

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-03 14:29:42

PAGE 1

REFERENCE NO: 206

This contribution was submitted to the National Science Foundation as part of the NSF CI 2030 planning activity through an NSF Request for Information, https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf17031. Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

Author Names & Affiliations

- 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 HQ
- David McCarren - Navy/CNMOC
- John Michalakes - UCAR
- Scott Sandgathe - APL/UW
- Timothy Whitcomb - NRL/MMD

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Earth system modeling and high-performance computing

Title of Submission

High Performance Computing Needs in Earth System Prediction

Abstract (maximum ~200 words).

The United States experiences some of the most severe weather on Earth. Extreme weather or climate events—such as hurricanes, tornadoes, flooding, drought, and heat waves—can devastate communities and businesses, cause loss of life and property, and impact valuable infrastructure and natural resources. The number and severity of extreme weather and climate events in the U.S. has risen since 1980, and is projected to continue rising this century. Growing populations in vulnerable areas create increased risks. If current trends continue, damages from extreme weather and climate events could grow four-fold by 2050(1).

Predictions and projections of weather and extreme events across time scales from weather to climate rely on sophisticated numerical models running on High Performance Computing (HPC) systems, which press the frontier of the Nation's HPC capability. The Nation's Earth system modeling community has a unique set of HPC requirements which differ from industry and visualization needs. Typically HPC advances are measured using computational peak performance metrics that are ill-suited to Earth system modeling applications. We advocate for a shift in processor design to increase emphasis on memory bandwidth, so Earth system models run more efficiently and better serve the public need.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

At operational and experimental weather prediction centers, improved weather prediction at higher spatial and temporal resolution with longer lead times comes as a result of improved initializations and inclusion of more processes (e.g., ocean, land, biogeochemistry, etc.) to more accurately simulate atmospheric and oceanic conditions. Improved model initialization is achieved by more accurately assimilating billions of observations collected by next-generation, high-resolution satellites, radars, and a myriad of in situ sensors worldwide. As an example, global weather forecast and assimilation models are driven by extreme requirements: each day of the model forecast must be calculated within eight minutes of wall-clock time to meet operational requirements.

Many operational centers have found that better physical parameterizations of sub-grid-scale phenomena typically improve Earth system predictions, particularly at extended time ranges. Better parameterizations require an advanced understanding of physical processes resulting from theory, targeted observational field studies, and modeling studies using cloud-resolving or cloud-permitting models, characterized by significantly greater resolution than current earth system models. Additionally, property exchanges between model domains (precipitation, radiation, soil moisture, watersheds, vegetation, surface waves, friction and momentum transfer, ice floes and fracture heterogeneity, and more) are critical to modeling extended range phenomena; these properties have significant horizontal variability(2). Appropriate modeling of these domains (e.g. eddy-resolving or eddy-permitting ocean, high-resolution ocean mixed layer, high-resolution land use and vegetation, high-resolution cryosphere) is needed for their commerce, defense, infrastructure, energy and resource implications themselves, as well as their improvements to atmosphere prediction.

Similarly, longer-term trend assessment is a result of running long-range Earth system models at higher resolutions to more accurately simulate physical processes, with predictions being pushed out to monthly-to-decadal time scales via the addition of initialization states. Increased computation speed is also important for longer forecast runs at the decadal scale, as shorter simulation times per model day aid in data evaluation and scientific development. Both weather and climate prediction exploit model ensembles (~100 members) to better quantify uncertainty. These models assimilate hundreds of gigabytes of data, compute over 10^9 horizontal grid points and over 100 vertical levels, and output terabytes of data each model-day.

Future Earth system models will need over 1000 times the computational power of today – and will need to use it much more efficiently. Current models run on petaflop (10^{15} flop/s) scale computers, but not within operational speed constraints. These models run at increasing resolutions but do not expand (scale) efficiently to the exaflop (10^{18} flop/s) systems expected in the 2020s.

Keys to efficient scaling are exploiting available parallelism and computational intensity, i.e. doing as many calculations as possible for each expensive access to the data. Models are comprised of different components presenting different challenges.

