LA-13341-C
Conference

UC-705
Issued: July 1997

Proceedings from the Conference on High Speed Computing

High Speed Computing and National Security

April 21-24, 1997

Compiled by
Kathleen P. Hirons
Manuel Vigil
Ralph Carlson

DISCLAIMER

Table of Contents

Conference Program

Monday, April 21, 1997

Keynote Session:

Keynote Address
Andy Heller

Tuesday, April 22, 1997

Session 1: Technologies/National Needs/Policies: Past, Present and Future

National Security Needs: Technologies & Policies
George Cotter, NSA

High Performance Computing's Information Future
John Toole, National Coordination Office for HPCC

Session 2: Information Warfare

A Perspective on the Evolution and Importance of Cryptography in the National Policy Debate
Mike McConnell, Booz, Allen & Hamilton

Update on Information Assurance Efforts and Current Starus
CAPT William Gravell, US Navy

Session 3: Crisis Management/Massive Data Systems

Every Town Hall Needs a Four Gigabit Hard Drive
Paul Fischbeck, Carnegie-Mellon University

Crisis Forecasting
Andy White, LANL

Session 4: Risk Assessment/Vulnerabilities

Computer Crimes 1997
Tom Tauller and Steve Nesbitt, NASA

Banquet

Dreams to Machines and Back
David Urie, Lockheed Skunk Works (Retired)

Wednesday, April 23, 1997

Session 5: Student Session

High Performance Superpiplined Design
Apporv Srivastava, University of Southern California

Parallel Programs from Constraint Specifications
Ajita John, University of Texas, Austin

Session 6: Internet Law/Privacy and Rights of Society

Defining The Rules of the Net: The Visions and Lessons of Self-Governance
Jeffrey Ritter, ECLIPS

Privacy in the Digital Age
Deirdre Mulligan, Center for Democracy and Technology

Thursday, April 24, 1997

Session 7: Challenges to Effective ASCI Programmatic Use of 100 TFLOP/s Systems

The TFLOPS Era is Here
Art Hale, SNLA

ASCI Applications Challenges
Ken Koch, LANL

Day-to-Day Programmatic Usage of 100 TFLOP/s Systems Demands Careful Balance in the Overall Computing Environment
Mark Seager, LLNL

Session 8: New Computing Technologies

Examining ASCI Computing Models
Karl-Heinz Winkler, LANL

Session 9: Future

Use of High Speed Computing in Manufacturing: Godzilla meets King Kong
Gene Meieran, Intel

Proceedings from the Conference on High Speed Computing

High Speed Computing and National Security

April 21-24, 1997

Compiled by
Kathleen P. Hirons
Manuel Vigil
Ralph Carlson

Abstract

This document provides a written compilation of the presentations and viewgraphs from the 1997 Conference on High Speed Computing. "High Speed Computing and National Security", held at Gleneden Beach, Oregon, on April 21 through 24, 1997.

Internet Backbone Observations and Prognostications

- Where is the Internet today and what are the primary forces of change
  - A look from the outside--the uses
  - The Internet today
  - A look from the inside--drivers of increasing bandwidth consumption within existing and new applications including a look at the technical, economic, social and legislative side
- Is NGI enough?
- A Quick Scorecard

The real question

- What is the size of the annual increase in overall bandwidth on the Internet Backbone?
  - 2x?
  - 5x?
  - 10x?
  - 20x?
  - 30x?
- Are there enough things happening to sustain that growth for the 3-5 year period?

General Categories of Usage

- Some big, general categories and some terms of endearment
  - net-balm--(spelling for airports only)
  - net-dregs--(thats de-reg for non-nerds)
  - net-ed-- are you?
  - net-ent
  - net shop--the mall effect
  - net spend--real money
  - net-work

NET-BO

- Super MPP on the network
- limited applications
  - ASCI
  - Large scale simulations
  - Engineering problems
  - Weather
  - War Games
  - LNL/LANL type problems
- requires 100's of gigabits to 10's of terabits/second

NET-Deregulation/Telecom 1994

- Competitive pricing along with performance, attachment and comm. improvements opening up massive growth in usage
- Much faster acceptance of INTERNET and WEB than anticipated
- Countries with protected public or private monopolies are rapidly falling behind the rest of the world. Europe (especially France) and Japan beginning to look like 3rd world
- States are entering the fray. Proposals for statewide networks. Reed Hunt look out

More on Deregulation

- If you own a pipeline there's more money in glass than gas
- USA T3 costs are about 1/6 Japan and result in more than 10x per capita penetration, but even Japan being forced to price low. In France, E3 install is more than 3.5 years of DIGEX USA T3
- Technology that enables usage of existing infrastructure evolving rapidly GB on tel. wires
- This is a business thing guys

NET-Ed(ucation)

- Snow-Rockefeller amendment to 1994 communications act
- Major state initiatives TIFF=$1.5B in Texas
- Will accelerate both # of hours and shift in content from dominant text to video, audio and graphics
- Rapid intro. of agent and x-caster technologies
- Use will be both during and after hours-homework
- Requires 10's of megabits/classroom, millions of classrooms in US

NET-ENT(ertainment) the family that Surfs together

- Moving the Internet from the home office and library to the living room and family room
- Major delivery improvements
  - x-DSL and Cable
  - ETHERNET to the home on the phone
- Moving the action to the big screen
  - Nintendo 64, Sony play station, WEB-TV and other network computers (small n&c)
  - Next generation on TV remotes, voice control and game controllers

More NET-ENT

- More focus on Video, teleconferencing, quality audio and graphics, group activities
- More advertising
- More multi-user interactive Network games not just downloaded games
- many more social (video and telephonic) activities
- Targeted advertising and materials delivery, legal to collect usage data in US but not in Europe

Net-Shopping

- Net commerce comes in the form of both shopping and making the buying decision and actually buying (net-spend) the item
- Both home and corporate models already visible
- Movement to much more complex interactions and much more video and graphic content
- Window shopping at home-the sunny day mall
- True comparison shopping and consumer ed
- Find exactly what you want at the best price
- Really let your fingers do the walking

Net-spending

- 1996 estimates on e-commerce range from $100m to $750m
  - ignores fundamentals of using the net to shop and other ways to buy
  - even the net-spend model will pass $150B on a to go basis by YE2000
- Security, clearing house, privacy issues all have solutions that are surfacing
- Complex corporate net shop and net spend applications already launched
- Home shopping for cars, books and computers

More Net-Spending

- Average time on commercial spending e-commerce WEB sites rising from less than 25 minutes to over 1 hour per visit
- Average number of HTML pages for business commerce sites moving from 2-3 for static informational sites to 15-20 for commercial e-commerce sites
- Many (most) commercial sites also have links to legacy DB systems and each HTML user page can require 10-1000 net packets to legacy systems

Net-Work

- Office in the home and Telecommuting
  - telecommuting becoming increasingly more attractive with more MIPS and more bandwidth
  - MPEG, video-teleconferencing and high speed file access rapidly increasing usefulness of office in the home and opening up interpersonal communications
- X-DSL and ETHERNET to the home will rapidly accelerate movement
- Horizontal business model built around WEB systems evolving with potential for increased outsourcing

Internet [illegible]minology and backgr[illegible]d

- ISP---Internet Service Provider
- NAP--- Network Access Point-a smart MAE
- FIX---Federal Internet Exchange
- MAE---MFS corporation network Access Exchange point-2 major 5 minor MAE sites
  - A MAE does no routing-the only devices connected to the MAE are routers generally provided by ISPs
  - MAE WEST sustaining over 180Mbps (6/96) Connected to NASA Ames FIX WEST
  - MAE EAST sustaining over 380 Mbps (6/96)

