SAND2020-5825C

SCREAM: a performance-portable, global cloud-resolving model based on the Energy Exascale Earth System Model

Presenting author: Benjamin R. Hillman *(Sandia National Laboratories)*

Contributors: Peter Caldwell, Andy Salinger, Luca Bertagna, Hassan Beydoun, Peter Bogenschutz, Andrew Bradley, Aaron Donahue, Chris Eldred, Jim Foucar, Chris Golaz, Oksana Guba, Ben Hillman, Rob Jacob, Jeff Johnson, Noel Keen, Jayesh Krishna, Wuyin Lin, Weiran Liu, Kyle Pressel, Balwinder Singh, Andrew Steyer, Mark Taylor, Chris Terai, Paul Ullrich, Danqing Wu, Xingqiu Yuan

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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.

Energy Exascale Earth System Model (E3SM)

- Global earth system modeling framework
- Collaborative project involving 8 U.S. DOE labs and 12 universities
- Development driven by energy and water issues over the next 40 years
- Key computational goal: performance on exascale computers
- Open source / open development
  - Website: www.e3sm.org
  - Github: https://github.com/E3SM-Project
  - DOE Science youtube channel: https://www.youtube.com/channel/UC_rhpi0lBeD1U-6nD2zvlB

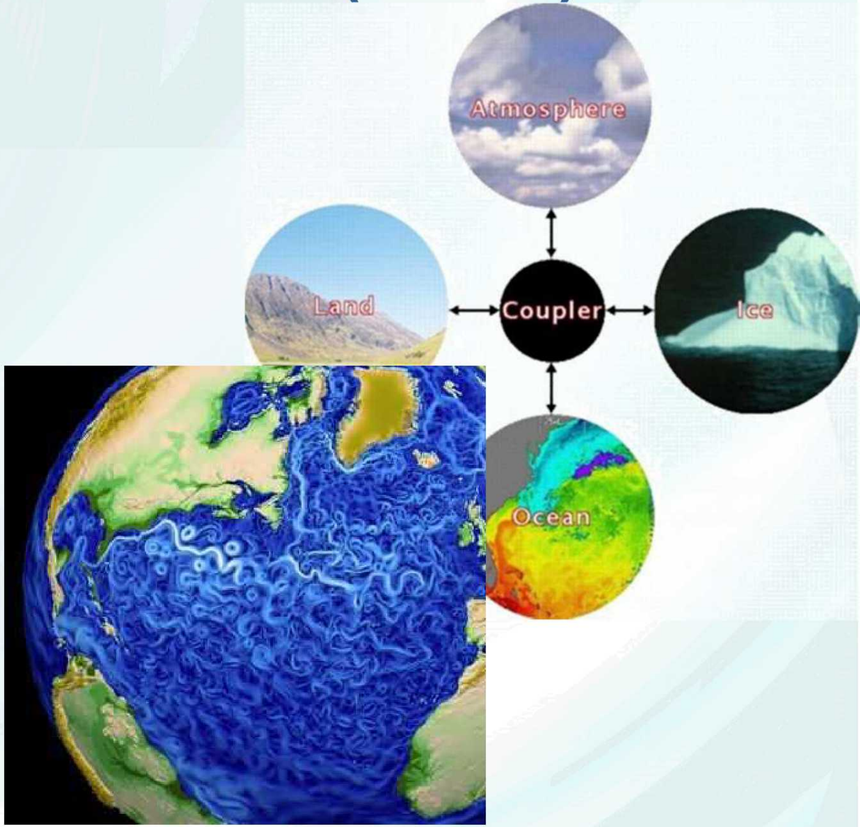

Context: the problem with coarse resolution

- Coarse resolution results in heavy reliance on complicated/empirical subgrid-scale parameterizations
- Uncertainty in subgrid-scale parameterizations can be a major source of uncertainty in large-scale models / climate projections

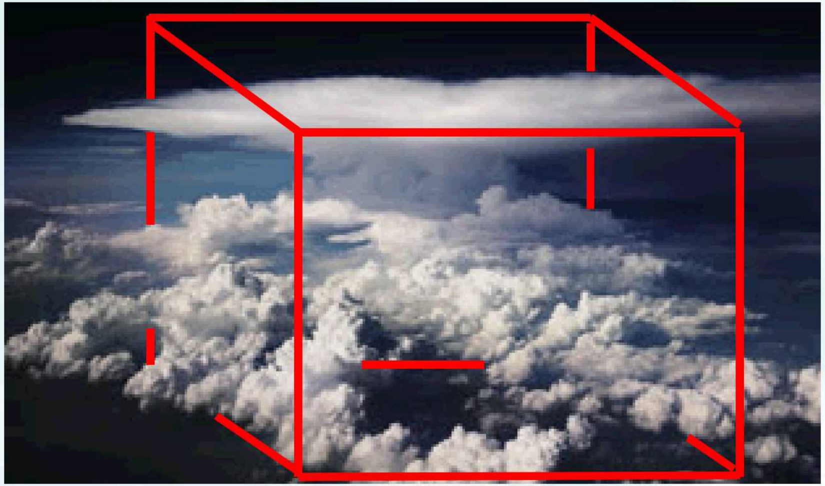

Problem: how to parameterize this subgrid-scale variability?

SCREAM - Simple Cloud-Resolving E3SM Atmosphere Model

- Goal is to keep code as simple as possible
  - You shouldn't trust results you don't understand physically…
  - Simplicity makes clean rewrite (needed for performance) possible
  - Resolving more makes complex parameterizations less important
- Not quite cloud-resolving (yet!), but makes for a cool acronym
  - Target dx = 3 km globally, 128 vertical layers with a top at 40 km
- E3SM: "Energy Exascale Earth System Model" (US Department of Energy coupled earth system model)
- E3SM ocean and sea-ice already work at these scales
  - Goal here is a *coupled* km-scale system, not just an atm model

Why 3 km resolution?

