DOE Final Technical Report

Award: DE-SC0006805 Period of Performance: 09/15/11-09/30/15
PI: Robert Walko

Project Title

Toward a Unified Representation of Atmospheric Convection in Variable-Resolution Climate Models

Motivation and Project Goals

The purpose of this project was to improve the representation of convection in atmospheric weather and climate models that employ computational grids with spatially-variable resolution. Specifically, our work targeted models whose grids are fine enough over selected regions that convection is resolved explicitly, while over other regions the grid is coarser and convection is represented as a subgrid-scale process. The working criterion for a successful scheme for representing convection over this range of grid resolution was that identical convective environments must produce very similar convective responses (i.e., the same precipitation amount, rate, and timing, and the same modification of the atmospheric profile) regardless of grid scale. The need for such a convective scheme has increased in recent years as more global weather and climate models have adopted variable resolution meshes that are often extended into the range of resolving convection in selected locations.

Variable-resolution computational grids can substantially improve the benefit-to-cost ratio in many environmental modeling applications, but they can also introduce unwanted and unrealistic numerical anomalies if not properly utilized. The PI and collaborators had previously found in exploratory simulations of horizontally homogeneous convective environments that resolved (non-parameterized) convection develops most quickly in the highest-resolution region of a variable grid, and progressively more slowly over regions of progressively lower resolution, even when all such regions adequately resolve mature convection. Furthermore, on variable grids that transition from resolved to parameterized convection, timing and intensity of the convection in both regimes is generally disparate unless special care is taken to tune the parameterization. In both cases, the convection that develops first (due to purely numerical reasons) tends to suppress convection elsewhere by inducing subsidence in the surrounding environment. This highly nonlinear competition, while desirable when induced by natural causes such as surface inhomogeneity, is highly undesirable when it is a numerical artifact of variable grid resolution and/or selective application of convective parameterization.

The following two examples illustrate the unwanted impact of variable-resolution grids on convection in idealized numerical simulations. We use the Ocean-Land-Atmosphere Model (OLAM) (Walko and Avissar 2008a, 2008b, 2011; Medvigy et al. 20011) which employs a global computational mesh of hexagons that can seamlessly transition to locally higher resolution.

The idealized simulations use a multi-resolution grid whose highest resolution region is shown in Figure 1 with grid spacings of 24, 12, 6, 3, and 1.5 km. Both simulations are initialized with horizontally homogeneous atmospheres that are motionless and conditionally unstable to deep convection, and a surface sensible heat flux is applied that is horizontally homogeneous (except for small random perturbations) and initiates convection within the first hour.

The first case is run with no cumulus parameterization anywhere, which means that the larger grid cells have no suitable mechanism for convection. However, deep convection can be reasonably resolved (or is at least “permitted”) in the finest 2 or 3 resolution zones. The top panel of Figure 2 shows accumulated precipitation after convection has developed, reached maturity, and mostly dissipated over the region shown (note that the region covered is slightly smaller than in Figure 1), while the bottom panel shows an azimuthal average of accumulated precipitation as a function of radial distance from the center of the region. Precipitation was found to develop earliest in the finest (1.5 km) mesh area, next in the 3.0 km grid cells, and subsequently in the 6.0 km cells. Precipitation maxima and minima occur at each transition in resolution, until coarse grid spacing suppresses convection altogether.

The second case is run with a conventional cumulus parameterization on the 6.0 km and coarser regions of mesh, but not on the 1.5 and 3.0 km regions. (The reasoning behind this choice is that it is preferable to directly model convection on those portions of the grid capable of resolving it in order to best represent both the causes and consequences of convection, whereas over lower resolution regions of the grid, there is no available alternative but to parameterize convection as a subgrid-scale process.) The results are shown in Figure 3. Convection was found to develop earliest in the larger grid cells because the cumulus parameterization happens to decide that environmental conditions are suitable for convection before resolved convection develops in the smallest grid cells. The latent heat released by the parameterized convection warms the atmosphere relative to the high resolution center region of the grid, which induces a solenoidal circulation with subsidence in the center region. This acts to stabilize the high resolution grid cells and suppress convection from developing there, although a small amount of resolved precipitation does occur in the 1.5 km grid cells. A sharp precipitation maximum occurs at the edge of the region of parameterized convection where abundant water vapor flows outward from the central subsidence region and supplies the convective parameterization in the adjacent grid cells.

The preceding results provide clear demonstration that variable resolution grids can cause anomalous convective behavior in atmospheric numerical simulations, and are similar to simulations that motivated this research. We proposed to carry out four research tasks with the goal of substantially reducing these anomalies, using OLAM and its variable mesh as a developmental testbed. These tasks are listed individually in the next section, along with the research accomplished for each.

Research Accomplishments

Task 1.

This task will examine the sensitivities of existing conventional convective parameterization schemes and explore ways to tune or adjust them for improved responses to their model environments. We will specifically examine the parameterizations by Kain and Fritsch (1990, 1992), Grell and Devenyi (2002), and Emanuel (1991), whose implementation in OLAM is already complete or is currently in progress. We will evaluate these schemes at different resolutions and in different physical environments including tropical and extra-tropical, orographically and synoptically forced, and idealized (such as in the examples shown above). We will consider whether each scheme could be improved by directly utilizing information on grid spacing and its gradient, and develop a means to reduce the sensitivity of the convective response to these grid factors.

We tested all cumulus parameterization schemes listed above, and in addition we implemented and tested two other schemes, one by Tiedtke (1989) and the other an early version of a scheme that was later revised and published in Grell and Freitas (2014). One large set of tests consisted of global climate simulations of a few to ten years, some at uniform resolution of 100 or 200 km and others with variable resolution reaching to 50 km over most of South America. These simulations were shared with another project under separate (NSF) funding and were part of a comprehensive evaluation of OLAM and all of its parameterizations. In the early stages of these tests, it became apparent that the choice of convective parameterization had a very large impact on simulated precipitation amounts and global distribution, as well as on surface temperature and radiative fluxes at the surface and top of the atmosphere. In addition, we found that although a particular convective parameterization might produce the best results when used in combination with one radiative transfer parameterization (we tested RRTMg and a scheme by Harrington) or boundary layer parameterization (we tested a Smagorinsky-Lilly formulation and a non-local PBL scheme), a different convective parameterization worked better if either or both the other parameterizations were changed. We thus conducted separate simulations for all possible combinations of these parameterizations. We also developed a new parameterization for sub-grid fractional cloudiness that is used by RRTMg and tested it with a range of parameter settings in combination with various sets of other parameterizations. Simulated surface temperature and precipitation, and surface and top-of-atmosphere radiative fluxes were compared with observed climatology (e.g., GPCP and CERES data), and we found that much better matches were obtained with certain combinations of parameterizations than with others.

Some of these global climate test results were presented and summarized in our request for a no-cost extension to this DOE project. The extension was necessary because the overall effort of these test simulations became a far larger task than anticipated. However, we considered it essential for all our research projects to use the most accurate configuration of OLAM possible, and that configuration needed to be determined before proceeding with the remaining research.

Among the 5 convective parameterizations that we tested in the global climate simulations, the Emanuel and Grell-Freitas schemes generally performed the worst, so we did not run additional tests or evaluations of them for this project. (We are currently testing and evaluating the revised Grell-Freitas parameterization in OLAM.) We conducted additional tests of the Kain-Fritsch, Grell, and Tiedtke parameterizations at several different resolutions. We encountered situations over the tropical ocean at medium resolution (tens of kilometers) where parameterized convection tended to be self-sustaining in some grid cells, acting to suppress convection in adjacent cells. Detailed analysis showed this to involve a positive feedback mechanism in which latent heating released by convection drove a solenoidal mesoscale circulation at grid scale that continuously fed moisture and energy to the convecting cell through boundary layer convergence. This is of course not unlike the mechanism that can sustain any convective cell, but it is unnaturally enhanced in the present case by the sensitivity of the convective parameterization to its environment. The effect was by far most pronounced with the Kain-Fritsch parameterization when used simultaneously with the Smagorinsky-Lilly PBL parameterization. However, certain other combinations of parameterizations appeared to occasionally produce the same effect, albeit much more weakly. To mitigate this effect, we developed and implemented a scheme that spreads a portion of the parameterized convective heating to the grid cells immediately adjacent to the convection. In the most extreme case (Kain-Fritsch with Smagorinsky-Lilly), each adjacent grid cell would receive up to 30% of the latent heat retained in the convective cell following the spreading, which means that the majority of the total latent heating (usually about 2/3) was removed from the convecting cell and distributed among the adjacent cells. The need for latent heat spreading was found to be less or nil for larger grid spacing (e.g., 200 km) so the fractional amount of heat transferred to adjacent cells was decreased linearly with grid spacing. For other combinations of convective and PBL parameterizations, the need for latent heat spreading was low enough that we did not undertake a series of tests to tune and optimize the scheme, instead opting to avoid its use altogether.

Through latent heat spreading, we were able to eliminate the positive feedback effect and obtain a more random and uniform distribution of convective-parameterized precipitation when averaged over time. The latent heat spreading also resulted in an overall reduction in convective activity, thus providing a means for tuning (moderating) the convective parameterizations. Our overall conclusion from tests of all convective parameterizations, however, was that the Kain-Fritsch scheme does not produce the best agreement with observations. Instead, the Grell scheme was the best-performing (in OLAM) among all the convective parameterizations.

Task 2.

This task will examine the relative responses of convective parameterizations and the combined system of resolved convection and the OLAM bulk microphysics scheme. This scheme was implemented in OLAM under current DOE support and is currently being evaluated for both wintertime orographic precipitation and convective cases by our collaborators at Colorado State University. In view of the sometimes disparate response between resolved and parameterized convection that we encountered in the exploratory experiments described in the Introduction, this task is particularly aimed at balancing these responses, particularly in timing. Most importantly, an obvious way to improve the coexistence of both schemes in different parts of the model grid is to delay the response of the convective parameterization by a time interval that is comparable to what is required for resolved convective motions to develop in the high resolution areas of the grid. To some extent, this task may involve a limited amount tuning of the bulk microphysics scheme itself, which contains a few parameters whose optimal values have not yet been well established.

Having concluded from Task 1 that two of the convective parameterization schemes that we tested did not perform well overall, we concentrated on the Kain-Fritsch, Grell, and Tiedtke parameterizations for Task 2. We simulated two real-world test cases for comparing these schemes on highly variable OLAM grids alongside selected regions where convection was resolved on the grid and convective parameterization was turned off. The first was a case from the TWP-ICE field experiment (19 January 2006). A principal objective (under a separate DOE project) was to resolve deep convection so that modeled dynamics and microphysics fields could be compared in detail with observations. We thus used substantial local refinement of OLAM's global grid to achieve this resolution with affordable computational cost. Starting with a grid cell size of 100 km over most of the globe, resolution was progressively doubled at radial distances of 3000, 2000, 1000, 800, 600, and 400 km from the TWP_ICE site, culminating in a grid spacing of 1.6 km over a circular area 800 km in diameter and centered on the site. Convective parameterization was activated in the coarsest resolution regions of the grid with grid cell sizes of 100, 50, 25, and 12.5 km, while it was switched off in the regions with 6.4, 3.2, and 1.6 km grid cell sizes. Selective use of convective parameterization is standard practice in nested grid models whose grids transition between areas that resolve convection to those that do not. It is well known that grid spacings of 5 or 10 km are too coarse to adequately resolve convective updrafts and too fine to fit the conceptual models on which the parameterizations are based. Therefore, the portions of the OLAM grid whose resolution is in that range do not have adequate representation of convection, and the choice of whether or not to activate convective parameterization there is challenging. In a simulation like the TWP-ICE case, this shortcoming can be mitigated by making the 1.6 km grid region very much larger than the focus region (the field site) so that the intermediate-sized grid cells are less influential on the results there. This is one of the main reasons for choosing the large extent of this region.

Figure 4 shows accumulated precipitation over a 24-hour period ending at 00 UTC on 19 January 2006 for a 2000 km x 2000 km area that is centered on the field site. It is immediately apparent that the precipitation amount and distribution pattern changes between the inner region where convective parameterization is switched off, that is, within a radius of 800 km of the center of the plot, and the outer region where convective parameterization is allowed to operate. In the inner region, narrow bands of higher precipitation are interspersed with paths that received much less precipitation or none at all. The precipitation bands were produced by convective cells that moved toward the east-southeast, and these systems were reasonably well resolved on the model grid, as seen in Figure 5 which zooms in on the center of the region depicted in Figure 4 and adds a grid overlay. Outside the radius of 800 km, precipitation is sharply higher. It is evident that the convective parameterization is a much more active producer of precipitation than the microphysics parameterization with well resolved convective dynamics. Figures 6 and 7 show individual contributions from convective parameterization and microphysics to the total precipitation shown in Figure 4. Precipitation from convective parameterization is especially heavy on the southeastern side of the central region, which is immediately downwind of the region where only microphysics is producing precipitation. The convective parameterization, being tuned to treat that environment as more highly convective than the model's own resolved dynamics and microphysics machinery, abruptly extracts more moisture from the air. By the same mechanism, the western and northwestern side of the 800 km radius circle are upwind, and air passes first through the zone where convective parameterization is allowed to operate and then into the central high resolution part of the grid where parameterized convection is switched off. The parameterization upwind has been excessive in its removal of moisture and convective energy from the air, leading to a rain shadow in the inner region, as seen in Figure 4, and the rain shadow extends for hundreds of kilometers downwind toward the center of the plot.

The rain shadow effect is particularly troubling because it is a purely numerical artifact caused by a mismatch in the behavior of parameterized and resolved convection in the model. Even though we configured a convection-resolving region of the grid that was far larger than the TWP-ICE field site and large enough to accommodate convective systems through their life cycles, the presence of the convective parameterization hundreds of kilometers away was projected much closer to the site. We considered this phenomenon to be highly important, not only for its impact on this project, but to modeling by the community in general because grid nesting or variable-resolution grids that transition from grid-resolved to parameterized convection are very commonly used. Moreover, this is a simulation of an actual case, not a highly idealized case such as those that motivated this current DOE project and were designed to highlight convective anomalies in connection with selective use of convective parameterization.