The Internet Today in the USA

- Europe 5 years behind the USA. As much digital commercial Bandwidth in Austin Texas as all of continental Europe???130K user YE96
- Even with OC12s and OC48s, insufficient bandwidth for current application load
  - Japanese curious about huge OC48 demand
- 35-40M users today adding >500K per month
- Not enough access points in the backbone
  - Many MAEs not profitable
  - NAPs just starting to emerge for other economic reasons than simply providing bandwidth-higher functionality

Changes in the Internet in the USA

- Shift from switched (circuit and virtual circuit) into routed links for better utilization (Except Bells)
- Emergence of more commercial NAPs and replacement of several MAEs
  - Economic viability of new NAPs broader than just backplane access
  - New NAPs take full responsibility for route optimization and protocol conversion
  - More profitable and efficient than old MAEs
- Emergence of Virtual Corporations / Internet based e-commerce as base of all activities

NGI The Next Generation Internet

- Funded for Fiscal 1998 at $100M
  - Split between DOD, DOE, NASA, NSF, COMMERCE
  - From GIGA to TERA bps
- Expect ongoing funding until at least 2004
- Money for pure Optical solutions, mixed optical-electrical solutions and electrical solutions
- Looking at backbone (TERA-BIT+), NAPs, FIXs,etc (100+ GIGABIT), and LANs (1-10 GIGABIT+)

NGI Goals

- 10 Universities/Natnl Laboratories connected together at 1000 times current maximum available bandwidth limits Approx. 1Tbps+
- 100 Universities/Natnl Laboratories connected together at 100 times current available bandwidth Approximately 100Gbps
- ESnet to be part of NGI
- Ties to Accelerated Strategic Computing Initiative (ASCI) may add source of technology and funding
- 100Gb project will likely get some (D)ARPA funds

What's driving growth of Internet traffic?

- User Community Growth
- Attachment speed
- Local Computing power and storage capacity
- Usage profiles
- E-Commerce
- Java
- Telephonics

What's dri g growth of Internet traffic?

- User Community Growth >450k/month
  - Universities and Research Labs
  - Engineering and Design shops
  - Programming institutions
  - Large businesses
  - Small businesses and home office
  - K-12 and home personal use
  - 1991/ <5M users, <5% home users
  - End 1996/ 35-40M users, about 25% home and home office use. Home/HO could double by 1998

What's driving growth of Internet traffic?

- Attachment speed
  - Growth of ETHERNET in the office (10Mbps)
  - Hi volume / low cost modems
    - 112 baud in 1968
    - 300, 1200, 2400, 4800, 9600 by 1990
    - 14.4 kbaud in early 1990s
    - 28.8 kbaud by 1995
    - 33.6 kbaud and X2(56 kbaud) by early 1997
    - Average Home/HO already over 14.4
    - Next step is 10Mbps or more using x-DSL or asymmetric (DBS), ETHERNET to the home

What's driving growth of Internet traffic?

- Desktop/LAN compute power and storage capacity
  - 1972 time sharing gave capacity of about 40KIPS to the user and less than 40kflops
  - Original PCs about the same capability through about 1981
  - Early workstations from 1 to 3 MIPS by mid 1980s
  - Average workstation over 100mips/mflop by 1996, desktop PCs over 50 mips, 20 mflops average by YE2000 moving to 400+mips, 200+mflops + DSPs--going from 1GByte to >10 Gbytes storage

What's driving growth of Internet traffic?

- Change in usage of the Internet
  - was dominantly e-mail with some FTP
  - the WEB (browsers) has dramatically increased both number of users on the Internet and the amount of data being moved per interaction
  - Starting to see emergence of more parallelism in usage of the Internet
  - E-commerce starting to become real-not just the same thing as before but done on the Internet. Beginning to see first implementations of "Virtual Corporations"

What's driving growth of Internet traffic?

- META Browsers
  - Natural language query
  - Search of probable WEB sites in parallel
  - Can generate thousands to tens of thousands of requests (agents) in parallel
- E-Commerce-Birth of the Virtual Corporation
  - Growing at very rapid rate-Strong enabling technology in security, authentication and clearing house support
  - Complexity of requests and size of returned data rapidly increasing-Large files being moved

What's driving growth of Internet traffic?

- Net-casting
  - Alerts
  - Information services
  - News
- Use of Internet to replace private nets as functions move from clerks to customers
- Will TELNET still exist as it is today
  - 2 packets/character in the backbone
  - build and tear down 2 virtual circuits per character

What's driving growth of Internet traffic?

- JAVA- Will cause rapid increase in average size of data transferred per WEB interaction- move both data and applet-increase of non-text content
- Improving Internet Telephonics
  - Not very good, extremely cheap
  - approx. 60mb/hr
  - Quality improving, getting to tolerable jitter
  - Cost about $0.01 on the Internet to $1.00 on standard long distance and oversea calling
  - Phone companies very unhappy with the situation-must learn to sell bandwidth

What will happen over the next 3-4 years on the Internet

- Massive growth of Infrastructure
- Major shakeout of ISPs-Merger mania
- E-Commerce carrying significant portion of cost burden-but major consumer
- Major inroads into K through 12 education
- 30-50 million Americans having access at their homes using PCs, NCs, Games, and WEB-TVs for information and commerce
- Massive displacement or functional adjustment of long distance carriers

SCORECARD

Function	YE96	YE00	Mult	Wgt	Impct
#of users~40m~120m		*3x*	*1.0*	*3*	
B/e-mail	<500	>60k	120x	0.8	300+
cast-k/u/d	<1	>500	500x	0.2	300+
metaBrs M/d	<.01	>20m/d	2000x	0.15	900+
av. modem	~15k	>1.5m	100x	0.6	180+
av.hrs/day	<.2	>2.5	12x	1.0	36
e-$*M	<$750	>150B	200x	4.0*	800+
tele %U	<0.1%	>25%	250x	1.0	250+
NC usage	<0.1	>15	450x	3.0**	1350‡

3xU*(20%)*5xpages+40 invisible accesses

**increased video/audio/conference gaming

NATIONAL SECURITY AGENCY

NATIONAL SECURITY NEEDS TECHNOLOGIES & POLICIES

THE CONFERENCE ON HIGH SPEED COMPUTING

GEORGE R. COTTER

22 APRIL 1997

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY NEEDS - TECHNOLOGIES & POLICIES

OUTLINE

- STATE OF THE HIGH PERFORMANCE COMPUTING INDUSTRY
- NEEDS OF THE NATIONAL SECURITY COMMUNITY
- PERSPECTIVE ON TECHNOLOGIES
- A STRATEGY FOR SURVIVAL

HIGH PERFORMANCE COMPUTING MAJOR CONCERNS

- **What are the Critical National Security Needs for High Performance Computing?**
- **Can the Industry Satisfy our Current and Potential Requirements? If not,**
  - **What High Performance Computing Technology Program Should the National Security Community Pursue?**
- **Are HPC Industry Market Trends consistent with our National Security Needs? If Not,**
  - **What is our Policy with Respect to Industry; its Survival, Leverage of Systems Design, Technology Cooperation?**
- **What Strategic Relationships within the Federal Government are Critical to Long-Term High Performance Computing Interests of the National Security Community?**

