

WRF DTC Visitor Programme, January 2009: EXPLORING CUT-CELLS

Sarah-Jane Lock

This short report is intended to summarise the work undertaken during my visit to NCAR in January 2009. It is hoped that this report will prove useful to the continuing work on cut-cells in both the WRF group and the group at Leeds University. It is intended as a working note rather than a formal report. If any relevant or useful information has been omitted from the report, it is a simple oversight and not intentional. I hope this report can form a basis for an on-going discussion on the potential of a cut-cell method.

Sarah-Jane Lock, July 2009

1 Introduction

During a 5-week visit to NCAR in January 2009, the cut-cell method already implemented in 2D “toy model” codes by Joe Klemp, Bill Skamarock and Dave Dempsey were explored and compared with the method used at Leeds University.

This report describes a number of models: “VHREM” refers to the code developed at Leeds University; “scmc.Rev1.f90” is Dave Dempsey’s FORTRAN90 version of Joe’s “shavedcell” code; occasionally, the term “WRF codes” is used to describe all three versions of cut-cell models being explored in the WRF group — the “shavedcell”, “scmc.Rev1.f90” and Dave Dempsey’s extended “scmc.Rev2.f90” codes. All results presented in this report from the “WRF codes” have been generated using the “scmc.Rev1.f90” version.

VHREM is a 3D model based on advection-form equations, which uses a time-splitting integration method (acoustic and gravity modes solved on the short time-step, $\Delta\tau$) with leapfrog time-stepping and 2nd-order centred spatial-differencing to compute model variables u , v , w , π' (the perturbation of the Exner pressure) and θ' (the perturbation of the potential temperature). Integrations are fully explicit (since the model is designed for microscale studies where $\Delta x \sim \Delta z$). Model variables are stored on a grid staggered in both the horizontal and vertical — π' is stored at the centre of a grid-volume, with u , v and w stored on the grid-volume’s eastern, northern and upper faces respectively; θ' is co-located with w . The orographic boundary is represented by piecewise continuous bi-linear surfaces, which connect at the interfaces between grid-columns (i.e. along grid-faces storing u or v). The cut grid-face areas/volumes that result from the intersecting orographic surface are computed *approximately* by summing over numerous very small rectangular surface areas/volumes. An approximate finite-volume approach is used to solve for flow through the irregularly-shaped cut-cells. The method is based on that proposed in Bonaventura (2000) for use with advection-form equations, such that only the divergence term from the continuity equation (found in the π' equation) takes account of the cut-cells, using Gauss’s theorem. All other fluxes are treated just as they would be on a regularly-shaped mesh.

The “scmc.Rev1.f90” code, based on the earlier “shavedcell” code, is a 2D model based on flux-form conservative equations, solving for u , w , ρ and θ . The model is also based on a time-splitting integration method with leapfrog time-stepping and 2nd-order centred spatial-differencing. Integrations are semi-implicit, using a tri-diagonal solver in the vertical for w , ρ and θ . Model variables are stored on a staggered grid, with ρ and θ co-located at grid-cell centres, and u and w on the eastern and upper grid-faces respectively. Orography is represented with a piecewise continuous linear surface, connecting at the u interfaces between grid-cells. The resulting cut cell-lengths and areas are computed exactly. A finite-volume representation using Gauss’s theorem calculates *all* fluxes through ρ -centred cells. (The “scmc.Rev2.f90” code aims to extend the finite-volume approach to include fluxes through u - and w -centred cells as well.)

2 Concerns/problems with the WRF codes

At the time of the visit to NCAR, there were a number of concerns and problems identified with the existing WRF codes:

- Spurious disturbances can be seen in the flow-fields upstream of the hill — in the flow aloft as well as near the boundary;
- Comparisons of the model results with either the linear solution or Long’s solution reveal differences, which appear greater at higher vertical levels;
- The cut-cell formulation in “scmc.Rev1.f90” only uses a finite-volume approach to compute fluxes through cut-cells centred on the density variable. For a complete finite-volume representation, the computations should be extended to compute the fluxes through the cell-volumes around u and w points as well (as described in Adcroft et al., 1997). The “scmc.Rev2.f90” code starts this extension, but is currently failing to produce good results.

3 Model tests

3.1 Tests: hydrostatic flow “Base Case” — a comparison of WRF and VHREM codes

To better understand the problems with the WRF codes, a hydrostatic flow set-up was chosen and replicated as closely as possible using the VHREM code. This set-up will be referred to as the “Base Case”, which later tests were compared against.

The set-ups for the Base Case for the “scmc.Rev1.f90” code and for VHREM are summarised in Tables 1 and 2 respectively. The VHREM set-up aimed to replicate as closely as possible the set-up in the scmc.Rev1.f90 code.

There are some basic differences between the VHREM and scmc.Rev1.f90 model codes:

- VHREM is a 3D code. To run a 2D test, the extent of the 3rd dimension (y -direction) is limited to just a small number of grid-spaces, and the model variables remain invariant in the y -direction (no Coriolis force is applied);
- VHREM only currently has the option for cyclic lateral boundary conditions. For the hydrostatic flow tests, the domain was made wide enough that a steady flow has developed before the perturbed downstream winds complete a full cycle and disturb the upstream flow. For the set-up described above, the

Model domain:		
Δx	2000m	horizontal resolution
Δz	250m	vertical resolution
NZ	81	number of vertical levels
$\Delta \tau$	5s	small time-step
Δt	20s	long time-step (i.e. $\Delta t = 4\Delta \tau$)
Upper level damping layer:		
ZD	8000m	height of bottom of layer
$XNUT$	0.05	damping coefficient
Cut-cell parameter:		
IDV	0	$\Rightarrow DVP = [0, 1]$
Plotting parameter:		
$LONG$	1	\Rightarrow plotted with Long's solution

Table 1: Parameter settings for the Base Case for the “scmc.Rev1.f90” code.