Topography at 25 km resolution

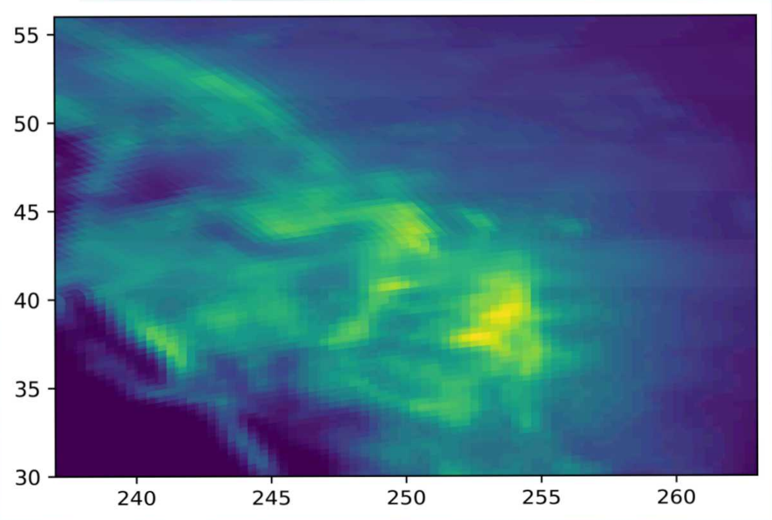

Topography at 3 km resolution

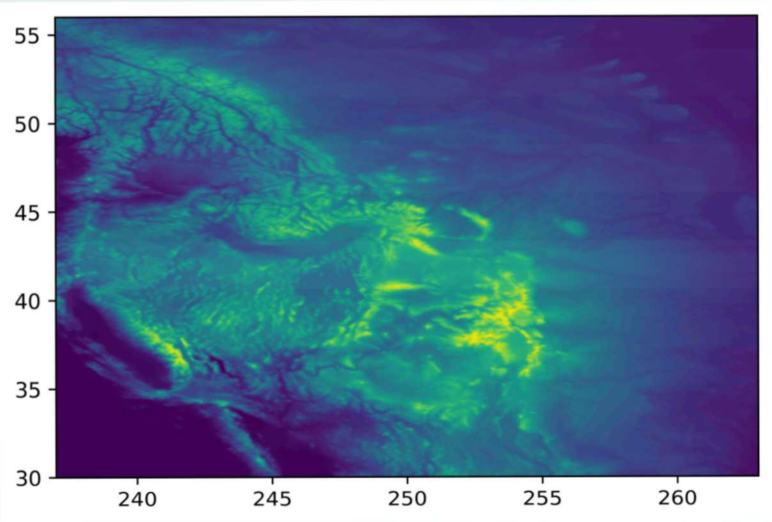

New 3 km resolution for SCREAM better resolves topographic features, as well as eliminates need for parameterization of processes unresolved at previous “high resolution” configuration of 25 km.

Why 3 km resolution?

- The impact of cold pools on DMS transport illustrates what's missing at coarser scales (see figure)

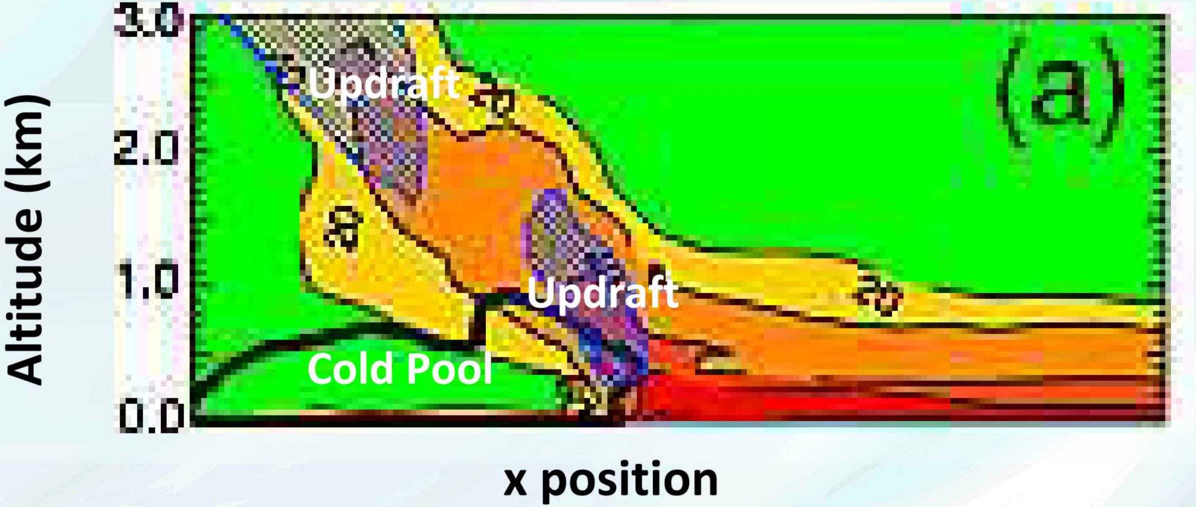

Fig: Snapshot of DMS concentration (green=low, red=high) in a 2d CRM of oceanic deep convection (domain = 256 km, Δx = 1 km). Neglecting the spatial pattern of DMS reduced convective transport by 50%. From Devine et al, 2006 GRL, pointed out by Ken Carslaw

Thesis: Climate Change Can be Understood from Short Runs

This thesis is critical to SCREAM because high-res requires short timesteps (for CFL stability), limiting simulation length

Reasonable because:

- Clouds are the main source of climate uncertainty
- Clouds respond rapidly to forcing change

  ⇒ Cleverly-designed short runs should tell us a lot

- Several clever short tests already exist
  - 5 yr Cess (prescribed SST increase) runs
  - 15 mo aerosol sensitivity tests nudged to observations

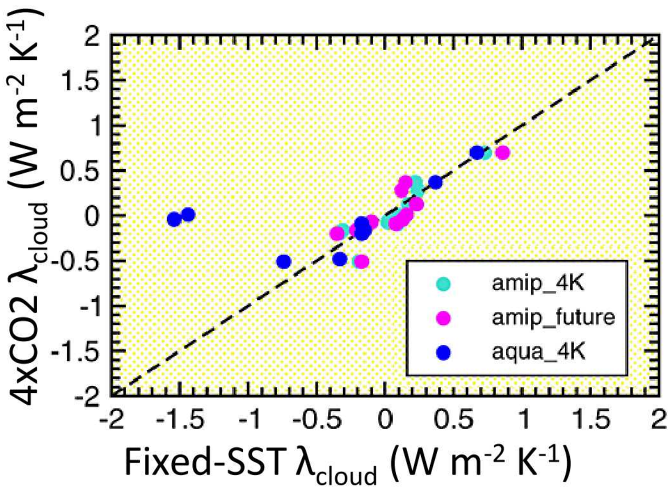

Fig: Cloud feedback from full-complexity (y-axis) versus fixed SST simulations in CMIP5. Adapted from Ringer et al, (2014 GRL).

