Progress toward an Earth System Model with Machine Learning Components

Dr. Richard Loft
loft@ucar.edu
Director, Technology Development
Computational & Information Systems Laboratory
National Center for Atmospheric Research
Boulder, CO

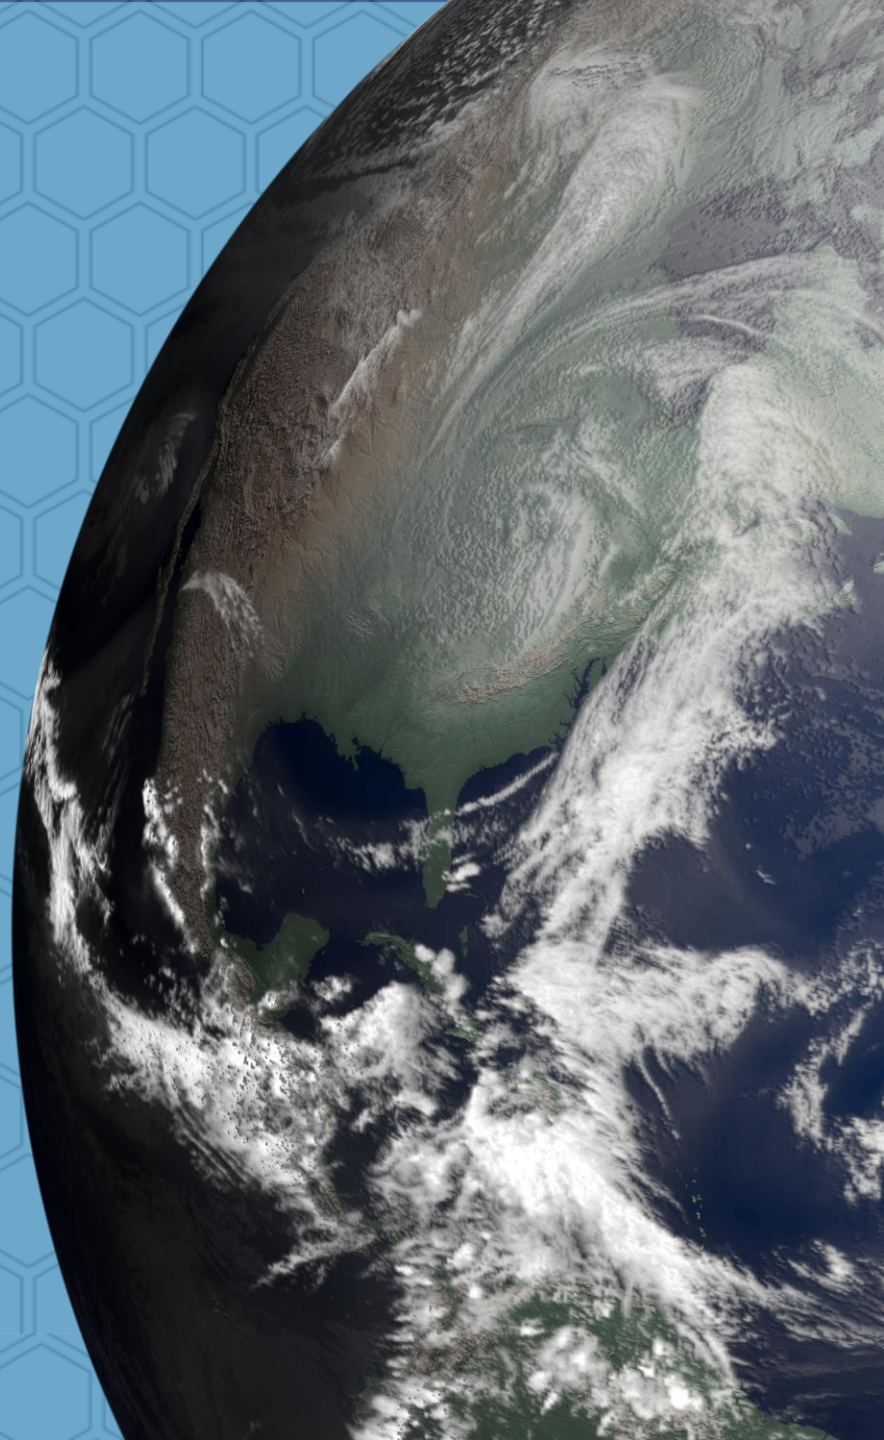

This material is based upon work supported, in part, by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.

Takeaways

- MPAS-A/GPU is the first operational GPU-accelerated *global weather forecast* model*.
- We have conducted scaling studies on the MPAS-A/GPU dynamical core with tracers on *700 AC922 nodes on Summit* (4200 V100's!)
- The dycore has achieved 6.5 days/hour at 3 km resolution. This is still slow for climate studies, but close to what's needed for seasonal forecasts.
- We have experimented with *overlapping calculations on heterogeneous nodes*.
- We've run a *machine learning component* in our climate model for nine simulated years with good results.

* MeteoSwiss has been running COSMO, a regional GPU model, in production for quite some time.

Global High-Resolution Atmospheric Forecasting System (GRAF)

Partnership between NCAR, IBM and the Weather Company

- NCAR's MPAS-A global weather model capable of regional refinement
- Untapped data sources
- GPU-based MPAS forecast model

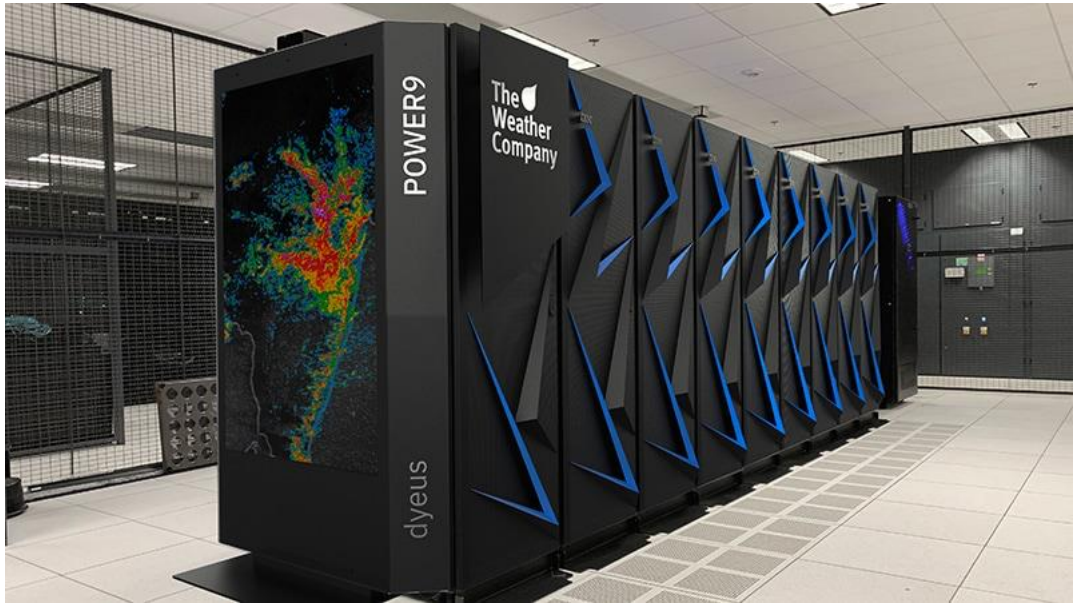

This talk will present details of the GPU acceleration of the MPAS-A model….

