LA-13474-C
Conference

Proceedings from the Conference on

High Speed Computing

LANL • LLNL

The Art of High Speed Computing

April 20–23, 1998

Algorithms
Architecture
Language

Salishan Lodge
Gleneden Beach, Oregon

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

Photocomposition by Wendy Burditt, Group CIC-1

Special thanks to Orlinie Velasquez and Verna VanAken for coordinating this effort.

LA-13474-C
Conference

UC-705
Issued: August 1998

Proceedings from the Conference on High Speed Computing

The Art of High Speed Computing

April 20–23, 1998

Compiled by
Kathleen P. Hirons
Manuel Vigil
Ralph Carlson

Table of Contents

Conference Program vii

Abstract ix

Keynote Address 1
Steve Wallach, CenterPoint Venture Partners

Stockpile Stewardship and Management Program 25
Larry Ferderber, LLNL

Predictability, and the Challenge of Certifying a Stockpile Without Nuclear Testing 41
Ray Juzaitis, LANL

100 TeraFLOPs and Beyond, an Industry View into the Future 59
Panel: Moderator—John Morrison, LANL; Mark Seager, LLNL; Panelists—Tilak Agerwala, IBM; Greg Astfal, Hewlett Packard; Erik Hagersten, Sun Microsystems; Richard Kaufmann, Digital Equipment Corp.; Steve Oberlin, SGI/Cray.

ASCI Aliances Program.. 105
Ann Hayes, LANL

The Next Fifty Years of Computer Architecture 123
Burton Smith, Tera Computer Company

Full Wave Modeling of Signal and Power Interconnects for High Speed Digital Circuits 133
Gary Haussman, University of Colorado

Simulating the Physical-Biological Factors Affecting Abundance of Calanus finmarchicus in the Gulf of Maine 141
Wendy Gentleman, Dartmouth College

The Next Generation Internet 143
Bob Aiken, DOE

Petaflops Computing: Opportunites and Challenges 151
David Bailey, LBNL

President's Information Technology Advisory Committee (PITAC): A Mid-Term Report 165
David M. Cooper, LLNL

Processing-in-Memory: Past and Present 177
Ken Iobst, IDA/CCS

High Volume Technology for HPC Systems 185
Justin Rattner, Intel Corporation

High Performance Computing and the NCAR Procurement—Before, During, and After 209
Bill Buzbee, NCAR; Jim Hack and Steve Hammond, National Center for Atmospheric Research

Economies of Scale: HPC into the Next Millennium 223
Gary Smaby, Smaby Group

The Other Side of Computing 241
William Trimmer, Belle Mead Research, Inc.

Crystalline Computing 249
Norm Margolus, MIT

Quantum Computing 257
Emanual Knill, LANL

List of Attendees 277

Conference Program

Monday, April 20, 1998

Keynote Session:

Keynote Address: Billions and Billions
Steve Wallach, CenterPoint Venture Partners

Tuesday, April 21, 1998

Session 1: The Stockpile Stewardship and Management Program

Stockpile Stewardship and Management Program
Larry Ferderber, LLNL

Predictability, and the Challenge of Certifying a Stockpile Without Nuclear Testing
Ray Juzaitis, LANL

Session 2: The Challenge of 100 TeraFLOP Computing

100 TeraFLOPs and Beyond, an Industry View into the Future
Panel Discussion: Moderator—John Morrison, LANL and Mark Seager, LLNL; Panelists—Tilak Agerwala, IBM; Greg Astfalk, Hewlett-Packard; Erik Hagersten, Sun Microsystems; Richard Kaufmann, Digital Equipment Corp.; Steve Oberlin, SGI/Cray.

Session 3: ASCI Alliance

ASCI Alliances Program
Ann Hayes, LANL

Session 3.5: Hardware Design

The Next Fifty Years of Computer Architecture
Burton Smith, Tera Computer Company

Banquet

Adversarial Inspection in Iraq: 1991 and Thereafter
Jay Davis, LLNL

Wednesday, April 22, 1998

Session 5: Student Session

Full Wave Modeling of Signal and Power Interconnects for High Speed Digital Circuits
Gary Haussmann, University of Colorado

Simulating the Physical-Biological Factors Affecting Abundance of Calanus finmarchicus in the Gulf of Maine
Wendy Gentleman, Dartmouth College

Session 6: News You Can Use

The Next Generation Internet
Bob Aiken, DOE

Petaflops Computing: Opportunities and Challenges
David Bailey, LBNL

President's Information Technology Advisory Committee (PITAC): A Mid-Term Report
David M. Cooper, LLNL

Thursday, April 23, 1998

Session 7: Chip Technology for Large Scale Systems

Processing-in-Memory: Past and Present
Ken Iobst, IDA/CCS

High Volume Technology for HPC Systems
Justin Rattner, Intel

Session 8: Reality Check

High Performance Computing and the NCAR Procurement—Before, During, and After
Bill Buzbee, NCAR; Jim Hack and Steve Hammond, National Center for Atmospheric Research

Economies of Scale: HPC in the Next Millennium
Gary Smaby, Smaby Group

Session 9: Future

The Other Side of Computing
William Trimmer, Belle Mead Research, Inc.

Crystalline Computation
Norm Margolus, MIT

Quantum Computing
Emanuel Knill, LANL

Proceedings from the Conference on High Speed Computing

The Art of High Speed Computing

April 20–23, 1998

Compiled by
Kathleen P. Hirons
Manuel Vigil
Ralph Carlson

Abstract

This document provides a written compilation of the presentations and viewgraphs from the 1998 Conference on High Speed Computing. "The Art of High Speed Computing," held at Gleneden Beach, Oregon, on April 20 through 23, 1998.

BILLIONS & BILLIONS

STEVE WALLACH
CENTER POINT VENTURES
WALLACH@CENTERPOINTVP.COM

ASPECTS OF BILLIONS

- Raised to the power (giga, tera, peta, exa)
- The inverse (nano, pico, femto, atto)
- In the computer industry they are closely related. From a technology and investment perspective
- US government policy must be consistent with industry trends. (the ultimate venture capitalist)

2

PRESENTATION OUTLINE

- Fundamental Laws- Physics
- Trends in Telecommunications
- Trends in Semi-conductors
- Trends in Computer Architecture
- Draw some conclusions
- US Government Policy

3

FUNDAMENTAL LAWS

- C - Speed of light
- Power Consumption
- Propagation Delay

4

POWER CONSUMPTION

$$P \cup C * V^2 * F$$

- C= capacitance
- V= voltage
- F= frequency

5

PROPAGATION DELAY

- Lossless Line

$$\text{Time} = \sqrt{LC}$$

- Lossly Line

$$\text{Time} = L * \sqrt{\varepsilon_r / c_o}$$

$$\varepsilon_r = \text{Dielectric Constant}$$

$$c_o = \text{Speed of Light}$$

6

OTHER CONSTRAINTS

- Cost of Investment - I (billions)
- Size of Market - M (millions)
- L'Hospital's Rule of Profit
  - – Profit = dM/dI
    - **as I approaches infinity**
    - **as M approaches K (sometimes 0)**
  - – result is { 1 (success) | 0 (failure)}
- The government uses different rules

7

TELECOMMUNICATIONS

- Advances in ***PHOTONIC*** (mainly ***WDM***) technology.
- ***TERAHZ*** (THz) requirements
- All optical networks (***AON)***
- Effect on digital computer architecture
- The next supercomputer topology
    - **www.ll.mit.edu/aon/**
    - **Lemott, et. al., " low-cost WDM", Aug. 97, IEEE summer topicals, Montreal.**

8

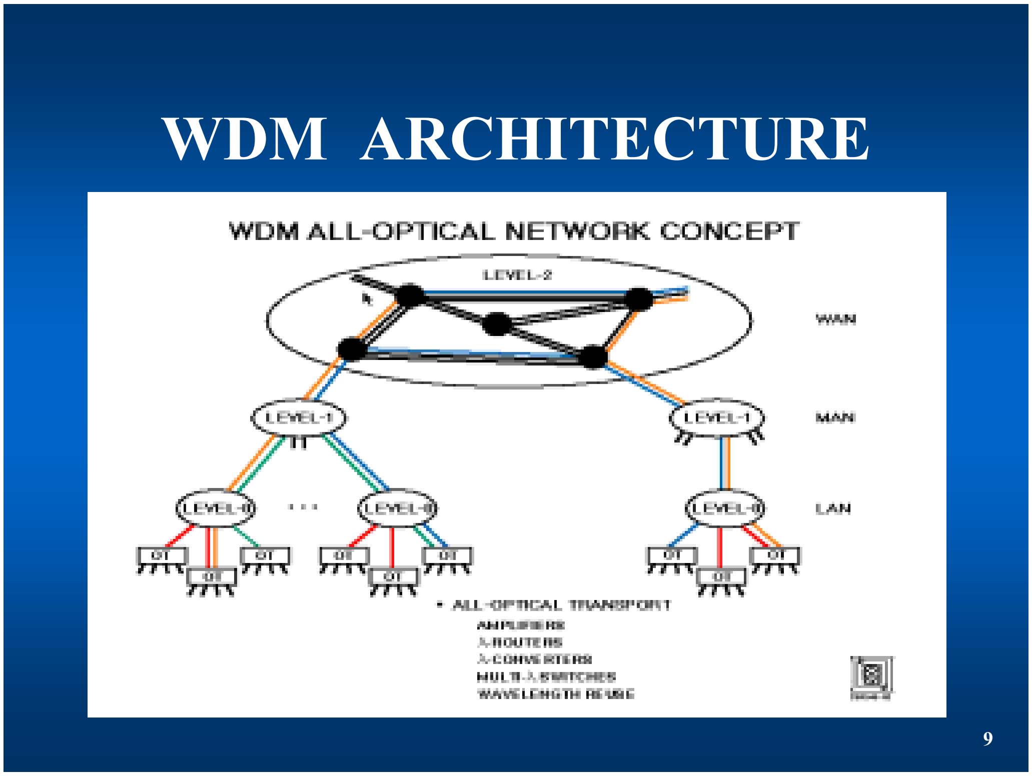
WDM ARCHITECTURE
WDM ALL-OPTICAL NETWORK CONCEPT
LEVEL-2
WAN
LEVEL-1
LEVEL-1
MAN
LEVEL-0
LEVEL-0
LEVEL-0
LAN
OT
• ALL-OPTICAL TRANSPORT
AMPLIFIERS
λ-ROUTERS
λ-CONVERTERS
MULTI-λ-SWITCHES
WAVELENGTH REUSE
9

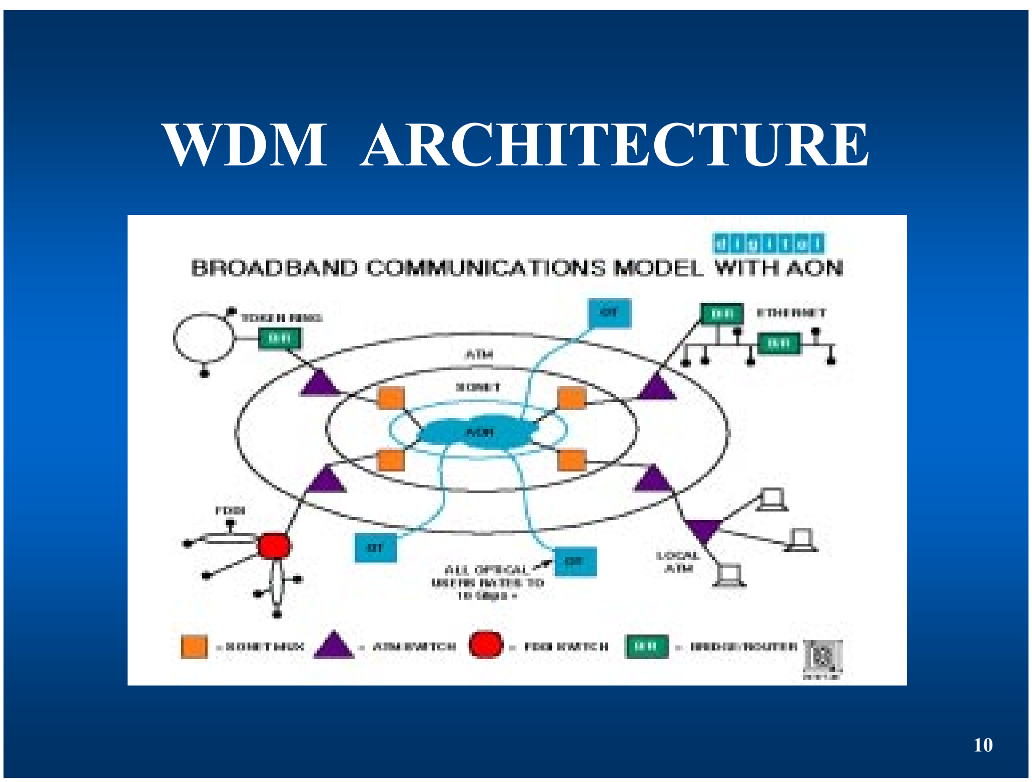
WDM ARCHITECTURE
digital
BROADBAND COMMUNICATIONS MODEL WITH AON
TOKEN RING
ETHERNET
ATM
SONET
AON
FDDI
OT
LOCAL ATM
= SONET MUX
= ATM SWITCH
= FDDI SWITCH
= BRIDGE/ROUTER
10

TECHNOLOGY

WAVELENGTH DIVISION MULTIPLEXING

$$c = \lambda * f$$

- EXPLOITS
  - ENORMOUS BANDWIDTH OF SILICA FIBER
  - HIGH-GAIN WIDEBAND OPTICAL AMPLIFIERS

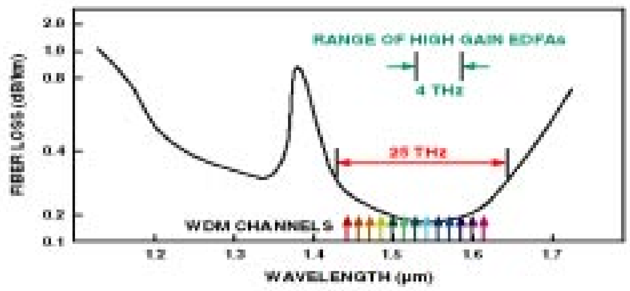

11

WDM IMPLEMENTATION

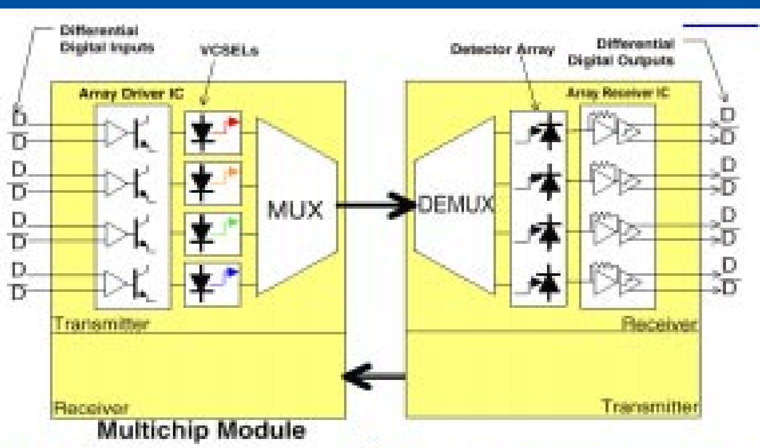

HEWLETT PACKARD

HP Labs / Communications and Optics Research Lab

12

SEMI-CONDUCTORS

- Lets examine what is driving the ***I*** (investment) in our equation for success.
  - Information from 1997 SIA report (www.semichips.org)

13

THE COST OF "FABS"

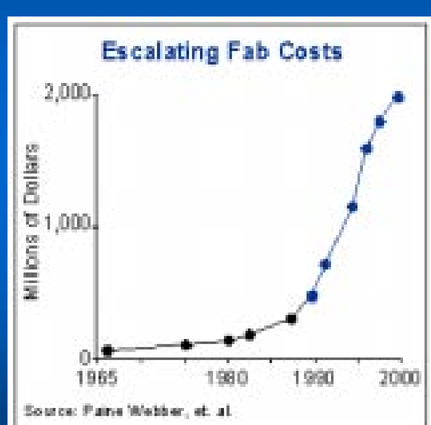

- 2 billion and climbing
- One per continent?
- Put on the moon?
- Only million piece design can be made?

