**4.9**

# HIGH RESOLUTION MODELLING OF CONVECTION OVER THE UK

Nigel M. Roberts *
Met Office, JCMM, UK

## 1 INTRODUCTION

Convection over the UK is extremely variable in character, ranging daily from small-scale scattered showers to more organised mesoscale clusters, frontal convection and deep summertime thunderstorms. An aim of the Met Office is to develop a high-resolution modelling system that is capable of dealing with this variability and still provide a reliable nowcasting tool for severe convective events. One of the main issues in high-resolution modelling is the use of and requirement for convection schemes.

Convection schemes were designed for global and climate models with large grid squares over which convection is not resolved. They are probably not appropriate for high-resolution models (gridlengths < 15-20 km). This is likely to be particularly true for models with gridlengths in the range 2 – 10 km in which convective storms are often only partly resolved. At gridlengths of 1 km or shorter most rain producing convection can be resolved and a convection scheme may not be needed for practical weather forecasting.

Simulations of convective events using the new semi-implicit, semi-lagrangian, non-hydrostatic version of the Met Office Unified Model have been run with a gridlengths of 12, 4, 2 and 1 km. The results were very sensitive to the resolution of the model (both horizontal and vertical), the way the convection scheme was used and the nature of the event itself. Results from two contrasting case studies will be briefly presented here to show the nature of the problem, which is particularly noticeable in 4-km gridlength simulations. The convection scheme was run in two different ways; firstly, with the standard closure used in the operational 12-km model and secondly with a modification to the closure that is thought to be more physically sensible in a high-resolution context.

## 2 THE PROBLEM WITH USING A CONVECTION SCHEME AT HIGH RESOLUTION

At the resolution of the Met Office global model (~60 km over the UK) it is not possible to resolve convection on the model gridpoints. Even the larger thunderstorms are too small to span a few grid squares. This means that some other way of representing convection is required and a convection scheme is used. Currently, the Unified Model (Cullen et al) uses an equilibrium mass flux scheme (Gregory and Rowntree 1990) which represents all the convection in a grid square as a single plume that is in equilibrium with any larger-scale tendency of the atmosphere to become convectively unstable.

At the resolution of the Met Office mesoscale model (~12 km gridlength) a convection scheme is still essential, but there are questions about whether a scheme that was designed for a much coarser resolution model is still appropriate. At higher resolutions (1 to 5 km gridlength), the theoretical basis of using an equilibrium convection scheme is even more questionable and it becomes much less clear when (or if) a convection scheme is needed.

Two questions that should be asked to start with as the resolution of numerical models is increased are:

1. Is it reasonable to use a convection scheme at high resolution that was designed for a coarse resolution model? If not, are there alternatives?
2. At what resolution is a convection scheme no longer required?

### 2.1 Assumptions

The convection scheme was formulated with a number of assumptions that are appropriate for a model with a gridlength of ~60 km, but become invalid in the range of gridlengths (1 – 12 km) that a storm-scale forecast system might use. Assumptions are: (Swann 2001)

1. Convection is in quasi-equilibrium with the forcing of instability over a grid square. This is not a good approximation for small grid squares in mid-latitudes when convection often responds to significant dynamical forcing that can be large, transient and act on

---

* *Corresponding author address*: Nigel M. Roberts, Met Office, JCMM, Meteorology Building, University of Reading, Reading, Berks, RG6 6BB, UK.
e-mail: nigel.roberts@metoffice.com

scales close to the gridscale of the model. For example, the passage of a frontal zone would make this assumption invalid.

2. The area of the updraughts in a grid square is assumed to be small compared to the grid square. This is clearly not a good approximation for small grid squares. A single updraught in a large thunderstorm might occupy an entire square if the square is small enough.
3. The convection is assumed to be in a steady state. This means that it is impossible for the convection scheme to represent any developing or decaying clouds - something we ought to be able to do in high-resolution models.

These assumptions mean that we have to be concerned about whether it is appropriate from a theoretical point of view to use the convection scheme in high-resolution models.

### 2.2 Limitations

In addition, we need to think about some other aspects of the behaviour of the scheme and what we expect from it. The convection scheme is supposed to represent the average effects of convection over a single grid column and does not know what other grid columns are doing. It can not propagate showers or develop convective organisation. This is not so much of a problem with large grid squares when we do not expect to see much convective organisation on the scale of the grid, but for grid squares that are a similar size to the area of a storm cloud it is not realistic for each grid column not to know what the adjacent columns are doing. The upshot of this is that the convection scheme will (if it is working correctly) produce a rainfall picture that is a smoothed average over an area rather than develop individual showers. This means that the precipitation will not look very much like a radar picture, which is fine if that is what is expected and required, but is not so useful for a high-resolution modelling system that is meant to simulate individual storms.

Another consideration for high-resolution modelling is how the convection scheme will interact with the model dynamics in situations when some convection is resolved by the dynamics and the convection scheme is also triggering. The only way to find out is to run experiments and see what happens.

It would seem then, that we are in a difficult position. It ought to be much more desirable to run a high-resolution modelling system without using the convection scheme because of the reasons mentioned above. On the other hand, it is likely that a convection scheme of some sort is still required, at least for gridlengths > 1 km, to prevent unrealistic grid-scale storms. A first step towards an answer the question of whether it is possible to get away with using the current convection scheme (or other schemes) at higher resolutions (gridlength of 1 – 5 km), or even run with no convection scheme, is to try these alternatives on some case studies.

## 3 CASE STUDIES

### 3.1 Squall line 2nd July 1999

Deep thunderstorms propagated rapidly across the southern half of the UK ahead of an upper-level vorticity anomaly and developed a squall line structure. Maximum values of CAPE in model simulations exceeded 3000 J/kg.