OFFICE OF THE CHIEF SCIENTIST

High End Market Flat; Strong Growth in the Mid Range

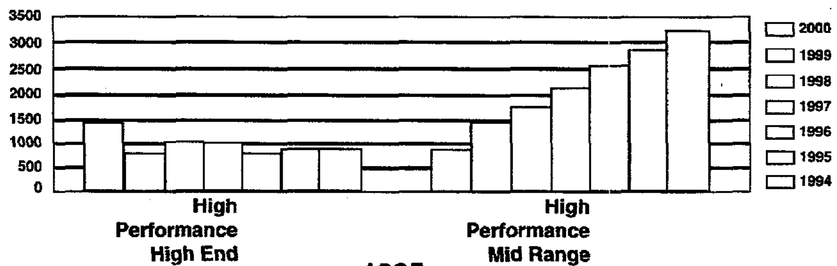

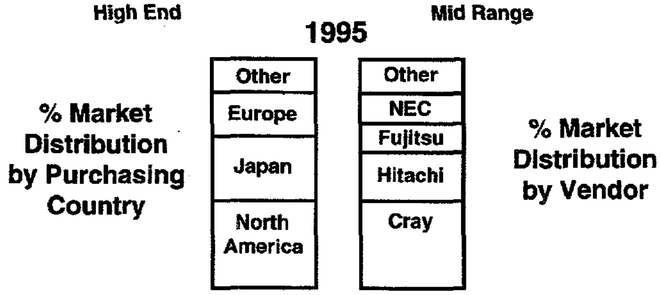

OFFICE OF THE CHIEF SCIENTIST

WORLD-WIDE TECHNICAL HIGH-PERFORMANCE COMPUTER REVENUE BY SEGMENT, 1995-2001 ($M)

	1995	1996	1997	1998	1999	2000	2001	1996-2001 CAGR (%)
High-Performance Midrange	1,343	1,938	2,422	2,835	3,255	3,666	4,129	16.3
Supercomputers	537	659	719	579	591	600	609	-1.6
Technical Parallel Processors	224	415	436	376	386	357	330	-4.5
Total	2,104	3,012	3,577	3,790	4,232	4,623	5,068	11.0

OFFICE OF THE CHIEF SCIENTIST

HPC Industry Has Had High Casualty Rate

Tried and Failed

- Alliant
- BBN
- CDC
- Denelcor
- Elxsi
- ETA
- FPS
- Goodyear
- Multiflow
- Myrias (Canada)
- Prime
- SCS
- Sequent
- SSI

Recent High End Failures

- CCC
- Intel (SSD)
- Kendall Square
- Maspar
- TMC
- NCube

Mergers	Others
SGI-CRI	IBM
HP/Convex	TERA

OFFICE OF THE CHIEF SCIENTIST

Comparison Japanese and U.S. Systems

High End Competitive PVPs

	FUJITSU	NEC	SGI/Cray
Machine	VPP700	SX-4	T90
Clock	7ns	8ns	2.2ns
Peak Performance (PE)	2.2GF	2.0GF	1.8GF
Size Range (PEs)	8 - 256	8 - 32	1 - 32
Memory/Type	SDRAM/ Dist	SSRAM/ Shared	SRAM/ Shared
PE Technology	CMOS/ Custom	COMS/ Custom	ECD/ Custom

High End Competitive MPPs

	HITACHI	NEC	SGI/Cray	IBM
Machine	SR2201	GENJU-3	T3E	SP-2
Clock	6.6ns	13.3ns	3.3ns	7.4ns
Peak Performance	.3GF	.05GF	.6GF	.48GF
(PE)	32 - 2048	8 - 256	16 - 2048	2 - 512 NODES
Size Range (PEs) Memory/Type	DIST/1 GB (PE)	DIST/64 MB (PE)	DIST/ Global/2 GB (PE)	DIST/2 GB (PE)
PE Technology	RISC/ Custom	RISC/ Custom	RISC/ Custom	RISC/ Custom

OFFICE OF THE CHIEF SCIENTIST

HPC INDUSTRY MARKET TRENDS ARE NOT CONSISTENT WITH NATIONAL SECURITY NEEDS

- Continued growth forecast only for mid-range market
- High end revenues are not sufficient to support essential R&D
  - Too many high-end systems chasing too little revenue:

SGI-CRI	IBM	JAPAN	NEW ENTRIES
T90	SP-2	SX-4	(MTA)
T3E		VPP700	(ORIGIN 2000)
		SR2201	
		CENJU-3	

  - <$100M per system (at high end)
- Little upward mobility evident for mid-range users
  - Architectural confusion?
  - Industry uncertainties?
  - ISV support lacking?
  - Adding mid-range increments?
- Decreasing U.S. government-funded HPC system R&D; acquisitions
- Competition for market share is intense
  - The NCAR Domino
  - Europe as the battleground

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY HPC REQUIREMENTS SHOULD SATISFY THE FOLLOWING FIRST PRINCIPLES

- **Importantly affect the nation's leadership in international security activities.**
- **Relate directly to major national security programs or capabilities: weapons development, intelligence, countermeasures.**
- **Meet established standards for critical Defense industrial base.**
- **Ensure clear superiority for U.S. military forces: training, support, operations.**
- **Be a critical element of the nation's economic security and competitiveness.**

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY NEEDS

- **Nuclear Weapons Stockpile Stewardship - DOE ASCI**
- **Imagery Processing and Correlation**
- **Cryptology**
- **High Performance Aircraft Design and Test**
- **Advanced Weaponry Design and Test**
- **Battlespace Modeling and Simulation**
- **Protection of the National Infrastructure**
- **Information Dominance**

