Non-hydrostatic dynamics with multi-moment discontinuous Galerkin methods

Los Alamos National Laboratory

Balu Nadiga, Jorge Urrego-Blanco

Sandia National Laboratories

Greg Barnett, Pete Bosler, Andrew Bradley, Oksana Guba, Nick Nelsen, Mark Taylor

SciDAC-4 PI Meeting, Rockville MD
July 24, 2018

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-NA-0003525. SAND2018-5822 PE.

Outline

- Introduction
- Semi-Lagrangian Multi-Moment (SLMM) methods
  - Cell-integrated (CISL)
  - Interpolation SL
- Performance results
- Background: Algorithm development
- **Quasi-Local Tree-based (QLT) density reconstruction**
- Partnerships
  - E3SM
  - SciDAC Institutes
- Ongoing science

NH-MMCDG Project

- **Primary goal:** Develop Semi-Lagrangian Mulit-Moment (SLMM) methods to create a fast dynamics solver with coupled transport scheme.
- **Intended result:** Atmospheric dynamical core tailored for heterogeneous computing environments, demonstrations of non-hydrostatic dynamics

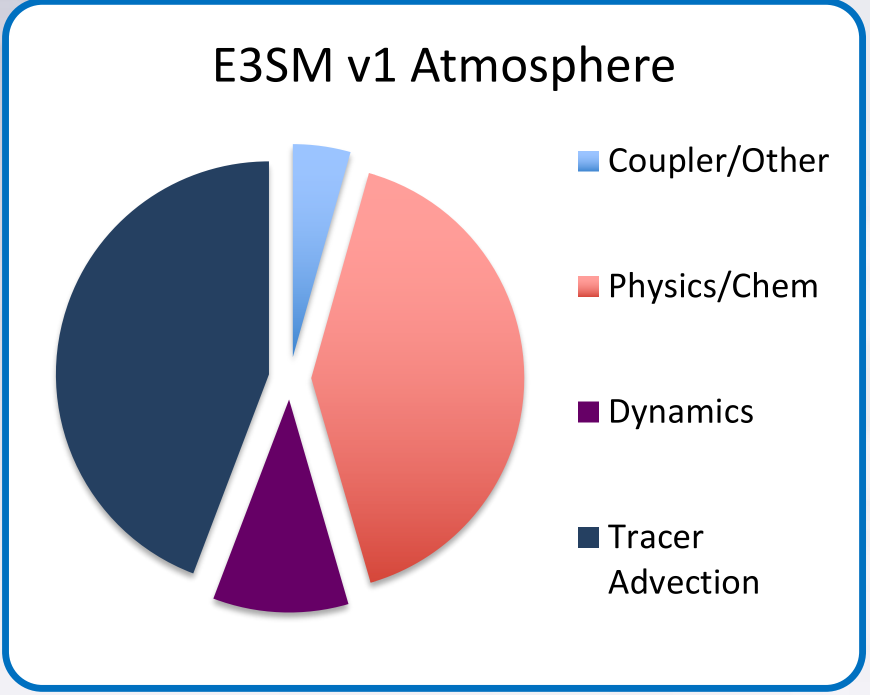

- P. A. Bosler, A. M. Bradley, M. A. Taylor. "Conservative multi-moment transport along characteristics for discontinuous Galerkin methods," submitted to *SIAM J. Sci. Comput*., 2018.
- A. M. Bradley, P. A. Bosler, O. Guba, M. A. Taylor, G. A. Barnett. "Communication-efficient property preservation in tracer transport," submitted to *SIAM J. Sci. Comput.*, 2018.

E3SM v1 Performance

- Atmosphere component has most work
- Atmosphere component scales best
- Transport is most costly

Semi-Lagrangian Multi-Moment (SLMM) methods

- Transport and Mass conservation on unstructured SEM grids
  - Semi-Lagrangian time step allows Cr > 1
  - Spectral element spatial discretization for compact data stencils
- Highlights:
  - Reduced MPI Communication rounds and volume
  - Second-order accuracy overall (with deformational flow, shape preservation)
  - Two cell-integrated (locally conservative) semi-Lagrangian variants
    - High order: allows for later development into higher order methods
    - Low order Jacobian-Combined Transport (JCT): maintains second order accuracy with lower number of nonlinear solves
    - **QLT** enforces shape preservation
    - **Speedup factor > 2.1 vs. current E3SM transport scheme with 40 tracers**
  - Pointwise (classical) semi-Lagrangian interpolation
    - Smallest possible comm. volume for a given SL discretization
    - **QLT** restores mass conservation and enforces shape preservation
    - **Speedup factor > 3.2 vs. current E3SM transport scheme with 40 tracers**

E3SM performance study

- Strong-scaling limit, 1 element per core
- Normalized transport time vs. number of tracers (lower is better)
- Eulerian vs. SL breakeven point < 10 tracers

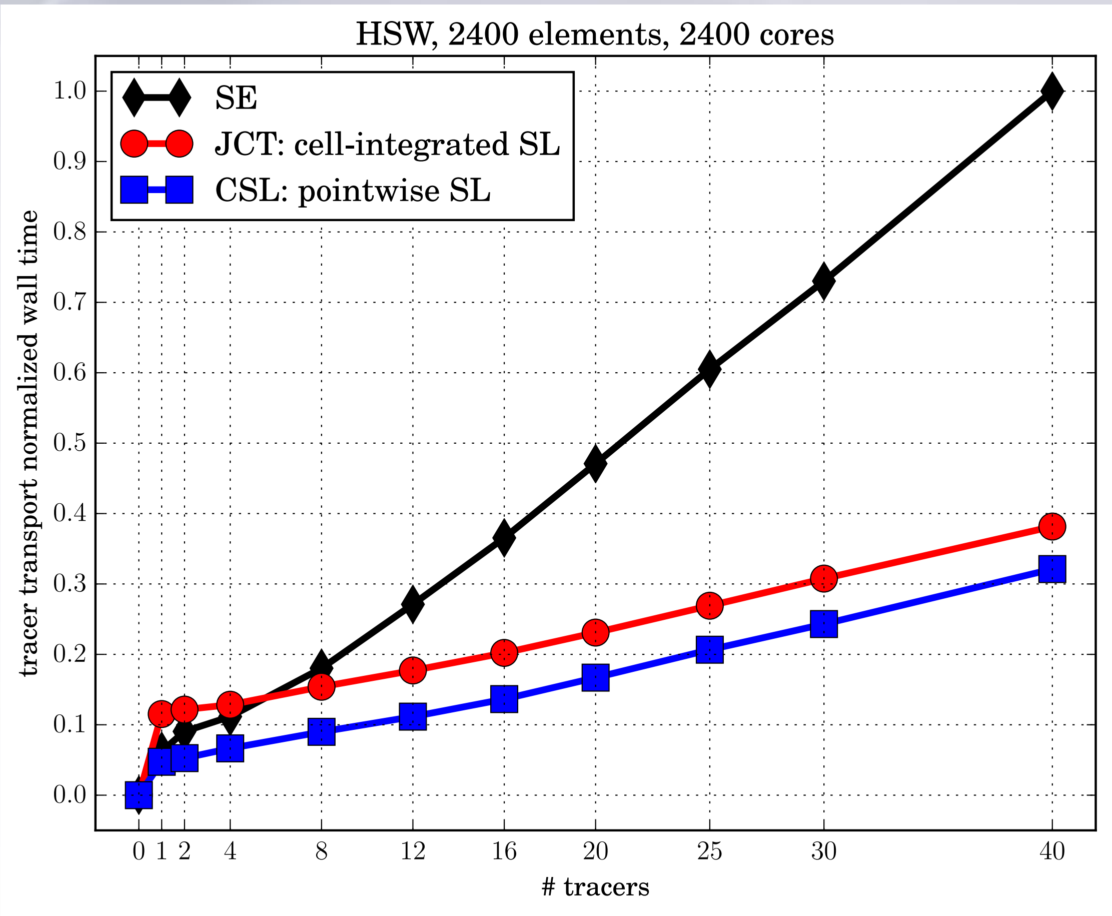

- Spectral element **(SE)**
- Semi-Lagrangian **(SL)**
- Cell-integrated **(JCT)**
- Classical interp. SL **(CSL)**
- **SE to JCT**
  - Speedup factor ~ 2.6
- **SE to CSL**
  - Smaller communication volume (basis-point)
  - Speedup factor ~ 3.1
- **MPI comm. still limits speed**

E3SM Atm. Dycore Performance

- SYPD (higher is better)
- Solid: Eulerian SE transport
- Dashed: Pointwise SL transport + QLT
- Red: Cori (KNL)
- Green: Edison (HSW)

3.2x speedup (Edison)