14

UNDERLYING REASONS

- 300 mm (12inch) wafers
- Billions to replace 8 inch fabs.
- Good news: keeps costs of chips down

15

UNDERLYING REASONS

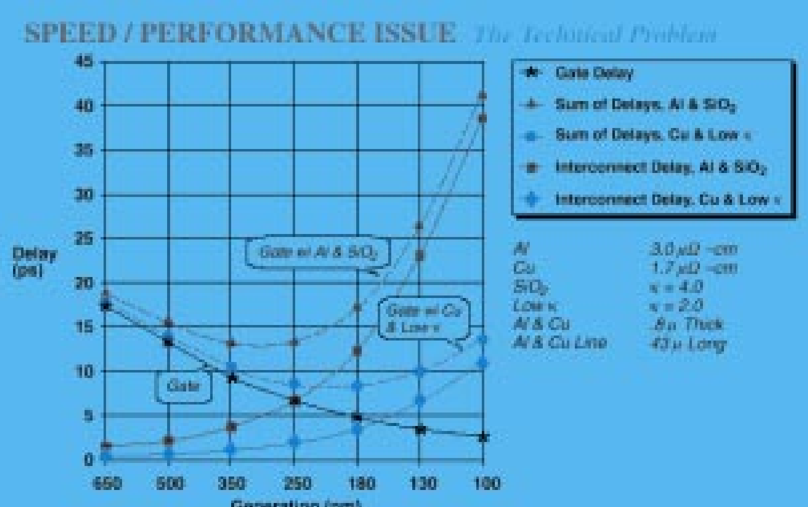

Figure 3 Calculated Gate and Interconnect Delay versus Technology Generation

16

HOW WE GET THERE

Table 63 Modeling & Simulation Technology Requirements (Continued)

Year of First Product Shipment Technology Generation	1997 250 nm	1999 180 nm	2001 150 nm	2003 130 nm	2006 100 nm	2009 70 nm	2012 50 nm
Numerical Methods							
Linear solvers—equations/minute	100k	150k	250k	250k	2.5M	5M	5M
Parallel speedup	—	4×	6×	9×	16×	30×	50×
Grid reliability (ppb)	300	180	120	90	26	14	7
MFLOPS* required	50	80	400	1000	4000	8000	8000
MC noise	NA	NA	NA	0.05	0.02	0.01	0.001
Simulation Environments							
Time needed for statistical sim.	10 weeks	6 weeks	4 weeks	2 weeks	2 weeks	1 week	1 week
Time needed for multi-tool initial problem setup	4 weeks	2 weeks	1 week	4 days	2 days	2 days	2 days
Correct data analyses per improvement cycle	0.1	1	1	1	2	4	10

Solutions Exist / *Solutions Being Pursued* / *No Known Solution*

* *MFLOPS—million floating point operations per second*

* *Number of linear equations generated by discretizing an increasing number of PDEs over a typical device grid of 5000 nodes in 2-D and later 50000 nodes in 3-D*

17

HOW WE GET THERE

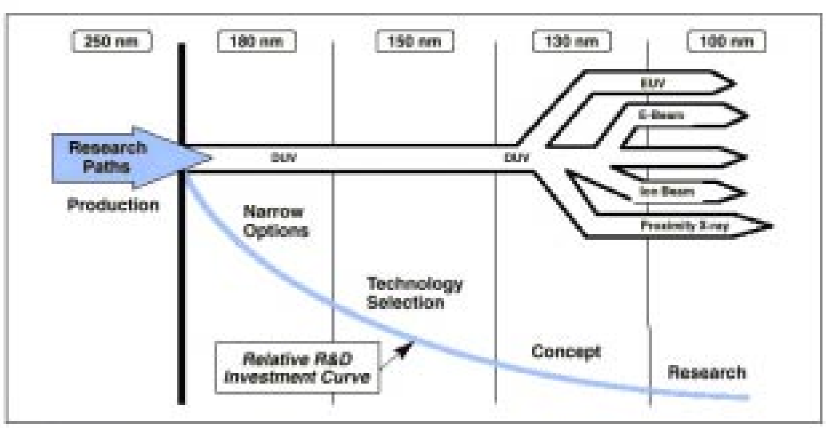

DUV—deep ultraviolet
EUV—extreme ultraviolet
E-beam—electron beam

Figure 1 Conceptual Illustration of Today's Research and Development Investments for Future Production Technologies

18

WHAT WE GET

Table 3 Performance of Packaged Chips

YEAR OF FIRST PRODUCT SHIPMENT	1997	1999	2001	2003	2006	2009
TECHNOLOGY GENERATIONS DENSE LINES (DRAM HALF PITCH) (nm)	250	180	150	130	100	70
ISOLATED LINES (MPU GATES) (nm)	200	140	120	100	70	50
Number of Chip I/Os						
Chip-to-package (pads) high-performance	1450	2000	2400	3000	4000	5400
Chip-to-package (pads) cost-performance	800	975	1195	1460	1970	2655
Number of Package Pins/Balls						
ASIC (high-performance)	1100	1500	1800	2200	3000	4100
MPU/controller, cost-performance	600	810	900	1100	1500	2000
Cost-performance package cost (cents/pin)	1.40–2.80	1.25–2.50	1.15–2.30	1.05–2.05	0.90–1.75	0.75–1.50
Chip Frequency (MHz)						
On-chip local clock, high-performance	750	1250	1500	2100	3500	6000
On-chip, across-chip clock, high-performance	750	1200	1400	1600	2000	2500
On-chip, across-chip clock, cost-performance	400	600	700	800	1100	1400
On-chip, across-chip clock, high-performance ASIC	300	500	600	700	900	1200
Chip-to-board (off-chip) speed, high-performance (Reduced-width, multiplexed bus)	750	1200	1400	1600	2000	2500

WHAT WE GET

Table 24 Product Critical Level Lithography Requirements

Year of First Product Shipment Technology Generation	*1997 250 nm*	*1999 180 nm*	*2001 150 nm*	*2003 130 nm*	*2006 100 nm*	*2009 70 nm*	*2012 50 nm*
Product Application							
DRAM (bits)	256M	1G	—	4G	16G	64G	256G
MPU (logic transistors/cm²)	4M	6M	10M	18M	39M	84M	180M
ASIC (usable transistors/cm²)*	8M	14M	16M	24M	40M	64M	100M
*Minimum Feature Size (nm)***							
Isolated lines (MPU Gates)	200	140	120	100	70	50	35
Dense lines (DRAM Half Pitch)	250	180	150	130	100	70	50
Contacts	280	200	170	140	110	80	60
Development capability (minimum feature size, nm)	140	120	100	70	50	35	25
Gate CD control (nm, 3 sigma at post-etch)**	20	14	12	10	7	5	4
Product overlay (nm, mean + 3 sigma)**	85	65	55	45	35	25	20

20

HOW WE USE IT

- TRENDS IN COMPUTER ARCHITECTURE

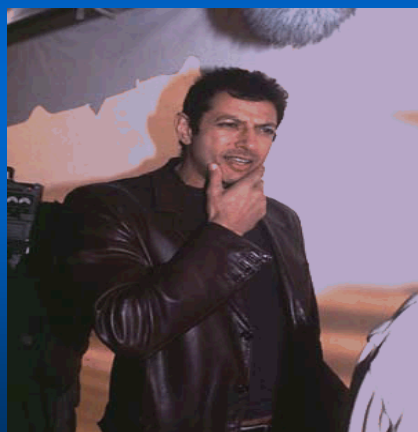

GENERAL VIEWS

2 TO 4 YEARS

10 YEARS (USING SIA STUDY)

21

PC/PERSONAL WORKSTATION

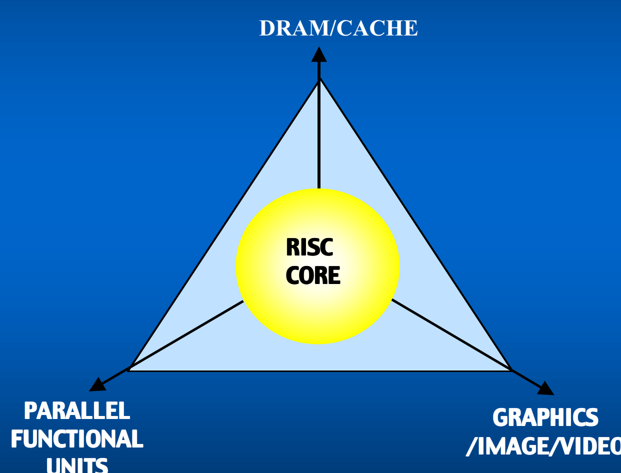

22

DIGITAL SIGNAL PROCESSOR

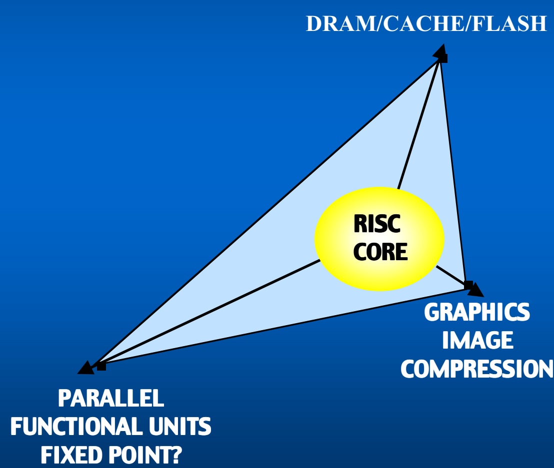

23

GRAPHICS/IMAGE PROCESSOR

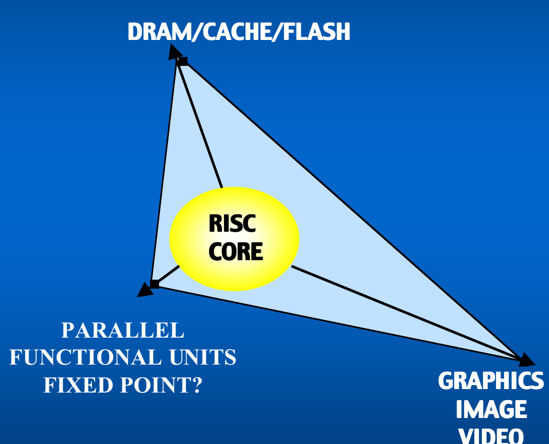

24

APPLICATION STRUCTURE ?

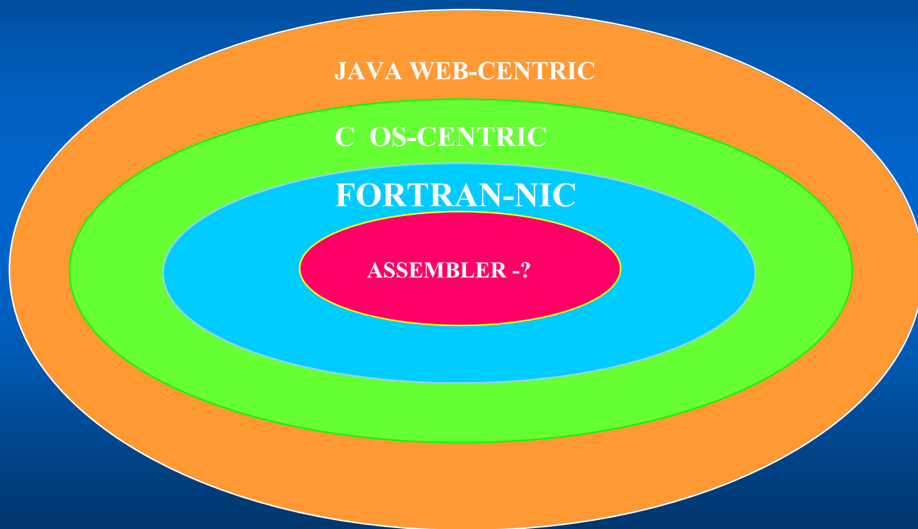

25

JAVA Vs C++

FEATURE	JAVA	C++
Memory Management	Garbage collected	Explicit Memory Freeing
Multi-threading	YES (Mesa-style)	NO
Inheritance Model	Simpler (separate sub-typing)	Complex
Exception handling	Supported	Sporadic
Parametric type	Does Not	Has template
Type casts	Checked Thus easier to write protected subsystems	Unchecked (pointer ⟷ integer)

26

SO WHAT HAPPENS?

- Fundamentally the following architecture evolves:
    - *PIM (processor in memory) or System-on-a-chip*
        - **more memory bandwidth**
        - **lower latency**
        - **consistent with PC pricing and technology curves**

27

WHICH APPROACH?

- SIMD
- MIMD
- MULTI-THREADED
- SUPERSCALAR
- VLIW

28

INTERCONNECT TYPE

- ***SOFTWARE***
  - – Message Passing
  - – Distributed Shared Memory (DSM)
  - – Cache Only (COMA)
  - – Object oriented
  - – Emulated DSM (e.g.. Threadmarks)

29

INTERCONNECT TYPE

- ***HARDWARE***
  - – Hierarchical - number of levels is a function of the number of cpu's.
  - – Physical - combination of copper and photonic. Ultimately ***WDM*** will play an important role in external chip interconnects.

30

NEXT ARCHITECTURES

- Short Term - 2 To 4 years- low performance - System-on-a-chip (SOAC)
- Long Term - 10 years (using SIA study)
  - High Performance
  - Supercomputing
- US Gov't R&D Policy

31

ARCH.- SHORT TERM

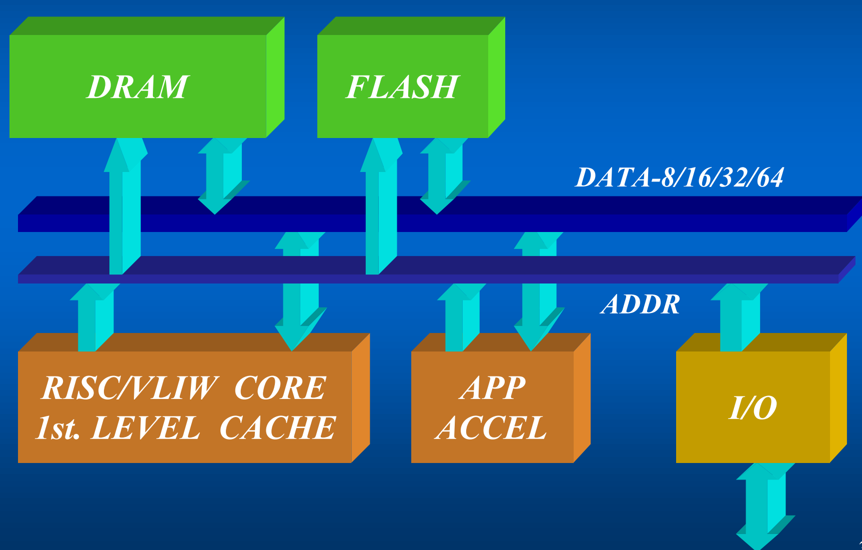

32

ARCH.- SHORT TERM

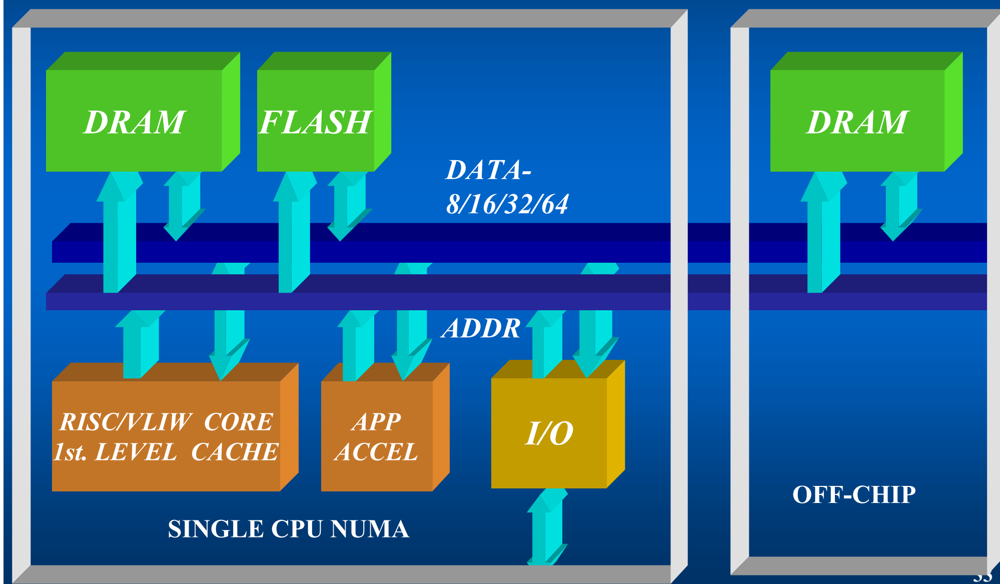

ARCH. - LONG TERM- 2009

- **THE SIA STUDY TEACHES US:**
  – 64 gbits of dram - (8 gbytes)
  – 8 gbits of sram
  – 520 million MPU transistors
  – 70 nm lithography, 2.54 cm on-a-side
  – 6 ghz clock within vliw/risc core
  – 2.5 ghz across die
  – 2500 external signal pins

