, such as analytical theory, numerical simulations, and laboratory experiments. Empirical rules derived from observations and experiments, and the physical laws of fluid motion, are used to obtain equations governing atmospheric flows, which can be used to investigate atmospheric motion using analytical methods and numerical models. These have a reduced observational requirement, as they only need observations to set up the initial conditions and boundary conditions. However, numerical models can be prohibitively expensive in terms of computing resources. A two-dimensional model, with no spanwise (cross-flow) component, reduces these costs, but it has limitations and cannot deal with complex three-dimensional terrain and the resulting flow patterns. Shutts *et al.* (1994) noted that direct simulation of gravity waves over complex terrain is now becoming feasible due to advances in supercomputing technology. Such simulations can be used in place of observations, since they provide far more comprehensive datasets than can be obtained using observations, but clearly they need to be validated against available observations in order to give confidence in their results. However, the true atmosphere is far more complicated than can be modelled, so various assumptions and

approximations are applied. These are dependent on the type of flow being investigated, and those used in this work are discussed in Chapter 3.

Motions in the atmosphere occur on a wide range of scales, from planetary waves down to molecular processes. Mason (1970) gives a useful overview of the different scales involved, and of where various weather phenomena fit in. Numerical Weather Prediction (NWP) models aim to forecast large scale weather patterns such as mid-latitude depressions and anticyclones, which are typically of order 500-5000 km. These are called synoptic scale processes. Mesoscale processes are typically of order 20-500 km, and include lee waves (which are the subject of this thesis), fronts and squall lines. Other atmospheric processes, such as showers, tornadoes and boundary layer eddies, occur on even smaller scales. Stull (1988) gives an overview of typical time and space scales both for mesoscale processes and for these smaller microscale ones. In general, global NWP models only resolve synoptic scales, because of the computational cost involved in solving the governing equations on a fine enough grid. Limited area models are often used to allow the resolution of mesoscale features for a particular region: for example, The Met. Office uses a mesoscale forecast model covering the United Kingdom and part of northern Europe. However, global models still need to represent sub-grid scale features (that is, features which are not resolved on the model grid), because they can have an important effect on the evolution of the synoptic scale weather pattern. This is achieved by using parametrization schemes, which attempt to include the effects of processes such as gravity wave drag (GWD), convection, boundary layer and surface processes (such as turbulent mixing and friction), radiation, clouds, precipitation, sub-surface heat fluxes in soil, and vegetative canopy effects. The background to the GWD scheme, which represents the drag caused by mountains, is discussed in the next chapter.

1.2 Structure of the atmosphere

The Earth's atmosphere is divided into several layers with different characteristics. Approximately 80% of the mass of the Earth's atmosphere is contained within the troposphere, with the troposphere and stratosphere between them accounting for 99.9% of the atmosphere. These two layers are described here.

The mean temperature in the troposphere decreases with height, primarily because of solar radiation, which heats the Earth's surface. This heat is then released into the lowest levels of the atmosphere. The unstable temperature profile leads to vertical mixing by convection and turbulence, transferring heat up through the troposphere. These vertical motions, and the presence of water vapour in the air, mean that most weather phenomena happen in the troposphere. Because of this, and because it is the layer closest to the Earth's surface, it is the most easily observed and best known part of the atmosphere.

The troposphere includes both the free atmosphere and the atmospheric boundary layer (ABL), a relatively shallow layer of the atmosphere whose depth varies from around 100 m to around 2 km, depending on the stability of the atmosphere and on the nature of the surface below. The ABL has no precise definition, but Garratt (1992) defines it as "the layer of air directly above the Earth's surface in which the effects of the surface (friction, heating and cooling) are felt directly on time scales of less than a day, and in which significant fluxes of momentum, heat or matter are carried by turbulent motions on a scale of the order of the depth of the boundary layer or less." The ABL is a very important layer, since most human activity takes place within it, and many animals and plants exist within it. Although shallow, it has a significant effect on larger scales, since the surface exerts a frictional drag on these motions, providing a momentum sink, and since mixing within the ABL spreads moisture and heat up from the surface to drive weather systems in the free atmosphere.

The boundary between the troposphere and the stratosphere is called the tropopause. In radiosonde ascents the tropopause can usually be identified by a temperature inversion, which is sometimes accompanied by strong winds known as a tropospheric jet. The tropopause is also marked by a sharp drop in the water vapour content, and by an increase in the ozone content. These strong concentration gradients indicate that there is very little mixing between the troposphere and the stratosphere. The position and temperature of the tropopause are a result of the balance between the convective and turbulent heating of the troposphere, and the radiative heating of the stratosphere, and they vary with latitude and synoptic conditions. The high surface temperature in the tropics means that convective mixing takes place to great heights, and adiabatic cooling of air reaching these levels makes it very cold, despite the original high surface temperature. Because of this, the tropical tropopause is generally cold and high (the equatorial tropopause is typically 16 km high and has a temperature of -80°C), whereas at higher latitudes the tropopause is warmer and lower (typically it lies between 10 and 12 km over the United Kingdom and Scotland during the winter months, whereas the polar tropopause has a typical height of 8 km, and a temperature of -40°C in summer or -60°C in winter).

The temperature at the tropopause stops falling significantly, and it generally starts to increase with height into the stratosphere, where vertical mixing is very limited. The resulting strong static stability acts as a lid on the tropopause, strongly inhibiting exchange of air between the two layers. The stratosphere receives its heat by direct absorption of solar radiation. The temperature structure of the stratosphere is dominated by absorption of ultraviolet radiation by ozone, with a temperature maximum at around 50 km above the surface, which defines the stratopause, marking the top of the stratosphere. This temperature maximum shows wide seasonal and latitudinal variation: it is approximately -30°C over the winter pole and $+20^{\circ}\text{C}$ over the summer pole.

The horizontal wind generally increases with height throughout the troposphere, reaching a maximum just below the tropopause, which is called the tropospheric jet. This reinforces the 'lid' effect on the troposphere, limiting exchange of air with the stratosphere above. The wind drops in the lower stratosphere, then rises again to a maximum just above the stratopause. The wind direction is very variable with latitude

and season, but the prevailing low level wind over the United Kingdom is westerly, and so the horizontally propagating wave trains described in the next chapter are most often observed to the east of mountain ranges such as the Pennines.

1.3 Content of this thesis

Chapter 2 gives a general introduction to the effects of mountains on the atmosphere, and presents the basic theory behind orographic waves. Other phenomena which are linked to lee waves, and their importance as an aviation hazard, are discussed. Chapter 2 also explains the representation of orographic waves in NWP models, and the motivation for this project. Chapter 3 covers the basic equations governing atmospheric flows, and some of the assumptions and approximations which are used in analytical and numerical modelling of these flows, as well as defining various parameters which are useful for the study of mountain waves. Chapter 4 describes the numerical model used for the simulations in this thesis. The results obtained using this model for several idealised cases are described in Chapter 5, with comparisons against published analytical results where possible. These are all two-dimensional simulations, with no boundary layer, in order to test the model's ability to predict adequately both upwardly and horizontally propagating waves. Although these simulations are designed to investigate inviscid flow, the numerical model requires viscosity, so a very low value is used, and this is explained further in Chapter 4. The next three chapters explain the various tests carried out in order to obtain a robust numerical configuration for the model. Grid independence issues are covered in Chapter 6, and the upstream and downstream boundary conditions are tested in Chapter 7. The remaining sensitivity tests are explained in Chapter 8: investigations were carried out into halving the model timestep, using a quasi-free slip surface boundary condition, changing the mountain shape, and reducing the initial noise caused by growing the mountain into the model domain at the start of a simulation. The results presented in these three chapters are drawn together and discussed in Chapter 9, along with suggestions and recommendations for future work.

2 The effect of mountains on the atmosphere

2.1 Introduction

The weather in mountainous areas is often very complex, varying greatly in time and space, and it can be surprising even to experienced observers. As mentioned in the previous chapter, mesoscale processes such as these can have a significant effect on synoptic scale flow patterns. Smith (1979) suggests that “if the Earth were greatly reduced in size while maintaining its shape, it would be smoother than a billiard ball”, illustrating that although mountains may seem large to someone standing on the Earth’s surface, from a global perspective they are a relatively small scale feature. However, they can affect the large scale flow, with mountain barriers such as the Rocky Mountains of North America causing changes to pressure patterns which can be observed up to thousands of kilometres downstream of the mountain, indicating the scale of influence. In addition, the drag exerted by mountains on the atmosphere is an important source of deceleration in the large scale atmospheric flow, and this is discussed in more detail in Chapter 3.

Air can flow over mountains, generating gravity waves which propagate upwardly and/or horizontally, or around them, leading to features such as lee vortices, vortex shedding and lee cyclogenesis. The air can also be blocked by the mountain and remain stagnant or re-circulating upstream of the obstacle. Local wind regimes around mountains have various names such as föhn, bora, mistral or chinook. The föhn effect is commonly studied in school geography classes: the basic concept is that moisture is removed as rain on the windward side as the air rises, so the air descending the lee slope is drier and warmer, and the weather is sunnier. This effect is commonly observed to

the east of the Welsh mountains, the Pennines and the Scottish Highlands, and is particularly striking over the Alps, where the name originates. Other orographic effects include anabatic (upslope) and katabatic (downslope) winds, breaking waves and downslope windstorms. The latter, often seen in Boulder, Colorado, can cause damage to crops, trees and buildings, and even loss of life. Much research has been carried out into this phenomenon: well known examples are Peltier and Clark (1979) and Neiman *et al.* (1988).

Rising air in waves can lead to condensation and formation of cloud bands parallel to and downwind of high ground, visible on satellite images, and to smooth lens-shaped clouds, known as lenticular clouds, which resemble flying saucers. Queney *et al.* (1960) and Carney *et al.* (1996) are both good sources of orographic cloud images. Mountains can also remove cloud: satellite images often show clear patches in stratocumulus due to downward motion in lee waves. Even mother-of-pearl clouds, at typical heights of 20-30 km, are believed to be due to orography (Förchtgott, 1952, Corby, 1954 and Queney *et al.*, 1960), indicating how far up into the atmosphere the effects are felt. Mountain ranges also have a profound effect on the distribution of precipitation, with windward slopes being wet and areas to the lee being dry. Orographic rainfall enhancement occurs when air rises over the windward side of the mountain, and this upward motion encourages convection, leading to an increased likelihood of rain, or heavier rain than over flat terrain nearby. The seeder-feeder effect is also very important, whereby the orographic cloud itself does not produce rain, but it acts as a seeder cloud because it is so moist, enhancing precipitation falling through it from a separate cloud (the feeder cloud) at upper levels. Mountains can also be responsible for flash flooding and landslides, and even storms. However, all the work in this thesis uses dry dynamics, and does not attempt to represent these moist effects, so the detail of the moist equations of motion is omitted in this work. Durran and Klemp (1982) examined the effects of moisture on trapped waves, and showed that in some situations it could have a significant impact on them, compared to those formed in a dry air stream with the same upstream wind and temperature profiles. The effects are complicated because condensation is not always reversible, but the presence of moisture can damp, amplify, untrap, or entirely destroy the trapped wave response. However, most analytical work in this field neglects moisture, partly because its effects are

difficult to include, but also due to the qualitative success of dry dynamics in reproducing lee wave flow (Shutts and Broad, 1993, for example, found that moisture was not a crucial factor in determining the lee wave structure for the cases they studied). The idealised cases presented in this thesis are designed to reproduce analytical results obtained using dry dynamics.

The rest of this chapter explains the basic theory behind orographic gravity waves, and then describes some of the turbulent phenomena which can be observed around mountains and which can present an aviation hazard. The representation of gravity wave drag (GWD) in NWP models is discussed, and the motivation for the work presented in this thesis is explained.

2.2 Orographic gravity waves

Internal gravity waves can occur in any stably stratified fluid and can be generated by a variety of mechanisms. In the atmosphere the main sources are convection, shear and frontal instability, and flow over orography; the latter is the subject of this thesis, and will be explained here. Vertical displacements of air parcels may be caused by air near the ground, which rises as it moves over the windward side of the mountain and descends over the leeward side. The stable stratification means that buoyancy forces act to restore the parcel to its original position, but its momentum makes it overshoot, and the air parcels oscillate as they move downstream, creating the wave motion.

The formation of gravity waves as a result of air flow over orography has been the subject of much observation and research over the last few decades, so what follows is a selective overview. Many measurements have been taken of gravity wave events, such as Shutts and Broad (1993), Shutts *et al.* (1994), or Vosper and Mobbs (1996). In-depth reviews of the topic can be found in Smith (1979) or Queney *et al.* (1960); the latter is a very comprehensive survey of work done up to that point, covering both observational and theoretical studies. Durran (1986) provides a general background overview of a

wide range of related subjects, and notes that despite the years of research already carried out, forecasting mountain waves is still far from easy. Wood (2000) provides a historical perspective of the progress and developments in modelling flow over complex terrain.

Orographic gravity waves can propagate upwardly or horizontally, depending on the mountain and on the upstream airflow. The equations governing mountain waves are discussed in the next chapter, and a simple idealised example is given. Waves which do not penetrate above a certain height, and which propagate downstream, are referred to as trapped lee waves, because their energy is trapped at low levels in the lee of the mountain. In the atmosphere, trapped lee waves typically have horizontal wavelengths of between 6 km and 30 km, but the wave train can extend far downstream of the mountain (up to 500 km in favourable conditions), and the waves remain almost stationary. Although the most common lee waves are set off by long ridges, isolated peaks can set off 'wake waves' which propagate in a 'V' shape rather like the wake of a ship (Bradbury, 1997). Lee waves develop when there is a strong wind component perpendicular to the ridge, which usually requires the mean wind direction to be within 30° of perpendicular to the line of the mountain ridge unless the wind is very strong; for example, the presence of a tropospheric jet may cause the wave energy to be reflected, leading to trapped waves. The wind speed must be sufficiently strong and increase with height, and a stable air mass (meaning an inversion or very stable layer) must be present near the mountain top with a deeper, less stable layer above. These conditions imply a decrease with height of the Scorer parameter (defined in Chapter 3), and this is generally accepted to be the condition for formation of trapped lee waves. Lee waves are more common in autumn and winter because the air is more stable and the winds are stronger. If the wind direction changes with time, or with height, to be parallel to the ridge line, or if the mean wind is zero at some height, the waves are trapped below that level, known as a critical level (Bretherton, 1969). If this occurs at the level of the mountain crest, waves disappear altogether (Gerbier and Berenger, 1961). Although mountain waves are more commonly associated with the higher ranges such as the Rockies or the Alps, they can be induced by far lower ridges, in some cases as low as 100 m (Allan, 1997).

Internal gravity waves can be observed in many ways, including radiosondes, wind profilers, lidars, microbarographs, satellite imagery, wind masts and constant volume balloons. All of these methods have advantages and disadvantages in terms of cost, accuracy and suitability to a particular situation. Radar observations, for example, are restricted by the location of radar sites, which are not always close to mountainous terrain, but they provide continuous measurements in time, so they can be useful for measuring travelling features. Vosper (1995) gives a discussion of several methods, and references much of the work done, including well known and comprehensive experiments such as ALPEX and PYREX (in the Alps and Pyrénées respectively). A disadvantage of intensive field experiments such as these is that they are restricted to one region and one period of time, so the results are unlikely to be representative of the mean global situation. However, the results of observational studies are usually reasonably consistent with gravity wave theory. Aircraft observations from commercial airline flights, such as those used by Nastrom *et al.* (1987) and Jasperson *et al.* (1990), are less useful for studying turbulent features than data obtained from flights planned specifically for observational purposes, because the air crews tend to modify their flight path wherever possible in order to avoid turbulence, so most of the turbulent episodes encountered on these commercial flights were close to airports, since the pilots have less choice of flight path on the ascent and descent. They also avoid areas of intense wave activity, where possible, so again, aircraft observations are more useful when obtained from field experiments designed with that in mind. Queney *et al.* (1960) mention some less obvious sources of observational data, such as crop damage and bird flight paths. A detailed review of observational studies is not covered here since this work is purely in numerical modelling; it is sufficient here to say that they are plentiful, and that it would be very useful to verify the results of work following on from this project (such as rotor simulations) by carrying out observational field campaigns in the region of interest.

2.3 Turbulent flow phenomena around mountains

Glider pilots rely on vertical air currents, using not only thermal currents and the upslope lift on the windward side of mountains, but also the rising motion in atmospheric waves. Consequently they are able to provide practical evidence about mountain flows, observing and exploring waves and vertical wind patterns using instruments on board their aircraft, as well as giving visual descriptions of the phenomena they see. Flying conditions are remarkably smooth in waves, which are normally well above the ground, although occasionally they reach down to the surface and cause sudden strong gusts.

While the smooth upward motion in lee waves can be beneficial to glider pilots, other mountain flow phenomena present a serious hazard to aviation, being turbulent and therefore harder to forecast than lee waves themselves. Queney *et al.* (1960) review a wide range of mountain effects from an aviation perspective. Knowledge of turbulence and vertical velocities is important to the aviation community, but these features are generally not resolved in NWP models, and although analytical theory gives a good representation of laminar steady flow, it is less good for unsteady turbulent flows (Corby, 1954). There is an extensive literature on these aviation hazards, most of which is made up of field studies, such as Küttner's (1938, 1939) work in the Alps, Manley's (1945) detailed observational study of the Crossfell Helm wind, and Förchtgott's (1949, 1952) observations of mountain ranges in the former Czechoslovakia. Queney *et al.* (1960) summarises all of these, and others. Förchtgott (1949, 1952) produced regime diagrams indicating which flow pattern is expected for given upstream conditions, and included lee waves as well as a range of re-circulating features, but Corby (1954) doubts the evidence for these stationary closed vortices, and dismisses Förchtgott's regime diagrams as a summary of the author's ideas rather than a useful forecasting tool, concluding that more work is required to give confidence in the results. However, the basic tenor of Förchtgott's (1949, 1952) results does fit with earlier findings, such as

Küttner (1938, 1939), and with much of the work done since, as discussed below; it also confirms predictions obtained from perturbation theory. Corby (1954) gives a useful overview of early work on this topic, and Carney *et al.* (1996) is a comprehensive summary of the current state of knowledge from an aviation perspective. This section is merely an introduction to potential applications of the work which forms the main body of this thesis.

There are many forms of mountain-induced turbulence which together represent a serious aviation hazard: in 1992 the United States General Accounting Office found that mountainous western states had a general aviation accident rate which was approximately 40% higher than that of all other continental states (Bedard and Neilley, 1998). Carney *et al.* (1996) quote an even more alarming statistic: the accident rate for airports in mountainous areas, compared with similar airports in non-mountainous areas, was 155% higher in a similar survey. In addition, Lamb and Baker's (1998) analysis of National Transportation Safety Board data revealed that 91 out of 202 deaths in aircraft crashes in the Aspen, Colorado region between 1964 and 1987 were due to accidents caused by the effects of flying over mountainous terrain. Carney *et al.* (1996) noted that the accidents were not limited to one place, altitude, time of year or type of aircraft, but seemed simply to be a case of being in the wrong place at the wrong time. Pilots have a selection of guidelines developed from their collective experience of flying near mountains (referred to as "the aviation version of the mountaineer's code" by Allan, 1997) which can help them to avoid the hazardous air flow often present. For example, United Airlines has developed a reliable set of forecast rules for domestic mountain wave turbulence as a result of four decades' study, and weather forecasters in mountainous regions use a checklist on a daily basis to determine the likelihood of turbulence, and issue alerts as necessary (Knable, 1995). However, most of these guidelines rely heavily on visual clues for the crew such as cap, lenticular and roll clouds (see, for example, Carney *et al.*, 1996), which are of course only present if the air is moist, which in itself is not a pre-requisite for lee waves or for turbulence of any description. Rotors can sometimes be made visible by dust or smoke movements they induce if they are very close to the ground (Scorer and Klieforth, 1959, and Queney *et al.*, 1960) but even this is not entirely reliable as a detection mechanism, and local knowledge of typical flow patterns is invaluable.

Unfortunately, even with appropriate training and with many hours' experience of mountain flying, pilots can still be taken by surprise by some of these phenomena, due to the difficulties involved in forecasting them. Improved guidelines and forecasting tools, based on scientific research and better communication of knowledge between meteorologists and the aviation community, would be an important contribution to air safety. Much work has been done to improve forecasting ability for various types of mountain-induced turbulence: for example, clear air turbulence and breaking waves have been extensively researched, and forecasting capability for these phenomena has improved. However, the difficulty in forecasting rotors remains, and so this phenomenon is discussed here in more detail. The work presented in this thesis forms further groundwork for future investigations into orographic features such as rotors.

Although the rotor phenomenon is relatively widely discussed in the literature, much of this is in terms of observational studies either in the real atmosphere (such as Förchtgott, 1952) or in the laboratory. Holets and Swanson (1988) observed a well-organised lee wave and rotor episode in California, which seemed to be a quasi-stationary feature, oscillating about 1 km in the along-wind direction and causing near-surface speeds to vary greatly within periods of 90 minutes or less, thus requiring averaging times of not more than 15 minutes for the wind field in order to follow the movements and dynamics of the wave and rotor. Another observation method which is potentially very useful for studying rotors is Doppler lidar, which can remotely measure radial velocity (towards and away from the instrument) at 1 minute intervals and at high spatial resolution, allowing the temporal evolution of the flow to be studied. Ralph and Neiman (1997), investigating the DC-8 crash over the Rockies in 1992, used lidar data to show that a region of reversed flow developed just above the tropopause due to overturning at a critical level, along with extreme vertical wind shear, at approximately the time and location of the incident. Lidar observations of flow reversal were also used by Ralph *et al.* (1998) when investigating stream-wise rolls, and Neiman *et al.* (1988) used lidar measurements to study a downslope windstorm event over the Rocky Mountains. Olivier and Poulos (1998) used them to show clearly the evolution and decay of a lee wave event and associated rotor flow near Colorado Springs Airport, with accompanying schematic diagrams. Bedard and Neilley (1998) hailed the lidar method as providing "some of the clearest documentation of rotors and flow reversals in the

vicinity of mountains” as well as giving a more detailed description of this and other observation methods appropriate to mountain flows. Real-time output such as lidar data would help aircraft controllers to respond to the evolution of these phenomena, but would be very expensive. However, lidar data could be useful in trials to validate forecast models. Most of the remaining work on rotors is idealised theoretical analysis (such as Scorer and Klieforth, 1959), and it appears that little work has been carried out into modelling and predicting realistic flows with a view to improving understanding of rotor formation. Poulos and Olivier (1998) discussed numerical simulations of the lee rotor event observed by Olivier and Poulos (1998) in the area of Colorado Springs Airport, but further investigation is still required.

The word ‘rotor’ is used in the literature to refer to many different phenomena involving flow reversal or overturning of some kind. The rotors discussed here are defined as stationary closed circulations which form underneath the crest of a trapped lee wave, whose axis of rotation is horizontal and parallel to the crest of the ridge. Figure 2.1 shows schematically the position of the rotor beneath the first lee wave crest, although of course the real situation is far more complex than the simplified version shown here. These lee rotors typically have a horizontal length scale in the range of 550 m to 3.2 km, and very strong vertical velocities between 10 and 25 m s⁻¹ (Knable, 1995).

Whilst the lee wave flow is generally smooth, the associated rotor flow tends to be highly turbulent, making it more difficult to forecast accurately, and specialised knowledge of the local terrain is essential. According to Knable (1995), “most of the severe turbulence [caused by orography] occurs near the first wave crest and associated rotor, typically 3-9 miles downstream of the ridge line”, as illustrated in figure 2.1, which shows the differing degrees of turbulence near the mountain. The twisted streamlines bring warm air under cold, leading to local convective instability with embedded cumulus clouds. As a warning sign, pilots frequently look out for this distinctive rough rotor cloud, which can resemble a cumulus cloud with its base at or below the ridge line, and often merges with the smooth lenticular clouds at the crest of the wave above, showing how narrow the boundary is between rough and smooth air. It is also called a roll cloud because it seems to rotate, with the top part moving forwards and the lower part moving back towards the mountains. As Knable points out, “rotors

can and do form in very dry, cloud free air”, so clearly clouds alone are not a sufficient warning sign. In addition, Carney *et al.* (1996) point out that one consequence of the progressive automation of observations is that cloud information is less readily available. Scorer and Klieforth (1959) mention an occasion during the Sierra Wave Project when a sailplane fell apart in severe turbulence near the upwind edge of the rotor cloud. The pilot descended by parachute, and at lower levels he drifted in the opposite direction from the background wind. The authors called this “a unique observation of the rotor circulation” but it is unlikely that the pilot shared their enthusiasm! The severe turbulence found in a rotor frequently includes a sideways rolling motion, which can be the most hazardous part as it often comes as an unpleasant surprise to even the most experienced pilots. Rotors present a danger not only in the air, but also on the runway, since the flow reversal means that the wind direction on the ground may be the opposite from that measured at another nearby airport, or even by a wind sock at another point on the same runway. These conflicting observations are another warning sign to look out for (Bradbury, 1997 and Findlater, 1984).




Figure 2.1. Schematic showing lee wave flow downstream of a mountain ridge, with position of rotor circulation and cloud formations, and different levels of turbulence (L = light, M = medium, S = severe). After Knable, 1995.

These rotors are a distinct phenomenon from that of wave breaking, which has been extensively modelled and researched. Castro and Snyder's (1993) towing tank experiments showed that the two can occur in close proximity and even merge together in some cases, but these results were for finite depth flows, and their validity in the atmosphere has yet to be researched. Figure 2.2a shows possible mean streamlines for steady two-dimensional flow containing wave breaking, and figure 2.2b shows the same for rotor flow. Although both occur when the lee wave steepens, and at the crest of the wave, these two phenomena are quite distinct. The wave breaking region is above the ground, and the flow within it circulates in both directions. The rotor, on the other hand, is believed from observational evidence to be formed by separation (at point A in figure 2.2b) and reattachment (at point B) of the flow at the surface, and the rotation is entirely in one sense (clockwise in this example). It should be noted that the re-circulating flow in both cases is highly turbulent, and that figure 2.2 only shows mean streamlines, so it does not attempt to represent the detail of the turbulent flow.

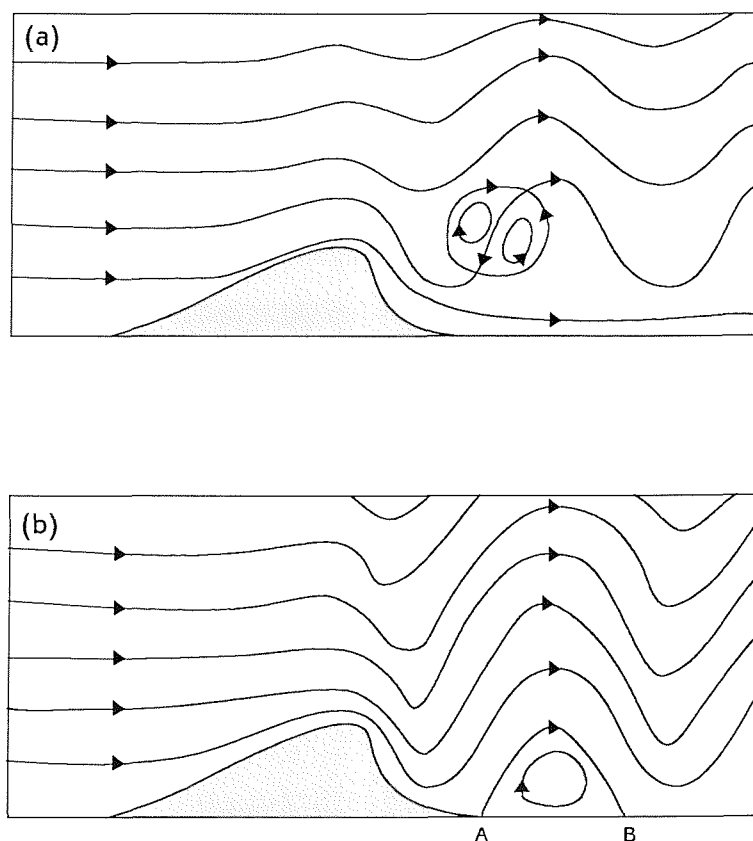


Figure 2.2. Postulated mean streamlines in two-dimensional steady flow containing (a) wave breaking; (b) rotor circulation. In the latter case, the flow separates from the surface at point A, and reattaches at point B.

In three dimensions the flow is more complex, since the streamlines are less restricted than in the two-dimensional case. For example, it becomes possible for the flow to 'roll up' in the plane perpendicular to the ridge, and then exit this roll along its axis of rotation, in the plane parallel to the ridge. Wave breaking in the atmosphere is very rarely observed, but is believed to be an unsteady phenomenon, since it breaks up the wave pattern, which then has to re-establish itself. Rotors, on the other hand, are postulated as a steady feature, linked to stationary, non-breaking, trapped lee waves. They will only vary in time if the wave pattern itself does so, and are unlikely to be unsteady when given steady upstream conditions. Further investigation will be needed to better understand the details of the flow patterns involved, and this is discussed in Chapter 9.

The conditions under which rotors will form are not well known, but one pre-requisite is commonly understood to be a steep lee slope, giving rise to large amplitude lee waves. Examples of this are given by Förchtgott's (1952) work in the former Czechoslovakia, Manley's (1945) study of the Helm Bar cloud, which can be observed when an easterly wind blows across the Pennines, and Küttner's (1959) work in the Sierra Nevada mountain range, which has a steep eastern escarpment favouring rotor formation in a westerly wind. The separation and reattachment mechanism of the rotor is similar to that of a separation 'bubble' on the lee side of the mountain, except that it occurs further downstream, underneath the first lee wave crest. Whether the flow separates on the slope or downstream depends on the relative wavelengths of the mountain and the lee wave. Hunt (1980) discusses the effect of both the Froude number and the steepness of the lee slope on the type of separation which occurs, and provides a useful background for exploring parameter space when attempting numerical simulations of rotors. This phenomenon is clearly important, and merits further investigation; Küttner (1959) goes so far as to suggest that rotors actually play a primary role in the development of the wave field, rather than the other way round.

2.4 Gravity wave drag parametrizations

As discussed in the previous chapter, mountain-induced gravity waves affect the global atmospheric momentum balance, so it is important to include their effects in NWP models. The waves can propagate to great heights, interacting with the mean flow and dissipating as they extract momentum from it, and the drag on the atmosphere which is associated with this wave activity may be greater than the boundary layer frictional drag, especially at high latitudes. The distribution of momentum flux divergence and atmospheric drag is important in NWP in order to correctly model momentum budgets in global models and to forecast the progression of weather systems. However, the waves are generally sub-grid scale for NWP models, because it would be too computationally expensive to run these models on a small enough scale to resolve lee waves and their associated phenomena.

Sawyer (1959) first recognised the importance of GWD in NWP models, noting that orographic gravity waves are able to transport momentum vertically downwards, decelerating the large scale motion at upper levels. He suggested a need to parametrize orographic gravity waves in large scale NWP models, since the wavelengths are generally too small to be accurately resolved. Since then there have been various attempts at GWD parametrization in NWP models and general circulation models (GCMs). These schemes can differ greatly from one model to another, but they all aim to diagnose the change with height of the vertical flux of horizontal momentum carried by the waves. Recently Mobbs (1994) has provided a review of the simple gravity wave theory related to current parametrization schemes, including features such as critical levels and breaking waves. Critical levels are not covered in any detail here, but they correspond to a singularity in the Taylor-Goldstein equation, discussed in the next chapter, and are an important aspect of parametrization schemes. They may absorb a large proportion of the wave energy and transfer it to the mean flow (see Booker and Bretherton, 1967, and Bretherton, 1969). In linear theory the wave energy is assumed

to be totally absorbed with no reflection, but in the real non-linear atmosphere, wave energy may also be trapped in the lower layers of the atmosphere, producing lee waves.

The model surface height is usually derived from a high resolution data set by averaging onto the resolution of the model, so that the height of the model surface at a gridpoint is equal to the mean height of the real surface over the area of the grid square. Although orography has profound effects on the weather at both local and synoptic scales, this smoothing process leads to the loss of many meteorologically important features, such as irregularities on the mountain surface, or narrow isolated peaks. Some models alleviate the effect of smoothing, which reduces the peak height, by using envelope orography, which adds an increment proportional to the standard deviation of the sub-grid scale variance. This means that the scheme is resolution dependent, so finer resolution leads to a decrease in the envelope increment, as would be expected. This scheme is more realistic, including the effects due to sheltered valleys and basins, and increasing the effective height, as well as increasing the area covered by high terrain, thus representing more accurately the physical barrier which mountain ranges present to the atmospheric flow. The increment is largest for rough terrain, which is generally more poorly represented than smooth terrain. Wallace *et al.* (1983) give more details; they generally found that envelope orography reduced errors, although there was some suggestion that the beneficial impact might be limited to synoptic scales or larger, since these features are more important in longer range forecasts, which is where the most marked improvements were seen. Clark and Miller (1991) found that envelope orography improved the accuracy of pressure drag calculations, equivalent to almost doubling the model resolution, but that it was less impressive with respect to Reynolds stress predictions. This led them to conclude that GWD parametrizations still have a useful role to play in augmenting the resolved stress profile. Palmer *et al.* (1986) noted that enhanced orography and GWD parametrizations are complementary to some extent, and that some combination of the two techniques may be optimum. If the model has fine enough resolution, the GWD parametrization is not considered necessary: for example, it is not used in the mesoscale NWP model at The Met. Office.

Since the introduction of GWD parametrizations in the mid 1980s, these schemes have improved the systematic westerly bias which was present in most forecast models, and

which was most marked over the northern hemisphere winter mountain ranges such as the Alps or the Rocky Mountains. There is less land in the southern hemisphere, so the circulation there is less affected by orography, which explains the lower impact seen in that region. Milton and Wilson (1996) showed evidence suggesting that deficiencies in orographic forcing cannot fully explain the errors seen in the mass and wind fields, but GWD parametrizations certainly play a large role in reducing those errors. Boer *et al.* (1984) were the first to implement a GWD parametrization in a large scale model of the lower atmosphere. They described their scheme as “highly simplified and rather crude” but found it to be beneficial. Palmer *et al.* (1986) implemented a parametrization scheme in an NWP model and in a GCM, and found it to be generally helpful in alleviating a bias towards strong westerly winds in the Northern Hemisphere during winter. McFarlane (1987) implemented a similar scheme in a GCM and also found it to have a positive impact. The NWP models had a general tendency to over-forecast upper level westerly winds, so the jet streams were too strong, and low pressure systems were often too low and too mobile. The largest errors were close to where the jet stream crossed the mountain, and systematic errors increased with forecast interval (Wallace *et al.*, 1983). Palmer *et al.* (1986) noted that these errors only became serious when increased resolution stopped models from underestimating poleward flux of momentum by large scale motion, since this error had previously cancelled out the underestimation of surface drag.

One problem with these parametrization schemes is that they tend to assume that the waves generated by the sub-grid scale orography (SGSO) are hydrostatic. Upwardly propagating waves are much better understood now than they used to be, but many questions remain on the subject of trapped waves, including where they dissipate, extracting momentum from the atmosphere and decelerating the mean flow. Keller (1994) obtained linear analytical results for gravity waves generated by idealised mountains, and demonstrated that non-hydrostatic effects can be very important when there is wind shear, even when the mountains are very broad. The assumption of hydrostatic flow can cause significant changes to the wave field and momentum fluxes, and non-hydrostatic flows such as those including trapped waves are now beginning to be represented in parametrization schemes. Two of the four cases simulated in this

thesis were specifically chosen to reproduce Keller's non-hydrostatic results; more details are given in Chapter 5.

2.5 Motivation for this work

Traditionally, boundary layer flows and gravity wave flows have been treated separately, but the boundary layer cannot be ignored in gravity wave cases, because it can significantly change the flow over complex terrain. Mountain-induced wave flows were always assumed to be frictionless, and the ABL was only considered important for flow over smaller orography. A loose distinction between 'mountain' and 'hill' was often used to distinguish between the two cases, and a large roughness length was often used in boundary layer cases to account for the sub-grid scale orography. There was a lack of boundary layer observations over mountains, because it was simpler to use flat ground. The numerical model used here was designed to simulate boundary layer flows, and although previous studies have included lee side separation, little work has been done to simulate flows including both gravity waves and a boundary layer, and the interaction between these phenomena is not well understood. However, recent work with other models has begun to combine these two areas (e.g. Peng and Thompson, 1998). The work presented in this thesis uses idealised cases without a boundary layer, as a test of the model, using a larger mountain and a larger numerical domain than before, with a view to providing a sound base for turbulent boundary layer simulations, including phenomena such as rotors. Upwardly propagating waves and wave breaking are relatively well understood and represented now, but work is ongoing to further understanding of the dynamics of trapped lee waves and associated phenomena, such as the rotors discussed above. This will lead to improvements in parametrizations and NWP performance. The simulations in this work are two-dimensional, because much of the analytical work is in two dimensions and because the computational cost is far lower than running in three dimensions. However, for the work following on from this, looking at rotors, three-dimensional simulations will be necessary, and some of the issues surrounding this are discussed in Chapter 9. This work does not attempt to

simulate flow over a realistic mountain, for comparison with observational data; that will become important in future work. The simulations presented here aim simply to reproduce the important processes within the flow for idealised conditions, and to compare the results obtained with analytical theory. The numerical model should be capable of simulating these idealised flows, as well as more realistic flows, and its use will enable an improved understanding of the formation of features such as rotors, with the potential to improve parametrizations and forecasts by including their effects. In addition, the work may lead to a local forecasting tool for aviation use, to reduce the hazard posed by these phenomena which are currently very difficult to predict. The work described in this thesis lays further ground work for future studies of this nature.

3 Equations of motion

3.1 Basic equations

3.1.1 Introduction and notation

This chapter explains the basic equations governing atmospheric flows, and some of the assumptions and approximations which are made in both analytical and numerical models of these flows. The atmosphere is regarded as a continuous fluid whose state can be specified in both space and time by unique values of the physical variables. The equations given in this chapter are all in the standard rotating frame of reference, denoted by Cartesian co-ordinates (x, y, z) , or (x_1, x_2, x_3) in tensor notation. Unit vectors in the x -, y - and z -directions are denoted by $(\underline{i}, \underline{j}, \underline{k})$ respectively. The standard notation for the mathematical operators is used here and is explained below, where ϕ represents any field.

$\frac{\partial \phi}{\partial t}$ denotes the Eulerian derivative: the local rate of change, at a fixed point.

$\nabla \phi \equiv \left(\frac{\partial \phi}{\partial x}, \frac{\partial \phi}{\partial y}, \frac{\partial \phi}{\partial z} \right)$ denotes the gradient of ϕ .

$\nabla \cdot \underline{\phi} \equiv \frac{\partial \phi_1}{\partial x} + \frac{\partial \phi_2}{\partial y} + \frac{\partial \phi_3}{\partial z}$ denotes the divergence of $\underline{\phi}$.

$\nabla^2 \phi \equiv \frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} + \frac{\partial^2 \phi}{\partial z^2}$ is the Laplacian operator (sometimes referred to as del-squared).

$\underline{u} \cdot \nabla \phi \equiv u \frac{\partial \phi}{\partial x} + v \frac{\partial \phi}{\partial y} + w \frac{\partial \phi}{\partial z}$ is the advection term and represents the advective rate of change at a fixed point, where $\underline{u} = (u, v, w) \equiv (u_1, u_2, u_3)$ is the velocity vector with components u , v , w in the x -, y - and z -directions respectively.

$\frac{D\phi}{Dt} \equiv \frac{\partial \phi}{\partial t} + \underline{u} \cdot \nabla \phi$ denotes the Lagrangian derivative: the rate of change following the motion of an air parcel. This is also known as the material, total or substantive derivative.

The basic equations governing flow in the atmosphere arise from three fundamental physical principles: conservation of mass, momentum and energy. The equations arising from these laws are, respectively, the continuity equation, the Navier-Stokes equations and the thermodynamic equation. For a more detailed background to the theory and derivation of the equations than is given here, there are many good sources, such as Batchelor (1967) and Garratt (1992).

3.1.2 Conservation of mass

The mass continuity equation is given by

$$\frac{D\rho}{Dt} + \rho \nabla \cdot \underline{u} \equiv \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \underline{u}) = 0$$

where ρ is the density.

There are two approximations which are commonly applied to this equation. The first is the anelastic approximation, and eliminates all acoustic waves, which have no meteorological significance and are believed to have a negligible interaction with waves of interest. They are the fastest non-electromagnetic waves in the atmosphere, and their removal allows more economical use of certain numerical methods, by relaxing the stability restrictions on the timestep (Pielke, 1984). The local rate of change of density is assumed to be negligible, i.e. $\frac{\partial \rho}{\partial t} = 0$, and the continuity equation becomes

$\nabla \cdot (\rho \underline{u}) = 0$. Horizontal gradients of density are also assumed to be negligible, and then the anelastic approximation is valid if the depth of the circulations considered is much

less than the density scale height, given by $H_\rho = -\left(\frac{1}{\rho} \frac{d\rho}{dz}\right)^{-1}$, which is typically around 8 km for the atmosphere.

The second approximation is one of incompressibility, and assumes that density variations from a constant background density profile are negligible in both time and space. Air is not an incompressible fluid, but if the flow is assumed to be approximately incompressible, the mass continuity equation reduces to

$$\nabla \cdot \underline{u} \equiv \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} = 0.$$

This is more strict than the anelastic approximation, and density no longer appears in the continuity equation. This approximation is justified for boundary layer models with shallow domains, as the density variations are then very small within the region of interest. However, it becomes less accurate for deeper domains, in which vertical motion leads to compressibility associated with the height dependence of the background density profile, which must be taken into account. The incompressible approximation is used in the rest of this thesis, because it is simpler to implement numerically. Although not strictly applicable for the deep domains considered in this work, the results obtained for cases 2 and 3 (described in Chapter 5) are in broad agreement with recent results obtained at The Met. Office for the same cases using a fully compressible model (Smith, private communication). The results are quantitatively very similar at low levels, becoming less so above 10 or 20 km, but the qualitative agreement at higher levels remains good. This agreement is sufficient to justify the use of the incompressible approximation for this work, although future studies will continue to use compressible simulations.

