Purpose-Built High-Performance Computing for Earth System Models

8th NCAR MultiCore Workshop (MC8)
September 18, 2018

Daniel Duffy
Computational and Informational Sciences and Technology Office (CISTO –Code 606)
NASA Center for Climate Simulation (NCCS – Code 606.2)
daniel.q.duffy@nasa.gov and @dqduffy

Special Thanks

- **Tsengdar Lee**
  - NASA High End Computing (HEC) Program Manager and Weather Lead, NASA Headquarters
- **Tom Clune**
  - Global Modeling and Assimilation Office (GMAO), NASA Goddard Space Flight Center (GSFC)
- **John Michalakes**
  - University Corporation for Atmospheric Research (UCAR)
- **Special Thanks**
  - Bill Putman
  - Matt Thompson
  - Mark Govett
  - Earth System Prediction Capability (ESPC) Working Group

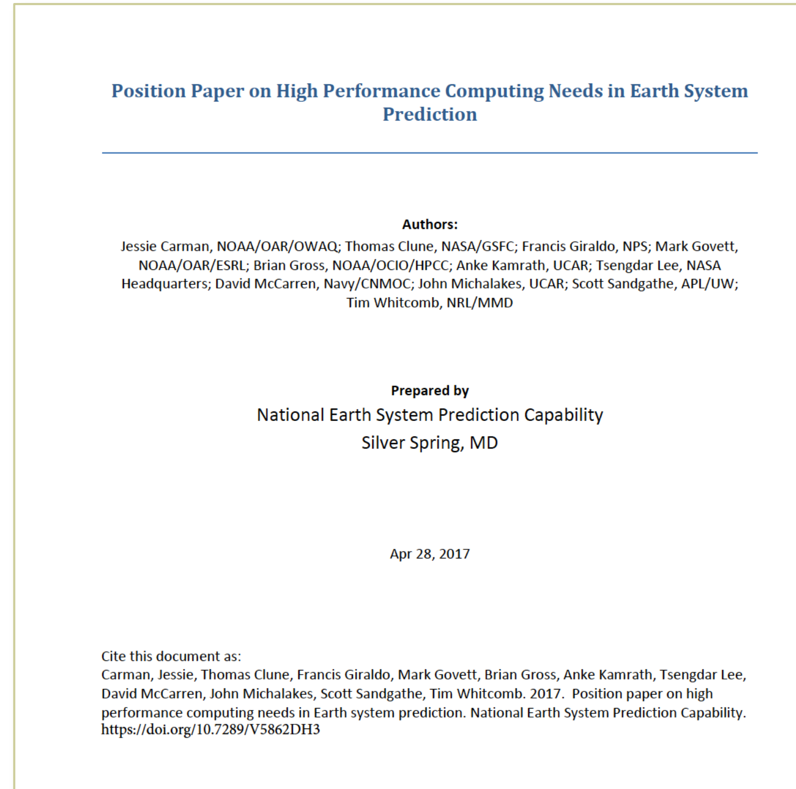

Position Paper on High Performance Computing Needs in Earth System Prediction

Authors:
Jessie Carman, NOAA/OAR/OWAQ; Thomas Clune, NASA/GSFC; Francis Giraldo, NPS; Mark Govett, NOAA/OAR/ESRL; Brian Gross, NOAA/OCIO/HPCC; Anke Kamrath, UCAR; Tsengdar Lee, NASA Headquarters; David McCarren, Navy/CNMOC; John Michalakes, UCAR; Scott Sandgathe, APL/UW; Tim Whitcomb, NRL/MMD

Prepared by
National Earth System Prediction Capability
Silver Spring, MD

Apr 28, 2017

Cite this document as:
Carman, Jessie, Thomas Clune, Francis Giraldo, Mark Govett, Brian Gross, Anke Kamrath, Tsengdar Lee, David McCarren, John Michalakes, Scott Sandgathe, Tim Whitcomb. 2017. Position paper on high performance computing needs in Earth system prediction. National Earth System Prediction Capability. https://doi.org/10.7289/V5862DH3

An Interagency Study on Purpose-Built HPC

- **The National ESPC HPC working group advocates for an interagency study investigating:**
  - The widening gap between earth system application requirements and currently evolving HPC
  - A hypothetical supercomputer designed with the singular purpose of running Exascale Earth system prediction models
- **This study will:**
  - Help identify the current needs of earth system prediction models
  - Determine whether or not a purpose-built earth system prediction computer is feasible, from several perspectives, including cost and efficiency
- **Birds of Feather session at SuperComputing 2018 to thoroughly discuss this study with the broader community**

NASA Center for Climate Simulation (NCCS)

Provides an integrated high-end computing environment designed to support the specialized requirements of Climate and Weather modeling.

- High-performance computing, cloud computing, data storage, and networking technologies
- High-speed access to petabytes of Earth Science data
- Collaborative data sharing, publication, and analysis services

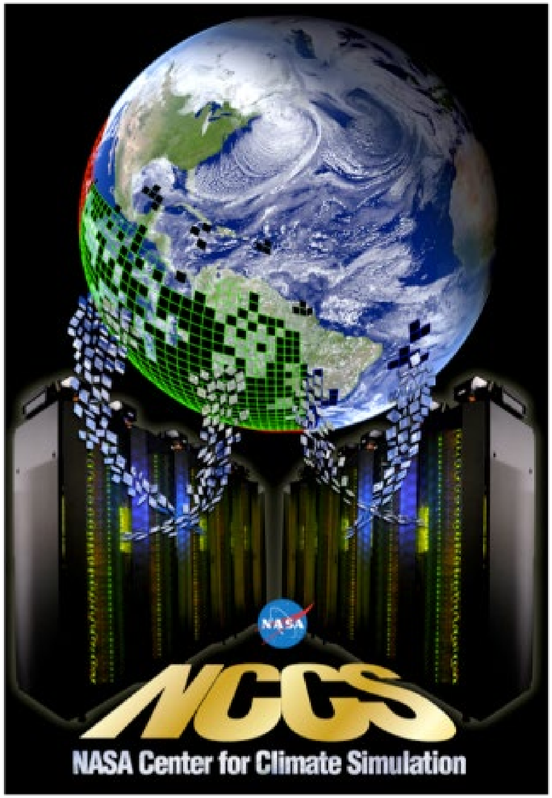

http://www.nccs.nasa.gov

Primary Customers (NASA Funded Science)

- Global Modeling and Assimilation Office (GMAO)
- Land Information Systems (LIS)
- Goddard Institute for Space Studies (GISS)
- Variety of other Research and Development (R&D) and Engineering
- ABoVE, HiMAT, CALET, WFIRST, ICESat2

High-Performance Science

- Funded by the High End Computing (HEC) program under SMD
- Dr. Tsengdar Lee, Program Manager
- Code 606.2 at NASA Goddard Space Flight Center in Greenbelt, MD.

1922, Lewis Fry Richardson Richardson's "Forecast Factory"

In 1922, Lewis Fry Richardson developed the first numerical weather prediction (NWP) system.

He divided the world into grid cells and applied finite difference solutions of differential equations.

His first attempt to calculate weather for a single eight-hour period took six weeks.

He proposed a "forecast-factory" of 64,000 people armed with mechanical calculators lead by a conductor to coordinate the forecast.

Yet even with this fanciful factory, Richardson would only be able to calculate weather about as fast as it actually happened.

