

Igniting Exascale

William Gropp

www.cs.illinois.edu/~wgropp



Panelists

- Pete Beckman
- Franck Cappello
- Al Geist
- Satoshi Matsuoka
- Thomas Sterling



Is there a Fear of Exascale?

- There is concern about
 - ◆ Faults
 - ◆ Scaling
 - ◆ Performance (latency)
 - ◆ Complex processing model (e.g., heterogeneous elements)
 - ◆ Cost (power, memory)
 - ◆ Impact on algorithms
 - ◆ Narrowness of application domain
 - ◆ Data handling for Exabyte data sets
- In addition, does Exascale imply a discontinuity in programming, algorithms, debugging, etc.?



How can We Overcome the Fear?

- Which fears are mistaken?
 - ◆ After all, many were convinced that petascale systems would be impossible without new programming models?
- Conversely, which problems apply at a smaller scale?
 - ◆ Thus can be addressed now and provide near-term benefits?
- Which problems are (nearly) unique to Exascale?
- How do we build/test/improve algorithms, software, and applications?
 - ◆ Do we need to build a much more sophisticated simulation environment?



How Can We Build Real Excitement?

- How do we provide evidence that Exascale systems will work well with applications?
- How do we demonstrate that Exascale systems can enable new application areas (after all, Exascale systems may be greatly different in architecture - will that be a virtue)?
- In all of the above, how do we move past qualitative statements to quantitative predictions?



How Much is Science Worth?

- Lakeside Technology Center
- 1.1 million ft²
- Over 100MW!
 - ◆ 2nd largest customer of ComEd (O'Hare is first)



How much is science that could transform our world worth?



Small (resource) Science vs. Large (resource) Science

- We have not made a compelling case that Exascale computing should be a priority
 - ◆ Why not data-centric computing?
Sensors? Privacy? Health?
- How to compare the different kinds of contributions?