…and provide thoughts about building an Exascale Earth System Model

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

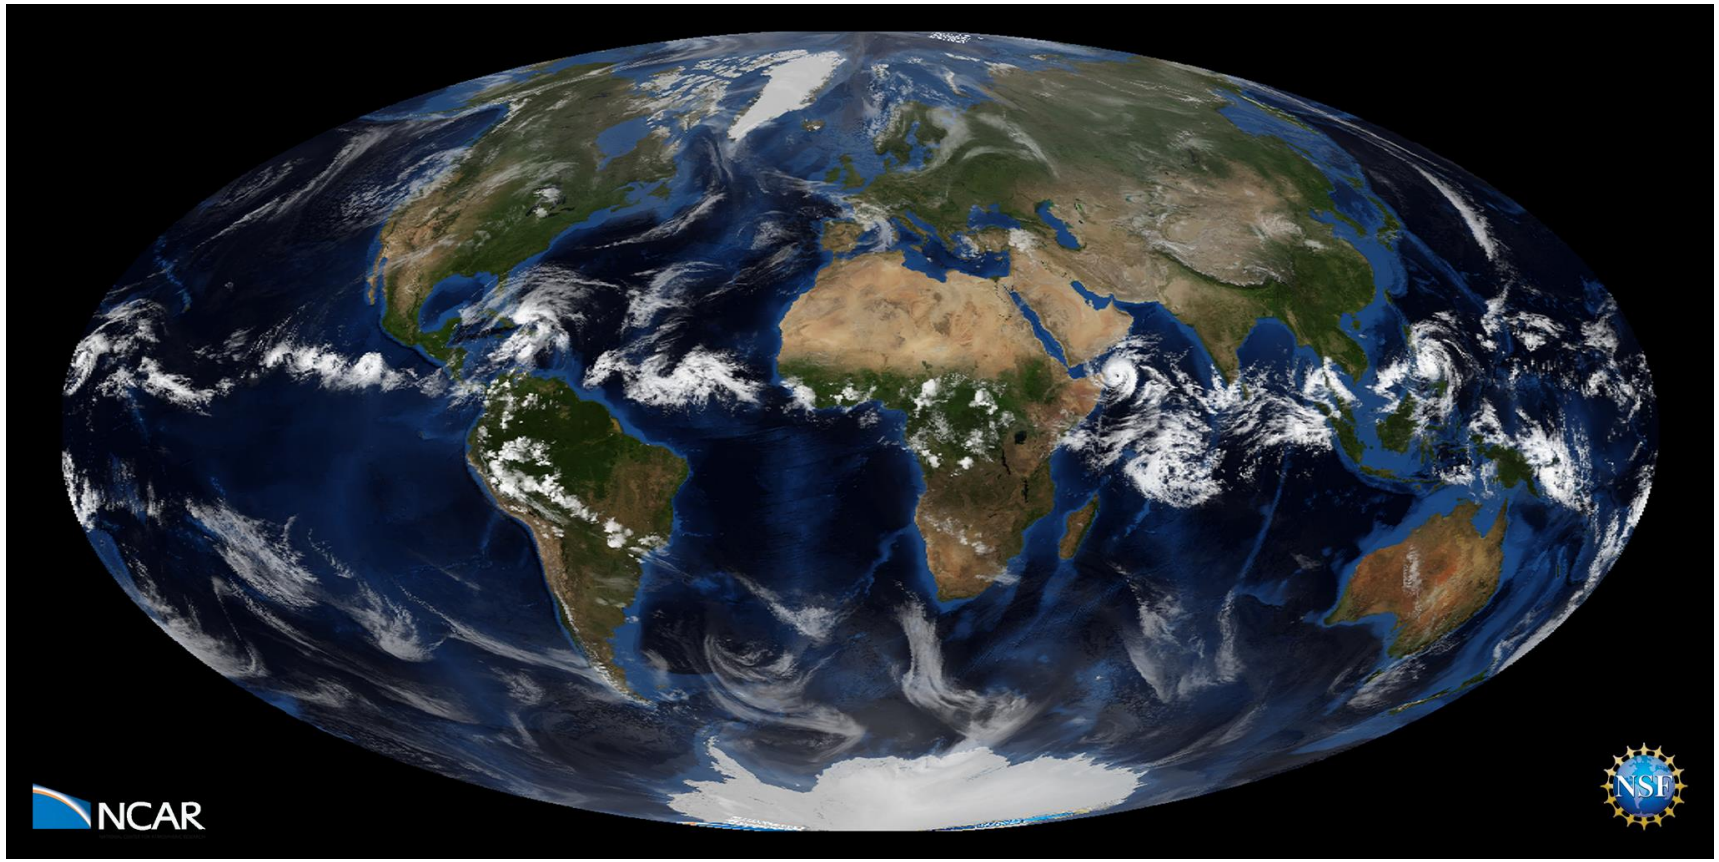

Simulation of 2012 Tropical Cyclones at 4 km resolution
– Courtesy of Falko Judt, NCAR

MPAS-A: the Algorithm

- **Fully compressible non-hydrostatic equations written in flux form**
- **Finite Volume Method on staggered grid**
  - The horizontal momentum normal to the cell edge (u) is sits at the **cell edges**.
  - Scalars sit at the **cell centers**
- **Split-Explicit timestepping scheme**
  - Time integration 3rd order Runge-Kutta
  - Fast horizontal waves are sub-cycled

MPAS is based on unstructured centroidal Voronoi (hexagonal) meshes using C-grid staggering and selective grid refinement.

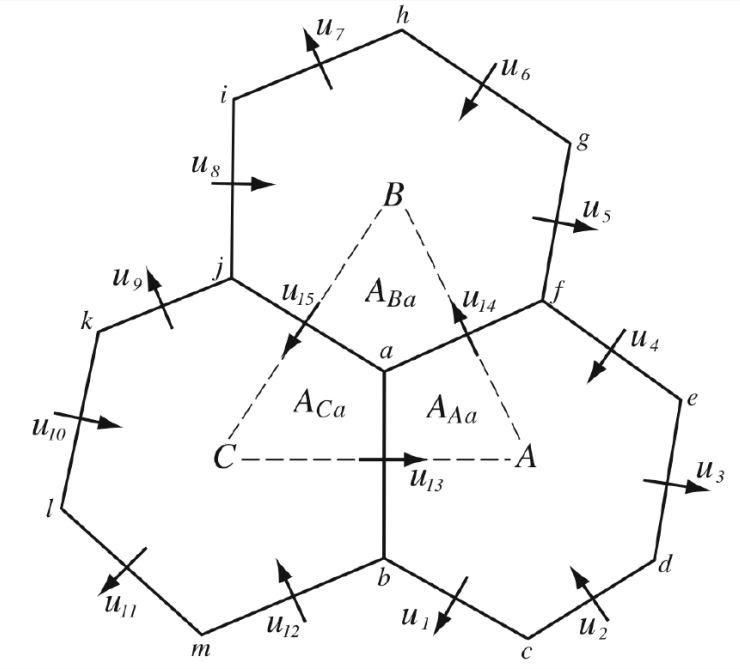

MPAS-A: The Grids

Horizontal

Local Refinement

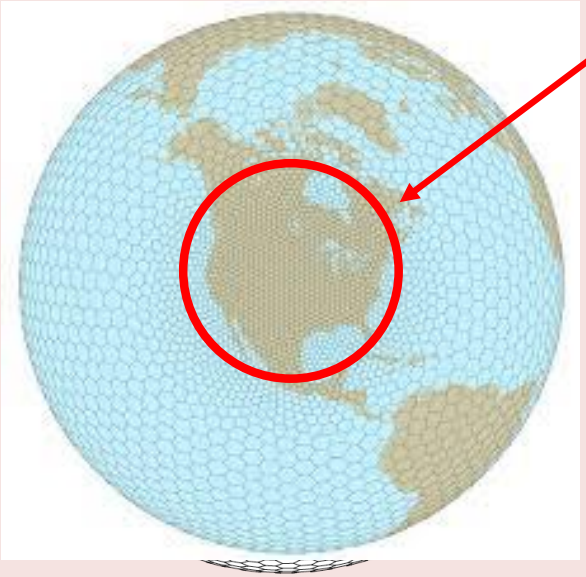

MPAS
Unstructured Voronoi
(hexagonal) grid

- Good scaling on massively parallel computers
- No pole problems

Vertical

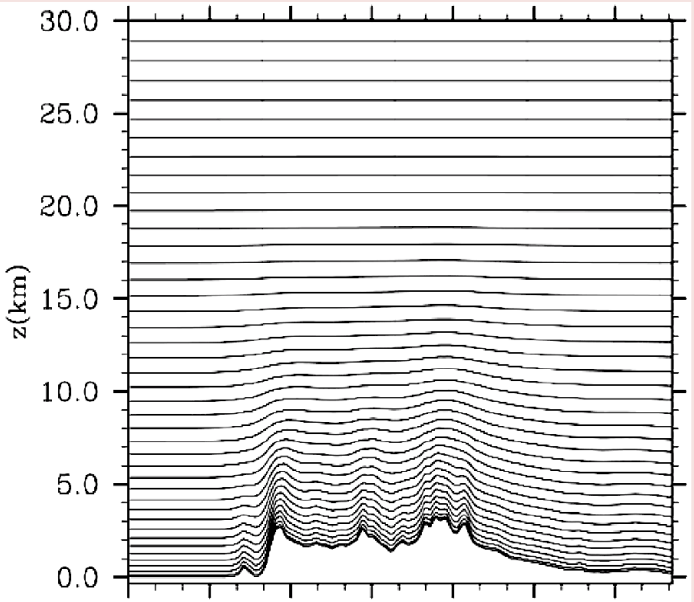

MPAS
Height-based hybrid smoothed
terrain-following vertical coordinate

- Improved numerical accuracy

Anatomy of an MPAS-A Timestep

There are ~350 halo exchanges /timestep!

Default time integration

Physics

Call physics

Moist Dynamics

Do dynamics_split_steps
    Do step_rk3 = 1, 3
        compute large-time-step tendency
        Do acoustic_steps
            update u
            update rho, theta and w
        End acoustic_steps
    End rk3 step
End dynamics_split_steps

Do scalar step_rk3 = 1, 3
    scalar RK3 transport
End scalar rk3 step

Physics

Call microphysics

Allows for smaller dynamics timesteps relative to scalar transport timestep and main physics timestep.

We can use any FV scheme here (we are not tied to RK3)
Scalar transport and physics are the expensive pieces in most applications.

WEAK and STRONG Scaling of MPAS-A moist dynamical core on *Summit*[1] and STRONG scaling on *Cheyenne*[2] (dual, Intel 18c v4 Xeon)

