Mosaic: a key abstraction for model coupling, nesting and chaining

V. Balaji
Princeton University and NOAA/GFDL

22 September 2006

What is a grid mosaic?

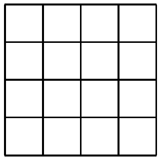

On the left is a basic 4×4 ***tile***; on the right are examples of grids composed of a mosaic of such tiles. The first is a ***continuous grid***, below is a ***refined grid***.

Most current software only supports what we call ***grid tiles*** here. The **`grid mosaic`** extension will allow the development of more complex grids for next-generation models.
First in our (GFDL's) sights is the ***cubic sphere***, primarily targeted at a next-generation finite-volume atmospheric dynamical core, but potentially others as well.
Further developments will include support for irregular tiling (e.g of the ocean surface following coastlines), and for refined, nested and adaptive grids.
Also, regular grids where an irregular decomposition is needed (e.g for a polar filter) can use mosaics to define different decompositions in different regions.

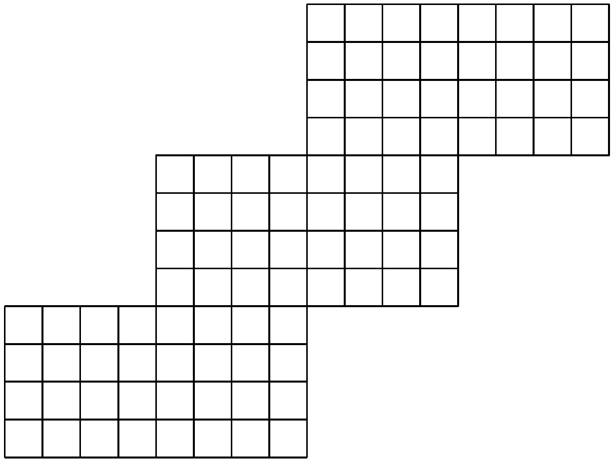

Regular grid mosaic.

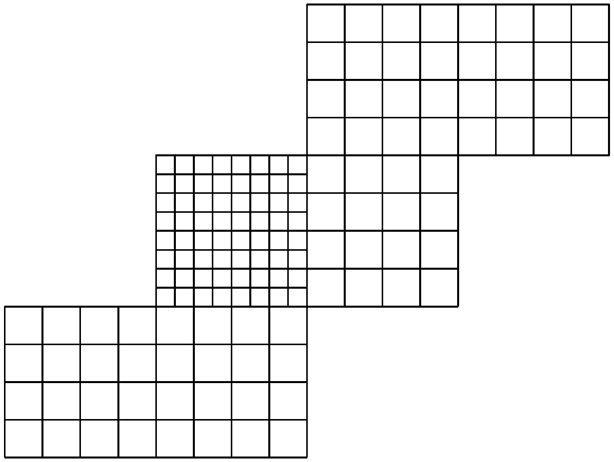

Refined grid mosaic.