Model domain:		
$\Delta x, \Delta y$	2000m	horizontal resolution
Δz	250m	vertical resolution
L	81	number of vertical levels
$\Delta \tau$	0.2s	small time-step
Δt	20s	long time-step (i.e. $\Delta t = 10\Delta \tau$)
Hill dimensions:		
$ZH0$	400m	hill height
ZAM	10000m	hill half-width
Background state:		
$ITH0$	2	stably stratified
$TH00$	300K	$\bar{\theta}$ at model bottom
BVF	0.0001	buoyancy frequency, N
$U00$	10ms^{-1}	mean wind
Upper level damping layer:		
$ZRDBT$	8000m	height of bottom of layer
$ALPRAY$	0.01	damping coefficient

Table 2: Parameter settings for the Base Case for VHREM.

domain length in the x -direction was $N = 200$ (i.e. 400km). Results from VHREM are compared with WRF code results at time $t = 24,000$ s (runs on larger domains suggest that the solution at $t = 24,000$ s is approx. a steady state solution);

- VHREM uses a much shorter time-step than the WRF models. In VHREM, a stable $\Delta\tau$ is computed from the CFL condition for the smallest *effective* grid-length (defined to be the smallest ratio of cell volume to cell face area, i.e. $\Delta V_p / \Delta S_x$, or “ DVP/DXW ” in terms of the WRF codes) in the model domain, and the long time-step is given by $\Delta t = N\Delta\tau$, where $N=10$. For this set-up, $\Delta\tau=0.2$ s. In the WRF codes, the long time-step is scaled by the hill dimensions and takes no account of the geometry of the cut-cells. For this base case, the WRF codes use $\Delta\tau = 5$ s, $\Delta t = 20$ s. VHREM could not produce stable results with $\Delta\tau \geq 0.4$ s;

Figure 1 shows results for the Base Case from the `scmc.Rev1.f90` code at model time $t = 36,000$ s and from VHREM at model time $t = 24,000$ s. For both models, plots illustrate the computed fields for the potential temperature ($\theta = \bar{\theta} + \theta'$), the perturbation in the horizontal wind field ($U' = U - U_0$) and the vertical wind field.

Results from the `scmc.Rev1.f90` model reveal disturbances in the upstream solution, which are particularly apparent in the vertical wind field. Plots (a) and (b) also include Long’s solution for comparison. Results from `scmc.Rev1.f90` generally compare well with Long’s solution in terms of magnitude and position; but, in addition to the upstream disturbances, the model solutions appear to differ more from Long’s solution at higher model levels.

There is no evidence of similar upstream disturbances in the VHREM results. To date, VHREM results have not been plotted against Long’s solution (or the linear solution). However, it can be seen that the wave in the u and w fields from VHREM is smaller in magnitude than that from `scmc.Rev1.f90`. (The contour lines that appear to overlay the lower boundary in the u plot are a consequence of not mastering how to force NCAR Graphics to only plot contour lines within the model domain, i.e. above the lower boundary!)

3.2 Tests: Thin-wall approximation (DVP=1.0)

The `scmc.Rev1.f90` code includes a simple switch (IDV) to compute the model solutions using the cell volumes for uncut cells — the “thin-wall approximation” described in Steppeler et al. (MWR, 2002). The Base Case outlined in Section 3.1 was repeated but applying the thin-wall approximation — in `scmc.Rev1.f90`, setting $IDV = 1$ ($\Rightarrow DVP = 1.0$ for all grid-cells); and with the equivalent alteration to the VHREM code.

Results for the thin-wall approximation test from the `scmc.Rev1.f90` model at time $t = 36,000$ s and VHREM at model time $t = 24,000$ s are displayed in Figure 2. For both models, the results from the thin-wall approximation show no discernible difference (by eye) from those for the Base Case.

The thin-wall approximation has recently been explored a little further using the VHREM code. The advantage with exploiting the thin-wall approximation is that it should enable a longer stable time-step to be used for $\Delta\tau$. By choosing a 2D hill case, VHREM has been run twice: once, with its default approach of using the cut-cell volumes (labelled “`orog2B`” in the following results) and then, repeated with all volumes set equal to 1.0 (labelled “`orog2B_DV1`”). In each case, the model’s built-in algorithm for computing a stable time-step has been used.

The model set-up is based on that for the case presented in Gallus & Klemp (2000) for a hill with half-width 1km and height 400m.

Results for the potential temperature and vertical wind fields at model time $t = 2,880$ s are shown in Figure 3. For both fields, plots are included for the case using the cut-cell volumes (“`orog2B`”) and for the thin-wall