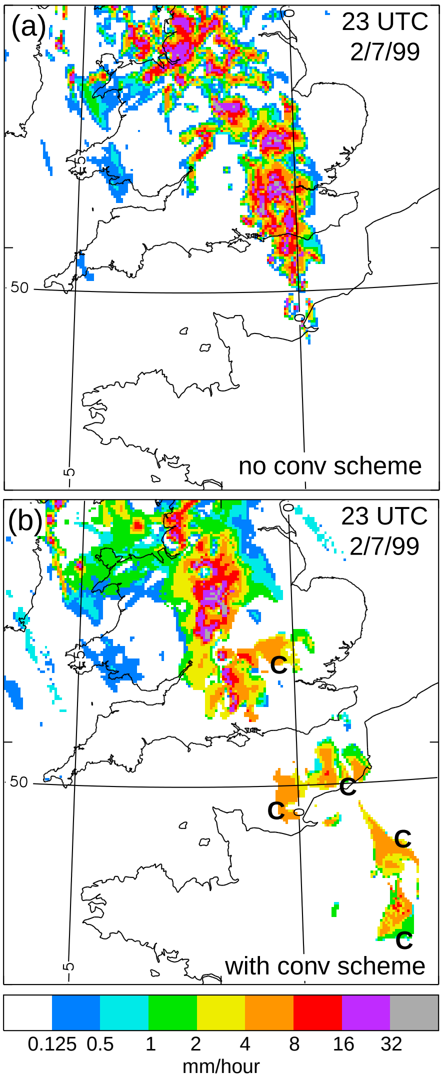

**Figure 1. Snapshots of the rainfall rate from two 4-km gridlength simulations of storms on 2nd July 1999 (a) with no convection scheme, (b) including the convection scheme (as 12-km operational)**

Figure 1 shows the very substantial difference between a 4-km gridlength model run with the convection scheme and a run without. An analysis of the forecasts (Roberts 2001) concluded that the run with no convection scheme was much better. Weisman et al also concluded that 4 km is sufficient to reproduce most of the structure produced by a 1km-squall line simulation. The run which used the convection scheme had a serious problem with the formation of bands or arcs of precipitation from the convection scheme (labelled C) that propagated through the domain. The bands were self-sustaining because of an interaction between the convection scheme and the dynamics. As well as producing rain in completely the wrong place, they led to the removal of the convective instability that was required to trigger storms over the correct region.

The probable mechanism of the dynamics/convection scheme interaction is depicted in Figure 2. The convection scheme was firstly triggered at gridpoints where there was sufficient moisture and local ascent to make the profile conditionally unstable. The convection scheme then cooled the profile below the convective plume. If convection through the scheme continued, then further cooling generated a low-level cold pool and the dynamics responded with a region of convergence and ascent ahead of the cold pool as a local frontal zone or density current structure developed. As the density current became established the ascent ahead of the cold pool acted to destabilise the profile in that location and trigger further convection which in turn cooled the region ahead of the cold pool and propagated forward the cold pool and ascent region. The model was, in fact, responding in a reasonable way, but the timescale of the response was too fast since in reality a convective cloud will have no downdraught until it is ~20mins old. Evidence for this mechanism is shown in Roberts (2001).

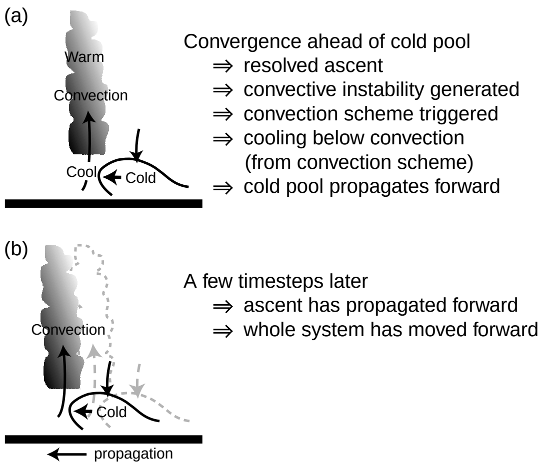

**Figure 2. Schematic of the mechanism described in the text.**

Clearly, it is not desirable to generate these convective bands in a forecast and they need to be removed, but in a way that does not have an adverse effect on the overall performance of the model. The most obvious way to remove the bands and still run with the convection scheme is to reduce the activity of the convection scheme or switch it off altogether at 4km.

### 3.2 Scattered convection 3rd May 2002

This case is presented because it is significantly different from the one just discussed. Rather than being deep and organised, the convection was mostly in the form of smaller scattered showers and thunderstorms that developed during the day. Maximum values of CAPE were around 300 J/kg. A convection scheme should be essential in a 4-km gridlength model in this situation because many of the showers were too small to be properly resolved.

