EarthWorks Update

EarthWorks is a five-year project, in which CSU and 3 NCAR laboratories are working together as partners.

The goal is to develop a *global convection-permitting coupled model* based on the CESM.

Nevertheless, EarthWorks is not a project of CESM.

Earthworks is supported by NSF/CISE. The CSU/NCAR split is 60/40.

The EarthWorks Vision

- Using CESM as a base, create a coupled model with the same 3.75-km geodesic grid for atmosphere, ocean, and land surface.
  - Partially explicit deep convection and gravity waves,
  - Resolved stratosphere,
  - Regional refinement option.

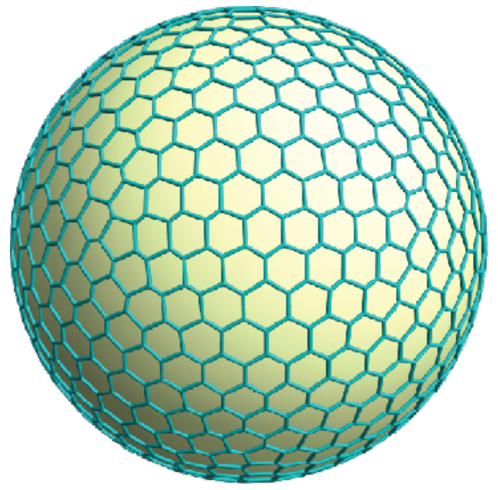

- Use the model to study both weather and climate, on time scales ranging from days to years.
- Use the model to understand deficiencies of lower-resolution versions of CESM (and other models).
- Use the model to create training data sets for machine learning.
- Develop improved parameterizations that work well with both 3.75-km and 120-km grids.
- Use lower-resolution versions of the same model to study century-scale climate change.

EarthWorks consists of:

- The MPAS non-hydrostatic dynamical core, with a resolved stratosphere and CAM-ish physics
- The MPAS ocean model, developed at Los Alamos
- The MPAS sea ice model, based on CICE
- The Community Land Model (CLM)
- The Community Mediator for Earth Prediction Systems (CMEPS)
- The Community Physics Framework, when ready

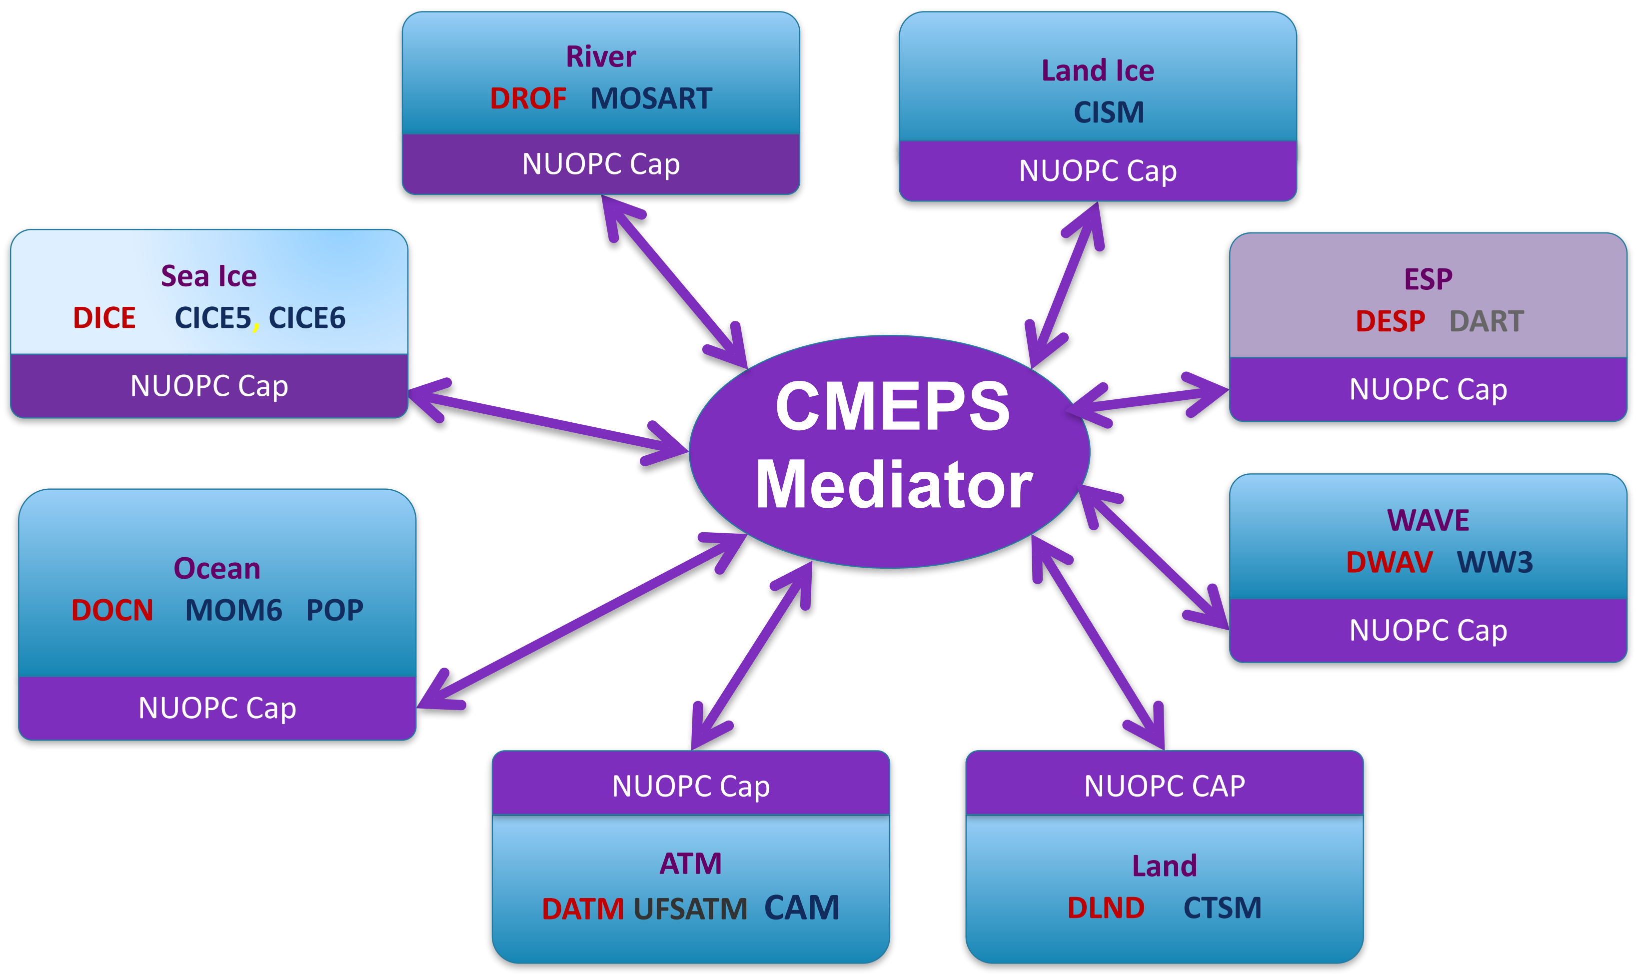

Goals

Science Goals

- The model will resolve mesoscale storms. ocean eddies, mountains, and large lakes and rivers.
- The deep convection and gravity-wave drag parameterizations will be eliminated.
- The stratosphere will be resolved.
- The model will be used to study the interactions of mesoscale weather systems with larger spatial scales, on time scales of days to years.

Computational Goals

- The model will run on CPUs for low resolution experiments and for short tests with ultra-high resolution.
- The ocean and atmosphere will run on GPUs for ultra-high resolution simulations.
- Our 2025 performance goals for a version of EarthWorks with 3.75 km global grid spacing are:
  - *Half a simulated year per day in atmosphere-only simulations with a resolved stratosphere, and*
  - *One simulated year per day in coupled simulations with fewer stratospheric layers.*

  This would be on a DOE "leadership class" machine.

Issues already found and fixed

A multi-institutional collaboration: CSU, NCAR, TACC, ESMF Core Team, AER

- Initialization (abnormally long times, high memory use)
  - Traced to an issue in ESMF, resulted in a patch release.
  - Impact: 5.7x speedup, 2x reduction in memory use during initialization.
- 100x slowdown in history I/O bandwidth
  - Traced to the ROMIO MPI-IO implementation in PnetCDF, resulted in a problem report and workaround.
  - Impact: Expected history I/O performance restored
- Run after restart errors
  - Traced to an issue with the PIO2 (parallel I/O) infrastructure in CESM, resulted in a patch release.
  - Impact: Correct model restarts restored
- Differing results between the CPU & GPU implementations of AER's Radiative Transfer Model (RRTMGP)
  - Traced to an incorrect argument being passed in the GPU version.
  - Fed back to the developers.
  - Impact: GPU-based results now match CPU version

