Algorithms and Software for Fast E3SM Atmosphere Tracer Transport

P. A. Bosler, A. M. Bradley, O. Guba, M. A. Taylor

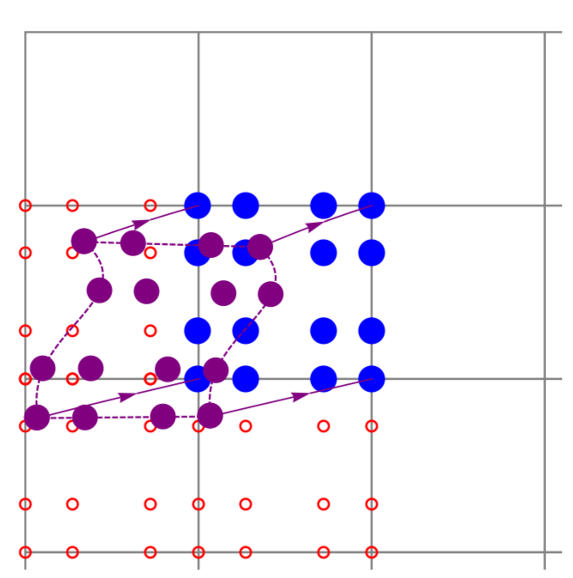

Figure 1: Schematic for semi-Lagrangian tracer transport. The target element with solid blue GLL points is advected backward in time to the deformed element having the solid purple points. Red open dots and blue dots represent GLL points of source elements. In remap-form cell-integrated semi-Lagrangian (CISL) transport, the departure element is integrated against the basis functions of the source elements. In remap-form interpolation semi-Lagrangian (ISL) transport, each source element interpolates to the departure points contained in that element. Both CISL and ISL can take long time steps, potentially providing speedup over Eulerian methods. ISL processes approximately 4x smaller data volume than CISL, providing still more speedup. Communication-efficient density reconstructors (CEDR) enable remap-form SL transport and motivate using the particularly efficient ISL method.