The above simulation, which used the Kain-Fritsch convective parameterization, was next repeated for the Grell and Tiedtke parameterizations. These two were found in Task 1 to be progressively less active than Kain-Fritsch, and the same was found here. Results using Tiedtke, the least active of the parameterizations, are shown in Figures 8, 9, and 10, which correspond to Figures 4, 6, and 7, respectively. Convective precipitation is much less with Tiedtke (Figure 9; notice the nonlinear precipitation scale) and the rain shadow effect is much weaker and less extensive (Figures 8 and 10). In addition, precipitation amounts are generally higher in the southeastern quarter of the inner 800 km diameter circle, showing that the difference in activity of the convective parameterizations has an impact hundreds of kilometers downwind.

For our second test case, the convection-resolving (or permitting) cells of the OLAM grid were located over a large portion of the equatorial Indian Ocean in an attempt to replicate MJO convective activity and structures that were observed around 20 November 2011 and afterwards. As in the TWP-ICE test case, the MJO simulation was first run using the Kain-Fritsch convective parameterization. Compared with observations, simulated precipitation was relatively weak over the equatorial Indian Ocean and self-sustaining MJO-like structures did not occur. A plot of mean rainfall rate over a 12-hour period from this simulation is shown in Figure 11. The region where convective parameterization is turned off is less evident than in the TWP-ICE simulations, but can nevertheless be roughly discerned in the plot. It extends approximately from -2000 km to +1600 km in the east-west direction (on the equator) and from -1000 km to +1000 km in the north-south direction, and is rounded on the eastern and western ends. Total precipitation is generally more abundant where parameterization is turned on, again indicating that the Kain-Fritsch parameterization is overly active compared with convection that is represented by resolved microphysics and dynamics on the smallest grid cells.

In view of the solenoidally-induced subsidence mechanism that we had previously identified and analyzed, we hypothesized over-activity of the convective parameterization to be a possible cause, or at least a contributing factor, for the relatively weak convection and failure to simulate any MJO-like organized structure in the high-resolution part of the grid. We investigated this further by re-running the final 12 hours of the preceding simulation three times, once with the Grell parameterization, once with the Tiedtke parameterization, and once with no convective parameterization at all, the results from which are shown in Figures 12, 13, and 14, respectively. Each case produces successively more active precipitation on the resolved grid, with horizontally-averaged precipitation rates of 3.84, 4.36, 4.55, and 5.01 mm/hr respectively, in the four simulations. These differences gradually increased over time as the Grell and Tiedtke simulations were continued for another 48 hours or more and became more distinct from the original Kain-Fritsch simulation from which they were re-started (results not shown). We are encouraged that these results now more closely resemble the observed MJO event, and we are continuing our simulations of the MJO event independently

In summary, these simulations demonstrated that by running OLAM with less-active convective parameterization schemes, the suppression of resolved convection in the non-parameterized region is substantially reduced. Additional experiments in which resolved microphysics was toggled between single- and double-moment formulations produced comparatively small changes in resolved convective activity, and therefore were not pursued further. In spite of the partial success achieved from Tasks 1 and 2 toward reducing the disparity between the activities of parameterized and resolved convection, the disparity remains unacceptably large. Furthermore, it is case-dependent, which implies that further reduction by these methods will often require independent tuning for new cases or tests under a wide range of conditions and much trial and error in pursuit of a universally adaptable scheme. We thus think that our proposed Task 3, which is the main new innovation in this project, holds the greatest promise in achieving the goal of eliminating grid scale effects on the bulk properties of simulated convective activity.

Task 3.

We will implement and test a new method of modeling convection on a temporary limited-area high resolution grid using an approach that is popularly known as a "superparameterization". The conceptual idea behind this approach is that instead of applying a conventional convective parameterization to coarse areas of the model grid, one can instead insert a small cluster of high resolution grid cell columns for resolving convection with far greater accuracy. The inserted cluster, or "convection grid" is introduced to any model grid cell where convection is remotely possible. (The bar is deliberately set very low so that insertion of the convection grid does not itself determine whether convection occurs. Insertion is avoided only in cases where convection is clearly impossible, and this avoidance is for no other reason than to conserve computational resources.) Once inserted, the convection grid is initialized with the mean properties of the parent grid cell, and is additionally perturbed about its own mean state such that some of its columns are slightly more buoyant and/or moister while others are less buoyant and/or drier, in order to represent expected natural subgrid-scale variability that would exist in the coarse parent grid cell. This variability may be determined, in part, from knowledge of underlying topographic or other surface features. The introduced perturbations develop convectively in the convection grid if and only if the environment is sufficiently unstable, the same as it would in high resolution areas of the parent grid.

Our proposed convection grid has the form illustrated in Figures 15 and 16. It is a horizontally unstructured grid consisting of a few cells of convection resolving scale, surrounded by others that are progressively coarser. The high resolution cells are intended to accommodate the most intense upward and downward convective motions, while the larger cells cover an area wide enough to encompass the primary mesoscale subsidence region. The mesoscale subsidence is thus well represented, even if the parent grid cell is too coarse to do so. The unstructured design of this convection grid is identical to that of the variable-resolution parent OLAM grid, and it will employ the same computational code. Thus, the convection grid is merely a small copy of the parent OLAM grid, and its implementation is a relatively simple task of bookkeeping and communication. The convection grid will be represented in the Fortran code as a defined data type, in which form an essentially unlimited number of copies can easily be spawned for any regions of the OLAM grid that require them. Each convection grid cluster will have the freedom to move horizontally with the mean propagation speed of the convective circulation, thus maintaining the intended alignment between convection and grid.

The effects of convection are eventually fed back to the parent grid, but this need not – and perhaps should not – be done immediately as the convection develops and runs through its life cycle. Instead, it may be best to keep the mesoscale compensating effects sequestered within the convection grid in order to maintain intact the immediate environment of the convection without dilution. In the latter stages of convection, however, the mean convection grid fields can be gradually mixed with parent grid fields. The best means for performing this will be a *principal topic to investigate for this task. We anticipate that the best means for recombining the convection grid and parent grid solutions will depend on the resolution of the parent grid. For example, Figure 16 shows how the convection grid might appear superimposed on a medium-resolution portion of the parent grid. The fact that the convection grid overlaps multiple parent grid cells means that the parent grid itself could represent the mesoscale subsidence region, so mixing of the solutions could commence sooner. It may be unnecessary for the convection grid to represent its own mesoscale region in this case.*

We first proceeded with Task 3 as planned by implementing the convection grid with the proposed configuration as illustrated in Figures 15 and 16. The OLAM code was structured so that replicas of the convection grid could be spawned as needed as a simulation progressed. Each replica would be initialized horizontally homogeneously with a vertical profile sampled from a column (on a coarse grid cell) or local group of columns (on mesoscale grid cells) on the parent grid where convection was possible, and small temperature perturbations were added. Then, forward integration was carried out on each replica for a period up to 1 hour to determine the convective response, if any. Horizontal averages of momentum, potential temperature, and moisture variables were taken and stored at frequent intervals from each integration. Finally, the averaged results were applied to the parent grid column (or group of columns) as integration on the parent grid commenced.

We soon discovered that we had overlooked a key problem in our initial design of this method: The convection grids themselves, possessing variable resolution, suffered from some of the same problems as we were trying to solve in the variable-resolution parent grid. That is, resolved convection tended to develop first in the smallest grid cells of the convection grid, followed by delayed convection in the slightly coarser cells, leading to grid-induced convective anomalies. Although we had originally envisioned that convection should develop on the smaller cells and that the mesoscale cells would serve to represent mesoscale subsidence, we did not anticipate the degree of grid imprinting on the resolved convection that we eventually encountered on the convection grid.

As a consequence, variable resolution on the convection grid, which had been intended to reduce the required number of vertical columns, was abandoned and replaced with a uniform mesh of high resolution. Furthermore, periodic lateral boundary conditions were implemented on the convection grid so that convection would not be unduly constrained in the horizontal by artificial lateral boundary conditions. Finally, the notion that each convection grid would follow the mean motion of the convective cells was dropped as it was no longer required: Convective cells could be free to propagate across a convection grid replica as it had neither variable resolution nor actual (non-periodic) lateral boundaries. The only drawback of the new configuration of the convection grid replicas was that they required a larger number of grid cells. At the same time, however, they offered much greater freedom for convection to develop that was now comparable to that of the highest resolution regions of the parent grid and therefore much more likely to yield a nearly identical convective response.

The new configuration of the convection grid replicas and their interaction with the parent grid is illustrated by the example in Figure 17. A portion of the variable-resolution parent grid is shown on the left, with the smallest grid cells occupying the circular grayish region at the center surrounded by a zone of grid cells of double width followed by another zone of grid cells having quadruple width. A single replica of the convection grid is shown on the right, along with the region on the parent grid it communicates directly with. At regular intervals during the time integration of the parent grid, such as every half hour, integration pauses and the convection grid procedure is carried out. A number of convection grid replicas are spawned with vertical and horizontal resolutions identical to those on the highest resolution portion of the parent grid, and each is initialized horizontally homogeneously from a vertical profile sampled from the parent grid. An individual sample is obtained from a region that is at least as large as the horizontal extent of the convection grid itself. In the illustration, a sampled region is any of the clusters of grid columns that share the same color. On the coarsest parent grid cells (not shown here), a sample would normally be obtained from a single grid column. Following the addition of small temperature perturbations, the atmosphere is integrated in time for up to an hour on each convection grid using essentially the same dynamic, thermodynamic, and microphysics machinery as employed on the parent grid. Radiative and surface fluxes are not evaluated on the convection grids but rather are copied from the most recent values on the parent grid. Horizontal averages of all prognostic variables are computed at frequent intervals (e.g., every 2 to 5 minutes) on each convection grid replica and stored. When the convection grid integration is completed, these horizontal averages are further processed in order to obtain a convective response that will be applied to the parent model grid. Forward integration then resumes on the parent grid, and the convective response is applied gradually over the same time period as the integrations on the convection grid replicas.

Another aspect of this procedure is that for each convection grid replica that is spawned and integrated, a corresponding “reference grid” is also spawned and integrated using the same initial condition, radiative, and surface fluxes from the parent grid. The reference grid differs from its corresponding convection grid only in that its grid cells are an exact replica of those from the sample region of the parent grid (shown on the left of Figure 17). Thus, they have lower resolution than the convection grid. The purpose of the reference grid is to determine what the convective response would be with the lower resolution but otherwise identical conditions used on the convection grid. The difference between the convection and reference grid solutions is the actual quantity that is used by the OLAM parent grid to represent the convective response. This is an important choice because the normal integration will be performed on the parent grid and it is essential to only add the difference that would arise from higher resolution.

In summary, we have developed a modified version of the proposed superparameterization and successfully implemented and tested it in OLAM.

Task 4

We will perform an evaluation of both approaches (Tasks 1-2 and Task 3) in selected case studies for which adequate data describing primary characteristics of convection are available. In particular, cases will be selected for which we can obtain suitable observations of the timing and horizontal scale of the convection, and reasonable estimates or measurements of the net latent heating and precipitation produced. We will select appropriate case studies in the course of this project, but anticipate using TRMM and other satellite products for estimating convective properties, as well as rain gauge and other products. Comparison with observations will be particularly important for establishing the solution toward which both methods should converge.

Case-study tests of the conventional convective parameterizations were described in the Task 1 and 2 sections of this report. Further testing of the new convective superparameterization developed under Task 3 is described here.

A major task in the development of this convection scheme is to determine the best methods of processing the differences from the convection and reference grid integrations and applying them to the host model simulation. The examples that follow were obtained from OLAM simulations of the TWP-ICE case and illustrate some of the wide variety of convection results that may be obtained. For illustration purposes, the convection simulations are carried out for 2 hours, but normally their duration would correspond to the convection time scale (1 hour or less).

The upper left panel of Figure 18 shows the change in total water, summed over all forms (vapor+ liquid + ice) for a high resolution (1 km) convection grid simulation (color contours) and the corresponding (4 km) reference simulation (line contours). It is evident that convection develops about half an hour earlier with the higher resolution, but the 4 km simulation nevertheless eventually produces a similar displacement of water. The difference between them (convection grid minus reference grid) is shown in the upper right panel, and the corresponding tendencies of the respective fields are shown in the 2 panels beneath the first two.

Figure 19 shows cloud water for both the 1 and 4 km simulations (upper left panel), the change in cloud water during the convection simulations (upper right panel), and the difference in cloud water change between the two (lower left panel). Also shown is the time-dependent total precipitation of each simulation and their difference (lower right panel).

Figure 20 shows corresponding results for potential temperatures, their differences, and the corresponding trends of both.

Figure 21 repeats the results shown in Figure 18 except that the reference simulation has a horizontal grid spacing of 20 km. We see that convection is effectively prevented in the reference simulation and that consequently a strong difference persists over time between the two cases. This helps to illustrate the importance of the reference simulation: It is necessary to take into account the possibility that the host model grid is fine enough for convection to develop, at least partially, although as we have seen it might be delayed. Determining whether and how much of the convection results to use must depend on these considerations.

Figure 22 shows the evolving cloud water field in a 1 km convection grid simulation and its 4 km reference grid simulation, along with their difference and tendency, for a different atmospheric profile. In this case, evolution and effects of convection are very similar. A reasonable decision to make for this case is to apply no convective tendency at all to the host model but to allow it to develop convection on its own.

