PROGRESS TOWARDS A REGIONAL CAPABILITY IN MPAS-ATMOSPHERE

Michael G. Duda[1], William C. Skamarock[1], Soyoung Ha[1], and Sang-Hun Park[2]

[1]National Center for Atmospheric Research*,
Mesoscale and Microscale Meteorology Laboratory
[2]Yonsei University, South Korea

*NCAR is funded by the National Science Foundation

OVERVIEW

Weather Research & Forecasting Model for Prediction Across Scales

How do we define limited-area domains in MPAS?

What do LBCs and LBC fields look like?

How expensive is limited-area MPAS?

What does the current workflow look like for running limited-area MPAS?

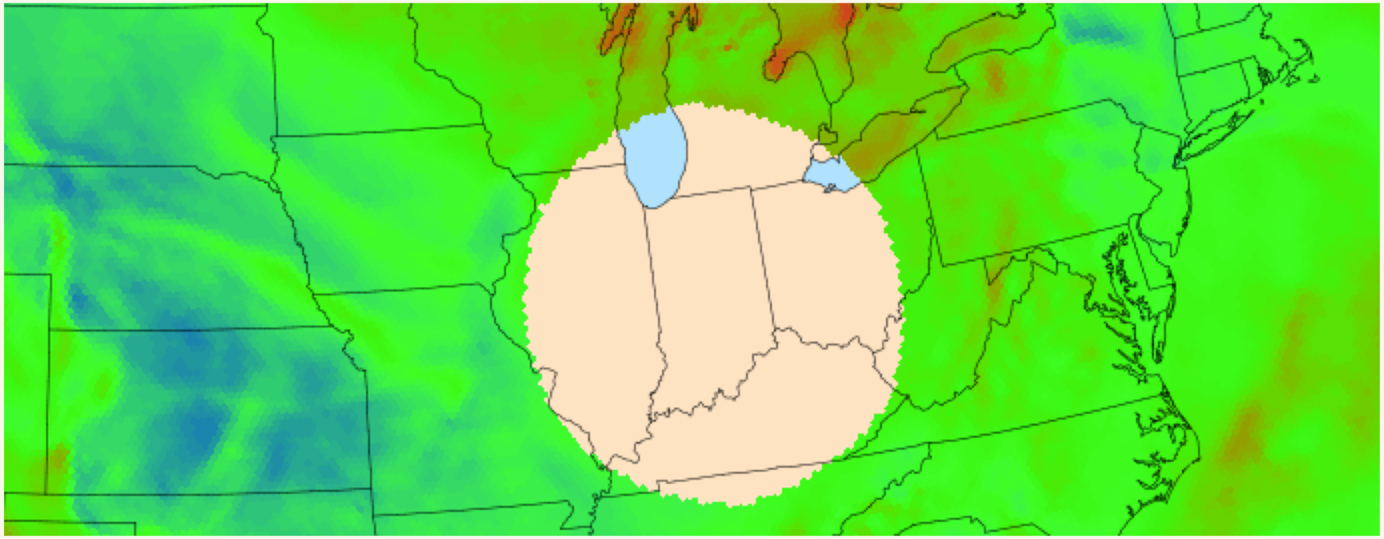

CREATING LIMITED-AREA MESHES

Rather than constructing regional centroidal Voronoi tesselations from scratch, regional domains are defined as subsets of existing MPAS meshes

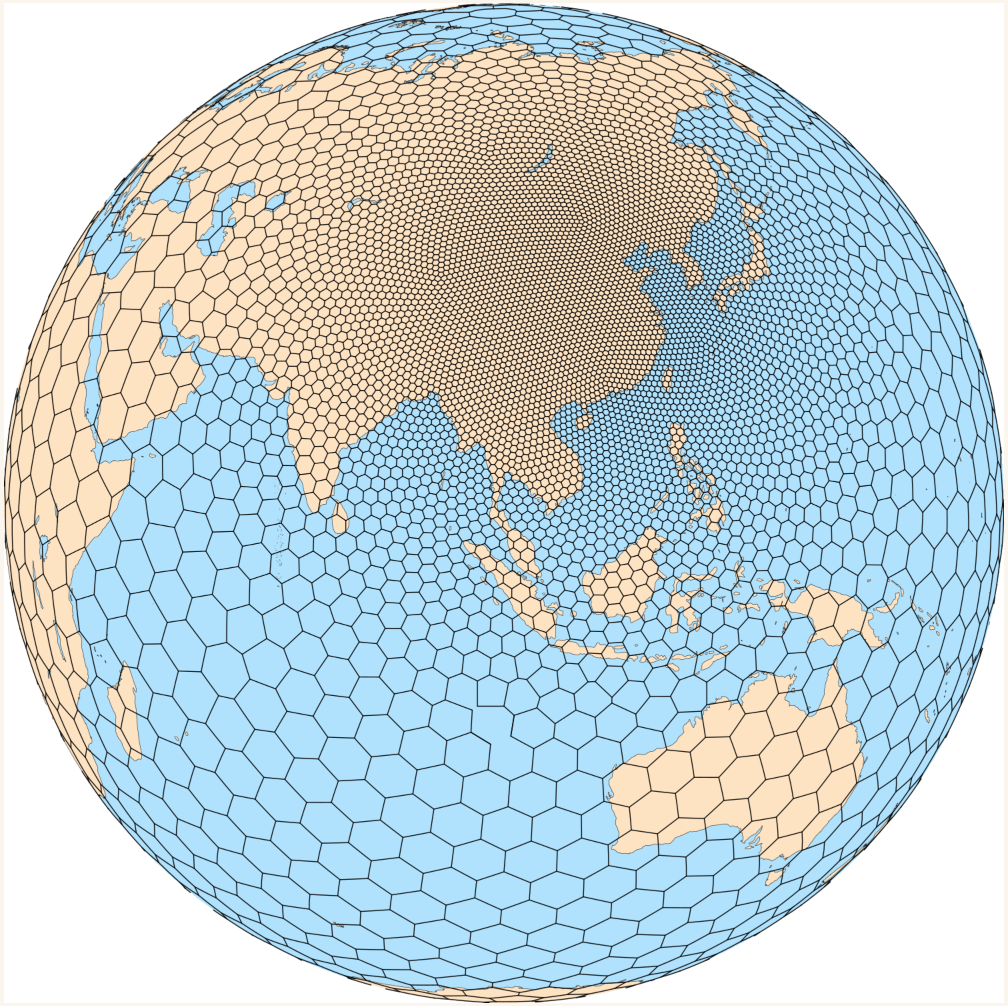

CREATING LIMITED-AREA MESHES

Six easy steps:

1. Take any existing (global) mesh, rotate the refinement as needed

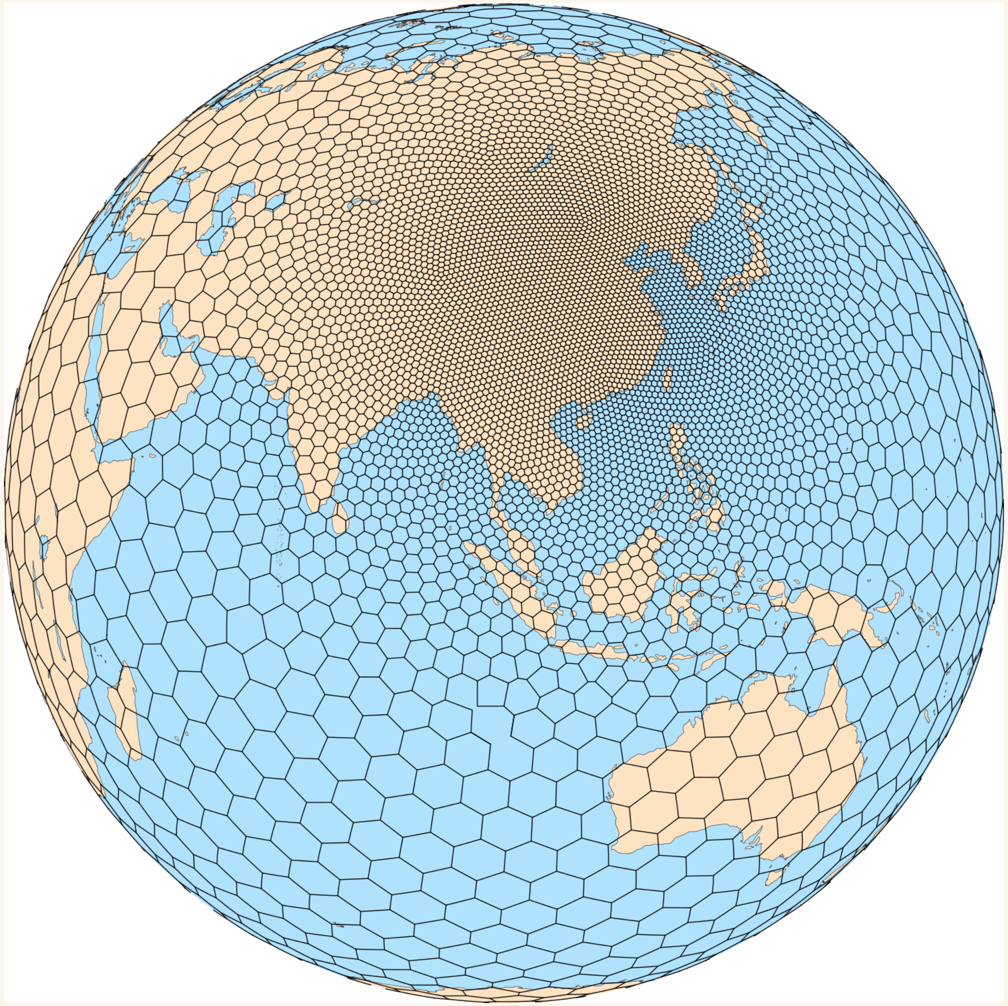

CREATING LIMITED-AREA MESHES

Six easy steps:

1. Take any existing (global) mesh, rotate the refinement as needed

2. Identify a region of interest

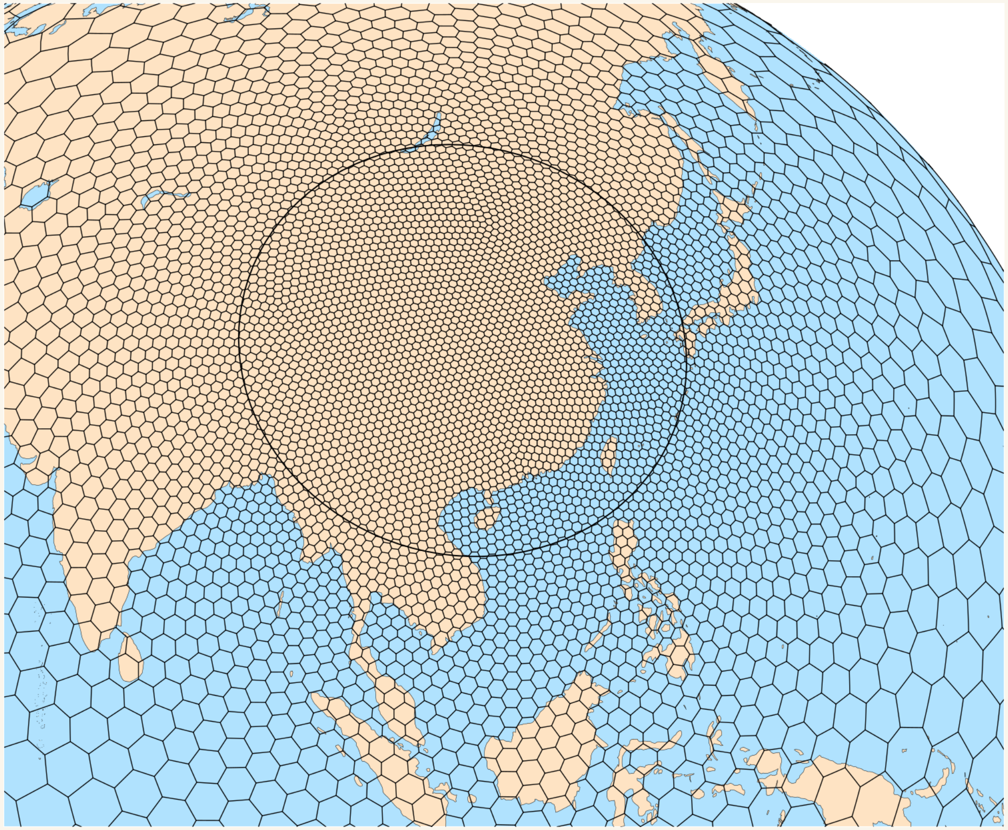

CREATING LIMITED-AREA MESHES

Six easy steps:

1. Take any existing (global) mesh, rotate the refinement as needed

2. Identify a region of interest

3. Let the regional domain consist of all cells whose generating point (center, roughly) lies within the region

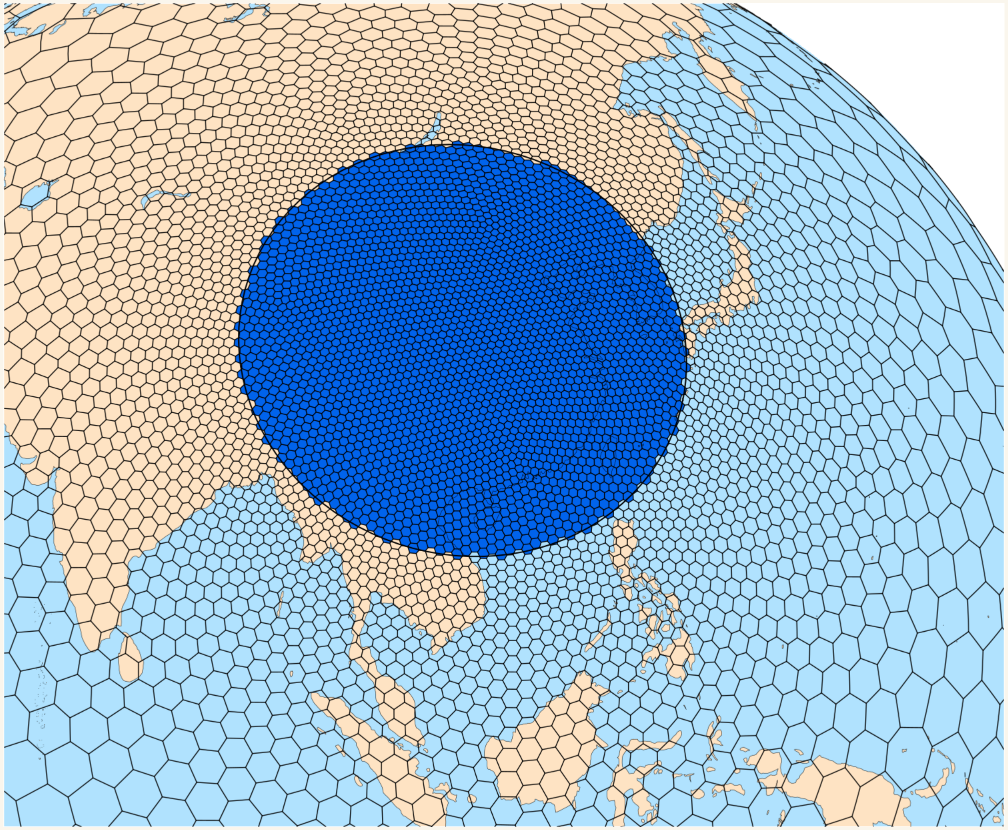

CREATING LIMITED-AREA MESHES

Six easy steps:

1. Take any existing (global) mesh, rotate the refinement as needed

2. Identify a region of interest

3. Let the regional domain consist of all cells whose generating point (center, roughly) lies within the region

4. Build layers of "relaxation zone" cells (we've been testing with 5)

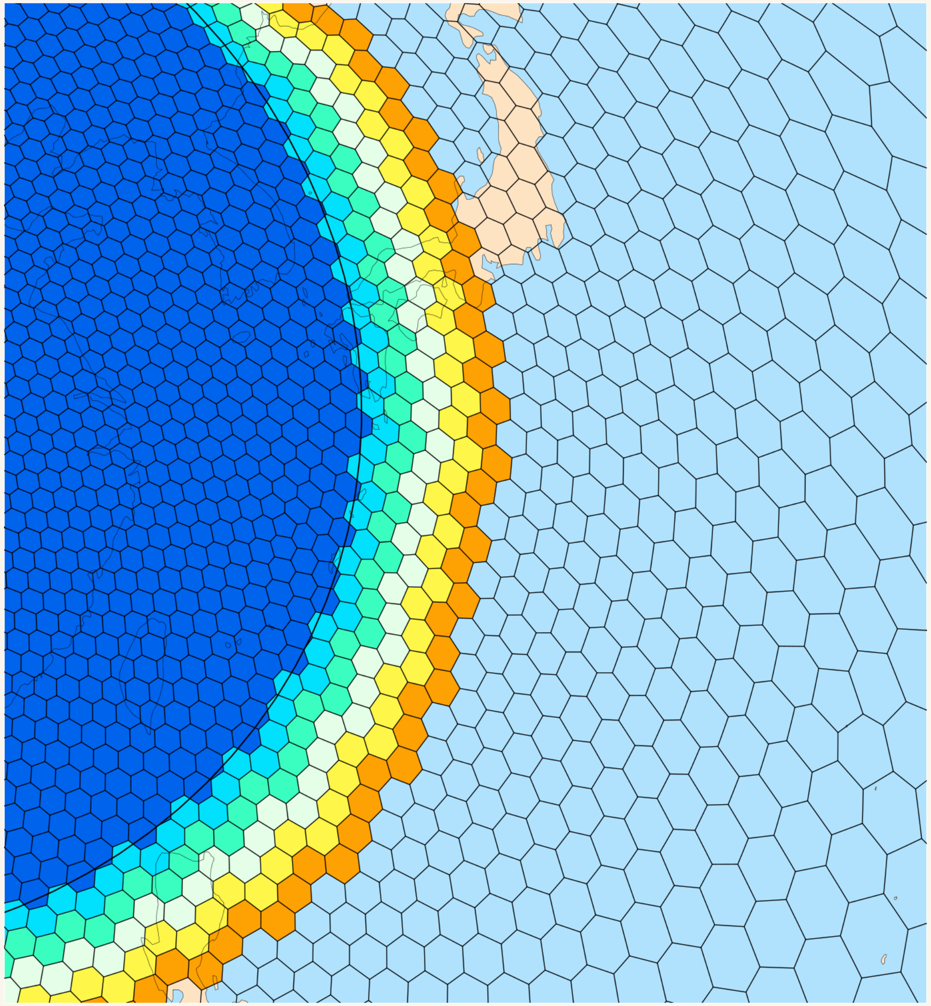

CREATING LIMITED-AREA MESHES

Six easy steps:

1. Take any existing (global) mesh, rotate the refinement as needed

2. Identify a region of interest

3. Let the regional domain consist of all cells whose generating point (center, roughly) lies within the region

4. Build layers of “relaxation zone” cells (we’ve been testing with 5)

5. Add two layers of “specified zone” cells

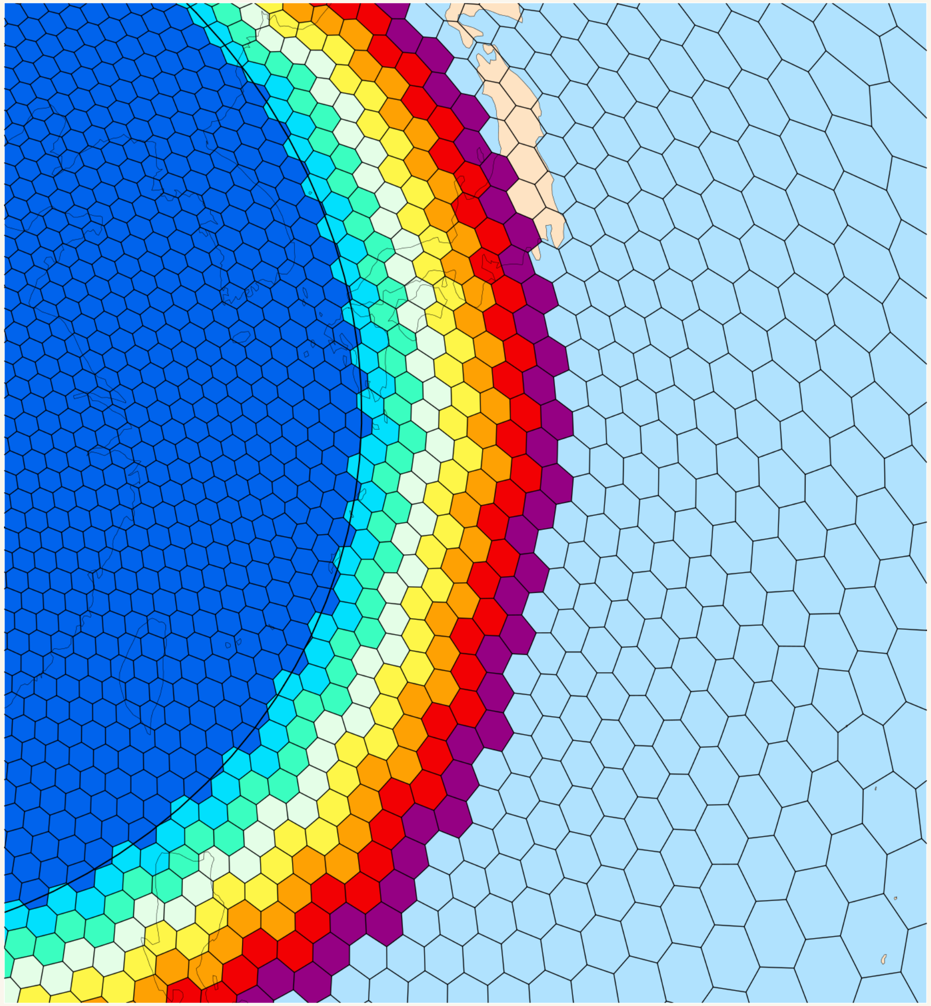

CREATING LIMITED-AREA MESHES

Six easy steps:

1. Take any existing (global) mesh, rotate the refinement as needed

2. Identify a region of interest

3. Let the regional domain consist of all cells whose generating point (center, roughly) lies within the region

4. Build layers of “relaxation zone” cells (we’ve been testing with 5)

5. Add two layers of “specified zone” cells

6. Cull any remaining cells and re-index cells/edges/vertices to a contiguous range of indices

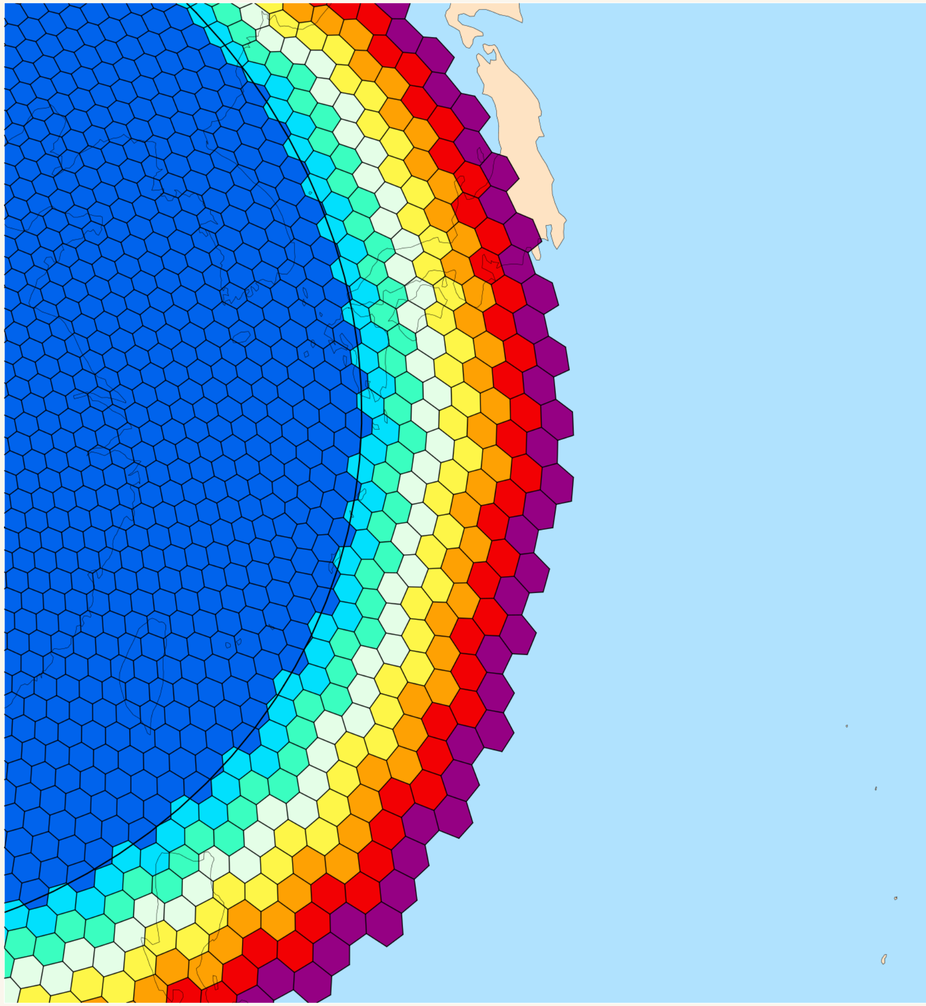

CREATING LIMITED-AREA MESHES

The same method works for regions defined by a set of connected points

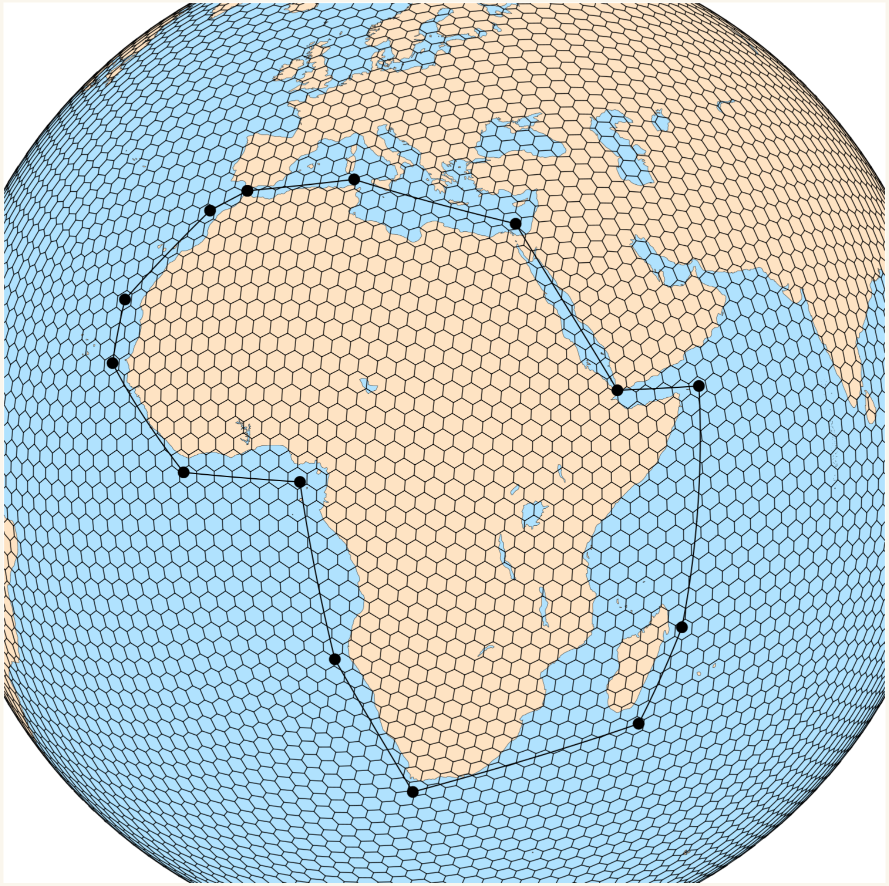

CREATING LIMITED-AREA MESHES

The same method works for regions defined by a set of connected points

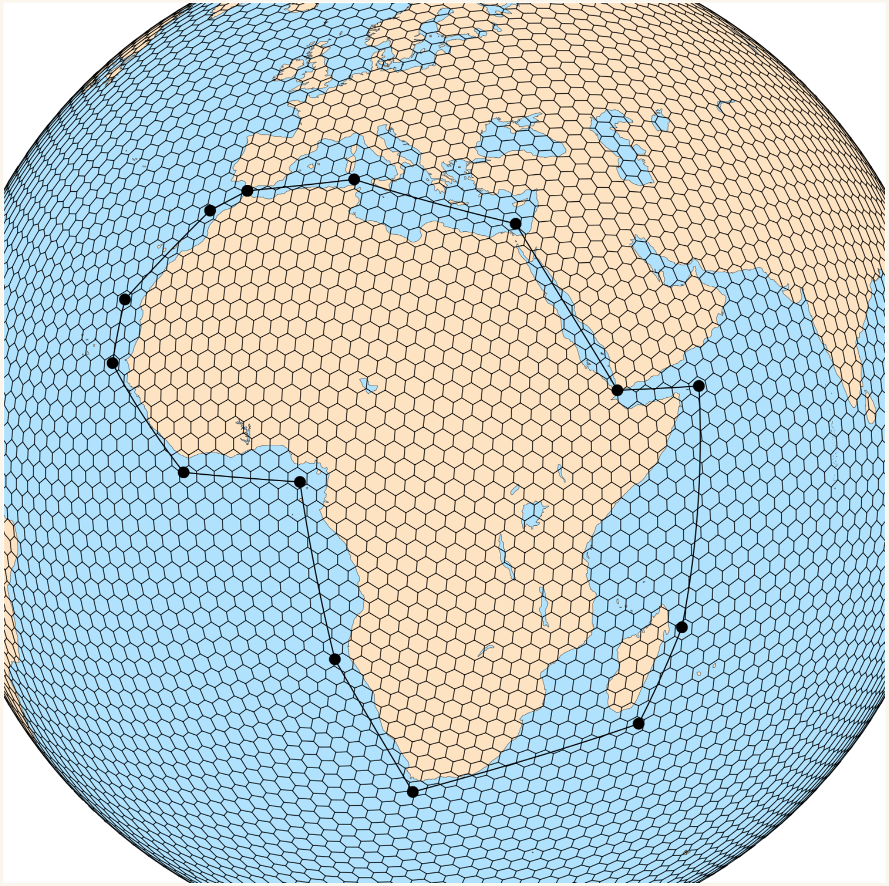

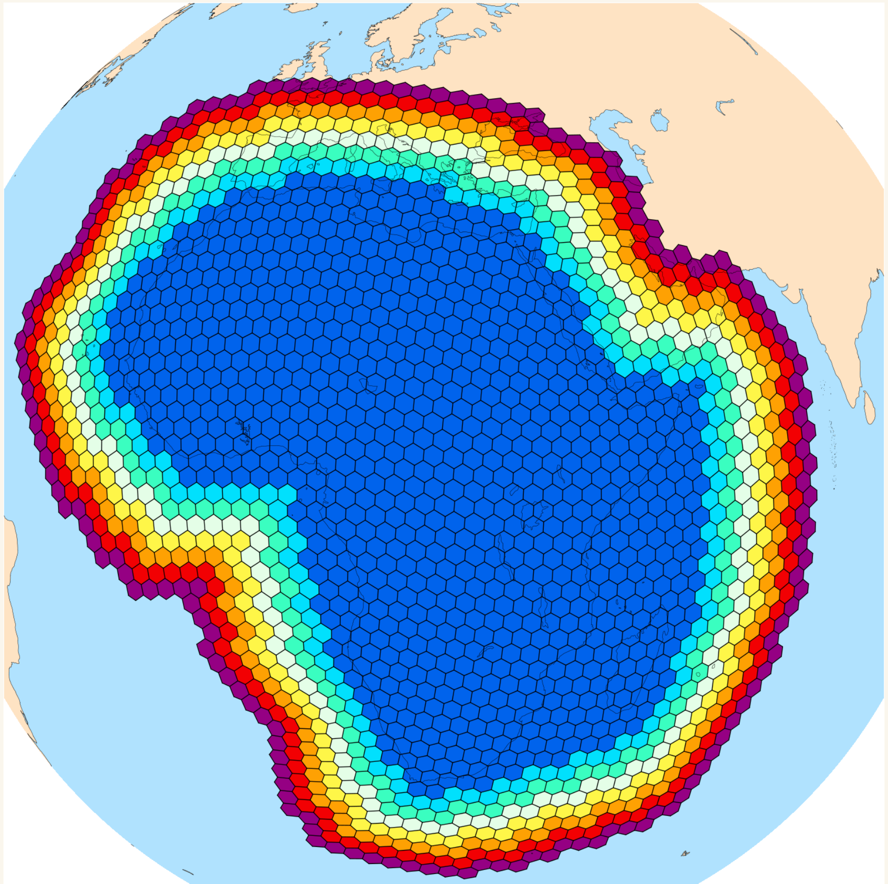

CREATING LIMITED-AREA MESHES

With unstructured meshes, the creative user may define exotic simulation domains:

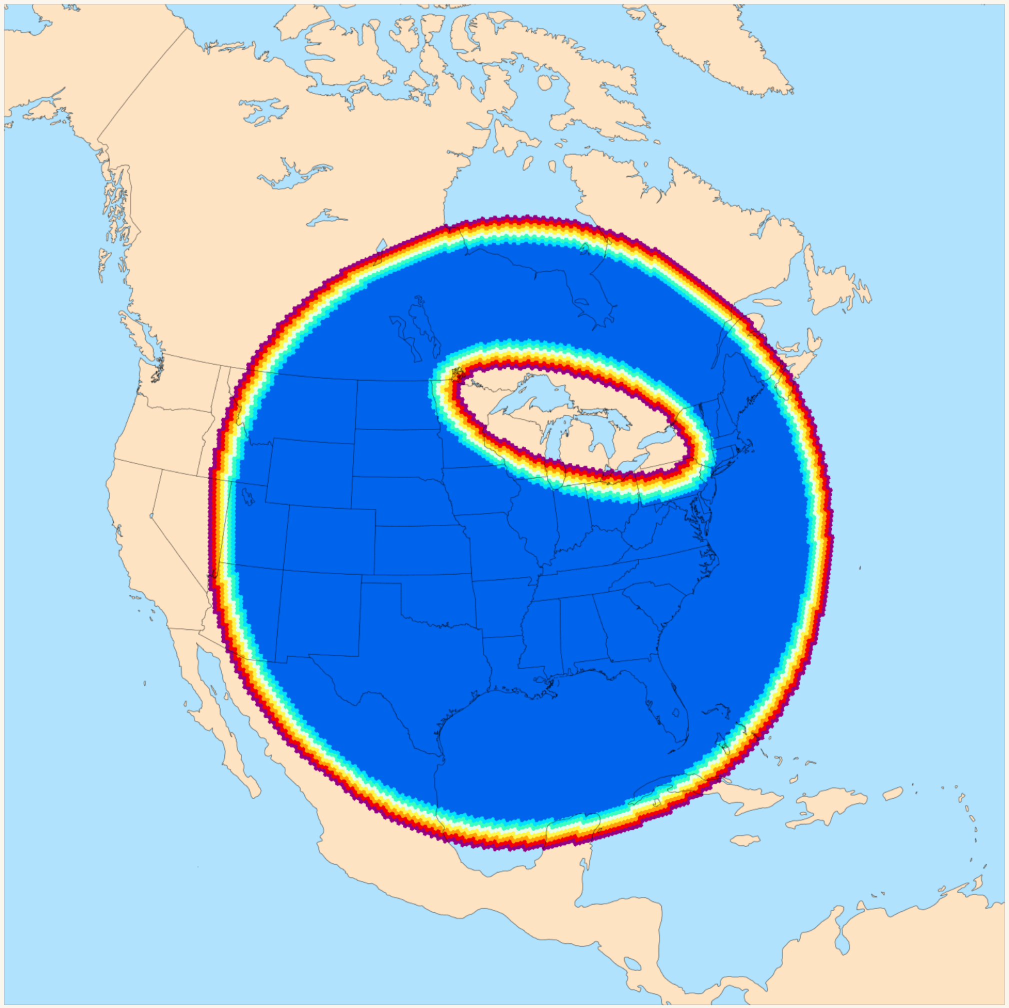

There's nothing at all in our regional implementation that precludes holes in domains…

CREATING LIMITED-AREA MESHES

With unstructured meshes, the creative user may define exotic simulation domains:

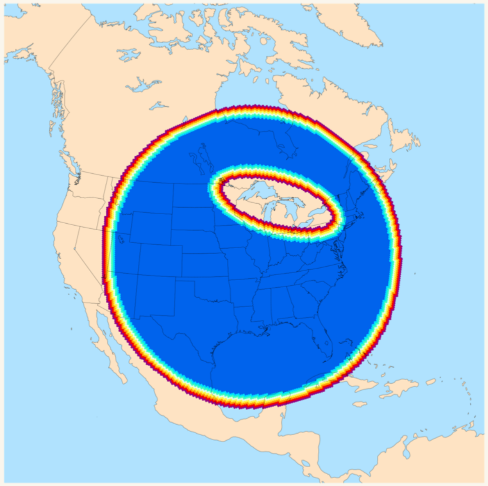

There's also no reason why domains must be connected…

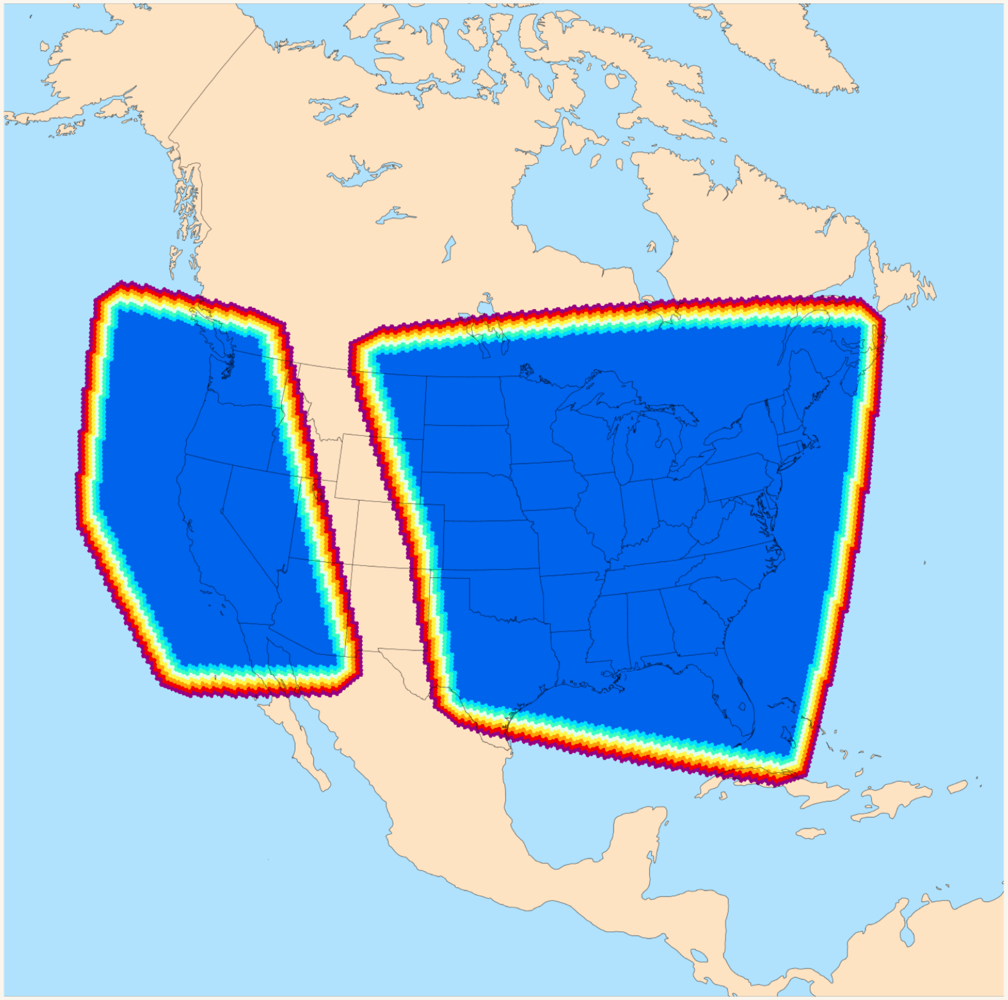

CREATING LIMITED-AREA MESHES

With unstructured meshes, the creative user may define exotic simulation domains:

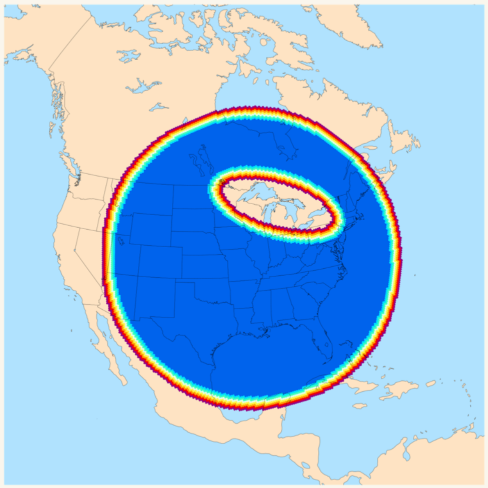

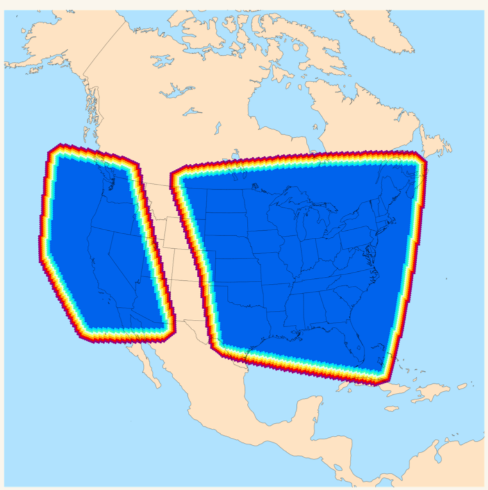

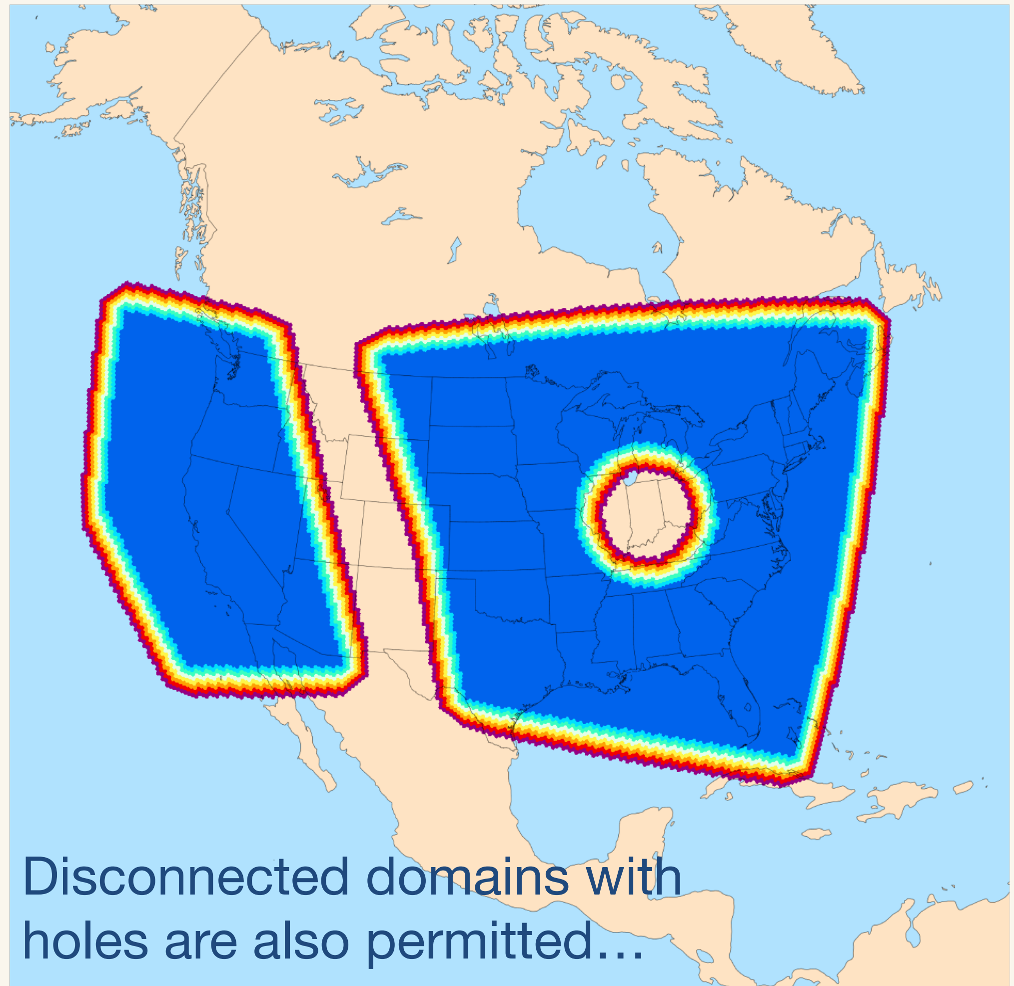

PARTITIONING OF REGIONAL MESHES

Partitioning limited-area meshes works the same as for global meshes

- Consider the connectivity graph of the cells in the mesh
- Divide the nodes in the graph into *k* partitions while minimizing the *edge cut* of the partitioning

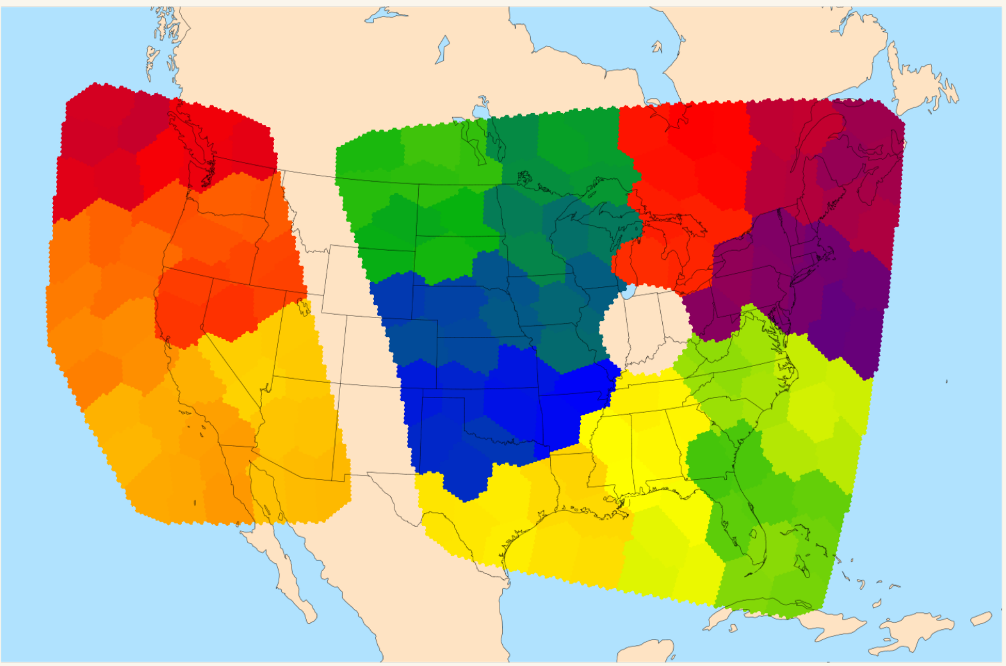

Partitioning of a regional MPAS mesh by Metis into 144 partitions.

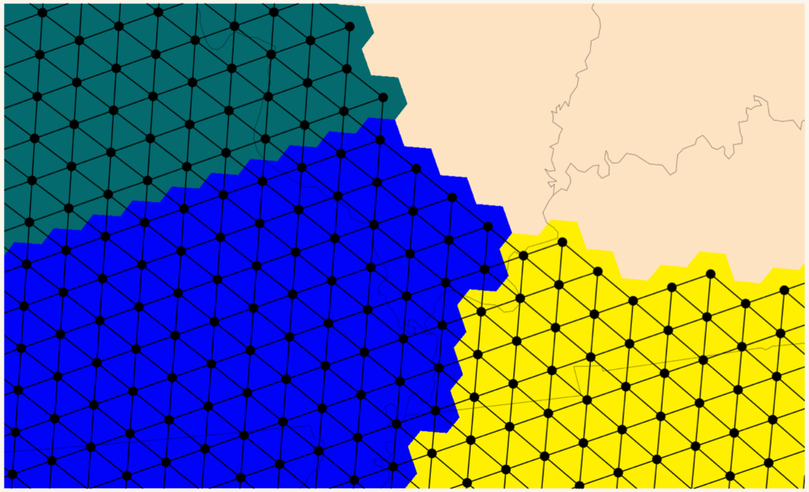

The connectivity graph of cells near the south-west of Indiana.

LBCS FROM THE USER'S PERSPECTIVE

Lateral boundary data are specified as uncoupled state (u, w, θ, ρ, q_x) at: T_0, $T_0+\Delta t_{LBC}$, $T_0+2\Delta t_{LBC}$, $T_0+3\Delta t_{LBC}$, …

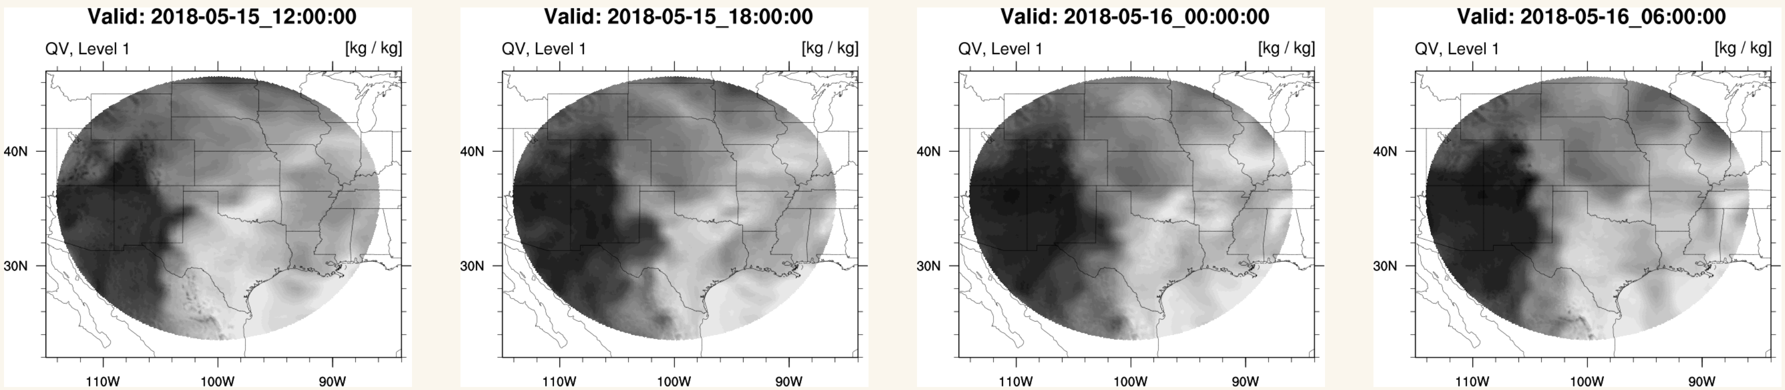

At present, complete fields (rather than just boundary elements) are stored and read

- Eventually, a capability to perform I/O on subsets of fields should be developed to boost efficiency of reading LBC data

HOW DOES THE COST OF REGIONAL MPAS COMPARE WITH WRF?

