The University of Reading

Met Office

# Numerical Weather Prediction

## The impact of a change to the use of the convection scheme to high resolution simulations of convective events
## (Stage 2 report from the storm scale numerical modelling project)

Forecasting Research Technical Report No. 407
Joint Centre for Mesoscale Meteorology Report No. 142

Nigel Roberts

email: nwp_publications@metoffice.com

PROJECT DOCUMENTATION

PRODUCT DESCRIPTIONS

***The impact of a change to the use of the convection scheme in high-resolution simulations of convective events. Stage 2 report from the Storm-Scale numerical modelling project***

Release: Draft
Date: 9th April 2003

PRINCE2

Author: Nigel Roberts

Owner: Nigel Roberts

Client:

Document Number: Version 1.2

# 1 Product Descriptions History

## 1.1 Document Location

This document is only valid on the day it was printed.
The source of the document can be found at:
\\Meto04\nwp_workgrps\Forecasting Research \JCMM\MesoscaleModelling\DEFRA_Storm_Scale_NWP

## 1.2 Revision History

Date of this revision:
Date of Next revision:

| Revision date | Previous revision date | Summary of Changes | Changes marked |
|---|---|---|---|
| | 13/03/03 | First issue, version 1.0 | |
| | 03/04/03 | First complete draft, version 1.1 | |
| 08/04/03 | | Revision to include discussion of HRTM parameter testing, more discussion of diffusion and other small changes inlight of comments by Brian Golding. Version 1.2 | |

## 1.3 Approvals

This document requires the following approvals. Signed approval forms are filed in the Quality section of the PCB.

| Name | Signature | Title | Date of Issue | Version |
|---|---|---|---|---|
| Brian Golding | *Brian Golding* | Hd(FR) | 09/04/03 | 1.2 |

## 1.4 Distribution

This document has been distributed to

| Name | Title | Date of Issue | Version |
|---|---|---|---|
| | | | |

# Contents

# 2 Overview

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 1 – 10 km in which convective storms are often partly resolved. When gridlengths get down to 1 km or shorter much of the rain producing convection should be resolved and a convection scheme may not be needed, but this cannot just be assumed.
Simulations of convective events using the Unified Model version 5.3 have been run with a gridlength of 4 km to show that neither including or removing the convection scheme at that resolution will produce satisfactory results. The use of the convection scheme can sometimes lead to the development of spurious rainbands and can also restrict the model dynamics from representing showers it should be able to resolve. Switching off the convection scheme is somewhat better, but can result in the formation of unrealistically intense showers and even single gridpoint (or 'grid-row') storms.
In an attempt address this problem, a change to the use of the convection scheme has been examined. The activity of the convection scheme has been restricted in order to try and allow the model dynamics represent the showers it should be able to resolve and leave the convection scheme to deal with the rest. This approach is shown to produce better results.

# 3 Introduction

This is the second interim report to be delivered from the Storm Scale Numerical Modelling project. The project was set-up to investigate the ability of a storm scale configuration of the Met Office Numerical Weather Prediction (NWP) model to predict flood-producing rainfall events. In this context a storm-scale model has a horizontal grid spacing of about 1 km and receives information at the boundaries from a model with a 4-km grid spacing. It is recommended that the previous report (Roberts 2003) is read before this one in order to gain familiarity with the case studies.

This report will concentrate on the difficulty of simulating/representing convection in high-resolution models. The focus will be on a gridlength of 4-km, which falls into the middle of the range (1 – 10 km) in which the question of what to do about the use of a convection scheme becomes most debatable and the problem can be particularly noticeable. Results from case-study simulations have fallen into line with expectations in showing that changes in the use of the convection scheme can have a very large impact on the outcome of forecasts. It is for that reason, that the use of the convection scheme (or not) has been singled out for special attention and a change to the way the convection scheme is used has been devised.

Although the focus is on a 4-km gridlength here, the problem of simulating convection properly does not go away at 1 km. Despite such a short gridlength there will still be convective clouds that are not resolved by the model dynamics. Even the larger convective systems are comprised of rapidly evolving smaller-scale clouds - with local mixing between these clouds and the environment, which can not be represented properly on the grid. The lessons learnt at 4 km still apply at 1 km. In particular, any change to the use of the convection scheme that has a positive impact at 4 km has the potential to do the same at 1 km, though tuned differently. The other implication for a 1-km model is that the forecast is sensitive to the information that passes through the boundaries from the 4-km gridlength domain, so the way convection is represented at 4 km directly feeds into the 1-km forecast.

It is recognised that, in addition to the use of the convection scheme, many other aspects of parameter selection and tuning also require investigation. They include: the amount of diffusion, the critical value of relative humidity for cloud (RH-crit), the sensitivity to surface parameters such as soil moisture and vegetation, microphysical parameters and boundary layer parameters. Changes to the vertical resolution, timestep, boundary updating, diffusion, convection scheme and critical relative humidity (RH-crit) have been systematically examined in the High-Resolution Trial Model (HRTM) project (Lean 2003). The aim of the early stages of that project was to arrive at a configuration for a high-resolution modelling system, which could be used as a standard from which further developments would take place. Much of that configuration has been adopted as the basis for the convection experiments described in this report and results have also fed back into the HRTM project. At this stage, it was not thought appropriate to repeat tests performed in the HRTM project until a suitable way ahead in dealing with the fundamental problem of modelling convective clouds had been found. However, some aspects, particularly diffusion will have to be re-visited and therefore some mention of diffusion has been made in this report. A review of the HRTM experiments and the configuration that was adopted here is now presented.

# 4 A review of parameter selection following testing in the HRTM project

Six parameters were tested (see Lean 2003 for details):

1. Vertical resolution. 76 or 38 vertical levels were tested. The standard HRTM configuration has 76 levels at both 4 and 1-km gridlength because it was thought best to have a large number of vertical levels in conjunction with short horizontal gridlengths. Convective case studies showed little sensitivity to vertical resolution. The convection scheme experiments described in this report have been performed on both 38 and 76 levels for the two main case studies. Although some of the detail was different, the trends and overall results were not sensitive to vertical resolution. Only 38 levels for case 1 and 76 levels for case 2 have been presented. The first report (Roberts 2003) did show sensitivity to vertical resolution in the triggering of convection in a 1.2-km gridlength simulation.
2. Timestep. The standard HRTM configuration has a 1-minute timestep for the 4-km gridlength simulations and a 30 seconds for the 1-km runs. There was no need to change these values. A shorter timestep should be more accurate, but also more costly in terms of run time and these values produce sensible results.
3. Diffusion. Adding diffusion is a way of damping the small-scale waves that can make a model produce spurious results or even go unstable and fail. Diffusion can be applied to moisture, temperature and winds and the amount can be controlled by setting the diffusion coefficient at each model level. In addition, adding diffusion can be a way of applying mixing that could be used to represent the turbulence associated with sub-grid-scale clouds. In theory, diffusion is not required, but at high-resolution is thought to be beneficial. The HRTM testing found that using a value of diffusion that was half the theoretically allowed limit was the best compromise. The addition of more diffusion produced rainfall that was too smoothed out and weak and delayed the triggering of convection. Much less or no diffusion resulted in more of the gridscale structure we probably wish to avoid. Some cases showed little sensitivity to the amount of diffusion. The experiments presented in this report used the same diffusion as the standard HRTM configuration in case study 1 and almost no diffusion in case study 2 except

where stated. The deep convection in case study 1 required some diffusion to control gridpoint behaviour. The addition of diffusion in case study 2 also had some impact that is shown. Although diffusion is clearly important, it does not seem to make anything like as much difference as the more significant impact of modifying the use of the convection scheme.

4. Frequency of updating the lateral boundaries. The standard HRTM configuration updates the 4-km gridlength runs every 30 minutes from 12-km runs and the 1-km every 15 minutes from 4-km runs. This has also been adopted here. More frequent updating is more accurate and a five minute update for 1 km will become the new HRTM standard.
5. Convection scheme. The standard HRTM configuration uses the convection scheme with the CAPE dependent CAPE closure that is the subject of this report (values of t=1200, c=10.0, see later sections) for 4-km gridlength runs, and no convection scheme for 1-km gridlength. This choice was based on early results from tests of the CAPE-dependent CAPE closure and sensitivity studies within the HRTM project. In the light of this report, new values of t and c will be required for the 4-km configuration, in combination with further testing of diffusion and other approaches outlined in this report and the HRTM report (Lean 2003). At 1-km it is still not clear that excluding the convection scheme is the best that can be achieved and further testing of the CAPE-dependent closure with very small values of c should be tried – again with further diffusion testing and other approaches outlined in this and the HRTM reports.
6. RH-crit. This is the minimum relative humidity at which cloud is diagnosed to start forming in the model and is there to take account of variability within a grid square. Operationally (12-km gridlength) this is set to 80% for most model levels. In the HRTM testing phase the outcome of using 95% in both the 4 and 1-km runs was examined. This was considered reasonable because there should be less cloud variability within smaller squares. It was anticipated that a change to this parameter might have a large and worrying impact on simulations of convective events, but it turns out that the runs were largely insensitive to changing this parameter, except in a random way to the finer structure. Since there was little trend in the impact, the standard HRTM configuration kept the operational values and the same numbers have been used here. It is possible that changing RH-crit could have a bigger effect than has yet been seen, so it might require futher investigation in future.

# 5 The problem of simulating convection in high-resolution models

At the resolution of the 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. The convection scheme is a section of code that is designed to calculate changes to the gridpoint values of temperature and humidity which are due to convection that the model is unable to resolve. Currently, the Unified Model (Cullen et al) uses an equilibrium mass flux scheme (Gregory and Rowntree 1990, Gregory et al 1999) 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. Each plume represents the effect of clouds of different sizes that reach different heights within a grid column (Figure 1). If a convection scheme was not used the model would develop too much convective instability at single gridpoints and might become unstable and fail.

