Requirements for ASCI

Participants:

Henry Abarbanel
Michael Brenner
John M. Cornwall
Bill Dally
Alvin Despain
Paul E. Dimotakis
Sid Drell
Douglas M. Eardley
Bob Grober
Raymond Jeanloz
Jonathan Katz
Steven Koonin
Darrell Long
Dan Meiron (Consultant)
Rip Perkins
Roy Schwitters (Study Leader)
Christopher Stubbs
Peter Weinberger

What is ASCI?
Advanced Simulation and Computing

- Mission: Provide the means to assess and certify the safety, performance and reliability of nuclear weapons and their components.
- Goal: Deliver predictive computer codes based on multi-scale modeling, verification and validation of codes, small-scale experimental data, nuclear test data, engineering analysis and expert judgment.
- Supports people, hardware and contracts to the greater scientific and computing communities
- Started in 1996; approximately 1/8 of SSP budget

What does ASCI cost?

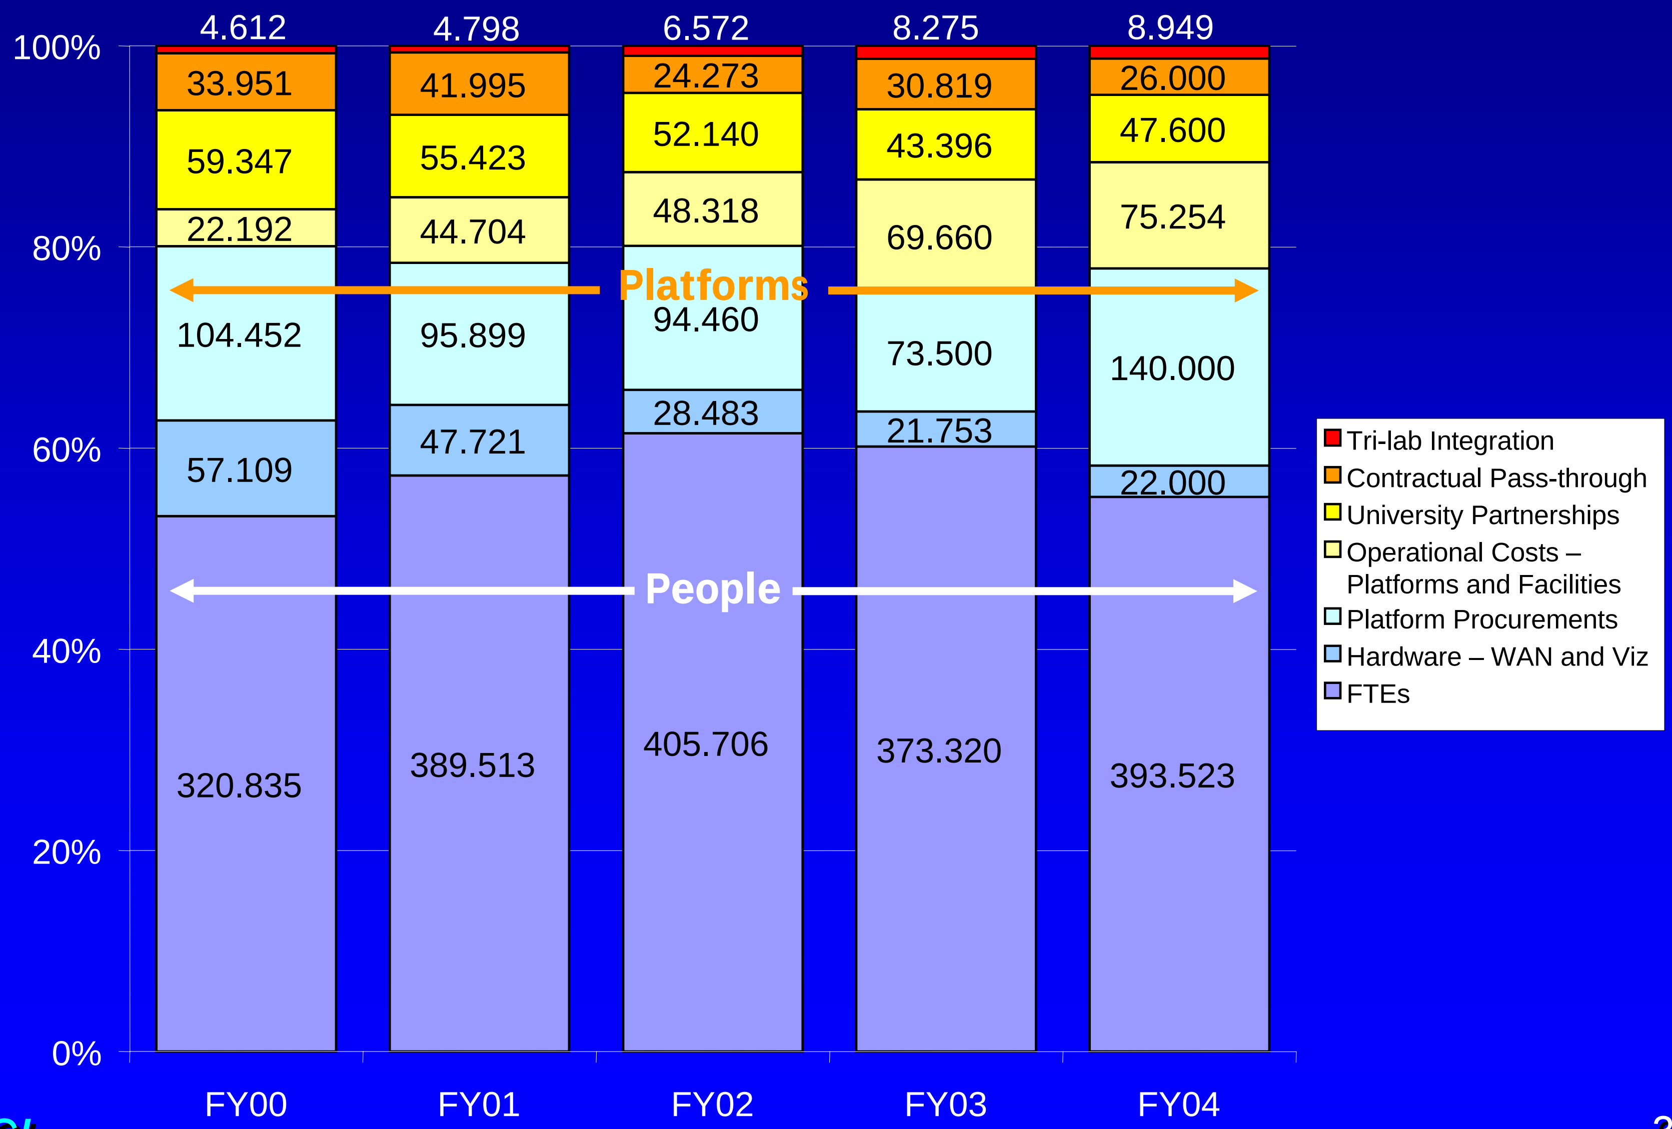

Charge to JASON

- Identify the distinct requirements of the stockpile stewardship program and its relation the ASCI computer acquisition strategy
    - Confidence in simulation
    - Balance in demands for capacity
    - Bases for sustainable and credible program
- Evaluate the increased risk to stockpile stewardship and to the scientific program that it supports, as a result of delaying acquisitions to advance capability.

Context

From the Senate report on FY03 Appropriations:

"   While the Committee recognizes the central importance of the ASCI program to the success of stockpile stewardship, *the Committee remains unconvinced that the NNSA's platform acquisition strategy is driven by identified requirements*, rather than a well intentioned, but insufficiently justified, desire to aggressively acquire larger and faster computing assets on an accelerated time-scale."

"   The NNSA is directed to commission two related studies, the first to be performed in collaboration with the Department's Office of Science and *the second focused solely on issues relevant to the stockpile stewardship program*."

From the current Senate markup of the FY04 request:

"The Committee recommendation includes $725,626,000, an amount that is $25,000,000 below the budget request.   The recommended reduction is without prejudice and the Committee expects to revisit the appropriate level of funding at conference with the benefit of the National Academies' and JASONs' reports."

Preview of JASON's conclusions

- ASCI has become *essential* to Stockpile Stewardship
  - Contributes to achieving technical milestones
  - Enables new capabilities with better science
  - Training cadre of experts; good young people entering program
- Distinct technical requirements place valid computing demands on ASCI that exceed present and planned computing capacity and capability

Outline

- Description of summer study
- Performance metrics
- Stockpile stewardship requirements and achievements
- Platform acquisition scenarios
- Role of research
- Conclusions & Recommendations

Summer Study

- Informal lab visits
  - One-day visits to LANL, SNL, LLNL during Spring
  - Sat down with designers/code experts
    - *How* they do their jobs
    - *What* they need
- 5 ½ days of formal briefings, discussions with lab experts on requirements, performance and science
- Briefings/comments by outside computer experts

Many thanks to all the briefers and to:
- Labs & staff for hosting us and for responding to queries.
- Dimitri Kusnezov, Hans Ruppel and lab ASCI "execs" for organizing and carrying out a *unified* set of briefings.

Capability and *Capacity*

- Terms of art in ASCI world
  - *Capability*: the maximum processing power that can be applied to a single job
  - *Capacity*: the total processing power available to run ASCI jobs
- No good metric for either (as we shall see)
  - We will use peak single-processor floating-point operations/s for both , usually in TeraFlops (TF)
- *Capability* $\Rightarrow$ *Capacity*
  - Capacity added
  - Capability machines can be configured to run multiple smaller jobs

ASCI “most capable” platforms

- Today
  - ASCI “White” at LLNL (12.3 TF)
  - ASCI “Q” at LANL (20 TF – reduced from 30 TF)
- Next procurements
  - “Red Storm” at SNL (40 TF)
  - “Purple C” at LLNL (100 TF)
    - Procurement includes “Blue Gene/L” (180/360 TF, potentially)
    - BG/L viewed as new-technology test bed