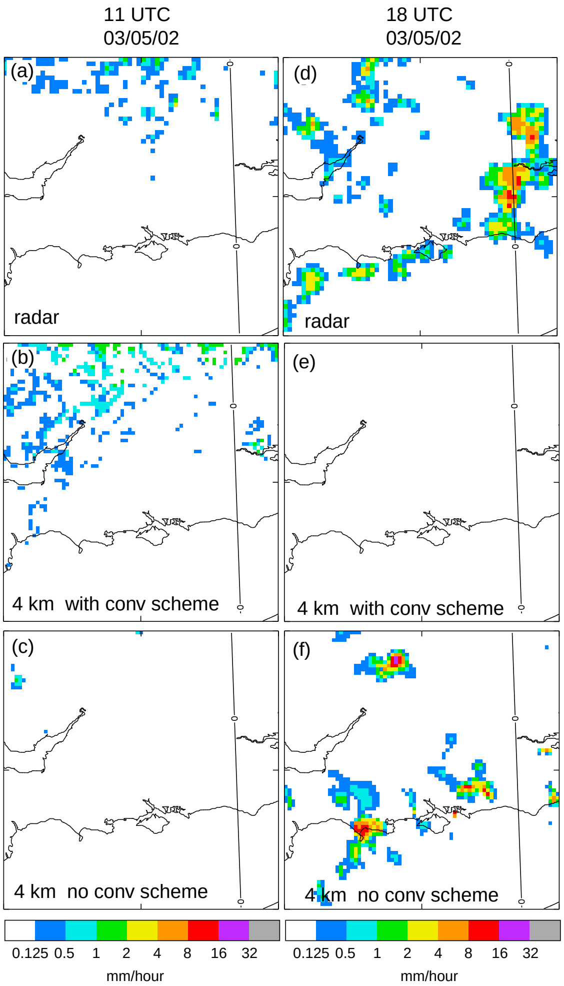

**Figure 3. Precipitation rates at 11and 18 UTC 03/05/02 from radar and from 4km-gridlength forecasts starting at 01 UTC with and without the convection scheme. Viewing 300x300 km area.**

Two snapshots of the showery day are shown in Figure 3 to compare the behaviour of the 4-km runs with and without the convection scheme. At 11 UTC the showers were just starting to develop. The run with the convection scheme had developed scattered precipitation at this stage of a similar intensity to the radar, though more widespread. The run without the scheme had hardly triggered any showers at 11 UTC. By 18 UTC the showers had become larger and more organised. The run with the scheme was no longer producing any precipitation, but the other run did have showers and evidence of organisation.

### 3.2.1 Rain rates, triggering and organisation

Figures 4 and 5 reveal that the run with the convection scheme has produced rates that are far too low. They do not even exceed 3 mm/hour and although this would equate to reasonable peak rates if the convection occupied around 10% of a gridsquare, the radar picture (also at 2-km gridlength, not shown) reveals that the model precipitation should have been representative of considerably more than 10% of many of the gridsquares.

The run without the scheme has produced rates that are far too large. Some rates are extreme with values exceeding 150 mm/hour. Such unrealistically high values in the run without the scheme emphasise why a convection scheme is used. The problem here is that we are no better off using the convection scheme if it means going to the other extreme of producing unrealistically low rates instead. An alternative might be to run with the convection scheme tuned to be more active, but, although it may help solve this particular problem, it could be catastrophic if applied to the previous case.

The times of initial shower development can be seen in Figure 6. Significant shower activity started around 10.00 UTC. In the run with the convection scheme it was around 15 minutes earlier and in the run without the scheme around an hour later. The delay of 1 hour in the no-scheme run is significant, although it is not surprising that this should happen because the smallest scale of the showers that can be generated by the model dynamics is determined the model gridlength. Triggering will not occur until the convective instability has become sufficiently large for showers of a gridlength or larger to form. The convection scheme does not suffer from this problem as it is attempting to represent showers on all scales.

Figures 3 and 6 both show that showers continued through the late afternoon. During this period, the precipitation in the run with the convection scheme gradually died out instead of persisting, and had entirely gone by 17.30 UTC. All of the precipitation came from the convection scheme in this run. The average rain rates were close to that observed in the period up to 15.00 UTC because the convection scheme was operating in equilibrium with the larger-scale forcing, which in this case was the solar heating. After 15.00 UTC, the solar heating became too weak to support the same level of convection, so showers could only be maintained through convective organisation, but the convection scheme is incapable of such organisation and the precipitation died away.

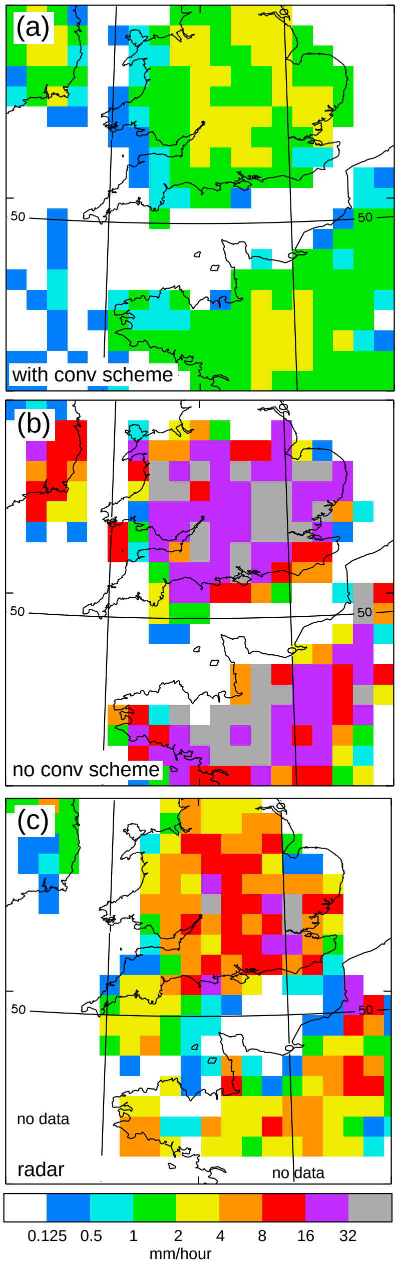

**Figure 4. (a)-(c) Maximum precipitation rates to occur over 40x40km squares in the period 10 to 18 UTC 03/05/02, sampling every 15 minutes from 4-km gridlength forecasts starting at 01 UTC (a) with convection scheme included, (b) no convection scheme, (c) network radar (5km grid).**

