The New Frontier of Computational Science: *Applications to High Leverage Decisions*.

ASCR Advisory Committee

Dr. Dimitri Kusnezov
National Nuclear Security Administration

November 6, 2007

What is changing?

- Technology:
  - Enabling a new way of thinking
  - No longer as constrained in the grand challenges we might be able to pursue
- The Needs: Problems of Societal Import where informed decisions need guidance
  - Energy
  - Security
  - Health
  - Infrasctructure
  - …

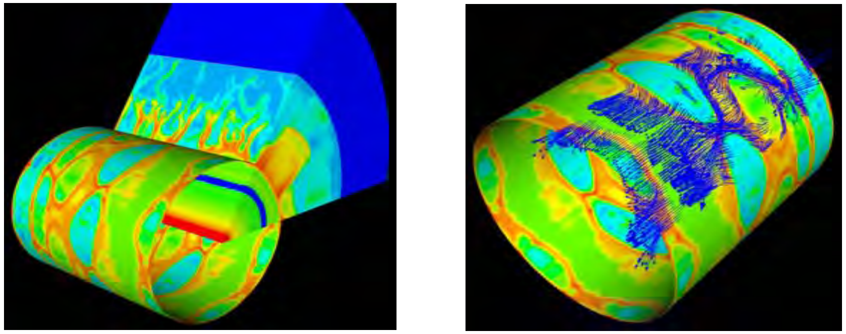

Z-pinch dynamics are 3-dimensional: current can flow around axial gaps in the plasma, which significantly affects the radial current and density distributions

I believe our biggest challenge today is developing the methodology to define our trust in predictions. But today the promise of computational science still exceeds our ability.

(D.Kusnezov, "How Big Can You Think?", Computing in Science & Engineering, Sept/Oct 2007, p.62-67)

Prediction is Simple…

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

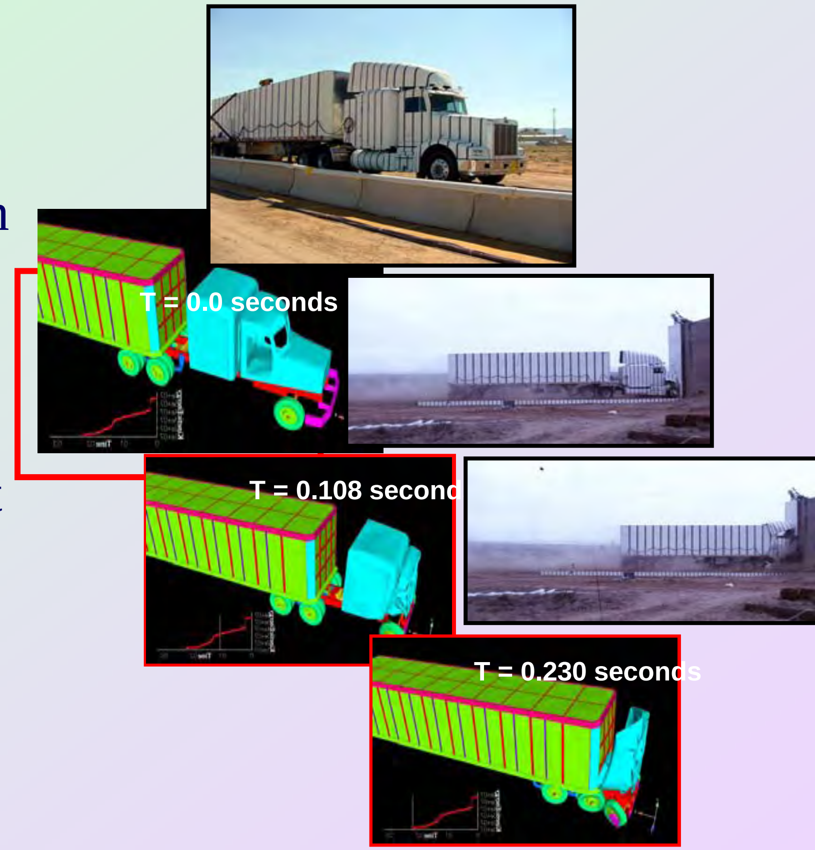

But consequences of poor predictions are not all the same.

