IDEAS productivity

Responding to the Software Crisis in DOE Scientific Computing

Introduction

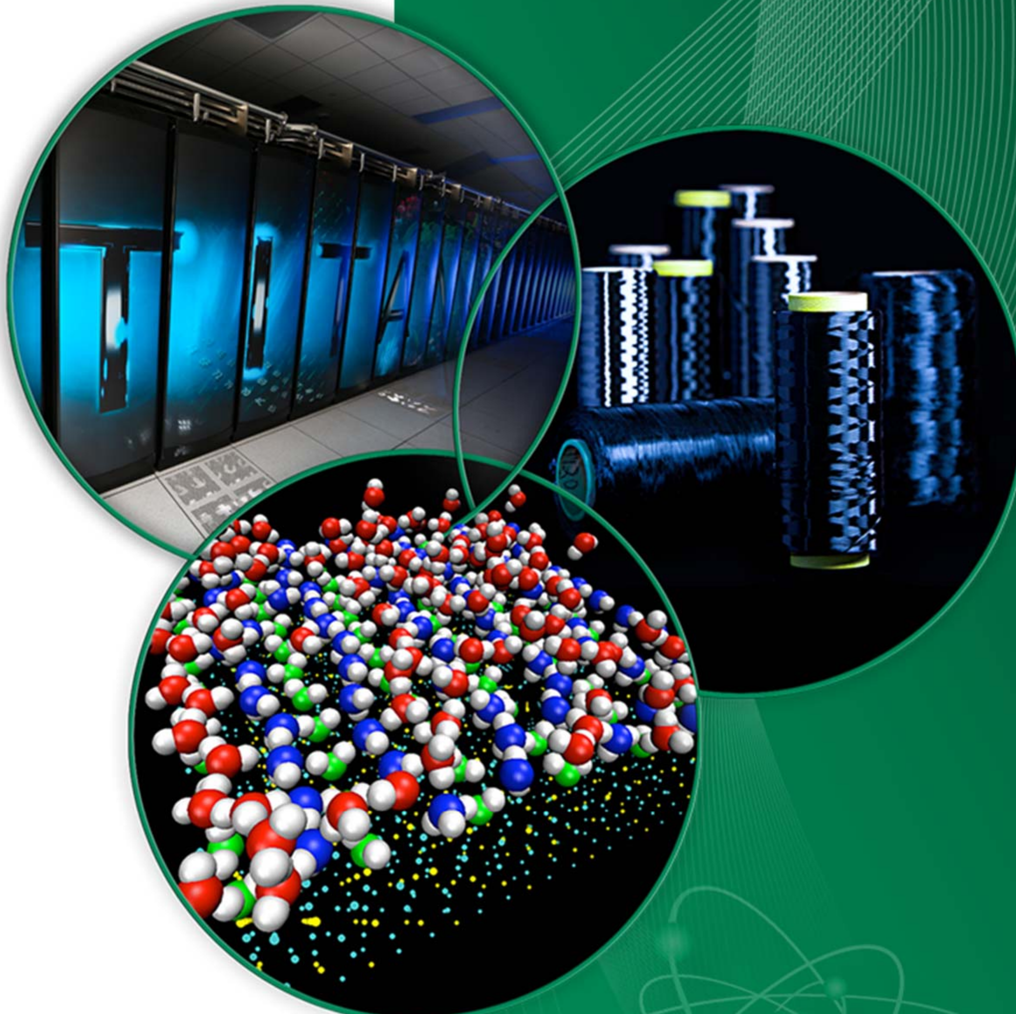

David E. Bernholdt
Oak Ridge National Laboratory

for the organizers

U.S. DEPARTMENT OF ENERGY

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE U.S. DEPARTMENT OF ENERGY

Reaching Crisis Proportions

- Science and engineering depend increasingly on computing to analyze, understand, and predict, even inform policy
  - DOE leads the way in many areas of science
- Increasing computational capability raises expectations for what computational science can deliver
  - E.g., recent rise in multiphysics and multiscale coupling
  - E.g., push for predictive simulation
- Computational science applications are major investments
  - Large code bases, developed over long periods
  - Under continual scientific development
- HW architectural trends disruptive to applications
  - After decades of stability, major changes are coming, increasing diversity
  - Uncertainties about exascale HW → uncertainties for apps

Productivity at Risk

- Scientific productivity in a computational context
  - Software productivity
    - Effort, time, and cost for software development, maintenance, and support
  - Execution-time productivity
    - Efficiency, time, and cost for running scientific workloads
  - Workflow productivity
    - Effort, time, and cost for the overall cycle of simulation and analysis
  - Value of computational results
- Challenges (just a sampling…)
  - Need to refactor and evolve codes, can't afford to re-write
  - Few computational scientists have training in software engineering
  - Little tooling to support refactoring needs
  - Hardware diversity & uncertainty makes performance a bigger challenge
  - Number of simulations and magnitude of data expanding rapidly

Recent DOE Activities: Exploring Crisis in SW Productivity for Extreme-Scale Science