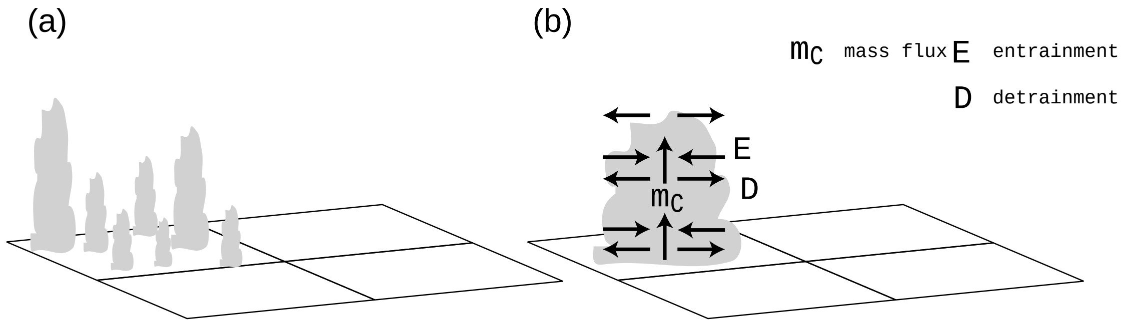


Figure 1. A schematic showing in (a) a number of convective clouds within a grid column and (b) a convection scheme representation of these clouds with a single modelled plume that entrains air in from the environment and detrains air out from the plume at different levels to account for the different sized clouds in the column.

At the resolution of the 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. There may be occasions when large thunderstorms are resolved by the model dynamics (particularly at ~1km) and the convection scheme is not essential. There may be other situations when small showers predominate that are not resolved by the model dynamics and a convection scheme of some sort is definitely required. It may be that with a gridlength of 1 km or shorter there is never a need to use a convection scheme for practical weather forecasts.

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?

## 5.1 Is it reasonable to use a convection scheme at high resolution

### 5.1.1 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 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.

### 5.1.2 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.

Yet another factor is that of cost. As resolution is increased, the model timestep has to get shorter, which means that the convection scheme will be run more frequently in simulations when it is probably required less.

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. One the other hand, it is likely that a convection scheme of some sort is still required, at least for gridlengths > 1 km, to stop the model producing unrealistic storms at single gridpoints and even failing. The encouraging thing is that the mesoscale model (gridlength of ~12 km) has already been running for several years with the convection scheme switched on and produced acceptable results. The first step towards an attempt to answer the question of whether it is possible to get away with using the current convection scheme at higher resolutions (gridlength of 1 – 5 km), or even run with no convection scheme, is to try these alternatives on some case studies.

# 6 Examples of running both with and without the convection scheme

Pictures from model simulations with a horizontal gridlength of 4 km and the convection scheme either switched on or off will be shown for two case studies. Both case studies have already been presented in the Storm-scale numerical modelling stage 1 report (Roberts 2003), so will not be introduced again in much detail here. The reason for choosing to look at 4-km gridlength simulations is that at this resolution we expect to encounter situations when convection is only partly resolved by the model dynamics, and it is also the resolution that would most likely be used to drive a 1-km gridlength model.

## 6.1 Case 1 02/07/99

Figure 2 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 (stage 1 report) concluded that the run with no convection scheme was much better. Weisman et al also showed success in 4-km simulations of squall lines with no convection scheme.
The run which used the convection scheme had a serious problem. The problem was the formation of bands or arcs of precipitation from the convection scheme (labelled C) that propagated northward and southward through the domain. The bands were self-sustaining because of an interaction between the convection scheme and the model dynamics. As well as producing rain in the wrong place (by hundreds of kilometres), they led to the removal of much of the convective instability that was required for the dynamics to trigger storms over southern and central England.

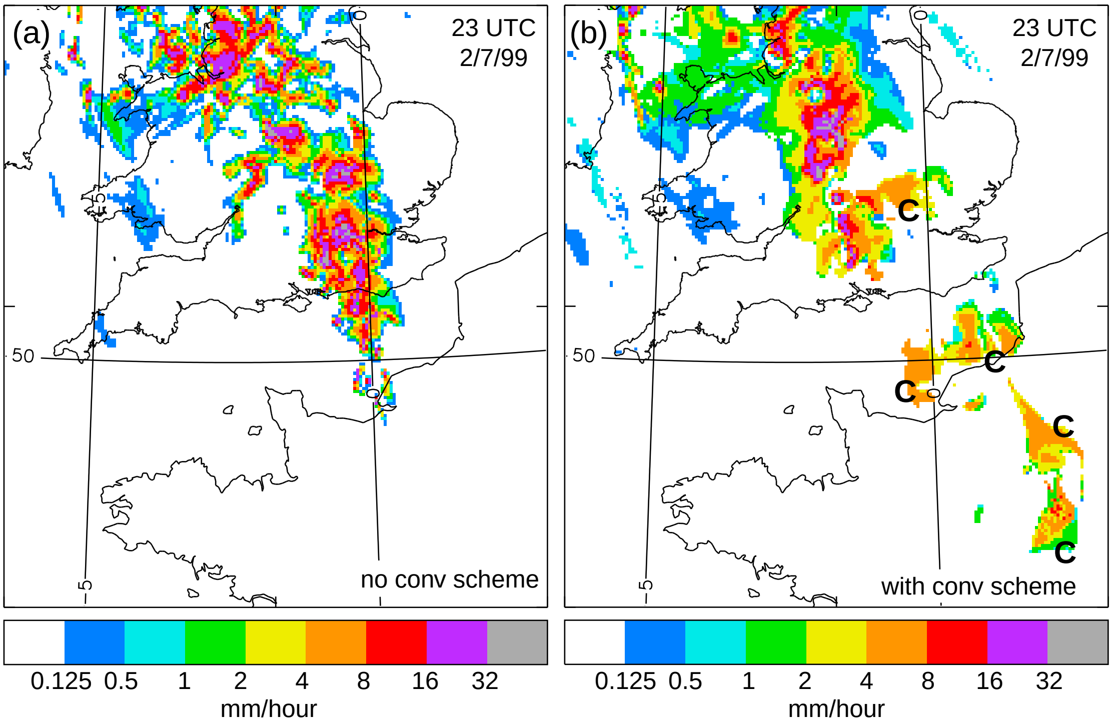


Figure 2. Precipitation rates at 13 UTC 02/07/99 from 8-hour, 4-km gridlength model forecasts (following a 6-hour 12-km forecast). (a) Forecast with no convection scheme. (b) Forecast with the convection scheme. The 'C's show regions of precipitation from the convection scheme that are discussed in the text.

The probable mechanism of the dynamics/convection scheme interaction is depicted in Figure 3(c) and (d). The convection scheme was firstly triggered at gridpoints where there was sufficient moisture and local ascent at around 750hPa 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. A self-maintaining system then developed as long as the surrounding environment had sufficient unrealised potential instability. The cross section (Figure 3(b)) shows that a band of precipitation from the convection scheme (labelled A in Figures 3(a) and (b)) was collocated with a region of ascent shown by the dark shading and a low-level cold pool shown by the bending up of the potential temperature contours. In an experiment (not shown) in which the convection scheme was switched off at 10 UTC - the ascent region and cold pool disappeared within 30 minutes.

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 evidence suggests that the most obvious way to remove the bands and still run with the convection scheme is to reduce the intensity of the convection scheme. Maybe the best solution is to switch off the convection scheme altogether at 4km.

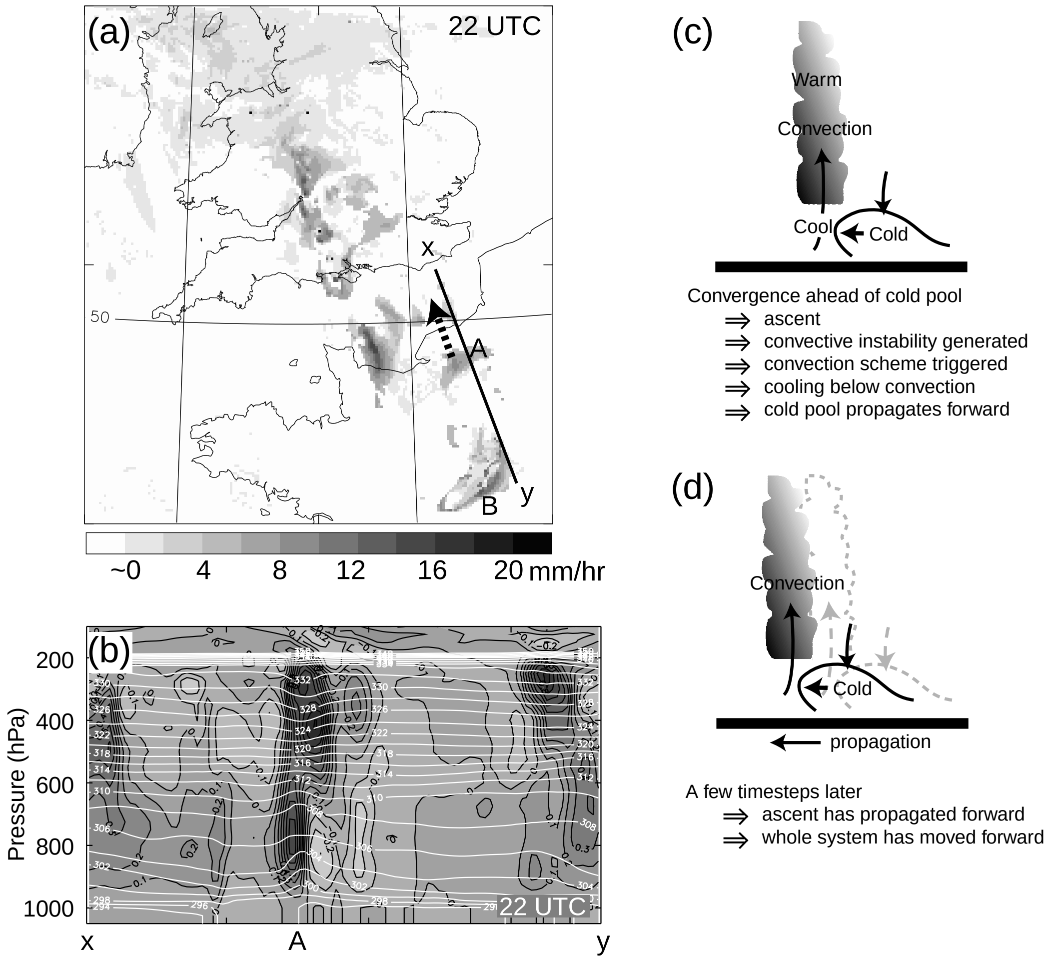


Figure 3. (a) Precipitation rates at 22 UTC 02/07/99. A 7-hour, 4 km gridlength forecast (following a 6-hour, 12 km forecast). A and B mark bands/arcs of precipitation from the convection scheme. The dashed arrow shows the direction of propagation of band A. (b) A cross-section along the line xy in (a). Shading shows vertical velocity (dark is ascent), white contours are potential temperature. (c) and (d) Schematic diagrams of the probable mechanism associated with the convective bands A and B – discussed in the text.