Figure 23 shows the evolving ice-liquid potential temperature, which is the prognostic energy variable in OLAM, and the (diagnosed) potential temperature in the upper right for a 1 km convection grid simulation and a 20 km reference grid simulation, with differences between simulations in the 2 lower panels. The coincidence of color and line contours in the upper panels and near zero differences in the lower panels show that these two cases evolve almost identically, implying that convection is essentially nonexistent. The large temporal change in the upper left panel is due to large-scale removal of precipitation particles that exist in the host simulation. Correct interpretation leads to the conclusion that no separate convective tendency should be added.

While some of the decisions that must be made in the application of the convection grid/reference grid differences to the parent grid are straightforward, situations are encountered where decision making is more challenging. This is a complex topic that requires further study. Nevertheless, the approach developed here completely avoids the limitations of conventional convective parameterizations, replacing them with the same convective processes as used on convection resolving grids.

At this point, an obvious question is "What is the point of a variable-resolution parent grid if convection is going to be resolved everywhere by high resolution convection grid replicas? Isn't the superparameterization method at least as expensive as a uniformly high resolution global grid, and possibly more so?" The answer to this would be "yes" except that the following method is used to substantially boost efficiency. Not all of the samples from the colored regions in the left panel of Figure 17 are used for initializing and integrating convection and reference grid pairs. In fact, only a small fraction are. This is made possible by the fact that environmental conditions very often are quasi-uniform or slowly varying over large geographic regions, especially over the oceans. This means that many of the samples lead to nearly identical convection grid results and therefore need not all be integrated on convection grids. By sampling only a relatively few strategically selected locations, convective results are obtained for a variety of environments that reasonably well span parameter space. A principal component analysis is performed over both the samples and the convection results in order to construct interpolation functions that map convection results to sample regions that are not integrated on convection grids. For the TWP-ICE case, over a surrounding region exceeding 3 x 10^7 km^2 and sampled at only 200 locations, we found that the first 4 principal components represented about 90% of the total variation of accumulated convective precipitation. With each convection grid containing about 1000 grid columns, this represents a far lower computational cost than uniform coverage of the area with convection grid resolution. Furthermore, pre-screening of samples from the parent grid can eliminate stable profiles from further consideration and avoid unnecessary integration on convection grids.

Because this method is still considerably more expensive than conventional convective parameterizations, its use is not expected to be preferred in many cases. However, given the problems that we have demonstrated in transitioning from resolved to parameterized convection, this method may prove most useful in providing a buffer region or transition between fully resolved convection and conventional parameterization, keeping the latter very far away from the high resolution region that is normally where interest is focused.

Results from this work have been presented in conferences (AGU) and seminars (including Colorado State University, UC Davis, Cal Tech, and Chapman University). Two papers that present the problems of representing convection on variable-resolution grids and detail our convective parameterization approach are in preparation for submission to Monthly Weather Review.

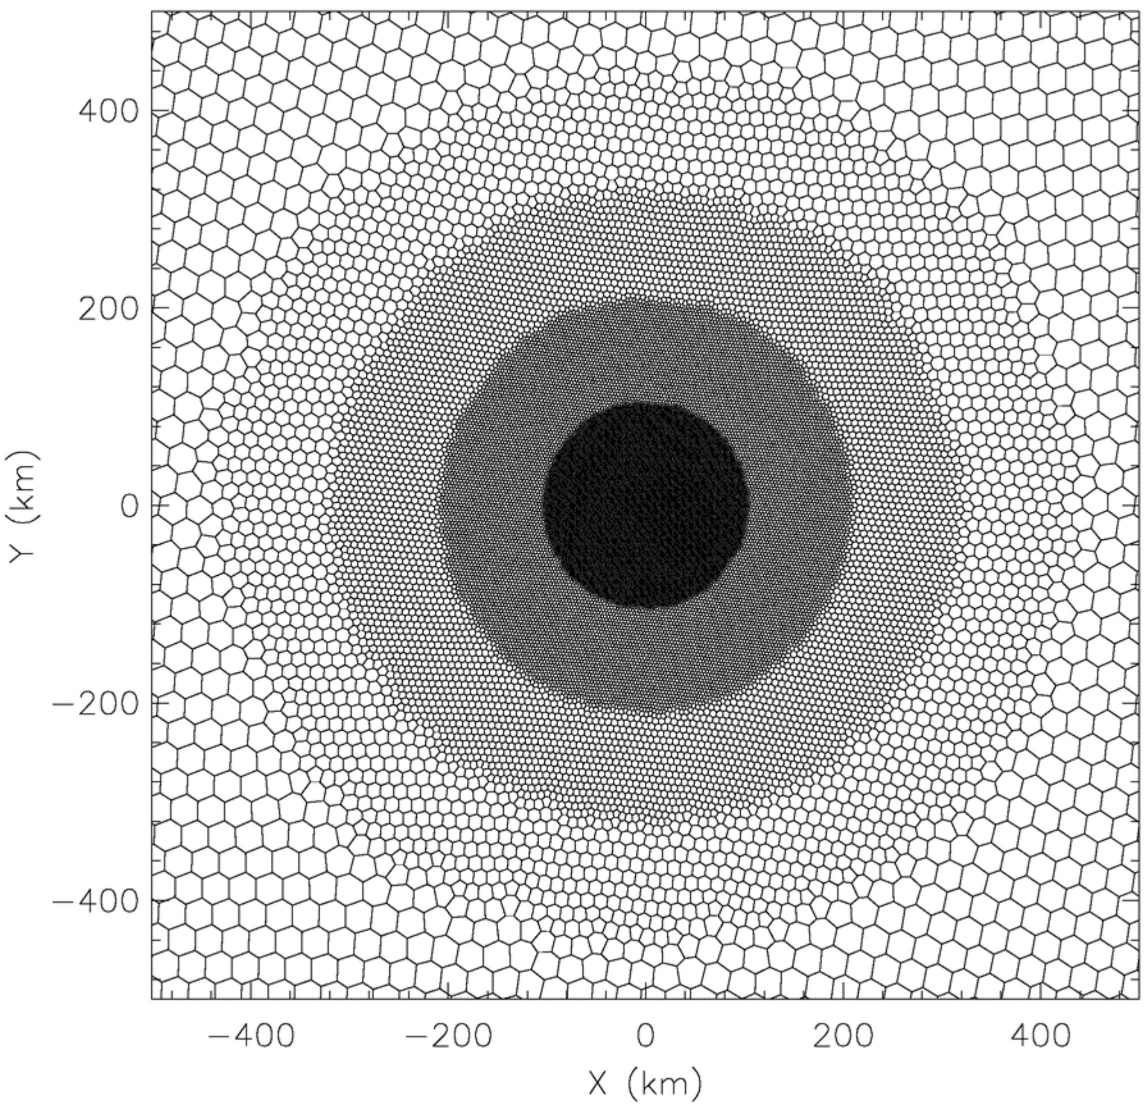

Figure 1. Variable-resolution grid used in idealized convection simulations.

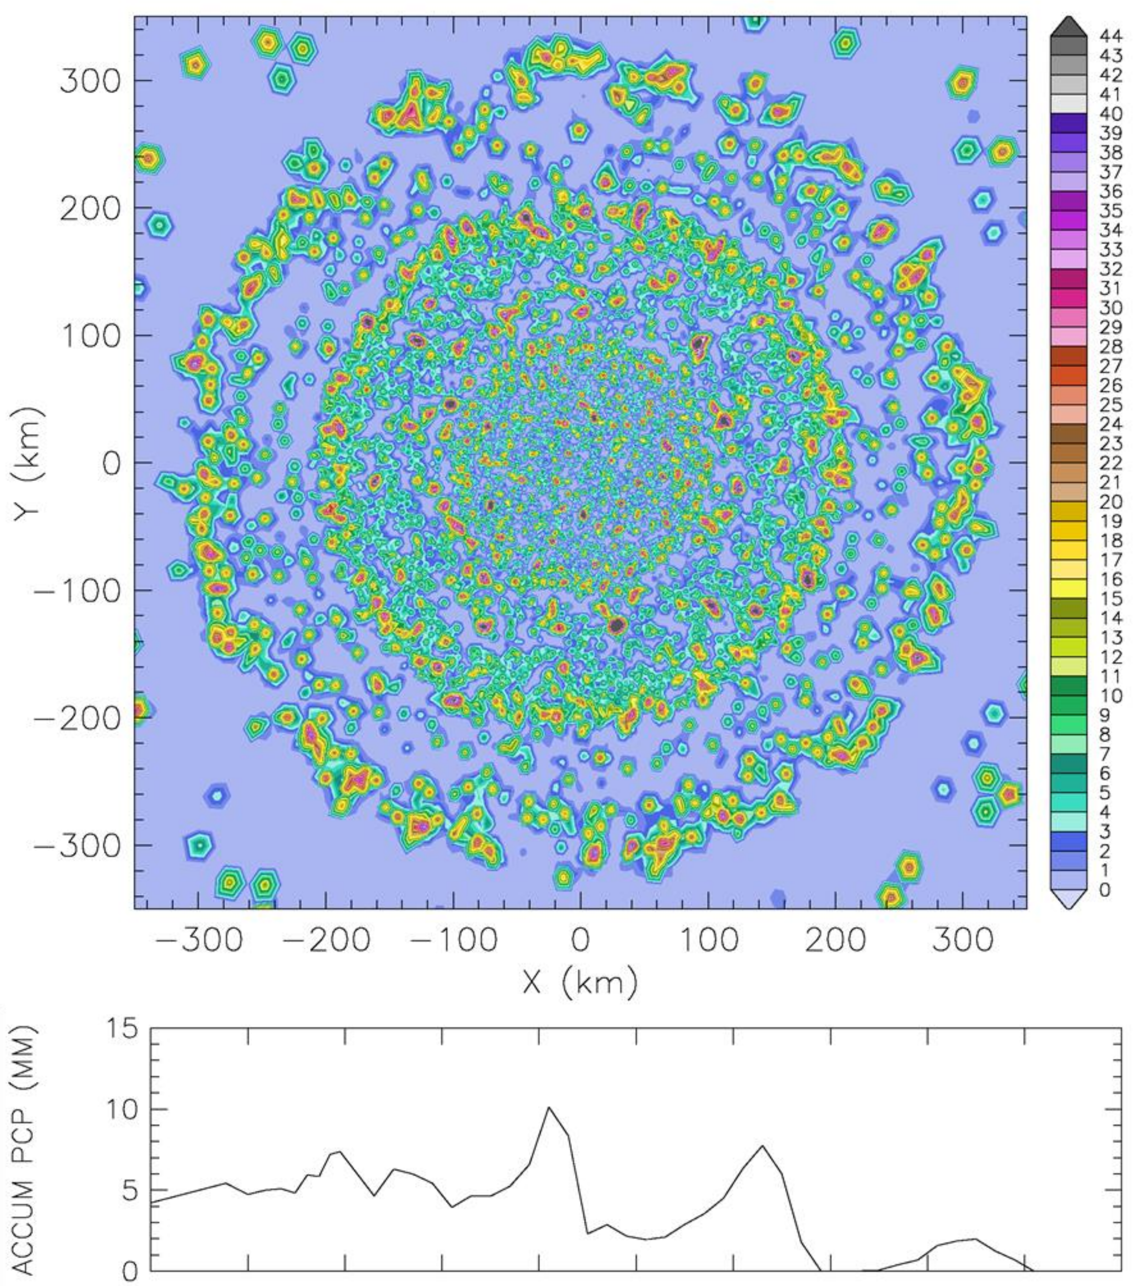

Figure 2. Horizontal distribution and radial profile of accumulated precipitation in idealized convection experiment with no convective parameterization.