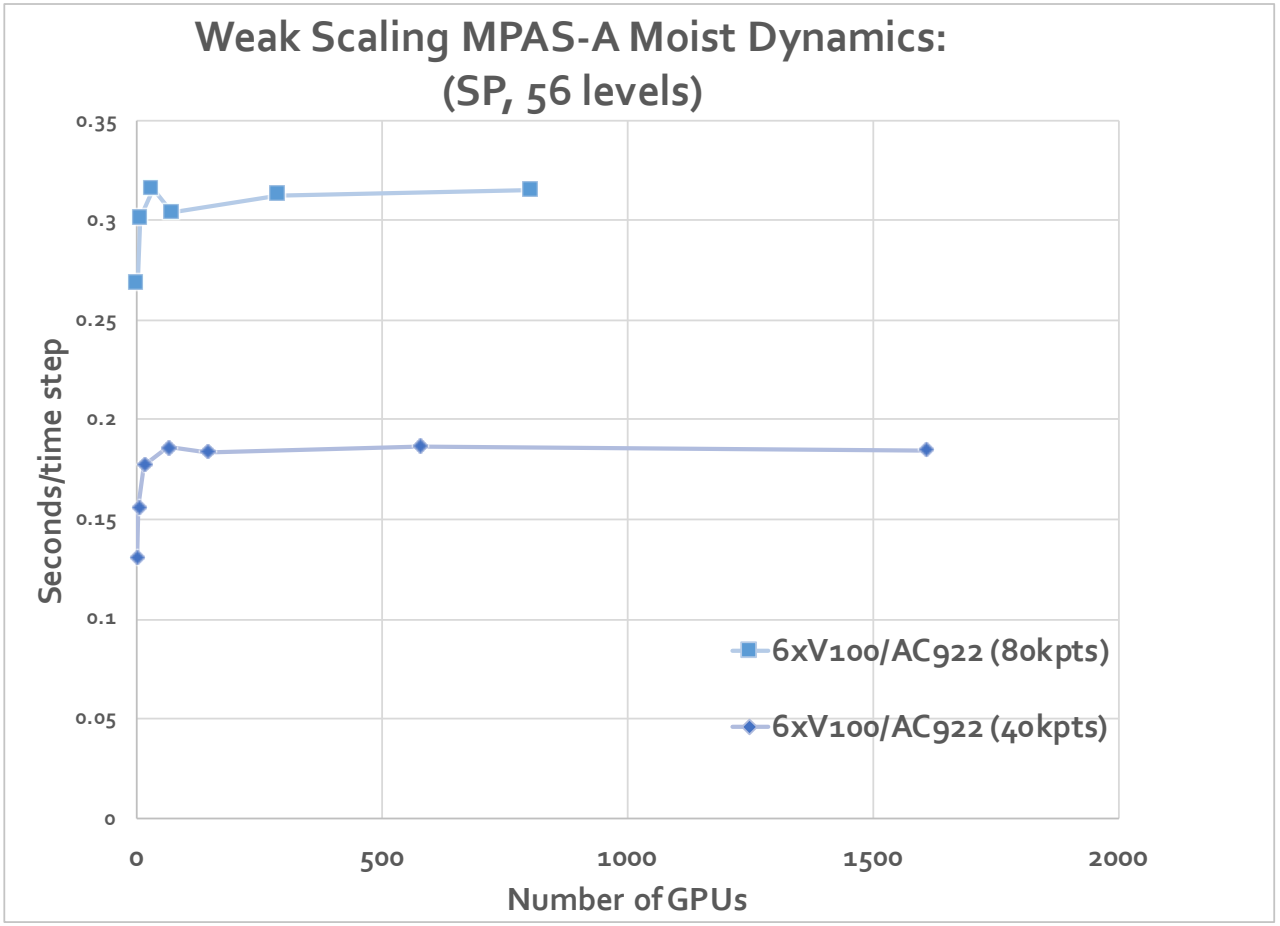

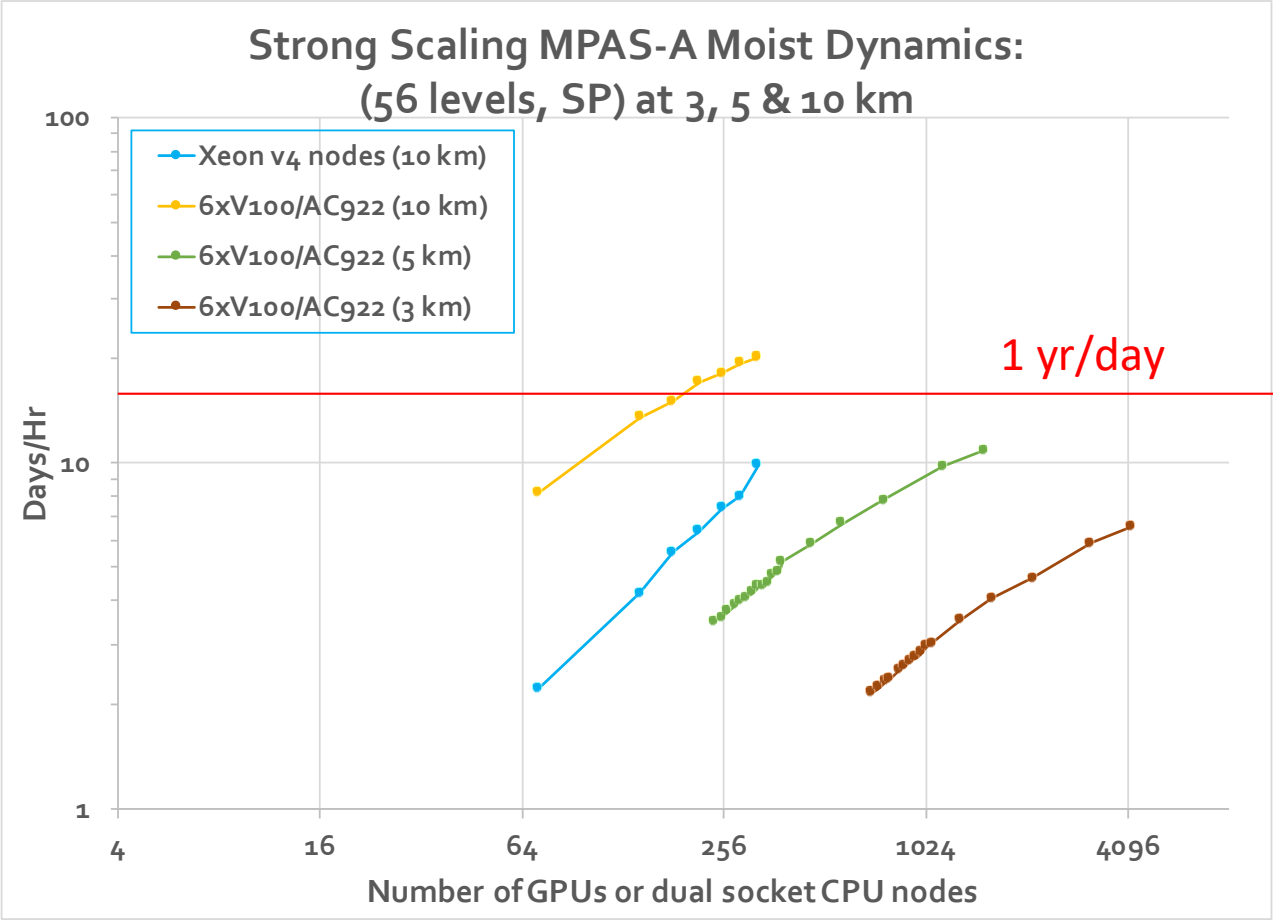

[1]Benchmarking on Summit supported by DoE via an OLCF Director's Discretionary Allocation

[2]Cheyenne is a 5.4 PF, 4032-node HPE system with EDR interconnect operated by NCAR

MPAS-A/GPU moist dycore scaling efficiency (Summit)

