Sandia National Laboratories

Exceptional service in the national interest

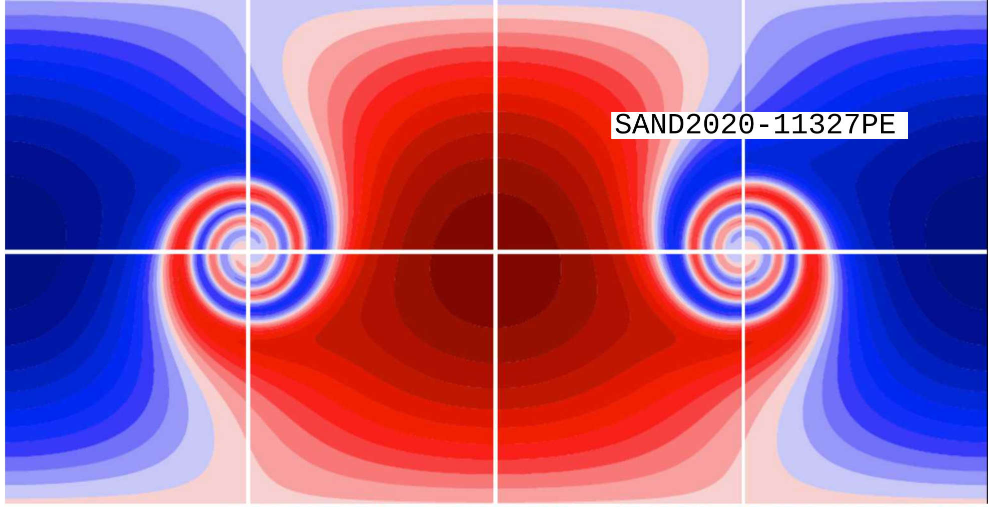

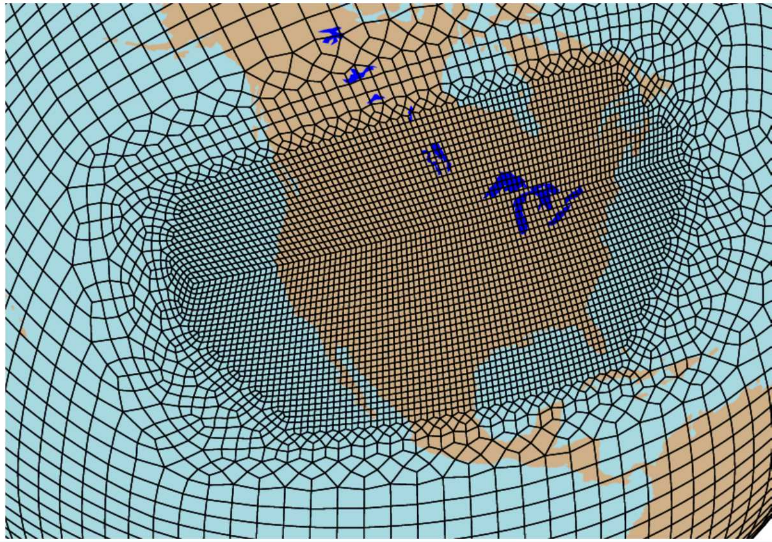

Computing challenges in weather and climate modeling

PRESENTED BY

Peter A. Bosler, Ph.D.

WITH ANDREW BRADLEY, OKSANA GUBA, AND MARK TAYLOR + MANY OTHERS

October 29, 2020

CCR Center for Computing Research

U.S. DEPARTMENT OF ENERGY | Office of Science

Outline

- Overview of various challenges:
  - Physical
  - Computing
  - A sea story
- Climate projection vs. weather forecasting
- Paths forward require better algorithms
  - Our approach: Compact, high order data; large time steps
  - 2 Semi-Lagrangian methods tailored for DOE's Exascale Earth System Model's (E3SM) Atmosphere component
  - A new shape preservation filter: Communication-Efficient Density Reconstruction
  - Upwind communication patterns
  - Reduce computations based on effective resolution: Grid coarsening
- **Combined effect:** We've **doubled** the computation speed of the E3SM Atmosphere Model (EAM)

Key point: Models do well

- Translating model error to the general public is quite difficult.
- Sometimes millions of people live between two model gridpoints.
- **Easy to focus on modeling errors and imperfections, but the trend is toward improved fidelity**

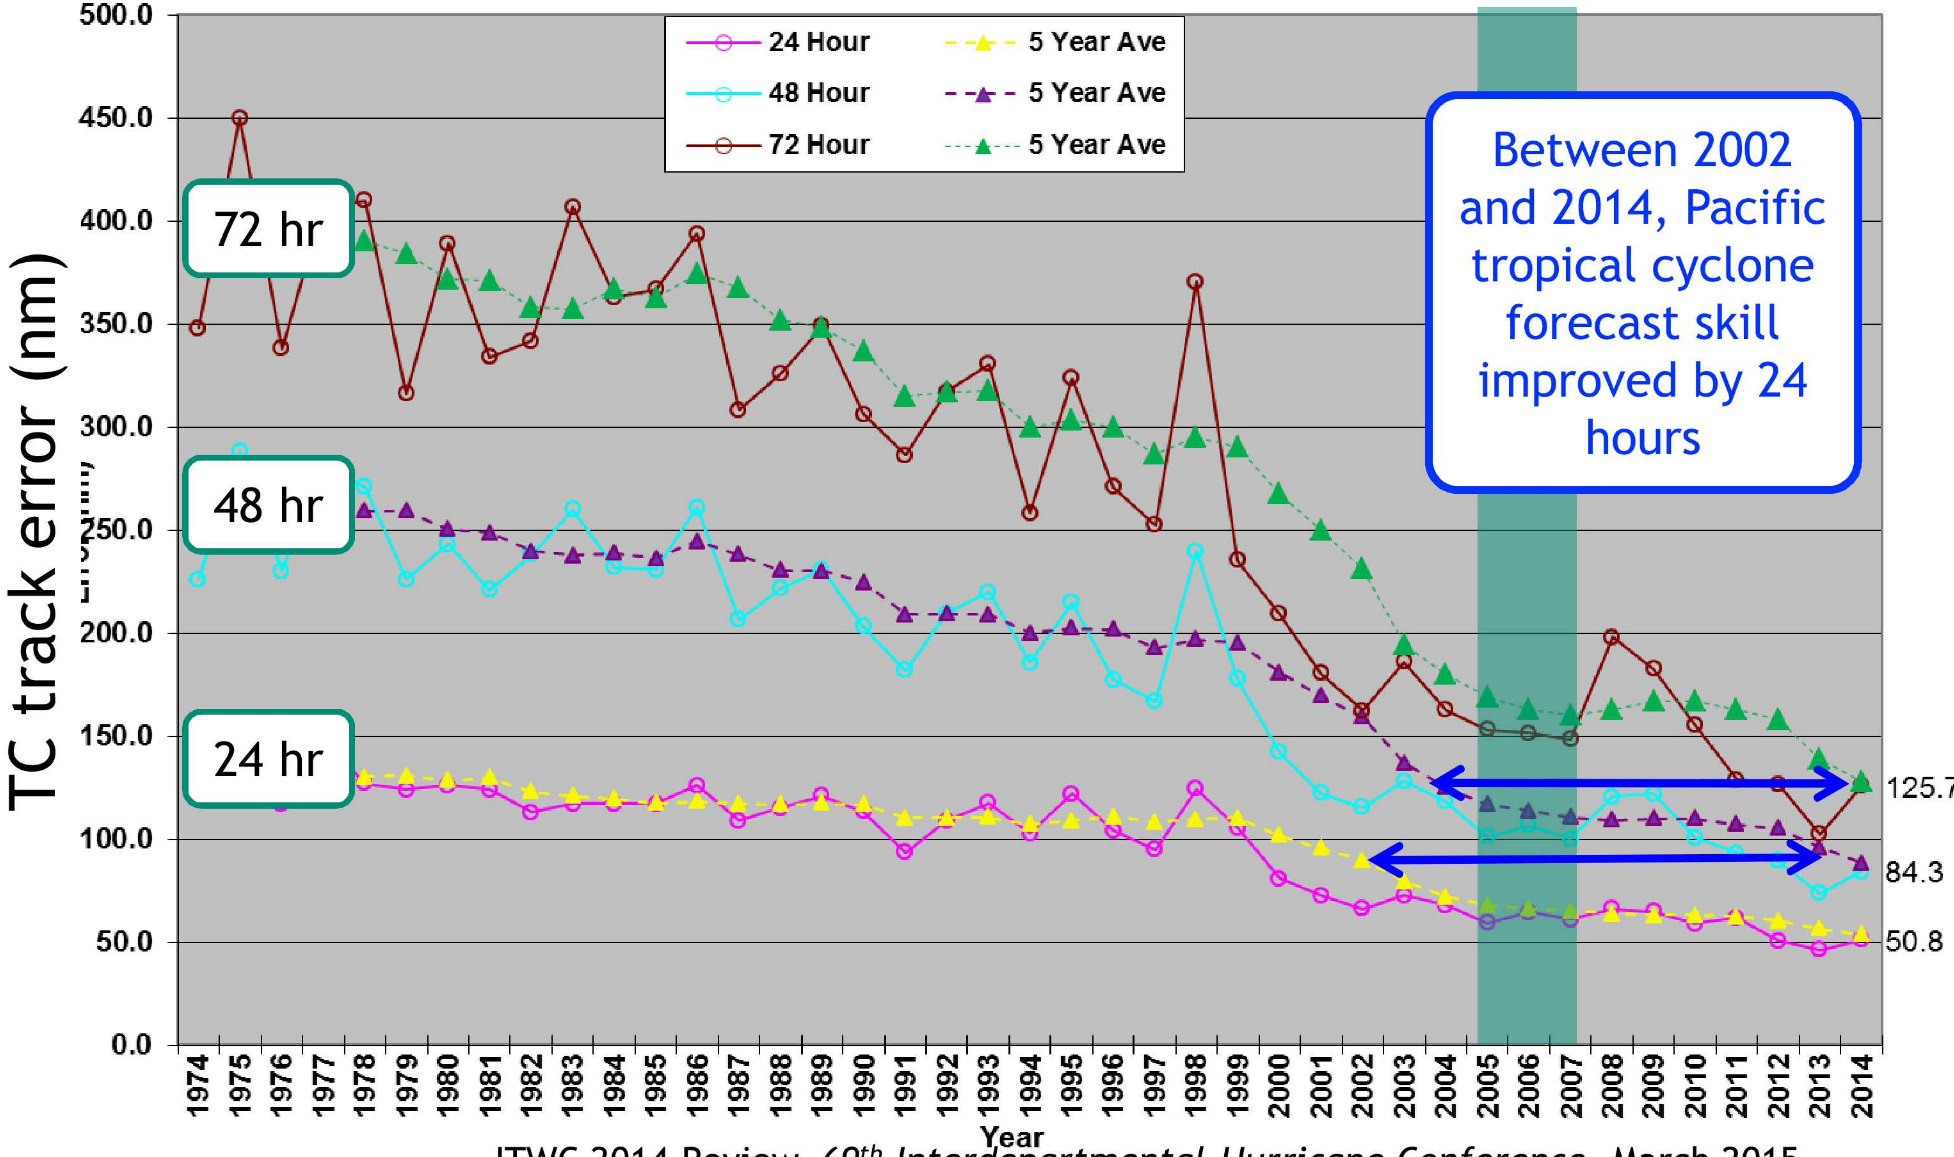

JTWC 2014 Review, *69th Interdepartmental Hurricane Conference*, March 2015

But challenges remain: A sea story

- 7-day forecast: typical tropical weather
  - Spotty showers, light NE trade winds
- 0630: NE wind, 10 kts, overcast
- 1130: W wind, 10 kts, clear
- Centuries of maritime lore: **"WATCH OUT!"**
- Satellites and state-of-the-art numerical models: **"Nothing to worry about."**

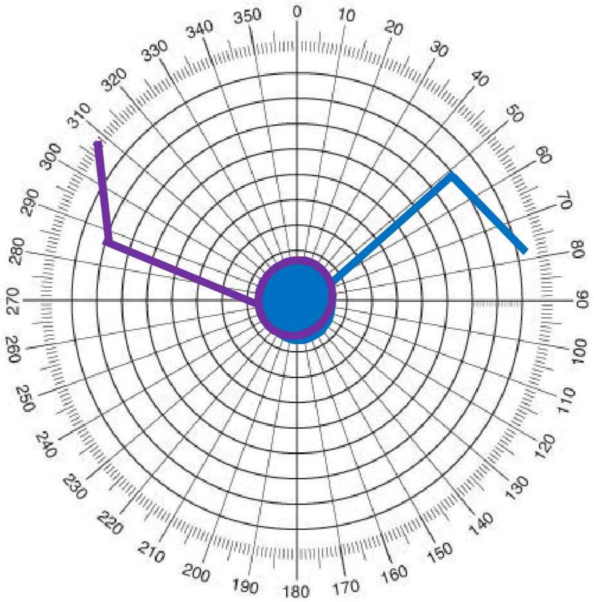

USS SHOUP (DDG 86)
Fall 2004, Eastern Pacific Ocean

Note: Neither **observations** (satellite) nor the **models** showed any indications of bad weather!

Challenge: Nonlinear multiscale dynamics

- Coastal mountains induce small-scale vortices in the boundary layer
- Some vortices find **favorable** local environments and grow
- Beginning stages are not resolved by models (sub-grid scales)
- Without clouds, they are invisible to satellites

Favorable conditions for tropical cyclones

1. High sea surface temperature (> 26 C)
2. Latitude > 5°
3. Low-level circulation
4. Upper-level divergence
5. Low vertical wind shear

Nonlinear multiscale dynamics

- 2D turbulence theory
  - KE cascade to large scale
  - Enstrophy cascade to small scale
- 3D turbulence theory
  - KE cascade to small scale
- Large scale atmospheric flow is well-approximated as a 2D fluid
  - ~ 80% KE goes up-scale, 20% goes down
  - ~ 20% enstrophy goes up, 80% enstrophy goes down
- Small scale flow: 3D fluid
- Transition region dx ~ O(100 km):
  - Nonlinear interactions between waves and vortices

Skamarock, et. al., *J. Atmos. Sci.*, 2014

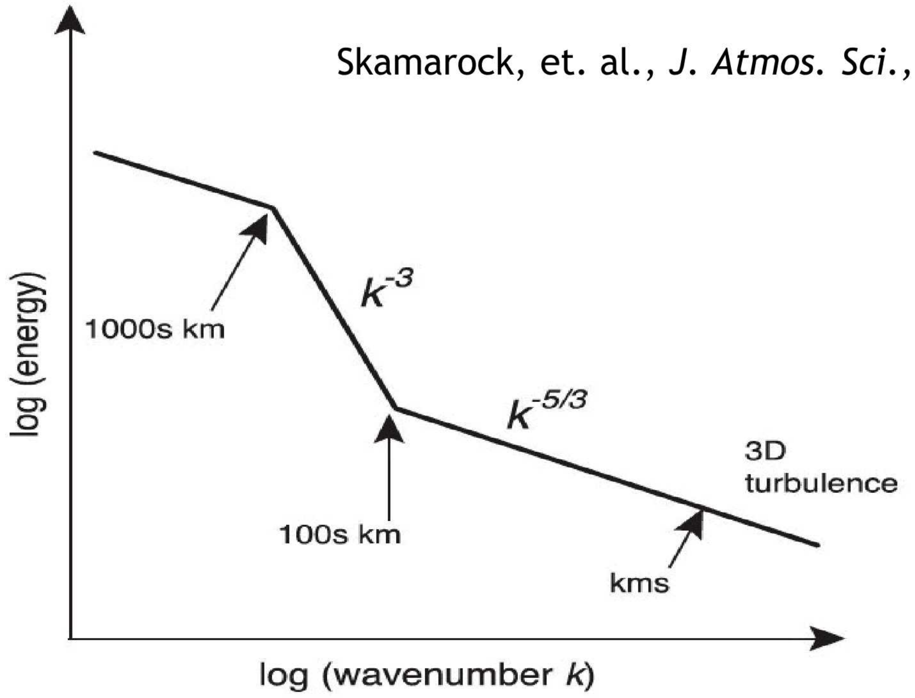

And: so far, this is just physics!

"The microphysical properties, even the macrophysical forms, of clouds are significantly affected by the chemicals in the air."

Lamb & Verlinde, *Physics & Chemistry of Clouds*, 2011.

Challenge: Chaotic dynamics

- Solutions lie on a ``fuzzy manifold,” or ``attractor,” of lower dimension than the whole phase space
- Sensitive dependence on initial conditions
  - Atmosphere Lyapunov time ~ 10-14 days
- **Weather forecasting goal:**
  - Deterministic solutions
- **Climate projection goal:**
  - Statistically describe attractors
  - Current method: Ensembles
  - **Challenge:** How to quantify ``statistically equivalent” model climates?

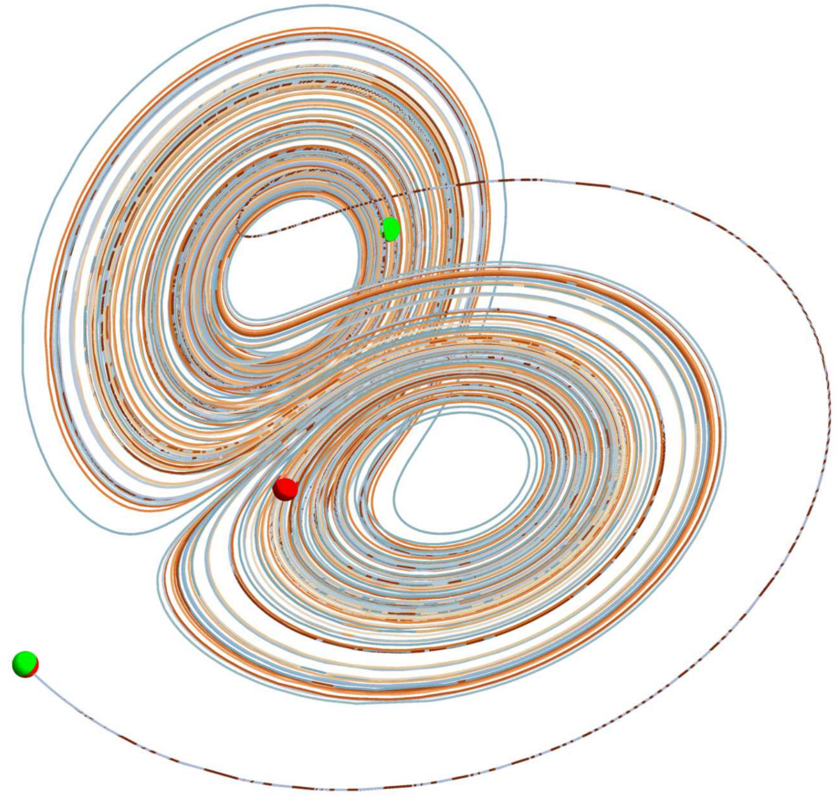

$$\dot{x} = \sigma(y - x)$$
$$\dot{y} = x(\rho - z) - y$$
$$\dot{z} = xy - \beta z,$$
$$\sigma = 3,\ \rho = 26.5,\ \beta = 1$$

E. Lorenz, *J. Atmos. Sci.*, 1963

2020: Numerical weather prediction

Global forecast:
- Spatial resolution: 1-10 km
- Time scale: 10-14 days

Atmosphere model

Resolved
Wind
Temperature
Cloud systems
Tracer advection

Unresolved
Turbulent mixing
Cloud microphysics
Aerosols/chemistry
Gravity wave drag

Boundary conditions

Resolved
Vegetative processes
Land fluxes and soil moisture
Land ice
Sea ice

Ocean model

Resolved
Currents
Temperature
Mesoscale eddies
Tracer advection

Unresolved
Turbulent mixing
Vertical wave breaking
Wind waves
Aerosols/chemistry/ fluxes

2020: Numerical weather prediction

Global forecast:
- Spatial resolution: 1-10 km
- Time scale: 10-14 days

Atmosphere model

Resolved
Wind
Temperature
Cloud systems
Tracer advection

Unresolved
Turbulent mixing
Cloud microphysics
Aerosols/chemistry
Gravity wave drag

Boundary conditions

Resolved
Vegetative processes
Land fluxes and soil moisture
Land ice
Sea ice

Ocean model

Resolved
Currents
Temperature
Mesoscale eddies
Tracer advection

Unresolved
Turbulent mixing
Vertical wave breaking
Wind waves
Aerosols/chemistry/ fluxes

2020: Climate projection

Global simulation:
- Spatial resolution: 25 - 100 km
- Time scale: 10-200 years

Current DOE research (next 5 years):

- 3 km Atmosphere model ("convection permitting")
- 25 km fully coupled model
- Robust 40-year projections: "Actionable information"
  - Water cycle
  - Cryosphere
  - Biogeochemistry

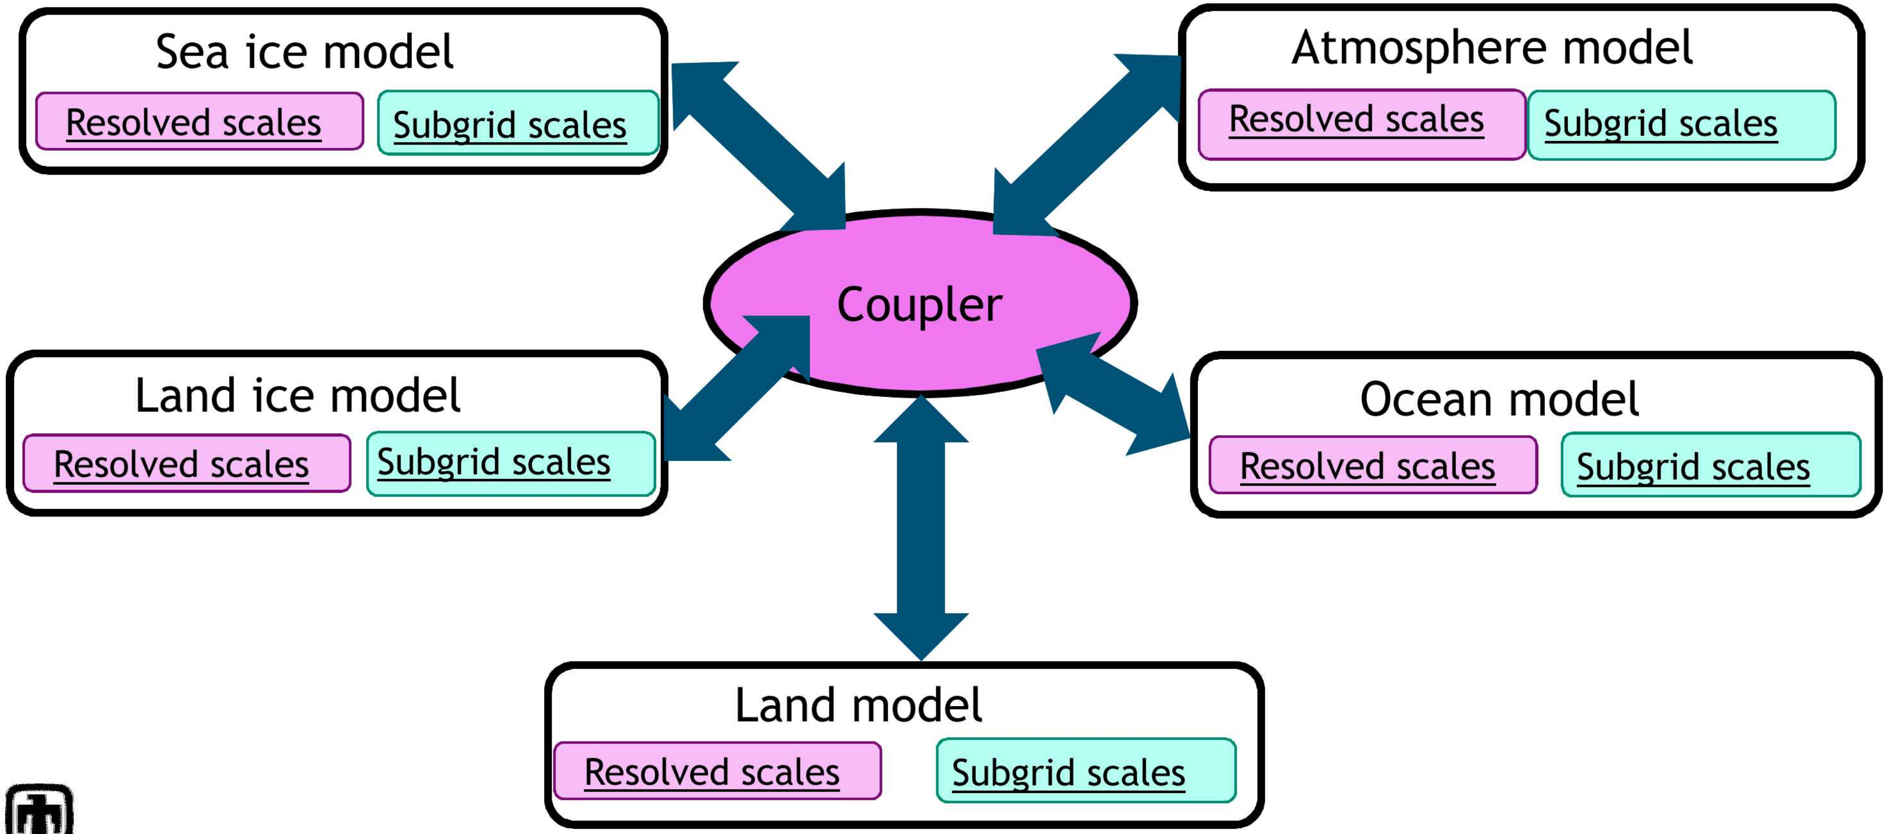

E3SM Energy Exascale Earth System Model

SciDAC Scientific Discovery through Advanced Computing

Computing challenges

- Required algorithm traits
  - Accuracy
  - Conservation
  - Tracer-continuity consistency
  - Shape preservation
  - Computational efficiency
- Throughput requirements
  - At odds with PDE structure
- Science goals need more resolution
  - At odds with accelerators and minimal data movement
- Heterogeneous computing architectures
  - Many-core vs. GPU
  - Programming models and software maintenance
  - Steep learning curves for non-CS folks