“True” refinement

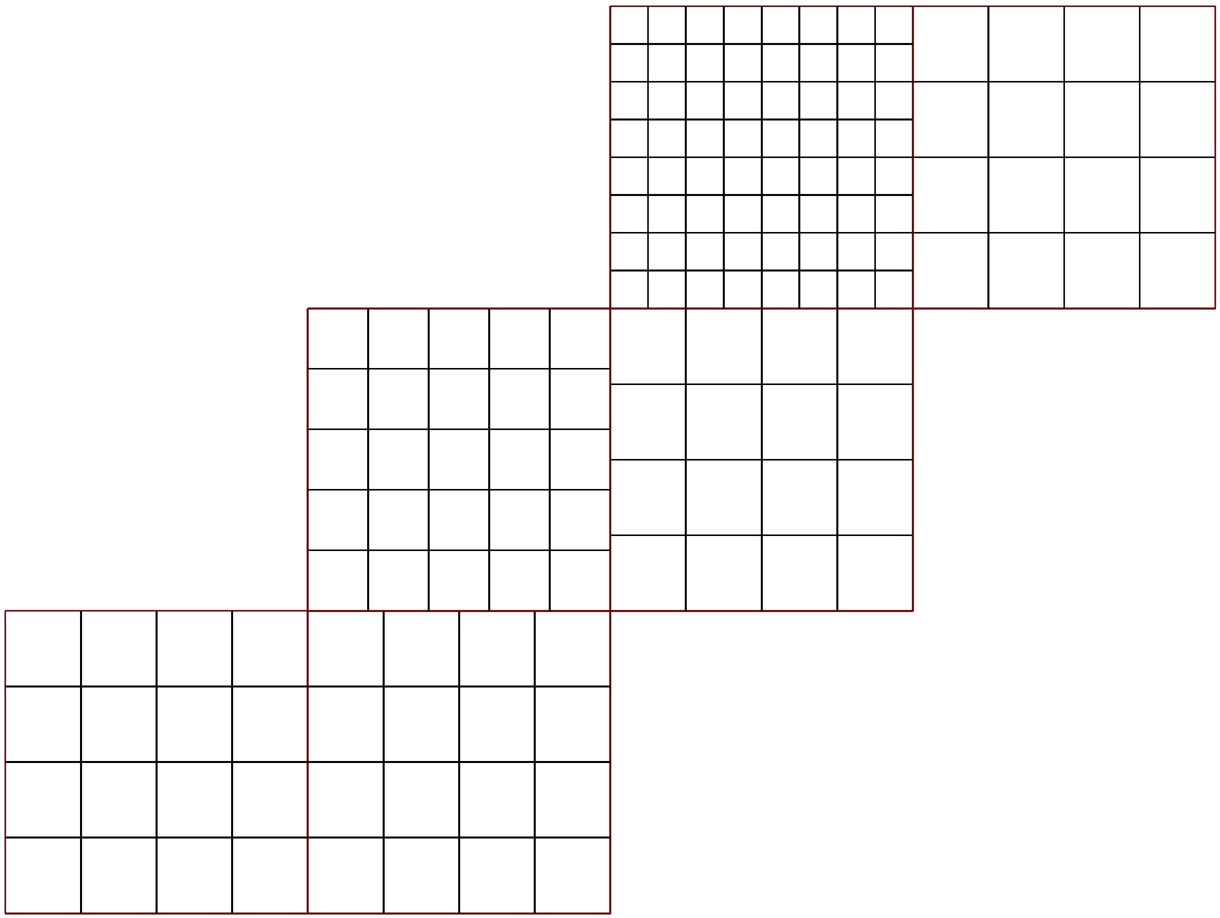

The tile at **`(8,8)`** shows ***true*** refinement; the tile at **`(4,4)`** does not.

Boundaries and contact regions

Aside from the grid information in the grid tiles, the grid mosaic additionally specifies connections between pairs of tiles in the form of ***contact regions*** between ***pairs*** of grid tiles.
Contact regions can be ***boundaries***, topologically of one dimension less than the grid tiles (i.e, planes between volumes, or lines between planes), or ***overlaps***, topologically equal in dimension to the grid tile. In the cubed-sphere example the contact regions between grid tiles are 1D boundaries: other grids may contain tiles that overlap. In the example of the ***yin-yang*** grid (Kageyama et al 2004) the grid mosaic contains two grid tiles that are each lon-lat grids, with an overlap. The overlap is also specified in terms of a ***contact region*** between pairs of grid tiles. Issues relating to boundaries are described below. Overlaps are described in terms of an exchange grid, more on which below.