Today we are turning to simulation for increasingly serious problems…

(see for example Dan Meiron , Paul Dimotakis et al, JASON Summer Study for the NNSA, JSR-05-335)

Two Classes of Problems are Emerging

- Those for which a mathematical theory of error exists
  - For example, scaling a single-physics application to the maximum scale the computer can handle; well defined problems in controlled approximations.
  - The most comfortable and conventional approach
- Those which we are applying to high-leverage decisions; where the promise of supercomputing is driving us to address issues of national importance, but where approach to prediction requires significant development.

Discovery lies here, where we are limited by our imagination and the risk we are willing to take.

There is a need for ‘Multi-scale Intellectuals’

- The questions we are being enabled to ask lie well outside our comfort zone
- Running what we are familiar with, only faster, bigger and/or longer is not visionary
- Multi-disciplinary teams are often inconsistent with current academic structures
- Integration and coordination of long-term activities and investments is required from hardware, software, code development to science, engineering and applied math.

• Current approach to Computational Science is not suitable
• One must approach it from the point of view of the entire ecosystem: people, resources, needs,…
• Not even Uncertainty Quantification is ready for application to multi-disciplinary problems

Challenges are everywhere

- Data Rich
  - Real-time streaming data
  - Sensor arrays, weather,…
  - Distinguishing causative for correlative signals
  - Solving the inverse problem
- Data Poor
  - Model based
  - Integral data that is not model specific
  - Don’t necessarily want more data either
  - Not possible to instrument experiments

Today's Nuclear Stockpile

High Leverage Decision:
Can you guarantee not having to test again?

Central Problem: *Replacement of underground testing with a rigorous scientific methodology with which to assess and maintain confidence without resorting to testing.*

Time Urgencies: *Supporting national policy with respect to the maintenance of the nuclear stockpile requires that we be able to certify annually to the Secretaries of the Departments of Energy and Defense that the stockpile is safe, reliable and secure.*

National Program: *Planned and coordinated across the three Defense Program Laboratories with partnerships with academic centers and industry.*