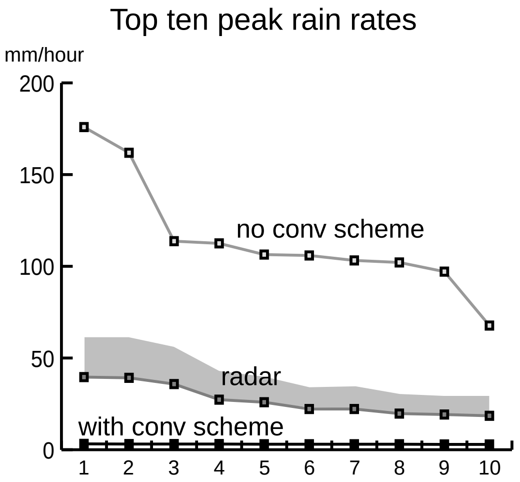

**Figure 5. Graph of the top ten peak rain rates taken from Figure 4.**

The run without the scheme could produce organisation – and did – but shower activity still died away too quickly because convection had become far too intense in the middle of the afternoon and removed too much of the convective instability. Figure 6 shows that the run without the scheme had produced an average rainfall rate of twice that observed.

### 3.3 Issues raised by the case studies

It is clear that, for a gridlength of 4 km, neither using the convection scheme or switching it off is satisfactory.

Switching off the convection scheme was the best choice for simulating the large storms that could be resolved on the model grid but was a poor choice for representing the smaller-scale scattered convection that could only be partly resolved. In the case of the scattered showers, the convection was triggered too late, the storm cells, when they did develop, became too large and intense before dying out too quickly. At the early stages some unrealistically intense single-grid-point cells developed. Even in the severe case the rainfall intensity was too high and the first cells that formed aliased onto the gridscale and produced extreme rainfall rates.

When the convection scheme was used in the simulation of the severe convection case, an unrealistic interaction between the convection scheme and the model dynamics developed.
In the scattered convection case, the use of the convection scheme produced a reasonable forecast at first. However, because all of the precipitation came from the convection scheme, the cells could not become organised and the rain died out far too quickly.

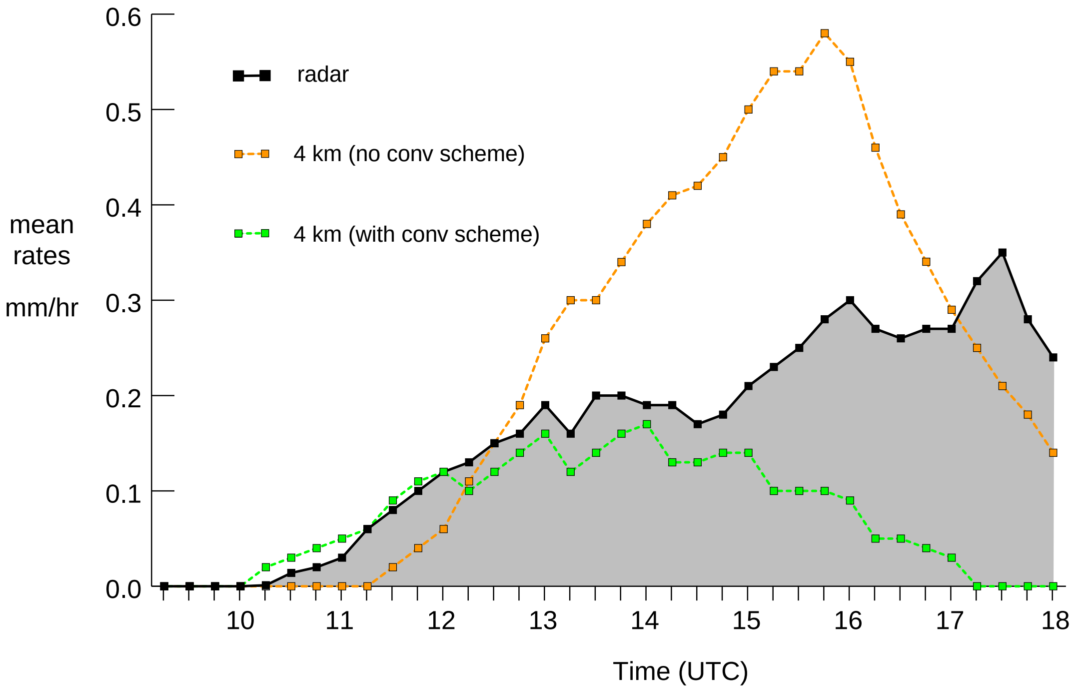

**Figure 6. Graph of the mean rainfall rates within the area shown in Figure 3 against time for the two forecasts discussed in the text (coloured lines) and the radar (black line enclosing grey shading).**

## 4 A CHANGE TO THE CONVECTION SCHEME CLOSURE

The ideal scenario for high-resolution simulations of convection is for the model to explicitly resolve all the convection that it should be able to resolve and leave the rest to a well-behaved convection/turbulence scheme. If the convection scheme does too much, we are not getting all of the benefit we should from a higher resolution model and might as well be running at coarser resolution. All the clouds that are large in comparison to a model gridlength (> 3 gridlengths) should be simulated by the model dynamics and the small clouds represented by the convection scheme. In practise we have seen that this does not happen. The current convection scheme is not scale selective in that way. This is because the intensity (cloud base mass flux M) of a convective plume is tuned by a single number called the CAPE Closure Timescale $\tau$ (CCT), which is defined as the timescale over which the Convectively Available Potential Energy (CAPE (J/Kg)) in an atmospheric profile is reduced to zero (relationship (1)). This means that the convection in the scheme is always more intense when the CAPE is larger, regardless of whether the model dynamics should be able to resolve the convection or not.

$$M \propto CAPE / \tau \qquad (1)$$

That is why the convection scheme is more active in more convectively unstable situations (first case study) and therefore more likely to generate spurious convective rainbands and inhibit the development of resolved convection. The way to stop the spurious convective bands from developing is to lengthen the CCT (increase $\tau$). This has been tried – and works – but the problem is that in order to reduce the intensity of the convection scheme in high-CAPE regions it has to be reduced to very little indeed in low-CAPE regions. By doing this, the convection scheme is not able to sufficiently represent the smaller clouds – and these are precisely the clouds we want to represent with a convection scheme in a high-resolution model. Evidence from idealised convection simulations (Cohen 2002) has indicated that the convective timescale should be related to inter-cloud spacing and shorter for smaller clouds that are closer together.