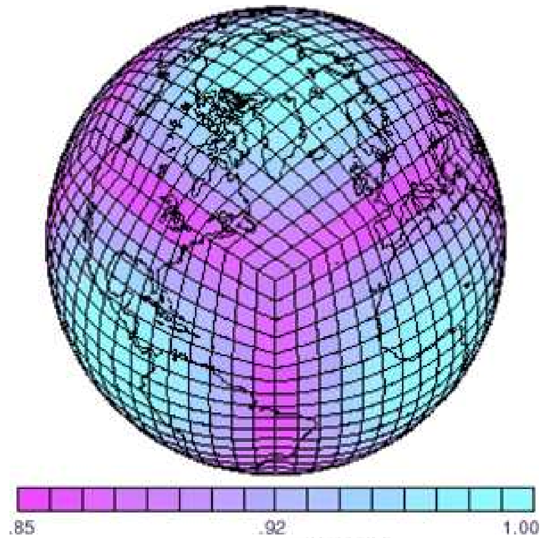

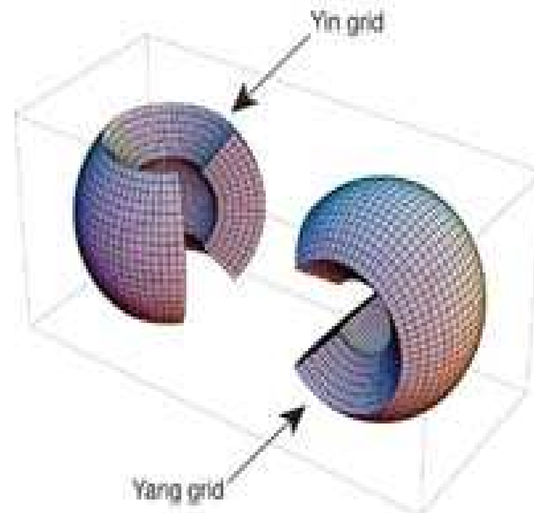

Grid mosaic definition

A ***grid mosaic*** is constructed recursively by referring to child mosaics, with the tree terminating in leaves defined by ***grid tiles***.
(There is a very useful analogy to be made between mosaic hierarchies and model component hierarchies).

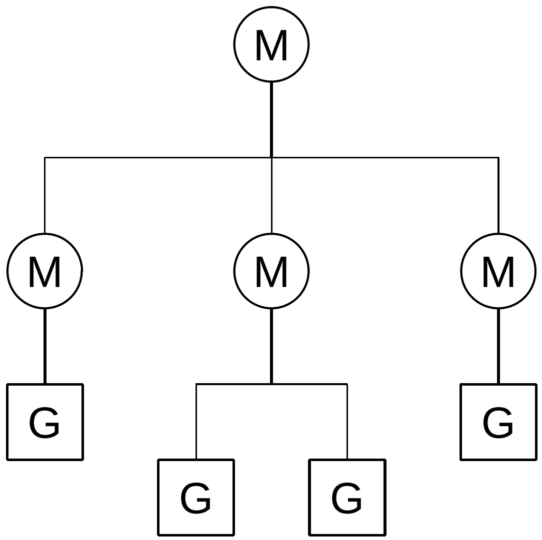

It is not necessarily possible to deduce contact regions by geospatial mapping: there can be applications where geographically collocated regions do ***not*** exchange data, and also where there is implicit contact between non-collocated regions.

Applications of grid mosaics

The grid mosaic is a powerful abstraction making possible an entire panoply of applications. These include:

- the use of overset grids such as the yin-yang grid;
- the representation of nested grids (e.g Kurihara et al 1990);
- the representation of reduced grids (e.g Rasch 1994). Currently these typically use full arrays and a specification of the “ragged edge”. A reduced grid can instead be written as a grid mosaic where each reduction appears as a separate grid tile.
- An entire coupled model application or dataset can be constructed as a hierarchical mosaic. Grid mosaics representing atmosphere, land, ocean components and so on, as well as contact regions between them, all can be represented using this abstraction. This approach is already in use at many modeling centres including GFDL, though not formalized.
- Finally, grid mosaics can be used to overcome performance bottlenecks associated with parallel I/O and very large files. Representing the model grid by a mosaic permits one to save data to multiple files, and the step of ***aggregation*** is deferred. This approach is already used at GFDL to perform distributed I/O from a parallel application, where I/O aggregation is deferred and performed on a separate I/O server sharing a filesystem with the compute server.

Boundary spec for a cubed sphere

Boundaries for LRG tiles are specified in terms of an ***anchor point*** and an ***orientation***.

An anchor point is a boundary point that is common to the two grid tiles in contact. When possible, it is specified as integers giving index space locations of the anchor point on the two grid tiles. When there is no common grid point, the anchor point is specified in terms of floating point numbers giving a geographic location.