3.1.3 Conservation of momentum

Newton's second law (e.g. Batchelor, 1967) gives

$$\frac{Du_i}{Dt} = F_i + \frac{1}{\rho} \frac{\partial \sigma_{ij}}{\partial x_j}.$$

F_i represents the body forces (gravity and the Earth's rotation, in this case), and

$\sigma_{ij} = -p\delta_{ij} + \mu \left(e_{ij} - \frac{2}{3} \frac{\partial u_k}{\partial x_k} \right)$ is the stress tensor, where p is the pressure, δ_{ij} is the

Kronecker delta, μ is the dynamic viscosity and $e_{ij} = \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}$. The standard

summation convention is used in terms such as $\frac{\partial u_k}{\partial x_k}$ which contain a repeated index, so

$$\text{that } \frac{\partial u_k}{\partial x_k} \equiv \frac{\partial u_1}{\partial x_1} + \frac{\partial u_2}{\partial x_2} + \frac{\partial u_3}{\partial x_3} \equiv \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}.$$

For an incompressible rotating fluid with constant viscosity, the momentum equations arising from Newton's second law, known as the Navier-Stokes equations, are

$$\frac{D\underline{u}}{Dt} = -\frac{1}{\rho} \nabla p - 2(\underline{\Omega} \times \underline{u}) + \underline{g} + \nu \nabla^2 \underline{u},$$

where $\underline{g} = -g\underline{k}$ is the gravitational acceleration, $\nu = \frac{\mu}{\rho}$ is the kinematic molecular

viscosity, and $\underline{\Omega}$ is the angular velocity of the Earth, usually taken to have magnitude $\Omega = |\underline{\Omega}| \approx 7.27 \times 10^{-5} \text{ s}^{-1}$. The terms represent, from left to right, the acceleration of the air (due to inertial forces), the pressure gradient force, the effects of the Earth's rotation, the gravitational force (which includes a component due to the rotation of the Earth) and the viscous stresses.

The term $2(\underline{\Omega} \times \underline{u})$ represents the effects of the Earth's rotation. Its expanded form is $(-(2\Omega \sin \phi)v + (2\Omega \cos \phi)w)\underline{i} + ((2\Omega \sin \phi)u)\underline{j} - ((2\Omega \cos \phi)u)\underline{k}$, where ϕ represents latitude here. Traditionally, the $(2\Omega \cos \phi)w$ term in the x -component is neglected using scale analysis, since $w \ll u$, and the z -component is neglected because it is much smaller than other terms in the w -momentum equation. However, these terms can be important, and are retained in some NWP models (e.g. White and Bromley, 1995). Since the simulations in this thesis are all two-dimensional, and Coriolis forces are neglected entirely (as discussed later), this issue will not be pursued here. The 'traditional' approximation is used instead, where the terms in $\cos \phi$ are neglected, and $2\Omega \sin \phi$ is denoted by f and called the Coriolis parameter, giving

$2(\underline{\Omega} \times \underline{u}) \approx -fv\underline{i} + fu\underline{j}$. The f-plane approximation takes f as a constant (typically 10^{-4} s^{-1} in mid-latitudes), and neglects the curvature of the earth.

The Navier-Stokes equations are then reduced to

$$\frac{Du}{Dt} = -\frac{1}{\rho} \frac{\partial p}{\partial x} + fv + \nu \nabla^2 u$$

$$\frac{Dv}{Dt} = -\frac{1}{\rho} \frac{\partial p}{\partial y} - fu + \nu \nabla^2 v$$

$$\frac{Dw}{Dt} = -\frac{1}{\rho} \frac{\partial p}{\partial z} + g + \nu \nabla^2 w$$

and the equations are used in this form in the next section of this chapter.

3.1.4 Conservation of energy

Businger (1982) gives much more detail of the full thermodynamic equation than is covered here, and discusses the assumptions which are commonly made. The full equation for conservation of energy is complicated, so some simplifications are made here in order to arrive at the thermodynamic equation given below. In a gas, conduction is an extremely inefficient heat transfer mechanism in comparison with convection, and outside the boundary layer, radiative transfer (not to be confused with the use of the term ‘radiative’ to refer to the open downstream boundary condition later in this thesis) only becomes significant on time scales of days, so a useful and valid approximation is to assume that the processes responsible for temperature change in an air parcel are adiabatic. Given that the boundary layer is not included in the simulations presented here, and that condensation is not relevant (since, as discussed in the previous chapter, all these simulations use dry dynamics), this is a valid approximation. Entropy is denoted by S and defined as $S = c_p \ln \theta + \text{const.}$, where c_p is the specific heat capacity for air at constant pressure (usually taken as $1005 \text{ J kg}^{-1} \text{ K}^{-1}$). In the absence of radiative transfer or phase changes, the motion is adiabatic, so there is no change in entropy, and the flow is said to be isentropic. This means that potential temperature is constant along streamlines, and the thermodynamic equation can be expressed as the conservation of potential temperature, θ , following the fluid motion:

$$\frac{D\theta}{Dt} = 0.$$

Potential temperature is defined as the temperature a parcel of air would have if its pressure were altered adiabatically (meaning that no heat enters or leaves the system) to some reference pressure p_r , which here is the surface pressure, usually taken to be 1000 hPa:

$$\theta = T \left(\frac{p_r}{p} \right)^\kappa.$$

Here T is the temperature (in Kelvin), p is the pressure (in hPa), and κ is defined as $\frac{R}{c_p}$, where R is the universal gas constant for dry air (usually taken as $8.31 \text{ J mol}^{-1} \text{ K}^{-1}$).

If the gradient with height of the background potential temperature profile is greater than, equal to or less than zero, the flow is said to be statically stable, neutral or unstable respectively. The buoyancy frequency, or Brunt-Väisälä frequency, N , is defined as

$$N^2 = \frac{g}{\theta} \frac{\partial \theta}{\partial z},$$

so the stability condition can also be expressed in terms of N^2 , where a

larger buoyancy frequency means that the air is more stable. The atmosphere is always stably stratified on the synoptic scale, because unstable regions, where denser air is above less dense air, are quickly stabilised by mixing. In general, the faster the air cools with height, the less stable it is. A stable air mass resists vertical motion, so any vertical displacement will tend to return to its original position. This can set up an oscillation and lead to waves such as those discussed in section 3.5.

3.2 Flow parameters

3.2.1 Introduction

The nature of the flow can be described using non-dimensional flow parameters, which represent the ratio of various terms in the basic equations. This allows some terms to be compared in magnitude, and the smaller terms discarded as not significant.

Non-dimensional notation is used, where $\hat{\phi}$ represents a non-dimensional variable.

Various scales are defined as follows:

L and H are the horizontal and vertical length scales, so $x = L\hat{x}, y = L\hat{y}, z = H\hat{z}$;

U and W are the horizontal and vertical velocity scales, so $u = U\hat{u}, v = U\hat{v}, w = W\hat{w}$;

L/U is the advective time scale, so $t = \frac{L}{U}\hat{t}$.

The pressure scale is dependent on the type of flow considered. For steady flow with no viscosity, Bernoulli's theorem gives $p \sim \rho u^2$. Another alternative is to use the geostrophic equations, which are obtained by neglecting the acceleration and viscous terms in the horizontal momentum equations, so the pressure gradient term is comparable to the Coriolis term:

$$-fv = -\frac{1}{\rho} \frac{\partial p}{\partial x}, fu = -\frac{1}{\rho} \frac{\partial p}{\partial y}.$$

The former is more relevant for the steady, laminar cases considered here, and the pressure scale arising from Bernoulli's equation is ρU^2 , giving $p = \rho U^2 \hat{p}$.

The non-dimensional u -momentum equation is given below, with the terms labelled for future reference:

$$\frac{U^2}{L} \left[\underbrace{\left(\frac{\partial \hat{u}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{u}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{u}}{\partial \hat{y}} + \frac{W}{U} \frac{L}{H} \hat{w} \frac{\partial \hat{u}}{\partial \hat{z}} \right)}_A - \underbrace{\frac{fL}{U} \hat{v}}_B = - \underbrace{\frac{\partial \hat{p}}{\partial \hat{x}}}_C + \underbrace{\frac{\nu}{UL} \left(\frac{\partial^2 \hat{u}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{y}^2} + \frac{L^2}{H^2} \frac{\partial^2 \hat{u}}{\partial \hat{z}^2} \right)}_D \right] = 0.$$

The following subsections introduce different non-dimensional parameters based on this equation.

3.2.2 The Reynolds number

The viscous terms in the momentum equations represent molecular diffusion, the process by which contiguous fluids mix. The Reynolds number, Re , is the ratio of the advection terms, A, to the viscous terms, D, and is defined as $Re = \frac{UL}{\nu}$. For

atmospheric flows, typical values are $U \sim 10 \text{ms}^{-1}, L \sim 10^4 \text{m}$ and $\nu \sim 10^{-5} \text{m}^2 \text{s}^{-1}$, where the length scale given is typical for the flows considered here (but there are atmospheric

flows on many different scales). This gives a typical Reynolds number of order 10^{10} , showing that the molecular viscosity term in the equations of motion can be neglected. For Reynolds numbers smaller than about 2000 (Acheson, 1990), turbulent flow is not sustained, and the flow is laminar, but typical Reynolds numbers in the atmosphere tend to be far larger than this, as discussed above. The molecular viscosity is thus negligible with respect to mixing by turbulent eddies (eddy diffusion), being too small to dissipate turbulence except at very small scales. The viscous terms are neglected in the rest of this thesis, although a very low viscosity ($\nu = 10^{-33} \text{ m}^2 \text{ s}^{-1}$) is necessary in the numerical simulations, as discussed in Chapter 4.

3.2.3 The Rossby number

The Rossby number, Ro , is the ratio of the advection terms, A , to the Coriolis term, B , and is defined as $Ro = \frac{U}{fL}$. An alternative definition is to use Ω instead of f , which is equivalent in both dimensional and scaling terms, since $f = 2\Omega \sin \phi$, as discussed above. This parameter indicates how important Coriolis effects are over the time scale of the flow being considered. For atmospheric flow at mid-latitudes, typical values are $U \sim 10 \text{ m s}^{-1}$, $f \sim 10^{-4} \text{ s}^{-1}$ and $L \sim 10^4 \text{ m}$, giving a typical Rossby number of order 10. A large Rossby number, which is much greater than unity (Pielke, 1984), means that the time scale of the flow being modelled is so short that deflection of the flow due to Coriolis effects is negligible, and so terms involving f can be neglected from the equations. Although the numerical model used here retains them, as do the three-dimensional equations discussed in the rest of this thesis, the two-dimensional simulations and equations in this thesis neglect the Coriolis terms, since they deal with flow over an infinite ridge which is narrow enough that the Rossby number governing the flow is large.

3.2.4 The Froude number

While not obtained from the u -momentum equation in the same way as the Reynolds and Rossby numbers, the Froude number is an important parameter in orographic flows.

It is defined as $F_h = \frac{U}{Nh_m}$, where h_m is the height of the mountain crest, and U and N take typical near-surface values here. It is worth noting at this point that some authors define the Froude number as the inverse of that given above, sometimes referring to it as the non-dimensional mountain height, and that the parameter defined here is sometimes referred to as the ‘inverse Froude number’. The definition given here, however, does seem to be the most widely used. Another potential source of ambiguity is that the Froude number defined here is sometimes referred to as the vertical Froude number; the horizontal Froude number is defined as $F_\alpha = \frac{Uk}{N}$, where k is the dominant horizontal wavenumber. This is discussed further in Chapter 5 for the four cases studied here. The term ‘Froude number’ in this thesis will be used to refer to the vertical Froude number, F_h , as defined at the start of this subsection.

The Froude number provides a measure of linearity, and is used in the discussions of the four cases described in Chapter 5. In general, if the Froude number is greater than unity, the flow is linear and the air flows over the mountain, whereas if the Froude number is less than unity, the flow is non-linear and can be much more complicated, including features such as upstream blocking, wave breaking, vortex shedding and flow around the mountain. The division between these high and low Froude number flow regimes is, of course, not clear-cut. It should be noted that in the two-dimensional numerical simulations studied here, the air cannot flow around the mountain, but non-linear effects can still be seen in some cases. Smith (1989a) includes a review of analytical investigations into low Froude number flow, including the early non-linear theory of Drazin (1961). The dynamics are complex, so much of the research carried out has involved laboratory experiments and numerical modelling. Laboratory examples of investigations into the effect of Froude number on the character of flow over obstacles include Hunt and Snyder (1980), Castro *et al.* (1983), Baines and Hoinka (1985) and Boyer *et al.* (1987). Snyder *et al.* (1985) and Smith (1989a, 1989b) discuss the concept of a dividing streamline between the flow over and around the mountain. High resolution non-linear simulations can be found in Hanazaki (1988), Smolarkiewicz and Rotunno (1990) and Miranda and James (1992).

3.3 Turbulent flow

Turbulence is briefly covered here for completeness, because the numerical model described in the next chapter does include it, but the work described in this thesis simulates laminar flow, as discussed in Chapter 4. This section explains the main differences between laminar and turbulent flow, then derives the turbulent governing equations from those obtained in the previous sections. An introduction to the turbulence closure problem follows at the end of this section.

Laminar flow is predominantly smooth, and small elements of fluid retain their identity without mixing with their surroundings. Momentum, heat and matter are transferred by the interaction of individual fluid molecules, and although the flow can be unsteady, a typical particle path is repeatable (i.e. stable to small perturbations). In turbulent flows, the particle paths become irregular, random oscillations, known as eddies, which typically occur on a range of different time and space scales. Because of the turbulent motions, momentum, heat and matter are transferred much more efficiently than in laminar flow, with an effective diffusivity which is usually many orders of magnitude greater than the molecular diffusivity of the equivalent laminar flow. It is turbulence which allows the boundary layer to respond rapidly to changing surface forcings, and most naturally occurring flows are turbulent.

Direct numerical simulation aims to resolve all scales of turbulence by using a very fine grid. It can provide useful results for fluid flow problems at low Reynolds numbers, which cover a small range of scales. However, it becomes computationally impractical at Reynolds numbers approaching realistic atmospheric values, because the length scales of interest range from several kilometres to small scale turbulence. DNS results can be applied to high Reynolds number atmospheric flows in some cases, but this requires care (Coleman *et al.*, 1992). An alternative approach is large eddy simulation (LES), which allows the numerical model to resolve the larger energy-containing turbulent eddies within the flow, and parametrizes only the unresolved smaller scale

eddies (ideally those in the inertial sub-range). This produces flow fields which reflect the turbulent nature of the flow, due to the resolved eddies, and it can be very computationally expensive because of the high resolution and long simulation times required to obtain stable statistics. Stull (1988) and Garratt (1992) provide more information on this approach, and Wood (2000) provides an introduction to LES as the next major step in modelling flow over complex terrain. Dörnbrack and Schumann (1994) compared DNS and LES results for breaking gravity waves below a critical level, and observed more efficient turbulent mixing in the latter case. In the numerical model described in the next chapter, all of the turbulence in the flow is parametrized, and this method is described below. The model does have an LES option, but this was not used in the work presented here.

Turbulent profiles can be measured using anemometers, such as in the field campaigns of Mason and King (1985), and Mickle *et al.* (1988), but the non-linearity causes random motions, which make deterministic description difficult. The analytical solution is mathematically intractable, and so statistical methods are used. As discussed above, turbulent flows fluctuate on a wide range of length and time scales, and this large range means that all the scales of motion cannot be modelled explicitly, so averaging is required. The Reynolds averaging method used here assumes two peaks in the temporal and spatial distribution: the mean flow, on the scale of interest, and the turbulence, on much smaller scales. The ‘spectral gap’ is used to separate these two parts of the flow, but in reality it is not a clear split and the two do interact. More detail of this method is given in Garratt (1992). The process is independent of the quantity being averaged, so ϕ is used to denote any one of u , v , w , or θ .

For an incompressible fluid, the continuity equation is $\nabla \cdot \underline{u} = 0$, and this means that the material derivative can be expressed in flux form as

$$\frac{D\phi}{Dt} \equiv \frac{\partial\phi}{\partial t} + \underline{u} \cdot \nabla\phi + \phi \nabla \cdot \underline{u} = \frac{\partial\phi}{\partial t} + \frac{\partial(\phi u)}{\partial x} + \frac{\partial(\phi v)}{\partial y} + \frac{\partial(\phi w)}{\partial z},$$

where a flux is defined as the transfer of a quantity per unit area per unit time.

The averaging operator separates each field into a mean part, denoted by a bar, and a turbulent fluctuation, denoted by a prime: $\phi = \bar{\phi} + \phi'$. By definition, the mean of the

turbulent fluctuation is zero, since $\bar{\phi} = \overline{\bar{\phi} + \phi'} = \bar{\phi} + \bar{\phi}'$, implying that $\bar{\phi}' = 0$. The mean part can be a time average, a space average or an ensemble average. The latter is obtained by averaging data from a repeated set of experiments, observations or numerical simulations. It is still a function of both space and time, so it has an advantage over the other two. It is also preferable in the laboratory, because it minimises the impact of random experimental errors by allowing the experiment to be repeated several times. In the atmosphere, however, it is rare to observe reproducible weather events, so ensemble averages are more difficult to use (Stull, 1988).

Substituting this decomposition into the above equation gives

$$\frac{D\phi}{Dt} = \frac{\partial \bar{\phi}}{\partial t} + \frac{\partial \phi'}{\partial t} + \frac{\partial}{\partial x} [(\bar{\phi} + \phi')(\bar{u} + u')] + \frac{\partial}{\partial y} [(\bar{\phi} + \phi')(\bar{v} + v')] + \frac{\partial}{\partial z} [(\bar{\phi} + \phi')(\bar{w} + w')].$$

The whole equation is then averaged using the standard technique, of which details can be found in Garratt (1992), to give

$$\frac{D\bar{\phi}}{Dt} = \frac{\bar{D}}{Dt}(\bar{\phi}) + \frac{\partial(\bar{\phi}'u')}{\partial x} + \frac{\partial(\bar{\phi}'v')}{\partial y} + \frac{\partial(\bar{\phi}'w')}{\partial z}.$$

The first term on the right hand side of this equation represents advection of the large scale flow by the large scale flow, where $\frac{\bar{D}}{Dt}(\bar{\phi}) \equiv \frac{\partial \bar{\phi}}{\partial t} + \bar{\underline{u}} \cdot \nabla \bar{\phi}$ is the material derivative with advection by the mean flow. The remaining terms, called the Reynolds stress gradients, represent the contribution to momentum (or heat) transport by smaller scale turbulent fluctuations in velocity or temperature. These terms are also referred to as the shear stresses or the eddy fluxes. Turbulent flux covariances, that is terms such as $\overline{\phi'u'}$, act as source terms, and they are a direct result of the non-linearity in the advection terms.

Applying the above method to the incompressible equations of motion derived in section 3.1, neglecting the laminar viscous terms as discussed in the previous section, gives:

$$\frac{\bar{D}}{Dt}(\bar{u}) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x} + f\bar{v} - \left[\frac{\partial(\overline{u'u'})}{\partial x} + \frac{\partial(\overline{u'v'})}{\partial y} + \frac{\partial(\overline{u'w'})}{\partial z} \right]$$

$$\frac{\bar{D}}{Dt}(\bar{v}) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial y} - f\bar{u} - \left[\frac{\partial(\overline{v'u'})}{\partial x} + \frac{\partial(\overline{v'v'})}{\partial y} + \frac{\partial(\overline{v'w'})}{\partial z} \right]$$

$$\frac{\bar{D}}{Dt}(\bar{w}) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial z} + g - \left[\frac{\partial(\overline{w'u'})}{\partial x} + \frac{\partial(\overline{w'v'})}{\partial y} + \frac{\partial(\overline{w'w'})}{\partial z} \right]$$

$$\nabla \cdot \bar{\underline{u}} = 0$$

$$\frac{\bar{D}}{Dt}(\bar{\theta}) = - \left[\frac{\partial(\overline{\theta'u'})}{\partial x} + \frac{\partial(\overline{\theta'v'})}{\partial y} + \frac{\partial(\overline{\theta'w'})}{\partial z} \right]$$

and a corollary of the mean continuity equation is that the divergence of the turbulent fluctuations is also zero: $\nabla \cdot \underline{u}' = 0$.

The turbulent momentum fluxes are often represented by $\tau_{ij} = -\bar{\rho}(\overline{u'_i u'_j})$, the Reynolds stress tensor, and the turbulent heat fluxes by $H_j = -\bar{\rho}(\overline{\theta' u'_j})$. The momentum and heat equations can then be written as follows:

$$\frac{\bar{D}}{Dt}(\bar{u}) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x} + f\bar{v} + \frac{1}{\rho} \left[\frac{\partial \tau_{11}}{\partial x} + \frac{\partial \tau_{12}}{\partial y} + \frac{\partial \tau_{13}}{\partial z} \right]$$

$$\frac{\bar{D}}{Dt}(\bar{v}) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial y} - f\bar{u} + \frac{1}{\rho} \left[\frac{\partial \tau_{21}}{\partial x} + \frac{\partial \tau_{22}}{\partial y} + \frac{\partial \tau_{23}}{\partial z} \right]$$

$$\frac{\bar{D}}{Dt}(\bar{w}) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial z} + g + \frac{1}{\rho} \left[\frac{\partial \tau_{31}}{\partial x} + \frac{\partial \tau_{32}}{\partial y} + \frac{\partial \tau_{33}}{\partial z} \right]$$

$$\frac{\bar{D}}{Dt}(\bar{\theta}) = \frac{1}{\rho} \left[\frac{\partial H_1}{\partial x} + \frac{\partial H_2}{\partial y} + \frac{\partial H_3}{\partial z} \right]$$

and the overbars will be omitted from the averaged quantities from now on, except in the turbulent covariance flux terms.

These equations cannot be solved in their present form, because the turbulent flux terms are unknown. Although it is possible to construct equations for these second order terms, this process introduces further unknowns in the form of third order terms. The situation continues like this, and is known as the closure problem: in order to close the

equations it is necessary to stop at some point and parametrize the unknown higher order terms. This closure problem is associated with the non-linear characteristics of turbulence, and remains one of the unsolved problems of classical physics; Businger (1982) gives a more detailed discussion. The closure scheme described in Chapter 4 and used in the numerical model presented in this thesis is first order, because it solves for the first order terms and parametrizes the second order terms. By analogy with molecular diffusion, the equations are closed using eddy diffusivities to relate the turbulent fluxes to the mean flow using flux-gradient relations of the form

$$\overline{\phi'u_j'} = -K_\phi \frac{\partial \bar{\phi}}{\partial x_j}.$$

For momentum, K_ϕ is called the eddy viscosity or momentum diffusivity, and for heat, K_ϕ is called the eddy conductivity or eddy thermal diffusivity. It measures the rate of diffusion effected by atmospheric eddies. In contrast with the molecular case, the eddy diffusivity is not a property of the fluid, but of the flow, and may be a function of many quantities, including height, stability and flow velocity. The principle behind first order schemes is to derive K_ϕ from known quantities, so first order schemes are also known as K -closure schemes. There are various methods for calculating K_ϕ but no detail is given here, since there is no turbulence in the results presented in this thesis. Garratt (1992) and Stull (1988) give a more detailed overview of this type of scheme, and others.

3.4 The Boussinesq approximation

Various approximations and assumptions can be applied to make the equations less complicated and easier to solve, but which ones of these are appropriate depends on the nature of the flow being considered. This section explains the Boussinesq approximation, which is used in the simulations described in this thesis, and is attributed to Boussinesq (1903); more details of the assumptions made, and of their validity and limitations, are given by Mahrt (1986) and Businger (1982).

The mean values of the thermodynamic variables (pressure, potential temperature and density) used in the previous section, here represented generally by ϕ , are separated into a background profile, ϕ_o , which is assumed to vary with height only, and a perturbation, ϕ'' , which is assumed to be of small amplitude such that $\phi'' \ll \phi_o$:

$\phi(x, y, z, t) = \phi_o(z) + \phi''(x, y, z, t)$. It should be noted that ϕ here represents the mean value (denoted by $\bar{\phi}$ in the previous section), and that ϕ'' is not a turbulent fluctuation like ϕ' in the previous section, but a perturbation from the background state. This perturbation is on an intermediate time and length scale, between the scales of the mean value and the turbulent fluctuation.

In the vertical component of the momentum equation for the background state (but not for the main flow), the vertical acceleration terms are neglected, and this leads to the hydrostatic equation, which states that the background pressure is proportional to the weight of the fluid above:

$$\frac{\partial p_o}{\partial z} = -\rho_o g.$$

This is valid on synoptic scales, where perturbations vanish because of the large spatial scale and long time scale, and is a good approximation for wave modes with large horizontal wavelength, where the vertical accelerations can be assumed to be small. It is only used for the background profiles here; no assumptions have been made about the main flow. The hydrostatic approximation filters out upwardly propagating sound waves, which are of no meteorological interest; the remaining horizontally propagating acoustic waves, called Lamb waves, are eliminated by using the anelastic approximation, as discussed above. The hydrostatic approximation is used as a basis for computations of height from vertical soundings, and is usually sufficient for this purpose, since the horizontal scale of motion is much larger than the vertical scale. It is a good approximation to the vertical distribution of background pressure, and only breaks down for small scale phenomena such as thunderstorm updrafts, or some orographic flows, where the background vertical acceleration cannot be neglected. These phenomena, however, are not explicitly resolved in global NWP models, where the hydrostatic approximation is often used and is valid.

The Boussinesq approximation has many forms, and these are discussed by Spiegel and Veronis (1960), Dutton and Fichtl (1969) and Mahrt (1986), among others. The form used here is known as the shallow convection form, because vertical motion is limited by the buoyancy term; it requires the vertical scale of motions to be much less than the density scale height. For a compressible fluid, a better approximation is to use the continuity equation in the form $\nabla \cdot (\rho \underline{u}) = 0$, and this is the deep convection form of the Boussinesq approximation.

In the shallow convection Boussinesq approximation, the flow is taken to be incompressible, and the background profile is taken to be hydrostatic. It is assumed that density variations from the constant background state can be neglected except in the buoyancy term, where they are coupled to the gravitational acceleration, since the buoyancy term is of approximately the same order of magnitude as the vertical acceleration term. The governing equations derived in the previous section then become

$$\begin{aligned} \frac{Du}{Dt} &= -\frac{1}{\rho_0} \frac{\partial p''}{\partial x} + fv + \frac{1}{\rho_0} \left[\frac{\partial \tau_{11}}{\partial x} + \frac{\partial \tau_{12}}{\partial y} + \frac{\partial \tau_{13}}{\partial z} \right] \\ \frac{Dv}{Dt} &= -\frac{1}{\rho_0} \frac{\partial p''}{\partial y} - fu + \frac{1}{\rho_0} \left[\frac{\partial \tau_{21}}{\partial x} + \frac{\partial \tau_{22}}{\partial y} + \frac{\partial \tau_{23}}{\partial z} \right] \\ \frac{Dw}{Dt} &= -\frac{1}{\rho_0} \frac{\partial p''}{\partial z} - \frac{\rho''}{\rho_0} g + \frac{1}{\rho_0} \left[\frac{\partial \tau_{31}}{\partial x} + \frac{\partial \tau_{32}}{\partial y} + \frac{\partial \tau_{33}}{\partial z} \right] \\ \nabla \cdot \underline{u} &= 0 \\ \frac{D\theta}{Dt} &= \frac{1}{\rho_0} \left[\frac{\partial H_1}{\partial x} + \frac{\partial H_2}{\partial y} + \frac{\partial H_3}{\partial z} \right] \end{aligned}$$

where the buoyancy variable, $-\frac{\rho''}{\rho_0} g$, which now appears in the vertical momentum

equation, involves density fluctuations, which are hard to measure. The ideal gas equation and the potential temperature equation can be linearised to give, to a good approximation,

$$\frac{\theta''}{\theta_0} \approx \left(1 - \frac{\rho''}{\rho_0}\right) \left(1 - \frac{p''}{p_0}\right)^{\frac{R}{c_p} - 1} - 1.$$

Dutton and Fichtl (1969) showed that in the shallow convection Boussinesq approximation, the relative pressure fluctuations from the mean state are much less than

the density and temperature fluctuations, and can be neglected: $\frac{p''}{\rho_0} \ll \frac{\theta''}{\theta_0}, \frac{\rho''}{\rho_0}$. In this

case the above equation gives an approximate buoyancy variable $+\frac{g\theta''}{\theta_0}$, and the

vertical momentum equation becomes

$$\frac{Dw}{Dt} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial z} + \frac{g\theta''}{\theta_0} + \frac{1}{\rho_0} \left[\frac{\partial \tau_{31}}{\partial x} + \frac{\partial \tau_{32}}{\partial y} + \frac{\partial \tau_{33}}{\partial z} \right].$$

For ease of reference, the governing equations for turbulent, incompressible Boussinesq flow derived in this chapter are given below in full, and these are the equations which are used in the numerical model described in the following chapter.

$$\frac{Du}{Dt} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial x} + fv + \frac{1}{\rho_0} \left[\frac{\partial \tau_{11}}{\partial x} + \frac{\partial \tau_{12}}{\partial y} + \frac{\partial \tau_{13}}{\partial z} \right]$$

$$\frac{Dv}{Dt} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial y} - fu + \frac{1}{\rho_0} \left[\frac{\partial \tau_{21}}{\partial x} + \frac{\partial \tau_{22}}{\partial y} + \frac{\partial \tau_{23}}{\partial z} \right]$$

$$\frac{Dw}{Dt} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial z} + \frac{g\theta''}{\theta_0} + \frac{1}{\rho_0} \left[\frac{\partial \tau_{31}}{\partial x} + \frac{\partial \tau_{32}}{\partial y} + \frac{\partial \tau_{33}}{\partial z} \right]$$

$$\nabla \cdot \underline{u} = 0$$

$$\frac{D\theta}{Dt} = \frac{1}{\rho_0} \left[\frac{\partial H_1}{\partial x} + \frac{\partial H_2}{\partial y} + \frac{\partial H_3}{\partial z} \right]$$

3.5 Linear mountain waves for a simple idealised case

In general, the fully non-linear equations of motion governing flow over orography cannot be solved analytically. Under certain restrictive assumptions, the non-linear equations become linear even for finite amplitude disturbances (e.g. Long, 1953 and 1955; Huppert and Miles, 1969), but generally the equations need to be linearised and reduced to a steady state. Examples of early analytical work are found in Lyra (1940, 1943) and in Queney (1948), which is a summary of lee wave theory in stratified flow. Reviews of two-dimensional hydrostatic steady solutions include Miles (1969), Nicholls

(1973), Baines (1995) and Smith (1979, 1989a). The latter two reviews by Smith also discuss three-dimensional linear hydrostatic solutions, and Wurtele *et al.* (1996) provide an overall review of lee waves in the atmosphere. Analytical solutions using asymptotic methods for linear non-hydrostatic flows, such as those described by Crapper (1959) or Janowitz (1984), can provide solutions which are valid far from the mountain.

In order to simplify the governing equations derived in the previous sections and allow their analytical solution for an idealised case, several assumptions are applied in addition to those described above. The first of these is that the flow is assumed to be two-dimensional in x and z , and the mountain ridges are assumed to be narrow, so that the Rossby number is large and Coriolis forces may be neglected, meaning that the flow is truly two-dimensional in x and z . Secondly, the flow is assumed to be in a steady state, so there is no time dependence. Next, the mountain ridges are assumed to be of small height, so that the Froude number is large and the flow is linear. The wind components are split into a background profile (dependent on height only) and a perturbation, in the same way as the pressure, potential temperature and density were treated for the Boussinesq approximation. As before, these perturbations are not turbulent fluctuations; the flow is assumed to be laminar for this example. The background horizontal wind, u_0 , is assumed to be normal to the ridge, and the background vertical wind, w_0 , is assumed to be zero. Here, the static stability is given

by $N^2 = \frac{g}{\theta_0} \frac{d\theta_0}{dz}$, which is a function of height only, since it depends on the background

potential temperature profile, itself a function of height only. The perturbation quantities are assumed to be small when compared to the background values, which permits linearisation of the equations. This assumption is usually valid as long as $H \ll L$, so that the mountain slope is small.

The following simple analytical example for two-dimensional steady flow is based on the working given in Durran (1986). Starting from the two-dimensional, steady state, laminar version of the equations derived at the end of the previous section, with the assumptions described above, linearising about the background wind profile gives

$$u_0 \frac{\partial u''}{\partial x} + w'' \frac{du_0}{dz} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial x}$$

$$u_0 \frac{\partial w''}{\partial x} = b'' - \frac{1}{\rho_0} \frac{\partial p''}{\partial z}$$

$$\frac{\partial u''}{\partial x} + \frac{\partial w''}{\partial z} = 0$$

$$u_0 \frac{\partial b''}{\partial x} + w'' N^2 = 0$$

where $b'' = \frac{g\theta''}{\theta_0}$ is the buoyancy.

Combining these equations gives

$$\frac{\partial^2 w''}{\partial x^2} + \frac{\partial^2 w''}{\partial z^2} + l^2 w'' = 0,$$

where $l^2 = \frac{N^2}{u_0^2} - \frac{1}{u_0} \frac{d^2 u_0}{dz^2}$ is the Scorer parameter (e.g. Scorer, 1949). This equation is

known as the Taylor-Goldstein equation, and has numerous variations depending on the simplifications made to the basic equations and to the background flow. It is obviously more complicated in the three-dimensional unsteady case (e.g. Coulter, 1998) but this simple example is sufficient to illustrate the point here.

For a background wind profile with constant shear, the horizontal Froude number,

defined above as $F_\alpha = \frac{Uk}{N}$, is the ratio of the Scorer parameter to the dominant

horizontal wavenumber. It also arises from the ratio of the advection terms to the buoyancy term in the vertical momentum equation, using a scale analysis similar to that carried out in section 3.2 for the u -momentum equation.

Analytical solutions exist for some special cases, but generally if the background wind and stability profiles vary with height, then the equation is impossible to solve analytically, so numerical methods must be used. In the case of cosine-squared ridges (sometimes known as ‘‘corrugated iron’’ ridges) with height $h(x) = h_m \cos(kx)$, and where the background wind and static stability are uniform with height, there is an

analytical solution, and this is given here. For this case the Scorer parameter is reduced

$$\text{to } l^2 = \frac{N^2}{u_0^2}.$$

In this simple case where the background wind speed and buoyancy frequency are independent of height, a dispersion relation can be easily found for the waves. The method used here follows Gossard and Hooke (1975), although the working here is in two dimensions rather than three, to fit with the examples given in this section, and with the work presented in the body of this thesis. In the unsteady version of the equations given above, each perturbation field ϕ'' , which represents any one of u'' , w'' , b'' or p'' , can be written as $\phi'' = \text{Re} \{ A \exp(i(kx + mz - \omega t)) \}$, where A is a complex constant, k is the horizontal wavenumber (defined as $\frac{2\pi}{\lambda_h}$ where λ_h is the horizontal wavelength) and m is the vertical wavenumber (defined similarly). The angular frequency, ω , is given by $\frac{2\pi}{T}$, where T here represents the period. Then the dispersion relation is given by

$$\hat{\omega}^2 = \frac{N^2 k^2}{k^2 + m^2}$$

where $\hat{\omega} = \omega - u_0 k$ is the intrinsic frequency, that is the frequency of waves observed from a frame of reference moving with the background fluid motion. The phase speed, c , is given by ω/k , and is the rate at which crests and troughs move in the x -direction. If the phase speed is independent of the wavenumber, then the propagation speed is independent of the wavelength, and a disturbance of arbitrary shape will move at the phase speed without change of shape. The waves are then non-dispersive. If the phase speed depends on the wavenumber, then an arbitrary disturbance will disperse: its shape will change with time, since different Fourier components travel at different speeds. The waves are then dispersive, and the group velocity gives the speed and direction of the envelope (shape) of the resultant wave.

For waves which have a steady forcing, such as the orographically generated waves considered here, $\omega = 0$, and the intrinsic frequency is given by

$$\hat{\omega} = \frac{-Nk}{\sqrt{k^2 + m^2}}$$

where the positive root has been discarded here, since $\hat{\omega} = -u_0 k$ and u_0 is positive in this example.

Rearranging this gives an expression for the vertical wavenumber:

$$m^2 = \frac{N^2 k^2}{\hat{\omega}^2} - k^2 = \frac{N^2 k^2}{u_0^2 k^2} - k^2 = l^2 - k^2.$$

This says that if $k^2 < l^2$ then the vertical wavenumber is real, and waves can propagate upwards. If $k^2 > l^2$, then the vertical wavenumber is imaginary, and the physically realistic solution is for disturbances to decay exponentially in the vertical direction, away from the source. The Taylor-Goldstein equation thus has two possible solutions, depending on the sign of $(l^2 - k^2)$.

When $l^2 > k^2$ the solution is $w''(x, z) = -u_0 h_m k \sin(kx + mz)$. The vertical wavenumber is real, so these long wavelength hydrostatic waves are able to propagate freely in the vertical. The solution for this case is long waves which are upwardly propagating, with phase lines slanting windward with height and energy transported upwards. These waves have non-zero vertical mean momentum flux and are shown in figure 3.1. They are favoured by strong stability, low wind speed and wide mountain ridges, such as the one in case 4, discussed in Chapter 5.




Figure 3.1. Upwardly propagating waves for the case where $l^2 > k^2$. After Durran (1986).

Waves in the atmosphere transport little heat, humidity and other scalars such as pollutants, but they are effective at transporting momentum and energy (Stull, 1988).

Under the above assumptions, the wave energy propagates with group velocity \underline{c}_g given

$$\text{by } \underline{c}_g = \left(\frac{\partial \omega}{\partial k}, \frac{\partial \omega}{\partial m} \right) = \left(u_0 - \frac{Nm^2}{(k^2 + m^2)^{\frac{3}{2}}}, \frac{Nkm}{(k^2 + m^2)^{\frac{3}{2}}} \right),$$

and hence the vertical component of the group velocity is upwards for the waves discussed here, with phase lines given by $kx + mz = \text{const.}$ The alternative solution, with phase lines given by $kx - mz = \text{const.}$, was rejected because the energy propagation is downwards in that case, towards the source, which is unphysical.

For $l^2 < k^2$ the solution is $w''(x, z) = -u_0 h_m k e^{-m_* z} \sin(kx)$ where m is now imaginary, and so m_* is defined as $m_* = im$ so that $m_*^2 = k^2 - l^2$ and m_* is real. The imaginary vertical wavenumber means that the waves cannot propagate freely in the vertical. The solution for this case is short waves whose amplitude decays with height, with vertical phase lines, and group velocity arguments show that the wave energy is transported downstream. These waves have zero vertical mean momentum flux and are shown in figure 3.2. These are known as evanescent waves, because the wave amplitude decays exponentially with height. They are likely to occur in weak stability or strong winds, over narrow mountain ridges.




Figure 3.2. Horizontally propagating waves for the case where $l^2 < k^2$. After Durran (1986).

It is also useful to construct a simple two layer model, where the background velocity and stability profiles have a two layer structure as shown in figure 3.3, so that the Scorer parameter decreases with height: $l^2 > k^2$ in the lower layer and $l^2 < k^2$ in the upper layer. This decrease in the Scorer parameter is a requirement for trapped lee waves, and in simple terms it can be achieved by strong wind shear or an inversion, both of which are common features in reported wave cases (Corby, 1954).




Figure 3.3. Potential temperature and wind profiles for the two layer case, showing the change in Scorer parameter with height. After Durran (1986). The left hand plot is a tephigram, which is a standard meteorological tool, with decreasing pressure (labelled in millibars, which are equivalent to hPa) on the left hand axis rather than increasing height. The dotted grid lines going from top left to bottom right are lines of constant potential temperature (isentropes or dry adiabats), labelled in Kelvin. The solid lines which are almost perpendicular to the isentropes are lines of constant temperature (isotherms), labelled in degrees centigrade. The dashed lines are the saturated adiabats (the same as dry adiabats, but for saturated air) and are not relevant for the dry cases considered here. Scorer (1978) gives a fuller description of the tephigram's construction and main uses.

In the upper layer, the wave amplitude decays with height. In the lower layer there are superimposed upwardly and downwardly propagating waves, where the latter are caused by reflection of the upwardly propagating waves at the layer boundary. This results in vertical phase lines, and the wave energy is trapped in the lower layer. This wave is known as a trapped lee wave, and in this idealised case, with no viscosity, it extends an infinite distance downstream. The wave pattern is shown in figure 3.4.




Figure 3.4. Wave pattern for the two layer case, showing the trapped lee wave at lower levels. After Durran (1986).

Because of the necessary assumptions and approximations involved, analytical methods produce conceptually simple results such as those described in this section. These are generally straightforward to interpret, because the set of parameters used in the simplified equations is relatively small. However, the obvious limitation is that these equations can only describe a simplified system, and cannot represent the more complicated, detailed aspects of the flow. Although analytical results can be obtained for more complicated cases (sometimes involving numerical integration to obtain solutions), more accurate modelling requires the solution of a more detailed non-linear set of equations, with no known analytical solution. In order to model these more complex cases, numerical simulation is required. The next chapter describes the numerical model used in this thesis.

3.6 Momentum fluxes and drags

As mentioned earlier, orography exerts a drag on the flow which can be greater than the drag which would be present if the terrain was flat. A variety of methods have been used to investigate the drag over mountains and hills, including modelling work (e.g. Wood and Mason, 1993), linear analysis (e.g. Belcher *et al.*, 1993), towing tank experiments (e.g. Davis, 1969 and Castro *et al.*, 1990) and field measurements (e.g. Smith, 1978 and Vosper and Mobbs, 1997). However, more work is required in this area in order to gain a better understanding of atmospheric drag, and to improve parametrization schemes.

Orographic drag can be conveniently divided into two distinct parts, both of which result from perturbations to the surface pressure field. These are the drag due to turbulent non-wavy flow, known as the form drag, and the drag due to orographic waves, known as the wave drag. The mathematical detail of the form drag is not included here, since the simulations in this thesis do not include a boundary layer or turbulence, but it can be complicated, depending on wind shear in the upstream profile (Belcher *et al.*, 1993) and on the mountain shape (Wood and Mason, 1993). In addition, it has been shown to be sensitive to the closure scheme used in numerical models (Sykes, 1980 and Wood and Mason, 1993).

Mason (1985) and Wood and Mason (1993) suggest that the form drag due to the SGSO can be parametrized in large scale models via an effective roughness length for momentum, which includes form drag effects. Xu and Taylor (1995) review several proposals for parametrization of orographic form drag in turbulent flow over hills, including the significant impact of the turbulence closure scheme.

When waves are present, the perturbation pressure lies out of phase with the mountain, and linear inviscid theory can give a good estimate of the wave drag. However, it does not include form drag, so it will under-predict the total drag when no or few waves are

present. In order to investigate how the drag varies with the flow and the orography, it is essential to understand how the perturbation pressure behaves under these conditions. Miranda and James (1992) used a three-dimensional numerical model to examine wave breaking and flow splitting with a bell-shaped mountain at low Froude number, and found that for Froude numbers greater than about 2, the drags predicted by the linear and non-linear theory were very similar. For Froude numbers less than 2, non-linear effects became important, and the surface drag differed from that predicted by linear theory. Vosper (1995) also found that a linear model could accurately describe the wave field for Froude numbers greater than about 2, and that the linear drag was progressively smaller than the non-linear drag as the Froude number decreased. This indicates that linear theory is not sufficient to predict the drag in more complicated flows.

There is still some discussion as to when the form drag is likely to dominate the wave drag, and vice versa. Belcher and Wood (1996) conducted experiments with the numerical model used in this thesis, for varying stratification, and found that the relative magnitudes of these drags were strongly dependent on the mountain shape and on the boundary layer depth. In the laminar simulations presented here, with no boundary layer, only the wave drag is significant.