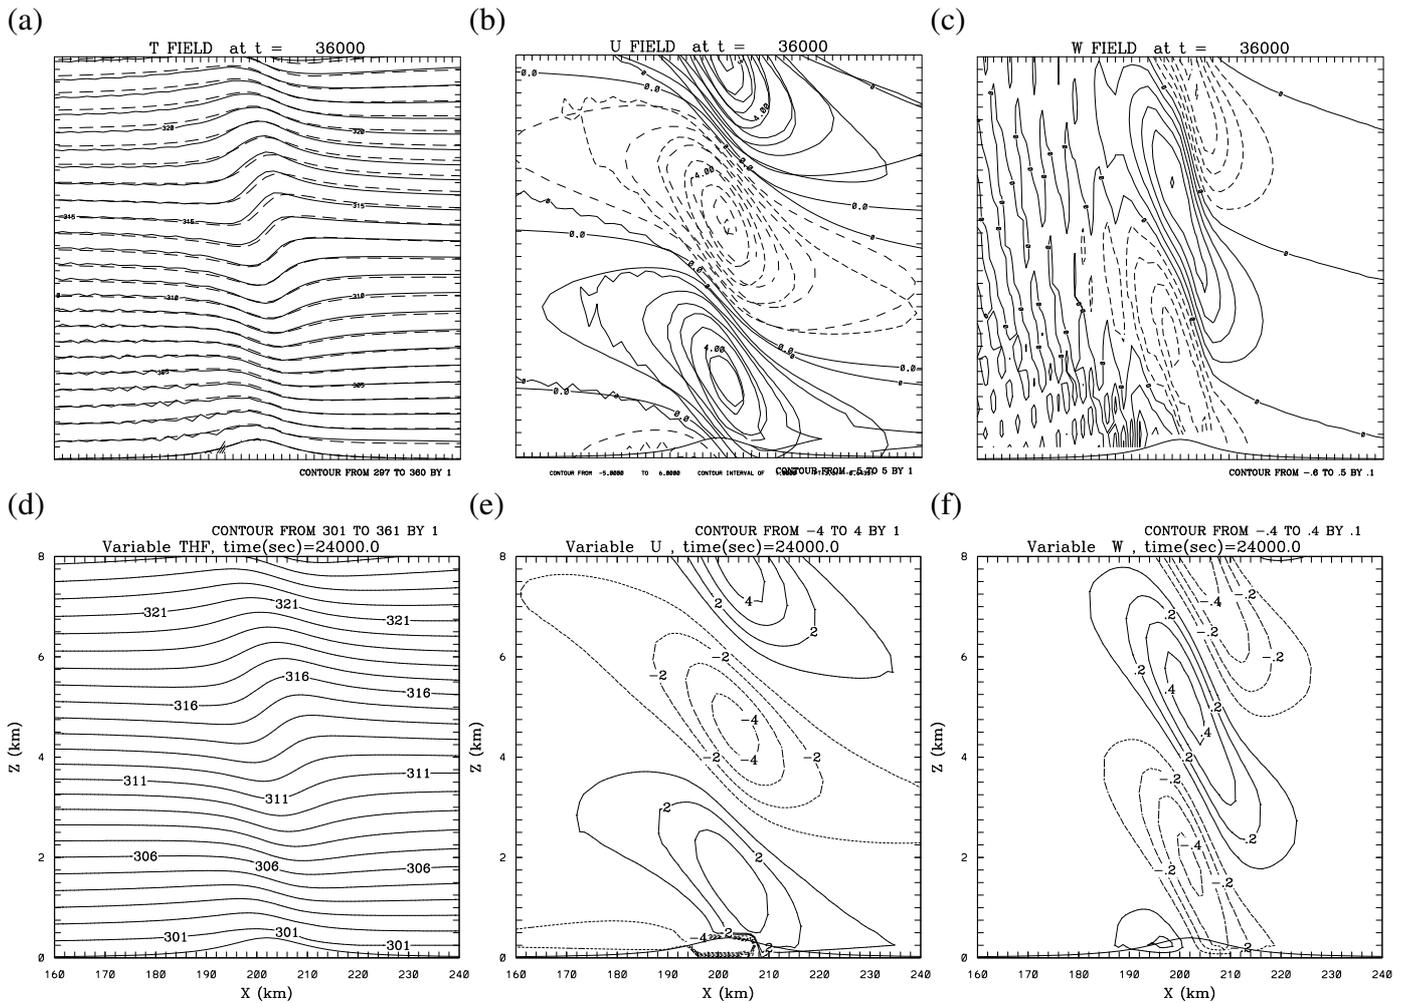


Figure 1: Results for the “Base Case” test: (a), (b) and (c) show the potential temperature field, horizontal wind field U' as a perturbation from the mean wind and vertical wind field respectively, from scmc.Rev1.f90 at model time $t = 36,000$ s (heavier lines in (a) and (b) are model solutions, fainter lines indicate Long’s solution); (d), (e) and (f) show plots for the equivalent model fields from VHREM at time $t = 24,000$ s.