We need to *predict* a broad set of properties related to nuclear weapons

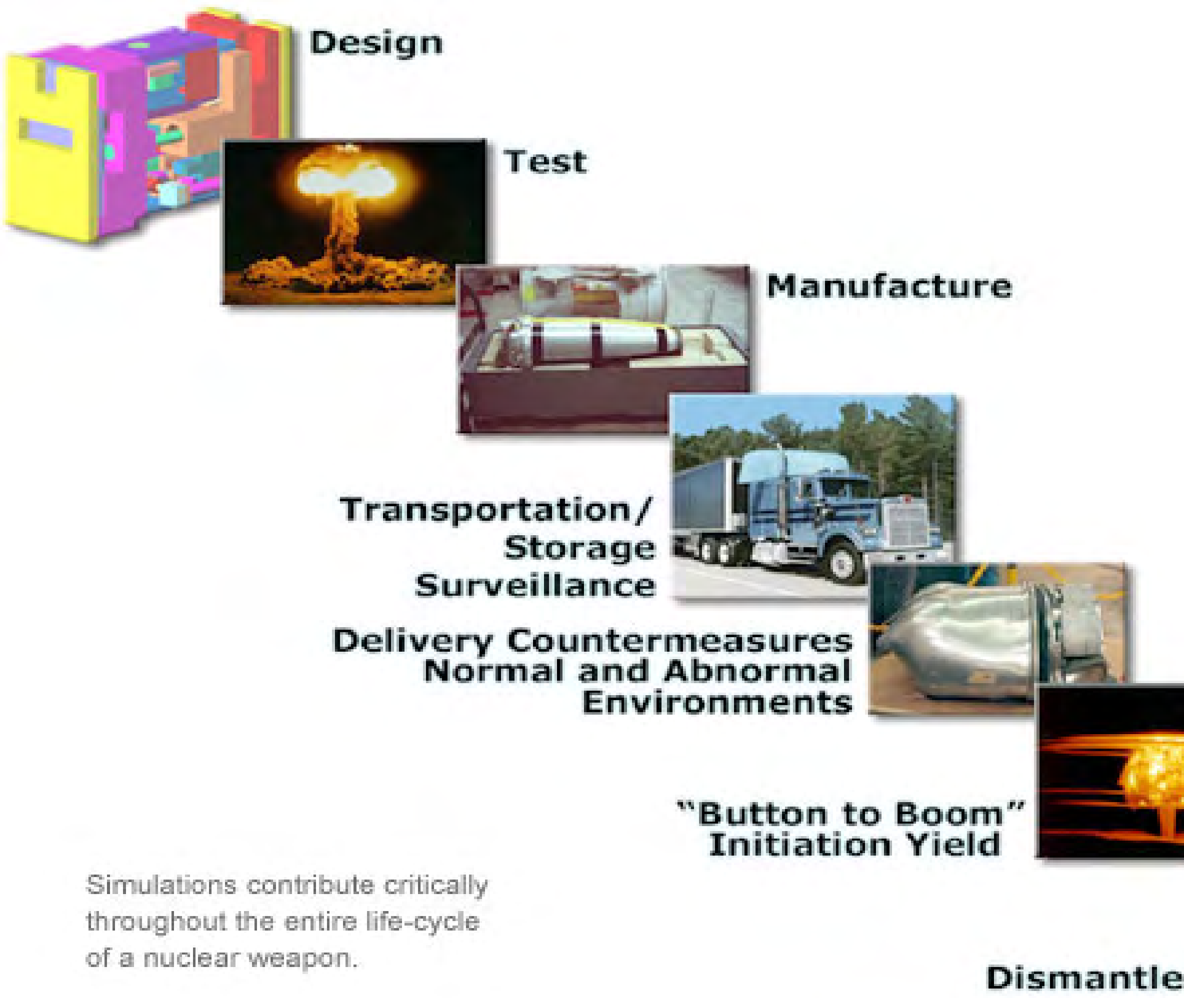

- Forensics
- Attribution
- Prevention
- Aftermath
- ….

Every aspect of the Nuclear Weapons enterprise is modeled before it is done.
A necessary approach requires a balance of phenomenology
and fundamental science.

Computational Science supporting Policy Decisions

Provide the science-based capability to assess and certify the safety, performance and reliability of nuclear weapons and their components without nuclear testing

- Create tools for annual certification & assessment
- Maintain a credible deterrent with a "zero-yield" nuclear test ban
- Support the President's Comprehensive Test Ban policy
- Ensure the effectiveness of science-based stockpile stewardship

Types of Issues:

- *As-built issues*
- *Aging issues*
- *Replacement of materials*
- *Improved safety, security*
- *Emerging threats*
- *Attribution/Identification*
- *Improved technology*

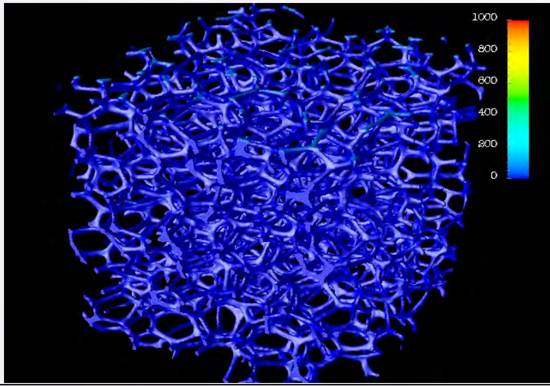

Computational science is at the intersection of leading-edge science and national policy, on a schedule to meet national nuclear security needs

Computational Science Critical to Scientific Discovery

One can begin to address problems that involve competing complex physical processes including:

- Ranges of length scales
- Ranges of time scales
- Extreme conditions and complexity not experimentally accessible
- Competing, complex physical processes and dynamics

