MPAS-A and Muram porting and Optimization tale

Supreeth Suresh
Cena Miller

June 1, 2021 - NHUG meeting

Slides taken from various presentations by Rich Loft and Raghu Kumar

Model for Prediction Across Scales - Atmosphere (MPAS-A)
A Global Meteorological Model & Future Earth System Component

Simulation of 2012 Tropical Cyclones at 4 km resolution
– Courtesy of Falko Judt, NCAR

Our Approach to Refactoring for CPU/GPU Portability

- Use OpenACC standard directives to achieve performance portability
- Test driven development
- Profiling to prioritize refactoring targets

```
!$acc parallel
!$acc loop seq
do k=1,nz
 !$acc loop collapse(2)
 do J=Jsq,Jeq ; do i=is,ie
 BT_force_v(i,J) = BT_force_v(i,J) + wt_v(i,J,k) * bc_accel_v(i,J,k)
 enddo ; enddo ; enddo
 !$acc end parallel
```

```
VALIDATION RESULTS...

Density        : 1.0241467e-10   PASS
Temperature    : 1.0215635e-10   PASS
Velocity       : 3.2897487e-09   PASS
Energy         : 7.567654e-11    PASS
```

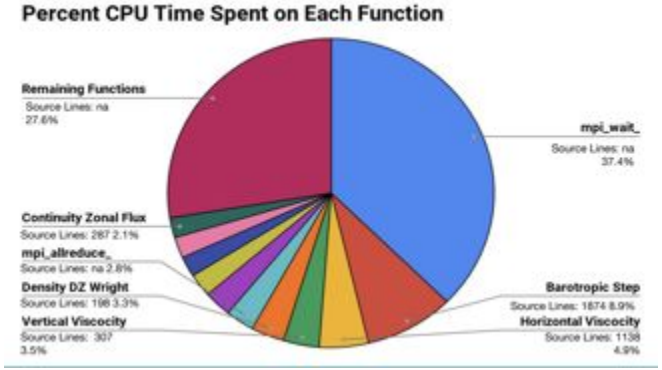

MPAS Generic Execution Method

*Slide by Rich Loft

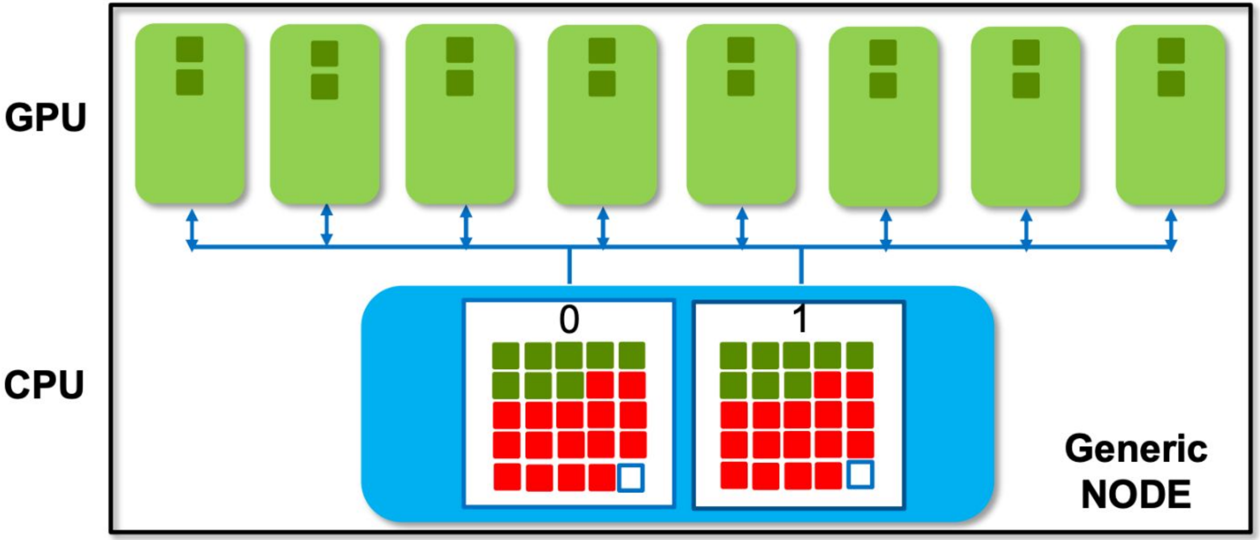

GPU/CPU Integration ranks

Radiation ranks

OS-I/O ranks

Take Offs:

- Integration ranks
- Radiation ranks
- OpenMP for Radiation
- MPS Enabled

MPAS Call Structure

Begin Timestep

1. Integration Setup
2. Moist coefficients
3. Physics tendencies
4. Vertically implicit coefficients
5. Dynamic tendencies
6. Small step
7. Acoustic solver
8. Divergent damping
9. Large step
10. Scalars
11. Solve diagnostics
12. Substeps
13. Scalars
14. Velocity Reconstruction