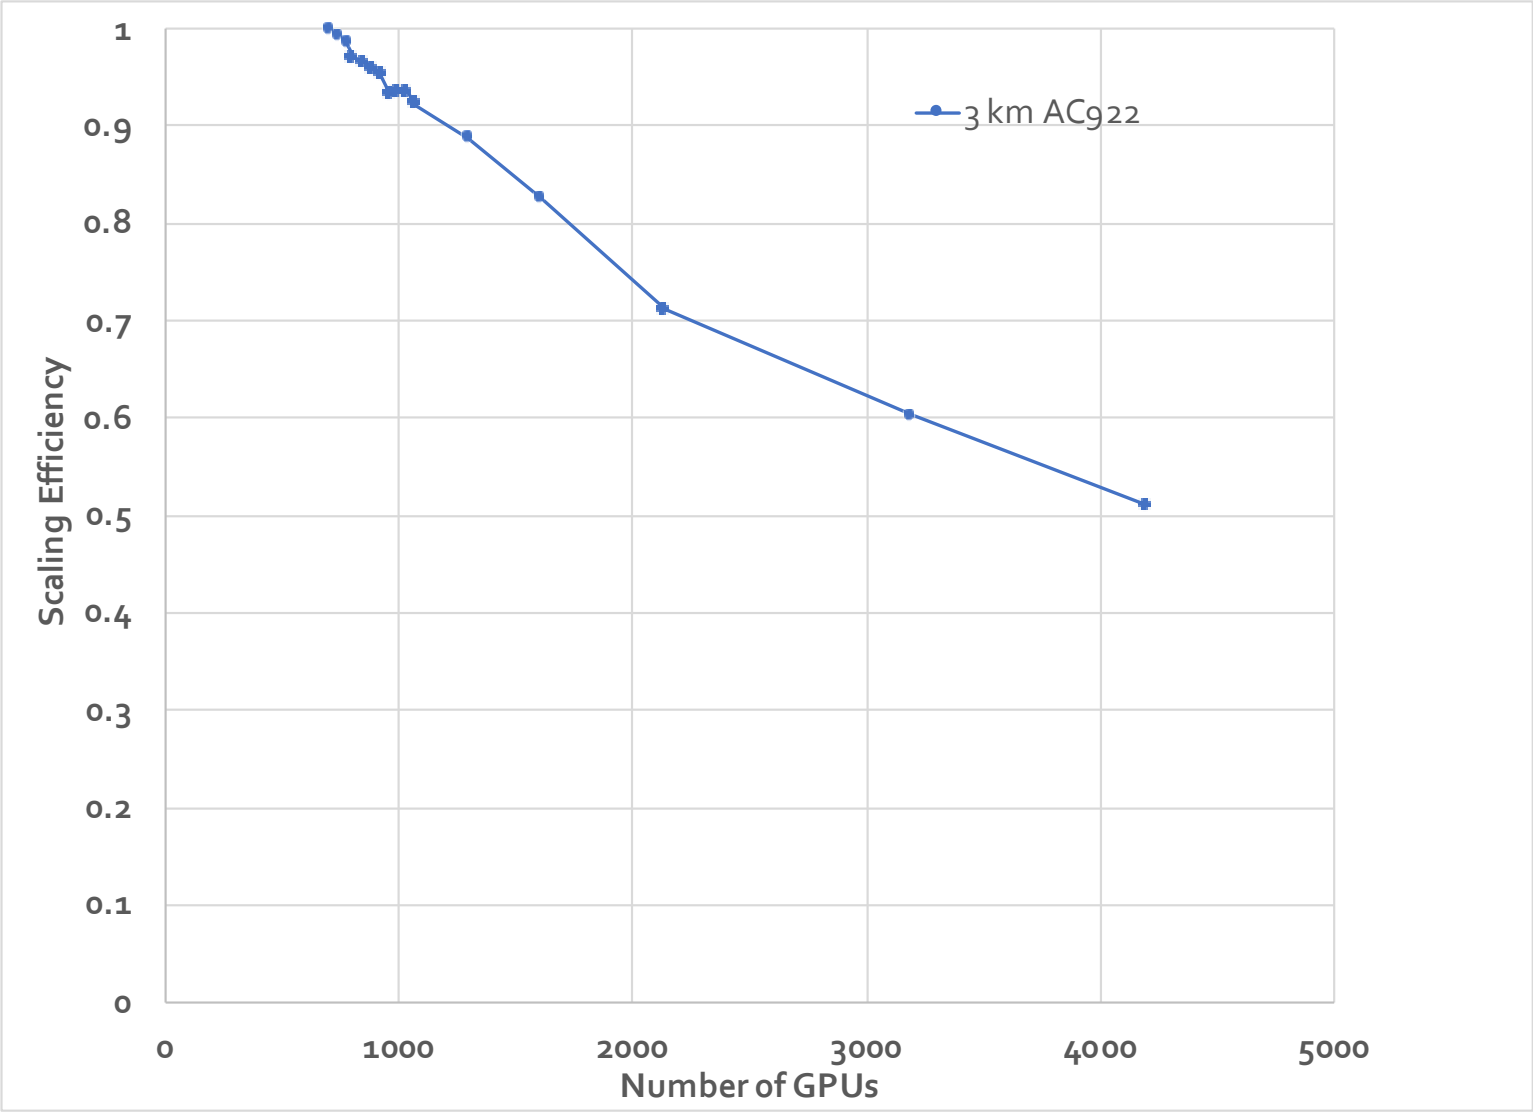

Current Design: Heterogeneous, Task-Parallel

Come see our in-situ vis demo in booth 401 (NCAR Booth)

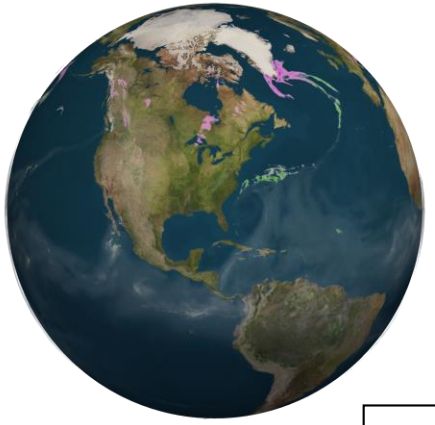

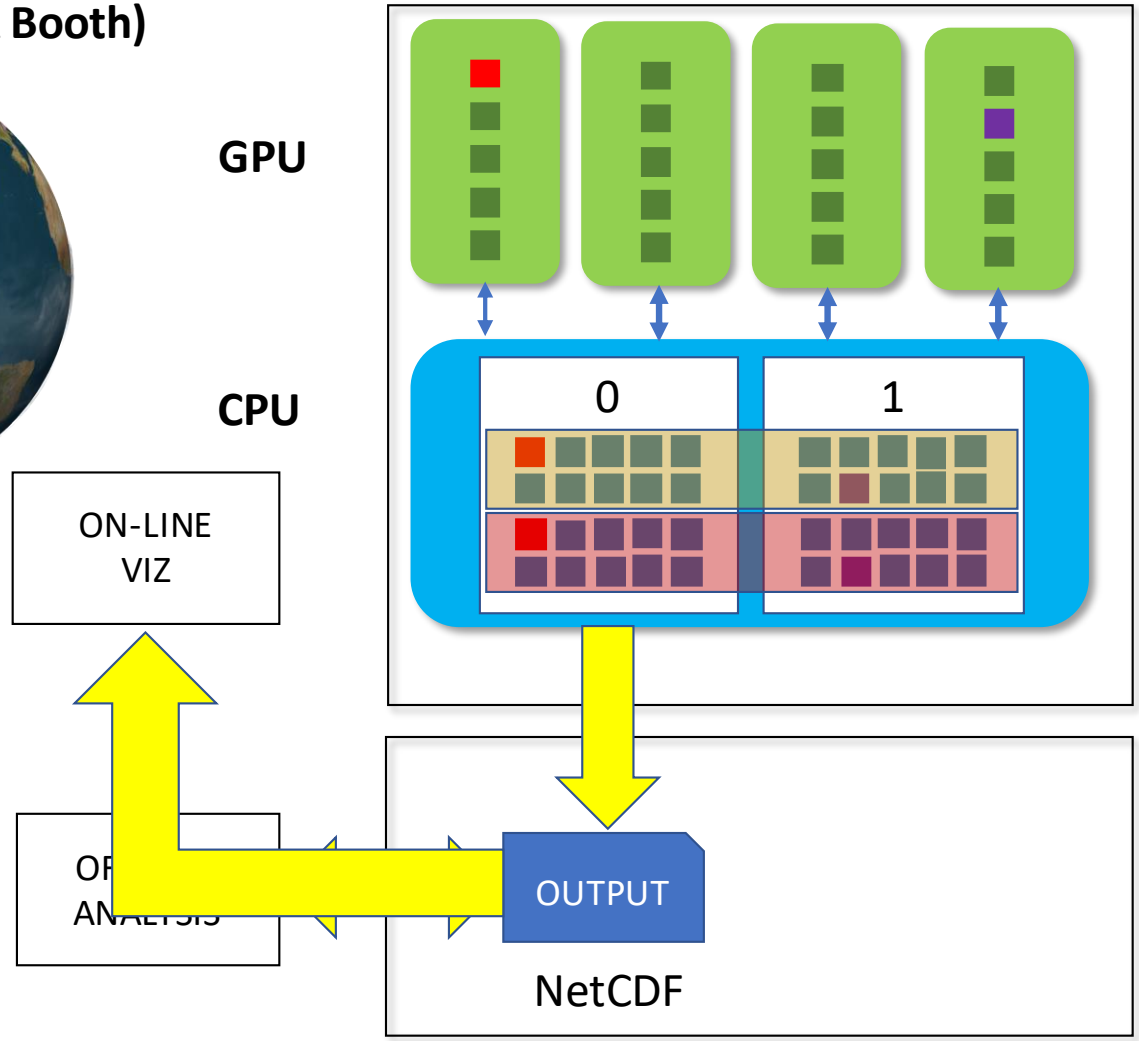

EXASCALE NODE

LAND SURFACE RADIATION

DISK STORAGE

GPU-ranks

CPU-ranks

Linked-ranks

Shown:

- **4 GPUs/node**
- **5 ranks/GPU**
- **20 ranks for LS & RT**

Full Physics MPAS-A on Summit (10 km, 56 levels, SP)

Configuration:

- 6 GPUs
- 40kpts/GPU
- 3 MPI ranks per GPU

Throughput

- 0.234 sec/step
- 3.2 days/hr @ 3 km (uniform)

Caveats

- Results for load-balanced RT
- Still room for improvement!
  - Halo exchanges (22%)