Components of a typical global atmosphere model

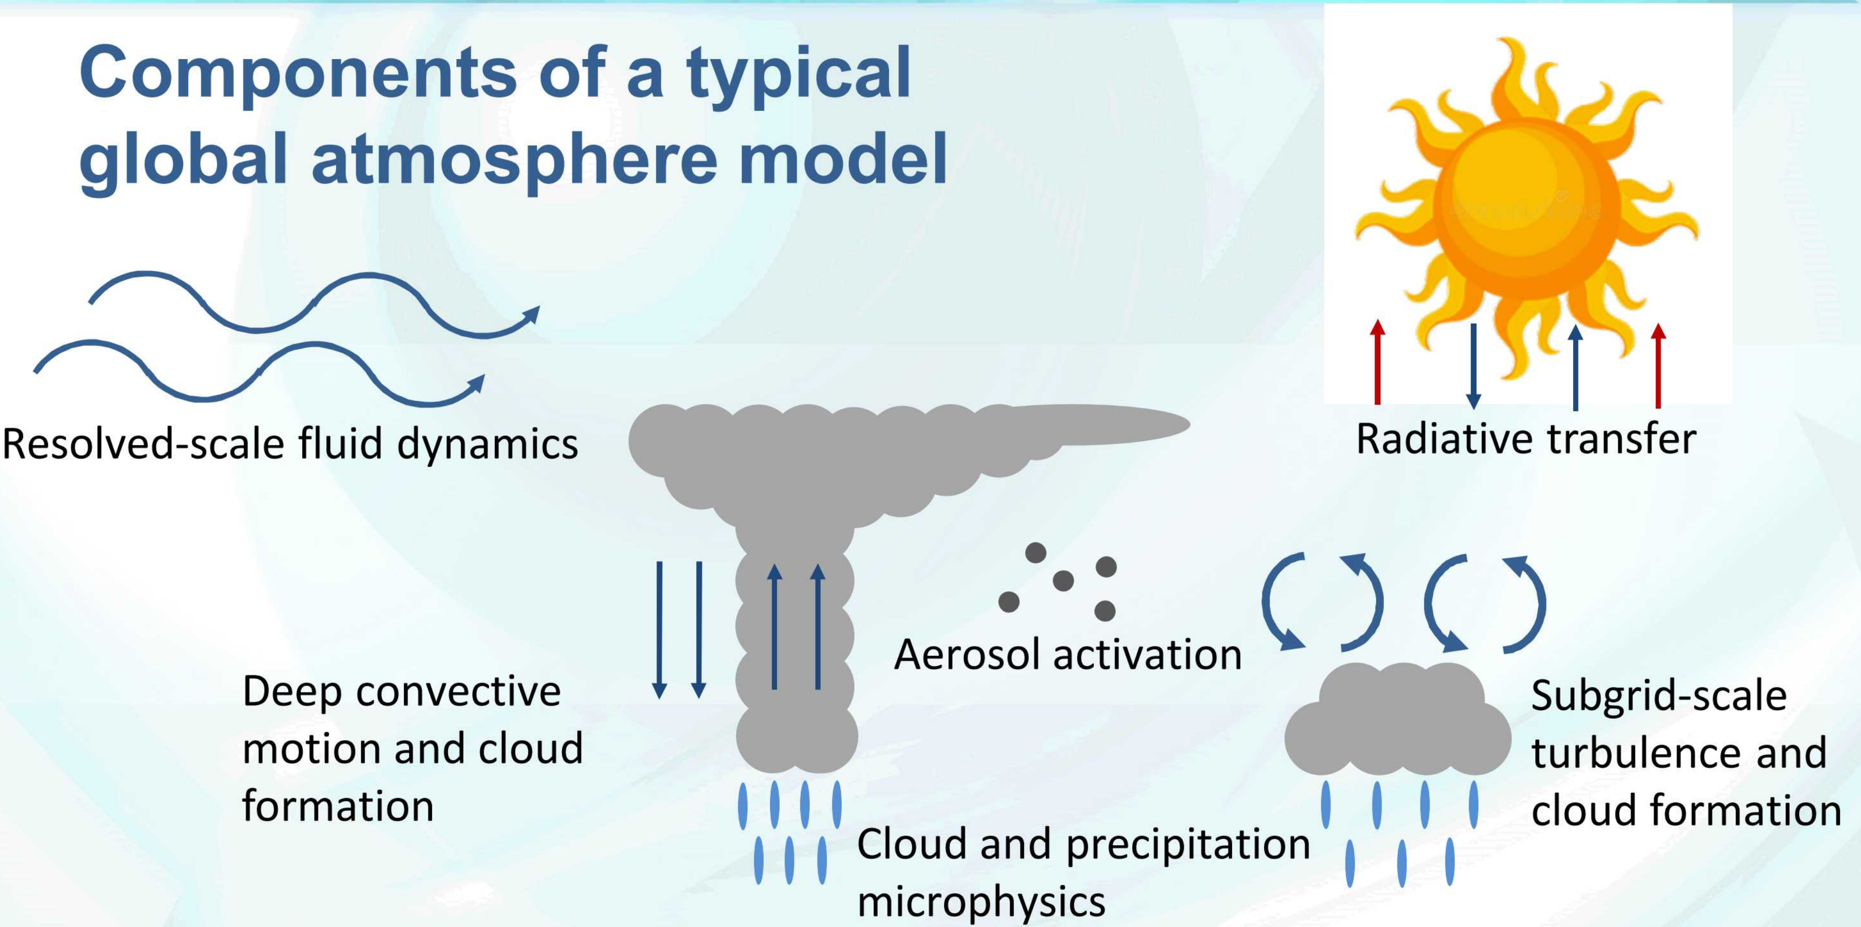

Components of a typical global atmosphere model

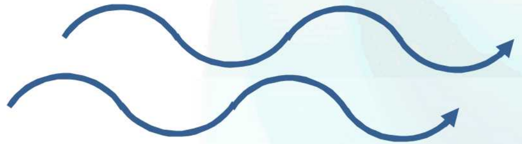

Resolved-scale fluid dynamics:
HOMME non-hydrostatic spectral element model; semi-lagrangian tracer transport

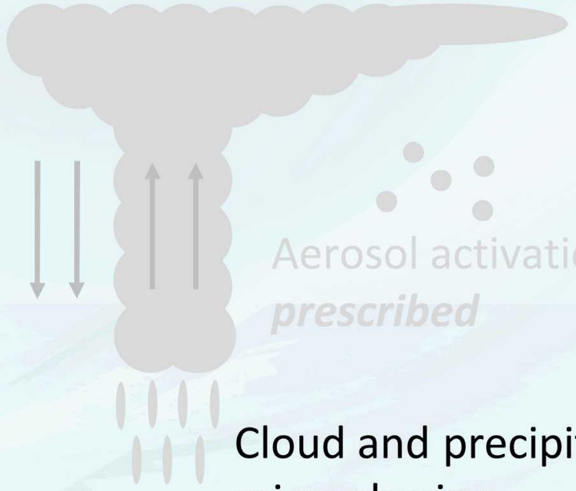

Deep convective motion is ***resolved***

Aerosol activation:
prescribed

Cloud and precipitation microphysics:
Predicted Particle Properties (P3)

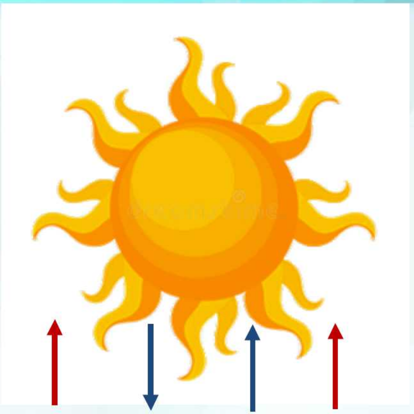

Radiative transfer:
RRTMGP

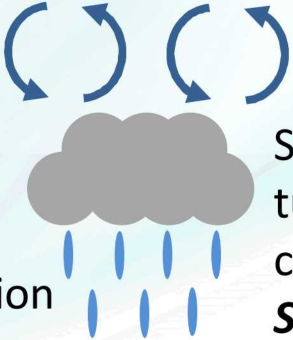

Subgrid-scale turbulence and cloud formation:
Simplified Higher-Order Closure (SHOC)

HOMME

- New non-hydrostatic dynamical core
  - Important for non-hydrostatic effects at 3 km and finer resolutions
- Reformulation of thermodynamic equation in terms of potential temperature
- IMEX time-stepping
- Semi-Lagrangian tracer transport
- C++/Kokkos for performance portability

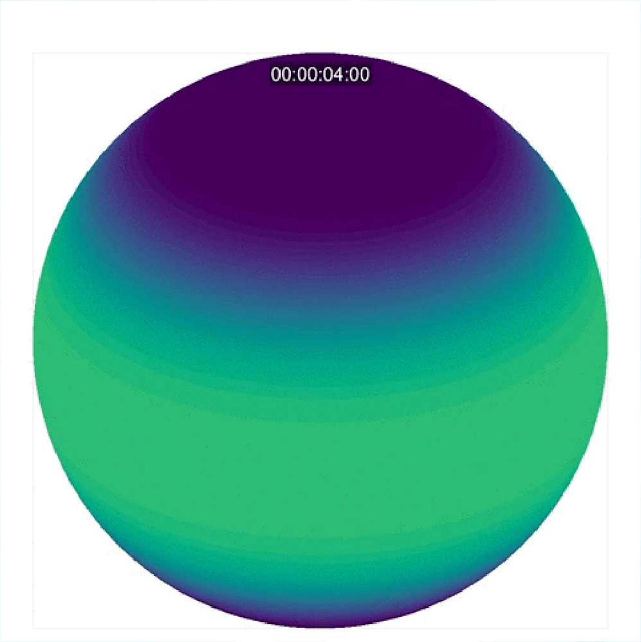

DCMIP2016 baroclinic instability test at 3 km (showing specific humidity)