Overhead Image - From Stereo Pair
Rosslyn, VA

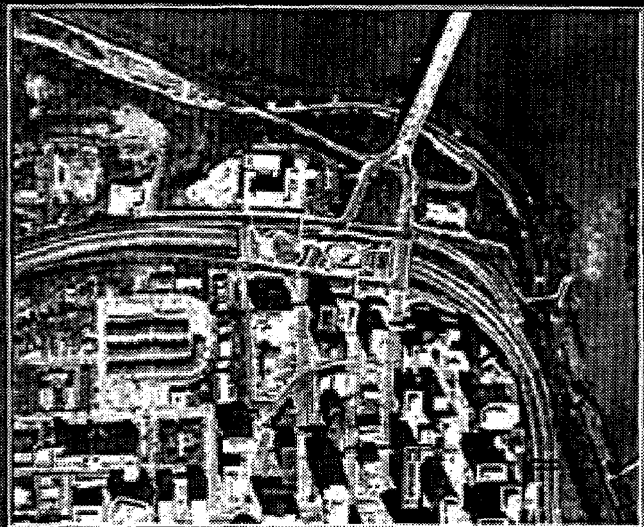

Feature Data
Collected from Stereo Pair

Terrain Data
Autocorrelated From Stereo Pair

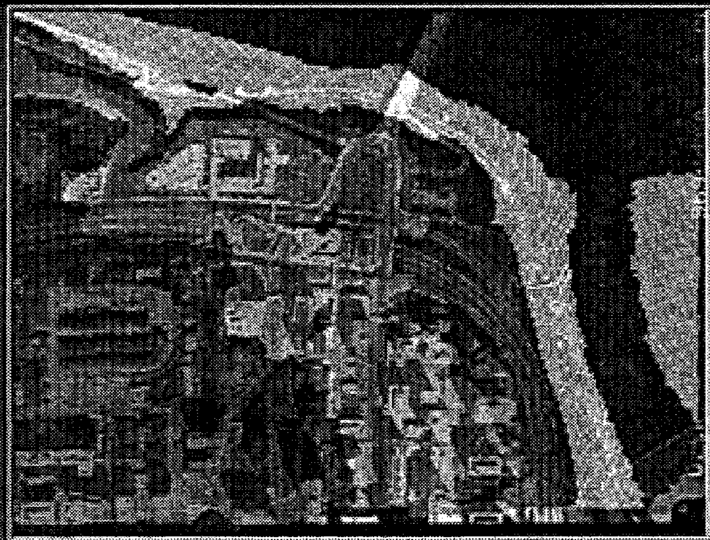

Demonstration of
Single Bounce and Diffraction

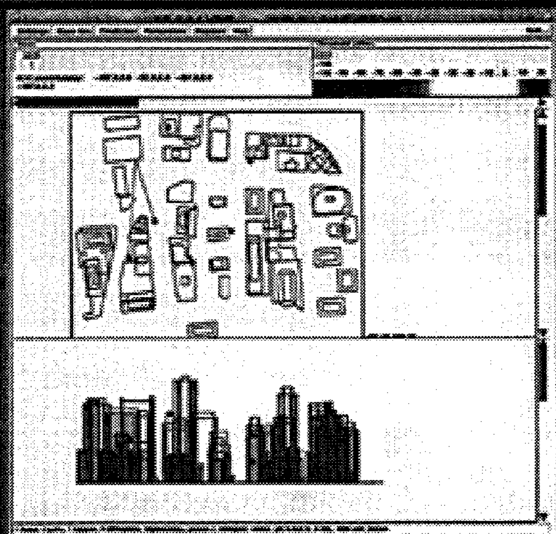

Single Bounce EM Wave Prediction Using Rosslyn Feature Data

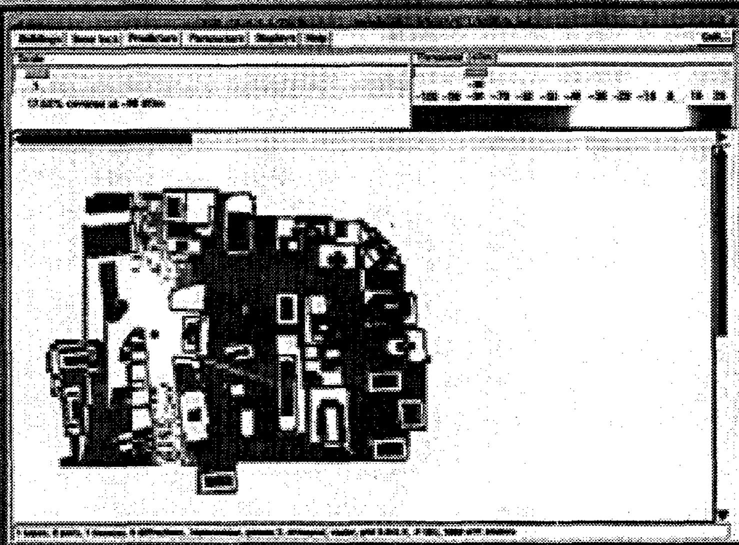

Demonstration of Multiple Bounces and Diffractions

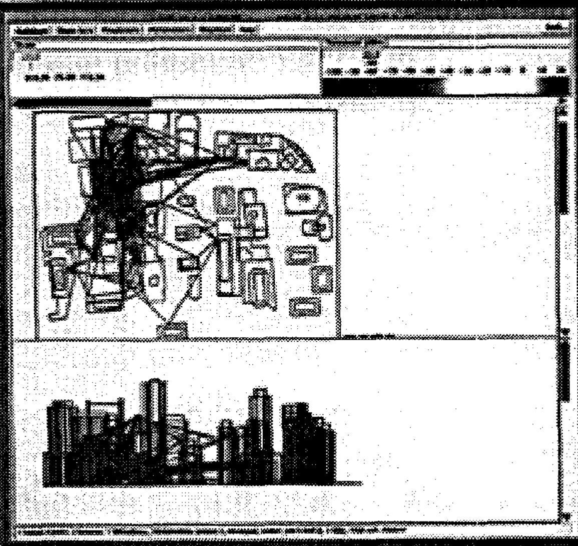

Multi-Bounce EM Wave Prediction Utilizing Rosslyn Feature Data

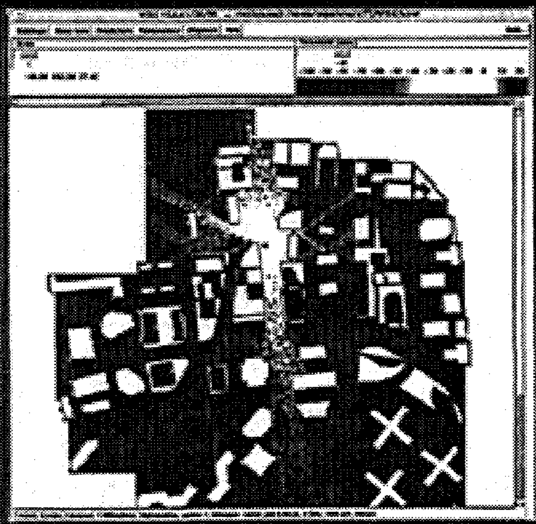

OFFICE OF THE CHIEF SCIENTIST

CRYPTOLOGY

- Cryptanalytic problems are among the most computationally intensive HPC applications to be found
- Relaxation of export controls, while inevitable, exacerbate National Security & Law Enforcement problems

AVERAGE TIME FOR BRUTE-FORCE ATTACK - - WORST CASE

Key Length	T3E	10TF	100TF	Petaflop
40	1.125 Min.	8.289 Sec.	.829 Sec.	.0829 Sec
56	1.86 Mo.	6.852 Days	16.44 Hrs	1.644 Hrs.
64	39.69 Yrs.	4.874 Yrs.	7.182 Mo.	21.54 Days

OFFICE OF THE CHIEF SCIENTIST

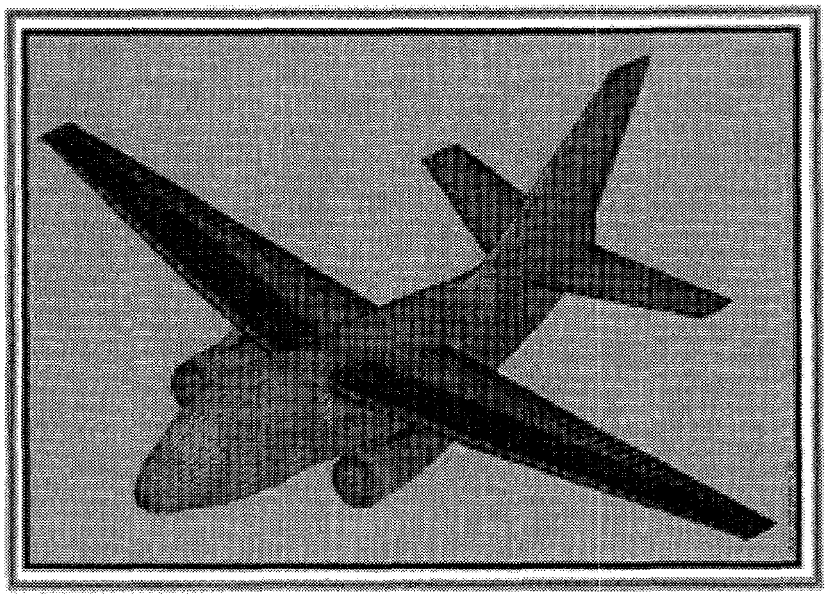

OFFICE OF THE CHIEF SCIENTIST

BATTLEFIELD DOMINANCE

- TARGET DISCRIMINATION
  - Combat Identification
  - Signature Extraction
  - Remote Sensing
  - Real-Time analytic support
- MODELING & SIMULATION
  - Synthetic theater of war
  - Command & Control
  - Mission Rehearsal
  - Decision Support
  - All-source data fusion, retrieval & display
  - Realistic Command Training & Development
- INFORMATION WARFARE
  - Manipulation or massive data sets
  - Dynamic adjustment to information infrastructure attacks
  - Active information operation
  - Integration with other conventional and unconventional warfare