The linearised u -momentum equation for laminar, two-dimensional steady flow is, as derived in the previous section,

$$u_0 \frac{\partial u''}{\partial x} + w'' \frac{du_0}{dz} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial x}.$$

Following Bretherton (1969), the perturbation quantities are assumed to vanish at $x = \pm\infty$, to permit integration by parts. Then if the surface is given by $z = h(x)$, the

horizontal force per unit area exerted on it is given by $p'' \frac{dh}{dx}$. Averaging this

expression over an infinite horizontal domain, the mean horizontal force per unit area exerted on the surface by the wave-induced pressure perturbations (known as the wave drag and denoted by F) is given by

$$F = \int_{-\infty}^{\infty} p'' \frac{dh}{dx} dx = \int_{-\infty}^{\infty} -h \frac{\partial p''}{\partial x} dx.$$

Multiplying the u -momentum equation by $\rho_0 h$ and substituting for the wave drag, F , gives

$$F = \rho_0 u_0 \int_{-\infty}^{\infty} h \frac{\partial u''}{\partial x} dx + \rho_0 \frac{du_0}{dz} \int_{-\infty}^{\infty} w'' h dx .$$

At the surface, $w'' = u_0 \frac{dh}{dx}$, so the second term vanishes, and a further integration by parts gives

$$F = -\rho_0 \int_{-\infty}^{\infty} u'' w'' dx = -\rho_0 \overline{u'' w''}$$

where the bar here indicates the integral over an infinite horizontal domain. Note that the term $-\rho_0 \overline{u'' w''}$ is not the same as the Reynolds stress, $\tau_{13} = -\rho \overline{u' w'}$, which was defined in section 3.3. That term involved the turbulent velocity fluctuations, and the bar did not necessarily represent a spatial average. The right hand side of this equation is the average vertical flux of horizontal momentum, so this equation says that the horizontal force on the surface is equal to the momentum transfer across that surface.

Equating the two expressions for F gives $\overline{p'' \frac{dh}{dx}} = -\rho_0 \overline{u'' w''}$, which simply states

Newton's third law: that the drag exerted on the orography must be of the same magnitude to and in the opposite direction from the force exerted on the atmosphere by the orography. This equation implies that one way of measuring the force exerted on the atmosphere is to directly measure the pressure perturbations at ground level, and this method has been used with some success. Vosper (1995) reports on drag measurements carried out this way, and gives an overview of work done by others using this method.

In parametrization terms, u_0 represents the large scale (resolved) horizontal velocity, and u'', w'' the sub-grid scale motions. In an NWP model, the large scale flow is of course time-dependent, although the analysis here is steady state. Following Mobbs (1994), if a thin layer of fluid lies between two material surfaces which are a distance δz apart, where the force exerted on the lower surface is F and that on the upper surface is $F + \delta F$, then the resultant force per unit area on this layer of fluid is

$$(F + \delta F) - F = \left(F + \frac{\partial F}{\partial z} \delta z \right) - F = -\frac{d}{dz} \left(\rho_0 \overline{u'' w''} \right) \delta z .$$

The mass of the layer per unit area is $\rho_0 \delta z$, and so if $\frac{\partial u_0}{\partial t}$ is the average acceleration per unit area of the layer of fluid, then the equation of motion for the layer is

$$\rho_0 \frac{\partial u_0}{\partial t} = -\frac{d}{dz} \left(\rho_0 \overline{u''w''} \right).$$

Eliassen and Palm (1960) showed that for a two-dimensional steady wave field in the absence of dissipation, and where the perturbation quantities decay to zero in the far field, $\frac{d}{dz} \left(\rho_0 \overline{u''w''} \right) = 0$, except at levels where the mean wind is zero (critical levels).

This states that the wave stress only causes a drag on the atmosphere (thus transferring momentum to the mean flow) near critical levels or where dissipative forces act. Broad (1995) extended this result to three dimensions, with further work by Vosper and Mobbs (1998). Shutts (1995) pointed out that if the wind vector turns with height, the momentum flux ceases to be independent with height because the critical level absorption occurs at every height for some part of the wave spectrum. The above equation states that in layers where the average vertical flux of horizontal momentum is divergent, there is a net transfer of momentum between the waves and the mean flow, and thus the mean flow is accelerated (or decelerated) by the orography. This is the equation used in most GWD parametrization schemes, and it is discussed further in Chapter 4.

Trapped waves are excluded from the above analysis, because the assumption that perturbations vanish far upstream and downstream of the mountain is no longer correct. In cases where significant trapped lee wave activity is present, there is a noticeable decrease in the magnitude of the momentum flux with height in the low to mid troposphere. Bretherton (1969) carried out linear analytical investigations into the contribution made by trapped waves to the vertical flux of horizontal momentum, and found that the lee wave drag was much less than the drag induced by upwardly propagating waves. Durran (1995), however, found that trapped lee wave drag can be significant, with magnitude which may be equal to or greater than to the drag due to upwardly propagating waves, and these effects are now incorporated in some GWD parametrization schemes (e.g. Gregory *et al.*, 1998). Trapped waves are a potentially important source of orographic wave drag, and the uncertainty remains as to where and how the waves are dissipated, and therefore where they deposit momentum.

4 The numerical model

4.1 Introduction

The BLASIUS model (Boundary Layer Above Stationary, Inhomogeneous Uneven Surfaces) was developed in the Atmospheric Processes Research branch of The Met. Office. It can be used to simulate either laminar or turbulent air flow, over a flat surface or over topography of moderate slope. The model is described in more detail by Wood (1992), Wood and Mason (1993) and Wood *et al.* (1999).

In the cases studied here, the stratification prevented a boundary layer from growing into the model domain, even though the surface boundary condition remained no slip for all the velocity components. A quasi-free slip lower boundary condition was tested on the cases described in the next chapter, but although this solved a few of the minor problems present in the no slip simulations, the results were generally discouraging. Section 8.2 gives more details of these tests. Due to the nature of the numerical model, truly inviscid simulations were not possible, so a very low viscosity was used ($\nu = 10^{-33} m^2 s^{-1}$). This enabled the no slip condition to be used in the work presented here.

BLASIUS was used initially to study neutral flow over hills (e.g. Wood and Mason, 1993 or Wood, 1992). The simulations described here aim to reproduce flows on a larger scale than most boundary layer studies, so the model is being used for a higher mountain and larger computational domain than in most previous work. In addition, only limited investigations had previously been carried out into the model's capability

to simulate orographic gravity waves, so the work presented here tests this more thoroughly with a view to producing guidelines on dealing with the numerical issues involved. An important aspect of this is the formulation of the upstream and downstream boundary conditions, which is discussed further below.

4.2 Model equations and initialisation

BLASIUS is a non-linear, three-dimensional, finite difference model which solves the ensemble-averaged time-dependent Navier-Stokes equations, simplified by applying either the anelastic or the Boussinesq approximation. For generality, BLASIUS is coded using the anelastic version of the governing equations, which allows for significant variation of background density with height over the depth of the domain. However, all the simulations presented in this thesis used the Boussinesq configuration of BLASIUS, which is described here. The starting point is the equations derived at the end of section 3.4; the assumptions and approximations described during that derivation will not be repeated here. The full three-dimensional equations are given here, but all the simulations discussed here are two-dimensional cases, where the background velocity is normal to the ridge. Taking ϕ to represent any field, then ϕ_0 is always a function of height only, but all the other model fields are functions of x , y , z and t unless stated otherwise.

The equations used in the model are those obtained at the end of section 3.4:

$$\begin{aligned} \frac{Du}{Dt} &= -\frac{1}{\rho_0} \frac{\partial p''}{\partial x} + fv + \frac{1}{\rho_0} \left[\frac{\partial \tau_{11}}{\partial x} + \frac{\partial \tau_{12}}{\partial y} + \frac{\partial \tau_{13}}{\partial z} \right] \\ \frac{Dv}{Dt} &= -\frac{1}{\rho_0} \frac{\partial p''}{\partial y} - fu + \frac{1}{\rho_0} \left[\frac{\partial \tau_{21}}{\partial x} + \frac{\partial \tau_{22}}{\partial y} + \frac{\partial \tau_{23}}{\partial z} \right] \\ \frac{Dw}{Dt} &= -\frac{1}{\rho_0} \frac{\partial p''}{\partial z} + \frac{g\theta''}{\theta_0} + \frac{1}{\rho_0} \left[\frac{\partial \tau_{31}}{\partial x} + \frac{\partial \tau_{32}}{\partial y} + \frac{\partial \tau_{33}}{\partial z} \right] \\ \nabla \cdot \underline{u} &= 0 \\ \frac{D\theta}{Dt} &= \frac{1}{\rho_0} \left[\frac{\partial H_1}{\partial x} + \frac{\partial H_2}{\partial y} + \frac{\partial H_3}{\partial z} \right] \end{aligned}$$

where BLASIUS assumes that the variables are ensemble averages, and takes the Coriolis parameter as constant.

There is a choice of two turbulence models in BLASIUS. The one used here is a first order mixing length closure, where the mixing length is representative of the scale of individual eddy motions in the boundary layer; the other is a 1.5 order closure, which carries the turbulent kinetic energy as a prognostic variable. A second order scheme would carry all six stress components, and would be expensive in terms of computation time and storage requirements. Although BLASIUS was designed to simulate boundary layer flows, the work described in this thesis simulates laminar flows, without a boundary layer. Because of this, the detail of the closure schemes is not given here, but it can be found in Wood *et al.* (1999), while Stull (1988) and Garratt (1992) provide a general overview of the various options.

The turbulent momentum and heat fluxes are modelled using a generalisation of the gradient transport theory, as discussed in section 3.3. Details of the theory are given by Smagorinsky (1963), among others. The momentum fluxes are modelled as

$$\tau_{ij} \equiv -\rho_0 \overline{u'_i u'_j} = \rho_0 K_m \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$$

and the normalised temperature fluxes are modelled as

$$H_i \equiv -\rho_0 \overline{u'_i T'} = \rho_0 K_h \frac{\partial T}{\partial x_i}$$

where K_m is the momentum diffusivity and K_h is the heat diffusivity. The model carries a normalised temperature variable, T , defined as

$$T = \frac{g(\theta - \theta_r)}{\theta_r}$$

and so the normalised turbulent temperature fluxes are given by

$$H_i = \rho_0 K_h \frac{\partial}{\partial x_i} \left(\frac{g(\theta - \theta_r)}{\theta_r} \right).$$

The reference temperature, θ_r , has an arbitrary value, but it is required in order to convert physical quantities such as surface temperature into the corresponding normalised variables. The default value of 288.15 K was used in all the simulations presented here.

The Brunt-Väisälä frequency profile, N , is input into the model and used to calculate the initial potential temperature profile, θ_{mit} . In the simulations used here, it is either constant with height throughout the whole domain, or has a two layer profile; more details are given in the next chapter. It is defined as

$$N^2 = \frac{g}{\theta} \frac{d\theta}{dz}$$

which gives the following equation for potential temperature, assuming that N is constant within a layer:

$$\theta_{mit} = \theta_r e^{\frac{N^2 z}{g}}$$

BLASIUS uses a first order approximation to this, defined as

$$\theta_{mit} = \theta_r \left(1 + \frac{N^2 z}{g} \right),$$

which is equivalent to the approximation of defining $N^2 = \frac{g}{\theta_r} \frac{d\theta_{mit}}{dz}$. This is acceptable

in boundary layer simulations which have a relatively small domain depth, but not really suitable for the simulations described here. The scale height for a typical simulation with $N = 0.02 \text{ s}^{-1}$ would be

$$H_\theta = \left(\frac{1}{\theta} \frac{d\theta}{dz} \right)^{-1} = \frac{g}{N^2} = \frac{10 \text{ ms}^{-2}}{0.0004 \text{ s}^{-2}} = 25 \text{ km}$$

whereas these simulations have a vertical domain height between 35 and 60 km (not including the Rayleigh damping layer, which is explained later), so some improvement is needed here. Preliminary tests of an exponential potential temperature profile have been carried out, but further testing would be required before implementing this.

Moisture can also be included in the model, but this option has not been used in any of the work described in this thesis, so it is not covered here. Details are given in the references at the start of this chapter.

BLASIUS is initialised with input profiles of horizontal velocity and Brunt-Väisälä frequency, as well as other physical and numerical parameters. A one-dimensional simulation is run in order to grow a boundary layer into the profiles, and these new profiles are then used to initialise a two- or three-dimensional simulation. In the cases

described here, without a boundary layer, the one-dimensional simulation was run for a nominal length of time (one second), and the original input profiles were used to initialise the two-dimensional simulations. The one-dimensional simulation had the same physical parameters as the two-dimensional simulation, but a flat surface.

In a two- or three-dimensional simulation, the initial domain is flat, with the input profiles of velocity and potential temperature used throughout. The mountain is then gradually grown into the domain, and this process is discussed further in section 8.4, which describes investigations into the sensitivity of results to the speed of mountain growth. The gradual introduction of the mountain is designed to minimise spurious noise in the initial fields; other methods of doing this include increasing the horizontal velocity gradually from zero to its initial value, and gradually increasing the gravitational acceleration from zero to its constant value (e.g. Durran, 1995). The model surface is smoothed to remove any large discontinuities in the initial profile, because smooth terrain helps the pressure solver to converge more rapidly. This in turn makes the model run more quickly, since the majority of its time is spent in the pressure solver. The pressure equation is solved in BLASIUS by transforming in both horizontal directions to give a second order differential equation, which is then solved by matrix inversion. There are two complications to be overcome when doing this. The first is that the co-ordinate transformation introduces extra terms into the source term which are dependent on the pressure field itself. These terms are described fully by Wood (1992). This problem is solved by iterating for the pressure until the maximum absolute difference between the new pressure field and the one obtained at the previous iteration is less than a prescribed tolerance value, chosen by the user. The accuracy of convergence in the pressure solver is measured by how accurately the model satisfies the continuity equation, which is an indicator of the overall accuracy of the model's numerical scheme, since it is effectively a conservation of mass condition applied over the entire numerical domain. A steeper slope means that the pressure solver takes longer, so accuracy must be sacrificed in order to reduce computational cost. BLASIUS is thus suited to simulations of topography with moderate slope. The second problem is that Fast Fourier Transforms cannot be used for a non-uniform grid, so the method described by Farnell (1980), which iterates for eigenvalues and eigenvectors, is used to overcome this. Farnell (1977) gives details of the theory involved, and complications

due to the horizontal boundary conditions and symmetry of the grids are covered in more detail by Farnell (1980).

The model uses a curvilinear terrain-following co-ordinate system, so that the surface is a co-ordinate surface. The transformation of the vertical co-ordinate is

$$z = \frac{\tilde{z} - Z_s(x, y)}{\left(1 - \frac{Z_s(x, y)}{H_D}\right)}$$

where z is the transformed co-ordinate, \tilde{z} is the original Cartesian co-ordinate, the surface is denoted by $Z_s(x, y)$, and H_D is the depth of the model domain. This means that the model surfaces are terrain-following near the surface, becoming flat (coincident with horizontal Cartesian surfaces) at upper levels, far from the surface, and mirrors what is done in many NWP models, simulating the fact that air motion tends to be terrain-following at low levels, but less so with increasing height. Uniform grids can be used, or stretching can be applied in one or more directions. More details of the stretching methods for both horizontal and vertical grids are given in Chapter 6. Figure 5.6 in the following chapter gives an example of the model grid, showing the bent mesh, and stretching in the horizontal; figure 6.6 shows a grid which is stretched in the vertical as well as the horizontal. It should be noted that a field which depends only on height in Cartesian co-ordinates becomes a function of all three co-ordinates on the transformed grid, and that only the co-ordinates are transformed, not the velocity components, whose directions remain horizontal and vertical.

The model grids are staggered using an Arakawa C grid in the horizontal and a Lorenz grid in the vertical, so the variables are held on different points, and this reduces interpolation and simplifies the computations involved. The grid is shown in figure 4.1. Capital letters U , V , W and P are used to represent the discretised model fields corresponding to the continuous variables u , v , w and p'' in the equations. Pressure is held on 'P-points', and $P(I, J, K)$ represents the pressure at the point $(XN(I), YN(J), ZN(K))$. All scalars and normal components of stresses are stored on these P-points, and other fields are held on their natural points. For example, the horizontal velocity is held on points which are horizontally half way between the P-points, because its momentum equation involves the horizontal gradient of pressure, so $U(I, J, K)$

represents the U -velocity at the point $(X(I), YN(J), ZN(K))$. Similarly the vertical velocity is forward-staggered by a half-gridpoint from the main grid where the pressure is held, so $W(I, J, K)$ represents the W -velocity at the point $(XN(I), YN(J), Z(K))$. All the other stress components are held on their natural points, so for example

$\tau_{13} = \rho_0 K_m \left(\frac{\partial u_1}{\partial x_3} + \frac{\partial u_3}{\partial x_1} \right)$ is represented by $TAU13(I, J, K)$ and held at the point

$(X(I), YN(J), Z(K))$. Wood (1992) and Wood *et al.* (1999) give more detail of the grid, including the transformed equations.

The grid staggering means that not all fields are held at the physical surface. Figure 4.2 shows the difference in the height of the lowest model level on Z levels and ZN levels, for a 500 m high mountain. The former are referred to in this thesis as full levels, since they include the physical surface as the lowest model level, labelled $Z(I)$. The latter are referred to as half levels, because they lie half way between the full levels, and here the lowest level, $ZN(I)$, is below the physical surface. This should be borne in mind when viewing plots of potential temperature later in this thesis, since this field is held on half levels, and so the ‘surface’ shown in those plots is actually below the physical surface, meaning that the mountain appears to be smaller than it really is, which can be misleading. Vertical velocity is held on full levels, so plots of that field show the true physical surface and the correct mountain height.


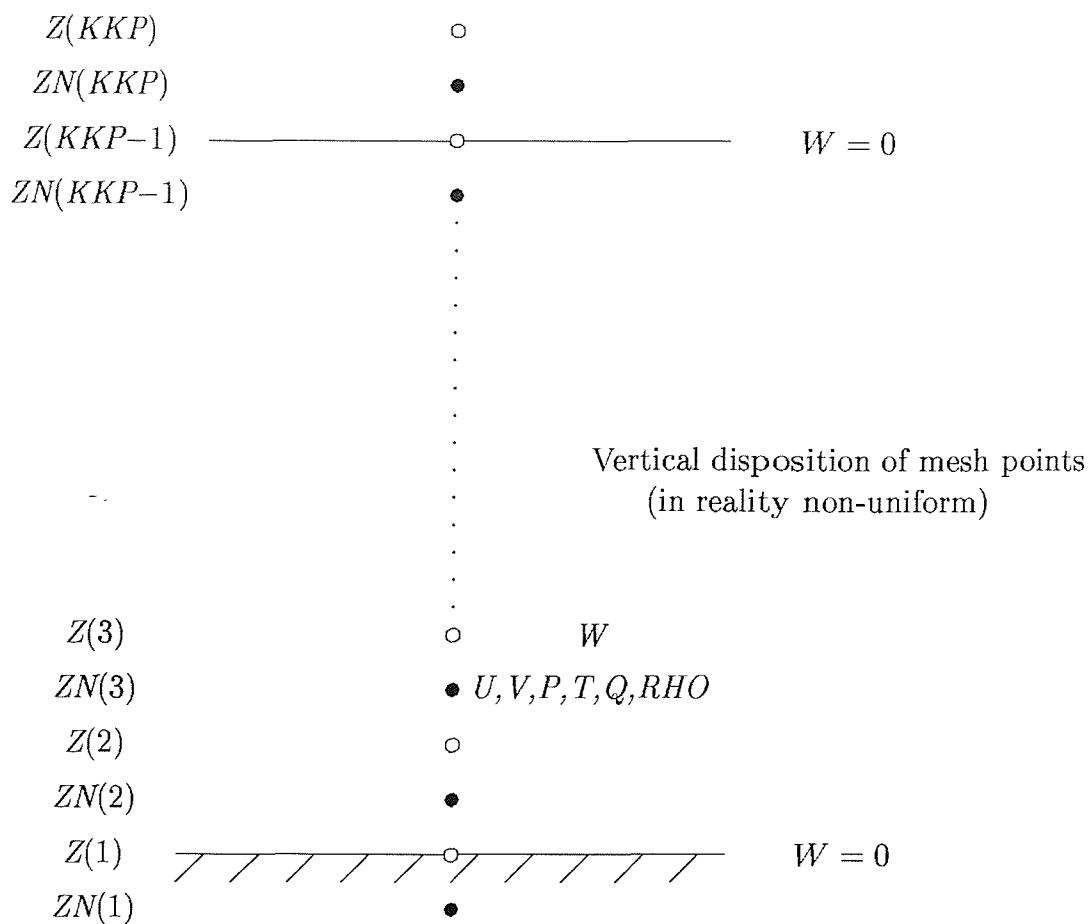


Figure 4.1(a). The upper diagram shows schematically the vertical grid used in BLASIUS, illustrating how the model fields are held on different levels. The shading represents the physical surface, and the solid line above represents the top of the numerical domain; the additional points are used for the boundary conditions. Q represents the moisture variable, which is not covered in this thesis; RHO represents the density. KKP is the number of vertical gridpoints in the model domain. The lower diagram illustrates the three-dimensional mesh arrangement. From Wood *et al.* (1999).





Figure 4.1(b). Schematic diagram of the BLASIUS grid at the upstream and downstream edges of the numerical domain, showing how the model fields are held on different points. The shading represents the edge of the numerical domain; the additional points are necessary for the boundary conditions. *IIP* is the number of horizontal gridpoints in the *x*-direction in the model domain; *JJP* is defined similarly in the *y*-direction. From Wood *et al.* (1999).




Figure 4.2. Lowest model level surfaces on Z levels (full model levels, shown by a solid line) and ZN levels (model half levels, shown by a dashed line), for a 500 m high mountain. The dotted line indicates the zero level. As explained in the text, the lowest full model level surface (e.g. where vertical velocity is carried) is the mountain surface, and the lowest half level surface (where e.g. pressure, temperature and horizontal velocity are carried) is below the mountain surface.

For the advection of momentum, the centred in time, absolutely conserving scheme of Piacsek and Williams (1970) is used, which conserves quadratic quantities (such as kinetic energy), whereas the other two schemes described below only conserve linear ones. This scheme is second order accurate and is known as the centred second order (CSO) scheme. For advection of the scalar quantities, BLASIUS gives the choice between this CSO scheme or two others, and the scheme chosen here is also used for temperature. The first alternative is the total variance diminishing (TVD) scheme described by Bull (1990), which is second order accurate where the solution is smooth, and first order accurate elsewhere. The second alternative is the third order accurate ultimate quickest (UQ) scheme described by Leonard *et al.* (1993). These three schemes were tested on the four cases described in Chapter 5, and the CSO scheme was preferred. Although the model fields were slightly smoother with the other two schemes, they both diffused the wave motion by an unacceptably large amount, and destabilised the solution. For case 4 (the case with the widest mountain: see Chapter 5 for a full description of all the cases) the CSO simulation was the only one that worked

at all, with the other two failing in the pressure solver right at the beginning (because the different advection scheme means that different coefficients go into the pressure solver matrix). These results are counter-intuitive, because these more diffusive schemes would be expected to act as smoothers, producing a more stable solution. Since the CSO scheme produced acceptable results, the problems seen with the other two schemes were not thoroughly investigated as part of this work. However, they may be explained by the instabilities which Bull (1990) discussed, where he concluded that the model should be set in a forward rather than a centred time differencing framework for best results. An alternative explanation has recently been suggested by Wood (private communication), who points out that the use of different schemes for the advection of momentum and of scalars means that in the cases studied here, diffusion is only applied to the temperature field. This could lead to a spurious phase shift of the temperature field with respect to the momentum fields, causing significant numerical distortion of the physical results.

The diffusive terms are solved using a centred in space, forward in time, explicit scheme. Section 8.1 gives more information on how the model timestep is determined during a simulation, including the adjustments made for numerical stability.

4.3 Upper and lower boundary conditions

The boundary condition on the velocity components at the surface is the no slip condition, as discussed in section 4.1 above. The vertical velocity, W , is held on full levels, so it is simply set to zero at the surface. For the horizontal velocity components, U and V , held at half levels above and below the surface, the boundary condition is $U(I,J,1) = -U(I,J,2)$ (and similarly for V) where $K = 1$ is the level just below the surface and $K = 2$ is the level just above. At the top of the domain the vertical component of velocity, W , is set to zero. For the horizontal components, in the simulations described here, the vertical derivative is set to zero at the upper boundary, so that $U(I,J,KKP) = U(I,J,KKP-1)$ (and similarly for V) where $K = KKP-1$ is the level

just below the top of the domain and $K = KKP$ is the level just above. The full and half levels were illustrated in figure 4.1.

There is a choice of lower and upper boundary conditions for the normalised model temperature. The simulations described in this document used a zero flux both at the surface and at the top of the domain.

In analytical work for large or infinite vertical domains, asymptotic boundary conditions are used far from the mountain. However, these are not appropriate for numerical modelling, so BLASIUS includes a Rayleigh damping layer at the top of the domain in these cases, which acts to prevent reflection of upwardly propagating wave energy at the upper boundary by the inclusion of terms of the form $-R(\phi - \phi_r)$ in the model equations, where ϕ_r is a reference state for the generic variable ϕ , and is taken to be the input profile. R is the damping coefficient, which has dimensions s^{-1} . Artificial absorbing layers such as this one have been used to good effect in several models, e.g. Durran and Klemp (1982), Miranda and James (1992). Reflection at the top of the domain can be avoided if any upwardly propagating waves are sufficiently damped by the absorbing layer, but the damping process itself may cause reflection, so the variables are gradually relaxed back to the reference state in this layer, over a time scale given by the reciprocal of R . The coefficient is input as a constant at the start of a simulation, and then used in the form

$$\frac{R(1 - \cos[z - z_b])}{2(z_t - z_b)}$$

where z is the height of the gridpoint in question, and z_b and z_t are the heights of the base and top of the damping layer respectively. This means that the damping effect increases from zero at the base of the layer to the full value of R at the top. The scheme needs tuning to minimise reflection from the base of the damping layer as well as from the top of the domain, and Chapter 6 describes the tests carried out.

4.4 Horizontal boundary conditions

Recent work at The Met. Office has introduced an option for inflow/outflow boundary conditions in BLASIUS, to allow simulation of cases with horizontally propagating waves forced by an isolated mountain. Previously the only option at the upstream and downstream edges of the domain was to use periodic boundary conditions, which are perfectly acceptable in a global model, but not always appropriate in a limited area model, particularly in cases with trapped lee waves, which propagate out of the downstream edge of the domain. Both options are discussed further in Chapter 7. All the simulations discussed in this thesis are two-dimensional x - z cases, so periodic boundary conditions were used in the y -direction throughout.

The inflow condition is relatively straightforward. A two- or three-dimensional simulation in BLASIUS is set up using initial profiles obtained from a one-dimensional simulation, as explained above. In the case of an inflow boundary, for the horizontal velocity components, these profiles are then simply fed in at the boundary for the entire simulation, representing a steady upstream flow. The same is true for passive scalars, whereas for the active field of temperature, a zero normal gradient condition is imposed. Pressure and vertical velocity have a zero normal gradient condition imposed.

There are three options for the outflow boundary condition: zero normal gradient, zero normal second derivative, or radiative. The latter was used for the simulations described in Chapter 5, so it is explained here; details of the other two boundary conditions are contained in Guilbaud (1997) and Wood *et al.* (1999). The aim of the radiative boundary condition is to allow disturbances to pass through the boundary without reflection.

The formulation of the radiative boundary condition calculates the outflow phase speed at each model level following the Orlanski (1976) scheme. This is based on the wave equation,

$$\frac{\partial \phi}{\partial t} + c \frac{\partial \phi}{\partial x} = 0,$$

where ϕ is any variable and c is the phase speed of the waves. The details of how this equation is manipulated to give a finite difference equation for the outflow phase speed are given by Orlanski (1976), and this equation has been used in several models (see e.g. Durran and Klemp, 1982). Raymond and Kuo (1984) extended Orlanski's work to include the derivative of ϕ in all three dimensions, rather than simply in the x -direction. For the two-dimensional test cases used, they found that their method was clearly superior to the various formulations of the one-dimensional version, as used here. Miranda and James (1992) used this three-dimensional scheme and found it to be very effective in preventing reflection at the boundaries. This method could be worth testing in BLASIUS, but would obviously increase the computational expense of the model.

The equation for the downstream radiative boundary condition in BLASIUS is

$$\phi_b^{n+1} = \frac{\phi_b^{n-1}(1-r_\phi) + 2r_\phi\phi_{b-1}^n}{(1+r_\phi)}$$

where r_ϕ is the non-dimensionalised outflow phase speed, obtained by applying this equation to interior gridpoints at preceding times:

$$r_\phi = \frac{(\phi_{b-1}^{n-2} - \phi_{b-1}^n)}{(\phi_{b-1}^n + \phi_{b-1}^{n-2} - 2\phi_{b-2}^{n-1})}.$$

The subscripts indicate the horizontal gridpoint where the variable is held: b represents a boundary value, and subscripts less than b represent interior values. Similarly the superscripts indicate the timestep used, where n is the current timestep. The radiative boundary condition is only applied to the normal velocity, vertical velocity and temperature; a zero normal gradient condition is imposed for the other variables.

Further information on the implementation of the scheme in BLASIUS is given by Wood *et al.* (1999), and can also be found in Guilbaud (1997), along with details of tests carried out for this and the other two outflow boundary conditions mentioned above. Orlanski (1976) showed analytically that the boundary condition gives no reflection for a single wave component, and also tested it for numerical simulations of non-linear cases where there are different components. Further tests of the radiative

boundary condition have been carried out using BLASIUS as part of this project, and these are discussed in Chapter 7.

The Orlanski boundary condition imposes a limit on the outflow phase speed for numerical stability, so that it lies between zero and $\frac{\Delta X}{\Delta t}$, where ΔX is the horizontal grid spacing and Δt is the timestep. The lower limit is imposed to ensure that the disturbance flows out of the domain, and the upper one arises from the CFL criterion, named after Courant, Friedrichs and Lewy (1928), and discussed further in section 8.1. The phase speed is adjusted as follows:

$$c = \begin{cases} 0 \\ c \\ \frac{\Delta X}{\Delta t} \end{cases} \text{ for } \begin{cases} c \leq 0 \\ 0 \leq c \leq \frac{\Delta X}{\Delta t} \\ \frac{\Delta X}{\Delta t} \leq c \end{cases}.$$

As part of the testing carried out in this work, the effect of this clipping was investigated to see whether it was linked to some of the problems experienced at the boundary.

Further details of these problems, and of the tests, are contained in Chapter 7. Other formulations of the outflow boundary were tested, but all of these introduced problems in at least one of the cases studied, so the default scheme was retained.

After imposing this stability limit on the phase speed, BLASIUS produces a single, vertically weighted average phase speed which replaces c in the above equation for subsequent model calculations. This vertically constant value, \bar{c} , is defined by considering the vertically integrated wave equation,

$$\int \frac{\partial \phi}{\partial t} + \int c \frac{\partial \phi}{\partial x} = 0.$$

The average phase speed should ideally satisfy the equation

$$\int \frac{\partial \phi}{\partial t} + \bar{c} \int \frac{\partial \phi}{\partial x} = 0,$$

and comparing these two versions of the wave equation leads to the definition of \bar{c} as

$$\bar{c} = \frac{\int c \frac{\partial \phi}{\partial x} dz}{\int \frac{\partial \phi}{\partial x} dz}.$$

The tests described in Chapter 7 included a modification where the outflow phase speed retained its vertical variation, but this was rejected.

The outflow phase speed obtained above is used to calculate the outflow velocities at the next timestep. A mass flux correction is then applied to these outflow velocities, in order to conserve mass flowing into and out of the model domain. As an additional test of the outflow boundary scheme, a new phase speed was calculated from the adjusted outflow velocities for all the cases in this work, and this was compared with the original phase speed. A large difference between these two values would indicate a problem with the implementation of the outflow boundary condition. Only very minor differences were found when using the original scheme, but larger differences were found with some of the alternatives described in Chapter 7.

4.5 Momentum fluxes and drags

The surface pressure force described in the previous chapter is diagnosed in BLASIUS. For the two-dimensional x - z case, the surface pressure force, F_x , averaged over the full horizontal domain per unit length, is given by

$$F_x = \frac{1}{L_x} \int_{x_-}^{x_+} p'' \frac{dh}{dx} dx$$

where the domain is from x_+ to x_- , L_x is the width of the domain, given by $x_+ - x_-$, and the rest of the notation is as previously defined. It should be noted that the surface pressure force is only non-zero over the mountain, since $\frac{dh}{dx}$ is zero where the surface is flat.

This means that averaging over the full horizontal domain is the same as averaging over the mountain, except that the diagnosed surface pressure force in BLASIUS is divided by the domain width rather than by the mountain width.

In the equation given above for the surface pressure force calculated in BLASIUS, F_x is the surface pressure force per unit length in the x -direction, so it must be multiplied by the domain width to give the drag force for the full domain, denoted by F_D . The non-dimensional drag coefficient, C_D , is defined in terms of the drag force, F_D , as

$$C_D = \frac{F_D}{\frac{1}{2} \rho_0 U^2 h_m}$$

where C_D is dependent on the obstacle shape and on the Reynolds number (e.g. Baines, 1995 or Smith, 1978) and h_m is the maximum mountain height. For the simulations considered here, the density is unity everywhere, and U is 10 m s^{-1} at the surface (or slightly larger at the height of the mountain crest in cases 2 to 4, but this difference is not significant here; more details are given in the next chapter).

Another important diagnostic to be considered is the horizontally averaged vertical flux of horizontal momentum, M , which was defined in the previous chapter as

$$M = -\rho_0 \overline{u''w''}$$

where the bar indicates a horizontal average. Wave momentum flux diagnostics have been calculated from previous control simulations performed at The Met. Office using the BLASIUS model, without a boundary layer, for flow over an isolated two-dimensional ridge. Post-processing to obtain the momentum flux defined above was coded by Guilbaud (private communication), with the horizontal averages computed over the full model domain using linear interpolation. As part of this project, these calculations have been incorporated into the main code, and all horizontal averaging is now done in Cartesian co-ordinates, rather than in model co-ordinates or in a mixture of the two. Guilbaud's work included calculations of the momentum flux for sub-domains (limited area, upstream, above the mountain, downstream), but these were based on a specific horizontal grid, and have not yet been fully incorporated into the main code for the more general case. The limited area sub-domain removes a portion at the upstream and downstream edges of the domain, thus reducing the impact of any errors in the numerical formulation of the inflow/outflow boundary conditions.

In addition, further diagnostics have been coded as part of this project in order to investigate all terms in the full two-dimensional u -momentum budget. The u -momentum equation used in BLASIUS in two dimensions is

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + w \frac{\partial u}{\partial z} = -\frac{1}{\rho_0} \frac{\partial p''}{\partial x} + \frac{1}{\rho_0} \left[\frac{\partial \tau_{11}}{\partial x} + \frac{\partial \tau_{13}}{\partial z} \right]$$

and the continuity equation used in BLASIUS in two dimensions is

$$\frac{\partial u}{\partial x} + \frac{\partial w}{\partial z} = 0.$$

The velocities are written as the sum of a background component (a horizontal average over any chosen domain from x_- to x_+ , indicated by a bar) and a perturbation (indicated by a double prime), so that

$$u = \bar{u}(z, t) + u''(x, z, t)$$

$$w = \bar{w}(z, t) + w''(x, z, t)$$

It should be noted that this is not the same as the decomposition into mean and turbulent parts in section 3.3, nor into background and perturbation parts in section 3.4, where the background profile was dependent on height only and was not a horizontal average. In analytical work with an infinite horizontal domain (and also in periodic model simulations) \bar{w} must be zero at all vertical levels, otherwise mass would accumulate. However, in the inflow-outflow simulations discussed here, it is possible to have a non-zero value, corresponding to a large scale circulation which is removed in order to allow parametrization of the smaller scale flow which remains.

Substituting these expansions into the two equations, after some algebra and horizontal averaging, gives the following equation, with each term labelled for future reference. L_x is the length of the chosen domain, given by $(x_+ - x_-)$. Some of the derivation is given by Durran (1991) with slightly different assumptions.

$$\frac{\partial}{\partial t} (\rho_0 \bar{u}) = -\frac{1}{L_x} [\rho_0 u^2 + p'']_{x_-}^{x_+} - \frac{\partial}{\partial z} [\rho_0 (\bar{u})(\bar{w})] - \frac{\partial}{\partial z} (\rho_0 \overline{u'w'}) + \frac{1}{L_x} [\tau_{11}]_{x_-}^{x_+} + \frac{\partial \tau_{13}}{\partial z}$$

A B C D E F

To minimise interpolation and differencing, which can amplify small errors, and to fit in with the rest of the code, the discretisation of this equation uses ‘ U -points’ on the model

grid, i.e. (X, YN, ZN) . This seems intuitively to be the most logical way to proceed, as it will most closely reflect the actual momentum budget obtained by the model. This could in principle be explicitly evaluated, but in bent co-ordinates this is a highly complex procedure. The validity of this decision to use U -points was verified by following through the analysis and discretisation methods for the basic model equations discussed earlier, and by experimenting with different gridpoints for the calculations presented here. The horizontal averaging is done on Cartesian levels which are at the same height as the half levels (ZN levels) over the flat parts of the domain, so interpolation is only required over the mountain itself. The lowest level used for averaging is equivalent to the first half level which lies above the height of the mountain crest, in order to avoid averaging through the mountain surface.

As discussed in section 3.6, terms A and D are the only ones which feature in the equations used by most gravity wave drag parametrizations, where term D, the vertical gradient of the horizontally averaged vertical flux of horizontal momentum, is diagnosed and used to derive term A, the rate of change of the horizontally averaged momentum. Term B is neglected under the assumption that perturbations should vanish at the edges of the domain, and term C is assumed to be identically zero because \overline{w} should be zero. Terms E and F are neglected under the assumption that the effects of the turbulent terms are small.

As discussed by Durran (1995), although the above assumptions are valid for the forcing of the globally averaged mean flow deceleration, where term A does generally mirror term D, not all of the other terms are necessarily negligible in analysing the results from high resolution limited area models such as BLASIUS. The perturbations do vanish at the edges of the domain in the periodic case, but not necessarily in the inflow/outflow case, where it is important to include term B in order to evaluate the local flow response correctly. As discussed above, \overline{w} is not always equal to zero in the simulations considered here, so term C makes a non-zero contribution. Term D was the one coded in BLASIUS by Guilbaud, and as part of this project the code has been modified to diagnose terms B, C, E and F in order to assess their relative importance in the calculation of term A. It would be useful to diagnose term A directly in order to validate the results obtained by this method, but this has not yet been done, and the code

for these new diagnostics is still being tested, so they are not used in the analysis of the results presented here. Typically, large increases are seen in the Rayleigh damping layer, but these can be ignored since unphysical dynamical changes are induced by the damping scheme, which slows down the flow, and changes the pressure field.

Preliminary results for cases 1 to 3 (see Chapter 5 for details of the test cases) are quite noisy, with the most pronounced noise in case 1 (with constant input wind speed and static stability), but this is unsurprising since there is no trapped lee wave and so the signal is expected to be small. This noise may be due to inadequate vertical resolution, so recent simulations with finer vertical resolution are in progress, and early results show that the u -momentum diagnostics have smaller amplitude oscillations, which is encouraging. Another potential explanation is that the noise may indicate a need to introduce vertical diffusion, so investigations will also be carried out into this aspect.

4.6 Artificially imposed horizontal diffusion

BLASIUS is designed to have minimal numerical diffusion in the model itself, because that would interfere with turbulent simulations. The centred second order scheme is absolutely conserving, so simulations like those discussed here should conserve quadratic quantities such as θ^2 . However, the centred timestep scheme means that the previous, current and next timesteps are all used, which can lead to numerical instability (computational modes), where the model flips between solutions from one timestep to the next. Time averaging is introduced to remove this, and in practice, this introduces some damping and numerical diffusion. In addition to this, previous work at The Met. Office has found that simulations without a boundary layer require some sort of artificially imposed horizontal diffusion in order to prevent the accumulation of energy at the grid scale (due to the downscale cascade of energy) and associated errors. Generally this is not required in boundary layer simulations as long as the horizontal grid spacing is small enough, but spurious noise was found in previous boundary layer simulations performed at The Met. Office with large horizontal grid spacing, and this was removed by introducing an artificial horizontal diffusion into the model, which is calculated at every timestep. Vichnevetsky (1987b) noted that imposed artificial

diffusion is significant only in motions with wavelengths which are smaller than the gridlength used in the model, and so it has little effect on the physical waves, damping out only the smaller scale numerical waves such as those caused by spurious reflection at the boundaries or as a result of a non-uniform grid. These reflections are discussed for the simulations studied here in Chapters 7 and 6 respectively.

In the vertical, there is more coherence and coupling than in the horizontal, and the resolution is generally better, so for a boundary layer simulation the effects of diffusion are larger in the vertical than in the horizontal. For this reason, the artificial diffusion is only required in the calculation of the horizontal fluxes of scalar quantities, and in the calculation of the horizontal components of the Reynolds stress tensor.

The aim of the artificial diffusion is to diffuse these grid scale signals before they are advected elsewhere in the domain, so the diffusion time scale needs to be smaller than the advection time scale, which gives the condition $\frac{(\Delta X)^2}{\nu} < \frac{\Delta X}{U}$, leading to $\frac{U\Delta X}{\nu} < 1$.

This is equivalent to requiring the computational cell Reynolds number, $\frac{U\Delta x}{\nu}$, to be less than unity. This non-dimensional parameter is also known as the grid Reynolds number or the mesh Reynolds number, and is denoted by Re_g . If it is too large, then the viscosity is too small, so there is not enough grid scale diffusion to damp out any spurious noise, which can lead to spatial instabilities. Roache (1972) discusses the implicit artificial diffusion introduced by some finite difference schemes, which is separate from the explicit, imposed artificial diffusion discussed here. The grid Reynolds number also appears in the analysis of implicit artificial diffusion, and in a later work, Roache (1976) suggests that the practical situation is not as restrictive as the formal accuracy requirement derived above, quoting analysis by Runchal and Wolfshtein (1969) which gave good results for maximum cell Reynolds numbers of approximately 20.

The scheme implemented in BLASIUS uses this value of 20 to provide a critical value, $Re_{g_{crit}}$, and applies the condition $\nu \geq \frac{U\Delta X}{Re_{g_{crit}}}$. This serves the purpose adequately,

diffusing the spurious noise mentioned above. The diffusion is along flat surfaces, as

are the horizontal fluxes. These intersect with the mountain surface, which is a disadvantage, but diffusing along the co-ordinate surfaces could produce artificial heat sources, by diffusing air to a level where the temperature is very different, for example. Ideally the diffusion would be along isentropes, but that would be very complicated to code, so flat surfaces are used instead. Clearly for fully three-dimensional simulations the artificial diffusion should depend on V and ΔY , as well as on U and ΔX , so the BLASIUS scheme should be modified to take account of this for simulations which use the y -direction. The work in this thesis is all two-dimensional in the x - z direction, so this is not an issue here.