? Parallelism means operations can be performed at the same time by multiple cores. However, the time dimension in Earth science applications does not scale, leading to a performance wall due to limited available parallelism in only the spatial dimensions. This is exacerbated by (a) stalled processor clock speeds, and (b) time-to-solution constraints imposed for real-time Earth system simulation.

? Dynamical calculations (equations involving fluid velocity, density, temperature, etc.) are parallel in 3 spatial dimensions (although models currently only exploit parallelism in 2 dimensions) but limited by data bandwidth to memory and between other components within the supercomputer. Dynamics involves inter-processor communication, which for newer numerical approaches is local (nearest neighbor) and scalable compared to non-local communications used for spectral transforms.

? Physical parameterizations have a high computational intensity but are only parallel in two spatial dimensions and limited by computation speed. Parallelism in the vertical is limited due to the extremely fast physical coupling, and is a topic of ongoing research. This structural factor is characteristic of the basic geophysical fluid dynamics problem.

HPC architectures are developing in the wrong direction for state-heavy, low computational intensity (CI) Earth system applications.

Present HPC systems are being optimized for problems that may require upwards of 50 flops/byte(3).

- The number one system on the November, 2016 Top 500 list, with SunWei processors on the 125 petaflop TiahuLight, are designed for problems requiring about 25 flops/byte(4).

- The number two system on that Top 500 list, with Xeon Phi-based Knights Corner processors on Tianhe-2, are designed for problems requiring 7 flops/byte.

- Knights Landing processors are designed for around 6 flops/byte.

In contrast, the most computationally intense components in today's Earth system models rarely reach two operations per byte and typically run less than one operation per byte over the full application.

Current models running on next generation architectures average less than 2% of peak performance, constrained by their ability to perform sufficient calculations for each expensive access to memory(5). Recent demonstrations of future model implementations and processor core-counts show the ability to achieve 6% of peak for only the fluid dynamics part of a weather or climate simulation(6,7). Much work remains for efficient implementation of physics, for balancing computational loads to prevent one slow task from delaying the whole job, and for minimizing the impact of communication overhead over many hundreds of nodes on a parallel computer.

Each processor design and system architecture requires specific coding structures optimized for that machine, forcing complete model redesign and rewriting for each subsequent and disparate hardware type. Architecture-agnostic programming could offer a possible solution to portability but may present a challenge to achieving performance across vastly different hardware.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

Environmental modeling has traditionally been a large consumer of high-performance computing capabilities. Societal vulnerability to extreme events increases demand for hazard prediction from hours to years and makes the Earth system prediction problem one of the driving computing issues of our day. The international community also struggles with this problem(8). Operational and research modeling are subject to pressures to complete larger runs in shorter times, while current modeling systems are able to use only a small percent of the peak capability of the large HPC systems available for these problems.

Presently, a limited number of vendors responding to market forces and Top-500 metrics drive hardware design, resulting in inefficiency and lost opportunity for the Earth system prediction enterprise. Optimal computing for efficient prediction across time scales will require co-design between the hardware technologists (chips, storage, communications and total system architecture) and the modelers creating software across the system (compute, storage, networking).

The Earth system modeling community sees a need for engagement (co-design activities) with vendors by the Earth system modeling community here and in coordination with other modeling centers (e.g. the European Center's ESCAPE(9)) to come up with designs that are most suitable for Earth system modeling.

The unique requirements of numerical simulation for Earth system modeling argue against patterning the next ten years of cyberinfrastructure strategy along the lines of general purpose HPC architectures. It is the recommendation of this working group that NSF should include in its strategizing the idea of purpose-built HPC architectures that are balanced to deliver efficient performance for applications such as Earth system modeling, which have abundant parallelism but limited CI. That might mean designing (and paying for) more thread concurrency (and additional local stores to support relatively heavy-weight threads), higher memory bandwidth, configurable computing, etc. Hardware and software support for fast I/O, resilience, partitioned global address space (PGAS) support, efficient reductions, and more should all be tailored to specific requirements and computational characteristics of Earth system modeling. The present high-water mark of 6% of peak performance achieved for a well-designed weather prediction model(10) falls far short of what is needed to advance weather and climate prediction in the next decade.