OFFICE OF THE CHIEF SCIENTIST

Performance Spectrum for High End National Security Requirements

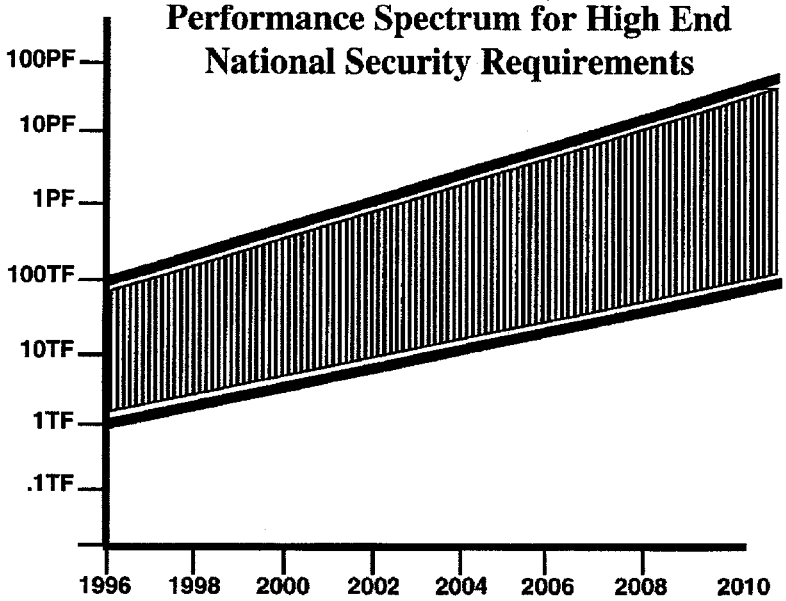

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY POLICY-DRIVEN NEEDS

- **High End HPC is a critical element of the National Security Industrial Base**
- **Essential to continuing world leadership in:**
  - **Weapons system design & production**
  - **Intelligence superiority**
  - **Information age battlefield dominance**
  - **Meeting major Treaty obligations**
  - **Development of new & critical Technologies**

OFFICE OF THE CHIEF SCIENTIST

TECHNICAL SHORTCOMINGS OF TODAY'S HPC SYSTEMS

- Peak Performance
- Deliverable Performance on Many Applications
- Systems Software (Compilers, Languages, O/S Tools)
- Latency Management (HW & SW Multithreading Support)
- Memory Bandwidth
- I/O Bandwidth
- Processor-Memory Speed Gap

OFFICE OF THE CHIEF SCIENTIST

HPCC WORKSHOPS

Apr 92	Pasadena I	Systems Software & Tools
May 93	Pittsburg	HPC Applications
Feb 94	PetaFlops I	Enabling Technologies
Jan 95	Pasadena II	Systems Software & Tools
Aug 95	PetaFlops	Applications
Apr 96	PetaFlops	Architectures
Jun 96	PetaSoft	Systems Software & Tools
Jan 97	PetaFlops	S/S Architecture Mode
Apr 97	PetaFlops	Algorithms

SUMMARY

PARTICIPANTS: Top HPC People from Academia, Industry, Labs, Govt

GOALS: Exhaustive Examination of Technology, Architectures, Systems SW, Algorithms , Applications Showstoppers

CONCLUSIONS: HW - Systems SW Gap Widening
Fresh, Focused Start Needed
- Fundamental Research
- Technologies
- Architectures

OFFICE OF THE CHIEF SCIENTIST

PETAFLOP POINT DESIGNS

- SP Hardware for Astrophysical Particle Simulations
- I-ACOMA: The Illinois Aggressive Cache-Only Memory Architecture for Multiprocessors
- Hybrid Technology Multi-Threaded Architecture
- MORPH: A Flexible Architecture for Executing component Software @ 100 TeraOPS
- Hierarchical Processors-and-Memory Architecture for HPC
- Scalable-Feasible Parallel Computing Implementing Electronic and Optical Interconnections
- Processors in Memory
- Architecture, Algorithms and Application for Future Generation Supercomputers

OFFICE OF THE CHIEF SCIENTIST

HTMT Architecture

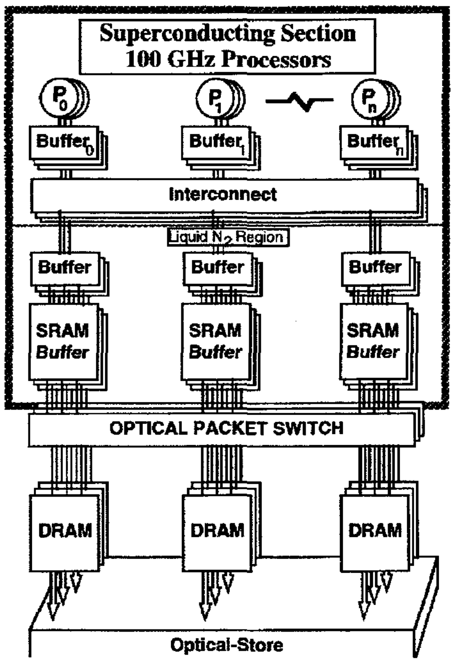

OFFICE OF THE CHIEF SCIENTIST

OVERALL ROADMAP TECHNOLOGY CHARACTERISTICS

Year of First DRAM Shipment Minimum Feature (µ m)	1995 0.35	1998 0.25	2001 0.18	2004 0.13	2007 0.10	2010 0.07	DRIVER
Memory Bits/Chip (DRAM/Flash	64M	256M	1G	4G	16G	64G	D
Logic (High Volume: Microprocessor) Logic Transistors/cm² (packed)	4M	7M	13M	25M	50M	90M	L(µP)
Logic (Low Volume: ASIC) Transistors/cm² (auto layout)	2M	4M	7M	12M	25M	40M	L(A)
Number of Chip I/Os Chip to package (pads) high perf.	900	1350	2000	2600	3600	4800	L, A
Number of Package Pins/Balls							
Microprocessor/controller	512	512	512	512	800	1024	µP
ASIC (high performance)	750	1100	1700	2200	3000	4000	A
Chip Frequency (MHz)							
On-Chip clock, cost-performance	150	200	300	400	500	625	µP
On-chip clock, high performance	300	450	600	800	1000	1100	
Chipt-to-board speed, high performance	150	200	250	300	375	475	L
Chip Size (mm²)							
DRAM	190	280	420	640	960	1400	D
Microprocessor	250	300	360	430	520	620	µP
ASIC	450	660	750	900	1100	1400	A
Maximum Power High performance with heatsink (W)	80	100	120	140	160	180	µP

OFFICE OF THE CHIEF SCIENTIST

MICROPROCESSOR - "THE MORTALITY OF SILICON"

