



The Goals for IESP

Rick Stevens \approx Argonne and University of Chicago

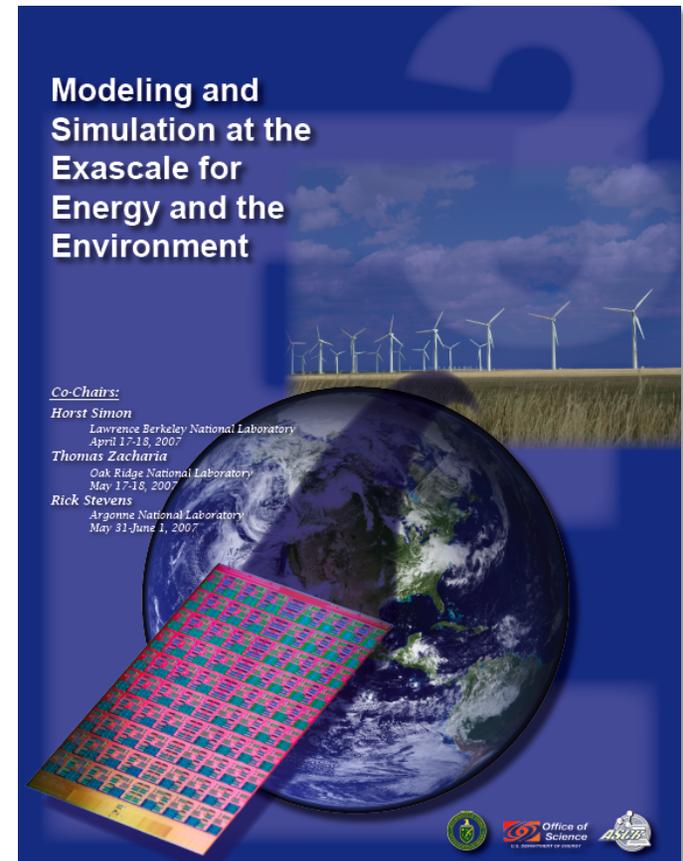
Overview of Planning for Exascale



- Since 2007 US Agencies (DOE, NSF and DARPA) have started planning for a *possible* Exascale Initiative
- The goal is to develop an exascale simulation and modeling capability by 2020
- Our vision includes significant international partnership
- The scope of the initiative concept includes
 - ▣ Applications Development
 - ▣ Systems Software and Programming Models
 - ▣ Multiple Hardware Platforms