Simplified Higher-Order Closure (SHOC)

- Bogenschutz and Krueger (2013)
- Represent subgrid-scale clouds and turbulence in cloud resolving models, but at reduced computational cost relative to comparable methods
- PDF-based tri-variate double Gaussian closure

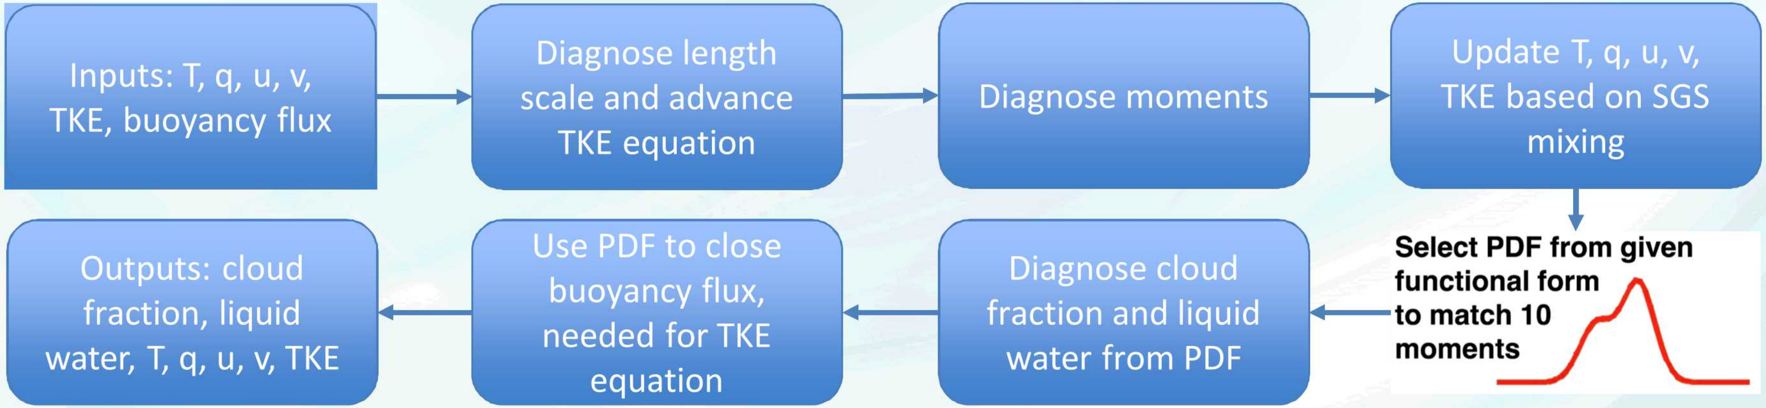

RTE-RRTMGP

- Pincus et al. (2019)
- Re-write of popular RRTMG radiative transfer package, designed with increased parallelism in mind
- Ported to C++/YAKL as part of Exascale Computing Project effort (Matt Norman)

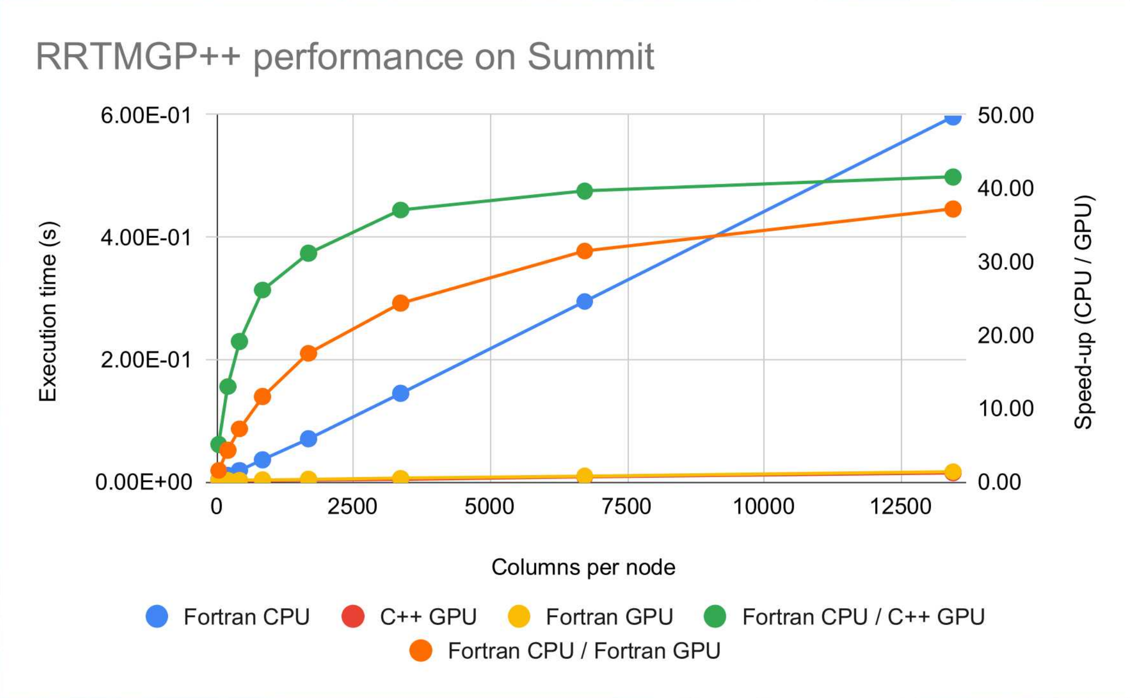

Figure: single-node performance on Summit for example problem (longwave only; similar results for shortwave)