Background: Shallow atmosphere approximation

- Relative to the size of the planet, the atmosphere is a vanishingly thin layer on the surface of a spherical Earth
- Scales of horizontal motion >> scales of vertical motion
- Dynamics equations
  - Horizontal spectral elements
  - Vertically staggered 2nd-order FD
  - Taylor et. al., JAMES, 2020
- HEVI splitting: Horizontally Explicit, Vertically Impli
  - Columns are treated independently of each other
  - Workload measured by horizontal resolution

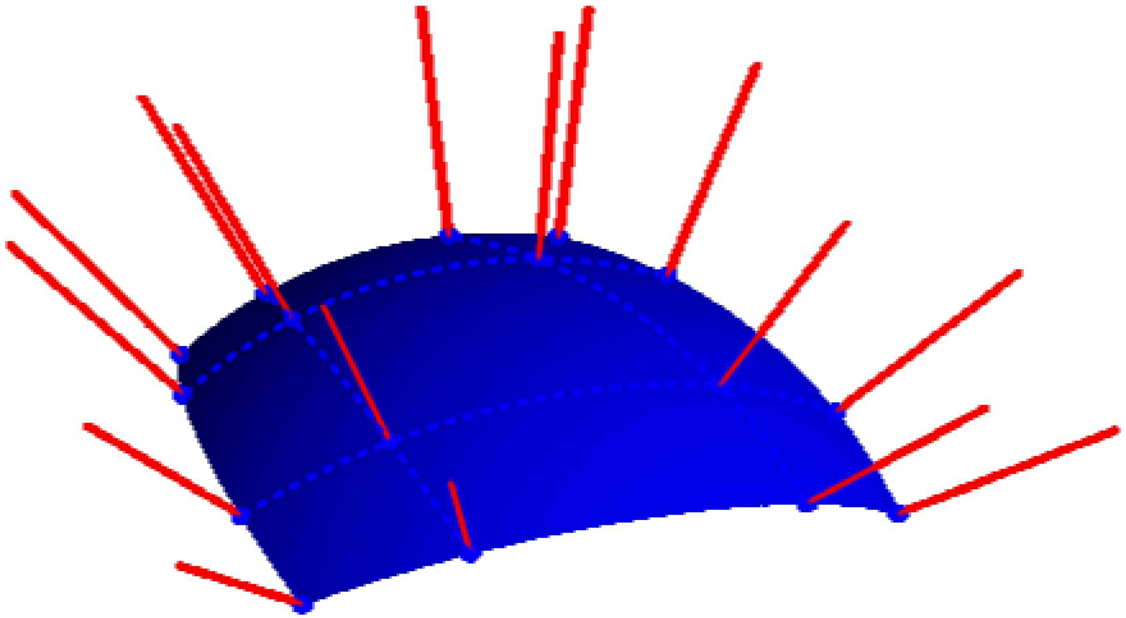

Challenge: Throughput

- Coupled simulation must run at appx. 2000 x real time (> 5 Simulated Years Per Day, SYPD)
- 2x spatial refinement -> 8x more work

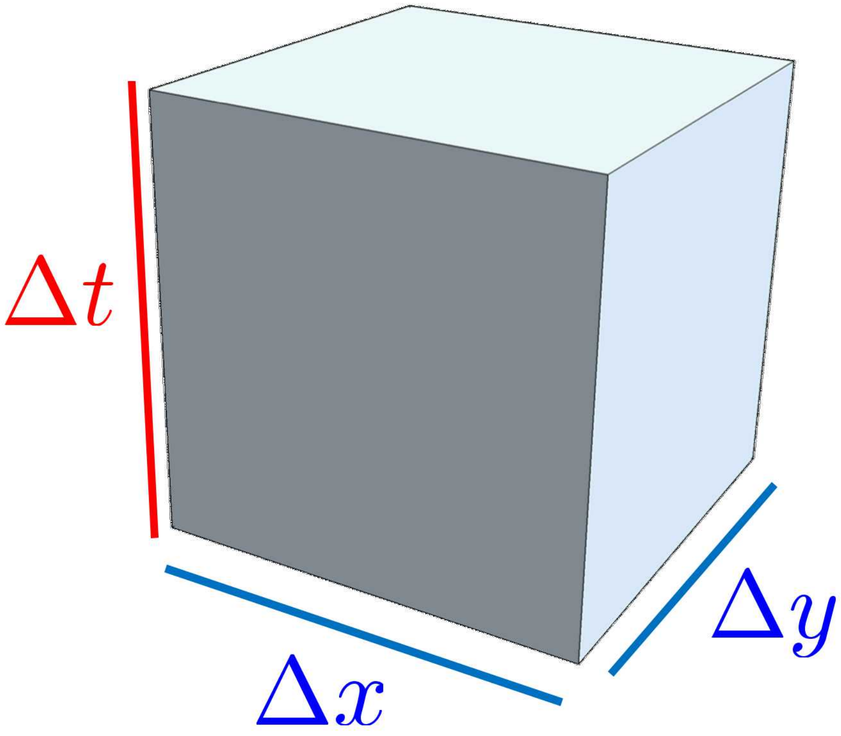

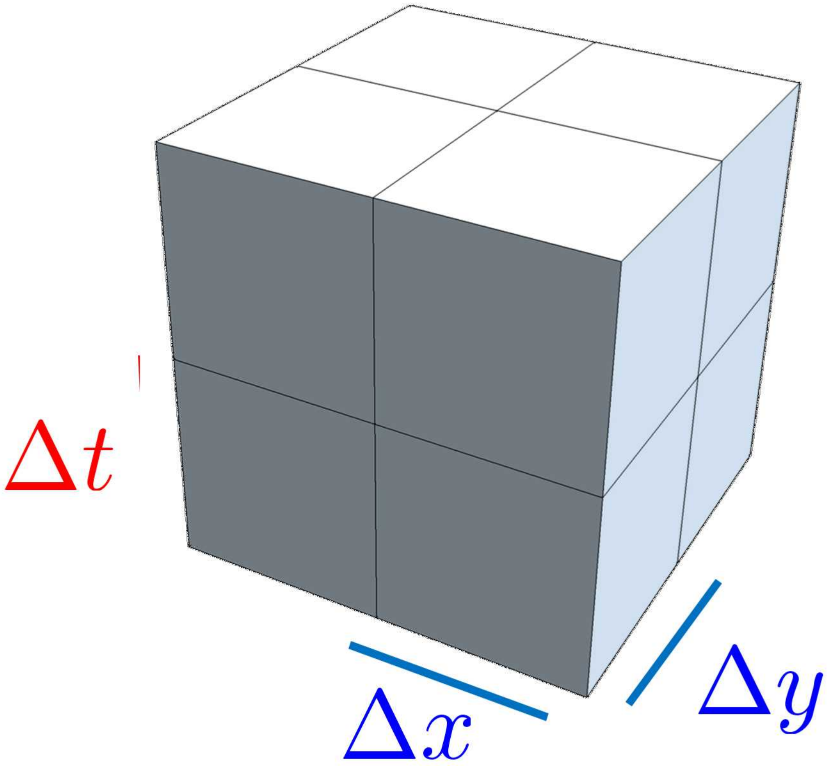

Courtesy of Matt Norman, ORNL

Challenge: Throughput

- Coupled simulation must run at appx. 2000 x real time (> 5 Simulated Years Per Day, SYPD)
- 2x spatial refinement -> 8x more work
- To keep same throughput, spread work over 8x more resources ("nodes")

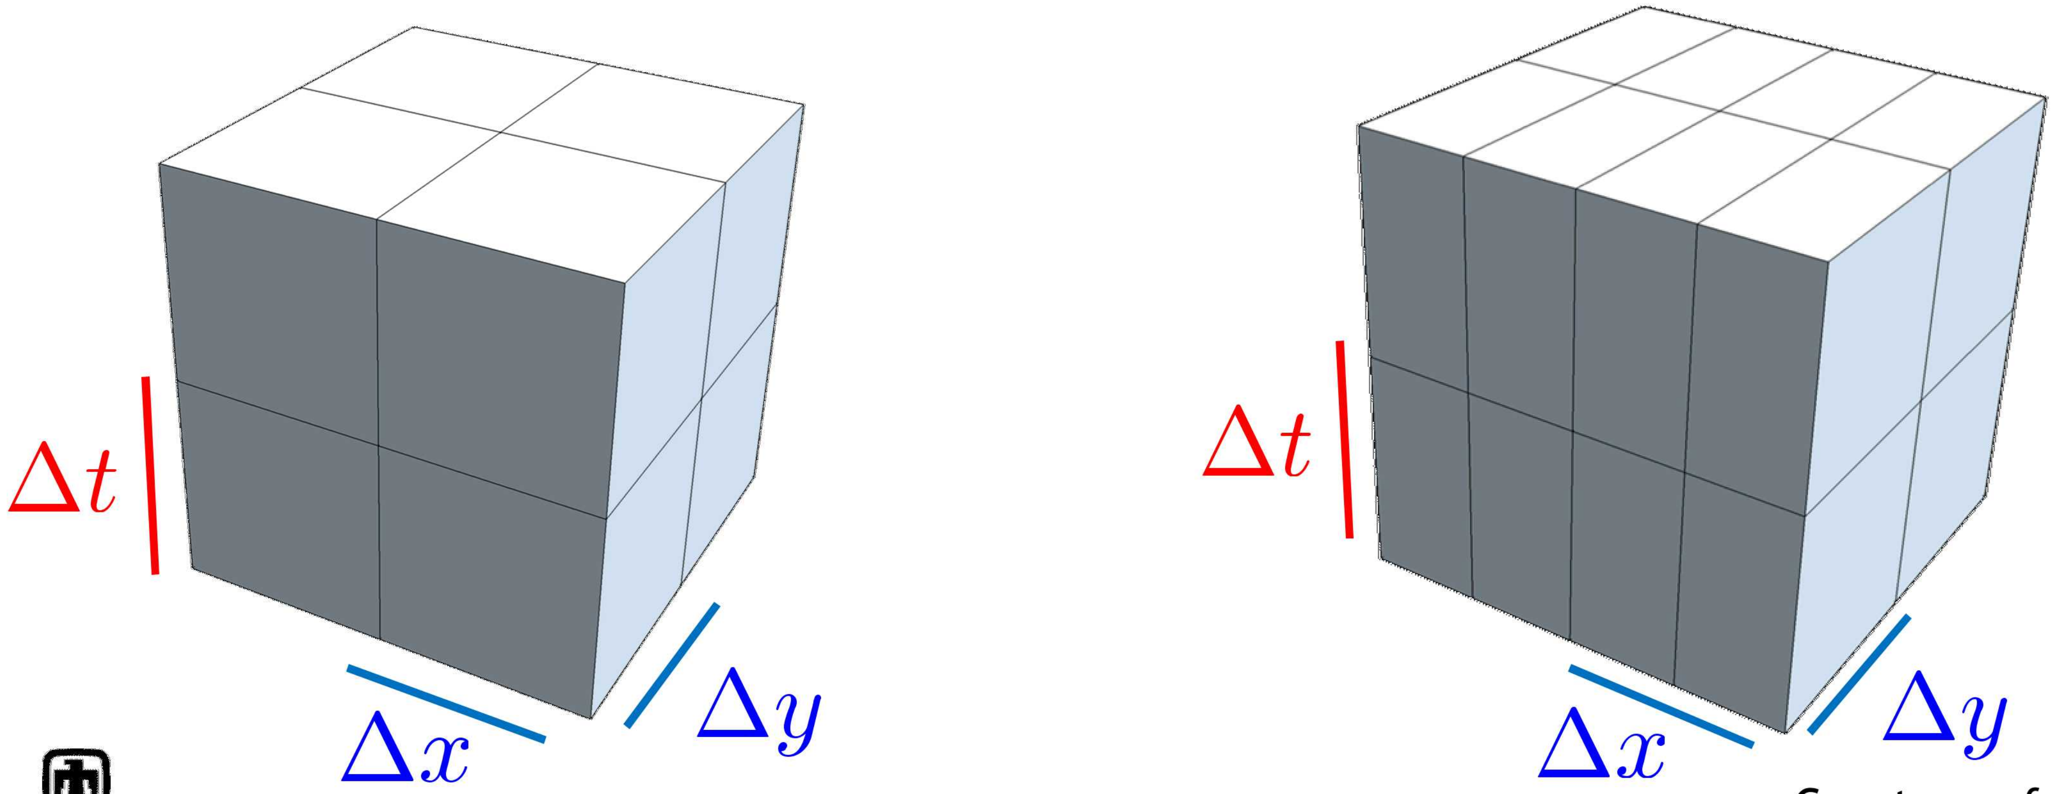

Courtesy of Matt Norman, ORNL

Challenge: Throughput

- Coupled simulation must run at appx. 2000 x real time (> 5 Simulated Years Per Day, SYPD)
- 2x spatial refinement -> 8x more work
- To keep same throughput, spread work over 8x more resources ("nodes")

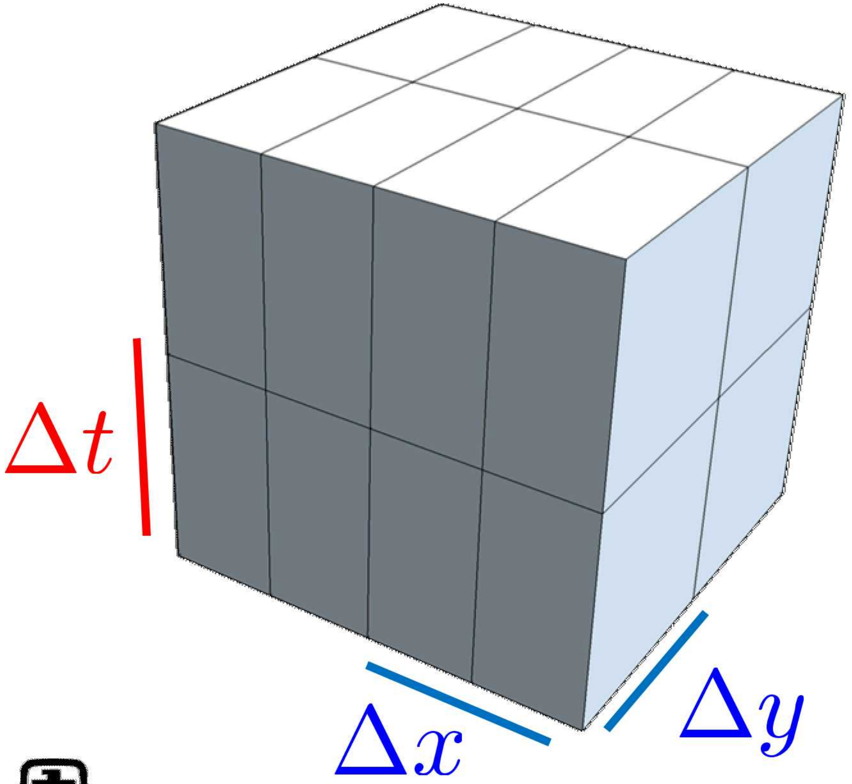

- **Work per node decreases by ½ with every 2x grid refinement**
  - Large MPI overheads
  - Decreasing workload for accelerators
- **Because of time step reduction, grid refinement is our friend**

Courtesy of Matt Norman, ORNL

Paths forward

- Algorithms for forward simulation with ensembles
  - Goals
    - Maximize “realism” per unit of data movement
    - Minimize Cost per parallelizable degree of freedom (parallelizable expenses are ok to add)
  - Strategies
    - High order algorithms with efficient limiters: **More resolution per**
    - **Large time steps**: Push the bounds of numerical stability
    - **Superparameterization***

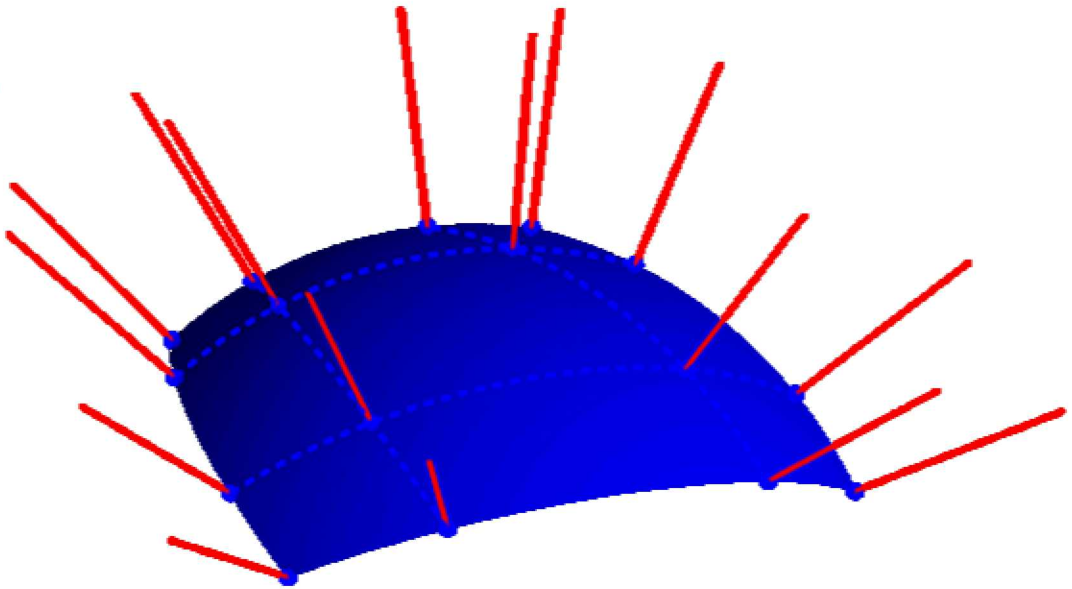

- Portable programming models
  - **Kokkos** metaprogramming model
    - Write code using the Kokkos c++ API
    - Compile for different architectures (e.g., CUDA, OpenMP, etc.)
- Other algorithms
  - Parallel-in-time
  - Reduced order models, machine learning

***Norman, Nair, *JAMES*, 2018**
***Hannah, et. al., *JAMES*, 2020**

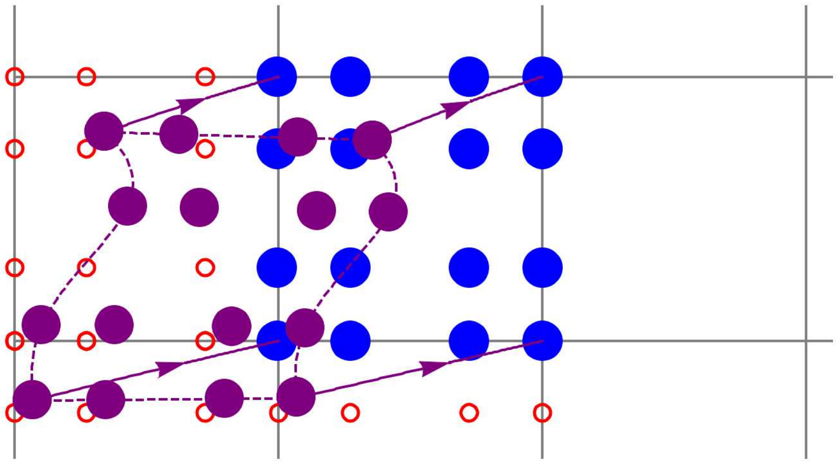

Semi-Lagrangian transport with spectral elements

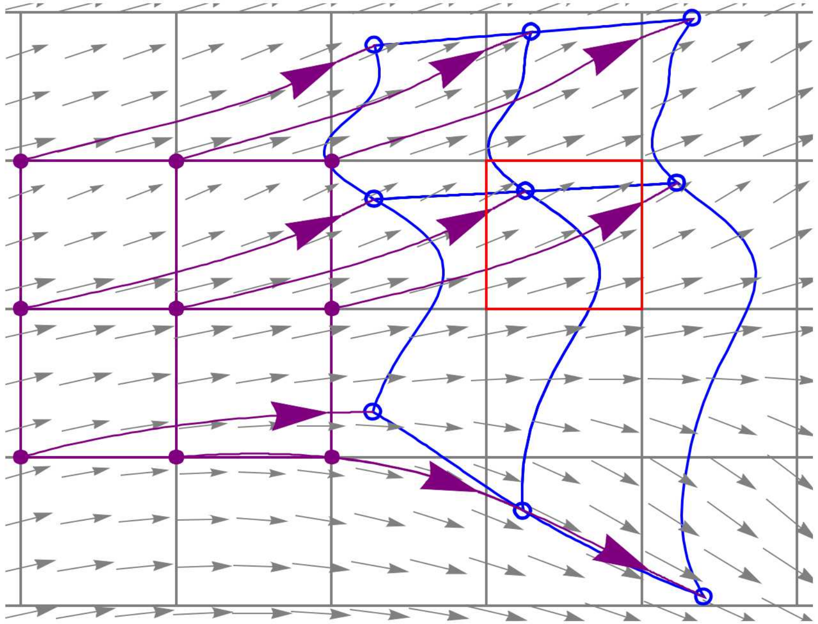

COMPact performance-POrtable SEmi-Lagrangian methods (COMPOSE)

- Spectral elements provide high-order accuracy with compact stencils
- Semi-Lagrangian (SL) time stepping permits Cr >> 1
- Highlights:
  - Second-order accurate in general flow, with shape preservation
  - Reduced MPI-communication rounds and volume vs. Eulerian transport scheme
  - Cell-integrated SL
    - Extensible into higher order (OOA 3) regimes
    - Locally mass conserving
    - **Speedup of ~ 2.6x over v1 Eulerian scheme**
  - Pointwise interpolation SL
    - Smallest possible communication requirements
    - Globally mass conserving
    - **Speedup of ~ 3.1x over v1 Eulerian scheme**