  Non-equilibrium systems, phase transitions,…

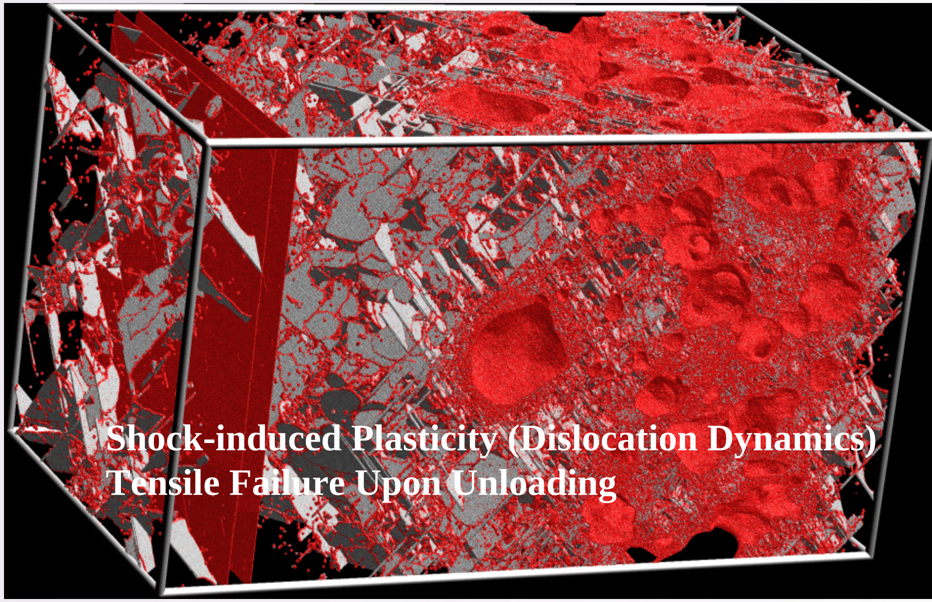

- Informs decisions, allows virtual testing/prototyping. Virtual testing increasingly replacing experimentation *(wind-tunnel evaluations of airfoils, crash tests)*

NNSA's ASC program provides a model of how to develop an appropriate ecosystem to foster intellectual development