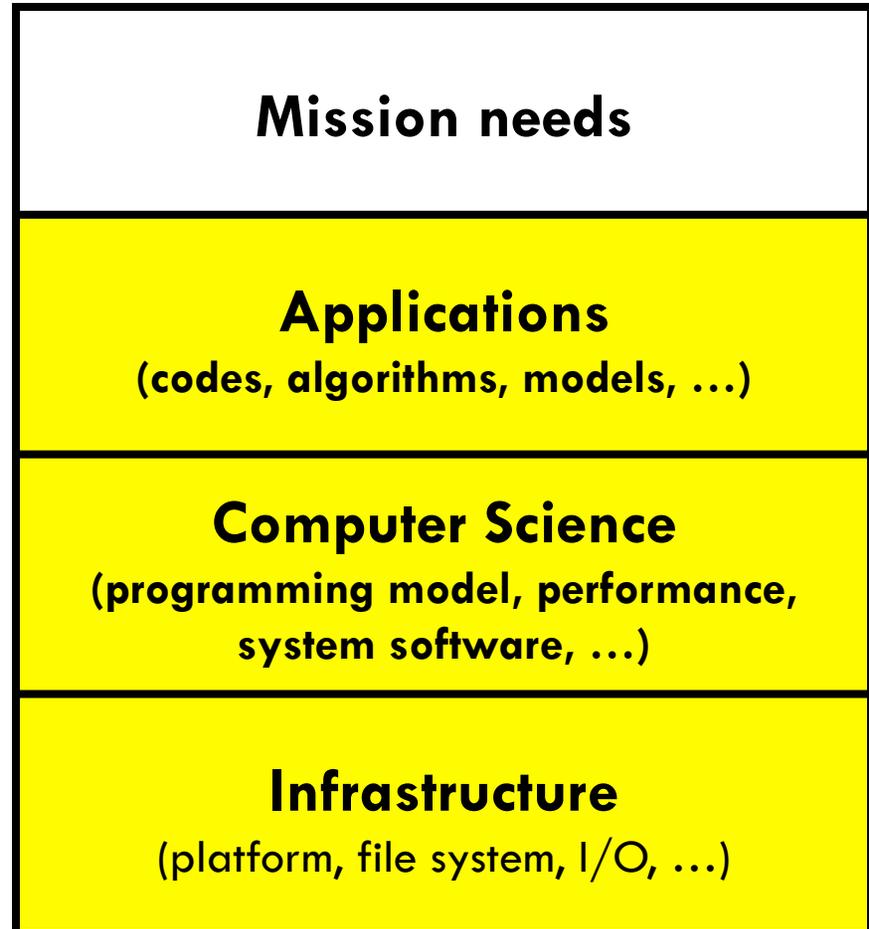
US Exascale Planning Activities

- Started in spring 2007
- DOE Lab Town Hall Meetings
 - ▣ Energy and Environment
- DARPA studies
 - ▣ Hardware and Software
- DOE and NSF Launch IESP
 - ▣ Roadmap for Software
- DOE Lab Steering Committee
 - ▣ Science Case (Applications)
 - ▣ Technical Roadmap



DOE's focus on mission needs in Energy, Environment, Science and Security

- DOE Exascale Steering Committee
 - ANL, LANL, LBNL, LLNL, SNL, ORNL + PNL, BNL
 - Charter: Decadal plan to provide exascale applications and technologies for DOE mission needs
- Workshops @ ~100 People
 - Climate Science (11/08)
 - High Energy Physics (12/08)
 - Nuclear physics (1/09)
 - Fusion Energy (3/09)
 - Nuclear Energy (5/09)
 - Biology (8/09)
 - Basic Energy Science (8/09)
 - Joint National Security (10/09)
 - Computer Science
 - Mathematics
 - Computer Architecture



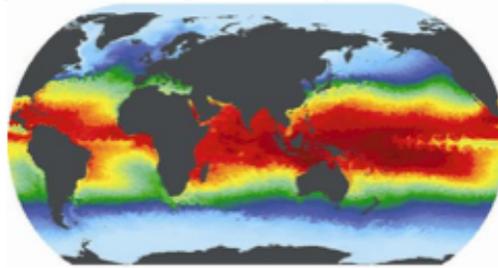
Applications Targets



- Energy
 - ▣ Solar, Grid, Wind, Storage
 - ▣ CFD and Combustion
 - ▣ Fission and Fusion
- Environment
 - ▣ Climate and Ecosystems
 - ▣ Carbon Sequestration and Storage
- Security
 - ▣ Stockpile Stewardship
 - ▣ Non-proliferation
- Basic Science
 - ▣ High-Energy and Nuclear Physics
 - ▣ Astrophysics and Cosmology
 - ▣ Biology, Chemistry, Materials
- Health
 - ▣ Neuroscience
 - ▣ Protein science
 - ▣ Genomics
 - ▣ Cells, Tissues and Organisms

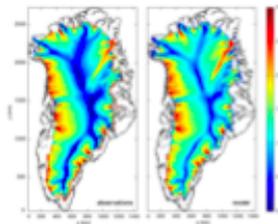
Simulations are critical for understanding climate change and predicting its impacts

- Sea level, sea-ice and ocean circulation
 - Develop a predictive land-ice model
 - Increase resolution of ocean < 1/10°
- Distribution and cycles of water, ice & clouds
 - Global, cloud resolving models
- Extreme events on local and regional scales
 - Hierarchical multiscale models
- Carbon, methane and nitrogen cycles
 - Biogeochemistry in oceans
 - Dynamic land surfaces changes