Where ASCI platforms fit into the world of high-performance computing

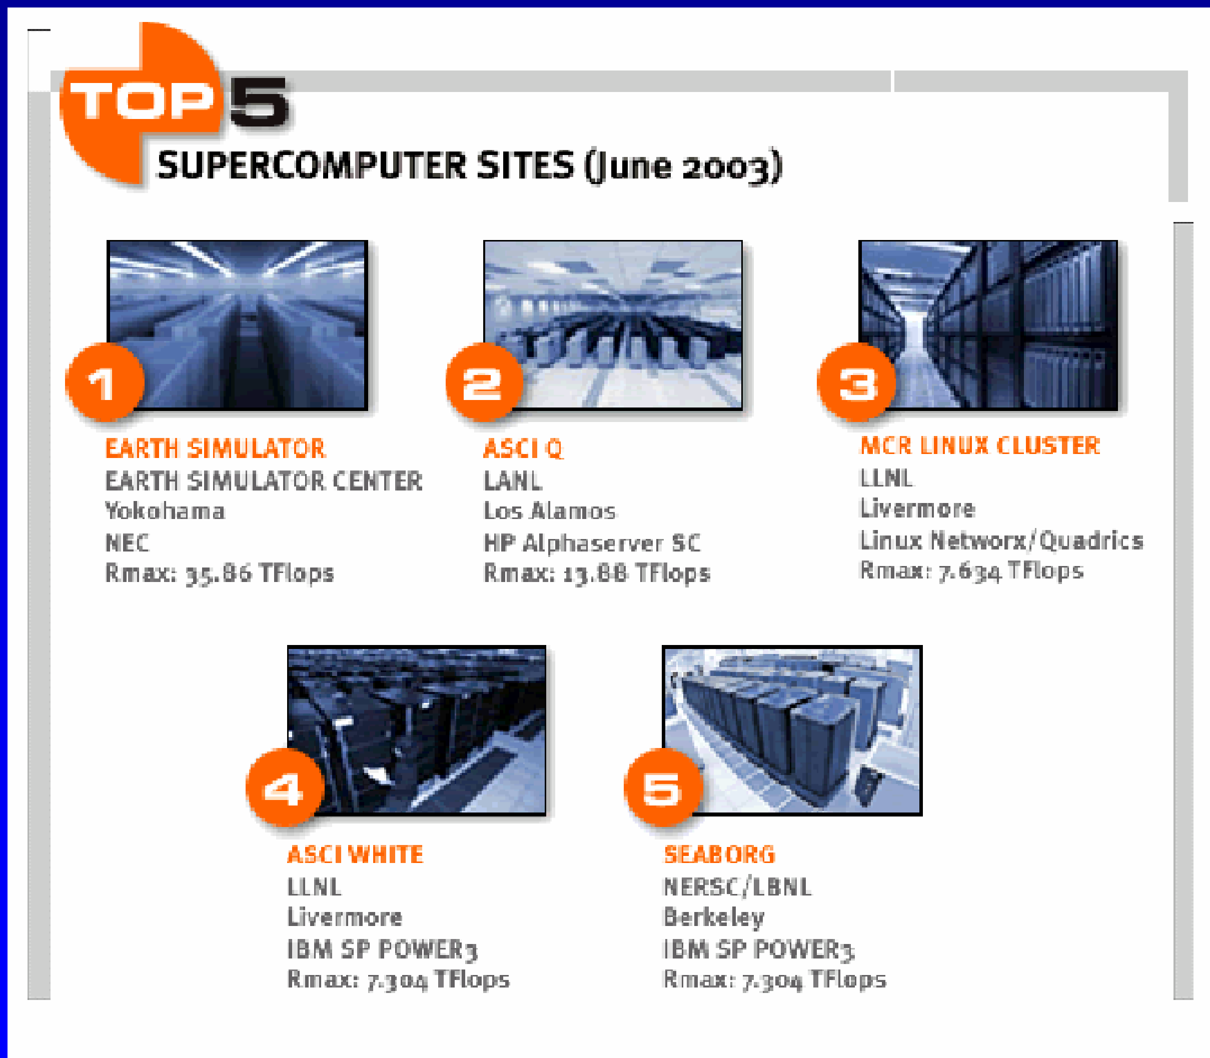

Performance metrics

- *Peak* TeraFlops (1 TF = 10^{12} floating-point operations/s) not truly representative of capability
- *Delivered* TFs depend on many things
  - Character of computational problem
  - Platform architecture
  - Compilers
  - Operating system, …
- Time-to-solution is the important metric to users
- Benchmarks should represent workload

ASCI platform performance

- Our considerations based on study by LANL performance analysis group
- Single processor performance
  - 0.5-15% of peak depending on particular ASCI kernel
  - Also observed in similar applications (e.g. University Alliances)
  - Efficiency is typical of applications requiring large numbers of memory references per operation
- Scalability
  - Unanticipated obstacles encountered at > 3K processors
  - All obstacles to date have been overcome or the required fix is understood:
    - Operating system issues – will require vendor response
    - Algorithm issues – being addressed by ASCI experts

ASCI performance analysis

- Relies on work of Hoisie, Kerbyson, Pakin, Petrini, Wasserman

$$T_{run} = T_{computation} + T_{communication}\ (-\ T_{overlap})$$

$$T_{run} = f\ (T_{1\text{-}CPU}\ ,\ \textit{Scalability function})$$

- Single processor performance obtained from hardware counters
- Multiprocessor performance from modeling
- Focused on ASCI workload
  - SAGE - hydro, AMR
  - ALE
  - PARTISN/SWEEP – rad transport
  - Monte Carlo

Performance of SAGE and PARTISN

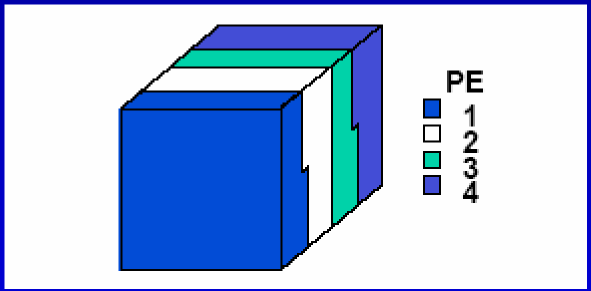

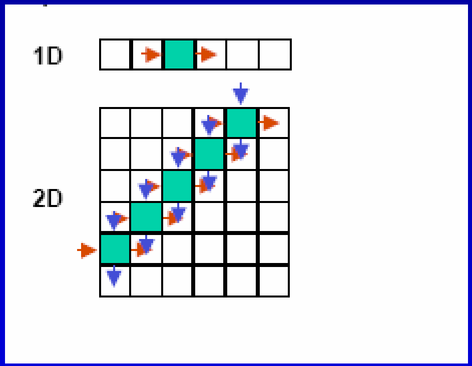

- Performance models can accurately predict how these codes will run on any architecture
- Typical characteristics
  - 3 memory references per flop
  - Leads to 13% of peak for PARTISN and 4% for SAGE (ASCI Blue Mountain)

But what about the dreaded Earth Simulator?

