DOE Exascale Initiative

Dimitri Kusnezov, Senior Advisor to the Secretary, US DOE
Steve Binkley, Senior Advisor, Office of Science, US DOE
Bill Harrod, Office of Science/ASCR
Bob Meisner, Defense Programs/ASC
Briefing to the Secretary of Energy Advisory Board, September 13, 2013

The Department of Energy is planning for a major computer and computational science initiative anchored in our mission challenges.

We use the term 'exascale' to capture the successful transition to the next era of computing in the 2020 timeframe. Change is coming in many ways, not just technology. For instance:

	World we were in	World we are going to
Processing/ Memory:	Processing expensive/ memory free	Processing free/ memory expensive
Ecosystem	US micro-electronics supply chain	Globalization of supply chain
	US government driven	Mobile device driven
Cyber	Guards/Fences – ‘local’	Sophisticated and organized adversaries – ‘global’

Traditional path of 2x performance improvement every 18 months has ended. The result is unacceptable power requirements for increased performance

U.S. DEPARTMENT OF ENERGY

Goals

The DOE Exascale Initiative will drive US leadership in several dimensions:

- *Deliver the next era of high-end computational systems*
- *Build-in cyber defenses including information protection*
- *Provide technologies to ensure supply-chain integrity*
- *Attack previously unsolvable problems critical to the nation's national security, science and innovation engine*
- *Deliver tremendous improvements in electronic component energy efficiency, reducing operations costs of HPC centers and carbon footprints*
- *Stimulate US competitiveness and intellectual property through aggressive technology development and adoption*

Exascale will be the culmination of a technological path defined and developed over a next decade.

It is about solving some of our nation's more critical problems while driving US competitiveness.

U.S. DEPARTMENT OF ENERGY

What actions are needed and when?

What is your confidence?

How do you bring science to bear into the decision process?

What are the risks?

What happened?

Can it happen again?

…

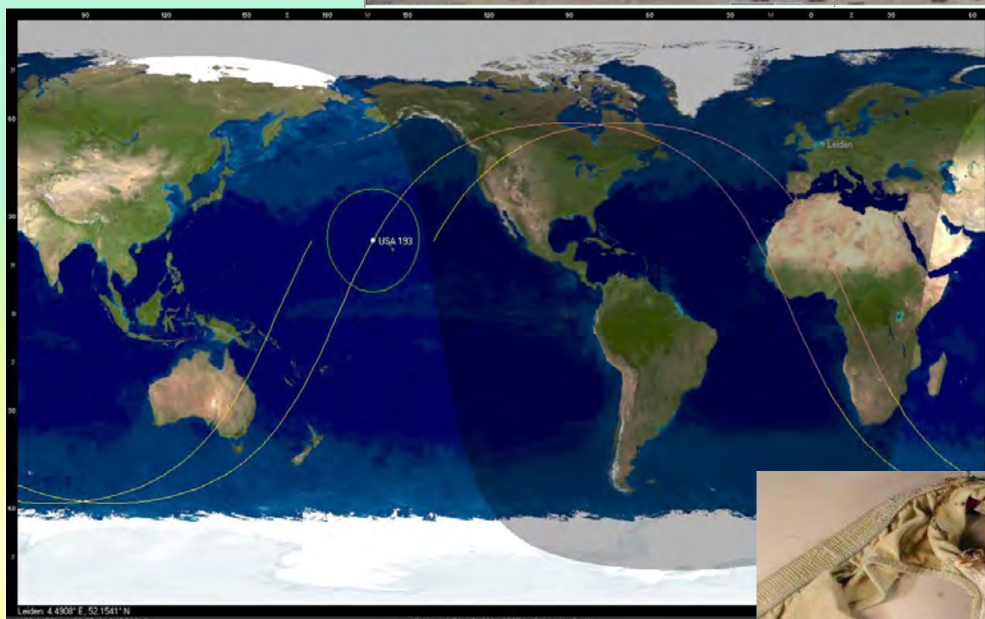

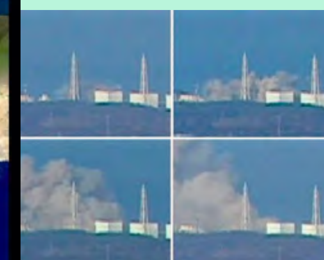

Prediction is Simple…

- Prediction examples surround us:
  - First prediction of sun eclipse by Thales of Miletus *(585 BC)*
  - Tide tables
  - Lifetime of first excited state of hydrogen
  - Laminar-/turbulent-flow drag of a thin plate, aligned with flow direction
  - Performance prediction for a new aircraft
  - National Hurricane Center forecasts
  - Climate prediction
  - 2014 World Cup winners
  - Astrology

