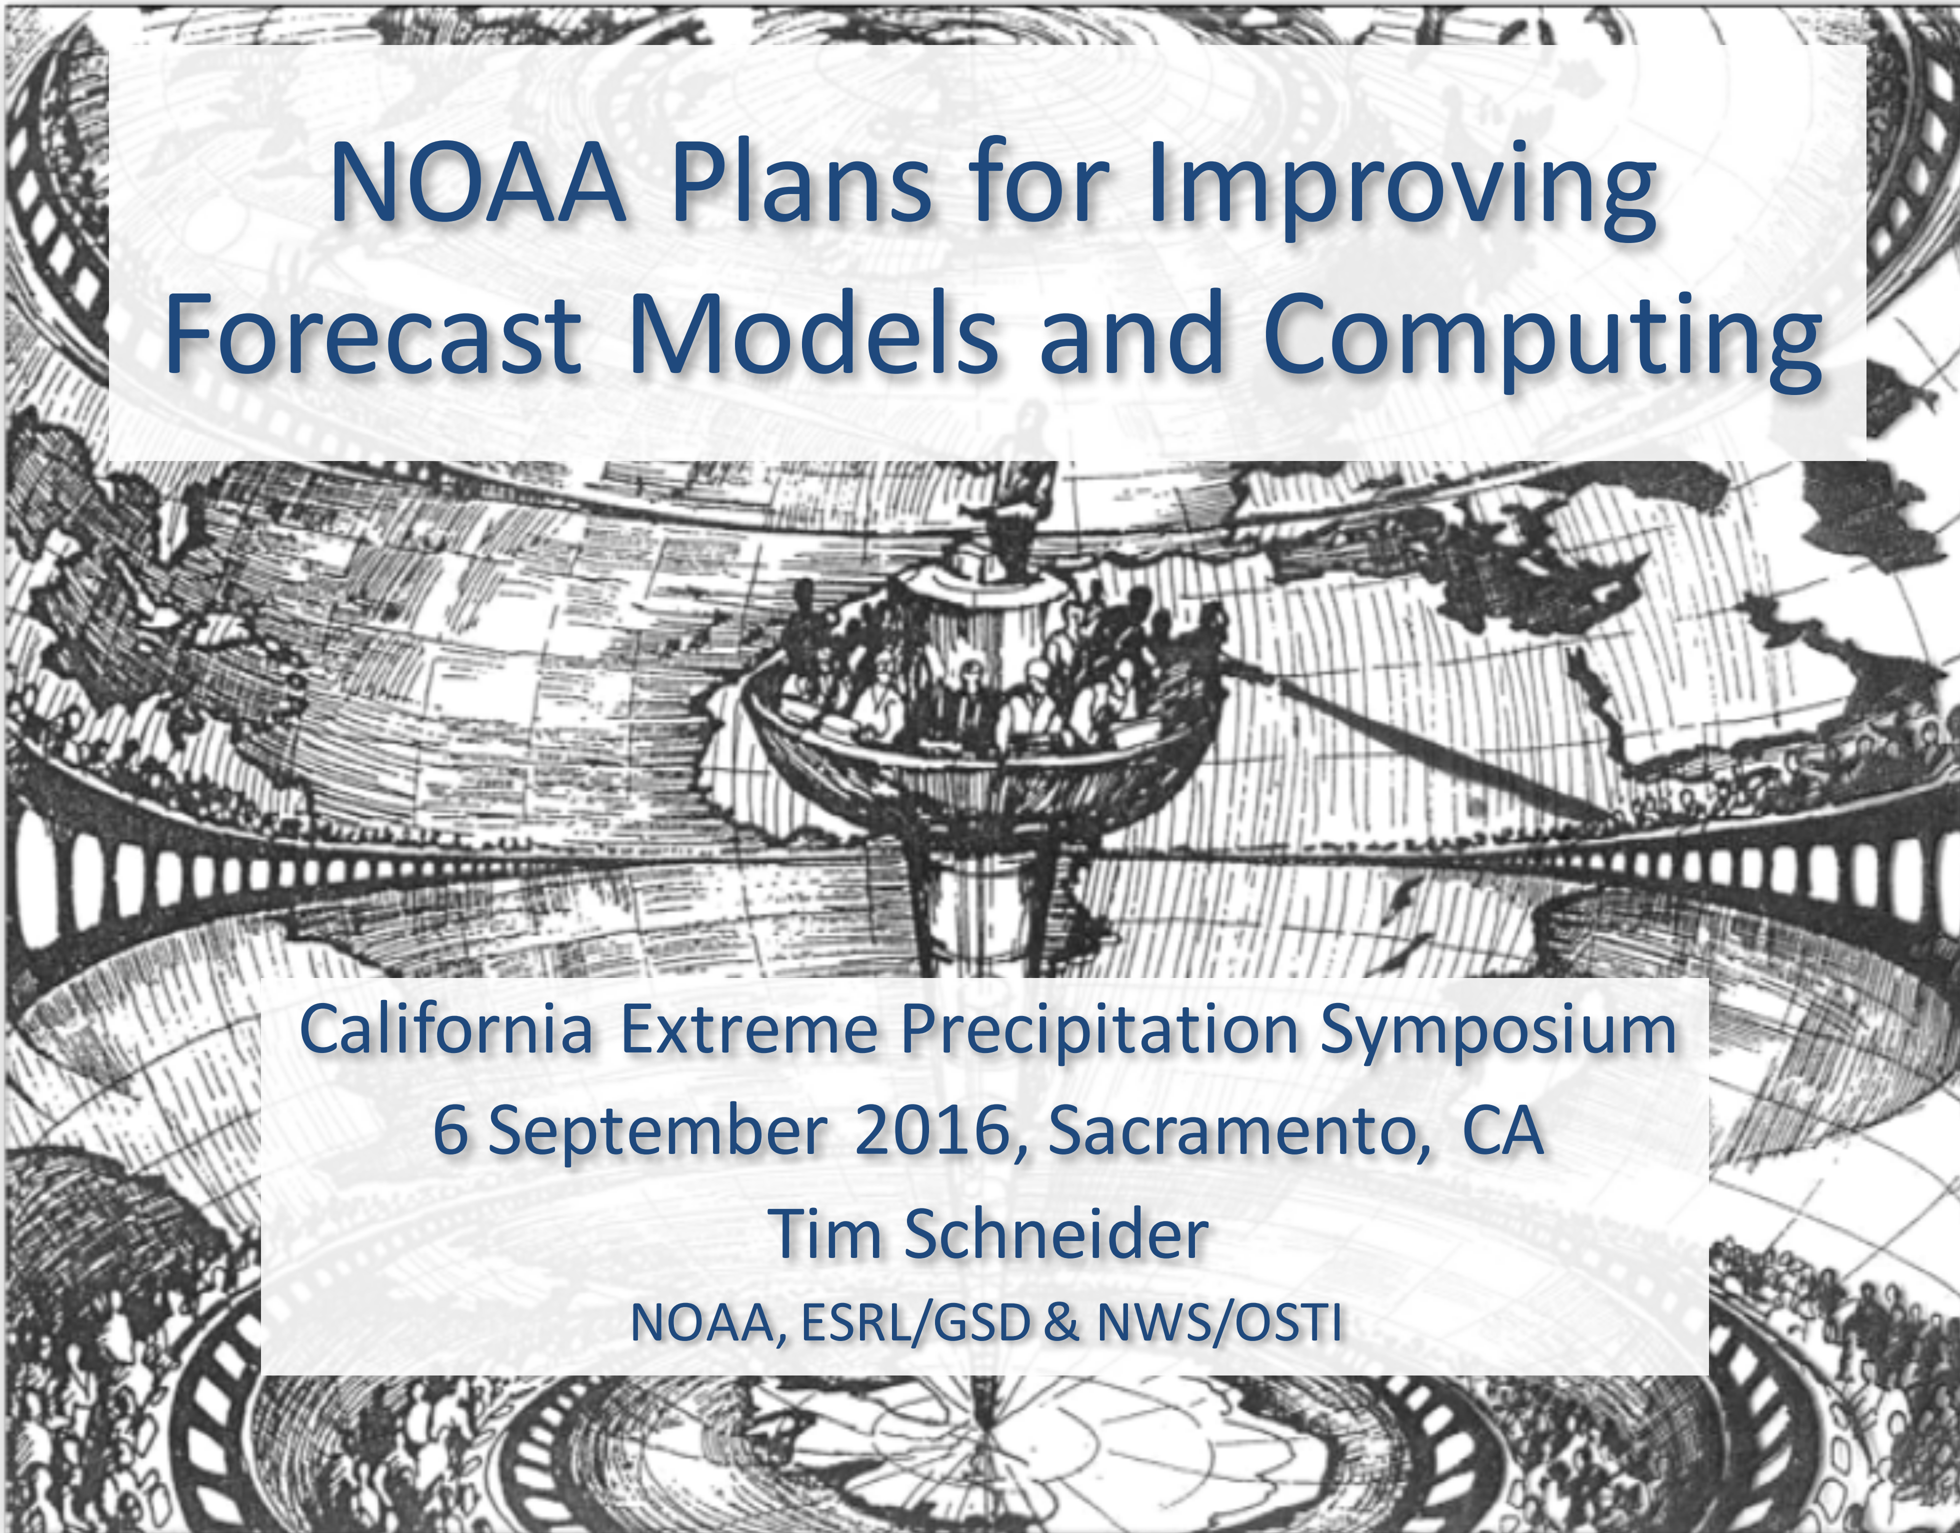

NOAA Plans for Improving Forecast Models and Computing

California Extreme Precipitation Symposium

6 September 2016, Sacramento, CA

Tim Schneider

NOAA, ESRL/GSD & NWS/OSTI

Abstract

By selecting a state-of-the-art "non-hydrostatic" dynamical core for its new global modeling system, the Finite Volume on a Cubed Sphere (or FV3), NOAA is poised to greatly advance and simplify its operational numerical weather prediction modeling suite. Non-hydrostatic models allow for explicit representation and resolution of clouds, precipitation, and storm dynamics and makes it possible to plan for a unification of global and regional modeling scales. FV3 is highly efficient and adaptable on high performance computers, and has demonstrated skill at both weather and climate time scales. It has the ability to work with refined meshes and/or nests, allowing for very high resolutions over limited areas. NOAA is also conducting research and development on advancing our numerical weather modeling on new massively parallel fine-grain computers. These systems employ hundreds of thousands of processors and take advantage of complex and powerful configurations that integrate CPUs with graphical processing units (GPUs) or many integrated cores (MICs).

The new global model based on FV3 will be implemented in a flexible and scalable modeling framework call the NOAA Environmental Modeling System (NEMS) and will include new physics and advanced data assimilation. The first instance of the new global system is anticipated in 2019.

Biographical Sketch

Tim Schneider is a Physical Scientist at the NOAA Earth System Research Laboratory in Boulder Colorado. His work has focused on observing and modeling the Earth's energy and water cycles, spanning weather, water and climate. Most recently he is overseeing collaborative research and development efforts of global weather models and on the optimal use of high performance "super" computers. He is also interested in connecting models of the atmosphere, the coastal ocean, and the land surface and sub-surface to predict how much water there is, where it is going, and water quality issues. Tim is currently on a temporary assignment to the National Weather Service to manage the Next Generation Global Prediction System (NGGPS) Program.