The artificial viscosity works well in a flat case, but is more complicated on the bent mesh and requires further testing. The imposed value should be as small as possible while achieving the desired effect. The process is an artificial means of dealing with numerical issues, so there is no “correct” way to do the calculation physically, but using a horizontal velocity value obtained from the model fields would be better than hard coding a fixed representative value to be used throughout the model domain (as is done currently, taking $U = 10 \text{ m s}^{-1}$). Possibilities include a maximum value, the value at the top of the domain, and an average value. If the imposed minimum viscosity is to be constant throughout the domain, the maximum horizontal grid spacing should be used. Tests should also be carried out into the sensitivity of Re_{crit} to other numerical and physical parameters.

The only test carried out in this work was on a case with constant input wind speed (10 m s^{-1}) and static stability (0.01 s^{-1}), using a 500 m high Witch of Agnesi mountain of half width 5 km. (This is the same as case 1 described in Chapter 5, but with the mountain twice as wide.) A test was carried out using $Re_{crit} = 15$, giving a stricter condition by imposing stronger diffusion. This change did not make a significant difference to the results obtained. Further investigations would provide a basis for recommendations on the optimum tuning of this scheme. An effective way to test the sensitivity of the results to the value chosen for Re_{crit} would be to try extreme values (e.g. 0, 1, 1000) where a significant impact would be expected. Tests should also be carried out for different velocity profiles, since the condition is dependent on horizontal velocity.

5 Idealised test cases

5.1 Introduction

BLASIUS was designed to model boundary layer flow on smaller scales than those required to simulate lee waves and their associated phenomena, such as rotors. This work aims to test the code for mountain wave flows by performing simulations without a boundary layer for a range of simple idealised conditions, as a pre-cursor to more complicated cases investigating situations including a boundary layer, such as rotor formation. The aim here is simply to verify that BLASIUS can reproduce adequately these simple gravity wave flows, by comparing the numerical simulations with published analytical results for these cases. Some numerical issues were encountered during this testing process, and several modifications were tested on the different cases described below. These issues, the modifications tested, and the conclusions drawn from the tests, are introduced in section 5.4, and described in more detail in Chapters 6 to 8. This testing procedure led to a satisfactory configuration of the code, which was used to produce the results presented in this chapter and is described below. This configuration seems to be reasonably robust numerically for the cases studied here, although there are some remaining problems with reflection at the outflow boundary in two cases. Potential solutions to these issues are discussed in the following three chapters and summarised in Chapter 9 with suggestions for future work. Further testing on other cases would help to verify the robustness of this configuration, but time constraints did not permit this.

Case 1, with vertically constant input profiles of static stability, N , (0.01 s^{-1}) and wind, U , (10 m s^{-1}) was chosen as the simplest case to compare with results in the literature. The profiles chosen for case 1 represent typical atmospheric values. Cases 2 and 3 were chosen to correspond to the troposphere-only and troposphere-stratosphere cases in the analytical work of Keller (1994). Previous work at The Met. Office had begun to test BLASIUS on these two cases, and the basic flow pattern had been reproduced, but the code was not thoroughly tested (Guilbaud, private communication). Case 4 was chosen to investigate how well the model would cope with a wider mountain in the troposphere-stratosphere case, since this seemed to be a potentially interesting test.

These four cases are described in section 5.2, and the first three are compared with published results in section 5.3. Their different physical parameters are shown in table 5.1, and the columns in table 5.1 are explained below.

Case	Input velocity (m s^{-1})	Static stability (s^{-1})	Mountain half width (m)	Mountain height (m)
1	10	0.01	2500	500
2	Trop-only	0.01	2500	100
3	Trop-strat	Trop-strat	2500	500
4	Trop-strat	Trop-strat	20000	500

Table 5.1. The four cases and their physical parameters (see text for an explanation of the columns).

The first column gives the input horizontal velocity profile. A constant value of 10 m s^{-1} means that the input velocity had that value throughout the vertical domain. “Trop-only” means a troposphere-only profile, which started from 10 m s^{-1} at the surface, rising with constant shear of 0.0025 s^{-1} throughout the domain. “Trop-strat” means a troposphere-stratosphere profile, which started from 10 m s^{-1} at the surface, rising with constant shear of 0.0025 s^{-1} to a height of 10 km (the tropopause), then remaining at this constant value (just over 35 m s^{-1}) for the rest of the domain. Figure 5.1 shows the initial velocity profiles used in these different configurations.

The second column gives the input static stability, or Brunt-Väisälä frequency (denoted by N), used to determine the initial potential temperature profile by the method described in Chapter 4. A constant value of 0.01 s^{-1} means that N had that value

throughout the domain. “Trop-strat” means that N was 0.01 s^{-1} below 10 km and 0.02 s^{-1} above 10 km, to simulate a troposphere-stratosphere configuration. Figure 5.1 shows the initial potential temperature profiles used in these two configurations.




Figure 5.1. Input profiles of wind speed and potential temperature for the four cases. See text for more details. The troposphere-stratosphere profiles are smoothed near the tropopause, as described in the text; this is illustrated more clearly in figure 5.2.

The troposphere-stratosphere profiles of both wind speed and potential temperature had sharp discontinuities at the tropopause, as shown in figure 5.2. In order to avoid any numerical problems arising from this discontinuity, the profiles were smoothed near the tropopause using a 1-2-1 filter, which was applied ten times. The filter is defined for any field ϕ as

$$\phi_k^{new} = \frac{(\phi_{k-1}^{old} + 2\phi_k^{old} + \phi_{k+1}^{old})}{4}$$

where the subscripts indicate vertical levels. Figure 5.2 shows both the unsmoothed and smoothed input profiles of wind speed and potential temperature.

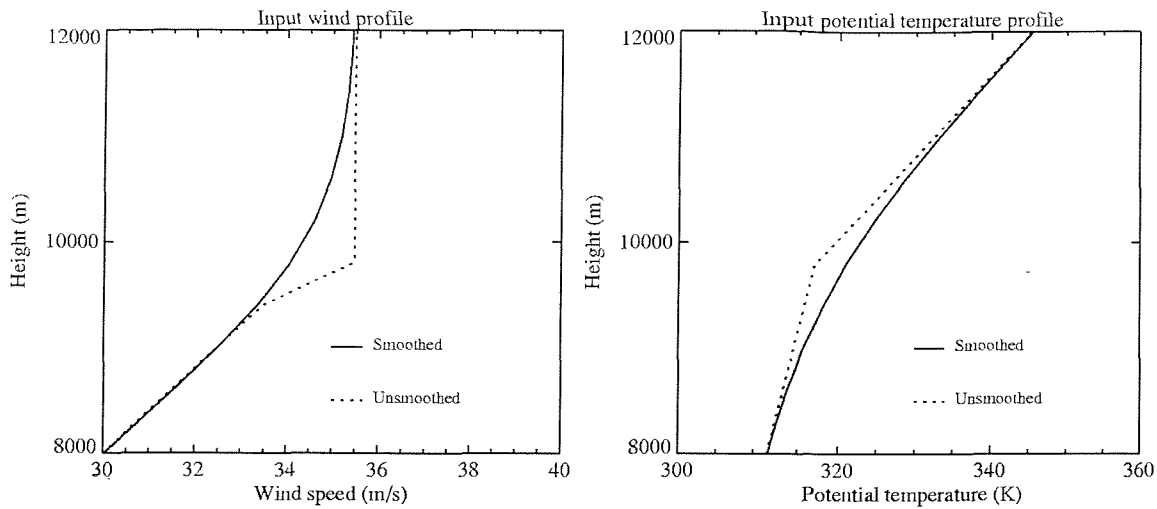


Figure 5.2. Close-up view of input wind speed and potential temperature profiles for the troposphere-stratosphere cases, which were shown in figure 5.1. Here both the smoothed and unsmoothed profiles are shown for a smaller domain surrounding the tropopause. The lower of the two kinks in the unsmoothed velocity profile occurs because the ‘tropopause’ is at slightly different heights in different vertical grid configurations, depending on the height of the grid level closest to 10 km. In order to maintain the same constant velocity in the stratosphere for all the troposphere-stratosphere simulations in this work, the velocity profile in this case was artificially adjusted just below the tropopause. As shown in this plot, the smoothed profile, which is actually used by the two-dimensional model simulation, does not retain any evidence of either kink.

The dominant horizontal wavenumber is given by a^{-1} , where a is the mountain half width as defined earlier, and represents a characteristic scale of the wavenumbers forced by the mountain (Durrán, 1986). This value can be compared with l , the Scorer parameter, discussed in Chapter 3 and defined as

$$l^2 = \frac{N^2}{U^2} - \frac{1}{U} \frac{d^2U}{dz^2}.$$

As discussed in Chapter 4, the input values of N do not truly represent the potential temperature profile used in the model, because of the approximation which is used in the code. This is equivalent to defining

$$N^2 = \frac{g}{\theta_r} \frac{d\theta}{dz}$$

where θ_r is a reference temperature; more details are given in Chapter 4. In addition, the potential temperature profile in the troposphere-stratosphere case is smoothed, as discussed above. The Scorer parameter profiles plotted in figure 5.3 are calculated

directly from the input wind and potential temperature profiles used in the model (and plotted in figure 5.1), as follows:

$$l = \sqrt{\frac{1}{U^2} \frac{g}{\theta} \frac{d\theta}{dz} - \frac{1}{U} \frac{d^2U}{dz^2}}.$$

Because of this, they reflect the change of potential temperature with height, and so they vary more with height than might be expected from a simple calculation using the input parameters. As an example, for case 1, $U = 10 \text{ m s}^{-1}$ and $N = 0.01 \text{ s}^{-1}$, so the Scorer parameter l is simply N/U ; however, the profile in figure 5.3 is not constant with height but has constant shear with height. Similarly, for case 2, the input parameters would give a Scorer parameter profile with constant shear, but there is more curvature in the actual profile. Vertical profiles of the Scorer parameter are shown in figure 5.3 for the four different cases, compared with the dominant horizontal wavenumbers. The ratio of horizontal wavenumber to Scorer parameter gives information on what type of wave pattern can be expected from the theory, and this is discussed below for each case in turn.




Figure 5.3. Input profiles of Scorer parameter (compared with the dominant horizontal wavenumber) for the four cases. See text for more details.

The last two columns of the table describe the mountain. The shape was a Witch of Agnesi mountain for all of these cases, although a cosine-squared mountain was used in some of the tests described in later chapters. The equation for the cosine-squared mountain is given in section 8.3, which discusses the different results obtained using the two different mountain shapes. The equation for the Witch of Agnesi mountain is

$$h(x) = \frac{h_m a^2}{a^2 + x^2}.$$

Here $h(x)$ is the mountain height as a function of distance from the crest, h_m is the height of the crest, and a is the half width of the mountain, as given in table 5.1, defined as the distance from the crest to the point at which the mountain height is $h_m/2$.

Figure 5.4 shows the mountains produced by these different configurations.

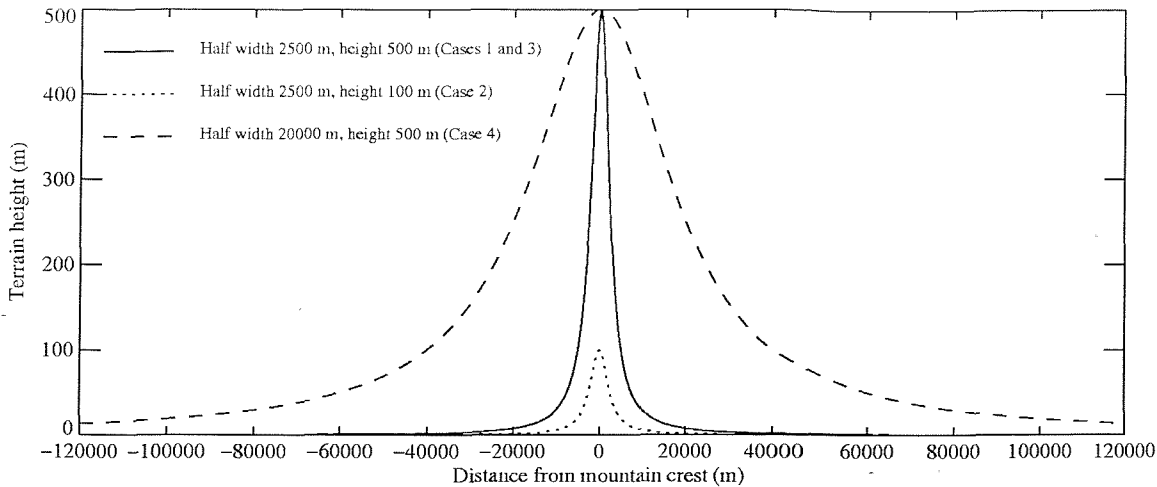


Figure 5.4. Mountain shapes for the different cases. Note that the vertical scale is stretched for clarity. See text for more details.

The structure of the grid used for these simulations was developed through consideration of the expected flow patterns, and through tests using grids with varying domain length, grid spacing and stretching. Details of the grid independence tests can be found in Chapter 6. The horizontal domain used here was 300 km wide, with the mountain crest in the centre of the domain. A stretched grid was used, with a total of 270 points in the horizontal. The grid was calculated using input parameters so that 35 points lay within 10 km either side of the mountain, and 60 points lay between 10 km and 70 km either side of the mountain, leaving the remaining 40 points in the last 80 km of the domain at either end. The grid was smoothed until the expansion ratio, defined as $\Delta X(I+1)/\Delta X(I)$, was a maximum of 1.05. This quantity provides a measure of how fast the grid spacing changes, and is discussed further in Chapter 6. After smoothing, the grid spacing was just over 286 m in the centre of the domain, increasing gradually to just under 2000 m at the edges of the domain. For case 4, where the mountain was eight times as wide as in the other three cases, the horizontal grid was scaled so that the gridlengths and domain width were all eight times as large. The grid was symmetrical about the centre of the domain. Figure 5.5 shows that both the horizontal gridlength and the expansion ratio varied smoothly and slowly, although the two stretch points can still be seen.




Figure 5.5. Horizontal gridlengths and expansion ratios for the grid described in the text (as used for cases 1 to 3; the grid used for case 4 was eight times as wide, as discussed in the text). The mountain crest is in the centre of the domain, and the grid is symmetrical about the centre of the domain. The expansion ratios in the upstream half of the domain are less than one because the gridlength is decreasing from left to right rather than increasing, as it does in the downstream half of the domain.

In the vertical, 127 gridpoints were used in a 50 km domain, with a constant grid spacing of 400 m. Two points are outside the physical domain due to the staggered grid and boundary conditions, of which more details were given in Chapter 4. The Rayleigh damping layer began at 35 km, so results are only presented up to that height.

Figure 5.6 shows the grid described above for a section of the model domain centred on the mountain. The full domain is not shown, since the grid would be hard to see. The Witch of Agnesi mountain of height 500 m and half width 2500 m, used in cases 1 and 3, is shown in this example plot; the grid used in case 4 is the same relative to the mountain as this grid. The surface elevation of the Witch of Agnesi mountain never

truly becomes zero, but asymptotes towards zero. The cosine-squared mountain of the same height and cross-sectional area has width 15700 m (as discussed in section 8.3), and there are 59 horizontal gridpoints on the surface of the Witch of Agnesi mountain within this reduced horizontal domain. In the case of a 500 m high mountain, there are two vertical grid levels below the mountain crest height, at the surface and at 400 m. The third vertical grid level is at a height of 800 m, above the mountain crest. Case 2 had a 100 m high mountain, so the first vertical grid level was at the surface, and the next one, at 400 m, was above the mountain crest in this case. The heights given for the vertical grid levels refer to the distance of the model level above the flat surface upstream and downstream of the mountain, and do not take account of the bent mesh.

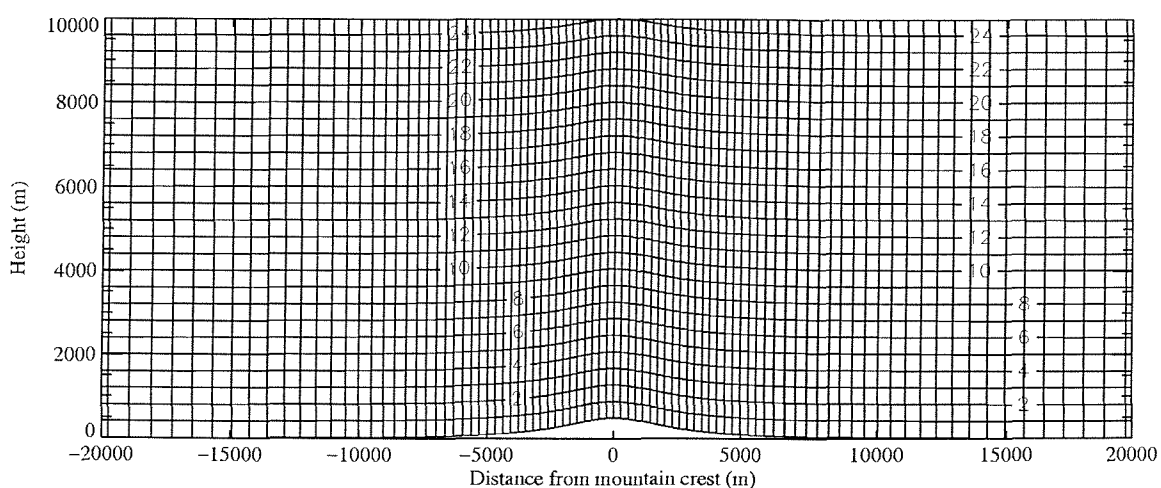


Figure 5.6. The grid described in the text, shown for a partial domain centred on the mountain. The mountain of cases 1 and 3 is used here to illustrate the grid. The numbers on the contours indicate the vertical model levels.

As discussed in Chapter 4, the simulations described here are all two-dimensional in the x - z plane, and Coriolis forces are neglected. The results for the four cases are presented below. The model diagnostics include time series of the horizontally averaged surface pressure force in the x -direction. Horizontal averaging was done in the x -direction only, since these are two-dimensional x - z simulations. Figure 5.7 shows the drag coefficient, which is obtained from the surface pressure force as described in Chapter 4, for all the cases. The time series show data collected at 125 second intervals throughout each simulation, up to the end of the simulation after 55000 seconds (approximately 15 hours). The noisy signal seen at early times is due to the large perturbations caused by

the introduction of the mountain, since the simulation starts with a flat surface and gradually grows the mountain into the domain. This ‘initial noise’ is evident in contour plots of vertical velocity at early times, but is usually quickly advected out of the model domain. The process of growing the mountain into the domain, and the initial noise, are discussed further in section 8.4.




Figure 5.7. Time series of the drag coefficient for all the cases.

The drag coefficient can be used to assess whether a simulation has reached steady state, since in that case the graph should flatten out to an approximately constant value. From figure 5.7 it is clear that cases 1 and 3 reached a reasonably steady state after approximately 30,000 seconds. Case 2 was less steady, and possible reasons for this are discussed in Chapter 6. Its final value was approximately 20% of the final value in cases 1 and 3, which may be explained by the dependence of the drag coefficient on the reciprocal of the Froude number (Hanazaki, 1988). Since the near-surface values of wind speed and static stability are the same in all of these simulations, this result reduces to a dependence on the mountain height. Although Hanazaki’s result was for flow past a sphere, it seems a plausible explanation for the difference seen in this case, since it relates the magnitude of the drag coefficient to the strength of the forcing, with a higher mountain providing a steeper forcing. The final value for case 4 was also lower than that for cases 1 and 3; this may be explained by a link to mountain width, since the wider mountain means the slope is more gentle, providing a weaker forcing.

Alternatively, this may be because this simulation is still spinning up, since case 4 did not run for as long as the others in terms of non-dimensional time. Figure 5.8 shows the drag coefficient for the four simulations as a function of non-dimensional time, t_{nd} ,

which is defined as $t_{nd} = (t - t_g) \frac{U}{L}$ where t is the time in seconds, t_g is the time (in seconds) taken to grow the mountain, U is a velocity scale (given by the near-surface wind speed) and L is a length scale (given by the half width of the mountain). This non-dimensional time is used instead of model time in the rest of the thesis, to enable easier comparisons between simulations with different physical parameters, and with mountains grown into the numerical domain at different speeds. These non-dimensional time units are denoted by ‘NDTU’ in the text.




Figure 5.8. Time series of the drag coefficient for all the cases, with non-dimensional time on the horizontal axis, as defined in the text. The vertical long-dashed line marks the point where the non-dimensional time is zero, at the end of mountain growth. It is clear that case 4, with the wider mountain, ran for far less non-dimensional time than the other three cases, and grew the mountain far more quickly.

A simulation of case 4 using the same horizontal grid as cases 1 to 3 was unsteady throughout the original 55000 seconds, and when it was run on for longer to obtain a comparable non-dimensional time, it became very noisy. The time series of vertical velocity (not shown here) indicated that this was most likely due to spurious noise at the outflow boundary, which was reflected into the domain and then reflected back and forth between the upstream and downstream boundaries. Growing the mountain more slowly (over a non-dimensional time which is comparable with the other 3 simulations)

and then running for a non-dimensional time which is comparable with the other 3 simulations did not ease the problem; periodic boundary conditions helped considerably. However, scaling the grid seems the most logical solution, since the numerical grid is then the same relative to the mountain in some non-dimensional sense, and produced by far the best results over the original simulation time of 55000 seconds. This case 4 simulation, with the wider horizontal domain, was run for longer, and did remain steady for longer than the one with the smaller horizontal domain. However, after 150000 seconds (65 NDTU), noise began to appear at the outflow boundary, and the drag coefficient increased considerably as this spurious noise was reflected back and forth across the domain, so the problem was delayed rather than removed, which is not an unexpected result. This indicates a problem with the outflow boundary condition, which is discussed further in Chapter 7; it also suggests that the results for case 4 may need to be interpreted with caution, although they do provide a useful test for the investigations reported in subsequent chapters.

In the descriptions below, contour plots of vertical velocity are used. These were produced at 5000 second intervals throughout each simulation, but only selected data are shown here, to illustrate the dominant features of each case. Potential temperature contours are used to represent streamlines, and this assumption of isentropic flow is valid for these steady state cases, as discussed in Chapter 3. The potential temperature plots are only shown for the lowest third of the model domain, to make the signal more visible in the plots.

5.2 Description of the results

For case 1, with constant wind speed, U , and static stability, N , the horizontal wavenumber is less than the Scorer parameter, so the theory predicts an upwardly propagating wave. The BLASIUS simulation did indeed produce an upwardly propagating wave for this case, showing that the model is able to simulate this type of flow pattern. The flow appeared from the vertical velocity plots to be steady after 60 NDTU, and this was confirmed by the drag coefficient, which settled to a constant value around the same time. The Froude number was 2 in this case, indicating that non-linear effects may be occurring, but there was no strong evidence of these in the results. Figure 5.9 shows the vertical velocity and potential temperature after 100 NDTU, at the end of the simulation.

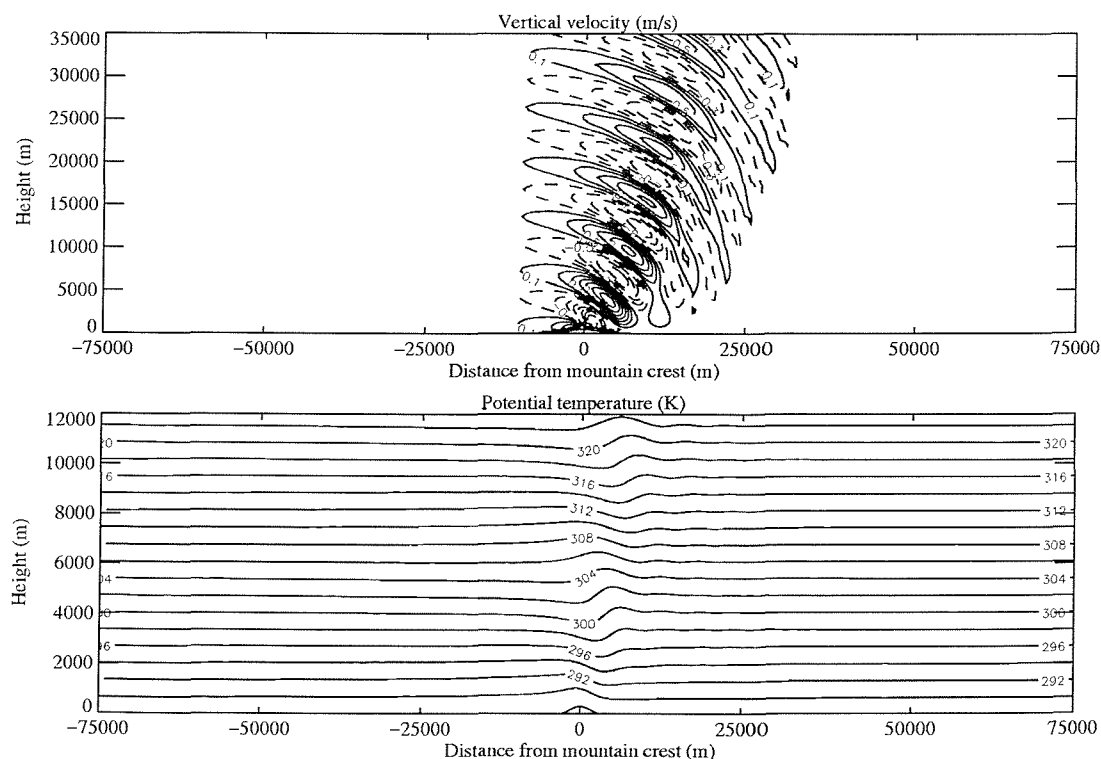


Figure 5.9. Vertical velocity (contour interval 0.2 m s^{-1}) and potential temperature (contour interval 2 K) for case 1 after 100 NDTU. The vertical velocity plot has a base contour of 0.1 m s^{-1} ; the zero contour has been removed from this plot, but is very smooth both upstream and downstream of the wave. The solid contours represent positive values and the dashed contours represent negative values; this convention will be used for all vertical velocity plots unless stated otherwise.

Case 2 was the troposphere-only simulation, and gave a horizontally propagating wave with an irregular structure, which fits with Keller's (1994) description of two superimposed waves of different wavelengths. The horizontal wavenumber is less than the Scorer parameter below 5 km, and the reverse is true above that height, so the theory predicts that the upwardly propagating wave will be trapped at lower levels and propagate downstream. The main component of the flow is a horizontally propagating wave, so this does fit with the theory. This simulation was designed as a test of the model against Keller's (1994) analytical results, and a comparison with these follows in section 5.3 below. The Froude number is 10 in this case, so the flow pattern should be linear, and the comparison with Keller's results indicates that this is so.

The wave developed in a stable manner with no apparent problem until the wave head reached the outflow boundary, which was some time between -45 and -25 NDTU (with negative non-dimensional time values indicating that the data were produced during mountain growth), as shown in figure 5.10.

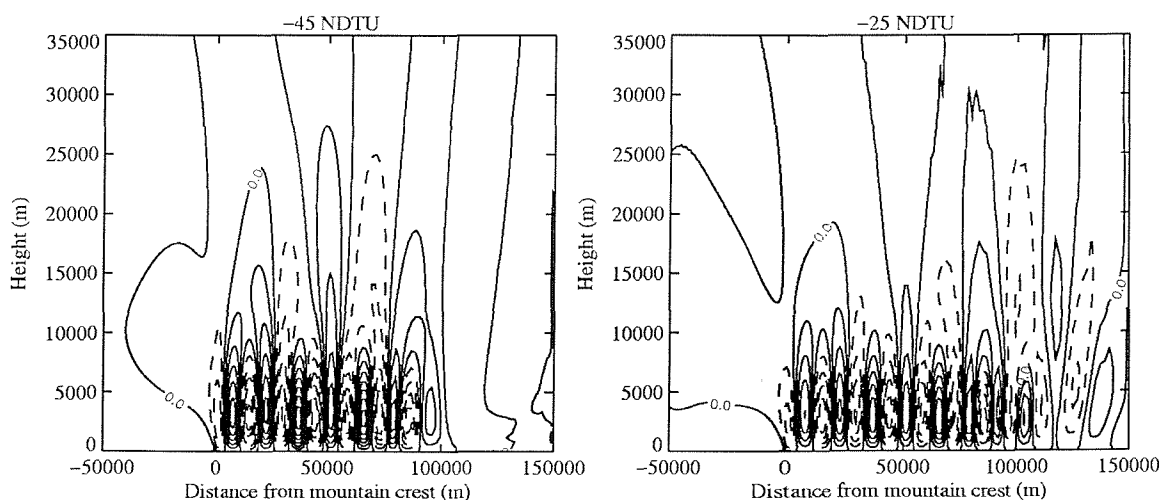


Figure 5.10. Vertical velocity (contour interval 0.2 m s^{-1}) and for case 2 after -45 and -25 NDTU (left and right hand plots respectively). The wave head reached the outflow boundary between these two times.

It is clear from figure 5.10 that the wave became less well resolved as it entered the coarser outer region of the horizontal grid, which is a deliberate numerical damping effect as an attempt to minimise reflection problems at the outflow boundary, and is discussed further in Chapter 7. However, this simulation still suffered from some

reflection at the downstream boundary, and the flow was perturbed by this, with upstream noise appearing after 0 NDTU, at the end of mountain growth, which is presumably due to the arrival of the reflected signal in the upstream domain. This signal, shown in the left hand plot of figure 5.11, fits with a large peak in the drag coefficient just before 0 NDTU, which is likely to be because the reflected noise was passing over the mountain at that point and affecting the pressure field.




Figure 5.11. Vertical velocity (contour interval 0.1 m s^{-1}) and for case 2 after 0 and 35 NDTU (left and right hand plots respectively). Spurious noise is evident upstream and above the wave after 0 NDTU, at the end of mountain growth, but the noise is almost gone by 35 NDTU.

From 35 NDTU the noise began to settle, and by the end of the simulation, after 135 NDTU, the wave pattern was reasonably smooth again, as shown in the right hand plot of figure 5.11, although the damping after the second grid stretch was still evident. The drag coefficient signal was relatively noisy, which is explained by the perturbations discussed above, resulting from reflection by the outflow boundary condition.

Figure 5.12 shows the vertical velocity and potential temperature after 135 NDTU.




Figure 5.12. Vertical velocity (contour interval 0.1 m s^{-1}) and potential temperature (contour interval 2 K) for case 2 after 135 NDTU. The vertical velocity plot has a base contour of 0.05 m s^{-1} ; the zero contour has been removed from this plot, but does show a small amount of spurious noise upstream of the wave, as discussed in the text.

The troposphere-stratosphere simulation, case 3, was very noisy initially but settled by 35 NDTU to give a horizontally propagating wave which faded downstream, with a weaker upwardly propagating wave above. Figure 5.13 shows the vertical velocity fields after 0 NDTU (at the end of mountain growth) and 35 NDTU.

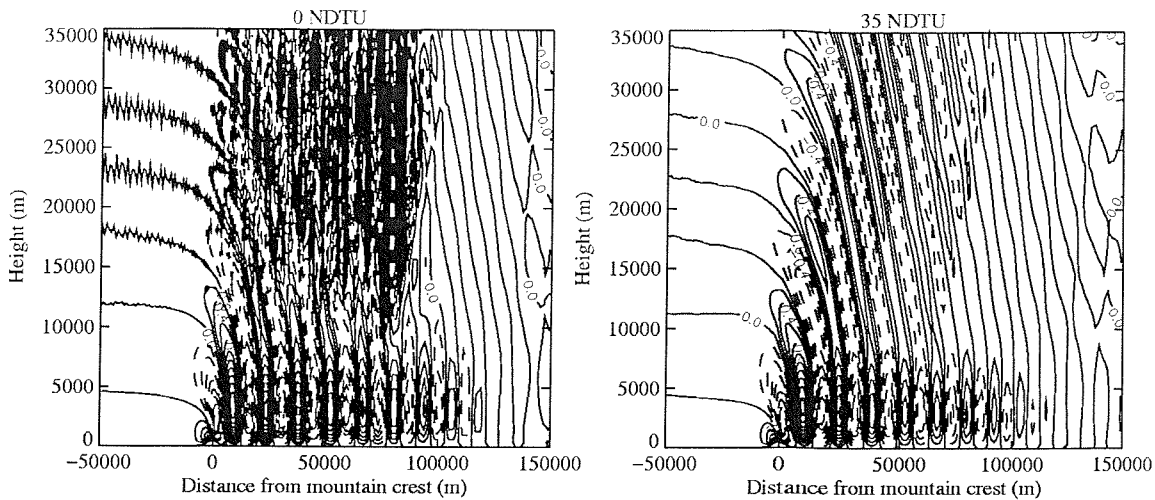


Figure 5.13. Vertical velocity (contour interval 0.2 m s^{-1}) and for case 3 after 0 and 35 NDTU (left and right hand plots respectively). There is a lot of initial noise after 0 NDTU, at the end of mountain growth, but it is almost gone by 35 NDTU.

The Scorer parameter profile is more complicated in this case than in the others, being greater than the dominant horizontal wavenumber below 5 km and above 10 km, but less in between. This indicates that upwardly propagating waves can be expected in the lowest layer. However, the tropopause acts to trap these waves, as discussed in Chapter 3, so that they propagate downstream, as evident in the model results for this case. When the dominant horizontal wavenumber is greater than the Scorer parameter, an evanescent wave with vertical phase lines is predicted, but above 10 km the reverse is true once more, and so the wave begins to propagate upwardly again, which explains the signal in the stratosphere. As in case 2, the second grid stretch is clearly marked in the vertical velocity signal because of the damping effect it had on both the initial noise and the actual solution. The damping region downstream worked well in this case, as both the upper and lower level waves stopped around 100 km downstream of the mountain, and there was no obvious sign of spurious reflection at the outflow boundary. The flow seemed reasonably steady after 75 NDTU, and the drag coefficient had settled to an almost constant value by then, with a slight oscillation. The Froude number is 2 here, and there is some evidence of non-linear interaction between the two wave trains, as discussed in section 5.3 below, where the results are compared with Keller's (1994) analytical results. Figure 5.14 shows the vertical velocity and potential temperature after 115 NDTU.




Figure 5.14. Vertical velocity (contour interval 0.2 m s^{-1}) and potential temperature (contour interval 4 K) for case 3 after 115 NDTU. The vertical velocity plot has a base contour of 0.1 m s^{-1} ; the zero contour has been removed from this plot, but is smooth both upstream and downstream of the wave.

The mountain width was multiplied by eight in case 4, giving a half width of 20 km and a hydrostatic upwardly propagating wave, as predicted by the theory, since the horizontal wavenumber is less than the Scorer parameter, as for case 1. Although the Rossby number is lower in this case than in the other three cases, Coriolis effects are not expected to be significant. The formulae given by Dörnbrack *et al.* (1999) suggest a vertical wavelength of around 11 km for this case (which fits well with the signal in figure 5.15), reduced by around 0.2% due to Coriolis effects. This change is small enough to justify the use of a two-dimensional simulation with a Coriolis parameter of zero. By the end of mountain growth, the vertical velocity signal showed no initial noise, and the wave was steady from then onwards. The Froude number in this case is 2, so there could be some non-linear effects, but there is no clear evidence of these in the results. The wave had greater horizontal extent than in case 1, due to the wider mountain which took up around one third of the horizontal model domain. Figure 5.15 shows the vertical velocity and potential temperature after 17.5 NDTU.




Figure 5.15. Vertical velocity (contour interval 0.05 m s^{-1}) and potential temperature (contour interval 4 K) for case 4 after 17.5 NDTU. The vertical velocity plot has a base contour of 0.025 m s^{-1} ; the zero contour has been removed from this plot, but shows the same spurious noise both upstream and downstream of the wave as is already visible in this plot; this is discussed in the text.

5.3 Comparison with published analytical results

Case 1, with constant input wind speed, U , and static stability, N , is widely discussed in the literature, and details can be found in the reviews listed at the start of Chapter 2. The potential temperature plot for case 1 can be compared to figure 20.2 of Durran (1986), which shows the steady analytical solutions for an isolated bell-shaped ridge with constant background wind speed and static stability. Durran's solutions depend on the relationship between a^{-1} and l , as discussed earlier. For the case 1 results presented here, the value of a is 2500 m, giving a^{-1} a value of 0.0004 m^{-1} . The Scorer parameter for case 1 has a value of 0.001 m^{-1} at the surface, decreasing linearly with height to approximately 0.00085 m^{-1} at a height of 35 km, as shown in figure 5.3. Hence

$a^{-1} = 0.4l$ at the surface and $a^{-1} = 0.47l$ at 35 km, so these results fall between those presented by Durran in his figure 20.2b (for $a^{-1} = l$) and 20.2c (for $a^{-1} \ll l$), which are reproduced as figure 5.16(i) here. The pattern of potential temperature for the numerical simulation of case 1 is shown again in figure 5.16(ii), with a smaller domain for ease of comparison with Durran's streamline plots. The numerical results should be closer to the $a^{-1} = l$ case, and in fact are qualitatively very similar to the analytical ones, showing that BLASIUS copes well with this simple case. It should be noted that some differences are to be expected between the inviscid analytical solutions and these low viscosity numerical simulations, which use different surface boundary conditions. A more thorough comparison could be carried out by working through Durran's analytical solution for the parameters of this simulation, but time did not permit this.

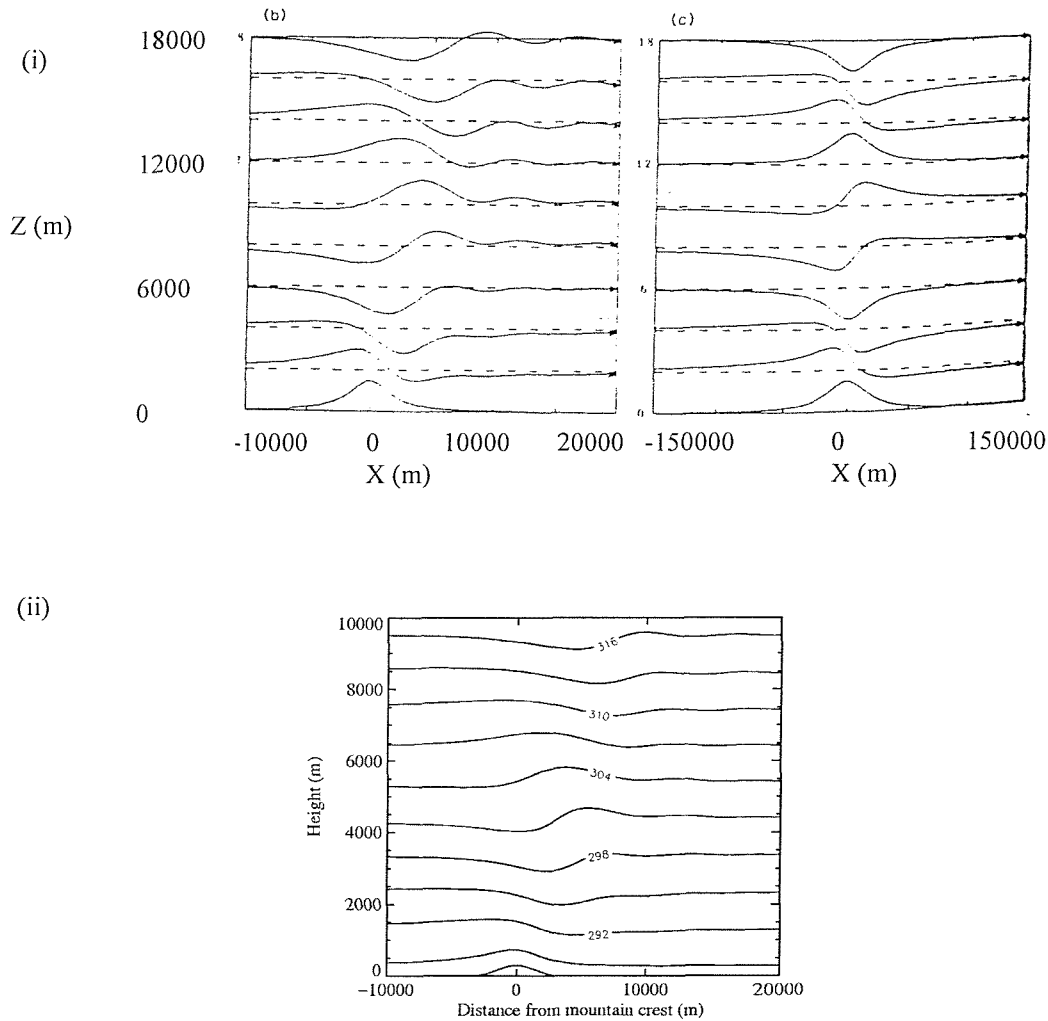


Figure 5.16. (i) Durran (1986) figure 20.2b for $a^{-1} = l$ and 20.2c for $a^{-1} \ll l$, as described in the text. (ii) Potential temperature for the BLASIUS case 1 simulation after 100 NDTU, shown for part of the model domain, with contour interval 3 K.

Linear two-dimensional steady state analytical solutions were derived by Keller (1994) for both hydrostatic and non-hydrostatic cases. The latter can be used as a test of the BLASIUS model, since it should produce similar qualitative results when given similar initial conditions. Cases 2 and 3 in this work are intended to reproduce Keller's results for the troposphere-only and troposphere-stratosphere cases respectively.

A difference between Keller's work and the numerical simulations presented here is that the analytical study had no viscosity, so was truly inviscid. Corby (1954), reviewing the assumption of frictionless flow in perturbation theory, suggested that viscosity would merely exert a small damping effect on the disturbances, but would not alter the results qualitatively. The BLASIUS simulations have some viscosity, as discussed in sections 4.1 and 4.6. However, this is for numerical purposes and has been shown not to have a significant physical or qualitative impact on the results, although as Corby suggested, it does act to damp the solution, as discussed below.

For the troposphere-only case, Keller observed that the solution consists of two trapped waves propagating horizontally downstream, with wavelengths of 15.2 and 35.9 km. Figure 5.17 compares the results from the BLASIUS simulation with Keller's figure 3. Qualitatively the results compare well, with very similar wavelength, although the damping effect due to the second grid stretch means that the comparison is unfavourable to the model downstream of 70 km (where the agreement is less good, as would be expected). Quantitatively the wave has only 25% of the magnitude of Keller's results, in terms of the maximum vertical velocity in the peaks and troughs. This can be at least partly explained by the implicit and explicit (imposed) diffusion in BLASIUS (of which further details are given in section 4.6), because this exerts a damping effect on the flow. Also, the no slip surface boundary condition may be expected to artificially retard the low level flow far downstream of the mountain, but no significant impact is expected on the general results. The quasi-free slip modification to BLASIUS (which is explained more fully in section 8.2) proved unsatisfactory, so it was not possible to use it in the model simulations presented here.