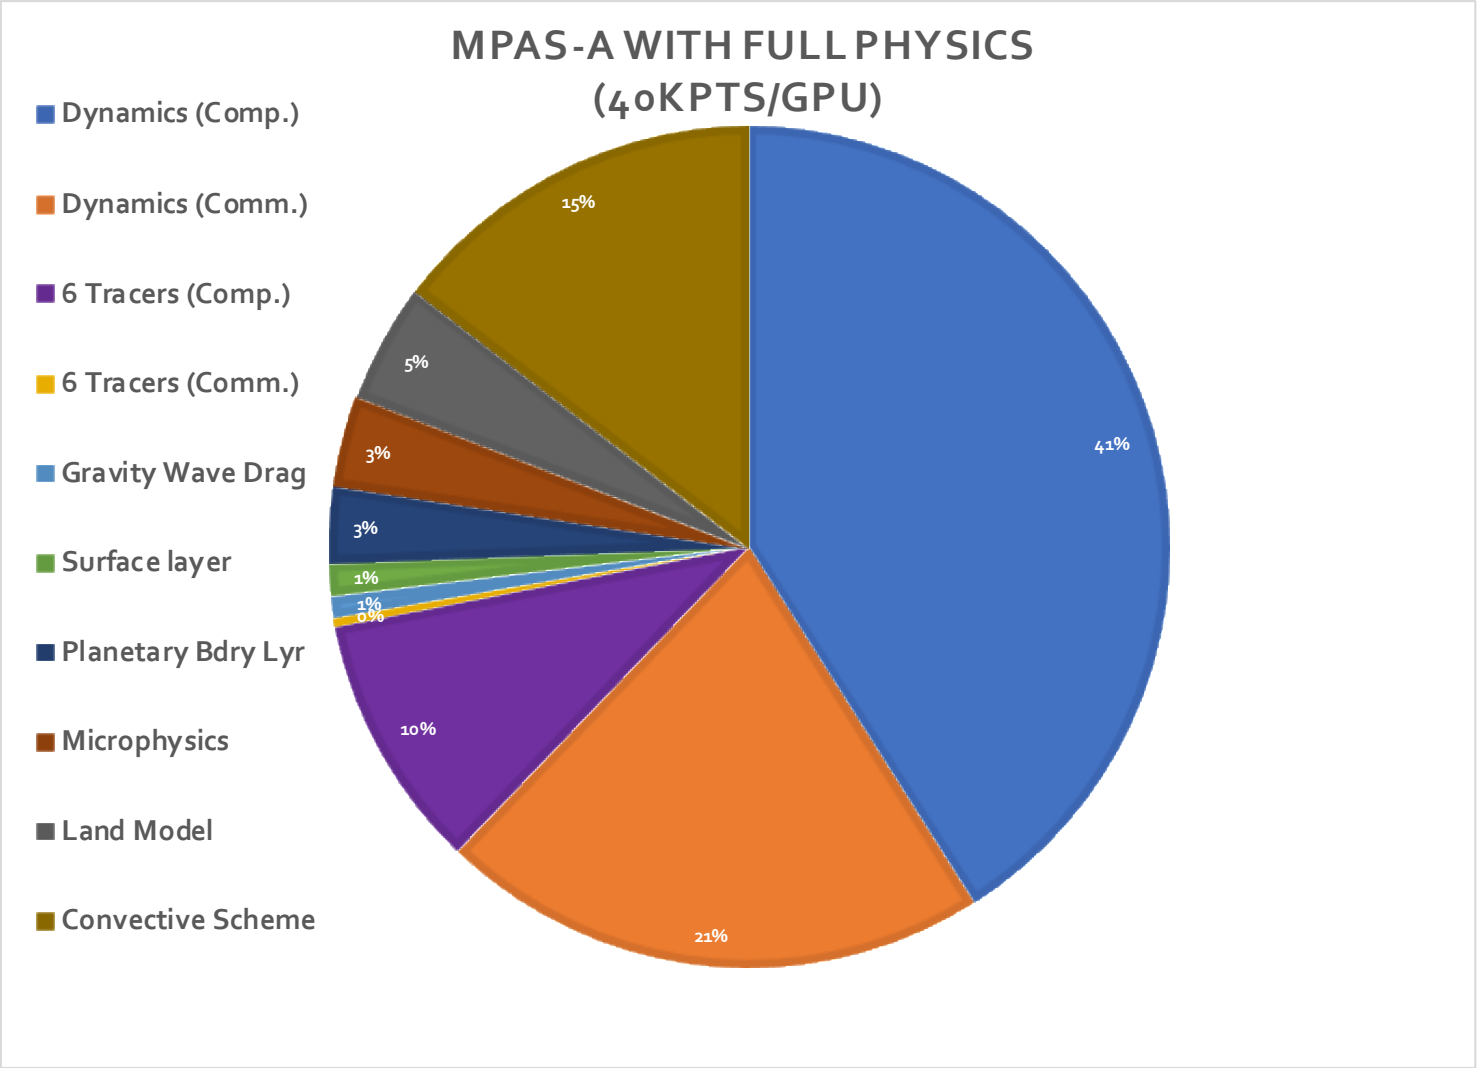

MPAS-A/GPU keys to success: *Partnership*

- Public-Private Partnership
- Open Source/Single Source
- Workforce development via student engagement

Come see our in-situ vis demo in booth 401 (NCAR)!

Key part of our team: University of Wyoming students and NVIDIA and PGI experts.

MPAS-A/GPU keys to success: *OpenACC*

- Exploiting **task parallelism** by
  - by overlapping CPU/GPU execution
- **Coalescing CUDA kernels** for MPI packs and unpacks in halo exchanges.
- **PCAST** auto-compare feature (PGI) for debugging GPU code.
- **Mixed execution** of OpenACC CPU threading and OpenACC GPU offload:
  - compile GPU files with "ta=tesla" flag;
  - compile CPU files with "ta=multicore" flag;
  - link all files together with "ta=tesla, multicore"

How close is MPAS-A to achieving a practical Cloud Resolving Model (CRM) capability?

$$N^{*} = S \times L \times F \times R$$

N = speedup required
S = ratio of desired to observed throughput
L = vertical adjustment (number of levels/56)
F = floating point precision adjustment (i.e. SP->DP)
R = resolution adjustment (Δx/1 km)

* A "Drake Equation" for CRM?

Credit: NRAO

How close is MPAS-A to achieving a practical Cloud Resolving Model (CRM) capability?

Goal: 1 simulated year per day in double precision and 100 levels at 1 km, then.

S = 4.75
L = 1.8
R = 3
F = 1.5

N = S x L x R x F ~ 40x

A "Drake Equation" for CRM?

What can we do besides just wait for 40x from HPC architectures?

Credit: NRAO

How can we tackle the throughput gap?
A multi-prong approach…

$$N = S \times L \times F \times R$$

Parallelism, Optimization & Faster HW

Machine Learning, Low-precision Physics

Mathematics, Solvers, etc.

A “Drake Equation” for CRM?

Various projects around the community are trying to tackle this:

- **ESCAPE2**
- **E3SM**
- **CLIMA**

Credit: NRAO

Exascale Earth System Model Vision: Heterogeneous, Task-Parallel, In Situ, Intelligent