Simulation Science: Microscale to Mesoscale

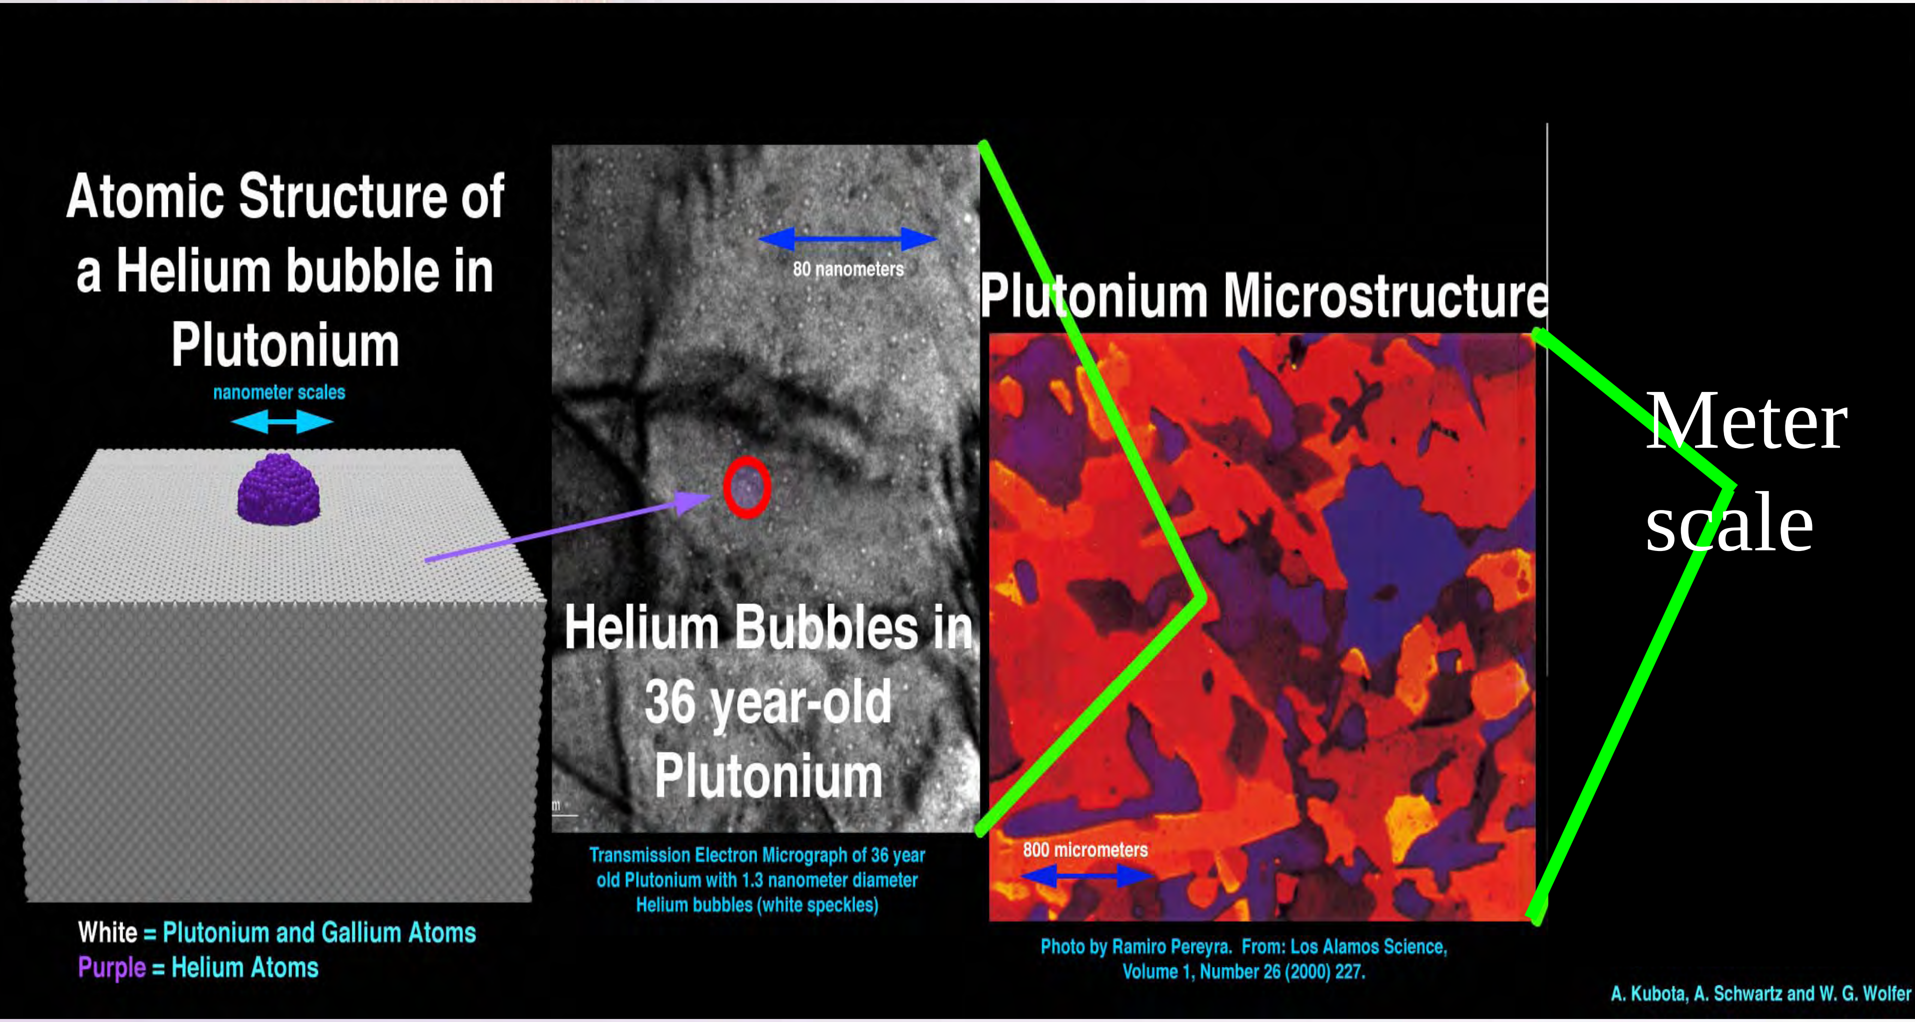

Not just structure, but dynamic effects are critical

New Performance Indicators

Performance Indicator	Endpoint Target
ADOPTION OF ASC MODERN CODES: The cumulative percentage of simulation runs that utilize modern ASC-developed codes on ASC computing platforms, as measured against the total of legacy and ASC codes used for stockpile stewardship activities	By 2013, ASC-developed modern codes are used for all simulations on ASC platforms…
REDUCED RELIANCE ON CALIBRATION: The cumulative percentage reduction in the use of calibration "knobs" to successfully simulate nuclear weapons performance	By 2018, the four major calibration knobs affecting weapons performance simulation have been replaced by science-based, predictive phenomenological models…
ASC IMPACT ON SFI CLOSURE: The cumulative percentage of nuclear weapon Significant Finding Investigations (SFIs) resolved through the use of modern (non-legacy) ASC codes, measured against all codes used for SFI resolution.	By 2013, ASC codes will be the principal tools for resolution of all Significant Finding Investigations (SFIs)…
CODE EFFICIENCY: Cumulative percentage of simulation turnaround time reduced while using modern ASC codes.	By 2013, achieve a 50% reduction in turnaround time, as measured by a series of benchmark calculations, for the most heavily used ASC codes…

Petaflops are no longer the challenge

- NNSA/DOE has several viable architectures for multi-petaflop systems as early as 2008.
- Two joint NNSA/Industry R&D efforts have led to new product lines capable of multi-petaflops:
  - Blue Gene/L [LLNL/IBM] ⇨Blue Gene/P ⇨Blue Gene/Q
  - Red Storm [SNL/Cray] ⇨ Cray XT3…
  - *RoadRunner [LANL/IBM] – testbed for heterogeneous systems?*

- Petaflop systems are now a matter of investment – the challenges lie beyond that as we move towards exaflops

Cost-Effective Platform Strategy used for NNSA Supercomputing needs

- **What is the ASC Platform Strategy?**
  - **Procure <u>Capability Systems</u> for the most demanding integrated weapons simulations**
  - **Procure <u>Capacity Systems</u> for stability, speed and turnaround time necessary for day-to-day needs of stockpile simulations**