### 6.2 Case 2 03/05/02

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. A convection scheme should be essential in this situation because many of the showers were too small to be properly resolved by the dynamics at 4km.

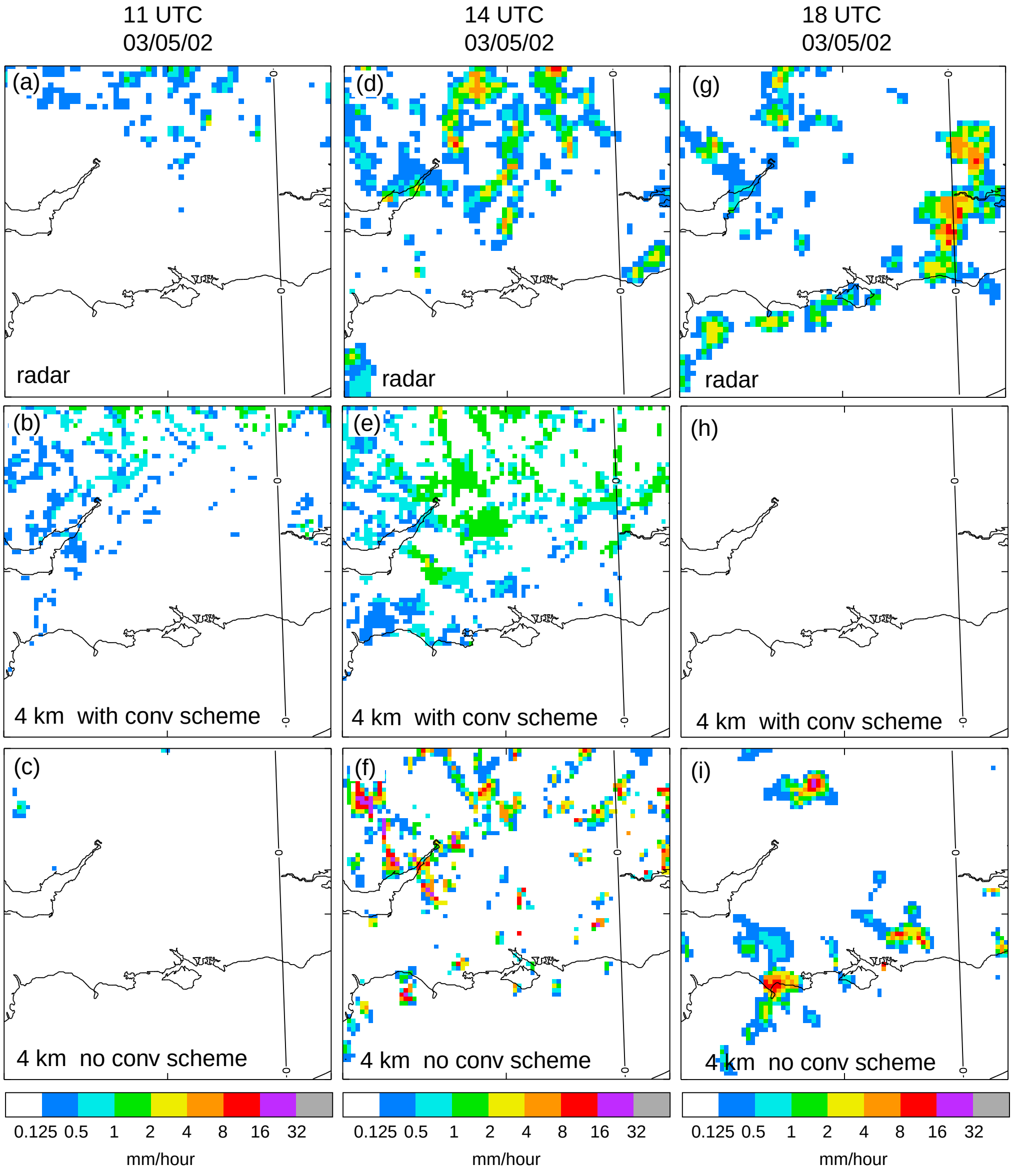


Figure 4. Precipitation rates at 11, 14 and 18 UTC 03/05/02 from radar and from 4km-gridlength forecasts starting at 01 UTC with and without the convection scheme. The area shown is a sub-area of the 4-km model domain.

Three snapshots of the showery day are shown in Figure 4 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 produced a reasonable forecast at this stage with scattered precipitation of a similar intensity to the radar, though more widespread. The run without the scheme was poor at 11 UTC because the model dynamics had hardly triggered any showers. Later, at 14 UTC, both simulations had extensive precipitation like the radar picture, but the rain was too light in the run with the scheme and too heavy in the run without. All of the precipitation in the run with the convection scheme came from the convection scheme. By 18 UTC the showers had become larger and more organised. The run with the scheme was very poor because all of the precipitation had died out, whereas the other run was much better because there were still showers and evidence of organisation.

There are three aspects of this case that need examination – the rainfall intensity, the initial triggering and the persistence of showers into the evening.

### 6.2.1 Rainfall intensity

The difference between the peak rainfall rates is revealed by Figure 5.

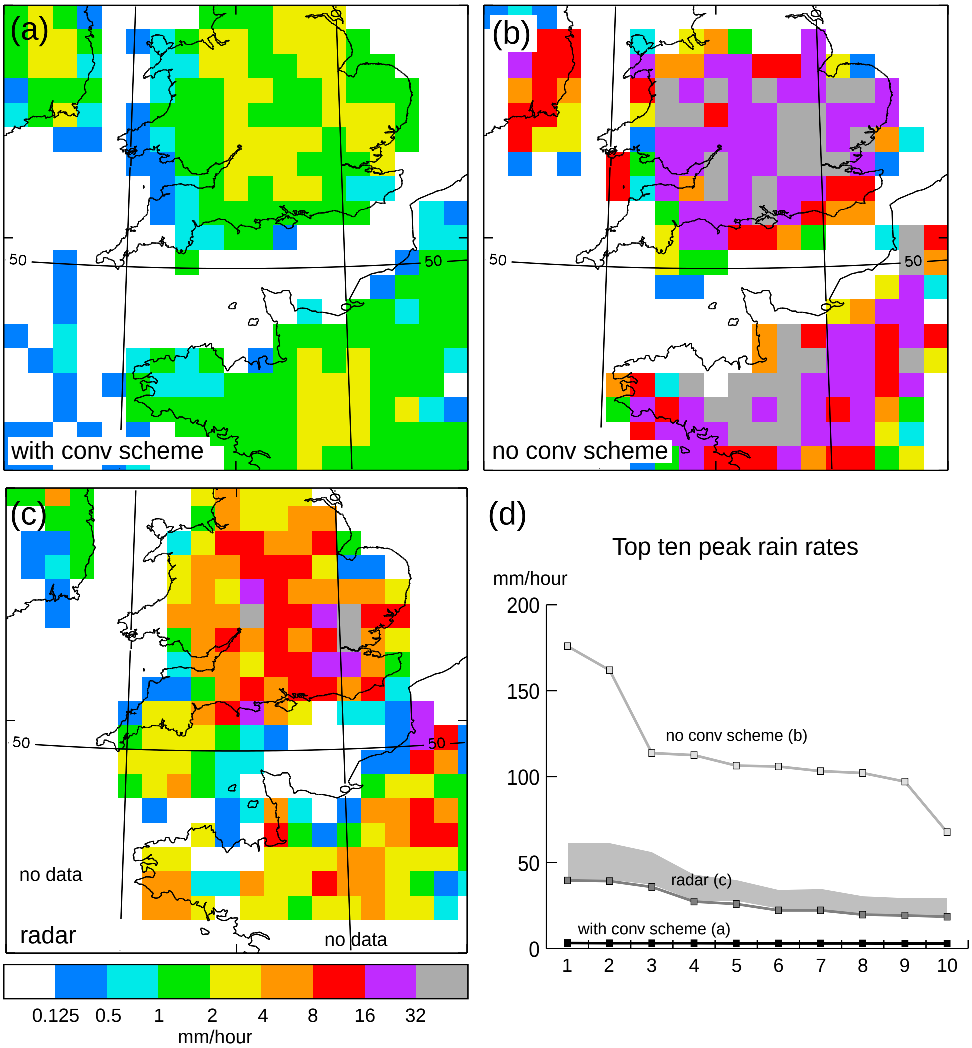


Figure 5. (a)-(c) Maximum precipitation rates to occur over 40x40km squares in the period 10 to 18 UTC 03/05/02, sampling every 15 minutes from (a) 9 to 17 hour, 4-km gridlength forecast with the convection scheme included, (b) 9 to 17 hour, 4-km gridlength forecast with no convection scheme, (c) Nimrod network radar (5km grid). (d) A graph of the top ten peak rainfall rates taken from pictures (a), (b) and (c). The shading shows the increase in radar rain rate that could occur when going from a 5km to 4 km square, if the rain is concentrated at a single point within the square.