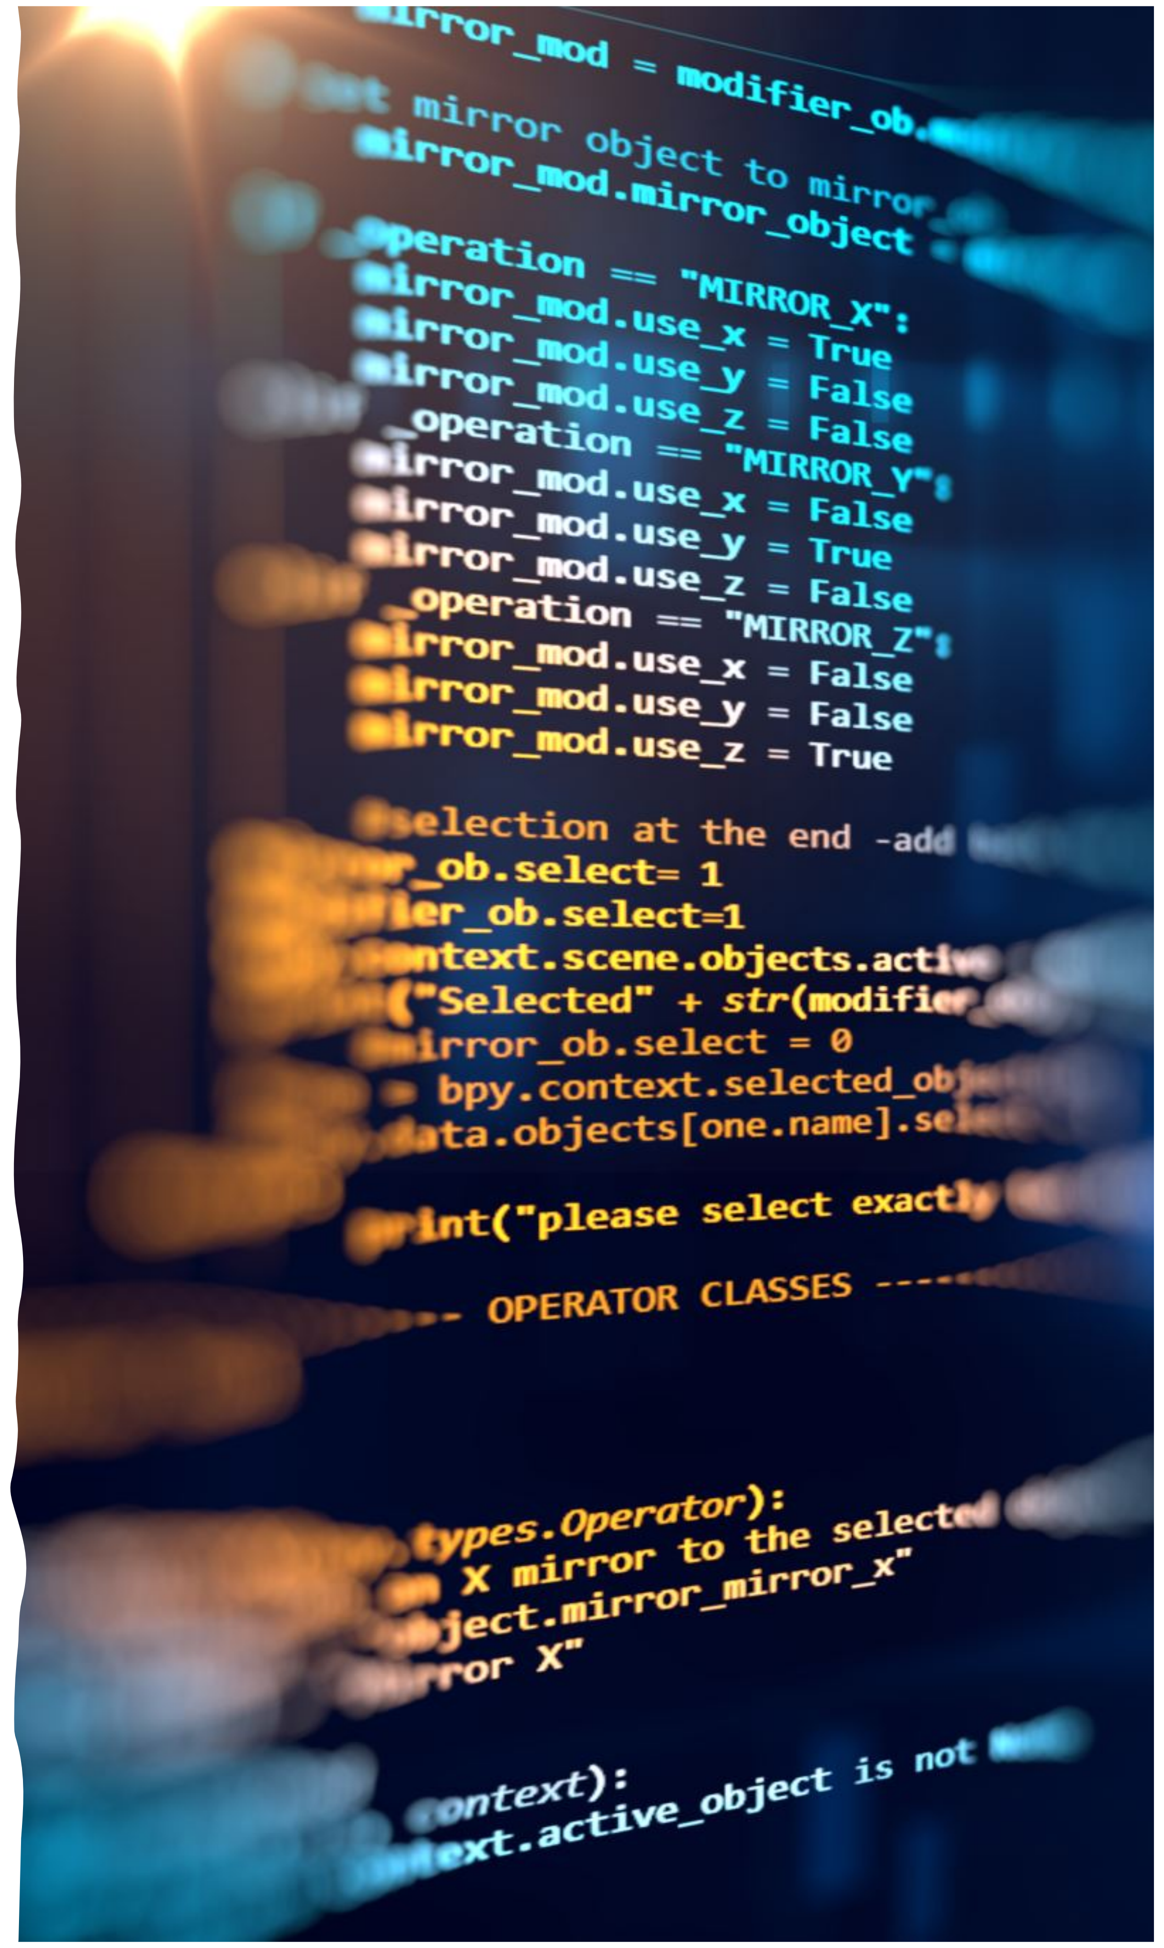

EarthWorks software roadmap

Version (Target Delivery Date)	Deliverables
Version 1.0 (Released 3/2023)	CPU-only 120,60,30 km @ 32 levels Compsets: FHS94,FKESSLER,Aquaplanet, F2000Climo, Fully Coupled Compilers: Intel classic/GNU
Version 2.0 (Released 11/2023)	+ 15 km resolution @ 58 levels + nvhpc compiler support + GPU PUMAS microphysics[1]
Version 2.1 (Planned 2/2024)	+ bug fixes + 30 km resolution (GPU) + GPU-offload MPAS dycore
Version 3.0 (Planned: 11/2024)	+ 7.5 km resolution, resolved stratosphere + Scalable diag tools v1 release + GPU offload for all target physics params[1]
Version 4.0 (Planned: 6/2025)	+ 3.75 km resolution + Scalable diag tools release + MPAS-Ocean GPU offload

[1] RRTMGP (radiative transfer), CLUBB (low clouds) , and PUMAS (μphysics)

\+ indicates a new feature/capability added cumulatively.

What's in EarthWorks SW Version 2.1

- **Version 2.1 does include the following improvements:**
  - Release of GPU-offload version MPAS-7.x dynamical core (30/60 km).
  - Code base is integrated with CAM6.3 tag 145 level (up from tag 124 in 2.0).
    - Temporary patch version of aerosol_optics_cam.F90 for nvhpc compiler.
  - MPAS Sea-Ice build issue under **ifx** (Intel OneApi fortran compiler) resolved.
- **Version 2.1 does not include the following elements:**
  - Fix to atmosphere run after restart issue under the nvhpc compiler.
  - GPU-offload versions of:
    - CLUBB (PR issued)
    - RRTMGP (PR under SW code review)

How do I get EarthWorks Version 2.1?

- GitHub access via git clone at:
  - https://github.com/EarthWorksOrg/EarthWorks.git
- .tar.gz or .zip files available here:
  - https://github.com/EarthWorksOrg/EarthWorks/releases/tag/release-ew2.1
- Questions/issues? Contact: Donald.Dazlich@colostate.edu

EarthWorks Software V2.1 Release: CPUs

Compiler	Nvhpc				Gnu				Intel (OneAPI)			
	120km	60km	30km	15km	120km	60km	30km	15km	120km	60km	30km	15km
FHS94 (32L)												
FKESSLER (32L)												
QPC6 (32L)												
F2000Climo (32L)												
F2000Climo (58L)												
Fully Coupled(32L)												

Runs

Doesn't restart

Not currently supported

EarthWorks GPU Offload Testing Roadmap

Compiler	OpenAcc				OpenMP Offload			
Offload Configuration	**120km**	**60km**	**30km**	**15km**	**120km**	**60km**	**30km**	**15km**
Dynamical Core (MPAS V7.x)								
cloud microphysics (PUMAS)								
Radiative Transfer (RRTMGP)								
Low Cloud Parameterization (CLUBB)								

Version 2
11/23

Version 2.1
2/24

Version 2.5*
6/24

Version 3
11/24

* The need for a 2.5 release is TBD.

PUMAS/MG3 Microphysics on GPUs

- Microphysics code has been fully ported onto GPUs and has been integrated into the CAM development code
  - Includes both OpenACC version and OpenMP offload versions
- Comparison done between one EPYC 7763(AMD Milan) and one A100 in the integrated within CAM (timing results contain data transfer time)
- Maximum speedup over CPUs is 5.8x (for OpenACC) and 3.4x (for OpenMP)