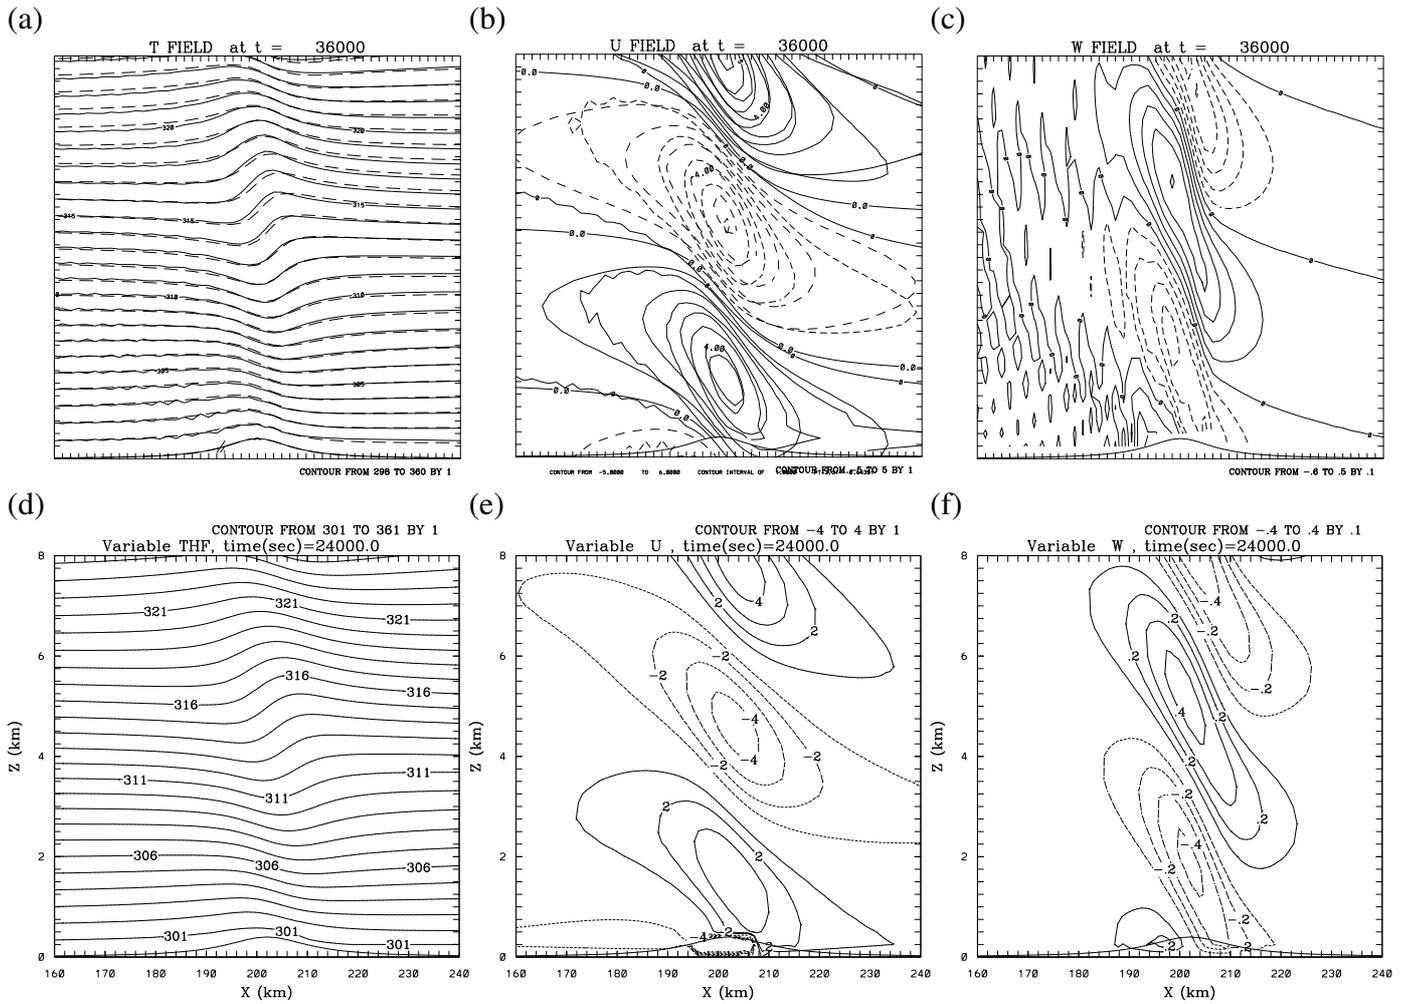


Figure 2: Results for testing the “thin-wall approximation”: (a), (b) and (c) show the potential temperature field, horizontal wind field U' as a perturbation from the mean wind and vertical wind field respectively, from scmc.Rev1.f90 at model time $t = 36,000$ s; (d), (e) and (f) show plots for the equivalent model fields from VHREM at time $t = 24,000$ s.

approximation (“orog2B_DV1”). In addition, the differences between the two plots are also displayed for contour intervals of 0.1K and 0.1ms^{-1} for θ and w respectively.