The run with the convection scheme has produced rates that are 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 – a build up of too much convective instability had been allowed, which was eventually released as small (1 or 2 gridpoint) intense storms. 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 option 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.

### 6.2.2 Initial triggering

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 not a good feature, 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.

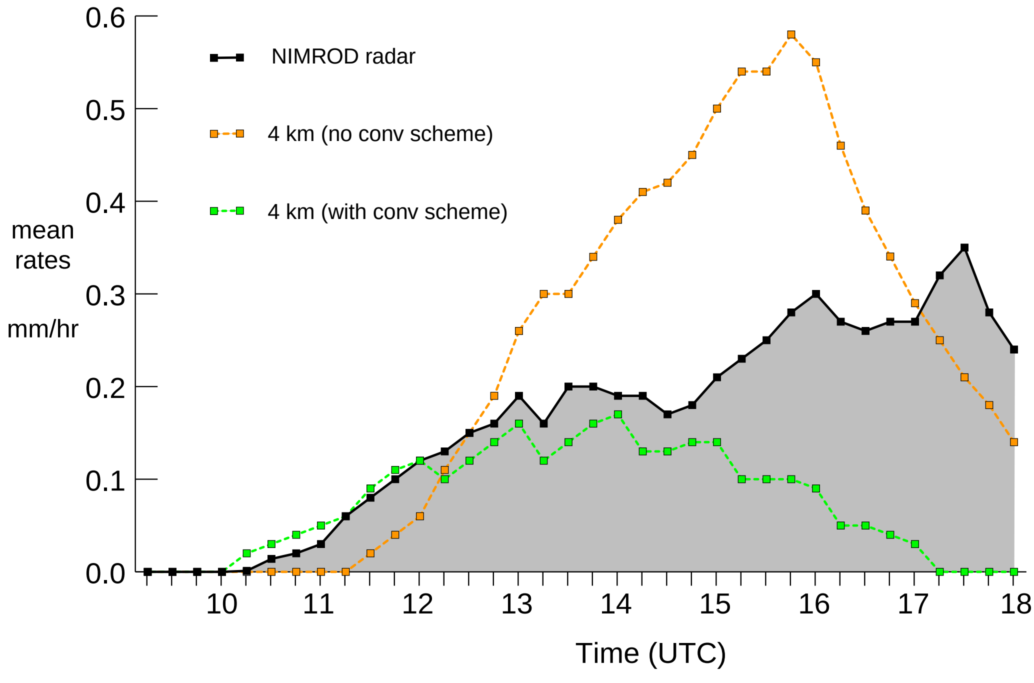


Figure 6. Graph of the mean rainfall rates within the area shown in Figure 4 against time for the two forecasts discussed in the text (coloured lines) and the radar (black line enclosing grey shading).

### 6.2.3 Persistence into the evening

Figures 4 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. The reason for this can be explained when in is known that all the precipitation in that run came from the convection scheme. 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.

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 7 shows that the run without the scheme had produced an average rainfall rate of twice that observed.

General

In this case, it is difficult to say whether it is best to run with or without the convection scheme. Both forecasts have positive and negative aspects. The negative aspects are a consequence of either the limitations of the convection scheme at this resolution or the restrictions associated with a gridlength that is incapable of resolving small showers properly.

## 6.3 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-square cells developed. Even in the severe-storm case the rainfall intensity was too high and the first cells that formed gave extreme rainfall rates from single grid squares.

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. Spurious, self-sustaining, bands of precipitation were generated by the convection scheme. These bands inhibited the triggering of convection by the model dynamics and produced rainfall in completely the wrong place. Similar bands have also been observed in other 2, 4 and 12-km gridlength simulations.

In the scattered convection case, the use of the convection scheme produced a reasonable forecast at first. The timing of the initial shower development was close to that observed and during the period when convection developed as a response to solar heating the convection scheme performed well. 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. In addition, the convection scheme produced rainfall rates that were too small.

They were small because the scheme was doing what it was supposed to do by generating an 'ensemble average' precipitation rate over each grid square, but because the convection scheme does not know about the size of the grid squares, it was generating rainfall rates that were not appropriate for a gridlength of 4 km. The convection scheme could be tuned to produce higher rainfall rates, but is then more likely to generate the spurious convective bands that have already been identified as a problem.

What are the alternatives? The first option is to make do with either the current convection scheme or no convection scheme and put up with the limitations. A different convection scheme could be tried, but any other currently available scheme would also suffer from the same inability to properly simulate time-varying and spatially-organised convection. Even so, it might be worthwhile to examine the behaviour of different schemes. The new shallow convection scheme that is being developed within the Met Office is a viable alternative that should be tested. The shallow scheme would only represent smaller convective clouds up to 2-3 km deep and the model dynamics would then deal with the rest. A moist turbulence scheme rather than a convection scheme may be worth considering for a gridlength of ~1 km but probably not a gridlength of ~4 km. Another approach is to modify the use of the current convection scheme in some way so that it is better suited to the requirements of a high-resolution modelling system. This is the approach that has been tried so far and will now be described.

# 7 A change to the use of the convection scheme

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.

$$M \;\alpha\; CAPE / \tau \qquad (1)$$

That is why the convection scheme is more active in more convectively unstable situations (case study 1) 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

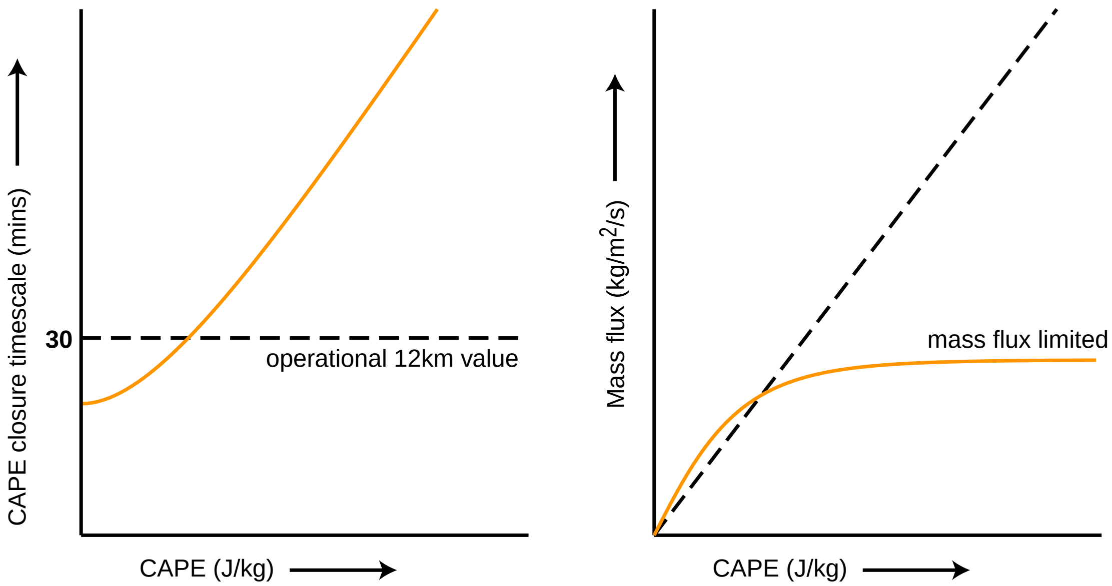


Figure 7. Graph of the function used to modify the CAPE closure timescale in the convection scheme and the subsequent behaviour of the cloud base mass flux with CAPE.

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 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. For the case studies reported on here 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 results from these tests will be shown.

# 8 Results from using the CAPE closure timescale function

Before showing some results, it is worth re-visiting the question of why the two case studies were chosen by looking at the CAPE in each case. Figure 8 shows the highest values of CAPE that were diagnosed during each event from 4-km gridlength simulations. They are very different – and it was considered important to examine two events with very different values of CAPE, especially when the CCT is by definition related to CAPE. The 02/07/99 storms grew in an environment with CAPE values of over 4000 J/kg – at least in the model simulation. Such values are extremely large, even in the USA, and will result in large violent storms. The 03/05/02 showers grew in CAPE of up to 200-400 J/kg, which is fairly modest and common.