End Timestep

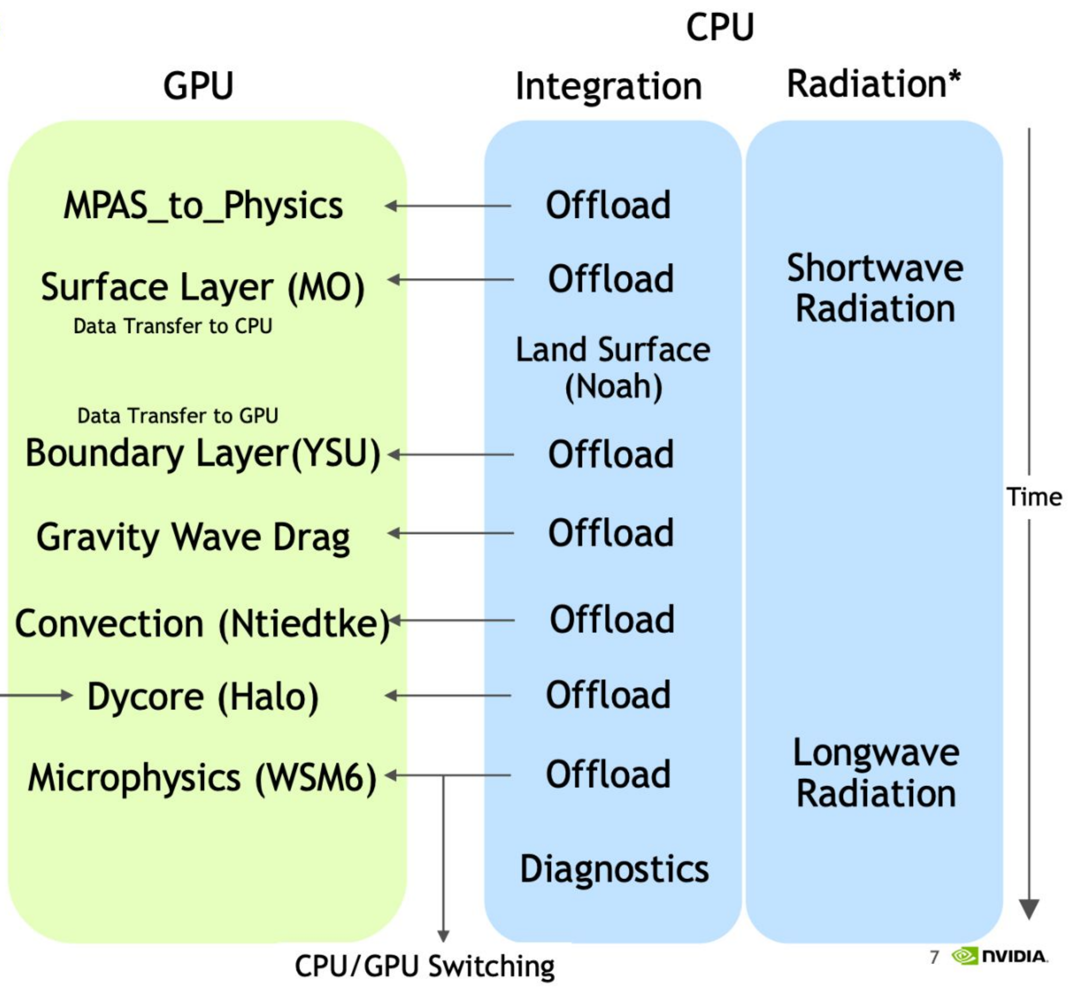

3 Years ago

: MPAS-A Dynamics on Summit[1] vs Cheyenne[2]