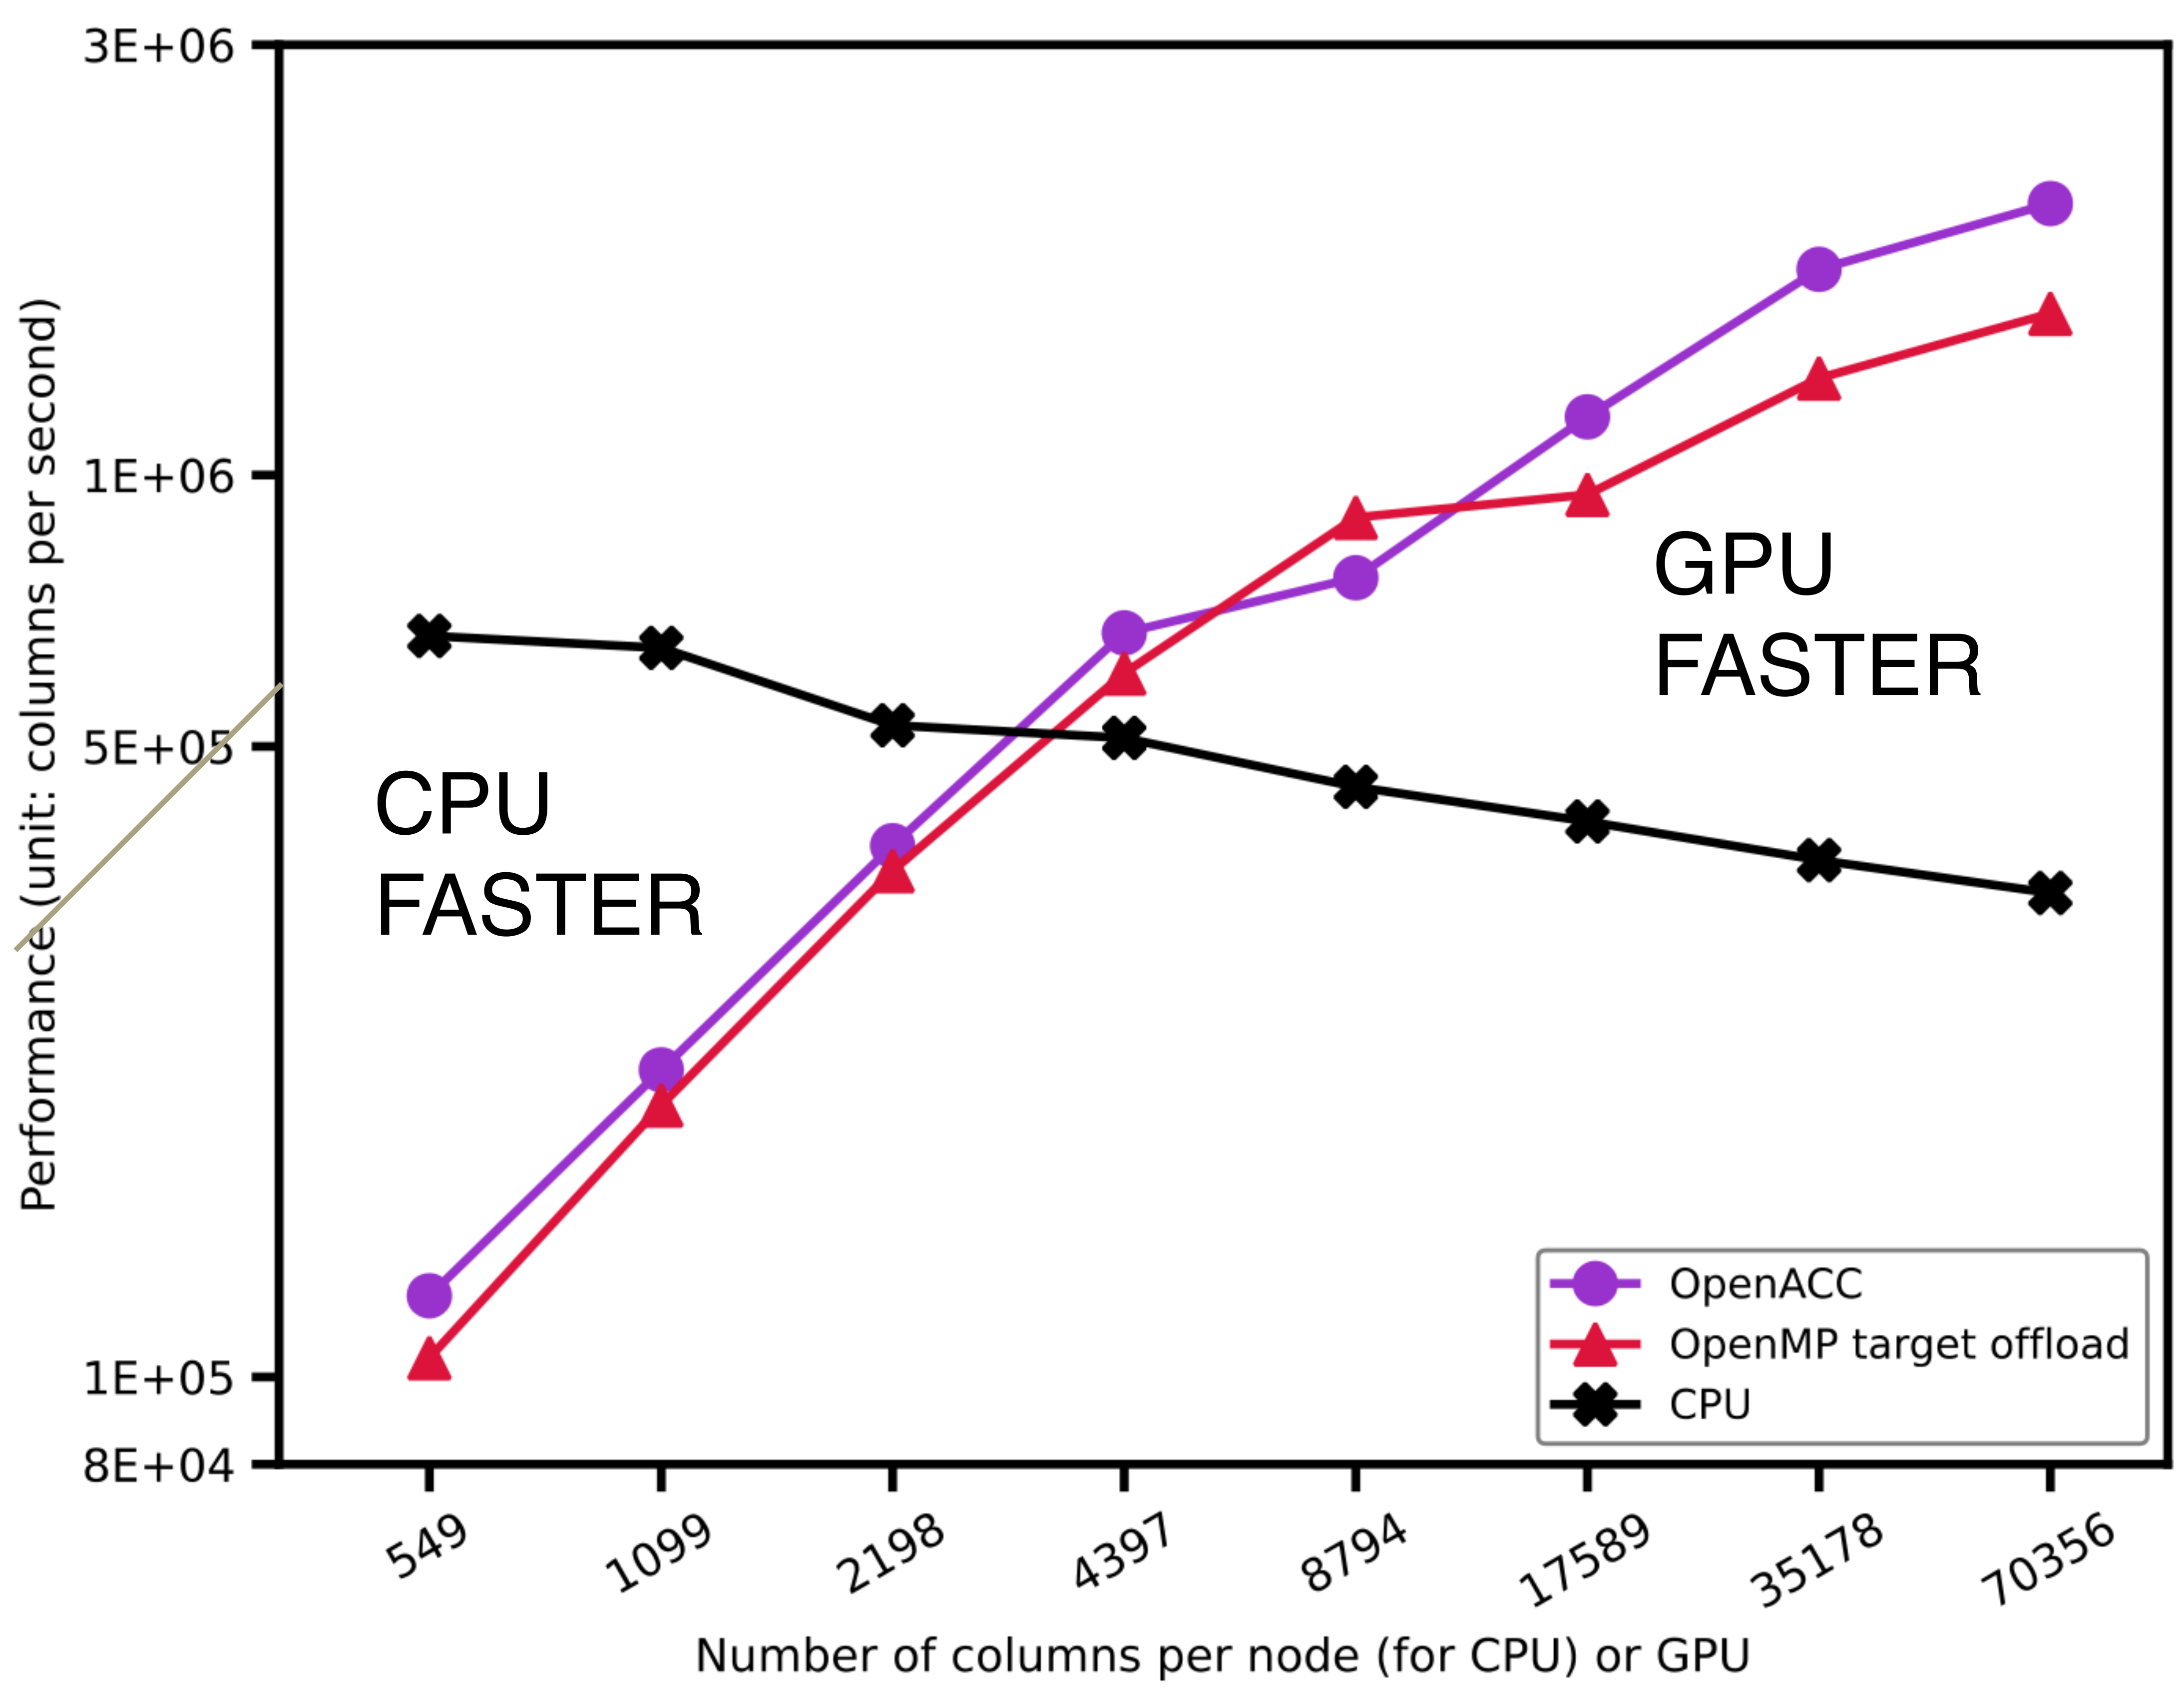

Results courtesy of Jian Sun, NCAR

PUMAS/MG3 Microphysics on GPUs

- **Goldilocks zone:**
  - **Blue:** too few columns per GPU, not enough parallel work to keep the GPU busy.
  - **Red:** too many columns, not enough memory on the GPU to hold the working data.
  - **Green:** Juuust right balance of parallel work and memory.
- **Example, the balance point is about:**
  - ~16 GPU's @ 30 km
  - ~64 GPU's @ 15 km
  - ~256 GPUs @ 7.5 km
  - ~1024 GPU's @ 3.75 km