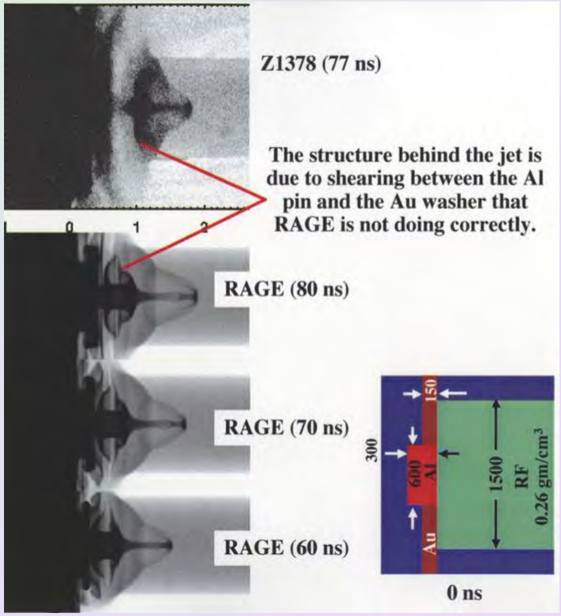

… But consequences of poor predictions are not all the same.

Today we are turning to simulation for increasingly serious problems and to push the frontiers of science and technology.

Alignment with Presidential Priorities

- Climate Action Plan
  - 'Resiliency' against natural and man-made conditions
  - DOE is the Sector-specific agency for Energy Infrastructure
  - Cross-sector dependencies of US critical infrastructures

- Nuclear Security Agenda (Prague, Berlin, NPR)
  - Preventing nuclear proliferation and nuclear terrorism
  - Reducing role of weapons
  - Maintaining strategic deterrence with reduced forces
  - Strengthening regional deterrence and reassuring allies
  - Ensuring a safe, secure and effective stockpile

Confidence in predictions here are critical. Today we make decisions with B$ impacts that we could not do just 2-3 years ago. But we need to do better.

Mission: Extreme Scale Science

Mission-Focused Applications

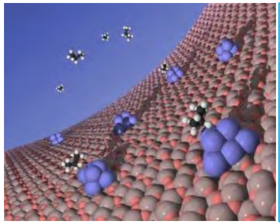

New Catalysts

10,000 atom quantum simulations

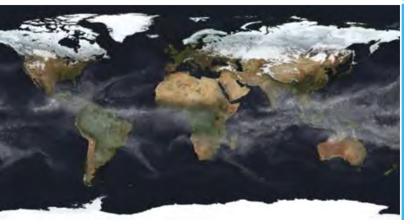

Climate Models

Large-ensemble multi-decadal predictions

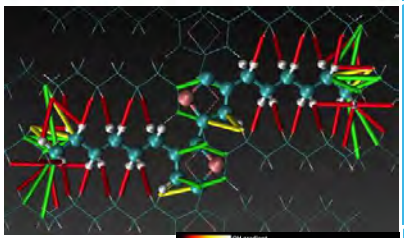

Materials Design

ab-initio electronic structure methods for excited states

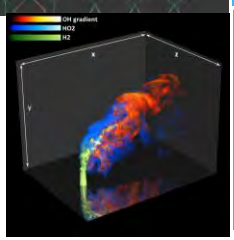

Combustion

Turbulent, chemically reacting systems

- **Improved energy technologies:** Photovoltaics, Internal combustion devices, Batteries
- **Manufacturing Technology**
- **Novel materials:** Energy applications, Electronics