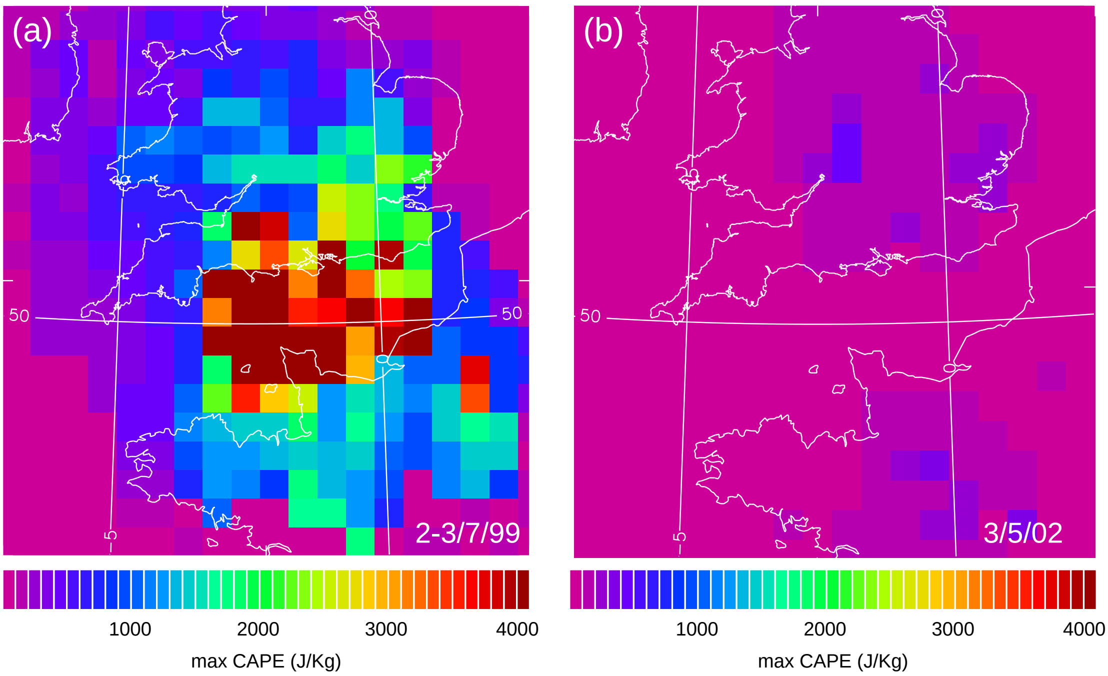


Figure 8. Pictures of the maximum CAPE to be diagnosed over 40x40km squares from (a) 16 UTC 02/07/99 to 01 UTC 03/07/99 from a 4-km gridlength forecast starting at 15 UTC (following 6 hours at 12km), (b) 10 to 18 UTC 03/05/02 from a 4-km gridlength forecast starting at 01 UTC.

## 8.1 Re-visiting case-study 1 02/07/99

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 some of these runs is shown in Figure 9. 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 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 or 10.0, Figures 9 (a) and (b)). That was the initial aim and it seems to have succeeded. When the CAPE dependent function put less of a restriction on the mass flux (c=100.0 or 250.0, Figures 9(c) and (d)) there were, once again, signs of the development of spurious convective bands.

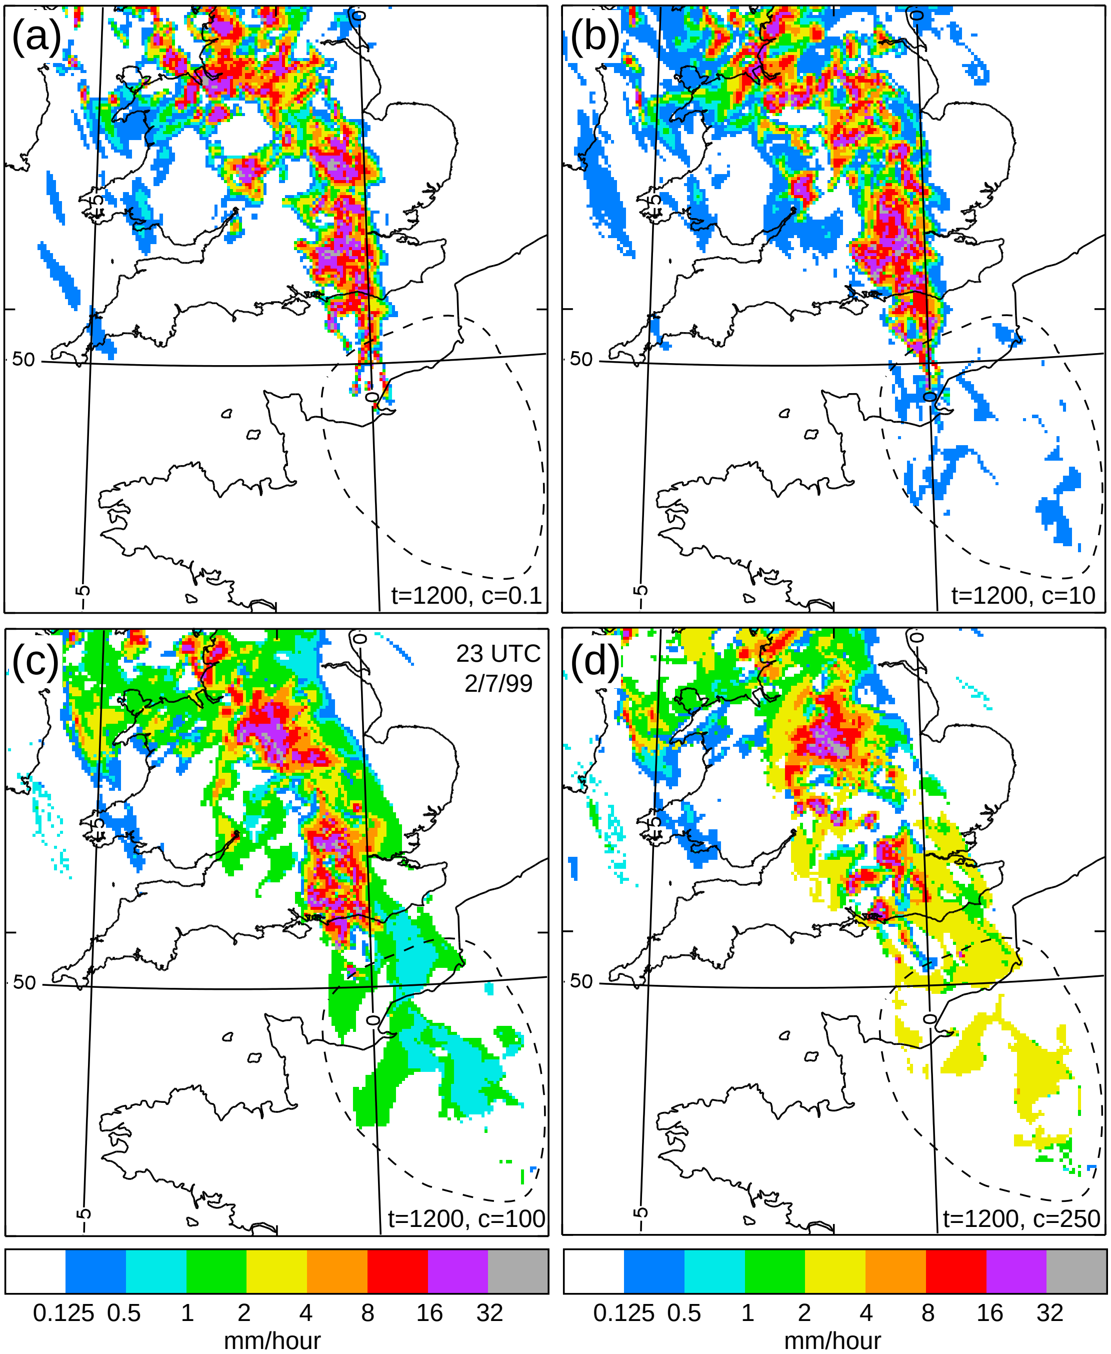


Figure 9. Rainfall rates at 23 UTC 02/07/99 from 4-km gridlength forecasts starting at 15 UTC (following a 6-hour forecast at 12km) with different values of the c parameter in the CAPE closure timescale function.

An interesting aspect of this event is the behaviour in the region enclosed by the dashed line in Figure 9. This is an area where the convection scheme produced precipitation widely (even in Figure 9(a) but below 0.125mm/hour), but the run without the convection scheme did not generate storms in most of that area. It is likely that, because the convection scheme triggered, the region was very sensitive and slightly different runs without the convection scheme might have produced a very different forecast to Figure 2(a). The advantage of using the CAPE dependent CCT is that the forecast behaves like the no-convection-scheme simulation, but the convection scheme adds information about the regions of uncertainty in the forecast.

### 8.1.1 Rainfall rates

Rainfall rates were very high in the simulations of this case. The question is, were they unreasonable for a 4km-gridlength model and does the CAPE dependent CCT make any difference? Figure 10 shows the peak rates, mean of the highest five rates and mean of the highest ten rates from each of the simulations discussed. Apart from the run with the constant CCT of 1800s all of the points on the graph exceed 200mm/hour. These values should be compared with observed rainfall rates that were measured by a tipping siphon rain gauge (Figure 11). The rain rate observed at Reading was between 150 and 200 mm/hour over a period of around 6 minutes and as this was only a 'random' point value it is almost certain that higher rates occurred elsewhere. The storm was travelling at 69 km/hour, so for that rain-rate to be sustained for 6 minutes the area of intense rain must have been wider than 4 km. It therefore does not seem too unreasonable for the 4-km forecasts to have generated the rainfall rates they did.

In comparison to the no-convection-scheme run, the CAPE dependent CCT did not make much difference to the peak rain rates for values of c=0.1 and 10.0. This is fine, because we think those values are reasonable. It is interesting that the peak rates leapt up with a value of c=100.0 when it might be considered that they would, if anything, be reduced. This would again, tend to suggest that a value of c=100.0 is not a good choice.