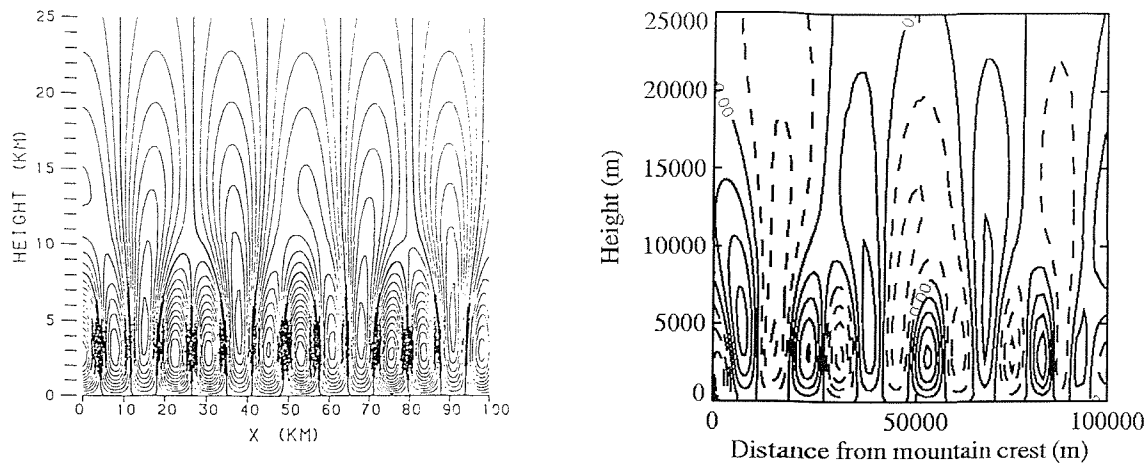


Figure 5.17. Vertical velocity for case 2, the troposphere-only case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 3, which is shown on the left. The contour interval in the numerical simulation is 0.1 m s^{-1} , half that used in Keller's plot. Note that in this case, all of Keller's contours are solid; the leftmost cell is negative in this plot.

For the troposphere-stratosphere case, figure 5.18 shows the results from the BLASIUS simulation, and may be compared with Keller's figure 5, with the same contour interval used in both plots. As for the troposphere-only case, the qualitative agreement is very good, with the wavelength matching well at both upper and lower levels, and the tilt of the upper level wave train reproduced well. Note that Keller's plot only extends 60 km downstream of the mountain in this case, so the second grid stretch (at 70 km), with its resulting damped region, need not be taken into account when comparing the results on this reduced horizontal domain. The magnitude of the vertical velocity in the peaks and troughs is very close at upper levels, but less so at lower levels, up to 30% stronger in the numerical simulation than in the linear calculation directly above the lee slope of the mountain, and then becoming progressively weaker downstream. In the final wave trough shown in Keller's plot, 60 km downstream of the mountain, the vertical velocity is 50% weaker in the numerical simulation. As discussed for the troposphere-only case above, part of the weaker signal is due to the diffusion in the numerical model. The more marked decrease in magnitude downstream (Keller's signal does decrease too, but far less dramatically) is most likely caused by the effect of the no slip surface condition in BLASIUS, which acts to retard the flow downstream by pulling the horizontal velocity back to zero rather than allowing a non-zero value at the surface. Limited tests of the quasi-free slip surface condition for horizontally propagating wave cases have been carried out for BLASIUS as part of this thesis (see section 8.2 for more details of

the boundary condition), and these do show less downstream damping of the wave than in the corresponding no slip cases, supporting the above hypothesis. The larger vertical velocity magnitude close to the mountain may be explained by a non-linear interaction between the two wave trains, which overlap at this point. Since their wavelengths are in phase, they would be expected to interfere constructively with each other, increasing the perturbation, as seen here. Keller's analysis is linear, so her results will not include any additive effects due to non-linear interaction between the two wave trains (although they will include linear additive effects). The Froude number is 2 in this case, so non-linear effects may be playing an important role, and the fact that this constructive interference effect is not evident in Keller's solution suggests that it is likely to be a non-linear feature. Repeating the simulation with a reduced mountain height would test this hypothesis, since the Froude number would then be lower, and any non-linear effects would be reduced, if not removed.

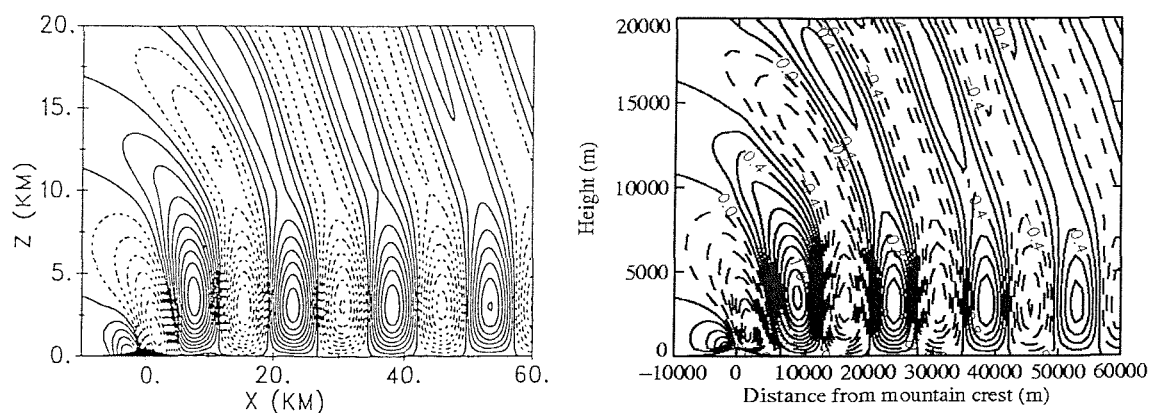


Figure 5.18. Vertical velocity for case 3, the troposphere-stratosphere case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 5, which is shown on the left. The contour interval in the numerical simulation is 0.2 m s^{-1} , the same as that used in Keller's plot.

For case 4, figure 5.3 shows that the dominant horizontal wavenumber was less than the Scorer parameter throughout the domain. The hydrostatic wave obtained in this case fits well with the theory. Although the input wind and potential temperature profiles were not constant with height here, a loose qualitative comparison with Durran's (1986) results, as shown in figure 5.16(i), indicates that the results are in reasonable agreement. Figure 5.19 shows the potential temperature for the numerical simulation of case 4, with a smaller domain for ease of comparison with Durran's streamline plots.




Figure 5.19. Potential temperature for the BLASIUS case 4 simulation after 17.5 NDTU, shown for part of the model domain, with contour interval 1.5 K.

Despite the differences described above, for which explanations have been suggested, the qualitative agreement for these three cases was sufficient to validate the results produced by BLASIUS.

5.4 Introduction to the numerical issues encountered

The results described in this chapter were produced after testing various different numerical configurations of BLASIUS. The aim of these tests was to find a configuration of the code which would produce stable results for all four cases, without the numerical problems encountered in the original simulations, and the tests are discussed in the following three chapters. Although the simulations should in theory be grid independent, there are some grid dependent changes in results, and these issues are discussed in Chapter 6. Sensitivity tests were carried out for both the horizontal and vertical grids, using stretched and uniform grids of different resolutions, as well as investigating the effect of changing the Rayleigh damping layer height and coefficient. Chapter 7 covers issues concerning the upstream and downstream boundary conditions, since many simulations had spurious noise at or near the outflow boundary, and it seems that the radiative outflow boundary condition described in Chapter 4 is not effective in all situations. The remaining issues are dealt with in Chapter 8: the model timestep, the surface boundary condition, the shape of the mountain, the initial noise caused by growing the mountain into the domain, and the formulation of the artificially imposed viscosity. Some cases produced reasonable results in the original configuration, with no numerical difficulties, and tests on those cases provided a control, to ensure that the changes did not introduce problems. For each modification, comparisons were made between pairs of simulations for the same case, to assess the impact of the change. Chapters 6 to 8 describe the modifications tested, and give illustrative examples of their impact.

6 Grid independence

6.1 Introduction

One of the many unavoidable approximations in numerical models of the atmosphere is the use of a discrete grid to represent the continuous space in which the fluid motions take place. Discrete time steps are also used, and this process is discussed further in section 8.1. The resolution of the discrete grid must be fine enough to resolve adequately the flow dynamics of interest, and the domain must be large enough to allow the flow to evolve. However, the amount of available computer memory and processing power can place limitations on the domain size and grid resolution which can be used. These competing issues need to be balanced in order to provide the best possible simulation of the flow in question. This chapter discusses investigations carried out, and results obtained, for different formulations of the horizontal and vertical grids in two-dimensional simulations using the BLASIUS numerical model. Plots of vertical velocity are used to illustrate most of the comparisons, since it is generally more sensitive to changes than is the potential temperature field. A uniform or stretched grid can be used in either direction, and both of these possibilities are considered, with varying gridlengths and different methods of stretching. In addition, the Rayleigh damping scheme described in Chapter 4 is tested for different damping coefficients and different layer heights. Upstream and downstream boundary conditions can also have an impact on the effectiveness of the grid and resolution used, and although these issues are mentioned here, they are discussed more fully in Chapter 7.

6.2 Horizontal grid dependence

6.2.1 Uniform grid

The simplest case considered was the case of constant input wind speed, U , and static stability, N , using a 100 m high Witch of Agnesi mountain with a half width of 10 km. A uniform grid was used in the horizontal with a grid spacing of 1 km. In the vertical, the uniform grid spacing was 350 m. The domain was 300 km wide (with the mountain in the centre) and the vertical domain was 35 km high with the Rayleigh damping layer starting at 20 km. The wave propagated upwards, so periodic upstream and downstream boundary conditions were used, and were not expected to have a significant impact on the flow pattern obtained. Figure 6.1 shows the vertical velocity field after 140 NDTU, although the simulation reached steady state much earlier. The horizontal resolution seems adequate here: the wave is smooth and well-resolved.




Figure 6.1. Vertical velocity after 140 NDTU for a 100 m high Witch of Agnesi mountain of half width 10 km. The contour interval is 0.01 m s^{-1} with a base contour of 0.005 m s^{-1} ; the zero contour has been removed from this plot, but is very smooth both upstream and downstream of the wave.

The mountain height was increased to 500 m, so the slope was five times steeper. The vertical velocity signal, shown in figure 6.2, was very similar, but its maximum magnitude in the wave crests was slightly more than five times larger. Potential

temperature is also shown for this simulation, but there was barely any discernible signal in the case with the smaller mountain. Again, the wave in this case is well resolved and the resolution seems adequate.

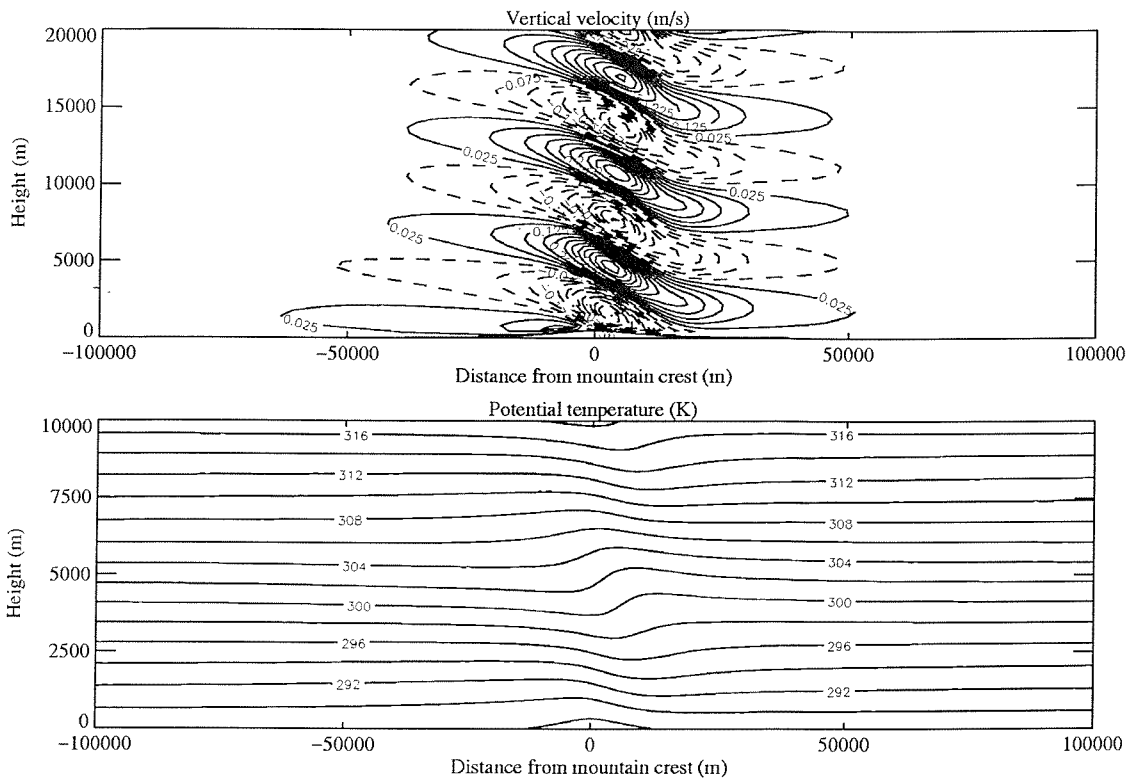


Figure 6.2. Vertical velocity and potential temperature after 140 NDTU for a 500 m high Witch of Agnesi mountain of half width 10 km. The contour interval for the vertical velocity is 0.05 m s^{-1} with a base contour of 0.025 m s^{-1} ; the zero contour has been removed from this plot, but is very smooth both upstream and downstream of the wave. The contour interval for the potential temperature is 2 K; note that potential temperature is plotted for half the vertical domain used for vertical velocity, in order to show the signal more clearly.

The mountain width was halved relative to that in the simulation shown in figure 6.2, while keeping the height at 500 m, and the vertical velocity and potential temperature were compared for half the horizontal domain and double the contour interval (see figure 6.3). The flow pattern relative to the mountain was qualitatively very similar in both simulations, with the vertical velocity maximum in the wave crests approximately 1.5 times larger. However, there was a small amount of noise in the vertical velocity signal for the wave induced by the narrower mountain, which did not amplify with time; one possible explanation for this may be that the horizontal resolution was not quite

good enough. The noise was not evident in the potential temperature plot, where the amplitude of the perturbation was the same as in the wider mountain case. Halving the mountain width further, to 2500 m, and keeping the horizontal grid spacing at 1 km, produced a very poorly resolved wave, demonstrating that a much smaller gridlength would be required for a mountain this narrow. Figure 6.4 shows how the resolution became less adequate as the mountain width decreased, at least by ‘chopping off’ the peak of the mountain.




Figure 6.3. Vertical velocity and potential temperature after 250 NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km. The contour interval for the vertical velocity is 0.1 m s^{-1} with a base contour of 0.05 m s^{-1} ; the zero contour has been removed from this plot, but is very smooth both upstream and downstream of the wave. The contour interval for the potential temperature is 2 K; note that potential temperature is plotted for half the vertical domain used for vertical velocity, in order to show the signal more clearly.




Figure 6.4. Horizontal gridpoints at the surface over 500 m high Witch of Agnesi mountains of half width 10 km, 5 km and 2.5 km, using a uniform horizontal grid spacing of 1 km.

Reducing the horizontal grid spacing from 1 km to 750 m produced a generally similar signal in the vertical velocity and in the potential temperature, although there were differences. Some initial noise is created in the model fields due to the growth of the mountain into the domain starting from a flat surface, and this process is discussed in section 8.4 along with investigations into the effect of growing the mountain at different speeds. The initial noise is usually visible in both the vertical velocity and the potential temperature fields, and in this case it was more marked than in the original case, and persisted throughout the simulation, although it faded with time; this is illustrated in figure 6.5, which shows the fields after 250 NDTU. The disturbance on the downstream edge of the wave was also more marked than in the original simulation, with a larger maximum magnitude in the wave crests and more pronounced distortion of the wave pattern. The mountain was better resolved, especially around the peak, but the flow was more noisy.

Reducing the horizontal grid spacing further, to 500 m, made the problem more noticeable, and a small wave became discernible on the potential temperature signal in the area affected by the noise. Horizontal profiles of vertical velocity showed a marked wave-like structure due to the downstream noise, with a wavelength of approximately 6 km (which equates to around eight 750 m gridlengths or twelve 500 m gridlengths). This, too, increased in maximum magnitude in the wave crests (but kept the same wavelength) as the gridlength was decreased.




Figure 6.5. Vertical velocity and potential temperature after 250 (in the left hand plot) and 230 (in the right hand plot) NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km. The contour interval for the vertical velocity is 0.1 m s^{-1} with a base contour of 0 m s^{-1} . The left hand plot has a horizontal grid spacing of 1 km; the right hand one has a horizontal grid spacing of 750 m. Data are not available at exactly the same non-dimensional time for these two simulations, but the difference is small enough not to be significant.

These results show that these simulations are not grid independent. Once the mountain and flow pattern are both adequately resolved, making the resolution finer should not change the solution, since the timestep is adjusted in order to keep the CFL number the same. The problems seen here could be due to initial noise, as discussed in section 8.4. A potential solution could be to make the CFL criterion more strict, as discussed in section 8.1. Rayleigh damping could be having a negative impact due to reflection, and this is discussed further in section 6.3.2. It should also be noted that the increase in noise as the gridlength decreases is consistent with the associated decrease in numerical dispersion and diffusion: stronger evidence is required in order to locate the source of the problems, but the above issues are discussed further in subsequent sections of this thesis. Horizontal artificial viscosity was discussed in section 4.6, but needs further investigation.

Earlier tests were carried out on simulations using the full timestep (see section 8.1 for an explanation of this; the current code uses a halved timestep), with a cosine-squared mountain (see section 8.3 for a discussion of the two different mountain shapes) and a stretched grid in the vertical (of which more details can be found later in this chapter). The tests were only carried out for cases 1 and 3, using inflow/outflow boundary conditions, and a uniform horizontal grid spacing of 500 m seemed to generally produce noisy flow patterns. At first this was thought to be due to inadequate resolution over the mountain, but the mountain was 15700 m wide so there were over 30 horizontal gridpoints on its surface, indicating that this was not to blame. Simulations with a uniform grid of 285 m resolution were very slow, and the first data available in each case showed very pronounced noise at the outflow boundary, looking very similar to the 500 m resolution simulations. This did not look likely to produce useful results, especially considering the increased computational cost, so the simulations were stopped at this point. Further tests using the halved timestep gave similar results, although the noise was less severe. Use of a very small horizontal gridlength does seem to create problems, but it is unclear why. The spurious noise seen in these simulations could be due to the initial conditions, or it could be a genuine instability which grows to a certain extent, but is then diffused by the artificial viscosity. One way to investigate this (which has not been done due to time constraints) would be to interpolate the steady state solution obtained for the 1 km horizontal grid onto a 750 m grid, for example, and then restart the 750 m simulation using these interpolated fields. Some transience in the fields would be observed due to the change in resolution, but if the noise is due to the initial conditions, it will not appear in this simulation; if, on the other hand, there is an instability, that will be seen in this case as well.

Many of the problems seen here can be alleviated by using a stretched horizontal grid, providing fine resolution over the mountain and gradually coarsening resolution further away, and this is discussed in the next section.

6.2.2 Stretched grid

The need to use a fine enough grid to resolve the mountain and the flow of interest, while having a large enough domain to allow the flow to develop, means that the

uniform horizontal grid can be expensive in terms of computer memory and processing power. A stretched horizontal grid can fulfil this need using fewer gridpoints, thus reducing the computational expense.

A stretched grid is less numerically accurate than a uniform one, because the stretching introduces errors (see Castro and Jones, 1987 for a discussion of these errors). As the wave moves into a less well-resolved region of the model domain, the loss of accuracy can act as a layer of artificial viscosity and damp the wave maximum magnitude in the wave crests; in some cases it can damp the wave motion completely. If the stretching is done too quickly, it can result in spurious reflection and/or noise. Vichnevetsky (1987a) carried out a detailed analysis of spurious reflections on irregular grids. These non-uniform grids generally lead to spurious wave-like solutions, which cannot escape from the computing domain, and which can only be eliminated by some form of artificial dissipation (Vichnevetsky, 1987b). Vichnevetsky's analysis suggests that as a sinusoidal wave propagates downstream and the grid becomes gradually coarser, the group velocity decreases. If the group velocity vanishes before the wave has left the computational domain through the downstream boundary, then the group velocity becomes negative, and there is internal reflection of the wave. Vichnevetsky also considers what he calls a "grid well", where the resolution is finest in the centre of the domain and coarser further from the centre, which is the type of stretched grid used in BLASIUS. He notes that waves can become trapped in this well if they are reflected at the same point both upstream and downstream of the centre, and suggests that waves trapped in an irregular grid may be removed by adding artificial dissipation, to eliminate smaller scale numerical waves while retaining the larger scale waves which form the physical solution. Some of these issues are investigated in the results presented here, and also in the following chapter.

The first type of stretched horizontal grid to be tested was a relatively simple one, using a central region with higher resolution than the rest of the model domain. The full domain was 240 km wide, with the mountain crest in the centre of the domain. There were 290 points in the horizontal, and the grid was calculated using input parameters so that 70 points lay within a 20 km region centred on the mountain, giving 55 horizontal gridpoints on the surface of the mountain itself. In order to make the jump from one

grid area to the next less abrupt, BLASIUS offers two options for smoothing the horizontal grid. The user can either specify how many times to apply the 1-2-1 smoother, or specify a maximum expansion ratio, as described in Chapter 5. The former method was used for the grid described here, with the smoother applied 100 times. This produced a grid spacing of 285 m in the centre of the domain, increasing gradually to 1 km in the outer region, more than 20 km from the mountain crest. The grid was symmetrical about the centre of the domain.

Figure 6.6 shows the grid described above for a section of the model domain centred on the mountain. The full domain is not shown, since the grid would be hard to see, but the increased resolution in the central region is clearly visible. The cosine-squared mountain of height 500 m and width 15700 m is used in this example. This grid is referred to here as the ‘one stretch grid’.




Figure 6.6. The ‘one stretch grid’ described in the text, shown here for a partial domain centred on the mountain. A cosine-squared mountain of height 500 m and width 15700 m is used here to illustrate the grid. The numbers on the contours indicate the vertical model levels for the stretched vertical grid described in section 6.3.1.

The results presented here for the one stretch grid use inflow/outflow boundary conditions, a stretched vertical grid, the halved timestep and the cosine-squared mountain pictured in figure 6.6. They are compared with those using a uniform grid, as discussed at the end of the previous section. The tests were carried out for cases 1 to 3, and the results were encouraging: the reduced accuracy of the stretched grid seemed to

damp out the spurious noise found on the uniform grid. An example for constant U and constant N is shown in figure 6.7.

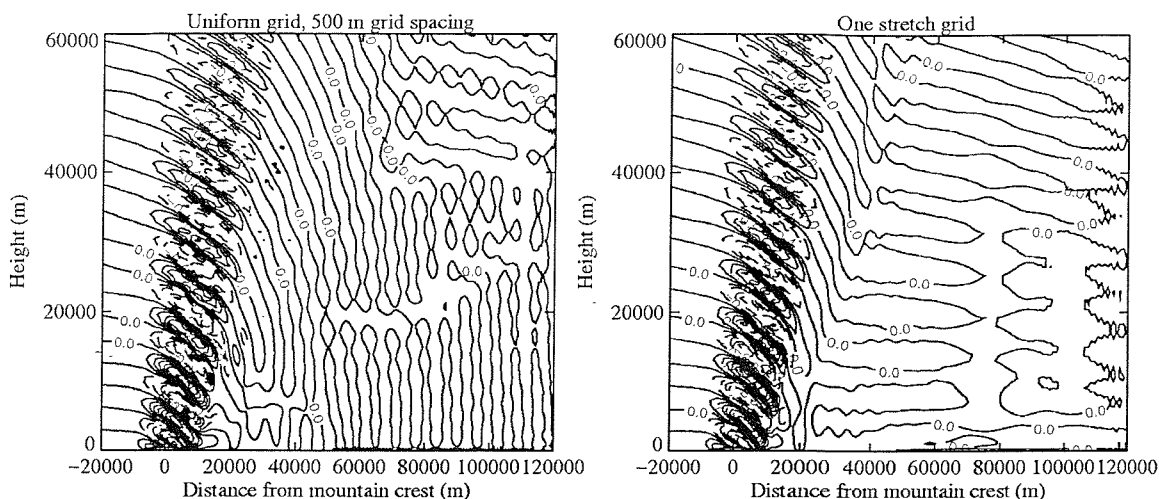


Figure 6.7. Vertical velocity plots for case 1 after 140 NDTU, using the uniform (500 m grid spacing) and one stretch grids. The contour interval for the vertical velocity is 0.2 m s^{-1} with a base contour of 0 m s^{-1} . The one stretch grid simulation does not have as much downstream noise as the uniform grid simulation. Both of these simulations used inflow/outflow boundary conditions upstream and downstream.

In case 2 the flow patterns were quite similar on both grids, but the maximum magnitude of the vertical velocity in the wave crests on the one stretch grid was larger near the mountain (with finer resolution) and smaller far from the mountain (with coarser resolution). This suggests that finer/coarser resolution resolves more/less of the wave pattern respectively. The drag coefficient was less noisy on the one stretch grid, reflecting the evolving pattern in the uniform grid simulation compared to the steadier flow on the stretched grid.

A comparison was carried out between these two grids for case 1 using periodic boundary conditions, to see how much of the spurious noise seen on the uniform grid was linked to the outflow boundary condition. The periodic simulations were far less noisy than their inflow/outflow counterparts, although the uniform grid simulation was still slightly noisier than the one stretch grid simulation. This indicates that the sensitivity to the horizontal grid was mainly, although not entirely, linked to the outflow boundary condition, and this is discussed further in Chapter 7.

The expansion ratio of a stretched grid was defined in Chapter 5 as $\Delta X(I+1)/\Delta X(I)$. The further this is from unity, the more rapid the stretching, and the more likely it is that there will be a negative effect on the accuracy of results, for example due to internal reflection. Castro and Jones (1987) have carried out analysis into the dependence of the magnitude of error terms on the expansion ratio. In the one stretch grid, the expansion ratio reaches a maximum of 1.15. Cases 1 to 3 were re-run with additional smoothing applied to the horizontal grid, so that the expansion ratio never exceeded 1.1. The expansion ratios and gridlengths for both these grids are shown in figure 6.8. The results were virtually identical, suggesting that the high expansion ratios are not the cause of the problems seen, although they may still be a contributing factor.




Figure 6.8. Horizontal gridlengths and expansion ratios for the one stretch grid (dashed line) compared with those for the additionally smoothed grid (solid line). The mountain crest is in the centre of the domain, and the grid is symmetrical about the centre of the domain. The expansion ratios in the upstream half of the domain are less than one because the gridlength is decreasing from left to right rather than increasing, as it does in the downstream half of the domain.

A set of tests was performed where the stretch point was moved from 10 km to 20 km away from the mountain crest, so the central region was 40 km wide in this case. Generally the flow patterns obtained were very similar to those using the one stretch grid, with wavelengths remaining the same. However, the maximum magnitude of the vertical velocity in the wave crests was generally stronger in the better resolved region between 10 and 20 km away from the mountain crest. This seems to indicate that there was a definite damping effect as the wave moved from the central region, where ΔX was 285 m, to the outer region, where ΔX was 1 km. These results suggest that better resolution is required in order to adequately represent the wave motion in these cases, although the uniform grid simulations had more noise at higher resolution, so it does seem that the stretched grid damps the spurious noise which is seen on a finely resolved uniform grid. This noise, as illustrated in figure 6.7, is believed to be spurious because in case 1 the wave is upwardly propagating, and no downstream signal is found in the analytical solution.

In order to investigate further the effect of horizontal resolution, cases 1 to 3 were re-run with halved horizontal resolution relative to the one stretch grid, so ΔX was 2 km in the outer region and 570 m in the centre of the domain. Generally the signal was similar and weaker, but the wave pattern seemed noisier, which seems to be a grid scale feature, indicating insufficient resolution. It should be noted that the noise referred to here is not the spurious noise seen outside the wave on a fine grid (such as in figure 6.7), but appears in the wave signal itself, where the gridpoints are clearly visible in the wave signal in this case, indicating that the resolution is not fine enough. The horizontally propagating waves in cases 2 and 3 were clearly not resolved once they had left the central region, as shown in figure 6.9 for case 3. In case 2, for example, the wavelength is approximately 14 km, so it was reasonably well resolved on the 1 km grid but not so well on this 2 km grid. In both these cases, the drag coefficient was steadier, but generally the wave pattern did not seem to be resolved and was weaker. Results obtained on the one stretch grid were already weaker than Keller's (1994) analytical results for case 2, so perhaps this was due to inadequate resolution.




Figure 6.9. Vertical velocity plot (contour interval 0.2 m s^{-1}) for case 3 after 140 NDTU, using the one stretch grid with halved resolution. The wave pattern was damped as it left the central region; the horizontal gridlengths are shown below for reference.

Both these sets of results indicate that the resolution needs to be better than the one stretch grid in order to represent the wave pattern adequately: the maximum magnitude of the vertical velocity in the wave crests increases and decreases with finer and coarser resolution respectively, and there is a marked change at the stretch point.

Starting from the one stretch grid, a second grid stretching was applied 70 km from the centre of the domain, and ΔX increased from 1 km to 2 km after this point. The domain size was unchanged, so only 240 horizontal gridpoints were required, instead of 290. This grid is referred to here as the 'two stretch grid'. The only significant differences between results obtained on the one stretch and two stretch grids were found in cases 3 and 4.

In case 3, noise was introduced in the upper level wave around the second grid stretch, which was not present on the one stretch grid. This noise at the second grid stretch may be due to stretching too quickly, but the expansion ratio is lower than in the first stretch, as shown in figure 6.10.




Figure 6.10. Horizontal gridlengths and expansion ratios for the one stretch grid (dashed line) compared with those for the two stretch grid (solid line). The mountain crest is in the centre of the domain, and the grid is symmetrical about the centre of the domain. The expansion ratios in the upstream half of the domain are less than one because the gridlength is decreasing from left to right rather than increasing, as it does in the downstream half of the domain. The expansion ratios at the second grid stretch are lower than those at the first stretch.

Using the two stretch grid, the wave signal almost disappeared at the second grid stretch, whereas with the one stretch grid, this 'whiteout' was more gradual and looked less artificial, as if the wave had actually faded and failed to reach the boundary, rather than being damped out by some numerical process which acts over a very small area (the second grid stretch in this case). As discussed in Chapter 5, this gradual fading may be explained by the no slip surface boundary condition. The waves had different wavelengths, and different maximum vertical velocity magnitudes in the wave crests. Figure 6.11 illustrates these differences in the flow pattern. For both case 2 and case 3,



a closer comparison showed that the flow patterns obtained using the one stretch grid were slightly closer to Keller's (1994) analytical results than those obtained with the two stretch grid, for the domain shown in her results.




Figure 6.11. Vertical velocity plots for case 3 after 140 NDTU, using the one stretch and two stretch grids. The contour interval for the vertical velocity is 0.2 m s^{-1} with a base contour of 0.1 m s^{-1} ; the zero contour has been removed from these plots, but is very smooth both upstream and downstream of the wave. Noise is introduced in the wave pattern at upper levels around the second grid stretch (70 km downstream of the mountain crest), and the wave signal does not reach the outflow boundary (120 km from the mountain crest) in the two stretch grid simulation, seeming to 'white out' at the second grid stretch. The wavelength and maximum vertical velocity magnitude in the wave crests are slightly different in the two simulations. The one stretch grid solution is closer to Keller's (1994) analytical results.

The simulation using the two stretch grid for case 4 had a 'bull's-eye' pattern in the vertical velocity at the outflow boundary after 7 to 10 NDTU. This noise was absent in the simulation using the one stretch grid, as shown in figure 6.12. In dimensional terms, these simulations ran for a similar length of time as the other cases, but because of the wider mountain, the non-dimensional time is far smaller. Time did not permit longer simulations of this case, but future work should take account of this factor, as discussed in Chapter 5. The results presented here for case 4 should be interpreted with caution, since they may well be less significant than those presented for the other cases, which ran for much longer in terms of non-dimensional time.

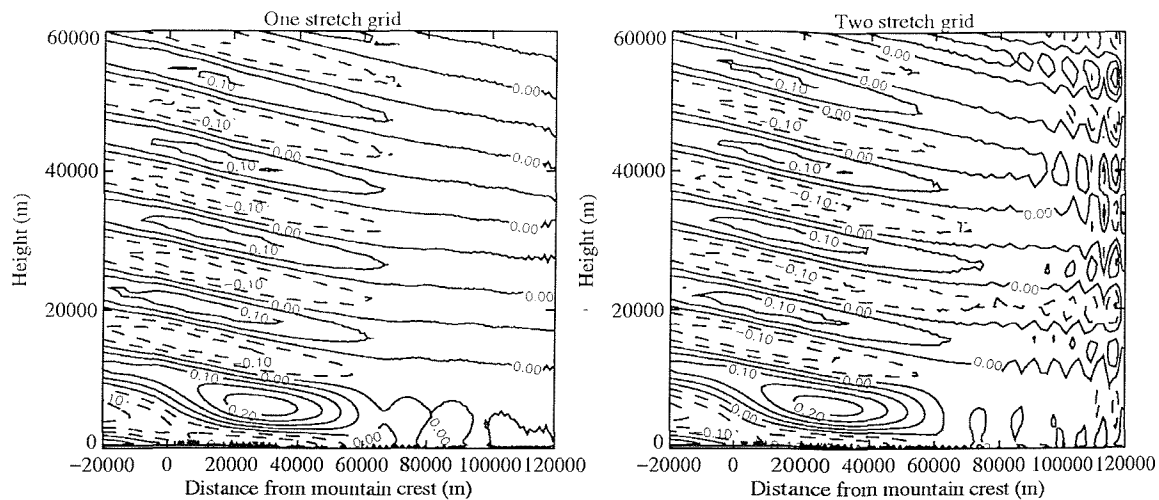


Figure 6.12. Vertical velocity plots for case 4 after 10 NDTU, using the one stretch and two stretch grids. The contour interval for the vertical velocity is 0.05 m s^{-1} with a base contour of 0 m s^{-1} . The ‘bull’s-eye’ signal at the outflow boundary in the two stretch grid simulation is absent in the one stretch grid simulation.

In some cases, the second grid stretch seemed to be causing spurious noise and/or numerical damping of waves. However, coarsening resolution, and reducing wave maximum magnitude in the wave crests, may be a useful device to mitigate problems at outflow, and this is discussed in Chapter 7. Figure 6.13 shows the drag coefficient for case 3 using different horizontal grids, and it is clear that the two stretch grid produced a much flatter, less noisy signal than the one stretch grid. This indicates that the simulation on the two stretch grid reached steady state, whereas the others did not. This may be because the second grid stretch seemed to act as a viscous layer, damping the wave so that it did not reach the outflow boundary, whereas on the one stretch grid, there may be a numerical reflection mechanism at the outflow boundary which causes the wave to come back to the mountain and alter the drag coefficient (since the surface pressure force is zero in the flat parts of the domain, so it only represents effects acting over the mountain itself). This is discussed further in Chapter 7.




Figure 6.13. Time series of the drag coefficient for the results presented in figure 6.11, with non-dimensional time on the horizontal axis, as defined in Chapter 5. The vertical long-dashed line marks the point where the non-dimensional time is zero, at the end of mountain growth. It is clear that the two stretch grid simulation was steadier than the one stretch grid simulation. The large oscillations at early times are most likely due to reflection of the initial noise by the outflow boundary condition, and this is discussed further in Chapter 7.

The 500 m high Witch of Agnesi mountain of half width 10 km was used in a further test, with constant U and constant N . This mountain is approximately four times wider than the cosine-squared mountain used above, so the central region of the one stretch grid was wider, extending 40 km either side of the mountain crest. Some noise appeared on the downstream edge of the wave, which seemed to be where the resolution changed, so this may indicate that the transition was not smooth enough and the expansion ratio should be smaller. Otherwise the flow was similar.

In view of the problems sometimes encountered around the stretching region, the one stretch grid was modified to use a single smooth stretch rather than two more abrupt ones either side of the mountain. The second grid stretch was re-introduced, because of the beneficial damping effect, but smoothed further to minimise the spurious noise. This grid was used in the simulations described in the previous chapter, with a 300 km domain (see Chapter 7 for tests of a wider domain with respect to the outflow boundary condition), and details were given there. The expansion ratios and gridlengths of this grid are shown in figure 6.14, along with those for the one stretch and two stretch grids. In general this grid produced smoother results than were obtained on the other horizontal grids, with less spurious noise.




Figure 6.14. Horizontal gridlengths and expansion ratios for the one stretch grid (dashed line) and the two stretch grid (dotted line), compared with those for the more smoothly stretched grid (solid line). The mountain crest is in the centre of the domain, and the grid is symmetrical about the centre of the domain. The expansion ratios in the upstream half of the domain are less than one because the gridlength is decreasing from left to right rather than increasing, as it does in the downstream half of the domain. The new grid uses a wider domain (300 km) than either of the other two grids (240 km).

6.3 Vertical grid dependence

6.3.1 Vertical resolution

The vertical grid can be uniform or stretched. When stretching the horizontal grid, the focus is on adequate resolution of the mountain. In the vertical, the boundary layer is the part which requires the finest resolution, since the flow there is on smaller scales and is more complex. Whilst this is not relevant in the simulations presented here, BLASIUS was intended to simulate these boundary layer flows, and so the vertical gridlength is smallest at the surface and largest at the top of the model domain. The input parameters, in addition to the domain height and the number of vertical gridpoints, are Z_l , the height of the lowest internal model level, and R_l , the expansion ratio (defined in the same way as for the horizontal grid). The heights of the model levels are then determined using the following method:

$$a_1 = -Z_l$$

$$a_2 = 0$$

$$a_3 = Z_l$$

...

$$a_n = a_{n-1} + R_l (a_{n-1} - a_{n-2})$$

so that the distance between successive vertical model levels increases in geometric progression from the bottom model gridpoint. This continues until the remaining distance from a_n to the top of the domain can be filled using the rest of the levels with a constant grid spacing $(a_n - a_{n-1})$. The Z_N levels are taken as a_n when n is odd, and the Z levels as a_n when n is even; more details of the staggered grid were given in Chapter 4. No smoothing is applied to the vertical grid.

A test was carried out using the constant U and constant N case with a 100 m high Witch of Agnesi mountain of half width 10 km. One simulation used a uniform vertical grid with $\Delta Z = 350$ m, and the other a stretched vertical grid with $Z_l = 15$ m and $R_l = 1.025$, giving $\Delta Z = 30$ m at the bottom of the domain and 633 m at the top. The

Rayleigh damping layer started at 20 km in both cases, and the horizontal grid had a uniform gridlength of 1 km. Figure 6.15 illustrates the two different vertical grids.




Figure 6.15. The different grids for the two simulations presented in figure 6.16. The top plot shows the uniform vertical grid, and the lower plot shows the stretched vertical grid. Note that these plots show a very small section of the vertical domain, so that the stretched grid levels are clearly visible.

The results were very similar: both simulations reached steady state fairly quickly and had a fairly flat drag coefficient signal. The vertical velocity signal was similar in both, but there were clear differences where the vertical resolution was most different: figure 6.16 shows that the stretched grid simulation had less smooth contours at upper levels, and there were also small differences at the lowest vertical levels, very close to the mountain, although these are barely discernible in figure 6.16. The vertical wavelength is approximately 7 km, so the uniform gridlength of 350 m gives around 20

points per wavelength, whereas on the upper part of the stretched grid, where the spacing is 633 m, there are only around 11 points per wavelength. Although this means that the wave is still reasonably well resolved, there are small differences in the results.




Figure 6.16. Vertical velocity after 75 NDTU for a 100 m high Witch of Agnesi mountain of half width 10 km. The contour interval for the vertical velocity is 0.01 m s^{-1} with a base contour of 0.005 m s^{-1} ; the zero contour has been removed from these plots, but is very smooth both upstream and downstream of the wave. The left hand plot has a uniform vertical grid spacing of 350 m; the right hand one has a stretched vertical grid, as described in the text. These simulations ran for the same dimensional time as most of the others, but a shorter non-dimensional time, because they used a wider mountain. However, since they reached steady state fairly quickly, this should not be significant.

Tests were also carried out using a 500 m high cosine-squared mountain of width 15700 m for cases 1 to 3. In the vertical, for cases 1 and 3, 80 gridpoints were used in a 75 km domain, and the grid spacing increased from 253 m at the surface to 1211 m at the top of the domain. This grid is depicted for part of the domain in figure 6.6, with the one stretch horizontal grid, which was used in these tests. The Rayleigh damping layer began at 60 km, so results are only presented up to that height. This grid gave two vertical levels below the mountain crest height, at the surface and at 253 m. The third vertical grid level was at a height of 519 m, just above the mountain crest. The vertical resolution of the mountain itself is less important here than in a boundary layer case, so there were few points below the mountain top. The heights given for the vertical grid levels refer to the distance of the model level above the flat surface upstream and downstream of the mountain, and do not take account of the bent mesh.

Case 2 had a different vertical grid from the other simulations, since it was intended to simulate the troposphere-only case of Keller, and had a smaller vertical domain. A total of 50 gridpoints were used in a 45 km domain, and the grid spacing increased from 253 m at the surface to 1744 m at the top of the domain. The Rayleigh damping layer began at 35 km. This grid is not shown, since it was similar to the one shown in figure 6.6. The first vertical grid level was at the surface, and the next one, at 253 m, was above the mountain crest in this case.

This vertical resolution was coarsened for these tests by doubling Z_1 and squaring R_1 , so that ΔZ was approximately twice as large throughout the domain. In case 1, the simulation with coarser vertical resolution had more marked noise in the vertical velocity, which grew with time. The wave pattern itself was similar, but individual vertical levels were visible in the wave pattern where it was only partially resolved. These can also be seen to a lesser extent in the original simulation (the two are compared in figure 6.17), indicating that even that resolution was not really good enough. The vertical wavelength was approximately 7 km, and the vertical resolution in the original grid increased to 1211 m, giving around six gridlengths (seven points) per wavelength at upper levels. This is close to the minimum required, which is generally considered to be five gridpoints (four gridlengths) per wavelength, as shown in figure 6.18. On the coarsened grid, the number of gridlengths per wavelength was reduced to three, which is clearly inadequate.

The drag coefficient increased significantly towards the end of the simulation, whereas it had almost levelled out in the original, indicating that coarsening the vertical grid made the solution less steady. Cases 2 and 3 showed less marked differences, but the inadequacy of the vertical resolution at upper levels was clear. These results indicate that the original vertical resolution is required in order to represent the wave pattern adequately, and that better vertical resolution is certainly required for case 1, if not for the others.