- Depends on single processor performance
- But for ASCI workload could be anywhere from equivalent to factor of 3 of ASCI's most capable current system (Q)
- Important thing is that the differences can be modeled

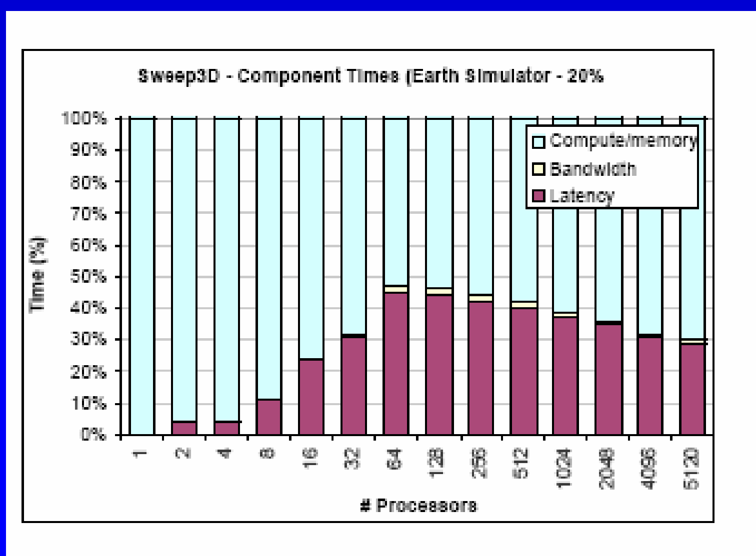

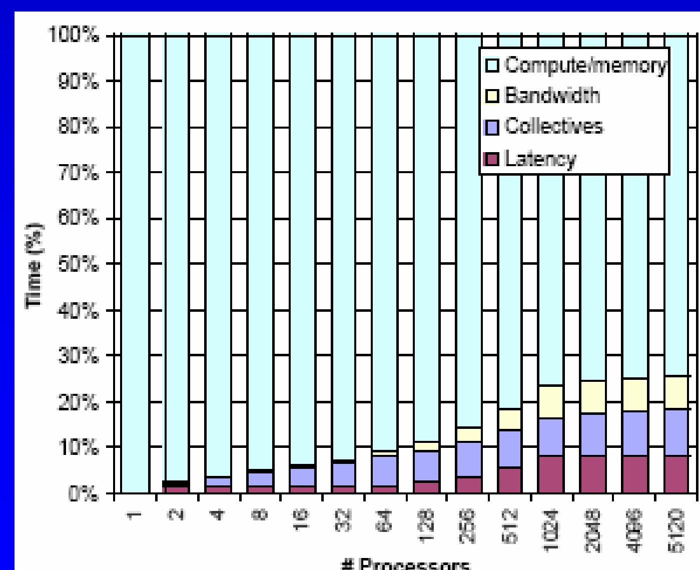

ASCI performance conclusions

- ASCI performance is good, appropriate to its mix of jobs
- ASCI has developed good analysis tools for understanding performance of relevant algorithms
- These tools can be (and should be) used to assess capability of future procurements
- Studies highlight importance of continuing to improve single-processor efficiency and balanced network bandwidth
  - Essential to future time-to-solution
  - JASON report suggests possible areas to be investigated
- Benchmarks need to be representative of ASCI workload
- Scaling to future capability requires development

But we should not declare victory…

- Commodity improvements may not get us to where we need to be
  - Dally – slowdown of Moore's law
  - Continued poor memory to flop ratio
  - Petaflop performance and beyond *will* be required
  - Scaling conventional solutions may lead to serious reliability problems
    - To get to a PFlop we must scale today's machines by factor of 100
    - Conventional microprocessors may only increase by factor of 4 in 2010
    - **Implies something like 300K nodes for a Petaflop**
- Possible solutions
  - Hardware
    - Vectors
    - Streaming
    - Electrical or optical high performance interconnection networks
    - Processor in memory
    - New chip architecture
  - Software
    - Reliable parallel OS and compilers
    - Automatic code optimization – ATLAS for ASCI
- **CS research must be supported in these areas**

Stockpile stewardship requirements and achievements

- Directed Stockpile Work (DSW)
  - Supports certification
  - Life-extension Programs (LEP)
  - Specific to weapon-type
- Campaigns
  - NW Science/Engineering
  - Cuts across weapon-types
- Significant Finding Investigations (SFIs)

- Baselining: adjusting models to UGT archives
- Safety: engineering studies of accident scenarios
- Stockpile-to-Target Sequence (STS) requirements: models of environments encountered during delivery of weapons
- Support to production
- Surety: use-control and other classified aspects