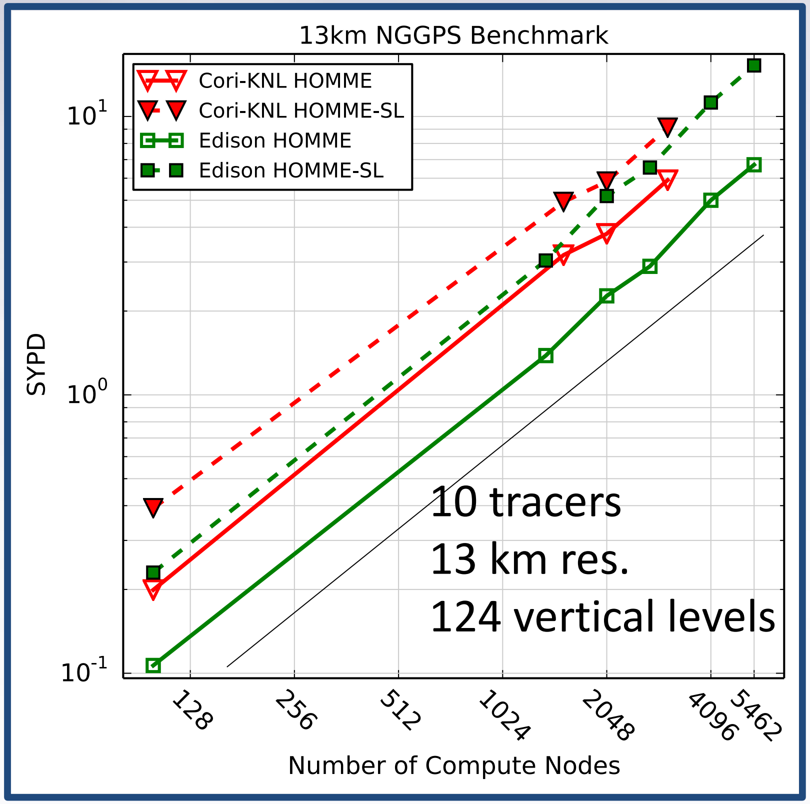

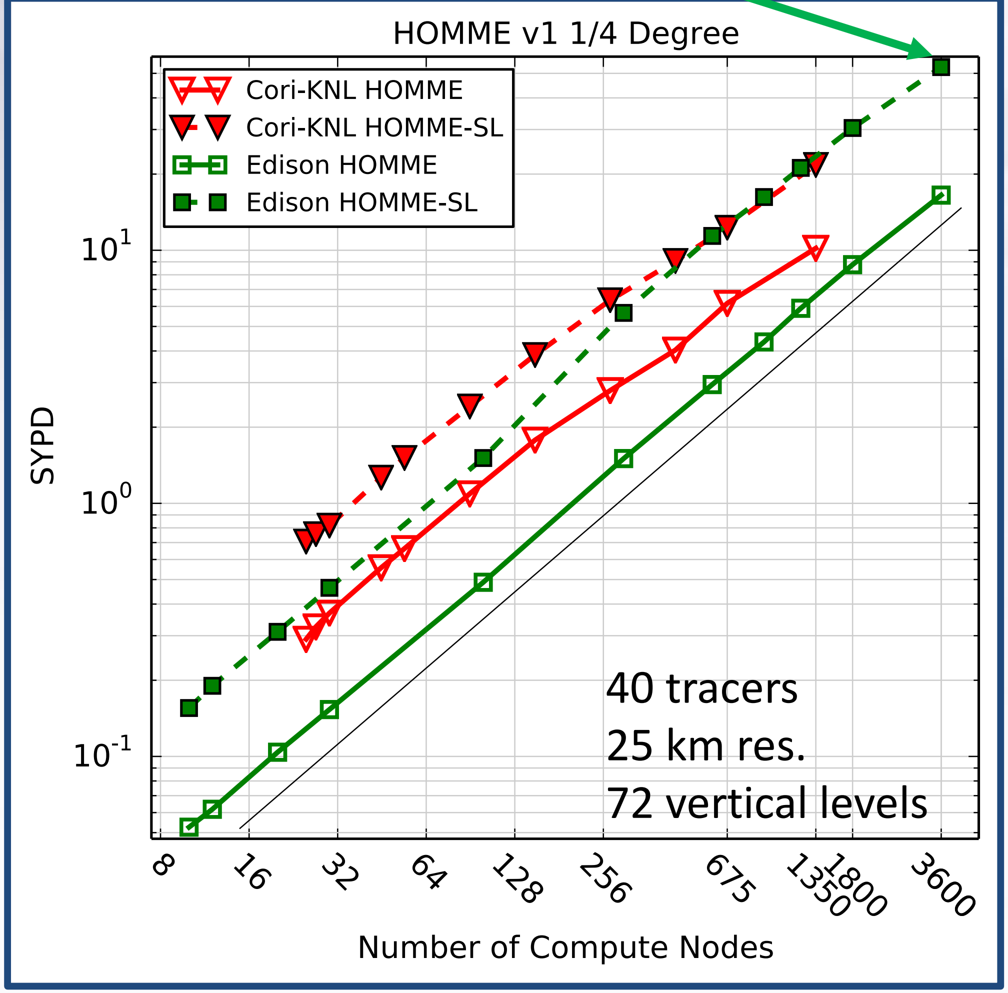

Cell-integrated transport

Given velocity $\boldsymbol{u}(\boldsymbol{x}, t)$ for all $\boldsymbol{x}$ and $t \geq 0$ and initial condition $q_0(\boldsymbol{x}) = q(\boldsymbol{x}, 0)$, solve transport equation for $q(\boldsymbol{x}, t)$ at $t > 0$.

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

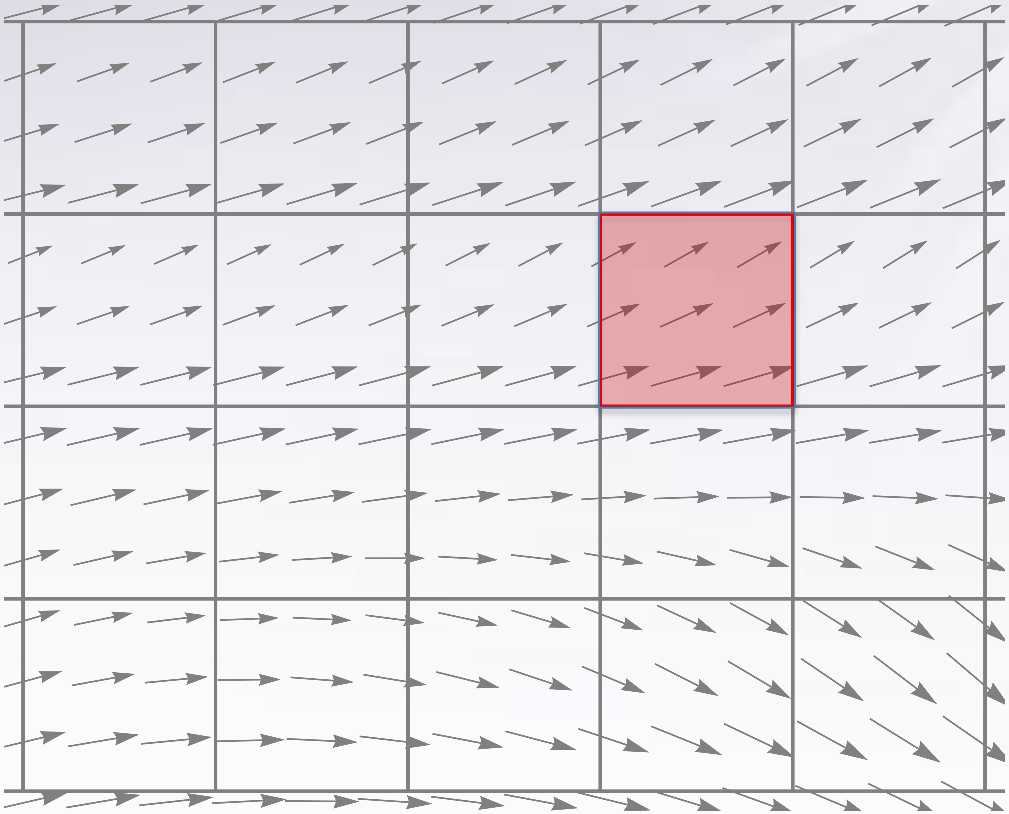

- Requirements
  - Conservation
  - Accuracy
  - Shape preservation
  - Consistency
  - Efficiency

Lagrangian flow map

- Flow characteristics labeled by Lagrangian parameter

$$\frac{d\boldsymbol{x}^*}{dt}(\boldsymbol{a},t) = \boldsymbol{u}(\boldsymbol{x}^*(\boldsymbol{a},t),t),$$
$$\boldsymbol{x}^*(\boldsymbol{a},t_n) = \boldsymbol{a}$$

