P: Advances in the application of parallel split physics and dynamics

Aaron S. Donahue and Peter M. Caldwell
Lawrence Livermore National Laboratory

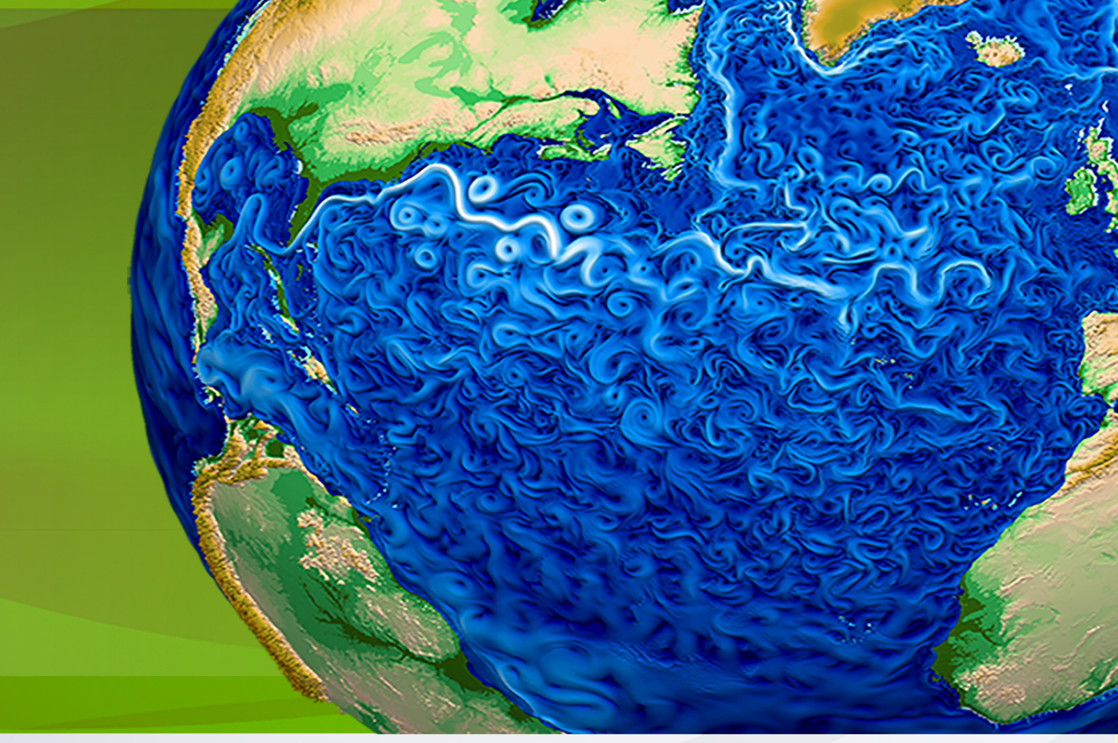

Opportunity

Improved performance through a change in coupling paradigm:
Dynamics and physics are loosely coupled in the ACME model through sequential-tendency splitting (STS). This has proven to be a limiting factor for scalability of the model, and application of parallel-splitting (PS) provides an opportunity to break past this limit and improve model performance.

Problem Description

Dynamics elements vs. physics columns:

- Dynamics and physics are solved on similar but different domains: a cubed sphere spectral element mesh for dynamics, and a set of individual columns for the physics mesh. See Figure 1.
  - The number of physics columns is greater than the number of dynamics elements by a factor of 9
- For standard run setups, dynamics and physics are roughly the same computational cost, see Figure 2.
  - New developments such as super-parameterization may make physics more expensive than dynamics in which case parallel-splitting could be critical for improved throughput.