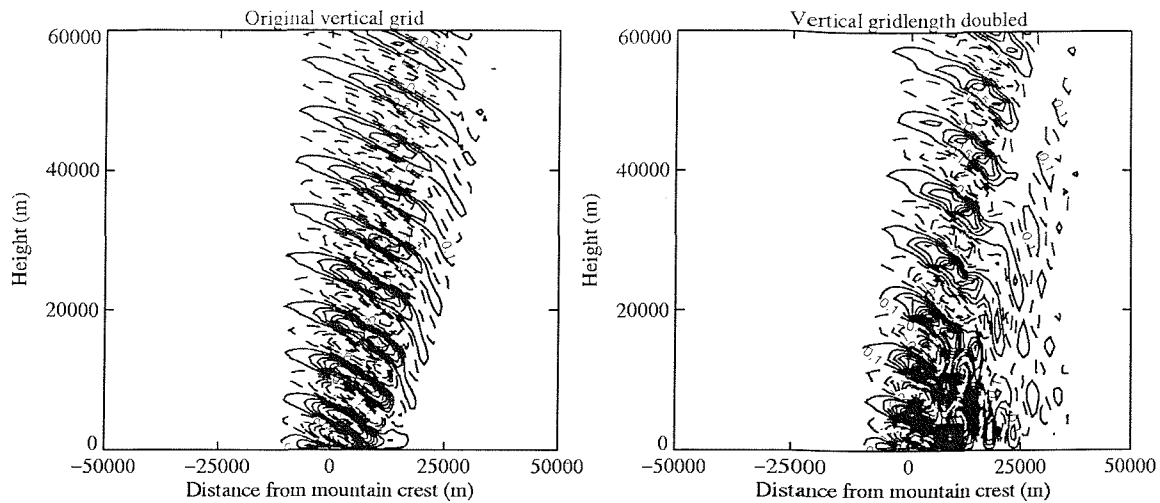


Figure 6.17. Vertical velocity after 140 NDTU for a 500 m high cosine-squared mountain of width 15700 m. The contour interval for the vertical velocity is 0.2 m s^{-1} with a base contour of 0.1 m s^{-1} ; the zero contour has been removed from these plots, but is very smooth both upstream and downstream of the wave. The left hand plot has the stretched vertical grid described in the text, and the right hand one has the same grid but with approximately double the gridlength.




Figure 6.18. Schematic representation of a wavelength in a discretised numerical model, using five gridpoints (four gridlengths). The solid line is the true wave, and the dashed line is the resolved wave, with the gridpoints marked by asterisks. It is clear that using fewer gridpoints would not adequately resolve the basic structure of the wave.

6.3.2 Rayleigh damping scheme

In order to simulate an infinite vertical domain, a Rayleigh damping layer is used at the top of the model domain to prevent reflection from the upper boundary. This layer gradually relaxes the model fields back to a reference state; more detail was given in Chapter 4. If the damping acts too quickly, it can lead to reflection from the base of the damping layer. Results presented in section 6.2.1 showed spurious noise in simulations with small ΔX , and one explanation for this may be reflection from the base of the Rayleigh damping layer, or indeed from the top of the domain. The vertical resolution was not varied between simulations in the tests of the horizontal grid, so it may be that any reflected signal was not resolved on the coarser horizontal grid.

In the simulation with 500 m uniform horizontal grid spacing, earlier data dumps were obtained during the mountain growth, to see whether this helped to isolate the cause of the noise. The initial noise reached right to the top of the domain after -70 NDTU, and spread sideways with time, gradually reaching less high up as the Rayleigh damping took effect. The negative non-dimensional time indicates that the mountain was still growing at that point in the simulation. The structure of the noise became finer with time, with more upward and downward cells in the vertical velocity field. However, it was hard to judge what was causing the problems.

In the tests of the coarsened stretched vertical grid described in the previous section, the Rayleigh damping layer took longer to damp out the wave motion when using a larger vertical gridlength. The Rayleigh damping scheme has no explicit dependence on ΔZ , but clearly the vertical resolution does have an indirect impact. A plausible explanation of this is that the coarser grid produced more marked noise in the vertical velocity field, meaning that the Rayleigh damping scheme had more work to do in order to damp out the signal. This implies that the Rayleigh damping scheme is implicitly dependent on ΔZ , due to the explicit dependence on ΔZ of the vertical velocity response.

Since problems were seen even with a 1 km horizontal grid spacing for the 500 m high Witch of Agnesi mountain of half width 5 km (as discussed in section 6.2.1), tests were carried out using different damping layers for that case. First of all, the damping layer

was moved up by 5 km, so that its base was at 25 km and the domain height was 40 km. The number of vertical gridpoints was increased to 117 in order to keep the same vertical resolution. The wave obviously reached higher into the domain with a higher Rayleigh damping layer, but the basic pattern, wavelength and maximum magnitude in the wave crests were the same in both simulations. The potential temperature and drag coefficient signals were also very similar. There was more initial noise in the higher Rayleigh damping layer simulation, but as it faded with time it looked more similar to that seen in the original simulation, and the noise on the downstream edge of the wave seemed more marked initially, but became less intense with time relative to the original simulation, and was more spread out, as shown in figure 6.19.

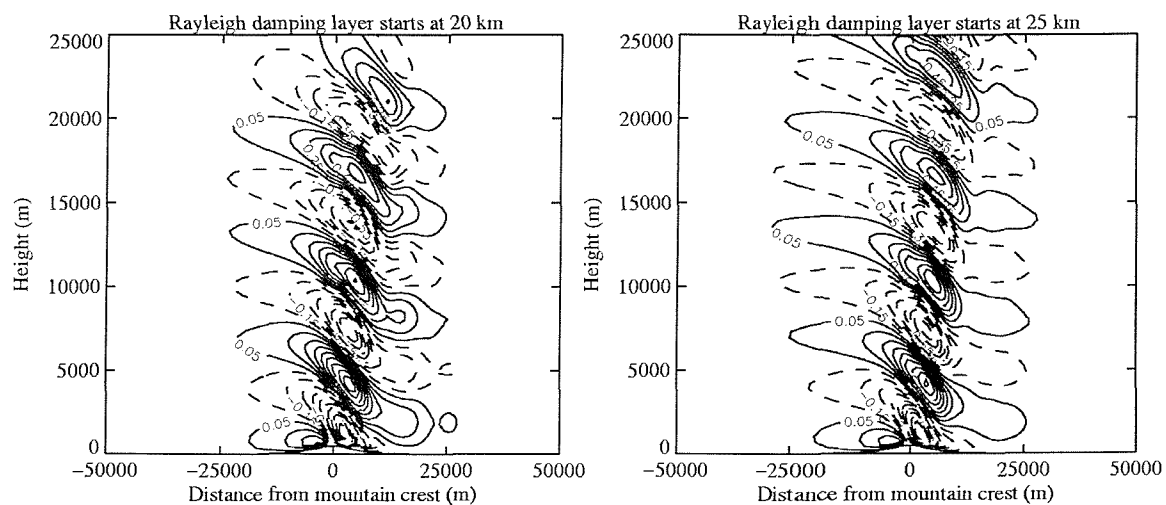


Figure 6.19. Vertical velocity after 250 NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km. The contour interval for the vertical velocity is 0.1 m s^{-1} with a base contour of 0.05 m s^{-1} ; the zero contour has been removed from these plots, but is very smooth both upstream and downstream of the wave. The Rayleigh damping layer is between 20 km and 35 km in the left hand plot, and between 25 km and 40 km in the right hand one.

A further test with the damping layer moved up a further 10 km, and $\Delta Z = 400 \text{ m}$ (a small increase on the original 350 m), gave generally a very similar wave pattern, except for the obvious difference that the wave extended higher up. Although there were some differences, they were not clearly positive or negative, so moving the damping layer still further away did not seem to change a great deal.

The Rayleigh damping coefficient (explained in Chapter 4) is set to 0.01 s^{-1} by default. Test simulations were carried out increasing and decreasing this value by an order of magnitude, and the results are shown in figure 6.20. In the case where the coefficient was 0.1 s^{-1} , the signal reached a little less into the damping layer but was basically the same. The noise on the downstream edge of the wave, present at all vertical levels below the height where the Rayleigh damping layer took effect, was less here than in the original. In the case where the coefficient was 0.001 s^{-1} , the signal reached the upper boundary throughout, and the downstream wave edge noise was less pronounced. The drag coefficient was initially noisier in both these simulations than in the original, but it increased in the same way and had a similar pattern, although the magnitude was slightly lower in both cases.




Figure 6.20. Vertical velocity after 250 NDTU for a 500 m high Witch of Agnesi mountain of half width 5 km. The contour interval for the vertical velocity is 0.1 m s^{-1} with a base contour of 0.05 m s^{-1} ; the zero contour has been removed from these plots, but is very smooth both upstream and downstream of the wave. The Rayleigh damping layer is between 20 km and 35 km in all cases; the damping coefficient is 0.001 s^{-1} in the left hand plot, 0.01 s^{-1} in the centre plot and 0.1 s^{-1} in the right hand plot.

Under certain circumstances partial reflection at the upper boundary can affect the surface drag. Vosper (1995) found that the upper boundary is truly radiative only when the mean wind does not vary with height for a considerable depth below the upper boundary. Although the upper boundary in BLASIUS uses an absorbing layer rather than a radiative boundary condition, a similar result may hold for the Rayleigh damping layer. This could explain why the drag coefficient in case 2 has larger amplitude variations than the other cases, as seen in figure 5.7, because the shear in the background wind speed continues all the way to the top of the model domain. A useful test (not done as part of this work because of time constraints) would be to relax the

background wind field to a constant profile in the Rayleigh damping layer, although this would make it more like a troposphere-stratosphere profile than a troposphere-only one, so some care would be required to ensure that the resulting wave pattern was not altered as a result. The upper boundary condition has less effect in trapped wave cases like this one, because less wave energy reaches the top of the domain, so this may not have a significant impact in case 2, but it merits investigation.

These sensitivity tests on the Rayleigh damping coefficient and on the height of the damping layer have so far proved inconclusive, so there may still be an issue with reflection from the base of the Rayleigh damping layer. Ideally further investigation will lead to a non-dimensional formula for choosing the heights of the damping layer and domain top, and the damping coefficient. The Rayleigh damping layer should, in theory, permit any domain height without affecting the solution, thus allowing good vertical resolution on a uniform grid, which can be prohibitively expensive in the horizontal due to the domain length required to resolve the wave motion.

6.4 Summary

These investigations have shown that a uniform horizontal grid with resolution fine enough to resolve the wave patterns studied here tends to introduce spurious noise into the flow. This noise seems to be linked to a sensitivity of the outflow boundary condition and/or the Rayleigh damping scheme to horizontal resolution, but further work is required to isolate the cause of the problem. Running the model for simulations like these but with no Rayleigh damping layer would be a good test.

A stretched horizontal grid is numerically less accurate, and this acts to damp the spurious noise mentioned above, allowing adequate resolution of the centre of the domain, and gradually coarsening resolution far from the mountain. Care must be taken not to introduce noise by stretching too quickly, and not to damp the wave motion itself

as the gridlength increases (unless this is desirable in order to avoid reflection by the downstream boundary; this is discussed further in the next chapter).

One way to reduce the computational cost of the horizontal grid would be to move the mountain from the middle of the domain. This is straightforward in the case of the uniform grid, although some testing would be required to assess sensitivity to upstream fetch, and this is unlikely to be a solution to the spurious noise found at fine resolutions. On the stretched grid, substantial re-coding would be required in order to have the most finely resolved section off-centre, over the mountain, but this could be worthwhile.

A uniform grid in the vertical presents fewer problems than in the horizontal. It seems possible to resolve the wave motion adequately without any of the spurious noise introduced by a fine horizontal grid. A stretched vertical grid can be used, and this will be more important in boundary layer simulations, where fine resolution is required at low levels. Care is required to ensure that the resolution remains adequate at upper levels, but no other problems have been found in these tests.

The Rayleigh damping layer would benefit from further investigation to determine the optimum tuning (layer height and depth, and damping coefficient) for a given flow regime. It is possible that the scheme is causing some problems due to reflection, especially on a finely resolved uniform horizontal grid, but further work is required to test this hypothesis.

7 Inflow/outflow boundary conditions

7.1 Introduction

The radiative outflow boundary condition was explained in section 4.4. This chapter describes the tests carried out to assess the accuracy of this boundary condition as implemented in BLASIUS. The inflow/outflow boundary conditions allow the model to simulate steady state cases with a single isolated mountain. Although in practice periodic boundary conditions can be used for some of these cases, the upstream flow is then not truly steady, and the problem being solved is one of an infinite series of widely spaced mountains, rather than a truly isolated one. Even in the case where the mountain excites a hydrostatic wave which does not propagate horizontally, and there is no discernible difference between the periodic and inflow/outflow solutions, the horizontal velocity profiles at the upstream boundary in the periodic case do show a slight wave pattern. Periodic boundary conditions are used in the y -direction for all of the two-dimensional x - z simulations in this thesis, so all mention of periodic or inflow/outflow boundary conditions in this chapter should be taken as referring to the upstream and downstream boundaries in the x -direction.

The radiative outflow boundary condition is designed to allow waves to propagate freely out of the model domain without being disturbed. In some of the simulations described in this thesis, the flow seems to be held up at the boundary, and does not exit the domain cleanly. In some cases there is evidence to suggest that the outflow boundary may be reflecting the wave back into the domain. Vichnevetsky (1987a) discusses analysis of such spurious reflections at both the upstream and downstream

boundaries. The strongest indication of this in the work presented here is seen with a horizontally propagating wave such as the one in the troposphere-stratosphere case, where a marked perturbation occurs in the drag coefficient during spin-up, shortly after the wave head reaches the outflow boundary.

Section 7.2 describes tests of the accuracy of the outflow boundary condition, by comparing with results obtained using periodic boundary conditions, and with results obtained using a wider horizontal domain. Section 7.3 describes alternative numerical schemes for the outflow boundary condition and their impact on the test cases. Section 7.4 summarises these results and discusses possible further investigations, and ways to minimise the effects of the inadequacies in the outflow boundary condition as it stands.

7.2 Tests of the radiative outflow boundary condition

7.2.1 Periodic boundary conditions

Periodic boundary conditions provide a test of the inflow/outflow boundary conditions, since the flow pattern should be the same in both cases until any horizontally propagating wave signal (or initial noise, which is discussed further in section 8.4) re-enters the periodic domain upstream and contaminates the solution. Differences in surface gridlength waves are ignored in the comparisons below, since the general tendency was for them to occur upstream in the periodic simulations rather than downstream in the inflow/outflow simulations. Tests were carried out for cases 1 to 3 using the one stretch horizontal grid with a domain width of 240 km, and the stretched vertical grid described in Chapter 6, with a cosine-squared mountain of width 15700 m. In cases 1 and 3 the mountain height was 500 m, and in case 2 it was 100 m. As expected, the greatest differences were seen in cases 2 and 3, where there is a horizontally propagating wave.

In case 2 the flow was generally less noisy in the periodic simulation than in the inflow/outflow one. The periodic contamination reached the mountain after 80 NDTU, and by the end of the simulation, after 140 NDTU, the wave pattern at lower levels was similar, but in the periodic case it was weaker and with shorter wavelength, whereas the upper section of the wave had almost disappeared. These differences, illustrated in figure 7.1, are most likely simply due to the effect of the upstream contamination, which seemed to interact destructively with the wave as it passed through the domain, with the result that both the wave and the initial noise were reduced in magnitude. The qualitative agreement with Keller's (1994) results after 60 NDTU (just before the periodic contamination affected the flow) was slightly better for the periodic simulation than for the inflow/outflow one.

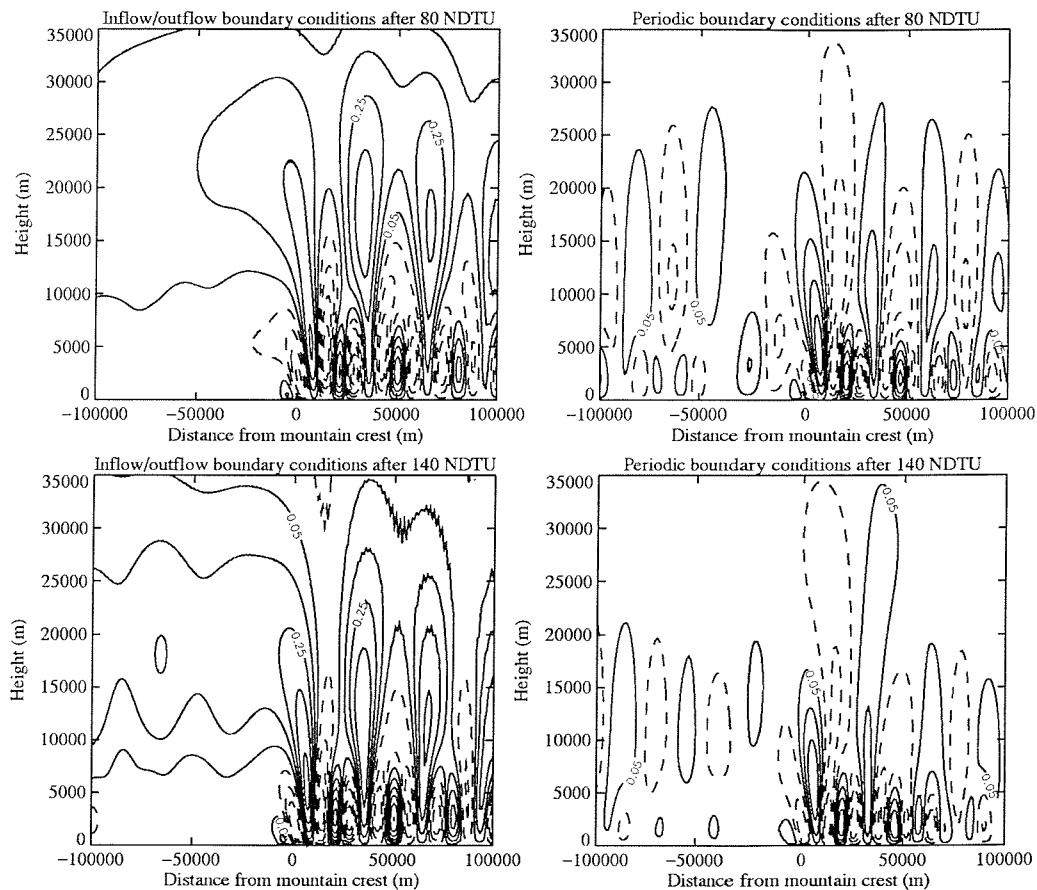


Figure 7.1. Vertical velocity after 80 NDTU (top plots) and 140 NDTU (bottom plots) for case 2. The contour interval is 0.1 m s^{-1} with a base contour of 0.05 m s^{-1} ; the zero contour has been removed from these plots, but the flow pattern is still clearly visible. The left hand plots used inflow/outflow boundary conditions, and the right hand plots used periodic boundary conditions. After 80 NDTU, the periodic contamination in the upstream domain reached the mountain, and the wave signal was weaker in the periodic simulation after 140 NDTU, especially at upper levels.

The periodic simulation for case 3 was also less noisy than the inflow/outflow version. At the end of mountain growth, the initial noise had almost left the downstream domain in the periodic simulation, but was held up in the inflow/outflow simulation, as shown in figure 7.2. In the periodic simulation, the initial noise re-entered the domain at the upstream boundary and reached the mountain after 20 NDTU (sooner than in case 2, because of the stronger forcing from the higher mountain). The upper wave was disrupted by this throughout the simulation, as this contamination passed through the domain several times. After 100 NDTU, the upper level wave had almost no signal in places, as shown in figure 7.2, and although it settled slightly towards the end of the simulation, after 300 NDTU, it was still very weak. The lower level wave was similar in both simulations, with similar wavelength. It slowly entered the upstream domain in the periodic simulation from 60 NDTU, and seemed to be reasonably steady from then on, fading much more gradually downstream than in the inflow/outflow case with a less marked drop in amplitude with time, and only extending about 25 km into the upstream end. This can be seen in figure 7.2, which shows the vertical velocity fields after 100 NDTU. The smaller wave amplitude at lower levels in the inflow/outflow case may be due to some form of interaction or reflection at the outflow boundary. The drag coefficient was very flat, without the large perturbation between 0 and 20 NDTU which was present in the inflow/outflow simulation; this flat signal is linked to the very steady lower level wave in this case. It was difficult to compare these results with Keller's, because the periodic contamination was significant in this case.

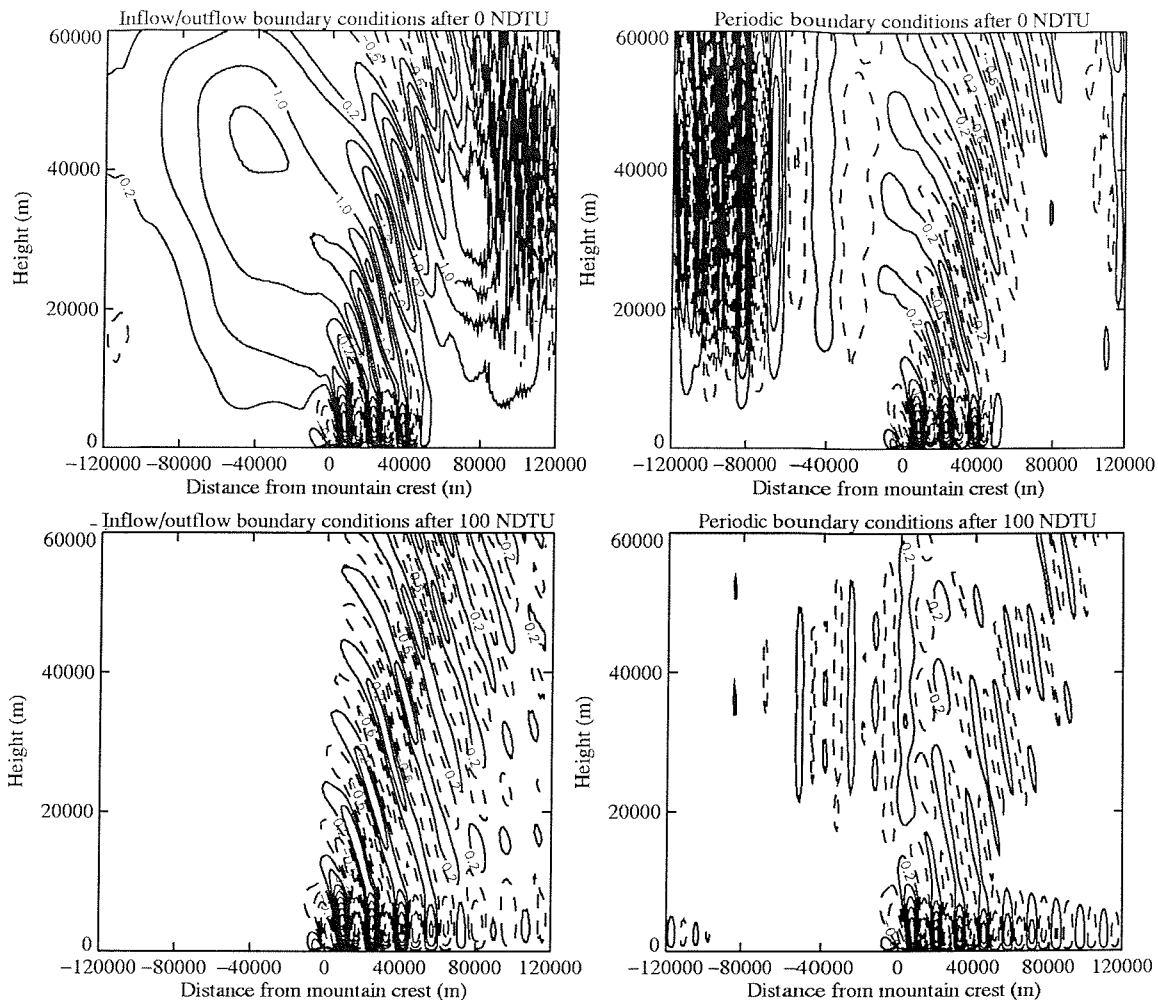


Figure 7.2. Vertical velocity after 0 NDTU (at the end of mountain growth; top plots) and 100 NDTU (bottom plots) for case 3. The contour interval is 0.2 m s^{-1} with a base contour of 0.1 m s^{-1} ; the zero contour has been removed from these plots, but the flow pattern is still clearly visible. The left hand plots used inflow/outflow boundary conditions, and the right hand plots used periodic boundary conditions. The initial noise seems to be held up at the outflow boundary at the end of mountain growth in the inflow/outflow simulation. After 100 NDTU, the wave signal at upper levels is weaker in the periodic simulation, and at lower levels the wave has a shorter wavelength and fades less downstream.

Periodic simulations for case 1 were carried out using both the one stretch horizontal grid (as for cases 2 and 3 above) and a uniform horizontal grid, with a grid spacing of 500 m. As for cases 2 and 3, the flow was generally less noisy in the periodic simulations than in the inflow/outflow ones, on both horizontal grids. In the one stretch grid case, there were only very small differences until 20 NDTU, when the upstream contamination began in the periodic simulation. The initial noise was different in the uniform grid simulations: by 20 NDTU it had entered the upstream periodic domain, but it seemed to be held up in the inflow/outflow simulation. The effect was similar to that

shown in figure 7.2 for case 3, although less marked. By 60 NDTU the upstream contamination was affecting the wave on both horizontal grids, and by 120 NDTU it had passed through the wave, and the wave looked similar in both the inflow/outflow and periodic simulations. The drag coefficient still increased slightly in the periodic simulations, but was generally flatter, without the large perturbation after 40 NDTU which was present in both the inflow/outflow simulations.

In the case 1 tests described above, the initial noise was held up at the outflow boundary in the uniform grid inflow/outflow simulation, but this was not the case in the one stretch grid inflow/outflow simulation, or in either of the periodic simulations. The uniform grid had a resolution of 500 m, whereas the one stretch grid had a resolution of 1 km in the outer region. These results could indicate a sensitivity of the outflow boundary condition to horizontal resolution, a hypothesis which was discussed in Chapter 6 following grid sensitivities found there. There may also be some sensitivity to the grid formulation (stretched versus uniform), since the change from one to the other had more effect in the inflow/outflow case than in the periodic case, and these two factors may be interacting.

The periodic simulations generally had a flatter drag coefficient signal than their inflow/outflow counterparts, indicating that they were more steady. The most notable thing about the drag coefficient is that the periodic simulations did not have the large amplitude perturbation after approximately 40 NDTU, shortly after the initial noise reached the boundary and was sometimes held up in the inflow/outflow simulations. This could indicate that the outflow boundary condition was causing the initial noise to reflect back into the domain, and a similar effect was seen when the horizontally propagating wave head reached the boundary in cases 2 and 3. This is discussed further later in this chapter.

A test was also carried out for the constant U , constant N case with a 500 m high Witch of Agnesi mountain of half width 10 km, and uniform grids in the horizontal and vertical direction with grid spacing 1 km and 350 m respectively. There was little difference between the inflow/outflow and periodic simulations for this case, which is unsurprising, since the forcing was quite weak and there was little initial noise.

7.2.2 Wider horizontal domain

Orlanski (1976) performed tests of the outflow boundary condition formulation by running the model with a wider horizontal domain, and then comparing the results for the original, smaller, domain between this simulation and the original, to see whether the outflow boundary condition had caused any problems in the original simulation. The same procedure was followed here in order to assess the accuracy of the outflow boundary condition in the BLASIUS model.

Tests were carried out for cases 1 to 3 in the same configuration as for the results described in the previous section. The wider horizontal domain, of width 300 km, was created from the one stretch horizontal grid, of width 240 km, by adding 30 km to each end, with a resolution of 1 km, so that the grid in the original domain was unchanged. Vertical velocity plots for the full domain gave information on whether the boundary condition was more effective for being further away from the mountain.

Wider domain simulations were carried out for both inflow/outflow and periodic boundary conditions, in order to isolate any effects which were directly due to the increased domain size rather than to the boundary conditions. The only significant differences in these tests were due to the longer time interval before the initial noise re-entered the periodic domain upstream. In theory, the initial noise should be exactly the same irrespective of where the downstream boundary is located, since it is generated by the process of growing the mountain into the model domain. In the periodic simulations this was indeed the case. In order to be confident that the outflow boundary condition is coping adequately, a similar result is desirable in the inflow/outflow simulations.

In case 1, using the wider domain, the inflow/outflow and periodic results were closer than they had been on the original domain, but the initial noise still seemed to be held up at the outflow boundary in both of the inflow/outflow simulations. They were also generally noisier than their periodic counterparts, and the noise at the outflow boundary seemed more severe on the original domain than on the wider domain (although it is not shown here, since it was generally restricted to the zero contour). The wave itself was

very similar in all the simulations, except when affected by periodic contamination. The downstream surface gridlength wave in the inflow/outflow simulations started later in the wider domain simulation, indicating that it may be caused or amplified by the outflow boundary condition. After about 75 NDTU, the drag coefficient signal levelled out far more in the periodic simulations than in the inflow/outflow ones, but still increased slightly, finishing at a lower value. The wider domain did not make a difference to the final value of the drag coefficient in either the inflow/outflow case or the periodic case, which is encouraging. Figure 7.3 shows the drag coefficient for all four simulations.

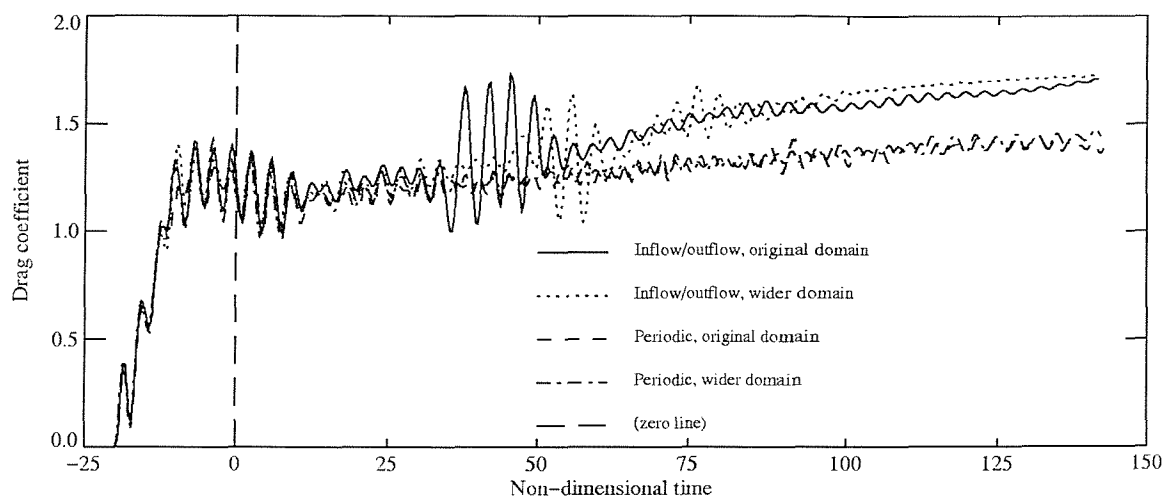


Figure 7.3. Time series of drag coefficient for case 1, using inflow/outflow and periodic boundary conditions, and the original and wider horizontal domain (as described in the text). The vertical long-dashed line marks the point where the non-dimensional time is zero, at the end of mountain growth.

The large perturbation seen in the drag coefficient in the inflow/outflow simulations was later in the wider domain simulation than in the original domain simulation. An approximate calculation of the expected time delay in this perturbation, assuming that it is caused by the reflection of the initial noise by the outflow boundary condition, is given by

$$\delta t \approx \frac{\delta X}{U}$$

where here δt is the time delay in the drag coefficient perturbation, δX is the extra distance travelled as a result of the wider domain (60 km, since the extra 30 km must be traversed once in each direction), and U is the propagation speed of the initial noise

(estimated by eye as 12 m s^{-1} , since the signal reached the edge of the original domain after approximately 10000 seconds). This formula gives an estimated time delay of 5000 seconds, or 20 NDTU, which fits well with the delay seen in figure 7.3. This large perturbation is not evident in the periodic simulations, although it might be expected to appear at around the same time, since the initial noise travels the same distance to reach the mountain through the upstream domain in the periodic case as it does through the downstream domain when reflected in the inflow/outflow case. A possible explanation for this is that the reflection is quite likely to be distorted or incomplete, with some components of the noise passing through the boundary and others being reflected, so the signal reaching the mountain in the inflow/outflow case is probably not the same as the one which reaches it in the periodic case.

In case 2 the results were all very similar, with the same wavelength and only small differences in wave amplitude. However, in case 3, the inflow/outflow simulations were generally noisier than the periodic simulations. At the end of mountain growth, the initial noise still seemed to be held up in the wider domain inflow/outflow simulation, but it had moved further downstream than in the original domain inflow/outflow simulation, indicating that the outflow boundary was acting to retard the flow, and was not allowing it to leave the model domain undisturbed. After 80 NDTU, the wavelength (measured along the wave train) at both upper and lower levels was shorter in the wider domain inflow/outflow simulation than in the original one, and the signal was generally weaker. The original simulation had an irregular structure in the upper wave downstream of the mountain, remote from both major wave trains. This was not noticeable in the wider domain simulation until a later time, and this delay may indicate that the phenomenon is due to reflection at the outflow boundary. These differences are shown in figure 7.4, and it is evident that the wave signal had not yet reached the outflow boundary in the wider domain simulation.

The drag coefficient was flatter towards the end of the wider domain inflow/outflow simulation than in the original domain simulation, although still increasing, as shown in figure 7.5. Again, the large perturbation appeared at a slightly later time in the wider domain inflow/outflow simulation than in the original domain case, and the above calculation for the delay in the perturbation gives 1000 seconds this time, because the

initial noise propagates much faster. This is equivalent to 4 NDTU, which again fits well with the delay seen in the plot of the drag coefficient.




Figure 7.4. Vertical velocity after 80 NDTU for case 3. The contour interval is 0.2 m s^{-1} with a base contour of 0.1 m s^{-1} ; the zero contour has been removed from these plots, but the flow pattern is still clearly visible. Both these simulations used inflow/outflow boundary conditions; the left hand plot used the original horizontal domain, and the right hand one used the wider one (as described in the text). The domain plotted here extends to the outflow boundary of the wider domain, 150 km downstream of the mountain crest. Note the irregular structure in the wave on the original domain, as if it is being reflected from the outflow boundary.




Figure 7.5. Time series of drag coefficient for case 3, using inflow/outflow and periodic boundary conditions, and the original and wider horizontal domain (as described in the text). The vertical long-dashed line marks the point where the non-dimensional time is zero, at the end of mountain growth.

Tests were also carried out using the two stretch horizontal grid, and here the wider domain simulations had 15 gridpoints added onto each end, with the same resolution (2 km) as the outer section of the existing domain, making the domain width 300 km as in the above tests. These tests were carried out on simulations using the full timestep.

In case 3 the major difference between the two simulations was a clear improvement in the flow near the second grid stretch in the wider domain simulation, which became more noticeable with time as the spurious noise worsened in the original case. These results are not shown, since they are very similar to those shown in the bottom row of figure 8.1. The drag coefficient did not have the oscillation which was present at the end of the original simulation, showing that it was closer to a steady state. The full domain vertical velocity plots were similar at the boundaries; notably, the whiteout still occurred in the same place, indicating that this phenomenon is more likely to be due to the change in horizontal resolution at the second grid stretch than to the outflow boundary condition, and this was discussed further in Chapter 6.

In case 4, the full domain vertical velocity plots showed very similar flow at both boundaries until 7.5 NDTU, then the original simulation developed outflow noise which did not appear in the wider domain simulation until 10 NDTU, as shown in figure 7.6. These results indicate that the boundary condition was coping badly in the original simulation, since the problem was still present in the wider domain simulation, but delayed because the flow had further to travel before reaching the boundary. It is more difficult to estimate an expected time delay for this case than for the delayed perturbations in the drag coefficient signals discussed above, because there is no clear signal in the earlier fields before the problem at the boundary becomes apparent. However, it is clear from figure 7.6 that the noise is very likely to be due to some sort of reflection issue at the outflow boundary. This case ran for much less non-dimensional time than the others, as already discussed in Chapters 5 and 6, so these results must be interpreted with caution.

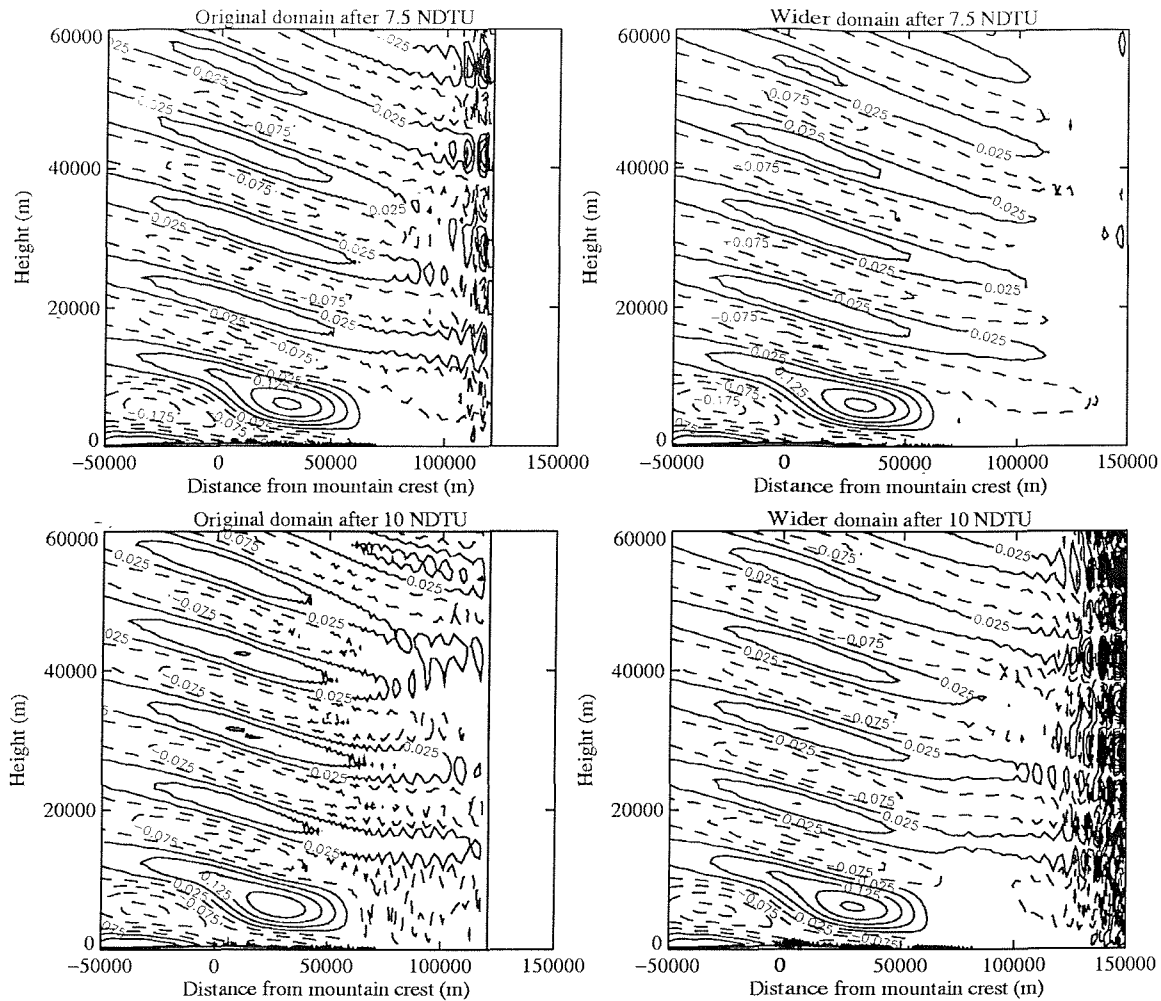


Figure 7.6. Vertical velocity plots after 7.5 NDTU (top plots) and 10 NDTU (bottom plots) for case 4. The contour interval is 0.05 m s^{-1} with a base contour of 0.025 m s^{-1} ; the zero contour has been removed from these plots, but the flow pattern is still clearly visible. The left hand plots used the two stretch horizontal grid and the right hand plots used the wider domain described in the text. The domain plotted here extends to the outflow boundary of the wider domain, 150 km downstream of the mountain crest. The noise at the outflow boundary occurs earlier in the smaller domain simulation.

7.3 Alternative schemes for the radiative outflow boundary condition

7.3.1 Introduction

Variations on the outflow boundary condition formulation explained in Chapter 4 were tested to see whether they produced more robust results, and these modifications are described below. All the simulations in this section had a halved timestep, and they all used the two stretch horizontal grid except where stated otherwise in the text. Case 4 was used as a preliminary test, since its original configuration had the most severe problems at the outflow boundary; some of the modifications made these problems even worse, and those changes were not tested on any other cases, since the boundary condition needs to be robust for a range of different cases. All the changes were rejected, since the original configuration gave the most reliable results: all of the modifications degraded the solution for at least one of the cases tested. In the light of previous comments about the significance of case 4 test results in terms of non-dimensional time, it may be worth re-running some of these tests, and allowing the simulations to continue for longer. However, the non-dimensional time issue was not fully appreciated at the time these tests were carried out, and time constraints did not allow them to be re-investigated for inclusion in this thesis.

7.3.2 Modifications to the restriction on the magnitude of the outflow phase speed

The formulation of the outflow boundary condition was explained in section 4.4, and the constant phase speed used in the outflow boundary condition was defined there. This constant phase speed is discretised in the code as follows, where IIP is the number of horizontal gridpoints, the sums are over all vertical gridpoints, and the other notation is as previously defined:

$$\frac{-}{c} = \frac{\int c \frac{\partial \phi}{\partial x} dz}{\int \frac{\partial \phi}{\partial x} dz} = \frac{\sum c \{ \phi(IIP) - \phi(IIP-1) \} \Delta Z}{\sum \{ \phi(IIP) - \phi(IIP-1) \} \Delta Z}.$$

In the two-dimensional case, ϕ is the horizontal velocity, U . This calculation in the control code uses absolute values of the difference between adjacent horizontal velocity values, presumably because negative phase speeds are clipped to zero. The first modification removed that restriction, but the flow pattern in case 4 was very noisy, as shown in figure 7.7. Overall this change degraded the solution and so was not tested further.