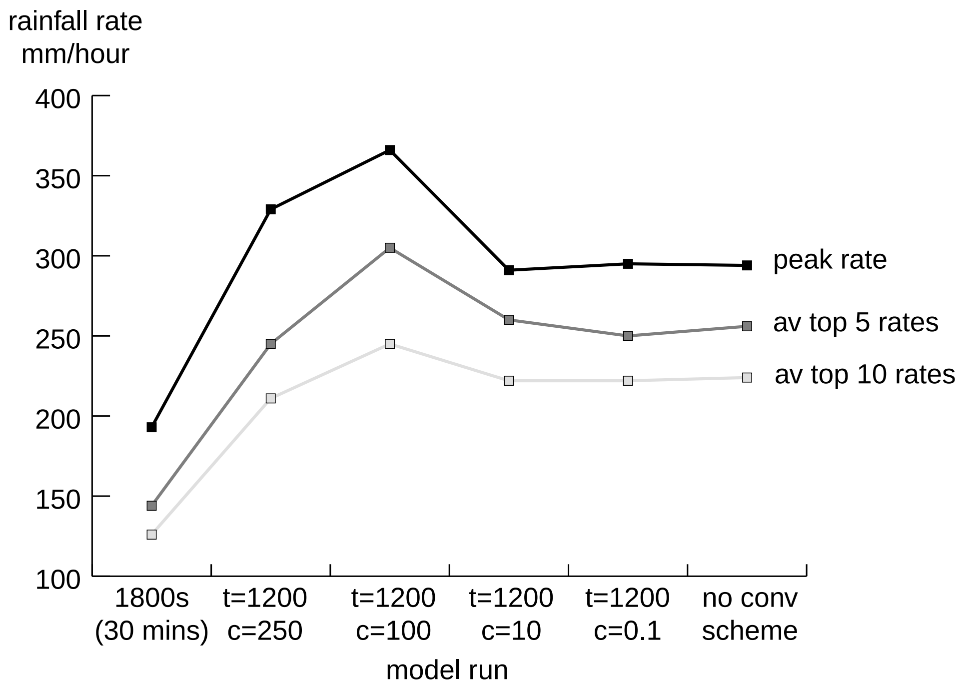


Figure 10. Graph of peak rainfall rates for different 4-km gridlength forecasts starting at 15 UTC 02/07/99 (following a 6-hour 12km forecast). 1800s refers to the run that used the standard CAPE closure timescale of 1800 seconds (30 minutes) and not a CAPE-dependent closure.

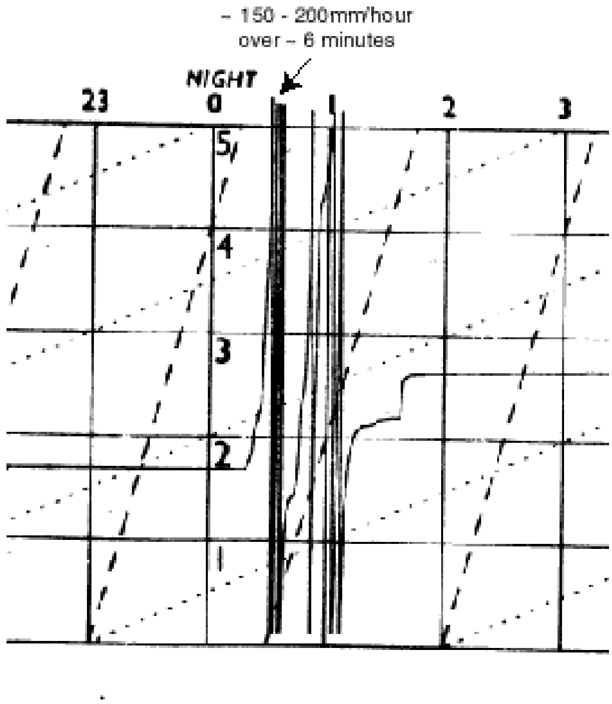


Figure 11. Tipping siphon measurement of rainfall rate at the Reading University field site on 3/7/99. Courtesy of Ken Spiers.

The CAPE dependent CCT has produced good results in this case study when the value of c=10.0 or 0.1.

## 8.2 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.

### 8.2.1 Peak rainfall rates

Figure 12 shows 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 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. It had the desired effect of reducing the very extreme single-point rainfall rates to values closer to the radar. When diffusion is added into a model, it allows some information to be passed between adjacent grid columns, which a convection scheme in incapable of doing directly. Diffusion could be thought to represent some of the mixing associated with the sub-grid-scale convection and in some ways behave like a convection scheme. The result shown here and from the HRTM testing indicate that the addition of some diffusion is probably desirable. The optimal amount may well be dependent on the way the convection scheme is used and will require further testing.

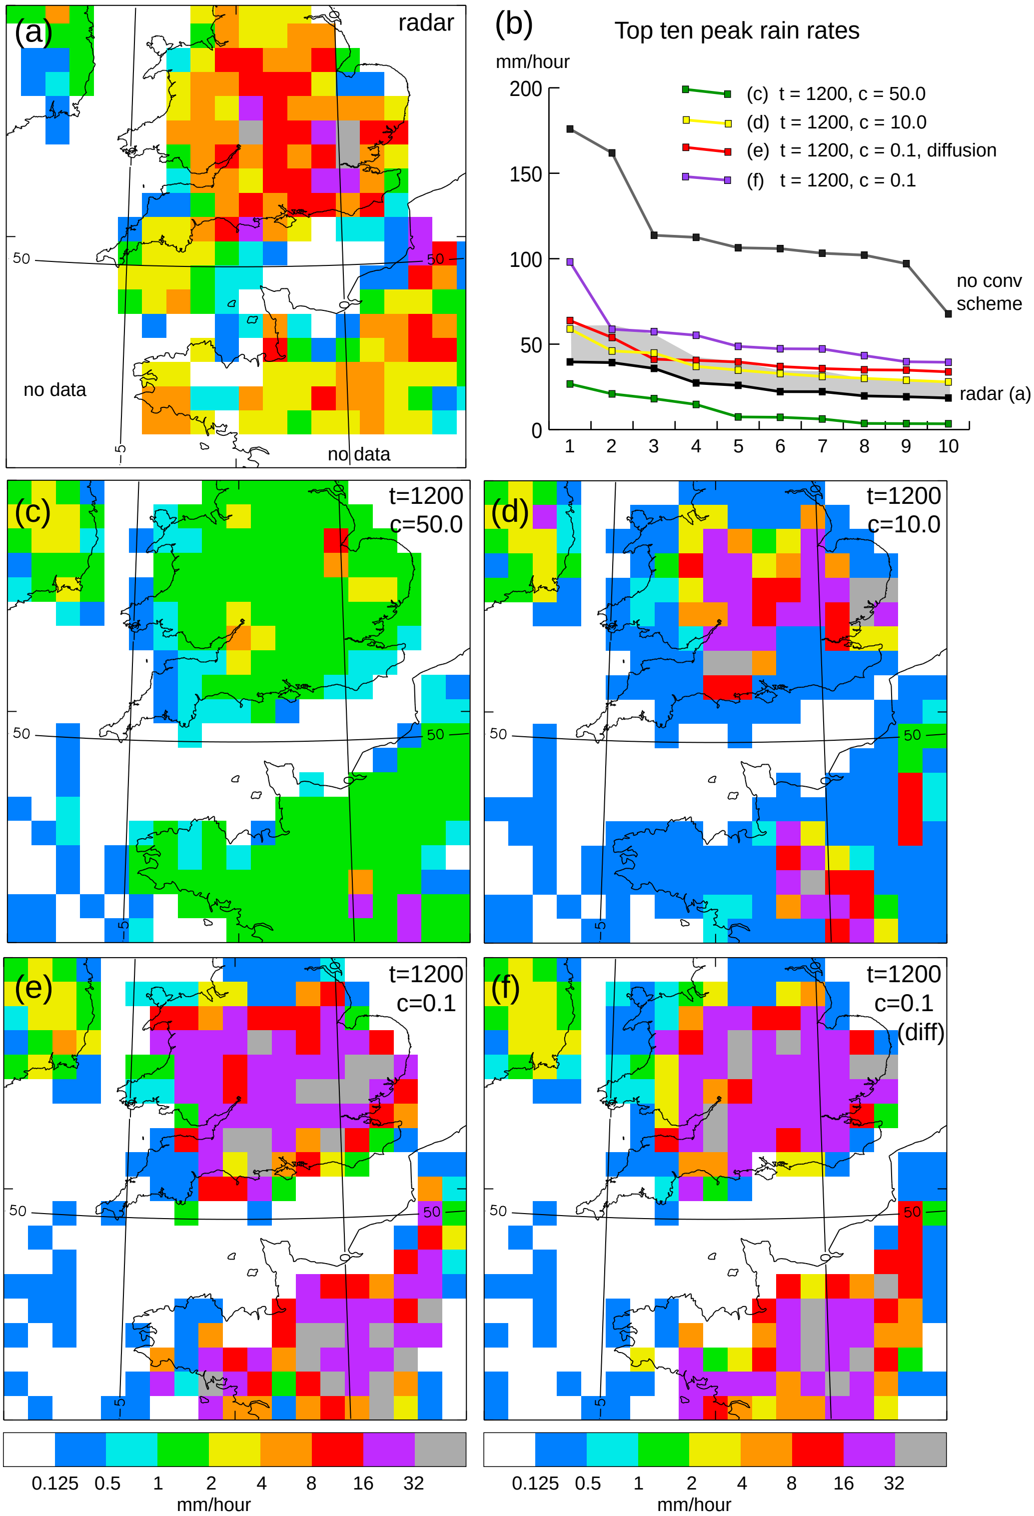


Figure 12. (a) Peak rainfall rates over 40x40 km squares within the period 10 to 18 UTC 03/05/02 from (a) radar and (c) to (f) 4-km gridlength runs discussed in the text with different values for the c parameter in the CAPE closure timescale function. (b) Graph of the top ten peak rainfall rates from (a) and (c) to (d) and Figure 5(b).

### 8.2.2 Initiation time

We have already seen that the showers in the run without the convection scheme started too late. Figures 13 and 14 show 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. Figure 13 shows a big difference of 2 hours between c=10.0 and c=2.5 for generating rainfall rates over 1.0 mm/hour. 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 problem of initiation will be discussed again later.