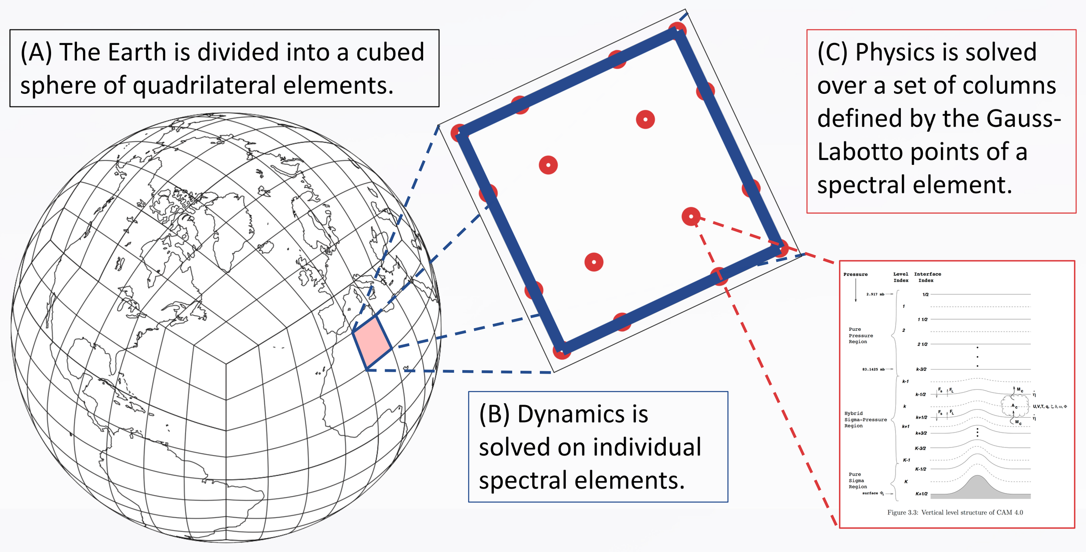

Figure 1: Dynamics and physics domains for the ACME model. (A) cubed sphere, (B) example spectral element, (C) example physics column. Image credit: Dennis *et al.* (2012) Int. J. of High Performance Computing Applications (A and B) and Neale *et al.* (2010) CAM 4.0 (C)

Thus, being able to use more processors than dynamics elements should improve performance!

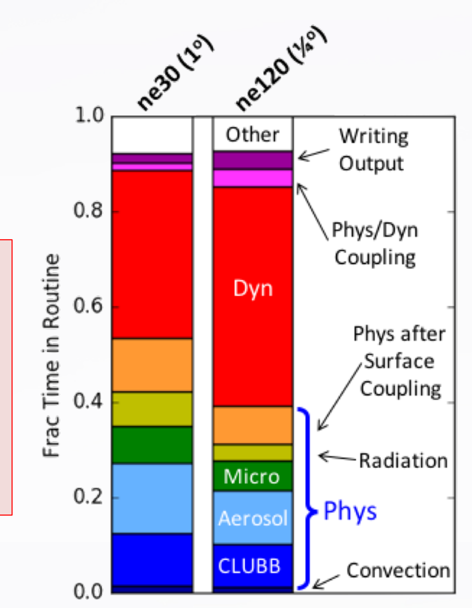

Figure 2: Fraction of ACME v1alpha7 integration time spent in various processes.

Scalability: to the elements…and beyond

Pushing core counts past element counts:
The ACME model scales up to the number of elements, see Figure 3. For our current suite of meshes this means:

- 5.4K cores for the 1 degree mesh (ne30), and
- 86.4K cores for the ¼ degree mesh (ne120).

As HPCs get larger so does the limit of available resources:

- ORNL-Titan has 299K cores
- NERSC-Edison has 134K cores
- NERSC-Cori has 659K cores

If we can scale to the number of physics columns we can use an order of magnitude more cores,. allowing us to utilize these and future resources more fully