34

ARCH. - LONG TERM - 2009 DESIGN ASSUMPTIONS

- 9 million transistors - vliw/risc core with first level cache.
- 2nd. Level cache - rule of thumb. 1/4 to 1/2 mbyte per 100 mflops peak.
- 96 mbyte 2nd. Level (6 Inst, 90 data)
- 170 watts
- .6 to .9 volts power supply

35

MAXIMUM PIN-USE EXTERNAL SMP- 6/8 CPU'S

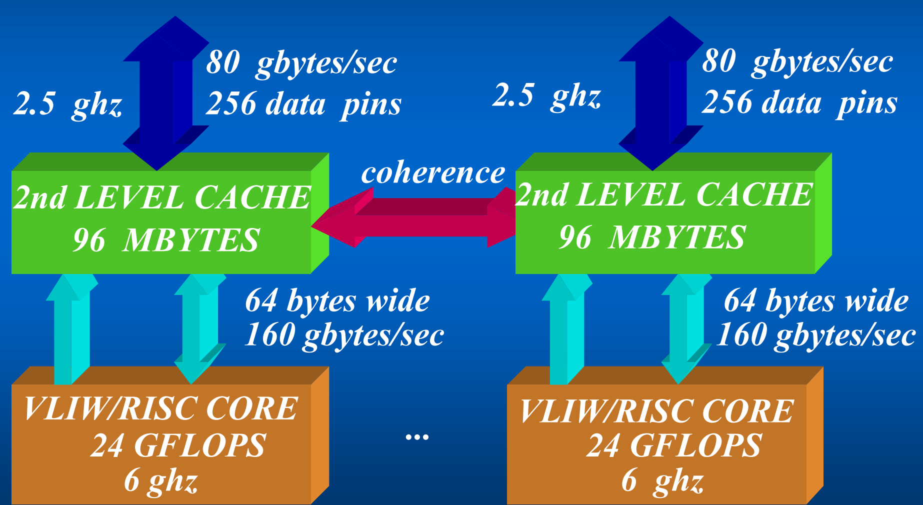

36

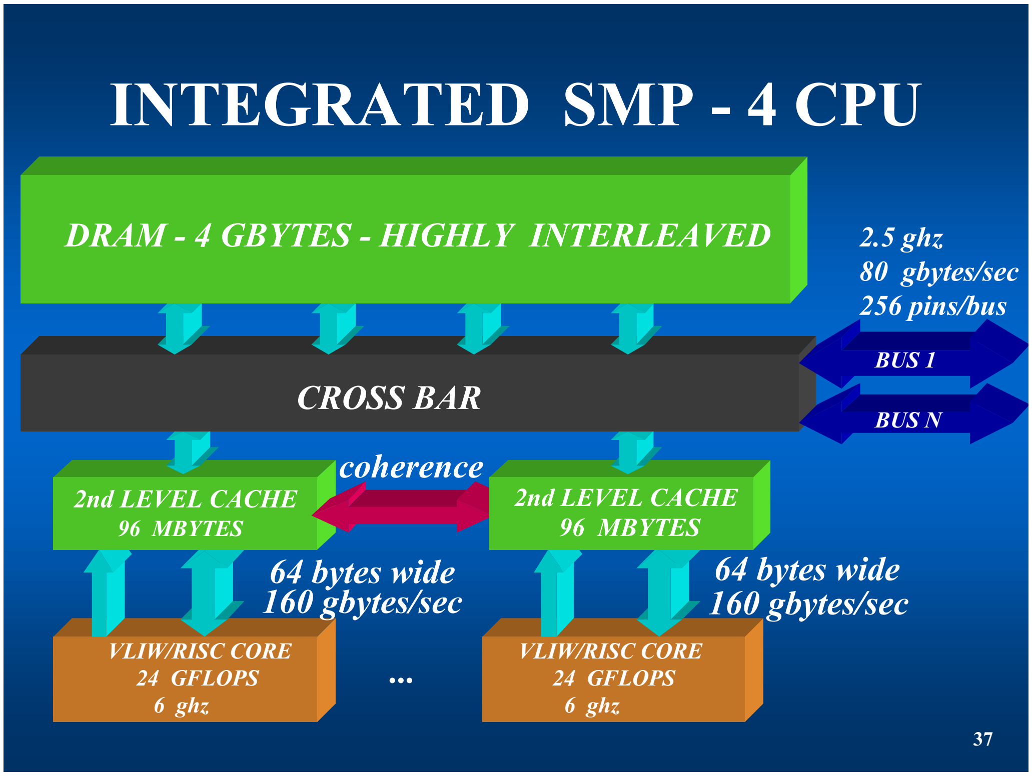
INTEGRATED SMP - 4 CPU
DRAM - 4 GBYTES - HIGHLY INTERLEAVED
2.5 ghz
80 gbytes/sec
256 pins/bus
BUS 1
BUS N
CROSS BAR
coherence
2nd LEVEL CACHE
96 MBYTES
2nd LEVEL CACHE
96 MBYTES
64 bytes wide
160 gbytes/sec
64 bytes wide
160 gbytes/sec
VLIW/RISC CORE
24 GFLOPS
6 ghz
...
VLIW/RISC CORE
24 GFLOPS
6 ghz
37

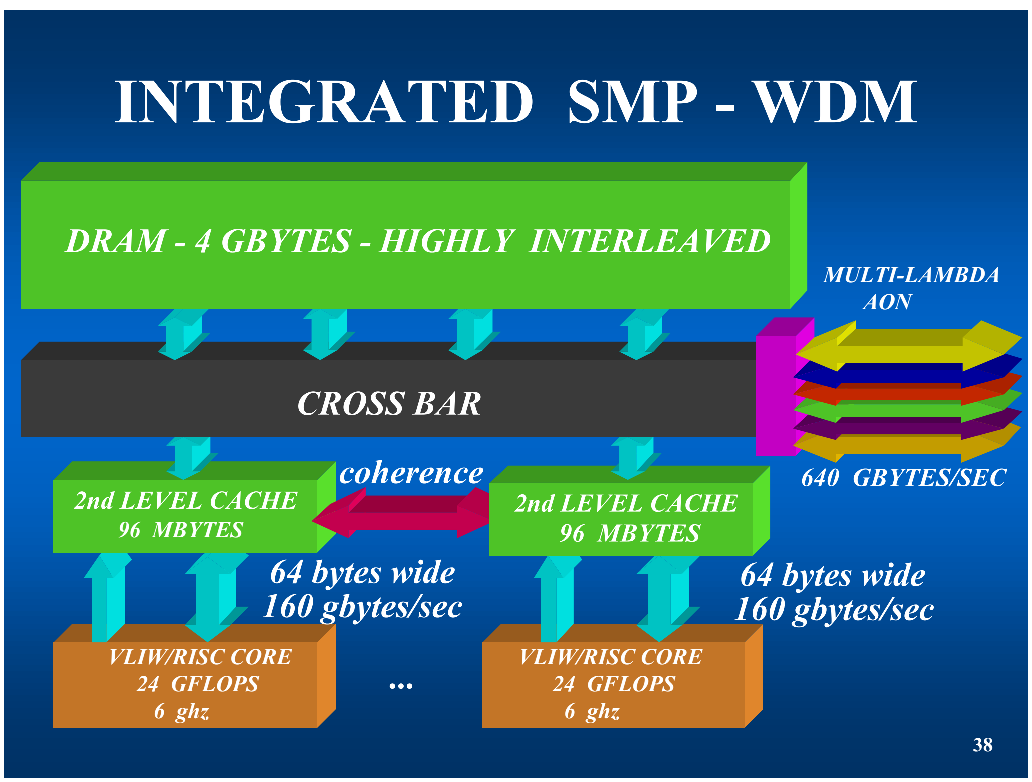
INTEGRATED SMP - WDM
DRAM - 4 GBYTES - HIGHLY INTERLEAVED
MULTI-LAMBDA
AON
CROSS BAR
640 GBYTES/SEC
coherence
2nd LEVEL CACHE
96 MBYTES
2nd LEVEL CACHE
96 MBYTES
64 bytes wide
160 gbytes/sec
64 bytes wide
160 gbytes/sec
VLIW/RISC CORE
24 GFLOPS
6 ghz
...
VLIW/RISC CORE
24 GFLOPS
6 ghz
38

US GOVERNMENT POLICY

- Examine the Past
- Use Tops 500 - LINPACK
- Observe Venture Capital Investments
- What should happen in the future

39

PROCESSOR SCALABILITY GENERAL PURPOSE-CPU

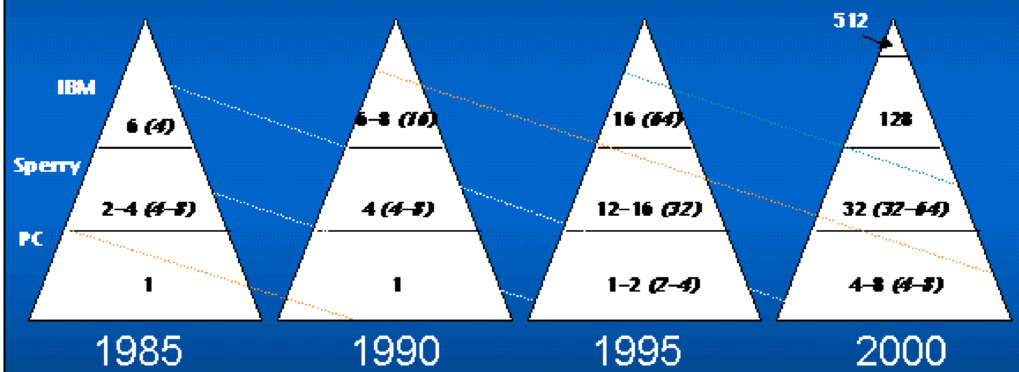

40

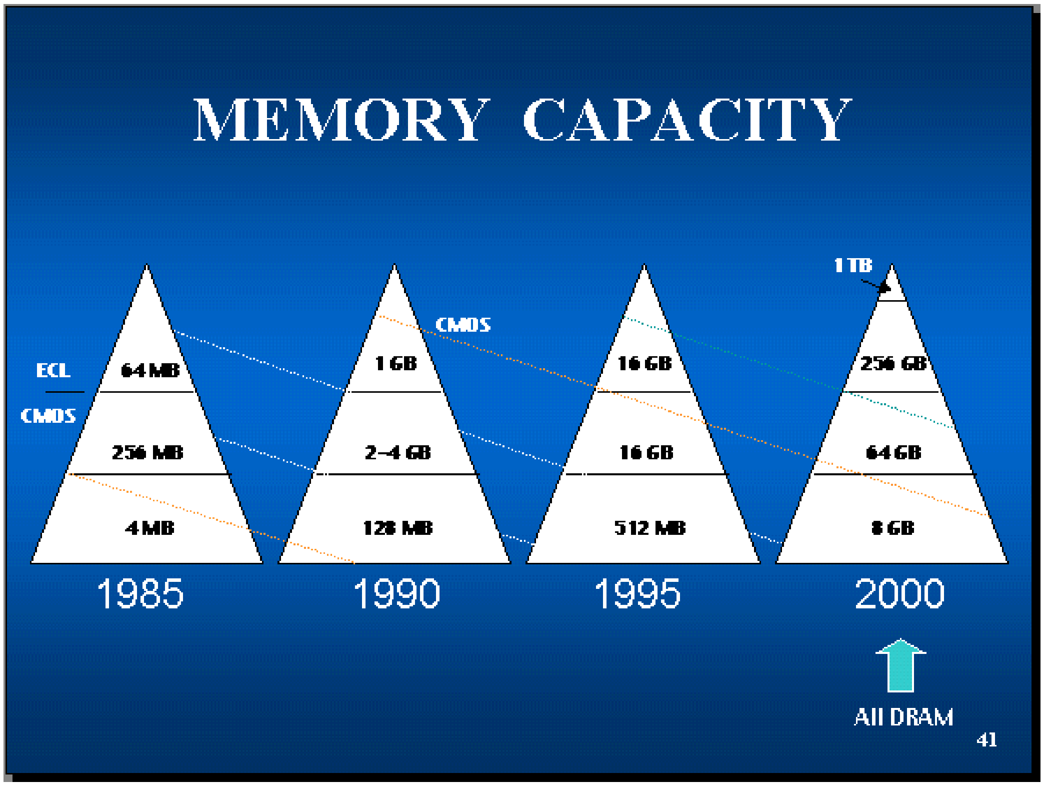
MEMORY CAPACITY
1 TB
CMOS
ECL
CMOS
64 MB
256 MB
4 MB
1 GB
2-4 GB
128 MB
16 GB
16 GB
512 MB
256 GB
64 GB
8 GB
1985
1990
1995
2000
All DRAM
41

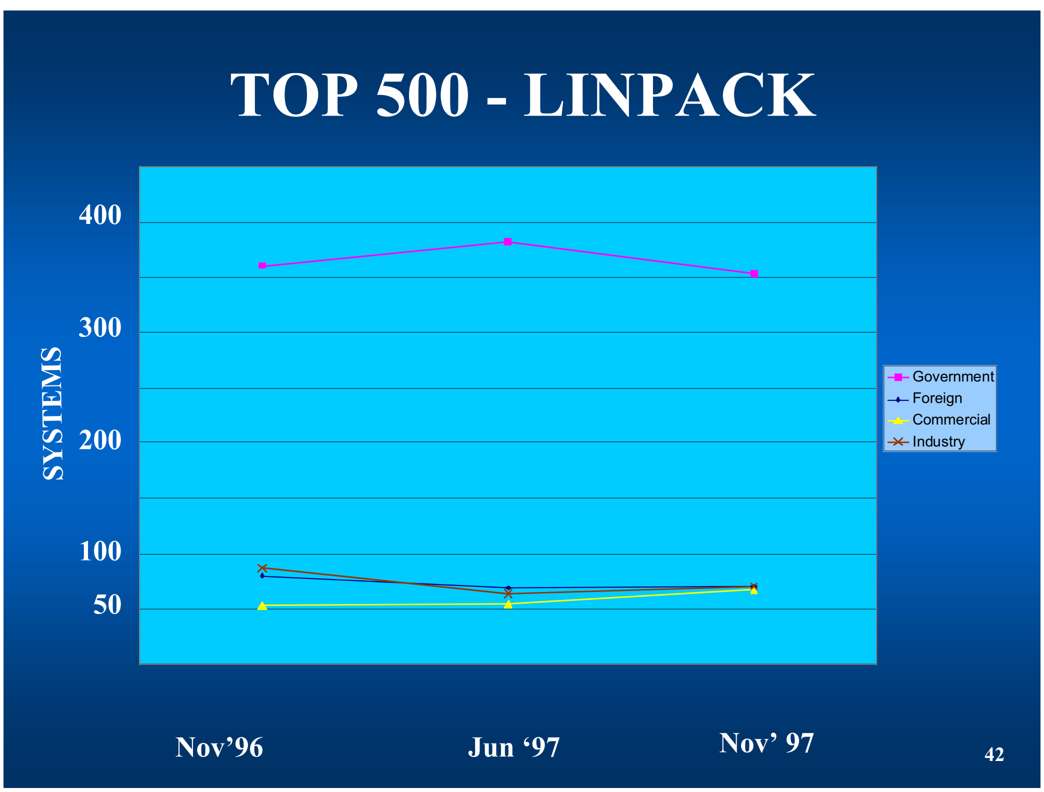
TOP 500 - LINPACK
SYSTEMS
400
300
200
100
50
Government
Foreign
Commercial
Industry
Nov'96
Jun '97
Nov' 97
42