Important contributions to this presentation were provided by teams at:

NOAA National Weather Service:

- Office of Science and Technology
- NCEP Environmental Modeling Center

NOAA Research:

- Geophysical Fluid Dynamics Laboratory
- Earth System Research Laboratory

NCEP Production Suite (*Simplified*)

ca. March 2016

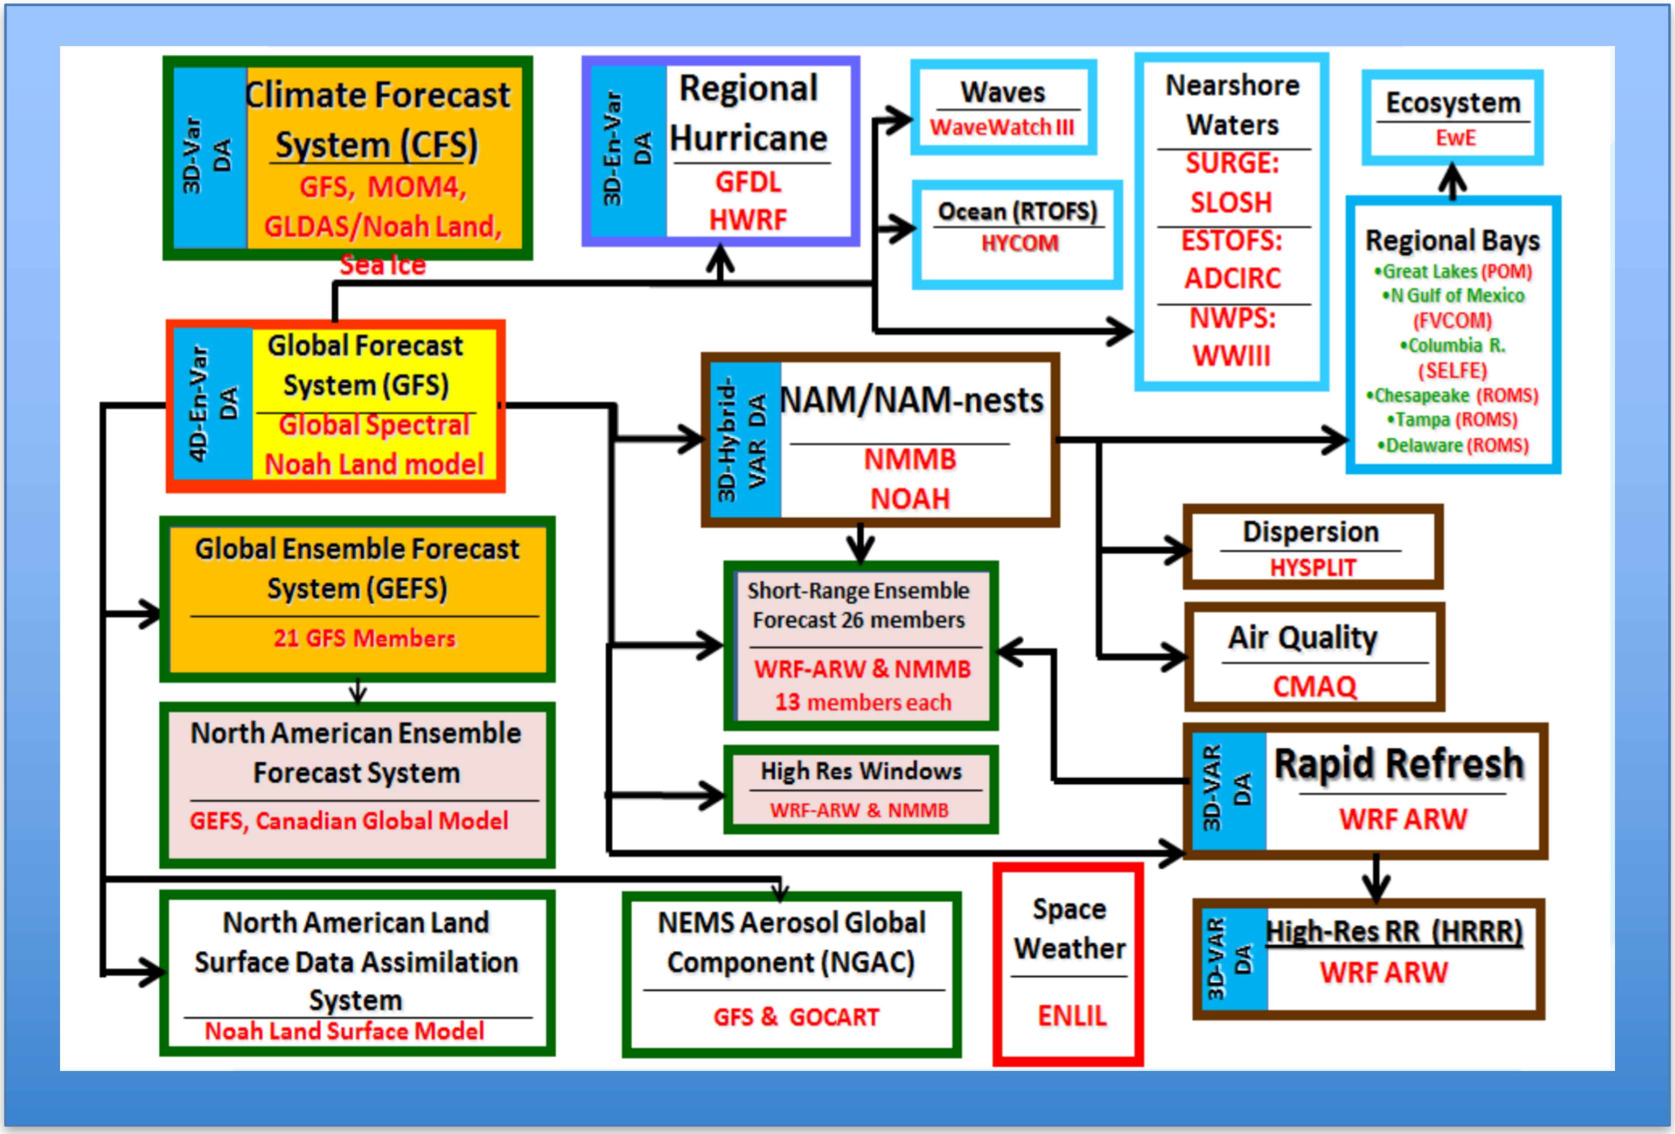

Production Suite Dependencies on GFS

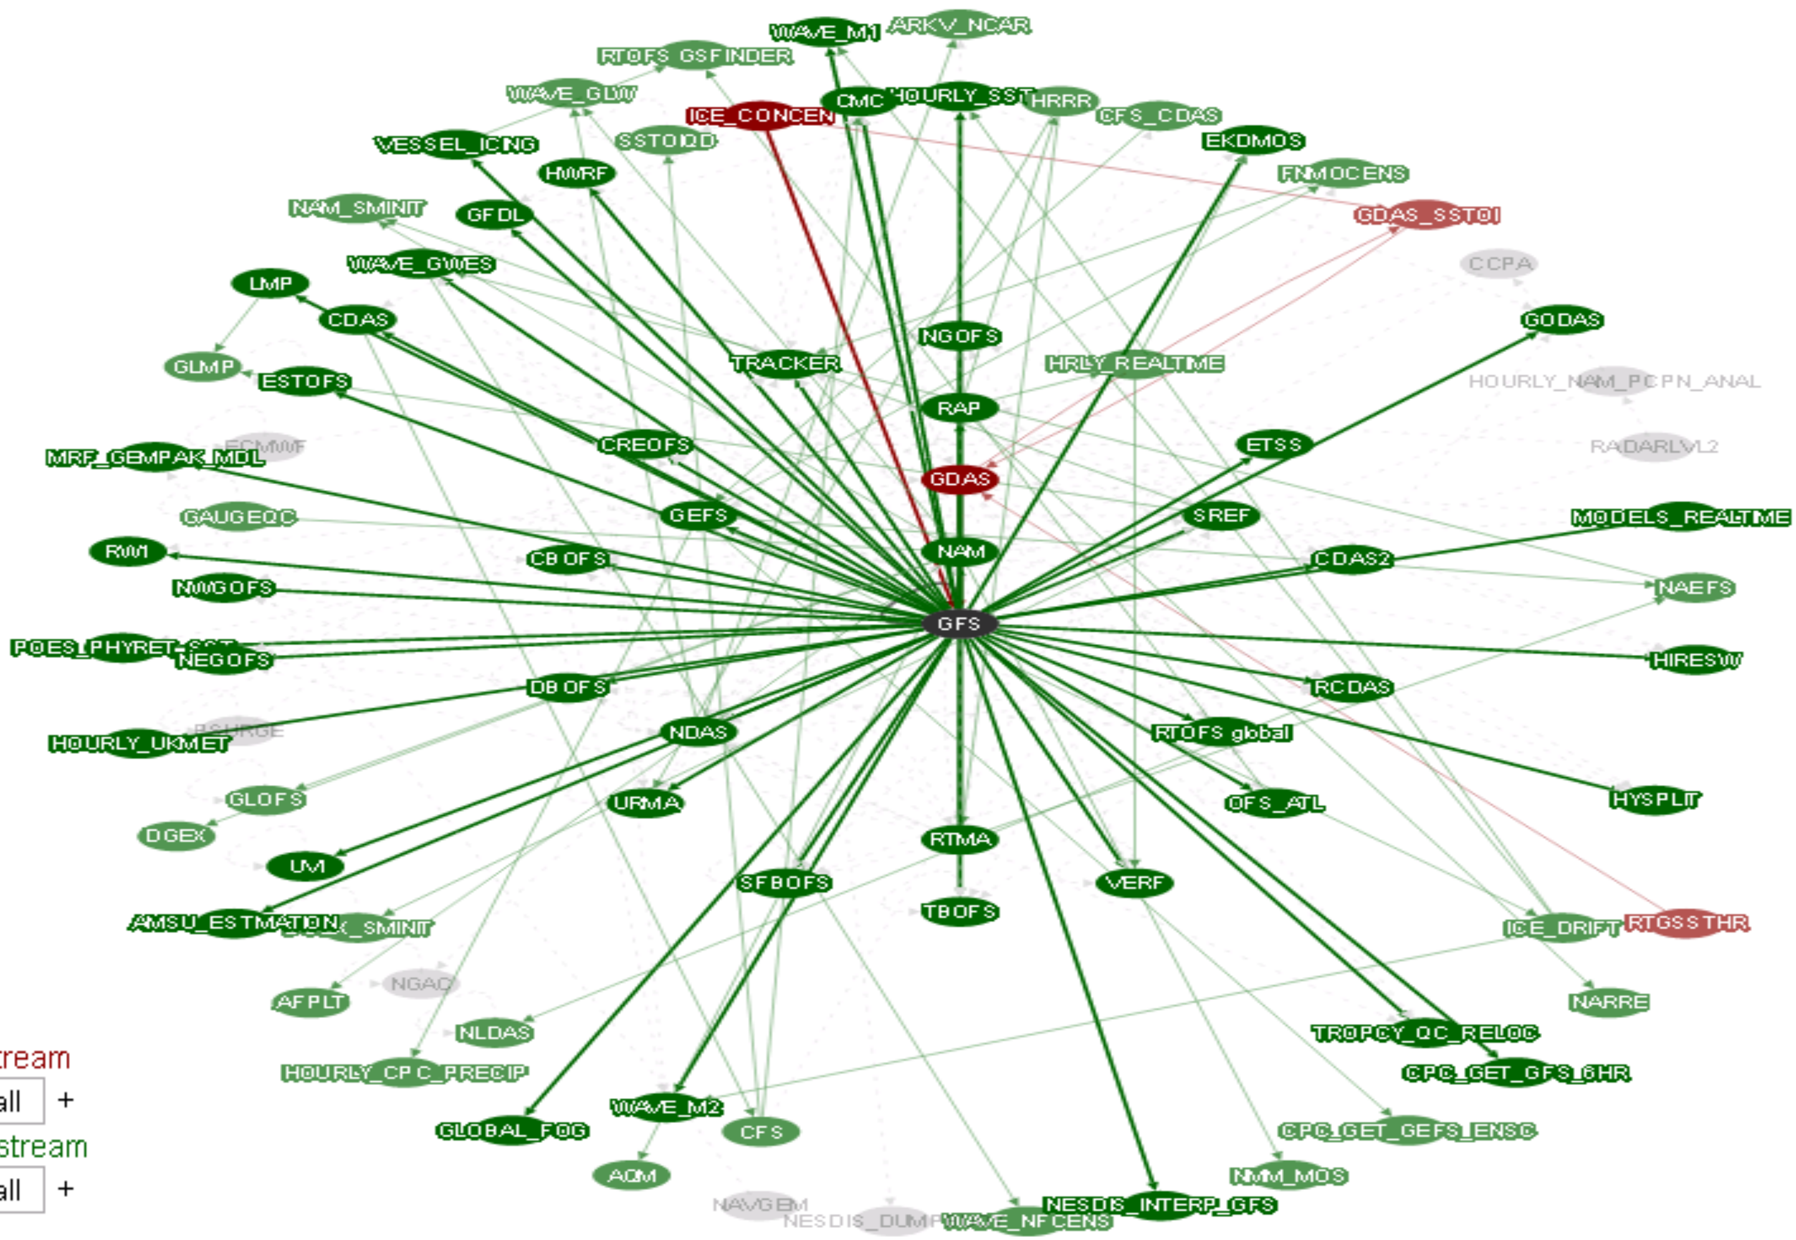

upstream

– all +

downstream

– all +

Goals

NOAA:

- To enable a Weather Ready Nation (models play an important part – but only a part)
- Moving towards evidence-based decision making
- Reducing the production suite complexity
- Moving to an Earth system modeling framework
- Engaging community

Talk Today:

- To provide glimpse into global model developments at NOAA: into what the future might hold

Outline

I. Introduction – A New Era
II. Next Generation Global Prediction System (NGGPS)
III. Finite Volume on a Cubed Sphere Model (FV3)
IV. Advances in High-Performance Computing (HPC)
V. Conclusions

Scale Matters…

Resolution <=> Fidelity

- Resolution
- Spatial
- Temporal
- Representation of physical processes

Resolution – Definition

From the Oxford Dictionary of English

3. "…the action of solving a problem…"
4. "…the process of reducing or separating something into constituent parts or components"
5. "The smallest interval measurable by a telescope or other scientific instrument; the resolving power."; the degree of detail visible in a photographic or television image: a high-resolution monitor"

For Example: Ocean-Land-Atmosphere Interactions: The Devil's in the Details

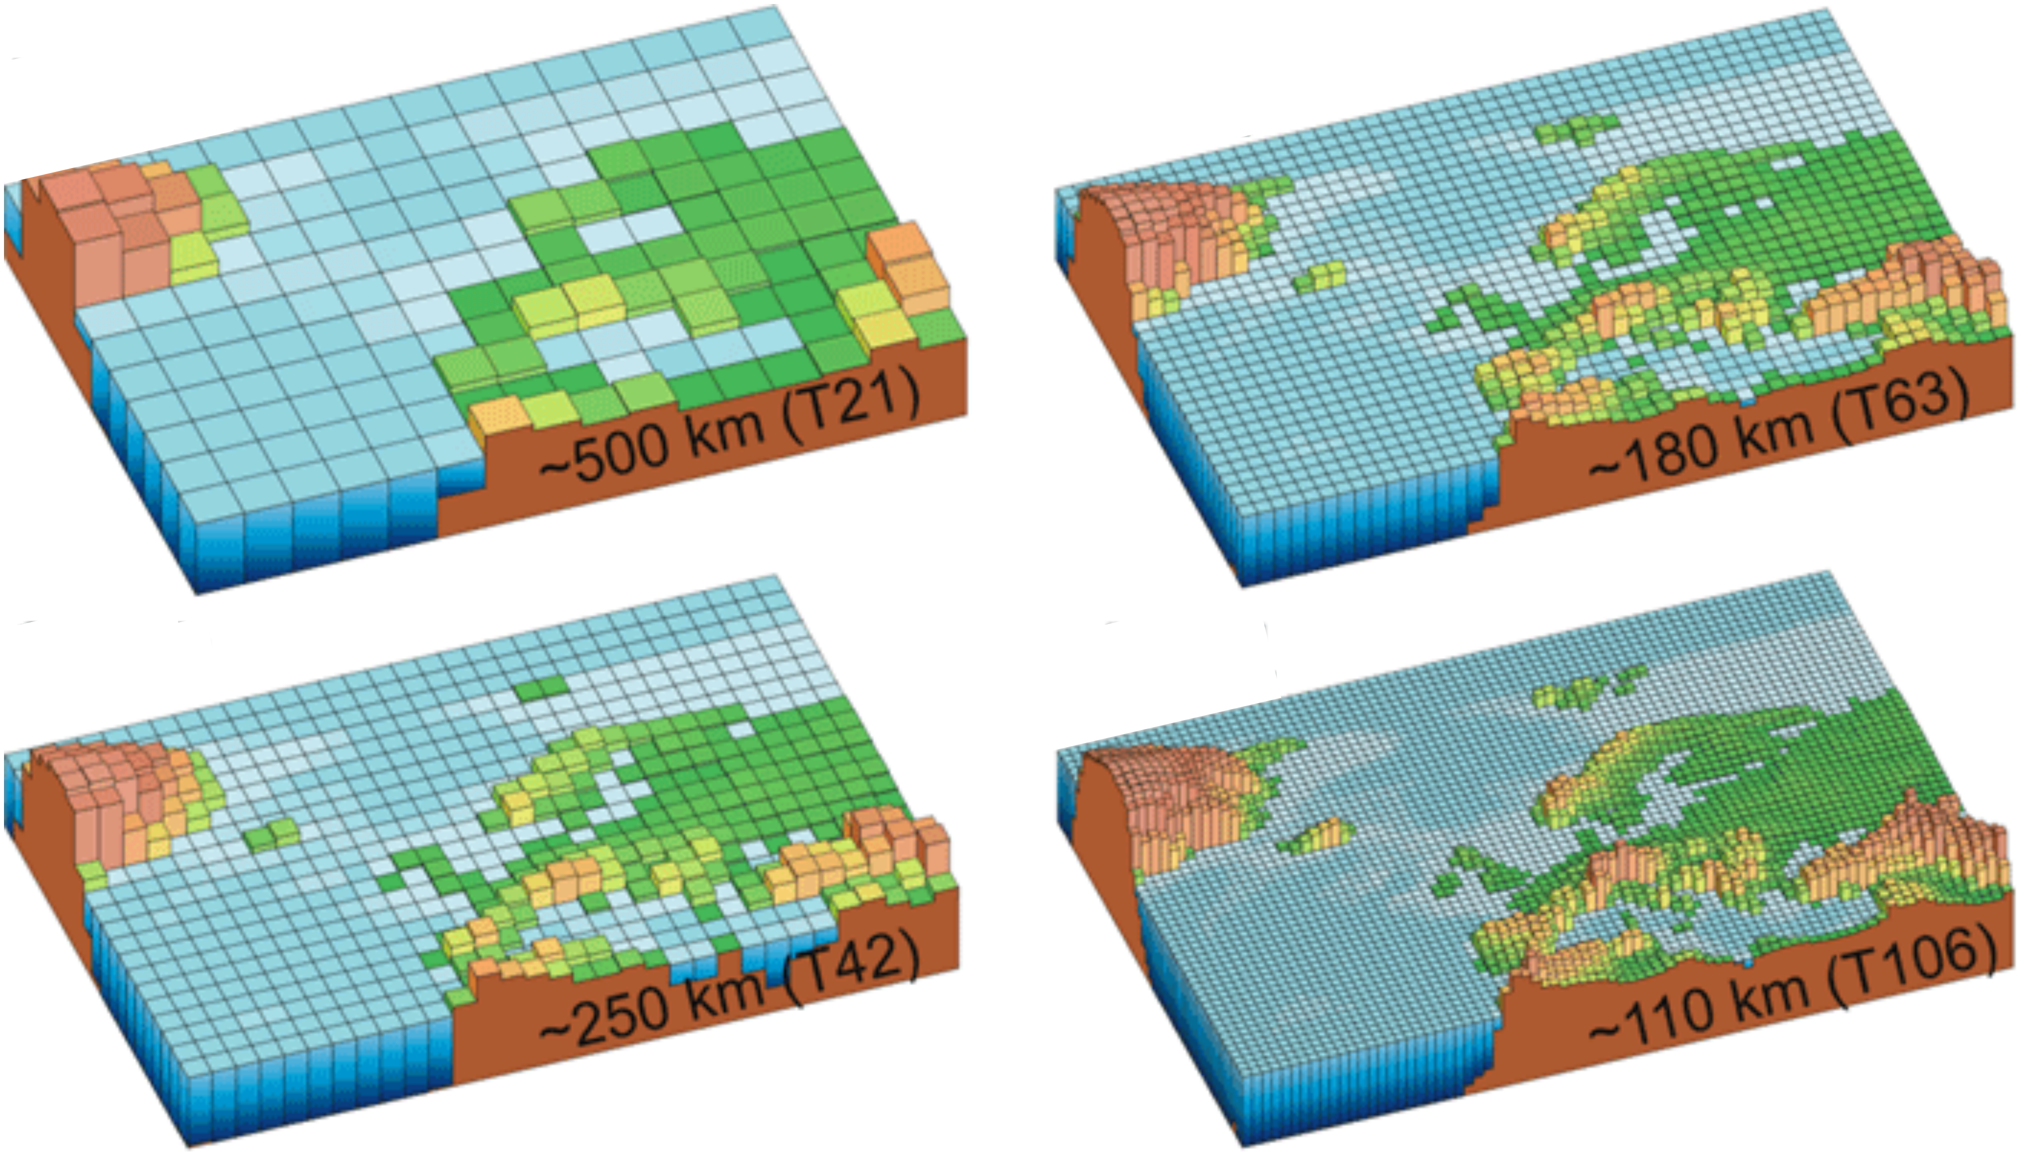

Or Simulation of a Tornado-Producing Super-Cell Storm

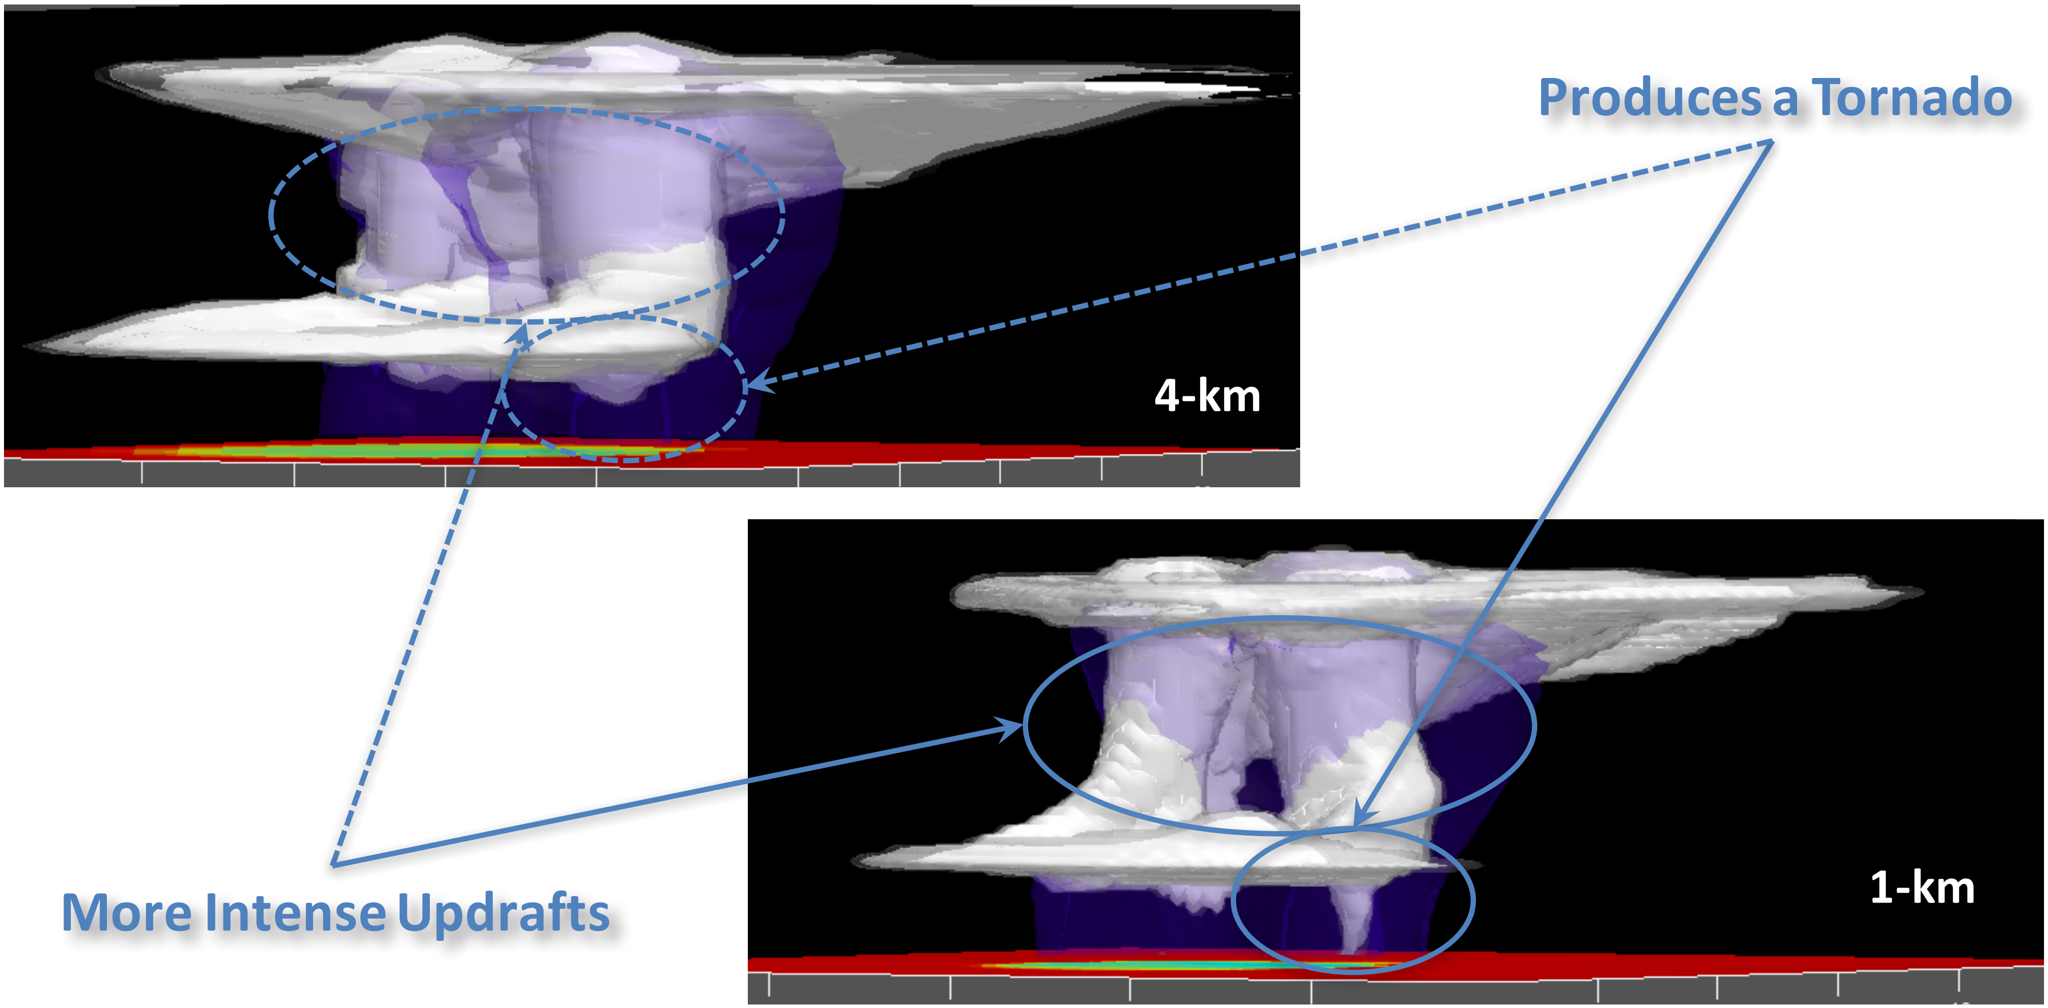

Simulations with GFDL's variable-resolution FV^3, non-hydrostatic (aka cloud-permitting) model. Courtesy of Lin and Harris (2015 manuscript)

Global and Mesoscales are Converging

- State-of-the-art operational *mesoscale* (or regional) models:
  - Hourly, 3 km resolution
  - 1-2 day forecasts
  - Continental US (“CONUS”) domain

- State-of-the-art operational *global* models:
  - Hourly – 6-hourly, 10-13 km resolution
  - 14-16 day forecasts
  - Global

II. NOAA's Next Generation Global Prediction System (NGGPS) [1]

'World's Best Global Forecast Guidance'

http://www.weather.gov/sti/stimodeling_nggps

[1]From NWS Budget Initiative proposal to OMB

NGGPS Goals and Objectives[1]

- **Design/Develop/Implement NGGPS global atmospheric prediction model**
  - Non-hydrostatic scalable dynamics
  - Accelerated physics improvement profile
- **Improve data assimilation**
- **Position NWS for next generation high performance computing**
- **Engage community in model/components development**
- **Reduce implementation time**
- **Increase effectiveness of product distribution**
  - Post-processing, assessments, and display

[1]From NWS Budget Initiative proposal to OMB

Implementation Plan: NGGPS Teams

Implementation Teams

Team are comprised of subject matter experts from across NOAA line offices/laboratories, Navy, NASA, and UCAR. Implementation teams provide input on the direction of team plans, and have an EMC lead or co-lead. Implementation team plans are listed below.

Implementation Team Plans	Team Lead(s)
Atmospheric Prediction - Dynamics	Jeff Whitaker(ESRL/PSD) Vijay Tallapragada (NCEP/EMC)
Atmospheric Prediction - Physics (pdf)	Jim Doyle (NRL Monterey) Bill Kuo (DTC/NCAR)
Aerosols and Atmospheric Composition	Ivanka Stajner (NWS/STI) Yu-Tai (NCEP/EMC)
Atmospheric Data Assimilation (pdf)	John Derber (NCEP/EMC)
Marine Prediction (incl ocean, waves, sea ice, and marine data assimilation) (pdf)	Avichal Mehra (NCEP/EMC)
Land Surface Prediction and land data assimilation (pdf)	Mike Ek (NCEP/EMC)
Nesting (includes hurricanes and convective systems) (pdf)	Vijay Tallapragada (NCEP/EMC)
Post-Processing (pdf)	Matthew Peroutka (NWS/MDL) Yuejian Zhu (NCEP/EMC)
Ensemble Development (pdf)	Tom Hamill (ESRL/PSD) Yuejian Zhu (NCEP/EMC)
Overarching System (architecture/integration incl NEMS/ESMF) (pdf)	Cecelia DeLuca (ESRL/CIRES) Mark Iredell (NCEP/EMC)
Infrastructure	Vijay Tallapragada (NCEP/EMC)
Verification and Validation (pdf)	Ivanka Stajner (NWS/STI) Glenn White (NCEP/EMC)
Testbeds (pdf)	Paula Davidson (NWS/STI) Mike Ek (NCEP/EMC)

http://www.weather.gov/sti/stimodeling_nggps_implementation

NGGPS Global Atmospheric Model Technical Strategy

- **Reduce implementation time and risk by separating dynamic core and model physics (modular framework)**
- **Identify and implement optimal core for global weather forecast applications**
  - Highly scalable
  - Non-hydrostatic
- **Accelerate evolution of model physics**
  - Develop/Implement Common Community Physics Package (CCPP)
    - Based on current GFS physics package
    - Integration of best of other existing physics packages
    - scale-aware
- **Integrate with advanced data assimilation**
  - Joint Effort for Data assimilation Integration (JEDI)
- **Develop a new community approach**