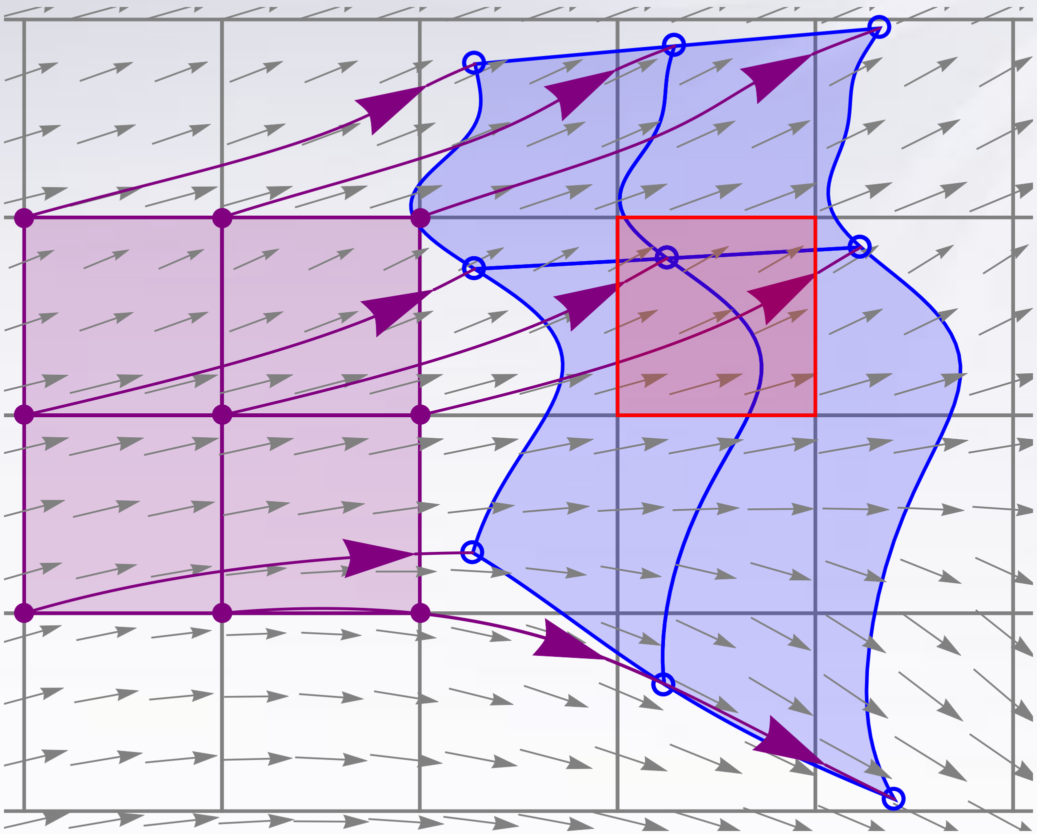

- Lagrangian flow map

$$\boldsymbol{x}^* : \boldsymbol{a} \mapsto \boldsymbol{x}^*(\boldsymbol{a},t_{n+1}),$$
$$\boldsymbol{a}^* : \boldsymbol{x} \mapsto \boldsymbol{a}^*(\boldsymbol{x},t_n)$$

Flux Form Semi-Lagrangian

- Each edge has ‘swept region’ to compute flux
- Flux added to one side, subtracted from other
- Automatic conservation

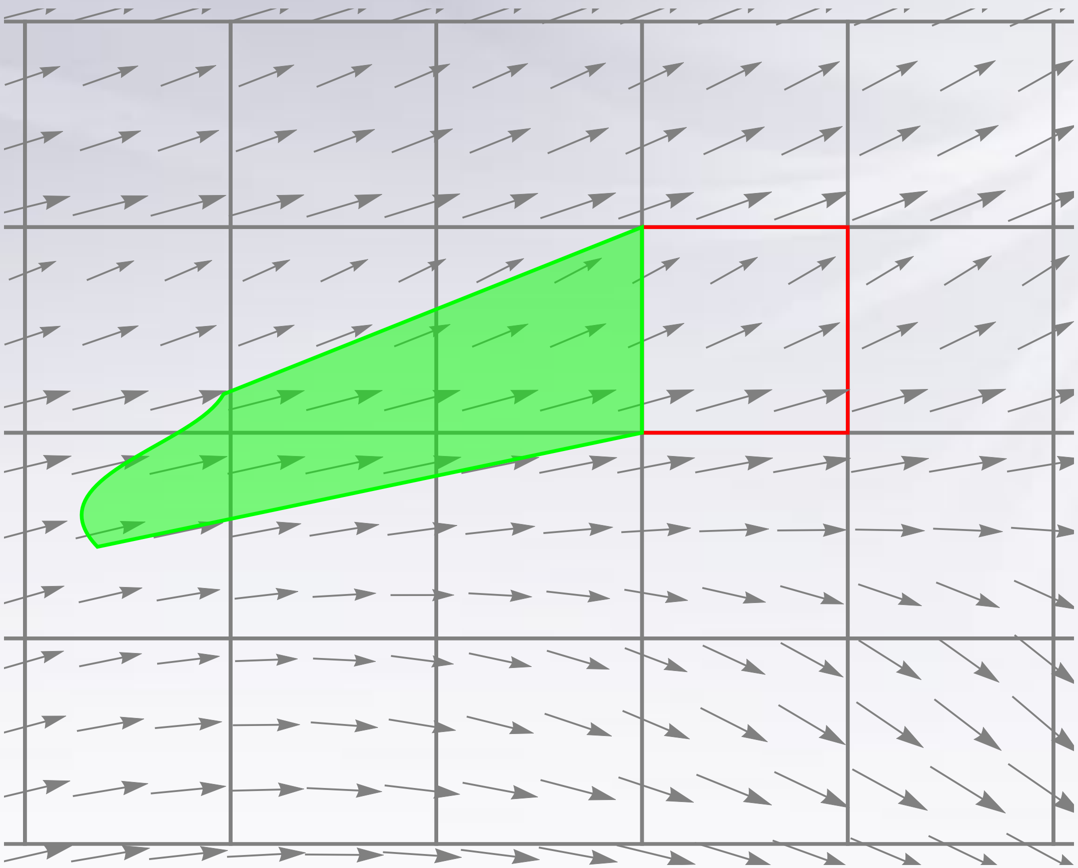

Flux Form Semi-Lagrangian

- Each edge has ‘swept region’ to compute flux
- Flux added to one side, subtracted from other
- Automatic conservation

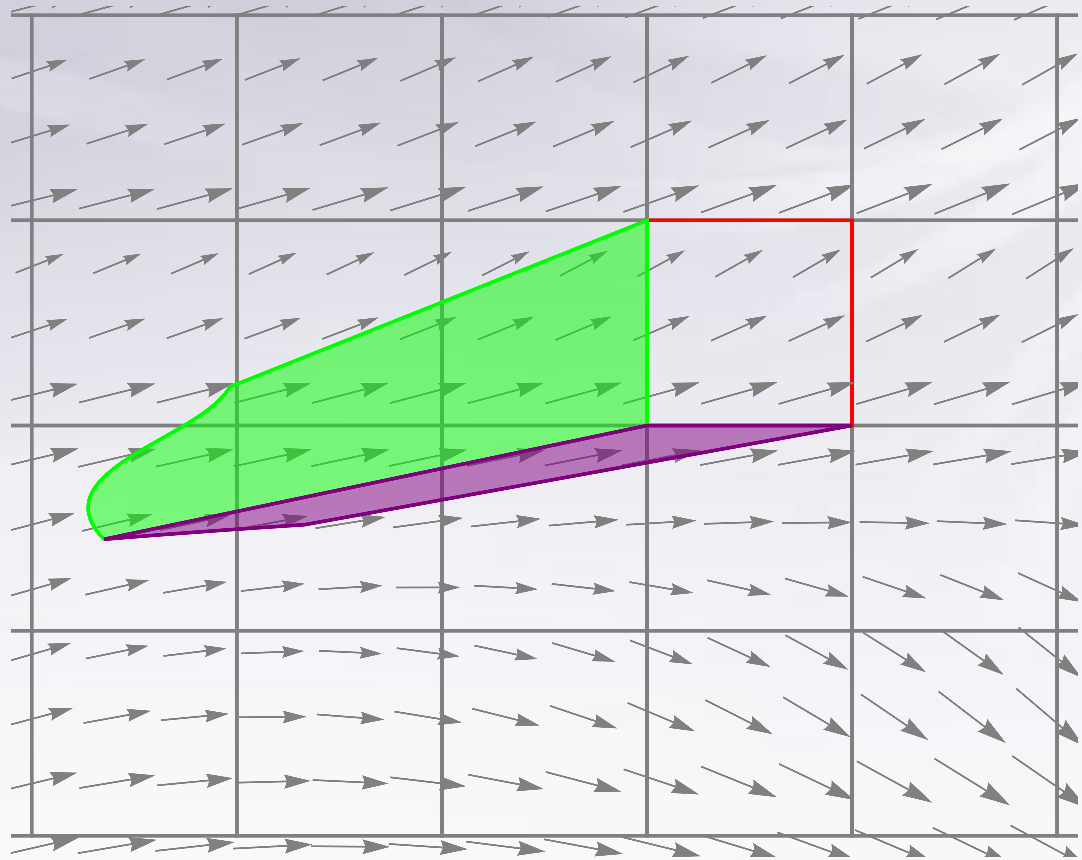

Flux Form Semi-Lagrangian

- Each edge has 'swept region' to compute flux
- Flux added to one side, subtracted from other
- Automatic conservation