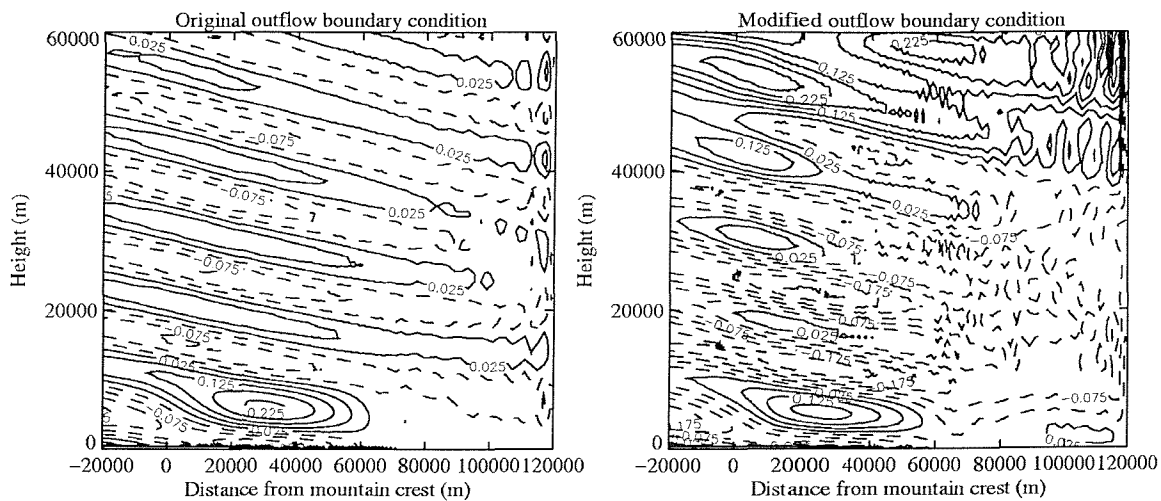


Figure 7.7. Vertical velocity plots (contour interval 0.05 m s^{-1}) for case 4 after 7.5 NDTU, using the original and modified outflow boundary conditions, as described in the text. The modified simulation is far more noisy than the original.

In the control calculation, the outflow phase speed is clipped to lie between zero and $\Delta X/\Delta t$ (for numerical stability purposes; more details are given in section 4.4) before the vertically weighted average is obtained, although these bounds are coded as zero and unity because the phase speed held in the code is non-dimensionalised. The second modification did not clip positive values until after that average had been obtained, making the condition less strict. However, negative values were still clipped to zero at this stage, and the use of absolute values in calculating the vertically weighted average was retained. In case 4, this simulation failed after 700 seconds with no data dumps, because the pressure solver did not converge. No further tests were carried out.

The third modification did no clipping at all until after the vertically weighted average was obtained, then clipped that single value in the same way as before, relaxing the restriction still further. No absolute values were used in calculating the average. The results of the tests were very mixed, and it was hard to decide whether this change was positive or negative overall. Case 1 showed no significant differences, cases 2 and 3 seemed more noisy with the modification, and in case 4 the change removed the outflow noise but introduced other minor problems. The spurious noise introduced in case 3 was most noticeable around the second horizontal grid stretch, especially at upper levels, where it caused the whole wave pattern to disintegrate towards the end of the simulation, as shown in figure 7.8. The problem may be caused by reflection at the outflow boundary, which would explain why it is seen in the interior domain, rather than at the outflow boundary itself. Whereas the signal in the original simulation was flat, the drag coefficient in the modified simulation became noisy at the same time as the spurious noise appeared in the vertical velocity signal. Since most of the difficulties seemed to be related to the second grid stretch, tests were carried out for all five cases with this modified outflow condition and the one stretch horizontal grid, but all of them failed before reaching 3000 seconds, because the pressure solver did not converge. This indicates that, overall, this modification was more negative than positive.

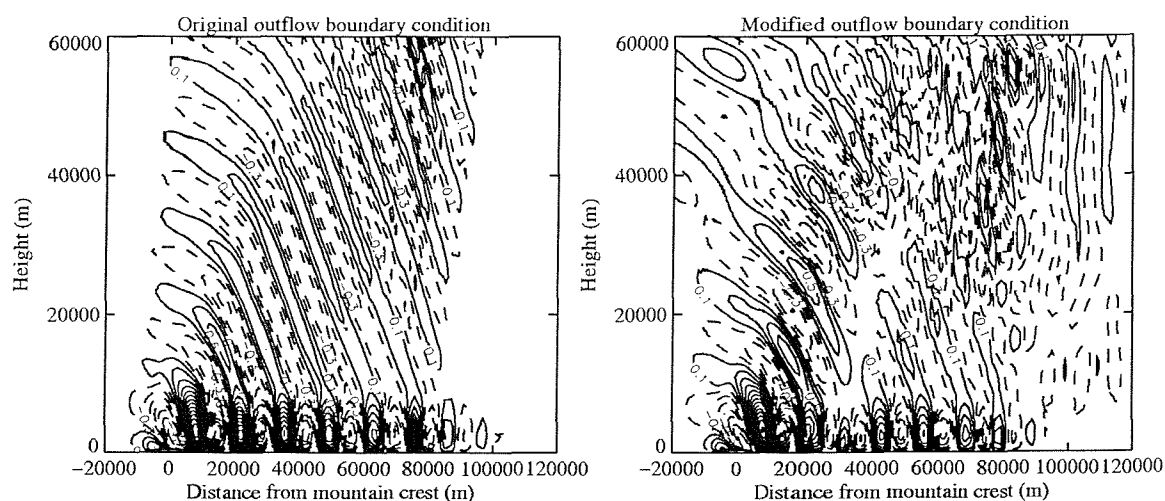


Figure 7.8. Vertical velocity plots for case 3 after 120 NDTU, using the original and modified outflow boundary conditions, as described in the text. The contour interval is 0.2 m s^{-1} , with a base contour of 0.1 m s^{-1} ; the zero contour has been removed from this plot, but the flow pattern is still clearly visible. The modified simulation is far more noisy than the original.

7.3.3 Vertically varying outflow phase speed, with different mass flux corrections

The restriction of using a constant phase speed (\bar{c} , as defined above) for the whole vertical domain is not very logical in physical terms, since one would not normally expect the phase speed to remain constant along the whole boundary. This modification allowed the outflow phase speed used in the model to vary with height, rather than calculating one constant value for the entire vertical domain. However, the vertical variation in the outflow phase speed was not great in general, although there were a few peaks. The phase speed was clipped as in the original code, but at each model level. Overall this change caused a deterioration in the simulations, although there were some positive impacts as well. The basic flow pattern was the same for both simulations in case 3, but at later times the noise around the second grid stretch looked slightly worse, and by 140 NDTU the flow was starting to disintegrate both here and at the intersection of the Rayleigh damping layer and the outflow boundary, although less dramatically than in figure 7.8. The negative effects seemed to be linked to the second grid stretch, but a simulation using this outflow condition and the one stretch horizontal grid was even worse. The situation was similar in case 4, but the vertical velocity plots in the modified simulation were very noisy throughout, and never settled into a steady flow pattern. The drag coefficient signal oscillated right to the end of the simulation, after 17.5 NDTU, whereas the original was flat from around 5 NDTU. Overall this change made the simulation much worse. A simulation using the one stretch horizontal grid with this outflow condition was slightly better than the modified original configuration, but the results still indicated that the change had deteriorated the solution.

A mass flux correction is applied in order to correct for differences in the mass flux into and out of the model domain. This is done by calculating the difference between the mass flux into the domain at the upstream boundary and the mass flux out of the domain at the downstream boundary, and then correcting for this by altering the outflow wind speed; this changes the effective outflow phase speed. Usually the effect on the outflow phase speed is small, but in cases where it is large, this may indicate a problem with the calculation of the outflow phase speed (that is, with the formulation of the outflow boundary condition). It is difficult to judge how best to apply the mass flux correction, since it is a numerical issue rather than a physical one. In the control code, the mass

flux correction is done by adding a constant increment to the outflow wind speed. This method was chosen because the outflow phase speed was constant, so it seemed a logical way to apply the correction. This method was used in all the tests described above. However, the use of a vertically varying phase speed opened up other possibilities for the method chosen for this calculation, and the results from testing some of these are presented below.

The first change to the mass flux correction scheme was that instead of adding a constant amount to the outflow wind speed at each height, the amount added was weighted on the vertical profile of the outflow phase speed, so that the increment was larger at levels with a larger outflow phase speed. The results of these tests were compared with those obtained using a vertically varying phase speed and the original mass flux correction scheme, in order to assess the effect of changing the mass flux calculation. Generally the flow was far more noisy in these simulations, so this modification was rejected.

The second change to the mass flux correction scheme was to apply a constant scaling factor to the outflow wind speed, rather than adding a constant amount. Since the horizontal velocity profile at the outflow boundary may have a wave-like structure with both positive and negative values, scaling those values makes more physical sense than simply increasing (or decreasing) them all by the same amount, as it preserves any wave-like structure present in the original profile. As above, the results of these tests were compared with those obtained using a vertically varying phase speed but the original mass flux correction scheme. Overall this was a negative change, introducing or amplifying spurious noise, and producing very large outflow phase speed values after the mass flux correction was applied. Using the two stretch horizontal grid, case 3 did improve slightly, but the results with the one stretch horizontal grid were less encouraging. The pressure solver failed to converge at the start of the two stretch horizontal grid simulation for case 4, and although the simulation using the one stretch horizontal grid did finish, the results were very discouraging.

7.3.4 Timestep adjusted to avoid clipping the outflow phase speed

The control scheme clips the outflow phase speed in order to maintain numerical stability. Rather than alter the phase speed, which has a physical value and may be expected to affect the flow solution, this modification calculated the timestep which was required in order to avoid clipping the phase speed at all, and used that to adjust the model timestep where necessary. Section 8.1 explains when and how the timestep is monitored and adjusted in BLASIUS. Apart from this change, this outflow condition used the control scheme, so the phase speed did not vary in the vertical and the original mass flux correction calculation was used.

The pressure solver failed to converge for both case 3 and case 4, and the vertical velocity plots were very noisy. The timestep was not changed many times by the stability condition on the phase speed, nor was it much smaller than in other simulations, but the drag coefficient was large and very noisy, and the mass flux correction had a marked effect on the outflow phase speed.

7.4 Discussion and options for how to proceed

The results presented in this chapter clearly indicate that the radiative outflow boundary condition does not always allow the flow to leave the model domain undisturbed. This indicates that there is a problem with the implementation of the Orlanski (1976) radiative boundary condition in BLASIUS. Some of the most marked difficulties are in dealing with initial noise, and these may be minimised by reducing the noise level, as discussed in section 8.4. After all, it is perhaps a little unfair to expect the boundary condition to deal with initial noise, which is not really a wave. However, problems have also been noted in dealing with a simple horizontally propagating wave, and attempts to find and fix such problems with the formulation of the boundary condition, as discussed in the previous section, have not proved successful.

Trying to find and fix inflow/outflow errors could be very time consuming, and may not have a high chance of success given the results presented in the previous section. One option could be to look for the optimal periodic configuration, which would avoid the problem of upstream contamination, but this does not solve the steady state isolated mountain problem, since there is a wave in the periodic horizontal velocity inflow profile. Although using a wider domain has proved to be of some use in reducing the impact of these problems, overall, these results indicate a need to fix, or at least find a way to minimise, problems with the outflow boundary condition formulation.

The outflow boundary condition does seem to be sensitive to the horizontal grid spacing at the downstream edge of the domain, and/or to whether the horizontal grid is uniform or stretched, as suggested by results presented in this chapter and in Chapter 6. Although in some cases, the second grid stretch seemed to be causing spurious damping of waves, and forcing or amplification of noise, it is clear that the two stretch horizontal grid produced a much flatter, less noisy drag coefficient signal for case 3 than the one stretch horizontal grid, as shown in figure 6.13; the same was true of the comparison with the uniform grid. This indicates that the simulation on the two stretch horizontal grid approached a much better representation of a steady state, despite the spurious noise seen at the second grid stretch, than the other configurations. This may be because the second grid stretch seemed to act as a viscous layer, damping the wave as it moved into the less well resolved part of the domain, so that it did not reach the outflow boundary. Coarsening the resolution in this way may be a useful device to mitigate problems at the outflow boundary, where there may be a numerical reflection mechanism causing the wave to come back to the mountain and alter the drag coefficient. It should be noted that the surface pressure force (and hence the drag coefficient) is zero in the flat parts of the domain, so if there is a reflection at the outflow boundary, the drag coefficient signal will not show up any perturbation caused by that reflection until the effect acts over the mountain itself. The drag coefficient perturbations discussed above could be due to an outflow boundary problem, which may be causing a pressure wave to reflect into the domain when the wave head reaches the outflow boundary. Stationary waves such as those discussed here have zero phase speed, and the wave crests themselves do not move in the steady case, so the most

important factor is the propagation speed of the wave head during spin-up, as discussed earlier.

Another option is to have a viscous layer at the downstream boundary in which the equations are not solved, but the solution is relaxed back to the inflow conditions. This would be similar to the Rayleigh damping layer used at the top of the model domain, which was described in Chapter 4. This would permit the use of periodic boundary conditions, while still allowing the desired profiles to be fed in at the inflow boundary. The radiative outflow boundary condition would no longer be used, but this configuration would allow the solution of the steady state isolated mountain case. Orlanski (1976) gave a brief discussion of this option, noting that the method wastes a significant number of gridpoints close to the boundary, and that there could still be reflection issues at the start of the viscous layer. Investigations would be necessary in order to tune the viscous layer, similar to those described in Chapter 6 for the Rayleigh damping layer. Vosper (1995) used damping regions at the upstream and downstream boundaries, and found that 10 to 20 horizontal points were required in each damping region in order to prevent reflection and instability at these boundaries when the horizontal resolution was high. He found this to be more effective than simply applying the radiative boundary condition with relaxation at the edges of the domain. Clearly it makes the simulations more computationally expensive, due to the extra points added onto the numerical domain.

Miranda and James (1992) used a variation of Orlanski's (1976) scheme, based on the work by Raymond and Kuo (1984), but with additional modifications to the numerical implementation, which substantially reduced spurious noise and reflection in their results. It could be worth investigating similar modifications to the implementation in BLASIUS. Orlanski (1976) suggested that the radiative boundary condition could be useful for cases where a weak mean flow is present and an initial inflow boundary condition could be changed to an outflow condition due to the presence of upstream propagating waves in the domain of integration. This merits further investigation for simulations such as those discussed here, since reflection of the upstream propagating waves (and initial noise present in the upstream domain) could well be having a detrimental effect on some of the solutions. Orlanski's scheme calculates the local

phase velocity at each time step, and then chooses an inflow or outflow condition depending on the sign of that phase velocity, but in BLASIUS either an inflow condition or an outflow conditions is imposed continuously at a particular boundary throughout the simulation, and not permitted to change character, which could explain some of the problems seen here.

An alternative would be to continue to use the inflow/outflow boundary conditions, but with a much wider domain. Testing would be required in order to have confidence in the results and to be sure that the boundary was far enough away not to have an adverse impact. It may be possible to find a formula for the minimum value of some non-dimensional parameter which is necessary in order to remove or avoid outflow problems. Suggestions for these parameters include:

- $\frac{U}{NX_d}$, where X_d is the downstream fetch;
- the non-dimensional domain length $\frac{X_d}{\lambda_h}$, where λ_h is the horizontal wavelength;
- $\frac{X_d}{c_{prop}t_T}$, where c_{prop} is the propagation speed of the wave head, and t_T is the total simulation time.

Subject to tests of the effect of upstream fetch, the mountain could be moved from the centre of the domain, as suggested in Chapter 6, in order to provide a greater downstream fetch with minimal increase in computational expense. This is straightforward on a uniform horizontal grid, but for a long downstream domain with a stretched horizontal grid, coding would be required in order to construct the stretched grid as normal, and then add extra points on the downstream end.

The drag coefficient signal is a useful indicator of how steady a simulation is. Figures 7.3 and 7.5, showing the signals for cases 1 and 3 respectively, are representative of the results obtained for these test cases. The large perturbation in the drag coefficient indicates a departure from the steady state. A steady solution does not necessarily exist for every case: Nance and Durran (1998), for example, have shown that non-linear dynamics can generate irregular, unsteady variations in a lee wave pattern even when the upstream flow is steady, which develop quite rapidly, and that

this variability increases with mountain height. However, cases 1 to 3 are designed to reproduce steady state analytical results, so numerical simulations of these cases should be expected to reach a steady state solution.

8 Other sensitivity testing

8.1 Halved timestep

The stability of an explicit numerical model is linked to the Courant number (also called the CFL number after Courant, Friedrichs and Lewy, 1928), α , which for the one-dimensional case is $U\Delta t/\Delta X$ where U is a velocity magnitude, Δt is the timestep and ΔX is the gridlength. This can be extended to the three-dimensional case in several different ways. Three examples are:

$$\alpha = \max \left| \frac{U\Delta t}{\Delta X}, \frac{V\Delta t}{\Delta Y}, \frac{W\Delta t}{\Delta Z} \right|$$

$$\alpha = \Delta t \sqrt{\frac{U^2}{\Delta X^2} + \frac{V^2}{\Delta Y^2} + \frac{W^2}{\Delta Z^2}}$$

$$\alpha = \left(\frac{U}{\Delta X} + \frac{V}{\Delta Y} + \frac{W}{\Delta Z} \right) \Delta t$$

where (U, V, W) are the magnitudes of the components of the three-dimensional velocity, and $\Delta X, \Delta Y, \Delta Z$ are the gridlengths in the three directions respectively. Neglecting viscosity (since the simulations described in this thesis are all laminar, without a boundary layer, as discussed earlier), the third one of these is the form used in BLASIUS. More details of how the stability criterion is implemented in BLASIUS are given by Wood (1989).

Depending on the numerical scheme used, the CFL number has an upper limit in order for the model to be stable. In BLASIUS, upper and lower bounds $CFLUPP$ and $CFLLOW$ are set, and the CFL number is calculated at each point in the model domain.

The maximum value is used to determine the timestep, which is changed if the maximum value of the CFL number, α , is outside the range $CFLLOW \leq \alpha \leq CFLUPP$. In that case, the new timestep is calculated as follows:

$$\Delta t_{new} = \frac{\alpha_{new}}{\alpha_{old}} \Delta t_{old} = \frac{\alpha_{new}}{\left(\frac{U}{\Delta X} + \frac{V}{\Delta Y} + \frac{W}{\Delta Z} \right)}$$

The default values are $CFLLOW = 0.2$ and $CFLUPP = 0.3$, and if the CFL number (α_{old}) lies outside this range then it is reset to 0.25 (α_{new}) by using the above equation.

The test described in this section halved the timestep compared to the way it was calculated in the original configuration, by halving the values of $CFLLOW$ and $CFLUPP$ to 0.1 and 0.15 respectively, and using a CFL number of 0.125 when adjusting the timestep. This might be expected to improve the numerical stability of the model. These simulations in theory go to steady state, so the timestep should not matter except for convergence to that steady state, and generally halving the timestep did speed up that convergence, with drag coefficient plots levelling out sooner and oscillating less. This modification was tested for various horizontal grids, as described in Chapter 6, and generally it reduced spurious noise in the vertical velocity signal, especially at the second grid stretch (where present) and near the outflow boundary. The smaller timestep meant that the outflow phase speed was subject to a less severe restriction, since it is not allowed to have a value greater than $\Delta X/\Delta t$ (more details can be found in Chapters 4 and 7). This may explain the benefits seen at the outflow boundary. This change had no impact on the interior flow pattern, and caused no adverse effects. In the descriptions which follow, ‘full timestep’ refers to the original configuration and ‘halved timestep’ to the modified one described above.

In case 3, using the horizontal two stretch grid, the introduction of the halved timestep meant that although the mountain grew over the same number of timesteps, it grew faster in the full timestep simulation (as discussed in section 8.4 below). Figure 8.1 shows the vertical velocity fields at the end of the mountain growth, and there is more initial noise in the full timestep simulation, where the mountain grew faster. After 40 NDTU, the initial noise had cleared, and the patterns were very similar, although the spurious noise around the second grid stretch was more noticeable in the full timestep simulation, and became even more so at later times, as shown in figure 8.1.

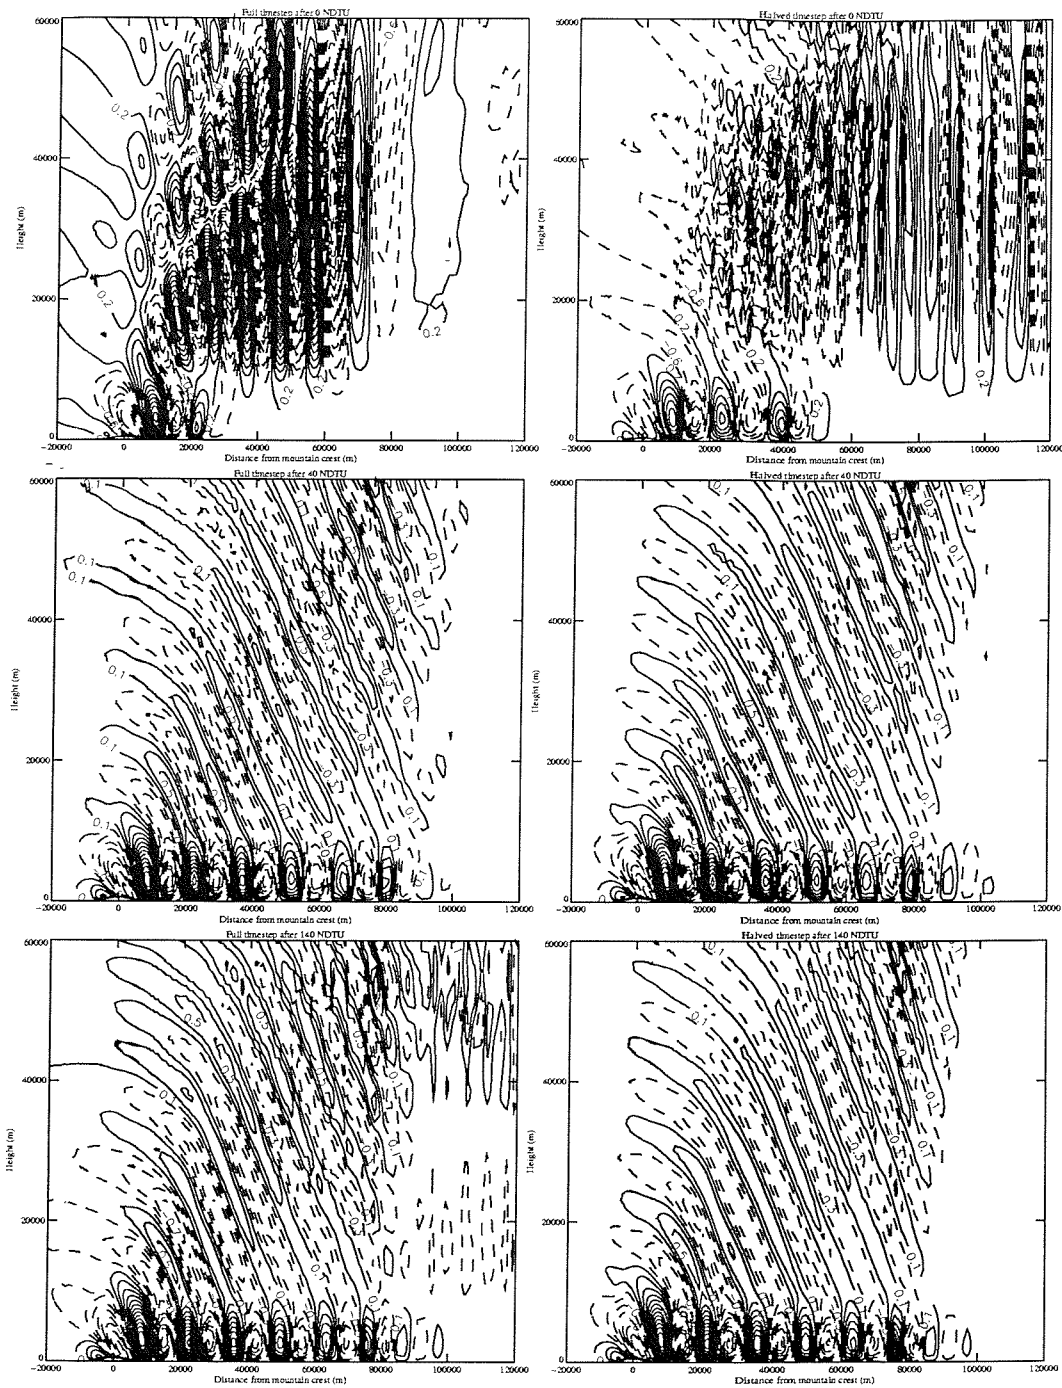


Figure 8.1. Vertical velocity plots for case 3 after 0, 40 and 140 NDTU (top, middle and bottom rows respectively), using the full and halved timestep (left hand and right hand columns respectively). The contour interval for the top two plots is 0.4 m s^{-1} with a base contour of 0.2 m s^{-1} , and for the remaining four plots, it is 0.2 m s^{-1} with a base contour of 0.1 m s^{-1} ; the zero contour has been removed from these plots, but the flow pattern is still clearly visible. There is more initial noise in the full timestep simulation, and more spurious noise around the second grid stretch (70 km downstream of the mountain crest).

Using the one stretch horizontal grid, the full timestep simulation for case 3 was notably more noisy than the one using the two stretch grid. Halving the timestep greatly reduced this noise, but did not remove it completely. Both these simulations were significantly different from their two stretch grid counterparts, and this sensitivity to horizontal resolution was discussed further in Chapter 6. However, overall the halved timestep simulation was much less noisy than the full timestep simulation. The same was true of case 4, where both the initial noise and the outflow bull's-eye noise (as shown in the right hand plot of figure 6.12) were much reduced.

This modification was retained in the final configuration, due to its positive impact in reducing spurious noise.

8.2 Quasi-free slip surface condition

Because the BLASIUS model was originally intended to simulate turbulent boundary layer flows, rather than laminar flows such as those described in this thesis, the surface boundary condition used is a no slip condition, as described in Chapter 4. Roache (1972) found that the surface boundary condition was the most important factor to take into account when attempting to simulate flows such as those in this work, which require a free slip condition. Modifications can be made to use a free slip surface boundary condition in BLASIUS, so that the velocity at the ground is everywhere tangential to the orography. This may be expected to remove the downstream surface gridlength wave (described below) seen in some simulations, and could have an impact on the flow elsewhere. The no slip boundary condition on the velocity sets all three components to zero at the surface. In the free slip case, the gradient of each horizontal velocity component perpendicular to the surface is set to zero, and the boundary condition on the vertical velocity component ensures that the velocity perpendicular to the surface is zero. The surface boundary condition tested here is referred to as the 'quasi-free slip' boundary condition, because the pressure solver still has a no slip boundary condition at the surface. Whilst this produces an inconsistent boundary condition, which may have a detrimental effect on the results, changing the pressure

solver would require a lot of coding and is not worthwhile since this modification is only used for laminar simulations without a boundary layer, which form a very small part of the overall work done with BLASIUS. A constant viscosity was used for the quasi-free slip simulations, rather than the first order closure described in Chapter 4, which was used for no slip simulations. This constant viscosity should be set to zero for a truly inviscid simulation, but as discussed in Chapter 4, a very low viscosity is necessary in the model in order to satisfy the no slip condition.

In some configurations, the vertical velocity plots for cases 1 and 4 showed a downstream surface gridlength wave in the vertical velocity plots, which was visible as a very shallow layer of noise near the surface downstream of the mountain, with peaks which matched exactly onto the horizontal gridpoints used in the model. It did not seem to have any detrimental effect on the solutions, with no obvious impact on the main flow solution. In general, the quasi-free slip surface condition removed the downstream surface gridlength wave, where present. In case 4, for the full timestep case, it also improved the spurious noise at the outflow boundary, as shown in figure 8.2. (The previous comments regarding the non-dimensional simulation time for case 4 should be borne in mind here.) However, using the halved timestep, the quasi-free slip simulation for case 4 crashed when the pressure solver failed to converge at the start of the simulation. In both the full and the halved timestep simulations, the quasi-free slip modification did sometimes have adverse effects, introducing noise both at the outflow boundary and at the base of the Rayleigh damping layer. This could be to do with the lack of a free slip modification in the pressure boundary condition at the surface, since any problems due to this would be likely to feed into the pressure solver, which is where these simulations failed. On balance, the quasi-free slip modification was rejected.

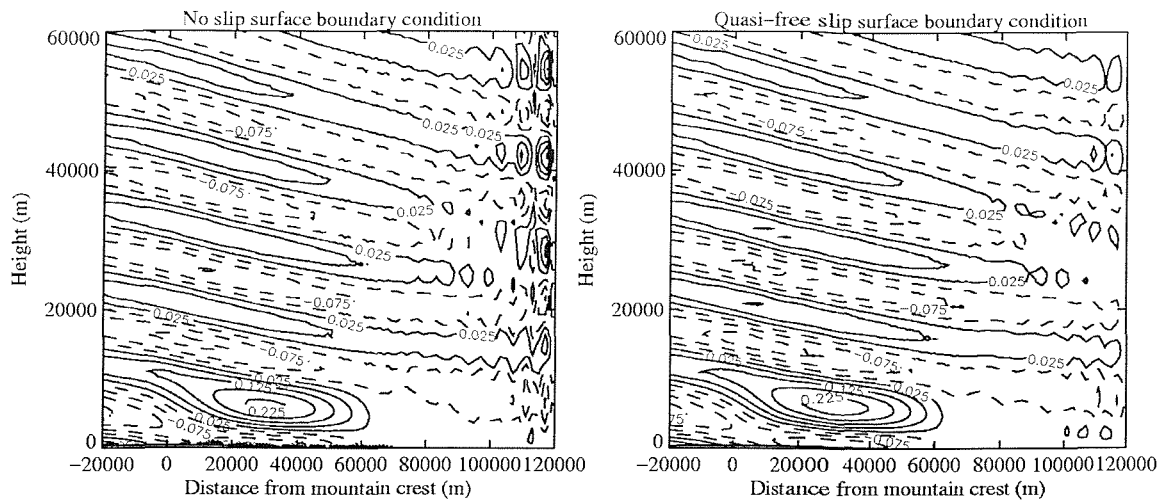


Figure 8.2. Vertical velocity plots for case 4 after 7.5 NDTU, using the no slip and quasi-free slip surface boundary conditions. The contour interval is 0.05 m s^{-1} with a base contour of 0.025 m s^{-1} ; the zero contour has been removed from these plots, but the flow pattern is still clearly visible. The ‘bull’s-eye’ signal at the outflow boundary in the no slip simulation is much reduced in the quasi-free slip simulation, and the downstream surface gridlength wave (visible as a very shallow layer of noise in the left hand plot) is removed.

8.3 Tests of different mountain shapes

The BLASIUS model offers a number of different mountain shapes, as well as the option to import a data set of terrain height values. Most analytical studies, including those presented by Durran (1986) and Keller (1994) which were discussed in Chapter 5, use a Witch of Agnesi mountain, because this has a straightforward Fourier transform. This mountain shape can be used in BLASIUS, or alternatively a cosine-squared mountain can be used to approximate the Witch of Agnesi one. Figure 8.3 shows an example of how the two mountain shapes compare for a mountain height of 500 m. The width of the cosine-squared mountain is chosen so that the cross-sectional area is the same as that of the Witch of Agnesi mountain.




Figure 8.3. Comparison of Witch of Agnesi and cosine-squared mountains of the same height (500 m in this example) and cross-sectional area. In this case the Witch of Agnesi mountain has half width 2500 m, and the cosine-squared mountain has width 15700 m.

The equation for the Witch of Agnesi mountain is

$$h(x) = \frac{h_m a^2}{a^2 + x^2}$$

where $h(x)$ is the mountain height as a function of distance from the crest, h_m is the height of the crest, and a is the half width of the mountain, defined as the distance from the crest at which the mountain height is $h_m/2$. This gives a cross-sectional area of $h_m a \pi$. The equation for the cosine-squared mountain is

$$h(x) = h_m \cos^2\left(\frac{\pi x}{L}\right)$$

for $-L/2 < x < L/2$, and zero elsewhere. Here $h(x)$ and h_m are as previously defined, and L is the width of the mountain. This gives a cross-sectional area of $h_m L/2$. For a Witch of Agnesi mountain of half width 10 km, a cosine-squared mountain of the same area has a width of 62.8 km. It is clear from figure 8.3 that the cosine-squared mountain has a much steeper slope than the Witch of Agnesi at the base, and that the Witch of Agnesi has the steeper slope of the two at the peak, although the difference in slope is less at the peak than at the base.

Tests were carried out using two Witch of Agnesi mountains of half width 10 km, with heights 100 m and 500 m, and their cosine-squared equivalents (i.e. those of the same

cross-sectional area), for the case with constant background wind speed and static stability. These simulations used a uniform grid in both the horizontal and the vertical directions, and periodic boundary conditions. Figure 8.4 shows the vertical velocity fields for the 100 m high mountain after 140 NDTU. In changing from a Witch of Agnesi mountain to a cosine-squared mountain, the maximum slope was reduced, and the wave signal became wider with smaller vertical velocity magnitude. The vertical velocity contours had more noticeable noise and the drag coefficient was larger (and less steady in the 500 m case). A further investigation, using a narrower mountain of height 500 m, whose Witch of Agnesi half width was 5 km, giving a cosine-squared width of 31.4 km, gave similar results to those described above for the wider mountain.

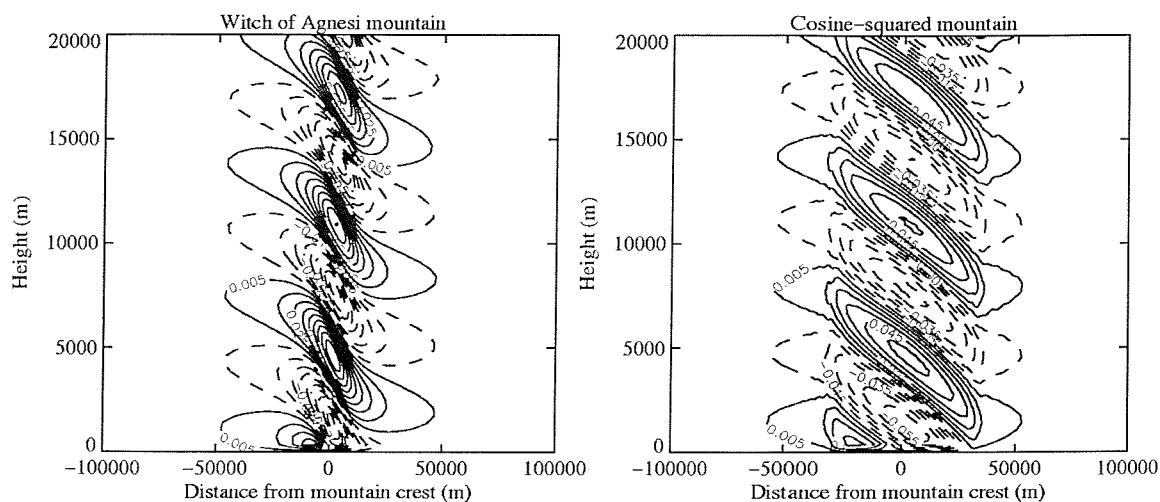


Figure 8.4. Vertical velocity after 140 NDTU for two 100 m high mountains of different shapes but equal cross-sectional areas. The contour interval for the vertical velocity is 0.01 m s^{-1} with a base contour of 0.005 m s^{-1} ; the zero contour has been removed from these plots, but is very smooth both upstream and downstream of the wave. The left hand plot is a Witch of Agnesi mountain of half width 10 km; the right hand one is a cosine-squared mountain of width 62.8 km.

Tests were also carried out for the troposphere-only and troposphere-stratosphere cases used in Keller's (1994) work. Figure 8.5 shows the results from the troposphere-only numerical simulation using a cosine-squared mountain, compared with Keller's figure 3, but note that the contour interval used in this plot is half that in Keller's. A quantitative comparison shows that the maximum vertical velocity values in the wave pattern produced by the simulation were at least 75% weaker than those obtained in the analytical results. This may be partly explained by the presence of numerical viscosity

in the simulation, which damps the wave motion (as discussed in Chapter 5), and partly by the result found above, that the cosine-squared mountain gives a smaller magnitude wave signal than the corresponding Witch of Agnesi mountain.

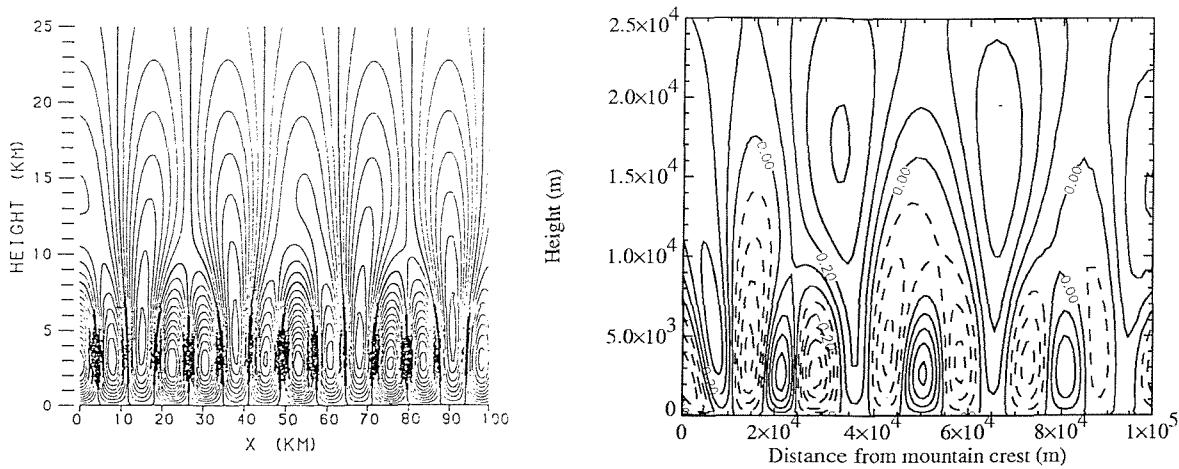


Figure 8.5. Vertical velocity for case 2, the troposphere-only case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 3, which is shown on the left. The contour interval in the numerical simulation, which uses a cosine-squared mountain, is 0.1 m s^{-1} , half that used in Keller's plot. Note that in this case, all of Keller's contours are solid; the leftmost cell is negative in this plot.

For the troposphere-stratosphere case, figure 8.6 shows the results from the BLASIUS simulation using a cosine-squared mountain, compared with Keller's figure 5, with the same contour interval used in both plots. Again, there was a difference in the intensity of the wave, but this time the vertical velocity signal near the mountain was much stronger in the numerical simulation than in the analytical result, increasing by over 80% at later times. This does not fit with the above observations about the cosine-squared and Witch of Agnesi mountains, or with the results found for the troposphere-only case. The peak in vertical velocity magnitude is in the trapped wave rather than in the upwardly propagating wave, and the above results indicate that the magnitude of the trapped wave should be smaller in the numerical simulation than in the analysis. This is the case further from the mountain, where the flow was slightly (15-30%) weaker, and the most likely explanation for the increase close to the mountain is that there is a non-linear interaction between the two propagating wave trains, as discussed in Chapter 5.

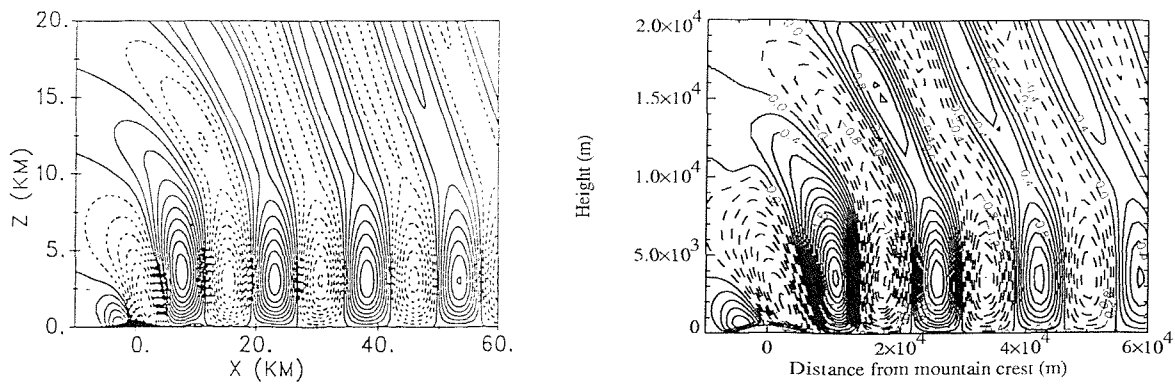


Figure 8.6. Vertical velocity for case 3, the troposphere-stratosphere case, shown for part of the model domain to allow easier comparison with Keller's (1994) figure 5, which is shown on the left. The contour interval in the numerical simulation, which uses a cosine-squared mountain, is 0.2 m s^{-1} , the same as that used in Keller's plot.

It was decided to use a Witch of Agnesi mountain rather than a cosine-squared mountain, since the results are less noisy and more robust. Much of the noise seen in the cosine-squared tests can be explained by the steep slope at the base of the mountain, which provides a strong initial forcing. Care is required in order to resolve adequately the sharp peak of the Witch of Agnesi mountain, as discussed in Chapter 6, but its use allows easier comparison with analytical results.

8.4 Initial noise and mountain growth

A two- or three-dimensional BLASIUS simulation starts with a flat surface and grows the mountain into the domain over a given number of timesteps; this number is input as a parameter. As discussed in section 8.1, the timestep is determined from the velocity components and the grid spacing, so the time taken to grow the mountain will not necessarily be the same in two simulations with different physical parameters, even if the number of timesteps chosen is the same. For example, a simulation where the background wind is constant with height, such as case 1, will have a larger timestep than one where the background wind increases with height, such as case 2. In reality,

the timestep can be quite different from what might be intuitively expected from the calculation described in section 8.1, and this is discussed below.

The process of growing the mountain introduces a forcing into the simulation, and in most cases the early vertical velocity data shows ‘initial noise’ due to this forcing. Figure 8.7 shows examples of this for two very different cases, where the mountain was grown over 1000 timesteps in each case. The left hand plot is for the constant U , constant N case, with a 100 m high Witch of Agnesi mountain of half width 10 km. There is hardly any initial noise in this case, which is what one would intuitively expect, since the mountain has a very shallow maximum slope of 0.0065, and grows relatively slowly, over 12500 seconds. The right hand plot is for the troposphere-stratosphere case, with a 500 m high cosine-squared mountain of width 15700 m. The initial noise in this case is much stronger, as one would expect, with the larger mountain (five times as high, and with a maximum slope of 0.1) providing a far stronger forcing, and also growing much faster, over 5000 seconds.




Figure 8.7. Vertical velocity at the end of mountain growth (0 NDTU). The left hand plot is for the constant U , constant N case, with a 100 m high Witch of Agnesi mountain of half width 10 km; here the contour interval is 0.01 m s^{-1} and the mountain grew over 12500 seconds; the Rayleigh damping layer started at 20 km. The right hand plot is for the troposphere-stratosphere case, with a 500 m high cosine-squared mountain of width 15700 m; here the contour interval is 0.4 m s^{-1} and the mountain grew over 5000 seconds; the Rayleigh damping layer started at 60 km, so the vertical domain plotted is three times that shown in the left hand plot. There is far more initial noise in the latter case.