In order to compare WRF and regional MPAS, make some adjustments to default namelists:

config_time_integration_order = 3
config_split_dynamics_transport = false
config_number_of_sub_steps = 6
config_dynamics_split_steps = 1

time_step_sound = 6
h_mom_adv_order = 3
h_sca_adv_order = 3
moist_adv_opt = 2
scalar_adv_opt = 2

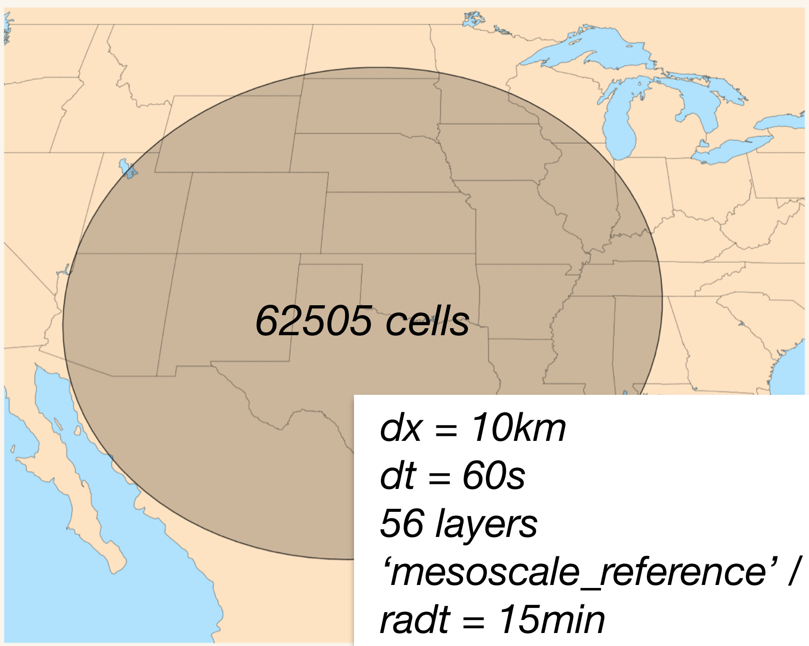

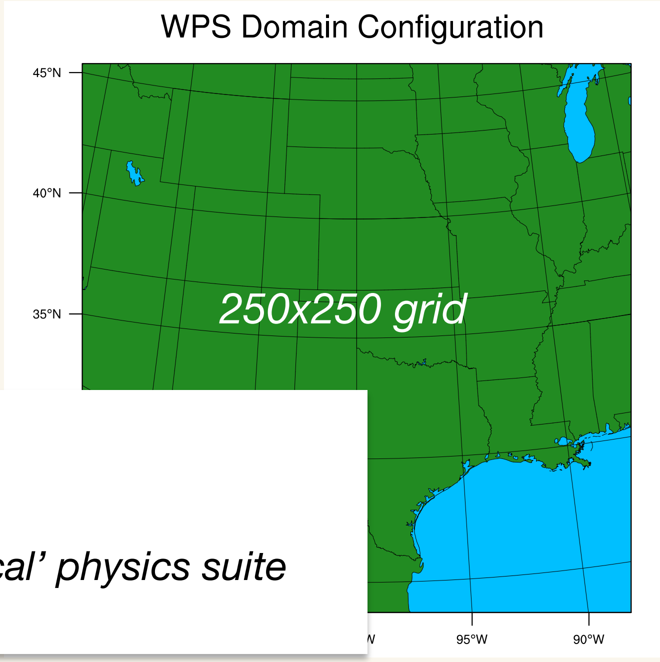

dx = 10km
dt = 60s
56 layers
'mesoscale_reference' / 'tropical' physics suite
radt = 15min

HOW DOES THE COST OF REGIONAL MPAS COMPARE WITH WRF?

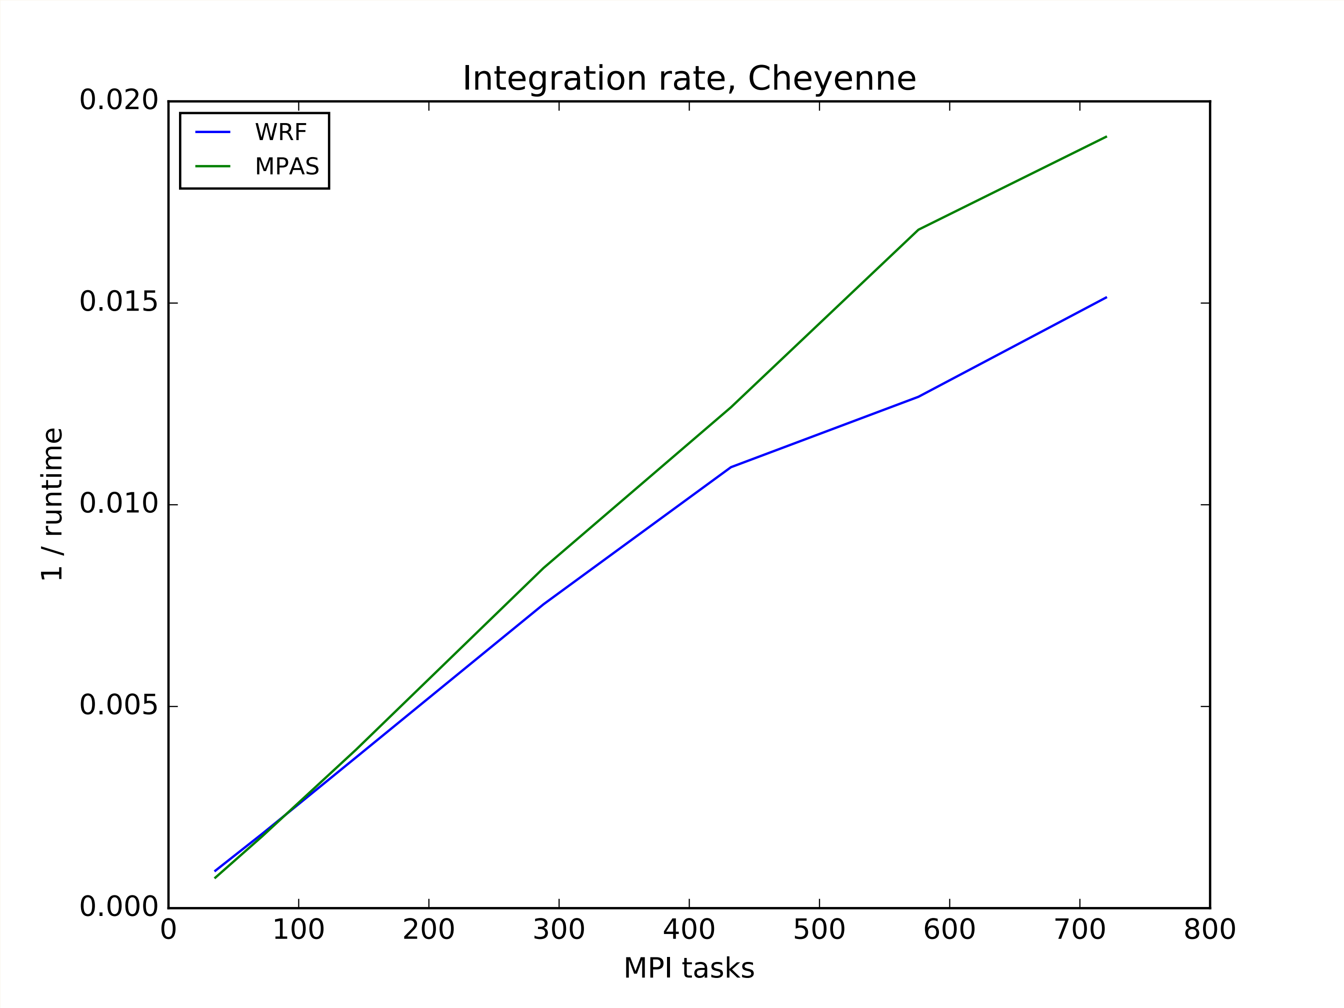

Integration rate (inverse of runtime) for 24-h WRF and MPAS simulations configured as on the previous slide on NCAR's Cheyenne machine.

HOW DOES THE COST OF REGIONAL MPAS COMPARE WITH WRF?

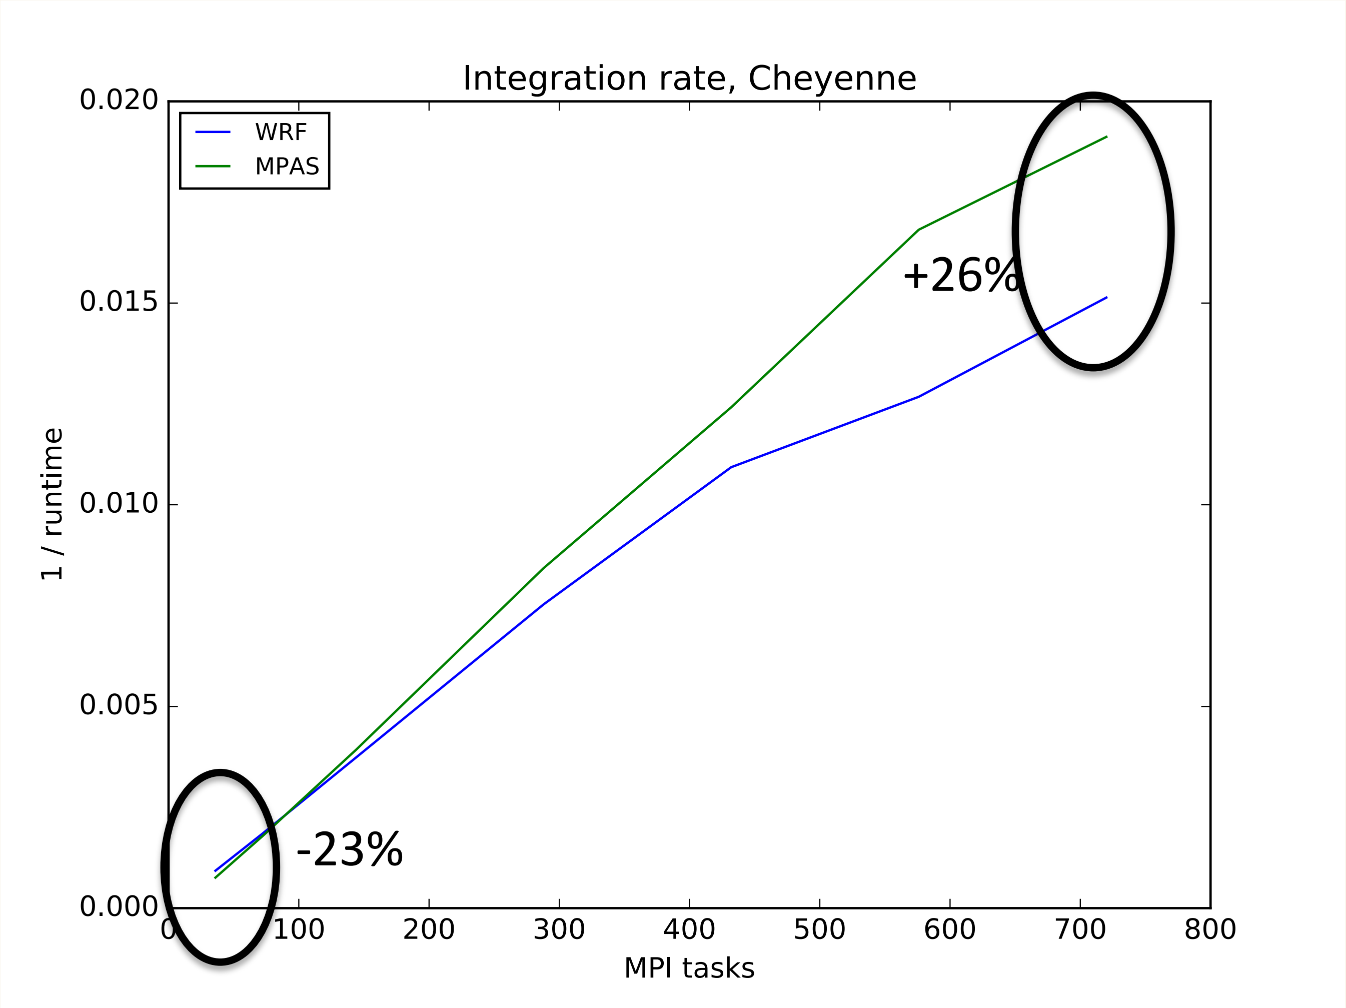

<u>Per column-timestep</u>, MPAS-Atmosphere can be roughly as efficient as WRF.