Cell-integrated SL:
Conservative mulitmoment transport along characteristics…
(Bosler, et. al., *SISC*, 2019)

New shape preservation algorithm:
Communication-efficient density reconstruction (CEDR)
(Bradley, et. al., *SISC*, 2019)

The transport problem

Given velocity, $\boldsymbol{u}(\boldsymbol{x},t)$, density, $\rho(\boldsymbol{x},t)$, and initial condition $q_0(\boldsymbol{x}) = q(\boldsymbol{x},0)$, solve for $q(\boldsymbol{x},t), \quad t > 0.$

$$\frac{\partial(\rho q)}{\partial t} + \nabla \cdot (\rho q \boldsymbol{u}) = 0$$

$$\frac{\partial q}{\partial t} + \boldsymbol{u} \cdot \nabla q = 0$$

Notation:

- Tracer mixing ratio: $q(\boldsymbol{x},t)$
- Tracer density: $Q(\boldsymbol{x},t) = \rho(\boldsymbol{x},t)q(\boldsymbol{x},t)$

Setting: Strong scaling limit

- 1 element per rank
- Density, velocity solved separately, “dynamics”

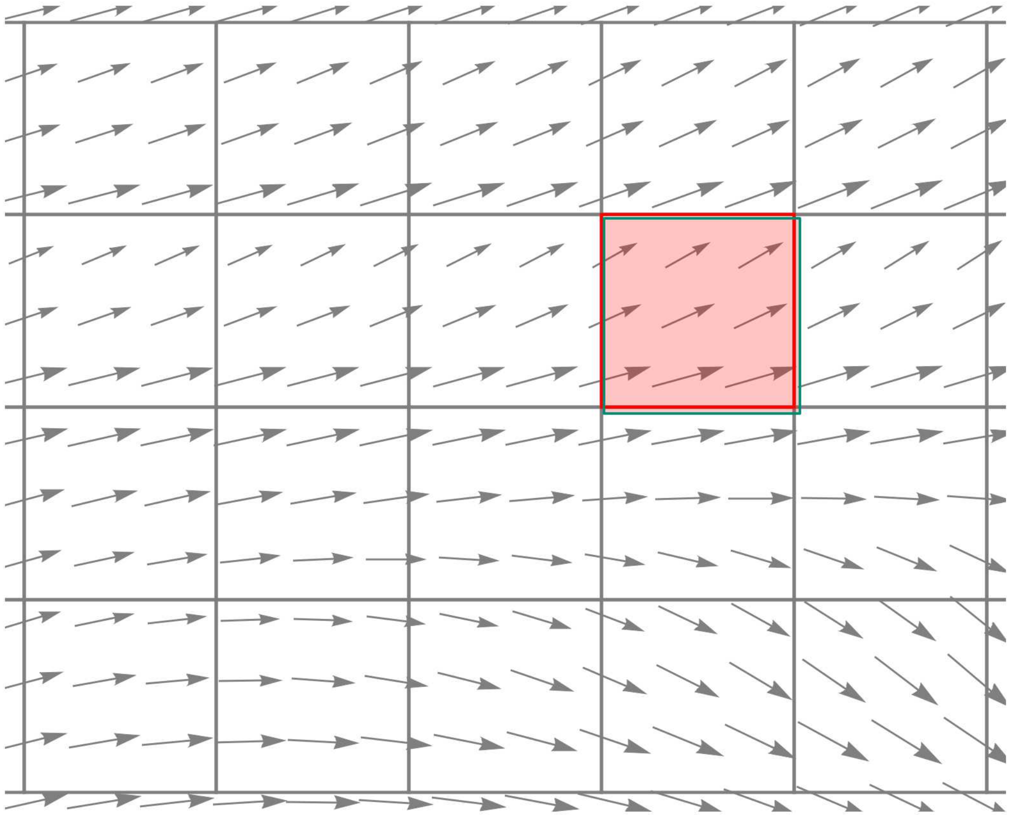

Algorithm requirements

- Conservation
- Accuracy & OOA >= 2
- Shape preservation
- Consistency: density equivalence between transport and dynamics
- Efficiency

Flux-form semi-Lagrangian methods

- Flux across each edge computed from its "swept region"
- Flux added to one side, subtracted from the other
- **Automatic conservation**

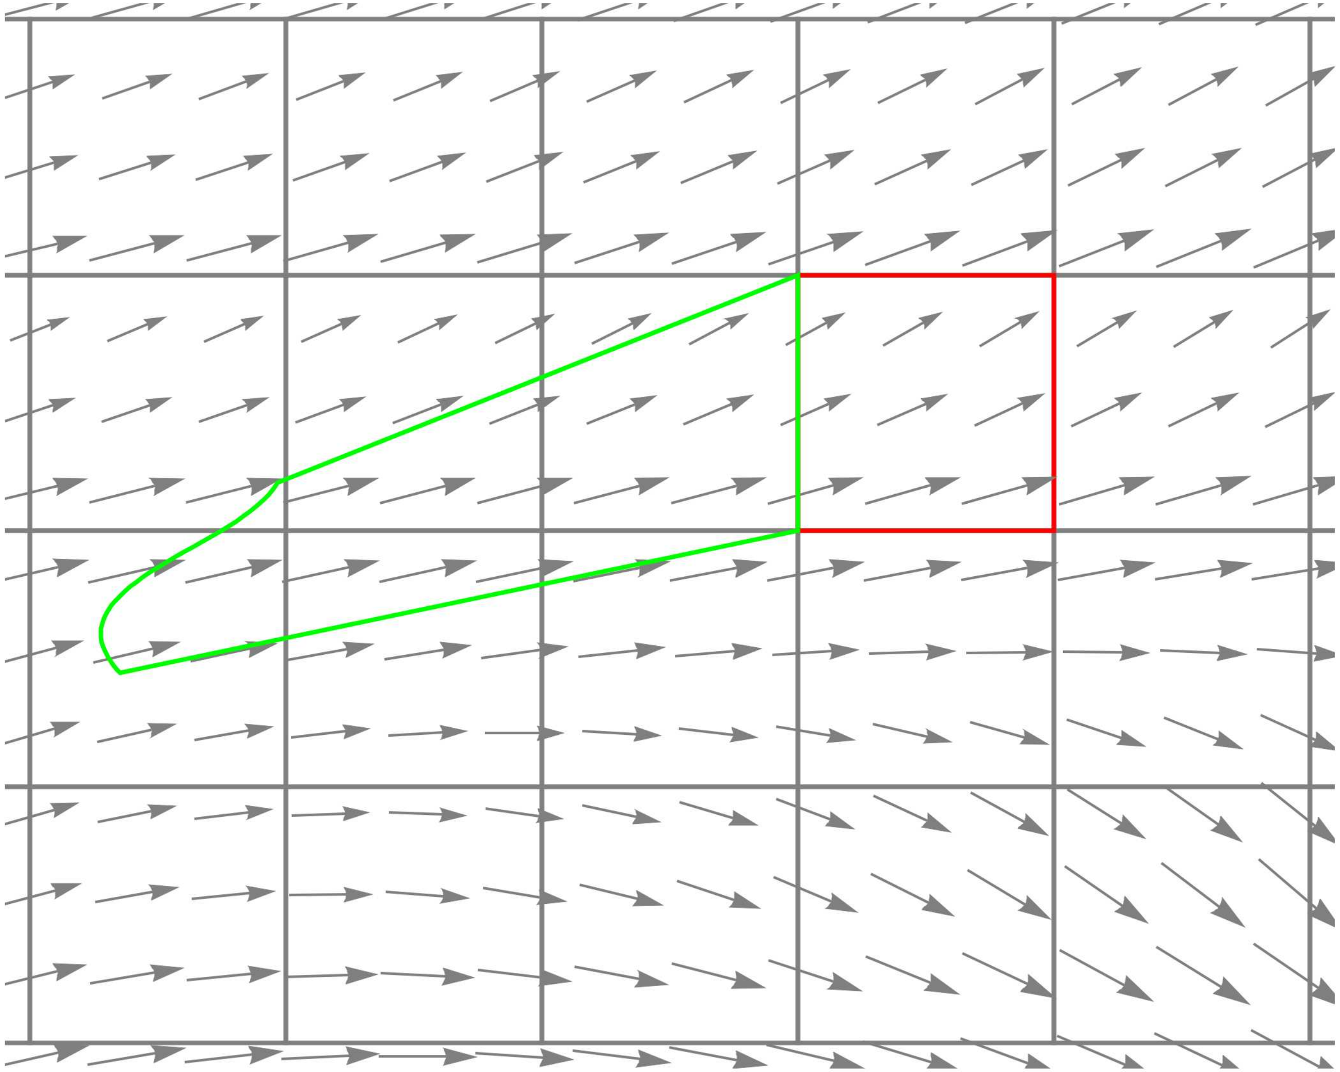

Flux-form semi-Lagrangian methods

- Flux across each edge computed from its "swept region"
- Flux added to one side, subtracted from the other
- **Automatic conservation**

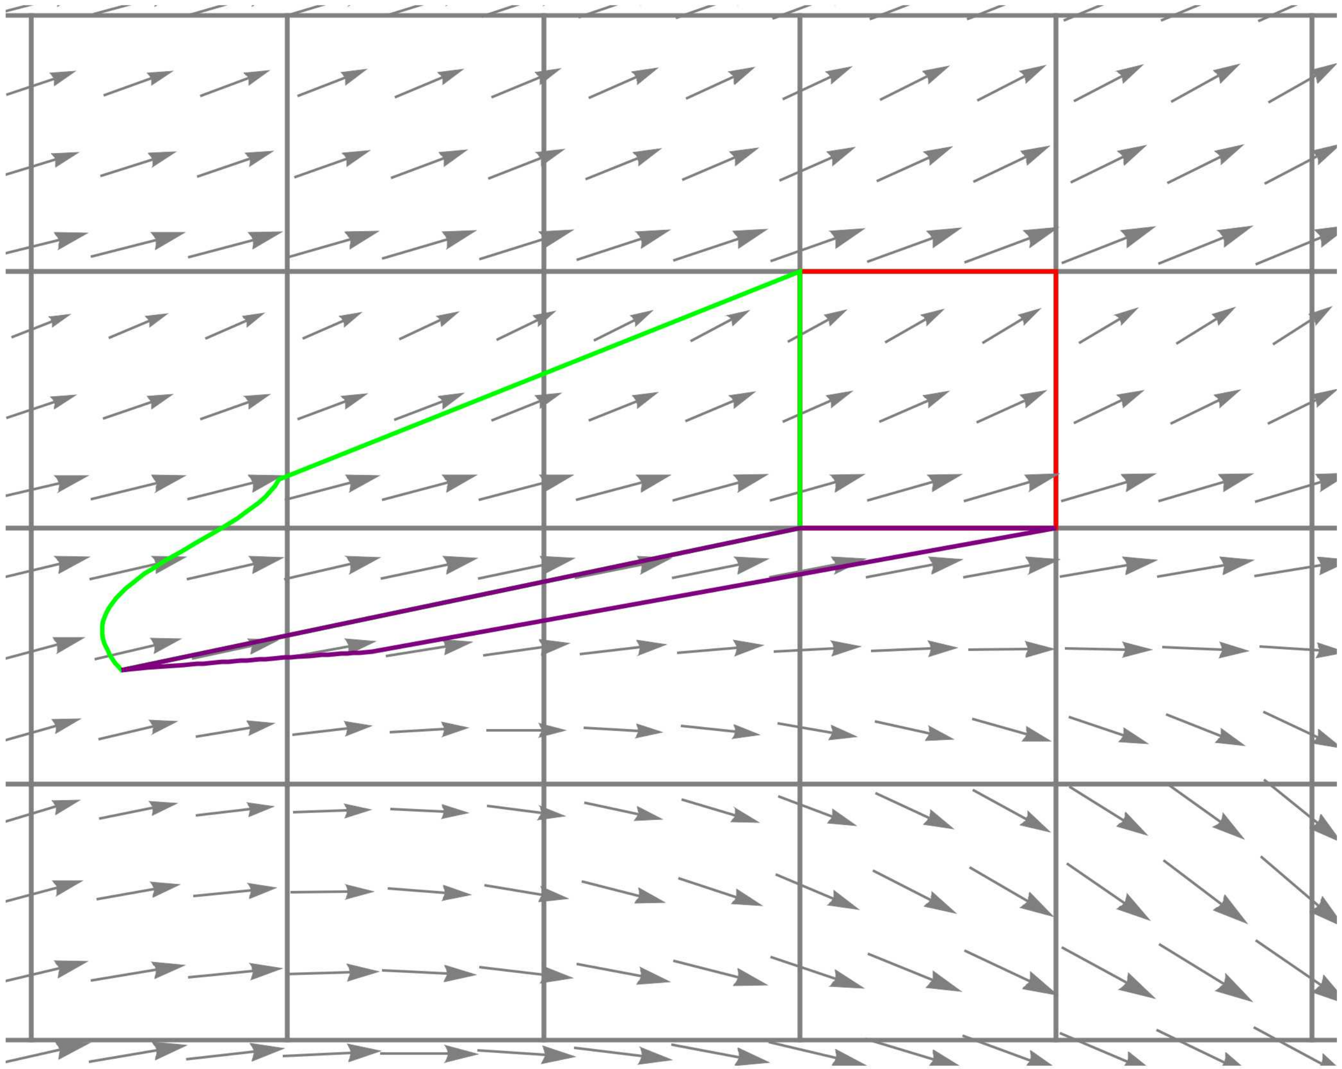

Flux-form semi-Lagrangian methods

- Flux across each edge computed from its "swept region"
- Flux added to one side, subtracted from the other
- **Automatic conservation**

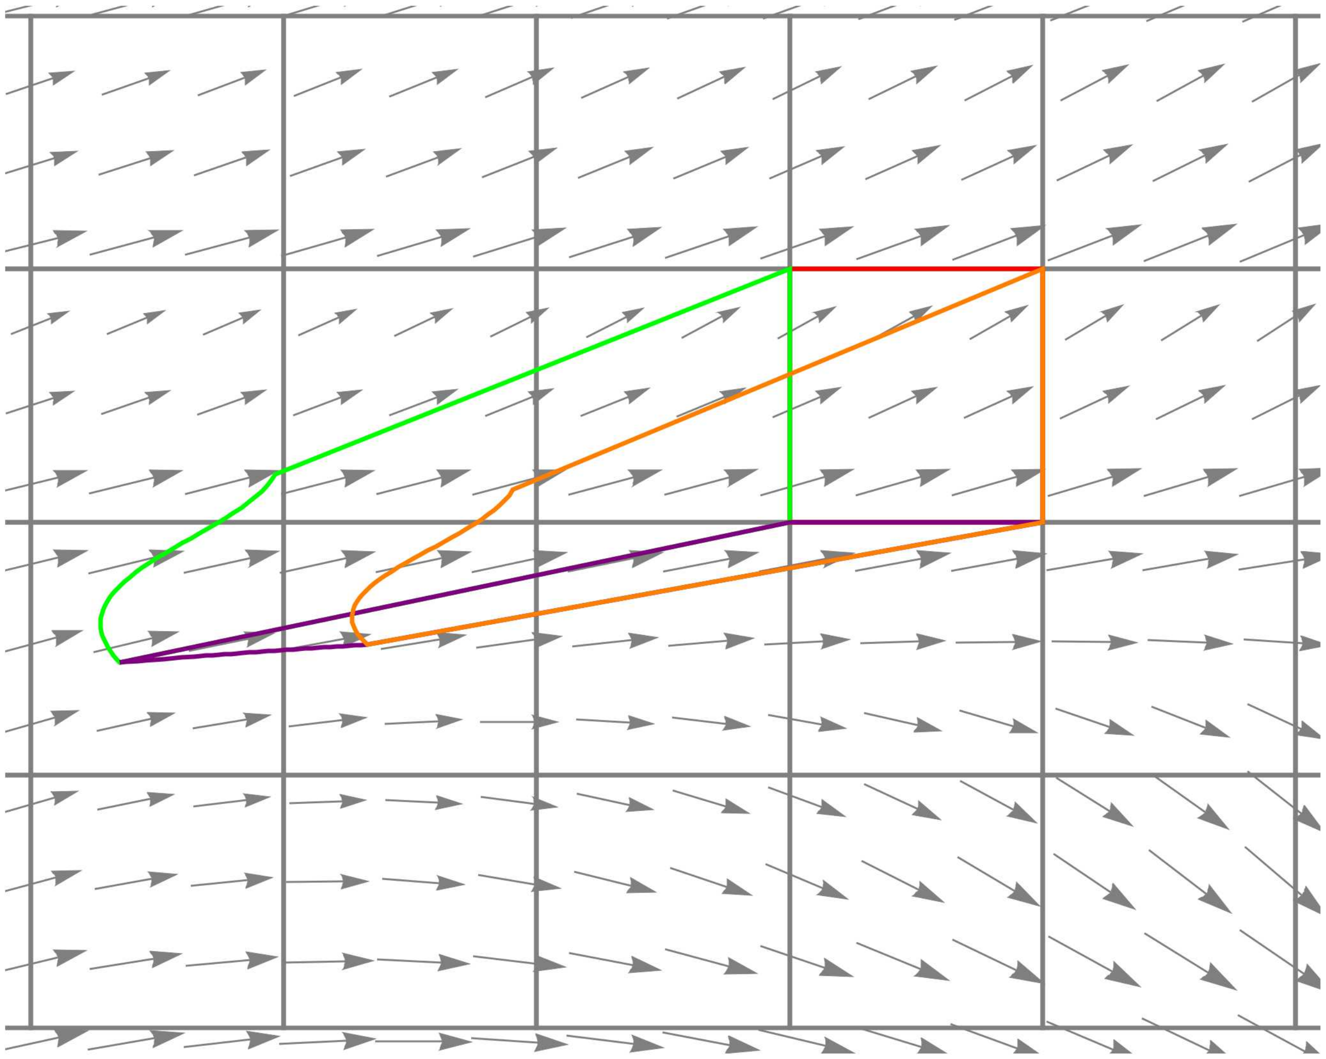

Flux-form semi-Lagrangian methods

- Flux across each edge computed from its "swept region"
- Flux added to one side, subtracted from the other
- Automatic conservation
- **Communication stencil grows with time step**

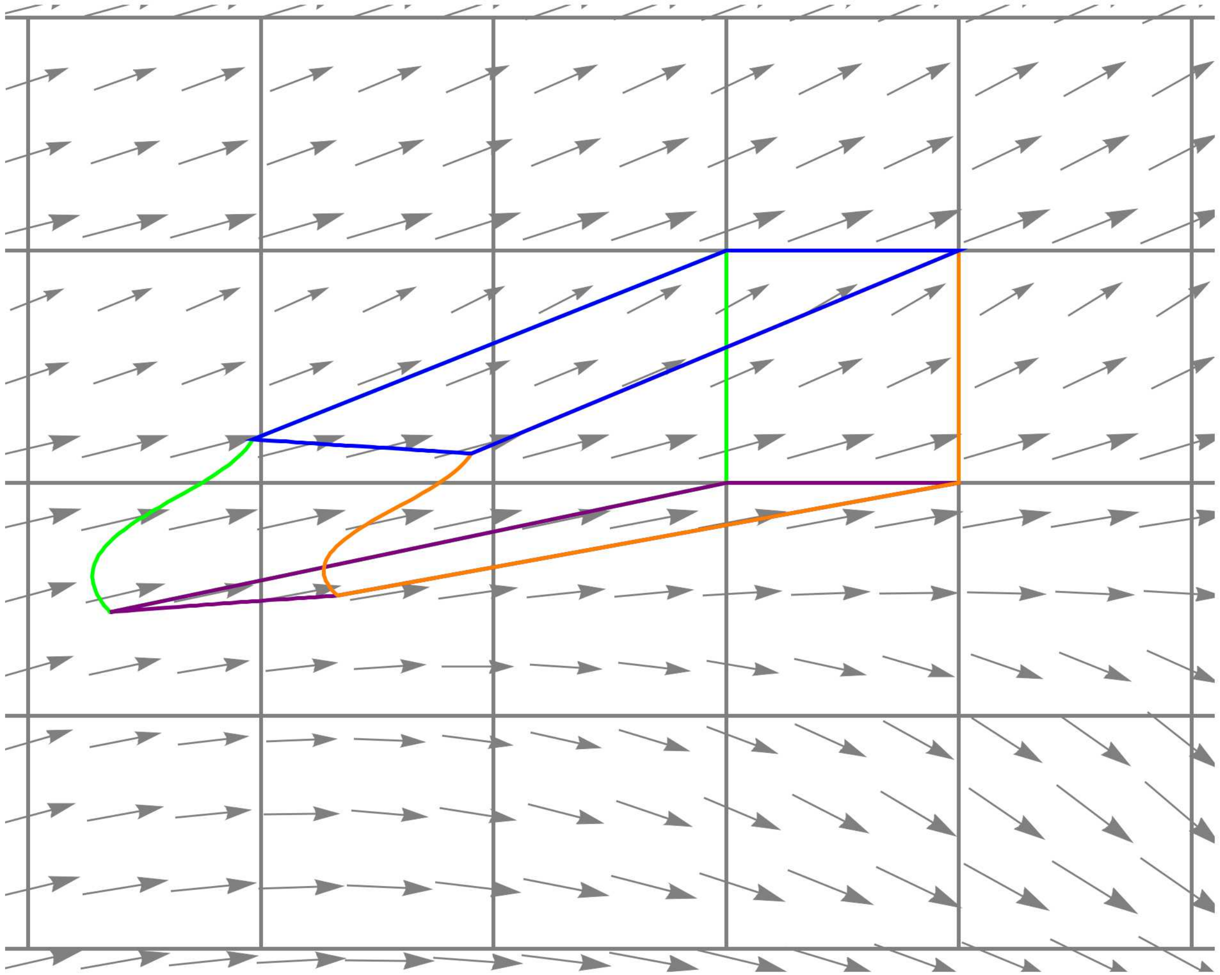

Remap-form semi-Lagrangian methods

- Elements are advected forward in time from t_n to t_{n+1} **(purple)**
- Distorted mesh at t_{n+1} provides ‘source’ data **(blue)**
- Eulerian mesh at t_{n+1} is ‘target’ **(red)**

