SAND2013-4813C

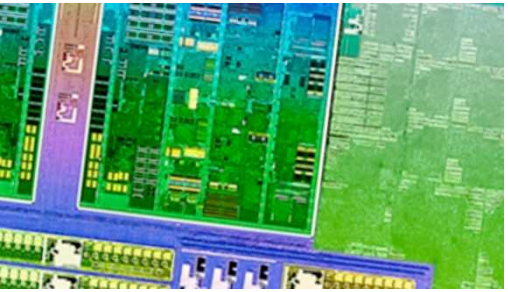

Sandia National Laboratories

Exceptional service in the national interest

Proxy and Proto miniApps for Exascale Co-design

James A. Ang, and Richard Barrett
Sandia National Laboratories
Albuquerque, NM USA

International Supercomputing Conference
Leipzig, Germany
June 17-20, 2013

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

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.

Our Exascale Efforts Focus on Co-design

Sandia National Laboratories

- Key Co-Design capabilities
  - Architectural Simulation Framework
  - Pre-production, First-of-a-kind Testbeds
  - Scalable R&D System Software
  - Mantevo miniApplications
- Exascale Implications
  - Sustained commitment of significant funding
  - Requires the Long View >5 years out

    —> time to *influence* Hardware Architectures

Which is Harder to change: HW or SW?

Sandia National Laboratories

- Conventional Wisdom
  - Hardware is difficult to change
  - Software is easy to change
- For the Long View, Conventional Wisdom is wrong!