The ***orientation*** of the boundary specifies the index space direction of the running boundary on each grid tile: the point just to the "west" of **`(5,6)`** is in fact **`(3,4)`**

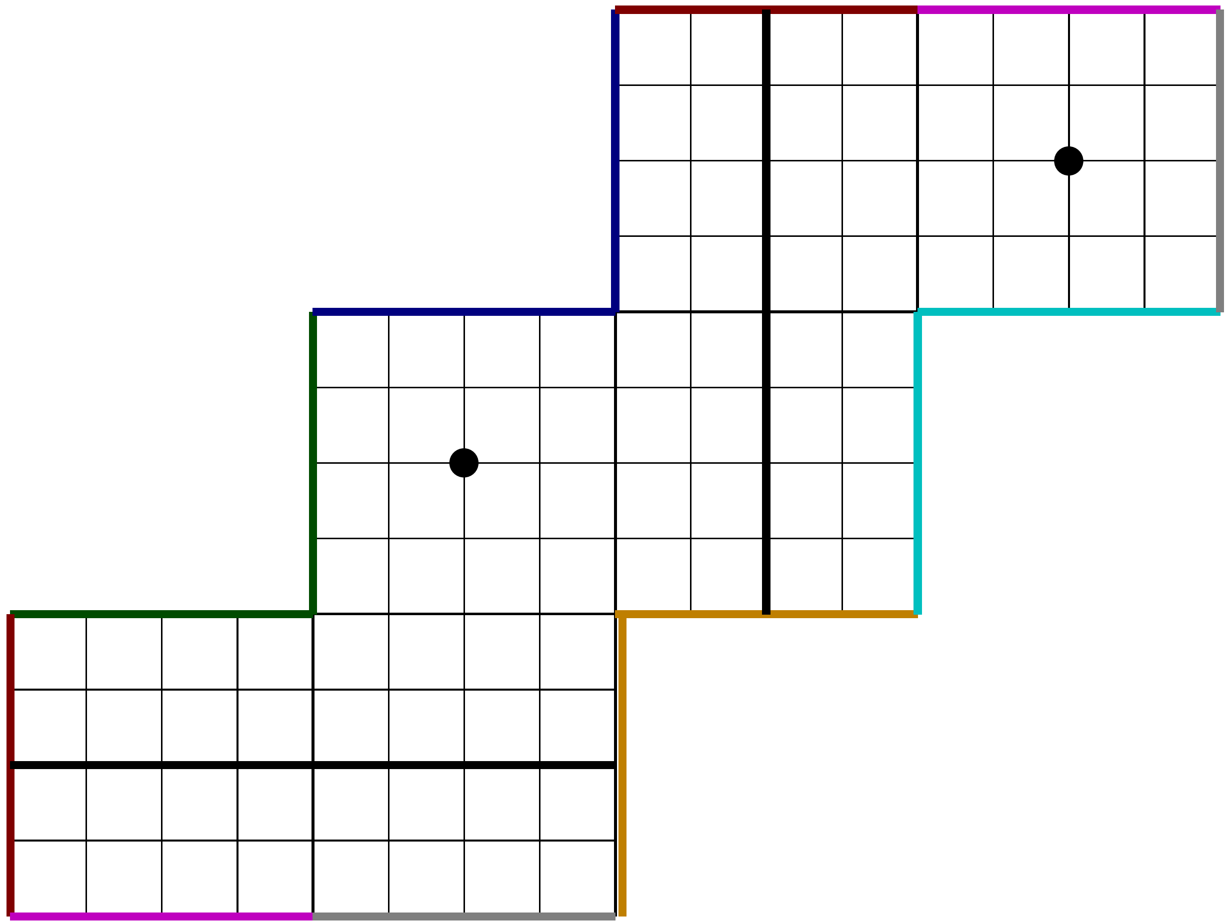

Overlap contact regions: the exchange grid

- A ***grid*** is defined as a set of ***cells*** created by ***edges*** joining pairs of ***vertices*** defined in a discretization.
- An ***exchange grid*** is the set of cells defined by the union of all the vertices of the two parent grids, and a ***fractional area*** with respect to the parent grid cell.

