NSF and Extreme Scale

William Gropp
http://wgropp.cs.illinois.edu

Disclaimer

- I do *not* represent NSF
  - And my presentation will be focused on my organization, the National Center for Supercomputing Applications (NCSA)
  - NSF supports other centers with significant capability, including TACC, SDSC, PSC, IU, as well as NCAR
- NSF supports a broad range of extreme scale computing research projects, including
  - Software Infrastructure for Sustained Innovation (S2I2)
  - Scalable Parallelism in the Extreme (SPX)
  - Big Data Regional Innovation Hubs (BD Hubs)
  - Critical Techniques, Technologies and Methodologies for Advancing Foundations and Applications of Big Data (BIGDATA)
  - Many others (search for “extreme computing” and “big data” at http://nsf.gov)

NSF and NSCI

- National Strategic Computing Initiative
- NSF is one of 3 “lead agencies” (DOE and DOD are the others)
- “NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development.”
  - Read this as NSF key in:
    - Use of extreme scale in science research of all types
    - Development of the software and algorithms ecosystem for a broad range of applications
    - Development of an HPC-expert workforce; this means both scientists who use HPC and staff that support HPC systems and applications

NSF's Plans for National Computing Resources

Current and future portfolio of NSF-supported National Computing Resources

Complements Larger Aggregate Investments from Universities and other Agencies

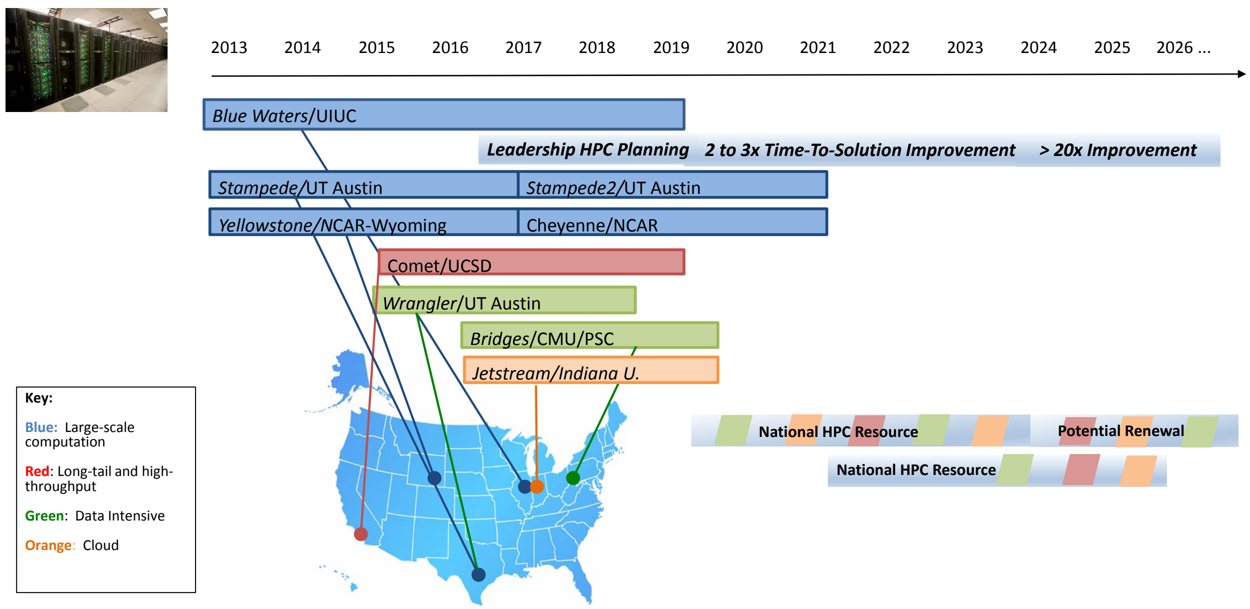

Blue Waters: NSF’s most powerful system

Largest U.S. system for open science and engineering research

- 13 PF peak performance
- 1.5 PB memory
- 1 TB/sec I/O bandwidth
- 26 PB disk
- 380+ PB near-line tape storage capacity
- Support includes experts for each science team

The Largest Supercomputers by Core Count

System	Cores
Sunway TaihuLight	10,649,600 (most SIMD; 40,960 nodes)
Tianhe-2	3,120,000 (most in Intel Phi)
Sequoia BG/Q	1,572,864
Blue Waters	792,064* + 1/6 acc (59,136 GPU stream proc)
Mira BG/Q	786,432
K computer	705,024
Stampede	462,462 (most in Intel Phi)
Julich BG/Q	458,752
Vulcan BG/Q	393,216
Titan	299,008* + acc (261,632 GPU stream proc)

* 2 cores share a wide FP unit

Geospatial Intelligence – ArcticDEM
Paul Morin, University of Minnesota

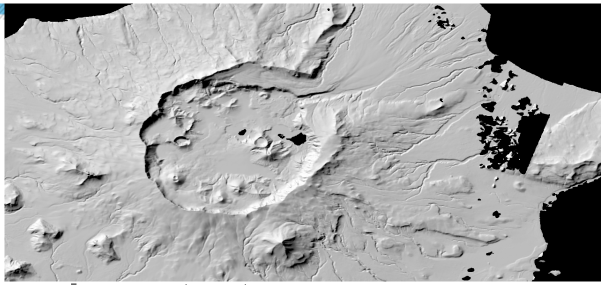

About the research

Creates 3D digital elevation maps (DEMs) at 2m-5m resolution using satellite imagery (current usable at 250m) President Obama in 2015 requested high quality elevation data for entire Arctic Circle

More info: www.pgc.umn.edu

Why Blue Waters was important

"Without Blue Waters would not have the necessary compute power to complete the project to the required scale in the allotted 2 year timeframe." Project has large footprint (18M square kilometers), requires over 19M node hours (600M core hours) to complete.