  - Employ Global Modeling Test Bed (GMTB)/Developmental Testbed Center (DTC) and Joint Center for Satellite Data Assimilation (JCSDA) to encourage and facilitate community interaction
  - Accelerate O2R & R2O

NGGPS Atmospheric Model Phased Implementation Approach

- **Phase 1 (FY15) – Identify Qualified Dynamic Cores**
  - Evaluate technical performance
    - Scalability
    - Integration of scheme stability and characteristics
- **Phase 2 (FY16) – Select Candidate Dynamic Core**
  - Integrate with operational GFS Physics/Common Community Physics Package (CCPP)
  - Evaluate meteorological performance
- **Phase 3 (FY17-19) – Dynamic Core Integration and Implementation**
  - Implement candidate dynamic core in NEMS
  - Implement Common Community Physics Package
  - Implement data assimilation (4DEnVar with 4D incremental analysis update and stochastic physics)
  - Implement community model environment

Testing Process: NGGPS Dycore Test Group (DTG)

- Ming Ji, Chair
  - Director, Office of Science and Technology Integration
- Fred Toepfer
  - NGGPS Program Manager
- Tim Schneider
  - Acting NGGPS Program Manager
- Bob Gall
  - Independent Consultant
- Ricky Rood
  - Independent Consultant
- John Thuburn
  - Independent Consultant

DTG operates by consensus
- Develop & approve plans
- Execute, review, assess tests & results

- Melinda Peng/Jim Doyle
  - Navy/NRL Monterey
- Ram Ramaswamy/SJ Lin
  - GFDL
- Hendrik Tolman/Vijay Tallapragada
  - NCEP/EMC
- Chris Davis/Bill Skamarock*
  - NCAR/MMM
- Kevin Kelleher/Stan Benjamin
  - ESRL/GSD
- Jeff Whitaker
  - NGGPS Test Manager
- John Michalakes
  - Chair, Advanced Computing Evaluation Committee

* NCAR ceased participation and withdrew from DTG on 20 May 2016

Testing Process: Candidate Models

Phase 1 Candidate Dynamic Cores*:

* Built upon HIWPP Non-hydrostatic Model Evaluation

- Non-hydrostatic Global Spectral Model (**GSM**) - **EMC**
- Global Non-hydrostatic Mesoscale Model (**NMM** & **NMM-UJ**) - **EMC**
- Model for Prediction Across Scales (**MPAS**) - **NCAR**
- Non-hydrostatic Icosohedral Model (**NIM**) – **ESRL**
- Navy Environmental Prediction System Using the NUMA Core (**NEPTUNE**) – **Navy**
- Finite Volume Model on the Cubed Sphere (**FV3**) – **GFDL**

- **FV3 and MPAS selected to advance to Phase 2**

Phase 2:

- ***July, 2016: NOAA selected the FV3 dynamical core***

NGGPS Phase 2 Test Plan

#	Evaluation Criteria
1	Plan for relaxing shallow atmosphere approximation (deep atmosphere dynamics)
2	Accurate conservation of mass, tracers, entropy, and energy
3	Robust model solutions under a wide range of realistic atmospheric initial conditions using a common (GFS) physics package
4	Computational performance with GFS physics
5	Demonstration of variable resolution and/or nesting capabilities, including supercell tests and physically realistic simulations of convection in the high-resolution region
6	Stable, conservative long integrations with realistic climate statistics
7	Code adaptable to NEMS/ESMF
8	Detailed dycore documentation, including documentation of vertical grid, numerical filters, time-integration scheme and variable resolution and/or nesting capabilities
9	Evaluation of performance in cycled data assimilation
10	Implementation Plan (including costs)

*** Results of evaluation for above criteria provided in background information**

Summary of Phase 2 Test Results

- Testing yielded sufficient information to evaluate both dynamic cores and produce a low risk recommendation without compromising performance or skill
- Summary of results:
  - Computationally, FV3 is more than twice as fast as MPAS with equivalent resolution
  - Full forecast experiments with GFS initial conditions and GFS physics showed significant differences between FV3 and MPAS, FV3 almost equivalent to GFS* (some stability issues with MPAS forecasts)
  - FV3 performs comparable to the GFS in cycled data assimilation test (without tuning, at reduced resolution)*, MPAS performance inferior to GFS
  - Effective resolution for both dynamic cores is found to be similar, and higher than GFS
  - High-resolution idealized and real-data simulations show qualitatively similar results in simulations of explicit moist convection
  - Cost to implement FV3 is significantly less than MPAS in terms of manpower and computational resources

* Dr. Andy Brown, Director of Science at the UK Met Office, called this a "remarkable result" and indicated that he's never seen this happen before over several instances of implementing a new operational model.