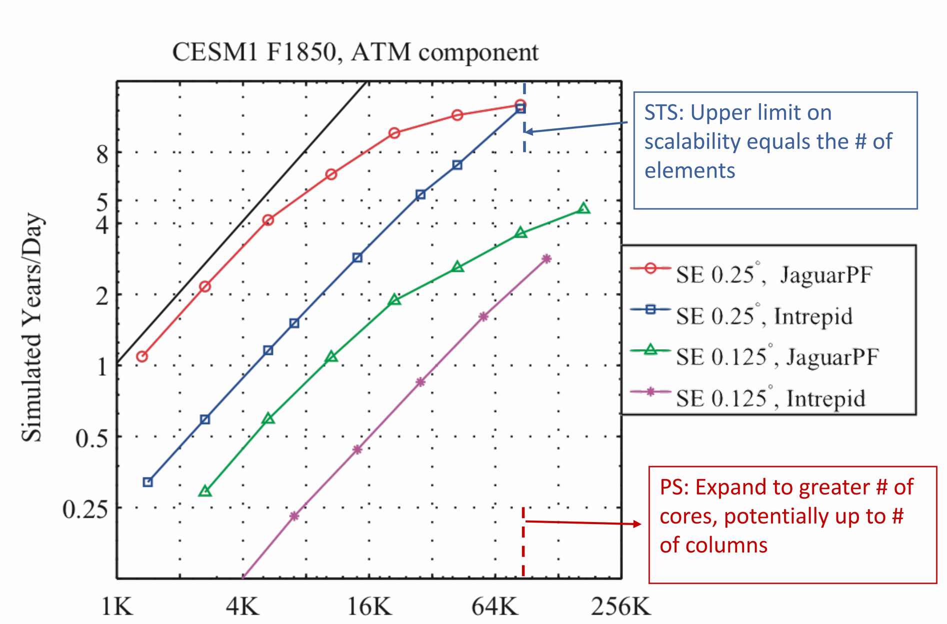

Figure 3: Scalability of CAM-SE, Dennis *et al.*, "CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model" (2012), Int. J. of High Performance Computing Applications.

Preliminary Issues and Solutions

Solution Fidelity:
Switching to a parallel-split paradigm will change the model solution. Does this result in a degraded solution?
No: Analysis of 10 year climatologies for STS and PS show very little change in the solution for the full suite of variables produced by the AMWG diagnostic toolkit, see example in Figure 4.

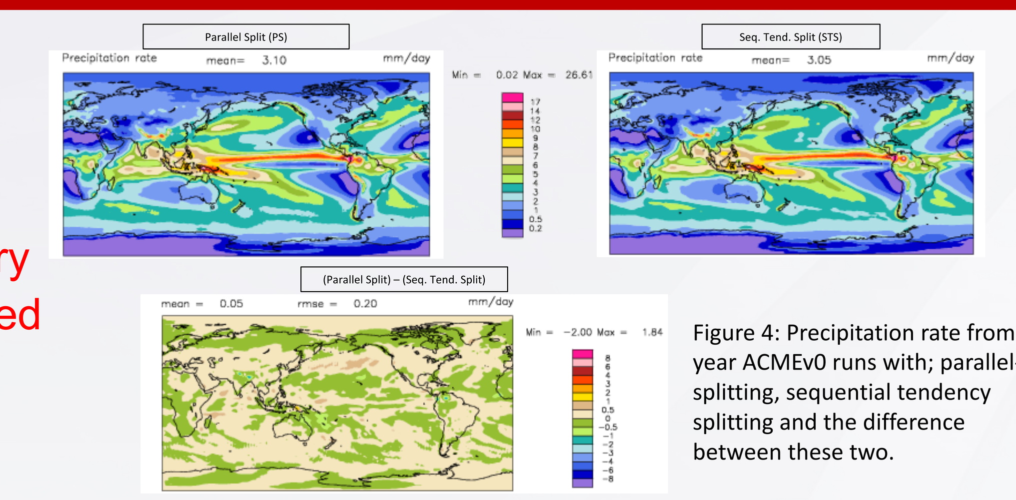

Figure 4: Precipitation rate from 10 year ACMEv0 runs with; parallel-splitting, sequential tendency splitting and the difference between these two.

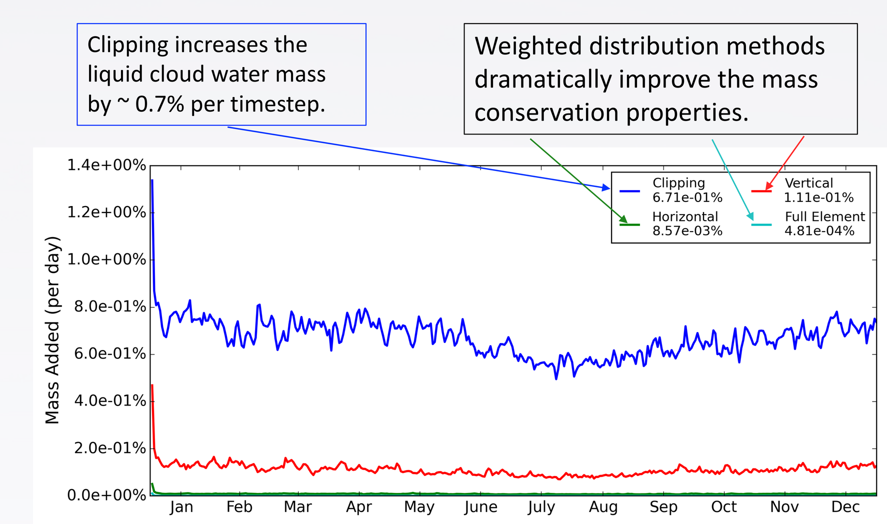

Figure 5: Daily average mass conservation corrections using parallel-split with mass conservation techniques of clipping and horizontal, vertical or full element mass redistribution.