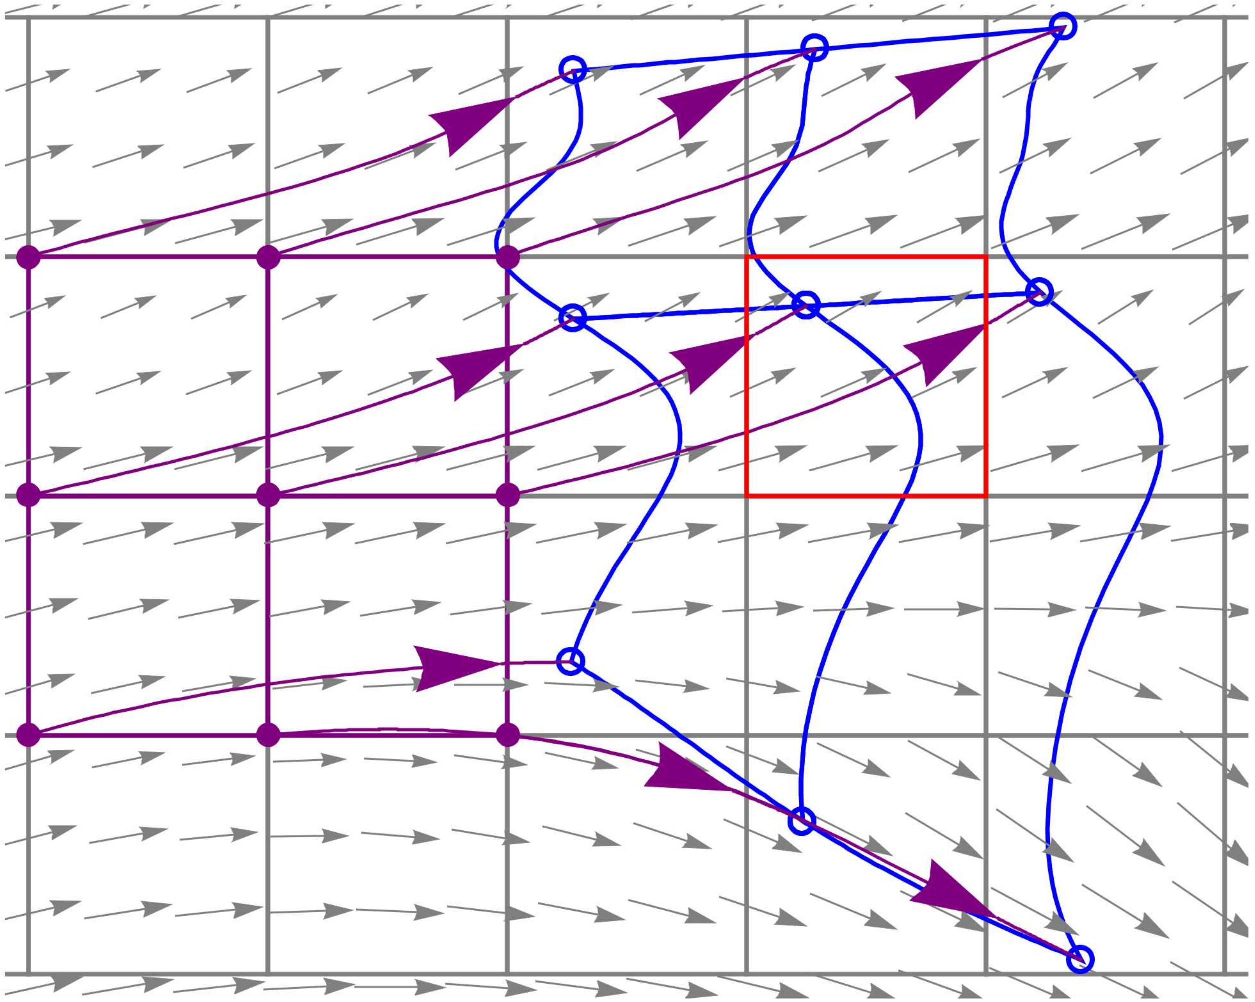

Remap-form semi-Lagrangian methods

- **Communication stencil (purple) is roughly independent of time step**
- Common refinement (Overlap mesh)
  - For each element E_k, the list L_k contains all intersecting distorted elements $E_l(t_{n+1})$

$$L = \{l \in [1, N_e] : E_l(t_{n+1}) \cap E_k \neq \emptyset\}$$

  - To each l in L_k there is an associated overlap region V_{kl} **(colors)**

$$V_{kl} = \{\boldsymbol{x} : \boldsymbol{x} \in E_k \text{ and } \boldsymbol{x} \in E_l(t_{n+1})\}$$

  - **Key development:** Global overlap mesh is not required
  - **Common refinement can be computed locally**

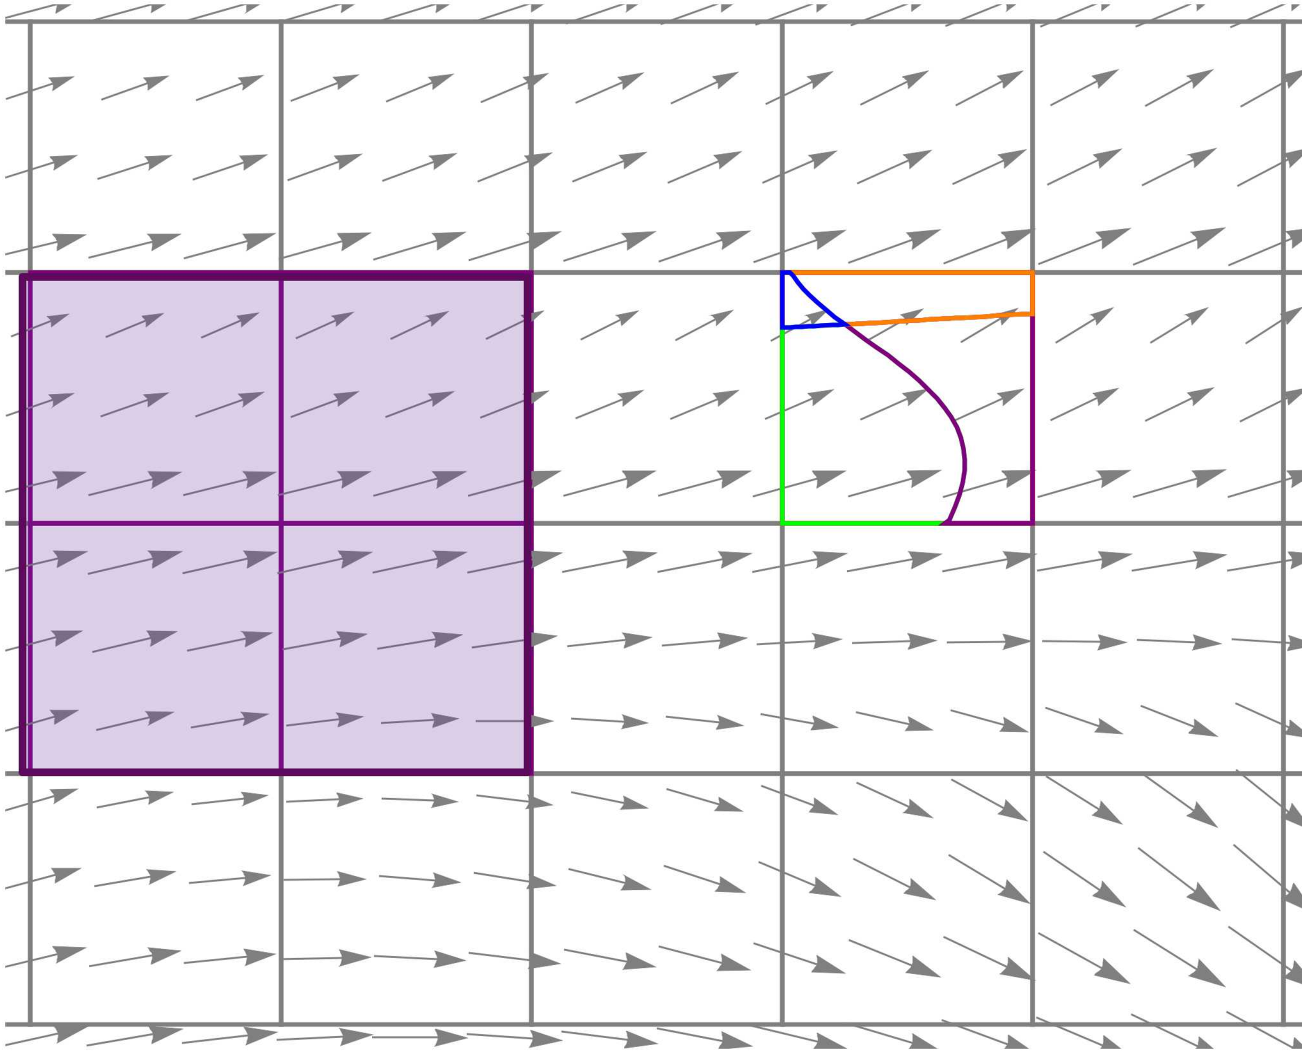

Bosler, Bradley, Taylor, SISC, 2019.

(Some) Related work

- Basic SL
  - Staniforth, Cote; *Mon. Weather Rev.,* 1991
  - McGregor; *Mon. Weather Rev*., 1993
  - **Many others…**
- SL Transport as incremental remap
  - Dukowicz, Baumgartner; *J. Comput. Phys.*, 2000
- Remap: Map data from one spatial discretization to another
  - SEM mesh to SEM mesh = Conservative L2 projection
  - Farrell et. al.; *CMAME*, 2009
- Cell-integrated SL: 'Lagrange-Galerkin' or 'Characteristic Galerkin' methods

Varoglu, Finn; *JCP*, 1980
Douglas, Russell; *SIAM J. Num. Anal.*, 1982
Morton, Priestley, Suli, *RAIRO*, 1988
Arbogast, Huang, *SIAM J. Sci. Comput.*, 2006
Priestly; *JCP*, 1993
Giraldo; *JCP*, 1997
Lee et. al.; *JCP*, 2016

E3SM Peformance study

- Strong scaling limit, 1 element per core
- Normalized transport time (lower is better) vs. number of tracers
- Breakeven of SL over Eulerian < 7 tracers
  - SL has some geometric computational overhead that Eulerian does not
- Cell-integrated SL is 2.6x faster
  - **Locally** mass conserving
- Pointwise interpolation SL is 3.1x faster
  - **Globally** mass conserving
  - Smaller communication volume (**basis-point** vs. **basis-basis**)
- MPI communications are still the limiting factor

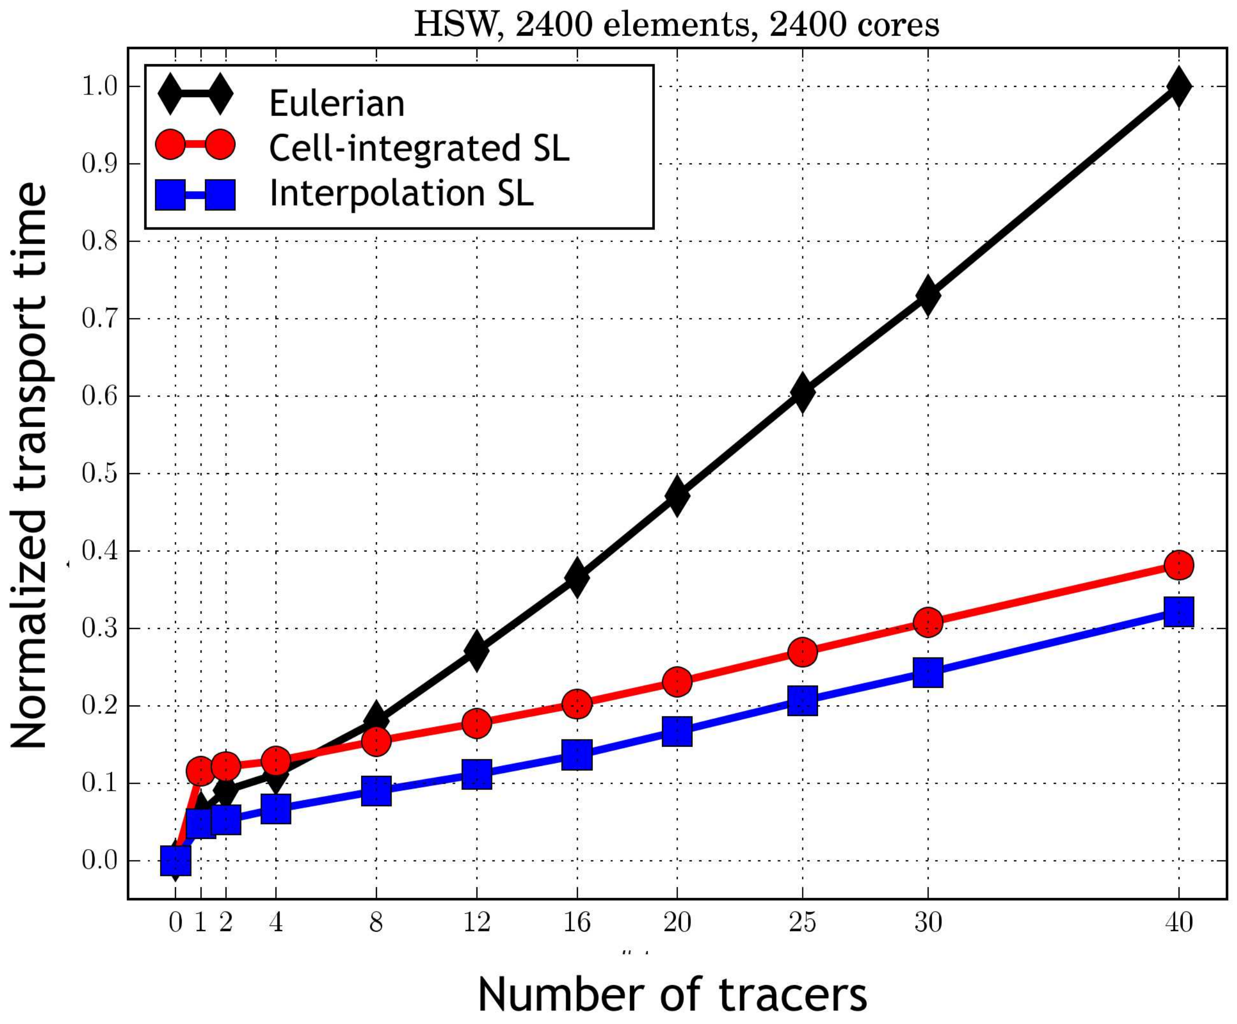

Property preservation: Definitions

- **Property:** A quantity that is required to be represented exactly (to machine precision) in an otherwise approximate numerical solution
  - Different discretizations and methods between coupled processes increase the difficulty of preserving properties
- **Static**: property does not depend on current solution; otherwise, **dynamic**
- **Global:** property is only relevant to the entire domain, Ω
- **Local:** property is defined by information in its domain of dependence, $\Delta\Omega(t)$

Example properties

Conservation (global, static) • Ensures physical conservation law	$$\int_{\Omega} f(\boldsymbol{x},0)\,d\boldsymbol{x} = \int_{\Omega} f(\boldsymbol{x},t)\,d\boldsymbol{x} \quad \forall t > 0$$
Conservation (local, dynamic) • Ensures physically realistic dynamics	$$\int_{\Delta\Omega(t_n)} f(\boldsymbol{x},t_n)\,d\boldsymbol{x} = \int_{\Delta\Omega(t_{n+1})} f(\boldsymbol{x},t_{n+1})\,d\boldsymbol{x}$$
Positivity (global, static) • Ensures physically realizable density	$$\rho > 0$$
Range (global, static) • Ensures physically realizable mixing ratios	$$0 \le q_i \le 1$$
Range (global, dynamic) • Ensures tracer consistency • **Safety problem:** Always feasible	$$\min_{\boldsymbol{x}\in\Omega} q_i(t) \le q_i(t_{n+1}) \le \max_{\boldsymbol{x}\in\Omega} q_i(t)$$
Range (local, dynamic) • Ensures physically realistic transport	$$\min_{\boldsymbol{x}\in\Delta\Omega(t)} q_i(t) \le q_i(t_{n+1}) \le \max_{\boldsymbol{x}\in\Delta\Omega(t)} q_i(t)$$

The property preservation problem

- Properties: Conservation, local dynamic range preservation
  - Define min and max tracer densities from domain of dependence

$$Q_{\min} \equiv \rho_{dyn} q_{\min}, \quad Q_{\max} \equiv \rho_{dyn} q_{\max}$$

  - Require:

$$Q_{\min} \leq Q \leq Q_{\max}$$

- Constrained optimization problem
  - Define set $\mathcal{Q}$ as set of all solutions that satisfy range preservation and conservation
  - Given a numerical solution $\boldsymbol{Q}^*$, we seek

$$\boldsymbol{Q} = \arg \min_{\boldsymbol{Q} \in \mathcal{Q}} \|\boldsymbol{Q} - \boldsymbol{Q}^*\|$$

- Related work: SLICE, CSLAM, HEL
  - Clip-and-assured sum
  - 2-norm minimization: Bochev et. al., *JCP* 2013, 2014
  - Other methods: Priestly, *MWR*, 1993; Bermejo & Conde, *MWR*, 2002

Feasibility

- Def: Cell mass boundedness: $$\sum_{i \in E_k} Q_i^{\text{min}} w_i \leq \sum_{i \in E_k} Q_i w_i \leq \sum_{i \in E_k} Q_i^{\text{max}} w_i$$
- **Condition:** Cell mass boundedness is necessary and sufficient to ensure feasibility of the shape preservation problem
- **Claim:** Compact, high order, semi-Lagrangian methods **cannot** assure cell mass boundedness
- **Implication:** Compact, high order, semi-Lagrangian methods require mass movement to ensure shape preservation
  - Global conservation requires necessarily non-local computation
  - Non-local methods are inefficient in parallel (all-to-all reductions, or "all-reduce")
  - **Goal:** Minimize the cost and number of all-reduces required to solve the shape preservation problem
- **Question:** What is the *smallest* number of all-reduces required to guarantee mass conservation, tracer consistency, and shape preservation?

Answer: 1
(Independent of the problem data)

Clip And Assured Sum (CAAS)

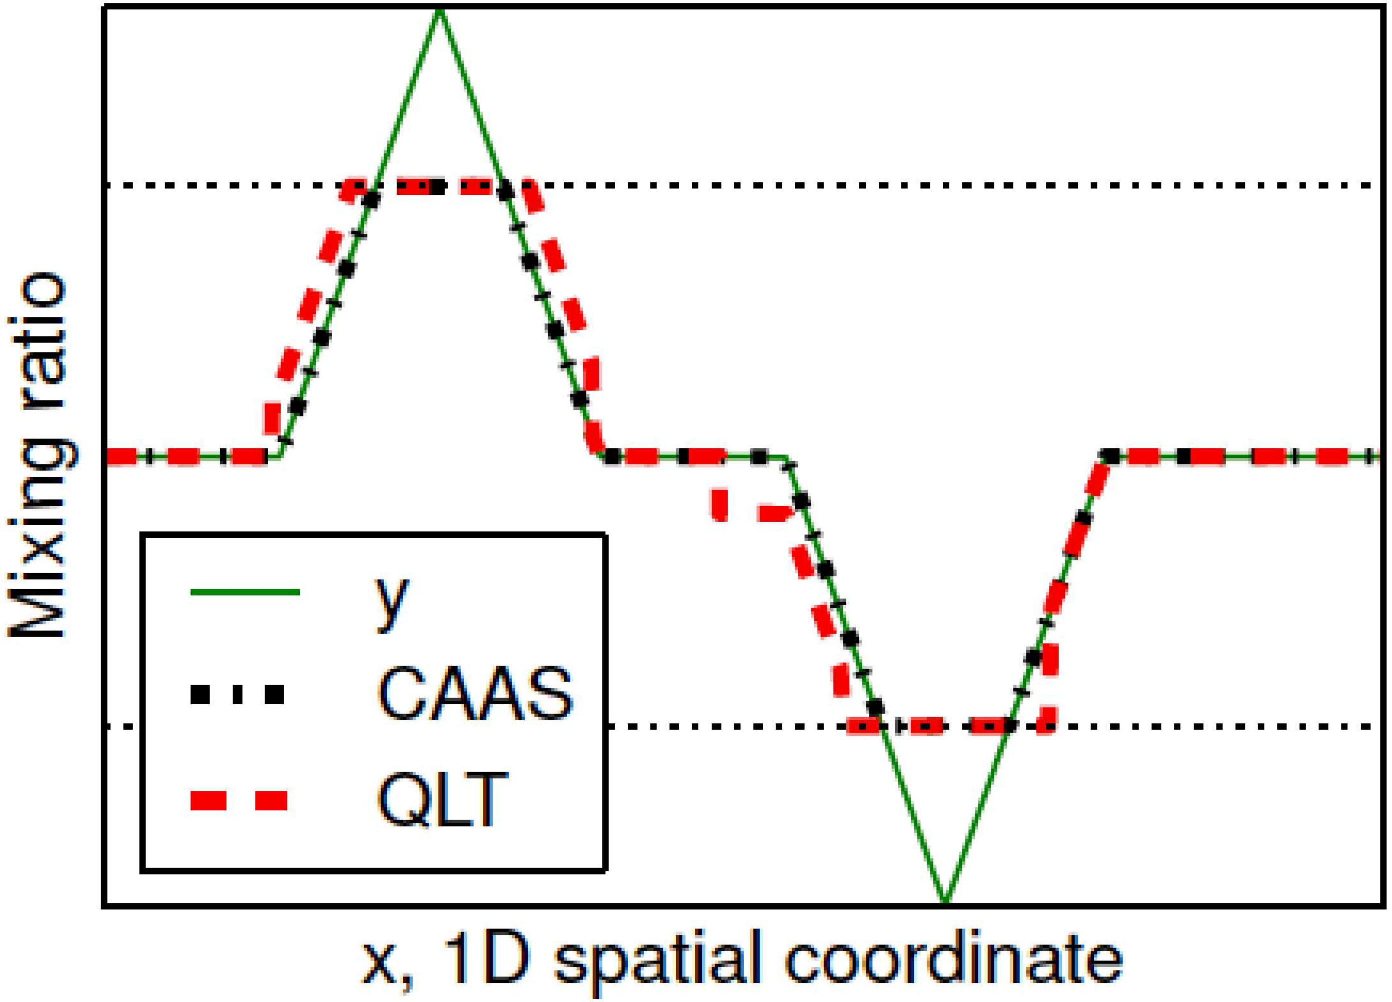

y: high-order solution **(green)**

CAAS : (black, dashed)

- Simple, 1 all-reduce
- Ensures conservation and shape preservation
- Mass movement is non-local, **"teleporting"**

QLT: Quasi-Local Tree **(red)**

- 1 all-reduce
- Ensures conservation and shape preservation
- Mass movement is **quasi-local**

QLT: Quasi-Local, Tree-based density reconstruction

- **Precompute step:** Build a tree over the mesh such that its leaf nodes are 1-1 with mesh cells
  - The tree breaks the global coupling of the shape preservation problem, at the cost of strictly local mass movement
  - Mass movement is now ``local" within the tree (hence the name, quasi-local)
- **Runtime: 2 step algorithm**
  - Leaves-to-root reduction
  - Root-to-leaves broadcast