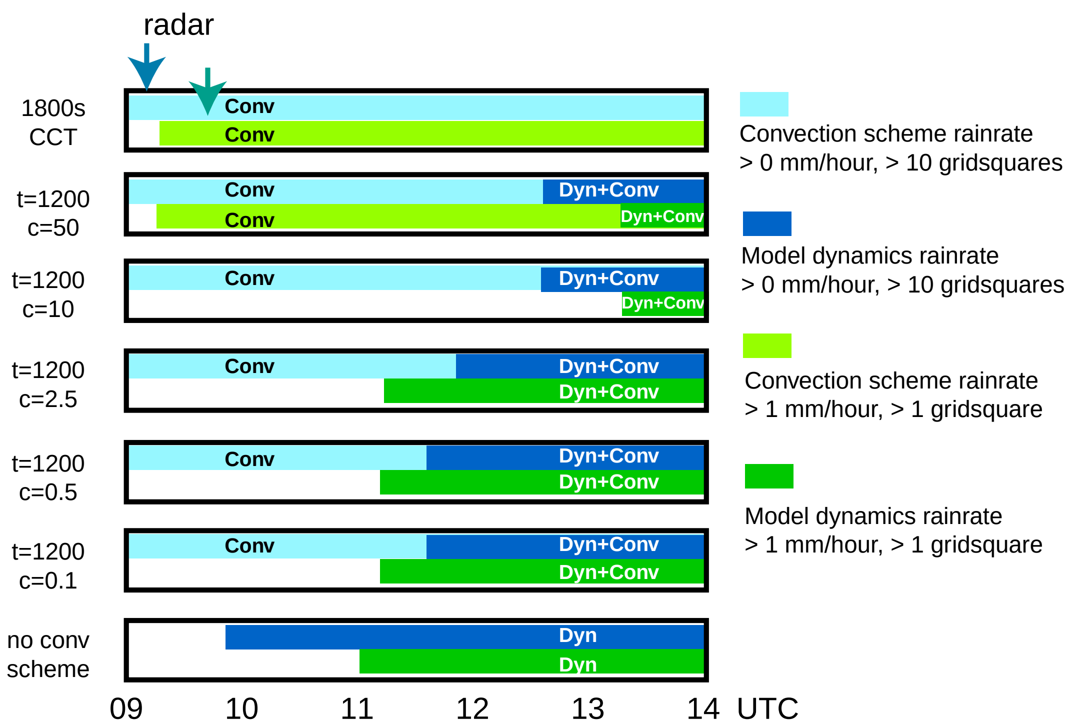


Figure 13. Diagram to indicate the onset of convection in different 4-km gridlength model runs. Light blue strips show when more then 10 gridsquares of precipitation from the convection scheme occurred, dark blue strips show when more than 10 gridsquares of precipitation from the model dynamics occurred. Light green strips – convection scheme > 1mm/hour (2 or more gridsquares), Dark green strips – model dynamics > 1mm/hour (2 or more gridsquares). 1800s refers to the run that used the standard CAPE closure timescale of 1800 seconds (30 minutes) and not a CAPE-dependent closure.

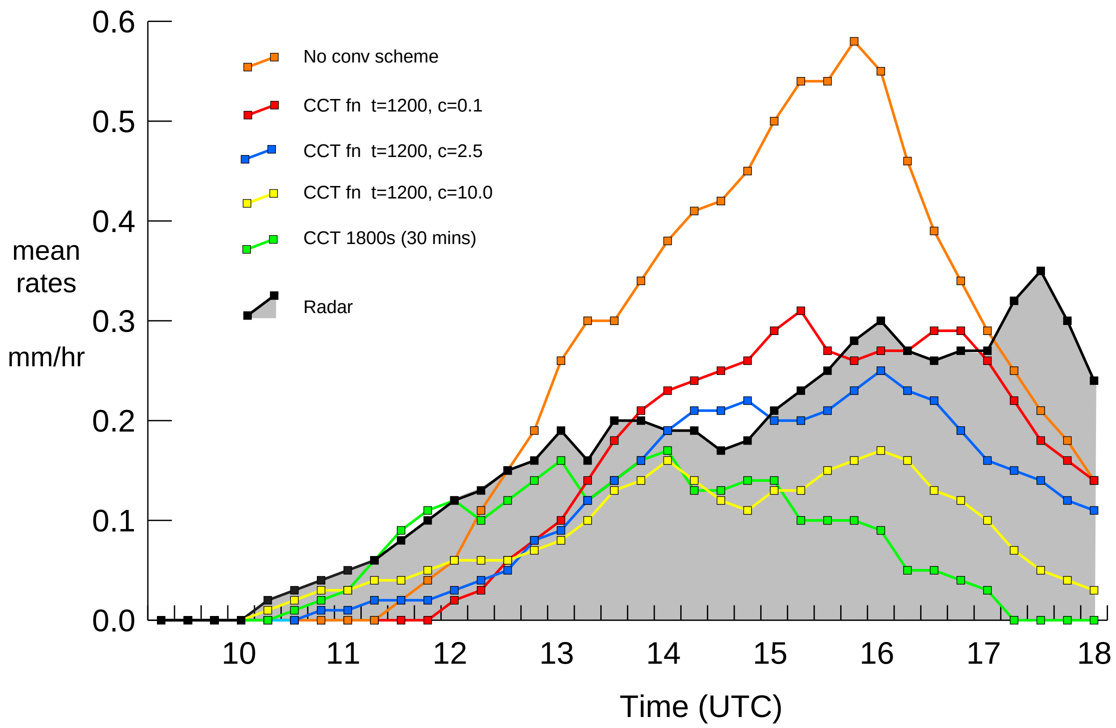


Figure 14. Graph of the mean rainfall rates within the area shown in Figure 4 against time for 4-km gridlength forecasts with different values of the c parameter in the CAPE closure timescale function (coloured lines except orange), 4-km forecast with no convection scheme (orange line) and radar (black 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.

### 8.2.3 Persistence into the evening

The use of the CAPE dependent CCT 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 14 shows that the 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. By 1800 UTC they were producing less rain than the radar, but still maintained significant shower activity. 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. 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 (Roberts 2003). This trend in behaviour with different values of c is encouraging because it fits with what was intuitively expected.

## 8.3 A third case, 11-12/10/2000

This is another case that was covered in the first report (Roberts 2003). Very high rainfall amounts (>70mm)were measured in a narrow band over parts southeast England.

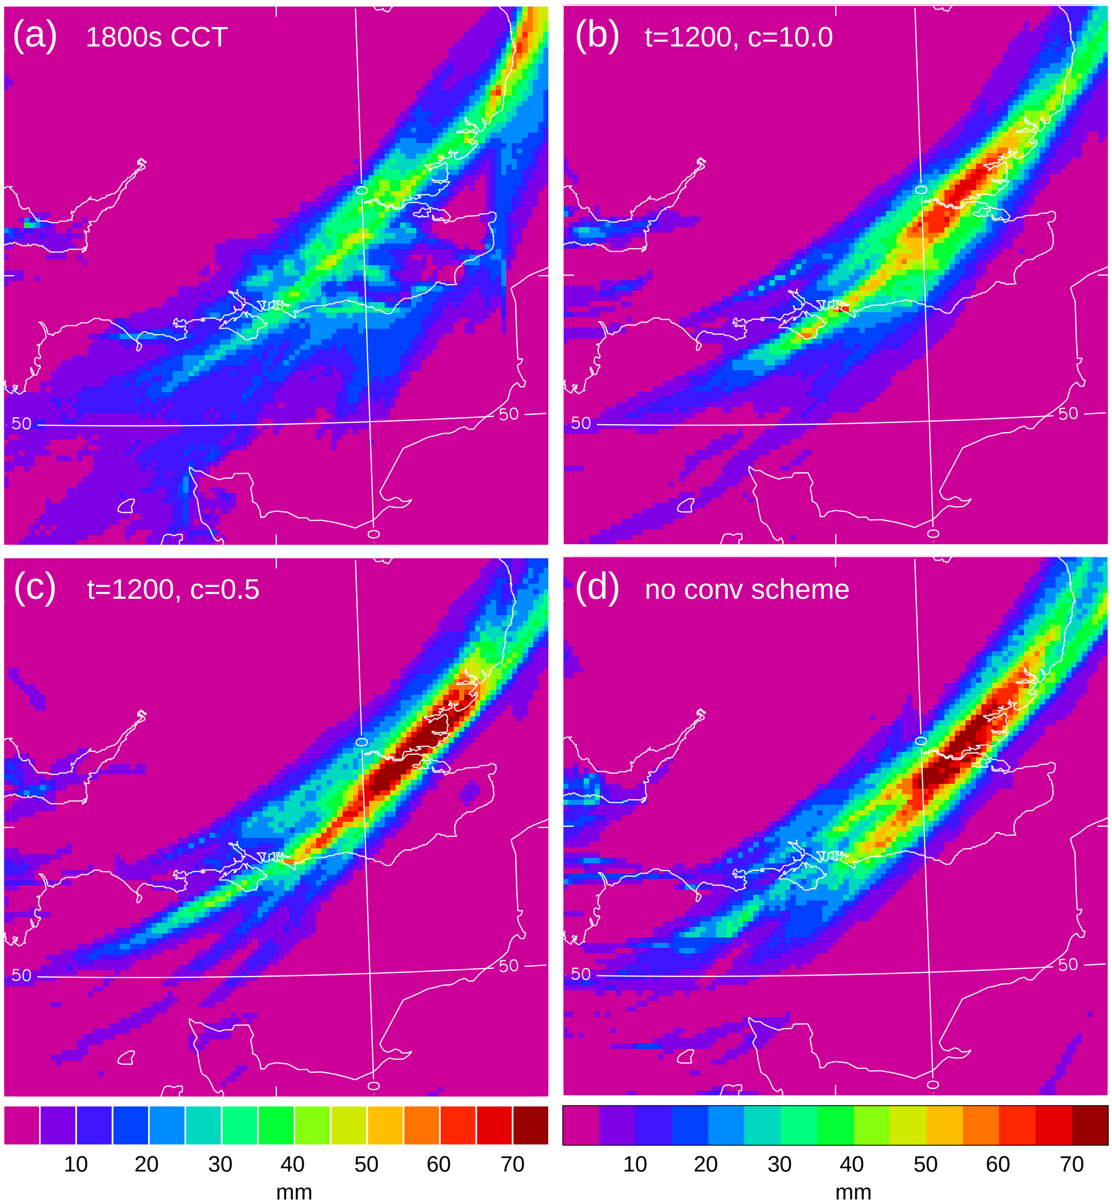


Figure 15 Rainfall accumulations over 18 hours from 18 UTC 11/10/00 to 12 UTC 12/10/00 from 4-km gridlength forecasts that started at 18 UTC 11/10/00. (a) Single CCT of 1800s, (b) CAPE dependent CCT t=1200, c=10.0, (c) t=1200, c=0.5 and (d) no convection scheme.