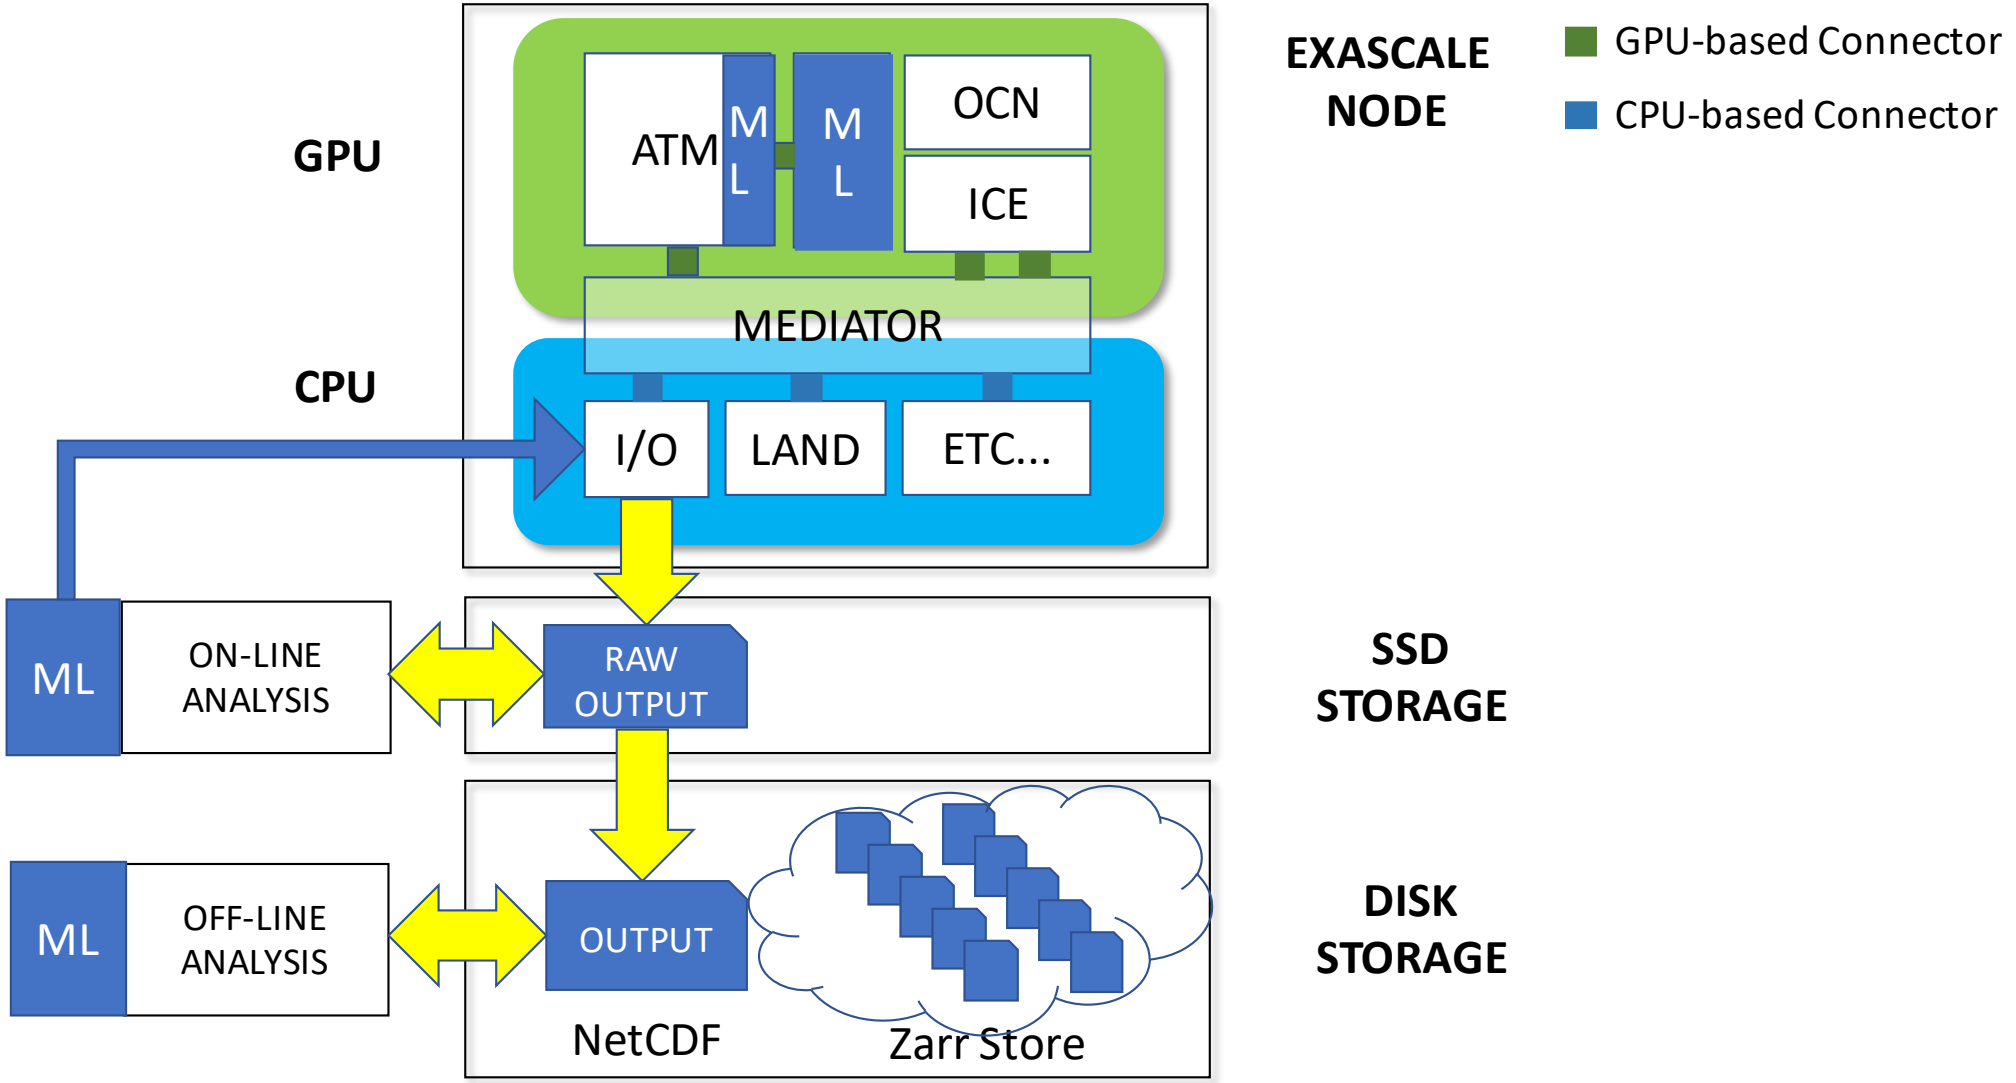

Replacing cloud for rain processes with NNs

ML TAU emulator:
3 classifier networks, 4 regression networks
82,327 weights total
Dense Neural Network Hyperparameters
- 4 layers
- 60 neurons per hidden layer
- 11,761 total weights
- Rectified Linear Unit (ReLU) activation fns

Bin Scheme (Tel Aviv University (TAU):
- Divide particle sizes into bins **(35)**
- Calculate evolution of each bin separately
- Better representation of interactions
- Much more computationally expensive than current bulk schemes used in climate models