- Pre-history:
  - DARPA-HPCS, DOE community meetings, SciDAC, SC, ICSE-CSE, etc.
  - Climate and environment NRC reports, FSP planning
- Summit & Whitepaper:
  - *Extreme-Scale Application Software Productivity*, Feb 2013
  - *Extreme-Scale Scientific Application Software Productivity: Harnessing the Full Capability of Extreme-Scale Computing*, Sept 2013
- Workshops:
  - *Software Productivity for Extreme-Scale Science*, Jan 2014
    - Whitepapers, relevant reading, report available via http://www.orau.gov/swproductivity2014
  - *Exascale Computing Systems Productivity*, June 2014
  - *Extreme Scale Productivity Workshop 2014*, August 2014
  - *Computational Science & Engineering Software Sustainability and Productivity Challenges*, October 2015
- Minisymposium at SIAM PP14:
  - *Software Productivity for the Next Generation of Scientific Applications* (8 presentations)

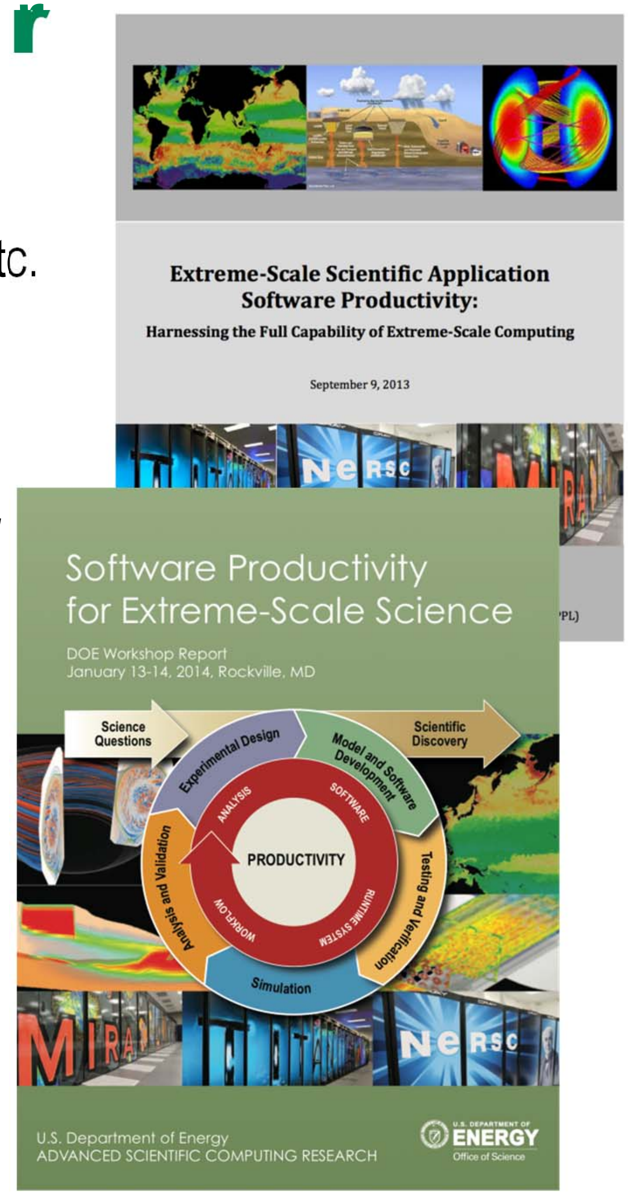

From the National Strategic Computing Initiative Fact Sheet

Improve HPC application developer productivity. *(Theme 3 of 5)*

Current HPC systems are very difficult to program, requiring careful measurement and tuning to get maximum performance on the targeted machine. Shifting a program to a new machine can require repeating much of this process, and it also requires making sure the new code gets the same results as the old code. The level of expertise and effort required to develop HPC applications poses a major barrier to their widespread use.

Government agencies will support research on new approaches to building and programming HPC systems that make it possible to express programs at more abstract levels and then automatically map them onto specific machines. In working with vendors, agencies will emphasize the importance of programmer productivity as a design objective. Agencies will foster the transition of improved programming tools into actual practice, making the development of applications for HPC systems no more difficult than it is for other classes of large-scale systems.

https://www.whitehouse.gov/sites/default/files/microsites/ostp/nsci_fact_sheet.pdf

[Emphasis added]

Goals of this Mini-Symposium

- Raise awareness of the looming software crisis in DOE scientific computing
- Look at ways software concerns have been included in past and present DOE programs/projects
- Identify some of the gaps where increased emphasis on software concerns could provide overarching benefits
- Generate discussion about how to respond

The Rest of this Morning's Agenda

Time	Title and Speaker
8:45am	*ASC Software Engineering Lessons Learned* Greg Pope, LLNL
9:10am	*High-Impact Modeling and Simulation: Realizing the Potential for the Applied DOE Offices and Industry* John Turner, ORNL
9:35am	*Observations on the SciDAC Software Experience* Bob Lucas, USC
10:00am	*A Facility's View of User Software Engineering Practices* Richard Gerber, NERSC
10:25am	BREAK
10:40am	*Preparing Scientific Software for Exascale* Jack Wells, OLCF
11:05am	*Overview of the IDEAS Software Productivity Project* Lois Curfman McInnes, ANL
11:30am	*Panel: Addressing DOE's Software Engineering Challenges* Panelists: Jim Davenport, BES; Thuc Hoang, ASC; Dorothy Koch, BER; John Shultz, EM; Moderator: Mike Heroux, SNL