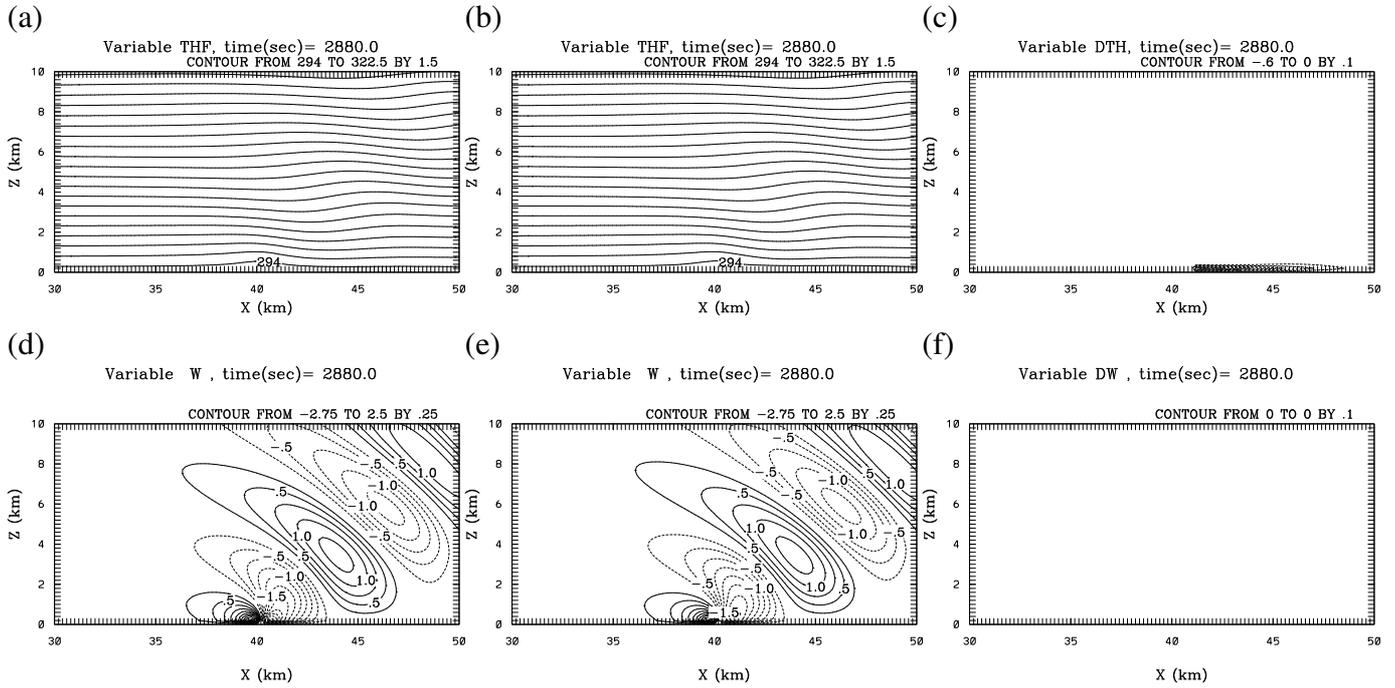


Figure 3: Results from VHREM, exploring the “thin-wall approximation” for a hill of half-width 1km. Results are displayed for model time $t = 2,880\text{s}$: (a) shows the potential temperature field (θ) for model solutions based on cut-cell volumes (experiment “orog2B”); (b) shows θ for the thin-wall approximation (“orog2B_DV1”); and (c) the difference between the two fields (“orog2B_DV1” minus “orog2B”); (d), (e) and (f) show the equivalent plots for the vertical wind field.

The differences in the vertical wind field between the two cases are less than 0.1ms^{-1} . In the potential temperature field, the only significant differences (up to 0.6K) are at the lowest level(s) in the lee of the hill. These are very early results which have not yet been investigated further.

Further exploration of the thin-wall approximation is intended with VHREM. In particular, the following aspects will be considered:

- A coherent investigation of the limits on the length of stable time-step for using cut-cell volumes and uncut volumes;
- A definitive check that the difference in time-step lengths between the two approaches can be directly related to the difference in cut-cell geometries, i.e. that the stable time-steps can be related by the CFL condition to the smallest effective grid-lengths in both cases;
- Further investigation of the resulting differences in model solutions between the two methods — why are there significant differences in the THF field at the lowest model levels but nowhere else?
- An exploration of whether the small effect of adopting the thin-wall approximation observed for the “orog2B” case can be expected for other orographic flow regimes.

3.3 Tests: shorter time-step

As mentioned in the description of the model set-ups for the Base Case (Section 3.1), to obtain stable results, VHREM must be run with a much shorter time-step than is used in the `scmc.Rev1.f90` code. For the Base Case, the short time-step used in VHREM was $\Delta\tau=0.2\text{s}$; whereas, in `scmc.Rev1.f90` $\Delta\tau=5\text{s}$. VHREM was not able to produce stable results for a short time-step $\Delta\tau \geq 0.4\text{s}$.

The CFL condition on the short model time-steps is determined from the ratio of the smallest effective grid-length in the model domain to the speed of acoustic waves ($\sim 300\text{ms}^{-1}$). In a cut-cell model, the smallest effective grid-length is given by the smallest ratio of grid-cell volume to grid-cell face area, which will be dependent on the shape and position of the hill with respect to the model grid (and could be arbitrarily small).

A major difference between VHREM and the WRF codes is that the WRF codes use an implicit solver for the vertical motions, meaning that the CFL condition is limited by the smallest effective *horizontal* grid-length, and need not take account of the smallest vertical grid-length. VHREM is fully explicit, so is limited by the vertical grid-lengths as well.

For the Base Case, the grid resolutions were $\Delta x = 2,000\text{m}$ and $\Delta z = 250\text{m}$. For such a grid, with *no* cut-cells, the CFL condition would imply $\Delta\tau \sim 2000\text{m}/300\text{ms}^{-1} \sim 6\text{s}$ for the WRF codes, but $250\text{m}/300\text{ms}^{-1} \sim 0.8\text{s}$ for VHREM; i.e. it should be no surprise that for the Base Case grid resolutions, the stable time-steps for VHREM and the WRF codes differ by about a factor of 10. However, it should be some surprise that the WRF codes are able to use a time-step based on an *un-cut* grid, despite having cut-cells which will make the effective horizontal grid-length smaller than $\Delta x = 2,000\text{m}$.

During the VHREM set-up of the cut-cells, the smallest effective grid-length is computed and is used directly to determine the length of the short time-step. In the `scmc.Rev1.f90` code, the length of the time-step is determined from the hill dimensions and *not* the resulting cut-cell geometry.

To explore the potential effect of using a shorter time-step in the `scmc.Rev1.f90` code, the long time-step was reduced to 10% of its original value — i.e. $\Delta\tau$ became 0.5s ($\Delta t = 2\text{s}$) — and the Base Case was repeated. Results for the θ , U' and w fields at model time $t=36,000\text{s}$ are presented in Figure 4.