Mass Conservation:
A drawback to parallel-splitting is the potential for local negative mass with dynamics and physics drawing at the same resource independently, the general approach of clipping these values to zero arbitrarily increases the global mass. See Figure 5.
Redistribution of mass over an element either horizontally, vertically, or over the full element provides a tool for addressing mass conservation issues.

Issues and Proposed Solutions

Mission accomplished? Not yet:
In our initial simultaneous dynamics and physics implementation, we do not observe improved scalability at high core counts, Figure 6. Why is this?

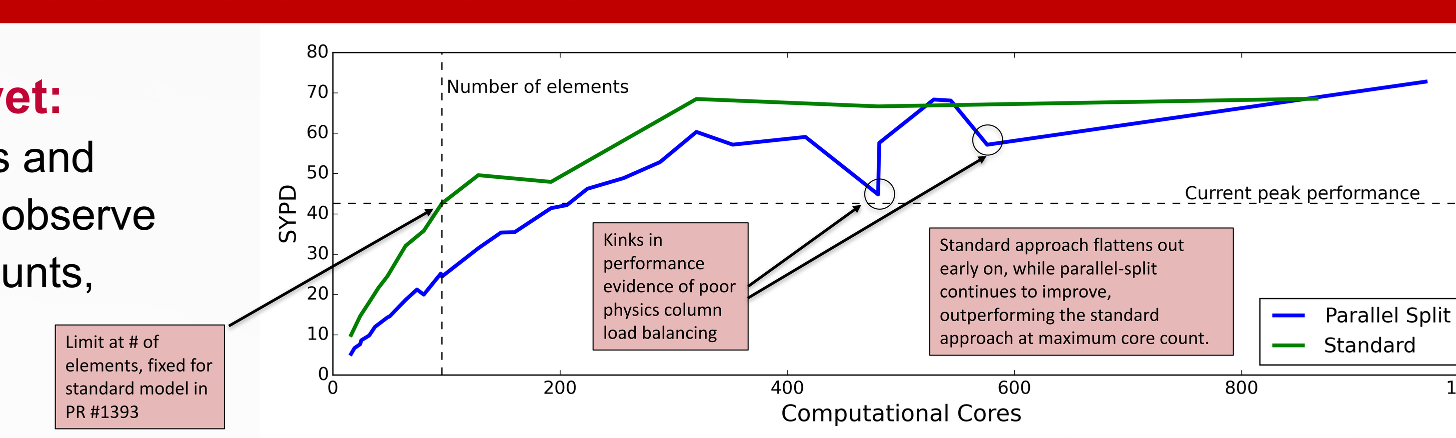

Figure 6: Solution timings for 10-day simulations on the 7.5 degree (ne4) mesh for the standard model and for parallel-split implementation.

Inter-core communication:
Improved performance in terms of solving physics and dynamics separately has been traded for increased communication between dynamics and physics cores, at a cost to scalability.

