

# A New Approach for Coupled Regional Climate Modeling Using More than 10,000 Cores

Marcus Thatcher, John McGregor, Martin Dix, and Jack Katzfey

CSIRO Ocean and Atmosphere Flagship, Melbourne, Australia  
{Marcus.Thatcher, John.McGregor,  
Martin.Dix, Jack.Katzfey}@csiro.au

**Abstract.** This paper describes an alternative method for coupling atmosphere-ocean regional climate models that communicates momentum, radiation, heat and moisture fluxes between the atmosphere and ocean every time-step, while scaling to more than 10,000 cores. The approach is based on the reversibly staggered grid, which possesses excellent dispersive properties for modeling the geophysical fluid dynamics of both the atmosphere and the ocean. Since a common reversibly staggered grid can be used for both atmosphere and ocean models, we can eliminate the coupling overhead associated with message passing and improve simulation timings. We have constructed a prototype of a reversibly staggered, atmosphere-ocean coupled regional climate model based on the Conformal Cubic Atmospheric Model, which employs a global variable resolution cube-based grid to model the regional climate without lateral boundary conditions. With some optimization, the single precision, semi-implicit, semi-Lagrangian prototype model achieved 5 simulation years per day at a global 13 km resolution using 13,824 cores. This result is competitive with state-of-the-art Global Climate Models that can use more than 100,000 cores for comparable timings, making CCAM well suited for regional modeling.

**Keywords:** geophysical fluid dynamics, MPI, regional climate modeling.

## 1 Introduction

Regional Climate Models (RCMs) are used to study how changes in global climate can affect meteorology at fine spatial scales. Such models are commonly used to dynamically downscale climate change projections from Global Climate Models (GCMs) to regional and local spatial scales appropriate for impact studies. There are four main components in the construction of an RCM:

- The dynamical core that simulates the atmosphere as a geophysical fluid.
- Parameterization of atmospheric physical processes, such as solar radiation, clouds, boundary layer turbulent mixing and aerosols.
- Parameterization of surface interactions at fine spatial scales, such as mountain ranges, forests, deserts, ice and oceans, which significantly influence the regional meteorology.

- An assimilation or nudging system to ingest large scale changes in atmospheric circulation and sea surface temperatures, as predicted by GCMs for different global warming scenarios.

The RCM integrates these four components to produce a consistent representation of the meteorology at regional spatial scales (i.e., 1 to 100 km). RCMs face a difficult computational challenge as they are required to simulate atmospheric processes that may last seconds to minutes, whilst integrating the simulations for time periods greater than one hundred years. Computational load-balancing issues arise as the various simulated atmospheric processes are not uniformly distributed over the simulation grid. Furthermore, assimilating climate data also requires the RCM to be reading and writing atmospheric datasets while running on hundreds to thousands of cores.

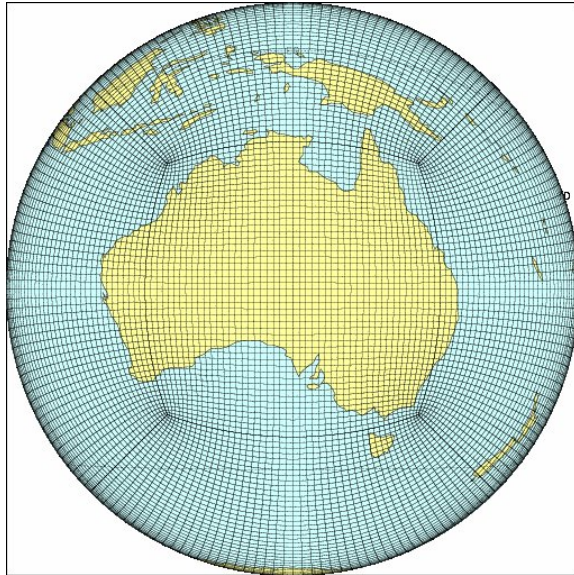
Traditionally the ocean is simply represented in an RCM by interpolating Sea Surface Temperatures (SSTs) from a GCM dataset. This approach neglects any feedback between the atmosphere and ocean and smoothes out variability arising from ocean surface eddies. Most attempts to construct an atmosphere-ocean coupled RCM are extensions of the GCM approach, where a ‘coupler’ is used to manage communication between atmosphere and ocean processors. Since the atmosphere and ocean usually use different grids, complex communication patterns can arise, resulting in considerable inter-node message passing that slows down the simulation. To minimize the impact of this communication bottleneck, couplers usually only exchange information between the atmosphere and ocean models every few simulation hours.