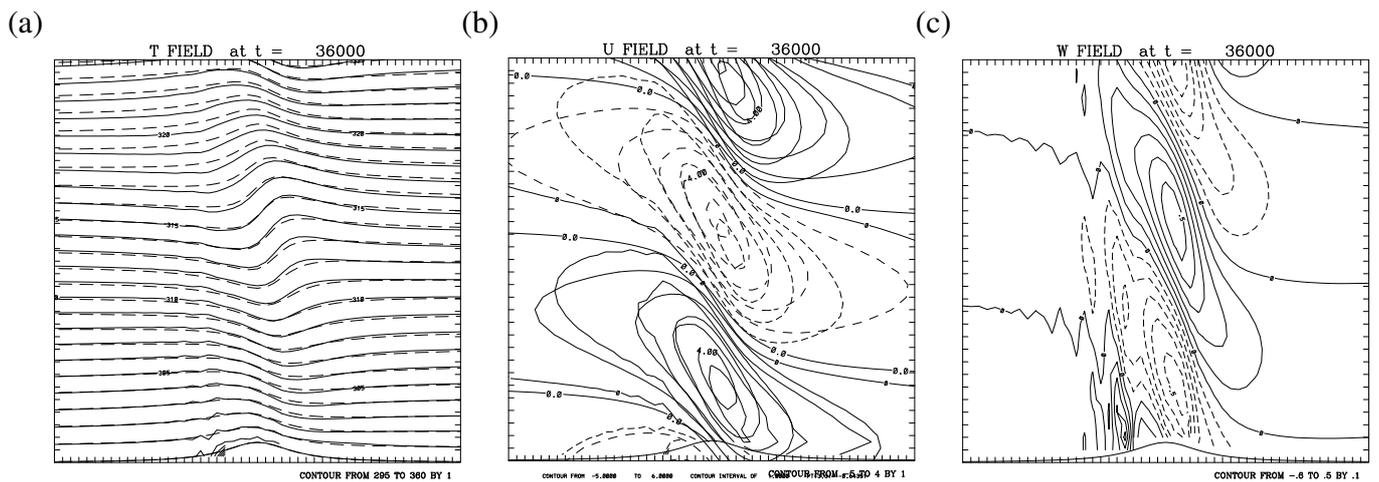


Figure 4: Results for testing a shorter time-step in the `scmc.Rev1.f90` model: (a), (b) and (c) show the potential temperature field, horizontal wind field U' as a perturbation from the mean wind and vertical wind field respectively at model time $t = 36,000\text{s}$.

Comparing the results for the shorter time-step shown in Figure 4 to those from the Base Case, it is apparent

that the disturbances upstream of the hill are greatly reduced in the velocity fields, with particular improvement seen near the lower boundary in the U' field. In the θ field, the difference is less clear-cut — there is still significant disturbance to the upstream flow at the lowest vertical levels using a shorter time-step. At higher levels and in the downstream flow, the solution differs from that for the Base Case, but is not noticeably closer to Long’s solution.

3.4 Tests: coarser vertical resolution

Finally, the models were run as for the Base Case, but with coarser vertical resolution. The spacing between vertical levels was chosen such that the first w level (above the model bottom) is higher than the top of the hill: $\Delta z = 500\text{m}$ for the hill height 400m . (For the `scmc.Rev1.f90` code, the time-steps were as for the Base Case: $\Delta\tau = 5\text{s}$, $\Delta t = 20\text{s}$.)

The results at model time $t = 36,000\text{s}$ from the `scmc.Rev1.f90` code are displayed in Figure 5. They show a large improvement on the results from the Base Case (with $\Delta z = 250\text{m}$). There are no obvious remaining upstream disturbances in the θ or U' fields, although some small oscillations are still apparent in the w field. Comparing the model solutions to Long’s solution (shown in Figure 5 in fainter contours), the general position and magnitude of the wave in the U' field continue to look good. However, the contours in the θ field appear to compare less favourably with Long’s solution using a coarser vertical resolution — the differences between the model field and Long’s solution are largest at higher vertical levels, both up- and downstream of the hill.

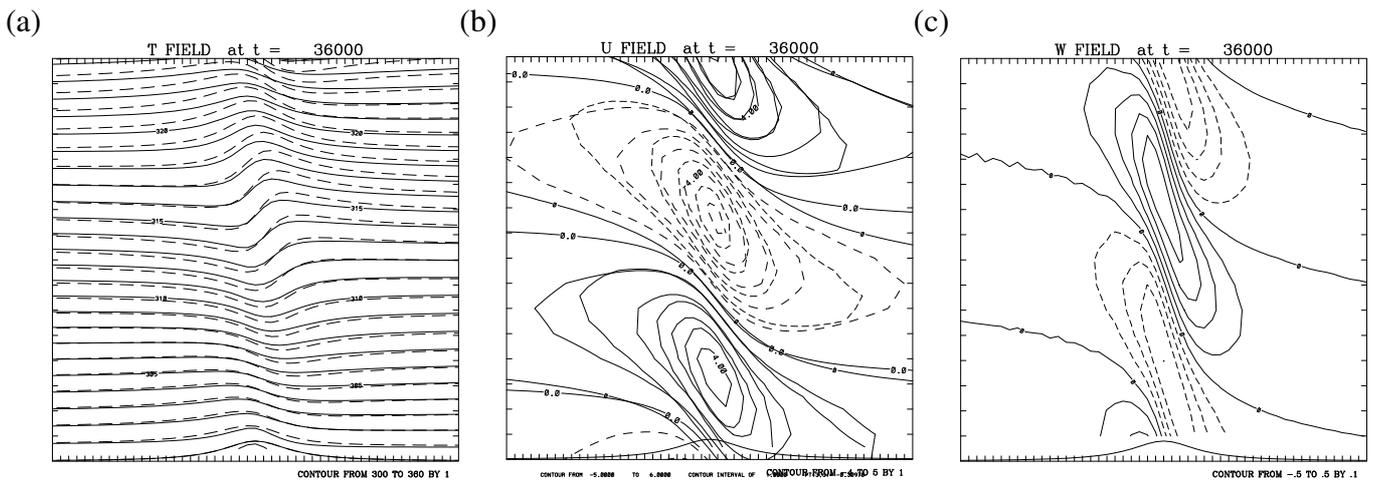


Figure 5: Results for testing a coarser vertical resolution ($\Delta z = 500\text{m}$) in the `scmc.Rev1.f90` model: (a), (b) and (c) show the potential temperature field, horizontal wind field U' as a perturbation from the mean wind and vertical wind field respectively at model time $t = 36,000\text{s}$.