- *This says nothing about the merits of nesting v. variable-resolution!*

SIMPLE WORKFLOW: MPAS DRIVING MPAS

Several new programs are needed to subset global meshes and to interpolate LBC fields at present

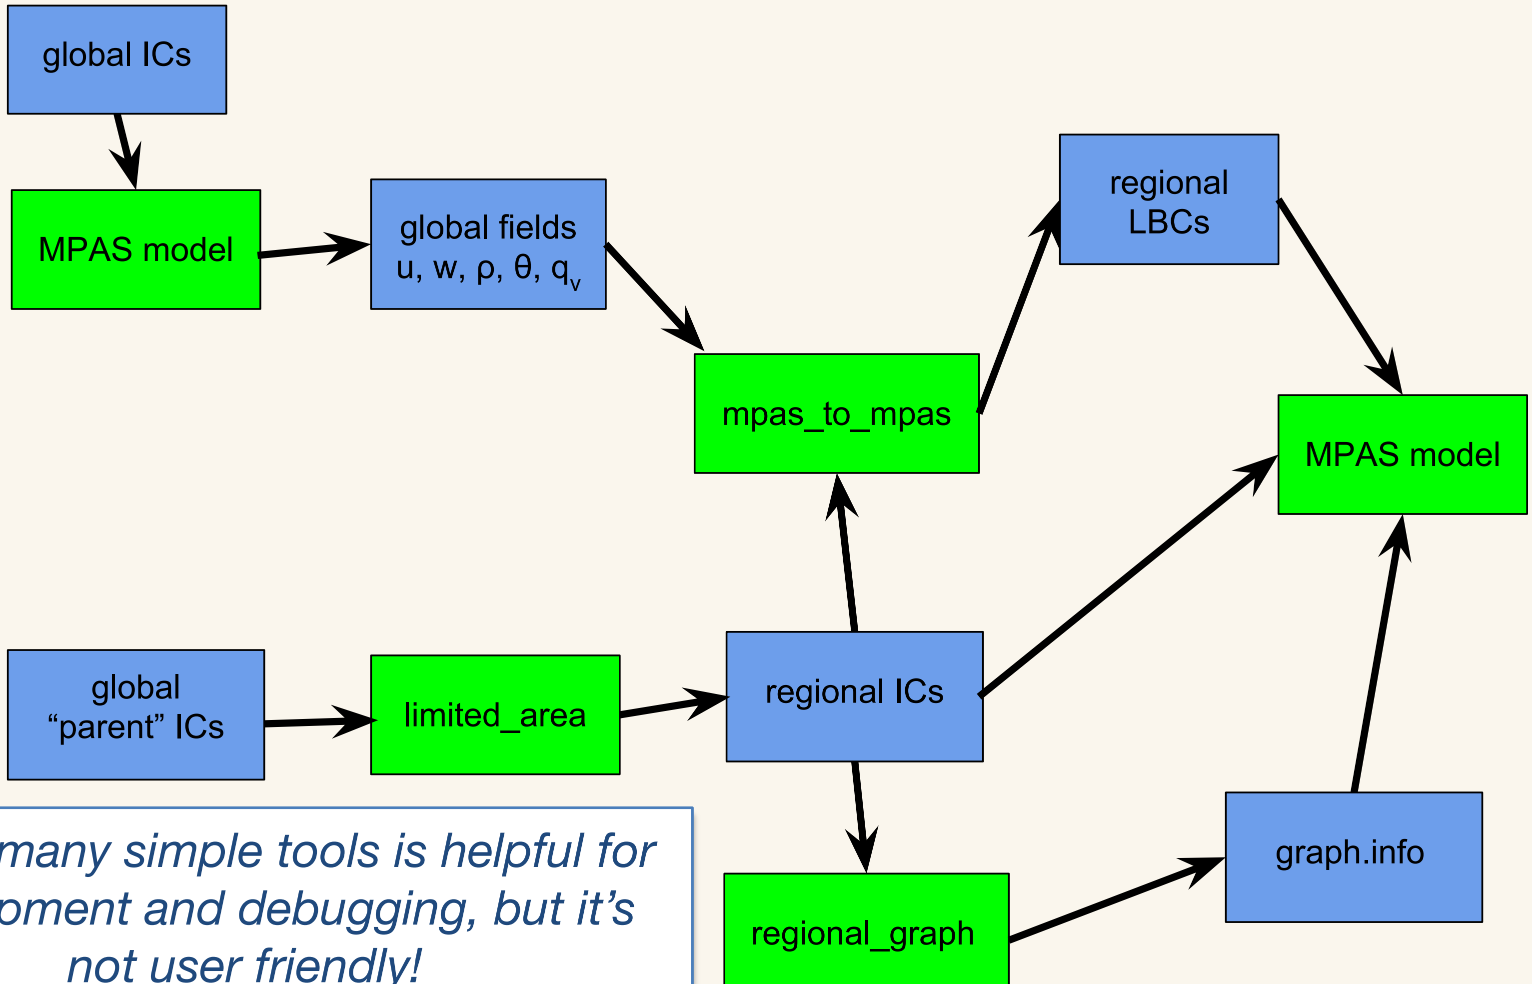

Having many simple tools is helpful for development and debugging, but it's not user friendly!

WORKFLOW: BLENDING BOUNDARY TERRAIN

Blending terrain along boundaries from the driving model is complicated by the Smoothed Terrain-following Coordinate (STF) used in MPAS

$$z = (z_t - A'h_s)\frac{\zeta}{z_t} + A'h_s = \zeta + Ah_s$$

(From Klemp (*MWR*, 2011))