Examples of work enabled by ASCI

Evaluating Engineering Margins Requires Very High Fidelity

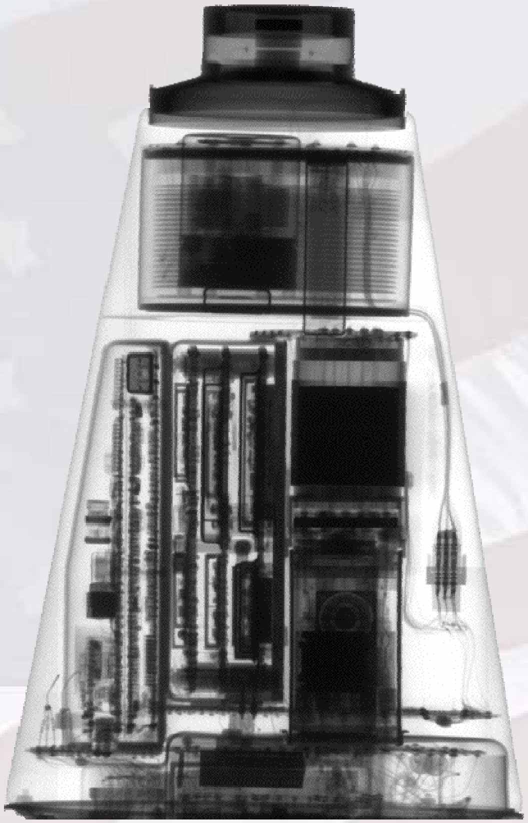

W76 AF&F
x-ray

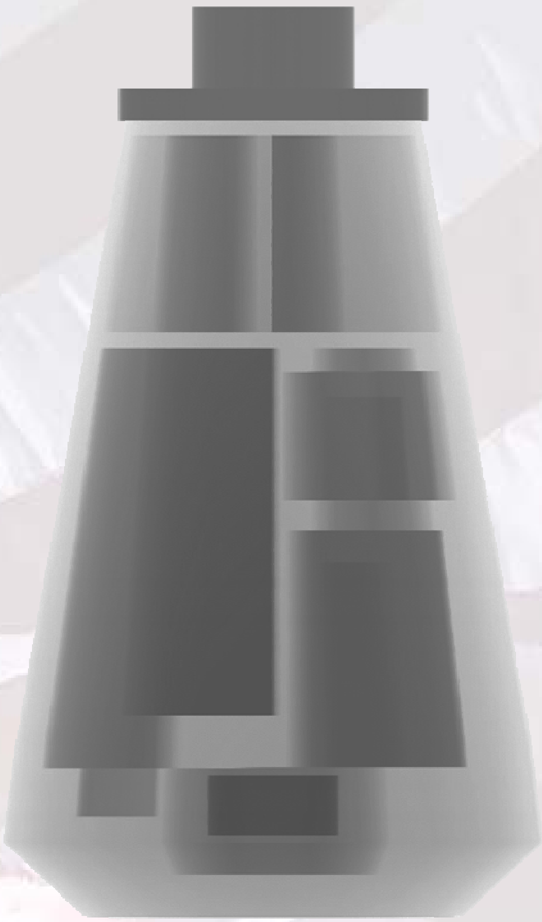

W76 AF&F
pre-1998 mesh
synthetic x-ray

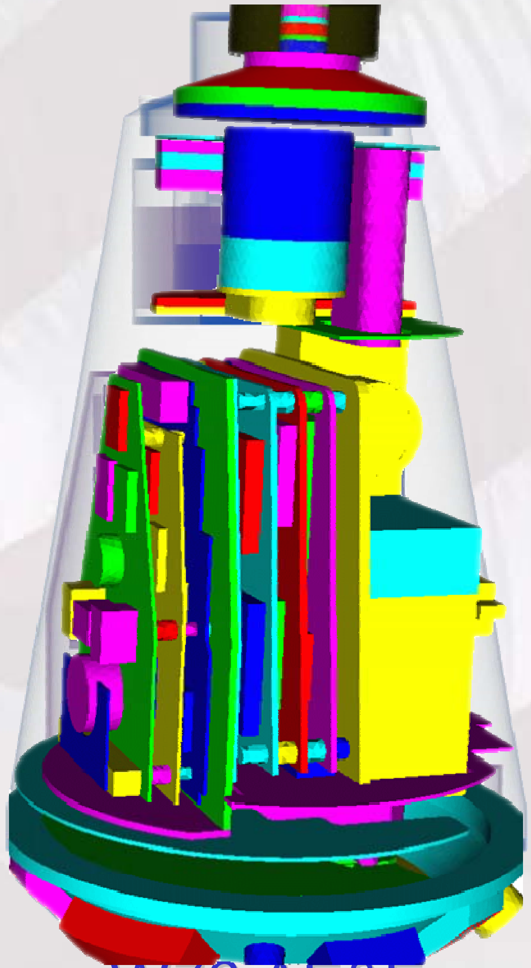

W76 AF&F
W76 AF&F
2001 mesh
2001 mesh
synthetic x-ray

NNSA National Nuclear Security Administration

Sandia National Laboratories

The JASON "S-matrix"

- JASON requested assistance from the labs to estimate computational complexity required to simulate the *science* representative of the distinct stages in a nuclear weapon
- We assessed the physics uncertainties of the different stages
- Labs were asked to describe both present-day and future requirements
- Used in our assessment of computational requirements

Example of present demand: W80 LEP Primary computing requirements

machine	fraction	number of days
2D: 10^7 White hours		
White	100%	51
	25%	203
Purple C	25%	26
3D: $3x10^7$ White hours		
White	100%	153
	50%	305
Purple C	50%	40
Surety: $3x10^7$ White hours		
White	100%	153
Purple C	50%	40

The current W80 computing needs can utilize the whole White machine for an entire year

Conclusions on computational load that follows from SSP requirements

- The S-matrix and lab responses helped sharpen our understanding of computational requirements
- Any reasonable "roll-up" of future demand is ≥ 2x projected capacity
- We concur with the labs' assessment that future capability requirements *exceed* 1 PF
- But, the path to 1 PF machines is not obvious
  - Scaling from experience problematic
    - Efficiency
    - Reliability
  - How to proceed? (NAS Committee, a national issue)
- There are hints that better science and phenomenology *may* ultimately point to a *sufficient* level of capability (beyond 1 PF)

JASON's assessment of alternative acquisition scenarios

- JASON was charged to assess risks of delaying procurement of new *capability* machines
- We do so mindful of substantial oversubscription in *capacity*
- Scenarios considered:
  - Current ASCI acquisition plan
  - Delay acquisition of new capability (Purple C and Red Storm) starting in FY04
  - "Requirement-driven" acquisition of capability and capacity

Assumptions entering risk assessment of procurement delay

- Assumed $34M cut (notional value)
  - Removed $25M from Purple procurement
  - Removed $8M from Red Storm procurement
- Assumed resulting delay in near-term platform delivery
  - Red Storm delayed by 1 year
  - Purple delayed by 1 year
- Assumed return of $34M but evened out large budget excursions in future years
  - LANL 200 TF delayed 1 year
  - SNL 150 TF possibly delayed 2 years

Purple Reduced 25M

Red Storm reduced 8M

8M returned to Red Storm

25M returned to Purple

Effect of FY04 procurement delay

Platform Procurement	FY1996	FY1997	FY1998	FY1999	FY2000	FY2001	FY2002	FY2003	FY2004	FY2005	FY2006	FY2007	FY2008	FY2009	FY2010	FY2011	FY2012
SNL Red	18	24	6	8													
LLNL Blue Pacific	13	27	49	8													
LANL Blue Mountain		43	10	34	35												
LLNL White				39.7	62.9	16.2	0	5	6	3							
LANL Q					3.5	71	90.5	0	8	8	8	4					
SNL Red Storm						7	8	16	18	22	14	2.5	2.5				
LLNL Purple								30	72	90	48	35	10				
LANL System											30	87	105	38	0	0	0
SNL system													13	27	59	0	0
total platform buget	**31**	**94**	**65**	**89.7**	**101.4**	**94.2**	**98.5**	**51**	**104**	**123**	**100**	**128.5**	**130.5**	**65**	**59**	**0**	**0**