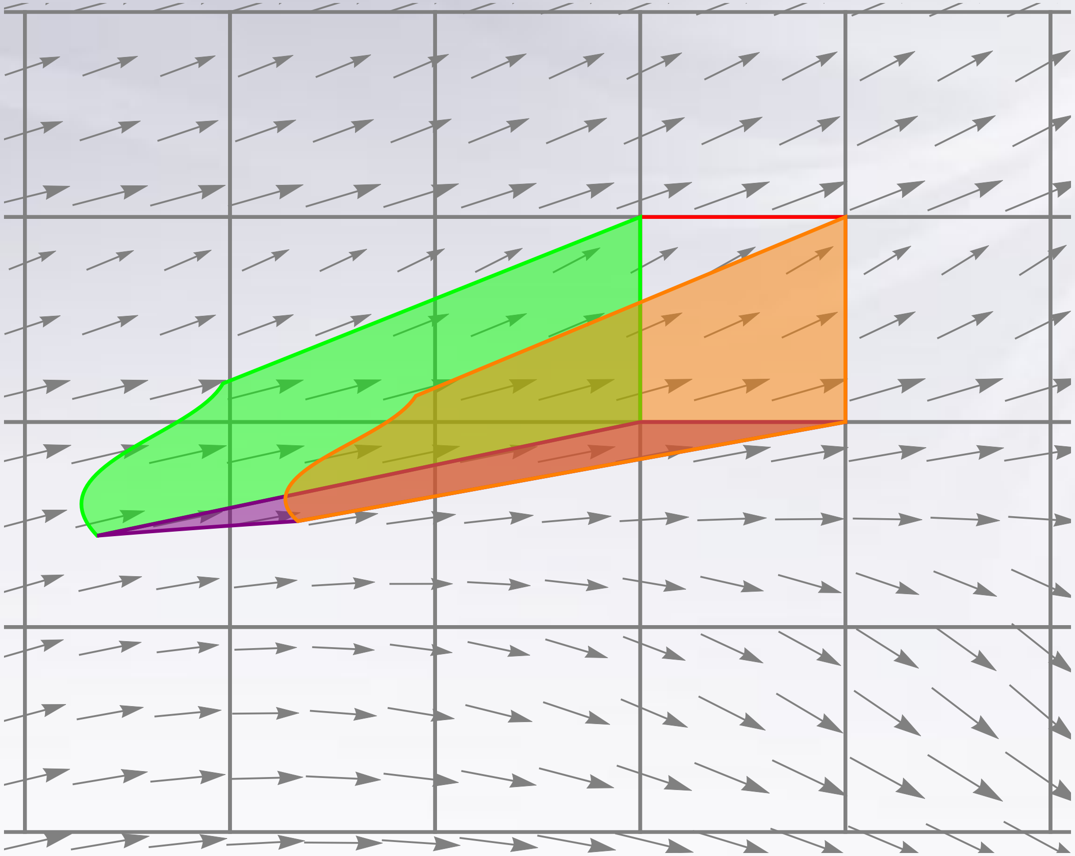

Flux Form Semi-Lagrangian

- Each edge has ‘swept region’ to compute flux
- Flux added to one side, subtracted from other
- Automatic conservation

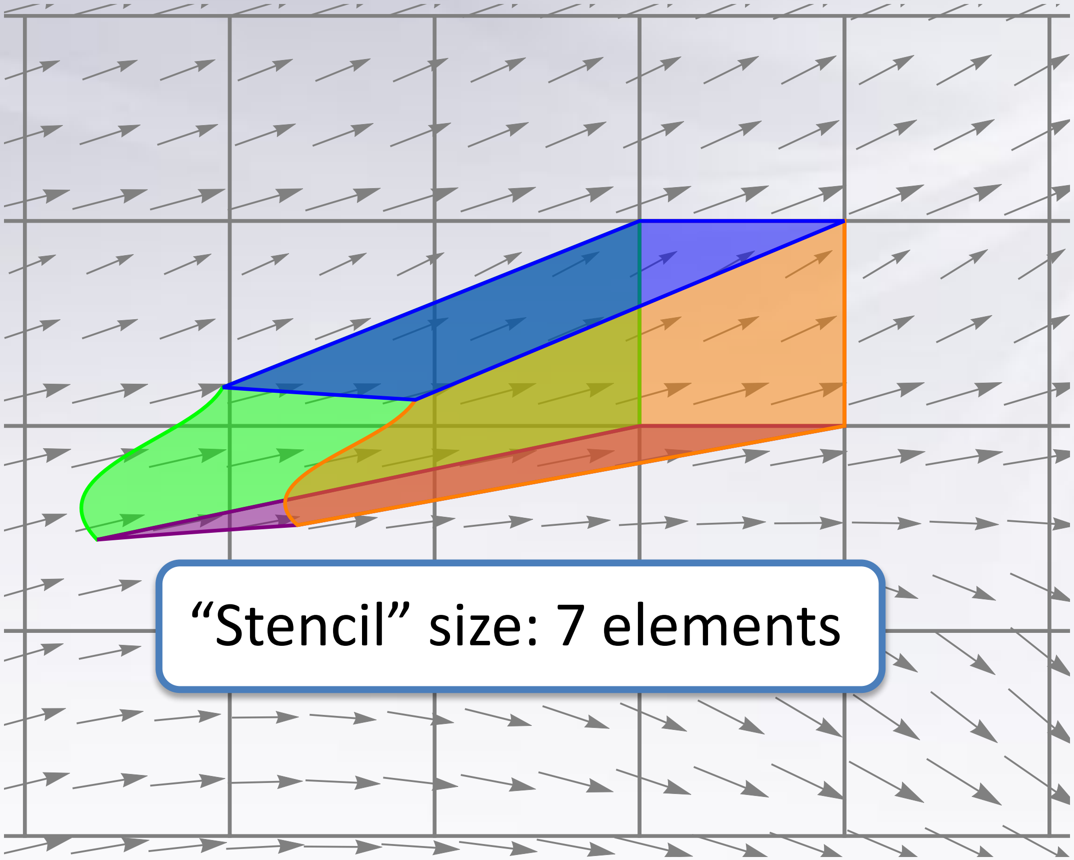

SLMM: CDG/IR

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

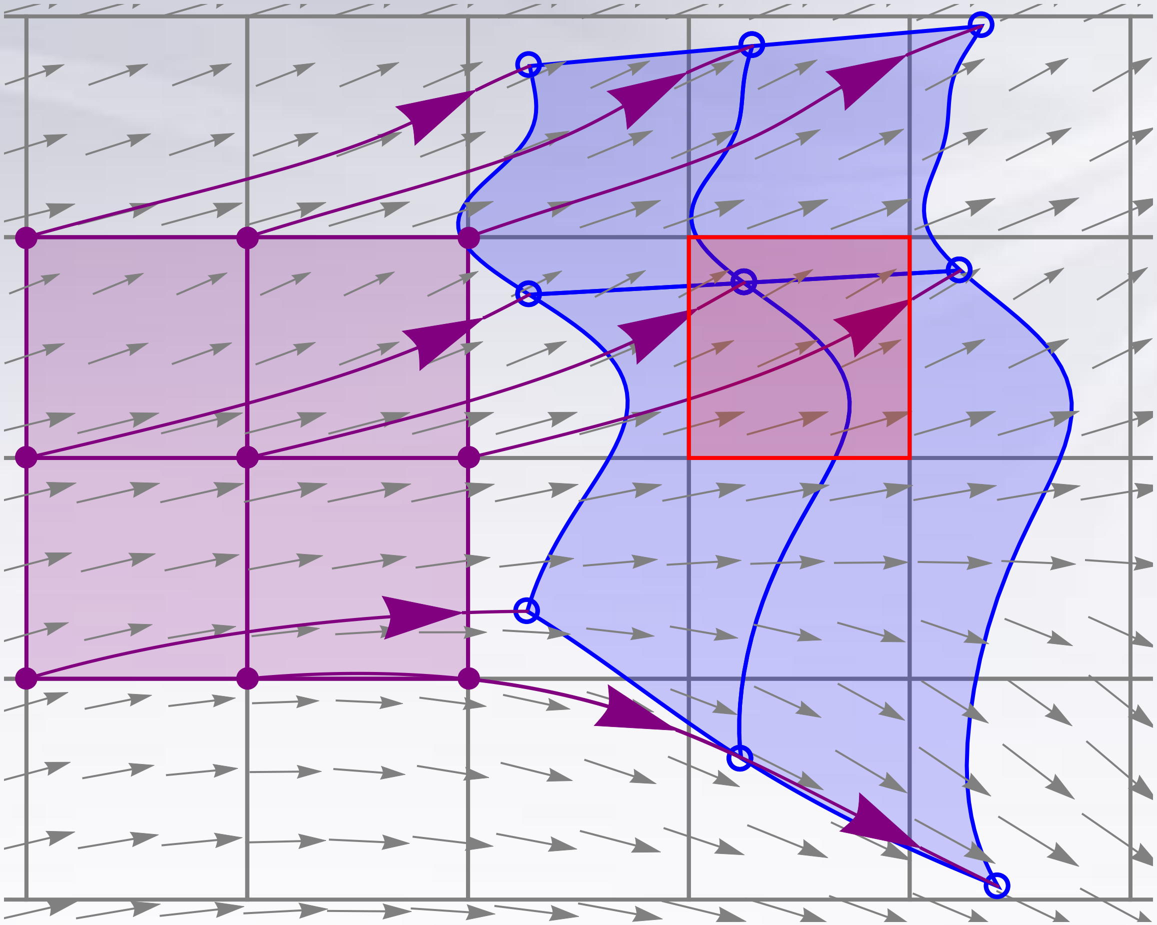

SLMM: CDG/IR

- Elements at t_n are advected forward in time
- Distorted mesh at t_{n+1} provides ‘source’
- Eulerian mesh at t_{n+1} is ’target’