[1-day per day]

http://celebrating200years.noaa.gov/foundations/numerical_wx_pred/theater.html
http://www.gfdl.noaa.gov/cms-filesystem-action/user_files/jrl/gcm/jrl_gcm_doc-history.pdf

1950, The First Numerical Weather Simulation

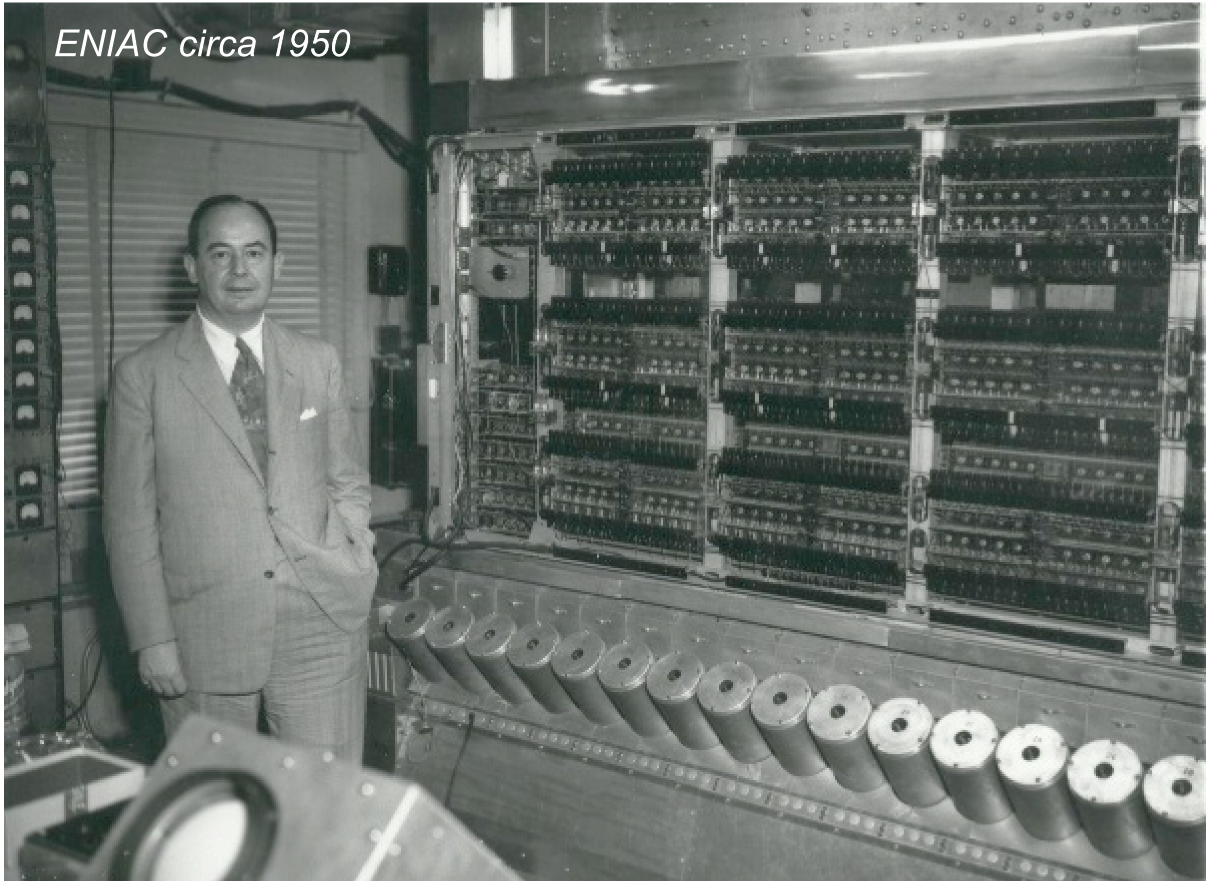

http://www.whyy.org/tv12/franklinfacts/oct1601_6.jpg

Jule Charney and John von Neumann completed a two-dimensional simulation on the ENIAC in 1950.

It covered North America with 270 points about 700 km apart. Starting with real weather data for a particular day, the computer solved all the equations for how the air should respond to the differences in conditions between each pair of adjacent cells.

It took so long between each run to print and sort punched cards that "the calculation time for a 24-hour forecast was about 24 hours, that is, we were just able to keep pace with the weather."

[1-day per day]

2015, 1.5km Global Simulation with GEOS-5 on SCU10 at NCCS

The highest resolution simulation performed with any US global model at that time

- Includes updated 2-moment micro-physics and interactive aerosols/chemistry with GOCART
- Executed on 1,055 nodes of the NCCS Discover SCU10 cluster (1,024 for compute: 31 for IO)
- Scalable Compute Unit 10 (SCU10) was the first 1 PetaFlop system at the NCCS – 1.5km GEOS-5
- ***Completed ~1 simulated-day/wallclock-day***

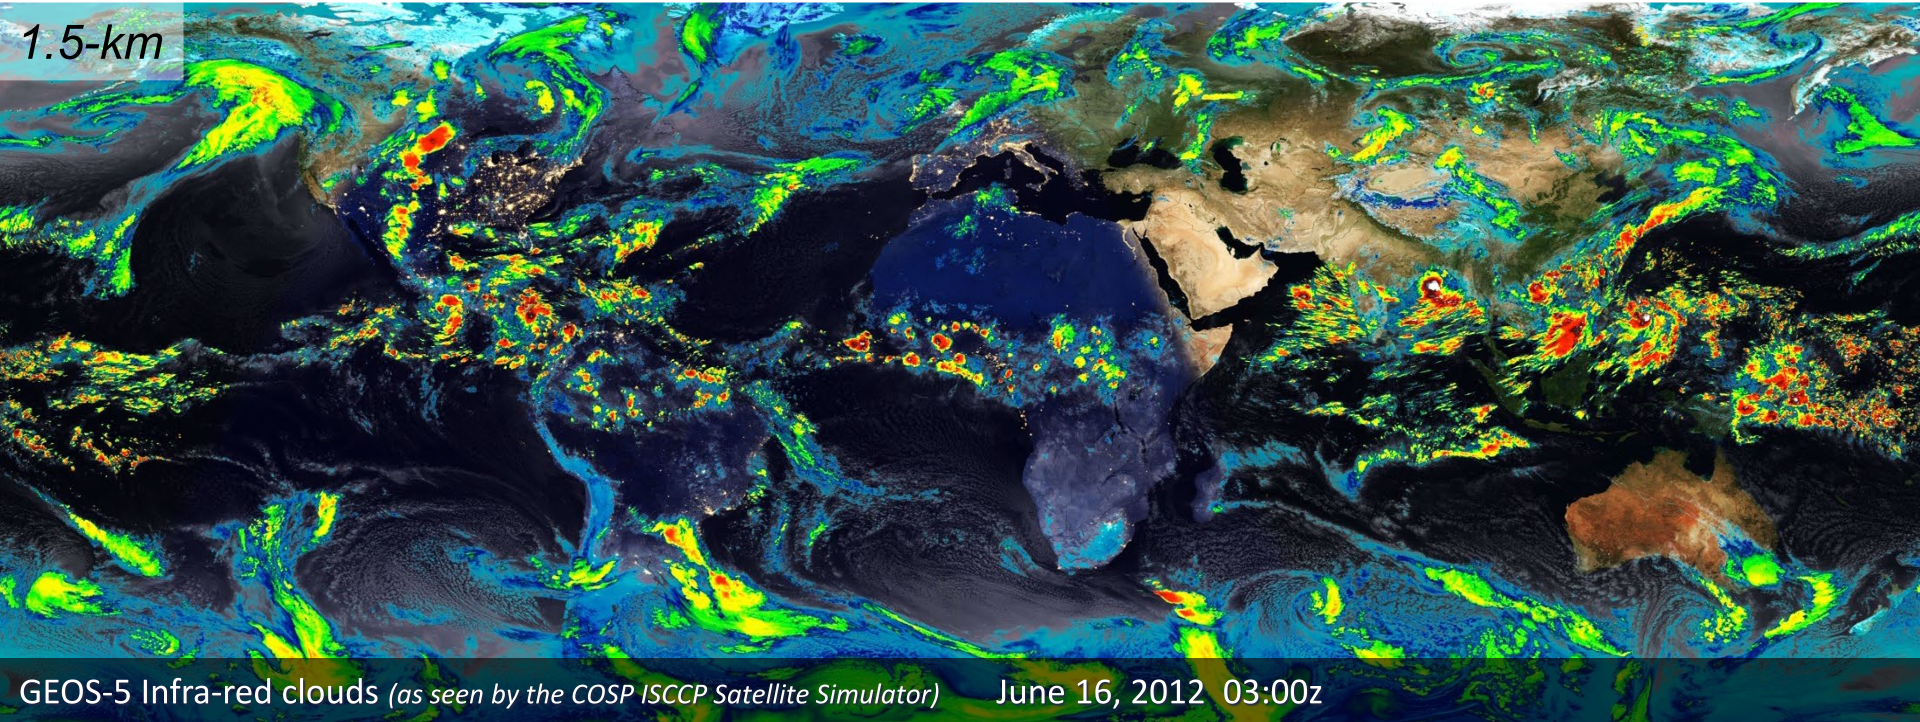

1.5km Global Simulation with GEOS-5 on SCU10 at NCCS

Wildfires throughout the western United States and industrial emissions in the East