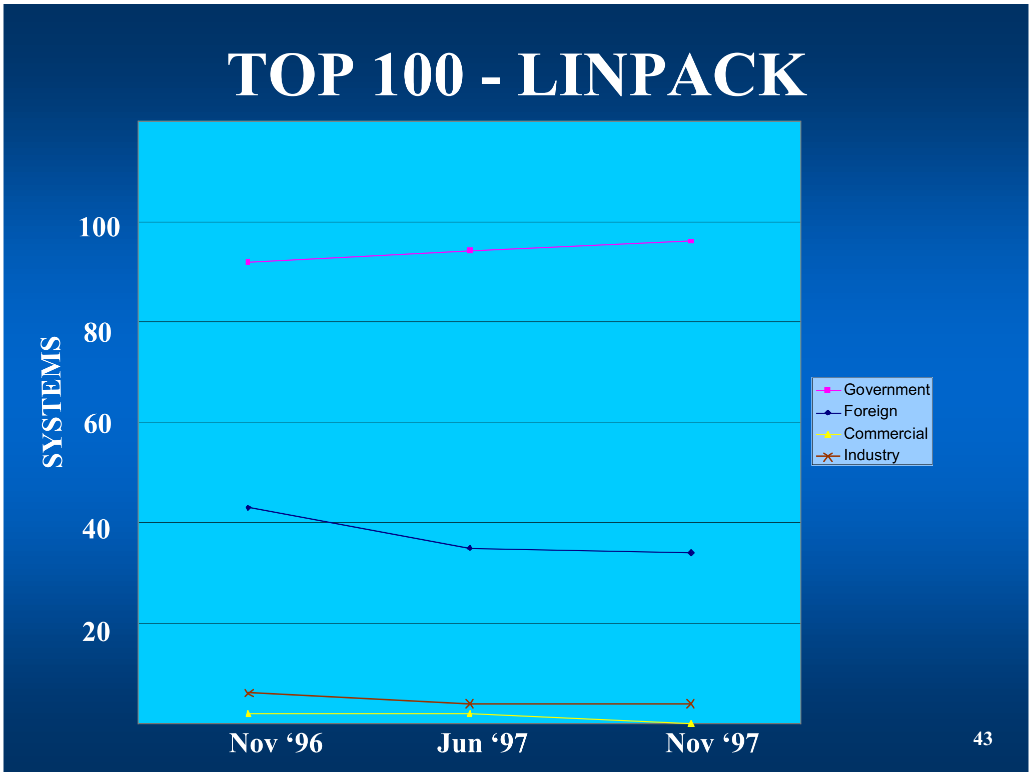
TOP 100 - LINPACK
SYSTEMS
100
80
60
40
20
Nov '96
Jun '97
Nov '97
Government
Foreign
Commercial
Industry
43

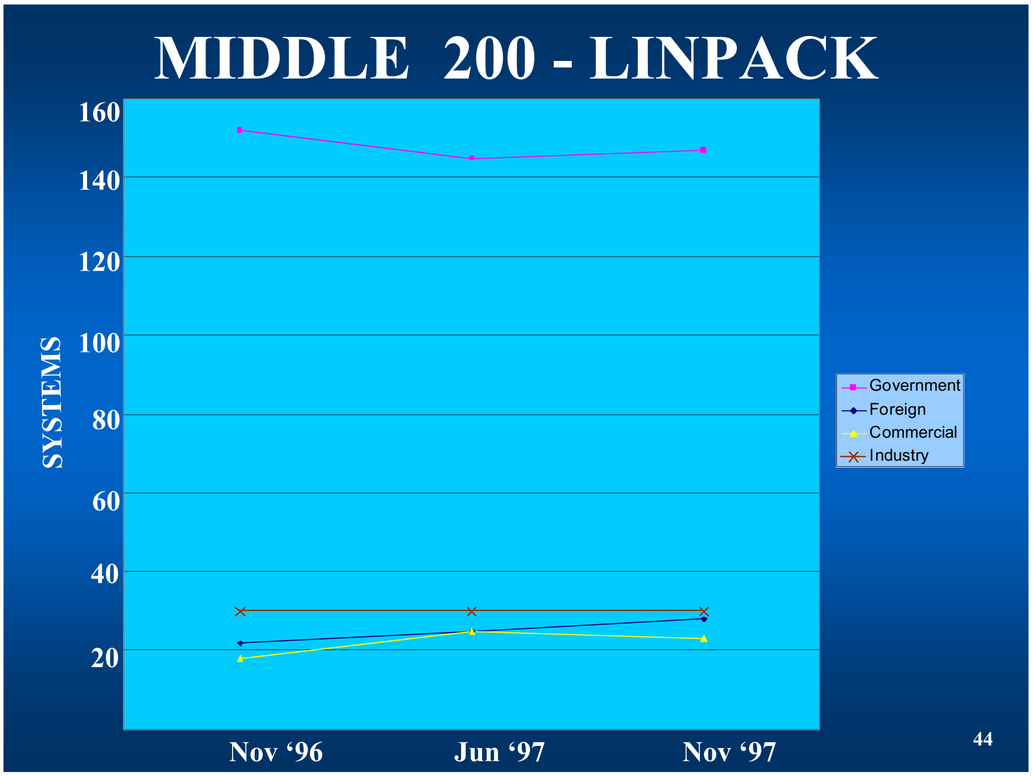
MIDDLE 200 - LINPACK
SYSTEMS
160
140
120
100
80
60
40
20
Nov '96
Jun '97
Nov '97
Government
Foreign
Commercial
Industry
44

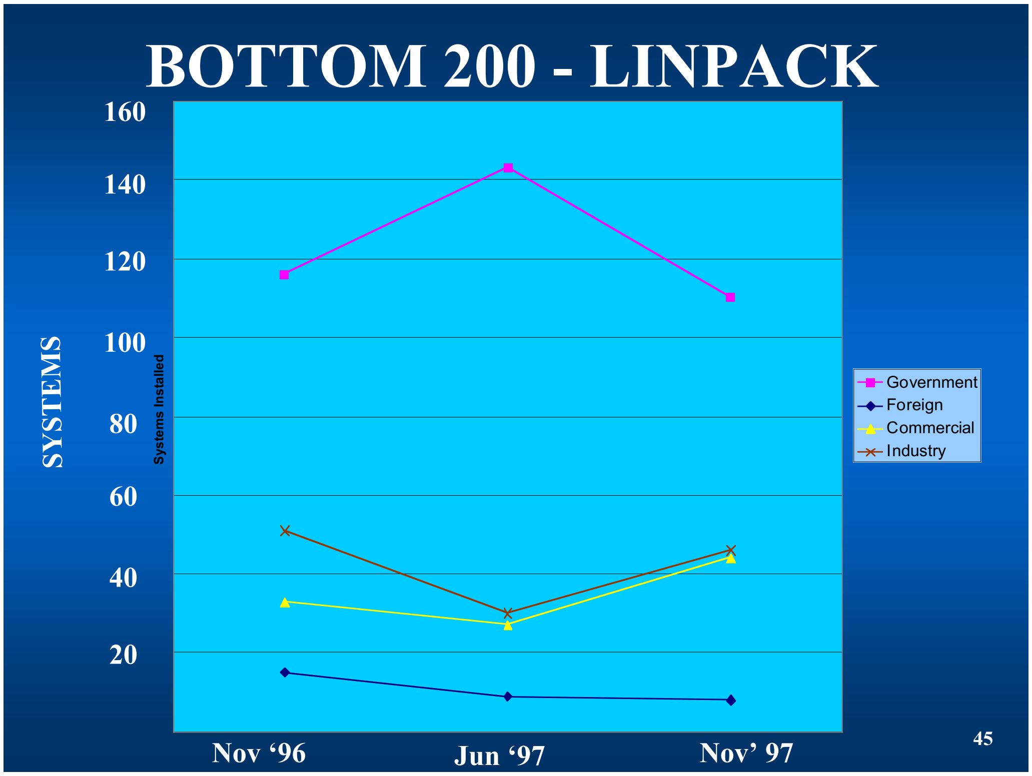
BOTTOM 200 - LINPACK
SYSTEMS
Systems Installed
160
140
120
100
80
60
40
20
Government
Foreign
Commercial
Industry
Nov ‘96
Jun ‘97
Nov’ 97
45

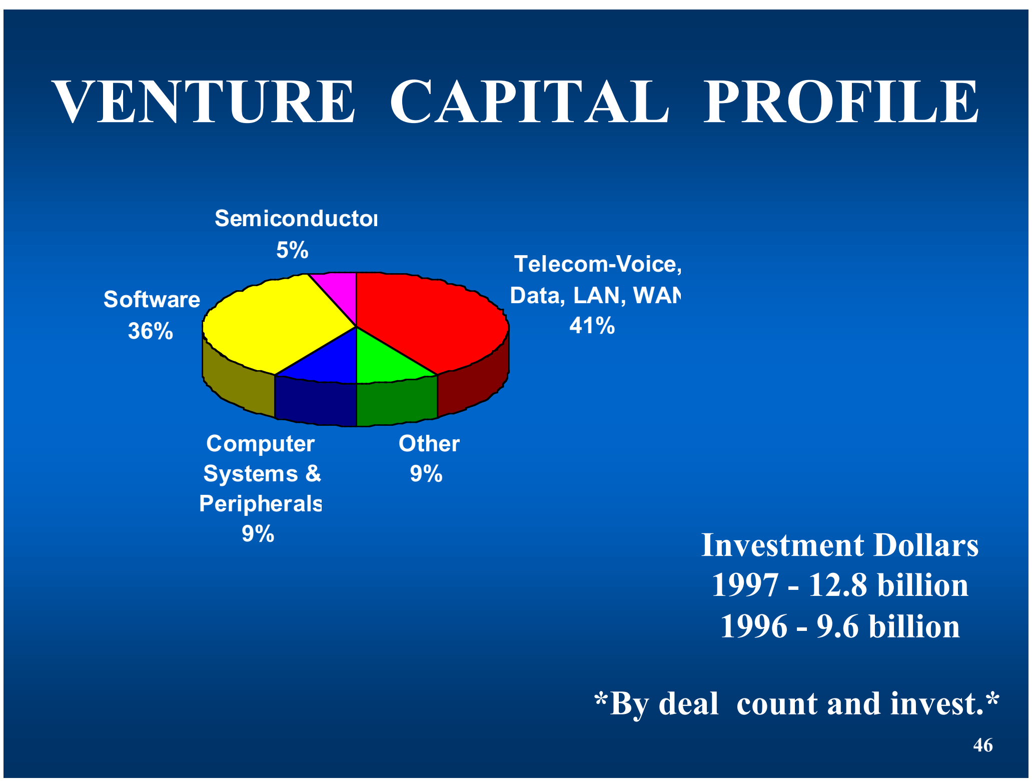
VENTURE CAPITAL PROFILE
Semiconductor
5%
Telecom-Voice,
Data, LAN, WAN
41%
Software
36%
Computer
Systems &
Peripherals
9%
Other
9%
Investment Dollars
1997 - 12.8 billion
1996 - 9.6 billion
By deal count and invest.
46

US GOVERNMENT POLICY

- Provide seed money - high risk/reward (darpa, nsf, dod, doe)
- Further national defense initiatives
- Begin the trickle down, technology xfer. What starts out as a US Gov't special becomes COTS after 1 or 2 generations
- Keep the US the most advanced and competitive in the world
- www.hpcc.gov/talks/petaflops-24june97

47

CONCLUDING

- **Convergence of telecommunications/computing**
  - **everything is digital**
  - **everything requires high bandwidth**
  - **voice is a digital packet (IP switching)**
  - **digital TV (a TV with a computer or a computer with a TV?)**
  - **overall system topology mirrors an AON**
- **Commodity Teraflop Computing**

Stockpile Stewardship Program (U)

1998 Conference on High Speed Computing
Gleneden Beach, Oregon
April 20-23, 1998

Lawrence J. Ferderber
Deputy Associate Director for National Security
Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94551

NS-98-031.1

The President tasked DOE to help maintain the nuclear deterrent through the Stockpile Stewardship Program

"... I consider the maintenance of a safe and reliable nuclear stockpile to be a supreme national interest of the United States."

"I am assured by the Secretary of Energy and the Directors or our nuclear weapons labs that we can meet the challenge of maintaining our nuclear deterrent under a Comprehensive Test Ban Treaty through a Science-Based Stockpile Stewardship program without nuclear testing..."

"In order for this program to succeed, both the Administration and the Congress must provide sustained bipartisan support for the stockpile stewardship program over the next decade and beyond. I am committed to working with the Congress to ensure this support."

"As part of this arrangement, I am today directing the establishment of a new annual reporting and certification requirement that will ensure that our nuclear weapons remain safe and reliable under a comprehensive test ban."

– August 11, 1995

NS-98-031. 2

Today the stockpile is safe and reliable, but we already require a Stockpile Stewardship Program to keep it that way

- Today's stockpile has a good "pedigree" based on
  - Nuclear tests
  - An experienced workforce
  - State of the art design (then)
- But
  - The stockpile is aging beyond our experience
  - Refurbished components will be made by new processes, in new plants by new people
  - Our experienced workforce is retiring
  - We have no nuclear tests to verify the validity of our decisions
- We need a program that will:
  - Attract and train a new workforce
  - Be able to assess the effect of changes in the stockpile
  - Certify that refurbished components are functionally equivalent to the original ones

NS-98-031. 3

Like every other technological object, a nuclear weapon ages and sometimes we are surprised when we test it

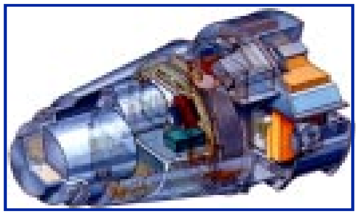

- One-point safety
- Performance at cold temperatures
- Performance under aged conditions
- Marginal performance
- Degradation of various key materials
- Pit quality control
- Metal components cracking
- Yield-select problems
- HE degrading
- HE cracking
- Detonators corroding
- Detonator system redesign
- Metal components corroding

NS-98-031. 4

The Stockpile Stewardship Program responds to these challenges via a few fundamental principals

- We have experimental data from nuclear tests which indicate that details matter – remanufactured components sometimes behave anonymously
- Current experimental and computational capabilities are not sufficient to preclude that these anomalies will occur in the future
- Without nuclear testing, we must take the conservative approach in proving our fixes are real fixes which do not introduce new problems
- We must also develop a strategy to deal with the "unknown anomalies" (e.g. Challenger O-ring) ... including residual design flaws that have not yet manifested themselves
- The stockpile will continue to age and we will be required to deal with changes to almost every components

The SSP approach is not without risk

NS-98-031. 5

Simple remanufacture is not a credible solution for highly optimized and complex products like nuclear weapons

- This point is illustrated by the Polaris A3 motor rebuild
  - The U.S. production line was placed on standby in 1963
  - Procedures were carefully documented
  - Nineteen years later, in 1982, it was found that the "replica" rebuild of the rocket motors required extensive full scale testing to get it right (four of the flight tests failed)
  - A recall of retired personnel was necessary

Replicating nuclear weapons would be more difficult (impossible) than replicating rocket motors

NS-98-031. 6

The SSP provides integrated capabilities to address DoD's near-term and longer-term issues

Four SSP Strategies:

- Surveillance
  - to monitor, maintain and predict the condition of the stockpile
- Assessment & Certification
  - of the consequences of change
  - that modifications and maintenance do not degrade warhead safety and reliability
- Refurbishment
  - design and manufacture of refurbished components
- Tritium replacement

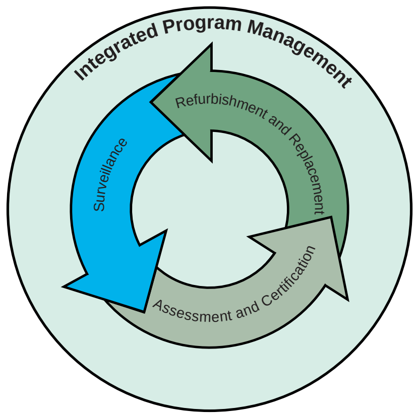

NS-98-031. 7

Our surveillance program is being expanded to meet the needs of an aging stockpile

- How do weapons age?
- What are the most likely issues?
- How will these issues affect performance and safety?
- When do components need to be refurbished?