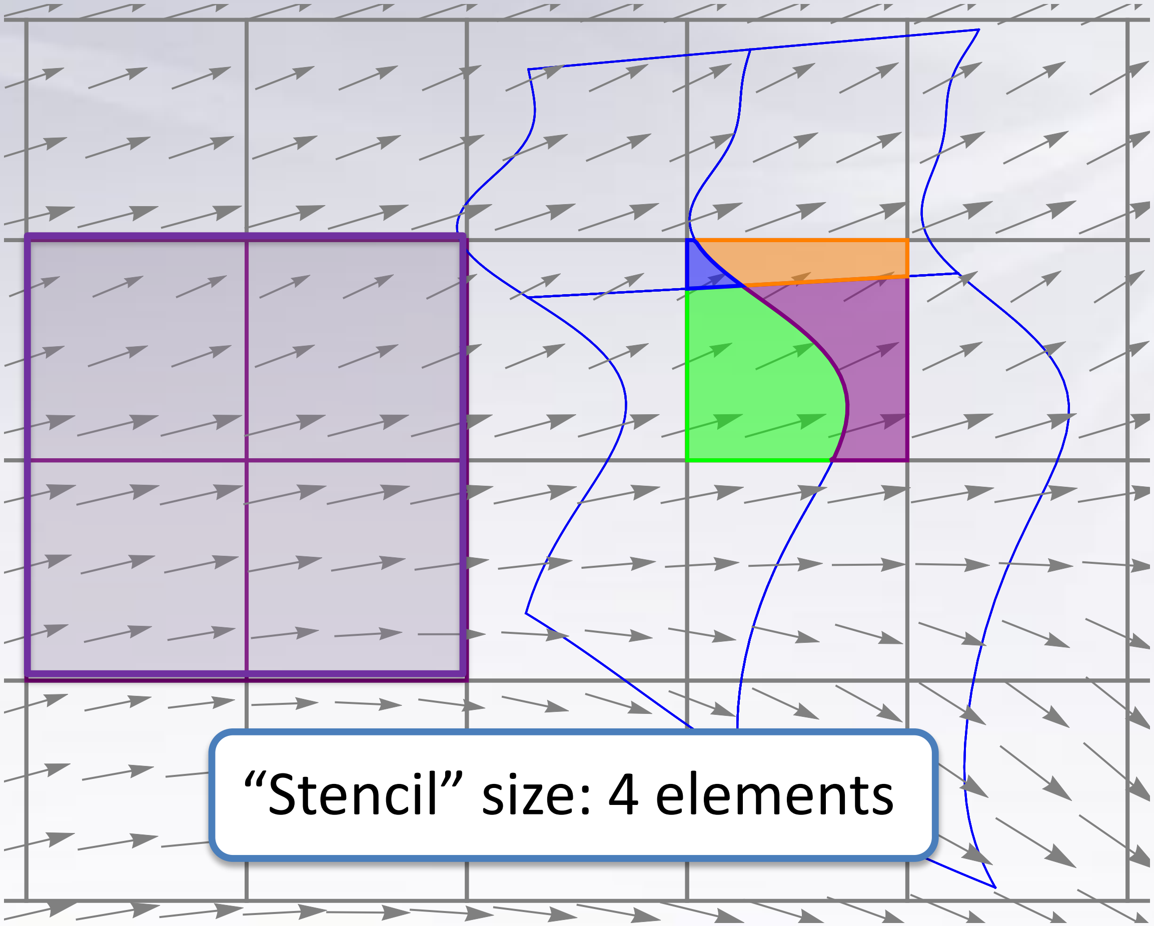

QLT: Communication-efficient Constrained Density Reconstruction (CDR)

- Property preservation
  - **mass conservation**, shape preservation, tracer consistency
- Minimize MPI communication: rounds, volume
- Highlights:
  - Algorithms with deterministic, non-iterative, data-independent performance
    - Roughly 1 MPI_Allreduce
  - Safe: Naturally solve safety problem if primary problem infeasible
    - Safety problem => mass conservation, tracer consistency, don't violate global extrema
  - Practically useful upper bound on mass redistribution
  - Quasi-Local Tree **(QLT)** algorithm redistributes mass for shape preservation and tracer consistency locally*