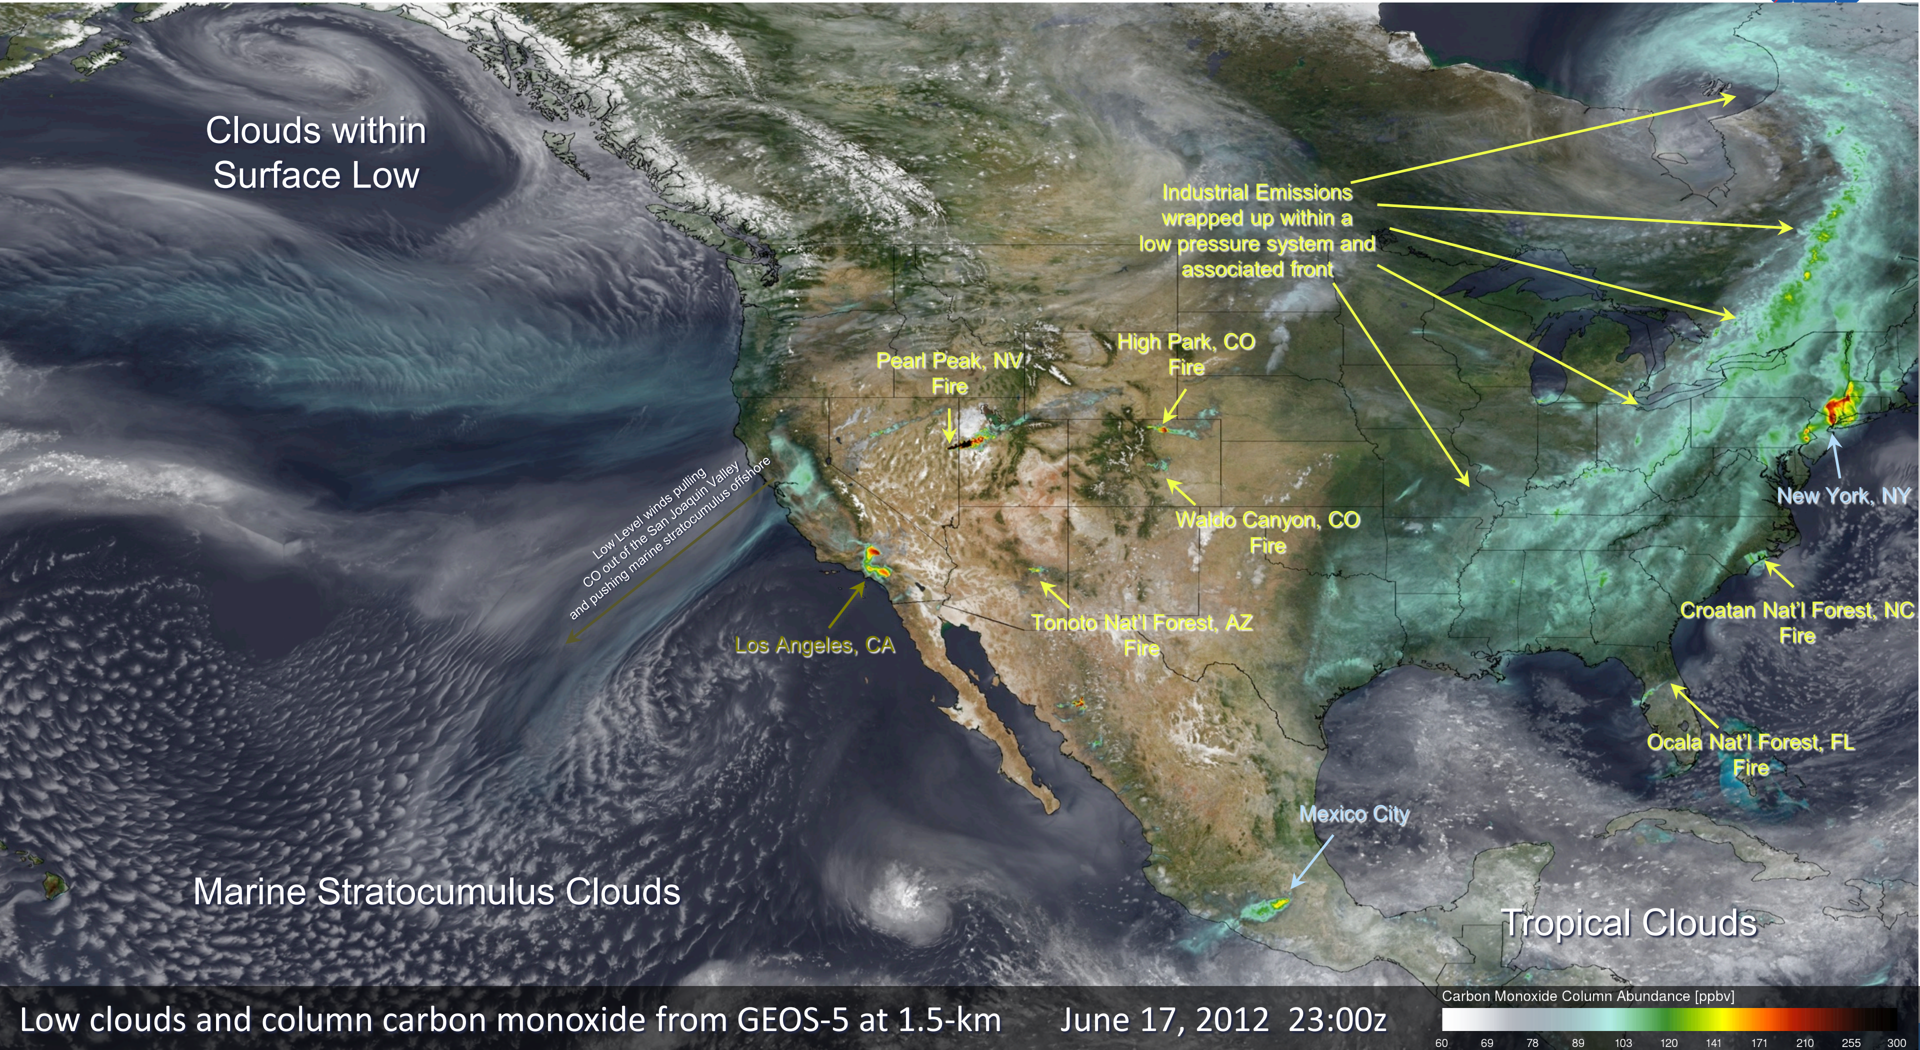

Low clouds and column carbon monoxide from GEOS-5 at 1.5-km June 17, 2012 23:00z

NASA Goddard Computing and Model Evolution

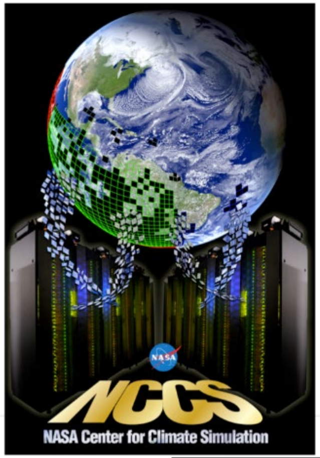

Over the past 10 to 15 years
- 1,000x computing capability
- 3,000x storage capacity