A way round this may be to use a CAPE dependent CCT. If the CCT is made longer wherever the CAPE is larger then it should be possible to limit the intensity of the convection scheme when we want the model dynamics to do more. That is what has been done by using Equation 2 to calculate the CCT. An assumption behind this is that the size of a convective cloud is related to the CAPE in the environment (large CAPE means big clouds), and that the model should therefore be allowed to explicitly resolve more convection in regions of high CAPE. This assumption is flawed to some extent because there are other factors in addition to CAPE that determine the size of convective showers, but it may not be so bad an assumption because the general trend will hold. Shallow convection is restricted to low-CAPE regions and large summer thunderstorms and mesoscale convective systems do develop in high-CAPE regions.

$$\tau = t/c * CAPE + t * \exp -(CAPE/c) \qquad (2)$$

t and c are tuneable parameters

Figure 7 shows a graphical representation of equation 2. The CCT increases exponentially with CAPE for low values of CAPE, then linearly with CAPE for high values of CAPE. This means that the maximum allowed cloud-base mass flux increases with CAPE for small values of CAPE and is then restricted to a limiting value wherever there are larger values of CAPE.

This provides an alternative view of what this CAPE related CCT is doing. It is putting a restriction on the convection scheme so that it can only represent the weak (hopefully shallow) clouds and therefore behave more like a shallow scheme. The hope is then, that assumptions used in the convection scheme that became invalid in a high-resolution model become reasonable because the scheme is once again only dealing with sub-grid-scale clouds. This is speculative and may well be in error because of the presence of larger clouds and dynamical interactions, but the reasoning has some merit. Another benefit of looking at the function in this way is that the validity of the assumption about cloud size and CAPE becomes unimportant since the mass flux does not vary with CAPE for most values of CAPE. In effect the CAPE dependence has cancelled out.

The parameters t and c are used to tune the function. The parameter t is in effect a CCT for very small CAPE. If it is less than the operational (12-km mesoscale model) CCT of 1800 seconds (30 minutes) then the allowed mass flux is greater than the operational for small values of CAPE – as in Figure 7. The limit on the cloud base mass flux M is determined by the value of t/c. The larger this value the more M is restricted. A value of 1200s was chosen for t, which seems sensible as we want a short CCT for small clouds. Several values of c have been tried and some results from the two case studies will now be shown.

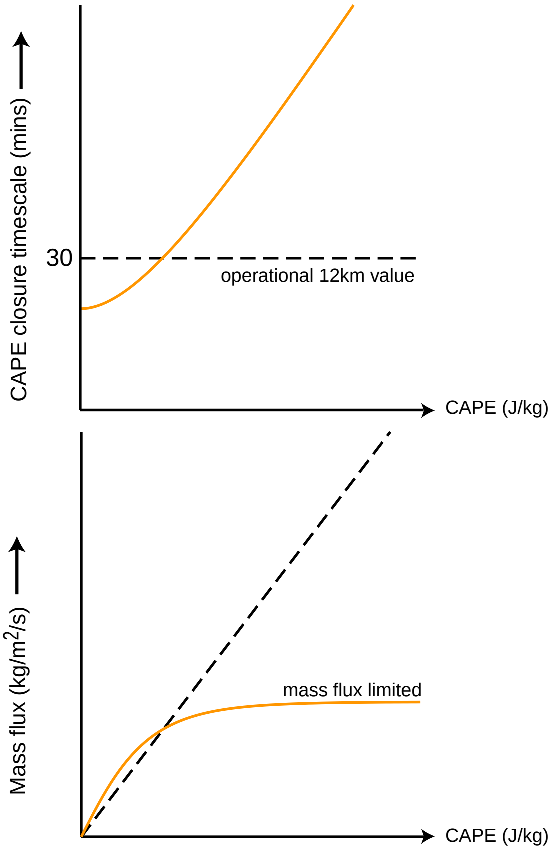

**Figure 7. Graph of the function used to modify the CAPE closure timescale in the convection scheme and the resulting behaviour of the cloud base mass flux with CAPE.**

## 5 RE-VISITING THE CASE STUDIES

### 5.1 Squall line 2nd July 1999

New 4-km gridlength simulations were run with everything unchanged except for the inclusion of the CAPE dependent CCT in the convection scheme. The runs used a constant value of t=1200 and several different values for c. Output from two of these runs is shown in Figure 8. A transition can be seen between the run with a value of c=0.1, which is close to the no-convection-scheme run shown in Figure 2(a) and the run with a value of c=250.0, which is behaving more like the constant CCT=1800s (standard) run shown in Figure 2(b). It is encouraging to see that it is possible to produce a solution with the CAPE dependent CCT that looks like that produced by the run with no convection scheme for this event (c=0.1, and also for c=10.0 not shown). That was the initial aim and it seems to have succeeded.