$$h_s^{(n)} = h_s^{(n-1)} + \beta(\zeta)d^2\nabla_\zeta^2 h_s^{(n-1)}$$

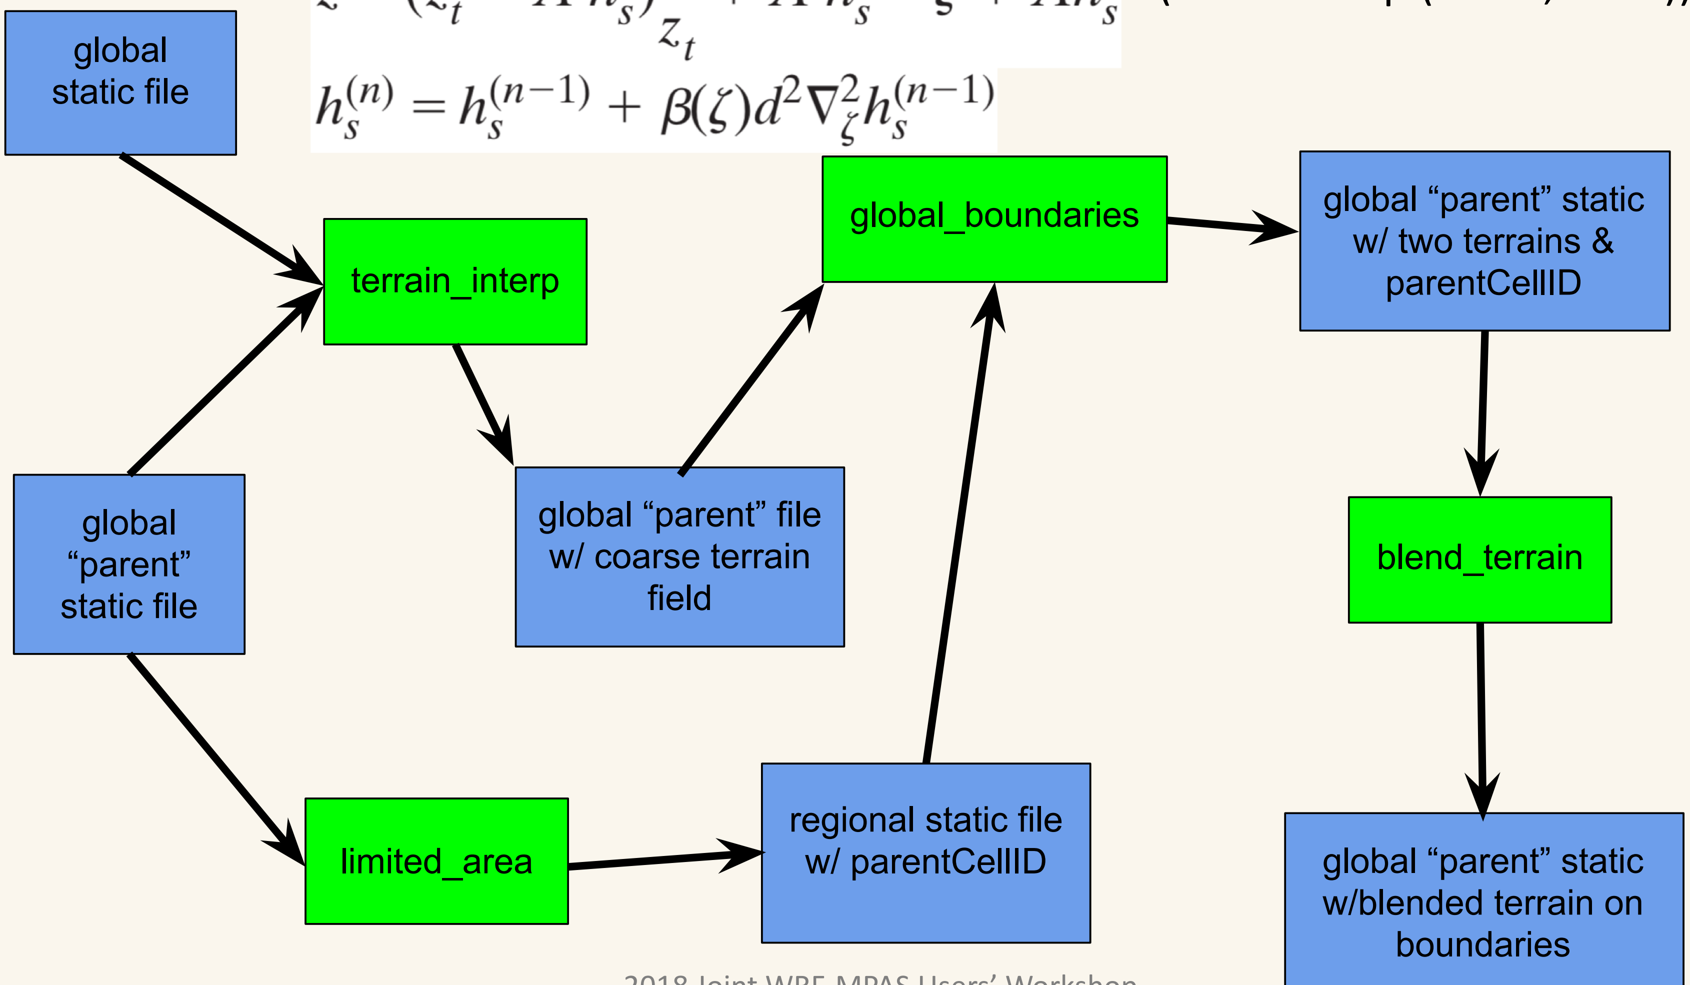

SOURCE OF LBCS? ANALYSES OR MPAS

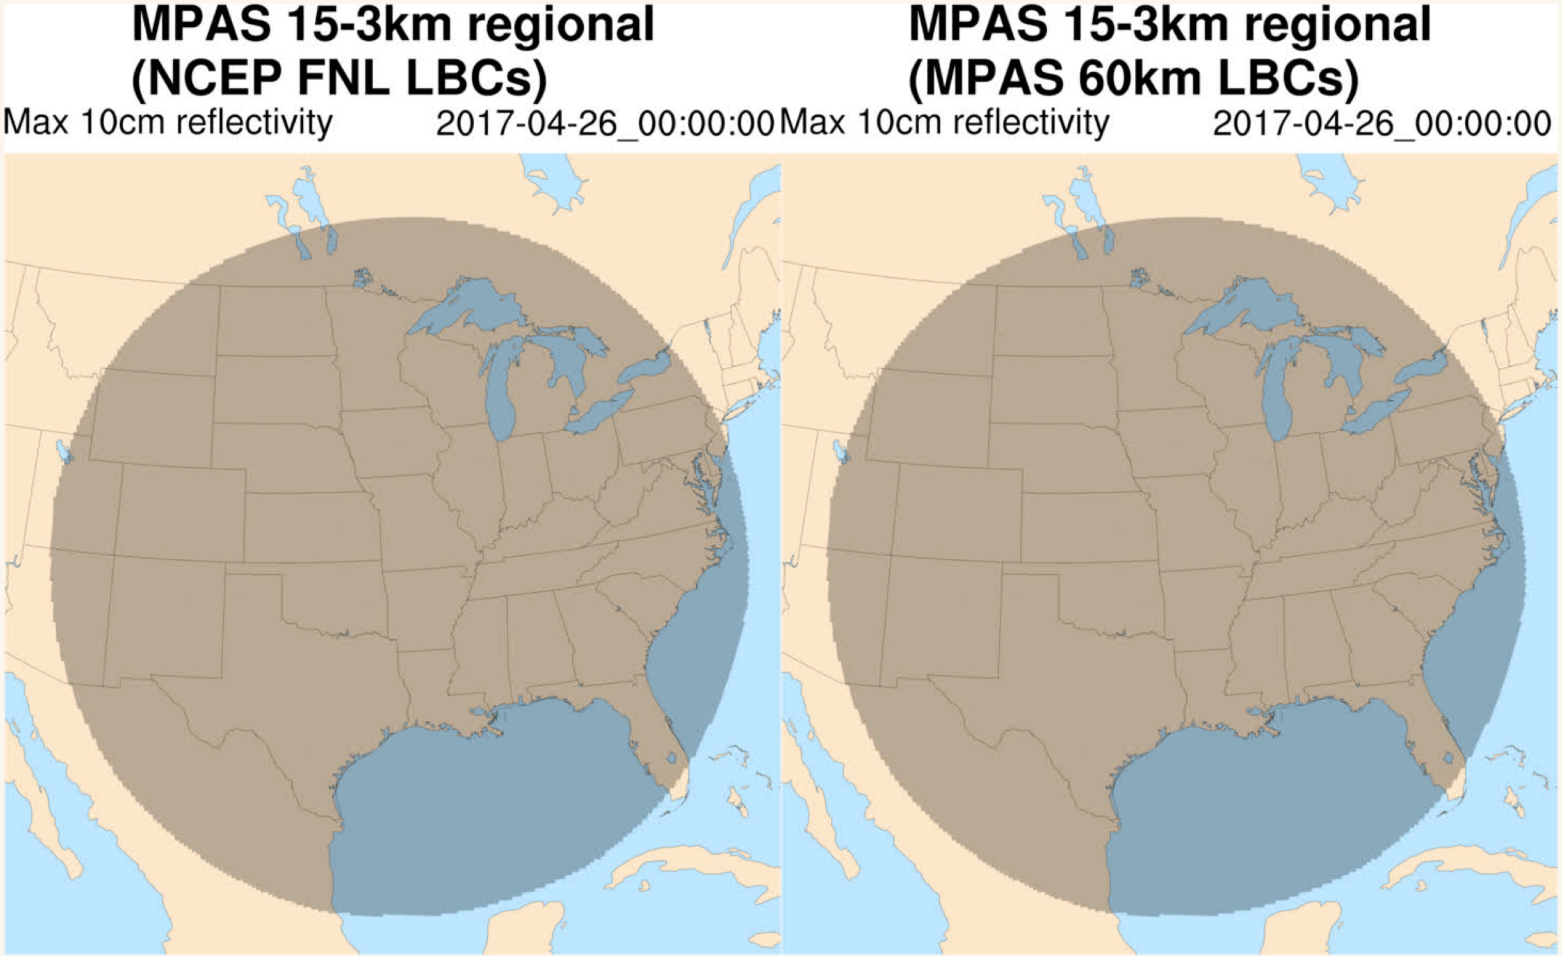

Lateral boundary conditions can be derived from analyses or from other MPAS simulations. Above, the FNL LBCs are updated every 6 hours, while the 60km MPAS LBCs are updated every 3 hours.

SUMMARY AND FUTURE WORK

Key points:

- Regional MPAS offers a tremendous amount of flexibility
- Cost of regional MPAS can be competitive with WRF
- Current working code is up-to-date with MPAS v6.0
- At this point, mostly code clean-up, documentation, and restructuring of pre-processing tools

- **We expect to make a publicly supported release of MPAS with a limited-area capability by the WRF-MPAS Joint Workshop next year**

Skamarock, W. C., M. G. Duda, S. Ha, and S.-H. Park, "Limited-Area Atmospheric Modeling Using an Unstructured Mesh", *Submitted to MWR*.

QUESTIONS?

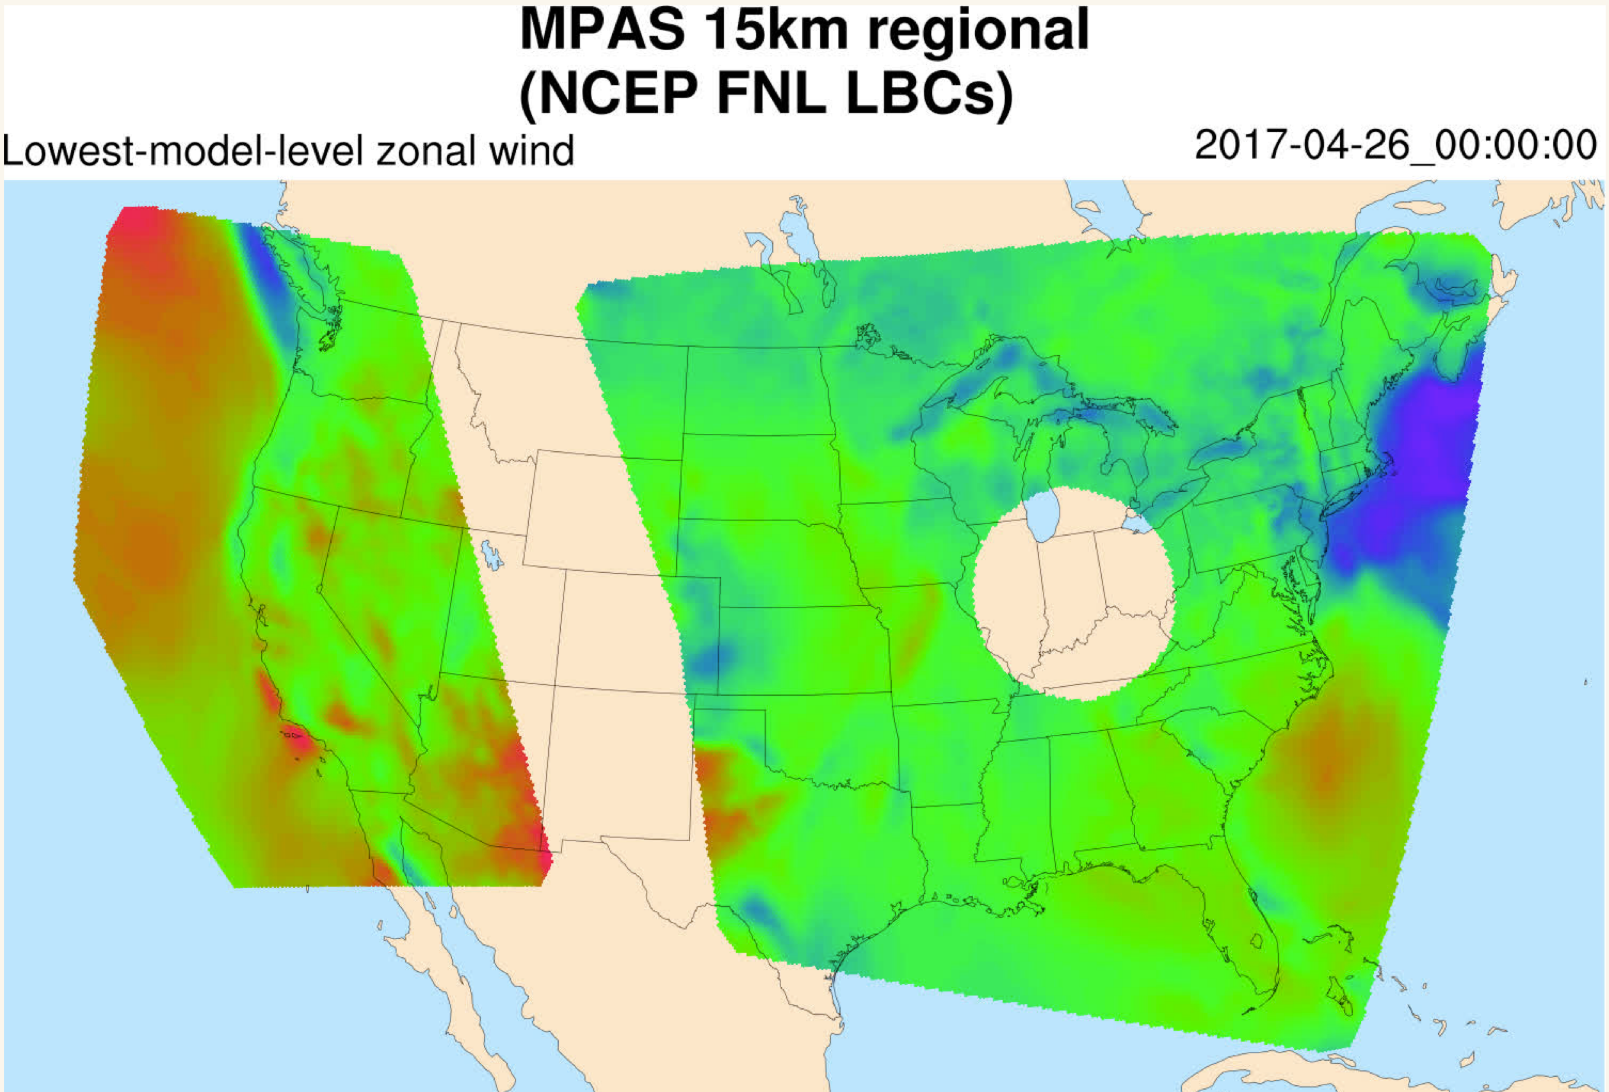
MPAS 15km regional
(NCEP FNL LBCs)
Lowest-model-level zonal wind
2017-04-26_00:00:00

EXTRA SLIDES

MOTIVATION

We already have WRF; why do we need a regional MPAS capability?

1. Allow less costly testing of MPAS at high resolutions
2. Leverage MPAS development for next-generation architectures for regional applications
3. Enable regional atmospheric applications within MPAS-enabled coupled modeling systems (e.g., CESM)
4. Employ variable resolution in regional applications to reduce LBC errors
5. Provide a consistent (equations, mesh) regional solver to complement global MPAS
6. Groups external to NCAR have asked for a regional capability in MPAS and are supporting its development (KISTI)

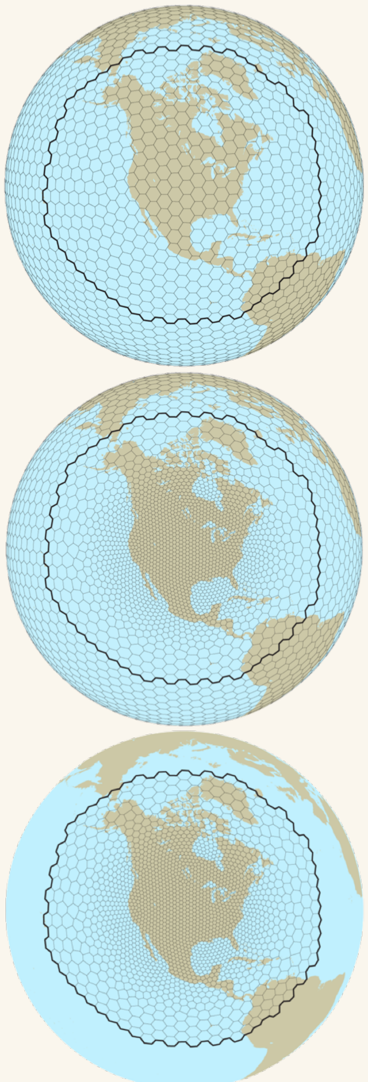

LBCS FROM THE USER'S PERSPECTIVE

Lateral boundary data are specified as uncoupled state at:

T_0, $T_0+\Delta t_{LBC}$, $T_0+2\Delta t_{LBC}$, $T_0+3\Delta t_{LBC}$, …

```
dimensions:
    Time = UNLIMITED ; // (1 currently)
    StrLen = 64 ;
    nEdges = 2410 ;
    nVertLevels = 41 ;
    nCells = 773 ;
    nVertLevelsP1 = 42 ;
variables:
    char xtime(Time, StrLen) ;
    float lbc_u(Time, nEdges, nVertLevels) ;
    float lbc_w(Time, nCells, nVertLevelsP1) ;
    float lbc_rho(Time, nCells, nVertLevels) ;
    float lbc_theta(Time, nCells, nVertLevels) ;
    float lbc_qv(Time, nCells, nVertLevels) ;
    …
}
```

LBCS FROM THE USER'S PERSPECTIVE

Boundary state is read at specified boundary interval from a standard MPAS *stream*, tendencies are internally computed

```
<immutable_stream name="lbc_in"
                  type="input"
                  filename_template="CONUS-LBC.$Y$M$D$h.nc"
                  filename_interval="input_interval"
                  input_interval=”3:00:00" />
```

Ideally, we could have a flag in the IC file that tells MPAS whether to activate the reading and enforcing of LBCs

- In our working code, this is currently hard-coded

COMPUTATIONAL OVERHEAD

Enforcing LBCs certainly comes with a cost

- Might there be significant load imbalance introduced by lateral boundaries (relaxation zone + specified zone)?

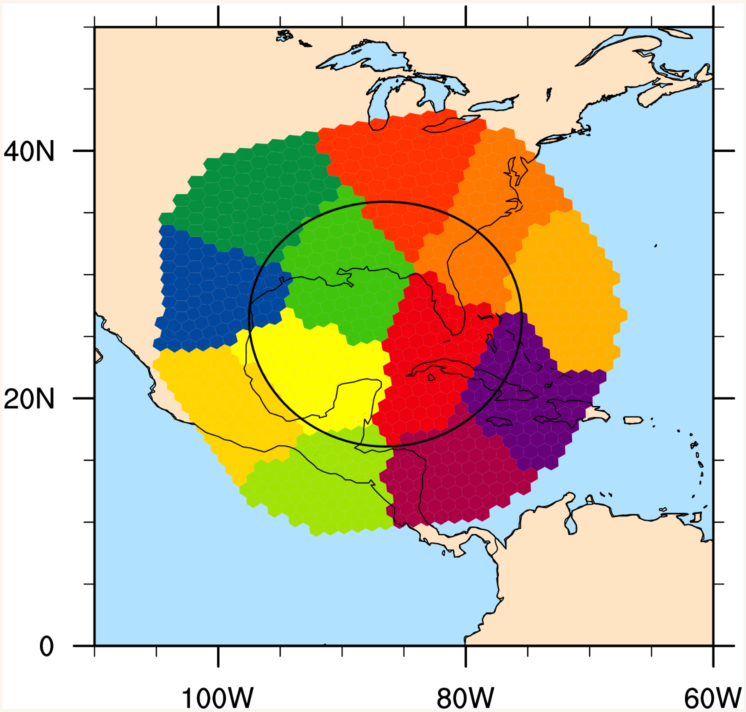

Partitioning of the mesh for parallel model execution is performed analogously to what is done for global MPAS meshes.