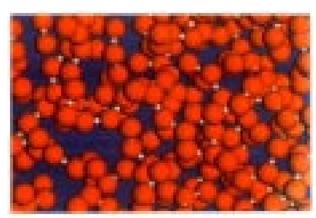

Interstitial helium

Assessment of disassembled components

Forensic surveillance techniques

NS-98-031. 8

The new complex must refurbish/replace components to counter age, performance, or safety degradation

Plutonium pits
Los Alamos, New Mexico

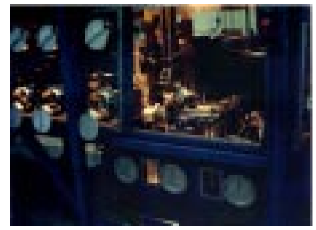

Integrated plutonium processing

Assembly expertise

These new plants, people and processes must be certified to be functionally equivalent those originally used

NS-98-031. 9

SSP requires dramatic advances in computational capabilities

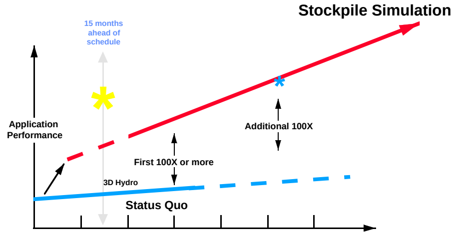

Accelerated Strategic Computing Initiative (ASCI)

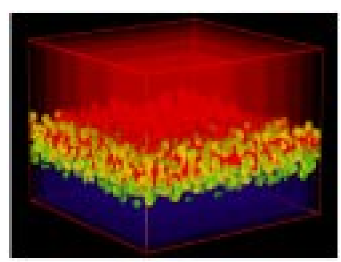

3D turbulent mix simulation

ASCI Blue Pacific SST, LLNL
3.3 TeraFlops
2.5 Terabytes

NS-98-031. 11

During nuclear testing, we depended on a design-test-build cycle to certify the performance of nuclear weapons

The computational capability was sufficient to provide reasonable assurance that the test would function properly (a cost issue)

The UNIVAC in 1953

LLNL's first "nuclear" test was designed on the UNIVAC and slide rules

- LLNL was less than one year old
- The device was placed on a 300-foot tower and the physicists stood far away, observing with dark glasses
- Upon detonation, only a small cloud of dust appeared
- When the dust cleared, the tower was still standing

Nowhere to go but up ... and 50 years later, the ASCI program

NS-98-031. 1

During design-test-build, our simulation codes normalized complex phenomena against test data

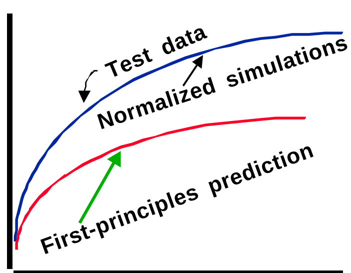

- Computers lacked speed and memory to run full problems
- Some nonlinear physical processes not understood
- Nuclear test data provided normalizing factors to make simulations accurate
- Normalization factors differed from system to system

NS-98-031. 2

At one time, the bulk of LLNL's work was done on four CDC 7600s, and everyone wanted their own 7600

CDC 7600

Apple "7600"

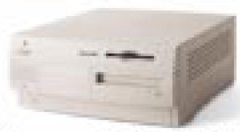

Memory: 16 MB (expandable to 512 MB)
Cycle time: 8.3 ns (upgradeable to 5 ns)
Performance: Livermore Fortran Kernels (LFK)

			Mega-Flops/Sec
Maximum	Rate	=	122.28
Quartile	Q3	=	34.34
Average	Rate	=	27.97
GEOMETRIC	MEAN	=	21.12
Median	Q2	=	18.97
Harmonic	Mean	=	16.54
Quartile	Q1	=	12.84
Minimum	Rate	=	4.81

Cost:
$3,000 (in 1996 dollars)

Memory: 4 MB
Cycle time: 27.5 ns
Performance: LFK
Geometric Mean = 4.0 MFlops

Cost:
$4,000,000 (in 1976 dollars)
$10,000,000 (in 1996 dollars)

NS-00-031. 15

LLNL Computer History

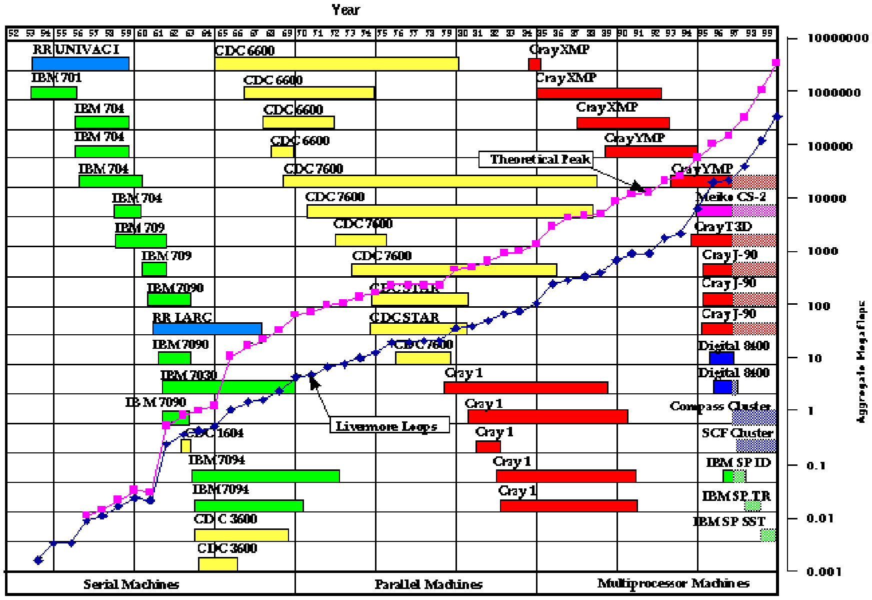

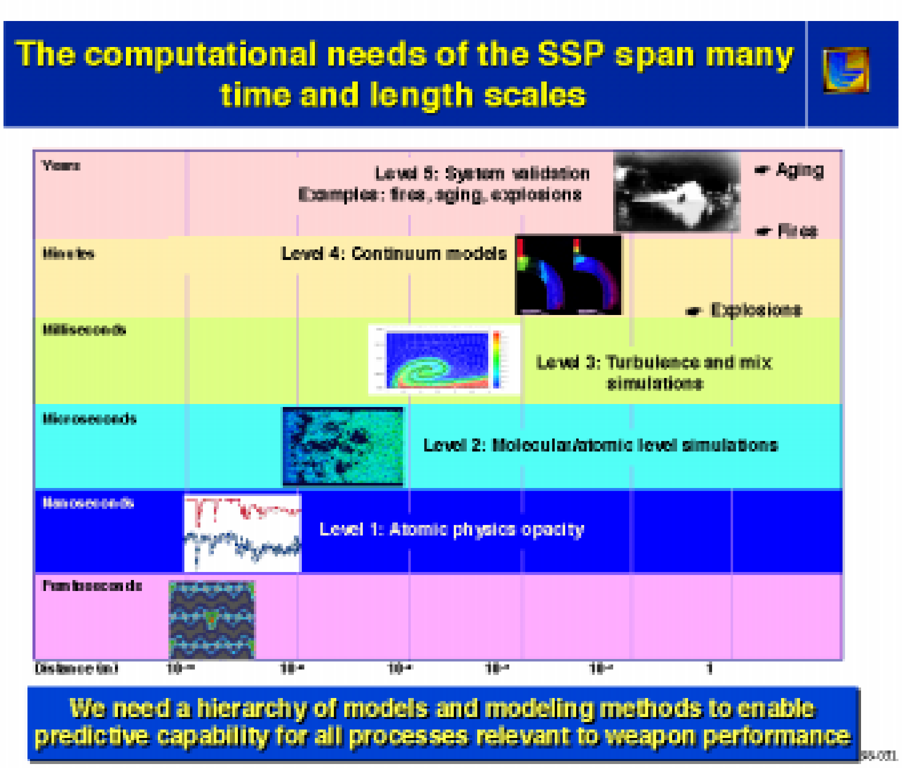

The ASCI applications strategy emphasizes both integrated simulation codes and the sub-grid physics they rely on

Full-Scale Integrated Codes

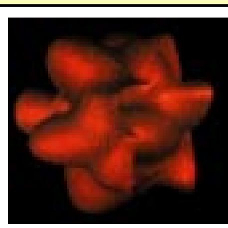

- 3D simulations
- High resolution
- Improved algorithms
  - — Accuracy
  - — Efficiency
  - — Scalability
- Applied mathematics
- Mesh generation
- Visualization
- Validation

Subgrid Models / Zonal Physics

Turbulence | **Materials Models**

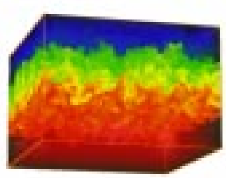

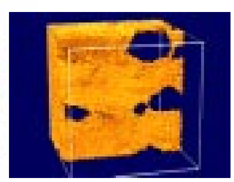

- Direct numerical simulations
- First principle approaches
- Predictive physics models
- Rigorous treatments of physical phenomena

NS-98-031. 1

Our full scale, integrated codes must support tradeoffs between dimensionality, resolution, and detailed physics

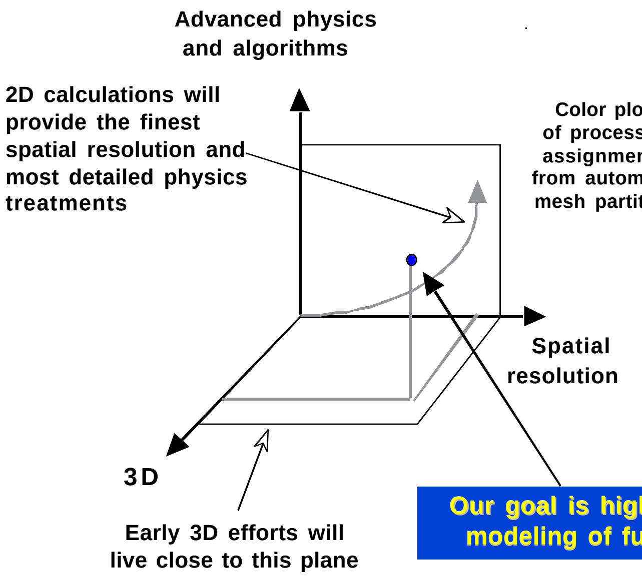

Color plot of processor assignments from automatic mesh partition

This ground shock calculation requires a million zones, 2 million element steps per second, and 256 processors

NS-98-031. 2

SSP structural analysis codes will develop mesh and boundary partitioning for a wide variety of integrated simulations

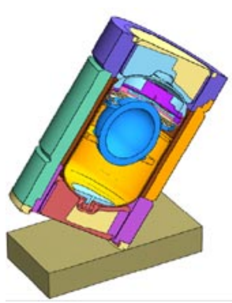

AT-400 shipping container drop test

50,000 elements
27 contact boundary conditions

code: ParaDyn

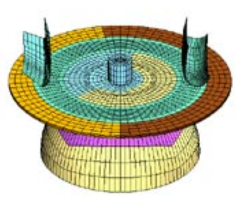

Spin plate metal forming application

4,800 elements
4 contact boundaries

code: ParaDyn

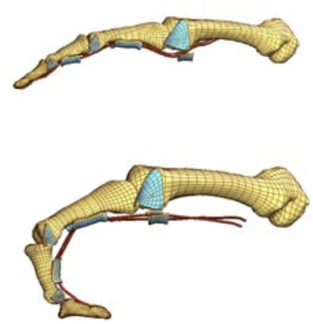

Human index finger under flexion

26,000 nodal points
21,000 8-node brick elements
15 sliding interfaces

code: Nike3D

NS-98-031. 3

Direct 3D numerical simulation code HYDRA provides detailed comparisons between NIF predictions and experiments

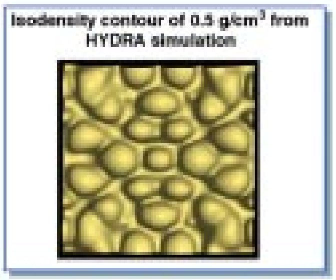

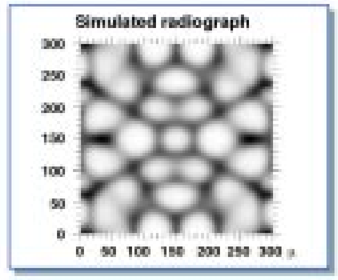

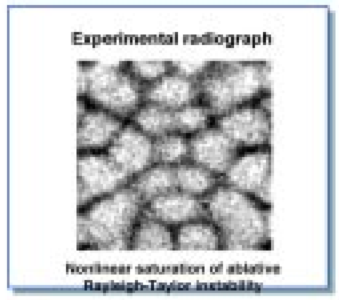

Experimental measurements of 3D multimode surface perturbations on an ablatively driven foil are directly compared to simulations

NS-98-031. 21

Analysis and visualization of terascale data sets places severe demands on all aspects of the computing environment

\>> 1 billion zone physics problem

Physics Application

\> > 100 GB visualization file every 3 minutes (667 Mb/s)

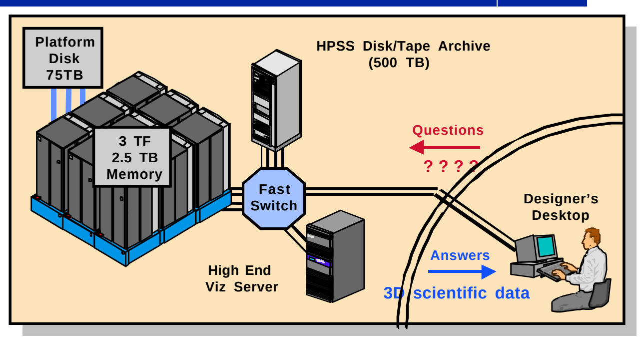

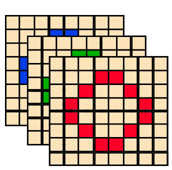

Geometry Engine Software

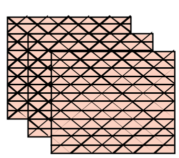

Vis Server

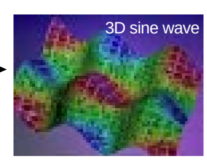

\>> 3GB Iso-surface dataset extracted... (20 MB/s)

\>> and 400 MB geometry data set generated (2.7 MB/s)

\>> 3MB raster image to desktop, fly-around at 10 fps (30 MB/s)