An interesting result from the first report was that 18-hour rainfall accumulations were more accurate, both in terms of location and intensity, in a 2-km gridlength simulation than a 4-km gridlength run (both using the convection scheme with CCT=1800s). The purpose here, is to examine the sensitivity of the rainfall accumulations to the use of the CAPE dependent CCT in the convection scheme in just 4-km gridlength simulations. The accumulations from the original run with CCT=1800s, two CAPE dependent CCT runs and a run without the convection scheme are shown in Figure 15. The run with no convection scheme and the two CAPE dependent CCT simulations produced larger rainfall amounts than the forecast with the constant CCT=1800s and also shifted the band further southeast. Both aspects agree better with radar. Again, there appears to be a trend in the CAPE dependent CCT runs – the smaller the value of c (larger t/c) the more the output looks like the no-convection-scheme run (reinforced by other runs not shown). The best forecast from the four shown was the t=1200, c=0.5. It did not have the glitches in the southeast part of the domain or the unrealistically high rainfall rates (not presented) that the run with no convection scheme had and was a considerably better forecast than the CCT=1800s run. It even made the output from the 4-km simulation look much closer to the 2-km output (shown in report 1, Roberts 2003) without the need for extra resolution. However, it should also be pointed out that the 2-km simulation could also be improved with the use of a CAPE dependent CCT (not shown here).

### 8.3.1 Diffusion

Figure 16 shows the impact on the rainfall accumulation of adding diffusion to a 2-km gridlength simulation. Neither of these runs had the convection scheme switched on. Adding the diffusion has had the desired effect of removing the nasty east-west grid-row glitch of higher accumulations over the Channel in the run without diffusion. However, it is also noticeable that the overall pattern of accumulations has not changed much, except for reducing a band of the largest totals exceeding 60 mm in the no-diffusion run. Even though the inclusion of diffusion is beneficial, it does not have the same impact on the general pattern that the convection scheme experiments do in Figure 15.

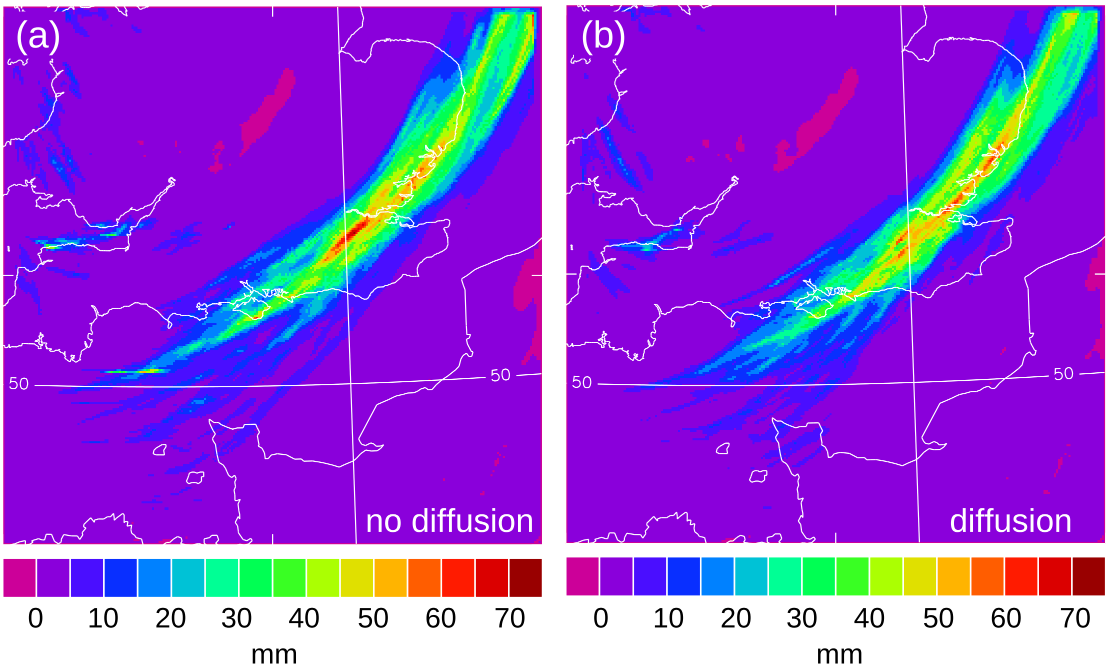


Figure 16. 18-hour Rainfall accumulations, 18 UTC 11/12/00 to 12 UTC 12/12/00, from 2-km gridlength forecasts with no convection scheme. (a) No added diffusion, (b) Diffusion added (del-4 diffusion, coefficient 5.276e3)

# 9 Conclusions

## 9.1 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 (case study 1), an interaction between the convection scheme and the model dynamics can generate spurious rainbands. In regions of lower CAPE (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 sensible 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 shown in Report 1 (Roberts 2003).

The best choices for the c and t parameters cannot be exactly known without endless experiments, 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. The function could also be beneficial in the operational 12-km gridlength mesoscale model with values of t and c set to restrict the convection scheme less (t/c smaller).

## 9.2 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 may be left with insufficient instability to trigger resolved showers. 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 (report 1, Roberts 2003). A 1-km model is supposed 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 (organisation, rain rates).

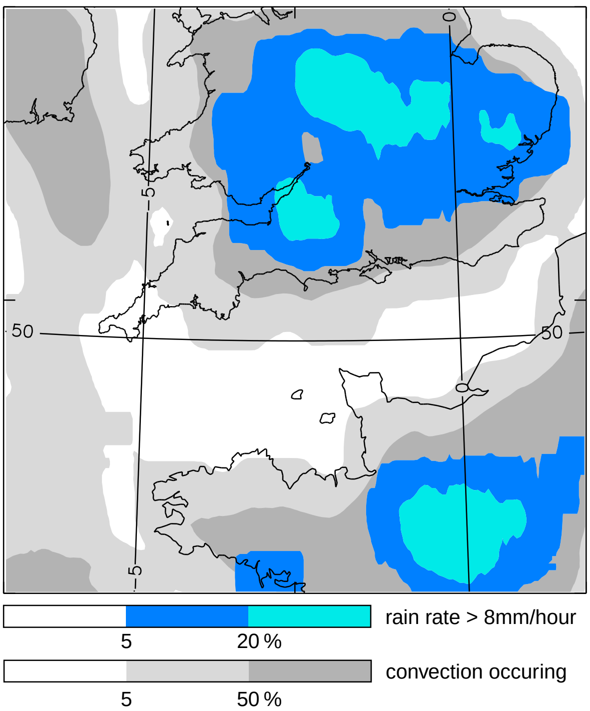


Figure 17. Example picture of a product that could be generated from a 4-km gridlength forecast with the CAPE closure timescale function in this case t=1200,c=0.5). Grey shades show the probability of convection occurring between 12 and 18 UTC 03/05/02 based on the proportion of gridsquares with convection-scheme triggering in 60x60km squares surrounding each gridpoint. Colours show the probability of rain rates above 8 mm/hour in the same period.

This can be an advantage if suitable precipitation diagnostics are produced because a deterministic forecast of the resolved convection also has an element of uncertainty attached – see Figure 17. It is not wise in any case to present raw output from high-resolution models because of the danger of believing fine-scale detail that is beyond the accuracy of the model.

### 9.3 Future work

Future work should consist of testing the CAPE closure function on more case studies and at ~1-km gridlength. This has begun. The High Resolution Trial Model (HRTM) project at JCMM (Lean 2003) has already been testing this approach with different values of c and t fixed at 1200. The function has been put into an idealised model framework (Roadnight 2003) and produced encouraging results. The impact of adding random temperature and humidity perturbations on triggering should be tested. The sensitivity to changes in other model parameters also needs to be examined. This includes the testing of a routine to add diffusion in a way that is dependent on vertical velocity (Terry Davies, personal communication) as well as documentation of the combined effect of using the standard method of adding diffusion with different CAPE closure functions. Ultimately, it is intended that this CAPE dependent closure function approach should be a stop gap method until superseded by more sophisticated techniques for dealing with sub-grid-scale clouds.

## 10 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*

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

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

Gregory, D., Inness, P. and Gregory, J. M., 1999: The convection scheme. *Unified Model Documentation Paper.,* **27**

Lean, H. W., 2003: High Resolution Trial Model project, results following the completion of stage 3. *JCMM internal report.*, In preparation

Roadnight, C. R., 2003: The representation of convection with a 4km gridlength: An idealised study. *JCMM internal report in preparation*

Roberts, N. M., 2003: Results from high-resolution simulations of convective events. *JCMM internal report.,* **140**. *(NWP Technical Report.,* **402***)*

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.

## 11 Acknowledgements

This work was partially funded by Defra under the Flood Forecasting and Warning Theme. I would like to thank Brian Golding for much needed comments on an earlier draft of this report. I would also like the thank Peter Clark, George Craig, Brenda Cohen, James Done Roy Kershaw and Carol Roadnight for discussions that have improved my understanding of this subject. I would also like to thank Humphrey Lean for technical help and for valuable discussions about this work.