Implementation strategy

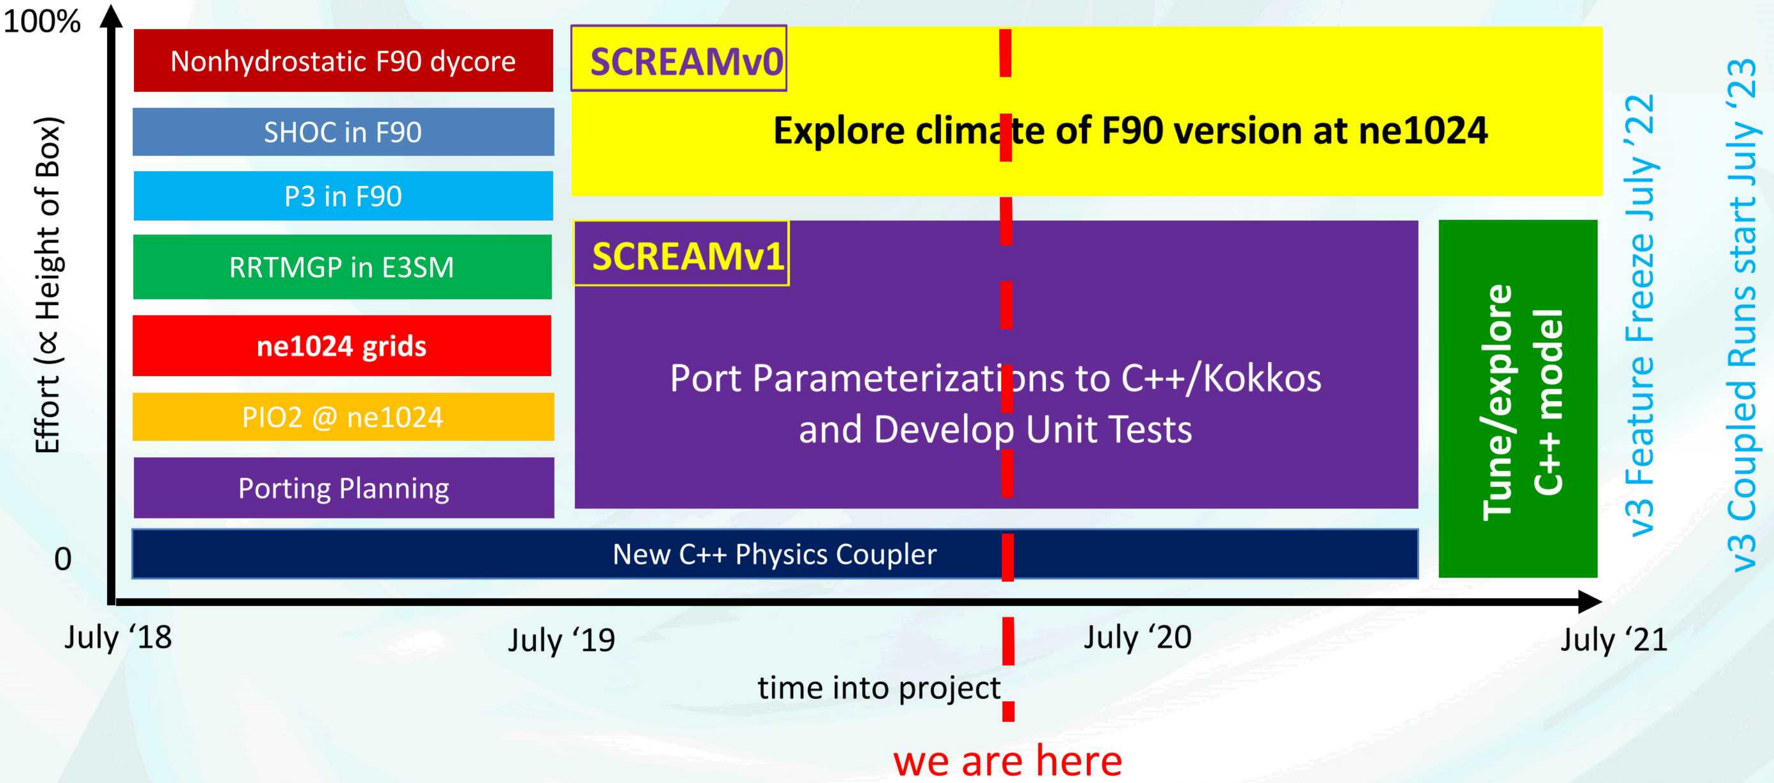

SCREAM v0 (F90 ve

- Goal = DYAMOND Phase 2 Intercomparison
  - Includes ~10 global storm-system r models (GSRMs)
  - 40 day run starting Jan 20, 2020
  - Results due Jan 1, 2021

- ne1024pg2 gets 5.2 simulated *d* wallday on 3072 nodes of cori-k
  - without performance optimization
  - ⇒ 40 day run costs 22.7M NERSC

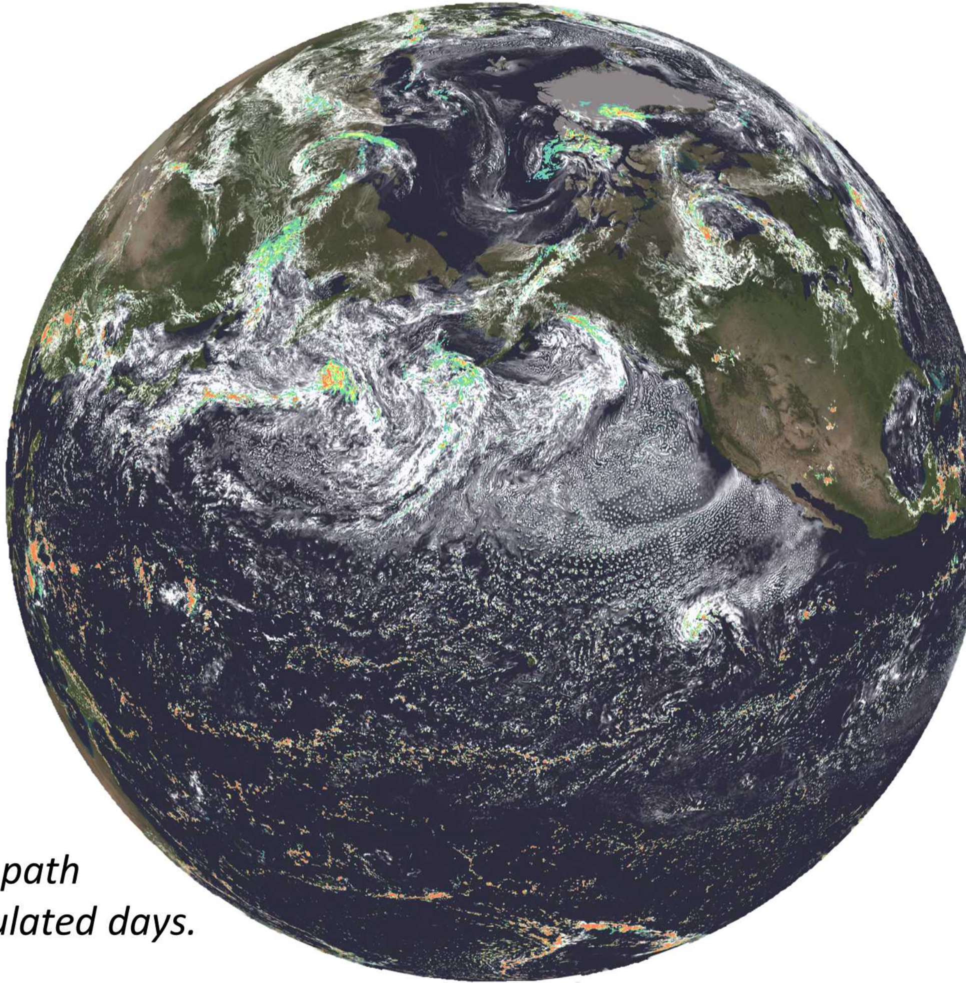

Fig: Snapshot of precipitation (color) and liquid water path (opacity with opaque white = 200 g m^{-2}) after 2.5 simulated days.