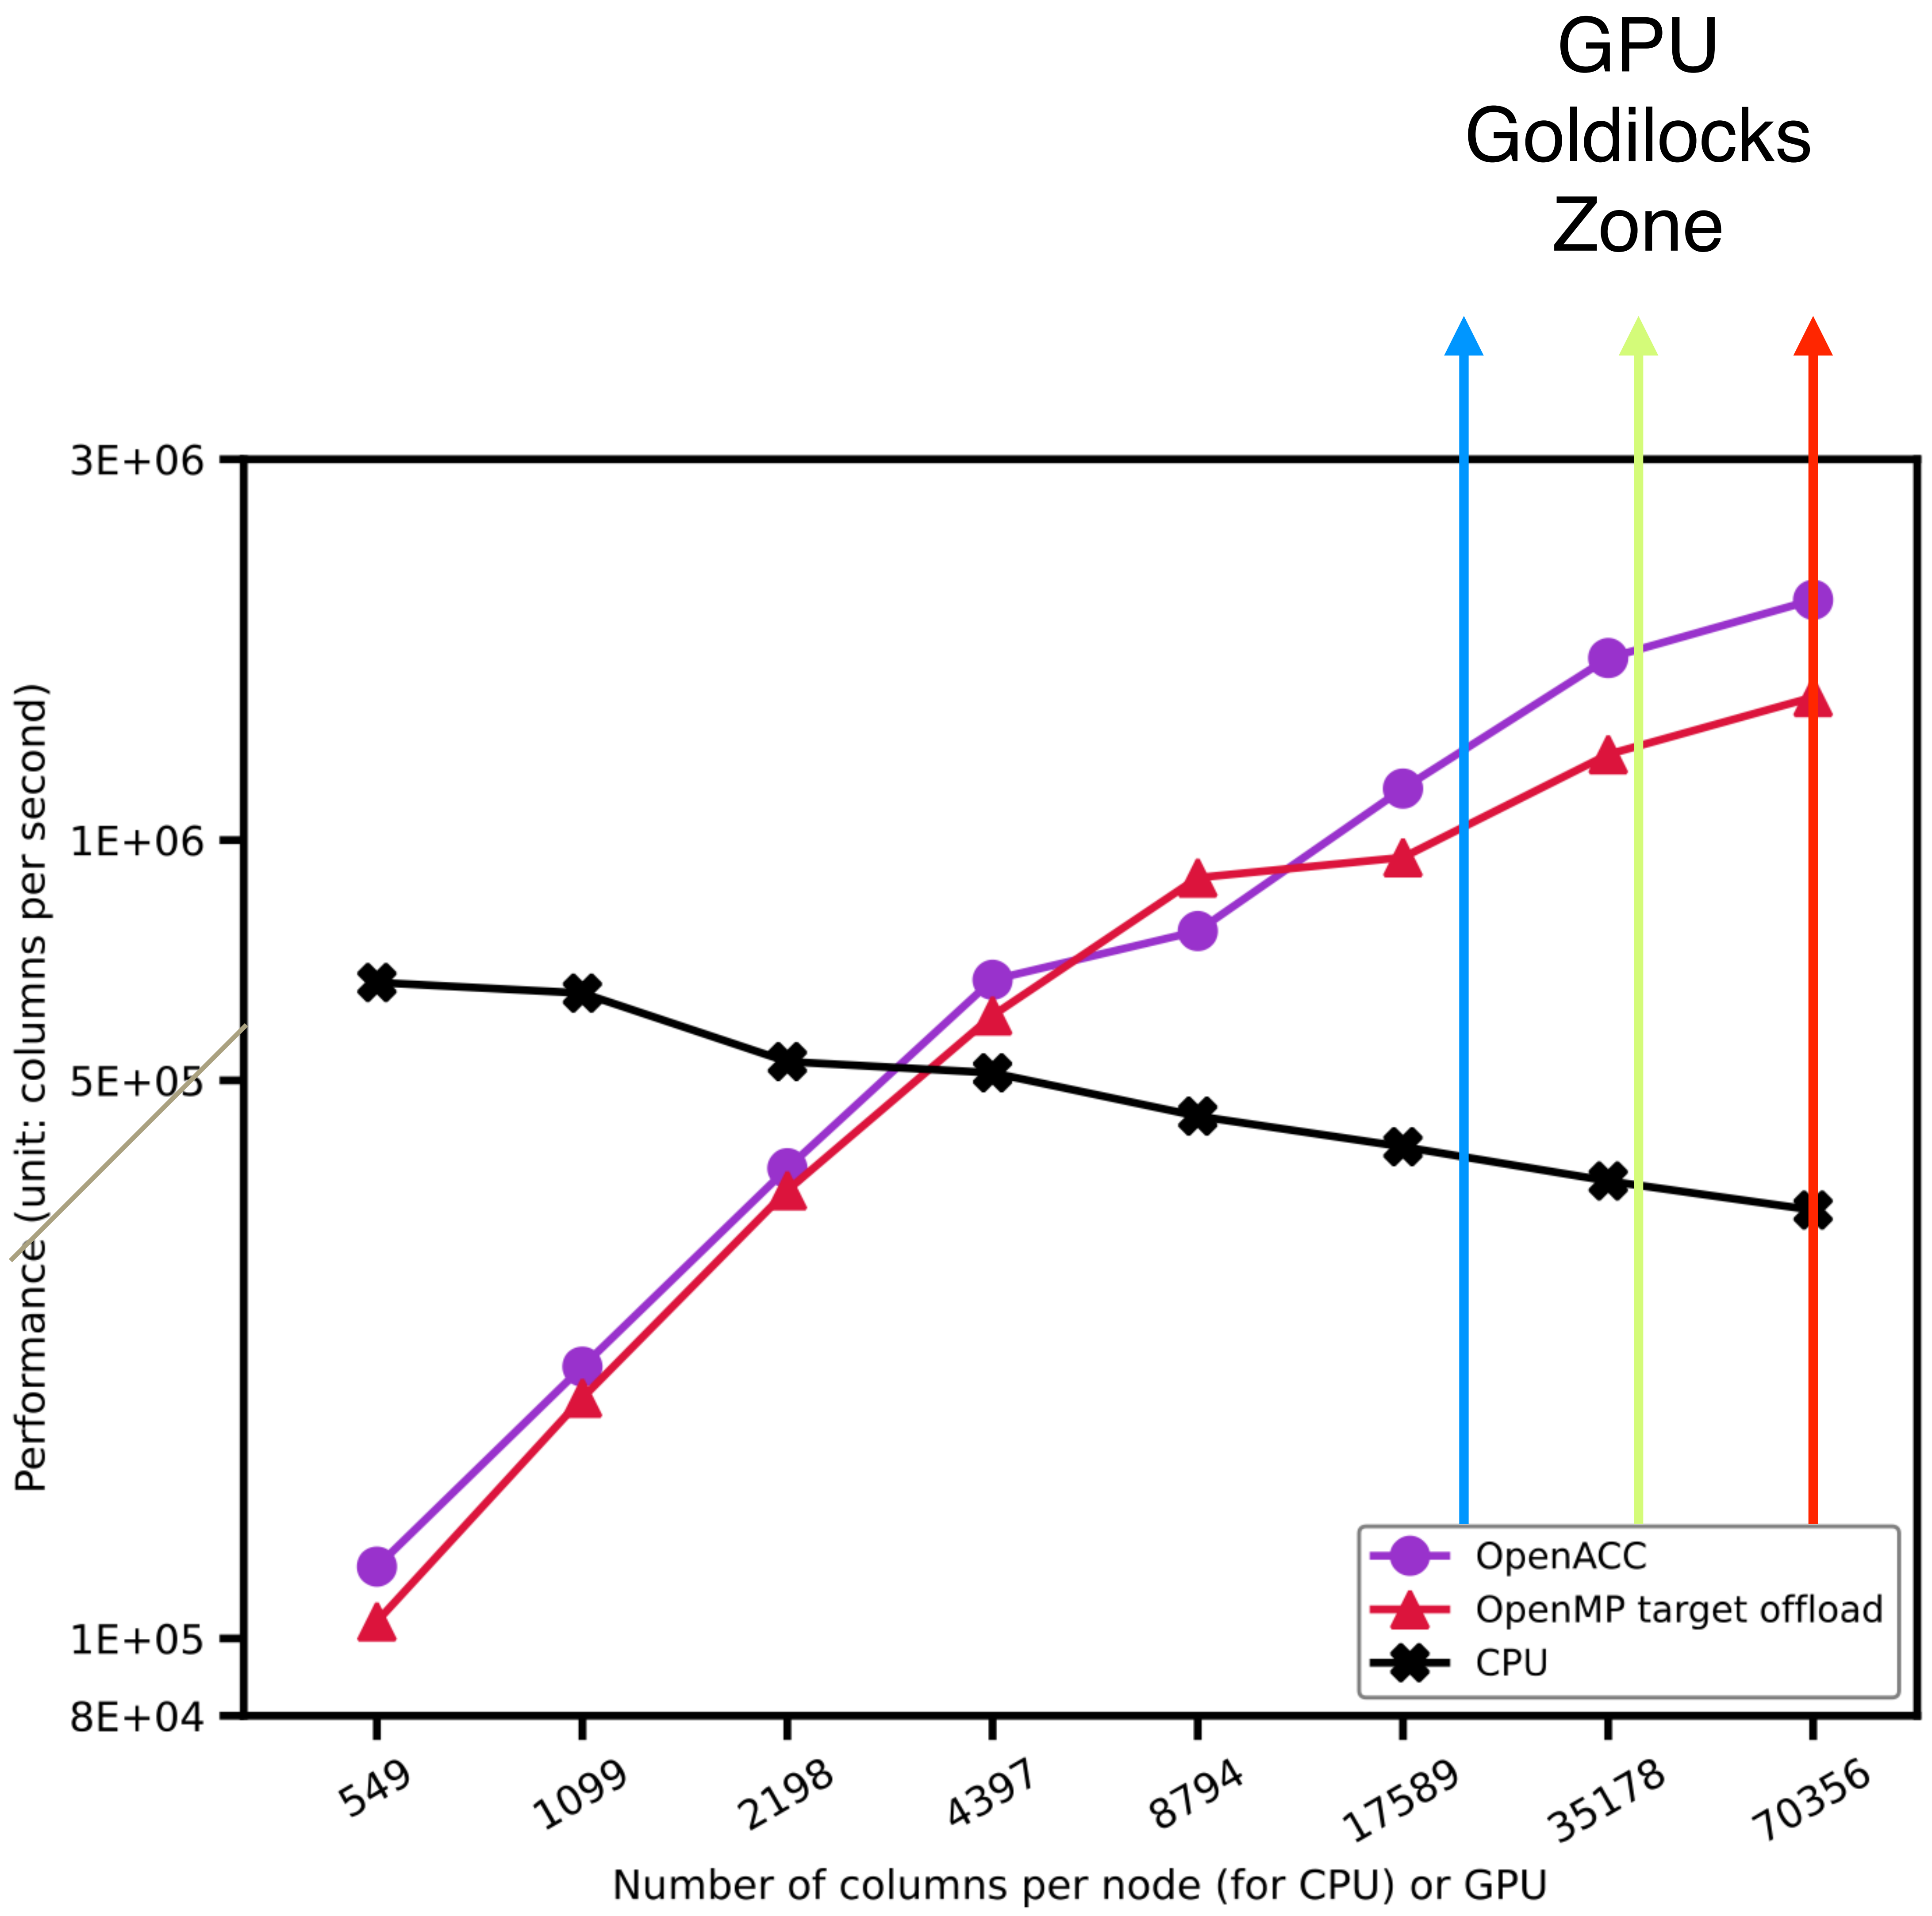

Results courtesy of Jian Sun, NCAR

Machine Learning the Warm Rain Process

Used for PUMAS microphysics in Earthworks

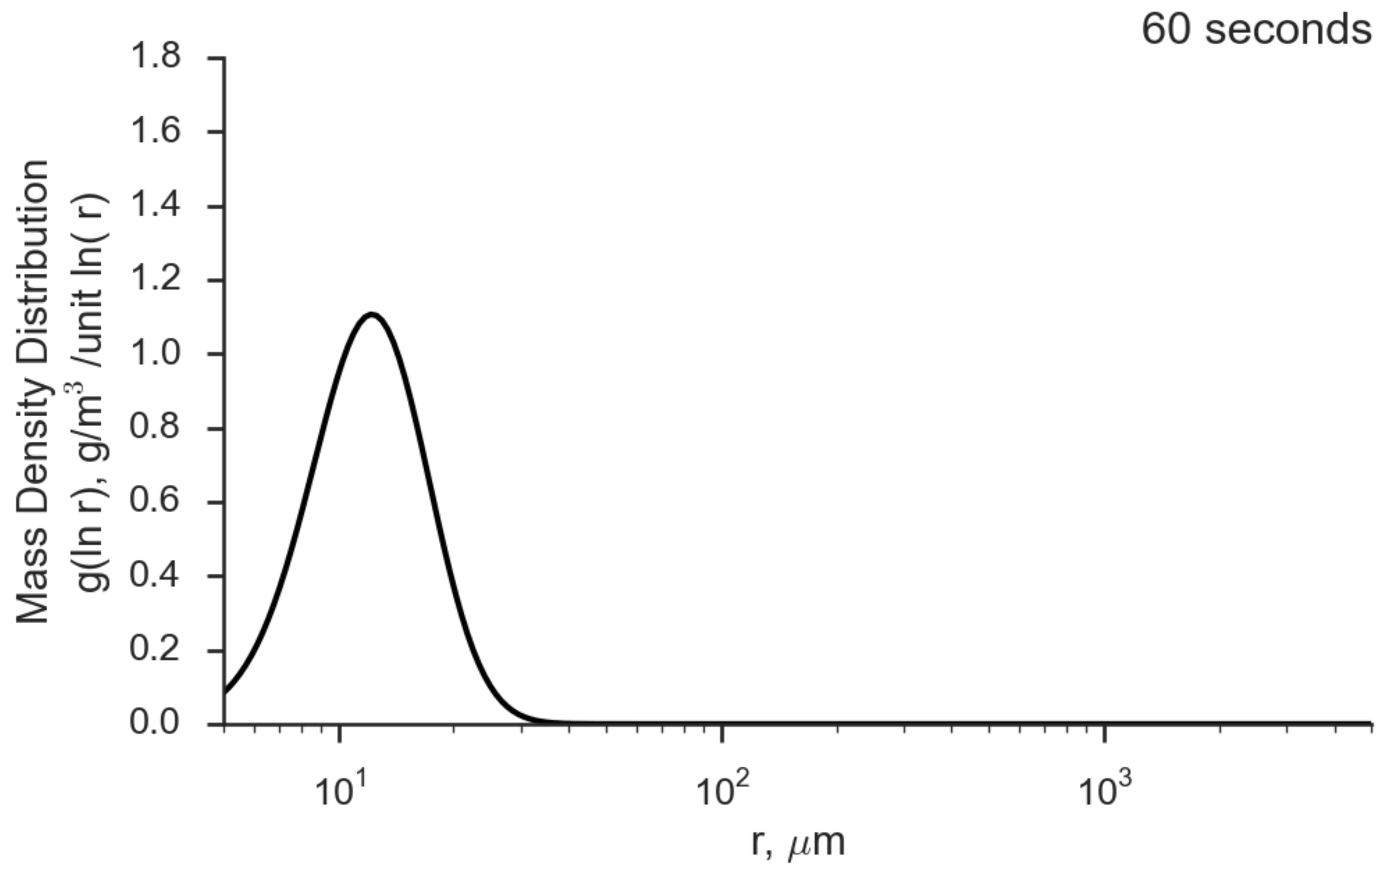

Replace traditional GCM bulk rain formation with a bin model formulation for stochastic collection (left). This is much too expensive for climate use, so we will emulate with a Neural Network.

NN Emulator reproduces detailed bin code

Emulator Performance

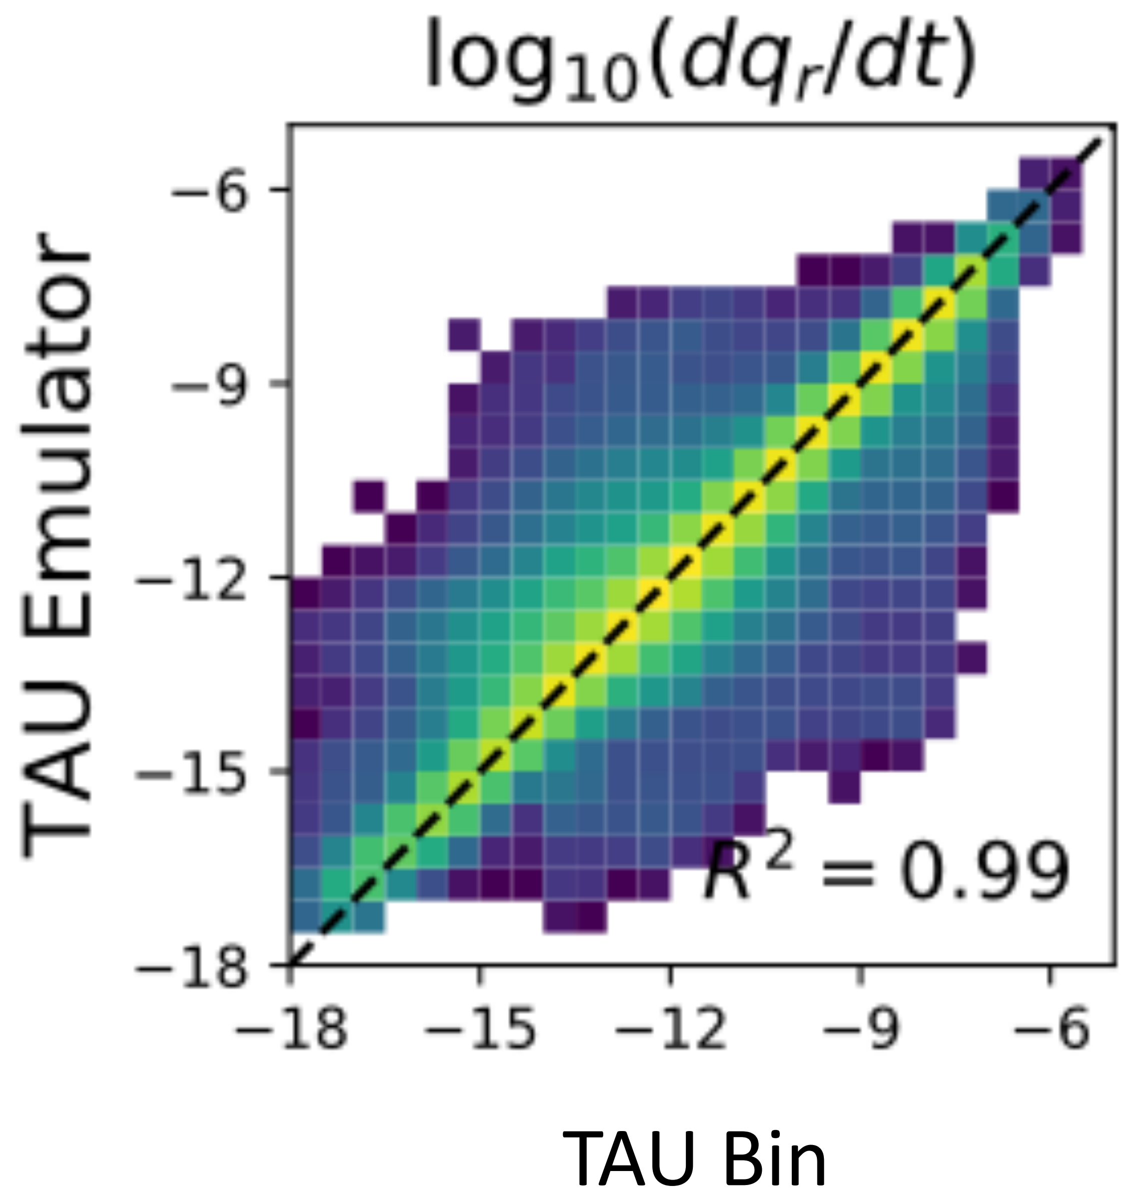

Large-scale precipitation frequency over the ocean

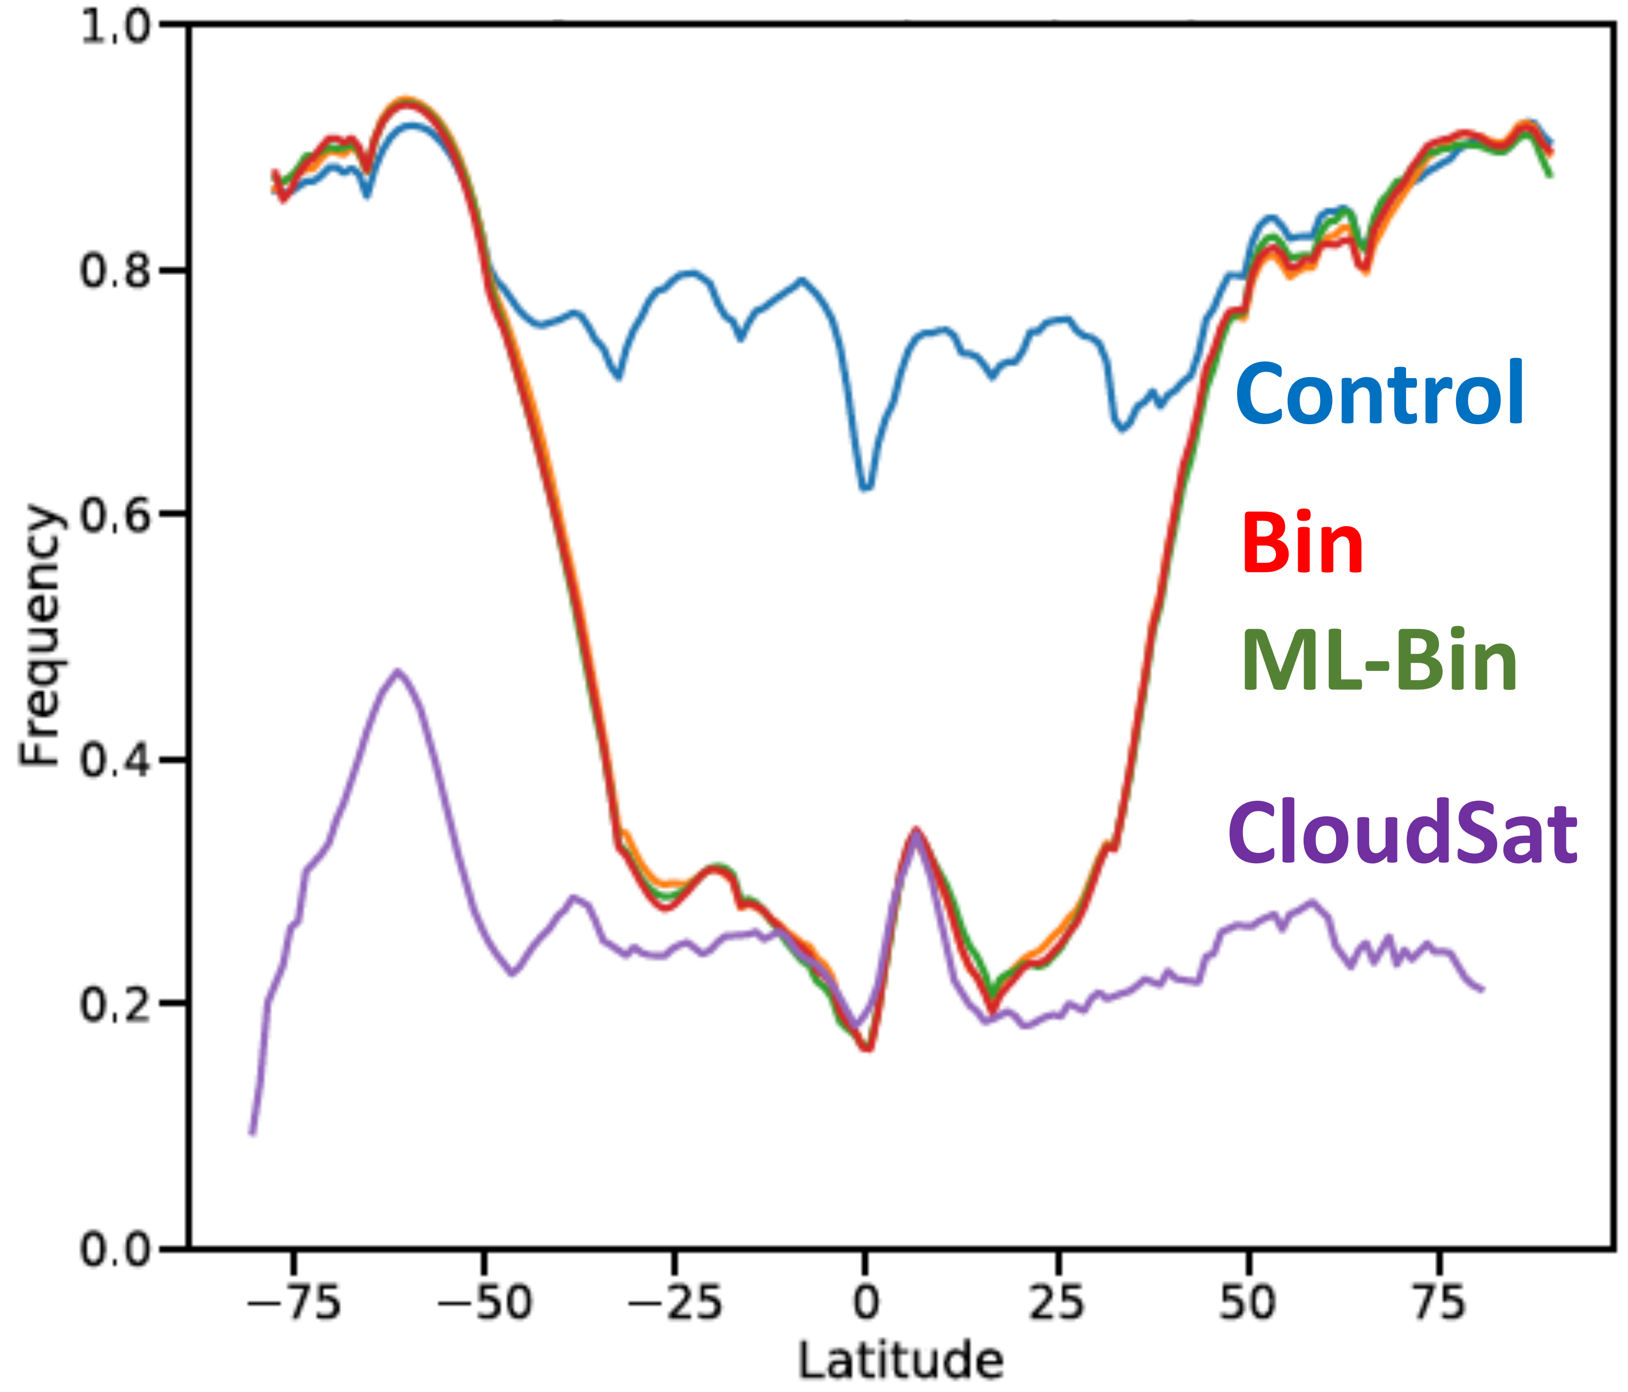

Both Bin and ML-Bin reduce rain frequency (frequent 'drizzle' problem in Control) in sub-tropics to match observations