Stockpile Science

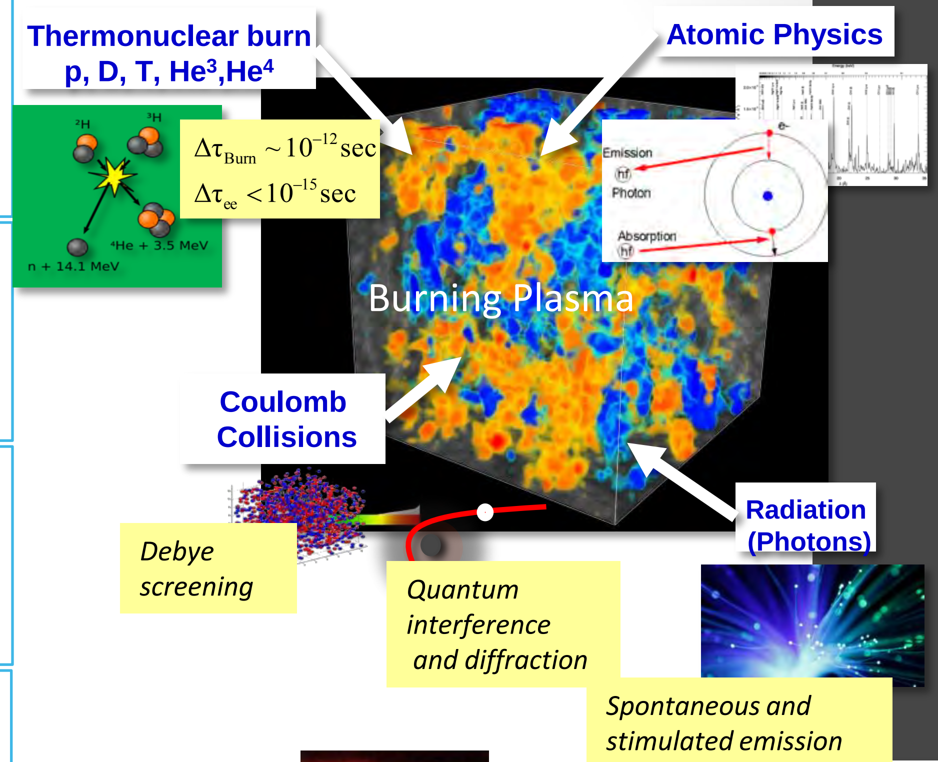

Non-Proliferation and Nuclear Counter Terrorism

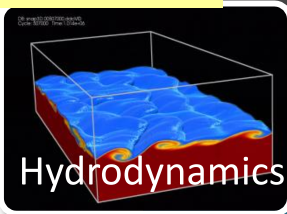

U.S. DEPARTMENT OF ENERGY

Advances in simulation is the common rate-limiting factor

Our Data is Exploding

Genomics

Data Volume increases to 10 PB in FY21

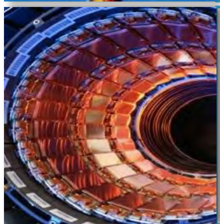

High Energy Physics (Large Hadron Collider)

15 PB of data/year

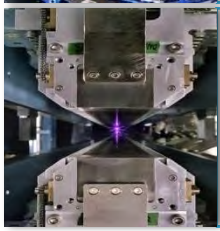

Light Sources

Approximately 300 TB/day

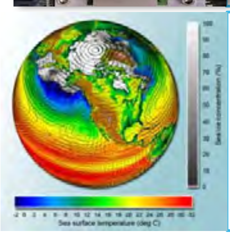

Climate

Data expected to be hundreds of 100 EB

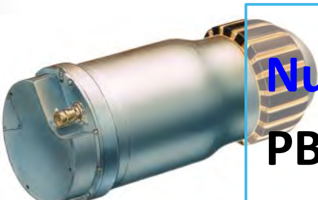

Nuclear Security

PBs per simulation

Driven by exponential technology advances

Data sources

- Scientific Instruments
- Sensors/ Treaty monitoring
- Scientific Computing Facilities
- Simulation Results

Big Data and Big Compute is the place we operate

- Analyzing Big Data requires processing (e.g., search, transform, analyze, …)
- Extreme scale computing will enable timely and more complex processing of increasingly large Big Data sets

Remarkable progress in last 15 years, but much more to do

- ASCI Red: Comissioned Dec 1996
- Just under 10,000 processors
- 1 TF
- 2500 sq ft
- 0.5 MW

- Intel: 1 TF chip (products within 2 years)
- Size of a pack of gum
- 100-250 W
- Prototype was run in Nov 2011

To get to exascale, one could imagine:

- 1 million, 1 TF processors
- 64 quadrillion bytes of memory
- high-performance interconnect
- software to make it work
- all must fit within 20MW envelope

Example of potential Exascale System	
System peak	1,000 PF
System power	~20 MW
System memory	≥ 64 PB
Storage	500 - 1000 PB
I/O aggregate bandwidth	60 TB/s
MTTI	~1 day

U.S. DEPARTMENT OF ENERGY

Memory Bandwidth & Data Movement Performance will be Energy limited

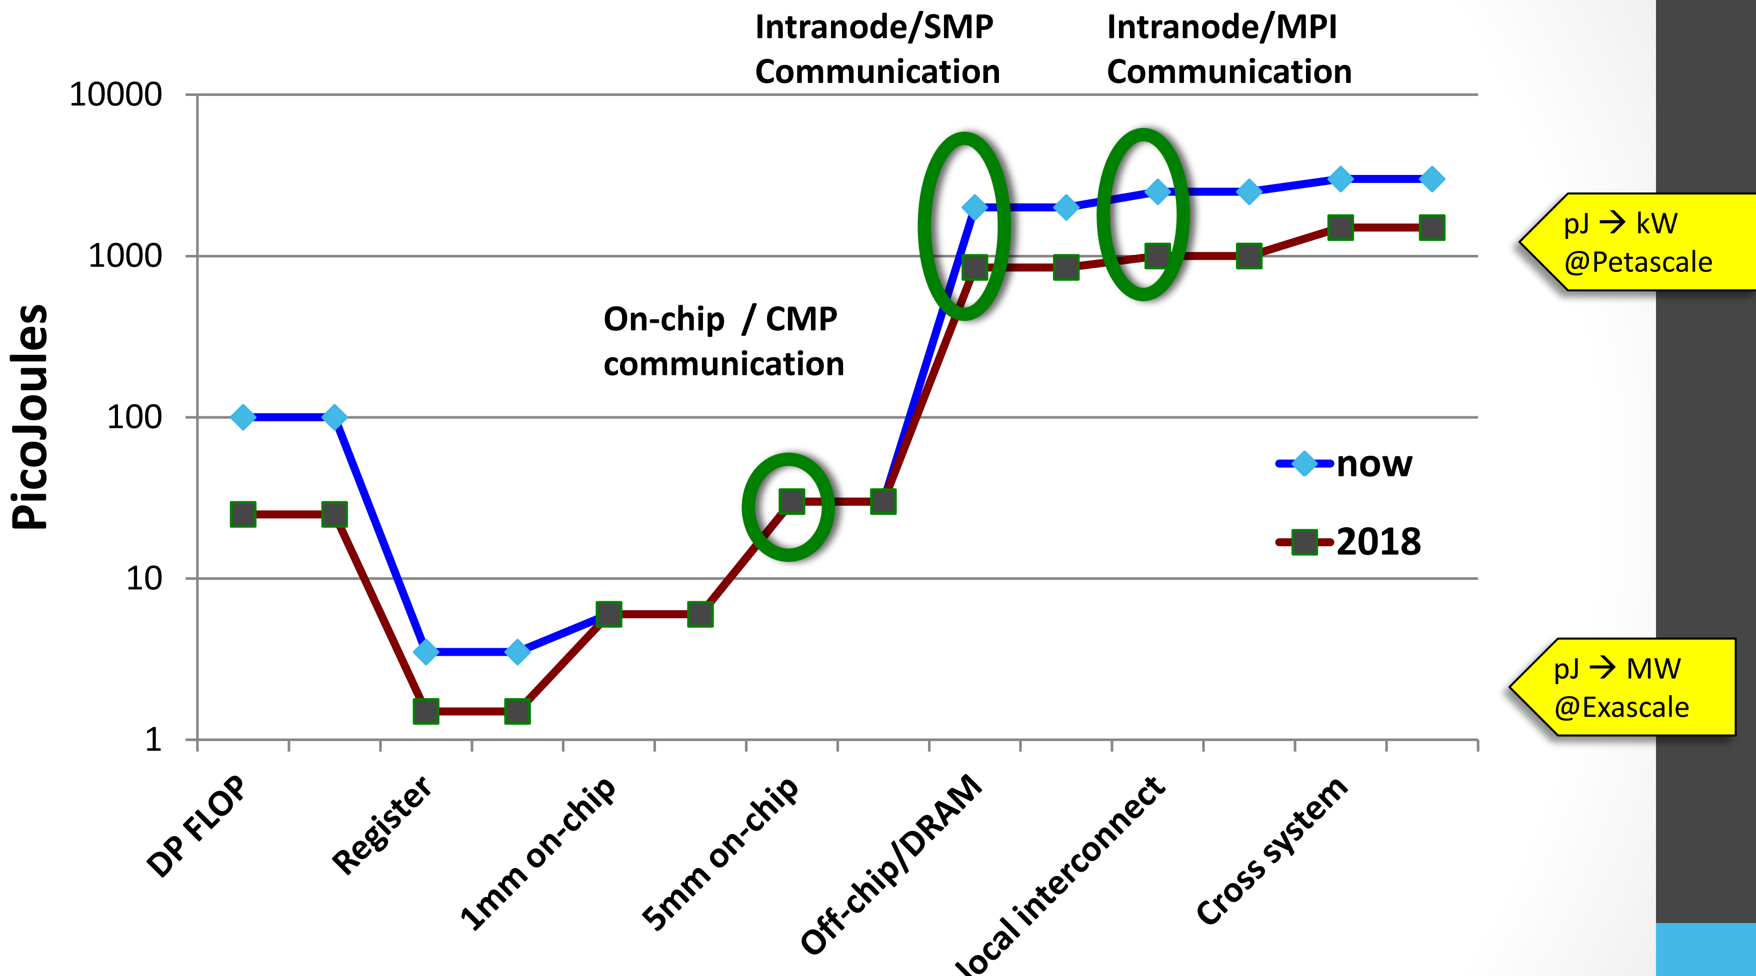

ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems, Peter Kogge, Editor and Study Lead, DARPA-TR-2008-13, 2008.

U.S. DEPARTMENT OF ENERGY

Some of our considerations in getting to the Exascale

- No obvious path to provide Exascale computing within a reasonable power envelope by 2020 timeframe
- We are thinking in the many dimensions. It includes:
  - Million processor systems, with each processor having 1000 cores $\rightarrow$ O(10^9) cores
  - O(10^9) threads with no clear path to programmability, debugging, optimization
  - New programming paradigms
  - Dealing with massive data sets:
    - Our TF systems created PBs of data; Might expect EF systems to create YB of data.
  - HW and SW design inexorably linked to attack problems of reliability
  - Sub-threshold logic
    - 10^{18} operations/sec x 10^{-12} J/operation = 1 MW. But today memory/cache access and average cost of operations is into the 100s of pJ.
  - The value proposition: What is your calculation worth?
- The ratio between flops and bytes of storage will continue to worsen.
- Application resilience and detection/recovery of soft-errors becomes a large problem: check-point/restart will not be practical.