SCREAMv0 Radiative Convective Equilibrium (RCE)

Like other models, SCREAM self-aggregates in RCE

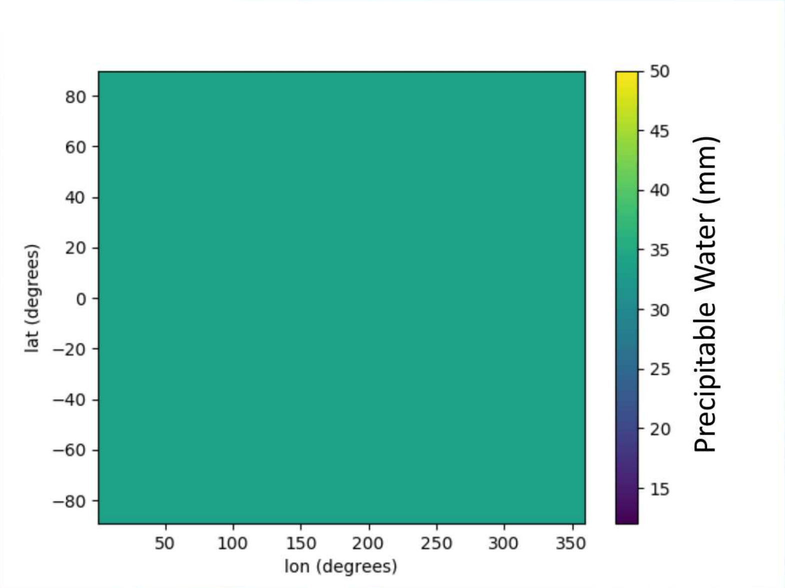

Running varying SST simulations to look at cloud response to warming as well as model physical soundness in an idealized setup.

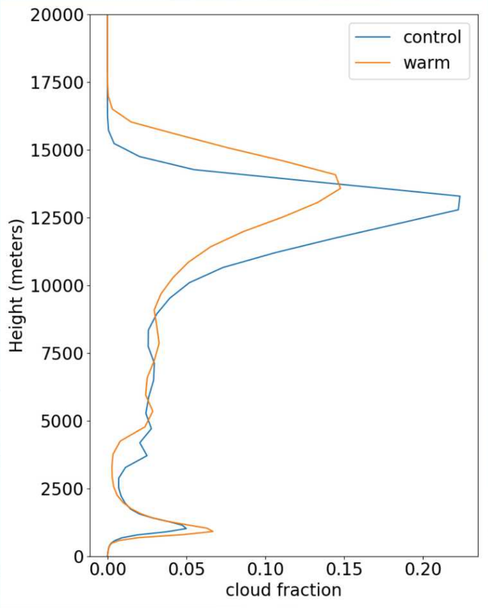

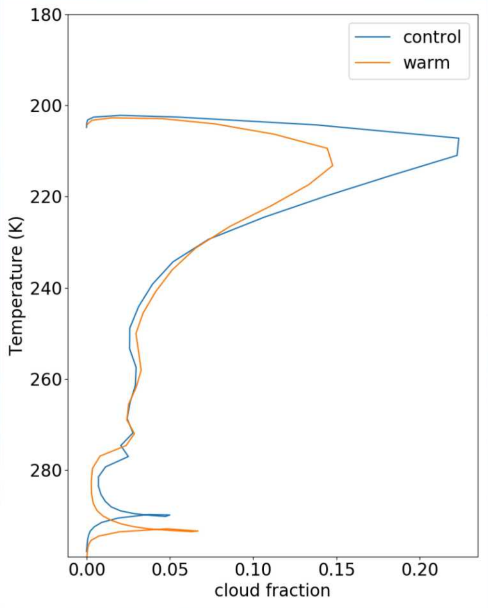

Coding

- SCREAM will be rewritten in C++ using the Kokkos performance portability library
  - Abstracts on-node parallelism
  - Single codebase runs efficiently with variety of hardware (CPU, GPU, etc)
  - Performance portability often comes at the cost of increased code complexity

Original F90

```fortran
kloop_sedi_c2: do k = k_qxtop,k_qxbot,-kdir

   qc_notsmall_c2: if (qc_incld(k)>qsmall) then
      !-- compute Vq, Vn
      call get_cloud_dsd2(qc_incld(k),nc_incld(k),mu_c(k),rho(k),nu,dnu,   &
           lamc(k),tmp1,tmp2,lcldm(k))

      nc(k) = nc_incld(k)*lcldm(k)
      dum = 1._rtype / bfb_pow(lamc(k), bcn)
      V_qc(k) = acn(k)*bfb_gamma(4._rtype+bcn+mu_c(k))*dum/(bfb_gamma(mu_c(k)+4._rtype))
      V_nc(k) = acn(k)*bfb_gamma(1._rtype+bcn+mu_c(k))*dum/(bfb_gamma(mu_c(k)+1._rtype))

   endif qc_notsmall_c2
   Co_max = max(Co_max, V_qc(k)*dt_left*inv_dzq(k))

enddo kloop_sedi_c2
```

Ported to C++/Kokkos

```cpp
Kokkos::parallel_reduce(
  Kokkos::TeamThreadRange(team, kmax-kmin+1), [&] (int pk_, Scalar& lmax) {
    const int pk = kmin + pk_;
    const auto range_pack = scream::pack::range<IntSmallPack>(pk*Spack::n);
    const auto range_mask = range_pack >= kmin_scalar && range_pack <= kmax_scalar;
    const auto qc_gt_small = range_mask && qc_incld(pk) > qsmall;
    if (qc_gt_small.any()) {
      // compute Vq, Vn
      Spack nu, cdist, cdist1, dum;
      get_cloud_dsd2<false>(qc_gt_small, qc_incld(pk), nc_incld(pk), mu_c(pk), rho(pk), nu, dnu, lamc(pk), cdist,
      nc(pk).set(qc_gt_small, nc_incld(pk)*lcldm(pk));
      dum = 1 / (pack::pow(lamc(pk), bcn));
      V_qc(pk).set(qc_gt_small, acn(pk)*pack::tgamma(4 + bcn + mu_c(pk)) * dum / (pack::tgamma(mu_c(pk)+4)));
      if (log_predictNc) {
        V_nc(pk).set(qc_gt_small, acn(pk)*pack::tgamma(1 + bcn + mu_c(pk)) * dum / (pack::tgamma(mu_c(pk)+1)));
      }

      const auto Co_max_local = max(qc_gt_small, -1,
                                    V_qc(pk) * dt_left * inv_dzq(pk));
      if (Co_max_local > lmax)
        lmax = Co_max_local;
    }
}, Kokkos::Max<Scalar>(Co_max));
team.team_barrier();
```

Testing

- Strive for *property tests* (check that code behaves physically) in addition to BFB testing (check that answers have not changed)
- Example of property tests: convergence in dt, dz, dx; Applying to SHOC revealed problems with:
  - Bretherton & Park (2009) shear production boundary condition
  - Blackadar (1984) turbulent length scale near surface
- Encapsulation of parameterizations makes unit testing straightforward

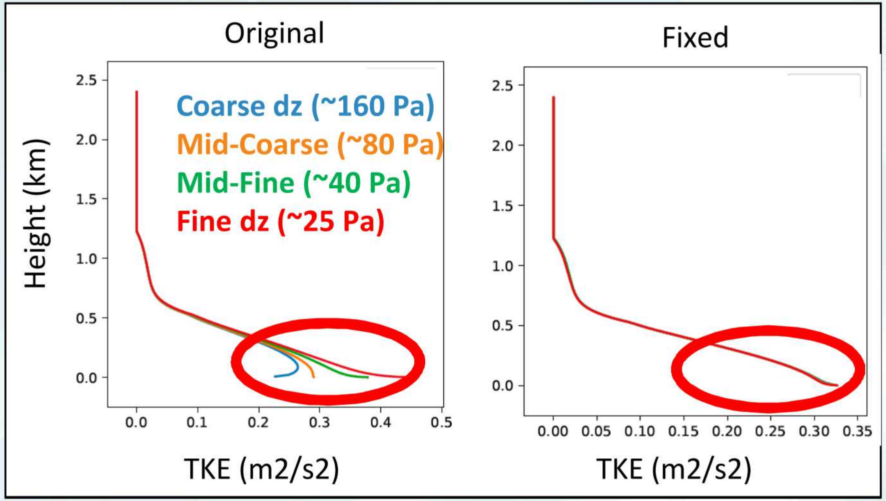

Fig: SHOC standalone simulations running the BOMEX test case (trade wind Cumulus) for 6 hrs with a variety of vertical resolutions. By Peter Bogenschutz