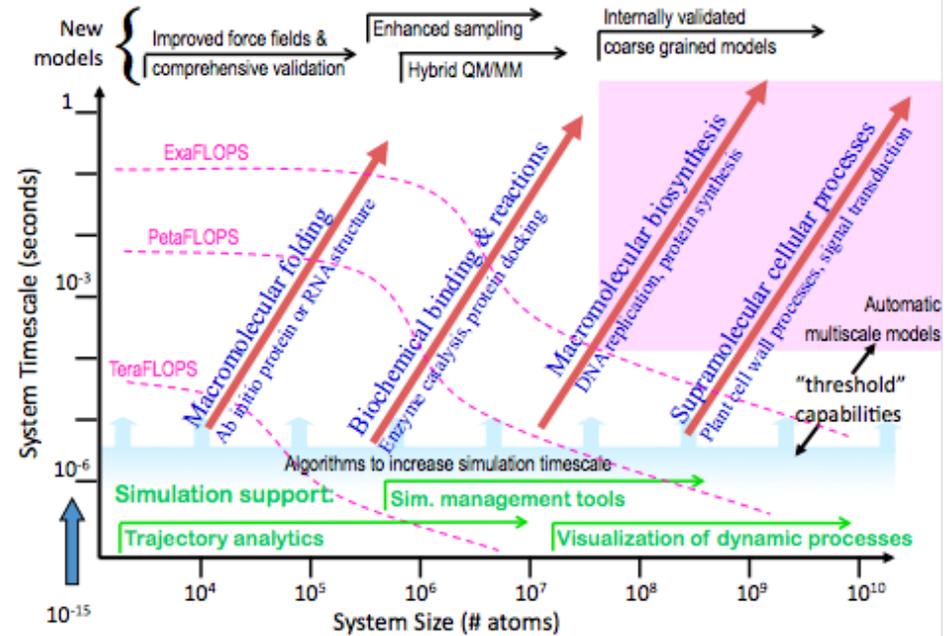
Eddy-resolving ocean simulations required to give accurate simulation of global ocean circulation

Advanced sea ice and ice sheet models

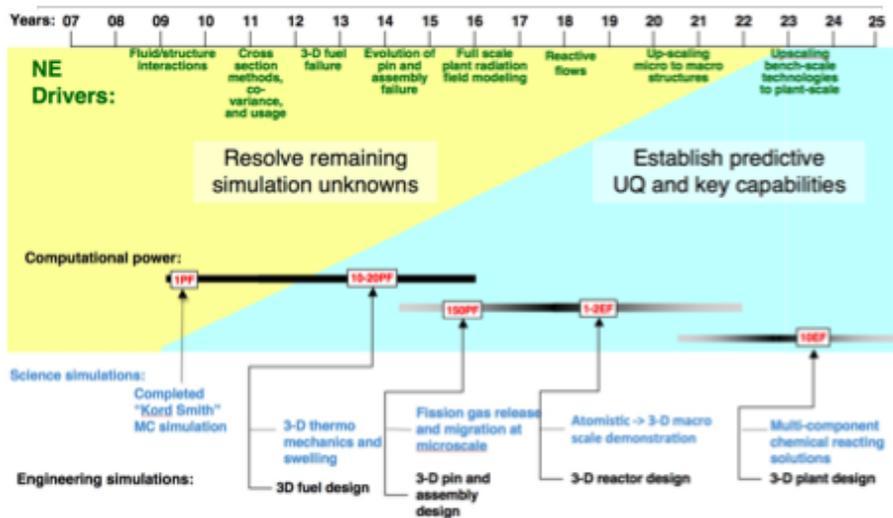


"Given these drivers ... it is clear that exascale computers and ultra fast networks, data systems and computational infrastructure will be required by 2020." *Challenges in Climate Change Science and the Role of Computing at Extreme Scale, November, 2008*