QLT: Tiny mesh example

Tiny mesh
- Quadratic elements
- Boundary conditions ignored

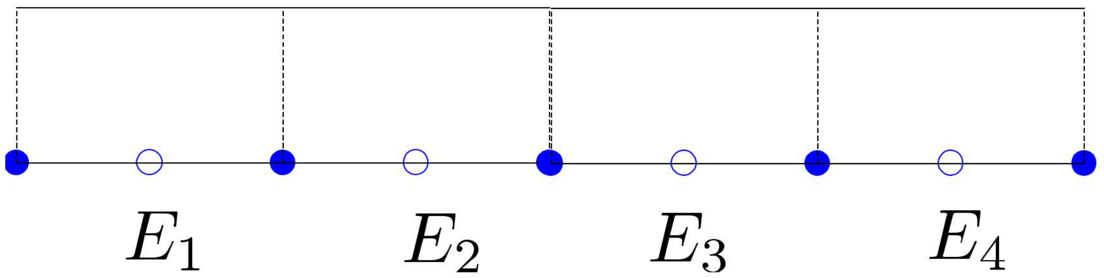

QLT: Tiny mesh example

Precompute
- Build a tree
- Leaves 1-1 with mesh elements

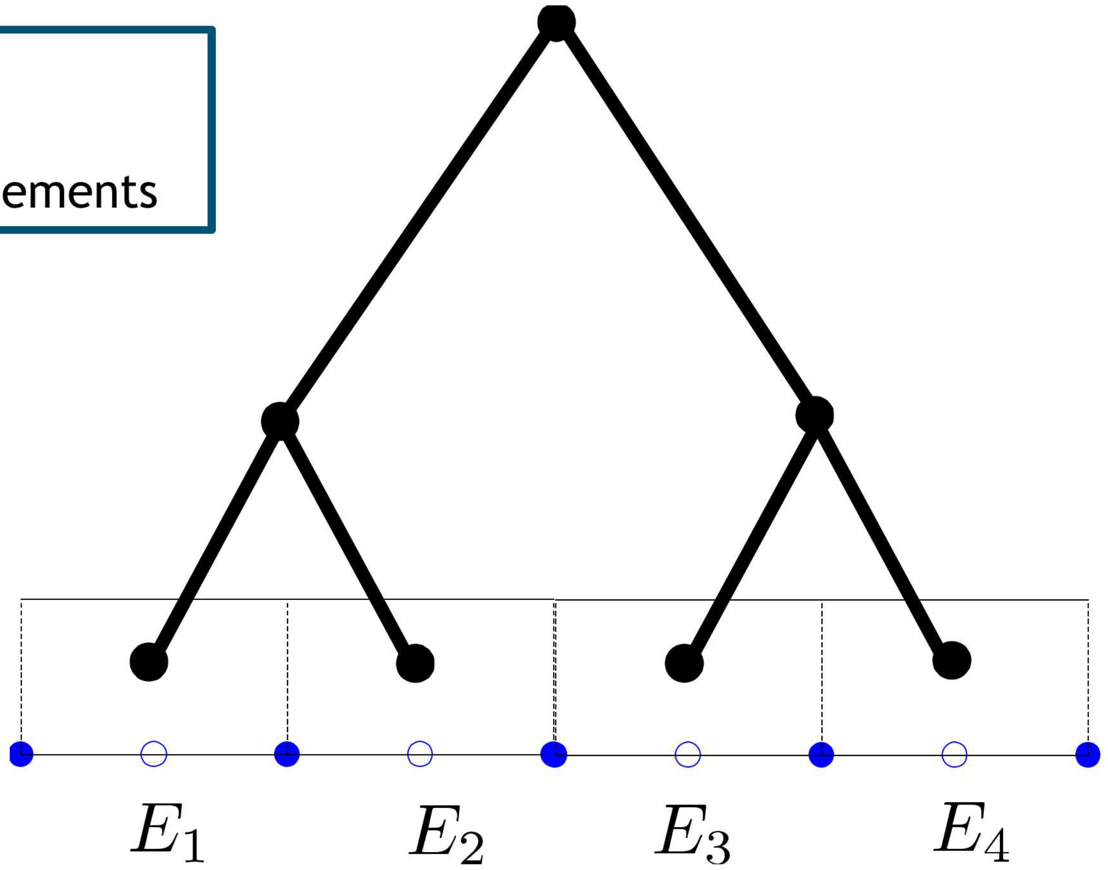

QLT: Leaves to root

<u>Each leaf node</u>
- Computes local bounds
- Sends bounds and mass to its parent

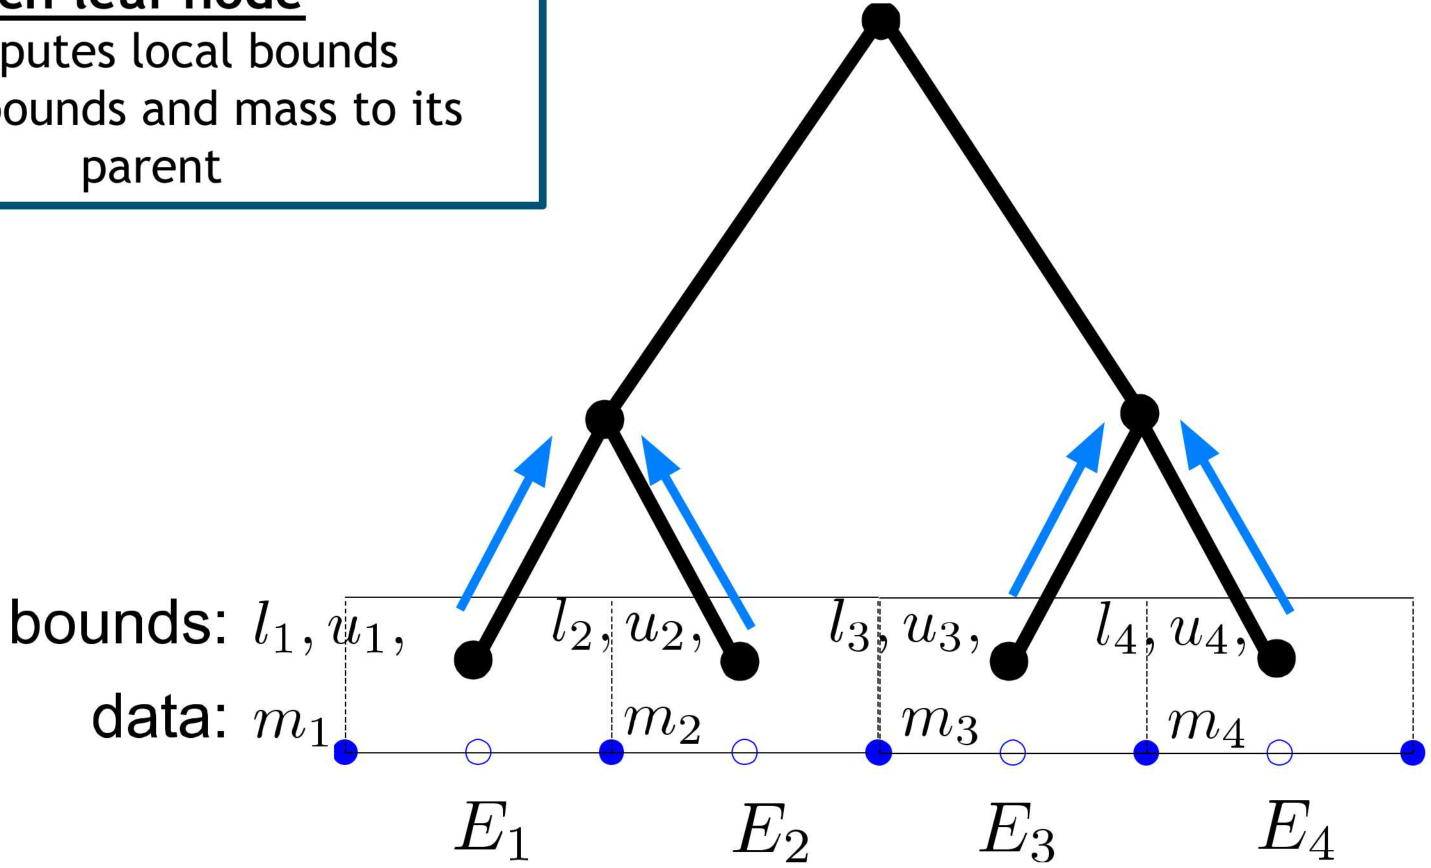

QLT: Leaves to root

Each internal node

- Sums bounds and mass from its kids
- Sends bounds and mass to its parent

$$l_{12} = l_1 + l_2,$$
$$u_{12} = u_1 + u_2,$$
$$m_{12} = m_1 + m_2$$

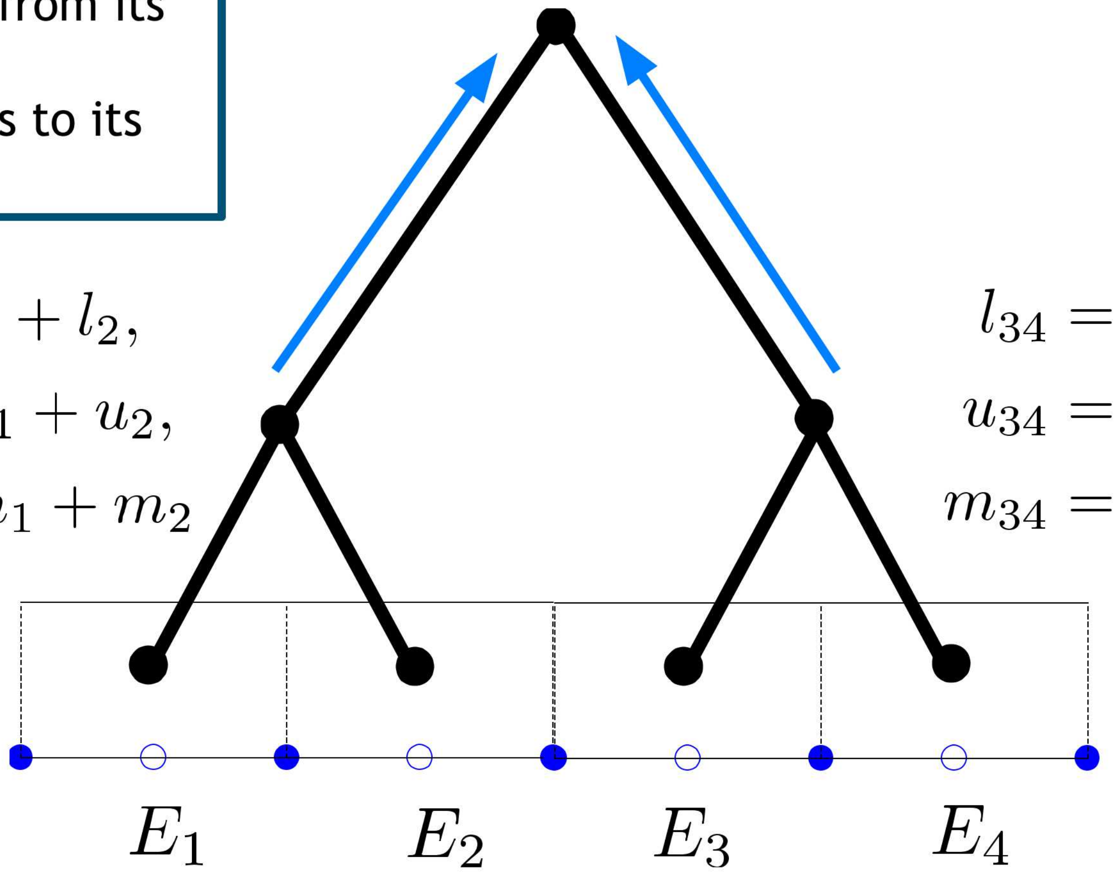

$$l_{34} = l_3 + l_4,$$
$$u_{34} = u_3 + u_4,$$
$$m_{34} = m_3 + m_4$$

QLT: Leaves to root

$$M_g = m_{12} + m_{34},$$

Cell mass boundedness: $l_{12} + l_{34} \leq M_g \leq u_{12} + u_{34}$

<u>Root node</u>

- Checks feasibility
- Global dynamic bounds also computed (same reduction step)
- If cell mass boundedness is not satisfied, global bounds are used instead
- Guaranteed feasible

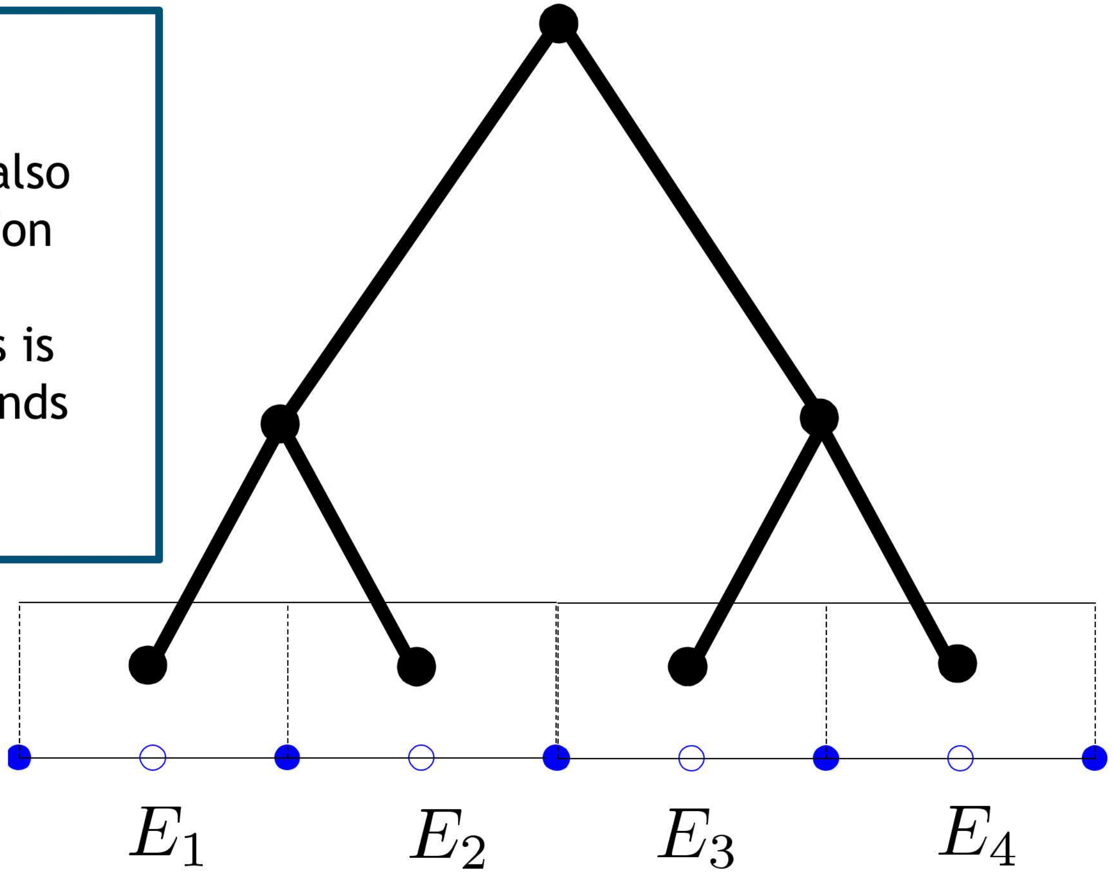

QLT: Root to leaves

$$\min_{m^*} \left\| \begin{matrix} m_{12} - m_{12}^* \\ m_{34} - m_{34}^* \end{matrix} \right\|,$$

$$m_{12} + m_{34} = M_g,$$

$$l_{12} \le m_{12}^* \le u_{12},$$

$$l_{34} \le m_{34}^* \le u_{34}$$

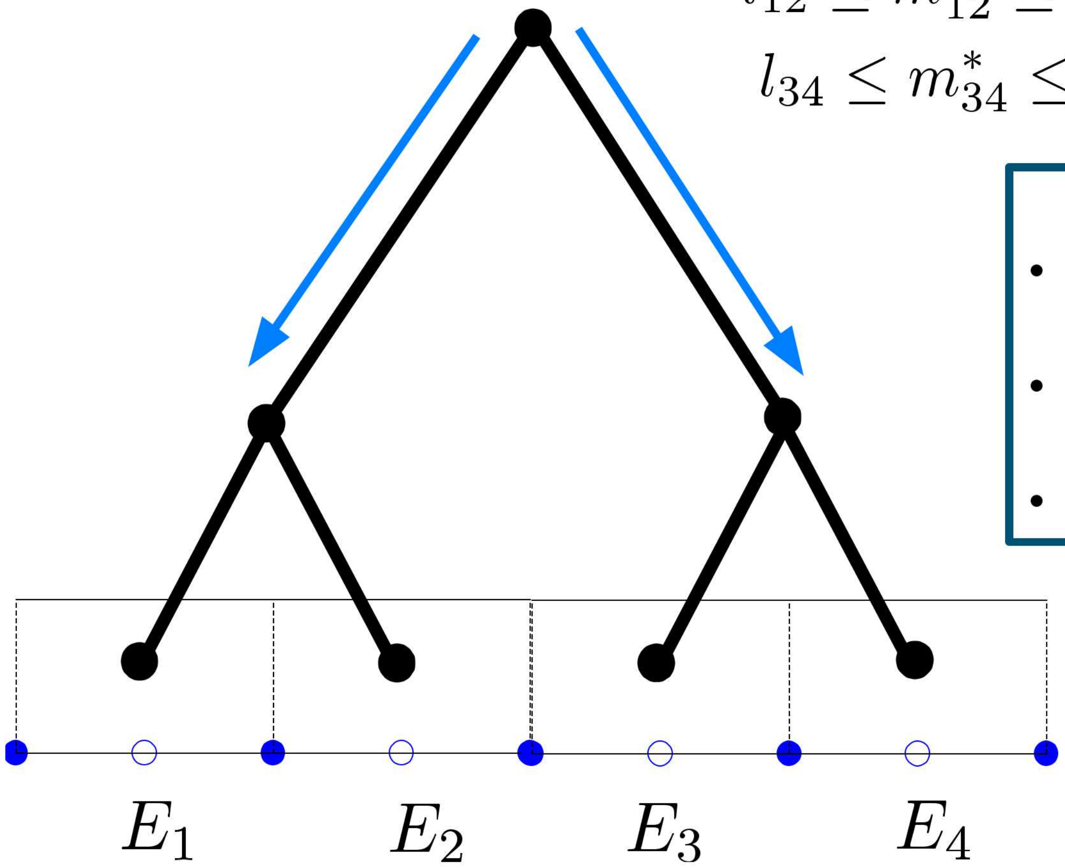

Root node

- Corrects for conservation (if necessary)
- Solves node-local optimization problem for global mass
- Sends results to its kids

QLT: Root to leaves

Internal nodes

- Correct for conservation (if necessary)
- Solve node-local optimization problems
- Send results to its kids

$$\min_{m^*} \left\| \begin{matrix} m_1 - m_1^* \\ m_2 - m_2^* \end{matrix} \right\|,$$
$$m_1^* + m_2^* = m_{12}^*,$$
$$l_1 \le m_1^* \le u_1,$$
$$l_2 \le m_2^* \le u_2$$

$$\min_{m^*} \left\| \begin{matrix} m_3 - m_3^* \\ m_4 - m_4^* \end{matrix} \right\|,$$
$$m_3^* + m_4^* = m_{34}^*,$$
$$l_3 \le m_3^* \le u_3,$$
$$l_4 \le m_4^* \le u_4$$

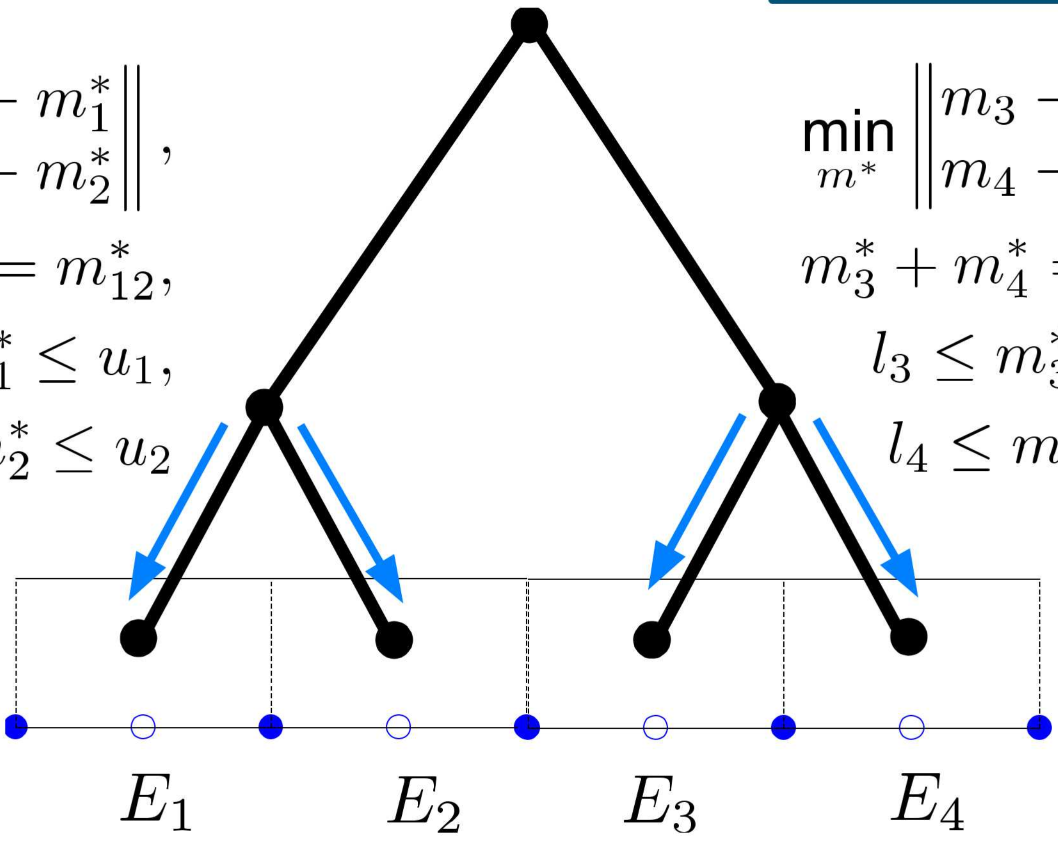

QLT: Root to leaves

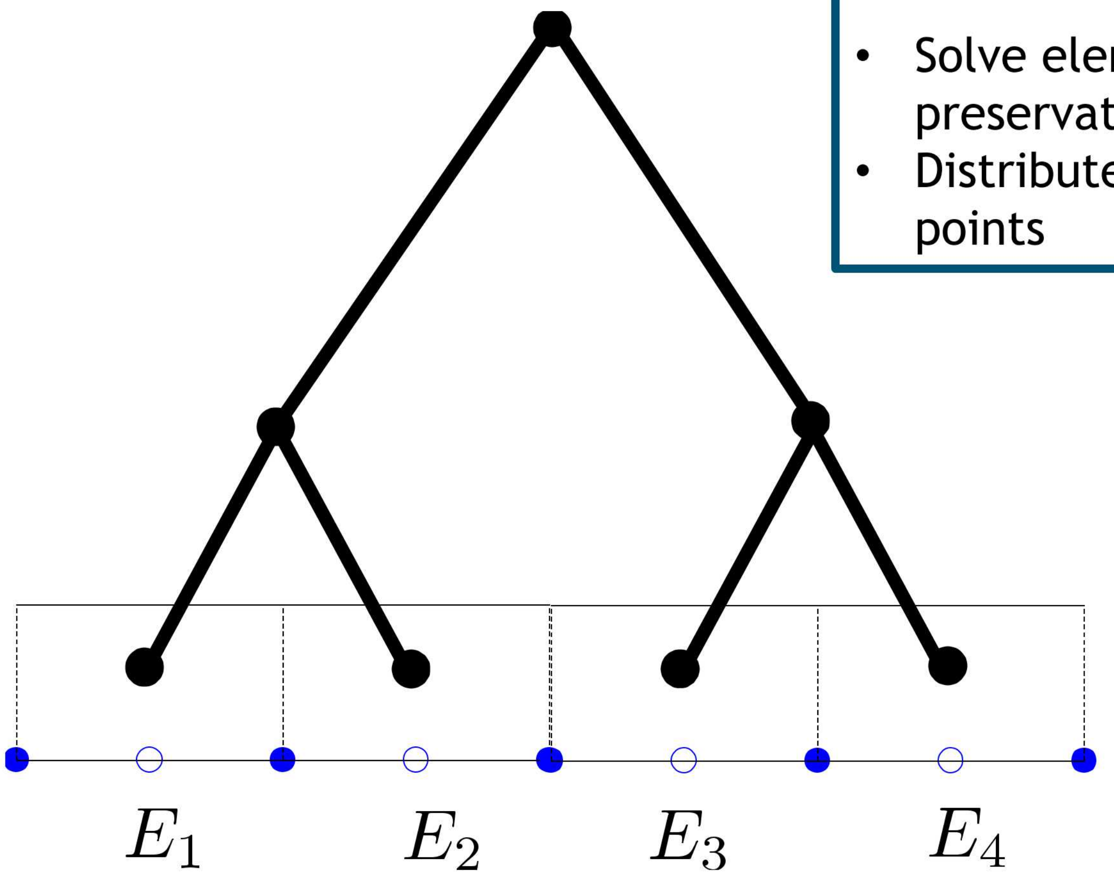

Leaf nodes

- Solve element-local shape preservation problems
- Distribute mass across its quadrature points

QLT: With a conservative numerical transport scheme

- If the numerical solution is already conservative (e.g., cell-integrated SL) QLT redistributes mass locally (within the tree)
- QLT is more efficient and less dissipative than any other constrained density reconstruction