Phase 3: Implementation Plan of FV3 Dynamic Core in NEMS GFS (FY17-FY19)

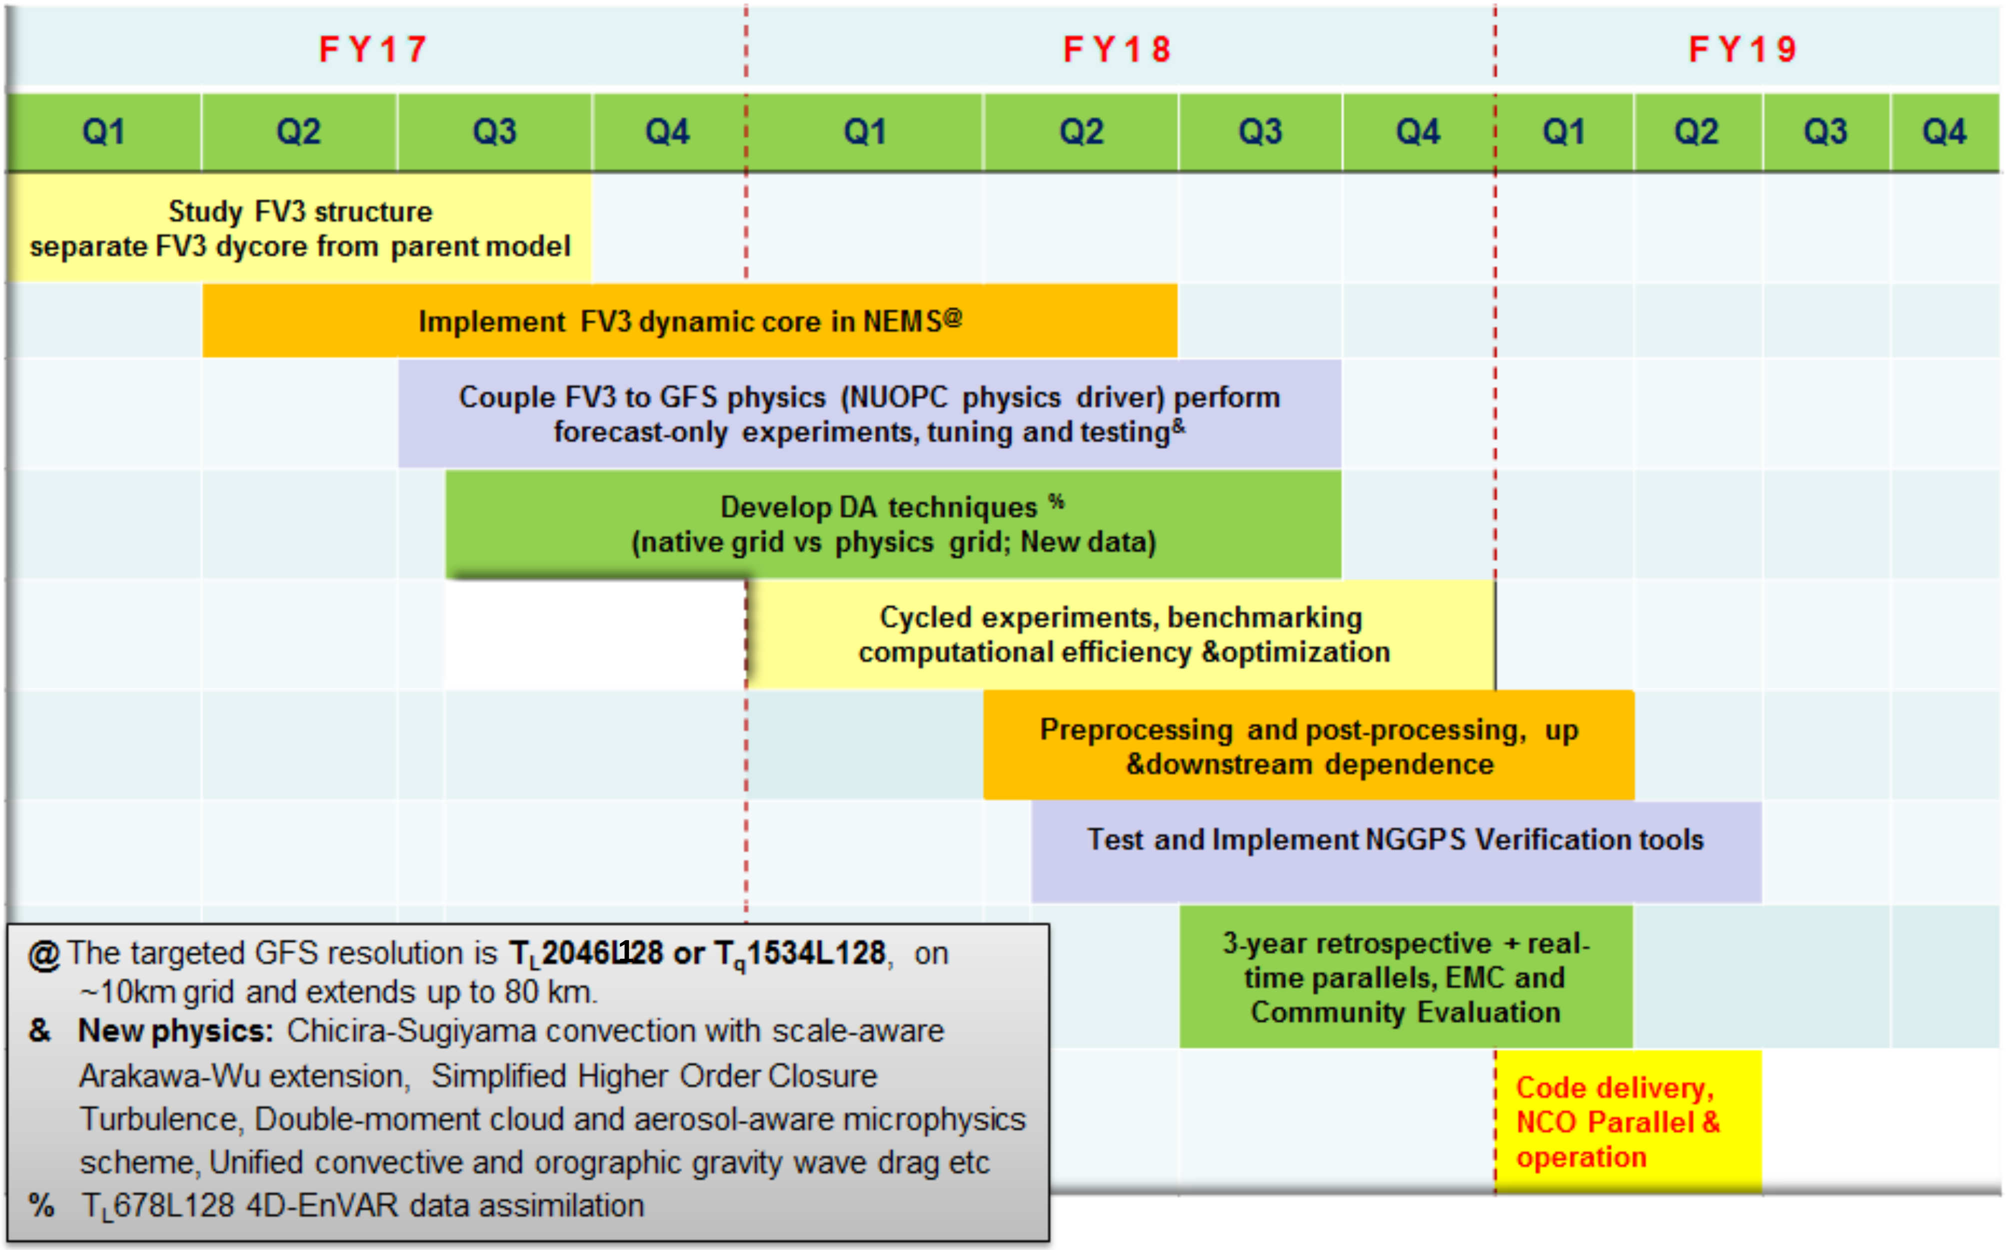

Schematic of NEMS GFS

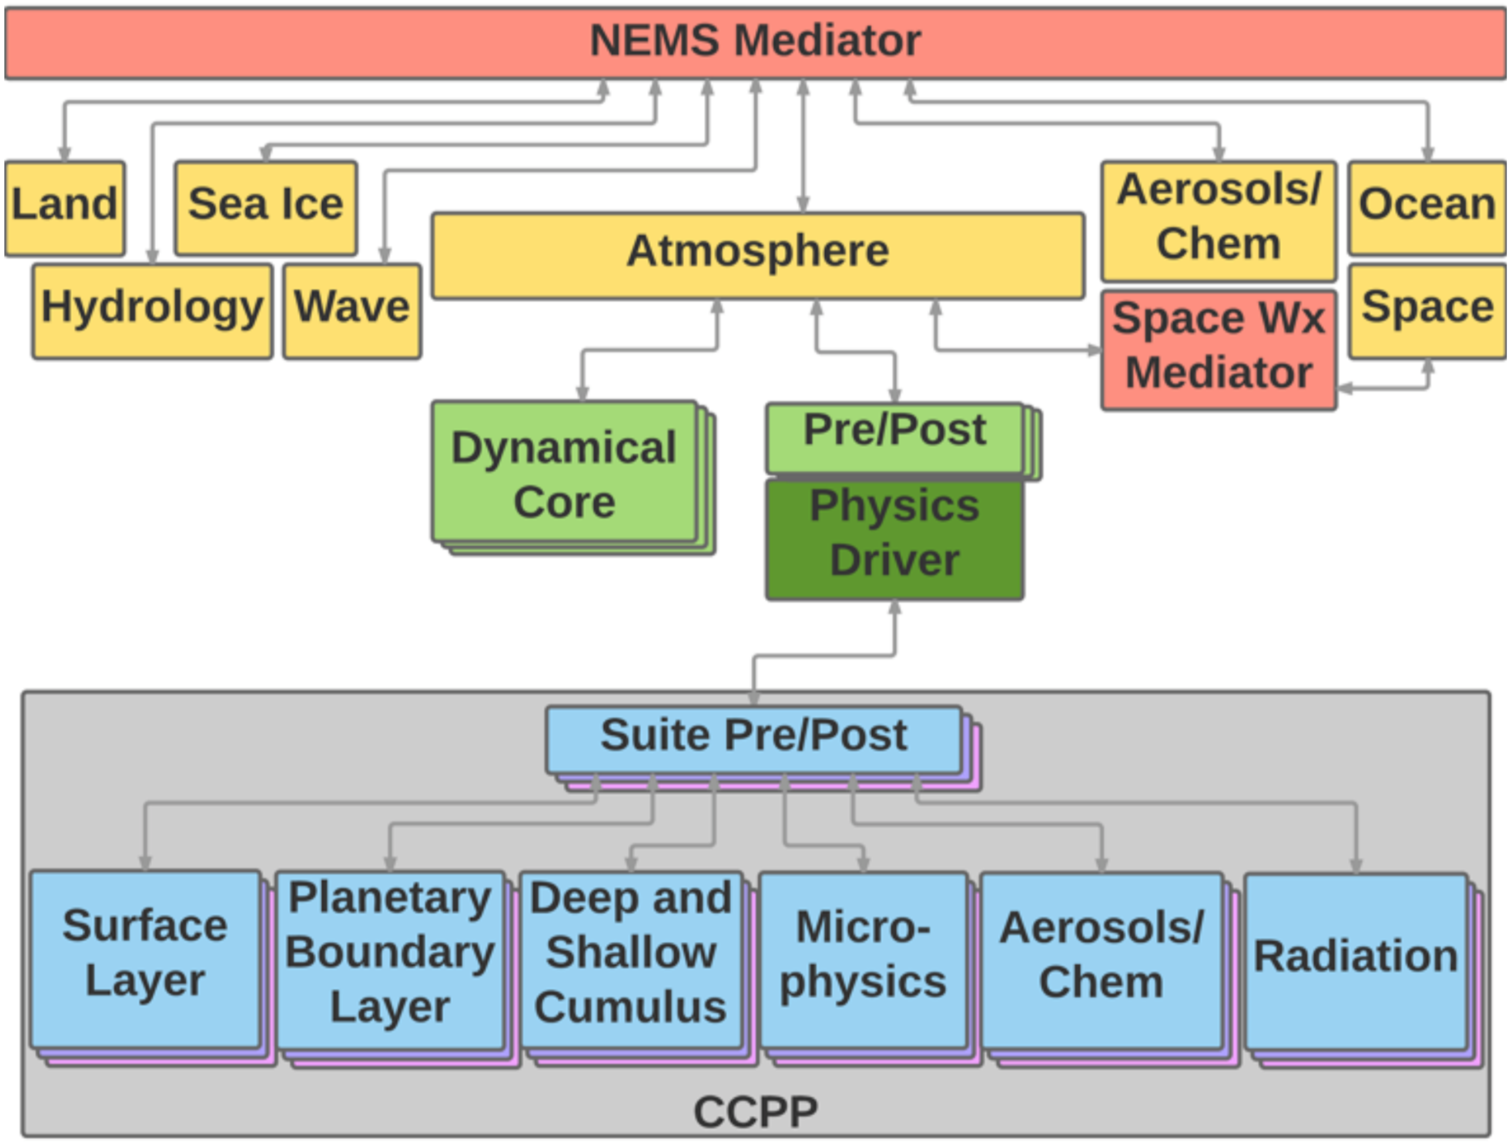

Diagram of the proposed Common Community Physics Package (CCPP)