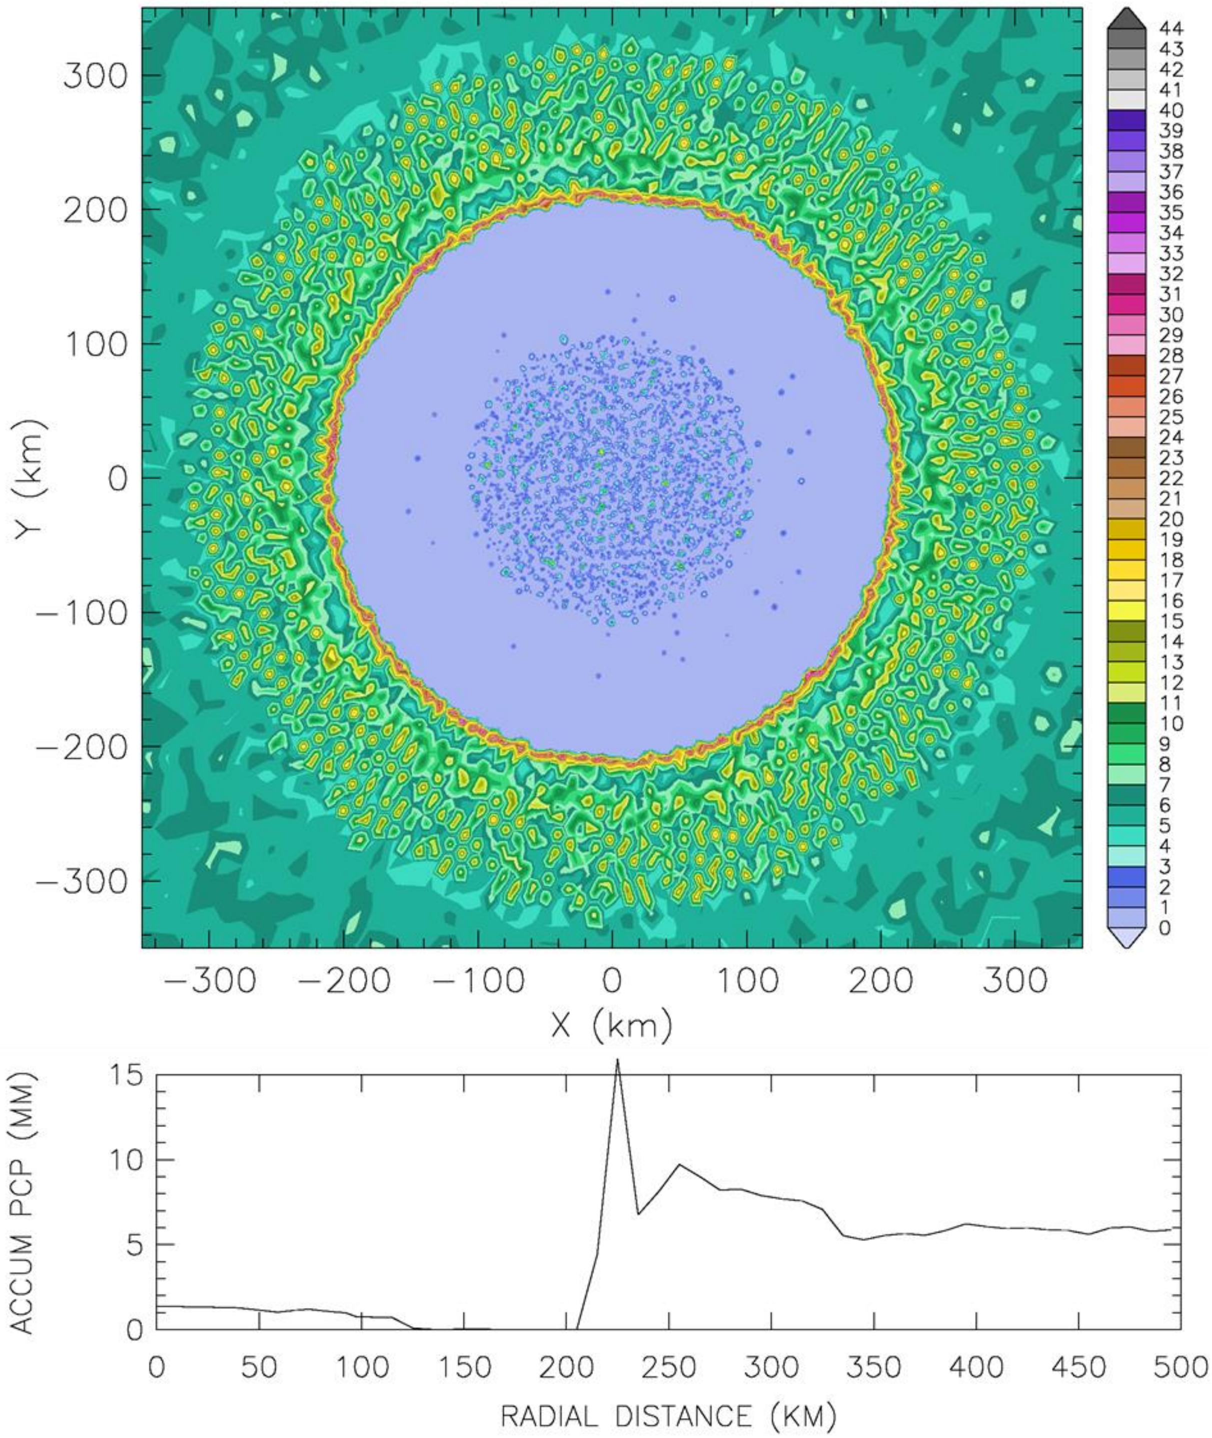

Figure 3. Horizontal distribution and radial profile of accumulated precipitation in idealized convection experiment with convective parameterization excluded only in 1.5 km and 3.0 km grid cells.

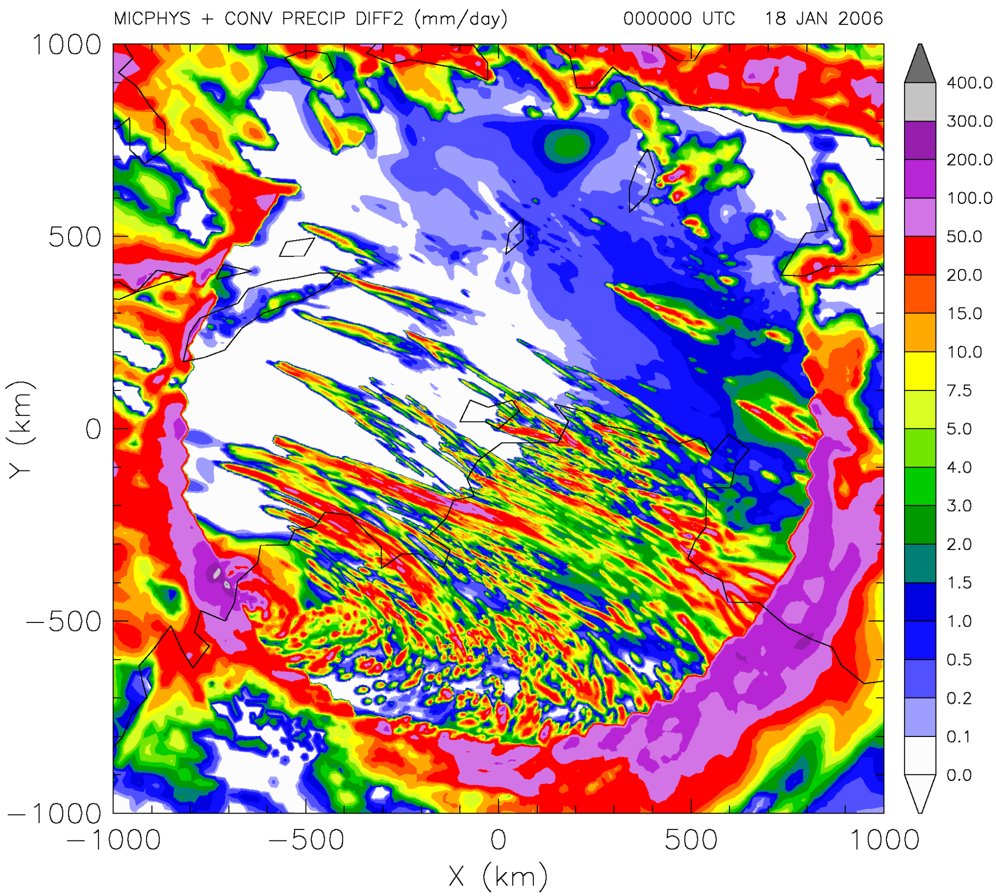

Figure 4. Combined microphysics and convective precipitation with Kain-Fritsch parameterization.

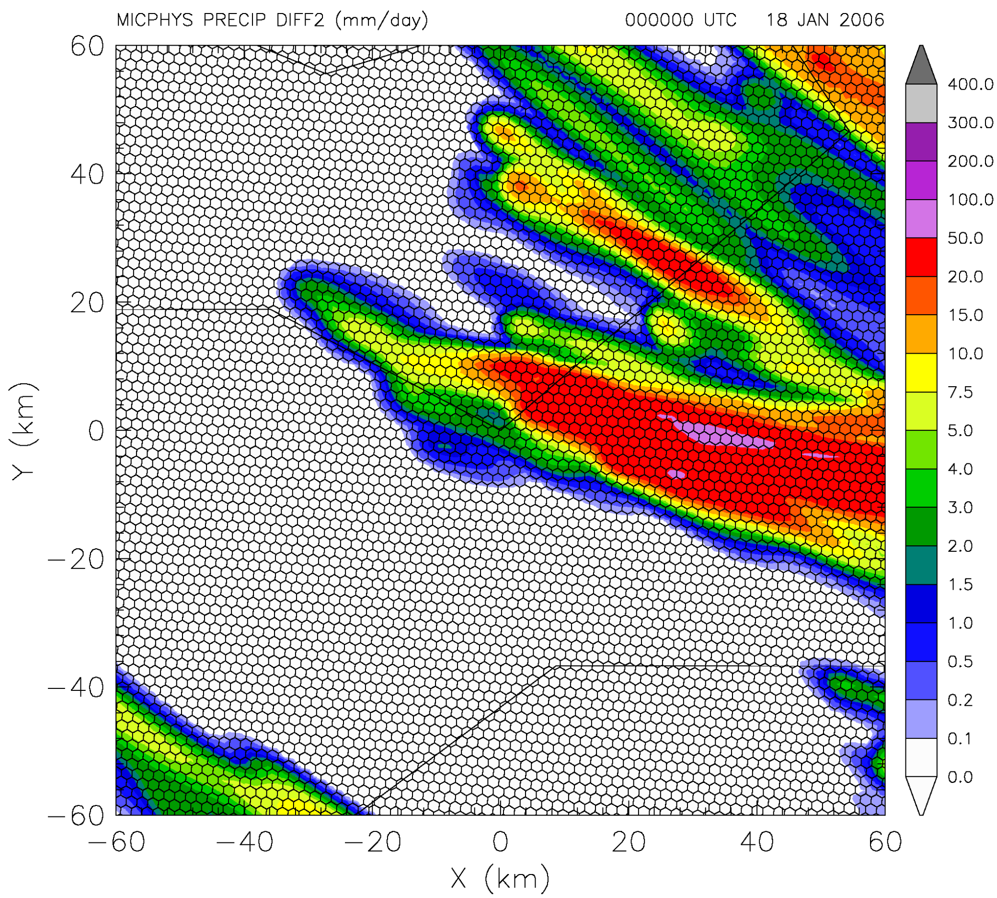

Figure 5. Microphysics precipitation with Kain-Fritsch parameterization used outside region shown.

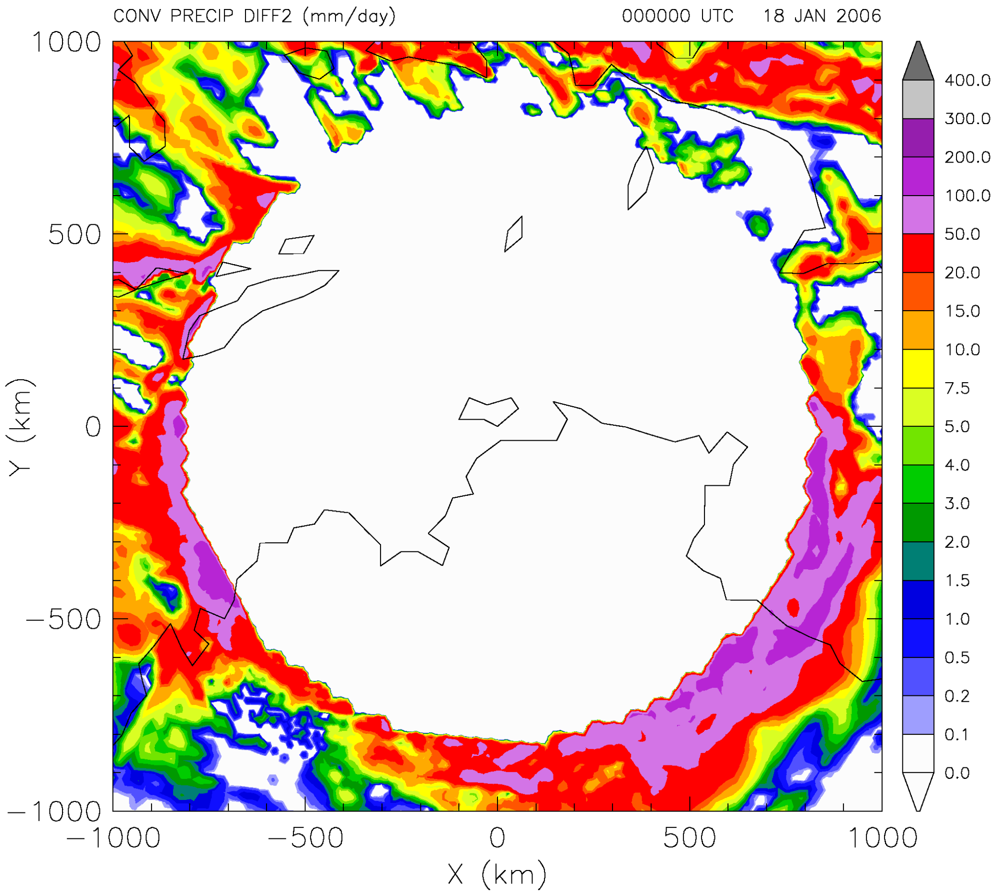

Figure 6. Convective-only component of precipitation with Kain-Fritsch parameterization.

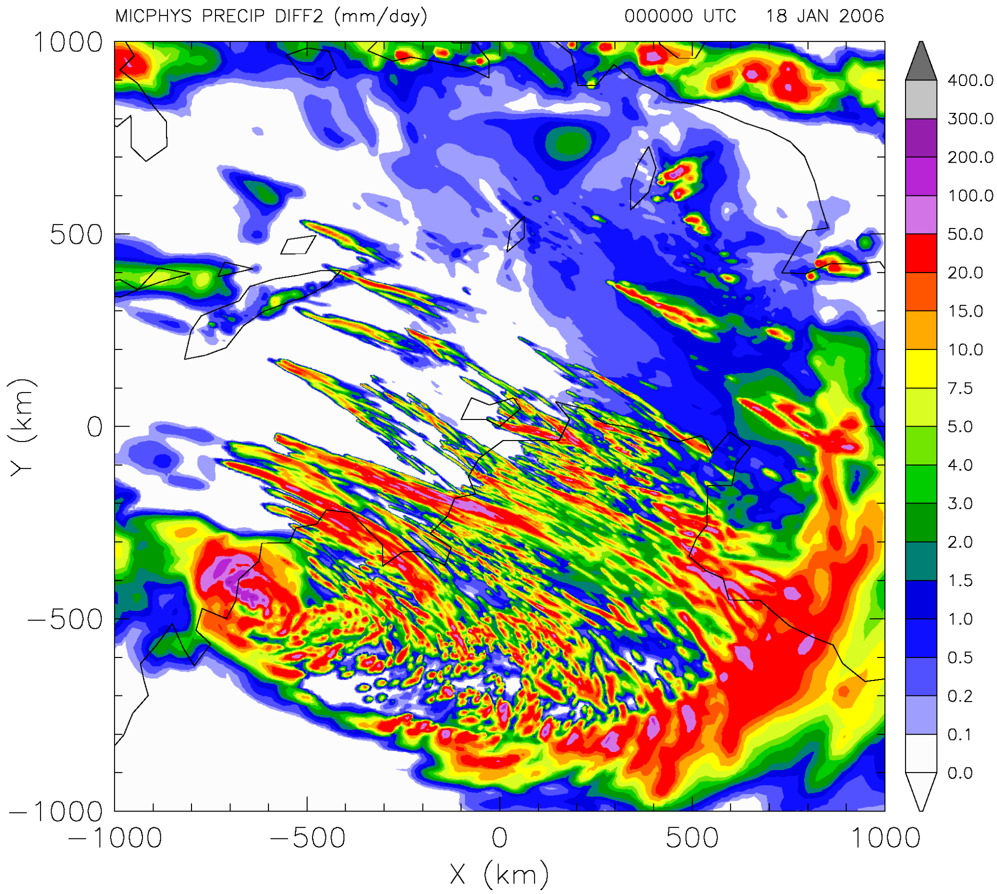

Figure 4. Microphysics-only component of precipitation with Kain-Fritsch parameterization.