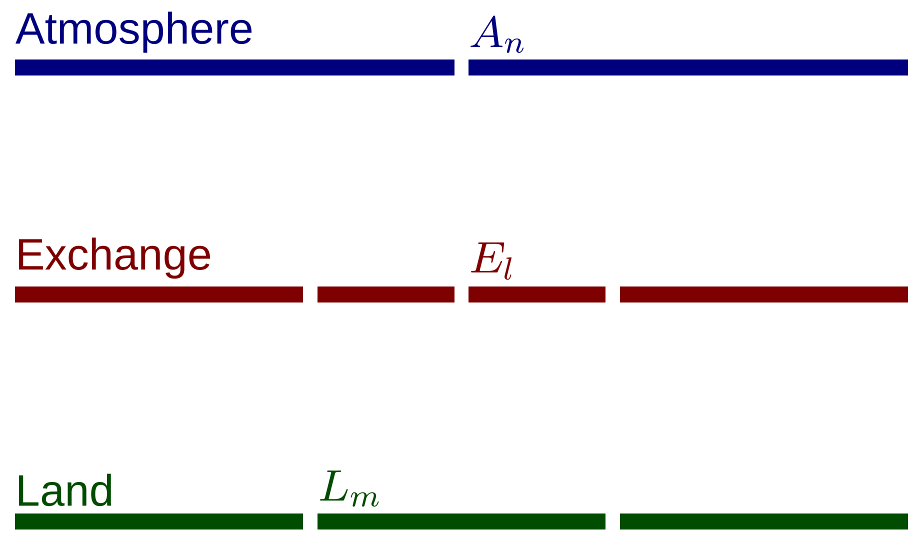

- Exchange: interpolate from source grid using one set of fractional areas; then average onto the target grid using the other set of fractional areas.
- Consistent moment-conserving interpolation and averaging functions of the fractional area may be employed.

Overlap contact regions: masks

Complementary components: in Earth system models, a typical example is that of an ocean and land surface that together tile the area under the atmosphere.

Land-sea mask as discretized on the two grids, with the cells marked **L** belonging to the land. Certain exchange grid cells have ambiguous status: the two blue cells are claimed by both land and ocean, while the orphan red cell is claimed by neither.

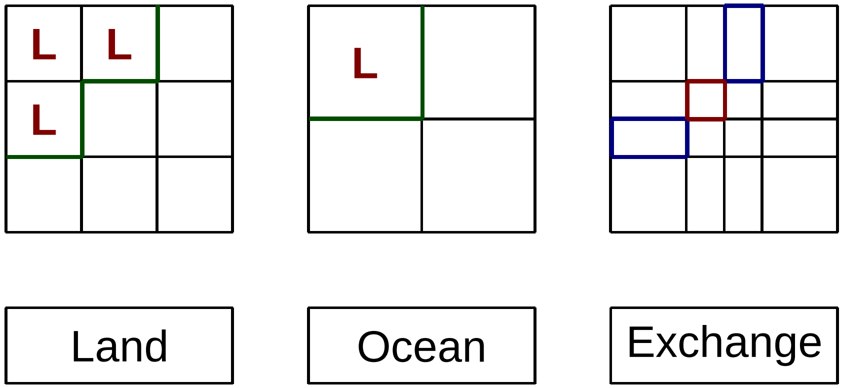

Therefore the mask defining the boundary between complementary grids can only be accurately defined on the exchange grid.

In the FMS exchange grid, by convention (and because it is easier) we generally modify the land grid as needed. We add cells to the land grid until there are no orphan "red" cells left on the exchange grid, then get rid of the "blue" cells by ***clipping*** the fractional areas on the land side.

Summary: approaches to grid coupling

- The generalization of ***grids*** to ***mosaics*** is a key development in GFDL's approach to nested grid models (Balaji, Adcroft, Zhi Liang).
- Mosaics are to be used in leading-edge model discretizations such as the ***cube-sphere*** (S-J Lin, Michael Herzog).
- ***Simple*** refinement will be used for two-way nesting of models within themselves, including adaptive grids (Adcroft, Hallberg, Herzog).
- Interpolation between arbitrary grids for model ***coupling*** and model ***chaining*** is accomplished using ***exchange grids*** (Anderson, Balaji, Winton).
- A ***standard*** for describing grids and mosaics is in draft stage and is expected to be adopted by the CF community (Balaji). This is expected to lead to the development of standardized regridding methods and tools to aid in the process of model nesting, coupling and chaining.