This paper proposes an alternative approach to developing an atmosphere-ocean coupled RCM that has good wave dispersion properties and allows frequent coupling without increasing simulation time. Our approach is based on a reversibly staggered grid [1], which was originally developed for the Conformal Cubic Atmospheric Model (CCAM) [2]. CCAM is a global atmospheric model based on a variable-resolution conformal-cubic grid so that it can simulate regional scales without lateral boundary conditions. The reversibly staggered grid handles horizontal velocity components and was introduced into CCAM to provide good gravity wave dispersion properties. However, the reversibly staggered grid has also been found to possess excellent dispersive properties when simulating the ocean [1], making it suitable as a universal grid for an atmosphere-ocean coupled model. Using a common grid for the atmosphere and ocean avoids the communication bottleneck associated with traditional coupling methods and improves the simulation timings. We have constructed such a reversibly staggered ocean model within CCAM to assess its potential for coupled regional climate modeling. After some optimization, the coupled model prototype was found to exhibit strong scaling beyond 10,000 cores while exchanging fluxes between the atmosphere and ocean models every time-step (e.g., 3 minutes).

CCAM and its variable-resolution conformal-cubic grid are described in Section 2. Section 3 outlines the CCAM software architecture and parallel programming issues. Section 4 explains how the regional ocean model was integrated into CCAM using the reversibly staggered grid. Section 5 describes the performance of the prototype atmosphere-ocean coupled regional model, including timings for different size grids and scaling issues. Conclusions and future work plans are presented in Section 6.

## 2 The Conformal Cubic Atmospheric Model

The Conformal Cubic Atmospheric Model (CCAM) is based on the conformal cubic projection [2] (see Fig. 1). CCAM was the first three-dimensional atmospheric model to use a cube-based grid, which provides significant computational advantages due to avoiding coordinate singularities while still locally resembling a Cartesian grid. The CCAM grid can be focused over a region of interest using the Schmidt transformation [3] and was part of the first generation of variable-resolution global atmospheric models [4]. By using a global, variable-resolution grid, CCAM avoids lateral boundary conditions and facilitates coupling between the global and regional spatial scales on the same grid. CCAM currently employs a non-hydrostatic, semi-implicit, semi-Lagrangian dynamical core that is efficient for regional climate modeling due to the relatively large time-step allowed by semi-Lagrangian methods. CCAM also has an extensive series of physical parameterizations to describe the behavior of the land-surface, aerosols, convection, cloud microphysics, radiation, gravity wave drag and boundary layer turbulence.

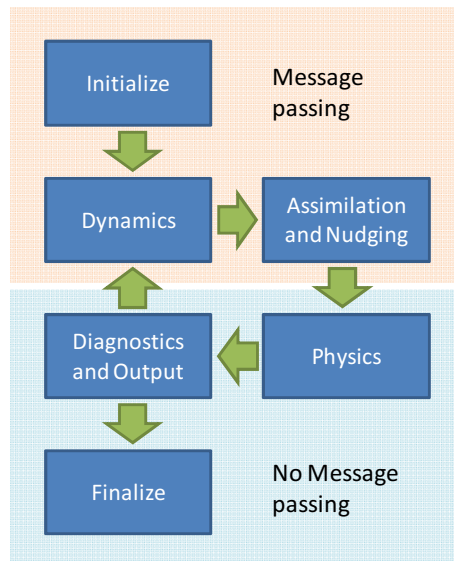


**Fig. 1.** Example of a stretched conformal cubic grid employed by CCAM, focused over Australia, using a Schmidt factor of 3 (i.e., the front panel of the cube has shrunk by a factor of 3)

For assimilation, CCAM employs a convolution-based scale-selective filter [6], that nudges the atmosphere at large length scales (e.g., greater than 3,000 km) towards a specified GCM projection of future climate. We base our filter on a convolution as it can be easily applied to CCAM's variable-resolution cubic grid and can accommodate irregular coastlines when assimilating ocean data in the atmosphere-ocean coupled RCM.