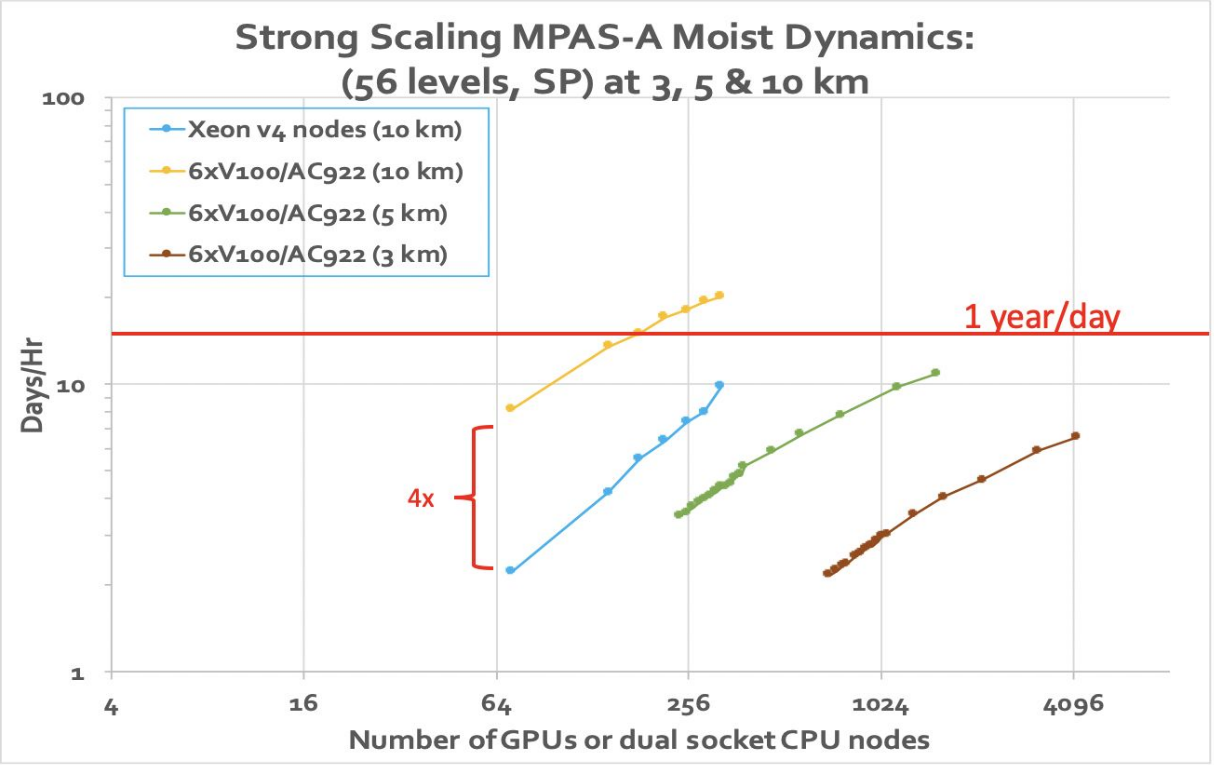

[1]Benchmarking on Summit supported by DoE via an OLCF Director's Discretionary Allocation
[2]Cheyenne is a 5.4 PF, 4032-node HPE system with EDR interconnect operated by NCAR

2 Years ago

MPAS full physics with lagged radiation

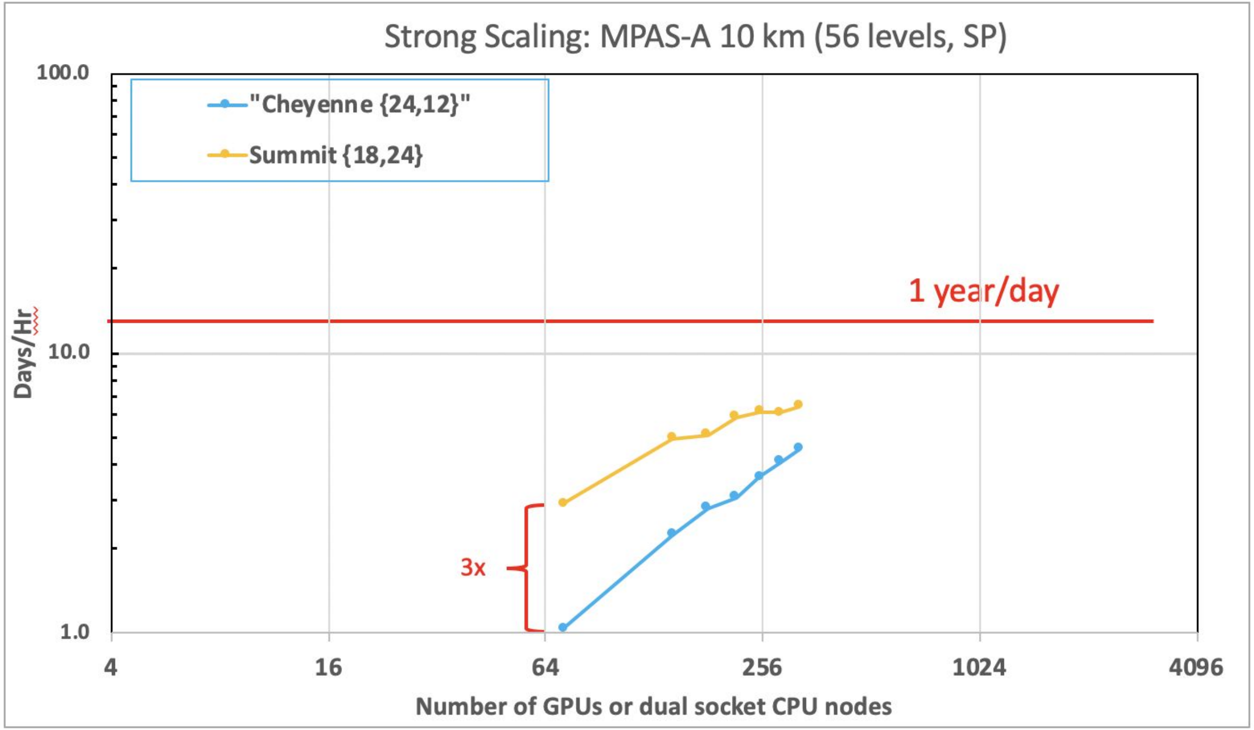

[1]Benchmarking on Summit supported by DoE via an OLCF Director's Discretionary Allocation
[2]Cheyenne is a 5.4 PF, 4032-node HPE system with EDR interconnect operated by NCAR