Other communities would also benefit from HPC technologies that emphasize orders-of-magnitude increases in memory bandwidth and I/O throughput over hardware that offers higher computational capabilities. Problems in geo-sciences (tectonics and magma flow), energy (power station design), aerospace (flow around wings and other objects), biomechanics (blood flow), solar physics (plasma circulation, coronal mass ejections) sectors provide specific examples.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-03 14:29:42

PAGE 4

REFERENCE NO: 206

Given the mismatch between the Earth system modeling problem and the HPC architectures, and the generational pace of both software and hardware upgrades, each machine upgrade carries a massive burden in revising model code. Presently, the community has difficulties training the proper computational scientists, appropriately knowledgeable in both Earth system modeling and in the compute-kernel languages required for these new systems (e.g., CUDA, OpenCL, OpenMP, etc.) either to understand the requirements for porting legacy codes to new hardware as it comes available or to restructure models to exploit new hardware. University courses on programming many-core architectures tend to attract only mathematics or computer science students (i.e., very few students from Earth Sciences), which causes a shortage of postdoctoral researchers able to transition Earth system model code.

Additionally, in DoE, efforts to create commonality across disparate hardware architectures rely on libraries, whereby each library handles many-core computing independently; most libraries rely on the same tool—a hardware-agnostic programming language called Kokkos. Programmers need to be well versed in the C++ programming language to use Kokkos; most operational weather models are based on Fortran resulting in Fortran being the core programming expertise of Earth system modelers. The community needs help in migrating away from the sole use of Fortran by creating a cadre of new scientists equally well-versed in C, Fortran, C++, and others as well as additional language support for programs like Kokkos.

A significant and long-term investment in computer scientists, software engineers, applied mathematicians, and statistics researchers to partner with Earth System researchers is needed to address all stages of the prediction process, including data assimilation, operation of high-resolution coupled Earth system models, and storage and management of results(11).

The wide range of time scales in Earth system prediction systems have similar computing needs, and National ESPC is providing a single voice to requirements for decision support across weather and climate scales. Many of the challenges are currently being explored by participating software groups. While their efforts will evolve to exploit the next generation computing capabilities, a two-way design process with the HPC industry will optimize computing to increase public safety and national economic resilience.

(1) <http://www.ncdc.noaa.gov/billions/> (CPI adjusted to 2013 dollars)

(2) NRC (National Research Council). 2016. Next Generation Earth System Prediction: Strategies for Subseasonal to Seasonal Forecasts. Washington, DC: National Academy Press.

(3) John Goodacre, Manchester U., at ECMWF, Feb 2017

(4) Jack Dongarra, ECMWF

(5) "Experiences Optimizing NWP for Intel's Xeon Phi, Knights Landing" John Michalakes (UCAR),

<http://www.ecmwf.int/en/learning/workshops-and-seminars/17th-workshop-high-performance-computing-meteorology>

(6) Yang, Chao, et al. "10M-core scalable fully-implicit solver for nonhydrostatic atmospheric dynamics." Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE Press, 2016.

(7) Muller, Andreas, et al. "Strong Scaling for Numerical Weather Prediction at Petascale with the Atmospheric Model NUMA", submitted to the International Journal of High-Performance Computing Applications, <https://arxiv.org/abs/1511.01561>

(8) <http://www.ecmwf.int/sites/default/files/elibrary/2014/13800-workshop-scalability-2014-working-group-summary-report.pdf>

(9) <http://www.hpc-escape.eu/>

(10) Muller, Andreas, et al. "Strong Scaling for Numerical Weather Prediction at Petascale with the Atmospheric Model NUMA", submitted to the International Journal of High-Performance Computing Applications, <https://arxiv.org/abs/1511.01561>

(11) NRC (National Research Council). 2016. Next Generation Earth System Prediction: Strategies for Subseasonal to Seasonal Forecasts. Washington, DC: National Academy Press.

Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>)."