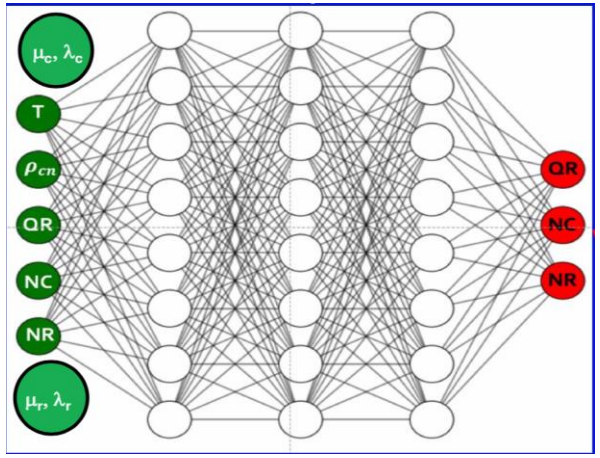

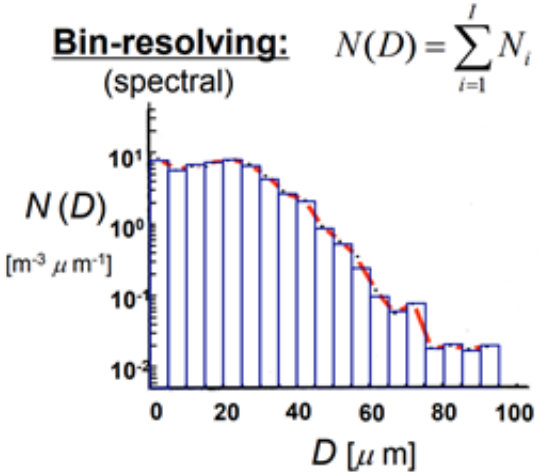

Results after 9 simulated years

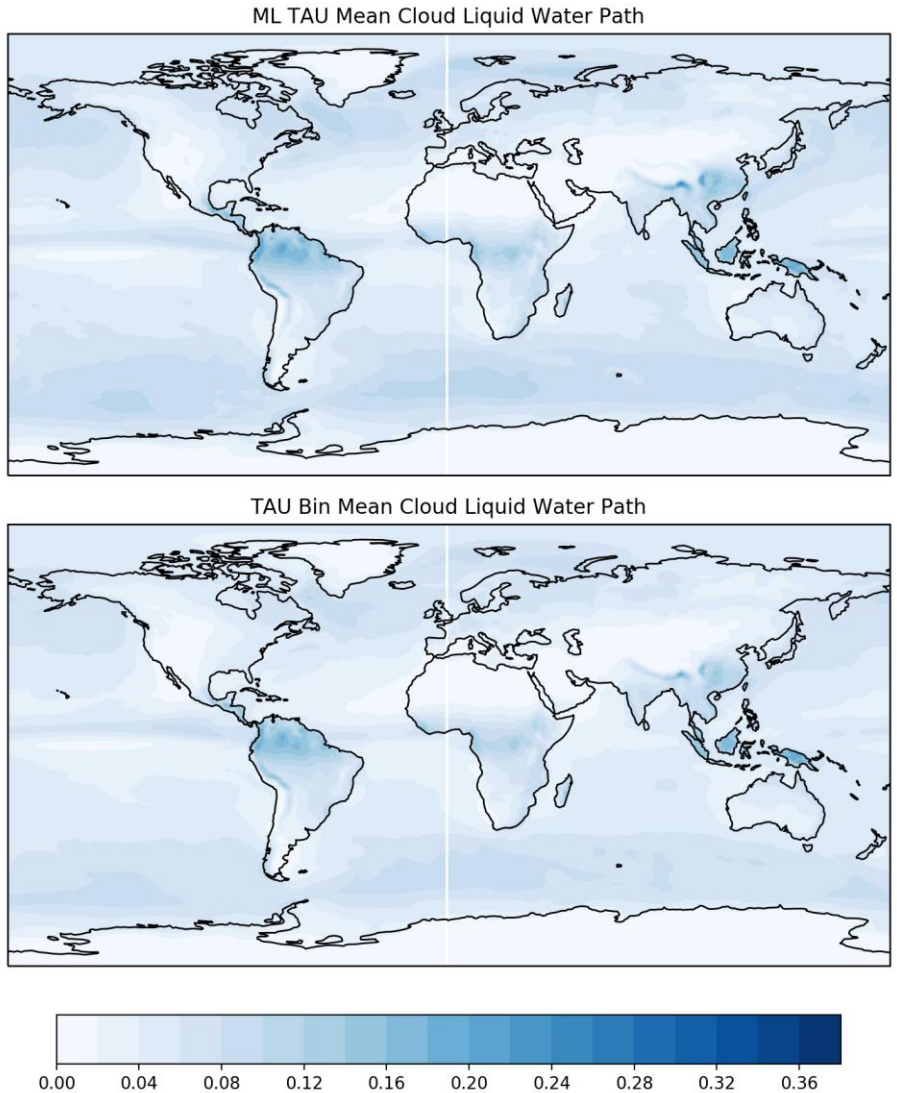

AI Everywhere

All Sky Imager

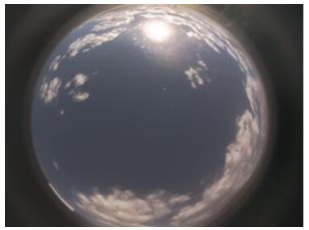

Richardson, et al. 2017

Surface Layer Parameterization

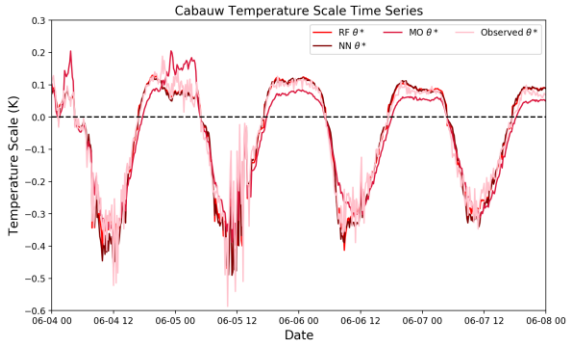

McCandless, Kosovic, Gagne

DICast® Model

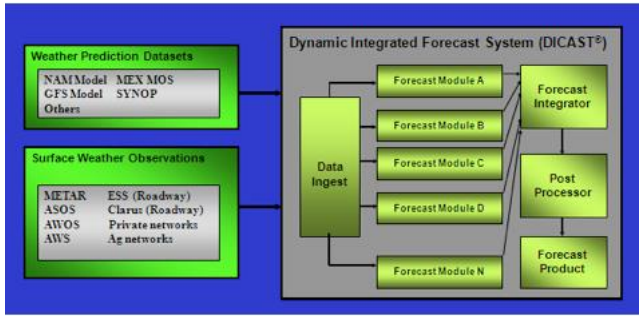

Haupt, et al. RAL

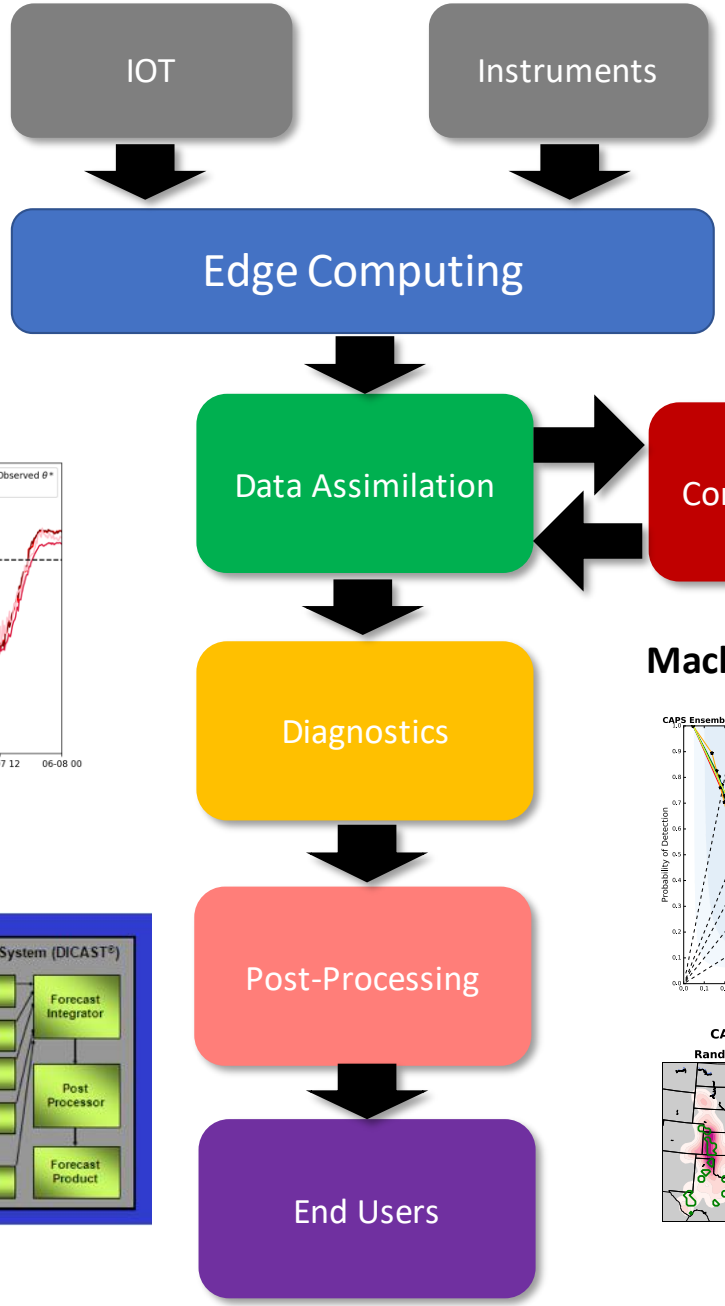

EOL "Big Data" Instruments

HOLODEC APAR

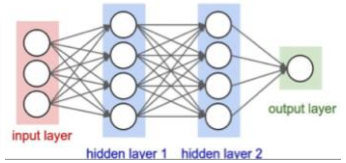

Machine Learning Hail Prediction

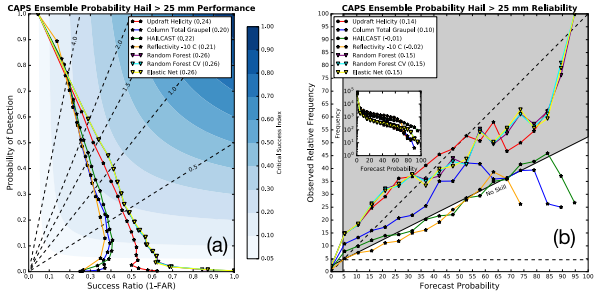

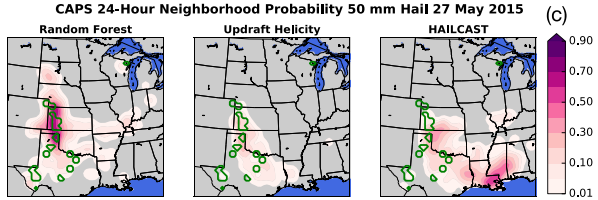

D. J. Gagne et al, 2017

Offload to GPUs is, perhaps, a good step to getting ready for downstream innovation…

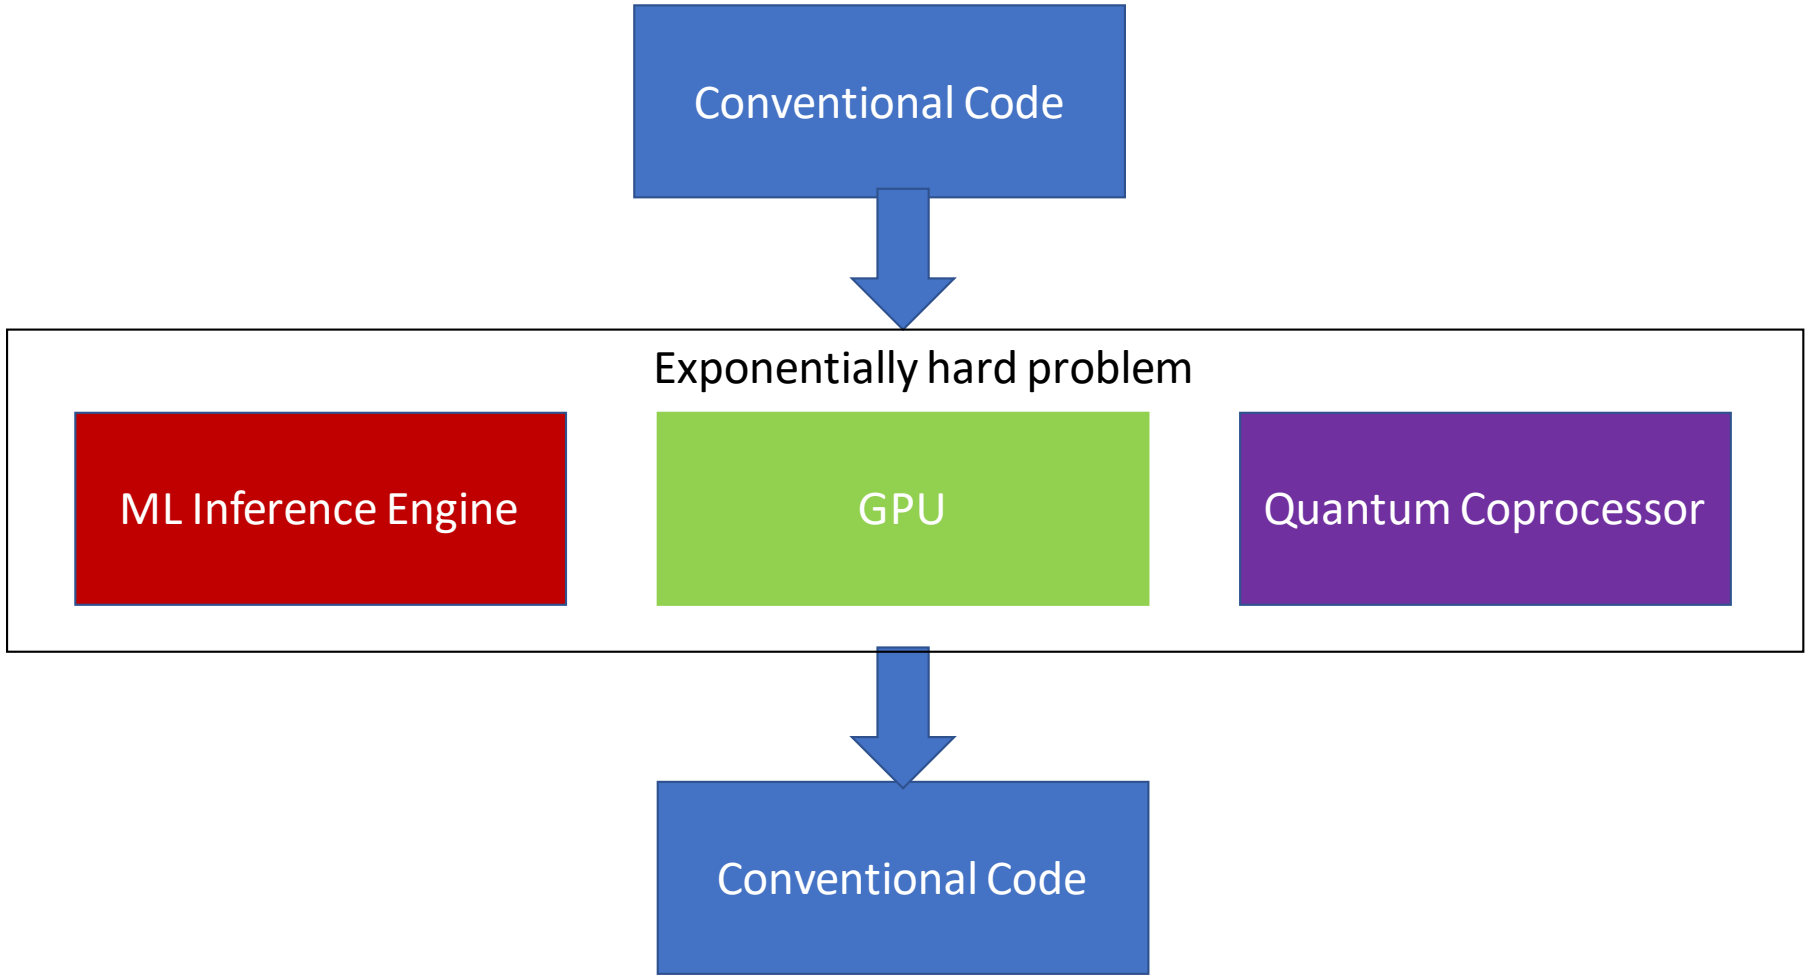

Thanks! Questions

Dr. Richard Loft
loft@ucar.edu
Director, Technology Development
Computational & Information Systems Laboratory
National Center for Atmospheric Research
Boulder, CO