MiniApps: Proxy and Proto

Conceptual Relationship among Full Applications, Proxy and Proto miniApps

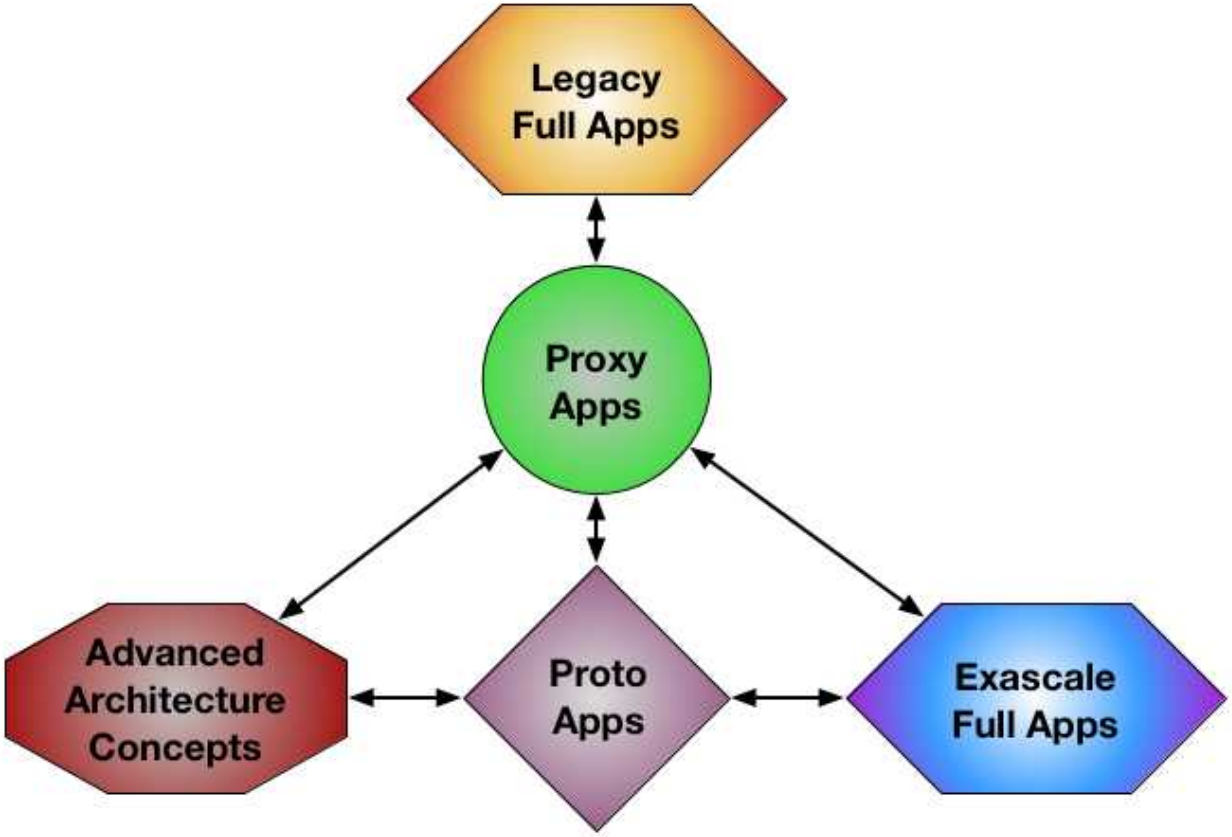

Representative Legacy app: CTH

- Eulerian multi-material modeling application.
- 3D, finite volume stencil computation.
- BSP with message aggregation (BSPMA).