### 3 Software Design

The software components of CCAM are represented in Fig. 2, where initial conditions and configuration data are read by the initialize component at the start of the simulation. Prognostic variables are transported around the model by the dynamics component. Assimilation perturbs CCAM towards a global GCM dataset. Additional or unresolved processes that need to be parameterized are included under physics. Prognostic and diagnostic data are periodically written to a history file from the output routines. Lastly, the finalize component stores all prognostic variables at the end of each simulation month in a restart file. The restart file is read as initial conditions for the next simulation month when CCAM is restarted with new configuration data. Each component in Fig. 2 can be further divided into atmosphere, ocean and land-surface sub-components.



**Fig. 2.** Schematic diagram of the CCAM software components, where the dynamics, assimilation, physics and output routines cycle every model time-step. Message passing between processors is only required in the upper half of the plot, as explained in the text.

CCAM is written in Fortran and uses the Message Passing Interface (MPI) system to communicate between parallel processors [5]. The conformal cubic grid is decomposed over the available processors into equal rectangular parts, with no partitioning in the vertical direction. This is because the atmospheric physics, with its faster processes, depends on the entire vertical column during a time-step. Communication between parallel processors is only required for the initialize, dynamics and assimilation components (i.e., the upper half of Fig 2). There is no communication between processors in the lower half of Fig. 2, as physics routines independently update the prognostic variables for each vertical column of grid-points and the output and finalize

routines write separate NetCDF files for each processor. CCAM output variables are selectively re-combined from the parallel output files using post-processing software, which also interpolates the simulation results to a latitude/longitude grid.

Fluxes are exchanged between the CCAM atmosphere and ocean models in the physics component, due to the use of a common reversibly staggered grid. This is because the fluxes between the atmosphere and ocean are independent for each grid-point on the common grid, in contrast to models where atmosphere and ocean grids are misaligned so that additional message passing is required to communicate fluxes. Atmosphere and ocean wave dispersion properties are not compromised when using a common reversibly staggered grid, as discussed in the Section 4.

To maximize the scaling properties of CCAM, most of the dynamics software routines use non-blocking, point-to-point communication between neighboring processors to exchange portions of atmosphere and ocean data at processor boundaries. However, CCAM's semi-implicit, semi-Lagrangian dynamical core relies on an implicit solution to the Helmholtz equation, with the model performance being largely determined by the numerical algorithm used. We found that a custom-designed geometric multi-grid solver for the Helmholtz equation demonstrated the best timings for more than 1,000 processors, although Conjugate-Gradient and Successive Over-Relaxation methods are also available.

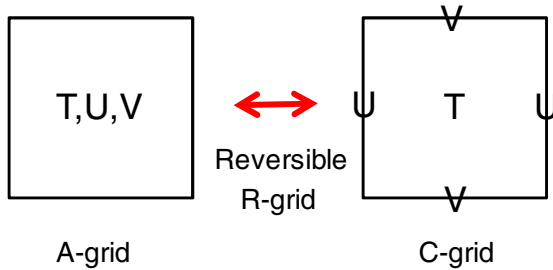
The scale-selective filter, which is used to assimilate host GCM data into CCAM at large spatial scales, naively requires  $O(N^2)$  computations where  $N$  is the number of horizontal grid points. However, the efficiency of the scale-selective filter was improved by approximating the two-dimensional convolution with two separate one-dimensional convolutions. Although this approach reduces the number of computations to  $O(N^{3/2})$ , processors still require data from all other processors whose grid-points lie along a common row or column on the CCAM cubic grid. We found that using MPI-2's Remote Memory Access (RMA) allowed us to efficiently pass the required information around the cubic grid, improving the filter timings by a factor of 5 compared to a naïve global 'gatherall' approach.