SCREAMv1

- C++ version of non-hydrostatic (NH) dycore done
  - Used for recent Gordon-Bell submission (see figure)
  - 0.97 SYPD using all of Summit
  - Not including semi-lagrangian advection, which gives a further speed-up
  - P3 port nearing completion
  - SHOC port starting now
  - RRTMGP port mostly done (starting interface now)

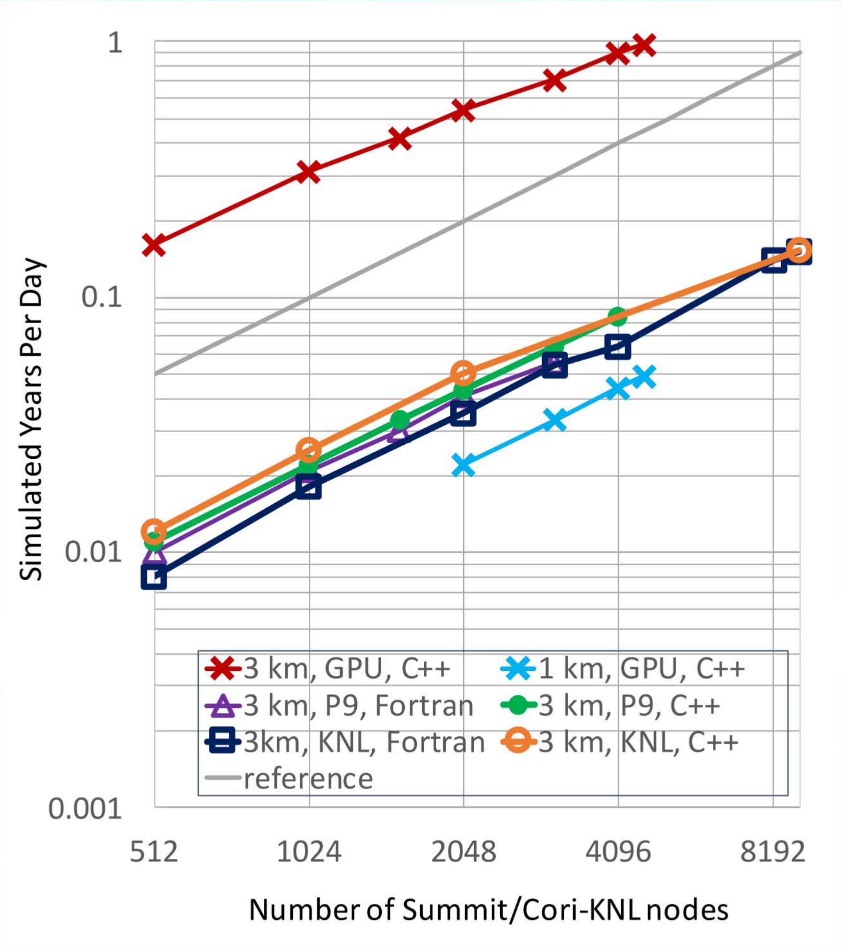

Fig: Nonhydrostatic C++ dycore-only NGGPS benchmark scaling (10 tracers).

Process coupling

- New atmosphere driver; all atmosphere processes are instances of a `atm_process` class
- Having all processes behave the same way:
  - Makes adding/reorganizing/parallelizing processes easy
  - Improves code readability
- Processes broken into:
  - SCREAM-specific **interface layers**
  - Model-agnostic **process implementations** that make:
    - Code easy to share with/implement in other models
    - Standalone process simulations straightforward (useful for testing)
- Processes communicate entirely through a **field manager** that provides interface layers with pointers to requested variables

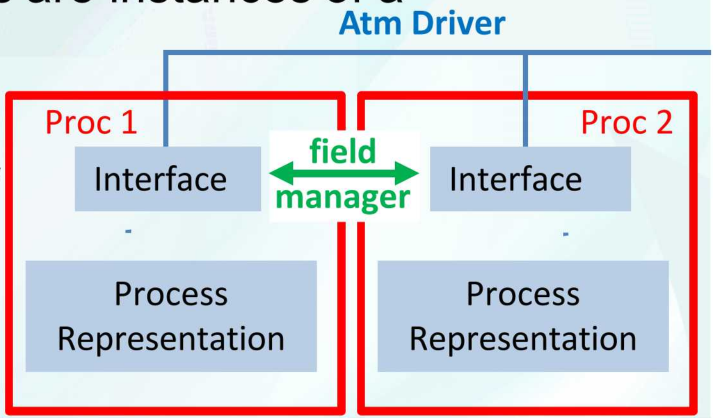

Fig: SCREAM coupler structure

Initial results

- DYAMOND Phase 1 configuration
  - Initialized from IFS reanalysis for 01 August 2016
  - Showing 2-day simulation from initialization
  - Prescribed SST/sea ice
- Results are *reasonable* after initial spin-up
  - Both shortwave and longwave fluxes are slightly above CERES daily average
  - Global mean precip rate slightly above GPM daily average
- Rigorous evaluation coming soon…

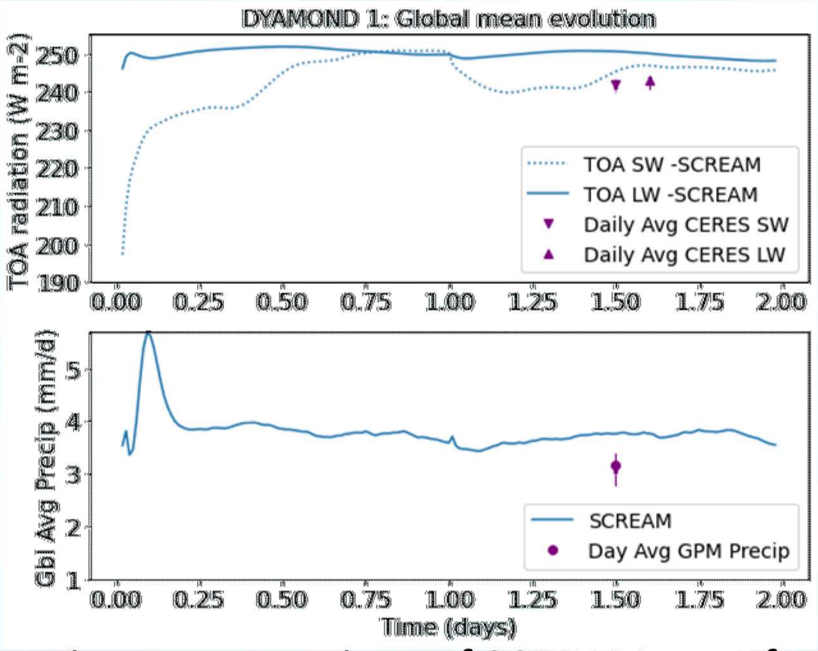

Figure: comparison of SCREAM top of atmosphere fluxes with CERES, and precipitation rate with GPM