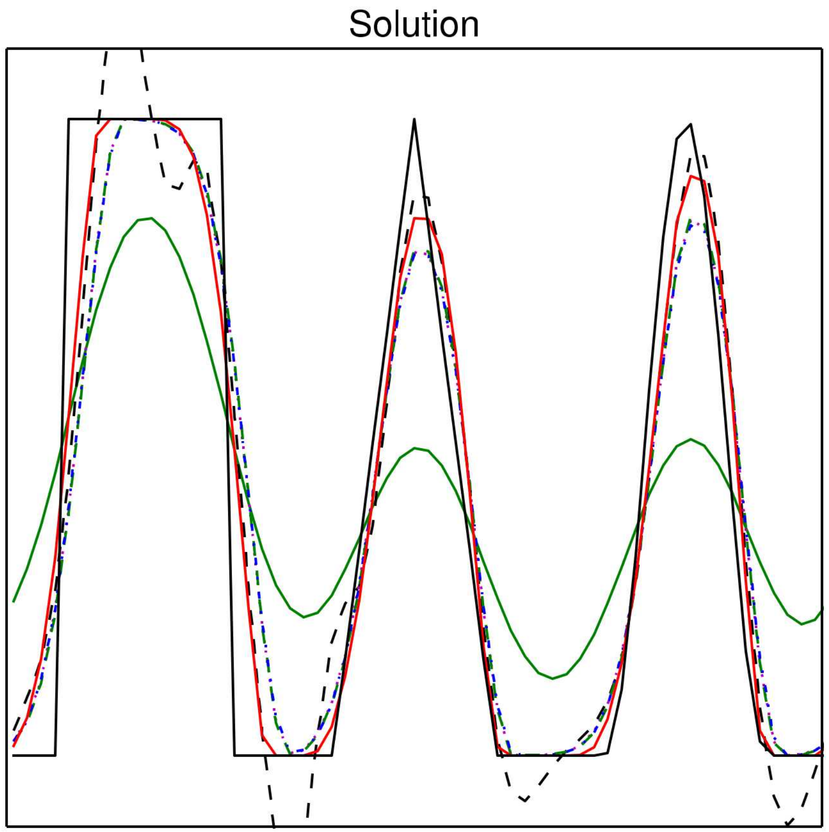

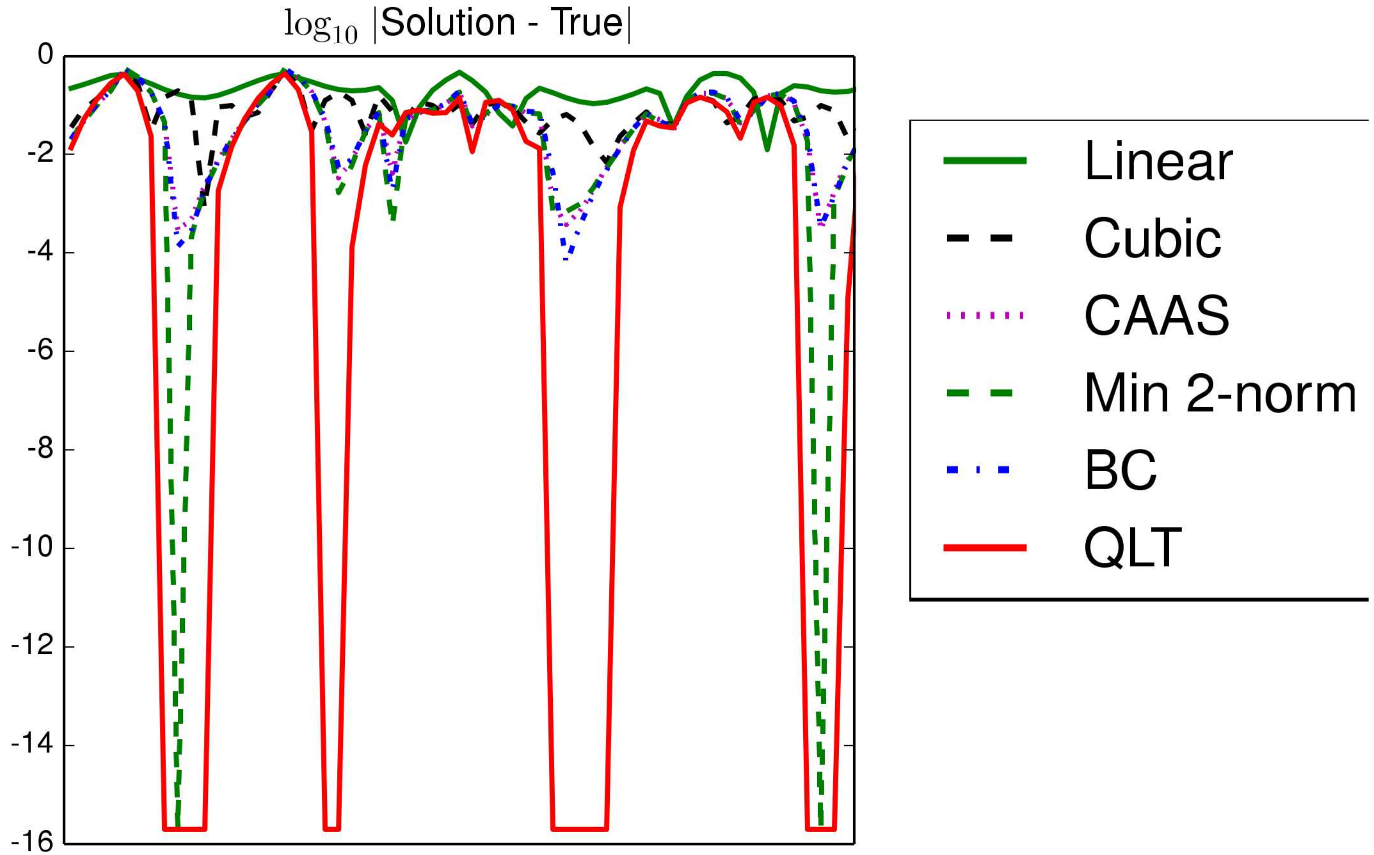

QLT: With a non-conservative numerical transport scheme

- If the numerical solution is not conservative (e.g., pointwise interpolation SL) a global mass fix is required, independent of the shape preservation problem
- QLT shape preservation still acts locally

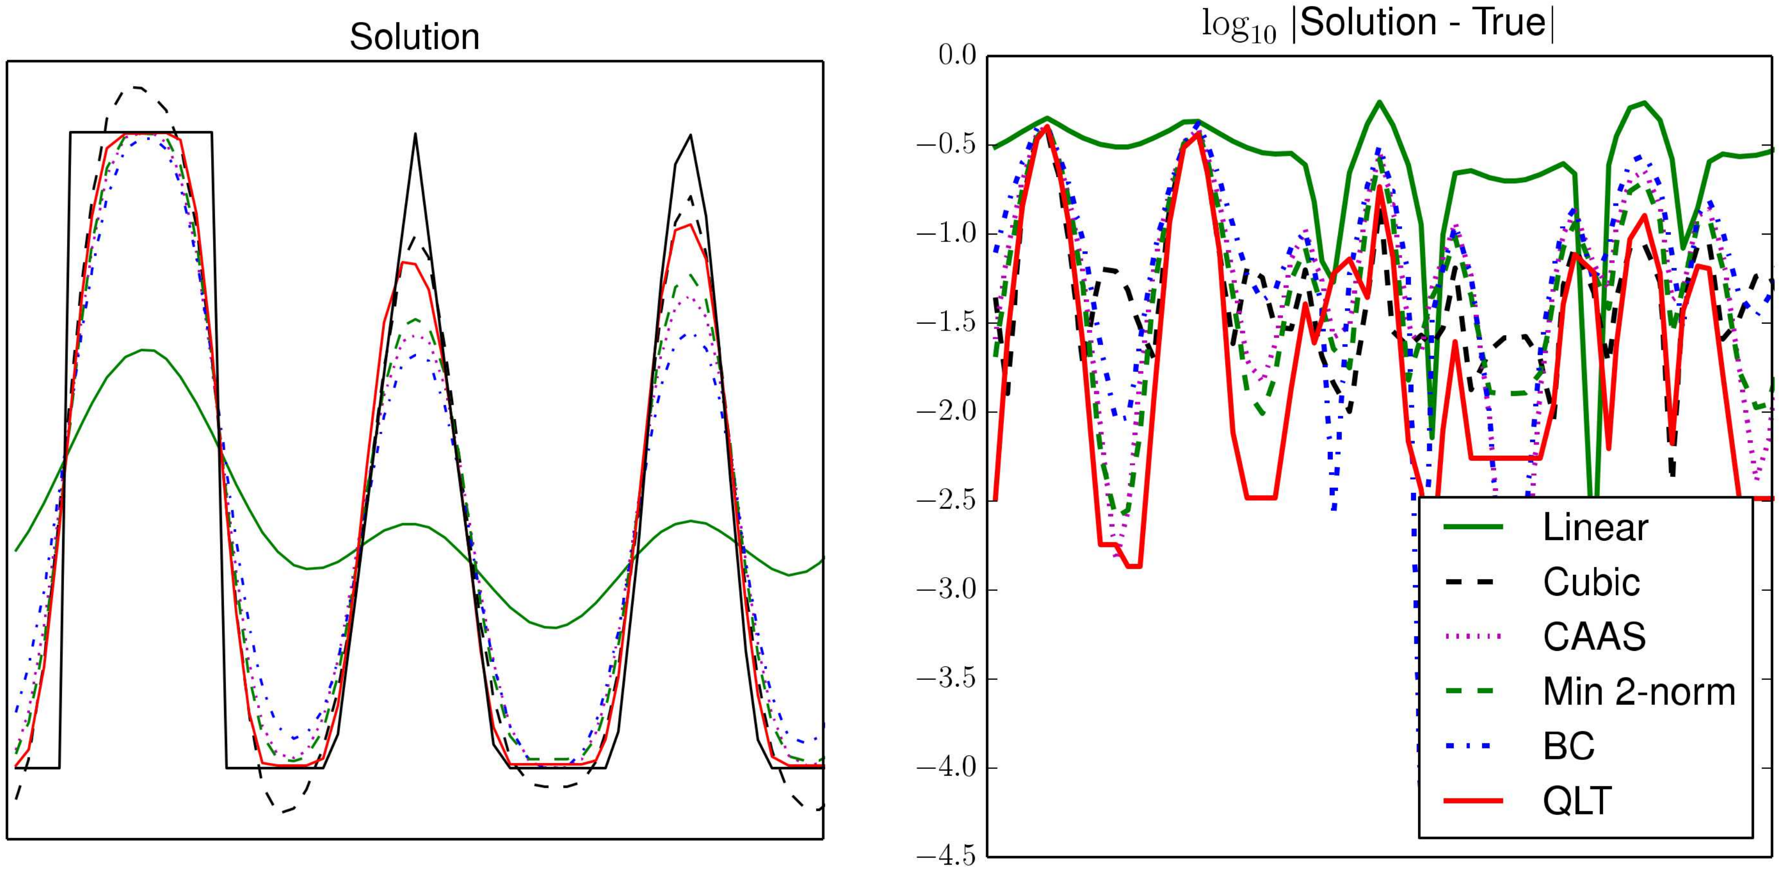

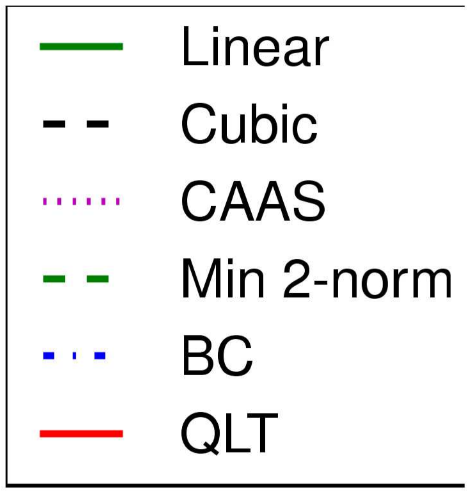

Check-in

2 SL methods + QLT

- Cell-integrated SL, **2.6x faster**
- Pointwise interpolation SL, **3.1x faster**

MPI still the bottleneck

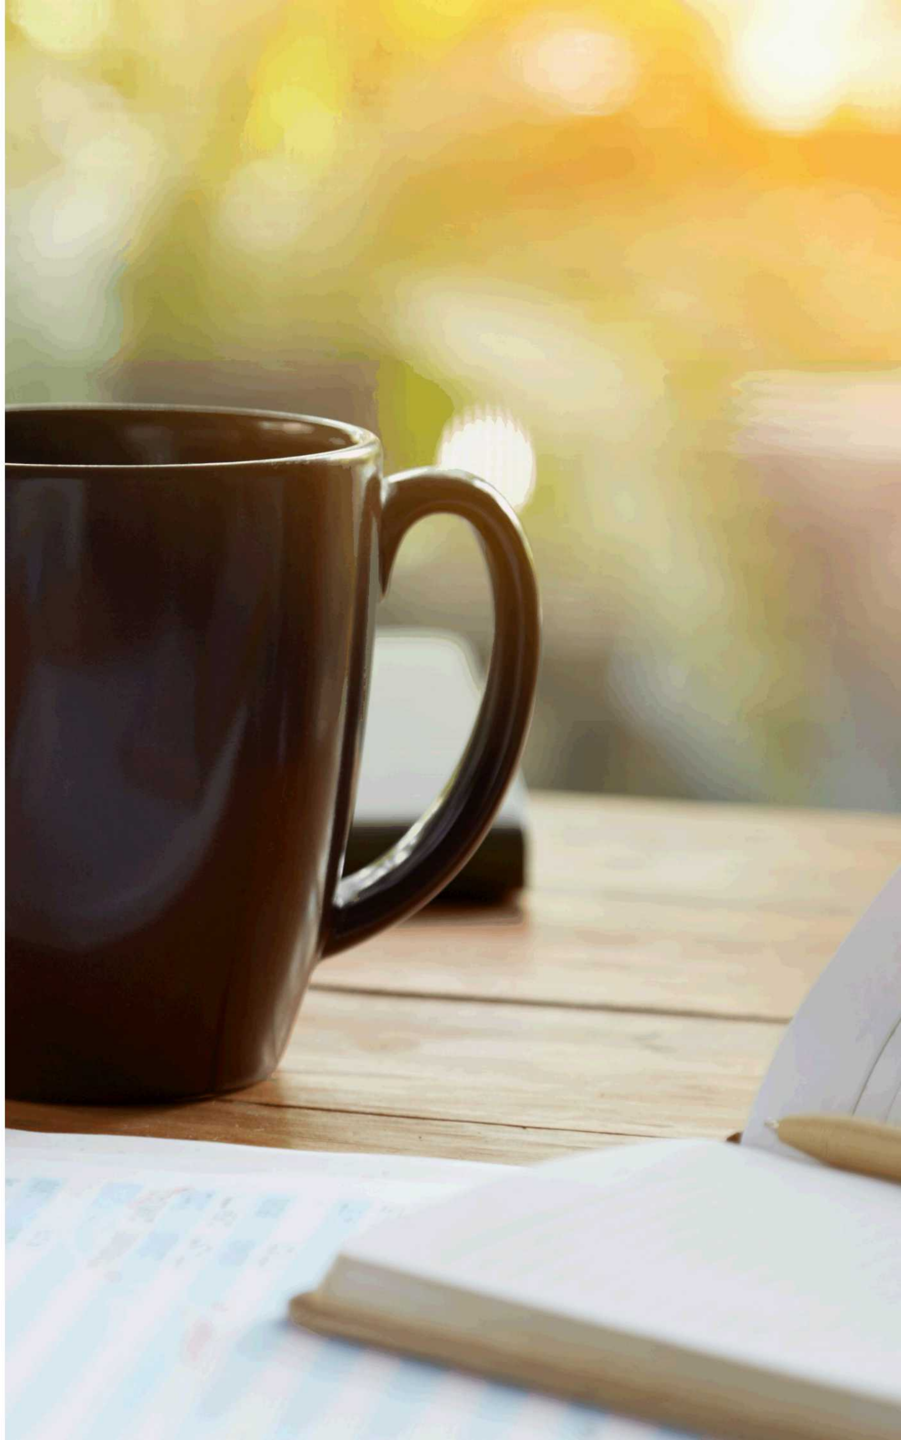

Halo-1 Communication patterns

- Trajectories computed locally
  - McGregor, *MWR*, 1993
- Data transfer: Full halo exchange
- Deterministic, constant
  - **Blue** receives data from **red**
  - Simple: send all
    - 8 x16 = **128 columns**
  - Optimal: send unique
    - 10 x 10 – 16 = **84 columns**

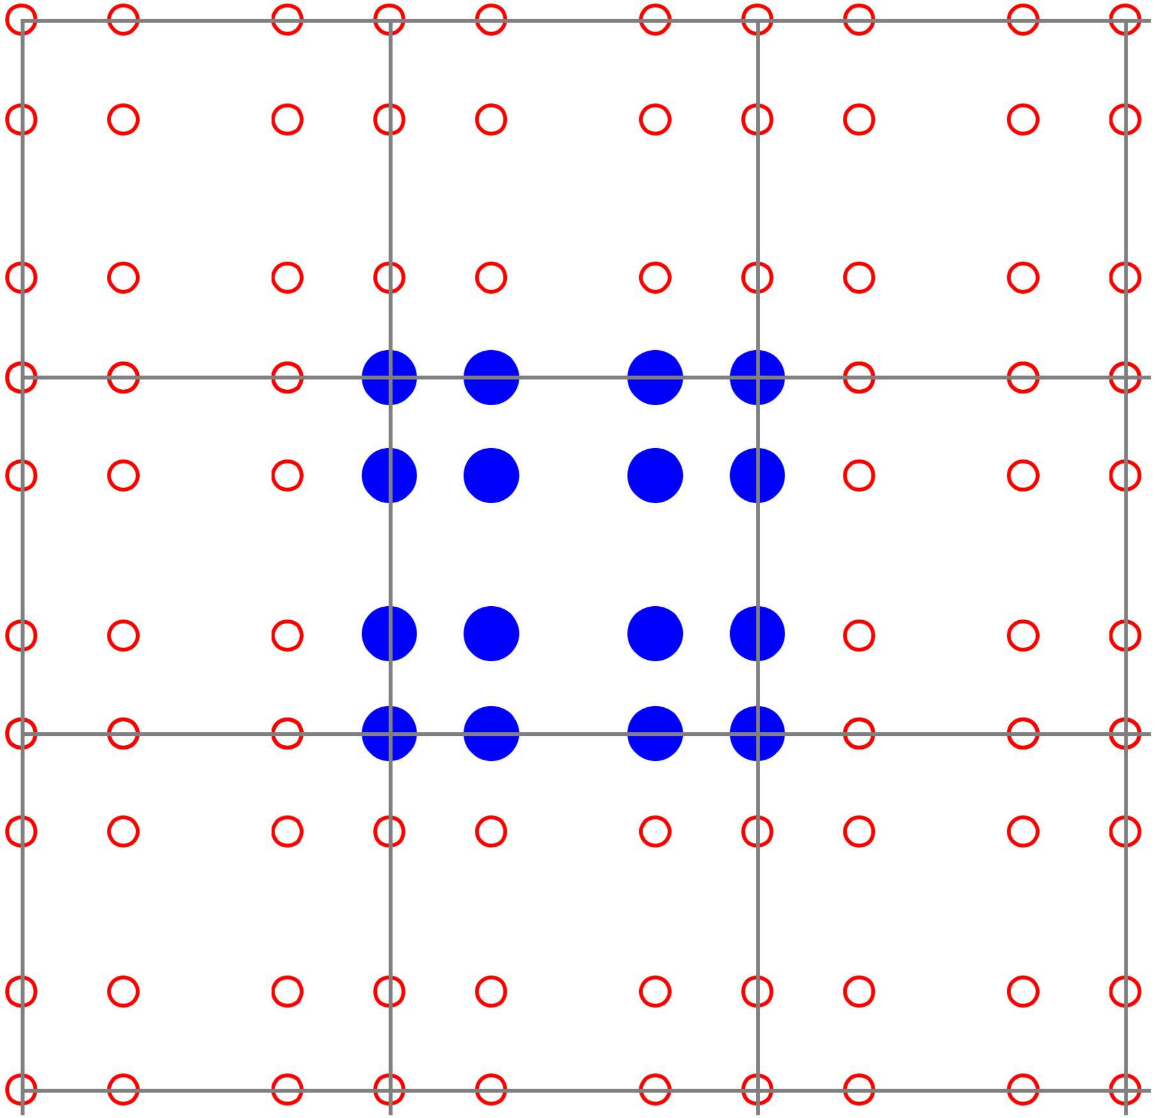

Halo-1 Communication patterns

- Trajectories computed locally
  - McGregor, *MWR*, 1993
- Data transfer: "Upwind"
- Time-varying, flow-dependent
- Step 1: Handshake with halo
  - **Blue** determines source **(red)** elements, sends trajectories
  - Asynchronous, negligible cost
- Step 2:
  - **Blue** receives data (elem. min/max and departure points) from **red**
  - **Illustrated:** 3 x 2 + 12 = **18 columns**
  - **Upper bound:** 8 x 2 + 16 = **32 columns**

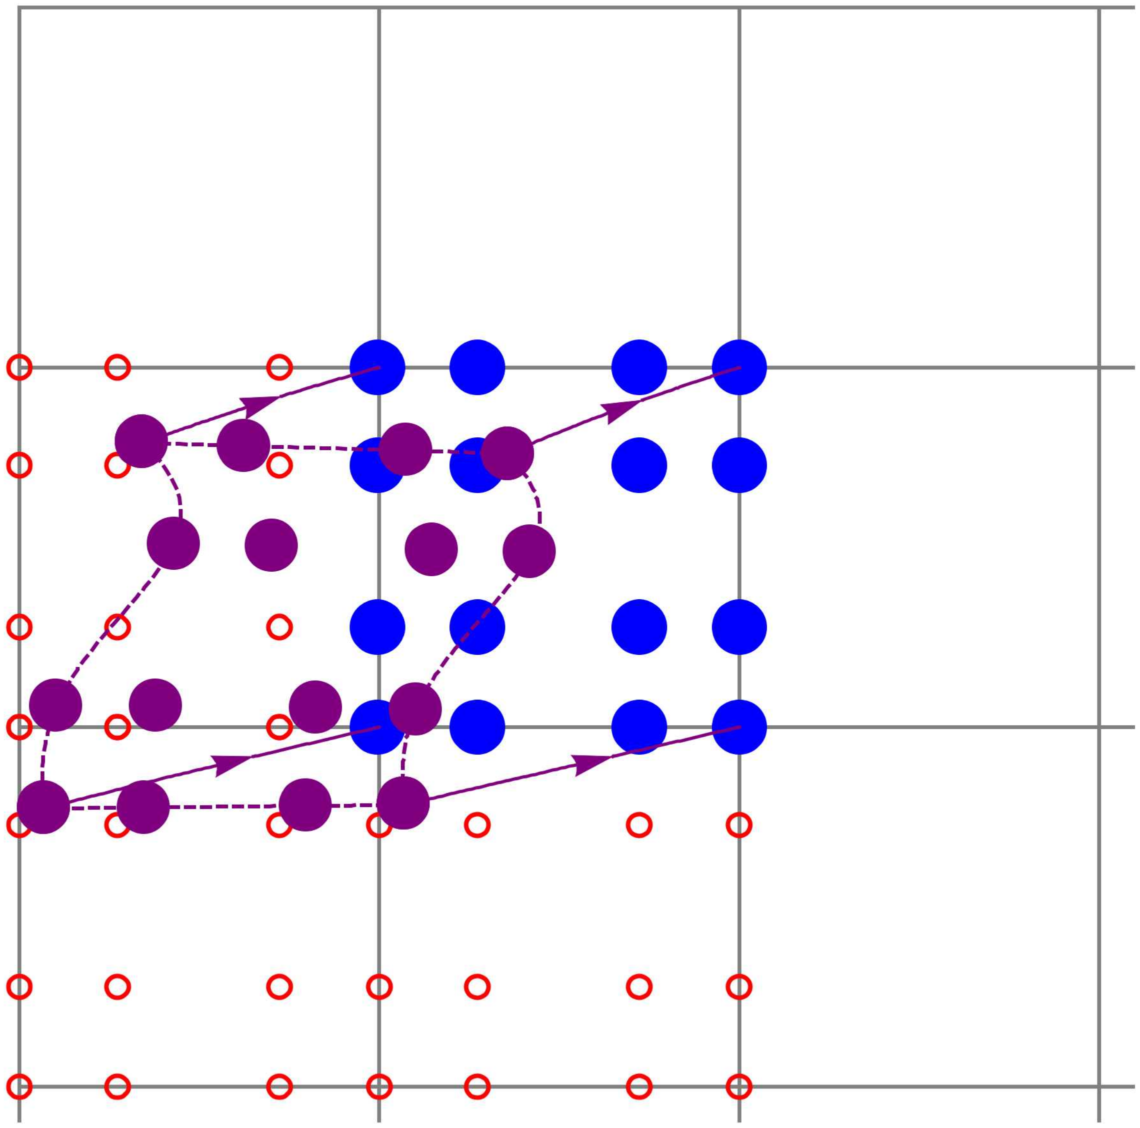

Upwind MPI performance

- Speedup over full halo (higher is better)
- Pointwise interpolation SL with QLT
  - **Colors**: Various resolutions, Courant numbers
- Haswell: Mulitcore, std. CPU
  - 12 nodes
  - 32 cores/node
  - Speedup levels off as workload increases
- KNL: Manycore accelerator
  - 54 nodes
  - 68 cores/node (64 used)
  - Speedup increases proportionally to amount of work