Platform Tflops	FY1996	FY1997	FY1998	FY1999	FY2000	FY2001	FY2002	FY2003	FY2004	FY2005	FY2006	FY2007	FY2008	FY2009	FY2010	FY2011	FY2012
SNL Red	0	0	0	3.15	3.15	3.15	3.15	3.15									
LLNL Blue Pacific	0	0	0	3.89	3.89	3.89	3.89	3.89									
LANL Blue Mountain		0	0	3.07	3.07	3.07	3.07	3.07									
LLNL White				0	12.3	12.3	12.3	12.3	12.3	0							
LANL Q					0	0	0	20	20	20	20	20	0				
SNL Red Storm						0	0	0	0	40	40	40	40	40			
LLNL Purple								0	26	26	100	100	100	100	100		
LANL System										0	0	0	200	200	200	200	200
SNL system											0	0	0	0	150	150	150
total platform Tflops	0	0	0	10.11	22.41	22.41	22.41	42.41	58.3	86	160	160	340	340	450	350	350
Program Requirements									106	343	565	631	923	1010	1633	2131	
Fraction oversubscribed									**1.8**	**4.0**	**3.5**	**3.9**	**2.7**	**3.0**	**3.6**	**6.1**	

SNL 150T purchase delayed 2 years

LANL 200T purchase delayed 1 year

(Return)