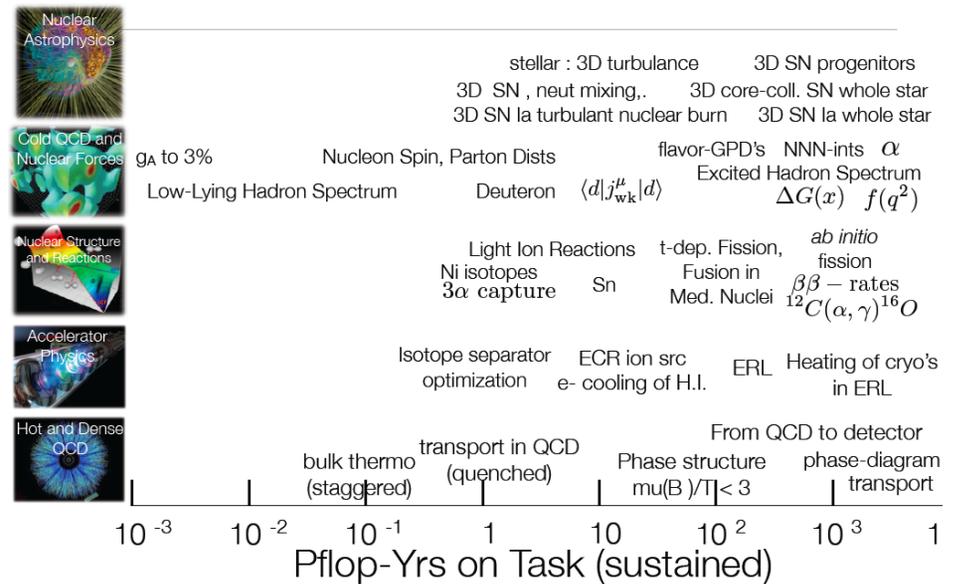
Towards Exascale Biochemical Simulations