Gettelman et al. 2021, JAMES

Machine Learning the Warm Rain Process

Used for PUMAS microphysics in Earthworks

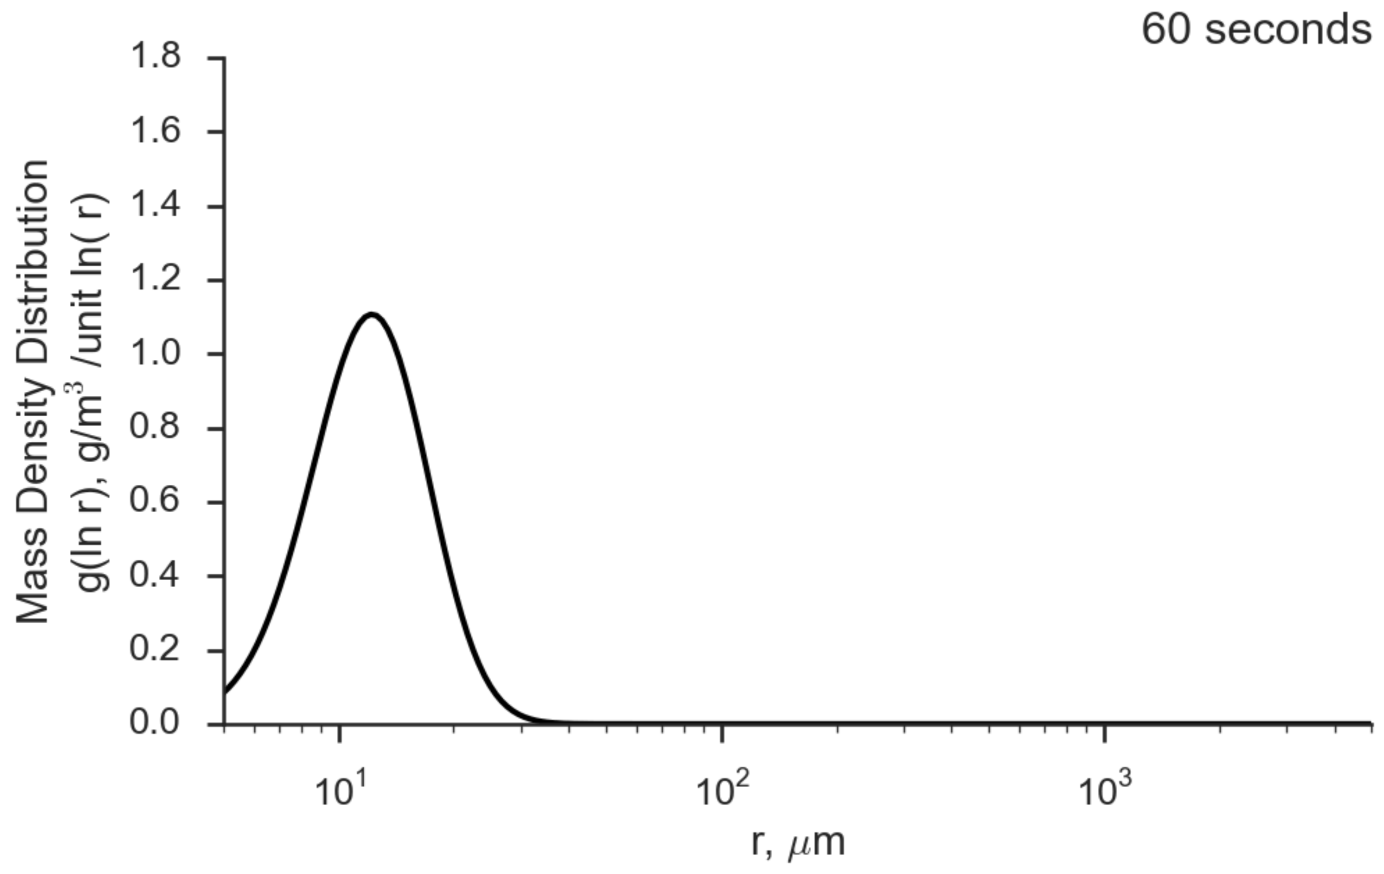

Replace traditional GCM bulk rain formation with a bin model formulation for stochastic collection (left). This is much too expensive for climate use, so we will emulate with a Neural Network.