NS-98-031. 1

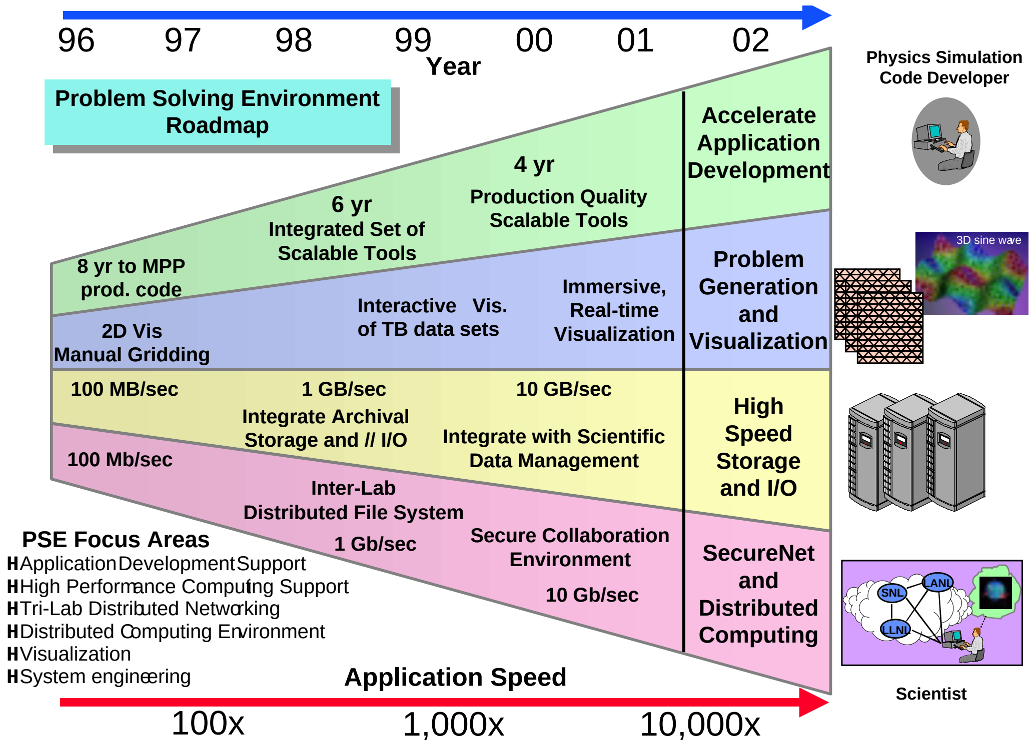
96 97 98 99 00 01 02
Year
Problem Solving Environment Roadmap
Physics Simulation Code Developer
4 yr
Production Quality Scalable Tools
6 yr
Integrated Set of Scalable Tools
8 yr to MPP prod. code
Accelerate Application Development
2D Vis Manual Gridding
Interactive Vis. of TB data sets
Immersive, Real-time Visualization
Problem Generation and Visualization
3D sine wave
100 MB/sec
1 GB/sec
Integrate Archival Storage and // I/O
10 GB/sec
Integrate with Scientific Data Management
High Speed Storage and I/O
100 Mb/sec
Inter-Lab Distributed File System
1 Gb/sec
Secure Collaboration Environment
10 Gb/sec
SecureNet and Distributed Computing
SNL
LANL
LLNL
Scientist
PSE Focus Areas
Application Development Support
High Performance Computing Support
Tri-Lab Distributed Networking
Distributed Computing Environment
Visualization
System engineering
Application Speed
100x
1,000x
10,000x

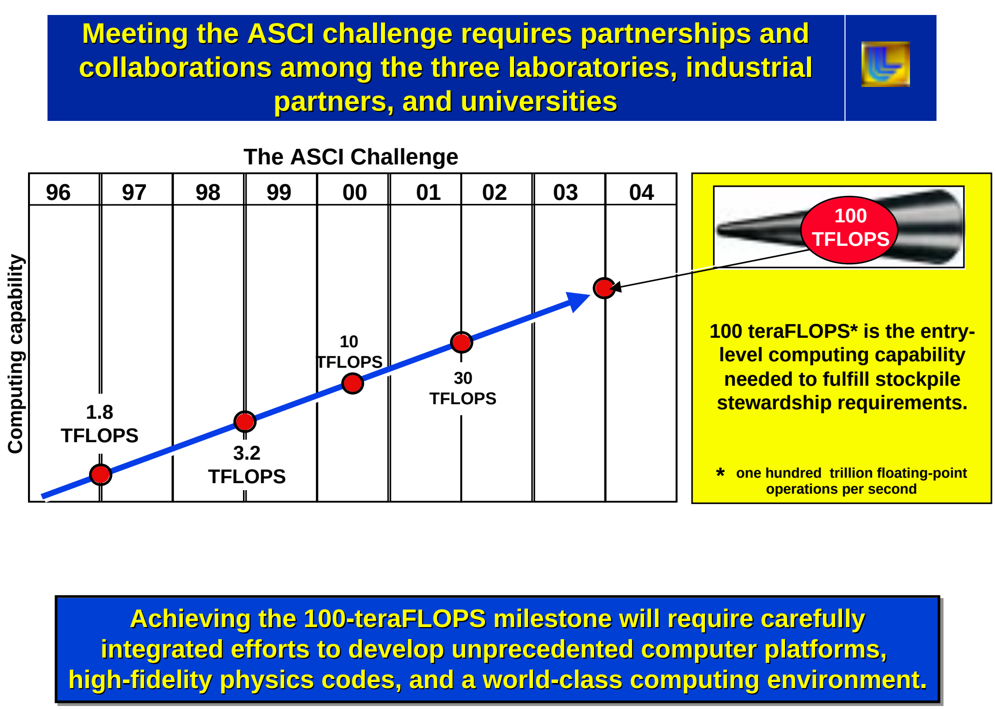
Meeting the ASCI challenge requires partnerships and collaborations among the three laboratories, industrial partners, and universities
The ASCI Challenge
96 97 98 99 00 01 02 03 04
Computing capability
1.8 TFLOPS
3.2 TFLOPS
10 TFLOPS
30 TFLOPS
100 TFLOPS
100 teraFLOPS* is the entry-level computing capability needed to fulfill stockpile stewardship requirements.
* one hundred trillion floating-point operations per second
Achieving the 100-teraFLOPS milestone will require carefully integrated efforts to develop unprecedented computer platforms, high-fidelity physics codes, and a world-class computing environment.

NS-98-031. 3

The ASCI machines are research partnerships with U.S. Industry

Intel
ASCI "Red" Computer
Sandia National Laboratories

SGI/Cray
ASCI "Blue Mountain" Computer
Los Alamos National Laboratory

ASCI "Blue Pacific" Computer
Lawrence Livermore National Laboratory
IBM

The SSP announced strategic academic alliances with five universities

- Stanford University
  - The Center for Integrated Turbulence Simulations
    - William C. Reynolds (wcr@thermo.stanford.edu)
- The University of Chicago
  - Astrophysical Thermonuclear Flashes
    - Robert Rosner (rrosner@oddjob.uchicago.edu)
- The University of Illinois an Champaigne, Urbana
  - Center for Simulation of Advanced Rockets
    - Michael T. Heath (m-heath@uiuc.edu)
- The University of Utah
  - Center for Simulation of Accidental Fires and Explosions
    - David W. Pershing (David.Pershing@dean.eng.utah.edu)
- The California Institute of Technology
  - Facility for Simulating the Dynamic Response of Materials
    - Daniel I. Meiron (dim@ama.caltech.edu)

ASCI is an essential part of the rapidly evolving Stockpile Stewardship Program

- **ASCI provides leading-edge, high-end simulation capabilities to meet weapon certification requirements**
- **ASCI integrates the resources of national laboratories, computer manufacturers, and academic institutions**
  - **national labs focus on application codes and related applied science**
  - **computer manufacturers develop technology and systems for 100 TeraFlops**
  - **Academic institutions research the basic science**

The ASCI codes will need to be continually evaluated against experimental data in the relevant regimes

NS-98-031. 2

We estimate it will take ten years to fully implement the SSP investment

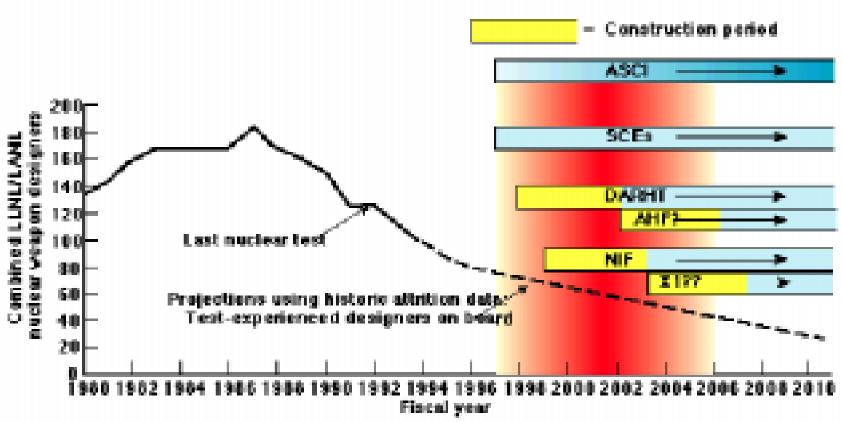

Time is critical because, in the transition, we will need to rely on the judgment of a diminishing number of nuclear-test trained weapon designers

NS-98-031. 3

Our future certification of the stockpile will rely on informed judgments

- Trained, knowledgeable people are required to assess and certify the stockpile
- A deeper understanding of the underlying science is required for practical weapon assessment capabilities
- New computational capabilities are needed to provide the integration formerly done with nuclear tests
- New experimental capabilities are required to provide detailed component level tests and validate the computation tools

Full implementation of SSP is required to sustain nuclear deterrence

NS-98-031. 29

There are many risks inherent in the SSP

- NASA did not accept the judgment of its engineers that the design was unacceptable and,
- As the problems grew in number and severity, NASA minimized them in management briefings and reports.
  - Reports of the Presidential Commission on the Challenger accident

"The contractor did not accept the implications of tests early in the program that the design had a serious and unanticipated flaw."

NS-98-031. 5

XD98-7044 4/20/98

Predictability and Stewardship

1998 Salishan Conference on High-Speed Computing
Gleneden Beach, Oregon
April 21, 1998

Raymond J. Juzaitis
Division Director, Applied Theoretical & Computational Physics (X)
Los Alamos National Laboratory

Applied Theoretical & Computational Physics Division — Los Alamos

---

XD98-7065 4/20/98

Stockpile stewardship and management must be a continuous, fully integrated process

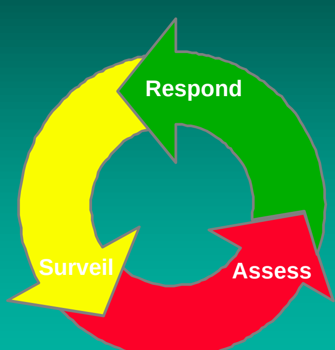

- An enhanced surveillance process, which is capable of diagnosing and predicting aging-related phenomena in stockpile weapons.
- A fundamental understanding of the consequences of the aging and manufacturing processes on weapons performance.
- The capability to repair and/or remanufacture and to revalidate stockpile weapons.

Applied Theoretical & Computational Physics Division — Los Alamos

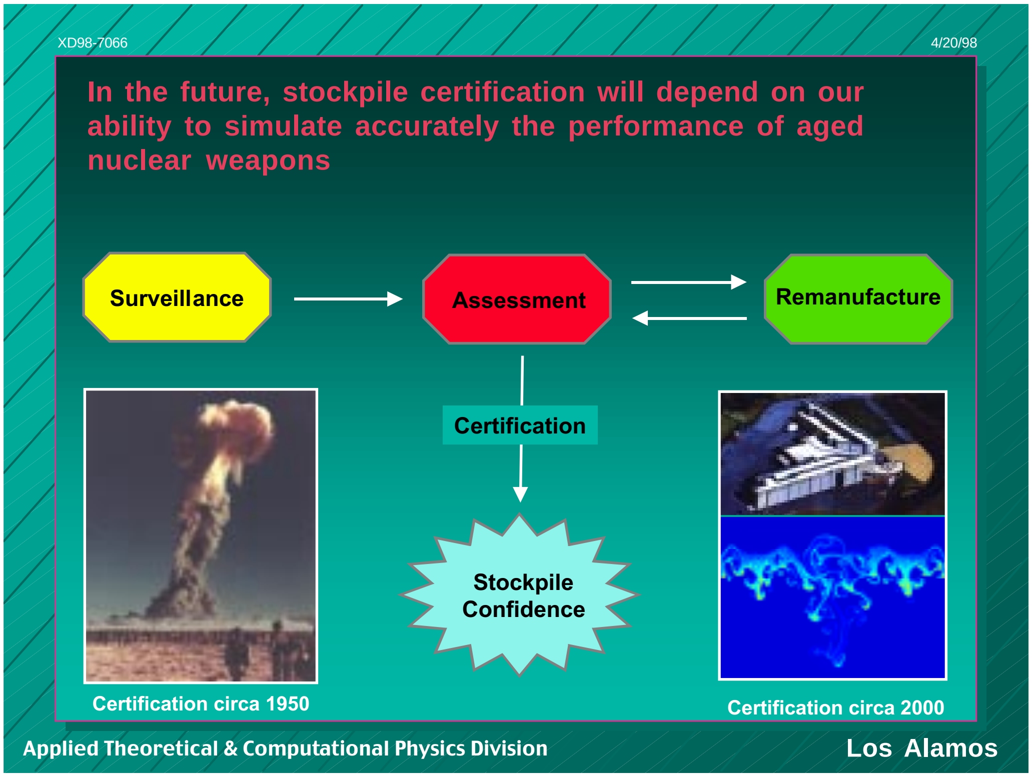
XD98-7066
4/20/98
In the future, stockpile certification will depend on our ability to simulate accurately the performance of aged nuclear weapons
Surveillance
Assessment
Remanufacture
Certification
Stockpile Confidence
Certification circa 1950
Certification circa 2000
Applied Theoretical & Computational Physics Division
Los Alamos


XD98-7067
4/20/98
Successful stockpile stewardship will require surveillance, assessment, and remanufacture
Surveillance—
regular destructive and nondestructive sampling of stockpile systems
Remanufacture—
refurbishment or replacement to meet safety, reliability, and performance requirements
Assessment—
investigation of surveillance observations and evaluation of refurbishment options
Applied Theoretical & Computational Physics Division
Los Alamos

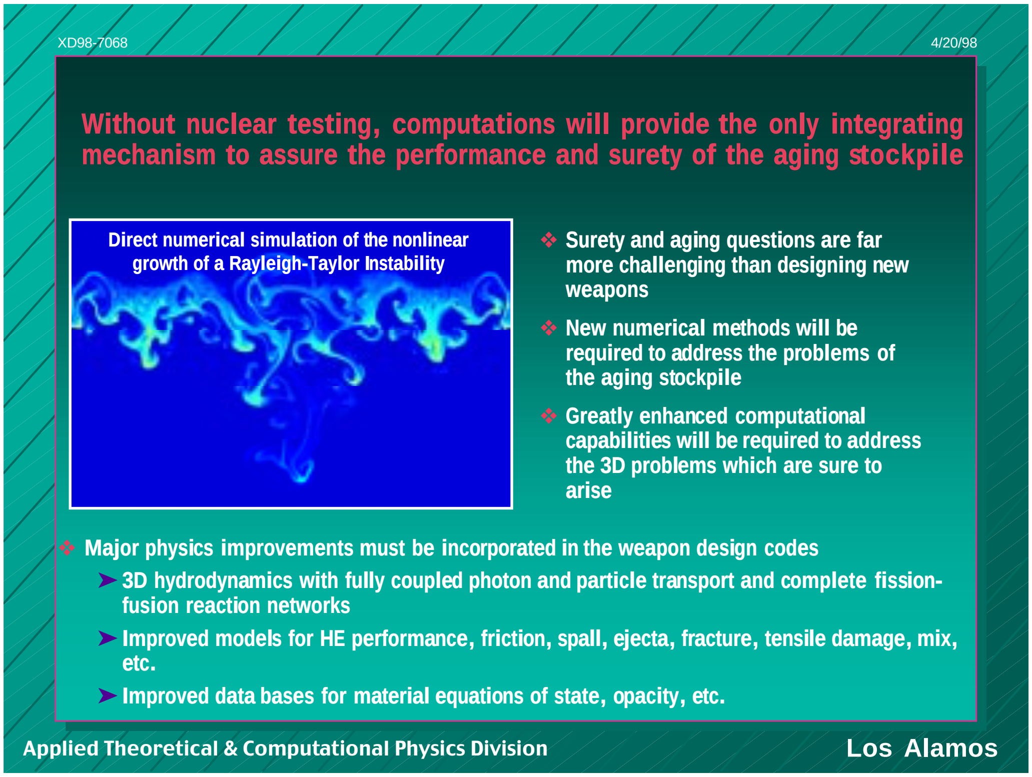
XD98-7068
4/20/98
Without nuclear testing, computations will provide the only integrating mechanism to assure the performance and surety of the aging stockpile
Direct numerical simulation of the nonlinear growth of a Rayleigh-Taylor Instability
Surety and aging questions are far more challenging than designing new weapons
New numerical methods will be required to address the problems of the aging stockpile
Greatly enhanced computational capabilities will be required to address the 3D problems which are sure to arise
Major physics improvements must be incorporated in the weapon design codes
3D hydrodynamics with fully coupled photon and particle transport and complete fission-fusion reaction networks
Improved models for HE performance, friction, spall, ejecta, fracture, tensile damage, mix, etc.
Improved data bases for material equations of state, opacity, etc.
Applied Theoretical & Computational Physics Division
Los Alamos