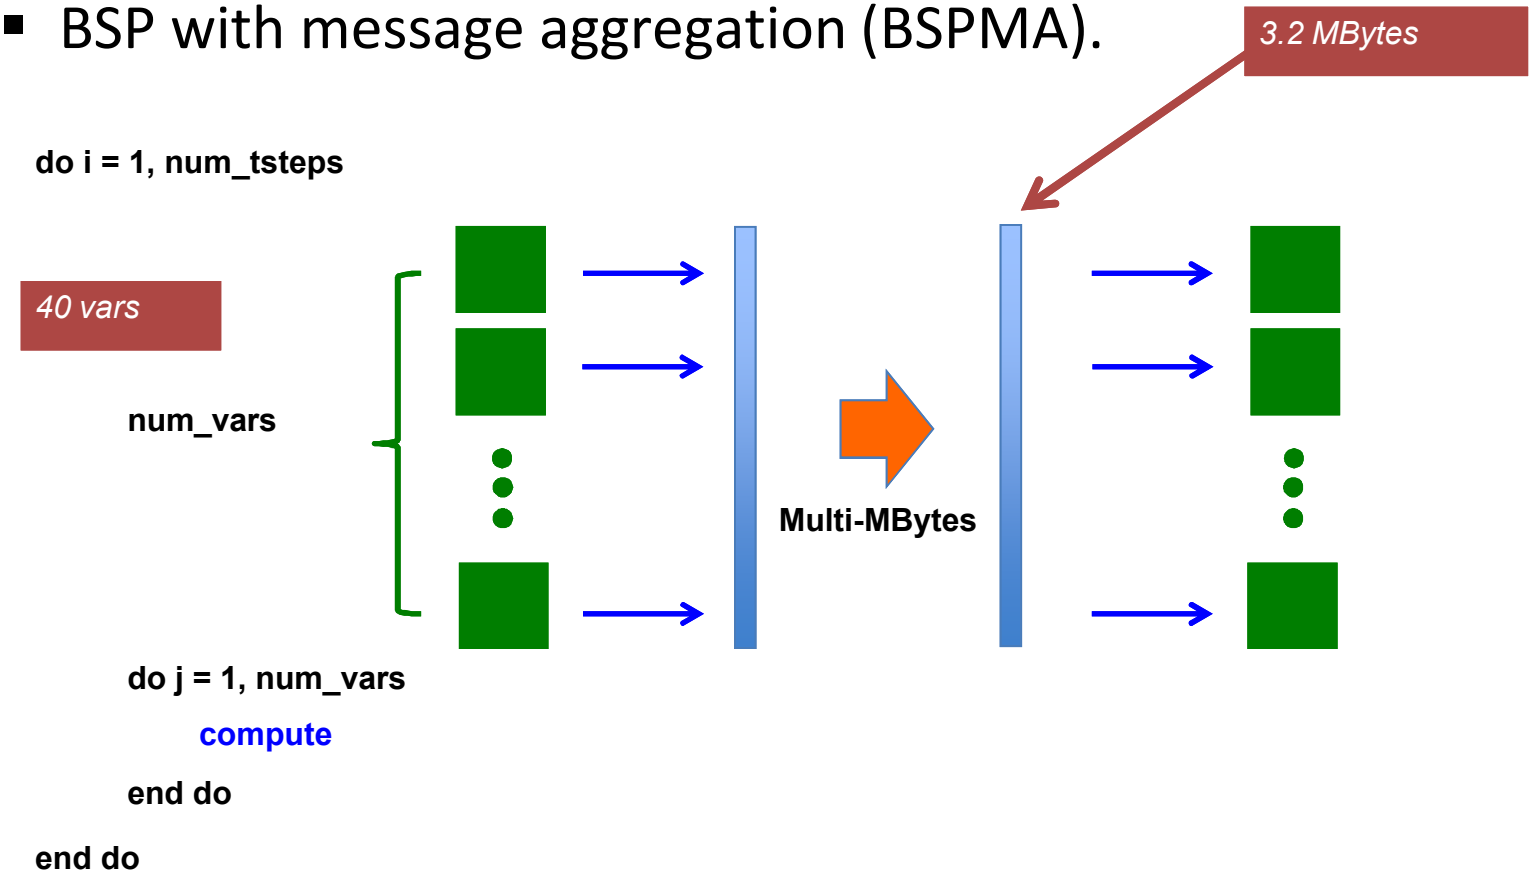

CTH and miniGhost performance

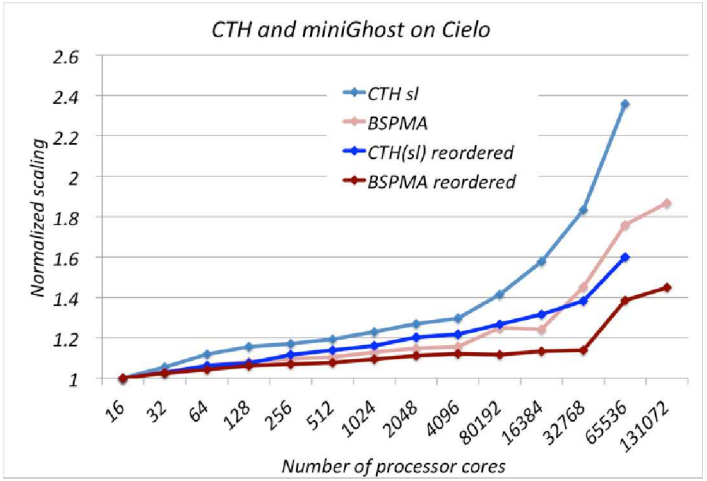

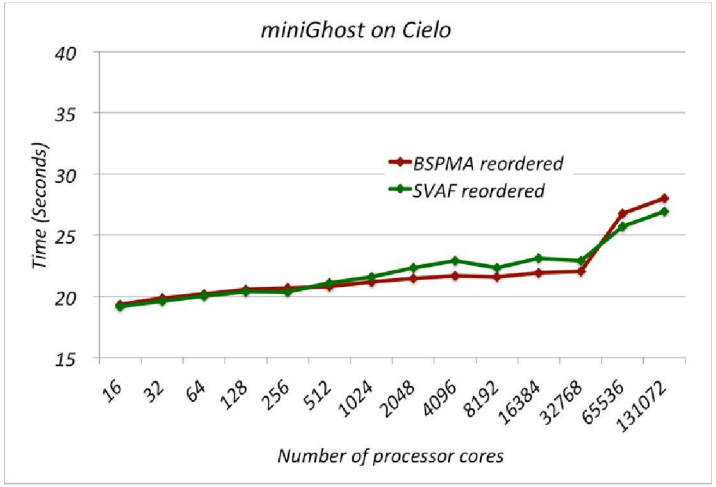

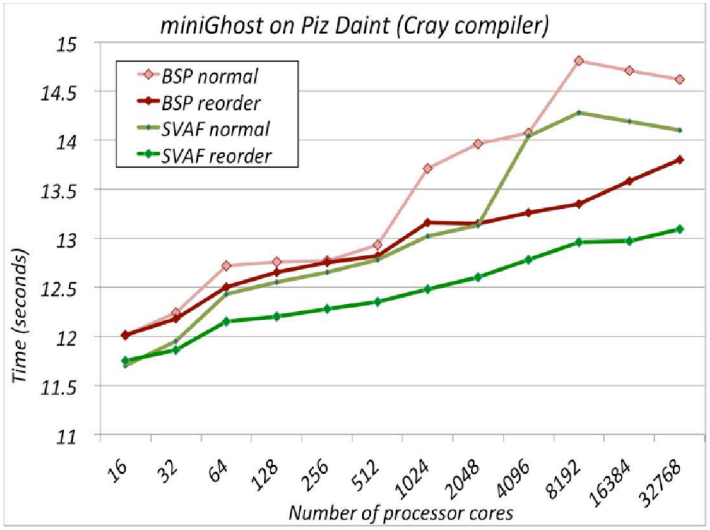

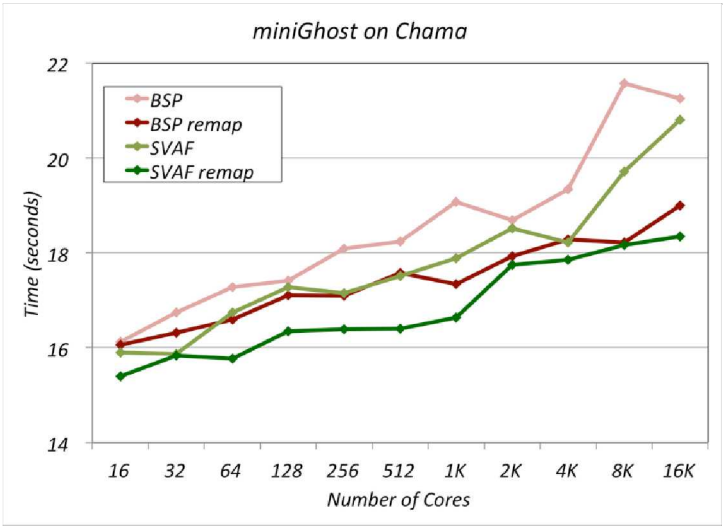

miniGhost: over-decomposition task parallel implementation

Sandia National Laboratories

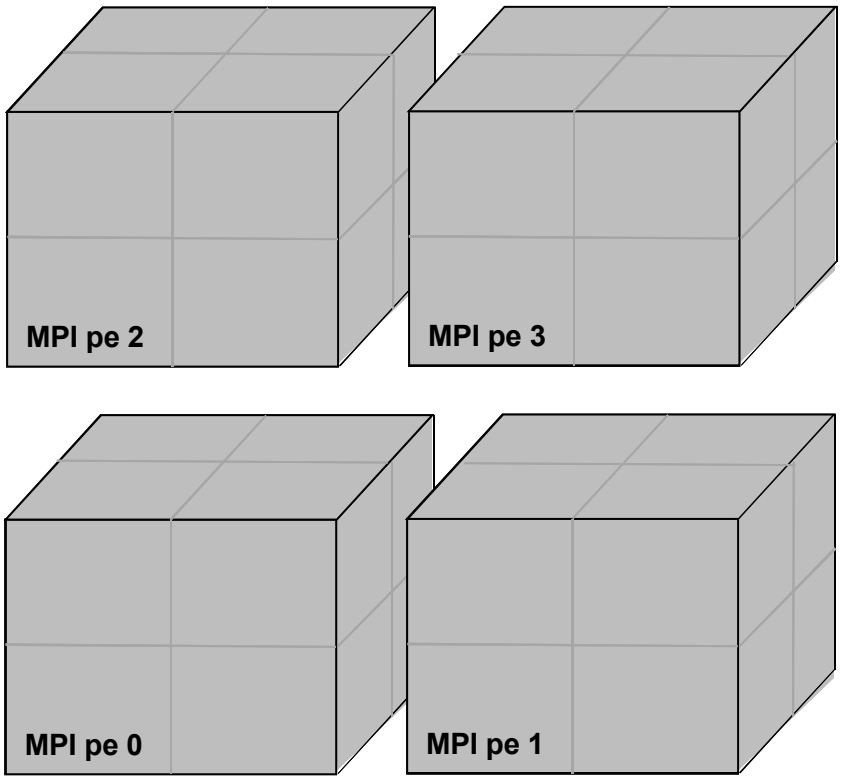

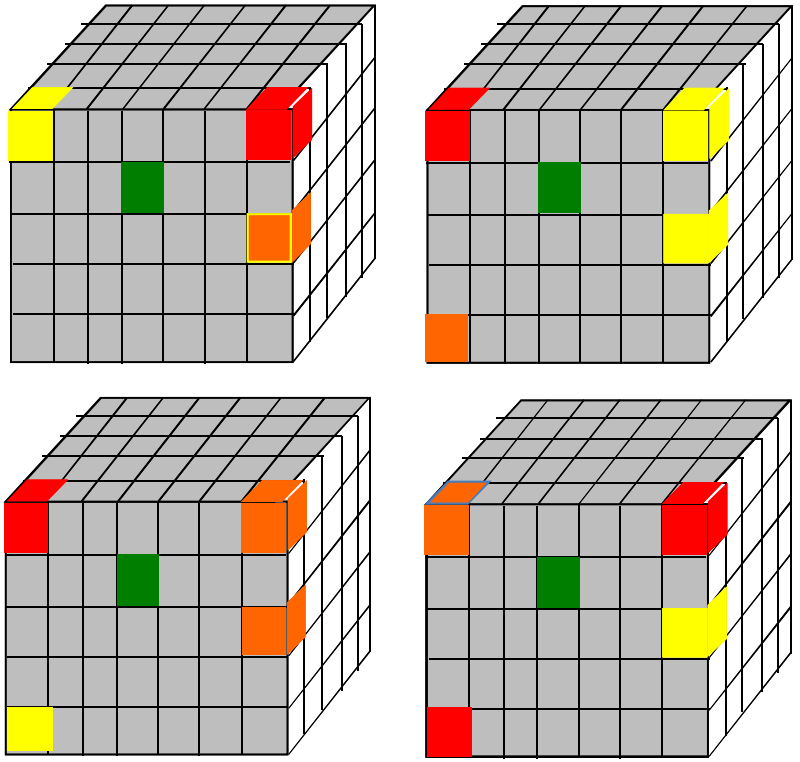

Data parallel

thread

Task parallel: some representative task workloads

Computation

Computation + MPI

Computation + BC

Computation + BC + MPI

miniGhost tp: Adding AMR

Sandia National Laboratories

- Diffusion over the 3D domain with random initial conditions and reflective boundary conditions.

  Two options:
  - uniform refinement, or
  - refinement based on the boundary or volume of an object being moved through the mesh and changing size. So, for example, a shock front can be simulated by refining based on a sphere which starts small and grows in size as the problem advances.
- Refinement within blocks.
- A block is refined into 8 blocks.
- Neighbors must be within one level of refinement.
- Computation is self-contained within a block.
- Communication aggregated to BSP model.
  - Excellent candidate for task parallelism version.