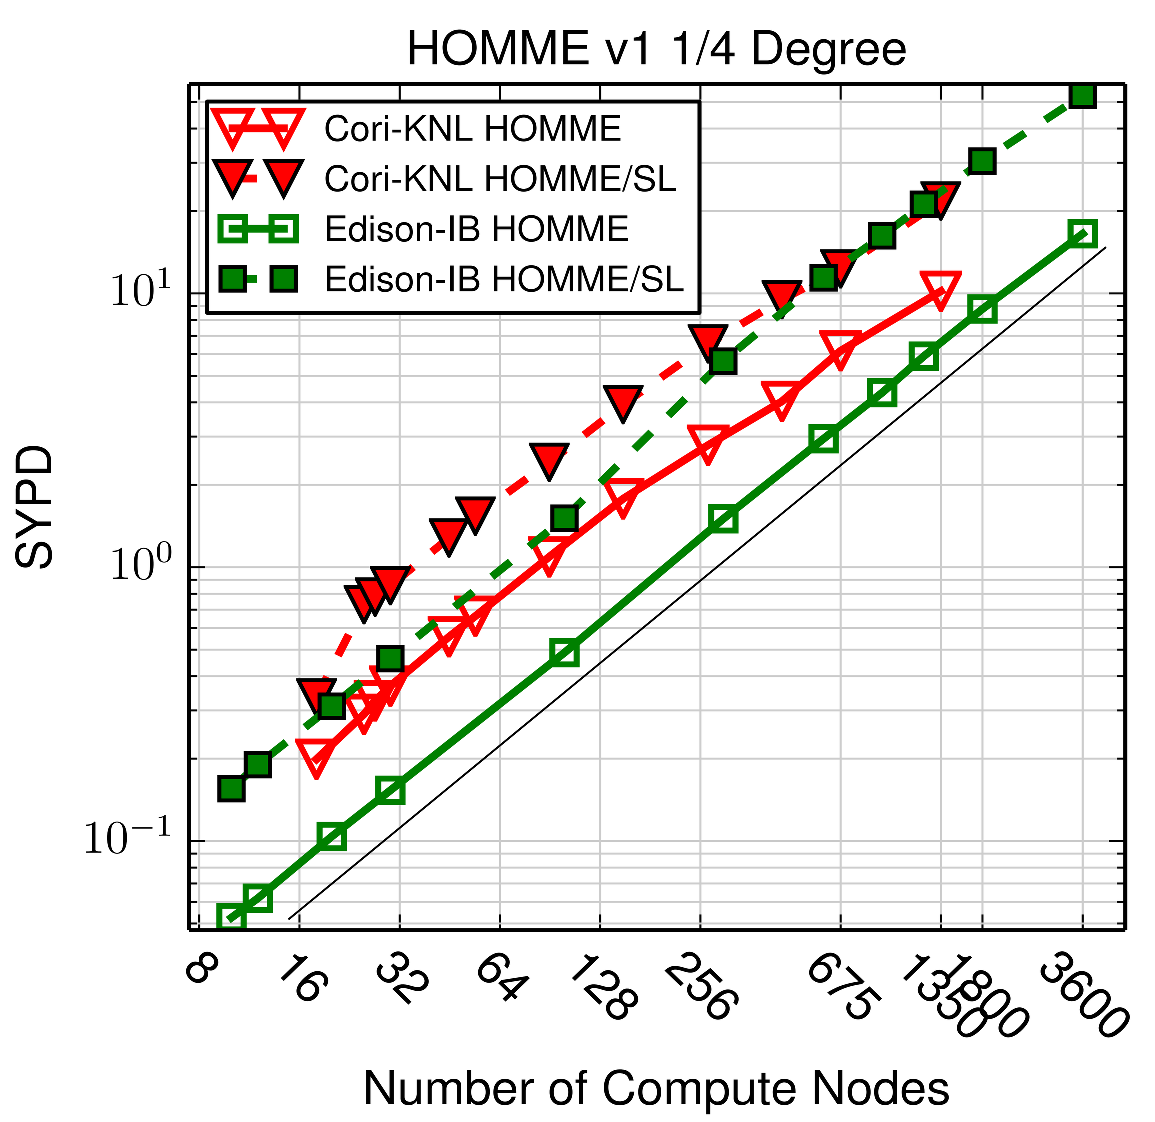

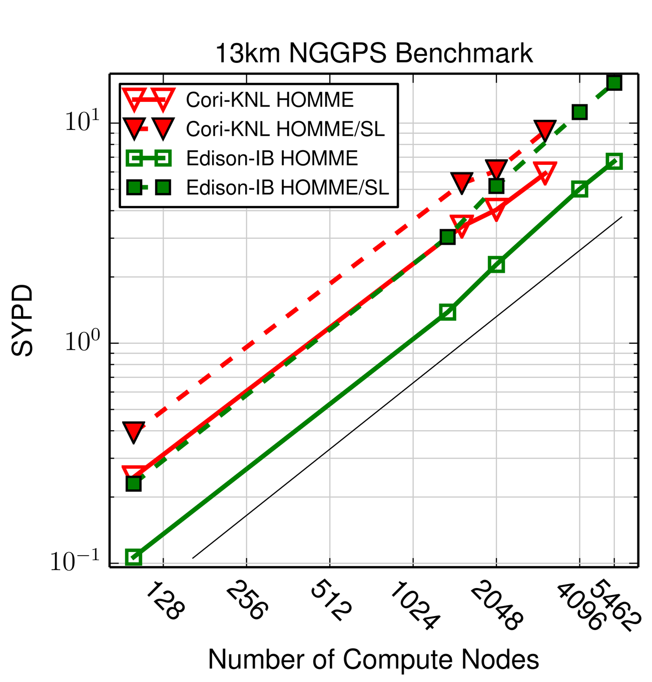

Figure 2 (Above) : Performance of the EAM spectral element dynamical core HOMME on the EAM version1 high-resolution configuration (1/4-degree, 40 tracers) with Eulerian transport (solid lines) and SL transport (dashed lines). y-axis is simulated years per day (SYPD); x-axis is number of compute nodes. Machines are Edison (green), which has dual-socket Ivy Bridge nodes, 12 cores/socket; and Cori-KNL (red), which has 68-core (64 used) Knights Landing nodes. Black line indicates perfect scaling. SL transport gives speedup uniformly at high workload per node (few compute nodes) and, more challenging, at low workload per node (many compute nodes). Speedup is over 2x on Cori-KNL and over 3x on Edison.

Figure 3 (Left): Performance of the EAM spectral element dynamical core HOMME on the NGGPS 13km benchmark problem, which has only 10 tracers.

Figure 4: Performance of tracer transport as a function of number of tracers on Haswell and KNL architectures. The SL transport time step is 6 times longer than the Eulerian time step.