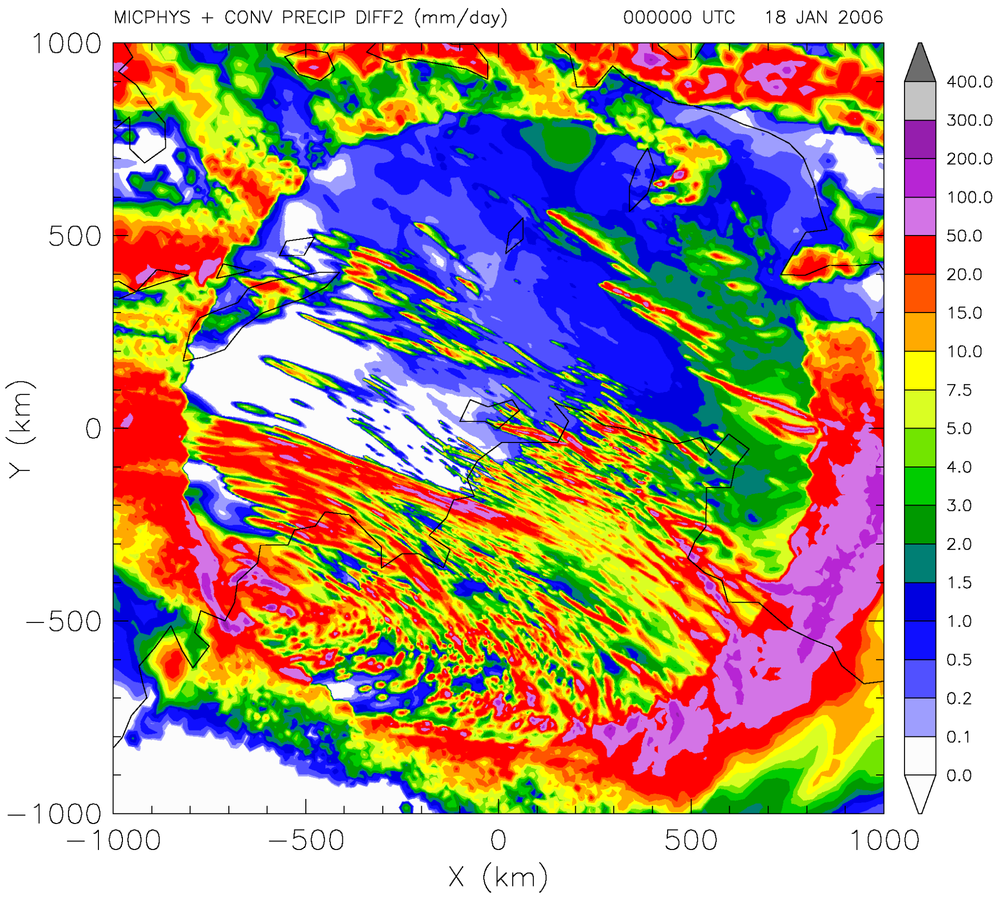

Figure 8. Combined microphysics and convective precipitation with Tiedtke parameterization.

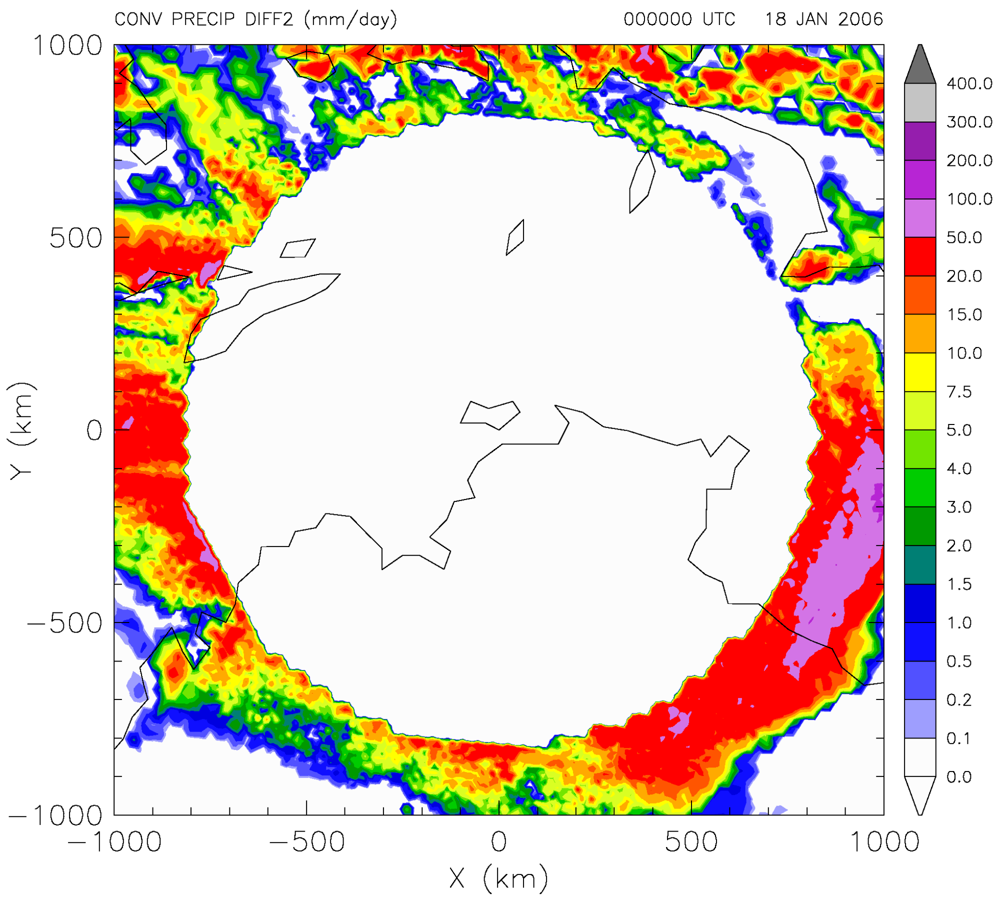

Figure 9. Convective-only component of precipitation with Tiedtke parameterization.

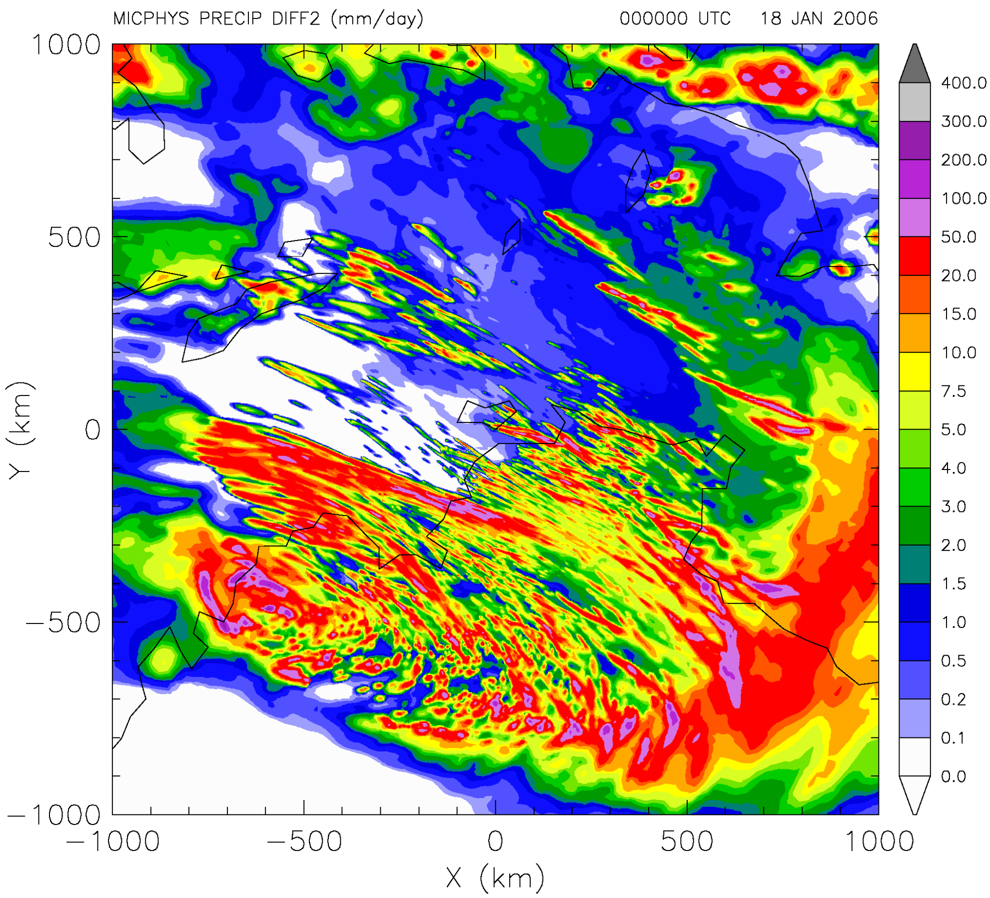

Figure 10. Microphysics-only component of precipitation with Tiedtke parameterization.

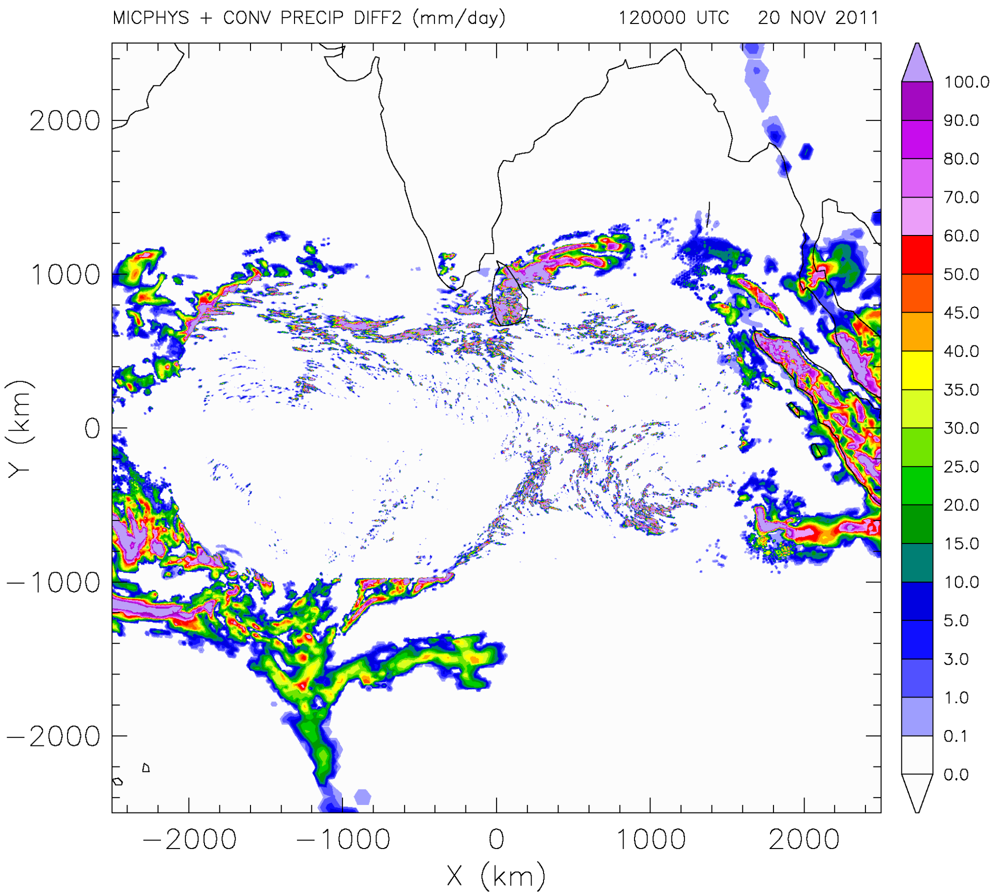

Figure 11. Combined resolved and parameterized precipitation rate in MJO simulation with Kain-Fritsch parameterization.