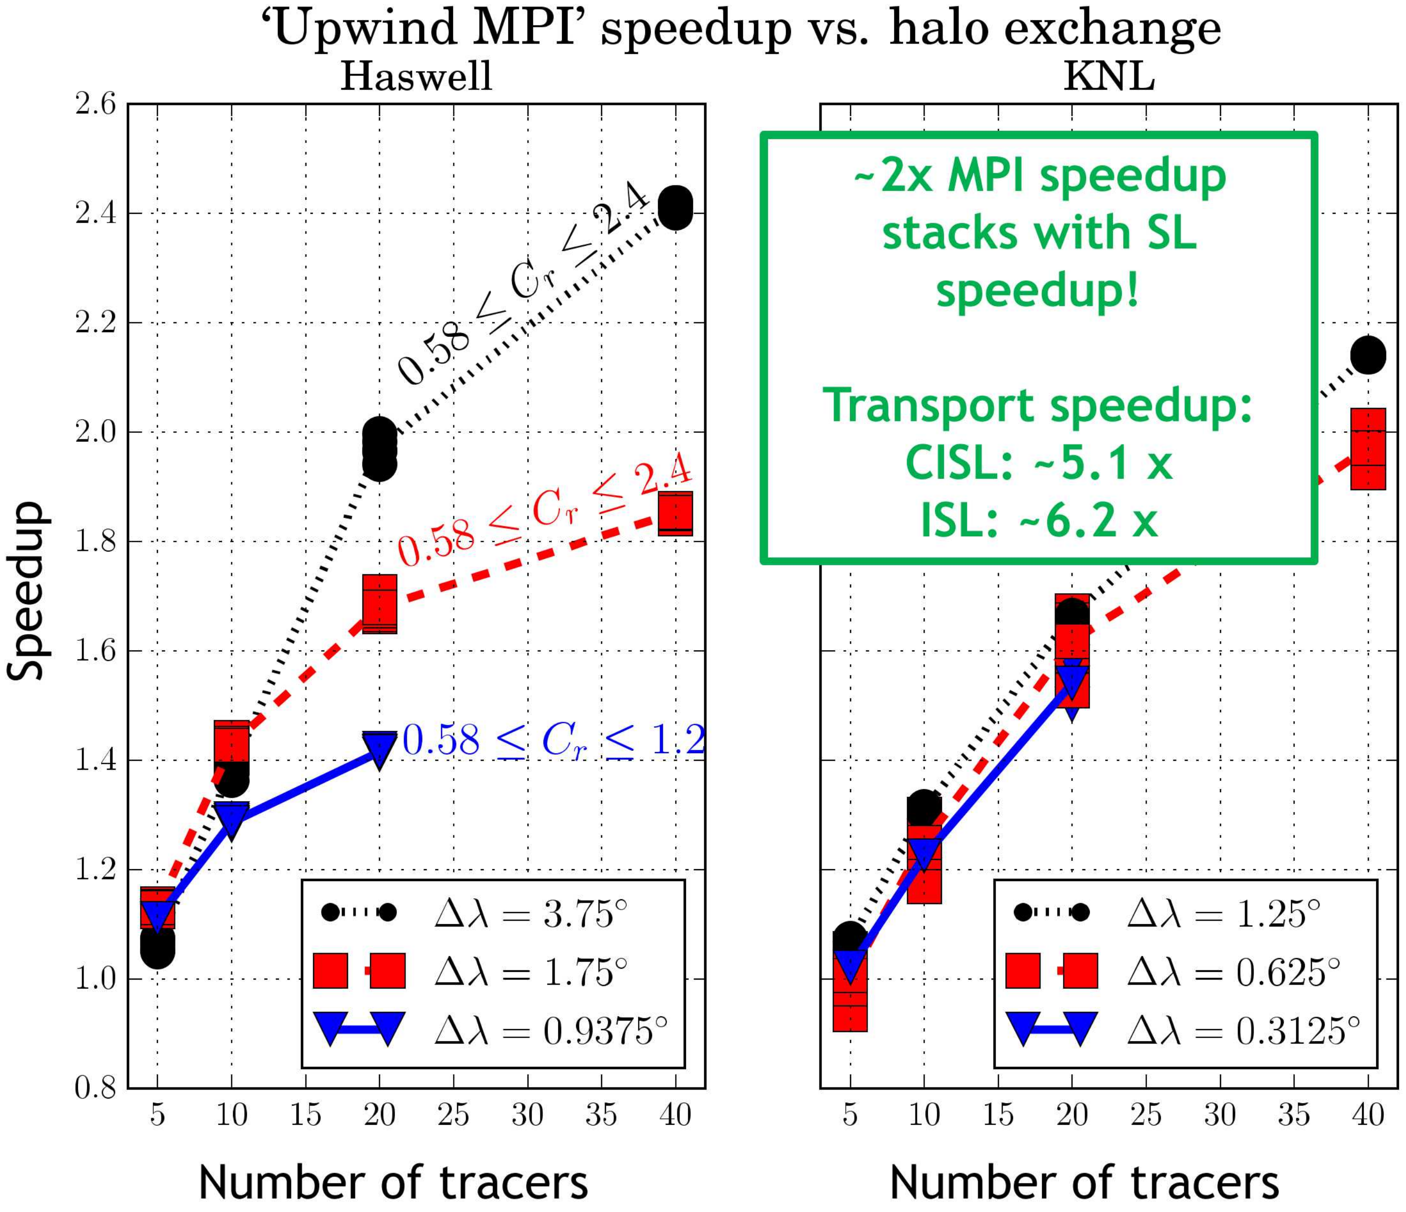

6x faster, and more accurate

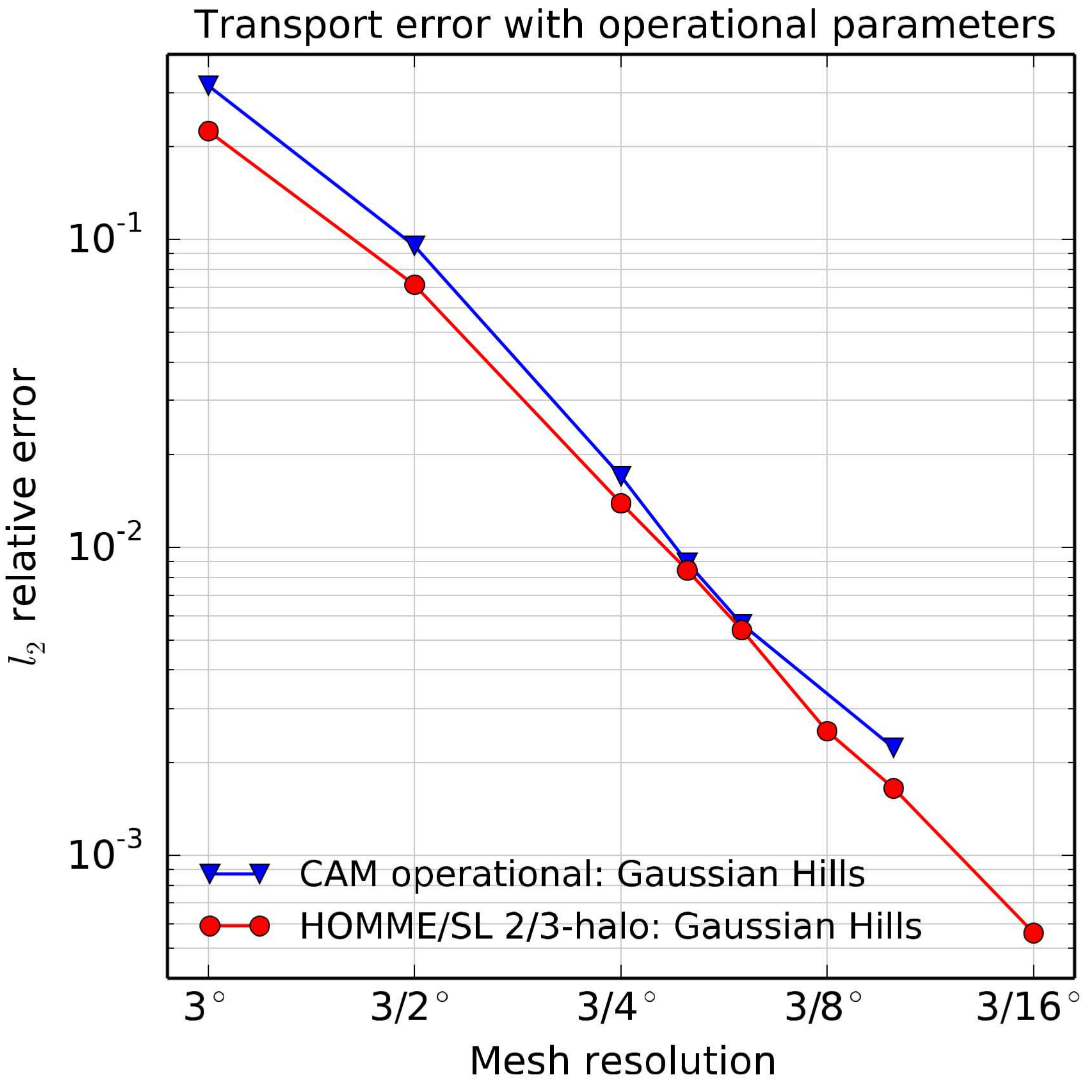

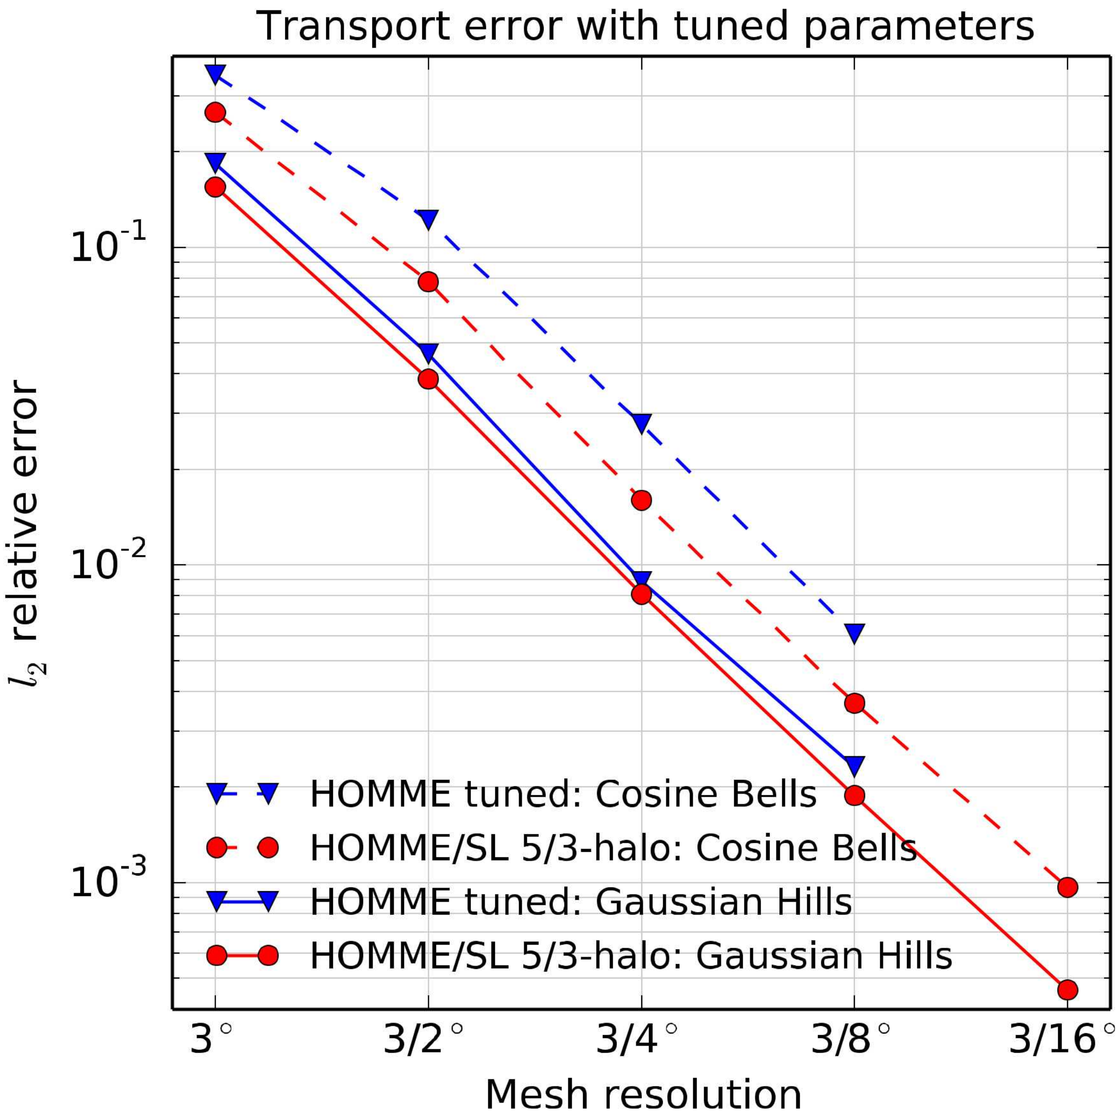

6x faster and less dissipative

Test case: DCMIP 2016 Moist baroclinic instability; day 30 specific humidity at appx. 500 hPa

Eulerian

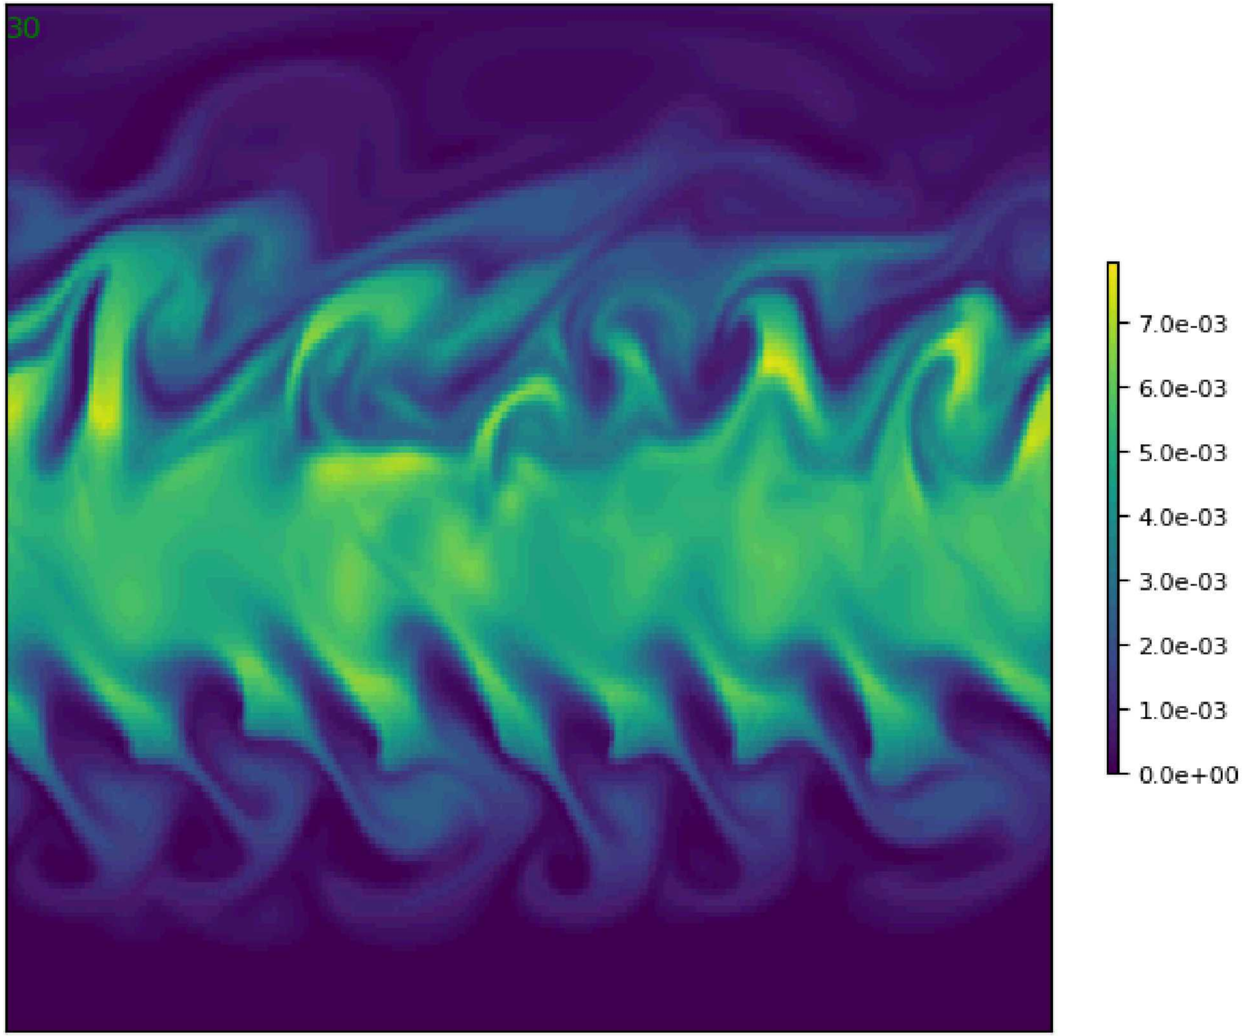

ISL/QLT

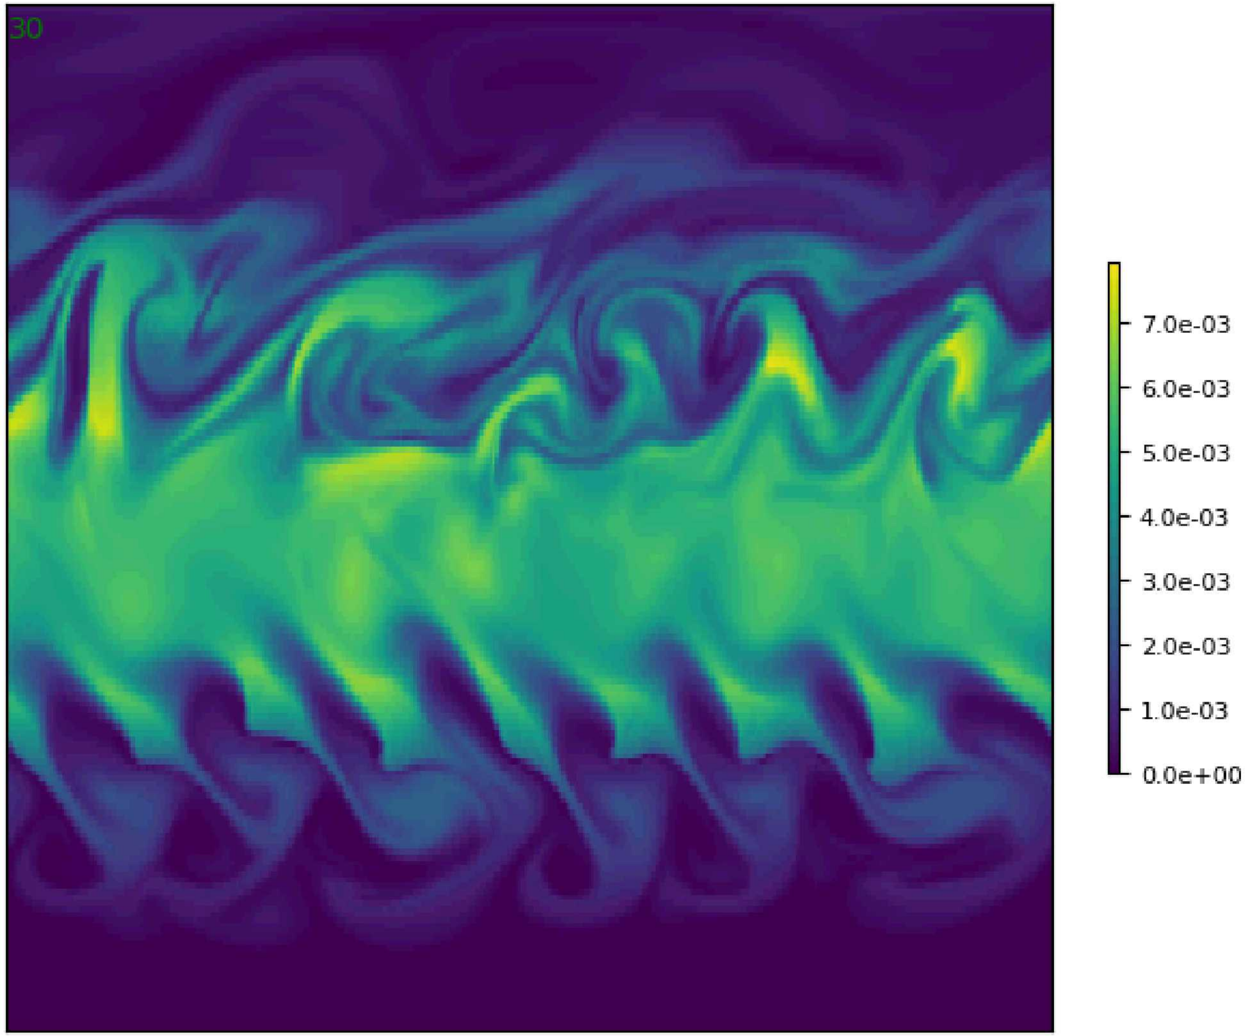

Atmosphere dynamical core performance

- SYPD (higher is better)
  - 0.25 deg global resolution
- Transport + dynamics
  - No physical parameterizations
- Solid: Eulerian SE transport
- Dashed: Pointwise SL transport + QLT
- Red: Cori (KNL)
- Green: Edison (HSW)
- Algorithmic speedups: Independent of architecture
- With SL transport, now physical parameterizations are most expensive part of model