***Locality with respect to mesh tree, not mesh elements**

QLT: Quasi-Local, Tree-based density reconstruction

Leaves-to-root

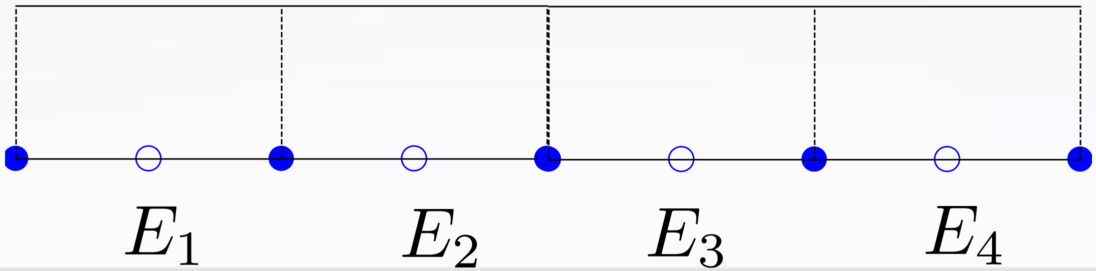

QLT: Quasi-Local, Tree-based density reconstruction

Leaves-to-root

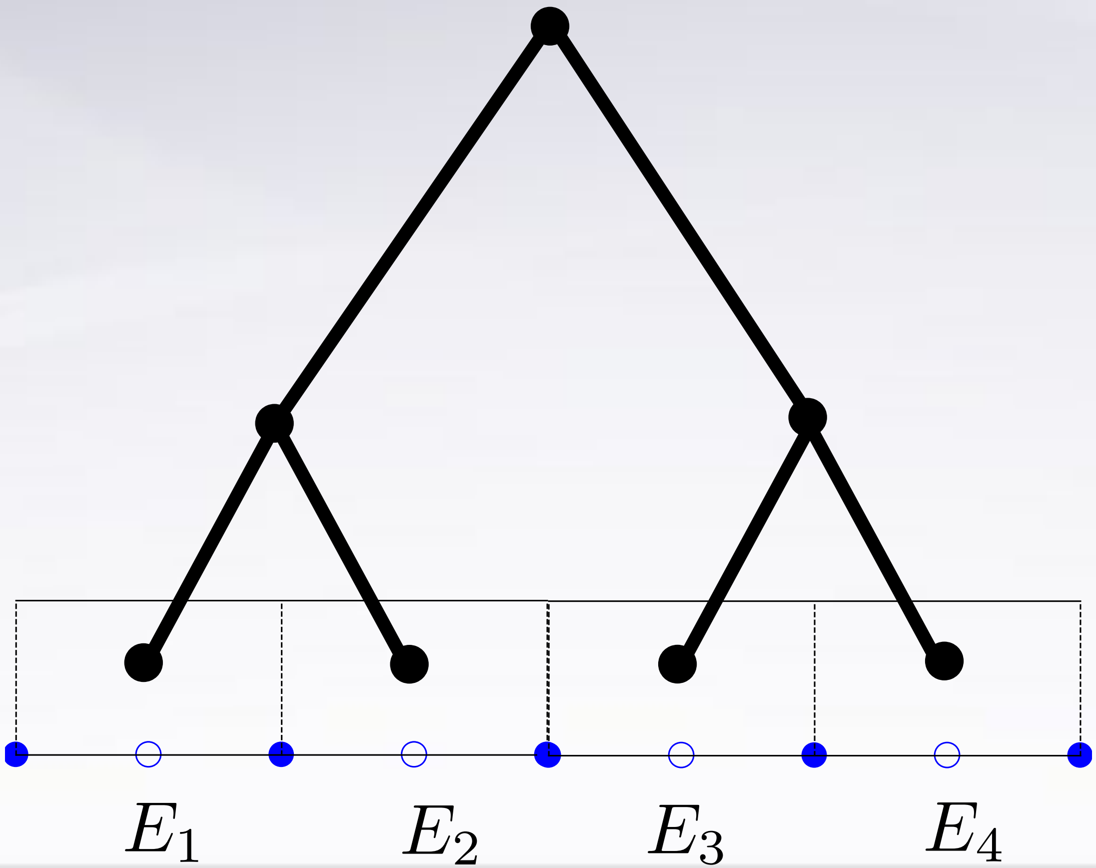

QLT: Quasi-Local, Tree-based density reconstruction

Leaves-to-root

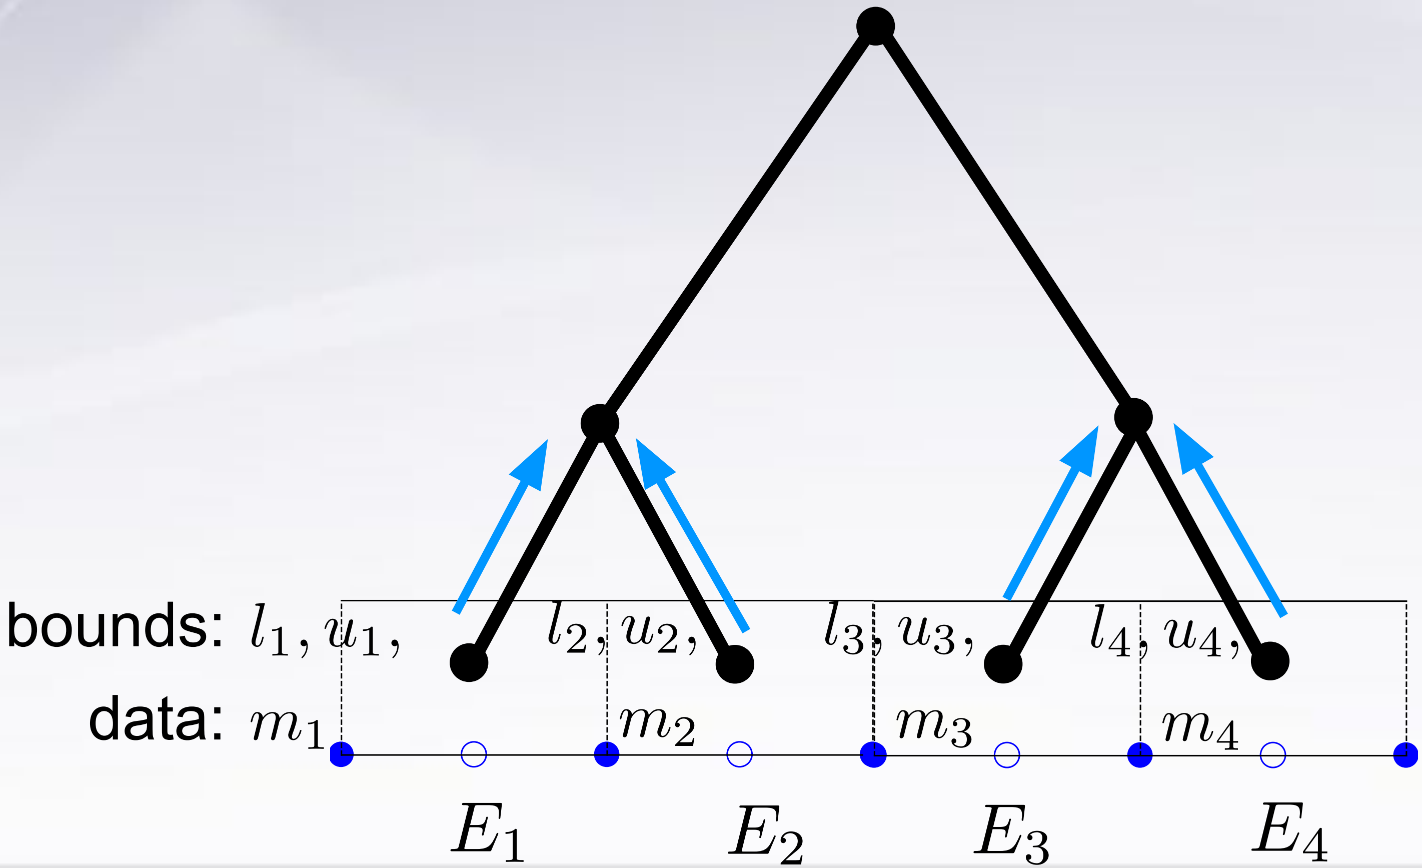

QLT: Quasi-Local, Tree-based density reconstruction

Leaves-to-root

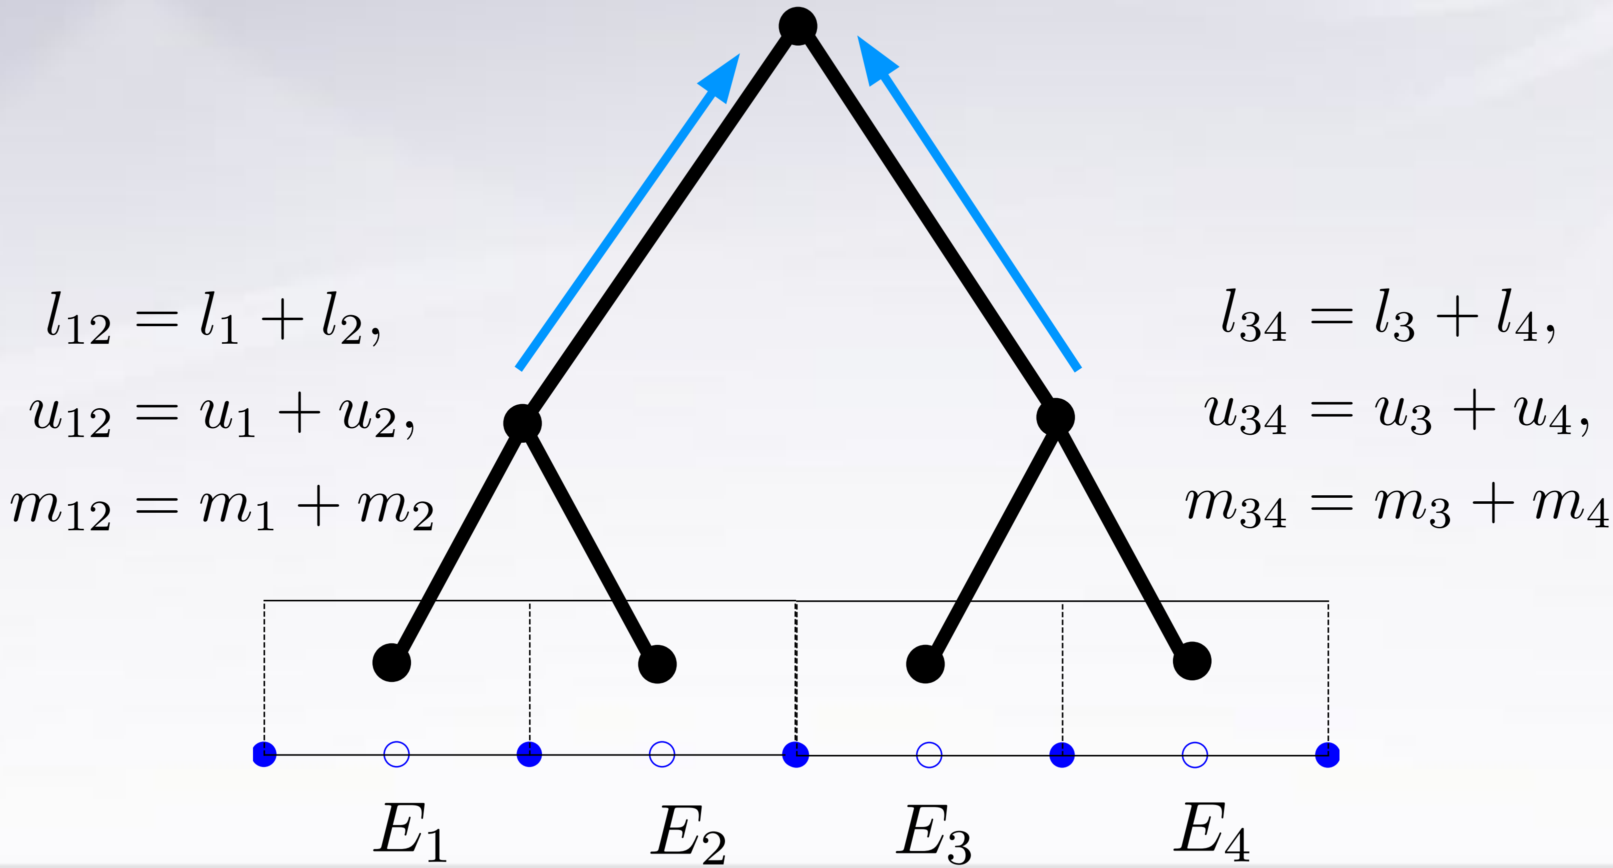

QLT: Quasi-Local, Tree-based density reconstruction

Leaves-to-root

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

NASC: $l_{12} + l_{34} \leq M_g \leq u_{12} + u_{34}$

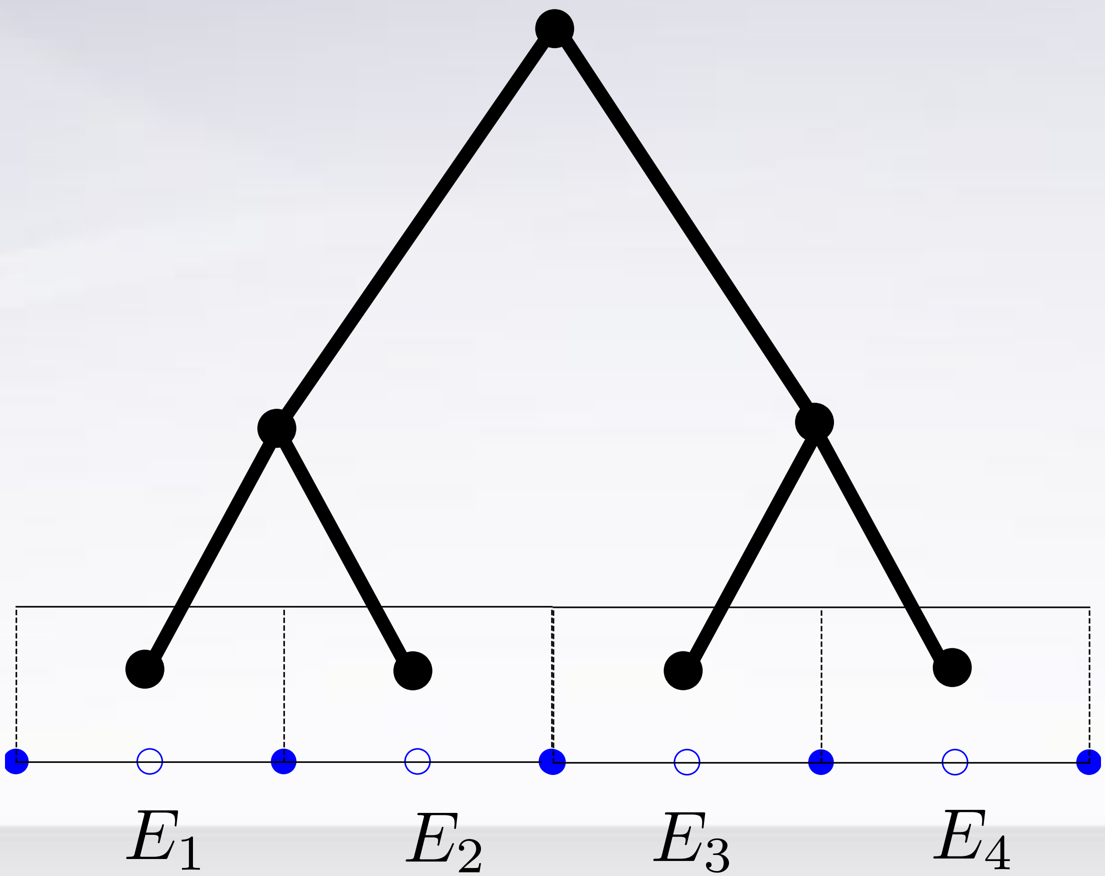

QLT: Quasi-Local, Tree-based density reconstruction

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

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

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