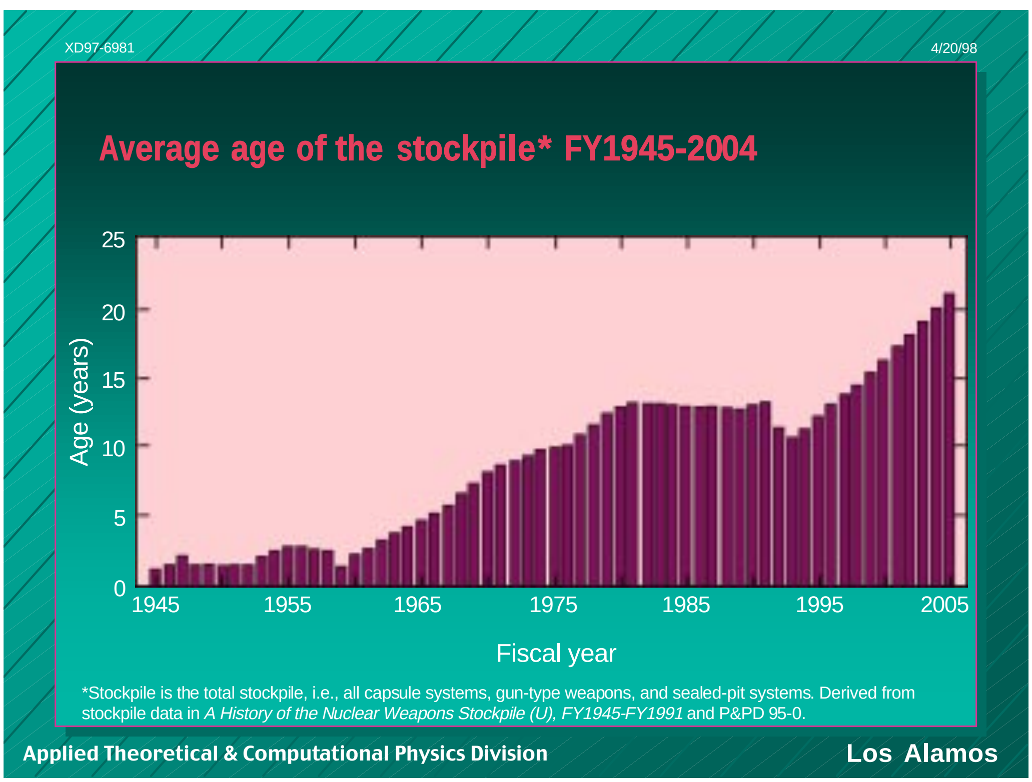
XD97-6981
4/20/98
Average age of the stockpile* FY1945-2004
Age (years)
0
5
10
15
20
25
1945
1955
1965
1975
1985
1995
2005
Fiscal year
*Stockpile is the total stockpile, i.e., all capsule systems, gun-type weapons, and sealed-pit systems. Derived from stockpile data in A History of the Nuclear Weapons Stockpile (U), FY1945-FY1991 and P&PD 95-0.
Applied Theoretical & Computational Physics Division
Los Alamos

CIC-1/96-1296b 7/96 4/20/98

Plutonium aging is one focus of enhanced surveillance

Concern: Structural instability caused by helium bubble nucleation and ingrowth of uranium, neptunium, and americium

Impacts:
- swelling/density
- phase stability
- alloy segregation
- surface structure

Tools:
- transmission electron microscopy
- extended x-ray absorption
- neutron diffraction/backscatter
- Auger electron spectroscopy

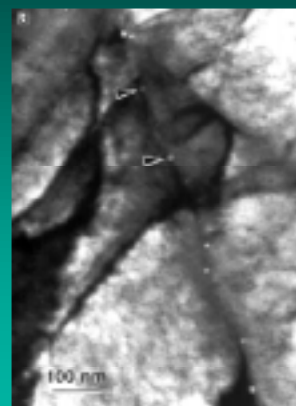

Helium bubbles in aged 20-yr old plutonium

Applied Theoretical & Computational Physics Division — Los Alamos

XD97-6977 4/20/98

Life cycle of manufactured systems

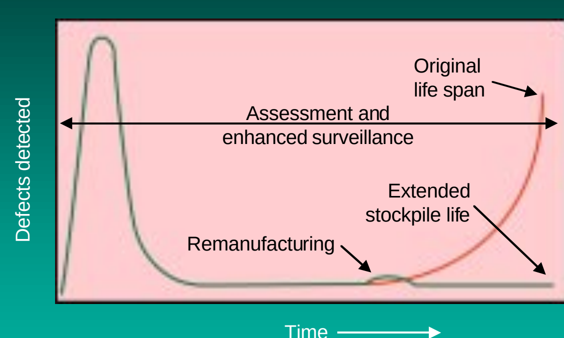

The stewardship program will extend the useful life of U.S. stockpile warheads through enhanced surveillance, assessment, and remanufacturing

Applied Theoretical & Computational Physics Division — Los Alamos

XD98-7069 4/20/98

The Accelerated Strategic Computing Initiative (ASCI) has four objectives

- *Performance:* create predictive simulations of nuclear weapon systems to analyze behavior and assess performance in an environment without nuclear testing.
- *Safety:* predict with high certainty the behavior of full weapon systems in complex accident scenarios.
- *Reliability:* achieve sufficient, validated predictive simulations to extend the lifetime of the stockpile, predict failure mechanisms, and reduce routine maintenance.
- *Remanufacturing and Renewal:* use virtual prototyping and modeling to understand how new production processes and materials impact performance, safety, reliability, and aging issues.

These objectives will be realized by implementing the five ASCI strategies

Applied Theoretical & Computational Physics Division — Los Alamos

---

XD98-7026 4/20/98

Policy community demands predictions of high-order accuracy to effectively address societal problems in a timely and resource-efficient way

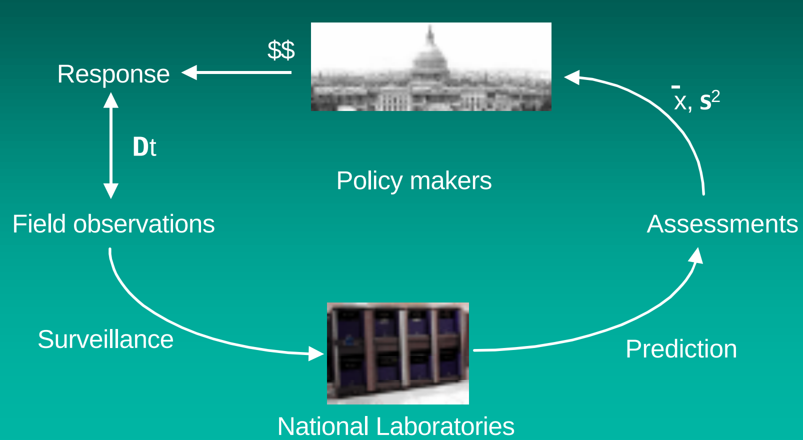

Applied Theoretical & Computational Physics Division — Los Alamos

XD98-7029 4/20/98

A broad class of nationally-important problems requires predicting response of complex systems outside the envelope of controlled experiments and direct reliable observation

- Science-based stewardship of the nuclear weapons stockpile
- Global climate predictions
- Nuclear reactor technology
- Virtual testing aerospace, auto, military technologies
- Natural disaster forecasting
- National infrastructure security

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7064 4/20/98

What does the policy community require to build confidence in "predictive" assessments?

- Rigor in technical work, formality in the process
- Technical integrity
- Some plan or approach to software quality assurance (SQA)
- Peer review
- Quantitative metrics for uncertainty

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7030 4/20/98

When "predictability" is the issue, quantification of error is just as important as solution itself

- Errors almost always imply some acceptance of "cost," in the response to initiating circumstances
- Confidence-building in the policy community
- Theoretical foundations are not always developed
- Increasing number of computed cycles is not always the answer

Applied Theoretical & Computational Physics Division Los Alamos

---

XD98-7023 4/20/98

Euler's taxonomy of "certainty"

- Perceptual certainty (direct experience)
- Demonstrative certainty (deduction, tools of logic)
- Moral certainty (knowledge "by faith")

Demonstrative certainty is attained through the process of scientific inquiry

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7024 4/20/98

The process of scientific inquiry is driven by our need to make sense of the things we see, so that we can predict the future course of events

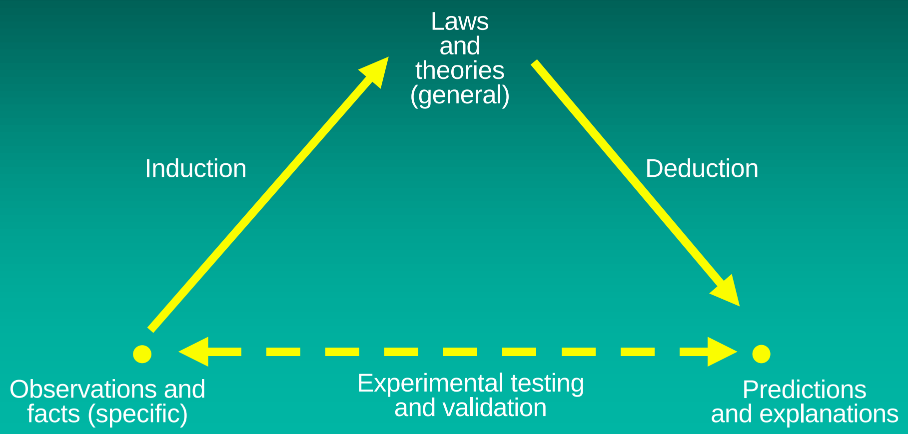

Applied Theoretical & Computational Physics Division — Los Alamos

---

XD98-7028 4/20/98

The nature of scientific prediction ...

- Governed by “rules,” “laws” tested by time
- Explicit schemes (clear and detailed rules, applied by “anyone”
- Publicly available, can be tested independently
- Can postulate unobservable “stats” or events as part of the explanation mechanism
- “Causality” does not imply “predictability,” e.g., celestial mechanics, earthquake prediction

Applied Theoretical & Computational Physics Division — Los Alamos

XD98-7025 4/20/98

Mathematics (modeling) gives a systematic, reliable procedure for generating new truths from old (rules of logical inference)

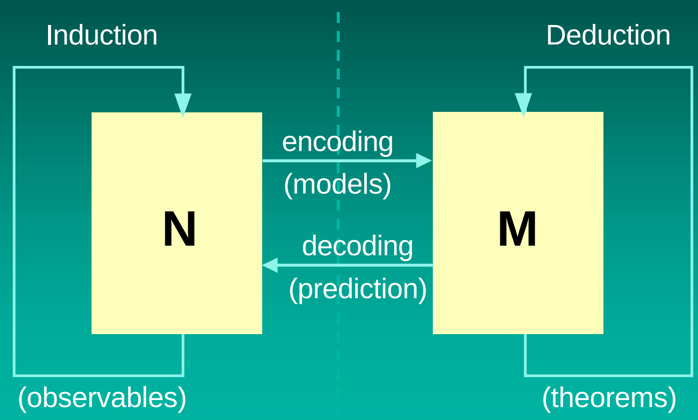

Applied Theoretical & Computational Physics Division Los Alamos

---

XD98-7022 4/20/98

Sources of uncertainty in numerical modeling

- Incomplete understanding of governing natural laws (unknown physics)
- Fundamental data limitations (initial and boundary specifications, physics data, and parameters)
- “Bugs” in computer codes
- Stochastic problem representations
- Multiscale approximating and fidelity across scales (representation of unresolved scales)
- Temporal evolution of dynamic systems

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7020 4/20/98

Strategic computing and simulation for stockpile stewardship and management...meeting "green book" requirements

Strategic Computing and Simulation Program

- **Accelerated Strategic Computing Initiative (ASCI)**
  - ASCI provides the leading edge, high-end simulation capabilities needed to meet weapons assessment and certification requirements without nuclear testing
- **Distance Computing and Distributed Computing for Weapon Simulation (DisCom²)**
  - DisCom² will develop and provide the technology needed to deploy an integrated environment that permits DP labs and plants to access computing (from desktops to TFlops) across thousands of miles
- **Stockpile Computing (SC)**
  - Conduct computing operations, models development, and code maintenance to support execution of the SSMP
- **Numeric Environment for Weapon Simulation (NEWS)**
  - A local computational environment for large numbers of designers to use high-end simulation capabilities to simultaneously address large numbers of stockpile issues
- **Validation an Verification (V²)**
  - Provide the tools, methodologies, and data to ensure high-end simulation capabilities reflect and predict the real world

Applied Theoretical & Computational Physics Division — Los Alamos

---

XD98-7045 4/20/98

Tools and techniques for verification

- Applying "best practices" and standards in software engineering
- Effective debuggers
- Configuration management
- Development of appropriate test suites and ????
- Use of analytical benchmarks
- Peer review, software review process

Applied Theoretical & Computational Physics Division — Los Alamos

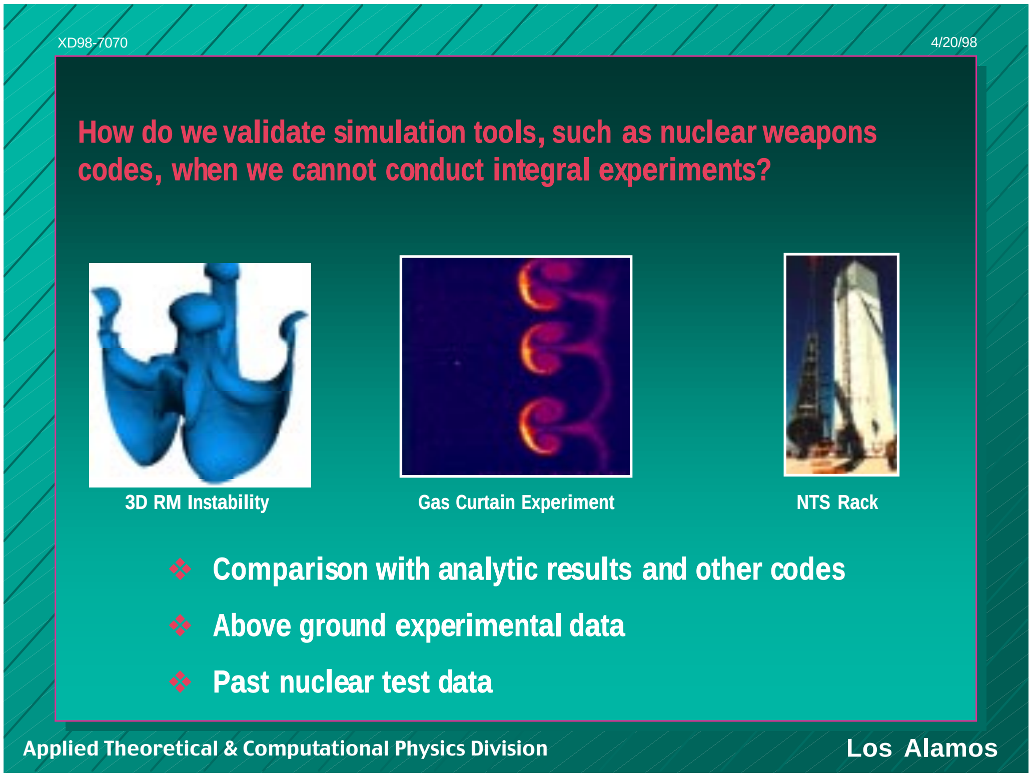
XD98-7070
4/20/98
How do we validate simulation tools, such as nuclear weapons codes, when we cannot conduct integral experiments?
3D RM Instability
Gas Curtain Experiment
NTS Rack
Comparison with analytic results and other codes
Above ground experimental data
Past nuclear test data
Applied Theoretical & Computational Physics Division
Los Alamos