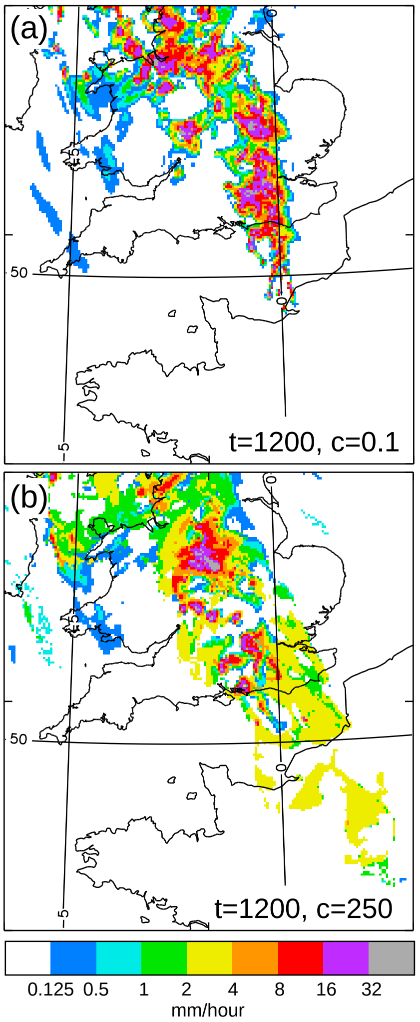

**Figure 8. Rainfall rates at 23 UTC 02/07/99 from 4-km gridlength forecasts starting at 15 UTC with different values of the c parameter in the CAPE closure timescale function.**

### 5.2 Re-visiting case-study 2, 03/05/02, with the CAPE dependent CCT

The best results in the previous case came from using a value of c=10.0 or 0.1 with t=1200s in the CAPE dependent CCT, so it made sense to use these values for this case and add c=50.0 and c=2.5. The new model runs are unchanged from before except for the inclusion of the CAPE dependent CCT or if stated otherwise. As before, three aspects of the case are examined – the peak rainfall rates, the initiation time and the persistence of showers into the evening.

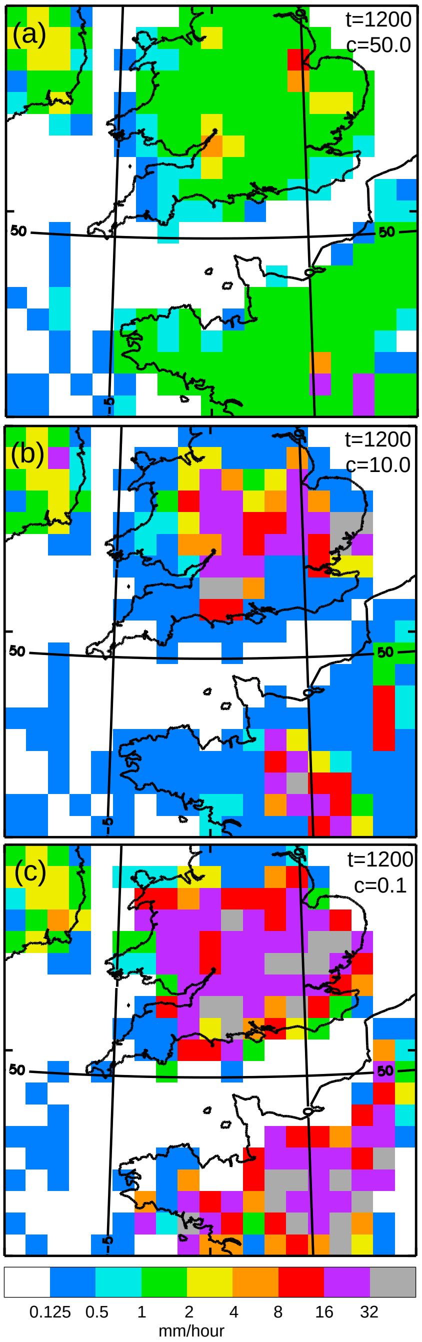

**Figure 9. (a) Peak rainfall rates over 40x40 km squares within the period 10 to 18 UTC 03/05/02 from 4-km gridlength runs with different values for the c parameter in the CAPE closure timescale function (a) c=50.0, (b) c=10.0, (d) c=0.1.**

Figures 9 and 10 show the sensitivity of the peak rainfall rates to the value of c. Unlike in the previous case, the peak rates from the run without the convection scheme were unrealistically high because of the development of single gridpoint storms. When a value of c=50.0 was used the peak rain rates became too small because the convection scheme largely inhibited the model dynamics from triggering. With a value of c=0.1 the peak rates were too high – though considerably less than the no-scheme run, but with a value of c=10.0 became much closer to the radar. The addition of extra diffusion into the model in the c=0.1 run caused the peak rates to become similar to the values in the c=10.0 run.

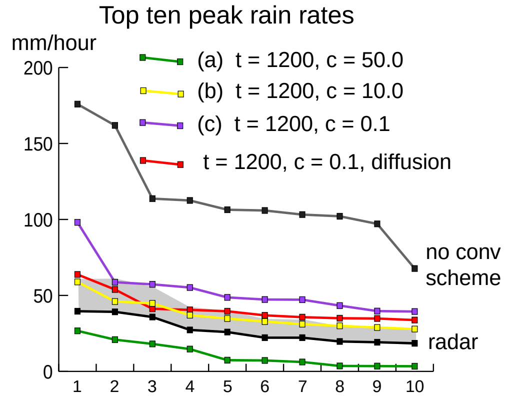

**Figure 10. Graph of the top ten peak rainfall rates from Figure 9.**

We know that the showers in the run without the convection scheme started too late. Figure 11 shows what happens with the inclusion of the CAPE dependent CCT. Whatever the value of c, the convection scheme triggered at approximately the correct time. However, the more the mass flux was restricted (smaller c) the less significant the rain from the convection scheme became. With a value of c=0.1 the convection scheme hardly produced any rain. In contrast, the initiation time of dynamically resolved showers was dependent on the value of c. The larger c, the later the dynamics triggered. The difficulty here is that it is impossible to have it both ways, either the convection scheme produces reasonable rain rates and the dynamics triggers too late (or not at all), or the dynamics triggers earlier (though still too late) and the convection scheme is too weak or not used. Unfortunately, the best result in terms of triggering convection at the right time, comes from the run with the single CCT=1800s in the convection scheme (standard setting), but we know that this run has other problems we wish to avoid.