Scientific and societal impact

- In Sept. 2016 released DEMs of Alaska, Novaya Zemlya (Russia) and Franz Josef Land (Russia)
- Changing how scientists work in the Arctic - better maps, lower cost; ArcticDEM 25-50 cents/km; other methods $55/km and up
- Enhanced data for disaster risk mitigation

Next steps

Will apply these same techniques to create high resolution DEMs for New Zealand and the Antarctic

Dark Energy Survey (DES)
Donald Petravick, University of Illinois at Urbana - Champaign

About the research

An international, collaborative effort to map hundreds of millions of galaxies, detect thousands of supernovae, and find patterns of cosmic structure that will reveal the nature of dark energy

DES produces ~1 TB of raw data nightly

Why Blue Waters was important

Robust I/O Lustre file system & node memory on XE nodes:

- Multi-epoch production campaign collects thousands of individual astronomical images
- Hundreds of jobs per night, average 2.3 TB data per night
- Each job transfers 3-10 GB into Lustre; generates 15 GB
- Each job uses ~10GB of RAM
- 3K (30%) jobs run on BW so far
- Large outbound network connectivity to DES outside data infrastructure (file store + DB); thus pushing HPC integration

Scientific and societal impact

Continuous and deep multi-color map of ¼ of the southern sky.

All data are available to a scientific collaboration of over 400 members in the U.S., Europe and South America.

All data will be available to the whole astronomical community through two major public data releases.

Next steps

- Run larger jobs that each take ~20 hours
- DES continues to work with BW for support of docker containerization on HPC (e.g., shifter)
- Continue usage of BW for massive DES yearly re-processing of campaigns
- Programmatically, DES is teaching us how to run large I/O processing astronomy jobs on HPC in preparation for LSST

Planning for the Future of Advanced Computing at NSF

- Four Major Goals Identified
  1. Position the United States for continued leadership in science and engineering
  2. Ensure that resources meet community needs
  3. Aid the scientific community in keeping up with the revolution in computing, and
  4. Sustain the infrastructure for advanced computing
- 7 Recommendations (with their own detailed recommendations)

- http://tinyurl.com/advcomp17-20 https://www.nap.edu/catalog/21886/future-directions-for-nsf-advanced-computing-infrastructure-to-support-us-science-and-engineering-in-2017-2020

Report Recommendations for Big Data

Recommendation 2 As it supports the full range of science requirements for advanced computing in the 2017-2020 timeframe, NSF should pay **particular attention to providing support for the revolution in data-driven science** along with simulation. It should ensure that it can provide unique capabilities to support large-scale simulations and/or data analytics that would otherwise be unavailable to researchers and continue to monitor the cost-effectiveness of commercial cloud services.

Recommendation 2.1 NSF should **integrate support for the revolution in data-driven science** into the Foundation's strategy for advanced computing by (a) requiring most future systems and services and all those that are intended to be general purpose to **be more data-capable in both hardware and software** and (b) expanding the portfolio of facilities **and services optimized for data-intensive** as well as numerically-intensive computing, and (c) carefully evaluating inclusion of facilities and services **optimized for data-intensive computing** in its portfolio of advanced computing services.

Recommendation 2.2 NSF should (a) **provide one or more systems** for applications that require a single large, **tightly-coupled parallel computer** and (b) **broaden the accessibility and utility** of these large-scale platforms by allocating high-throughput as well as high-performance workflows to them.

Recommendation 2.3 NSF should (a) **eliminate barriers** to cost-effective academic **use of the commercial cloud** and (b) **carefully evaluate the full cost** and other attributes (e.g., productivity and match to science workflows) of all services and infrastructure models to determine whether such services can supply resources that meet the science needs of segments of the community in the most effective ways.

Recommendations WRT NSCI and NSF

- High-performance computing (HPC) remains critical for science and industry; if anything, the need and value continues to grow. (NSCI Section 1)
- “Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.”(NSCI Section 2.2)
- Building on its successes in cyberinfrastructure, the National Science Foundation (NSF) has an important role to play in both providing HPC (including data and compute) for basic science and in development of the science needed to advance HPC, including the algorithms, software, and hardware for extreme scale computing. (NSCI Section 3a)
- NSF must also contribute to the development of an HPC workforce. (NSCI Section 3a)
- Public-private partnerships should be explored. (NSCI Section 1.2)
- HPC research must be transitioned into practice. (NSCI Section 1.4) This report’s recommendations to NSF echo this need; in particular, NSF needs both to perform research in support of HPC and support bringing that research into practice as needed by the NSF user community.
- Embrace an integrated approach to providing effective HPC, combining hardware, software, and algorithms, as well as addressing the development of an HPC-capable workforce and the whole of HPC, including the mid-range as well as the high-end. (NSCI Section 2.4)

Recommendations for NSF not addressed in NSCI

- Although convergence of data-intensive and compute-intensive systems is important and will address many needs, some applications require more specialized approaches that may emphasize compute or data. (NSCI Section 2.2 focuses on technology convergence)
- The demand for computing continues to outstrip supply; more needs to be done to (a) provide greater resources (especially systems and expertise in using them) and (b) to make the best use of these resources (NSCI makes no statements on budgets; efficient use of the ecosystem is mentioned but without specific coordination).
- A diversity of platforms and software will be needed to capture the long tail of science. NSCI calls for acceleration of the deployment of an Exascale class system but says nothing about the acceleration needed for future science needs at all scales.