CM3/ESM3 Planning
January 12, 2006

The goal of this document is to lay out a pathway for the next 5 years of global coupled model development at GFDL, for both the physical model (CM) and earth system model (ESM). This document is not intended to be overly prescriptive, but it is important to set general targets in terms of resolution, physics, and model timing, as well as a specific development timetable.

We have outlined a pathway in which ESM and CM development is tightly coupled within a unified code. It is understood that the code will have switches to allow one to turn off ESM features, such as interactive carbon and chemical cycles, that are too time-consuming for certain applications. These switches reflect a hierarchical approach to model development in which successively more complex models are built upon the stable platforms of acceptable solutions with less complex models. It is recognized that the generation and maintenance of multiple code opt ions will have to be mindful of the Laboratory's available human resources. The goal is to permit easy use of the same code for both physical climate studies and fully interactive ESM studies. We have scoped out resolution choices with the requirements of the fully-interactive option in mind. This necessitates tradeoffs, but the benefits of a coherent development pathway should more than offset those tradeoffs.

PATHWAY:

CM3/ESM3 should be completed by January, 2009. This date is determined both by the timing for the next round of the IPCC process (estimated for completion in 2013) and a healthy 4-5 year interval between major coupled model releases (CM2.1 was finalized in summer, 2004). To meet this time frame, the component models should be mature (generating satisfactory solutions in uncoupled mode) by December, 2007. This will allow 12 months for optimizing CM3/ESM3.

Component model development activity will continue using CM2.1 as a base for much of this period. An example is the development of ESM2.1, adding interactive carbon cycle/land vegetation to CM2.1. The timing as to when particular component development paths are merged into the main trunk model will be determined by the MDTs.

Model speed is a crucial consideration. We advocate a model speed of at least 5 years per day on a "reasonable" number of dedicated processors. A "reasonable" number is one which permits good turnaround when jobs need to move through the queues, and also satisfies the constraint that the number of simultaneously executing copies of the model needed for development and application utilizes an acceptable fraction of the system. For the current system, our estimate is that 120 processors can be viewed as a "reasonable" number, corresponding roughly to 5% of the total system for a single copy.

We realize that, in order to sustain frontline status in the world, e.g., for IPCC AR5, GFDL's principal climate model must have the capability to perform simulations with interactive chemistry (gas and aerosol) and an interactive carbon cycle.

We describe three streams of research on CM3/ESM3, but only the first of these streams currently meets these computational criteria.

Stream I:

(a) ATMOSPHERE: M45 (finite volume -- probably cubed-sphere) atmosphere, 40-50 levels, with new physics to be decided upon by the MDTs. This physics should include a more realistic treatment of convection, interactive chemistry (both tropospheric and stratospheric) including predicted ozone, and interactive aerosols, including the "indirect" interactions between aerosols and clouds.

(b) OCEAN: resolution similar to OM3 (1 degree, finer in tropics, approximately 50 levels), with ocean biogeochemsitry. Whether this is HIM based, MOM based, or a hybrid of the two will be determined by the MDTs.

(c) LAND: LM3, including dynamic vegetation, vertically-resolved soil hydrology and explicit rivers.

(d) SEA ICE: similar to CM2.

Our estimates are that the above configuration, running with interactive atmospheric and oceanic chemistry, would achieve 4 years per day on 150 processors on the current system (Jan 2006). We anticipate that such a configuration should meet the 5 years per day requirement on 120 processors on our next system. Crucial to the above plan is the need for efficient scaling of the M45 atmospheric code to at least 100 processors. Efforts to achieve this must be of very high priority. A small number of model versions need to exist, each generating acceptable solutions, with more or less interactivity. The details of these versions will be decided at a later date, but they could include versions with and without interactive atmospheric chemistry (that is, with prescribed ozone and aerosols), versions with and without ocean biogeochemistry, and versions with and without interactive vegetation. While it will be a challenge to manage, these "switches" are necessary given the requirements of our science applications, the reduction in model efficiency that results from some of the added interactive components, and the need to maintain a unified code.

This Stream I model needs to receive the bulk of the attention of the development process. The choice of resolution, particularly for the atmosphere, is determined by estimates of modest future computer increases. The model design also represents a consensus that, given our current position, the greatest scientific benefit would come from concentrating additional computational resources on 1) the inclusion of the oceanic geochemistry needed for an interactive carbon cycle, and on 2) the atmosphere, specifically on the need for higher vertical resolution to simulate stratosphere/troposphere interaction and to predict ozone, and, possibly, to improve cloud simulations, and on adding the capability to simulate climate/aerosol/chemistry interaction.

Stream II:

This is the same as Stream I but using an M90 resolution in the atmosphere. The Stream II model will be almost 4 times as expensive as the Stream 1 model. It is our hope that Stream 2 can become the workhorse version for CM3/ESM3, since the M90 atmosphere offers significant improvements. This may be possible with increases in computer power and/or code efficiencies. Efforts at model optimization are crucial if we are to increase the likelihood of replacing Stream I by Stream II.

It is expected that Stream II will be developed in part for S-I applications initially. The goal is to produce a coupled model that works well at both resolutions, but our estimates strongly suggest that choosing M90 rather than M45 for our workhorse at this point will make us dangerously dependent on unforeseeable increases in computational resources. The value of compromise resolutions between M45 and M90 should also be examined.

Stream III:

Stream III augments Stream II by increasing the ocean resolution to approximately 1/4 degree. This increase in resolution, while not eddy-resolving, is considered to present the opportunity for significant improvements in ocean circulation. It is assumed that this Stream will receive attention from the Ocean Group initially, with increasing attention and efforts for long-term climate use as computer resources and code efficiencies permit. We estimate that the atmospheric and oceanic components of Stream III will roughly restore balance between the computational requirements of the atmospheric and oceanic components of the model. Candidate ocean components scale well enough that the wall-clock speed of Streams II and III should be similar with concurrent coupling, using roughly 1.5-3 times as many processors for Stream III. We note that Stream III will have significant scientific value even before it is cost-effective for IPCC-style scenario runs, and that results from Stream III should have positive feedback on the development process for Streams I and II.

In summary, Stream I outlined above will be the primary initial pathway for CM3/ESM3 development. We will also put effort into Stream II, but at a lower level of intensity - until CP and/or code efficiencies permit Stream II to achieve the required speed. We expect work on Stream III to gain momentum in the latter part of this 5 year window, and that initial efforts in the direction of greater oceanic resolution will provide a firm foundation for follow-on efforts.
We hope that the multiple stream process outlined here will foster a continuous approach to coupled model development. Multiple streams will provide a hierarchy of tuned coupled models to meet the needs of a diverse scientific user base.