MPAS GPU: Summit Breakdown

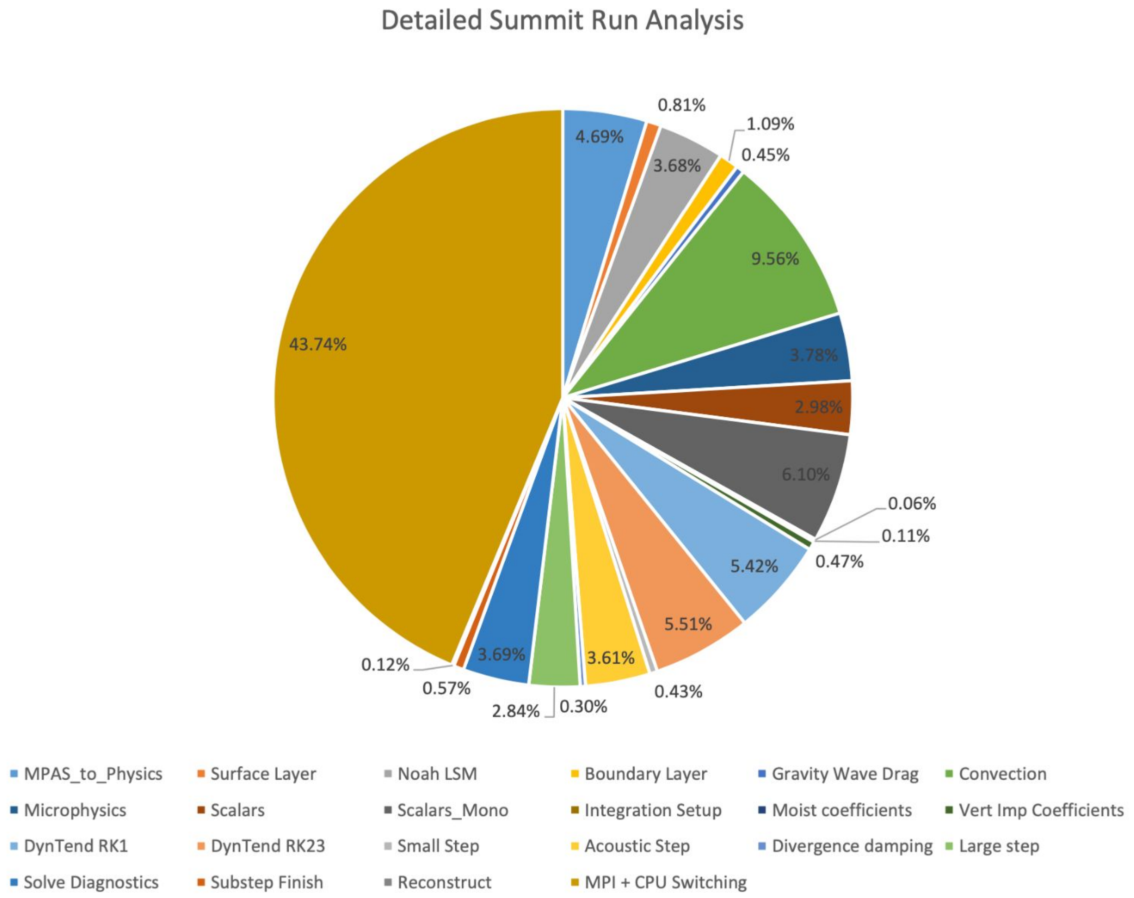

MPAS-A Strong Scaling Improvements: Moist Dynamics (56 levels, SP) at 10 km

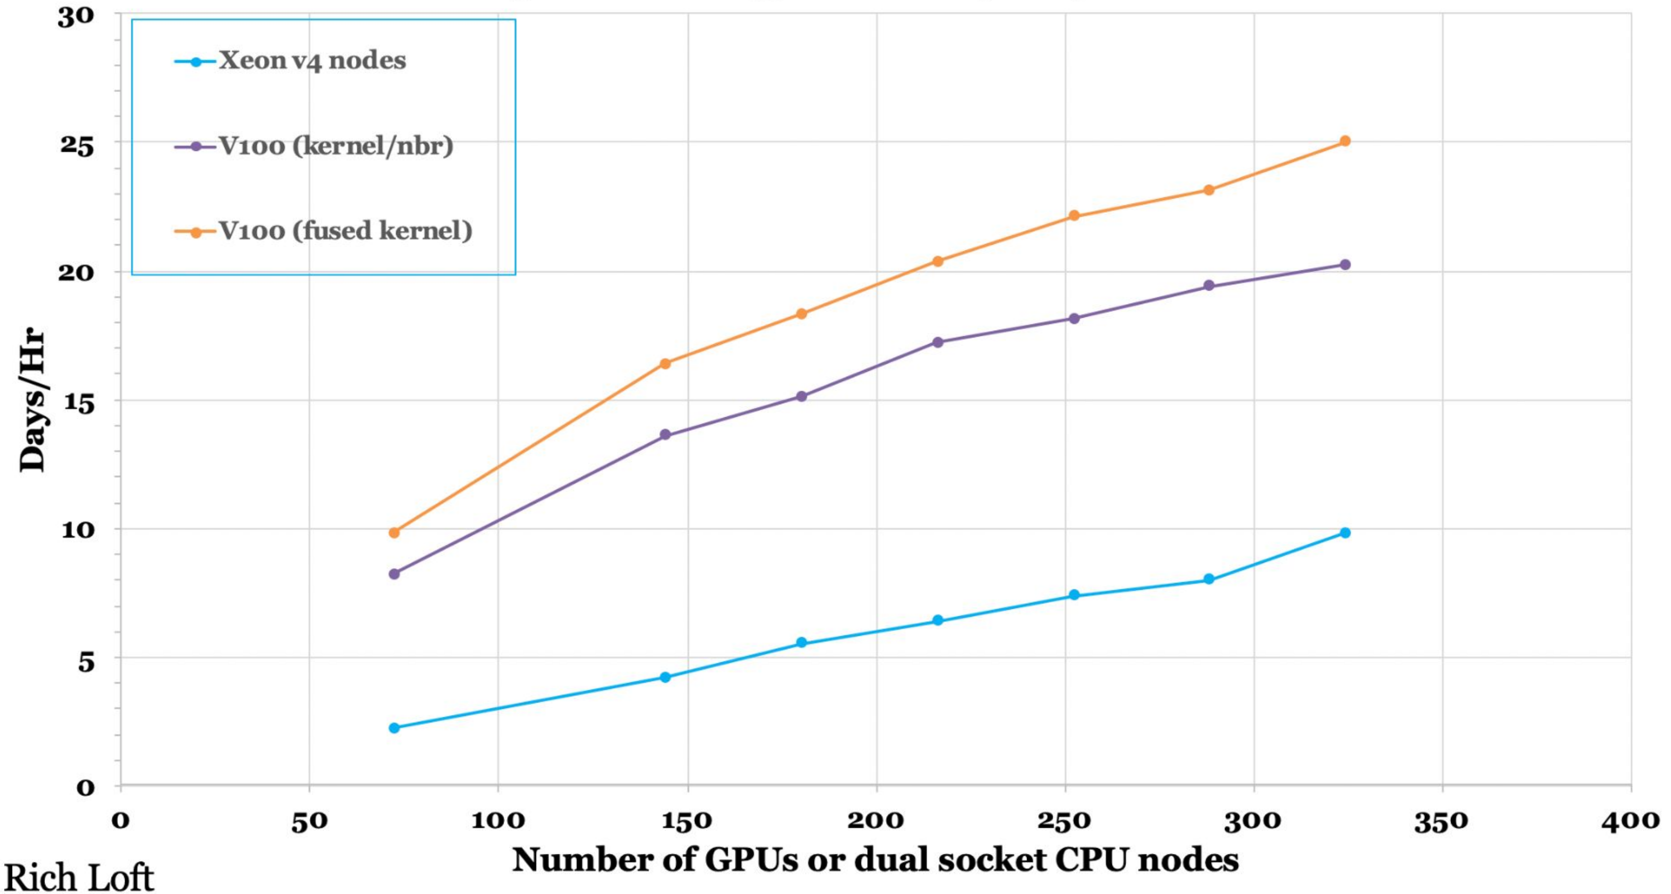

Slide by Rich Loft

MURaM (Max Planck University of Chicago Radiative MHD)

- Solar model for simulations of the Sun's atmosphere
- Jointly developed by HAO, the Max Planck Institute for Solar System Research (MPS) and the Lockheed Martin Solar and Astrophysics Laboratory (LMSAL)

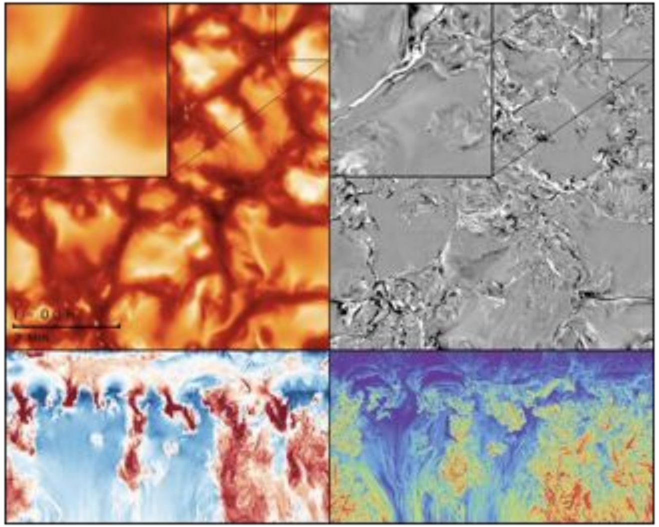

MURaM simulation of solar granulation

Refactoring MURaM with OpenACC: Challenges and Solutions

Eric Wright, CRPL
University of Delaware

- 3D Radiative Transport (RT) is the most expensive routine in MURaM.
- RT is so costly that it is typically run with one frequency bin (grey RT).
- The **Integrate** function (called along 24 different ”rays”) in the most expensive in RT.
- It has a dependency in the outermost loop that:
  - Creates hundreds of kernel launches per **Integrate** call
  - Limits data parallelism (occupancy)
  - Creates badly-strided memory accesses in 3D variables