Assessment of risk

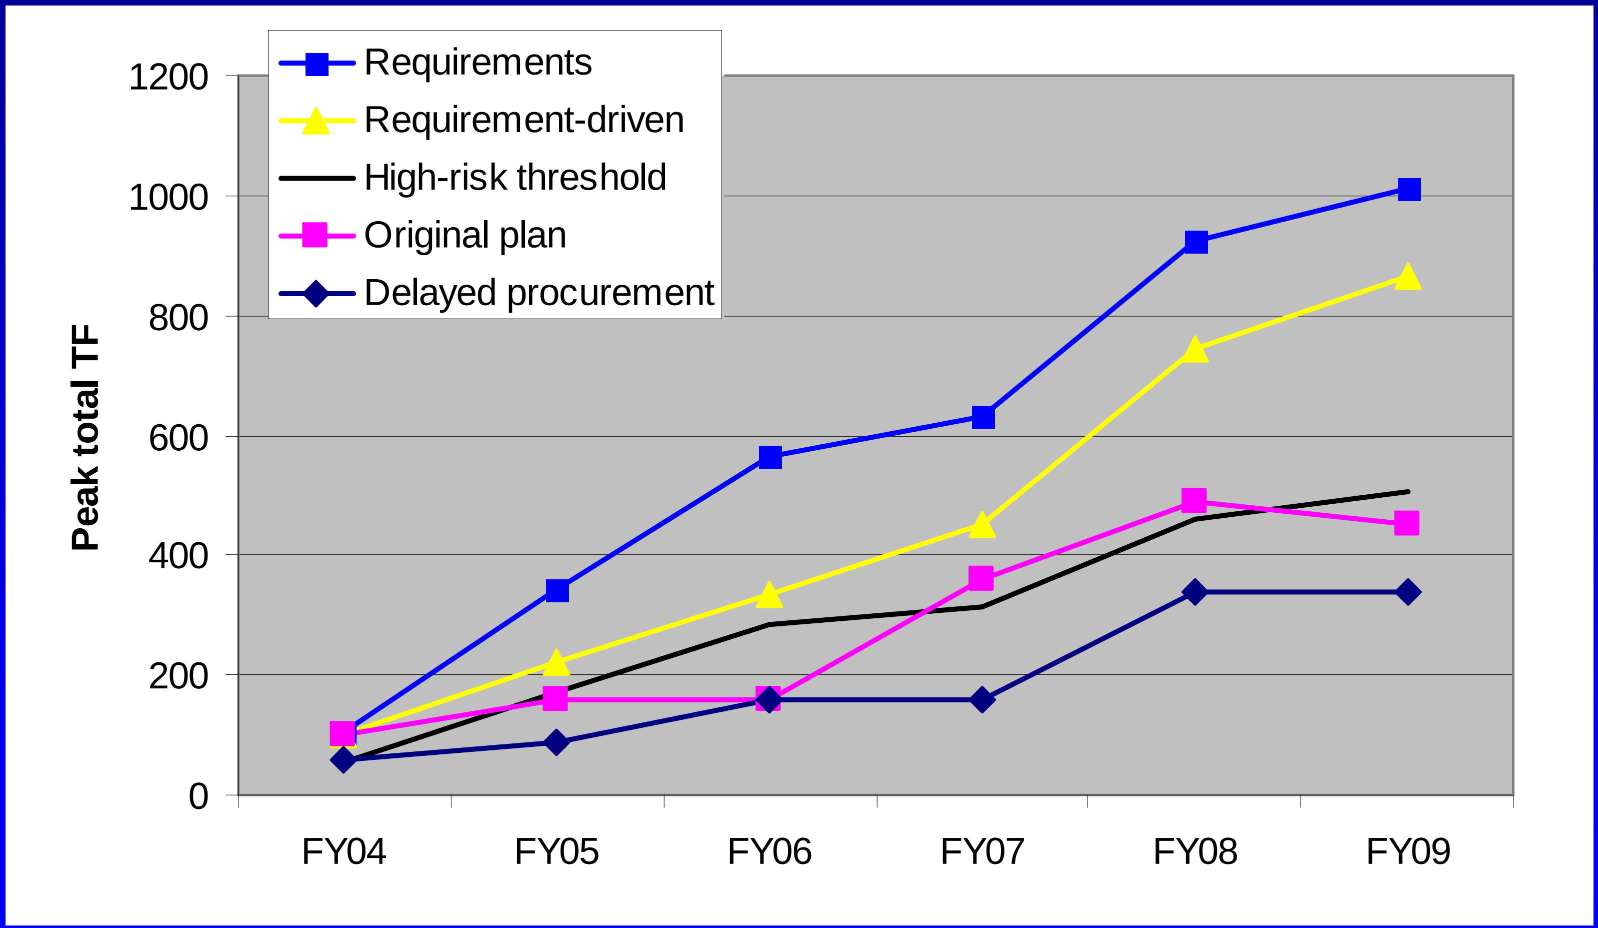

Alternative Scenario Assumptions

- Assumes Tri-lab acquisition and management of *capability*
- Assumes Tri-lab procurement of *capacity*
  - 500-2000 node clusters
    - Possibly Linux based
    - $1M per TFlop of capacity
- Assumes Purple and Red Storm procurements proceed
- Investment in capability exploration architecture to lead to 1PFlop capability in 2010-2011

Commodity capacity

Enhanced Capacity and Capability Scenario

	FY1996	FY1997	FY1998	FY1999	FY2000	FY2001	FY2002	FY2003	FY2004	FY2005	FY2006	FY2007	FY2008	FY2009	FY2010	FY2011	FY2012
SNL Red	18	24	6	8													
LLNL Blue Pacific	13	27	49	8													
LANL Blue Mountain		43	10	34	35												
LLNL White				39.7	62.9	16.2	0	5	6	3							
LANL Q					3.5	71	90.5	0	8	8	8	4					
SNL Red Storm						7	8	16	26	22	6	2.5	2.5				
LLNL Purple								30	97	70	68	10	10				
Tri lab capacity									0	30	40	30	20	20	20	20	20
Capability exploration phase 1										10	20						
Tri lab capability												30	60	80	80	30	10
Capability exploration phase 2												30	30	30	30	30	30
total platform buget	31	94	65	89.7	101.4	94.2	98.5	51	137	143	142	106.5	122.5	130	130	80	60

Platform Tflops	FY1996	FY1997	FY1998	FY1999	FY2000	FY2001	FY2002	FY2003	FY2004	FY2005	FY2006	FY2007	FY2008	FY2009	FY2010	FY2011	FY2012
SNL Red	0.0	0.0	0.0	3.2	3.2	3.2	3.2	3.2									
LLNL Blue Pacific	0.0	0.0	0.0	3.9	3.9	3.9	3.9	3.9									
LANL Blue Mountain		0.0	0.0	3.1	3.1	3.1	3.1	3.1									