NN Emulator reproduces detailed bin code

Emulator Performance

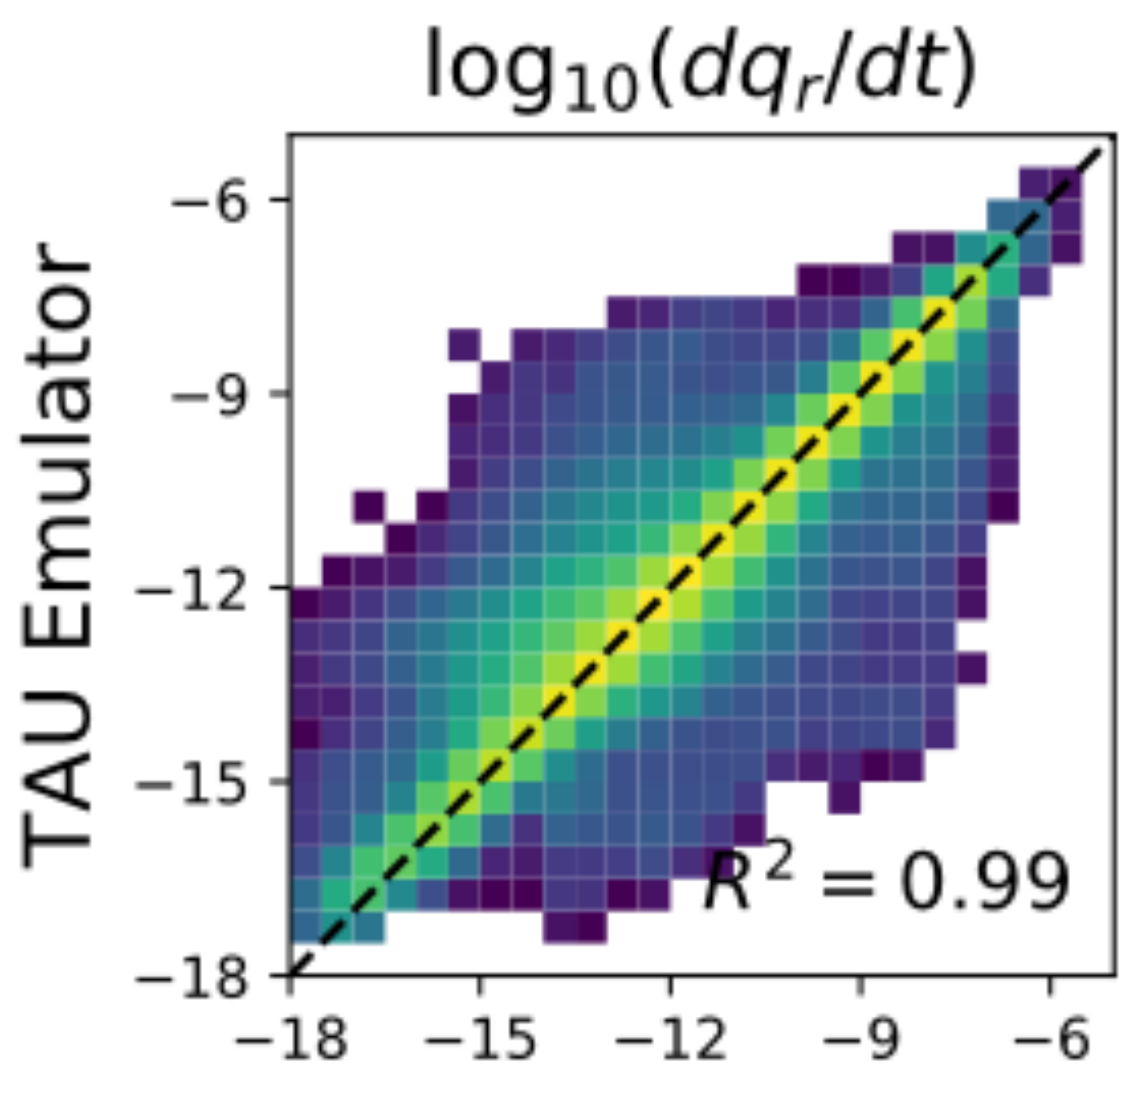

Large-scale precipitation frequency over the ocean

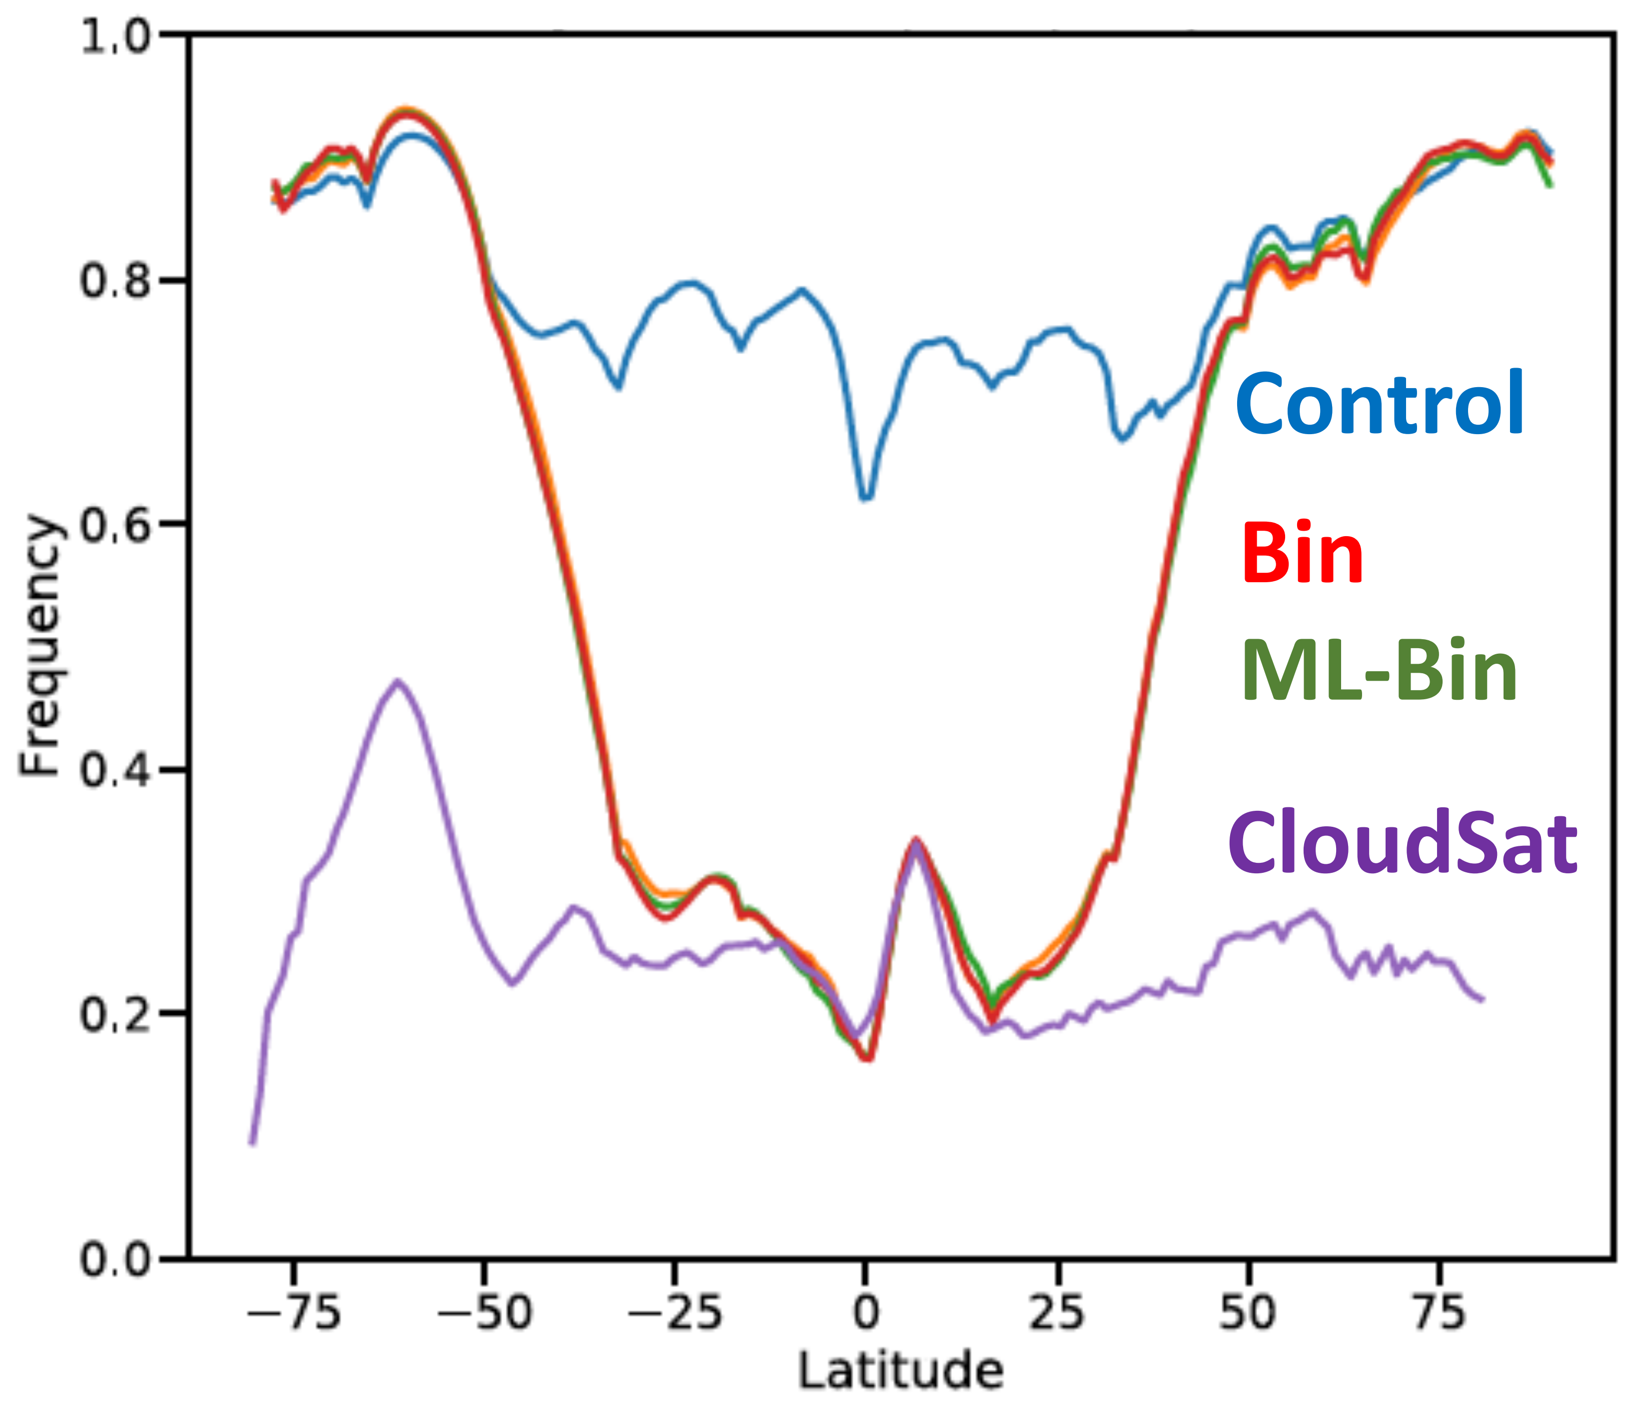

Both Bin and ML-Bin reduce rain frequency (frequent 'drizzle' problem in Control) in sub-tropics to match observations

Gettelman et al. 2021, JAMES

Thinking about EarthWorks phase 2

CSU and NCAR are currently discussing a possible follow-on proposal (start date 2025).

Backup Slides

Our approach to porting to GPUs: Directive-based offload

- We currently offload to GPUs using OpenACC directives.
- We see the best performance with OpenACC
- With a large Fortran code base, porting to a new language is not currently an option.
- Our plan includes a “cautious” pivot to OpenMP offload directives.
- We have had reasonable success with Intel’s OpenACC to OpenMP offload conversion tool.

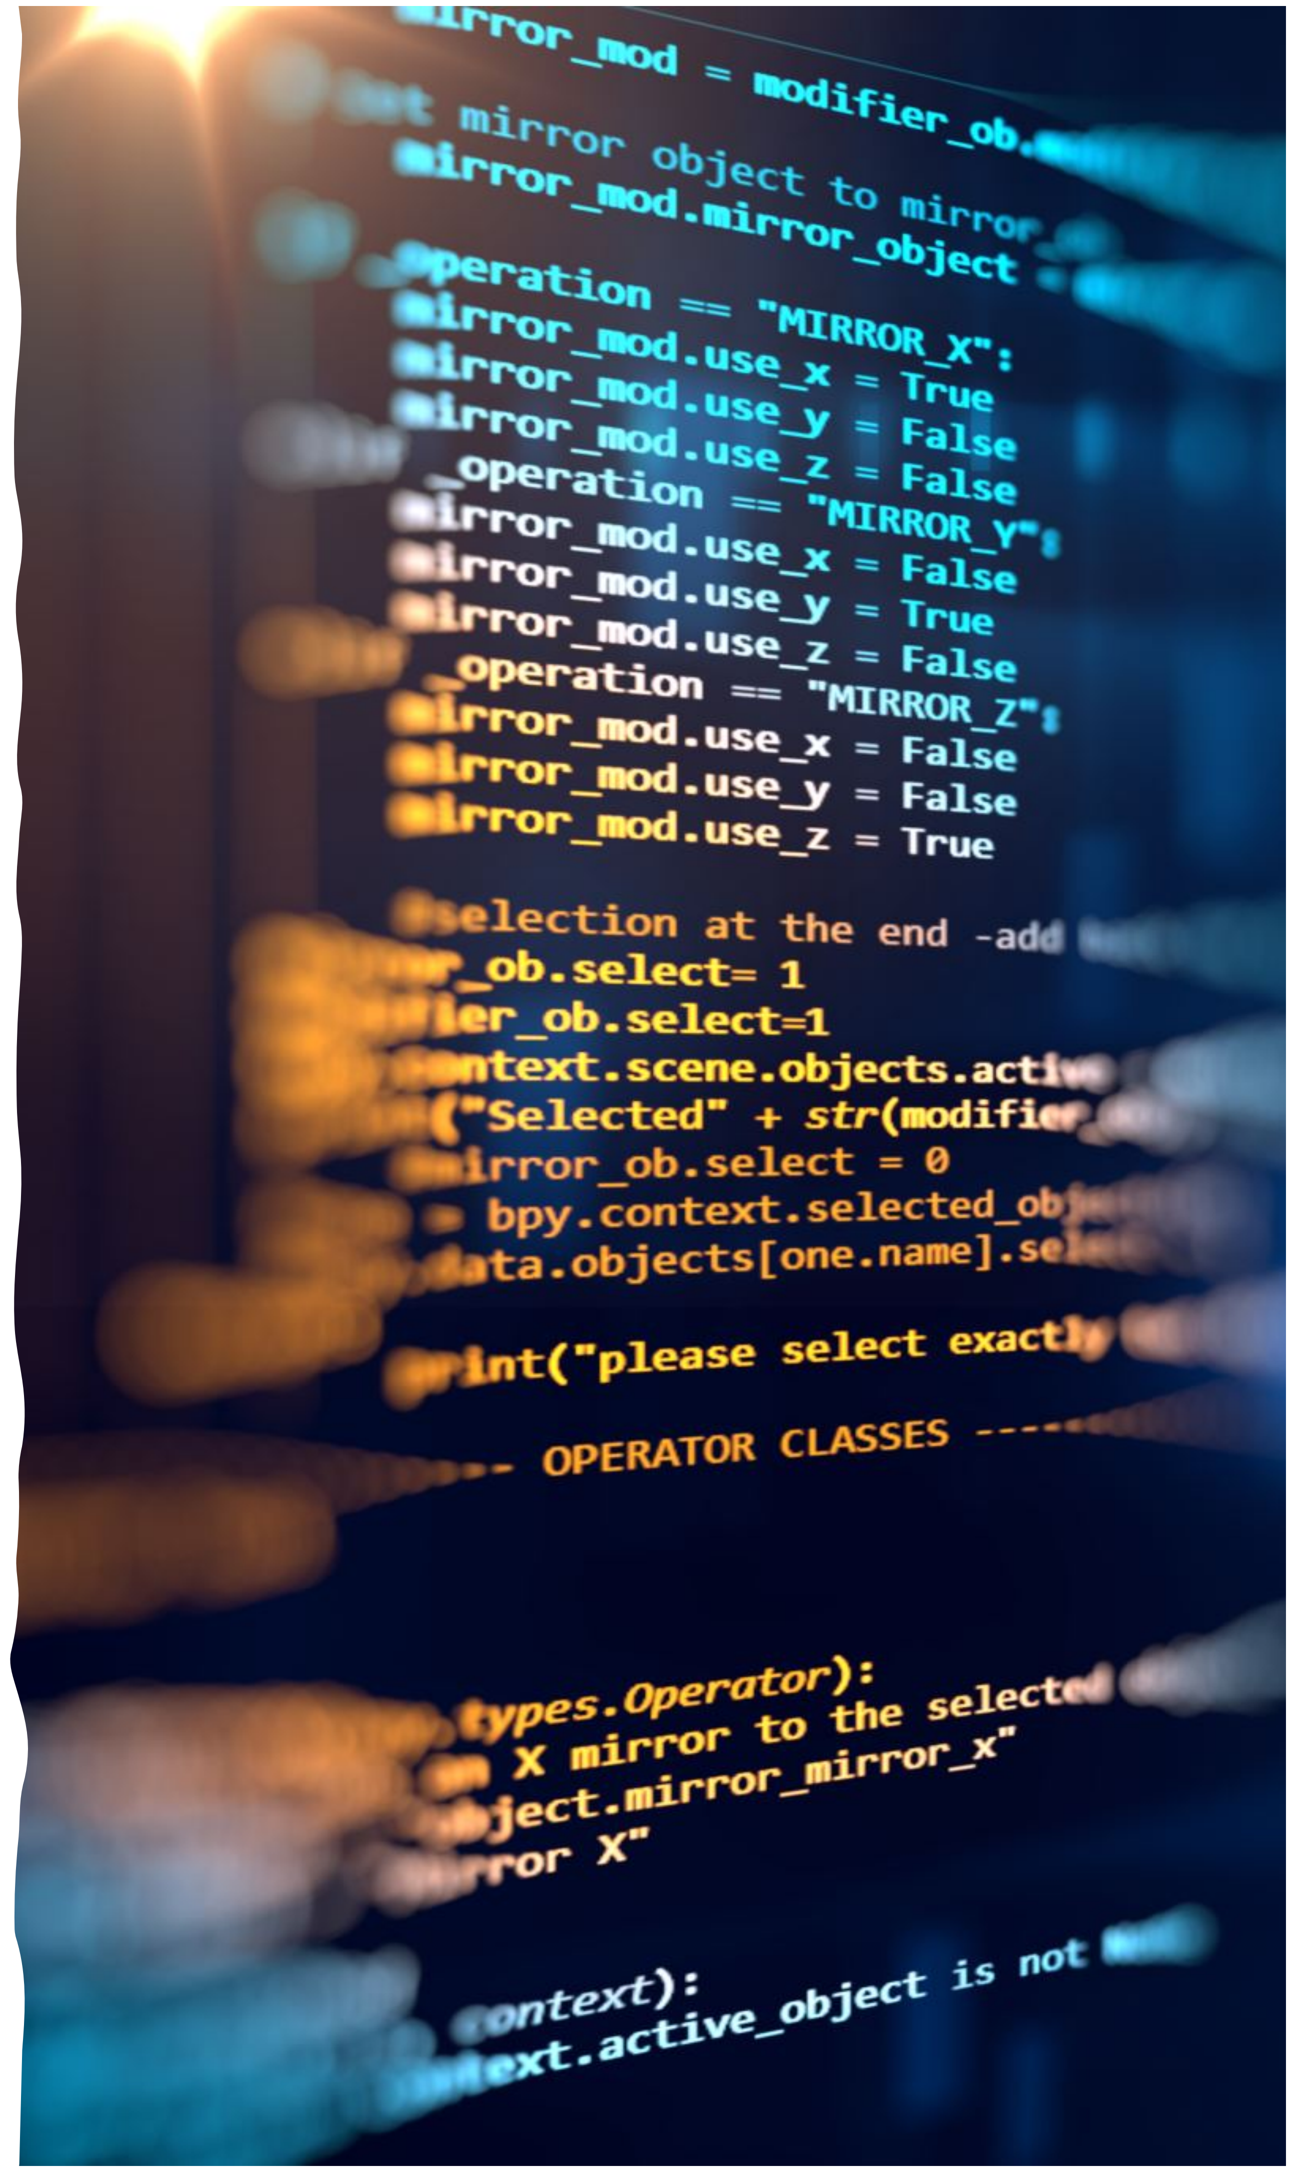

CLUBB

- CLUBB has been fully ported onto GPUs using OpenACC.
- There is also a version that uses OpenMP offload.
- We are currently working on incorporating CLUBB-GPU into CAM