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

Root-to-leaves

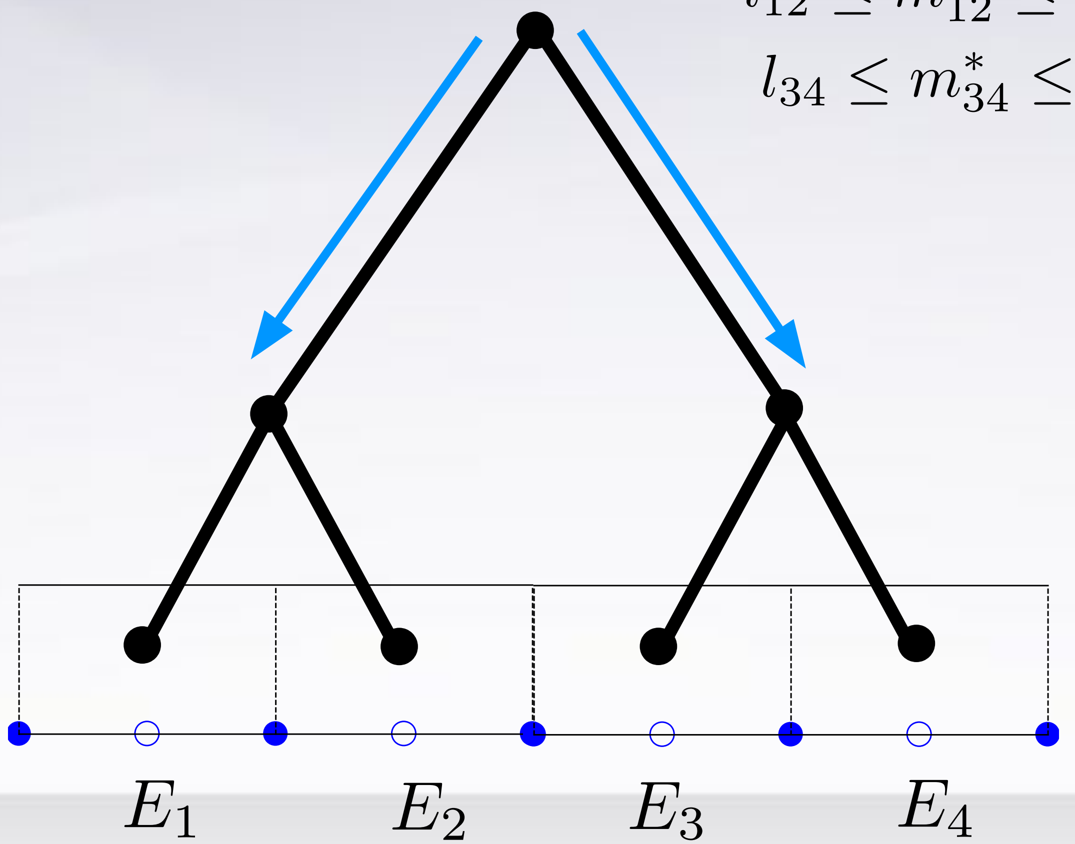

QLT: Quasi-Local, Tree-based density reconstruction

Root-to-leaves

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

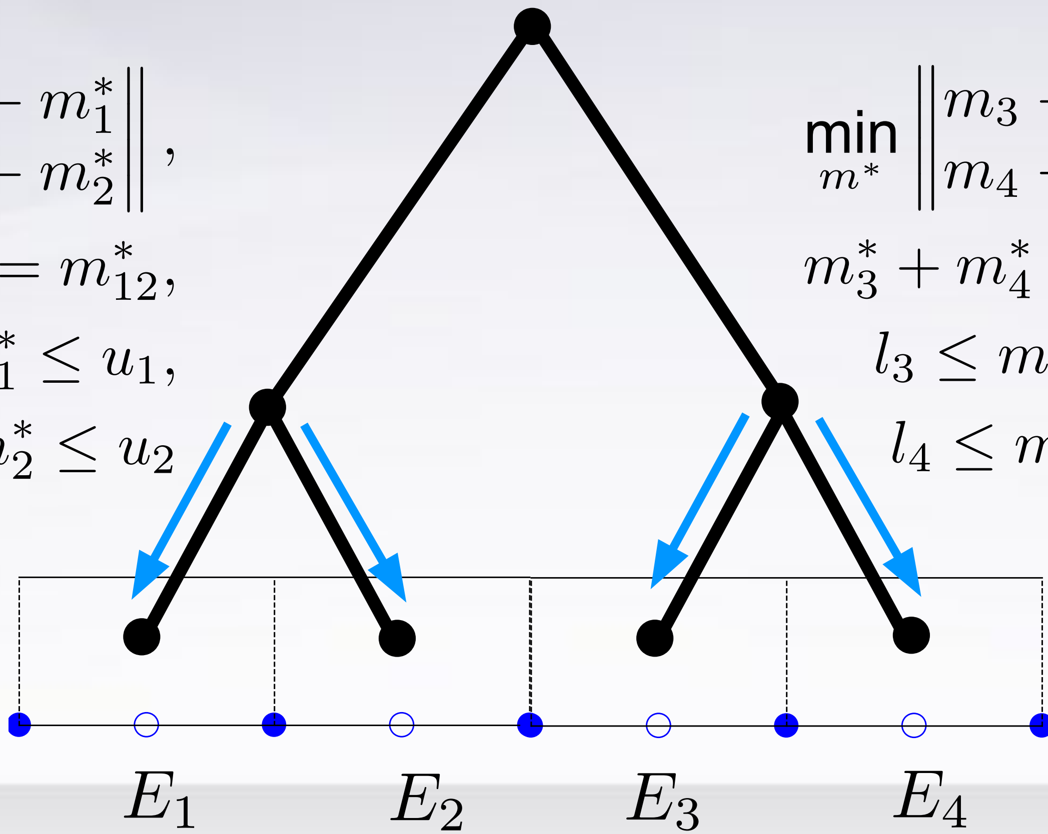

$$\min_{m^*} \left\| \begin{matrix} m_3 - m_3^* \\ m_4 - m_4^* \end{matrix} \right\|, \quad m_3^* + m_4^* = m_{34}^*, \quad l_3 \leq m_3^* \leq u_3, \quad l_4 \leq m_4^* \leq u_4$$

QLT: Quasi-Local, Tree-based density reconstruction

Root-to-leaves

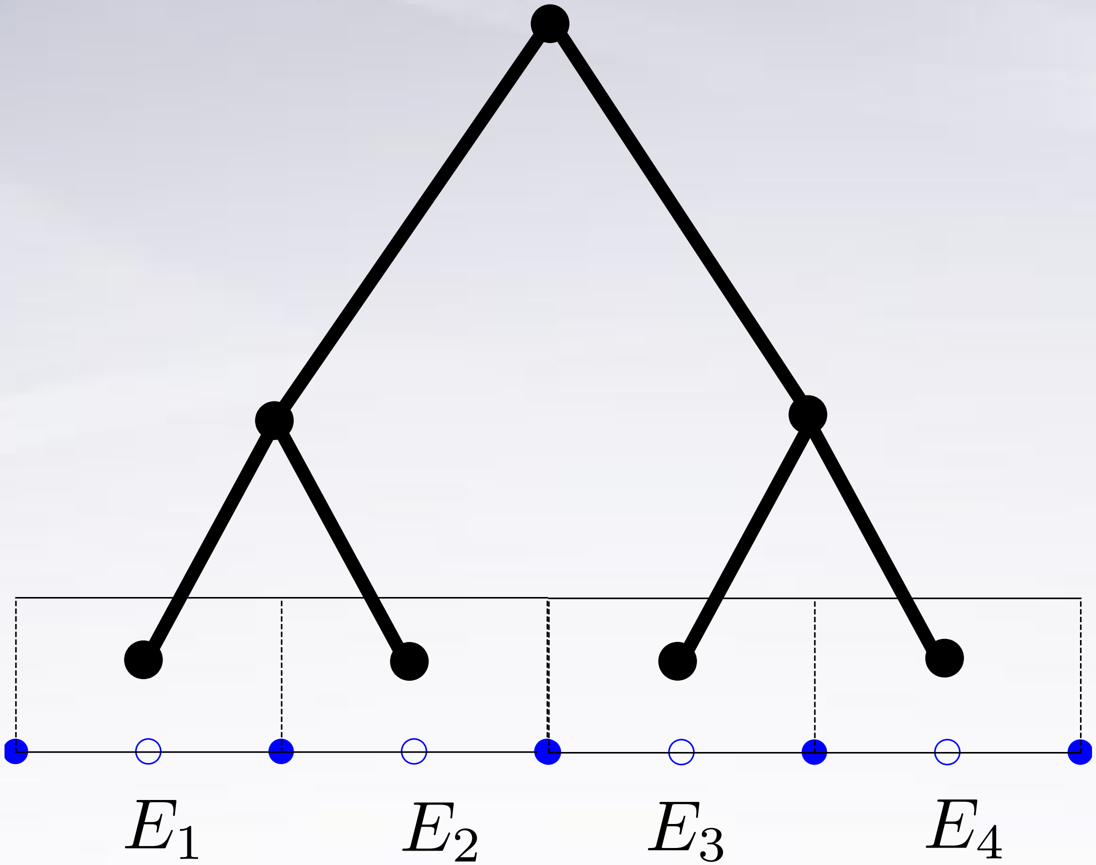

Final: element-local shape-preservation

Partnerships: E3SM and SciDAC

- **E3SM Master:** New limiter, new time step coupling & convergence test (see poster)
- **E3SM Next-Gen. Development:** Semi-Lagrangian MPI communication patterns
- **FastMath:** Multigrid solvers for semi-implicit dynamics (Helmholtz equation)
- **RAPIDS:** Compact, high-order elements for spatial discretizations

