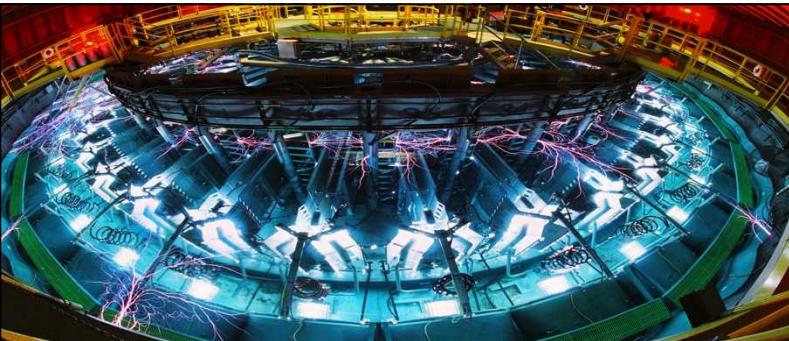


Exceptional service in the national interest



Challenges of Codesign

(Panel Discussion)

S.D. Hammond

Scalable Computer Architectures
Sandia National Laboratories, NM



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Example of Great Codesign

- **Example:** Development of optimized basic math kernels for Trilinos
- **Why?**
 - Initial benchmarking results showed we were not achieving the performance level we could achieve (performance analysts/test beds)
 - Mini-app work used to explore performance bottlenecks (performance analysts, hardware vendors)
 - Exploration of “fixes” in programming models, compilers and Kokkos (application/library developers, system software and performance analysts)
 - Analysis of bottlenecks in hardware with vendors now showing up in products (2015/6 onwards) (hardware specialists)
- **Impact:** Trilinos 12.0/2015 release will contain faster primitive mathematics kernels for some operations (designed for Trinity and CORAL)
- **Warning:** this took longer than we thought it might, maybe 2 years?



+ Carter Edwards and Matt Bettencourt

Successes (for Production?)

- I think there are many, not all obvious or things we can say without NDAs
- There are definitely changes in future hardware
- Kokkos and RAJA in my view look like incredible tools for the future
 - Not so sure we have an equivalent for Fortran but we are looking
- Vast number of changes in compilers and runtimes (particularly OpenMP)
 - Capabilities driving requirements for new features
 - Performance studies showing areas of concern
- Libraries are seeing huge innovations
 - New kernels, optimizations from vendors, some open source *etc*

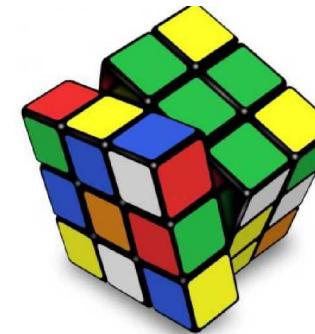
What are we still lacking?

- My gut feeling is that we still don't have enough applications developers/analysts in the loop, or at the very least enough buy in
 - Busy on adding physics and engineering codes
 - Busy with the day job of trying to do science
 - No real established programming model to rule them all
 - Hardware choices still seem in flux
- Poorly understood relationship of mini-applications to full applications
 - What is the coverage we *really* have?
- Poorly understood performance issues in applications
 - “My code is slow here .. Why? .. Because it was slow there 10 years ago when I last ran a profiler”

What are we still lacking?

- Not really sure I understand what is happening with tasking runtimes
 - Most of us agree we like the approach
 - Most of us can't agree we like any individual model
 - In the end this is a huge problem for NNSA production applications
- Still not sure I see attention to power or reliability bubbling to the top
 - These may take time to mature
- My feeling is almost out of runway in production computing
 - Remembering that we have to (re-?)develop codes, validate the outputs and get the analysts to be happy
 - Debugging and validating non-determinism because of in-node threading *could be very hard*
 - Are the tools *really* there to make this happen?

The Real Challenge



- ..isn't any single application
- We aren't focused enough on codesigning our workflow
- That means we need to understand what the users are doing from the moment they start work to the moment they go home
 - I/O? Analysis? Meshing? Creating input decks? Submitting jobs (and checking the queue every 5 minutes to see they are running)?
- Poor understanding of how we are going to see alignment with analytics and data sciences
 - But critical to our success?
 - The real driver in the industry?

National Computing Initiative

- **If NCI is funded, how will this change codesign efforts?**
 - Can expect to see vendors more engaged (funding) = more lab time?
 - Pressure to push harder, particularly into the unknown/higher risk
 - Likely to see more aggressive push on technology leading to great variation in hardware, programming models and system software
 - Greater opportunity to influence?
- **Changing Codesign Activities:**
 - Need to work out how to establish better links back to the application communities, in the end they own the world
 - More sophisticated tooling, profilers, performance analysis *etc*
 - Can we embed vendor staff in the labs? “open” up our applications?
- **Still need to understand NCI is small in comparison to industry as a whole**
 - Facebook, Google and Cell Phones will still dictate the broader market

What happens next?

- NCI and Exascale are not the end
- We need to be pipelining our research to ensure success isn't defined by it being used in Exascale machines
 - Some things are not going to deliver until later
 - Some things will become more important after Exascale
 - Reliability?





**Sandia
National
Laboratories**

Exceptional service in the national interest