IESP Exascale Challenge:

Co-design of Architectures and Algorithms

Al Geist (ORNL) and Sudip Dosanjh (SNL)

Historically, huge supercomputers were built and delivered with little or no software on them. The application developers were left with heroic efforts to get there simulations to run efficiently on these systems. In order to improve the effectiveness of peta and exascale systems we need to have a paradigm shift where architectures and algorithms are co-designed.

There is a large gap between the peak performance of supercomputers and the actual performance realized by today's algorithms. This architecture-algorithm performance gap will get even wider with the increase in computing power being driven by a rapid escalation in the number of cores incorporated into a single chip rather than increases in clock rate. The transition from massively parallel architectures to multi-core architectures will be as profound and challenging as the change from vector architectures to massively parallel computers that occurred in the early 1990's that enabled our Nation and the U.S. Department of Energy to break the teraflop barrier. To effectively bridge this architecture-algorithm gap and use the next generation of computers, we must solve a host of architectural challenges in hardware and software.

Hardware challenges:

- Moore's Law still holds, but clock speed is constrained by power and cooling limits
- Processors are shifting to multi/many core with attendant hierarchical parallelism
- Compute nodes with hardware accelerators create the additional complexity of heterogeneous architectures
- Processor cost is increasingly driven by pins and packaging, which means the memory wall is growing in proportion to the number of cores on a processor socket
- Supercomputer architectures must be designed with an understanding of the applications they are intended to run
- A supercomputer architecture that performs well on full scale real applications cannot be built from only commodity components.

Software challenges:

- Scaling limitations of present algorithms
- Hierarchical algorithms to deal with bandwidth across the memory hierarchy
- Software strategies to mitigate high memory latencies
- More complex multi-physics requires large memory per node
- Need for automated fault tolerance, performance analysis, and verification
- Innovative algorithms for multi-core, heterogeneous nodes

Promoting the integrated co-design of architectures and algorithms represents a fundamental shift from simply procuring and operating large scale systems. A key way to lower the risk of

adopting novel architectures and technologies is to demonstrate through paper studies, system simulation, and hardware prototypes the performance benefit of these technologies. The vision is that both the hardware designers and the software designers will compromise based on what the other group can do in a given timeframe. The evolution of the architecture and algorithms then becomes more aligned, which helps close the performance gap. Deploying small prototype systems will facilitate application, algorithm and system software development, prove the technology to industry, and lower the risk of adoption of advanced architectures. The metrics for success will be measured through changes to product roadmaps, and integration or adoption of co-designed technologies into next generation supercomputer systems.