Implications for Co-design at a System Architecture Level

Sandia National Laboratories

- Case Study
  - Application – Task parallel implementation
  - System Software – integrated support for asynchronous, adaptive threads
  - Node Architecture – integrated support for light-weight threading
  - System Architecture – Interconnect Fabric with high Radix routers

- We know that BSPMA Applications create congestion problems for high-radix topologies

- Hypothesis–If applications are Asynchronous, Task Parallel (ATP), supported through runtime system software and multithreaded processors, the congestion issues will dissipate

Concluding Thoughts

Sandia National Laboratories

- The Multiple Dimensions of Co-Design
  - HW: Node and System architecture
  - SW: Application and System Software
- We need to rethink COTS and system balance
  - Component performance
  - Investment
    - Platform costs
    - R & D investments

Sandia National Laboratories

Backup Slides

Exascale Hardware Challenges

- Left to the *Invisible Hand*
  - Industry follows an evolutionary path focused on Peak Flops
- In the Era of Dennard Scaling our *ad hoc* approach to integration of MPPs with COTS microprocessors was acceptable
- With the end of Dennard scaling, this is no longer able to meet DOE Mission Application Requirements

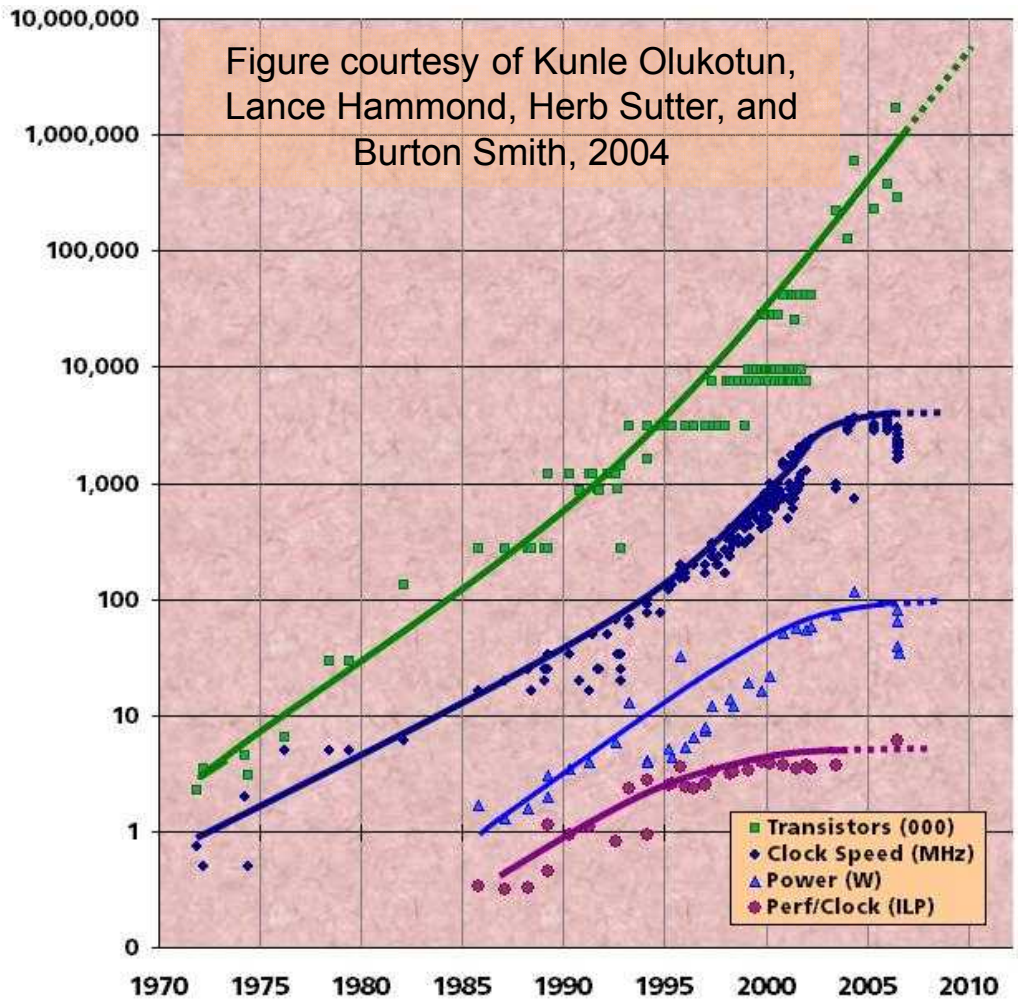

Define and Develop the Co-design Methodology for HPC

Sandia National Laboratories

- Key Co-design Capabilities
  - Mini Applications
    - Development and evolution to represent mission needs
  - HPC Architectural Simulators
    - Flexible to accommodate fidelity/speed tradeoffs
  - Proxy Architectures to explore advanced concepts
    - Abstract machine models
  - Advanced architecture testbeds
    - Evolving representation of vendor *state of the art*