Computational Requirements for Nuclear Energy Modeling



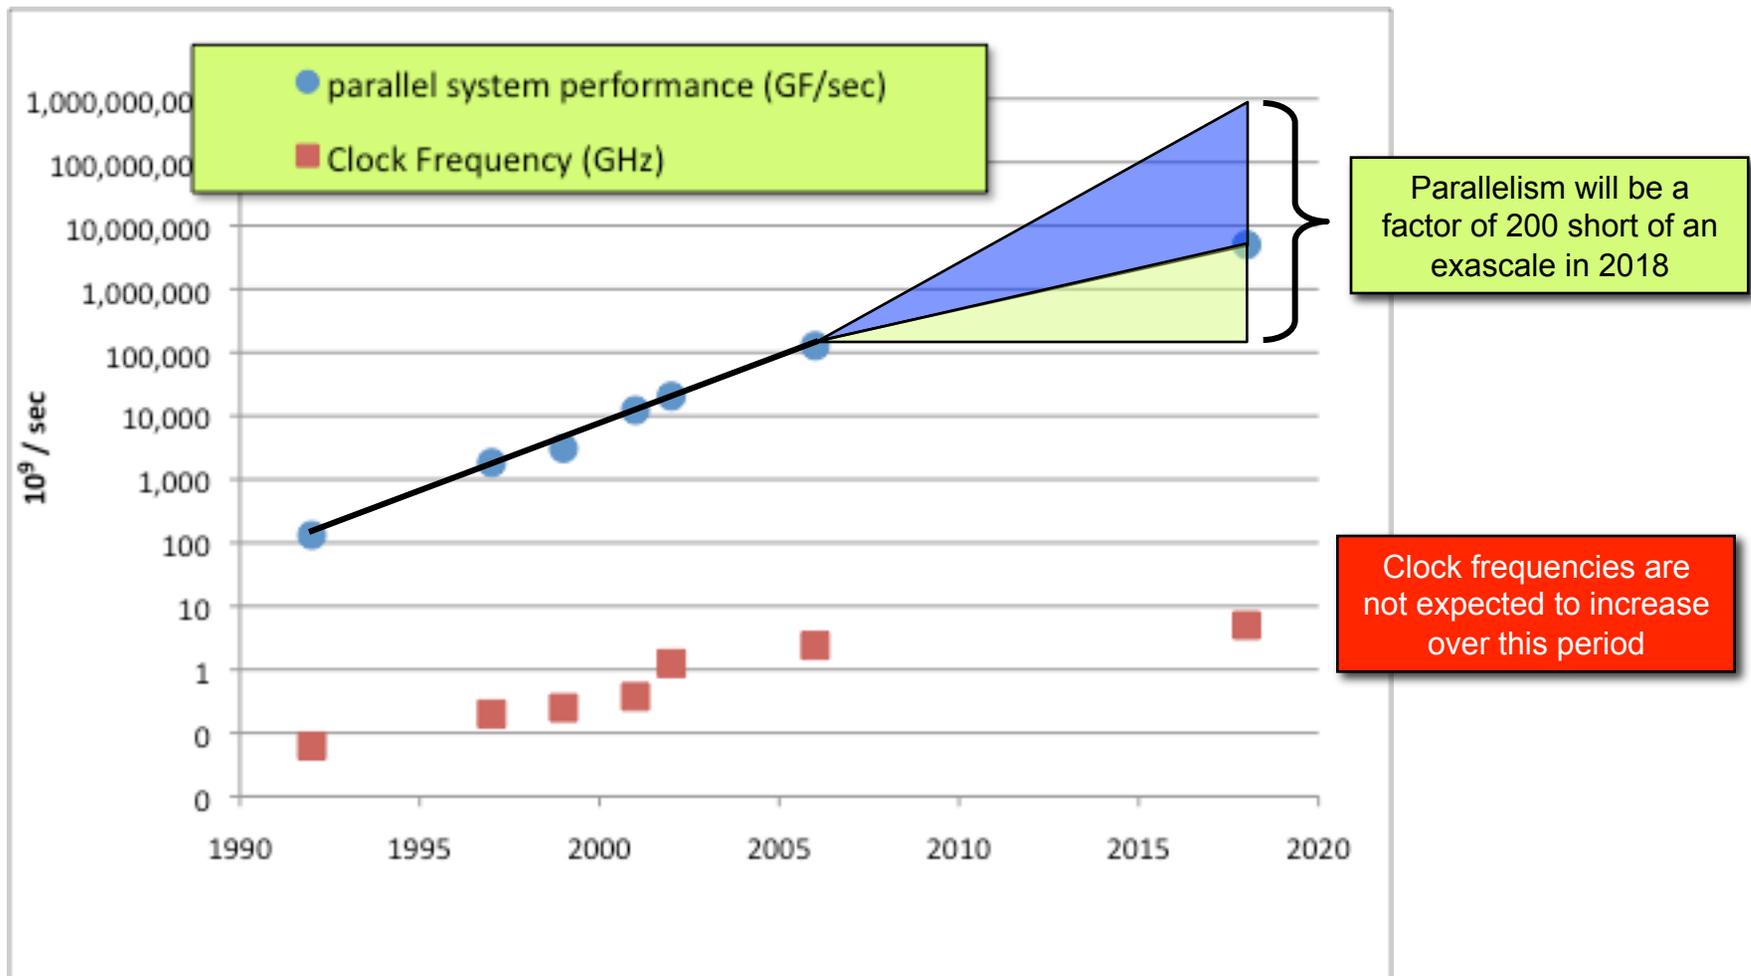
Nuclear Physics Requires Exa-scale Computation



Systems Scaling Projections

Begin Full System Delivery (Yr)	2004	2007	2012	2015	2019
Design Parameters	BG/L	BG/P	25PF	300PF	1200PF
Cores / Node	2	4	8-24	32-64-128	96-128-500
Clock Speed (GHz)	0.7	0.85	1.6-4.1	2.3-4.8	2.8-6.0
Flops / Clock / Core	4	4	8-32	8-32	16-64
Nodes / Rack	1024	1024	100-1024	256-1024	256-1024
Racks / Full System Config	64	72	128-350	128-400	256-400
MB RAM/core	256	512	1024-4096	1024-4096	1024-4096
Total Power	2.5MW	4.8MW	8MW-20MW	20MW-50MW	30MW-80MW
Flops / Node (GF)	5.6	14	128-640	640-2000	2000-6000
Flops / Rack (TF)	5.7	14	200-400	400-1200	1600-4800
LB Concurrency	5.E+05	1.E+06	10E6-64E6	100E6-1E9	1E9-10E9
Full System					
Total Cores (Millions)	0.13	0.3	.3M-1.5M	1M-50M	4M-200M
Total RAM (TB)	33.6	151	2,000-4,400	3,000-10,000	5,000-50,000
Total Racks	64	72	128-350	128-400	256-400
Peak Flops System (PF)	0.37	1	25	300	1200