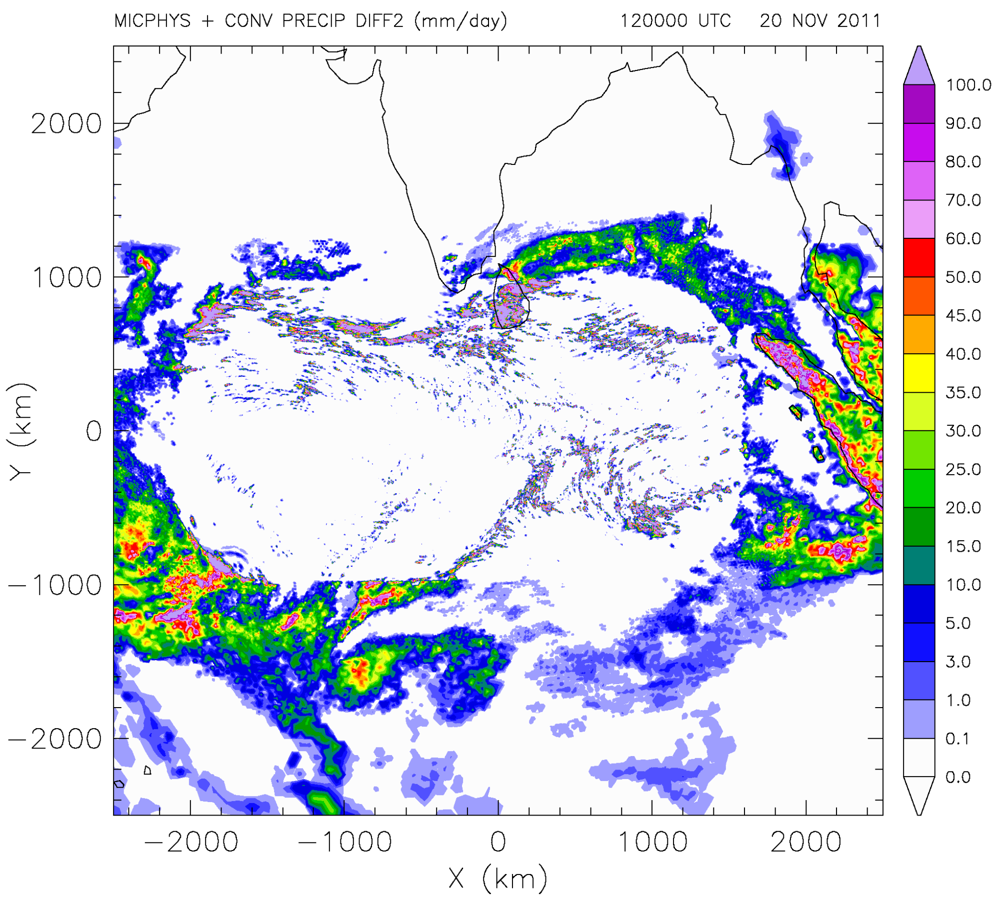

Figure 12. As in Figure 11 but with Grell parameterization.

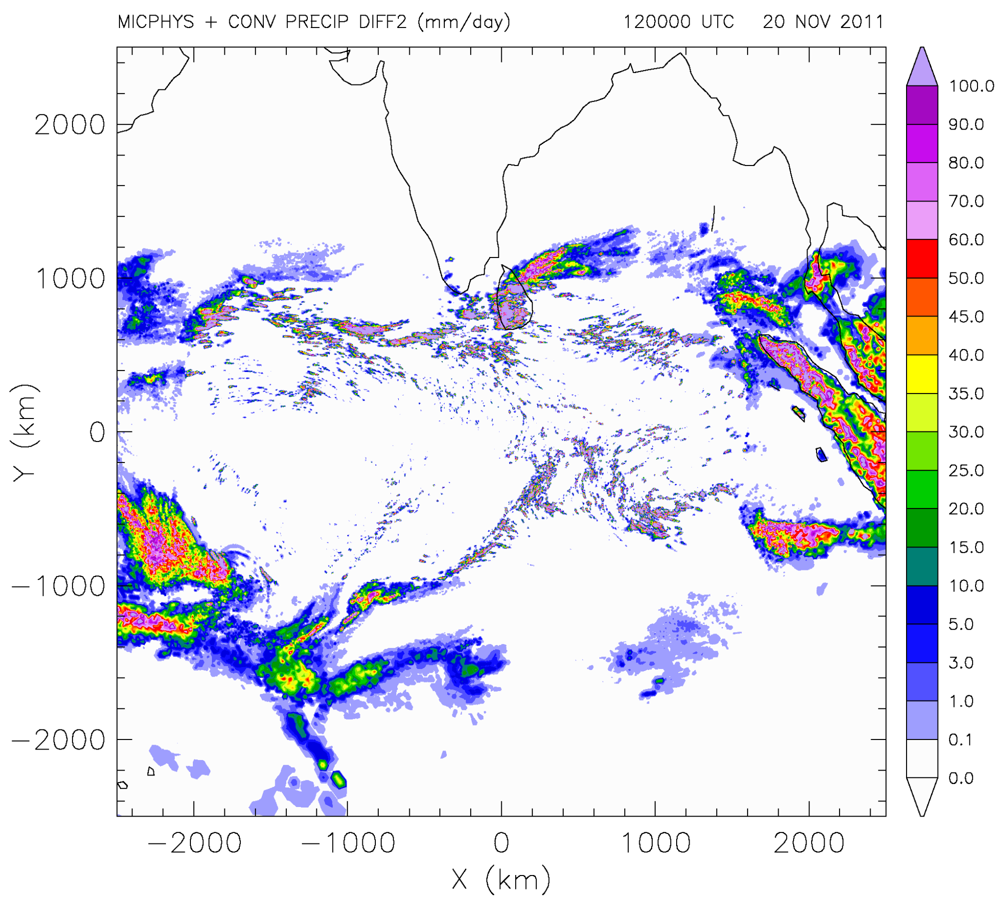

Figure 13. As in Figure 11 but with Tiedtke parameterization.

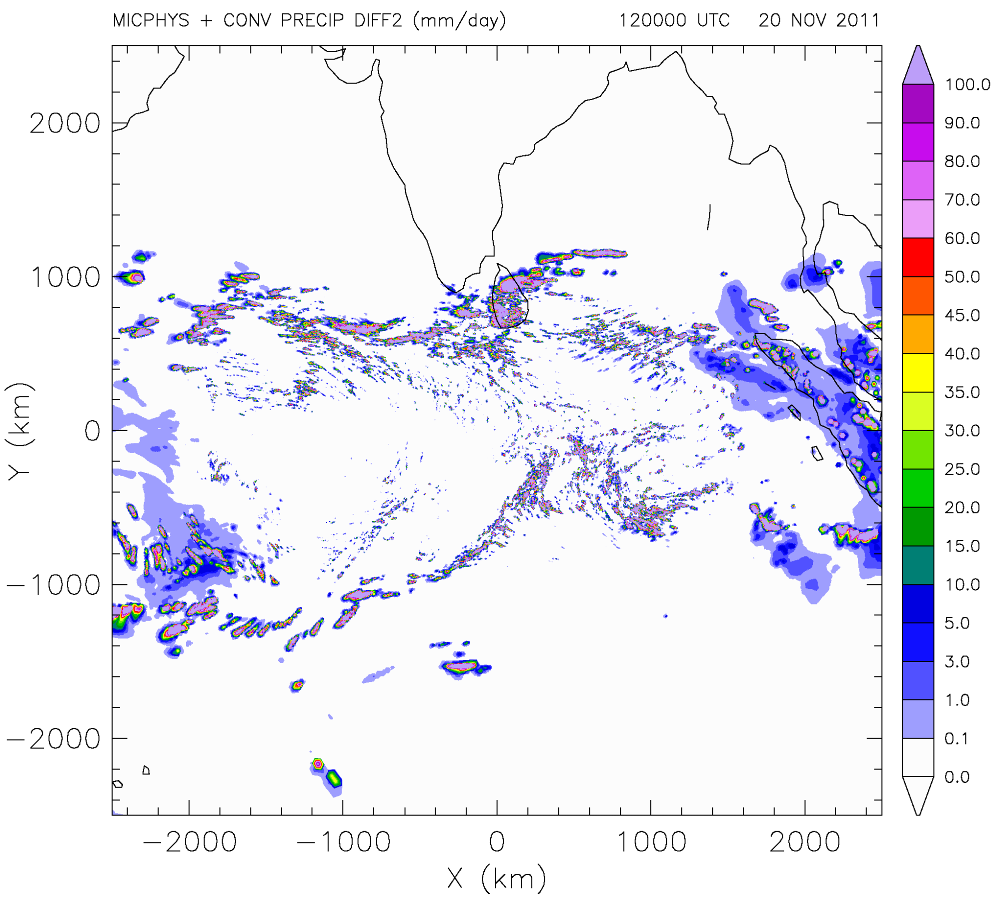

Figure 14. As in Figure 11 but with no convective parameterization applied anywhere.

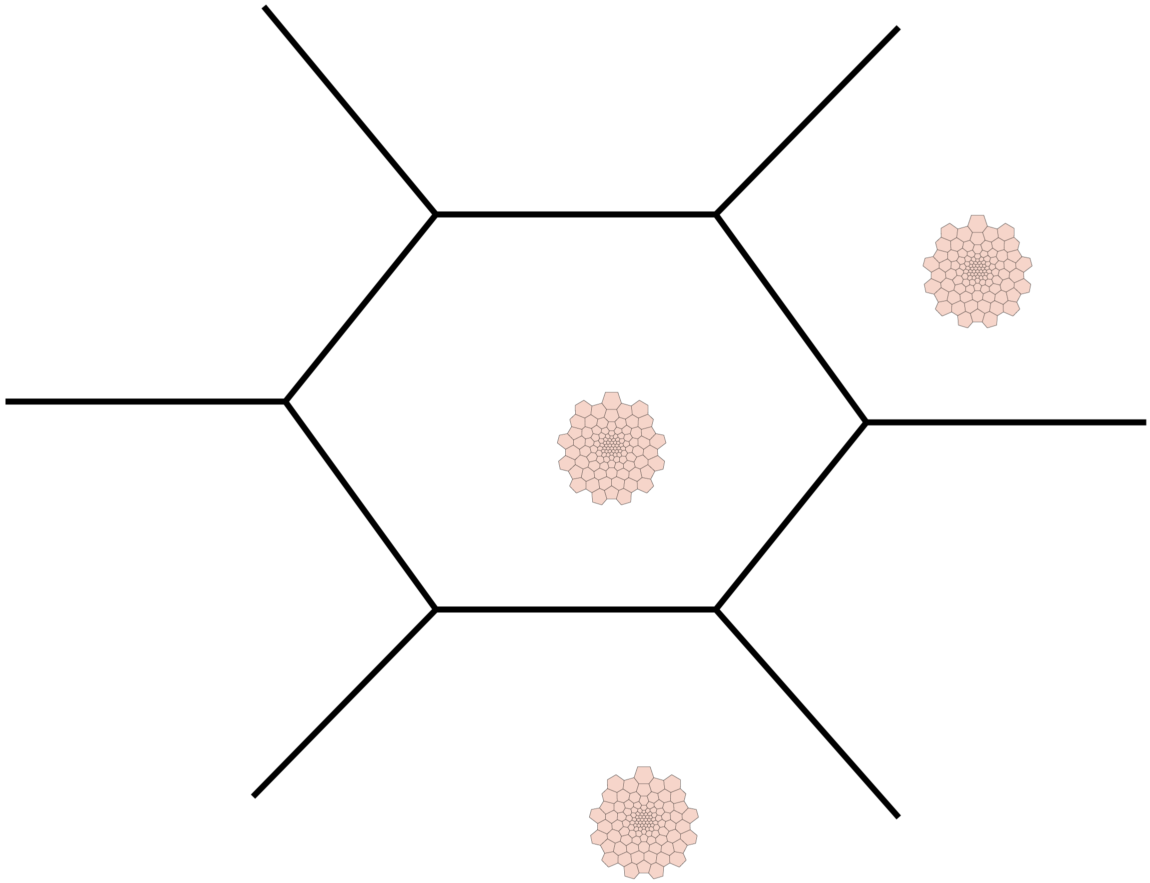

Figure 15. Proposed configuration for embedded “superparameterization” grid on coarse parent grid.

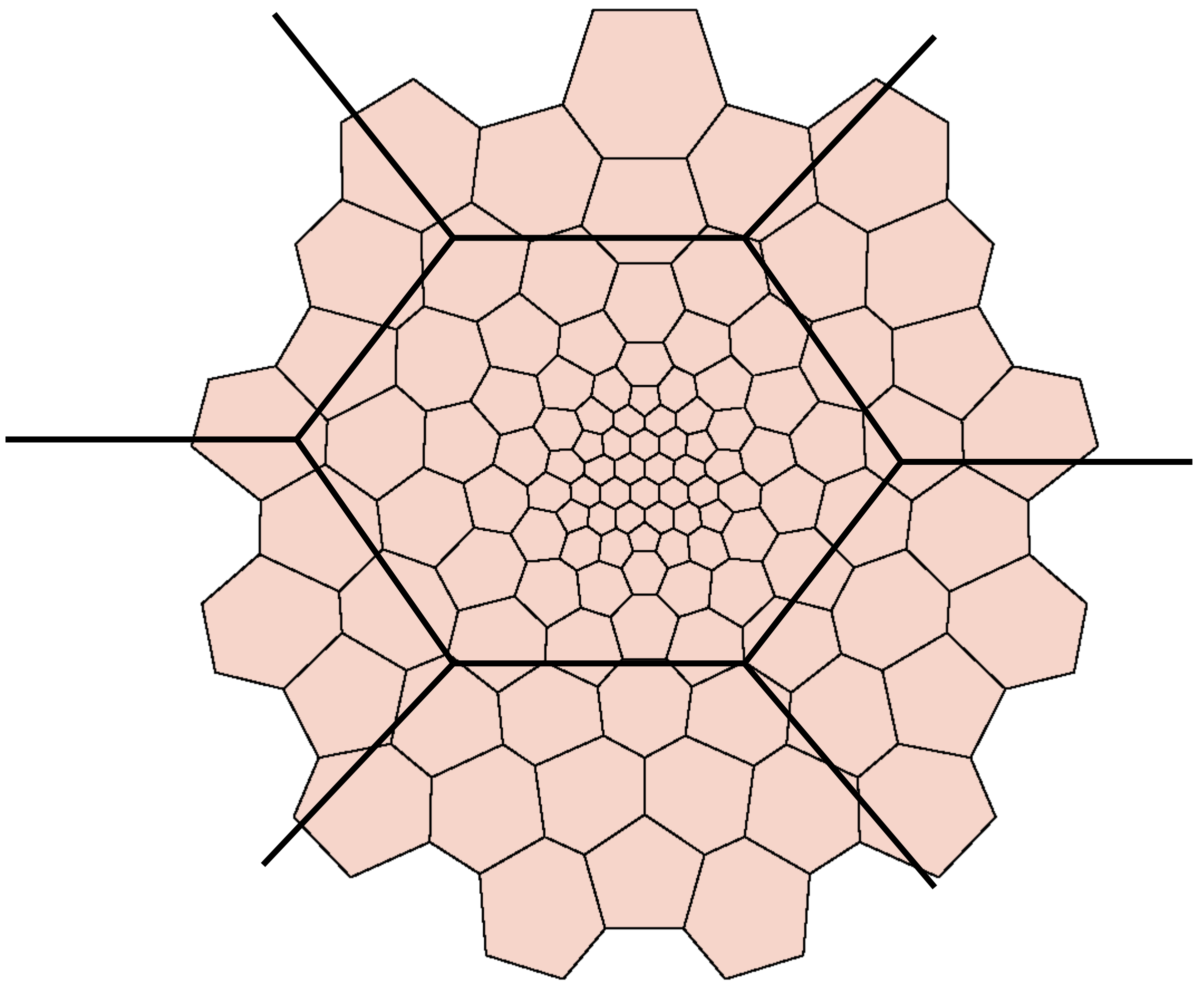

Figure 16. Proposed convection grid on medium parent grid.