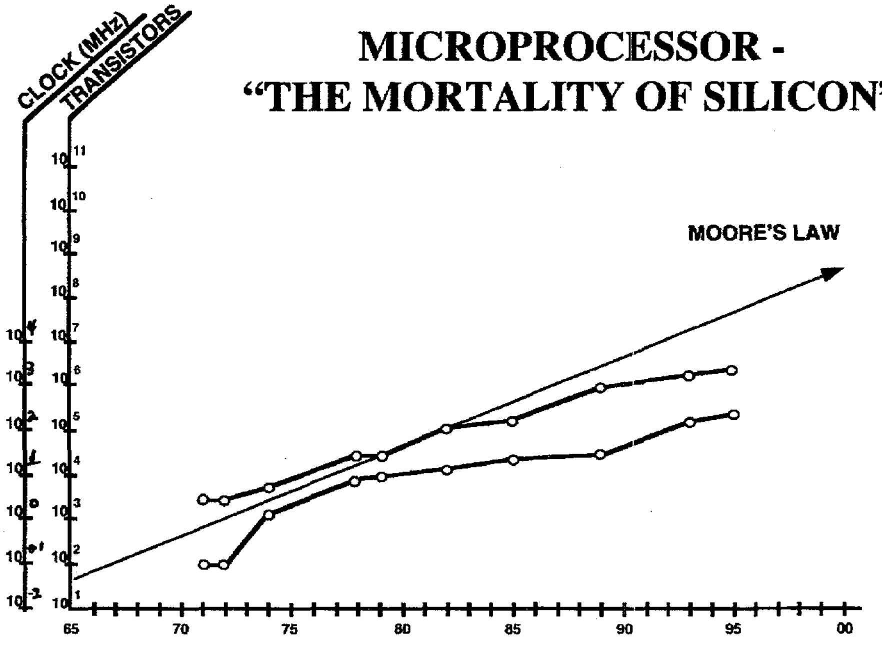

OFFICE OF THE CHIEF SCIENTIST

MICROPROCESSOR - "THE MORTALITY OF SILICON"

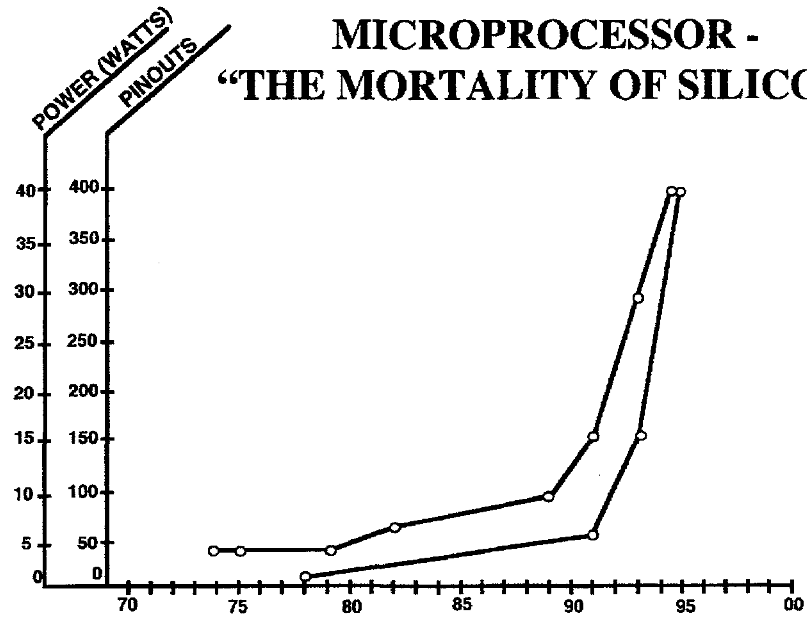

OFFICE OF THE CHIEF SCIENTIST

The Silicon-Based Microelectronics Patient May Be About To Go On Life Support . . . . .

ISSUES:

- Cost of Fabs ~$10B-Year 2000
- Growing Gap Between Processor & Memory Speeds
- Design Complexity - On Chip Integration
- Major Lithography Roadblock:
  - Extended Ultra Violet
  - Ion Beam
  - 1X Xray
- Impact on Tools
- Quantum Effects ~.05 Micron

WORKAROUNDS:

- Increased Speed; Power
- Multi-Chip Modules
- Microprocessor Parallelism
- 3D Arrays
- PIM

OFFICE OF THE CHIEF SCIENTIST

A STRATEGY FOR SURVIVAL

- Reinforce HPC as a Strategic National Asset
- Recognize High End of HPC is Primarily a Niche Market
- Focus HPC Research & Technology Programs in Critical Performance-limiting Areas
- Invest in Promising New Technologies
  - On critical petaflop path
  - Leverage work of others, e.g.:
    - Processing-in-memory
    - Superconducting processors
    - Optical networks
    - Advanced packaging concepts
    - Advanced interconnection techniques
    - Advanced storage concepts
  - Pursue Alternatives to Silicon aggressively
- Pursue Parallel Initiatives in System Software
- Continue, Enhance Long-term Architecture an System Design Initiatives
- Nurture the HPC Research Community

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY HIGH PERFORMANCE COMPUTER PROGRAM

OVERVIEW

- GOALS:
  - Survivability, ultimately leadership, for U.S. HPC industry
  - Aggressive architectural & technological evolution to meet future national security mission needs
- COMPREHENSIVE INTEGRATED LONG-TERM EFFORT TO ENSURE:
  - Availability of highest end systems from U.S. industry for national security missions
  - Essential technology undercarriage for high end computing
  - maximum leverage of work of others
  - In partnership with industry
- STRUCTURE OF THE PROGRAM
  - Research Element
  - Technology Base Element
  - Architecture and System Design Element
  - Development and Engineering in Collaboration with Vendors

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY HIGH PERFORMANCE COMPUTER PROGRAM

RESEARCH ELEMENT

PURPOSE: To challenge the research community to produce basic, enabling breakthroughs in materials, hardware and software, focused on national security system needs; 3-15 years out; Emphasis on proofs-of-concept.

STRATEGY: Maximally leverage work of other research consortia for critical systems software technologies
- Compiler research
- Language research
- Programming interfaces and libraries
- Tools research
- Algorithm research
- Operating systems research

OTHER CRITICAL HIGH-END RESEARCH ISSUES:
- Materials science
- Transition from silicon
- Superconductive processors and memories
- Support for massive multithreading
- Processor-memory imbalance
- Novel component (e.g., WSI, PIM)
- High-Speed interconnects, at several levels
- Advanced computing concepts (quantum, DNA, molecular nanotechnology
- HEC design and development tools

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY HIGH PERFORMANCE COMPUTER PROGRAM

TECHNOLOGY BASE ELEMENT

PURPOSE: To ensure creation of the crucial building blocks of technologies for hign-end computing to support all program activities. Provide proof-of-value as opposed to proof-of-concept

STRATEGY:
- Fund Technology development as opposed to research, focused 10 years out
- Partnerships across academia, designers, technology industries, and the HPC vendor communities
- Creation of a High-End Computing Research and Technology Center, managing System Software Research Consortia, Supporting Mission-Direct Development and Engineering (MDDE) and Architecture and Design Program elements, strongly linked to HEC industry.