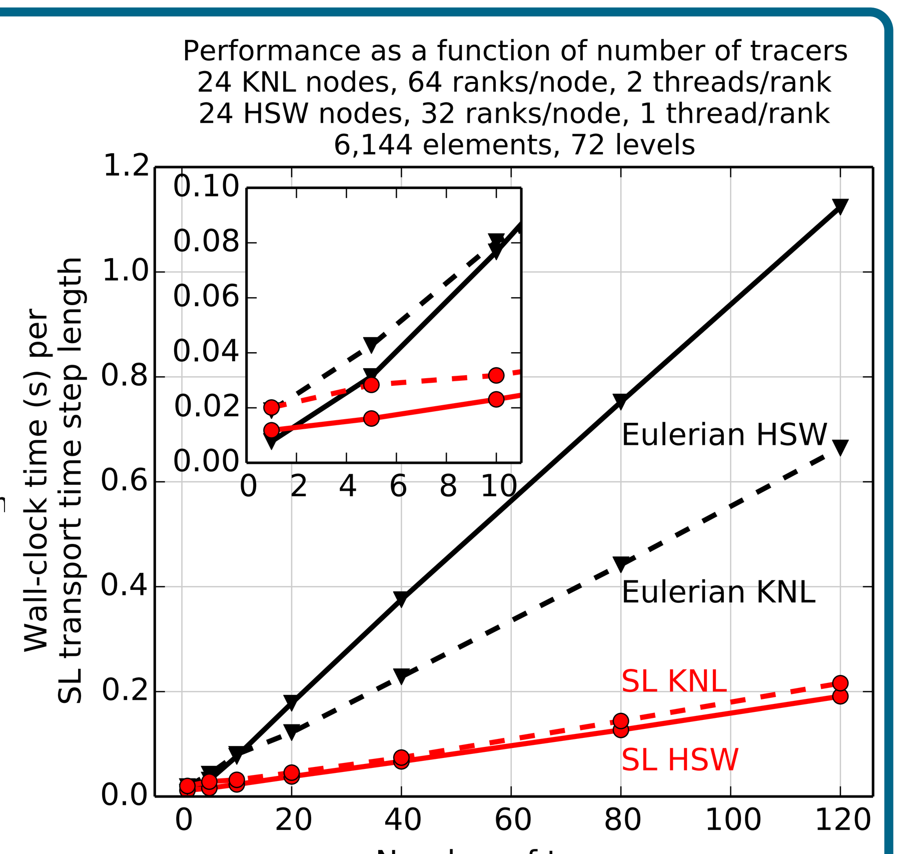

Impact

Semi-Lagrangian (SL) transport for E3SM Atmosphere Model

- SL algorithms speed up dycore by **2-3.5x,** on desktops and pre-exascale LCF machines
- Already integrated into **E3SM master** for version 2

Goal

Substantially increase tracer transport efficiency to enable faster simulations and/or many more tracers.

Approach

1. Develop semi-Lagrangian (SL) transport for E3SM Atm. model, HOMME (Fig. 1).
   - Key challenge of running well on modern and future supercomputers: minimizing communication rounds.
   - SL transport can take long time steps or, in other words, simulate more time per communication round, addressing above challenge.
   - 2 varieties: Cell-Integrated SL (CISL), Interpolation SL (ISL). ISL is most efficient.
   - CISL : P. A. Bosler and A. M. Bradley and M. A. Taylor, Conservative multi-moment transport along characteristics for discontinuous Galerkin methods, accepted to *SIAM J. Sci. Comput.*, 2019.
   - ISL : A. M. Bradley, O. Guba, P. A. Bosler, M. A. Taylor, Interpolation semi-Lagrangian tracer transport on spectral elements, in preparation.
2. Preserve properties using Communication-Efficient Density Reconstructors (CEDR).
   - Achieve transport-dynamics mass consistency, mass conservation, and shape preservation in one global communication round.
   - A good CEDR enables using ISL.
   - A. M. Bradley, P. A. Bosler, O. Guba, M. A. Taylor, G. A. Barnett, Communication-efficient property preservation in tracer transport, *SIAM J. Sci. Comput.*, 41(3), 2019, doi:10.1137/18M1165414.
   - Software release: github.com/E3SM-Project/COMPOSE