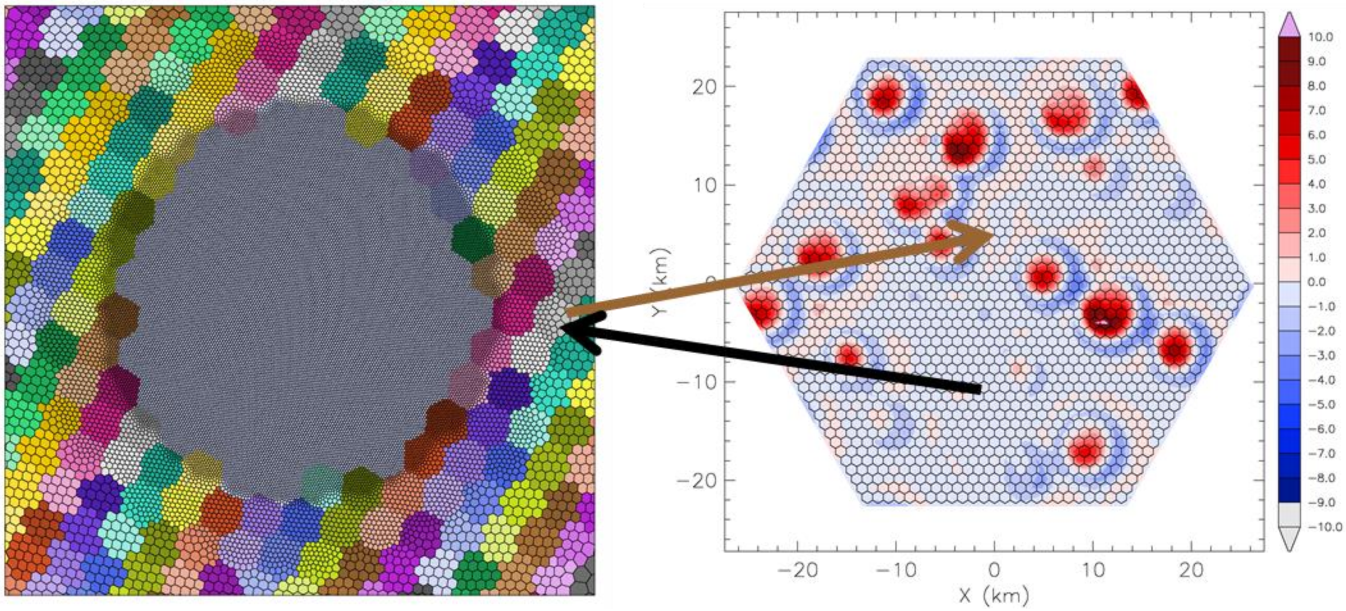

Figure 17. Schematic of variable-resolution parent grid (left), convection grid (right), and information transfers between them. Simulated convective vertical velocity is represented on the convection grid.

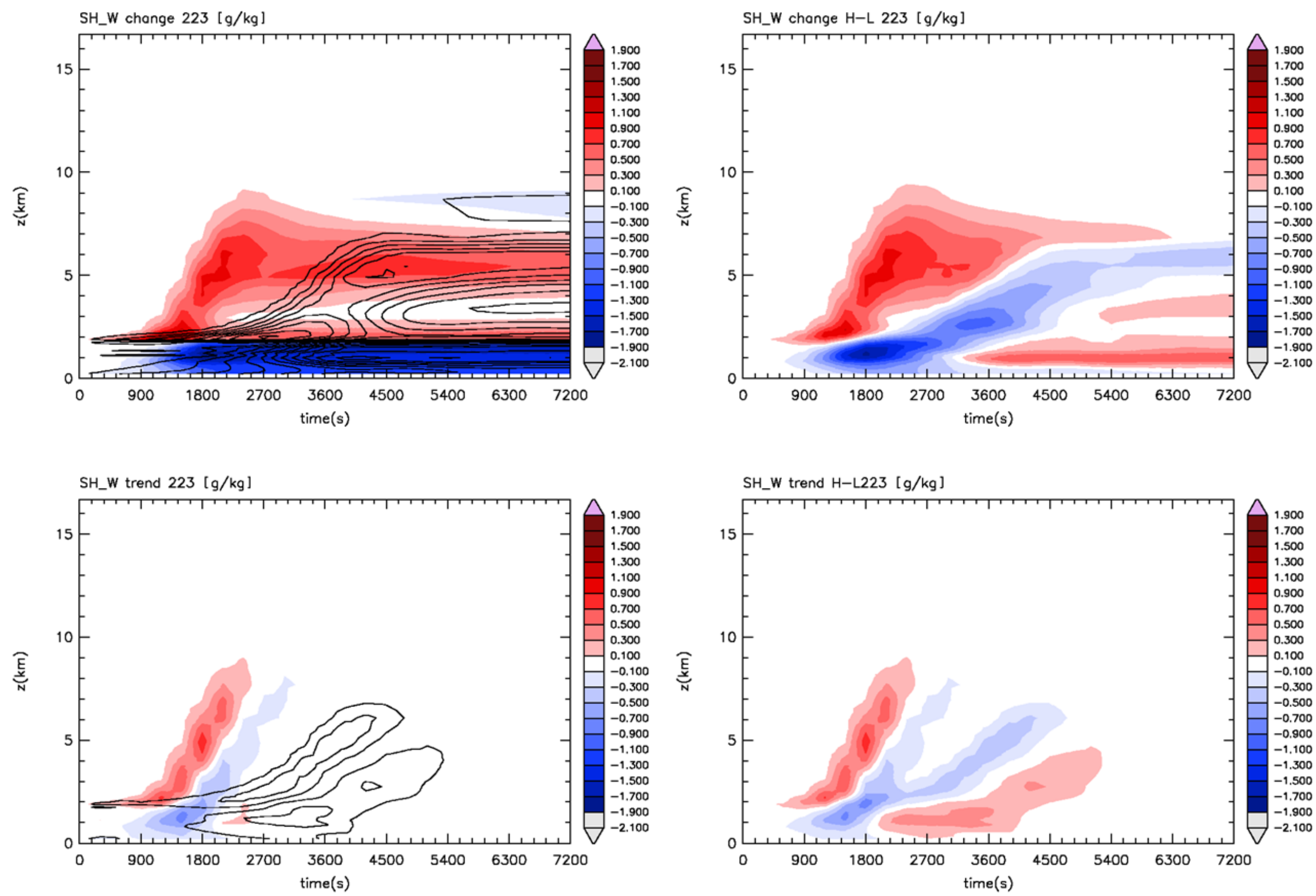

Figure 18. Time-height plots of horizontally-averaged total water change from convective grid (colors) and reference grid (contours) simulations and their differences and tendencies. The convective and reference grids have 1 km and 4 km spacing, respectively.

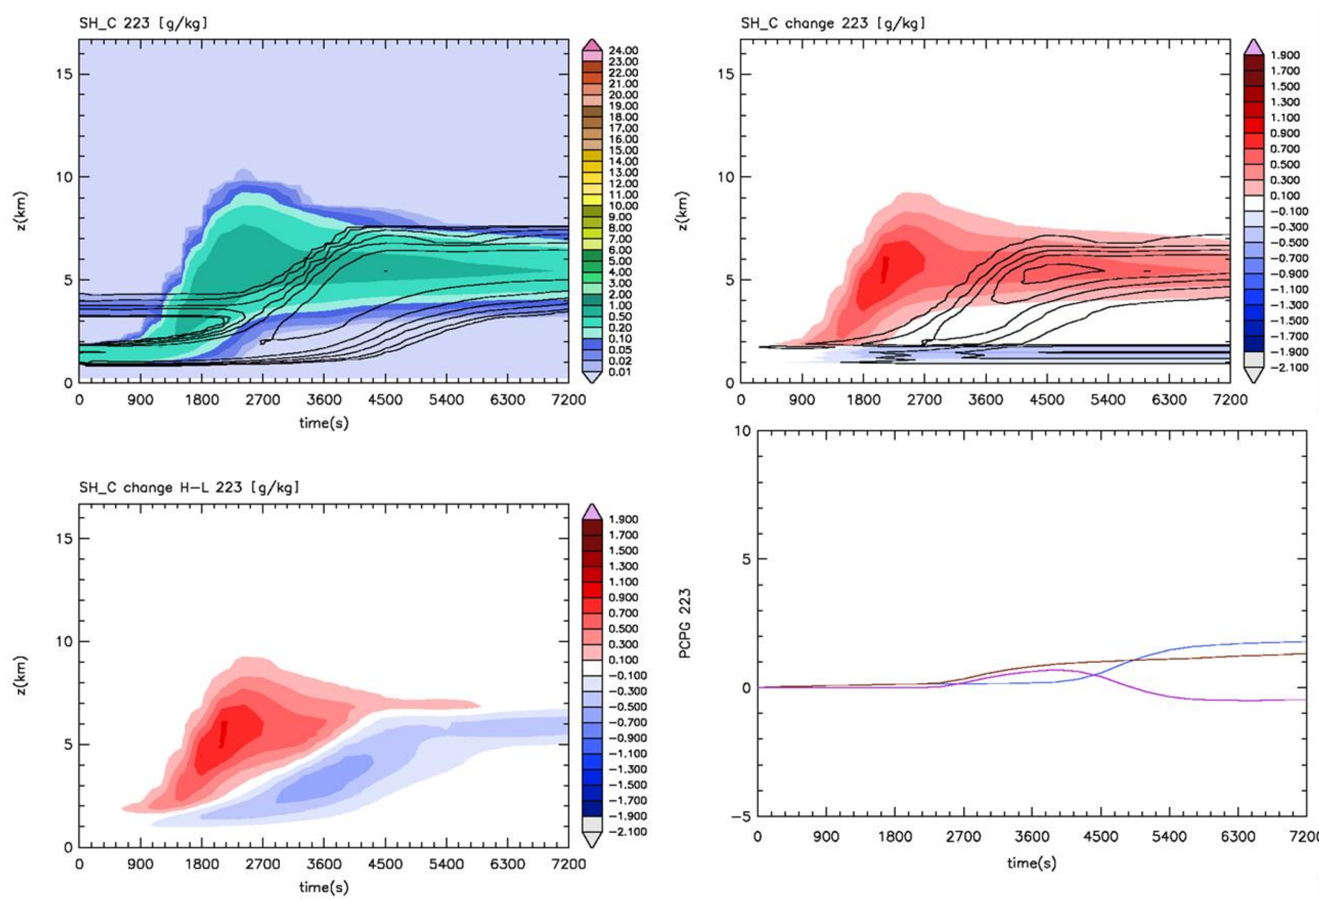

Figure 19. As in Figure 18 but for cloud water.