  - **Explore and exploit disruptive breakthrough technologies, optimized to run specific applications at extreme scales highly effectively – <u>Advanced Systems</u>**

- **Why do we need Capacity computing?**
  - **Capacity computing off-loads the less-demanding simulations from the expensive Capability supercomputers so that large computers can be used for large simulations**
  - **These cost-effective computers are capable of running the majority of stockpile calculations not requiring the extreme performance characteristics of the Capability and specialized systems**

What is the Tri-lab Linux Capacity Cluster (TLCC) Procurement?

TLCC is a headquarters initiative to reduce costs to the program through:

- 30-50% reduction in cost through “group buy”
- Less frequent acquisitions of large and costly supercomputers
- A single, multi-year contract with vendors for all LANL, LLNL and SNL capacity computing
- A standard, common software environment to improve user productivity
- Support for Capability system runs (i.e. Purple as a national user facility)

Labs Get Matching Computers

U.S. Cuts a Price Deal With California Maker

BY JOHN FLECK
Journal Staff Writer

The National Nuclear Security Administration has cut a deal with a California computer company for new nuclear weapons research supercomputers the agency says will yield substantial savings for the normally pricey machines.

The $26.1 million contract with Appro, of Milpitas, Calif., means Los Alamos, Sandia and Lawrence Livermore national labs will get identical supercomputer clusters, said Dimitri Kusnezov, director of the NNSA's Advanced Simulation and Computing program.

Since the end of nuclear testing in the early 1990s, supercomputers have been a key tool for designing and maintaining nuclear weapons. Last year, the program had a $618 million budget. Next year's proposed budget is $586 million, a 6 percent decrease.

“We struggle all the time with, ‘How do we save money in these programs,’” Kusnezov said.

In the past, each of the three nuclear weapons laboratories has been responsible for its own computer purchases. By pooling the buy, it was possible to save a substantial amount of money, Kusnezov said in a telephone interview.