The right hand case in figure 8.7 is an inflow/outflow simulation, so the initial noise flows out of the downstream domain and does not reappear upstream of the mountain. The left hand case is periodic, but since the initial noise is very slight and there is no horizontally propagating wave, no ill effects are seen. Using periodic boundary conditions in the right hand case, however, caused the strong initial noise seen in figure 8.7 to re-enter the domain upstream and destroy the upper level wave as it passed through, and this effect was discussed in the previous chapter. Durran and Klemp (1982) pointed out that a radiative outflow boundary condition downstream should allow the initial transients to pass out of the domain, so the model should not be highly sensitive to the initialisation procedure. They found that this was true for their model, but in BLASIUS the initial noise can cause problems at the outflow boundary, as discussed in Chapter 7. In any case, it is desirable to minimise this noise as far as possible, and one would expect to achieve this by growing the mountain more slowly. Investigations for the constant U , constant N case have revealed some simple formulae, linking the maximum mountain slope, the horizontal velocity and the horizontal grid spacing, which calculate with sufficient accuracy how many timesteps should be used to grow the mountain in order to keep the initial noise below a chosen level. Further investigations will need to include other cases, in order to produce a more general result.

The first of these empirical formulae relates the maximum initial noise (in terms of the magnitude of the vertical velocity signal), the maximum mountain slope, and the time taken to grow the mountain. The maximum initial noise over the whole model domain, I_{\max} (m s^{-1}), is non-dimensionalised by U , and this quantity $\frac{I_{\max}}{U}$ is plotted on the horizontal axis of figure 8.8. The time (in seconds) taken to grow the mountain, t_g , is non-dimensionalised by $\frac{U}{h_m}$, since the mountain height is a more relevant length scale here than the half width, as used to non-dimensionalise the simulation times elsewhere in this thesis. The vertical axis of figure 8.8 shows $\frac{t_g U}{S_{\max} h_m}$, where S_{\max} is the maximum mountain slope. Results are plotted for 24 simulations, which differ in mountain shape, height and width, in vertical domain height, in Rayleigh damping layer height and coefficient, in vertical and horizontal grid, in upstream and downstream boundary

conditions, and in the number of timesteps used to grow the mountain. The average point and the line of best fit exclude the simulations which produced no initial noise: these used gently sloping mountains (of maximum slope 0.005 or less), which provide a very weak forcing, so the mountain could be grown more quickly in these simulations and no initial noise would be expected until t_g became very small. A useful test would be to investigate the limiting case where the mountain is grown over one timestep: initial noise would be expected for any mountain in this case, and it would be interesting to see how quickly the flow would adjust to the instantaneous appearance of the mountain. The average point and the line of best fit also exclude the simulation with a very steep Witch of Agnesi mountain (of maximum slope 0.13), since this simulation was unrepresentatively noisy due to its inadequate horizontal resolution (which was discussed in Chapter 6).




Figure 8.8. The relationship between the time taken to grow the mountain, the maximum mountain slope and the maximum initial noise level throughout the model domain, as described in the text. Results are shown for 24 different simulations. Crosses indicate simulations using a Witch of Agnesi mountain; asterisks indicate those using a cosine-squared mountain. The average (excluding four simulations, as described in the text) is denoted by a diamond, and the line of best fit is marked for the same subset of simulations. The right hand plot is a close-up of the left hand one, to show the region of interest more clearly; it excludes the simulations which are not used in calculating the average or the line of best fit.

The line of best fit plotted in figure 8.8 has the equation

$$\frac{t_g U}{S_{max} h_m} = 8950 - 84025 \frac{I_{max}}{U} .$$

The total time taken to grow the mountain is then given by the following formula:

$$t_g = \frac{25h_m S_{\max}}{U^2} (358U - 3361I_{\max}).$$

This can be used to calculate how many seconds should be used to grow the mountain in order to achieve a chosen level of initial noise for a particular mountain slope.

Although a goal of no initial noise is desirable, in many cases this is not computationally practical, since it can require up to 40000 seconds to grow the mountain. In practice a non-zero initial noise level (up to around 0.1 m s^{-1}) can still provide satisfactory results.

The results shown in figure 8.8 used different numbers of timesteps to grow the mountains, ranging from the default value of 1000 up to 10836 timesteps in one case. Some of these simulations were carried out in order to test the formula obtained from earlier results, and they fitted in well with the pattern. Generally there was no difference between a pair of otherwise identical simulations where the mountain grew at different speeds, except that the spin up was slower in the one where the mountain was grown more slowly. This slower spin up was evident in the vertical velocity signal, which was weaker at earlier times, and in the drag coefficient signal, which took longer to climb initially. The initial noise was reduced as expected, and after the mountain had finished growing, the results were the same in both simulations.

Once this formula has been used to determine the number of seconds desired to grow the mountain, another relationship is required to calculate how many timesteps this will take. The timestep is not easily determined in advance of starting the simulation, because it can be different from what would be obtained by using the equation given in section 8.1. That equation, for the two-dimensional case, gives

$$\Delta t = \frac{\alpha}{\left(\frac{U}{\Delta X} + \frac{W}{\Delta Z} \right)}$$

where α is 0.125 here. For the purposes of this calculation, the above equation was

reduced to one dimension by neglecting the $\frac{W}{\Delta Z}$ term, since the vertical velocity is

identically zero in the basic state. Although this may not be strictly justified on a scale analysis of the two terms in the wave motion itself, consideration of the basic state is sufficient for the purpose here, since the calculation does not need to be highly accurate.

This gives an expected timestep of $\frac{\Delta X}{8U}$, but in practice the timestep is often much smaller than this. The expected time taken to grow the mountain, t_e , is defined as

$$t_e = \frac{n_g \Delta X}{8U}$$

where n_g is the number of timesteps taken to grow the mountain. The actual time taken to grow the mountain, t_g , can be expressed as a percentage of the expected time taken to grow the mountain as follows:

$$\text{percentage} = 100 \frac{t_g}{t_e} = \frac{800t_g U}{n_g \Delta X}.$$

This percentage gives a measure of how much the actual timestep differed from what was expected.

Figure 8.9 shows a graph of this percentage against the maximum mountain slope for the same range of simulations as figure 8.8. For simulations using a stretched horizontal grid, the smallest grid spacing was used, since that leads to the smallest (i.e. most restricted) timestep. The average excludes the same values as before, as does the linear line of best fit, which has the equation

$$\text{percentage} = \frac{800t_g U}{n_g \Delta X} = 79S_{\max} - 219.$$

However, in this case it is appropriate to include the simulations with the gentle mountain slopes, since one would intuitively expect the percentage to be close to 100 for these cases, which fits with the results obtained. A curve has been fitted by eye to all but the remaining excluded point, with the equation

$$\text{percentage} = \frac{800t_g U}{n_g \Delta X} = \frac{1}{0.8S_{\max} + 0.02} + 51.25.$$

This demonstrates how a steeper mountain leads to a non-linear decrease in the timestep, compared to what was expected. A plausible explanation of this is that there is a non-linear feedback process whereby the steeper mountain causes larger perturbations, leading to larger velocity values and hence a smaller timestep, and so the mountain grows faster.




Figure 8.9. The relationship between the time taken to grow the mountain, expressed as a percentage of the time it would be expected to take using the equation given in section 8.1, and the maximum mountain slope. The vertical axis starts at 45% rather than at zero, to make the results easier to see. Results are shown for the same 24 simulations as in figure 8.8. Crosses indicate simulations using a Witch of Agnesi mountain; asterisks indicate those using a cosine-squared mountain. The average (excluding four simulations, as described in the text) is denoted by a diamond, and the line of best fit (solid line) is marked for the same subset of simulations. The dashed line is a curve fitted by eye to all simulations except for one extreme value considered unrepresentative. More details are given in the text.

The above equation can be rearranged to give the following formula for the number of timesteps used to grow the mountain:

$$n_g = \frac{640t_g U (40S_{\max} + 1)}{\Delta X (1640S_{\max} + 81)}$$

In order to calculate how many timesteps are needed to grow the mountain for a chosen maximum initial noise level I_{\max} , this formula is combined with the one obtained earlier for initial noise, to give:

$$n_g = 16000 \left(\frac{h_m}{\Delta X} \right) S_{\max} \left(\frac{40S_{\max} + 1}{1640S_{\max} + 81} \right) \left(358 - 3361 \frac{I_{\max}}{U} \right).$$

This should be tested for other cases with varying upstream wind and temperature profiles, but it does hold to an acceptable degree of accuracy for a wide range of simulations with constant values of $U = 10 \text{ m s}^{-1}$ and $N = 0.01 \text{ s}^{-1}$.

9 Discussion and suggestions for future work

9.1 Summary and recommendations

Several numerical issues have been investigated in this thesis: grid independence, inflow/outflow boundary conditions, the model timestep, the surface boundary condition, the mountain shape, and the initial transience caused by growing the mountain into the model domain at the start of a simulation. No numerical model can provide a perfect representation of the real atmosphere, because there are always approximations involved, so it is inevitable that numerical issues such as those encountered here will arise to some extent. However, the results presented in Chapter 5, obtained after completion of the above investigations, show that the effects of these issues have been minimised adequately to allow satisfactory simulation of the cases studied here. In addition, suggestions have been made for further investigations where appropriate. This section summarises recommendations for future work using BLASIUS to simulate this type of flow.

It should be emphasised here that BLASIUS has been used successfully for boundary layer flows over hills for several years, and that there is confidence in the model; references for this type of work were given in earlier chapters of this thesis. The investigations presented here have tested the model for cases which go beyond its original intended purpose: as discussed in Chapter 2, BLASIUS was not designed with this type of gravity wave flow in mind. Because of this, some teething problems are to be expected during this preliminary testing phase. The main issue revealed by the present study is a sensitivity of the model results to the formulation of the upstream and

downstream boundary conditions, and the grid independence issues are not obviously unrelated to this. However, it should be borne in mind that other models have been used to simulate two-dimensional orographic gravity wave flows such as those considered here, without such problems (e.g. Durran and Klemp, 1982, among many others), which suggests that the issues raised here are simply linked to recent model changes and are not insurmountable.

The dominant issue by far is the outflow boundary condition, and several options for further investigation were suggested in section 7.4. More testing is required, but in the meantime recommendations for using the present code are given here. Clearly care is required in order to resolve the mountain adequately, especially any smaller scale features such as the relatively sharp peak on a Witch of Agnesi mountain. However, small horizontal gridlengths can cause spurious noise, especially in the downstream domain and most notably near the outflow boundary. The reasons for this are not entirely clear, and further investigation is required, but the problem does seem to be linked to the outflow boundary condition. Pending further testing, the recommendation for the model as it currently stands, with the outflow boundary condition coded in its present form, is to use the largest horizontal gridlength which will resolve both the mountain and the wavelength of the resulting flow pattern, and to use stretching so that the gridlength can become larger near the boundaries, to minimise problems there. The stretching needs to be done slowly and smoothly (maximum expansion ratios of 1.05 seem to produce satisfactory results) in order to minimise spurious internal reflections. The wave signal will obviously become less well resolved as the horizontal resolution decreases, so care is required to maintain adequate resolution in the region of physical interest, and the rest of the domain should be regarded as a numerical damping region, similar to the Rayleigh damping layer at the upper boundary. The domain needs to be wide enough for the mountain and the expected wave signal: whilst a horizontally propagating wave should leave the numerical domain undisturbed (whether through the downstream boundary or by numerical damping due to decreasing resolution), a vertically propagating wave should really be contained within the confines of the horizontal numerical domain.

Similar comments apply for the vertical resolution: the minimum should be four gridlengths (five gridpoints) per wavelength, but preferably more. A uniform vertical grid is satisfactory for laminar simulations such as those discussed here, but stretching will become necessary when a boundary layer is introduced, in order to obtain adequate resolution at low levels. This is discussed further in the next section, but one thing to bear in mind is the expansion ratio: the resolution can easily become too coarse at upper levels, and this can have a detrimental effect on the Rayleigh damping scheme, which is explicitly dependent on the vertical gridlength. The Rayleigh damping scheme itself could be more effective after further testing and tuning for gravity wave cases such as these, and suggestions for how to do this were given at the end of Chapter 6.

When comparing different simulations, non-dimensional time should be used in order to remove differences due to the speed of mountain growth, and in order to compare physically different simulations in a consistent non-dimensional manner. The formula derived in section 8.4 gives an idea of how quickly to grow the mountain. This needs testing for cases other than the constant U , constant N case for which it was derived (case 1), but provides a useful method to minimise the initial noise, thereby minimising reflection at the outflow boundary and general spurious noise during the simulation. The drag coefficient is a good indicator of whether a simulation has reached steady state; large perturbations can often be explained by looking for signs of reflection at the outflow boundary and calculating how long the initial noise or propagating wave head would take to reach the mountain in that case.

9.2 Suggestions for further work

Work is currently underway at The Met. Office (Smith, private communication) to reproduce cases 2 and 3 of the results presented here, which were based on Keller's (1994) analytical work, using a fully compressible model, and the results from that study are in broad agreement with those presented here, as discussed in Chapter 3. The next step in order to build on this groundwork is to introduce a boundary layer into these gravity wave simulations. Recent simulations at The Met. Office of flow over

mountains using the BLASIUS model have included a realistic representation of the boundary layer. In addition, preliminary simulations similar to case 3 described in Chapter 5, but with different generic types of boundary layer, have shown a sensitivity of the lee wave structure to the presence of the boundary layer (Guilbaud, private communication). Peng and Thompson (1998) demonstrated that the boundary layer on the lee slope acts to effectively change the terrain profile, and concluded that this caused the changes which they observed in the downstream flow in their results compared to the case with no boundary layer.

Future work using two-dimensional model simulations for three different boundary layer types (neutral, stable and convective) would be useful, in order to investigate interactions between the boundary layer and the free atmosphere. Following on from this, investigations into the effect of variations in the relative heights of the mountain and the boundary layer could be carried out, to assess differences in the flow when the mountain is not wholly contained within the boundary layer. Including the boundary layer requires better vertical resolution at low levels, and before commencing these simulations, the simulations described in Chapter 5 should be repeated with this 'boundary layer grid', to ensure that the increased resolution itself does not change the solution. The boundary layer simulations will be more computationally expensive than the ones discussed in this work, due to the increased resolution and the representation of turbulence.

One of the first tasks when commencing turbulent boundary layer simulations of phenomena such as rotors (i.e. separation and reattachment of the downstream flow) should be to examine which turbulence closure is most appropriate for this work, since it can affect the details of the separation which occurs in the model for a given case. Wood (1995) discusses this aspect, as well as giving a more general overview of separation for neutral turbulent flow over hills. Although stable stratification is required for the gravity wave flows studied here, the same broad principles are likely to apply, and will certainly provide a starting framework. Hunt (1980) provides a discussion of the various regimes, including lee side separation and rotor formation.

The issue of how to diagnose flow separation in the model simulations will need to be addressed. Coulter (1998) simply followed the streamline close to the surface, and if it detached, diagnosed separation. In two dimensions she found that this was usually dramatic, with a large detachment of the streamline from the surface, and a distinct closed separation bubble. In three dimensions the diagnosis of separation is more difficult than in the two-dimensional case, since the streamfunction cannot be defined, and Coulter (1998) describes two techniques. The first uses path trajectories as streamlines, following particles from close to the surface from different horizontal locations over the domain, looking for a significant increase in distance from the ground. An alternative method uses flow visualisation, tracing particles using the tangential surface stress field to represent the velocity, and looking for singularities where $\frac{\partial U}{\partial z}$ is zero at the surface. These appear as strong divergence or convergence of the particle paths, but are not a necessary condition for flow separation in three dimensions, so Coulter used both methods.

Investigations aimed at gaining an improved understanding of rotors should include the effect of variations in mountain shape, height, width and slope, to find out which lead to rotor formation, including using an asymmetric mountain to achieve the steep lee slope which is believed to be preferential for rotor formation. Preliminary work using an asymmetric mountain, where the downstream part was half the width of the upstream part, was carried out in this project, but much more investigation is required. In addition, the effect of variations in upstream basic state profile should be studied in each of these situations, to explore the parameter space further and find out which situations are most likely to lead to rotor formation. The next step will be to combine these and investigate the interaction between the two factors, as well as investigating the effects of boundary layer stability. Surface roughness can also change results, with a smoother surface acting to retard separation (Hunt, 1980), so this is another area worthy of investigation.

Idealised initial profiles have been used for the basic state in this thesis, but for rotor simulations, suitable basic state profiles from observed and documented cases of rotors should be given particular attention. Examples of these can be found in Dent and Dyson (1963), Bedard and Neilley (1998) and Olivier and Poulos (1998). The use of

realistic terrain profiles would also be beneficial, since these are usually very different from an isolated idealised ridge or mountain, and interaction between wave trains excited by neighbouring orography frequently becomes important.

After the two-dimensional investigations, three-dimensionality should be explored. Miranda and James (1992) used a three-dimensional numerical model to examine wave breaking and flow splitting with a bell-shaped mountain at low Froude number, and found that the results were not the same as those obtained in two dimensions. Hunt (1980) also found differences in lee side separation in two and three dimensions: the two-dimensional separation had closed streamline loops and reattachment at the downwind end of the separation, whereas the three-dimensional flow was much more complicated. He noted from both experimental and theoretical evidence that an isolated mountain tends to produce a longer separated region than a mountain followed by another mountain. Once rotor simulation has been achieved in two-dimensional simulations, and parameter space explored to find out when rotors form, validation of these results using three-dimensional simulations will be crucial. These simulations will be very computationally expensive and it may prove necessary to reduce the model domain and/or the resolution, performing tests similar to those carried out in this thesis, in order to find a suitable configuration. Initial three-dimensional simulations should include an infinite two-dimensional ridge, to compare the results with those obtained for the same case in two dimensions, and validate the model. Queney *et al.* (1960) reported that both theory and observations have shown that individual mountains are much less effective in exciting trapped lee waves than are long ridges. The waves tend to have smaller amplitude and less downstream extent, which is partly explained by the possibility for air flow around the mountain as well as over, and means that forecasters and pilots can afford to pay less attention to these isolated mountains than to long ridges of similar height. This finding may go some way towards explaining why discussions of rotor observations mainly talk of ridges, and the distinctive roll cloud which marks the rotor along the length of the ridge. However, even with seemingly two-dimensional orography such as this, three-dimensional flow effects can and do appear. Inclusion of moisture in the model might also be useful, as it would provide some indication of whether clouds are likely to be formed, and therefore whether the visual warning signs upon which pilots depend are likely to be present. Castro and Snyder's (1993) results

also indicate that emphasis on mountain slope can be misleading, and that spanwise width may be equally, if not more, important in determining the flow regime, so this is worth investigating. However, it should be noted that those results were obtained for a towing tank, and may not be directly applicable to the atmospheric simulations in this work.

References

Acheson, D.J. (1990) *Elementary Fluid Dynamics*, Oxford Applied Mathematics and Computing Science Series, Clarendon Press, Oxford, England

Allan, J. (1997) Mountain flying, *Pilot*, September 1997: 28-31

Baines, P.G. (1995) *Topographic Effects in Stratified Flows*, Cambridge University Press, Cambridge, England

Baines, P.G. and Hoinka, K.P. (1985) Stratified Flow over Two-Dimensional Topography in Fluid of Infinite Depth: A Laboratory Simulation, *Journal of the Atmospheric Sciences* 42 (15): 1614-1630

Batchelor, G.K. (1967) *An Introduction to Fluid Mechanics*, Cambridge University Press, Cambridge, England

Bedard, A.J., Jr. and Neilley, P. (1998) Project MCAT (Mountain Induced Clear Air Turbulence): Background, Goals, Instrumentation, and Methodologies, *Preprints, Eighth Conference on Mountain Meteorology* (American Meteorological Society, 3-7 August 1998, Flagstaff, Arizona): 163-171

Belcher, S.E., Newley, T.M.J. and Hunt, J.C.R. (1993) The drag on an undulating surface induced by the flow of a turbulent boundary layer, *Journal of Fluid Mechanics* 249: 557-596

- Belcher, S.E. and Wood, N. (1996) Form and wave drag due to stably stratified turbulent flow over low ridges, *Quarterly Journal of the Royal Meteorological Society* 122 (532): 863-902
- Boer, G.J., McFarlane, N.A., Laprise, R., Henderson, J.D. and Blanchet, J.-P. (1984) The Canadian Climate Centre spectral atmospheric general circulation model, *Atmos.-Ocean*, 22: 397-429
- Booker, J.R. and Bretherton, F.P. (1967) The critical layer for internal gravity waves in a shear flow, *Journal of Fluid Mechanics* 27 (3): 513-539
- Boussinesq, J. (1903) *Théorie analytique de la chaleur*, Vol. 2, Gauthier-Villars, Paris, France: 154-176
- Boyer, D.L., Chen, R. and Davies, P.A. (1987) Some Laboratory Models of Flow Past the Alpine/Pyrenees Mountain Complex, *Meteorol. Atmos. Phys.* 36: 187-200
- Bradbury, T. (1997) Mountain waves, *Pilot*, September 1997: 47-50
- Bretherton, F.P. (1969) Momentum transport by gravity waves, *Quarterly Journal of the Royal Meteorological Society* 95: 213-243
- Broad, A.S. (1995) Linear theory of momentum fluxes in 3-D flows with turning of the mean wind with height, *Quarterly Journal of the Royal Meteorological Society*, 121: 1891-1902
- Bull, J.M. (1990) *Recent Developments in Numerical Methods for the Advection Equation and their Application to Boundary Layer Modelling*, Met O (P) Turbulence and Diffusion Note No.195, The Met. Office, Bracknell, England (unpublished)
- Businger, J.A. (1982) Equations and concepts, *Atmospheric Turbulence and Air Pollution Modelling* (ed. F.T.M. Nieuwstadt and H. Van Dop), D. Reidel Publishing Company, Dordrecht, Holland: 1-36

Carney, T.Q., Bedard, A.J., Jr., Brown, J.M., McGinley, J., Lindholm, T. and Kraus, M.J. (1996) *Hazardous Mountain Winds and Their Visual Indicators*, Handbook, Dept. of Commerce, NOAA, Boulder, Colorado [Recently republished as US Federal Aviation Administration Advisory Circular 00-57]

Castro, I.P. and Jones, J.M. (1987) Studies in numerical computations of recirculating flows, *Int. J. for Num. Methods in Fluids* 7: 793-823

Castro, I.P. and Snyder, W.H. (1993) Experiments on wave breaking in stratified flow over obstacles, *Journal of Fluid Mechanics* 255: 195-211

Castro, I.P., Snyder, W.H. and Baines, P.G. (1990) Obstacle drag in stratified flow, *Proceedings of the Royal Society of London A* 429 (1876): 119-140

Castro, I.P., Snyder, W.H. and Marsh, G.L. (1983) Stratified flow over three-dimensional ridges, *Journal of Fluid Mechanics* 135: 261-282

Clark, T.L. and Miller, M.J. (1991) Pressure drag and momentum fluxes due to the Alps. II: Representation in large-scale atmospheric models, *Quarterly Journal of the Royal Meteorological Society* 117: 527-552

Coleman, G.N., Ferziger, J.H. and Spalart, P.R. (1992) Direct simulation of the stably stratified turbulent Ekman layer, *Journal of Fluid Mechanics* 244: 677-712

Corby, G.A. (1954) The airflow over mountains: A review of the state of current knowledge, *Quarterly Journal of the Royal Meteorological Society* 80: 491-521

Coulter, S.C. (1998) *Flow Over Forested Hills*, Ph.D. thesis, Environment Centre, University of Leeds (unpublished)

Courant, R., Friedrichs, K. and Lewy, H. (1928) Über die partiellen Differenzgleichungen der mathematischen Physik, *Math. Ann.* 100: 32-74

Crapper, G.D. (1959) A three-dimensional solution for waves in the lee of mountains, *Journal of Fluid Mechanics* 6: 51-76

Davis, R.E. (1969) The two-dimensional flow of a stratified fluid over an obstacle, *Journal of Fluid Mechanics* 36 (1): 127-143

Dent, L. and Dyson, B. (1963) Rotor Streaming Over the Pennines, *Meteorological Magazine* 92: 358-363

Dörnbräck, A., Leutbecher, M., Kivi, R. and Kyrö, E. (1999) *Mountain wave induced record low stratospheric temperatures above Northern Scandinavia*, Institut für Physik der Atmosphäre Report No. 115

Dörnbrack, A. and Schumann, U. (1994) Numerical simulation of breaking gravity waves below a critical level, in *Contributions to the first ERCOFTAC workshop on direct and large-eddy simulation, March 28th-30th 1994, Guildford, England*, eds. P. Voke, J.-P. Chollet and L. Kleiser, Institut für Physik der Atmosphäre Report No. 17

Drazin, P.G. (1961) On the steady flow of a fluid of variable density past an obstacle, *Tellus* 13: 239-251

Durrán, D.R. (1986) Mountain Waves, *Mesoscale Meteorology and Forecasting* (Ed. P.S. Ray), American Meteorological Society, Boston, Mass.: 472-492 (Chapter 20)

Durrán, D.R. (1991) Orographic Wave Drag on the Lower Troposphere: The Importance of Trapped Waves, *Preprints, Eighth Conference on Atmospheric and Oceanic Waves and Stability* (American Meteorological Society, 14-18 October 1991, Denver, Colorado): 377-380

Durrán, D.R. (1995) Do Breaking Mountain Waves Decelerate the Local Mean Flow? *Journal of the Atmospheric Sciences* 52 (22): 4010-4032

- Durrán, D.R. and Klemp, J.B. (1982) The effects of moisture on trapped mountain lee waves, *Journal of the Atmospheric Sciences* 39 (11): 2490-2506
- Dutton, J.A. and Fichtl, G.H. (1969) Approximate equations of motion for gases and liquids, *Journal of the Atmospheric Sciences* 26: 241-254
- Eliassen, A. and Palm, E. (1960) On the transfer of energy in stationary mountain waves, *Geofys. Publ.* 22 (3): 1-23
- Farnell, L. (1977) *Numerical integration of flow in a rotating annulus. III. Solution of Poisson's equation*, Met.O.21 Occ. Note 77/1, The Met. Office, Bracknell, England (unpublished)
- Farnell, L. (1980) Solution of Poisson equations on a nonuniform grid, *Journal of Computational Physics* 35: 408-425
- Findlater, J. (1984) Lee Wave and Rotor Turbulence at Inverness (Dalcross) Airport, *Flight Safety Focus* (Flight Safety Comm., London) No.3: 13-15
- Förchtgott, J. (1949) Wave currents on the leeward side of mountain crests. Translation of *B. Met.* (Inst. Mét., Prague) 3: 49-51
- Förchtgott, J. (1952) Mechanical turbulence. Translation of *Letecká Met.* (Aeron. Met., Prague) 114-161
- Garratt, J.R. (1992) *The atmospheric boundary layer*, Cambridge University Press, Cambridge, England.
- Gerbier, N. and Berenger, M. (1961) Experimental studies of lee waves in the French Alps, *Quarterly Journal of the Royal Meteorological Society* 87: 13-23
- Gossard, E.E. and Hooke, W.H. (1975) *Waves in the atmosphere*, Developments in Atmospheric Science 2, Elsevier Scientific Publishing Company, Amsterdam

Gregory, D., Shutts, G.J. and Mitchell, J.R. (1998) A new gravity-wave-drag scheme incorporating anisotropic orography and low-level wave breaking: Impact upon the climate of the UK Meteorological Office Unified Model, *Quarterly Journal of the Royal Meteorological Society* 124 (546): 463-493

Guilbaud, C. (1997) *Inflow-Outflow boundary conditions for Blasius: Description and tests*, Met O (APR), The Met. Office, Bracknell, England (unpublished)

Hanazaki, H. (1988) A numerical study of three-dimensional stratified flow past a sphere, *Journal of Fluid Mechanics* 192: 393-419

Holets, S.H. and Swanson, R.N. (1988) Observations of a lee wave and rotor episode in central California, *Boundary Layer Meteorology* 44: 97-103

Hunt, J.C.R. (1980) Wind Over Hills, *Workshop on the Planetary Boundary Layer* (ed. J.C. Wyngaard, American Meteorological Society, 14-18 August 1978, Boulder, Colorado): 107-146

Hunt, J.C.R. and Snyder, W.H. (1980) Experiments on stably and neutrally stratified flow over a model three-dimensional hill, *Journal of Fluid Mechanics* 96 (4): 671-704

Huppert, H.E. and Miles, J.W. (1969) Lee waves in a stratified flow. Part 3: Semi-elliptical obstacle. *Journal of Fluid Mechanics*, 35 (3): 481-496

Janowitz, G.S. (1984) Lee waves in three-dimensional stratified flow, *Journal of Fluid Mechanics* 148: 97-108

Jasperson, W.H., Nastrom, G.D. and Fritts, D.C. (1990) Further study of terrain effects on the mesoscale spectrum of atmospheric models, *Journal of the Atmospheric Sciences* 47 (8): 979-987

Keller, T.L. (1994) Implications of the Hydrostatic Assumption on Atmospheric Gravity Waves, *Journal of the Atmospheric Sciences* 51 (13): 1915-1929

- Knable, C. (1995) Atmospheric Turbulence, *Flight Deck* (British Airways Safety Services, Hounslow) No.16: 17-22
- Küttner, J. (1938) Moazagotl und Föhnwolle, *Beitr. Phys. Frei. Atmos.* 25:79-114
- Küttner, J. (1939) Zur Entstehung der Föhnwolle, *Beitr. Phys. Frei. Atmos.* 25: 251-299
- Küttner, J. (1959) *The rotor flow in the lee of mountains*, GRD research notes No. 6 (AF CRC-TN 58-626), USAF, Cambridge Research Center, Bedford, Mass.
- Lamb, M.W. and Baker, S.P. (1998) Flight Hazards of Microscale Mountain Weather, *Preprints, Eighth Conference on Mountain Meteorology* (American Meteorological Society, 3-7 August 1998, Flagstaff, Arizona): 182-185
- Leonard, B.P., MacVean, M.K. and Lock, A.P. (1993) *Positivity-Preserving Numerical Schemes for Multidimensional Advection*, NASA Technical Memorandum 106055, ICOMP-93-05
- Long, R.R. (1953) Some Aspects of the Flow of Stratified Fluids: I. A Theoretical Investigation, *Tellus* V (1): 42-58
- Long, R.R. (1955) Some Aspects of the Flow of Stratified Fluids: III. Continuous Density Gradients, *Tellus* VII (3): 341-357
- Lyra, G. (1940) Über den Einfluß von Bodenerhebungen auf die Strömung einer stabil geschichteten Atmosphäre, *Beitr. Phys. Freien Atmos.* 26: 197-206
- Lyra, G. (1943) Theorie der stationären Leewellenströmung in freier Atmosphäre, *Z. Angew. Math. Mech.* 23 (1): 1-28
- Mahrt, L. (1986) On the shallow motion approximations, *Journal of the Atmospheric Sciences* 43 (10): 1036-1044

- Manley, G. (1945) The Helm wind of Crossfell, 1937-1939, *Quarterly Journal of the Royal Meteorological Society* 71: 197-219
- Mason, B.J. (1970) Future developments in meteorology: an outlook to the year 2000, *Quarterly Journal of the Royal Meteorological Society* 96 (409): 349-368
- Mason, P.J. (1985) On the parameterization of orographic drag, *Physical parameterizations for numerical models* (ECMWF seminar proceedings, Volume 2, 9-13 September 1985, Reading, England): 139-165
- Mason, P.J. and King, J.C. (1985) Measurements and predictions of flow and turbulence over an isolated hill of moderate slope, *Quarterly Journal of the Royal Meteorological Society* 111 (468): 617-640
- McFarlane, N.A. (1987) The effect of orographically excited gravity wave drag on the general circulation of the lower stratosphere and troposphere, *Journal of the Atmospheric Sciences* 44 (14) 1775-1800
- Mickle, R.E., Cook, N.J., Hoff, A.M., Jensen, N.O., Salmon, J.R., Taylor, P.A., Tetzlaff, G. and Teunissen, H.W. (1988) The Askervein hill project: Vertical profiles of wind and turbulence, *Boundary Layer Meteorology* 43 (1-2): 143-169
- Miles, J.W. (1969) Waves and wave drag in stratified flow, *Contrib. Scripps Inst. Ocean.* 39: 1675-1702
- Milton, S.F. and Wilson, C.A. (1996) The impact of parametrized sub-grid scale orographic forcing on systematic errors in a global NWP model, *Mon. Weather Review* 124 (9): 2023-2045
- Miranda, P. and James, I.N. (1992) Non-linear three-dimensional effects on gravity-wave drag: Splitting flow and breaking waves, *Quarterly Journal of the Royal Meteorological Society* 118 (508): 1057-1081

Mobbs, S.D. (1994) The parametrization of sub-grid scale orography, *Parametrization of sub-grid scale physical processes* (ECMWF seminar proceedings, 5-9 September 1994, Reading, England): 305-324

Nance, L.B. and Durran, D.R. (1998) A Modeling Study of Nonstationary Trapped Mountain Lee Waves. Part II: Nonlinearity, *Journal of the Atmospheric Sciences* 55: 1429-1445

Nastrom, G.D., Fritts, D.C. and Gage, K.S. (1987) An investigation of terrain effects on the mesoscale spectrum of atmospheric models, *Journal of the Atmospheric Sciences* 44 (20): 3087-3096

Neiman, P.J., Hardesty, R.M., Shapiro, M.A. and Cupp, R.E. (1988) Doppler Lidar Observations of a Downslope Windstorm, *Monthly Weather Review* 116 (11): 2265-2275

Nicholls, J.M. (1973) *The airflow over mountains: Research 1958-1972*, WMO Technical Note 127

Olivier, L.D. and Poulos, G.S. (1998) Frontal Passage, Mountain Waves and Flow Reversals in the Vicinity of the Colorado Springs, CO Airport, *Preprints, Eighth Conference on Mountain Meteorology* (American Meteorological Society, 3-7 August 1998, Flagstaff, Arizona): 176-181

Orlanski, I. (1976) A Simple Boundary Condition for Unbounded Hyperbolic Flows, *Journal of Computational Physics* 21: 251-269

Palmer, T.N., Shutts, G.J. and Swinbank, R. (1986) Alleviation of a systematic westerly bias in general circulation and numerical weather prediction models through an orographic gravity wave drag parametrization, *Quarterly Journal of the Royal Meteorological Society* 112: 1001-1039

- Peltier, W.R. and Clark, T.L. (1979) The evolution and stability of finite-amplitude mountain waves. Part II: surface wave drag and severe downslope windstorms, *Journal of the Atmospheric Sciences* 36: 1498-1529
- Peng, M.S. and Thompson, W.T. (1998) Effect of Surface Friction on Mountain Waves, *Preprints, Eighth Conference on Mountain Meteorology* (American Meteorological Society, 3-7 August 1998, Flagstaff, Arizona): 364-367
- Piacsek, S.A. and Williams, G.P. (1970) Conservation Properties of Convection Difference Schemes, *Journal of Computational Physics* 6 (3): 392-405
- Pielke, R.A. (1984) *Mesoscale Meteorological Modelling*, Academic Press, Orlando, Florida
- Poulos, G.S. and Olivier, L.D. (1998) Mountain Wave/Frontal Interaction Dynamics in the Lee of Pike's Peak, *Preprints, Eighth Conference on Mountain Meteorology* (American Meteorological Society, 3-7 August 1998, Flagstaff, Arizona): 172-175
- Queney, P. (1948) The problem of air flow over mountains: a summary of theoretical studies, *Bull. Amer. Meteor. Soc.* 29: 16-26
- Queney, P., Corby, G.A., Gerbier, N., Koschmieder, H. and Zierep, J. (1960) *The airflow over mountains*, WMO Tech. Note No. 34
- Ralph, F.M. and Neiman, P.J. (1997) Lidar observations of a breaking mountain wave associated with extreme turbulence, *Geophysical Research Letters* 24 (6): 663-666
- Ralph, F.M., Neiman, P.J., Levinson, D., Clark, T. and Hall, W.D. (1998) Observations of upper-tropospheric streamwise rolls associated with airflow over mountains, *Preprints, Eighth Conference on Mountain Meteorology* (American Meteorological Society, 3-7 August 1998, Flagstaff, Arizona): 119-125

- Raymond, W.H. and Kuo, H.L. (1984) A radiation boundary condition for multi-dimensional flows, *Quarterly Journal of the Royal Meteorological Society* 110: 535-551
- Roache, P.J. (1972) On Artificial Viscosity, *J. Comp. Phys.* 10(2): 169-184 [Also reprinted as Appendix B of Roache (1976), below]
- Roache, P.J. (1976) *Computational fluid dynamics*, Hermosa Publishers, Albuquerque, New Mexico
- Runchal, A.K. and Wolfshtein, M. (1969) Numerical Integration Procedure for the Steady State Navier-Stokes Equations, *J. Mech. Eng. Sci.* 11 (5): 445-453
- Sawyer, J.S. (1959) The introduction of the effects of topography into methods of numerical forecasting, *Quarterly Journal of the Royal Meteorological Society* 85 (363) 31-43
- Scorer, R.S. (1949) Theory of waves in the lee of mountains, *Quarterly Journal of the Royal Meteorological Society* 75: 41-56
- Scorer, R.S. (1978) *Environmental Aerodynamics*, Ellis Horwood, Chichester, England
- Scorer, R.S. and Klieforth, H. (1959) Theory of mountain waves of large amplitude, *Quarterly Journal of the Royal Meteorological Society* 85 (364): 131-143
- Shutts, G.J. (1995) Gravity wave drag parametrization over complex terrain: the effect of critical level absorption in directional wind shear, *Quarterly Journal of the Royal Meteorological Society* 121: 1005-1021
- Shutts, G.J., and Broad, A.S. (1993) A case study of lee waves over the Lake District in northern England, *Quarterly Journal of the Royal Meteorological Society* 119: 377-408

- Shutts, G.J., Healey, P. and Mobbs, S.D. (1994) A multiple sounding technique for the study of gravity waves, *Quarterly Journal of the Royal Meteorological Society* 120: 59-77
- Smagorinsky, J. (1963) General circulation experiments with the primitive equations. I. The basic experiment, *Mon. Weather Review* 91 (3): 99-164
- Smith, R.B. (1978) A measurement of mountain drag, *Journal of the Atmospheric Sciences* 35 (9): 1644-1654
- Smith, R.B. (1979) *The Influence of Mountains on the Atmosphere*, Advances in Geophysics 21, Academic Press, Orlando
- Smith, R.B. (1989a) Hydrostatic airflow over mountains, *Advances in Geophysics* (ed. B. Saltzman) 31:1-41, Academic Press, San Diego, California
- Smith, R.B. (1989b) Mountain-induced stagnation points in hydrostatic flow, *Tellus* 41A: 270-274
- Smolarkiewicz, P.K. and Rotunno, R. (1990) Low Froude number flow past three-dimensional obstacles. Part II: Upwind flow reversal zone, *Journal of the Atmospheric Sciences* 47 (12): 1498-1511
- Snyder, W.H., Thompson, R.S., Eskridge, R.E., Lawson, R.E., Castro, I.P., Lee, J.T., Hunt, J.C.R. and Ogawa, Y. (1985) The structure of strongly stratified flow over hills: dividing-streamline concept, *Journal of Fluid Mechanics* 152: 249-288
- Spiegel, E.A. and Veronis, G. (1960) On the Boussinesq approximation for a compressible fluid, *Astrophysics Journal* 131 (2): 442-447
- Stull, R.B. (1988) *An Introduction to Boundary Layer Meteorology*, Kluwer Academic Publishers, Dordrecht, Holland

- Sykes, R.I. (1980) An asymptotic theory of incompressible turbulent boundary-layer flow over a small hump, *Journal of Fluid Mechanics* 101: 647-670
- Vichnevetsky, R. (1987a) Wave Propagation Analysis of Difference Schemes for Hyperbolic Equations: A Review, *International Journal for Numerical Methods in Fluids* 7: 409-452
- Vichnevetsky, R. (1987b) Wave Propagation and Reflection in Irregular Grids for Hyperbolic Equations, *Applied Numerical Mathematics* 3: 133-166
- Vosper, S.B. (1995) *Observation and modelling of orographic internal gravity waves*, Ph.D. thesis, Department of Applied Mathematical Studies, University of Leeds (unpublished)
- Vosper, S.B. and Mobbs, S.D. (1996) Lee waves over the English Lake District, *Quarterly Journal of the Royal Meteorological Society* 122: 1283-1305
- Vosper, S.B. and Mobbs, S.D. (1997) Measurement of the pressure field on a mountain, *Quarterly Journal of the Royal Meteorological Society* 123 (537): 129-144
- Vosper, S.B. and Mobbs, S.D. (1998) Momentum fluxes due to three-dimensional gravity-waves: Implications for measurements and numerical modelling, *Quarterly Journal of the Royal Meteorological Society* 124 (552): 2755-2769
- Wallace, J.M., Tibaldi, S. and Simmons, A.J. (1983) Reduction of systematic forecast errors in the ECMWF model through the introduction of an envelope orography, *Quarterly Journal of the Royal Meteorological Society* 109: 683-717
- White, A.A. and Bromley, R.A. (1995) Dynamically consistent, quasi-hydrostatic equations for global models with a complete representation of the Coriolis force, *Quarterly Journal of the Royal Meteorological Society* 121: 399-418

- Wood, N. (1989) *The Description of a Partial Implicit Scheme Applied to a Two-Dimensional Turbulence Model with a Linear Stability Analysis for both Explicit and Implicit Schemes*, Met O (APR) Turbulence and Diffusion Note No.192, The Met. Office, Bracknell, England (unpublished)
- Wood, N. (1992) *Turbulent Flow Over Three-Dimensional Hills*, Ph.D. dissertation, Reading University
- Wood, N. (1995) The onset of separation in neutral, turbulent flow over hills, *Boundary Layer Meteorology* 76: 137-164
- Wood, N. (2000) Wind flow over complex terrain: a historical perspective and the prospect for large-eddy modelling, *Boundary Layer Meteorology* 96: 11-32
- Wood, N., Hewer, F.E. and Hobson, J.M. (1999) *BLASIUS Version 3.03: Documentation of a Model of Flow over Hills*, Met O (APR), The Met. Office, Bracknell, England (unpublished)
- Wood, N. and Mason, P.J. (1993) The Pressure Force Induced by Neutral, Turbulent Flow Over Hills, *Quarterly Journal of the Royal Meteorological Society* 119: 1233-1267
- Wurtele, M.G., Sharman, R.D. and Datta, A. (1996) Atmospheric Lee Waves, *Annu. Rev. Fluid Mech.* 28: 429-476
- Xu, D. and Taylor, P.A. (1995) Boundary-layer parametrization of drag over small-scale topography, *Quarterly Journal of the Royal Meteorological Society* 121 (522): 433-443