An equivalent set-up was attempted with VHREM, but an initial attempt did not yield any meaningful results — after setting $\Delta z = 500\text{m}$, the flow-fields showed no sign of detecting a hill and remained unchanged from the initial values! (This test has not since been repeated or further explored.)

4 Observations

From the tests demonstrated in the previous section, a couple of observations and resulting questions were particularly apparent:

- In none of the tests demonstrated above (nor in any earlier cases) has VHREM been seen to suffer from the spurious upstream disturbances seen in the results from the WRF codes. What is causing the difference in the two models?
- VHREM uses a time-step scaled to the smallest effective grid-length — for the Base Case set-up, $\Delta\tau$ is around 25% of the theoretical limit on the time-step for the equivalent set-up with no cut-cells. The WRF codes use a time-step length appropriate for a domain of uncut cells to achieve a steady, albeit noisy, solution. However, by reducing the length of the time-step by a factor of 10, the solutions generally appear to improve.

5 Potential problems in the WRF codes

Alongside the test cases demonstrated in Section 3, the `scmc.Rev1.f90` code was explored for potential problems. One potential source of error was noted, and is summarised here.

Consider the advection terms for the flux-form equations for the winds:

$$\begin{aligned}\frac{\partial}{\partial t}U &= -\{\nabla \cdot (\mathbf{u}U)\} + \dots, \\ \frac{\partial}{\partial t}W &= -\{\nabla \cdot (\mathbf{u}W)\} + \dots,\end{aligned}$$

where $U = \rho u$, $W = \rho w$. The quantities U and W are advected by the winds \mathbf{u} . When accounting for fluxes through *cut*-cells, it is the flux of the winds (\mathbf{u}) across the cell faces that should be adjusted proportional to the fraction of the cell-face that lies above the orographic surface; it is *not* the advected quantities U and W that are adjusted.

On the staggered computational grid, the quantities U and W are represented by $\bar{\rho}^x u$ and $\bar{\rho}^z w$ respectively, where

$$\bar{\rho}^x \equiv 0.5 (\rho_{i-1/2,k} + \rho_{i+1/2,k})$$

and similarly for $\bar{\rho}^z$ — i.e. ρ must be averaged across the grid-cell to provide a value which coincides with the position of u or w , as illustrated in Figure 6.

In the set-up for `scmc.Rev1.f90`, the variables “RU” and “RW” are generated, to be used as the prognostic variables for the winds. The quantities are defined to be

$$\begin{aligned}\text{RU}(K,I) &= 0.5 \{RHO(K,I) + RHO(K,I+1)\} U(K,I) \text{DZU}(K,I), \\ \text{RW}(K,I) &= 0.5 \{RHO(K,I) + RHO(K-1,I)\} W(K,I) \text{DXW}(K,I),\end{aligned}$$

such that RU represents the quantity $\bar{\rho}^x u$, *premultiplied* by DZU, which represents the fractional grid-cell length that lies above the orographic surface for the face on which u is stored (and similarly for RW).

As outlined above, the fractional grid-cell lengths, DZU and DXW, should act on the advective winds at the cell-faces; and should not therefore be applied to the prognostic variables in this way. By pre-multiplying the prognostic variables by the fractional cell lengths, the order of differencing in the subsequent computations is incorrect.

For example, take the w -advection of the $U \equiv \bar{\rho}^x u$ variable (stored at location (i,k) in Figure 6), which should follow the form:

$$\frac{\partial}{\partial z} \{wU\}.$$

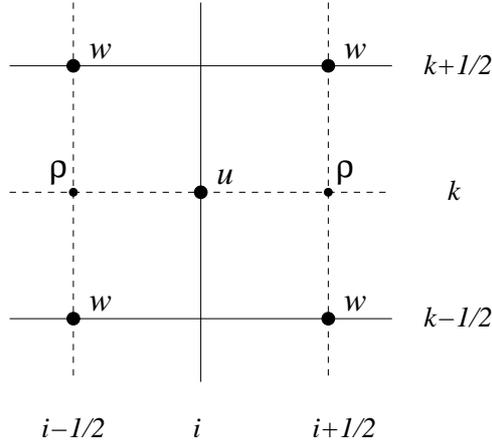


Figure 6: Illustration of the position of model variables on the staggered grid — variables u and w are positioned half a grid-length from ρ in the horizontal and vertical directions respectively.

On the staggered computational grid, differencing in the z -direction to compute this quantity requires several layers of averaging:

$$\begin{aligned} & \frac{1}{\Delta z} \left\{ (\overline{w^x U^z})_{i,k+1/2} - (\overline{w^x U^z})_{i,k-1/2} \right\} \\ & \equiv \frac{1}{\Delta z} \left\{ (\overline{w^x \overline{\rho^x u^z}})_{i,k+1/2} - (\overline{w^x \overline{\rho^x u^z}})_{i,k-1/2} \right\}. \end{aligned}$$

Including the effect on the advective flux w from the cut-cell faces, the computed quantity becomes

$$\frac{1}{\Delta z} \left\{ (\overline{\Delta x_w w^x \overline{\rho^x u^z}})_{i,k+1/2} - (\overline{\Delta x_w w^x \overline{\rho^x u^z}})_{i,k-1/2} \right\}, \quad (1)$$

where Δx_w denotes the fraction of the cell-face with centre w that lies above the orographic surface (represented by DXW in the model code).