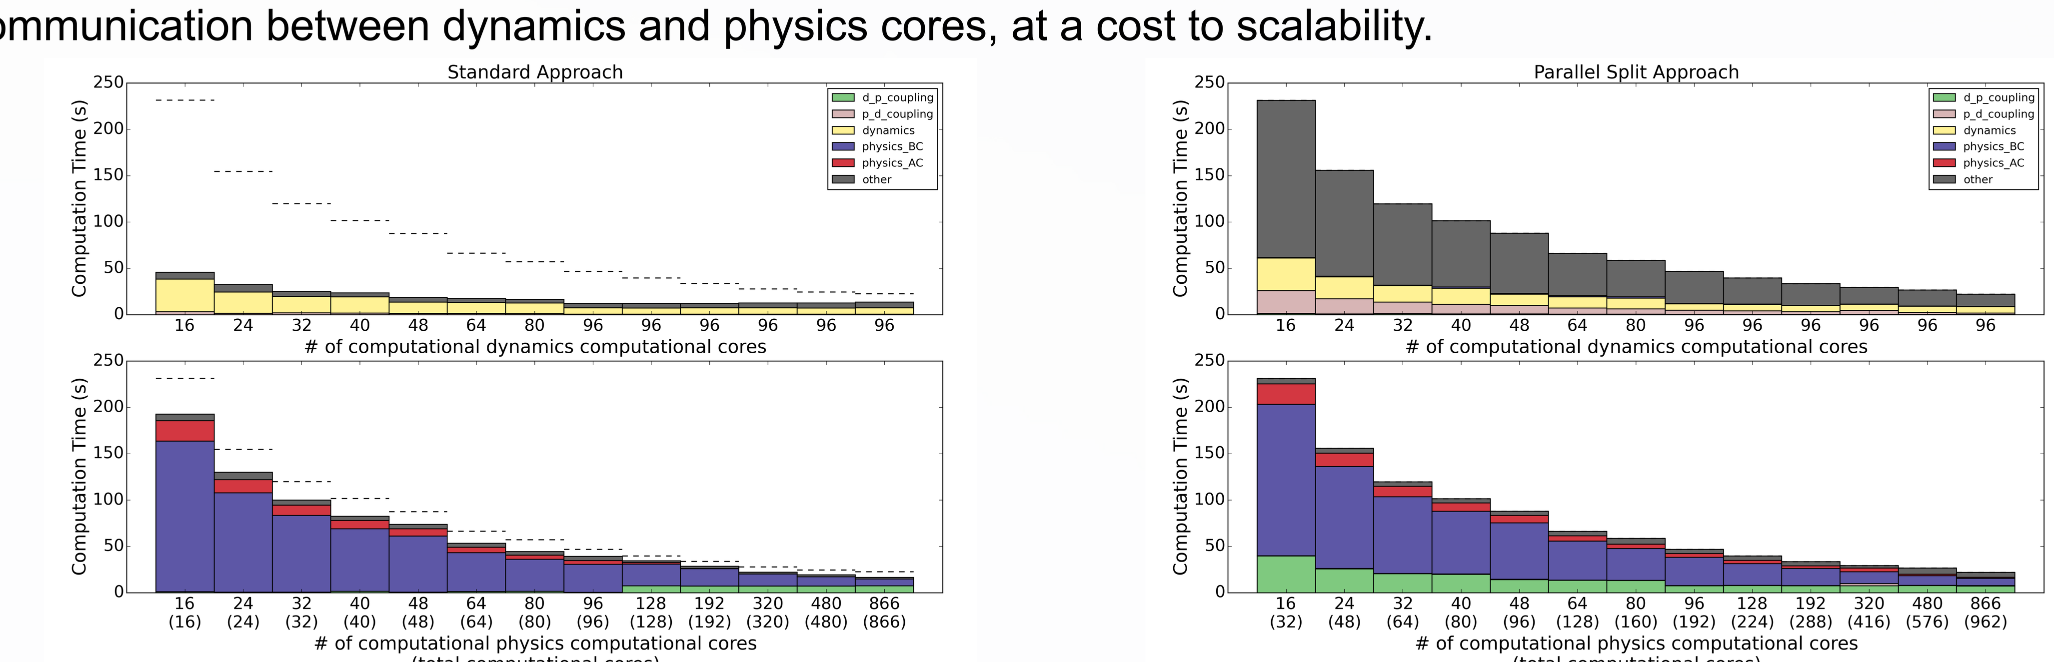

Figure 7: Average computational cost per core for dynamics (dynamics), physics (physics AC and BC) and dynamics-physics communication (d_p and p_d coupling) for the standard ACME model (left) and the parallel-split implementation (right). Top panels represent only the cores assigned to dynamics, bottom panels are cores assigned to physics.

Dynamics/physics balancing per node:
A proposed solution to minimize inter-core communication cost is to balance the assignment of a fraction of cores per node to dynamics and the rest to physics such that most communication between physics and dynamics is done within a compute node. This work is currently underway.

This work funded by the ACME Software Modernization CMDV project

ACME Accelerated Climate Modeling for Energy

For additional information, contact:
Aaron Donahue
Scientist
Lawrence Livermore National Laboratory
(925) 423-2572
donahue5@llnl.gov

climatemodeling.science.energy.gov/acme

U.S. DEPARTMENT OF ENERGY

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 IM Release Number LLNL-POST-732406