Results

1. The EAM dynamical core, or "dycore," with 40 tracers (v1 std.) is 2-3.5x faster, depending on problem and architecture (Figs. 2—5).
2. SL tracer transport has greater multi-tracer efficiency (Fig. 4) and is at least as accurate as E3SM version 1 with Eulerian transport (Figs. 6—8).
3. Key climate statistics, without changing v1 tuning for the new transport method, look good (Fig. 9).
4. CEDR & ISL already integrated to HOMME for E3SM v2

Future work

1. Apply this SL framework to MPAS-Ocean to speed up biogeochemical (BGC) simulations.
2. Using high-order formulas from our new ISL method, implement a capability to compute extremely high resolution tracers for fixed dynamics cost (Fig. 7).
3. Increase SL trajectory accuracy for yet more tracer accuracy.
4. For additional speedup, explore fully 3D SL transport or SL transport decoupled from the dycore's vertical remap.

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. SAND2018-12624 PE.

Sandia National Laboratories

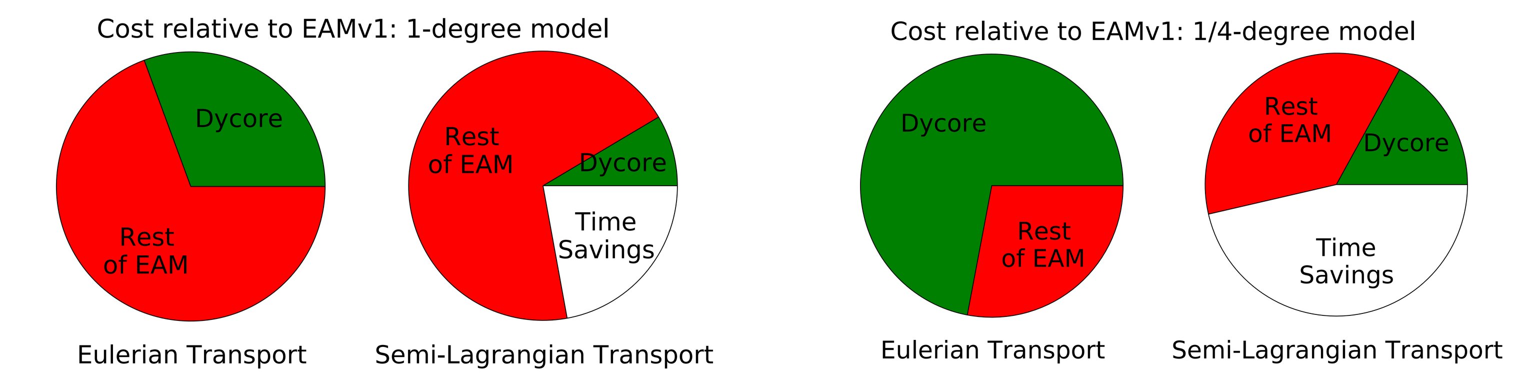

Figure 5: Performance of E3SM Atmosphere Model in version 1 low- **(left)** and high- **(right)** resolution configurations on 75 dual-socket Intel Xeon Broadwell nodes, 18/cores socket, of Anvil cluster, with and without SL transport. SL transport makes the dycore much faster. "Rest of EAM" accounts for physics parameterizations and dynamics-physics coupling. The high-resolution configuration has a relatively lower "Rest of EAM" cost; thus, overall EAM speedup is larger than in the low-resolution configuration.