The use of the CAPE dependent CCT has had a significant and positive impact on the behaviour of the showers from early afternoon onwards. Instead of either dying out entirely with the constant CCT=1800s option or becoming too active and then rapidly decaying with no convection scheme, the showers persisted into the evening and without producing too much rain. Figure 11 shows that the

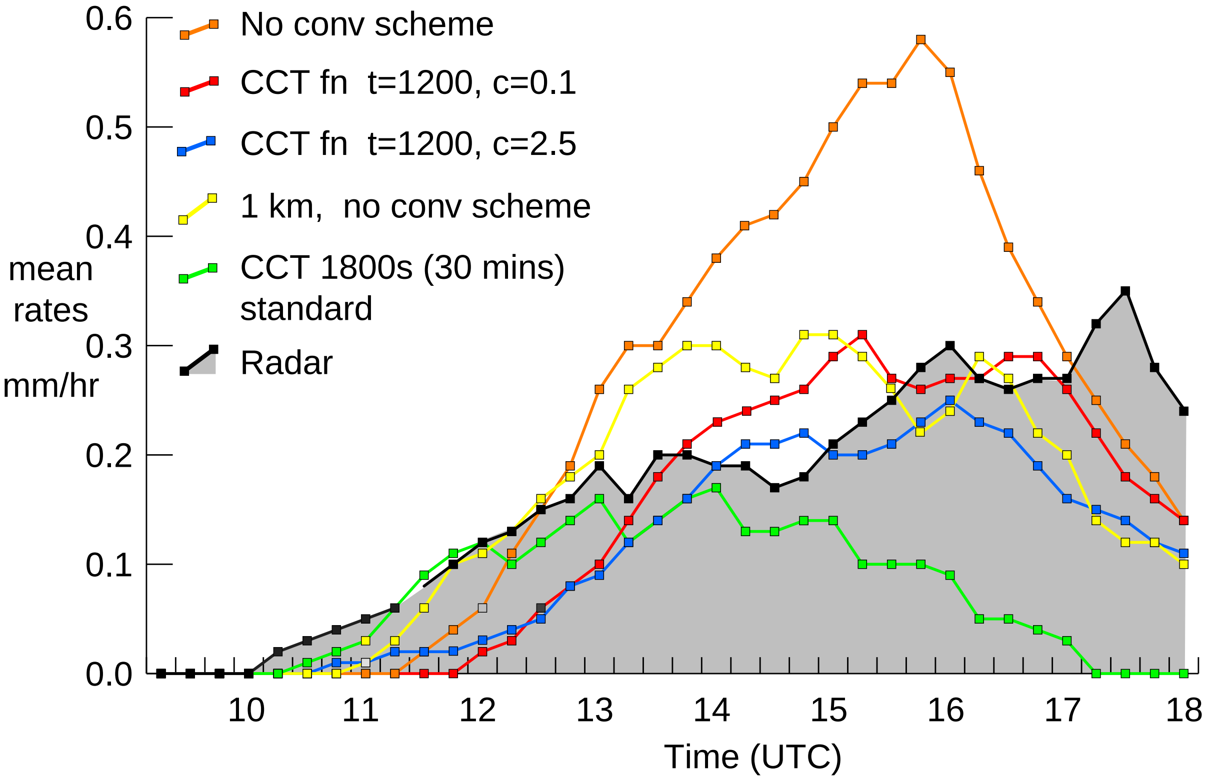

**Figure 11. Graph of the mean rainfall rates within the area shown in Figure 3 against time for 4-km gridlength forecasts with different values of the c parameter in the CAPE closure timescale function, 4-km forecast with no convection scheme (orange), 1-km forecast with no convection scheme (yellow) and radar (black line enclosing shading). CCT 1800s refers to the run that used the standard CAPE closure timescale of 1800 seconds (30 minutes) and not a CAPE-dependent closure.**

average rainfall rate after 14 UTC became consistently higher as the value of c became smaller. Between 1400 and 1630 UTC both the c=0.1 and c=2.5 runs were close to the radar. The run with c=10.0 had too little rain throughout, but was still better than the constant CCT=1800s run over this period.

The CAPE dependent CCT forecasts were more realistic because the restriction on the convection scheme allowed the model dynamics to generate showers that could then organise, yet removed enough instability through the convection scheme to prevent the resolved activity from becoming too large. After 13 UTC, the c=0.1 and c=2.5 runs were just as good in terms of mean rainfall rate as a 1-km simulation of this event with no convection scheme (Figure 11). This trend in behaviour with different values of c is encouraging because it fits with what was intuitively expected.

## 6 CONCLUSIONS

### 6.1 Results

Examples from case studies have shown that there are problems with the use of an equilibrium convection scheme in a model with a gridlength of 4 km. In regions of high CAPE (first case study), an interaction between the convection scheme and the model dynamics can generate spurious rainbands. In regions of lower CAPE (second case study 2) the convection scheme can prevent the model dynamics from developing showers and therefore stop any convective organisation from occurring. Switching off the convection scheme is not the solution; it only creates different problems. The scales of showers are determined by the resolution of the grid rather than by the natural scales of the event. This can lead to unrealistically high rainfall rates and the formation of single grid-point storms as well as causing a delay in the initial triggering (case study 2).