Compare that to what is currently computed in the scmc.Rev1.f90 code. The flux divergence of term ρu contributes to the slow modes computation (stored as FU in the code) by:

1. First, defining a quantity WDUZ:

$$\begin{aligned} \text{WDUZ(K)} &= 0.5*(\text{RW(K,I)} + \text{RW(K,I+1)}) && \longleftrightarrow \overline{\rho^z \Delta x_w w^x} \\ & \quad * \text{RDZ} && \longleftrightarrow 1/\Delta z \\ & \quad * 0.5*(\text{U(K,I)} + \text{U(K-1,I)}) && \longleftrightarrow \overline{u^z} \end{aligned}$$

i.e. referring to the grid-locations represented in Figure 6

$$\text{WDUZ(K)} \longleftrightarrow \frac{1}{\Delta z} \left(\overline{\rho^z \Delta x_w w^x u^z} \right)_{i,k-1/2};$$

2. and then, taking the difference of WDUZ at neighbouring vertical levels to complete the contribution:

$$(\text{WDUZ}(K+1) - \text{WDUZ}(K)) \longleftrightarrow \frac{1}{\Delta z} \left\{ \left(\overline{\rho^z \Delta x_w w^x \bar{u}^z} \right)_{i,k+1/2} - \left(\overline{\rho^z \Delta x_w w^x \bar{u}^z} \right)_{i,k-1/2} \right\}. \quad (2)$$

It can be seen that the flux divergence terms from (1) and (2) are not the same — in particular, the contribution from ρ differs due to differences in the order in which the averaging over grid-locations takes place. Ultimately, the difference in the order of calculations may make have a negligible effect on the model solutions. However, for completeness, the necessary adjustments to correct the code should be made – specifically, to remove the pre-multiplication of RU and RW by DXU and DZW respectively, i.e.:

$$\begin{aligned} \text{RU}(K, I) &\rightarrow 0.5 \{ \text{RHO}(K, I) + \text{RHO}(K, I + 1) \} U(K, I), \\ \text{RW}(K, I) &\rightarrow 0.5 \{ \text{RHO}(K, I) + \text{RHO}(K - 1, I) \} W(K, I). \end{aligned}$$

Under this approach, RU and RW are the true prognostic variables, which are acted on by each of the contributing factors in the integration. Where RU and RW are acted on by the fluxes of the winds U and W, adjustment factors DZU and DXW would need to be applied respectively (to represent the flux-divergence described by (1)).

The correction outlined above for RU and RW has not been completed in the `scmc.Rev1.f90` code, and therefore, its effect on the results has not been tested. From a quick glance at the code, it would seem that the correction is not required for the other prognostic variables — potential temperature and density. Although, where they are multiplied by the currently defined RU and RW (which currently include the factors of DZU and DXW respectively), a correction factor would need to be added.

6 Next steps

Based on the analyses above, there are a number of different avenues to explore with both the WRF codes and VHREM to better understand the potential of the cut-cell method:

- WRF codes
 1. Correct the averaging terms in the WRF codes (as described in Section 5) & re-test — does this improve the solutions: reducing the upstream noise/producing solutions closer to Long’s or the linear solutions? (The correction could first be explored for the partial finite-volume method in `scmc.Rev1.f90`; and then, for the full finite-volume model in `scmc.Rev2.f90`.)
 2. Further consideration of the length of the time-step used in the WRF codes — a simple algorithm would enable the length of time-step to be determined by the geometry of the smallest cut-cells. Does such a change result in the solution improving when $\Delta\tau$ is appropriate to the CFL limit for the given model set-up?
 3. Further explore the improvement seen for the coarser vertical resolution — is the improvement associated with the first vertical level sitting above the top of the hill or simply with a coarser resolution, i.e. test with other hill heights (and widths?) and other vertical resolutions.

- VHREM

1. Further explore the thin-wall approximation — the resulting accuracy of the solutions and potential efficiency gains from using a longer time-step (summarised in Section 3.2).
2. Go back and re-explore the case with coarser vertical resolution where the first vertical level sits above the top of the hill. In theory, there is no obvious reason for the model to fail — was there an error in the model set-up or is there a problem with the code?
3. Move towards a more complete finite-volume representation of the cut-cell method by re-formulating the advection terms as two flux terms:

$$\mathbf{u} \cdot \nabla \phi = \nabla \cdot (\mathbf{u}\phi) - \phi (\nabla \cdot \mathbf{u})$$

both of which can be handled using Gauss's theorem, without having to re-cast the entire model equation set in flux-form.

7 References

- Adcroft, A., Hill, C., Marshall, J., 1997. Representation of topography by shaved cells in a height coordinate ocean model. *Monthly Weather Review* 125, 2293–2315.
- Bonaventura, L., 2000. A semi-implicit semi-Lagrangian scheme using the height-coordinate for a non-hydrostatic and fully elastic model of atmospheric flows. *Journal of Computational Physics* 158, 186–213.
- Gallus, W. A., Klemp, J., 2000. Behavior of flow over step orography. *Monthly Weather Review* 128, 1153–1164.
- Steppeler, J., Bitzer, H., Minotte, M., Bonaventura, L., 2002. Nonhydrostatic atmospheric modeling using a z -coordinate representation. *Monthly Weather Review* 130, 2143–2149.