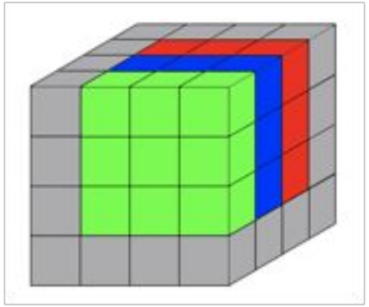

The serial dependency in **Integrate**

- Remedies
  - **asynch** programming to combine processing of multiple rays
  - Loop fusion to increase amount of work performed inside each kernel.
  - Replicate arrays in transposed format to reduce striding hit

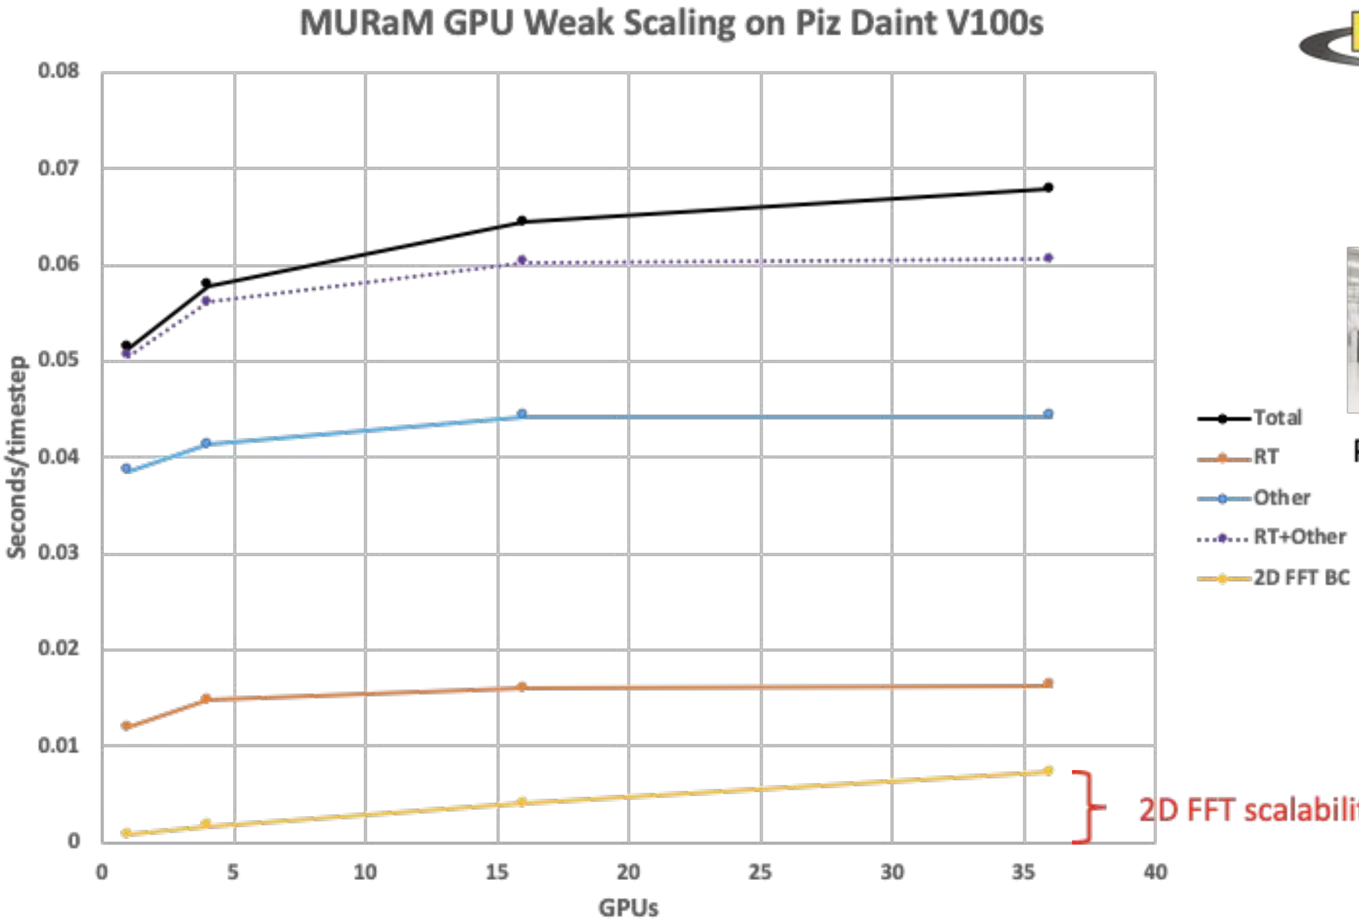

2D FFT scalability is hurting us!