CCAM can be compiled for single precision as well as double precision computations, where the use of the single precision mode can significantly reduce communication bandwidth problems. Care is taken to avoid any degradation to model performance when using single precision, such as always using double precision for the radiation code and improving the precision for calculating global sums.

## 4 A Reversibly Staggered Ocean Model for CCAM

Our atmosphere-ocean coupled regional climate model exploits CCAM's reversibly staggered grid so that a common grid can be used for both the atmosphere and ocean. The atmosphere and ocean models can then be coupled together in the physics component, avoiding additional message passing overheads associated with exchanging momentum, radiation, temperature and moisture fluxes between models. The reversibly staggered grid determines the location of the velocity fields on the model grid as shown in Fig. 3. Most atmospheric models are based on the staggered Arakawa C-grid due to poor wave dispersion properties associated with the unstaggered A-grid. However, the A-grid is suitable for computing Coriolis terms that couple the U and V

velocity components, related to the rotation of the Earth. The A-grid is also convenient for carrying out the physical parameterizations at coincident grid points. CCAM's reversibly staggered grid switches between the Arakawa A-grid and the C-grid, thereby exploiting the advantages of both grids. Importantly, the reversible transformation can switch between the staggered grids without any loss of information. It has been shown that the reversibly staggered grid approach provides excellent dispersive properties for both atmosphere and ocean geophysical fluid dynamics [1].



**Fig. 3.** Schematic representation of the reversibly staggered grid that switches between the Arakawa A-grid (left) and the C-grid (right).  $T$  represents a scalar field (e.g., temperature) at the center of the grid box, whereas the velocity components  $U$  and  $V$  can be reversibly staggered with respect to  $T$  to improve the dispersive properties of the simulation.

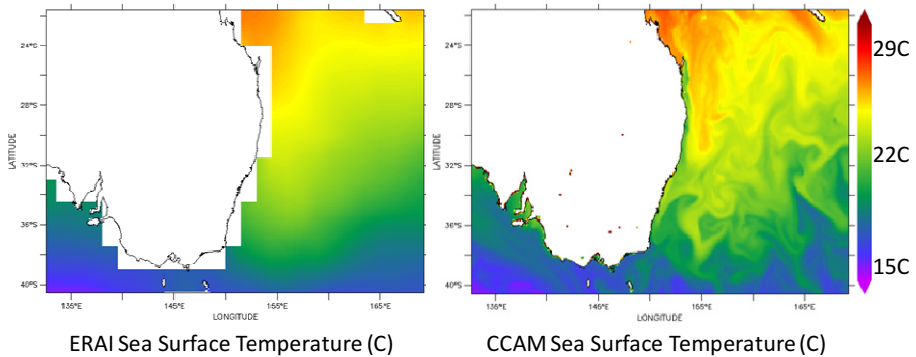
The CCAM ocean model has prognostic equations for the currents, potential temperature, salinity and the surface displacement height of the ocean. We also include a cavitating fluid model of sea-ice, although we do not discuss sea-ice further in this paper. Both the atmosphere and ocean models use a semi-implicit, semi-Lagrangian dynamical core, but with some differences such as different equations of state for density. The ocean model adopts bathymetry-following sigma- $z$  coordinates such that the vertical grid stretches and compresses as the ocean depth changes. Careful treatment of the pressure gradient terms was required to avoid spurious pressure gradients arising from discretization errors.

The Vandermonde formula, used for constructing the reversibly staggered grid in CCAM, is represented as a periodic tri-diagonal matrix that circles the globe. As there is a portion of the tri-diagonal matrix on each processor, we iteratively solve the tri-diagonal matrix to improve the scaling of CCAM, thereby avoiding global MPI communications. Typically six iterations are sufficient for the problem to converge. In the case of the ocean model, special treatment of the reversibly staggered grid is required along coastlines to be consistent with the non-slip boundary condition for currents. In practice, this problem is solved by reducing the order of the Vandermonde formula when approaching coastlines (i.e., reduced to a linear model) so that no information is lost when switching between the Arakawa A and C-grids.