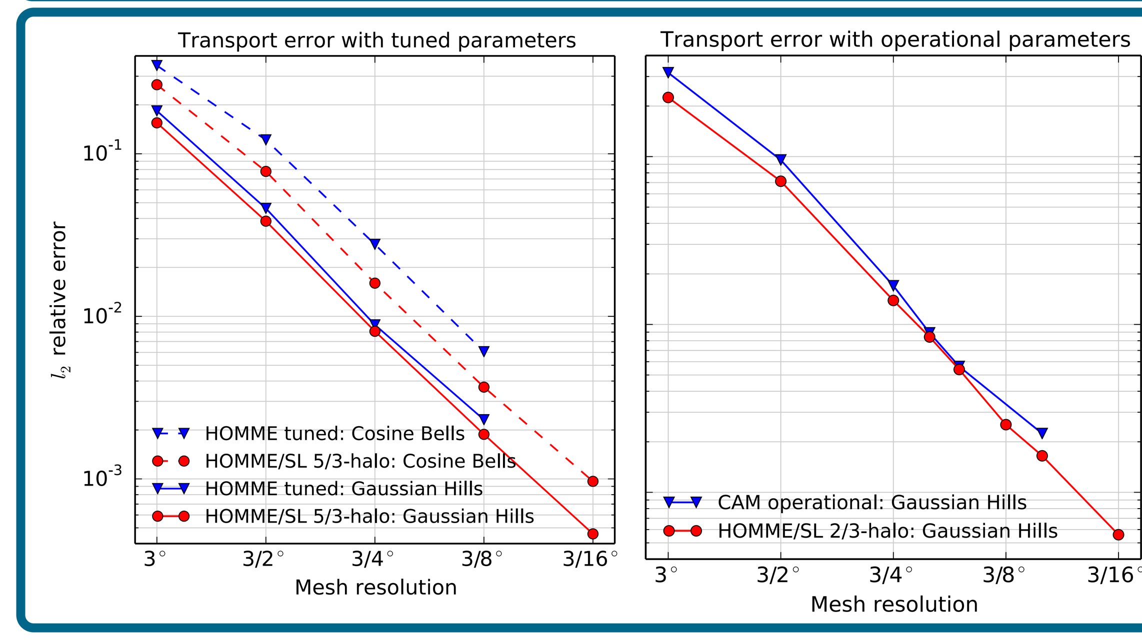

Figure 6: Benchmark accuracy test: L2 relative error as a function of mesh resolution. Tracer fields (Gaussian Hills, solid; Cosine Bells; dashed) in deformational flow. Blue is HOMME with Eulerian transport; red is HOMME with SL transport. Left plot is for parameters tuned for highest accuracy; right plot is for operational parameters. SL transport is uniformly at least as accurate as Eulerian.