The change to using a CAPE dependent CAPE closure timescale produced significantly better results provided that suitable values for the parameters t and c were chosen. In case study 1, spurious rainbands did not develop. In case study 2, the unrealistically high rain rates produced by the run with no convection scheme were greatly reduced and resolved convection developed that was allowed to organise and persist into the evening. The runs with t=1200,c=2.5 and t=1200,c=0.1 even produced mean rainfall rates after 14 UTC that were comparable to the 1-km simulation (Roberts 2003).

The optimal choices for the c and t parameters cannot be exactly found, but because predictable trends were apparent, a range of sensible values is known. Given a value of t=1200 in a 4-km gridlength model, c should be less than 20.0 or the convection scheme is too active in large-CAPE situations and greater than 0.005 or the convection scheme has too little effect in low-CAPE situations. Although tuning parameters are not usually a good thing to have in a numerical model, a benefit of having them here is that it is possible to make choices that are appropriate for the purpose of the model that is being run. If the aim is to have a model that is meant to be used primarily to forecast severe convective events at the expense of not representing smaller showers properly, then a low value of c should be chosen (t/c is large). If the aim is to have a model that ‘plays safe’ and represents most of the convection with the convection scheme at the expense of restricting the dynamics from generating some organised storms then a high value of c should be chosen. In practice, a compromise is sensible.

The results are also applicable to model gridlengths other than 4 km. Tests with a gridlength of 2 km have produced very similar results, though different values of t and c may be appropriate.

### 6.2 Issues

A problem that still remains however, is the delay in triggering resolved showers in situations with weak dynamic forcing. In case study 2, the 4-km runs with the CAPE dependent CAPE closure timescale triggered the resolved convection too late and although the convection scheme produced rain at the correct time, there was not enough. The only way to produce more rain from the convection scheme was to place less of a limit on the mass flux (make t/c smaller), but then the resolved convection was delayed even more. It could be argued that it is not a problem to delay the triggering of resolved showers if the convection scheme is doing a good job of representing the convection. This argument however is only valid if the convection scheme is not inhibiting the development of showers that should be resolved - but we know that it does. Ideally, the CAPE dependent CAPE closure timescale should allow the sub-grid-scale and near-grid-scale clouds to be represented by the convection scheme and leave the dynamics to simulate any larger showers. In practice, the dynamics will only trigger if the convection scheme is restricted to representing only the very small sub-grid clouds. If the convection scheme is allowed to represent the near-grid-scale clouds the dynamics will often not trigger. A way to encourage the model dynamics to initiate showers earlier, without restricting the convection scheme too much, might be to add random, low-level temperature and humidity perturbations wherever the convection scheme is active. The perturbations would represent the effect on the grid of sub-grid-scale variability associated with the unresolved convection and provide enough convective instability at a few points for the dynamics to trigger. Done (2003) has already shown that the addition of random perturbations to a 12-km gridlength model can have an impact on local triggering. The delay in triggering is less pronounced with a gridlength of 1 km than it is with 4 km, but is still a cause for concern (Roberts 2003). A 1-km model is intended to be more accurate over shorter time periods, so even a short delay in convective initiation could be significant.

There is another issue to do with the interpretation of precipitation forecasts. Precipitation output from a run using the CAPE-dependent function may not look very much like the usual precipitation output from numerical models. Convective precipitation will consist of uniform regions of very light precipitation from the convection scheme and more intense resolved showers. The light-precipitation regions show where the convection scheme has triggered and hence where there is a risk of convection, whereas the resolved showers reveal the nature of any convection once it has developed (i.e the organisation and rain rates). This can be an advantage if suitable precipitation diagnostics are produced because a deterministic forecast of the resolved convection will also have this element of uncertainty attached. It is not wise in any case to present raw output from high-resolution forecasts because of the danger of believing fine-scale detail that is beyond the accuracy of the model.

## 7 REFERENCES

Cullen, M.J.P., Davies, T., Mawson, M.H., James, J.A., Coulter, S.C. and Malcolm A., 1997, An overview of Numerical Methods for the Next Generation UK NWP and Climate Model. Numerical Methods in Atmospheric and Ocean Modelling. The Andre J.Robert memorial volume. Edited by Charles A Lin, Rene Laprise and Harold Ritchie 425-444

Cohen, B. G., 2001: Fluctuations in an ensemble of cumulus clouds. *PhD thesis, University of Reading, UK*

Done, J., 2003: Predictability and representation of convection in a mesoscale model. *PhD thesis, University of Reading, UK*

Gregory, D. and Rowntree, P. R., 1990: A mass flux scheme with representation of cloud ensemble characteristics and stability-dependent closure*. Mon. Wea. Rev.,* **118***,* 1483-1506

Roberts, N. M., 2001: Results from simulations of an organised convective event using the New Dynamics at 12, 4 and 2 km resolution. *FR Technical Report* **344***.* Unpublished paper available from the National Meteorological Library, Bracknell, Berks, RG12 2SZ, UK

Roberts, N. M., 2003: Results from high-resolution simulations of convective events. FR Technical Report **402***.* Unpublished paper available from the National Meteorological Library, Bracknell, Berks, RG12 2SZ, UK

Swann, H., 2001: Evaluation of the mass-flux approach to parametrizing deep convection*., Q.J.R.Meteorol.Soc,* **127***,* 1239-1260*.*

Weisman, M. L., Skamarock, W. C. and Klemp, J. B., 1997: The resolution dependence of explicitly modelled convective systems*. Mo..Wea. Rev.,,* **125***,* 527-548*.*