The new Appro machines will not compete with the high-end super-fast machines at the labs that do the biggest nuclear weapons design jobs. Those machines, including Blue Gene/L at Livermore and Red Storm at Sandia, typically rank among the fastest in the world.

In addition to those top-of-the-line machines, the labs also traditionally purchase large numbers of less capable computer clusters to do large numbers of smaller jobs. That is the niche at which the Appro purchase is aimed.

By having identical systems at all three labs, the NNSA will save on installation and maintenance costs, as well as the costs of modifying software to run on the machines, Kusnezov said.

Kusnezov described the machine as like a set of Lego bricks — easy to expand by simply tacking on more units to the cluster.

Livermore will get the first eight-unit cluster in early 2008, according to NNSA spokesman John Broehm. The next set will be sent to either Los Alamos or Sandia, Broehm said.

Purple National User Facility Contributing to Stockpile Knowledge

Bank	Application (U)	<10%	10 – 30%	30 – 75%	>75%	Lab
ccc2-1	RRW (WR-1)	7.1%	65.3%	27.7%	0.0%	LLNL
ccc2-2	B61	6.6%	65.1%	28.3%	0.0%	LANL
ccc2-3	NIC	9.8%	48.2%	42.0%	0.0%	LLNL
ccc2-4	QASPR	0.8%	2.1%	70.8%	26.3%	SNL
ccc2-5	TBI/NBI	8.6%	30.4%	61.0%	0.0%	LANL
ccc2-6	TBI/NBI					
ccc2-7	DARHT					
ccc2-8	W76/W87/W78					
LANLEXEC						
LLNLEXEC						
SNLEXEC						

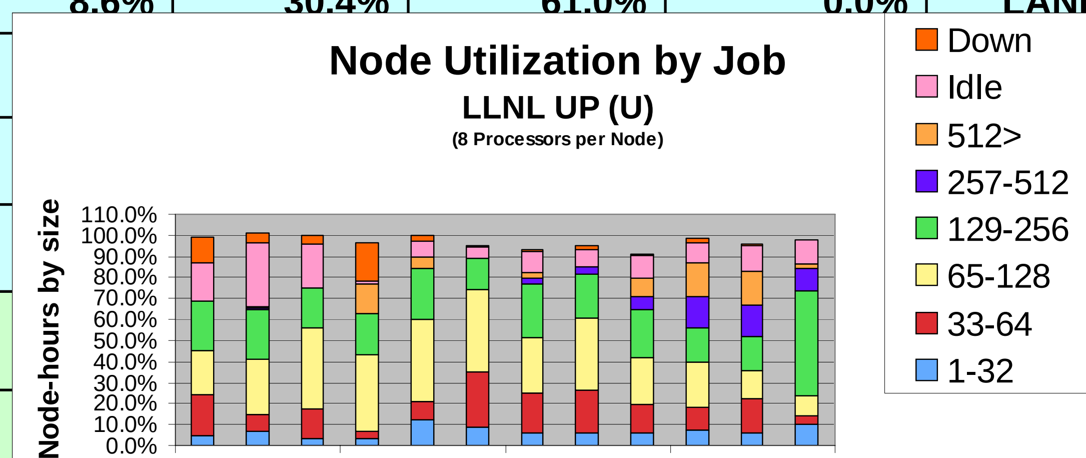

- Capability Computing Campaigns – prefers largest scale simulation needs
- Platform usage – gauges program pressures and capacity needs

Partnerships are Key to our Success

- **ASC Approaches**
  - Collaboration and Interdependence
  - Focused Research
  - Focused Metrics & Investments
  - Programmatic Integration
- **Partnership Profile**
  - National: Government, Industry, University
  - International: UK, France, Russia
- **Partnership Interests**
  - Simulation of Complex Problems
    - Multi-scale, coupled multi-physics applications
    - Extreme environments
    - Management of Large data sets