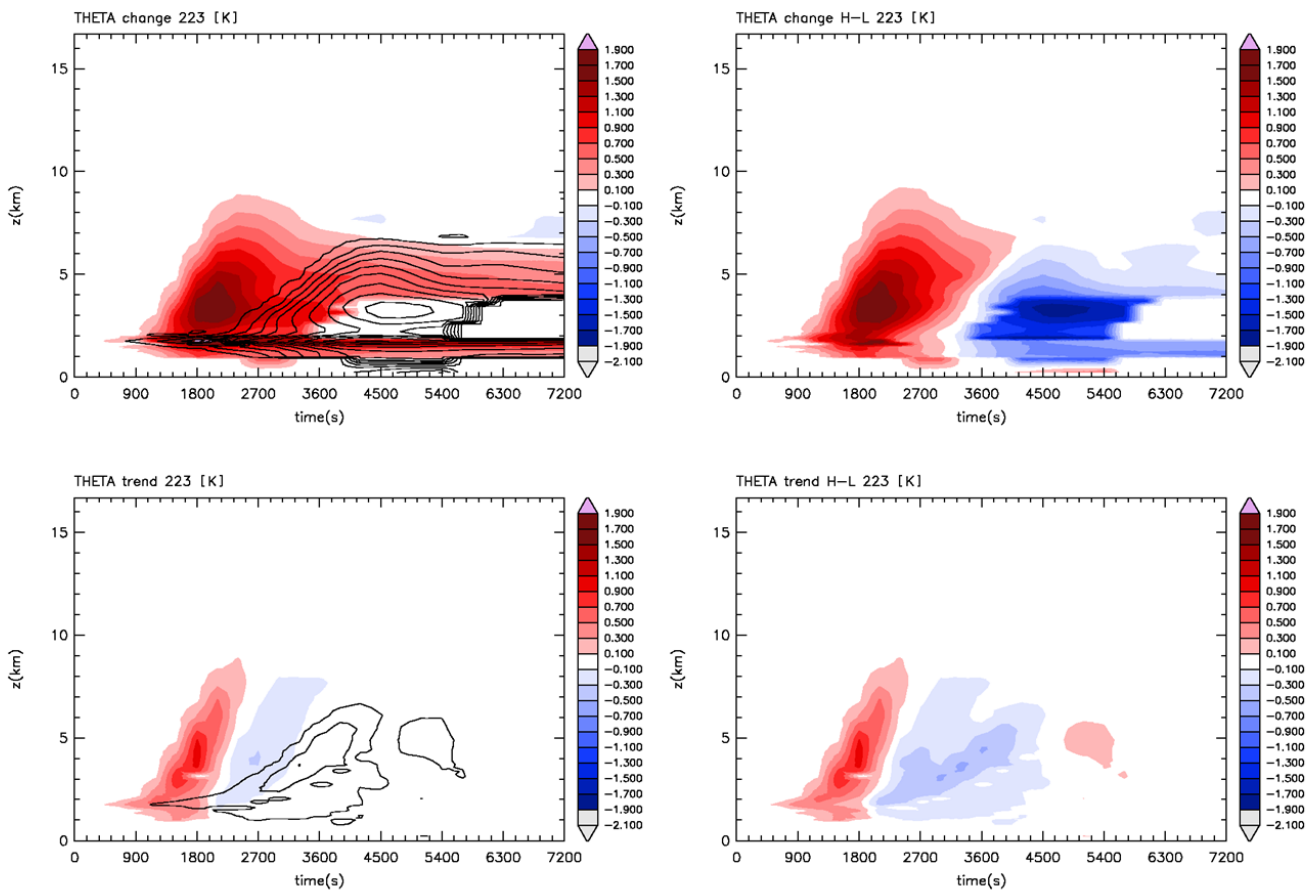

Figure 20. As in Figure 18 but for potential temperature.

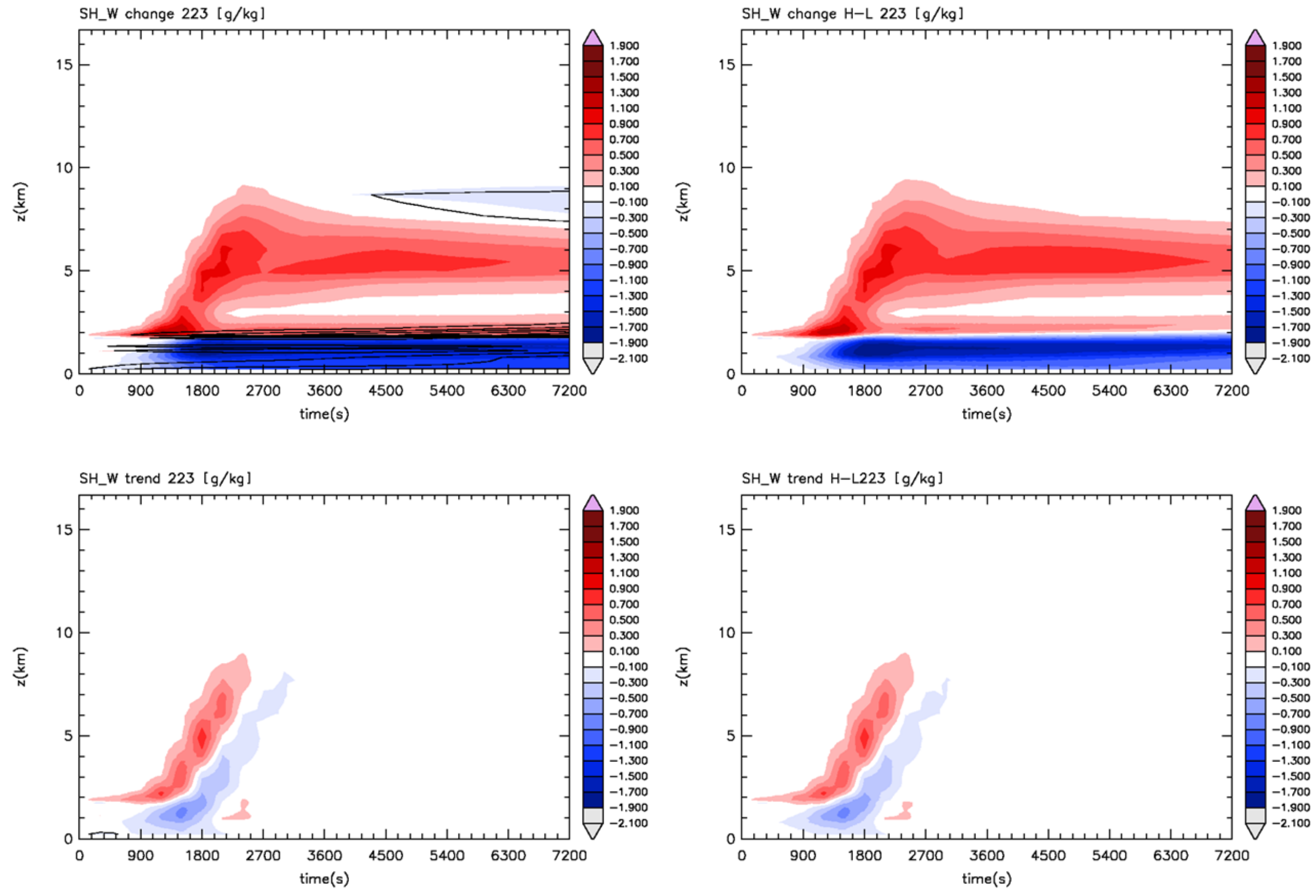

Figure 21. As in Figure 18 except for reference grid cell size of 20 km.

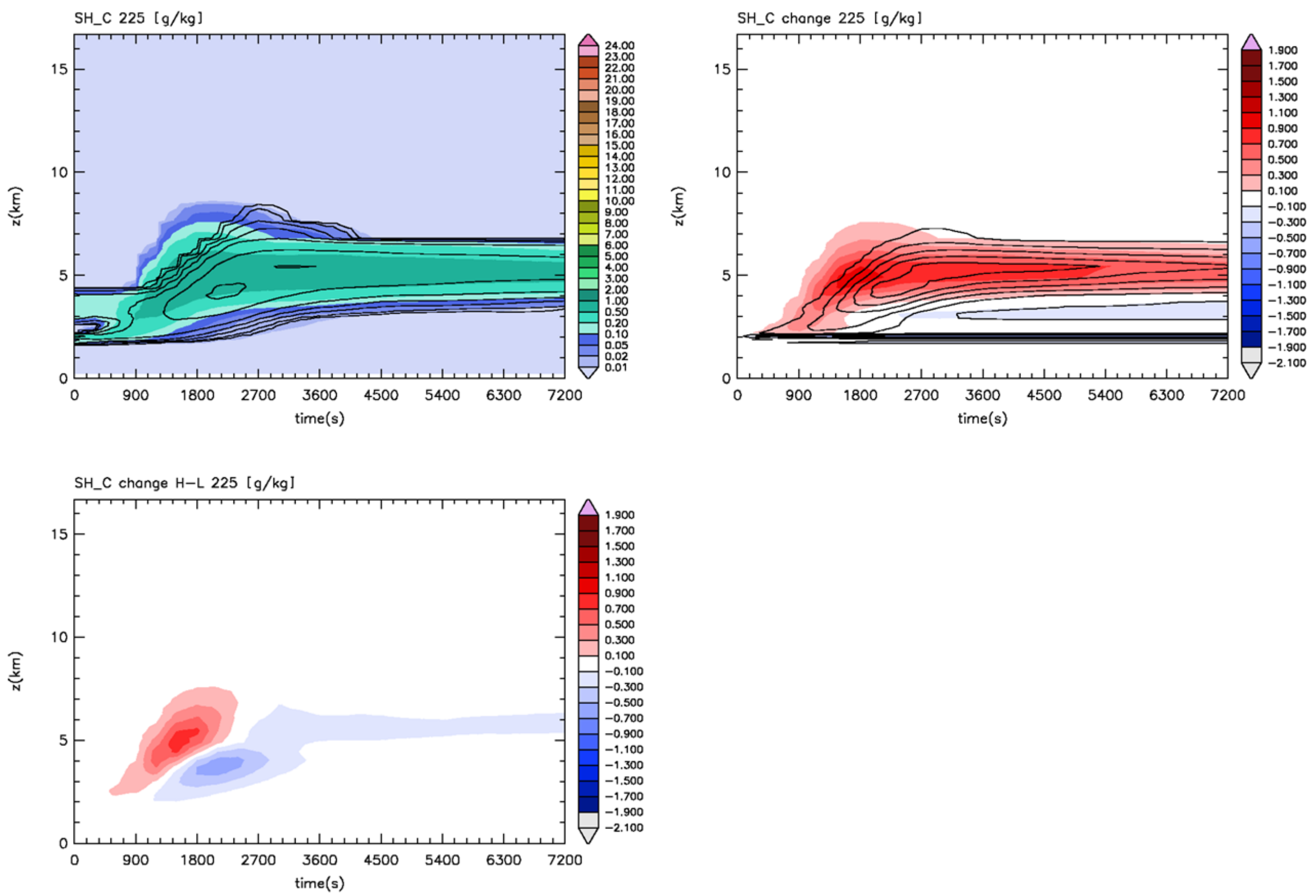

Figure 22. As in Figure 18 but for a different environment.

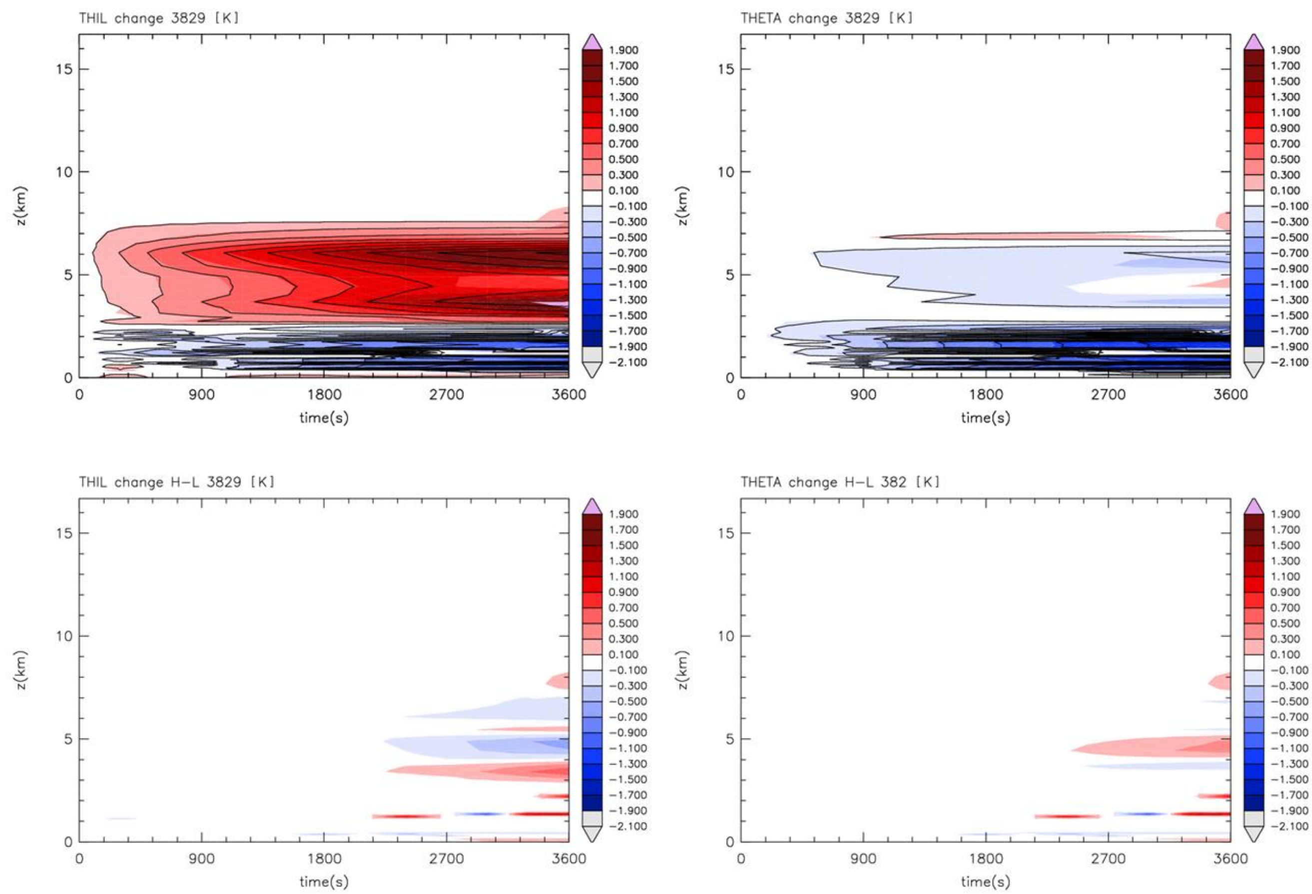

Figure 23. As in Figure 21 but for a different environment and for potential and ice-liquid potential temperatures.

Tiedtke, M.,1989: A comprehensive mass flux scheme for cumulus parameterization in large-scale models. *Mon. Wea. Rev*., **117**, 1779-1800.

Grell, G. A., and S. R. Freitas, 2014: A scale and aerosol aware stochastic convective parameterization for weather and air quality modeling. *Atmos. Chem. Phys.*, **14**, 5233–5250.