III. Finite Volume on a Cubed Sphere (FV3)

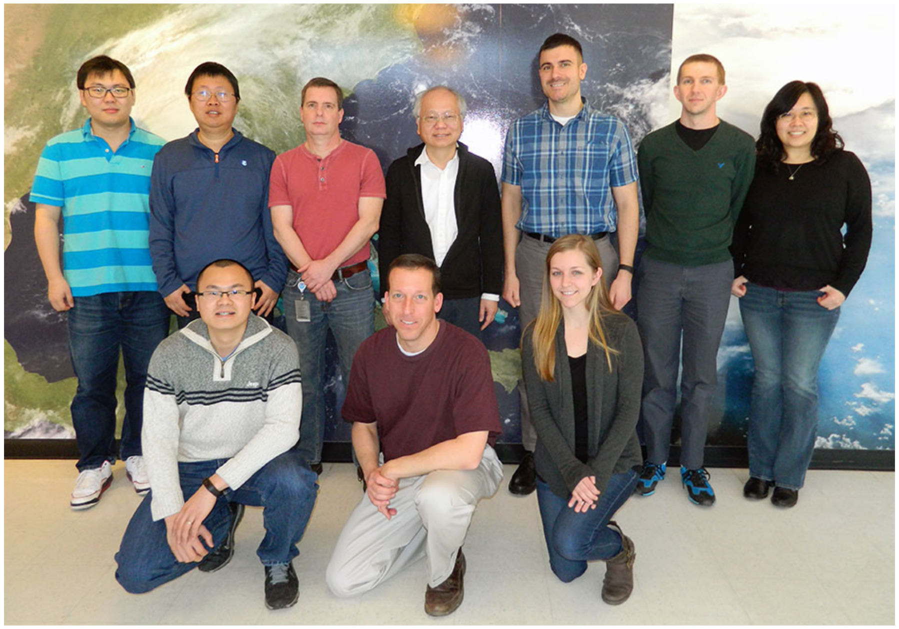

GFDL FV3 team

- FV concept for global models matured in NASA-Goddard in the late 1990s.
- GFDL advanced this core in 2004, developed first for “Climate” modeling, and rapidly finding applicability for “Weather” modeling.

The GFDL FV3: Finite-Volume Dynamical Core on the Cubed-Sphere

A unified regional-global modeling framework for weather & climate

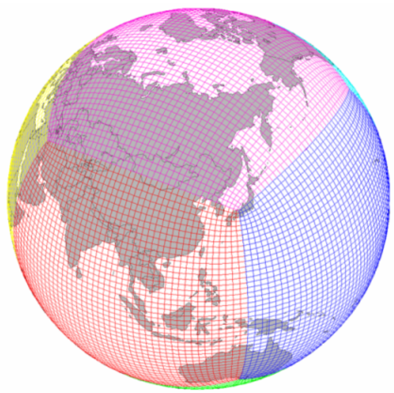

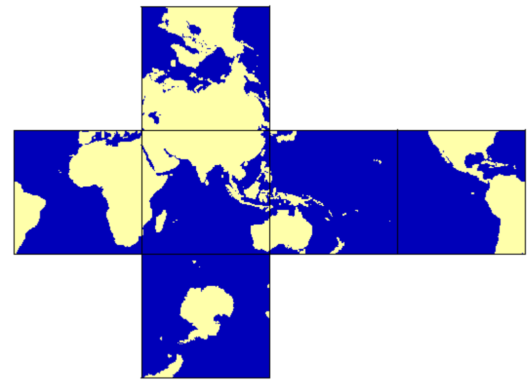

- **High-order Finite-Volume algorithms with a true Lagrangian vertical coordinate**

- Hydrostatic or **non-hydrostatic**: applicable to all scales from LES to climate

- **Designed for general fine-grained parallelism** (MPI, openMP, and GPU*); scalable to ~million cores (IBM BG) at cloud-resolving resolution

- **Configurable as a regional model or a global model** with two variable resolution options (two-way nested or stretched grid)

**A working GPU version has been developed by W. Putman (NASA/GMAO)*

Draft requirements for a NOAA unified modeling system:

Dynamical Core:

- Highly scalable & Conservative algorithm (mass and total energy)
- Monotonic transport – important for advanced cloud micro-physics
- Applicable from "low-resolution" climate simulations (~ 100 km) to cloud-resolving resolution (1 km)
- Flexible design for both regional and global applications

Physics:

- Scale-aware, applicable from 1 km to 100 km, and must be skillful for gray-zone (3-10 km): next frontier for medium-range NWP
- Interactive dust-aerosols, and cloud micro-physics with prognostic precipitation (rain, snow, and hail)
- Advanced PBL parameterizations (TKE, EDMF, SHOC, …) and gray-zone convection

"FV core" important attributes

- Conservation properties
- Fidelity extending to very high spatial resolutions
- Computational speed-up of the model

Building a NOAA Unified Weather-Climate Modeling-Prediction System

Simulations & predictions of high-impact weather-climate events from the storm-scale to the decadal scale (tornadoes, thunderstorms, hurricanes, MJOs, flood/droughts, ENSO) within a unified weather-climate modeling system

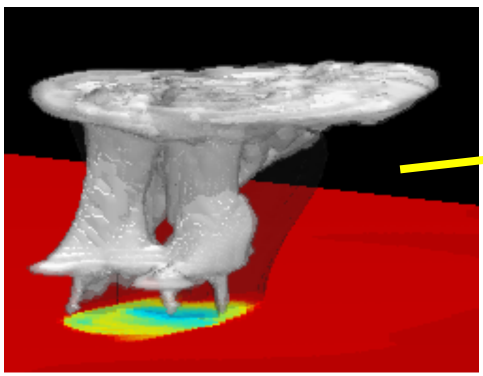

GFDL Super High Resolution Atmosphere Model (Super HiRAM, with FV3 core)

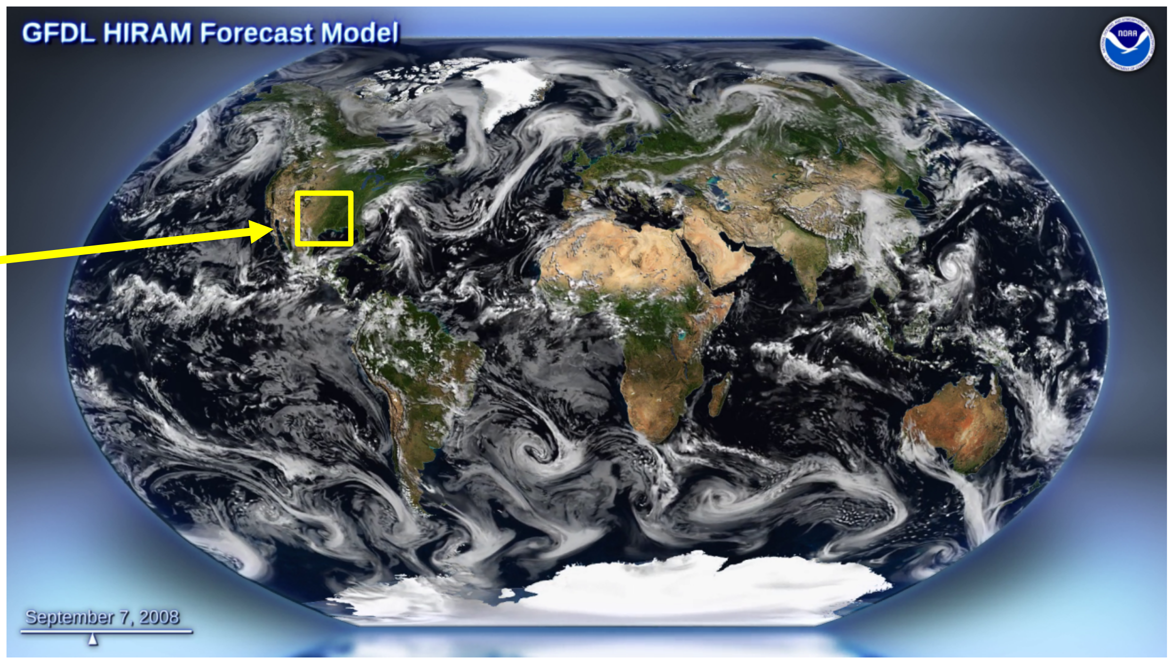

- **Most climate models do not have the non-hydrostatic dynamics nor the physical parameterizations to represent <u>all-scale</u>**
- **A unified weather-climate modeling-prediction system is being developed and evaluated at OAR/GFDL**

High resolution atmospheric modeling with FV3 (began in ~2008)