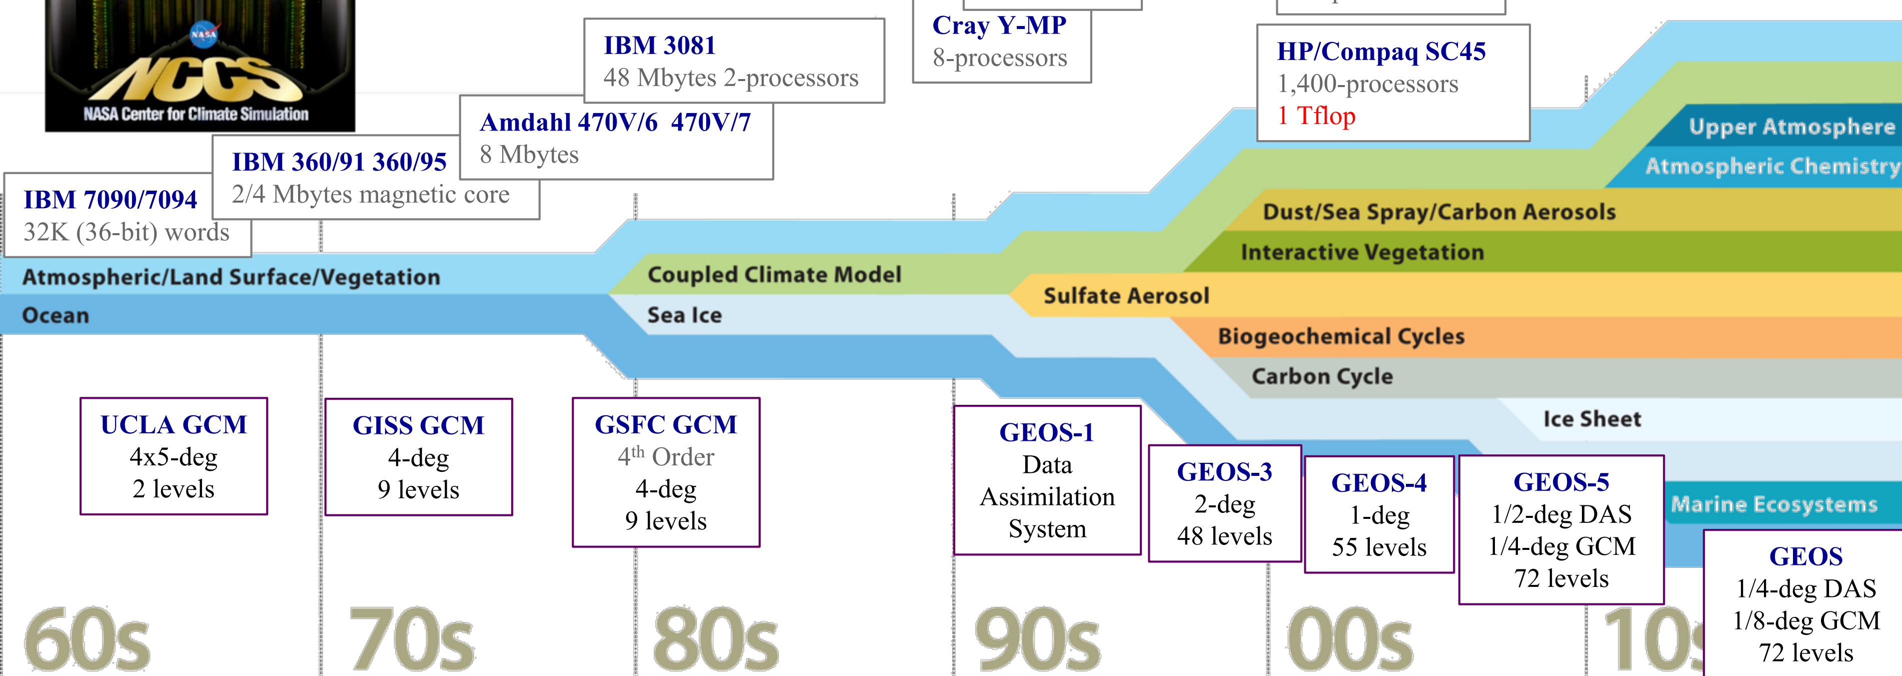

GOES-16

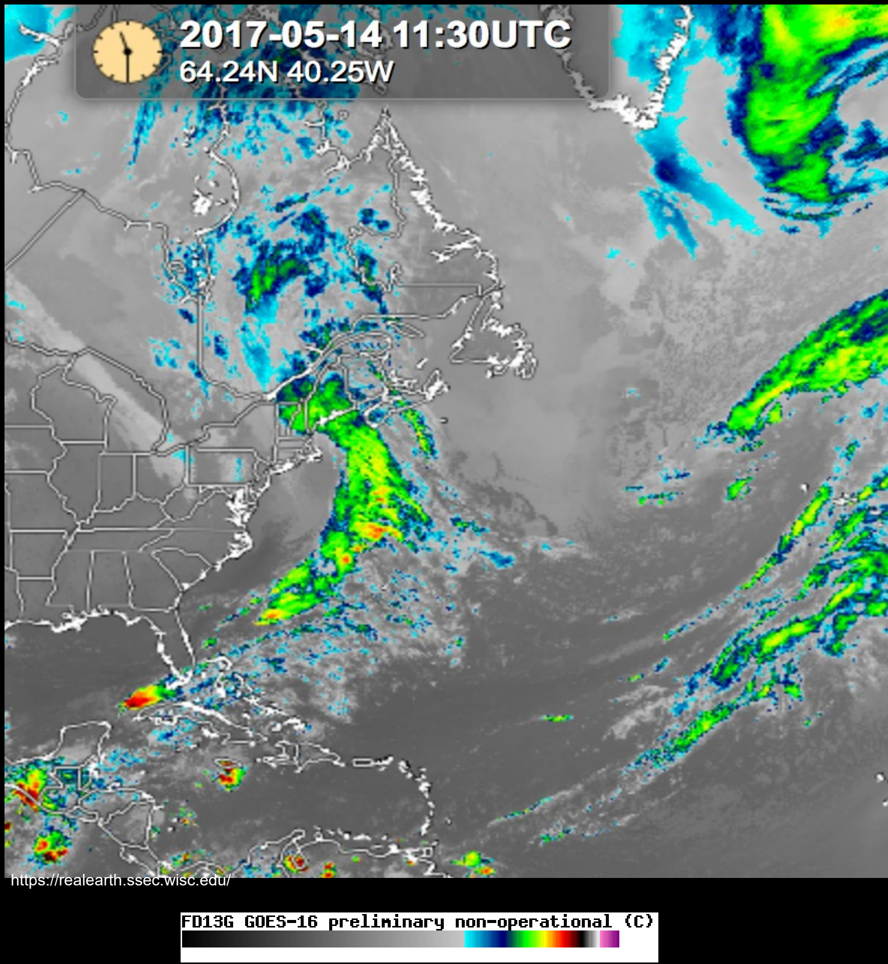

1.5-km GEOS Simulation

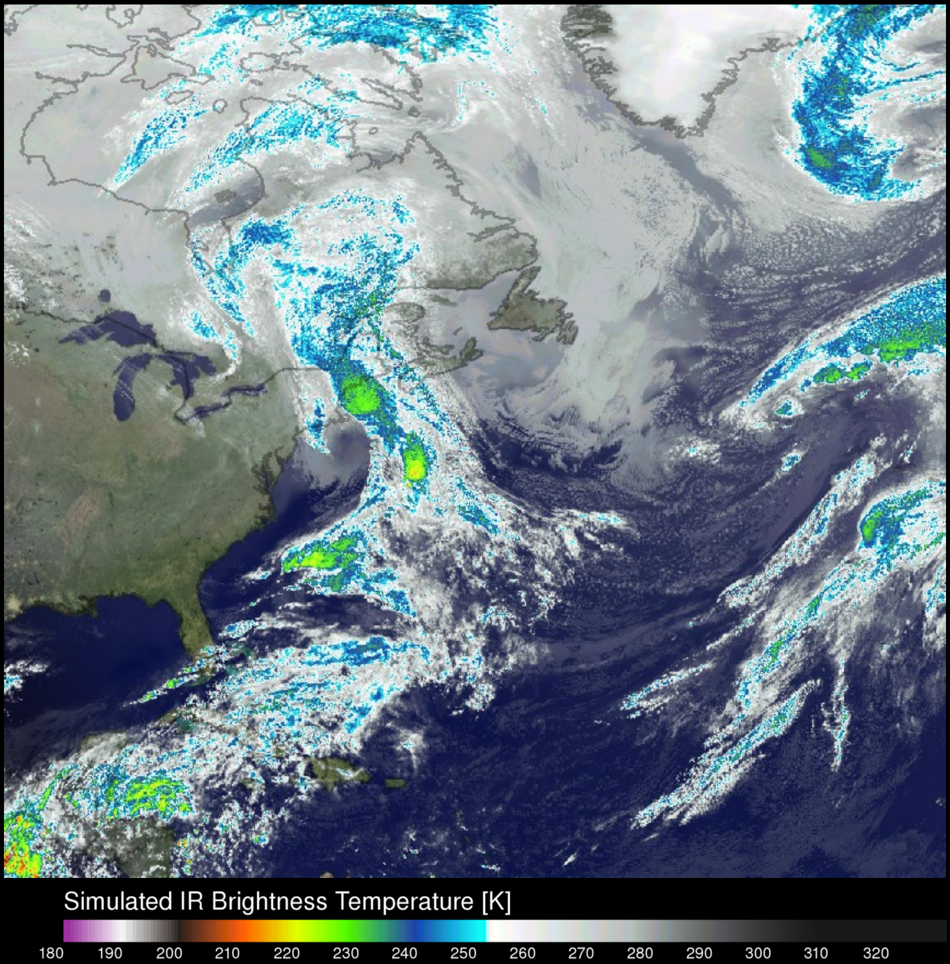

Science and Computing Required to Increase Resolvable Scales

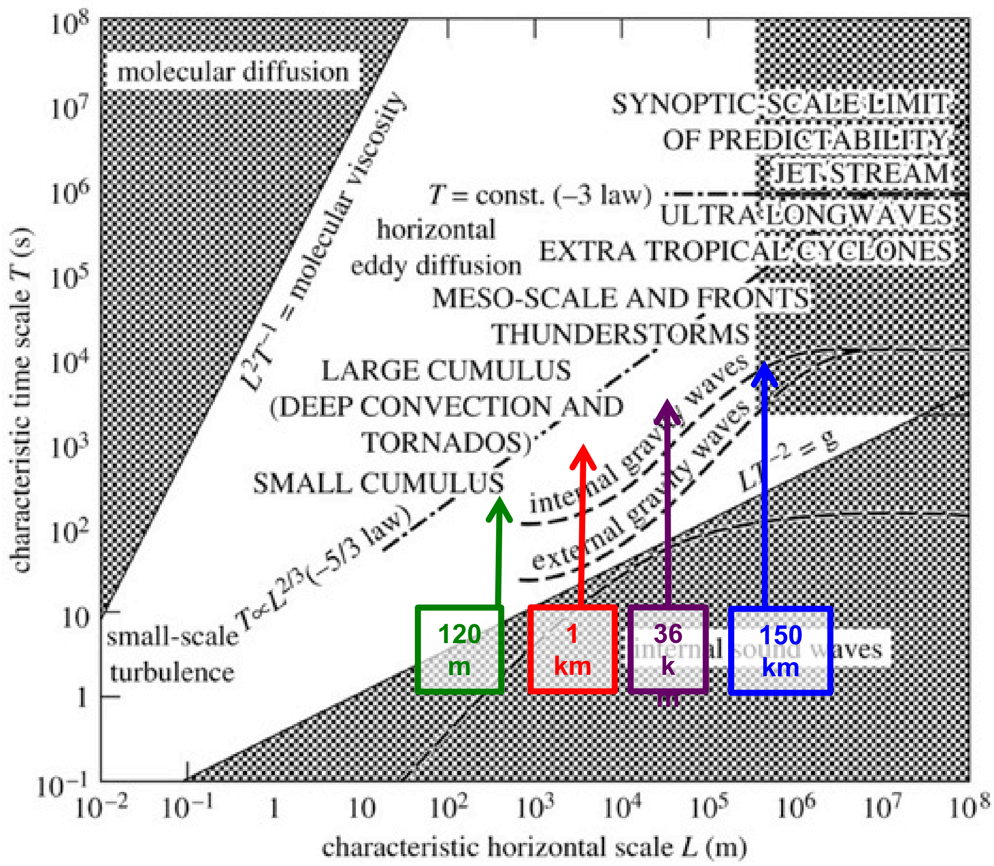

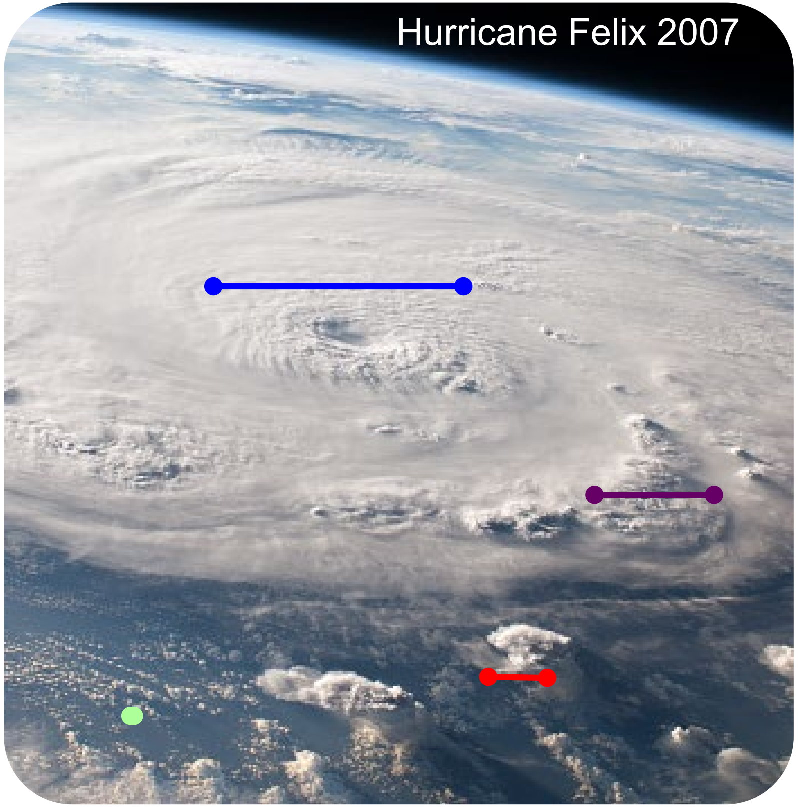

The Pursuit of Exascale

http://www.gfdl.noaa.gov/visualizations-mesoscale-dynamics

Tornado resolving global models at 1-km FV3

Resolution (km)	Resolvable ~12x (km)	Computing (Cores)
25.0	300	800
12.5	150	6,400
3.0	36	462,963
0.1	1.5	6,400,000,000
10 (m)	120 (m)	21,600,000,000,000,000

Adding Complexity to Global Models

Interactive Chemistry and Aerosols

Volcanoes

Photograph of ash plumes streaming from Mt Etna, Italy in 2001 taken by the Expedition 2 crew aboard the International Space Station

Ship Tracks

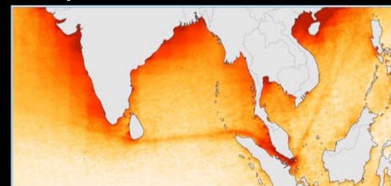

Narrow well established shipping lanes observed by Ozone Monitoring Instrument (OMI) on NASA's Aura satellite between 2005 and 2012

Smoke/Haze

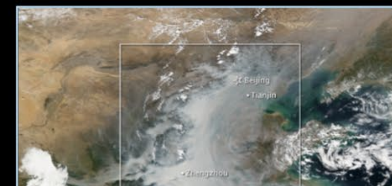

Haze over China in October 2010as seen by the Moderate Resolution Imaging Spectroradiometer (MODIS)

Dust

May 2001 Sea-viewing Wide Field-of-view Sensor (SeaWiFS) reveals a large, thick plume of aerosols blowing eastward over the North Atlantic Ocean