LLNL White				0.0	12.3	12.3	12.3	12.3	12.3	0.0							
LANL Q					0.0	0.0	0.0	20.0	20.0	20.0	20.0	20.0	0.0				
SNL Red Storm						0.0	0.0	0.0	40.0	40.0	40.0	40.0	40.0				
LLNL Purple								0.0	26.0	100.0	100.0	100.0	100.0	100.0			
Tri lab capacity									0.0	60.0	173.1	293.1	406.3	566.3	732.5	939.4	0.0
Tri lab capability											0.0	0.0	200.0	200.0	1000.0	1000.0	1000.0
total platform Tflops	0.0	0.0	0.0	10.1	22.4	22.4	22.4	42.4	98.3	220.0	333.1	453.1	746.3	866.3	1732.5	1939.4	1000.0
Program Requirements									106.0	343.0	565.0	631.0	923.0	1010.0	1633.0	2131.0	
Fraction oversubscribed									1.1	1.6	1.7	1.4	1.2	1.2	0.9	1.1	

Purple procurement proceeds on schedule

Red Storm procurement proceeds on schedule

Capability R&D delivers 1PFlop in 2010-2011

Conclusions on alternative acquisition scenarios

- Deferral of Purple and Red Storm increases risk substantially because of pressure on *capacity* and *capability*
- Alternative, requirement-driven scenario could lead to a more balanced program
  - Use of commodity clusters to increase capacity
  - Capability exploration program to enable 1 PF in 2010
  - Management of computing resources across the complex indicated

A cautionary tale:
The Livingston curve

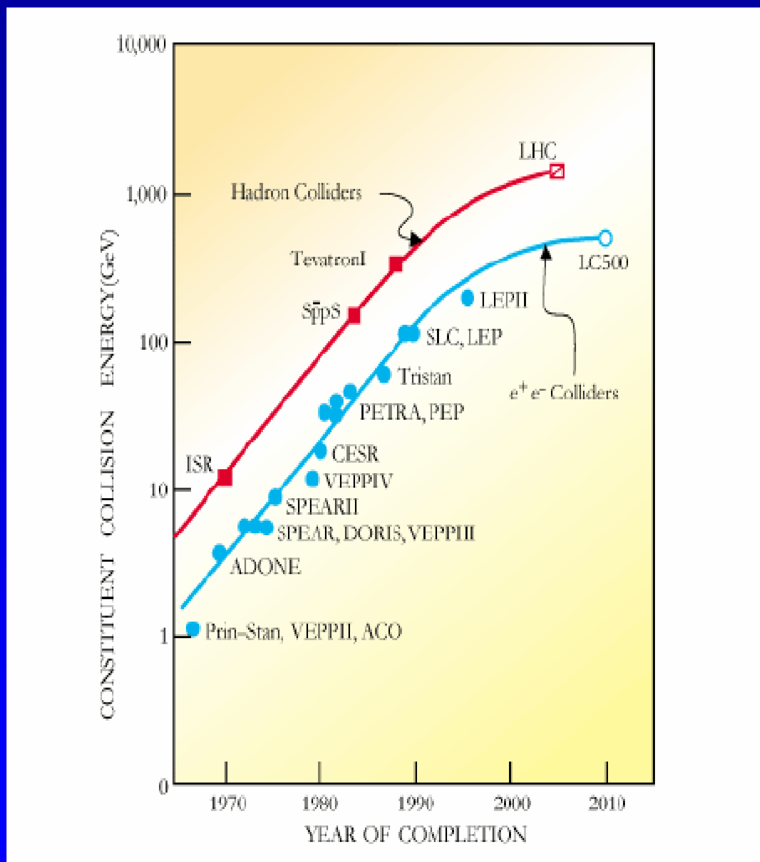

- Equivalent of Moore's law for accelerators
- Knee in curve is not due to physical limits (yet)
- Economics is the driver
- Accelerator community has responded by creating major shared facilities
- Comparison to HPC operation is strained but perhaps worth considering