Halo-1 Communication patterns

- Cubic element illustrations
- Trajectories computed locally
  - McGregor, *MWR*, 1993
  - **Restricts time step**
- Data transfer
  - Full halo
  - Upwind
- Full halo exchange
  - Deterministic, blue receives data from red
  - Simple: send all
    - 8 x 16 = **128** nodes (columns)
  - Optimal: send unique
    - 10 x 10 - 16 = **84** columns

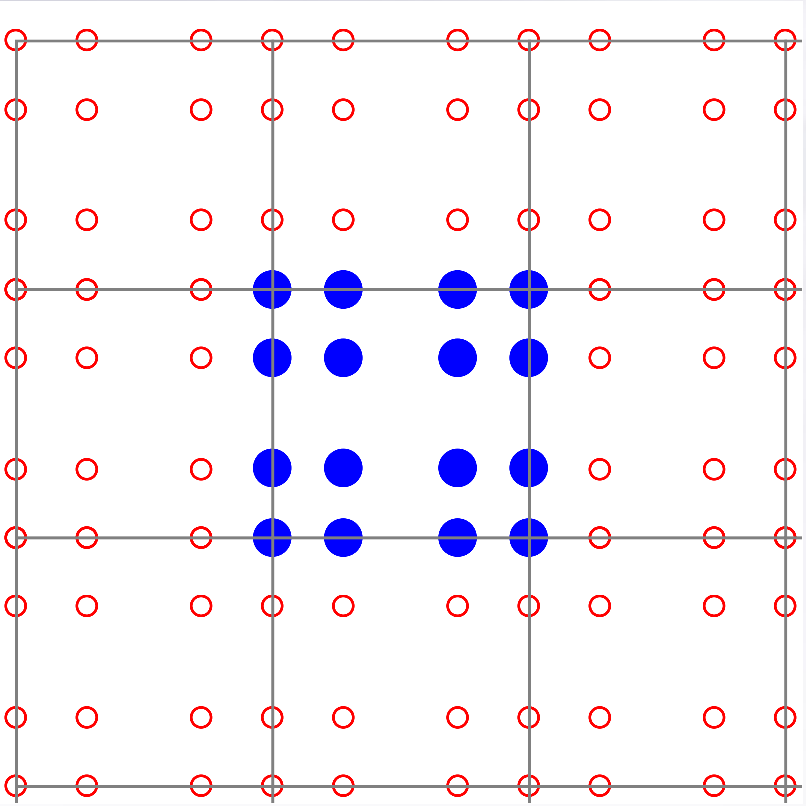

Halo-1 Communication patterns

- "Upwind," cell-integrated
- Step 1: Handshake with halo
  - Determine source (red) elements
  - Asynchronous, negligible cost
- Step 2:
  - Blue receives data from red
  - All red columns required for basis-basis interactions
  - Illustrated (avg. case)
    - Simple: **48** columns
    - Optimal: **33** columns

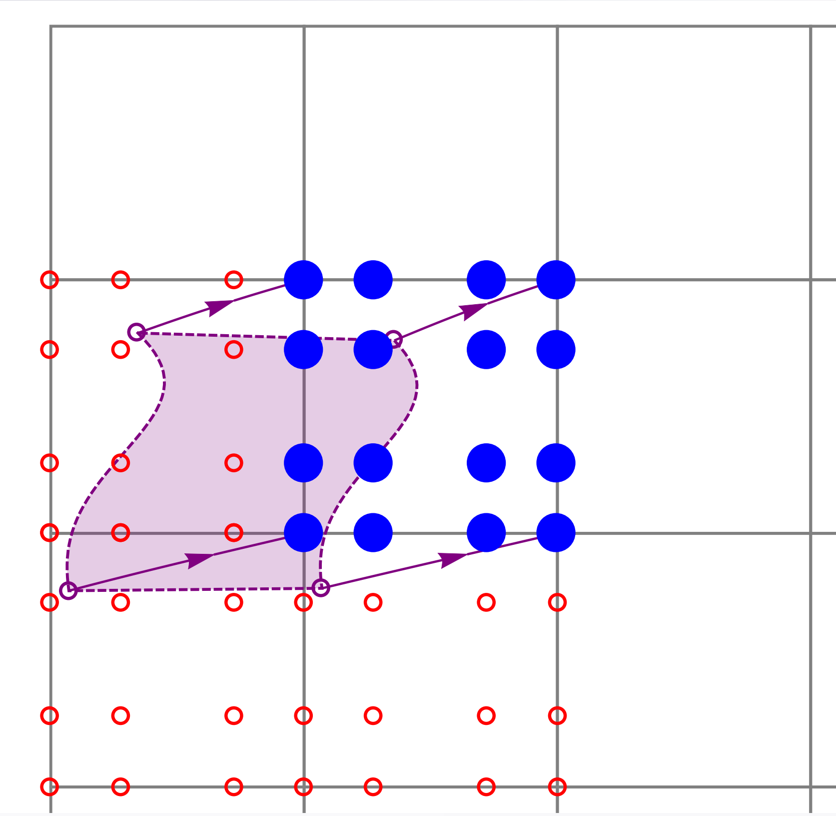

Halo-1 Communication patterns

- "Upwind," pointwise
- Step 1: Handshake with halo
  - Determine source (red) elements
  - Asynchronous, negligible cost
- Step 2:
  - Blue receives data from red
  - Only elem. min/max + dep. points for point-basis interactions
  - Illustrated:
    - 3 x 2 + 12 = **18** columns
  - Upper bound
    - 8 x 2 + 16 = **32** columns

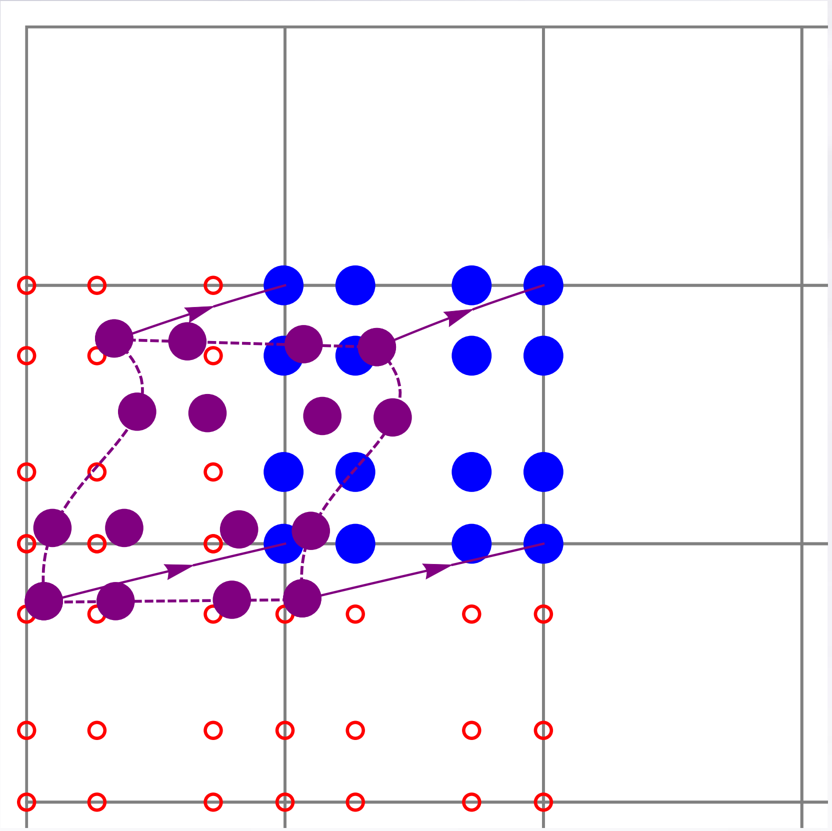

Summary

- E3SM: Science and performance goals for fully coupled model
- Transport: Simple + expensive = opportunity
- QLT: Data-independent property preservation with 1 all-reduce per time step
  - Locally cons. cell-integrated methods
  - Classical (interpolation) SL
- Speedups from transport, E3SM v1
  - SE to pointwise SL
  - Halo exchange to upwind MPI
- Ongoing:
  - SL Dynamics
  - E3SM optimization & development
  - Science investigations
    - "Physically correct" error