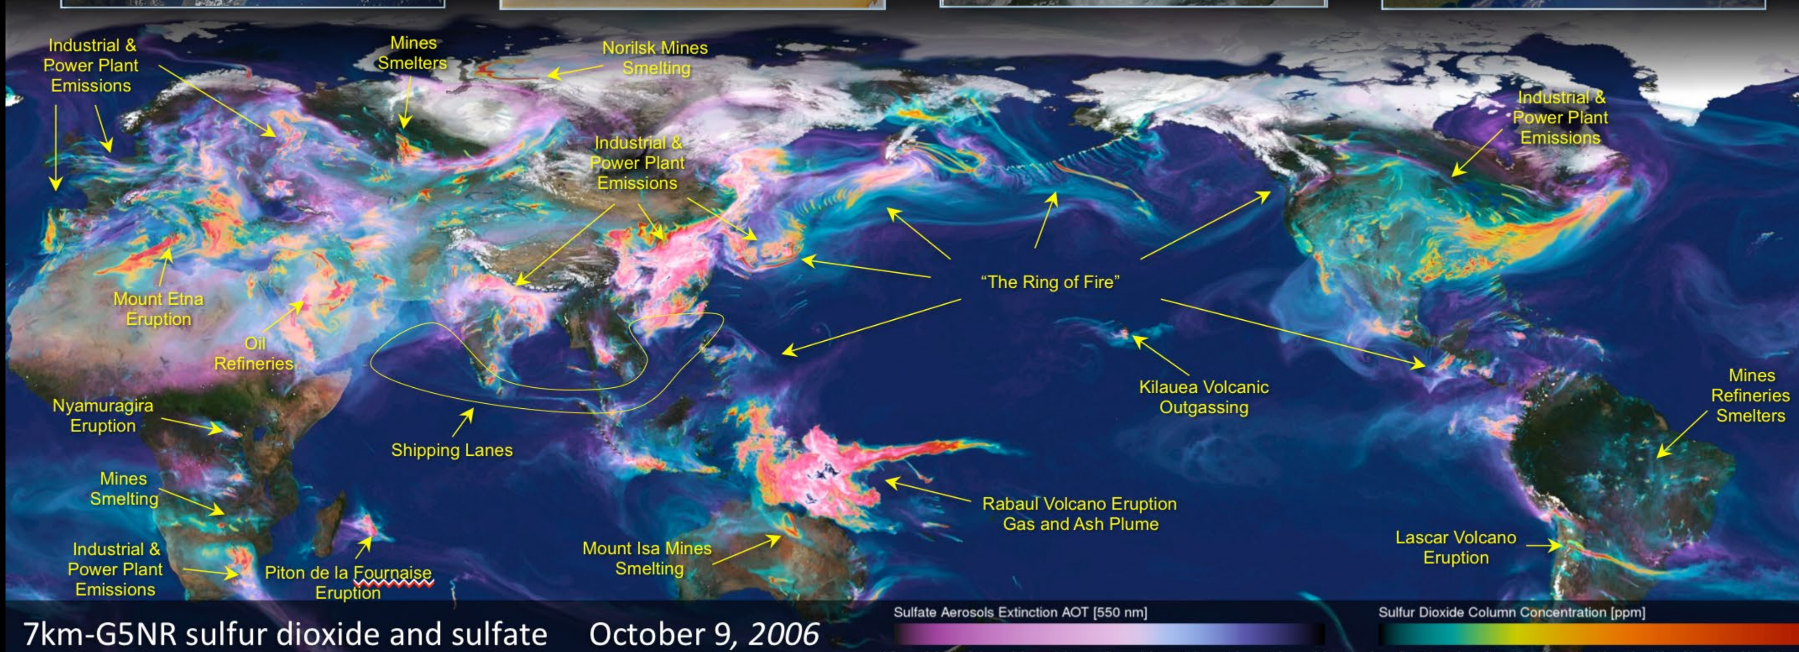

Looking Toward the Future

Mesoscale Convective Complex (MCC)

A mesoscale convective complex is a large thunderstorm

- Typical in spring over the Midwest US
- Progress over long distances
- Heavy rainfall, strong winds, frequent lighting, hail and often tornadoes.

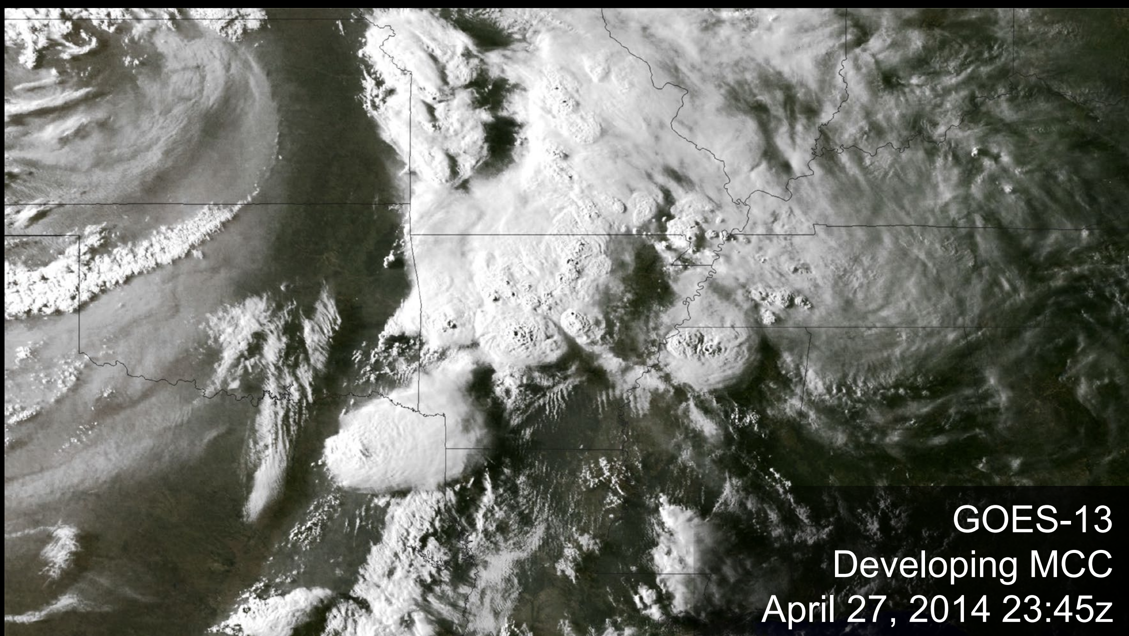

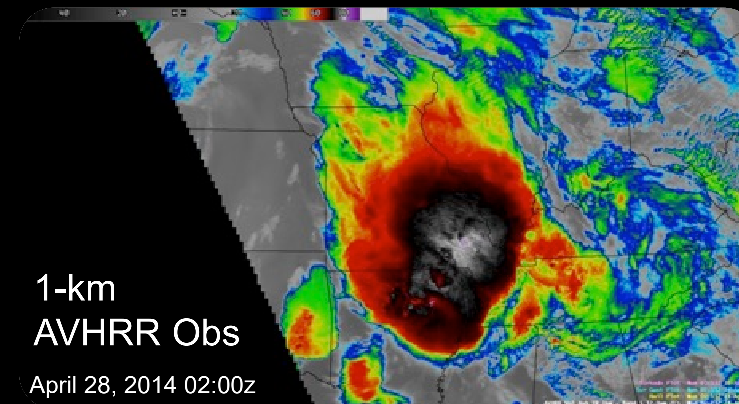

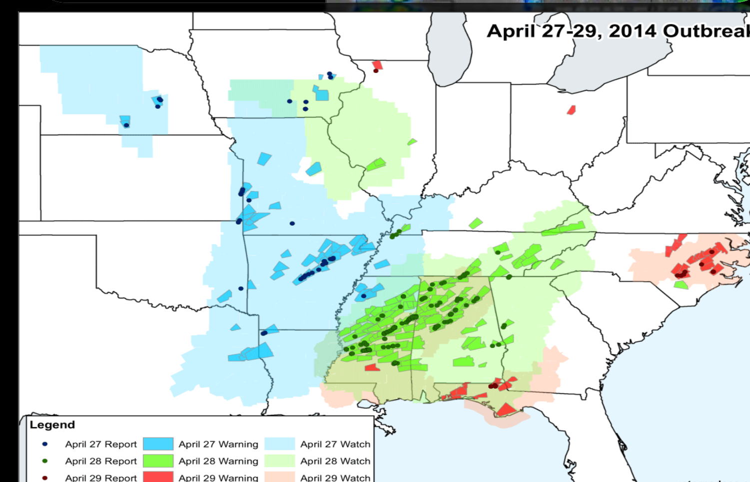

Looking Toward the Future

Supercell Thunderstorms

An isolated severe rotating thunderstorm

- Broad anvil cloud top
- Overshooting convective updrafts
- Damaging winds, large hail, tornadoes

GOES-14 Satellite Observations

- 1-km Resolution
- 1-minute Super Rapid Scan Operations for GOES-R
- Thunderstorms over southwest Texas, 19 May 2015