A Hybrid Programming Model Will Likely Be Required To Reach An Exaflop/Sec



The Concurrency Challenge

- **How will 1000x increase in concurrency be delivered to Extreme Scale Systems?**
 - Weak Scaling by increase in data size, W_S
 - Weak Scaling by increase in computation per datum, W_T
 - Strong Scaling by increasing parallelism in fixed computation, S
- **Over Provisioning required for Latency Hiding, L**
 - L is determined by latency hiding requirements --- memory access time, fraction of memory operations (parameterized by memory hierarchy level)
 - Assume $L \sim 100x$
- **How much additional parallelism must come from software, relative to today's applications?**
 - $W_S * W_T * S$ needs to be $> 10^5$

The Concurrency Challenge

- **W_S is limited by DRAM cost and memory-computation ratio**
 - Petascale ~ 1 byte/FLOPS
 - Exascale ~ 0.01 bytes/FLOPS ??
 - Improving memory-computation balance is expensive
 - 50-50 cost balance dictates 0.004 bytes/FLOPS
 - 90-10 cost balance dictates 0.04 bytes/FLOPS
 - Assume $W_S \sim 10x$
 - **W_T is limited by application domain and algorithm**
 - Assume $W_T \sim 10x$
- ➔ Software must deliver $S \sim 1000x$ increase in Strong Scaling to use full capability of Extreme Scale systems**

Four Goals for IESP



- **Strategy for determining requirements**
 - clarity in scope is the issue
- **Comprehensive software roadmap**
 - goals, challenges, barriers and options
- **Resource estimate and schedule**
 - scale and risk relative to hardware and applications
- **A governance and project coordination model**
 - Is the community ready for a project of this scale, complexity and importance?
 - Can we be trusted to pull this off?

Goals for IESP



- **Strategy for determining and refining “systems” software requirements**
 - Agencies need to understand how scope is being determined and how definitive that scope is
 - Need strong coupling to the applications science and engineering groups to better understand their problems
 - Need coupling to architecture development teams
 - Need to understand what the private sector will be able to deliver and when
 - Need to understand the role of research in generating options for the future

Goals for IESP



- **Develop a comprehensive community software roadmap for Exascale systems**
 - Identify those software capabilities that will be needed for fully functional exascale systems, what are the barriers and how can we overcome them
 - Determine which elements will occur naturally and which elements need R+D investment
 - Determine those components that have solid starting points and which that need *ab initio* efforts
 - Determine which components are suitable for an open community development model

Goals for IESP



- **Develop an estimate of the resources required and timeline needed to develop the required software**
 - Need to put the software element of exascale in appropriate budget and schedule context
 - Need to understand the risks (technical, schedule and organizational)
 - Need to distinguish between the applications software efforts and the systems software
 - The software timeline should be aligned with that of the hardware (and precede it where possible)

Goals for IESP



- **Develop a governance and coordination model that encourages international participation and open development and yet capable of meeting the deadlines for deploying on real systems**
 - The software aspect of Exascale is likely to be as difficult and expensive as the hardware and applications development
 - The agencies can not afford to try a new community model if they don't have confidence it will result in delivering the needed software on time and on budget

Final Words



- This is an important project
- Its success will have lasting impact on the planet and people all around the world and for generations into the future
- I believe our community is ready to step up to the challenge to engage in a worldwide effort to build the common software that will benefit all
- If we do this right it will also be fun!