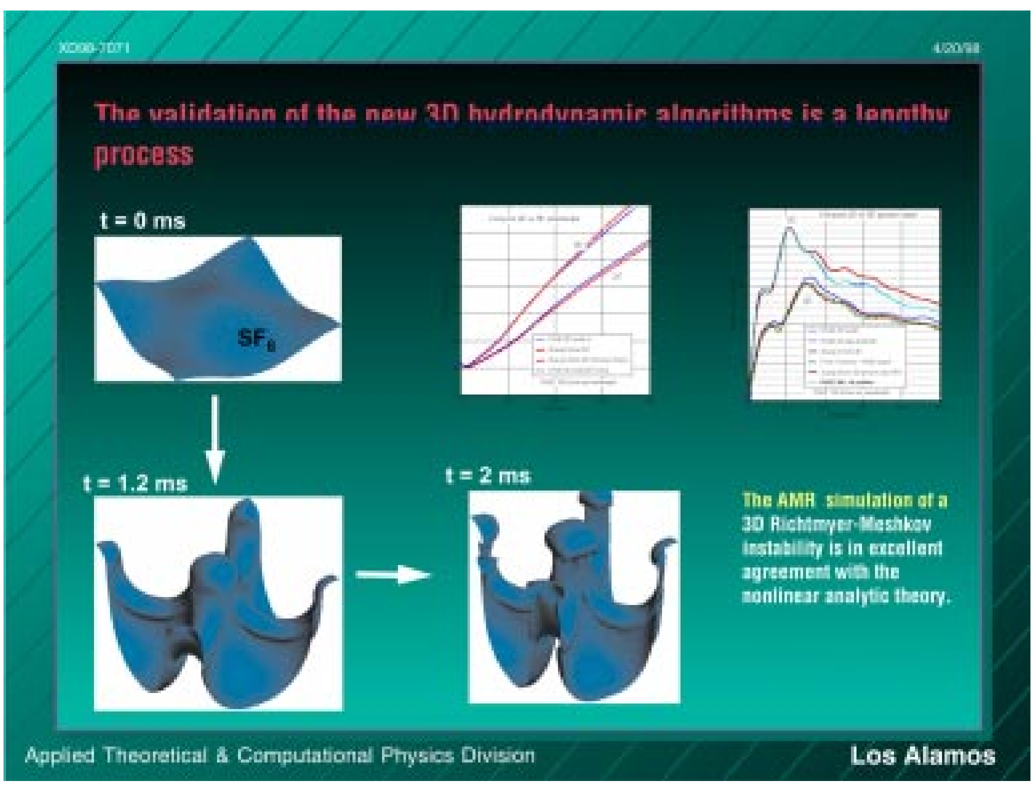
The validation of the new 3D hydrodynamic algorithms is a lengthy process
t = 0 ms
SF6
t = 1.2 ms
t = 2 ms
The AMR simulation of a 3D Richtmyer-Meshkov instability is in excellent agreement with the nonlinear analytic theory.
Applied Theoretical & Computational Physics Division
Los Alamos

XD98-7072 4/20/98

Laboratory experiments can be used to validate the algorithms for direct numerical simulation of hydrodynamic instabilities

- Los Alamos DX-13 shock tube data:
  - Gas curtain experiments.
  - High quality data for onset of turbulence
  - Mach 1.2 shock induced instability
- Experimental innovations include:
  - Laminar gas jet to produce interfaces
  - Laser Rayleigh scattering sheet to observe cross-section with high resolution

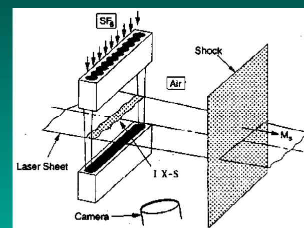

Direct numerical simulation of laboratory data provides confidence in the hydrodynamics that is "necessary but not sufficient"

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7073 4/20/98

The development of a Richtmyer-Meshkov instability is very sensitive to the initial conditions

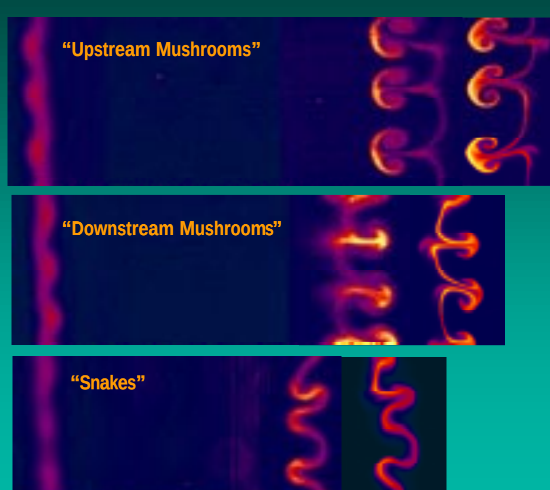

- Density plots
- The Benjamin and Budzinski data are shown on the left
- The RAGE Adaptive Mesh Refinement Eulerian simulations are shown on the right
- The RAGE calculations used initial conditions taken directly from the experimental data

Applied Theoretical & Computational Physics Division Los Alamos

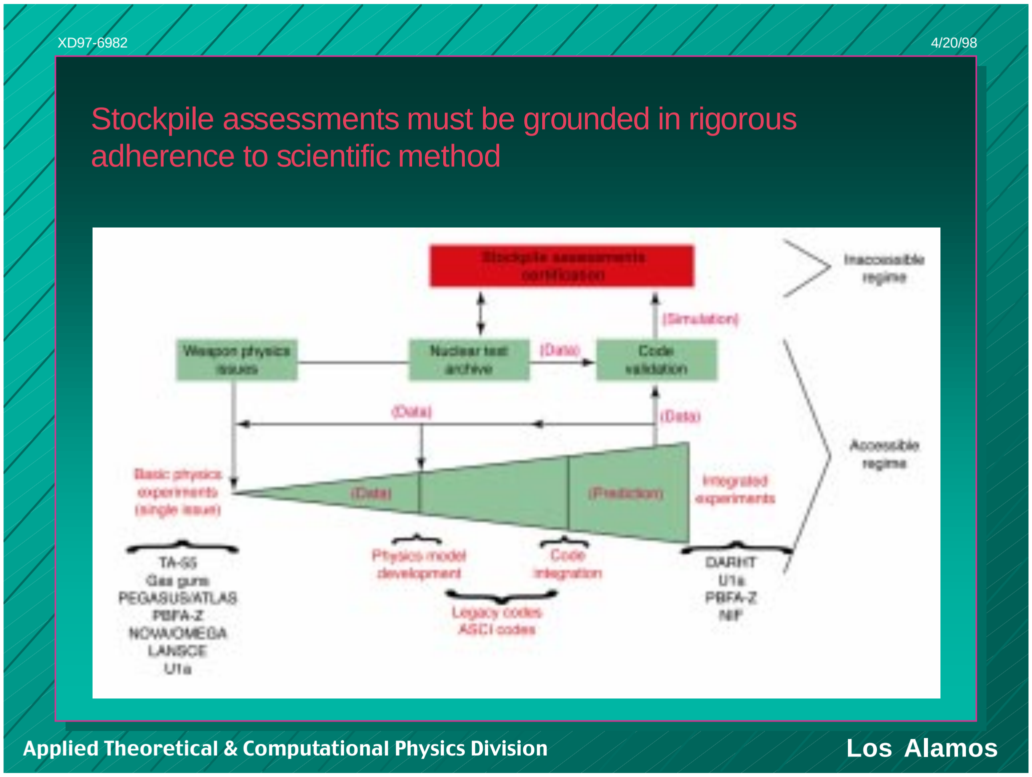
XD97-6982
4/20/98
Stockpile assessments must be grounded in rigorous adherence to scientific method
Stockpile assessments certification
Inaccessible regime
(Simulation)
Weapon physics issues
Nuclear test archive
(Data)
Code validation
(Data)
(Data)
Accessible regime
Basic physics experiments (single issue)
(Data)
(Prediction)
Integrated experiments
TA-55
Gas guns
PEGASUS/ATLAS
PBFA-Z
NOVA/OMEGA
LANSCE
U1a
Physics model development
Code integration
Legacy codes
ASCI codes
DARHT
U1a
PBFA-Z
NIF
Applied Theoretical & Computational Physics Division
Los Alamos

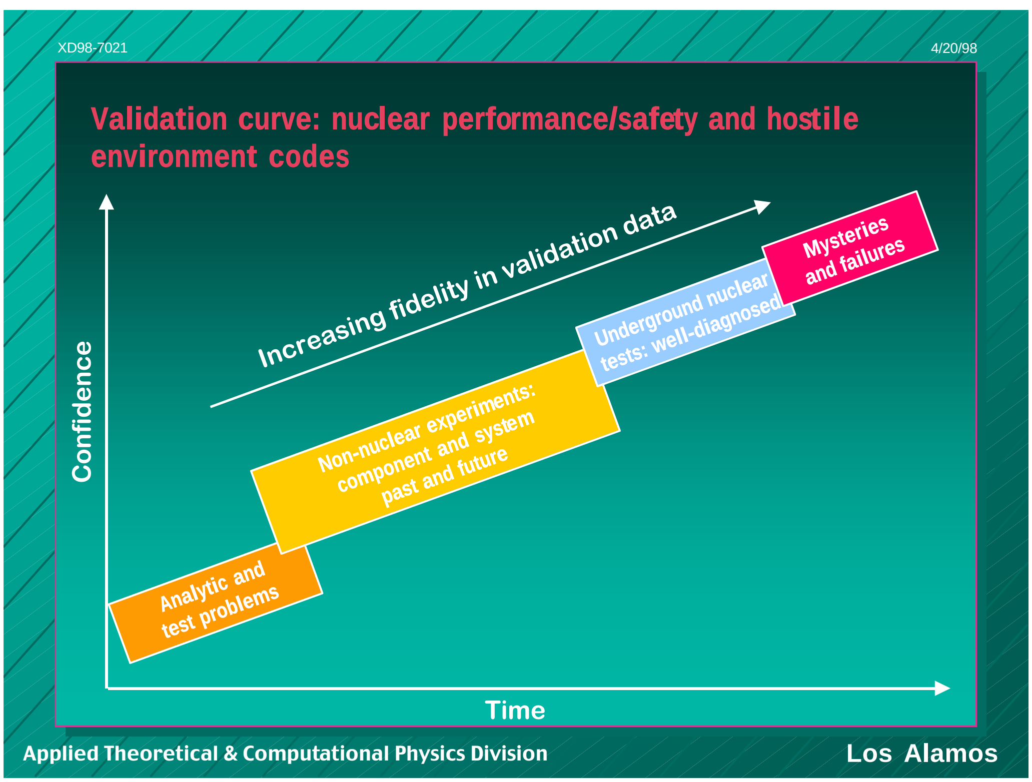
XD98-7021
4/20/98
Validation curve: nuclear performance/safety and hostile environment codes
Increasing fidelity in validation data
Mysteries and failures
Underground nuclear tests: well-diagnosed
Non-nuclear experiments: component and system past and future
Analytic and test problems
Confidence
Time
Applied Theoretical & Computational Physics Division
Los Alamos

XD98-7048 4/20/98

Data sources for validation

- Laboratory-scale and above-ground experiments in the most complete range of thermophysical conditions allowable
- Archive of underground nuclear test data
- Stockpile surveillance data

Establish, with increasing confidence, that a code is predictive. Can you confidently extrapolate beyond the range of experimental data a code was validated with?

Applied Theoretical & Computational Physics Division Los Alamos

---

XD98-7074 4/20/98

Probabilistic risk assessment (PRA) has been employed in major technology sectors to predict behavior of complex systems (e.g., quantify risk)

- 1974 Reactor Safety Study (WASH-1400)
- Major application in nuclear reactor safety analysis (core meltdown, containment releases, consequences)
- Uncertainty analysis employs
  - system-level *fault trees*
  - accident sequence *event trees*
  - appropriate failure data
- Subjective judgements may be used when data is sparse or unavailable (expert elicitation)
- Weapons safety applications—probabilistic weapon response models
- Methodology aids in identification of major contributors to uncertainty; provides insight for resource investment
- May not be a tractable tool in a highly coupled nonlinear dynamic system

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7050 4/20/98

SBSS assessment uncertainties should be quantified within a full probabilistic framework

- Successful prediction must account for uncertainty
- Two-part approach has strong potential for success
  - stochastic PDEs (SPDEs) for forward prediction
  - Markov Chain Monte Carlo (MCMC) and other modern statistical inference methods for inverse prediction

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7049 4/20/98

Predictive science can be regarded as essentially a two-step procedure *(Glimm, Sharp, et.al.)*

- Forward problem (forward prediction)
  - given model - equations, initial data, parameters
  - predict behavior
- Inverse problem (inverse prediction)
  - given observations
  - validate forward solution
  - improve (predict) model
- How can we do this in presence of uncertainties?

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7052 4/20/98

Role of stochastic PDEs

- SPDEs propagate uncertainty in input to determine uncertainty in output
- Stochastic PDEs combine
  - determinism of physical law (PDEs)
  - modeling of uncertainty (statistical analysis)

Applied Theoretical & Computational Physics Division Los Alamos

---

XD98-7054 4/20/98

MCMC is based on three ideas

- Bayesian inference
- Monte Carlo simulation
- Markov chain stochastic processes

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7047 4/20/98

Bayesian inference

- A “model,” *m*, specifies
  - all equations, initial data, parameters used to describe system
- A *prior* distribution is a probability *p(m)*, which represents our current knowledge of *m, including assessment of uncertainties*
- Further observations *O*
  - add information and constrain model
- Constraints expressed as probabilistic consistency
  - probability of *O* given *m* = *p(O|m)*
- *Posterior* distribution *p(m|O)* is updated knowledge of *m*, again including uncertainties

Applied Theoretical & Computational Physics Division Los Alamos

---

XD98-7053 4/20/98

Markov Chain Monte Carlo (MCMC)—what is it for?

- Statistical inference deals with the reality of incomplete information
- It is a tool to answer crucial question
  - How does adding or removing information affect uncertainty in the answer?
  - What is the uncertainty due to missing physics?
  - How do you quantify effect of solution errors?
- MCMC is a method to carry out Bayesian statistical inference in complex systems

Applied Theoretical & Computational Physics Division Los Alamos

XD98-7055 4/20/98

Blending responsible conservatism with a commitment to the methodical development of true predictive capability

- Establish modern baseline for each enduring stockpile weapon system
- Draw a “box” in parameter regime about each system, based on quantified uncertainty
- Don’t certify “outside the box”
- Enhance simulation fidelity (SBSS, ASCI)
- Use verification, validation, and stochastically-based “predictivity tools” to demonstrate an expanded range of design parameter space (go to second bullet)

Applied Theoretical & Computational Physics Division Los Alamos

---

XD98-7056 4/20/98

In summary

- The fundamental science-based stockpile stewardship paradigm joins a larger class of complex simulations which directly impact the policy/technical interface
- Quantified uncertainties in predictions are just as important as the assessments themselves—they are required for ensuring long-term confidence of the policymakers
- Methodical verification and validation of our ASCI codes is time-urgent and critical, but not sufficient for establishing their absolute “predictability”
- Parallel investments must be made in developing the appropriate theoretical foundations for the quantification of uncertainty in a probabilistic framework

Applied Theoretical & Computational Physics Division Los Alamos