CRITICAL TECHNOLOGY ISSUES:
- Compiler infrastructure, modular, extensible
- Tools and performance monitoring infrastructure
- Integrated set of National Security Benchmarks
- Parallel technologies
- Refinements/improvements of vendors' software
- HEC-specific security systems
- Critical hardware technologies (semi, superconductors, 3-d memories, packaging, PIM, high-speed interconnects, optical components, I/O systems and advanced storage architectures

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY HIGH PERFORMANCE COMPUTER PROGRAM

ARCHITECTURE AND SYSTEMS DESIGN ELEMENT

PURPOSE: Full-phase investigation of advanced architectures, systems design concepts and prototypes, employing the best-of-breed Research and Technology Base results; focused ahead of vendor products and product plans, up to 10 years out.

STRATEGY:

- Pursue 3 to 4 point designs competitively developed
- 3-year parallel investigative cycles to proof-of concept
- Teaming of architects, materials scientists, component developers, systems software experts and national security mission users
- Subsequent phases for full prototype development
- Bonding with industry for commercialization
- A PetaFLOP goal

FEATURES: Four Design Categories

Baseline approach (leveraging COTS)
Advanced hardware technology approach
Merged architecture approach
Other Novel approaches

Critical Software Issues

Immediate software development
Focused software development efforts
Layered software architectures
Targeting of key software show-stoppers

OFFICE OF THE CHIEF SCIENTIST

NATIONAL SECURITY HIGH PERFORMANCE COMPUTER PROGRAM

MISSION-DIRECT DEVELOPMENT AND ENGINEERING

PURPOSE: MDDE would be a critical link between a broad National Security HEC program and industry product-cycle efforts. It has the dual objective of flow-through of the results of the Program Research and Technology Base development efforts into the vendors' product lines <u>and</u> the leverage of vendor technologies into higher levels of performance.

STRATEGY:

- Leverage available commodity components with vendors
- Market results of Research & Technology Base programs into vendor product cycle
- Directly influence and accelerate vendors' systems designs
- Task other Program components for HEC technologies the market will not create.

CRITICAL TECHNOLOGY ISSUES:

- Processor access to memory
- High speed interconnects
- Scalable system software
- HEC tools and performance monitors
- Parallel/scalable I/O
- HEC storage system
- SPD interfaces/standards

SUMMARY

- HIGH END OF U.S. HPC INDUSTRY IS IN JEOPARDY
  - Cutthroat international competition
  - ROI does not support R&D
  - Research, advanced technology development is almost non-existent
  - One viable high end U.S. vendor survives
- NATIONAL SECURITY HEC MISSION NEEDS ARE BROAD; EXCEED VENDORS PLANS
  - Government market losing leverage
  - ASCI Program only bright note
  - Government R&D investment insufficient and not well focused.
- TECHNOLOGY FUTURES ARE FUZZY
  - Scalable low-mid range systems remains industry strategy for high end
  - HPC workshop results challenge the view
  - HPC will be first casualty of silicon slowdown
- COLLABORATIVE EFFORT BY NATIONAL SECURITY COMMUNITY APPEARS ESSENTIAL
  - Underpin Industry's efforts
  - Get in front in Research Technology Base & New Architectures
  - Very close collaboration with industry

High Performance Computing's Information Future

The Conference on High Speed Computing
April 22, 1997

John C. Toole
Director
National Coordination Office for
Computing, Information, and Communications

Caution........

- You can't believe everything you hear
  - There actually are fundamental commonalities and differences of technical and policy opinions, particularly in R&D
  - Marketing projections have been one of our most important and useless pieces of information to predict the future
- There's been some significant progress and changes in the last several years...expect more
- Support for science and technology is everyone's SECOND priority
- You, personally, can make a BIG difference
  - Look at people who (still) have a big dream!
  - You can even work for the government and have POSITIVE impact!

Overview

- Where do you find information?
- "Crises" and Opportunities
- Evolution of Federal Strategies
  - Brief History
  - Establishing the Presidential Advisory Committee
- Current Program Embassies
  - Overview
  - Emphasis on High End Computing and Computation, Large Scale Networking (NGI)
- National Security IPT Report
- What's important?

National Coordination Office for Computing, Information, and Communications

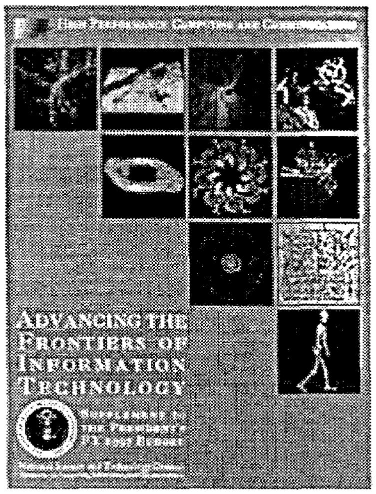

FY 1997 Blue Book

- Supports CCIC and its interagency R&D programs
- Coordinates preparation of planning, budget, and assessment documents
- Central point of CCIC contact to U.S. Congress, Federal agencies, state and local organizations, academia, industry, and the public
- All interagency publications and information available at

http://www.hpcc.gov

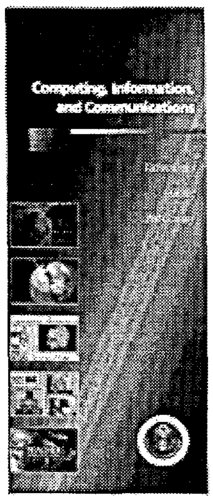

FY 1998 Brochure

"Crises" and Opportunities

- Downsizing, Funding Environments
- Vector Processing
- Parallel Processing
- Scalable Technology
  - Trickle down?
  - Trickle up?
  - Commodity Components
- Markets have seen:
  - Monopolistic Periods
  - Competitive, Startup Periods
  - Shakeouts
- International Competition & Trade

HPC Food Chain

- Basic materials, physics, math, etc.
- Microelectronics
- Component computing technologies (hardware, software)
- Computing system technologies
  - Systems
  - Compilers, debuggers, performance evaluation tools, etc.
- Integrated experimental systems
- A rich world of applications!

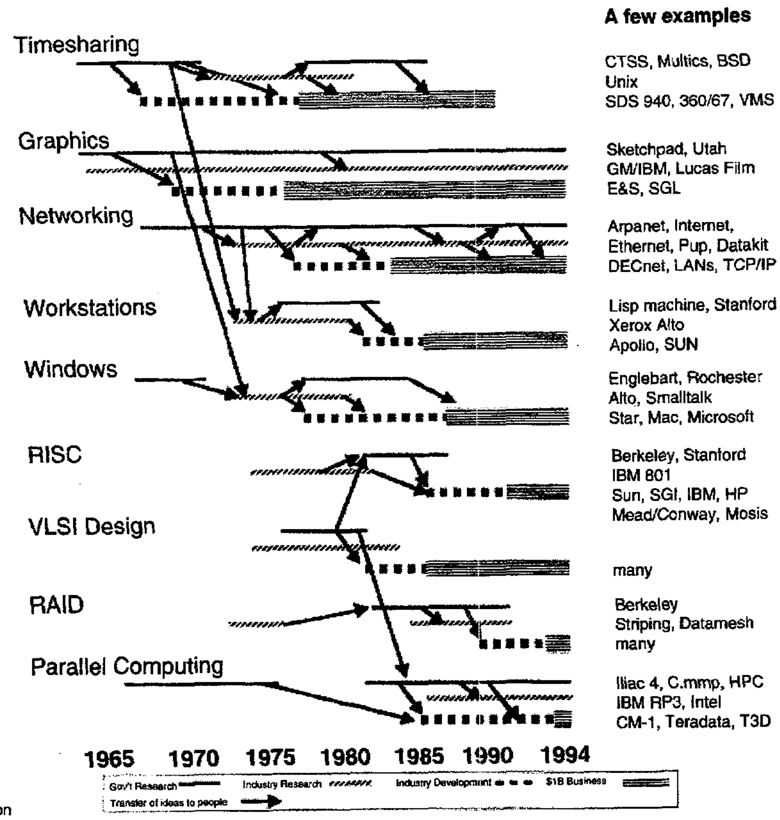

NRC report on "Evolving High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure," 1995, Figure 1.2