HPC Requirements for Earth System Modeling

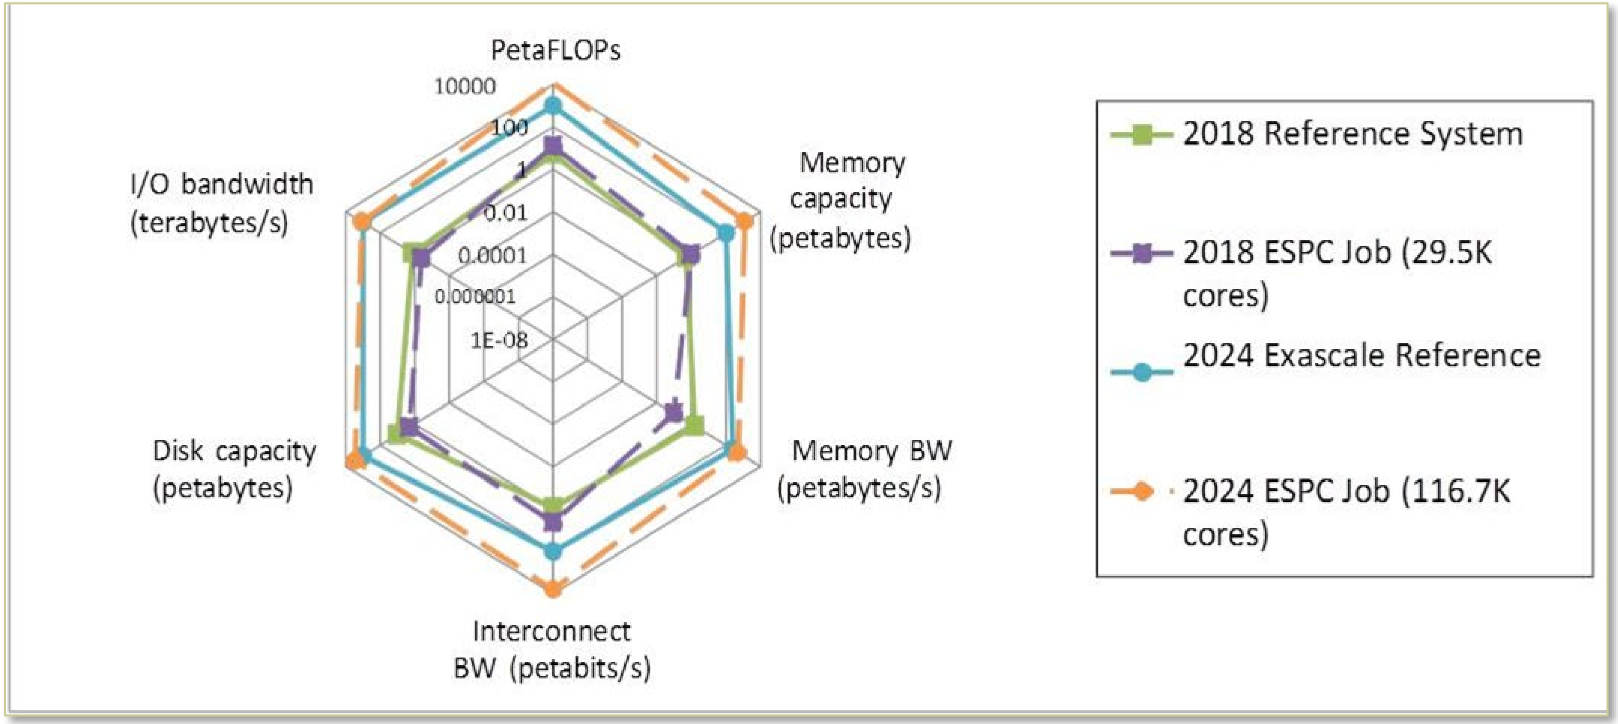

Internal report: *The Future of DoD Climate, Weather and Ocean High Performance Computing Requirements*, 15 Aug 2016, Figure 24

Source: Dave McCarren, National ESPC Project Manager: Purpose-Built HPC: Last Hope for Earth System Prediction

HPC Outlook

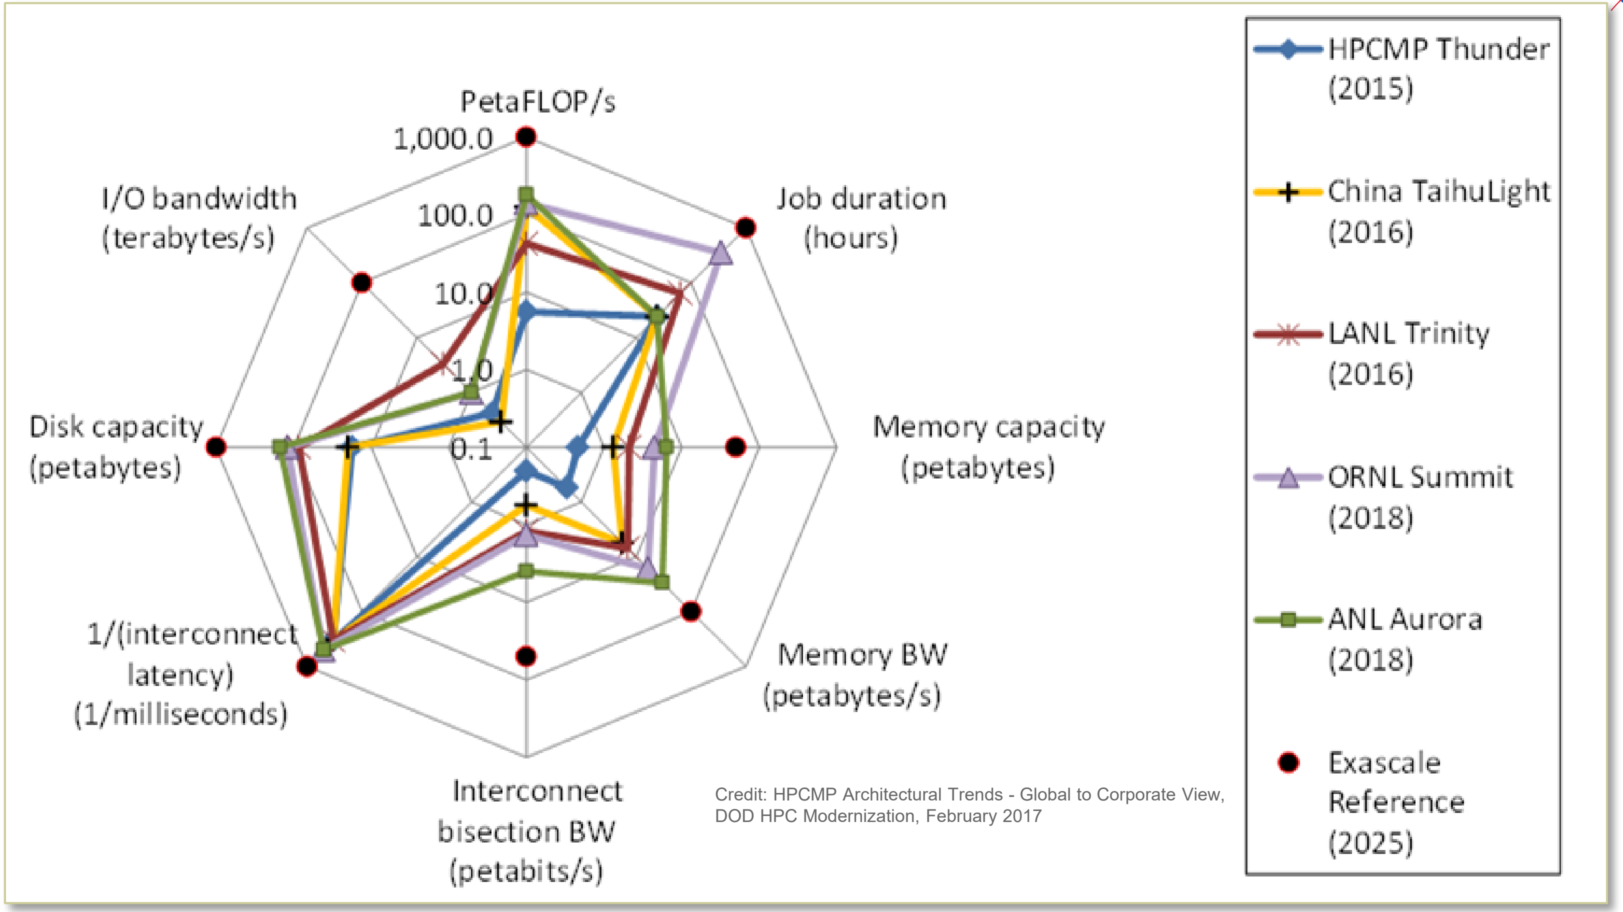

Source: Dave McCarren, National ESPC Project Manager: Purpose-Built HPC: Last Hope for Earth System Prediction

Intel Configurable Spatial Architecture (CSA)