Damian Przybylski

Piz Daint (XC40/XC50) @ CSCS

The Sun is big (!), so weak scaling represents increasing the area modeled while holding the resolution (and per-device patch size) fixed.

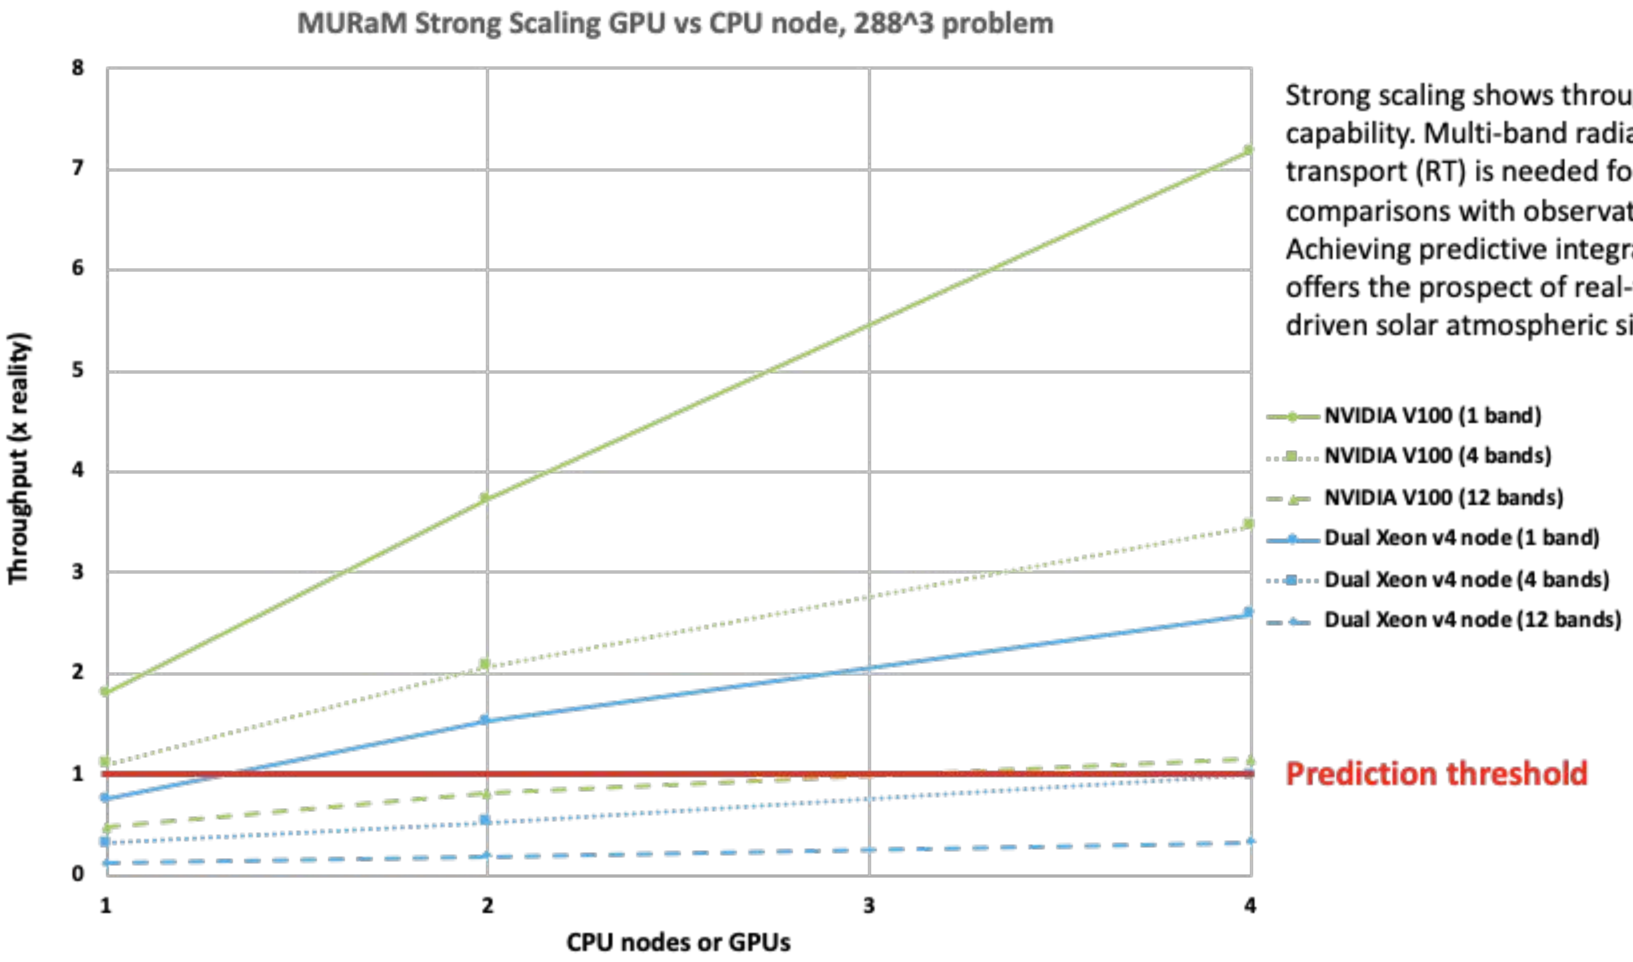

Strong scaling shows throughput capability. Multi-band radiative transport (RT) is needed for detailed comparisons with