Our atmosphere-ocean coupling method is particularly suited to RCMs, since the spatial resolution of the shared reversibly staggered grid is sufficient to resolve eddies in both the atmosphere and ocean. RCMs can also capitalize on the frequent exchange of information between the atmosphere and ocean models, which can be important for

simulating extreme weather events such as cyclones. Although different numbers of land and ocean points on different processors can result in load imbalance issues, this problem already existed with the simple interpolated SST approach used by most RCMs and is not exacerbated by our coupling method.

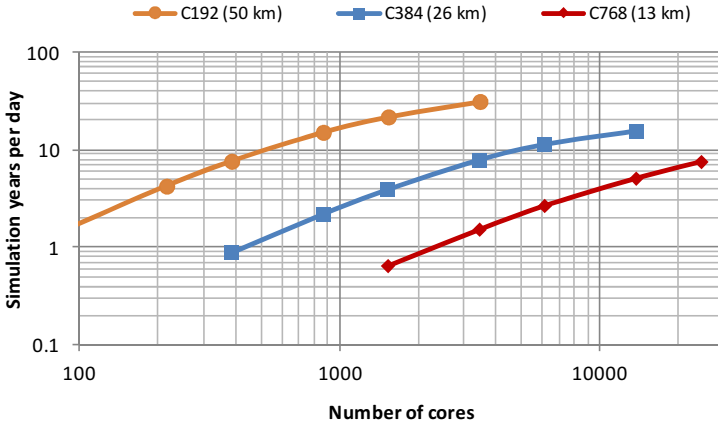
An example of how the reversibly staggered ocean model modifies the structure of SSTs can be seen in Fig. 4, where ocean surface eddies were spun up by the coupled model and can be seen in the small-scale structure of the simulated SSTs.



**Fig. 4.** Comparison between interpolated SSTs from ERA-Interim ‘observations’ (left) and spun-up SSTs by the reversibly staggered ocean model (right) at 10 km resolution over the target domain. SSTs are in units of degrees Celsius.

## 5 Performance of the Atmosphere-Ocean Coupled Model

Figure 5 shows the timings for (unstretched) global, atmosphere-ocean coupled CCAM simulations as a function of the number of processors for various grid resolutions. These simulations were performed on the Pawsey Centre’s Magnus, which is a Cray XC40 system using Intel Xeon “Haswell” processors. The CCAM simulations employed 35 vertical levels over a 40 km height for the atmosphere and 30 levels over a 5 km depth for the ocean, including a minimum of 10 levels in the top 200 m of the ocean. We used CCAM in single precision mode, which reduced the required inter-node message passing bandwidth. Finally, we applied the scale-selective filter every six simulation hours to perturb the air temperature, winds, surface pressure and SSTs towards ERA-Interim reanalyses (i.e., an atmospheric dataset that is constrained by observations). For the case of a 13 km resolution simulation, the size of the raw output files is approximately 10 GB per month, although the output is split into separate parallel output files for each processor. The 13 km resolution single precision result of 5 simulation years per day with 13,824 cores is competitive with state-of-the-art models such as the double precision NCAR GCM that produced 2.6 simulation years per day for the same horizontal resolution, but required 200,000 cores [7]. Hence CCAM’s design allows it to simulate the regional climate with considerably less computing resources than typically required for RCM experiments.



**Fig. 5.** Scaling results for single precision, atmosphere-ocean coupled CCAM simulations on IVEC's Magnus for three different global grid resolutions. Time-steps of 900 s, 360 s and 240 s were used for the 50 km, 26 km and 13 km resolution simulations, respectively.