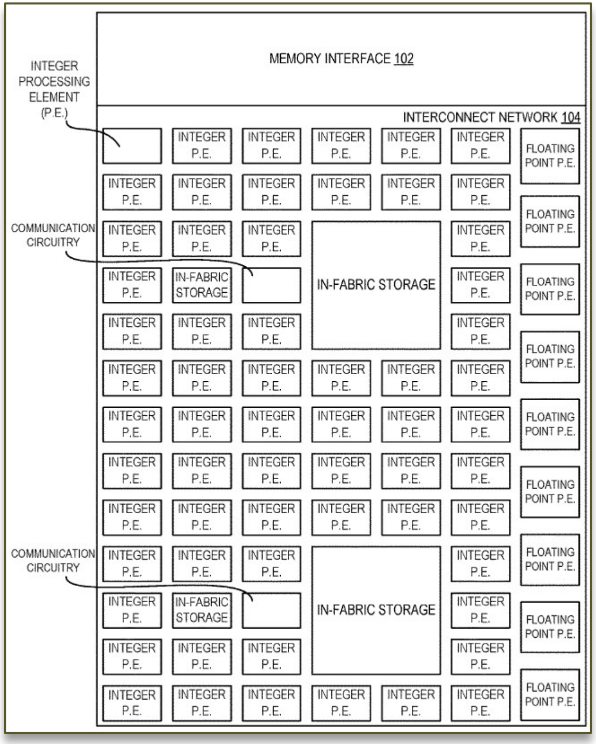

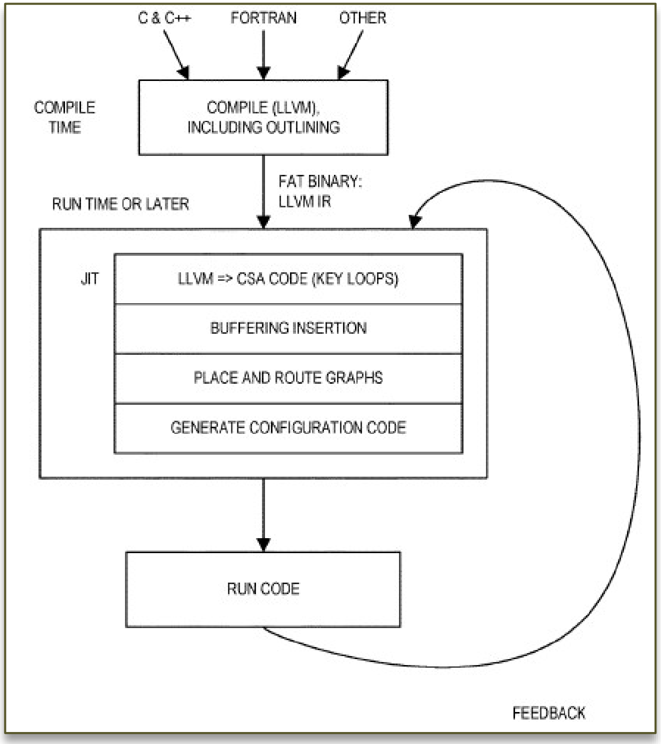

To be used for Aurora at Argonne National Laboratory and installed around 2021.
Map and place the dataflow of the program onto a processor with a large amount of logic units. Different CSA processors can be designed based on classes of applications, such as climate and weather.

Reference: https://www.nextplatform.com/2018/08/30/intels-exascale-dataflow-engine-drops-x86-and-von-neuman/

Task Based Software Approach

NASA Small Business Project with EMPhotonics

- Acknowledge that HPC codes are running on increasingly varied, heterogeneous systems
- Need a better way of efficiently utilizing all the resources available to an application
- Scientists not experts in performance optimization

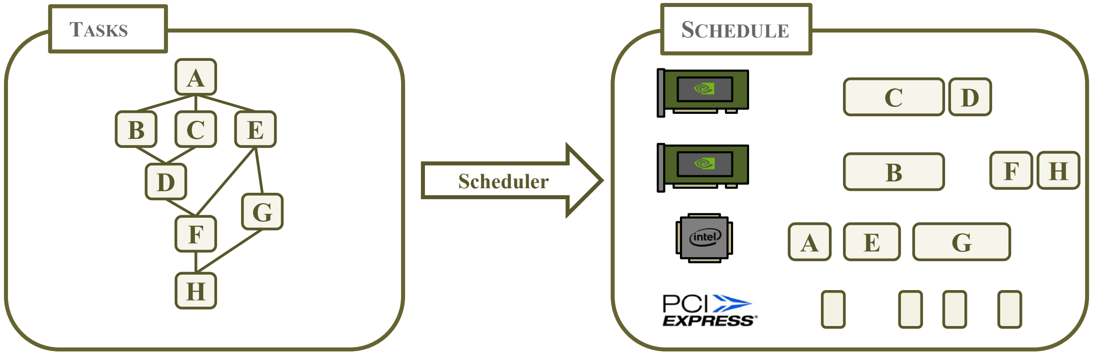

- Analyze dependencies between tasks
- Dynamically determine best hardware for each task
- Manage data movement between devices

Other similar approaches include Legion (Oak Ridge) and HTGS (UMBC).

Machine Learning and Artificial Intelligence at GSFC

- **Models**
  - Land Information Systems (LIS)
  - Improving the skill of remotely sensed snow water equivalent (SWE) retrievals using Deep Learning
  - Trained model components for Goddard Earth Observing System
  - Counting trees from space in high resolution imagery
- **Partnerships with Universities**
  - Phil Yang at GMU – Comparing dynamic downscaling of weather and climate data to a machine learning model
  - Milt Halem at UMBC – Potential gap filling of observation data to increase weather forecast skill and reduce skill dropouts
- **Goddard Strategic AI Group**
  - Dan Mandl and Jacqueline Lemoigne-stewart
  - NASA Goddard Workshop on Artificial Intelligence
  - Dates: Nov. 27-29, 2018
- **NASA Small Business Subtopic**
  - S5.03 Machine Learning and Deep Learning for Science and Engineering
  - Phase 1 project with Sivananthan Laboratories: Advanced Hyperspectral Imaging through Integrated Compressive Sensing/Inpainting via Machine Learning
  - Phase 2 project with Continuum Analytics (Anaconda) to create an Open Source Parallel Image Analysis and Machine Learning Pipeline

Library of "Scalability" Unit Tests

- Leverage the concept of unit tests
- Build relatively simple applications that represent major components of weather and climate applications
    - Dynamics, Physics, Chemistry, I/O
- Applications that are 10K lines of code or smaller and not 100K or greater
- Use these instead of full application benchmarks; does not take the place of full applications
    - To test at Exascale
    - Use for CSA architecture
- Available to HPC vendors and others in the community
- Community driven and maintained
- Needed within the next few years to better understand our approach for the next decade
- ***Help us to better understand the question: Can we get there from here or do we have to fundamentally rethink our algorithms and applications?***

What are our options?

Continue to evolve applications:

- Mixed mode parallelism
- PGAS languages

Add on to the box:

- Offloading
- GPUs
- FPGAs
- CSA

Open up the box for different uses:

- Task based programming approach
- Machine learning
- Neuromorphic

Build a different box:

- Purpose built HPC for weather
- Custom Hardware
- Co-design
- Quantum Computing

Purpose-Built High-Performance Computing for Earth System Models
NASA Funded Study

- ***What are the different options for building a custom high performance computer targeted for weather and climate applications?***
- To be funded by Dr. Tsengdar Lee (NASA High-End Computing Lead – Headquarters)
- Key aspects of the study
  - Components for customization
    - » Processors, memory, bandwidths, I/O, storage, software stack, etc.
  - Potential performance improvements
  - Vendors that could potentially create such a system
  - Total level of effort required by both vendors and application teams
    - » Ease of use by applications; assess the investment needed to use the system
  - Total cost of development and potential system costs
  - Pros and cons of all approaches
- Will be contacting HPC centers running weather and climate applications
  - Very appreciative of input
- Will share the study with interested centers
  - Time frame on the order of 6 months

Concept: NASA Exascale Applications Workshop

- Idea being discussed in NASA – not a plan of record at this point
    - ***Bring together domain scientists, applications engineers, HPC software engineers, and HPC system architects to discuss how to get NASA applications to Exascale***
- Possible partnership between the Earth Science Technology Office (ESTO) and the NASA High-End Computing (HEC) Program
- Outcome of the workshop: identify the following
    - Grand challenge NASA applications; candidates for Exascale
    - Gaps and challenges for NASA applications moving to Exascale
    - Potential NASA investments that can be made: both funded and directed research
    - Partnerships that can be made to accelerate NASA applications toward Exascale
    - Plan for the next 3 to 5 years
- Would like to gauge people's interest for participation

Classic Riddle

You reach a fork in the road where both paths are blocked by a guard. One guard always tells the truth and one always lies. You may ask one question to only one guard in order to determine which path leads to your destination. What do you do?

ANSWER: Ask one of the guards what the other would do, and then do the opposite.

Exascale Riddle

You reach a fork in the road where both paths are blocked by an HPC Vendor. Both vendors only tell lies. You may ask as many questions as you want in order to determine which path leads to your destination. What do you do?

ANSWER: ???? Need a plan.

Thank You! Questions?

Please feel free to contact me: daniel.q.duffy@nasa.gov

Fundamentals of Weather/Climate Models

Physics: Radiative transfer
Heat exchange
Water cycle
Land surface

Dynamics: Grid decomposition
Momentum and heat fluxes
Moisture transport