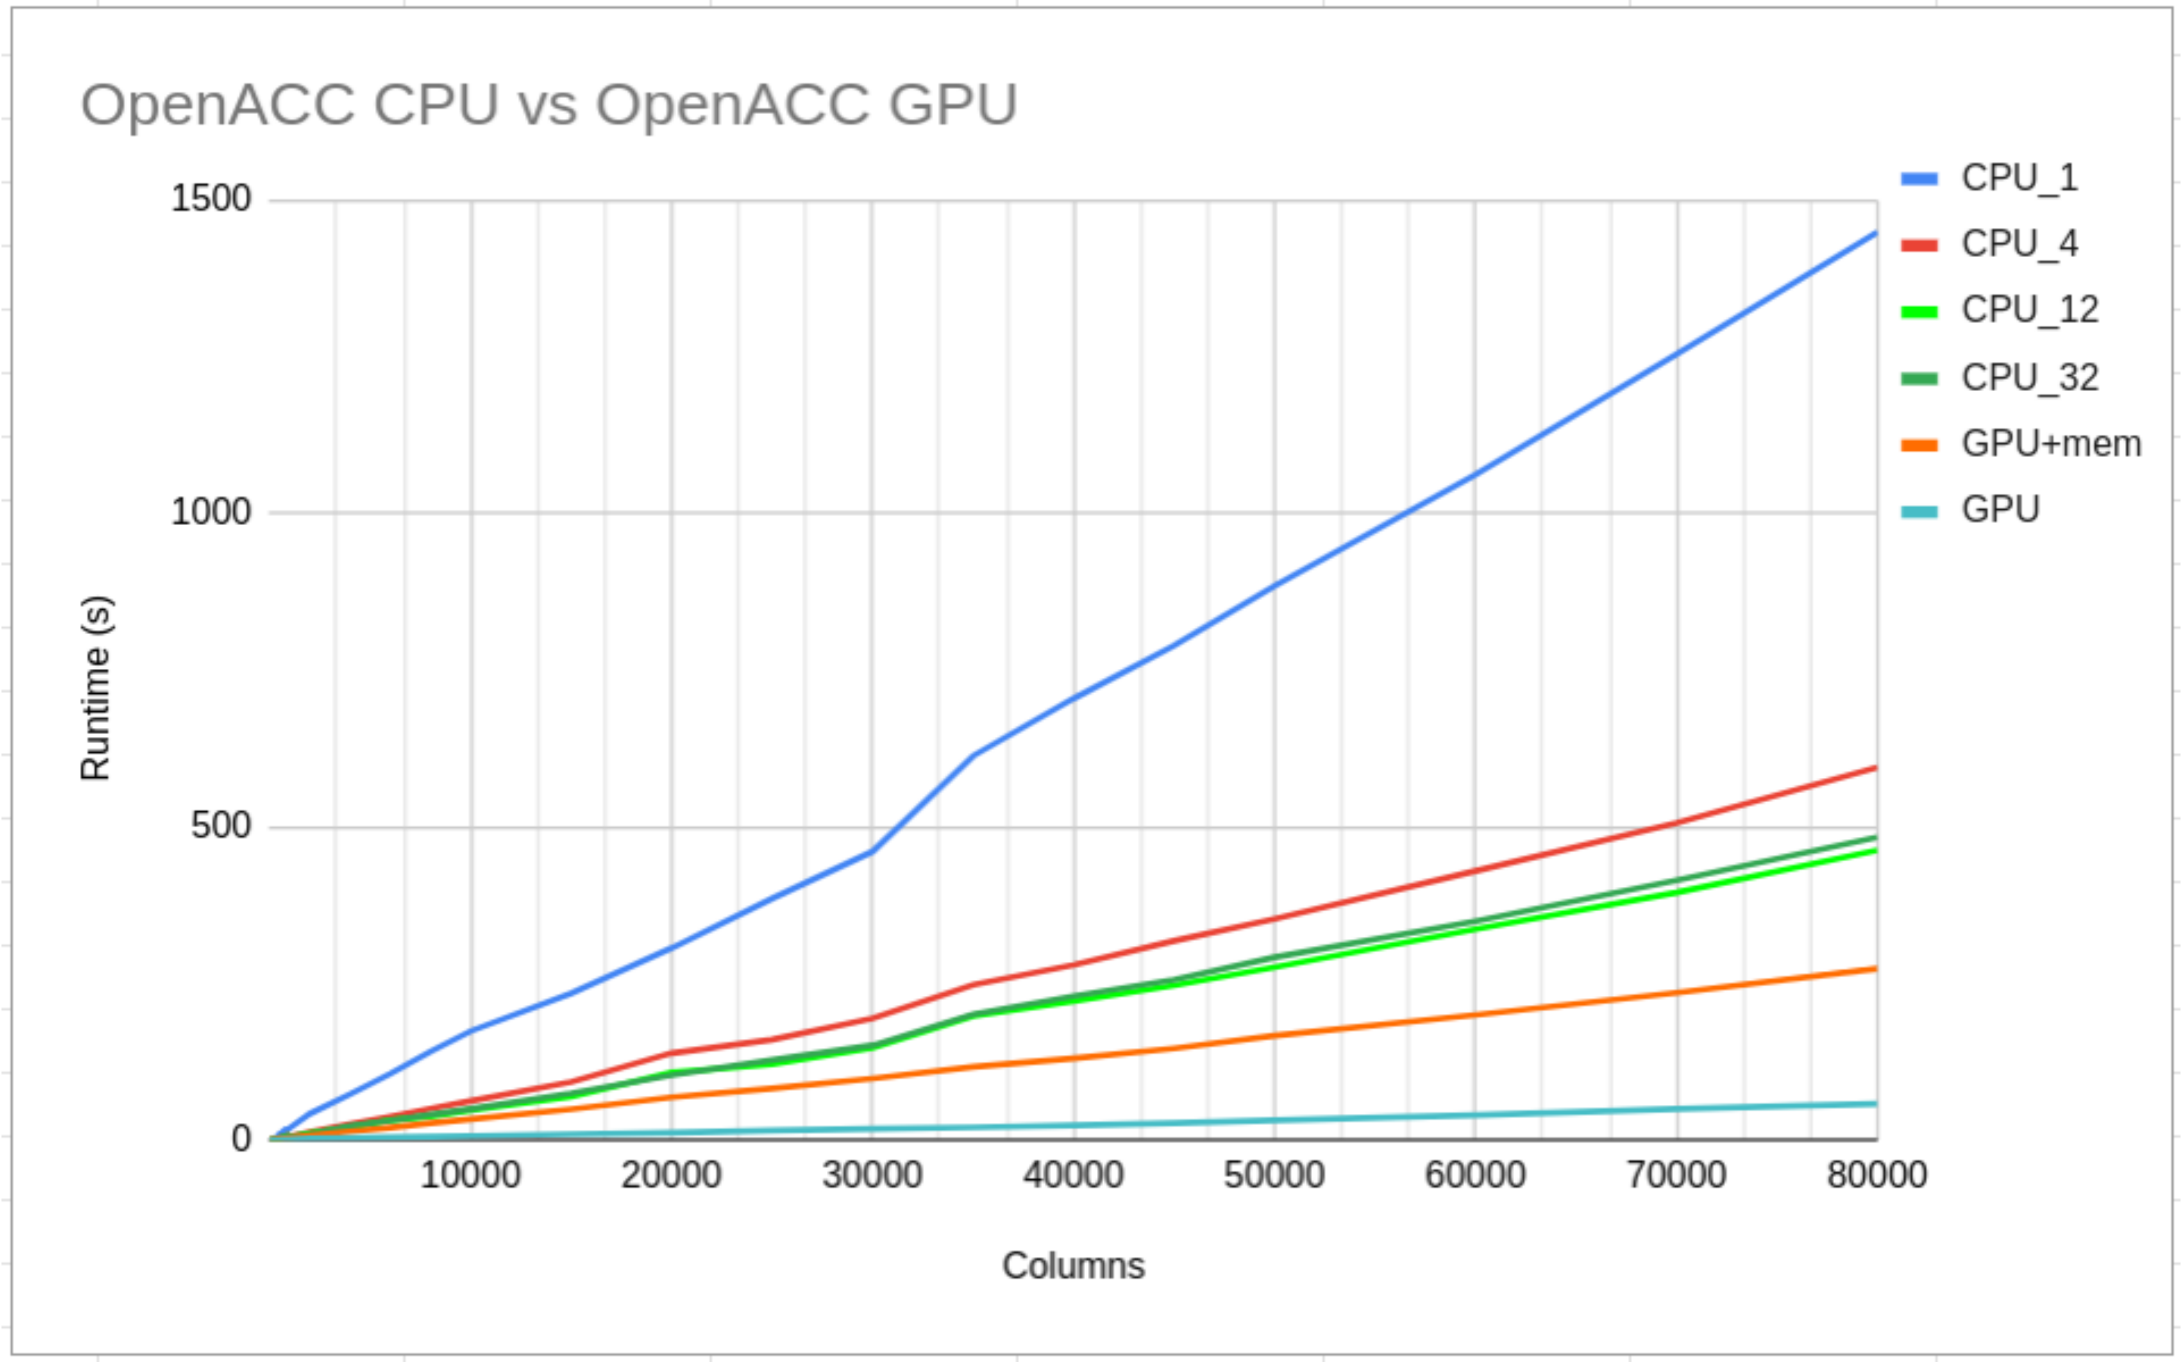

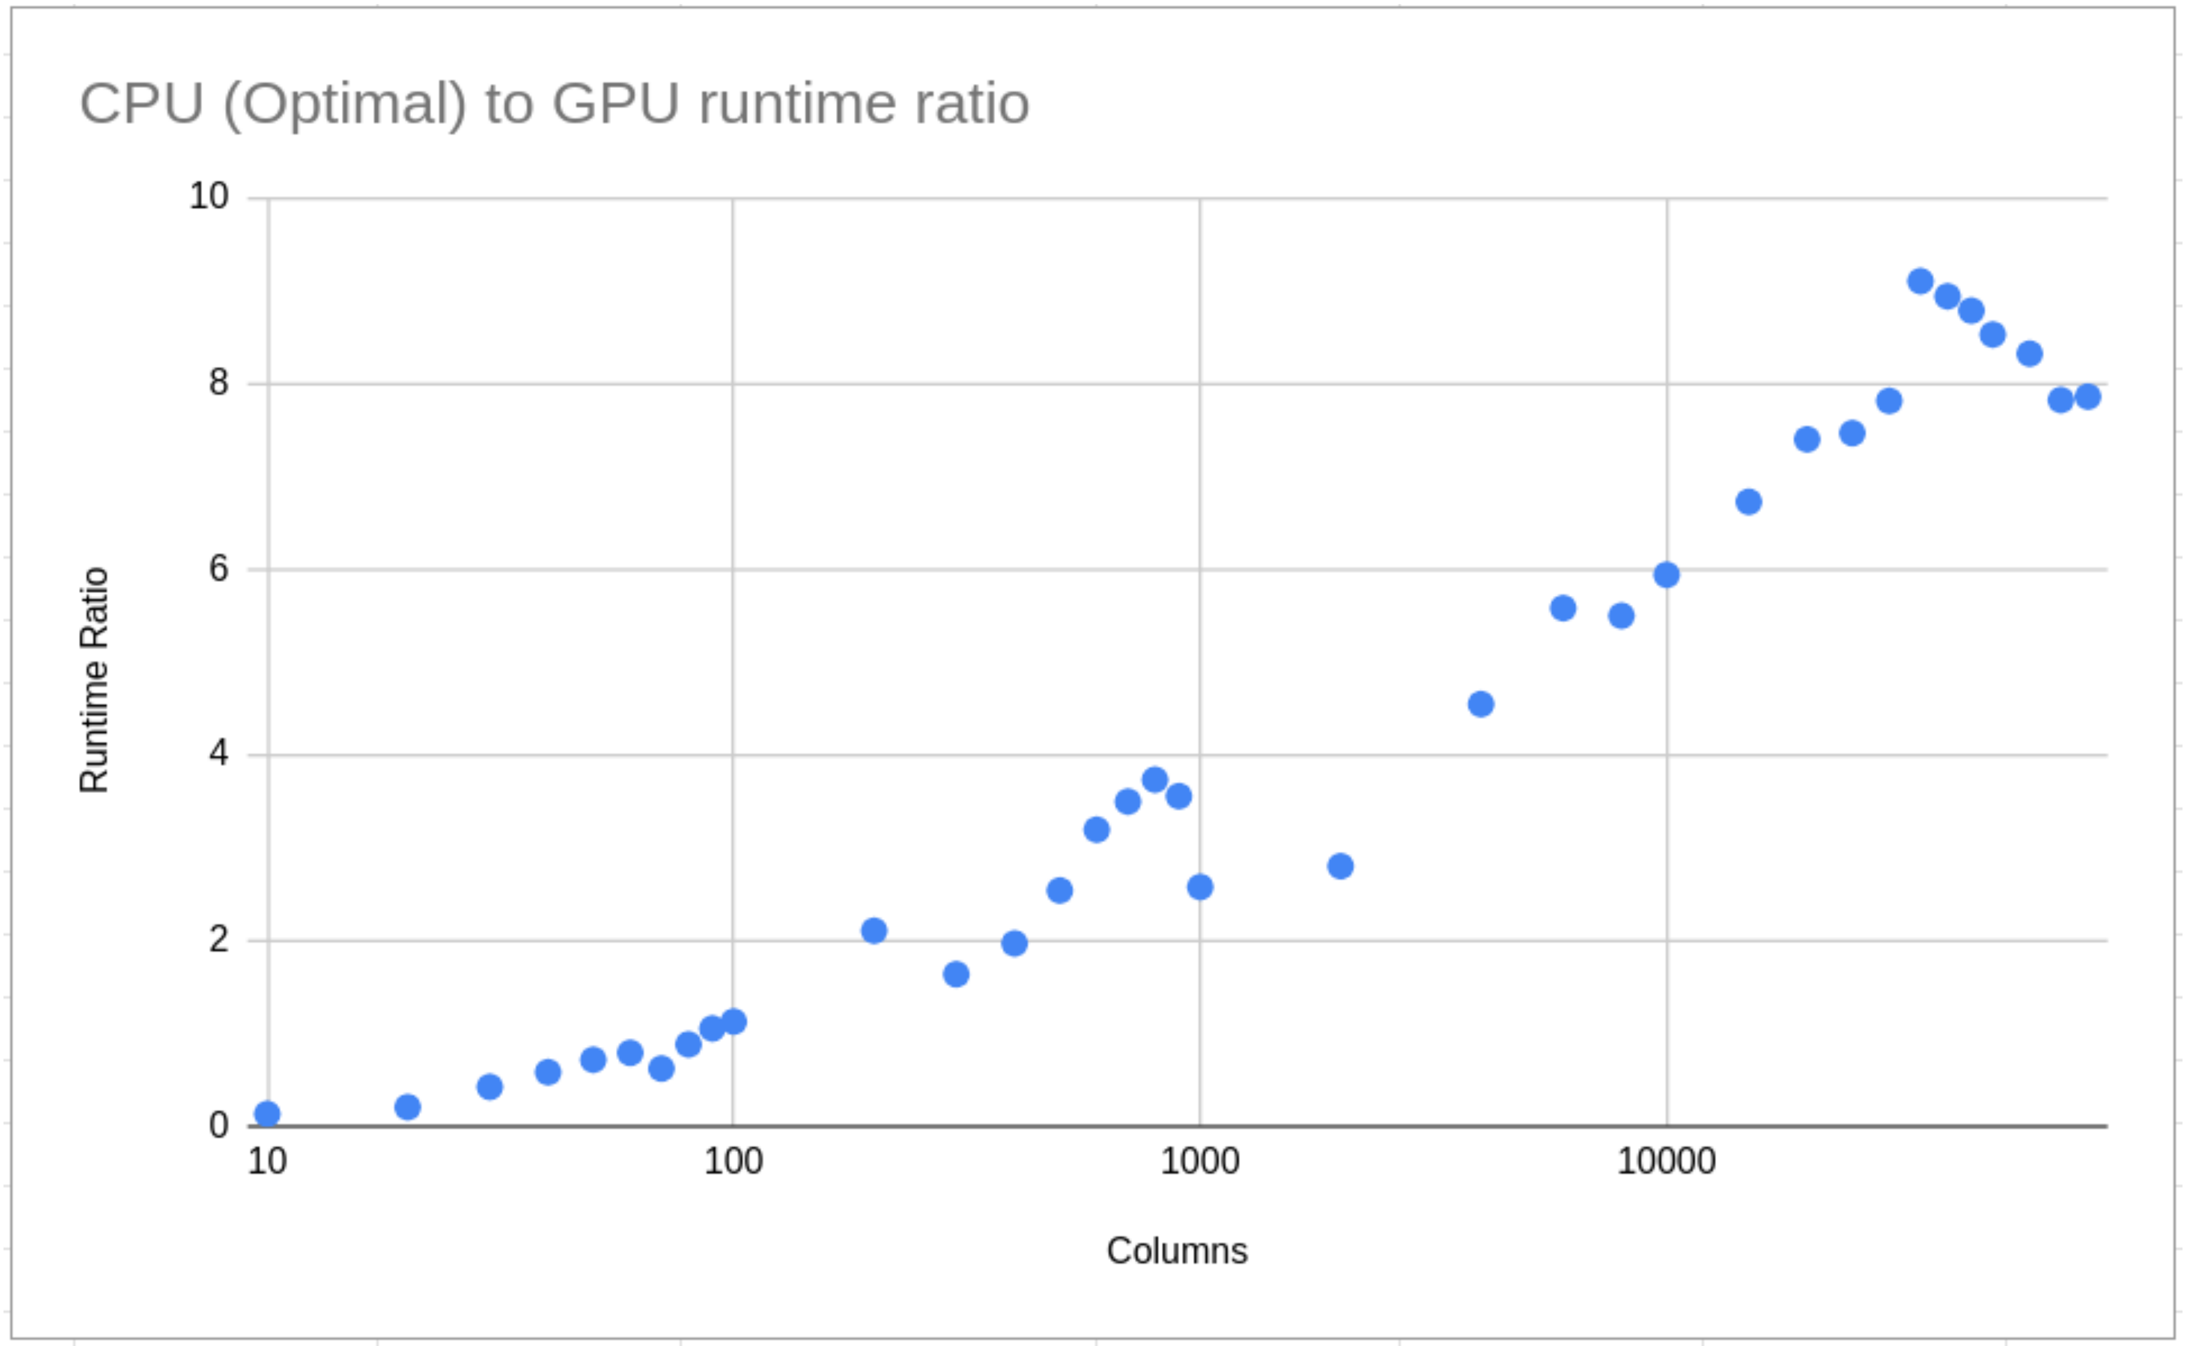

*Comparisons were run one EPYC 7763(AMD Milan) and one A100, comparisons done in standalone CLUBB

Results courtesy of Gunther Huebler, University of Wisconsin, Milwaukee

MPAS dycore performance

- **Experiment**: MPAS-7 (5.9M cell mesh; 56 levels; FP32) ran dry baroclinic test case for 10 simulated days
- **Equipment**: Selene supercomputer; nodes = AMD Dual socket EPYC 7742 "Rome" CPUs with 8x NVIDIA A100 GPUs; 10 HDR links/node.
- **Resources**: Benchmark of 128-core ROME CPU node vs A100 GPU
- **Takeaways**:
- Early scaling looks impressive - and 3.5x faster than CPU node.
- Slowdown of MPAS-7 compute (m) was recently isolated to not declaring new variables GPU resident.

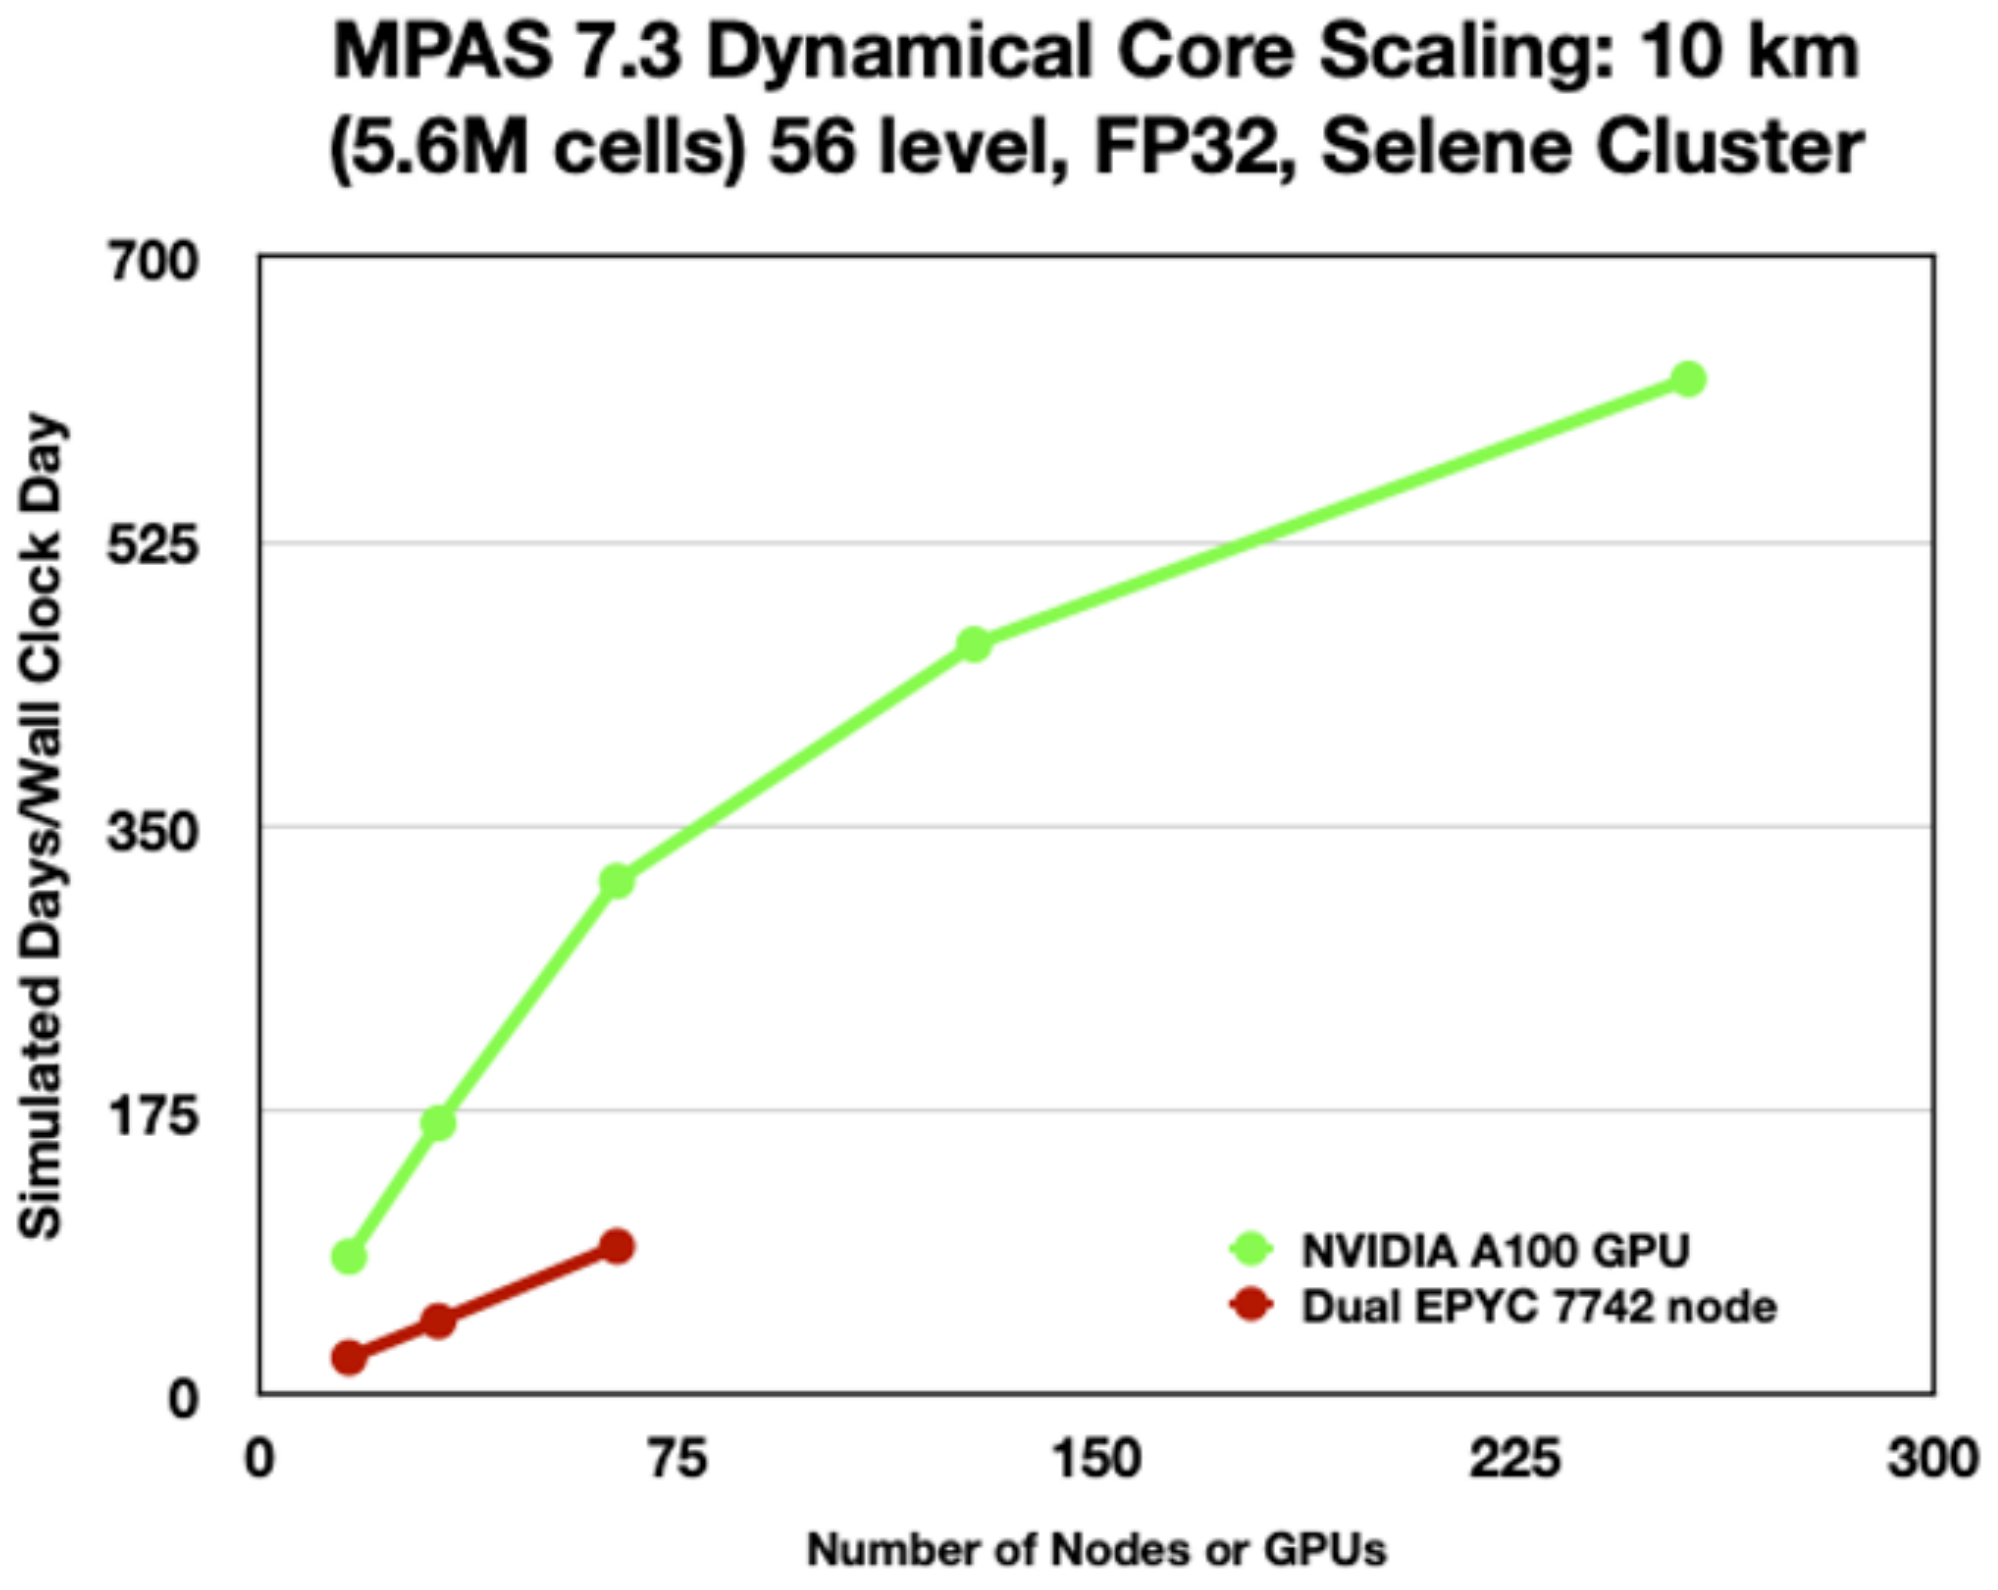

Results courtesy of Raghu Raj Kumar of NVIDIA

RRTMGP

- Utilizing code from Robert Pincus, et al
  - https://github.com/earth-system-radiation/rte-rrtmgp
- We have incorporated RRTMGP into the latest CAM development version
  - It took some time to debug
  - Verified that answers match between CPU and GPU
- Preliminary results in standalone tests show about a 10x speedup on GPUs (without data transfers).