- Power mitigated by users and by the introduction of special-purpose processors, generally vector processors of fixed (but arbitrary!) length.

Exascale Challenges:

- *Starting to build in self-awareness of state and activities*
  - *analytics based monitoring, fingerprinting for provenance, monitoring for rogue operations, software control layers, resiliency to attacks,…*

Security Advisory Board

- *Delivering a useable exascale systems by early 2020s at about 20MW*
- *Co-design of simulation tools in key scientific areas*
- *Systems software, compilers, operating systems*
- *Uncertainty quantification engines*
- *Big data analytics*
- *Requires efficiencies in memory, storage, data movement, code architectures (eg memory (~2pJ/bit); data motion (< nJ/bit, 10x))*
- *Resiliency and reliability*

Technology Advisory Board

- *US ownership of IP - eg supporting technologies and tools, software stacks, network fabrics, interconnects,…*
- *Working with business leaders and going beyond technology diffusion – ROI, business schools, energy efficiency,…*

Business Advisory Board

U.S. DEPARTMENT OF ENERGY

Summary

- The Exascale Initiative is about the innovative capacity and insight that will result from potent combination of predictive simulation & big data analytics in an environment that is more secure from cyber, supply chain and resiliency issues.
- The Exascale Initiative is about pioneering technology in computing that will have an enormous effect on the rest of the the computing marketplace.
- It will require strong partnerships among government, industry, and universities

U.S. DEPARTMENT OF ENERGY