Initial results

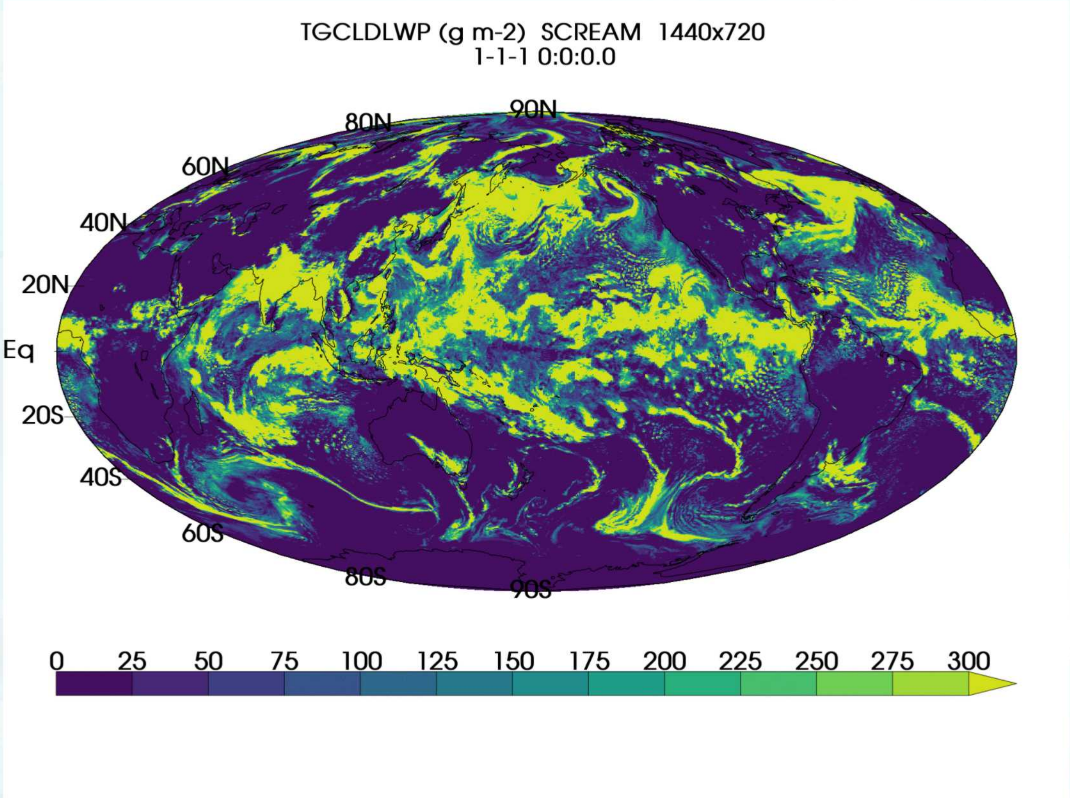

Figure: animation of total gridbox-mean cloud water path from initial simulation at 3 km resolution

Challenges

- 3 km global resolution makes for *huge* grids
  - Very expensive (~5 simulated *days* per wallclock day)
  - I/O is a problem (restart files for atmosphere alone are > 3 TB in size)
- Debugging
  - Difficult if not impossible to debug at scale
  - Bugs specific to coupling at high resolution not always reproducible at lower resolution

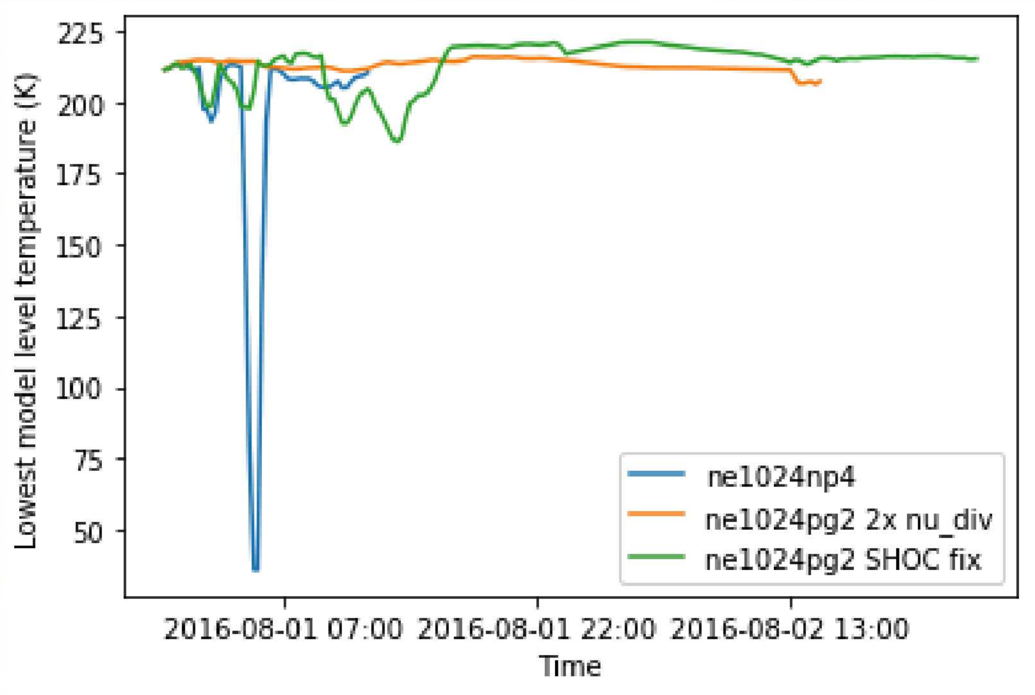

Fig: temperature instability arising only when coupling physics and dynamics at very high resolution

Next steps

- DYAMOND Phase 2
- Continue porting to C++/Kokkos for performance portability
- Migrate to pre-exascale and exascale machines

References and further reading

- SHOC: Bogenschutz, P. A., and Krueger, S. K. (2013), A simplified PDF parameterization of subgrid-scale clouds and turbulence for cloud-resolving models, *J. Adv. Model. Earth Syst.*, 5, 195–211, doi:10.1002/jame.20018.
- HOMMEXX: Bertagna, L., Deakin, M., Guba, O., Sunderland, D., Bradley, A. M., Tezaur, I. K., Taylor, M. A., and Salinger, A. G.: HOMMEXX 1.0: a performance-portable atmospheric dynamical core for the Energy Exascale Earth System Model, Geosci. Model Dev., 12, 1423–1441, https://doi.org/10.5194/gmd-12-1423-2019, 2019.
- RRTMGP: Pincus, R., Mlawer, E. J., & Delamere, J. S. (2019). Balancing accuracy, efficiency, and flexibility in radiation calculations for dynamical models. *Journal of Advances in Modeling Earth Systems*, 11, 3074– 3089 https://doi.org/10.1029/2019MS001621
- P3: Morrison, H. and J.A. Milbrandt, 2015: Parameterization of Cloud Microphysics Based on the Prediction of Bulk Ice Particle Properties. Part I: Scheme Description and Idealized Tests. *J. Atmos. Sci.*, **72**, 287–311, https://doi.org/10.1175/JAS-D-14-0065.1