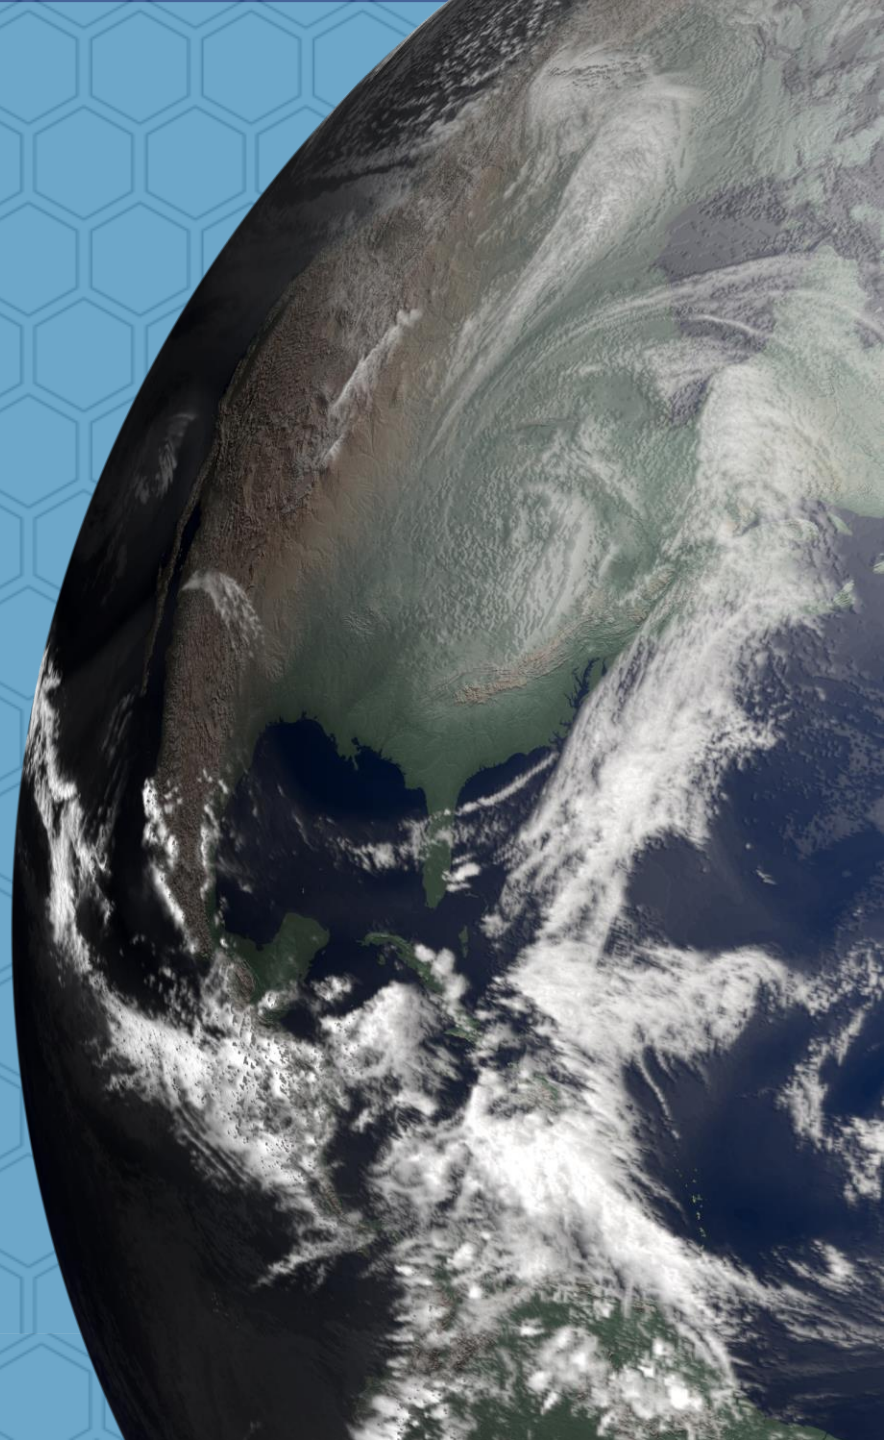

This material is based upon work supported, in part, by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.

Cloud to Rain Processes

Cloud water mixing ratio (q_c)
Cloud water number concentration (N_c)
Rain water mixing ratio (q_r)
Rain water number concentration (N_r)

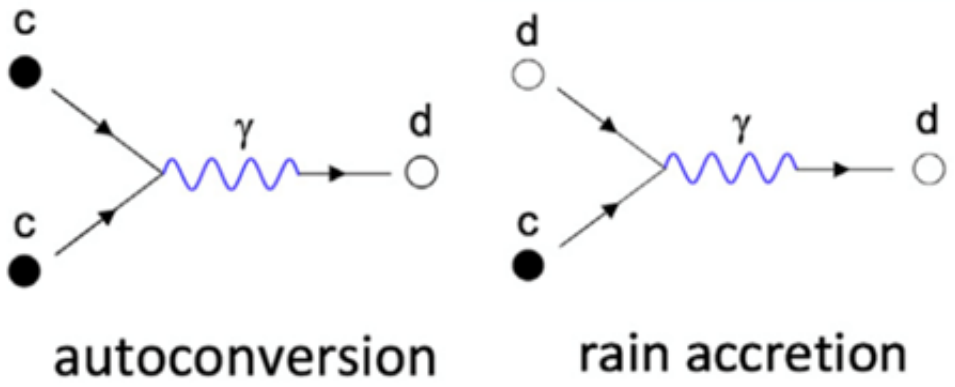

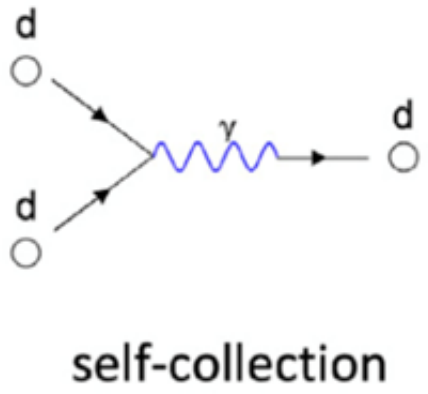

d: rain drop
c: cloud droplet
CCN: cloud condensation nuclei

Cloud droplets grow into rain droplets through 3 processes:

Autoconversion: cloud droplets collide in a chain reaction to form rain drops
$dq_c/dt < 0$, $dq_r/dt > 0$
$dN_c/dt < 0$, $dN_r/dt > 0$

Rain Accretion: rain drops collide with cloud droplets
$dq_c/dt < 0$, $dq_r/dt > 0$
$dN_c/dt < 0$, $dN_r/dt = 0$

Self-Collection: rain drops collide with other raindrops
$dq_c/dt = 0$, $dq_r/dt = 0$
$dN_c/dt = 0$, $dN_r/dt < 0$

Microphysics Emulation Recipe

- Run CESM2/CAM6 for two years and obtain instantaneous hourly output
- Filter and subsample data to find grid points with realistic amount of cloud water
- Transform and normalize inputs and outputs
- Train classifier deep neural networks to classify zero and non-zero
- Train regression deep neural networks to predict non-zero values
- Evaluate and interpret neural network predictions.

Extending Microphysics Emulator Work

Even more complex droplet process models exist – e.g. Superdroplet (right)

Such detailed codes (right) are ridiculously expensive for large scale models, so ML may be the only path forward.

Other complex interaction networks exist in Earth system models, including:

- Chemical reaction networks
- Ecological networks

Can lessons learned with droplet formation models be applied more broadly?

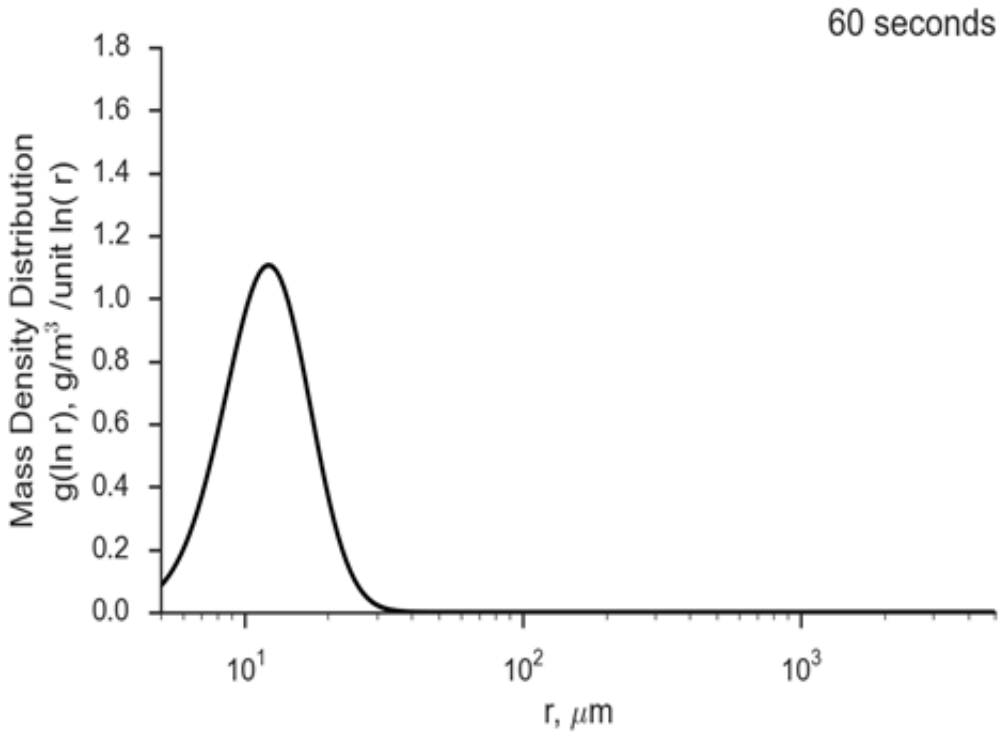

Superdroplet model output animation
Credit: Daniel Rothenberg