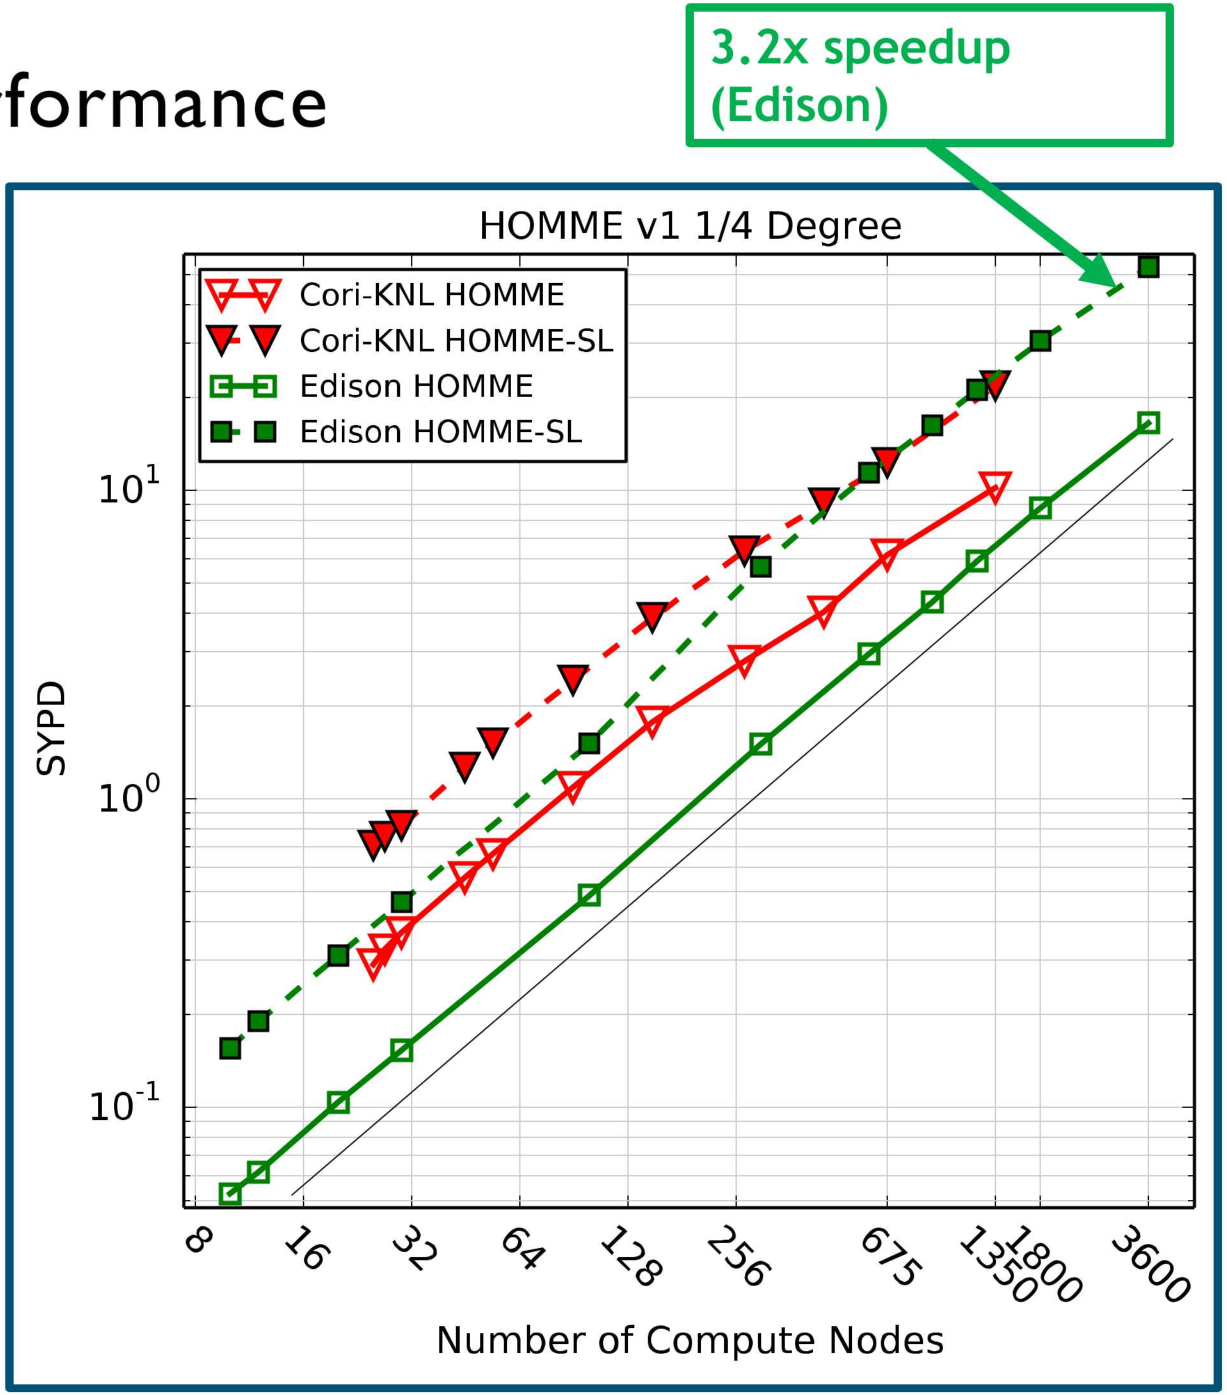

Improving performance of physics

- What is the “effective resolution” of the atmosphere model?
  - Many definitions; all imply that it’s coarser than the GLL mesh
- Idea: Use a coarser grid for column physics
  - Herrington et. al., *MWR*, 2019
  - Berthet et. al., *JAMES*, 2019
  - Hannah, et. al., 2020 (in prep).
- Physical parameterizations are naturally expressed in finite volume form
  - Define finite volume subcells **(red)** of each spectral element **(green/blue)**
  - ``PG2” (2 physical cells per dynamics element) has 4/9 as many columns
  - > 2x computational efficiency
  - Effective resolution argument implies that the answer is approximately the same, at half the cost

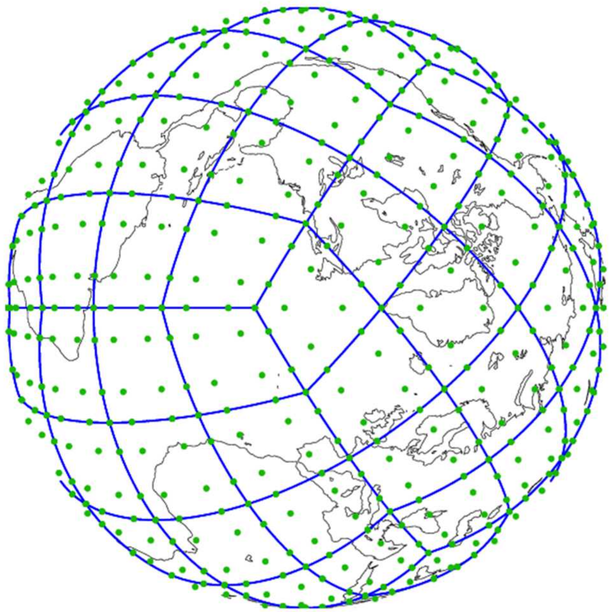

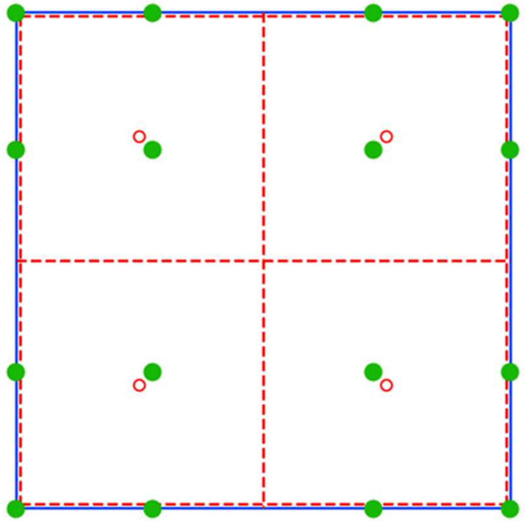

Courtesy of Andrew Bradley

Physgrid remapping algorithms

- Notation:
  - Dynamics variables on dynamics grid $\boldsymbol{d}$
  - Physics variables on physics grid $\boldsymbol{p}$
  - Dynamics variables on physics grid $\boldsymbol{d}'$

- Linear operator requirements
  1. Mass conservation
  2. Remap is element-local
  3. If $A^{p \to d}\boldsymbol{p} = \boldsymbol{d}$, then $A^{d \to p}\boldsymbol{d} = \boldsymbol{p}$
  4. If $\boldsymbol{p} = A^{d \to p}\boldsymbol{d}$ and $\boldsymbol{d} = \mathcal{I}^{d' \to d}\boldsymbol{d}'$, then $A^{p \to d}\boldsymbol{p} = \boldsymbol{d}$

- Reasoning
  - Requirement 2 implies no additional communication
  - Requirements 3 and 4 specify limited forms of idempotence; these help minimize dissipation from remap
  - Requirement 4 assures the remap operator order of accuracy is high as permitted by the physics grid
  - Mathematically, the remap problem is nearly equivalent to the cell-integrated SL algorithm (only a different basis used here)

Courtesy of Andrew Bradley

Physgrid remap operators

- Dynamics to physics, $A^{d\to p}$
  - Simply average GLL density over the physics subcell
  - Satisfies requirement 1, 2
    - (Conservative, no extra comm.)
- Physics to dynamics, $A^{p\to d}$
  - Definition of $A^{d\to p}$ and requirements 2, 4 uniquely specify $A^{p\to d}$ and this satisfies requirement 3
- Add nonlinearity
  - Mass-conserving local limiter
- Communications
  - None in dynamics to physics
  - Limiter requires min/max communications
  - Final DSS to restore continuity

Test case: Remap from dynamics to physics and back

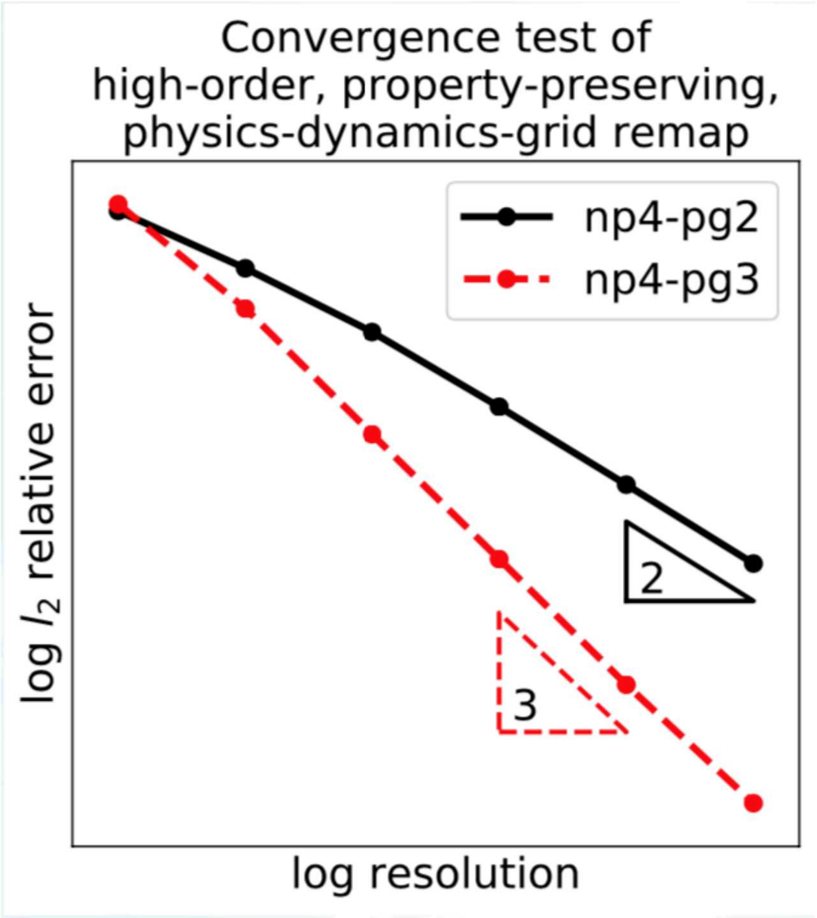

Courtesy of Andrew Bradley

Combined SL transport and PhysGrid

Tracer mass conservation over 1 year full atmosphere simulation

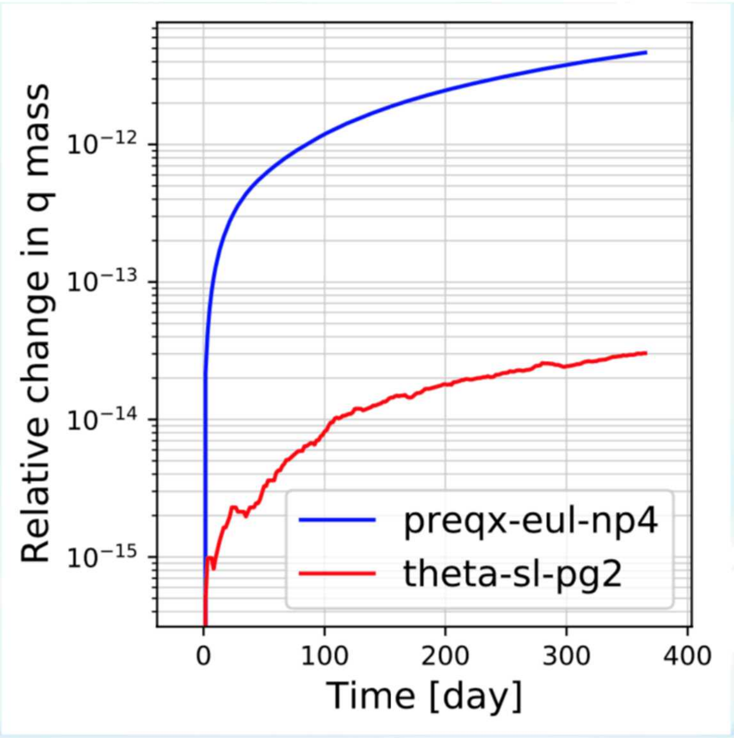

<u>Runtime results (lower is better)</u>

Low-res, 68 Compy nodes

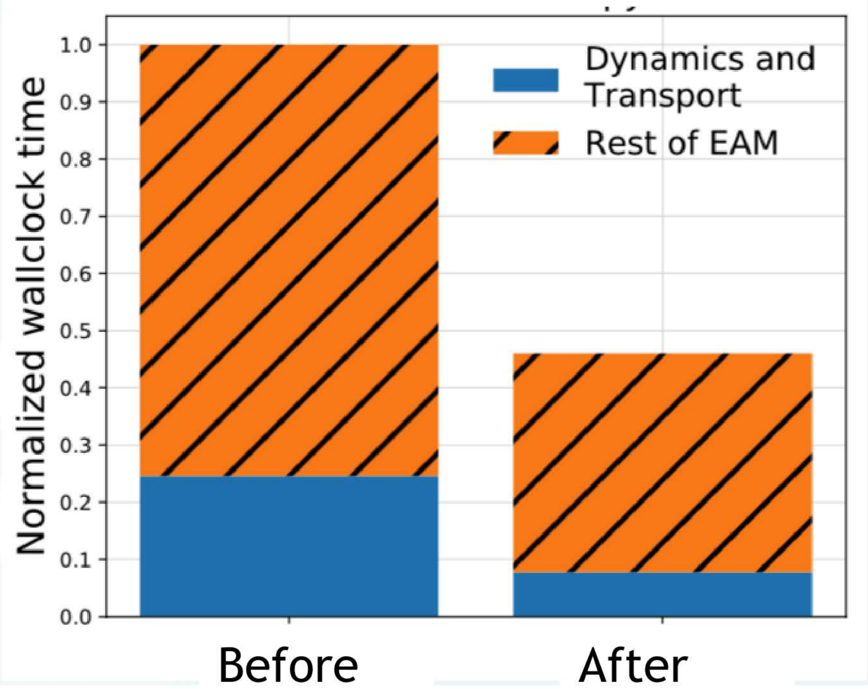

Variable-res, 113 Cori-KNL nodes

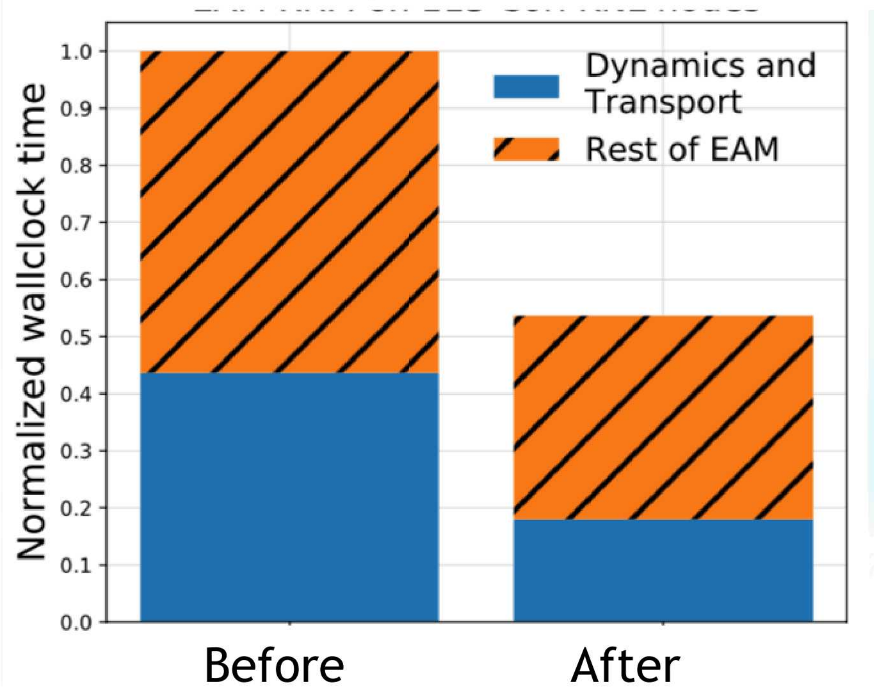

Courtesy of Andrew Bradley

Combined SL transport and PhysGrid

- DCMIP 2016 Moist baroclinic instability; day 25 specific humidity at appx. 600 hPa
- Variable resolution mesh (CONUS)
- Left: Eulerian flux-form transport with physics on the dynamics grid
- Right: SL transport with PG2

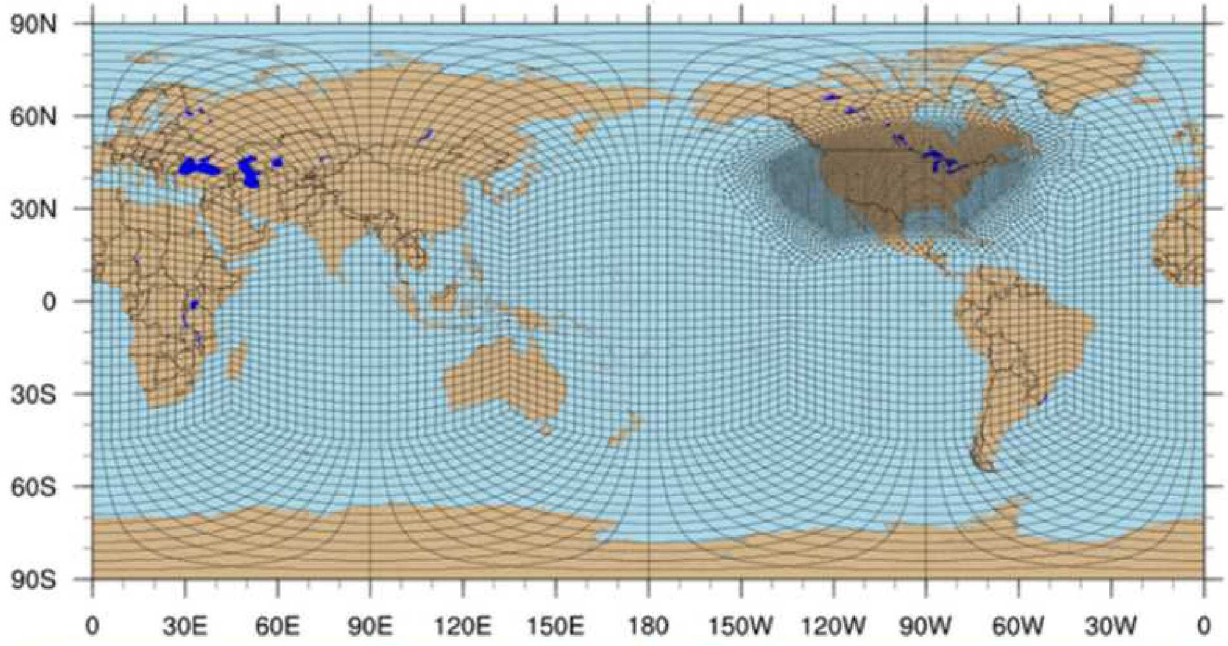

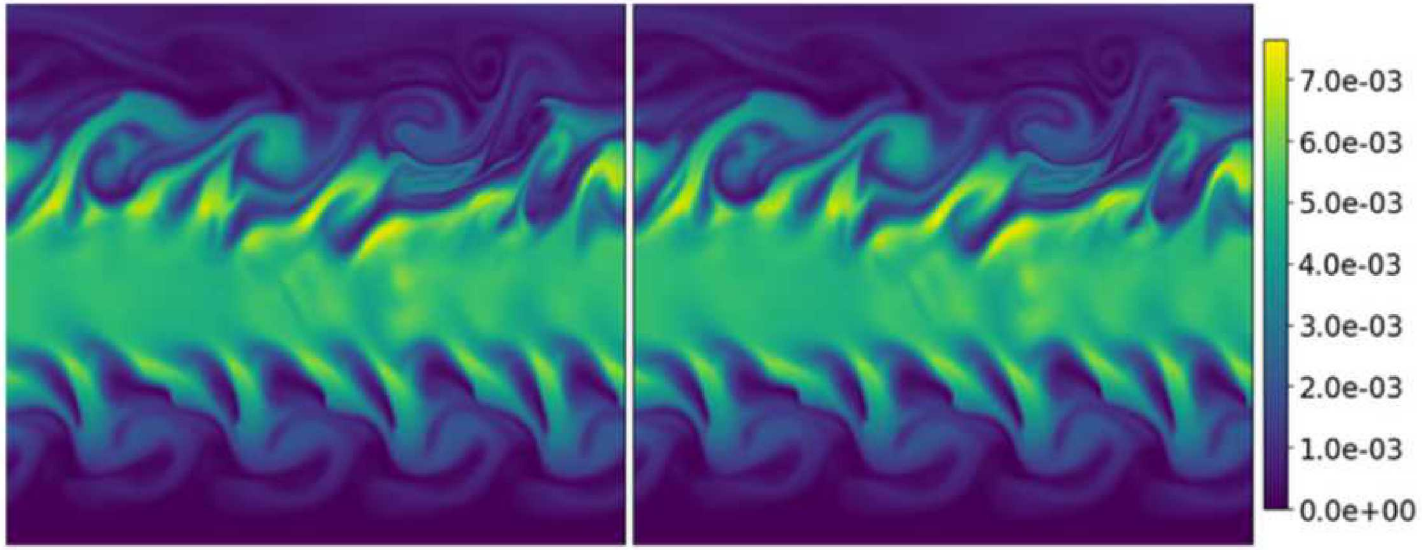

Current and future work

- Implement similar algorithmic improvements in the MPAS-Ocean model
  - Semi-Lagrangian transport for biogeochemistry science campaign
  - Flexible time step coupling methods for super-cycling physical parameterizations
  - Physgrid (Berthet et. al. *JAMES*, 2019)
- Ultra-accurate tracers: Can increase tracer accuracy by up to 100x
  - Combine ideas above, in other direction: Ultra-high order tracer mesh (e.g., 9^{th})
  - Interpolate velocity from dynamics
  - Compute transport on high order mesh
  - Remap tendencies back to dynamics
- Simple, Cloud-Resolving, Exascale Atmosphere Model (SCREAM): 3km global model
  - Aerosol parameterizations

Sandia National Laboratories

Exceptional service in the national interest

Thank you

U.S. DEPARTMENT OF ENERGY NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.