Collaborations with other Agencies

- **Activities with Office of Science**
  - ANL/LLNL co-funding $34M IBM Blue Gene P/Q R&D
  - SCIDAC-2 co-funding ($4M annually, $20M total)
  - Software tools and environments for peta-scale applications
  - Draft MOU
- **Interagency Activities**
  - DARPA High Productivity Computing System (HPCS)
    - ASC participating in advisory role only; not as a funding agency
  - OSTP's Networking and Information Technology Research & Development Subcommittee
  - NSA Expert Council and procurement reviews
  - DoD High-performance Computing Modernization Office procurement reviews and benchmarking
- **Fellowship Programs**
  - DOE Computational Science Graduate Fellowship
- **Council on Competitiveness**
  - High-Performance Computing Advisory Group

Collaboration with SC: Scientific Discovery through Advanced Computing (SciDAC-2)

Application Area	Proposal Title	University
Materials Science and Chemistry	A Multi-Scale Many-Body Science Application for Strongly Correlated Materials	University of Cincinnati
	Hierarchical Petascale Simulation Framework for Stress Corrosion Cracking	USC
	Quantum Simulations of Materials and Nanostructures	UC Davis
Turbulence	Simulations of Turbulent Flows with Strong Shocks	Stanford
Astrophysics	Computational Astrophysics Consortium: 3 Supernovae, Gamma-Ray Bursts and Nucleosynthesis	UC Santa Cruz
Nuclear Physics	Nuclear structure and low-energy reactions	University of Washington

Major Hurdles in Effective Use of Technology: Uncertainty Quantification Issues

How do you convince others that your predictions are trustworthy?

Three main issues exist:

- **Characterization of uncertainties as inputs to simulations:**
- **Propagation of uncertainties (aggregation/convolution) through codes to quantify uncertainty in some integrated performance parameters.**
- **How calibration works within a structured, disciplined UQ methodology and for what purpose.**

Connection to experiment through a validation program is vital.

"Prediction is very difficult, especially if it's about the future." A. Einstein

New Predictive Science Academic Alliance Program (PSAAP)

Defense Programs

- **Focus on multi-scale, multi-disciplinary, unclassified applications of NNSA interests**
- **Demonstrate validated simulation capability for prediction**
- **Produce new methodologies on:**
  - Verification
  - Validation
  - Uncertainty quantification
  - Tight integration of experiment and simulation

- **Proposal Peer Review: Aug 21-22, 2007**
- **Site Visits: September-October 2007**
- **Projected Selection Announcement: by March 2008**

Where do we go from here?

Leadership computational science will continue to push the frontier

*Broad applicability of computational science methods outside the weapons program results in a highly **transferable** knowledge base to changing needs*

(nuclear energy, climate, cyber-security,…)

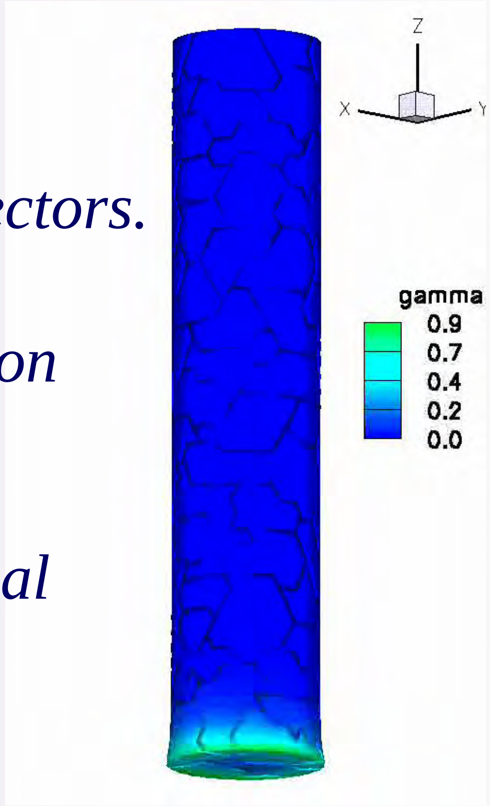

This will pay-off in the commercial and academic sectors.

Virtual testing increasingly replacing experimentation

(zero wind-tunnel evaluations of airfoil; new oil sources,…)

Multidisciplinary challenges – collaboration essential

A new era of great promise is emerging