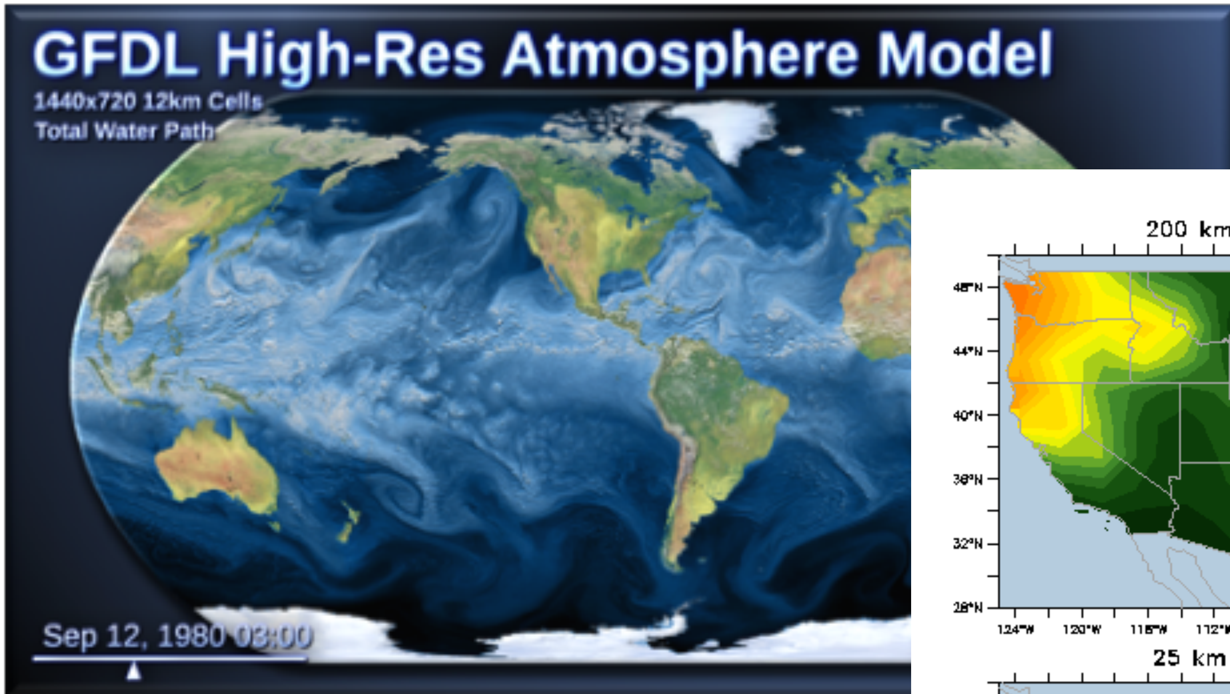

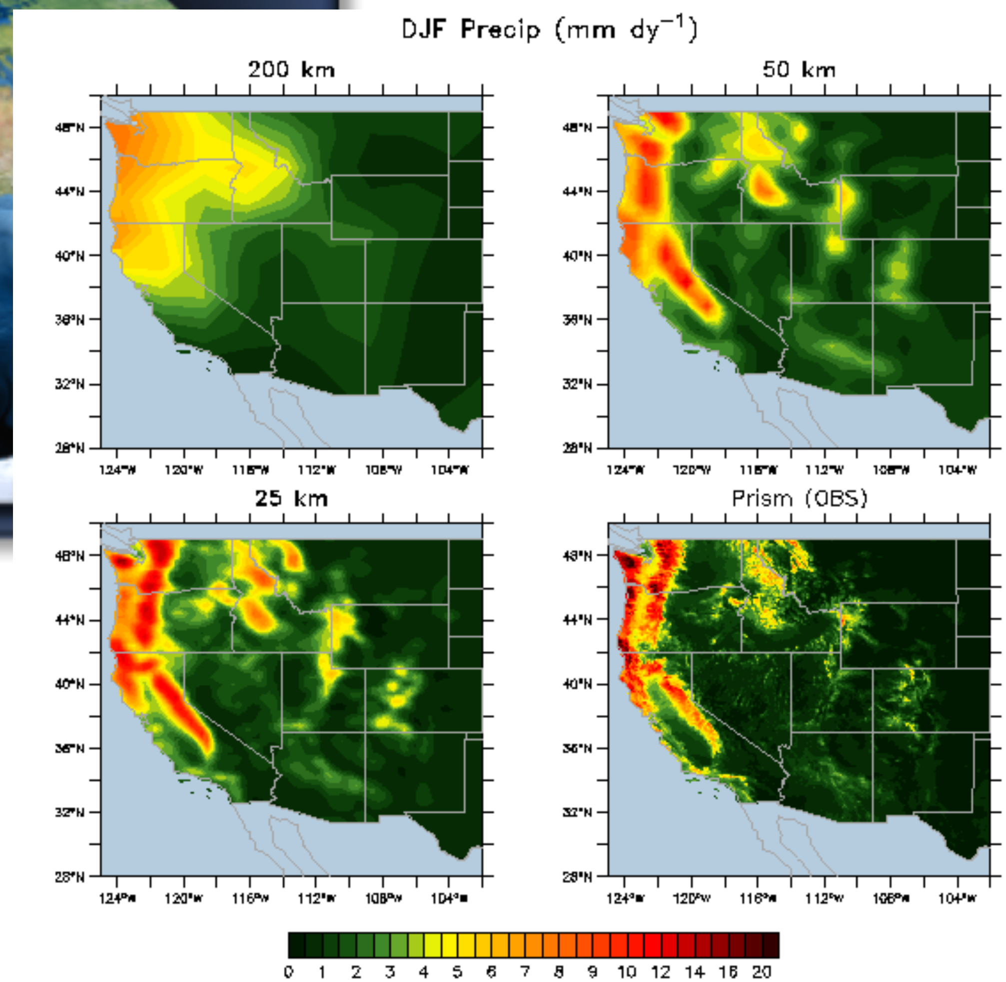

C720 (12.5 km)

GFDL's affordable ultra-high resolution regional-global model

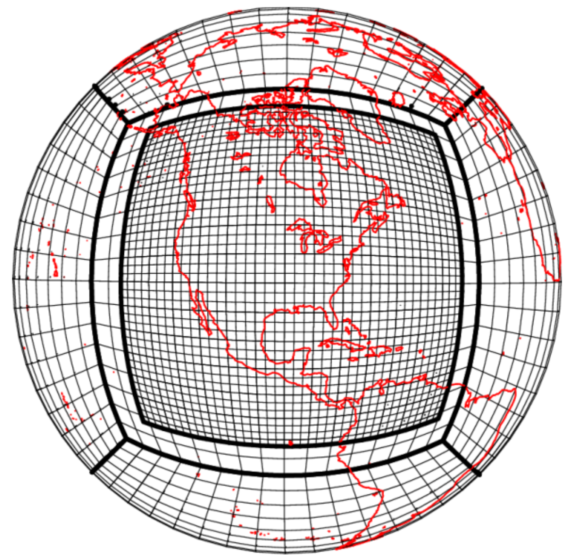

A. Nested regional-global model

- Regional model can be run independently (one-way) or two-way nested with the global model
- Suitable for weather & climate

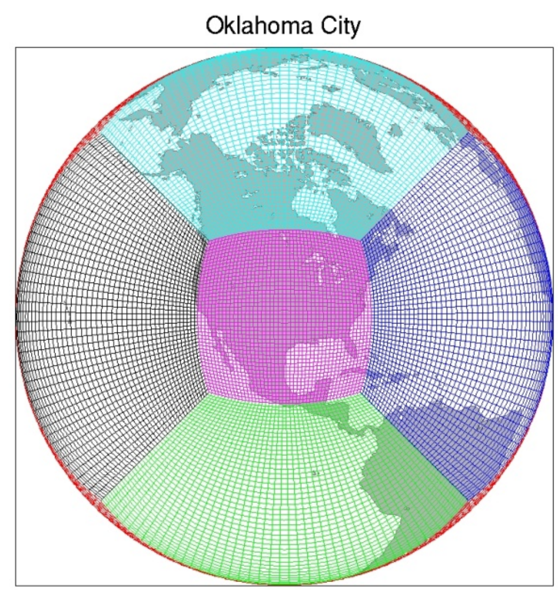

B. Variable resolution global model

- Single model framework with smooth transition in resolution (e.g., 3.5 km over CONUS stretched to 30 km over Indian ocean)
- Suitable for weather & climate

Integrated regional-global prediction with a 3-km nest hurricane Patricia (2015)

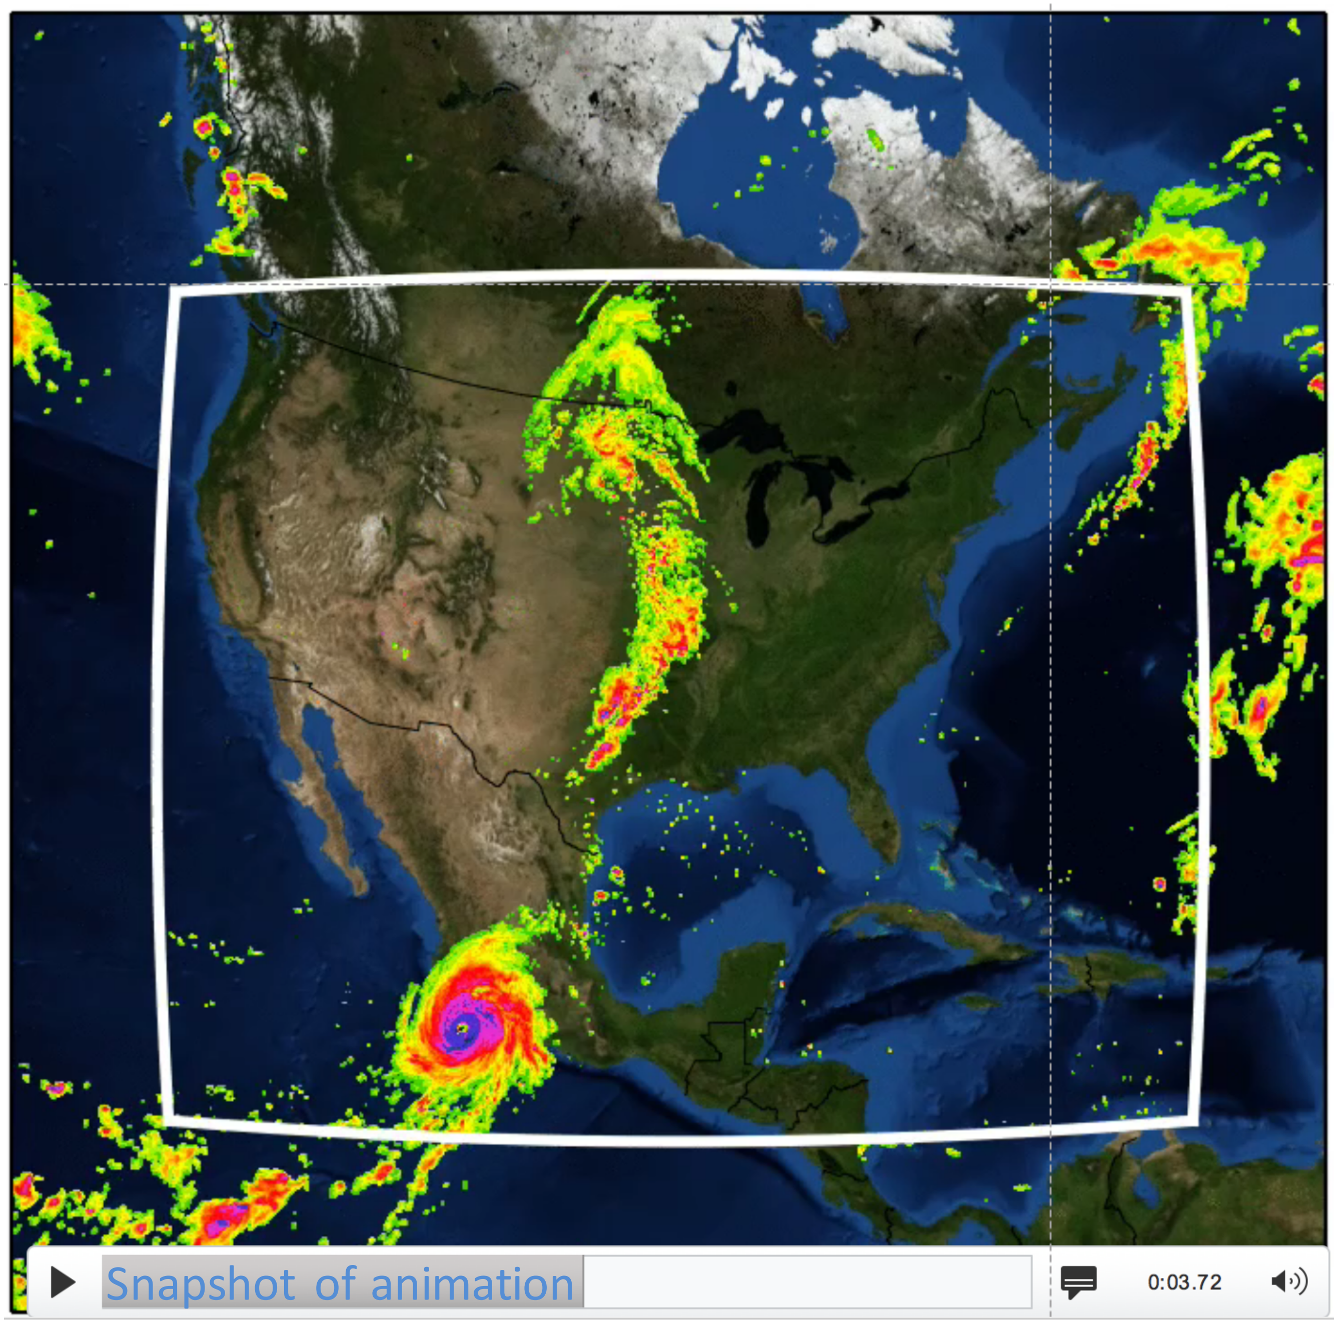

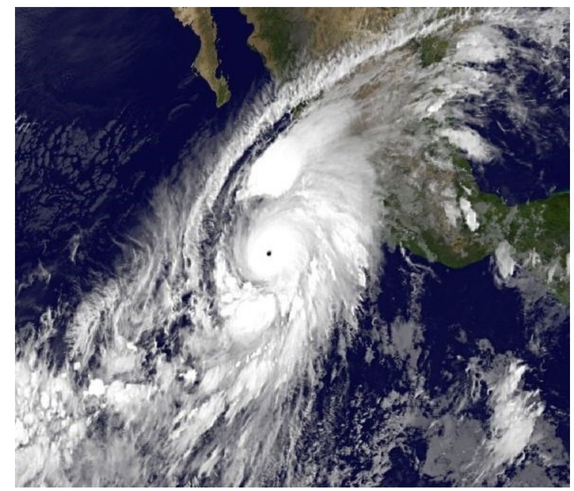

Hurricane Patricia was the most intense tropical cyclone on record in the Western Hemisphere. It intensified from a tropical storm to a cat-5 hurricane in 24 hours.