observations. Achieving predictive integration rates offers the prospect of real-time, data-driven solar atmospheric simulations.

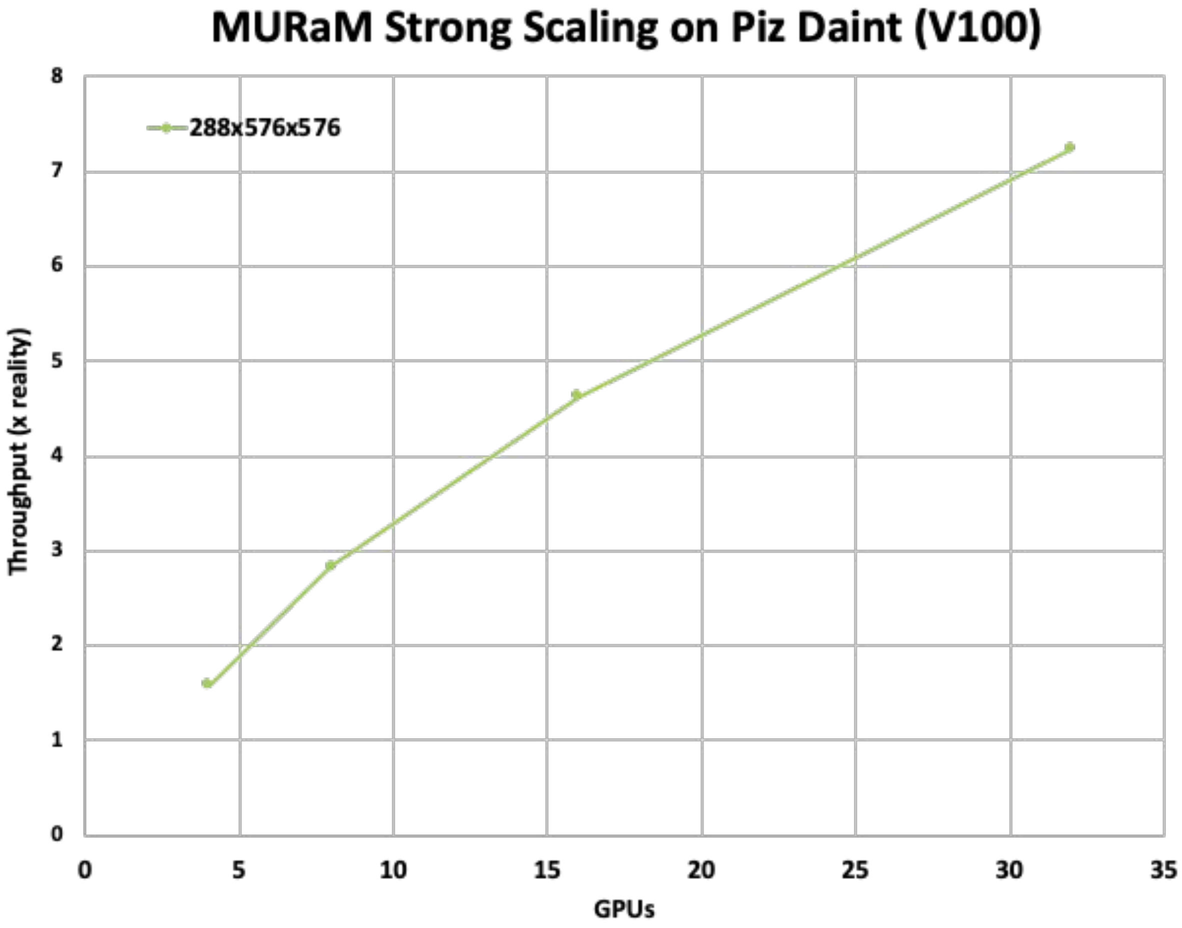

Piz Daint (XC40/XC50) @CSCS

Putting the two scaling studies together, the strong scaling throughput remains good on larger domains using more GPUs.

Prediction threshold

Thank you!!!!
and
Questions?