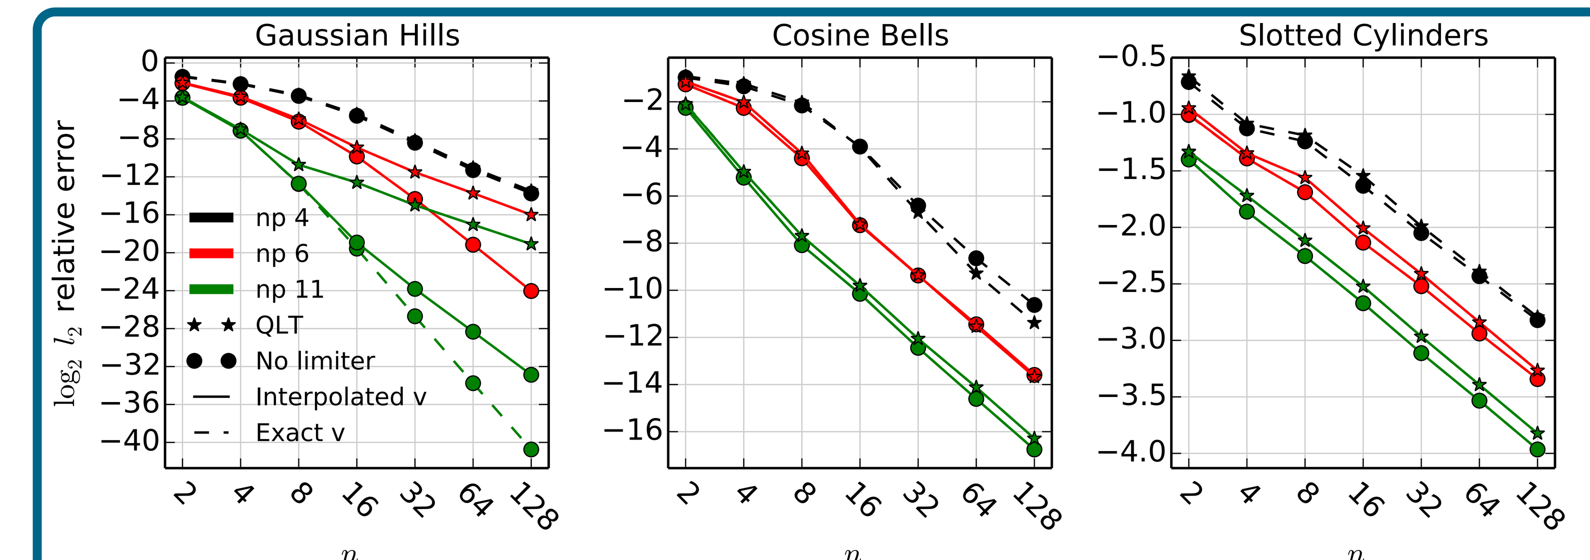

Figure 7: p-refined tracer transport increases the resolution of the tracers for fixed dynamics grid and cost; it is efficient and scalable, with no time step penalty. It uses higher-order formulas from our stabilized interpolation semi-Lagrangian method. Plotted are error (y-axis, note log2) vs. cubed-sphere mesh parameter ne (number of elements is 6*ne*ne) for the benchmark divergent flow problem. Velocity is computed on the dynamics grid with GLL np=4. SL departure points are interpolated to the p-refined grid, and tracers are computed on this grid. For the same dynamics cost (grid and time step), p-refinement can increase tracer accuracy by over 10x.

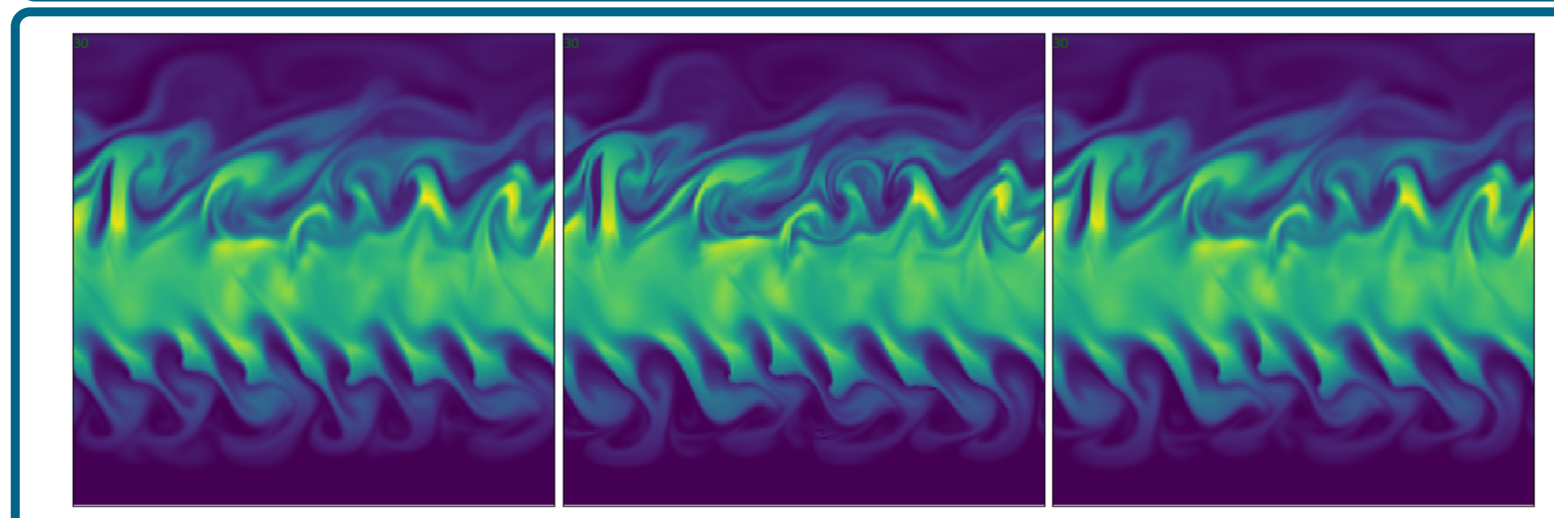

Figure 8: Specific humidity at approximately 500 hPa, day 30 in DCMIP 2016 moist baroclinic instability test. Left, the tracers are advected with Eulerian transport, which requires hyperviscosity; middle, with SL transport and no hyperviscosity; right, SL transport with hyperviscosity. Large-scale structure is very similar, but SL transport provides higher resolution when hyperviscosity is not used.

Figure 9: Relative differences in climatology statistics between EAM with SL transport and EAMv1. Red dots: relative difference in spatio-temporal means. Blue triangles: relative difference in root mean square error (RMSE). Values are made relative by dividing by reference model's maximum minus minimum global values. Temporal mean is either annual (default) or northern hemisphere winter (indicated by DJF). TREFHT is near-surface land and ocean temperature. PSL is pressure at surface. T is temp, U zonal wind. RESTOM is top of model radiative flux. NETCF, LWCF, SWCF are net, long wave, and short wave cloud forcing.