Simulation of of Tornado-like vortices using GFDL's variable-resolution global model (with FV3)

Computational cost:

3-hour simulation needs ~ 1 hour (wall clock) using only 384 CPUs – computationally trivial

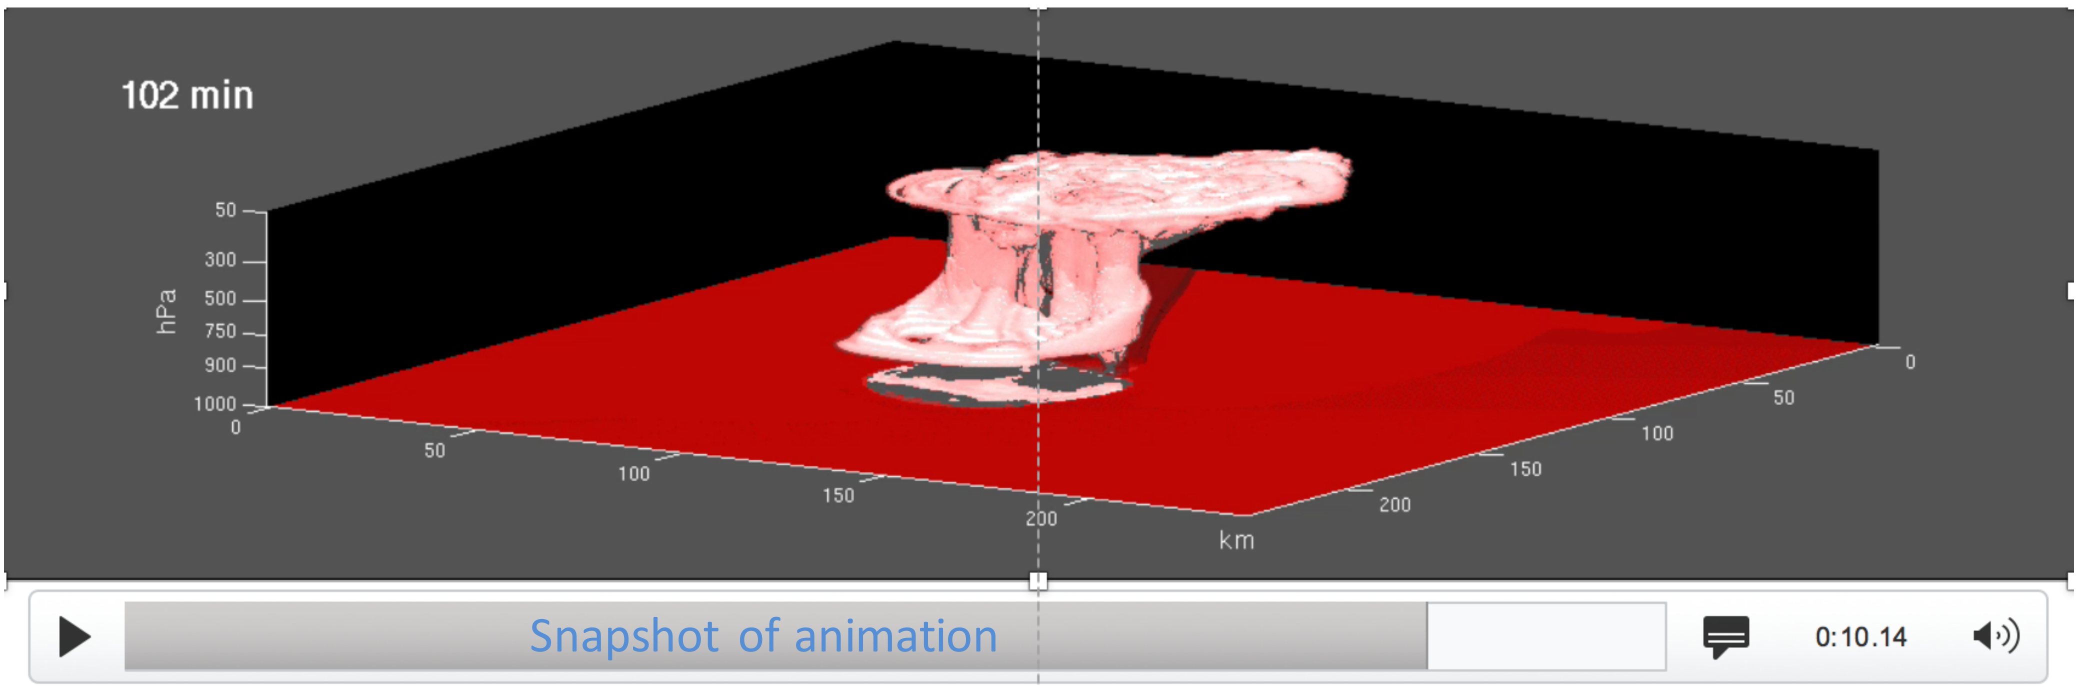

Darker shade: *rain water;* Lighter shade: *cloud liquid water*

Bottom: lowest layer air temperature (illustrating cold pool)

Animation by Lucas Harris, NOAA/GFDL

III. Advances in High-Performance Computing (HPC)

Enabler: High Performance Computing

New supercomputing funded in part by the Sandy Supplemental Disaster Relief Act:

- Enabling Sandy Supplemental R&D
- 30% Increased Capacity

Theia

- Weather & Climate Supercomputing System
- 2 petaflop increase

Operations

- Accelerator technology
- Exploratory project

Fine Grain

Is Moore's Law Finally Abandoning Us?

(and can we afford the energy/cooling costs for HPC?)

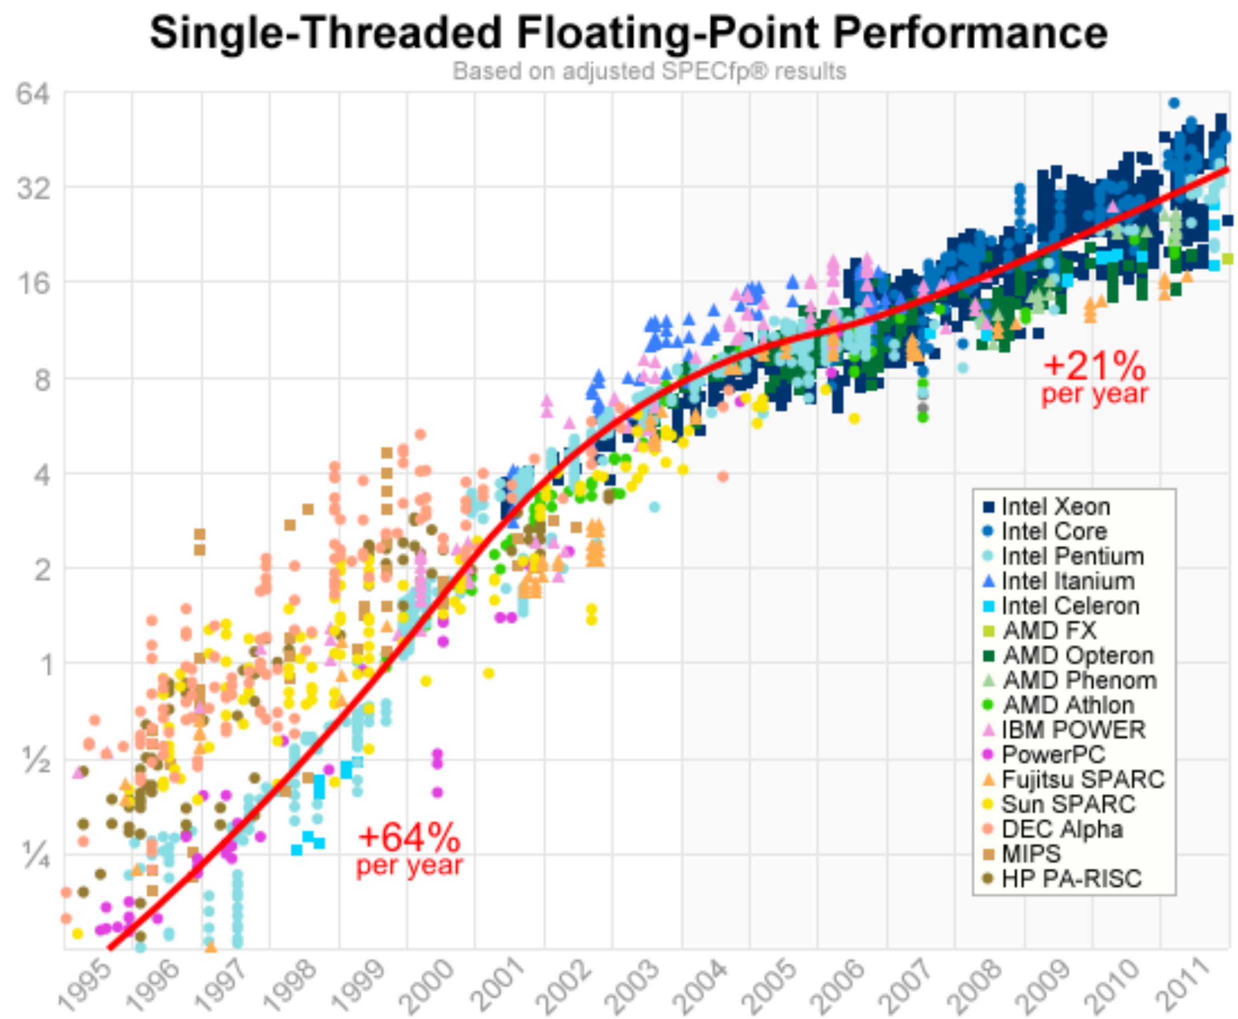

http://preshing.com/20120208/a-look-back-at-single-threaded-cpu-performance/

Massively Parallel Fine Grain (MPFG) - Defined -

Govett, et al., Parallelization and Performance of the NIM Weather Model on CPU, GPU and MIC Processors, In review, BAMS

- Massively Parallel – refers to systems with hundreds of thousands of compute cores
- Fine-Grain = refers to the ability for applications to execute millions of calculations every clock cycle

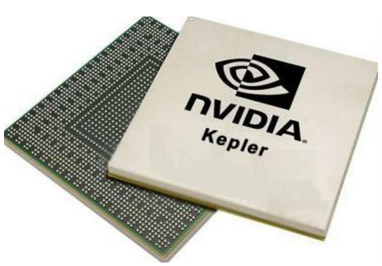

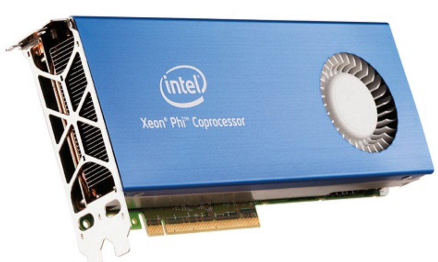

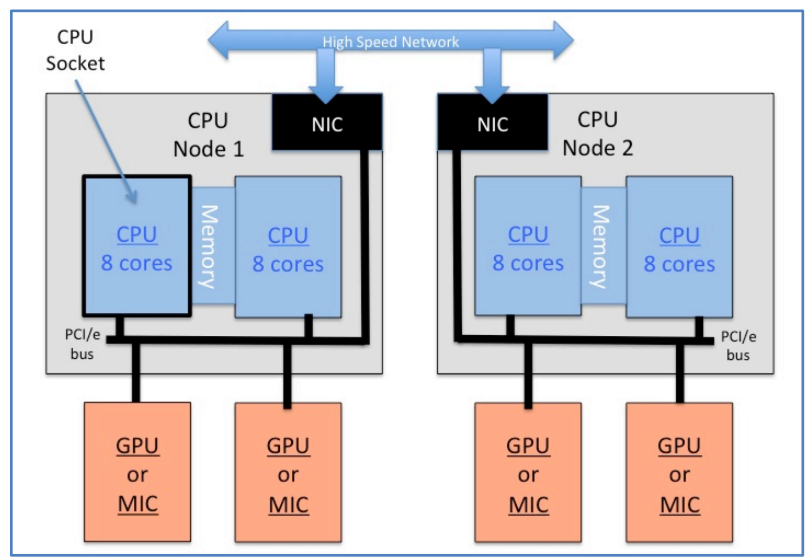

What would it take …

To run the selected NGGPS model at 3km resolution globally?

- Using Non-hydrostatic Icosahedral Model (NIM) dycore performance as a guide
- And assuming perfect scaling:
  - 3KM resolution need ~1500 NVIDIA K80 GPUs to run in 1.6% of real-time
  - Add ~40% more for physics => ~ 2200 GPUs
  - Double to get ~0.8% real-time 4400 GPUs
  - Factor a 2X application benefit for Pascal (or MIC KNL)

~ 2200 GPUs (or MIC) processors

(or ~3200 processors for 8 mins/day)

HPC: the Path Forward:

- Fine-Grain computing is the future (at least for the next decade)
- Applications may need to be modified to run efficiently on GPU and MIC processors
- GPU programming is getting simpler
  - Similar to OpenMP programming for the CPU
- Single source code for CPU, GPU and MIC should be a goal, and is achievable
  - But there are challenges (structuring parallelism; inter-process communications; I/O)
- New NOAA R&D MPFG System in October 2016

Thank You