Ref: M. Tigner Phys. Today Jan. 2001

ASCI is a tool for managing risk

- Matches knowledge, including uncertainty, of weapons systems to customer requirements
  - Naturally entails a great many “what if” calculations to span uncertainties
  - Growth in demand is inevitable
    - Learning more all the time about nuclear weapons science and how to exploit ASCI capabilities
    - SFIs, ageing, new concepts, … increase requirements
- Consequences of not demonstrating confidence in meeting customer requirements can be large
  - Failure to certify
  - Decisions to modify a weapon system or process can cost 100’s of $M
- Risks to ASCI’s availability to inform decisions must be viewed in context with the potential cost of overly conservative decisions

Recommendations to mitigate risk in *present* acquisition plan

- Platform Acquisition:
  - Plan *now* to acquire additional *capacity* platforms
  - Lay groundwork for future *capability:* 1PF by 2010
- SSP Requirements:
  - Set priorities and assign ASCI resources accordingly.
  - Review STS requirements in light of current and anticipated US security needs
- ASCI Operations:
  - Be flexible with access to ASCI "Most-Capable" systems
  - Invest in effort to improve computational efficiency, including allocation of dedicated machine time
- Encourage the advance of NW science at every opportunity

Enhancing Scientific Credibility

- Neither feasible nor necessary to have "full-up" — quarks to mushroom clouds — simulations as long as "sub-grid" models or "phenomenology" are *understood*
  - Physical basis
  - Range of validity
- Notable examples from ASCI
  - Energy balance (O. Hurricane)
  - Test problems relevant for verification (B. Moran)
- Some JASON thoughts
  - Turbulent mixing: possibility of better mix phenomenology?
  - Search for scaling laws to compare with experiments

ASCI is an important tool in resolving important open research issues in weapons science

- EOS of weapons materials
- Constitutive properties of weapons materials
- Aging
- Radiative cross sections
- Nuclear reactions
- Detonation
- Dynamic response of materials
- Interface dynamics
- Radiation transport
- Hydrodynamics of multiphase materials
- Instabilities, turbulence and mixing
- Fast charged particles in plasma
- Interaction of radiation with matter

Research in these areas leads to more refined ASCI requirements

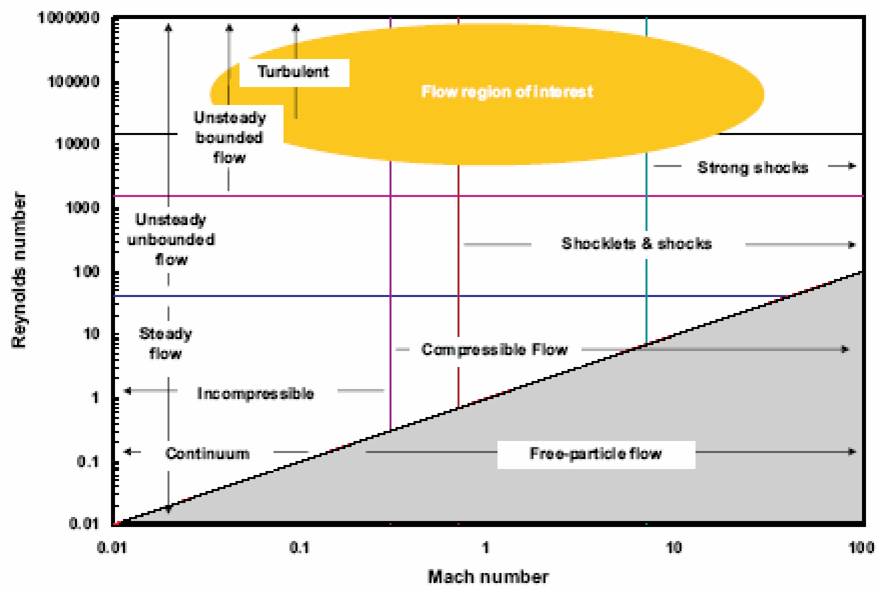

Figure 4-1: Compressible turbulent flow regimes

Understanding the relevant number of scales can provide guidance for where to simulate and where to model

Great Virtue in “Toy” Models

- Simplified, usually analytic model of some physical process
  - Capture the essential symmetries, dynamics
  - Tractable
- Compare analytic results with computations
  - Verification of codes
  - Study mesh/time-step convergence
- Provide insight into relevant scaling laws
  - Quantitative comparison with experiments
  - Metrics for assessing margins

ASCI should be the vehicle to enhance NW science

- Validation of ASCI models by *quantitative* comparisons with experiments
  - Metrics for radiography, subcrits, NIF
  - Scaling laws from models verified by ASCI
- Community “bulletin board” for resolving outstanding issues
  - Understanding phenomenological “knobs”

Summary Conclusions

- ASCI has become *essential* to Stockpile Stewardship
  - Contributes to achieving technical milestones
  - Enables new capabilities with better science
  - Training cadre of experts; good young people entering program
- Distinct *technical* requirements drive acquisition needs
- Present acquisition plan has areas of *substantial* risk
  - Capacity oversubscribed by ~2x
  - Lack of a credible road map to acquiring next-generation of capability which needs ~1 PF
- Delaying FY04 procurements judged to have *high* risk