We note that the CCAM timings are a consequence of a number of factors including the use of single precision, the semi-Lagrangian dynamical core, the geometric multi-grid algorithm, as well as the reversibly staggered coupling strategy. The multi-grid solver and the scale-selective filter were found to be the weakest components of CCAM with respect to scaling, as the atmosphere-ocean coupling has no communication overhead. The scale-selective filter was found to scale proportionally to the square root of the number of processors when using RMA, but is only called every six simulation hours. Therefore the relative cost of the filter remains the same portion of the total run time when the resolution is increased and the model time-step is decreased. We are then left with the geometric multi-grid solver as the primary limitation on model scaling. Although our custom-designed geometric multi-grid solver has competitive timings compared to other multi-grid solvers reported in the literature, it can be difficult to fully exploit the available parallel processing with multi-grid methods. For this reason, we are investigating explicit methods for the dynamical core as described in Section 6.

## 6 Conclusions

This paper describes an alternative approach to atmosphere-ocean coupled regional climate modeling that is based on a reversibly staggered grid. The excellent dispersive properties of this grid are suitable for simulating both atmosphere and ocean geophysical fluid dynamics, making it appropriate as a common grid for integrating the two models. The message passing overhead associated with exchanging fluxes between the atmosphere and ocean models can then be eliminated. The approach is valid when both atmosphere and ocean eddies are resolved on the common grid and the frequent exchange of fluxes between the models is important, making the method



well suited for RCMs. We have constructed a prototype of the reversibly staggered, atmosphere-ocean coupled regional model using CCAM and show that the single precision version can produce timings of 5 simulation years per day using 13,824 cores. Factors contributing to this result include the use of single precision, the semi-Lagrangian dynamical core, the geometric multi-grid algorithm, as well as the reversibly staggered grid coupling strategy. The CCAM timings are comparable to a state-of-the-art GCM using 200,000 cores [7], although CCAM requires considerably fewer cores making it well suited to regional climate modeling.

Although entirely adequate for petascale computing resources currently available in Australia, CCAM scaling will eventually be limited by the implicit geometric multi-grid solver for the Helmholtz equation. This issue is being addressed with the development of a split-explicit, flux-conserving dynamical core that will use an equi-angular, gnomonic cubic grid. Nevertheless, the current semi-implicit, semi-Lagrangian prototype successfully demonstrates the viability of the reversible staggering approach for atmosphere-ocean coupled regional climate modeling. It may also be possible to use the reversibly staggered grid for other types of geophysical fluids, although the feasibility of our approach will depend on the dispersive properties of the reversibly staggered grid for the modeled system.

**Acknowledgements.** The authors wish to thank Pawsey Centre's Petascale Pioneer program for providing computing resources for the simulations described in this paper. Thanks to Aaron McDonough and Paul Ryan for their technical advice. We also acknowledge the constructive feedback on the manuscript from Peter Dobrohotoff and the two anonymous reviewers.

## References

1. McGregor, J.L.: Geostrophic adjustment for reversibly staggered grids. *Mon. Wea. Rev.* 133, 1119–1128 (2005)
2. McGregor, J.L.: C-CAM: Geometric aspects and dynamical formulation. CSIRO Marine and Atmospheric Research Technical Report 70 (2005), [http://www.cmar.csiro.au/e-print/open/mcgregor\\_2005a.pdf](http://www.cmar.csiro.au/e-print/open/mcgregor_2005a.pdf)
3. Schmidt, F.: Variable fine mesh in spectral global models. *Beitr. Phys. Atmos.* 50, 211–217 (1977)
4. Fox-Rabinovitz, M., Cote, J., Dugas, B., Deque, M., McGregor, J.L.: Variable resolution general circulation models: Stretchedgrid model intercomparison project (SGMIP). *J. Geophys. Res.* 111, D16104 (2006)
5. McGregor, J.L., Dix, M.: An updated description of the conformal-cubic atmospheric model. In: Hamilton, K., Ohfuchi, W. (eds.) *High Resolution Simulation of the Atmosphere and Ocean*, pp. 51–76. Springer (2008)
6. Thatcher, M., McGregor, J.L.: Using a scale-selective filter for dynamical downscaling with the conformal cubic atmospheric model. *Mon. Wea. Rev.* 137, 1742–1752 (2009)
7. Worley, P., Craig, A., Dennis, J.: Performance of the community earth system model. In: *2011 International Conference High Performance Computing, Networking, Storage and Analysis (SC)*, Seattle, WA, pp. 1–11 (2011)