COMPUTATIONAL OVERHEAD

In the worst case (i.e., an MPI task contains all boundary points), we can expect less than ~6% increase in runtime for limited-area MPAS

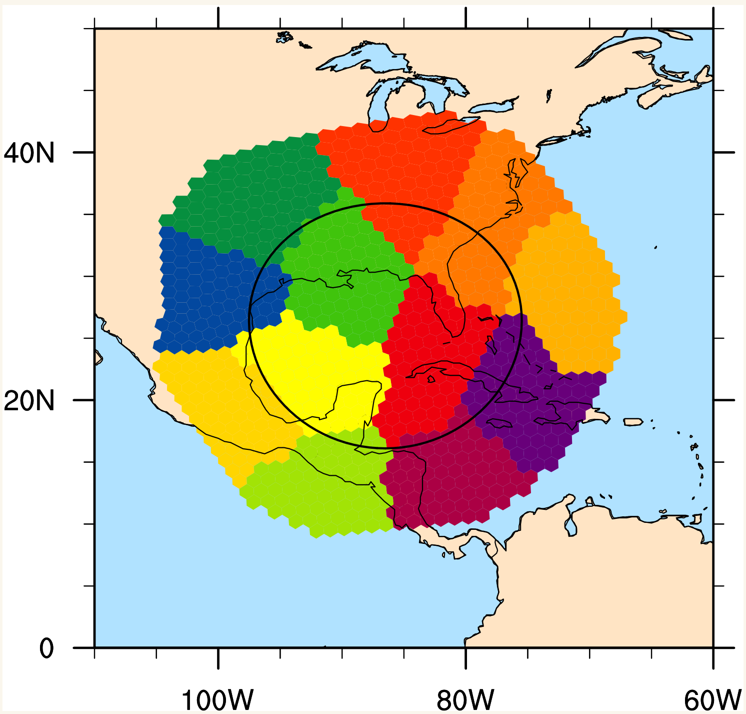

Partitioning of the mesh for parallel model execution is performed analogously to what is done for global MPAS meshes.

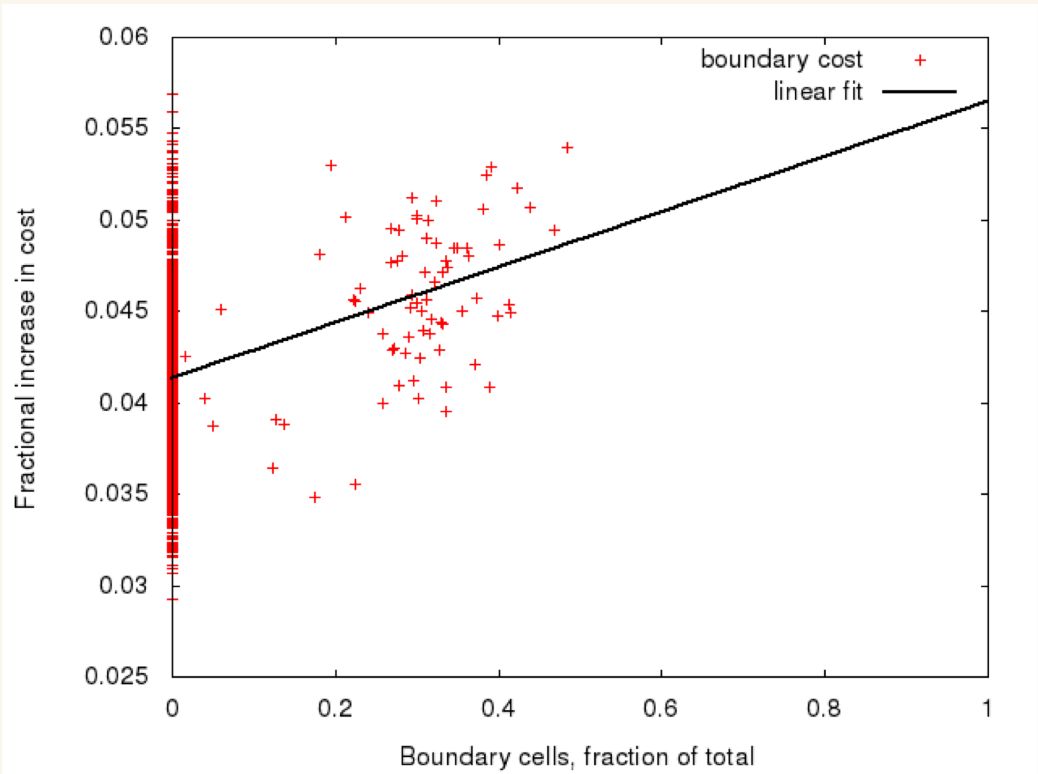

Measurements of relative cost of enforcing boundary conditions from an example simulation

COMPUTATIONAL OVERHEAD

In the worst case (i.e., an MPI task contains all boundary points), we can expect less than ~6% increase in runtime for limited-area MPAS

- We should be able to improve this significantly for the common case…

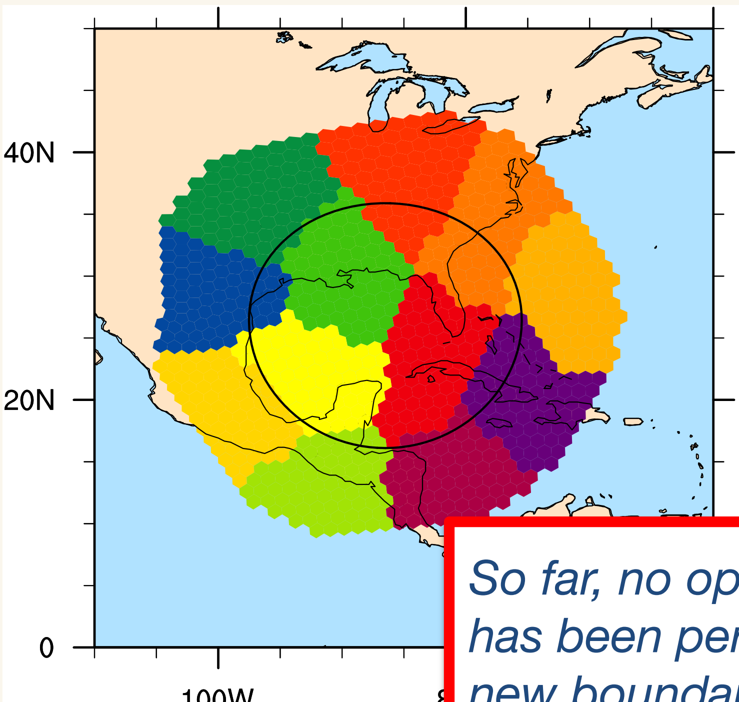

Partitioning of the mesh for parallel model execution is performed analogously to what is done for global MPAS meshes.

So far, no optimization at all has been performed on the new boundary code!

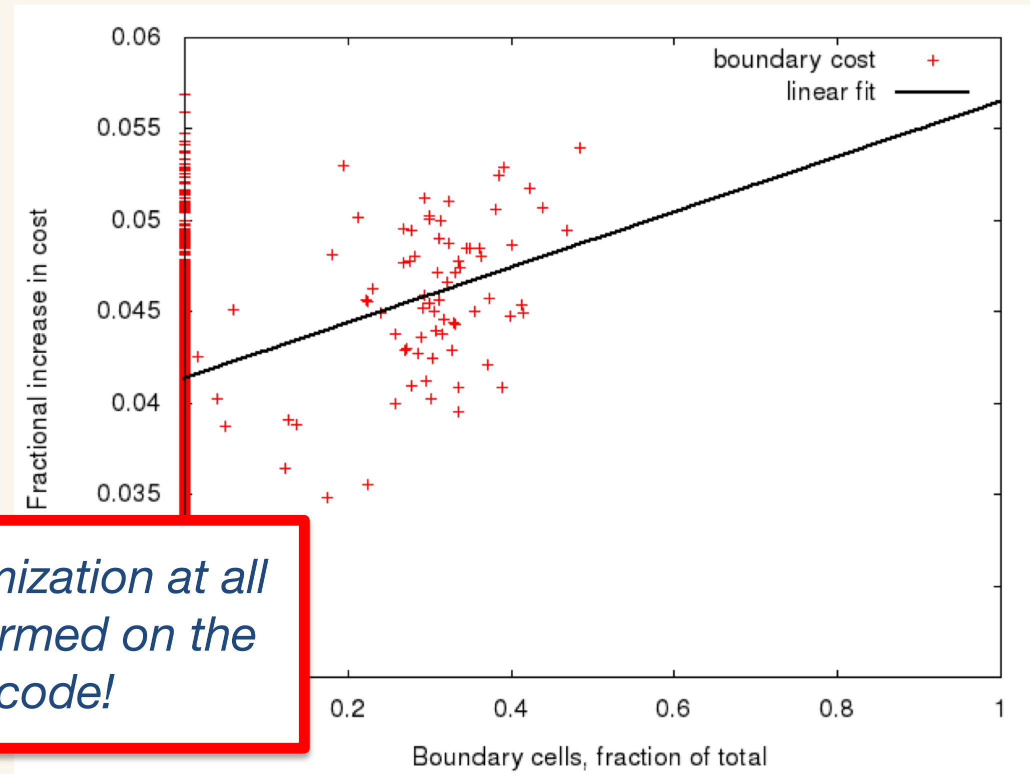

Measurements of relative cost of enforcing boundary conditions from an example simulation

HOW DOES THE COST OF REGIONAL MPAS COMPARE WITH WRF?

Lower MPAS performance for small MPI task counts may be explainable by cache effects

- Cells in MPAS meshes are essentially randomly ordered
- When each MPI task "owns" a large block of cells, we see poor cache reuse

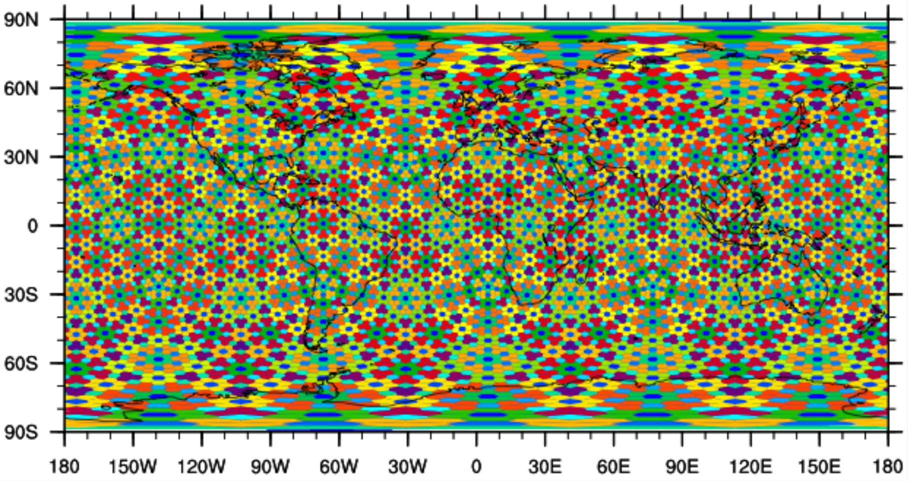

Cells of a 10242 mesh colored according to their global index

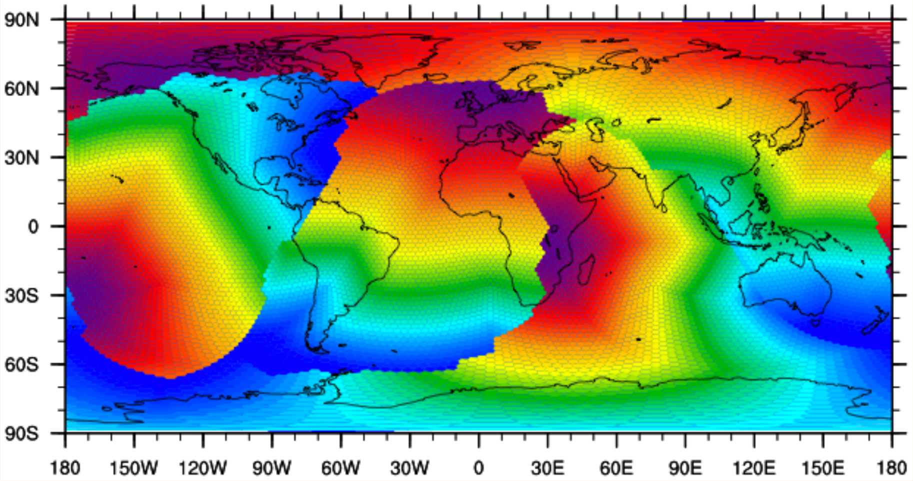

Cells of a 10242 mesh decomposed across 4 MPI tasks and colored according to their re-ordered local index