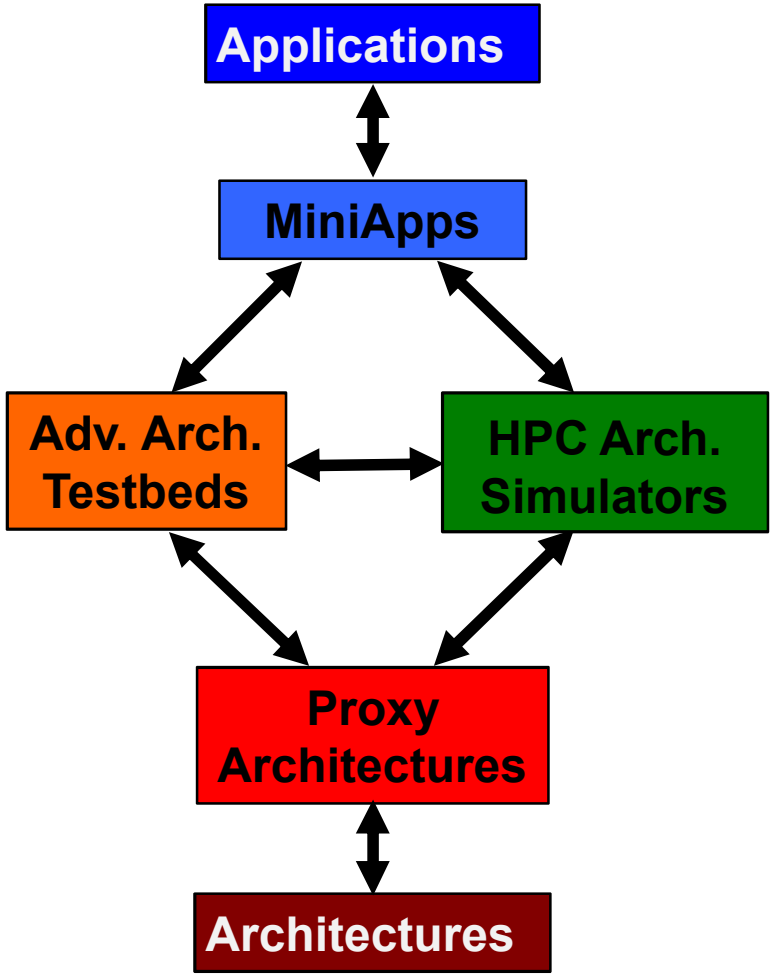

Paths to *Influence* COTS Development

Sandia National Laboratories

- Fund R&D Projects with Industry
  - Initiate Fast Forward R&D Projects with Industry
    - Patterned after Original ASCI Path Forward Program but improved
    - National Laboratory Staff are assigned to collaborate with Industry Partners via Co-Design activities, Proxy Applications, Proxy Architectures, system software, etc.
  - DOE is establishing R&D Projects with Micron Technology
- Explore SoC options for development of HPC COTS processors
  - Active discussions with ARM Holdings and several SoC companies, including some “traditional” companies such as Nvidia and AMD

Exascale Hardware Challenges

- We need to Motivate and *Influence* Architectural Changes
  - Processor Architectures
  - System Architectures
- Our Investments are not only in Architectures
  - We cannot just develop new Exascale Architectures and *Throw it over the wall* to our application developers
  - We need Hardware/Software Co-design
  - Later talks describe DOE application investments
- The transition of the DOE Legacy Code base is another important challenge
  - Also addressed by Applications talks
  - Challenge will influence hardware thru co-design