The Spectrum of Computer Systems: 1980 vs. Today

A Gap

1980s: Micro Computers — High End Mainframes & Workstations ... Supercomputers

10-100

Today: Micro Processor — Gigaop System — Teraop System

Scalable Machines & Applications with Supporting Software

Related Federal Programs

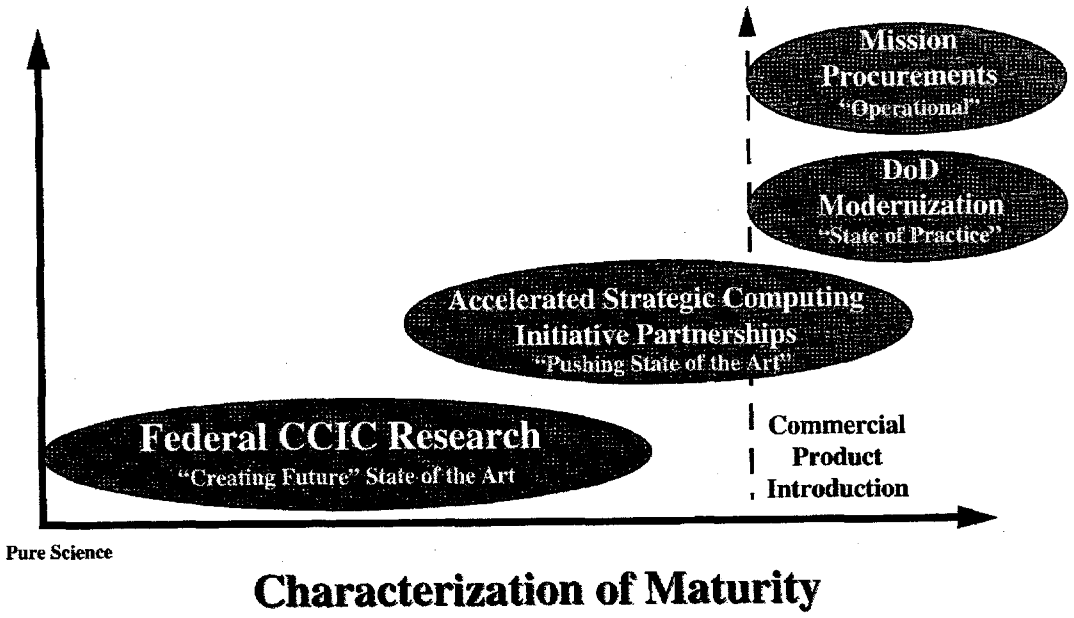

What Should You Know About?

- Advanced Strategic Computing Initiative (ASCI)
- Partnership for Advanced Computational Infrastructure (PACI)
- DoD Modernization Program
- NASA HPCC and IT Evolution
- DARPA Programs and Directions
- Interagency Efforts:
  - CIC R&D Evolution
  - ALL of the component areas - but High End Computing and Computation (HECC) and Next Generation Internet (NGI)

Stand Back and Take a Look...

- Commercial Industry has blossomed across the Information Technology Sector
- Federal investments supported timesharing, computer networks, programming languages, parallel systems, speech and vision systems, ...etc.
- "Today, the economic benefits of all these innovations account for 10 percent of the world's industrial economies, nearly $ 2 trillion a year worldwide. Not a bad return - 100,000 % - on the $ 1 billion ARPA spent on computer research during its early years to fund "half" of these innovations!"

From "What Will Be," by Michael Dertouzos

NSF and DOE (ASCI) Projected Upgrade Paths

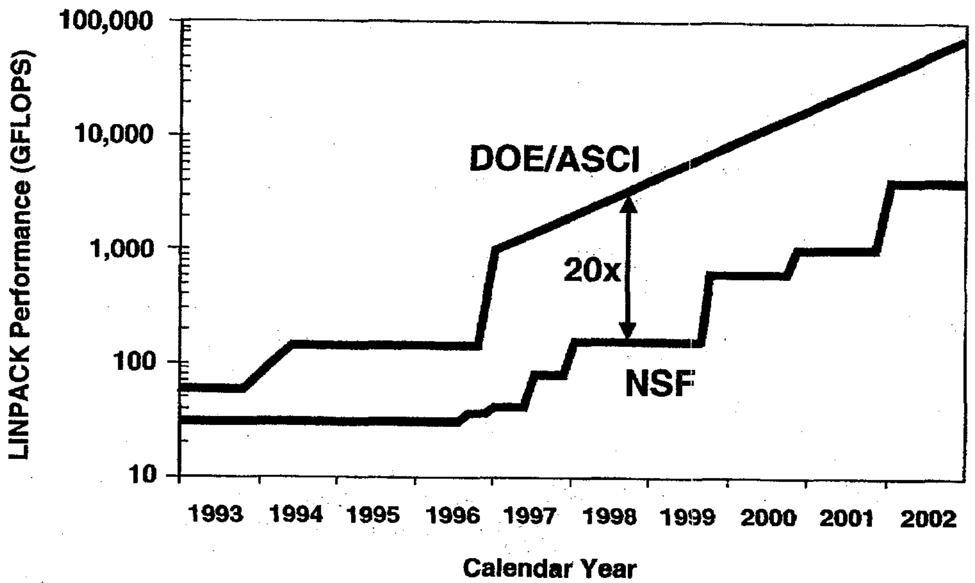

NCSA
University of Illinois at Urbana-Champaign

Processor cycle times

- The reduction in processor cycle time for three "classes" of machines.

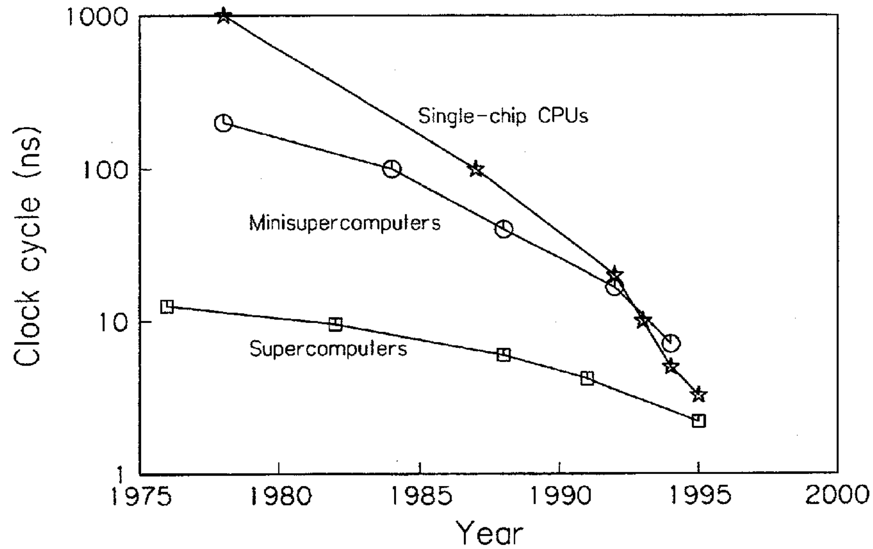

Continuing Exponential Growth in National User Demand

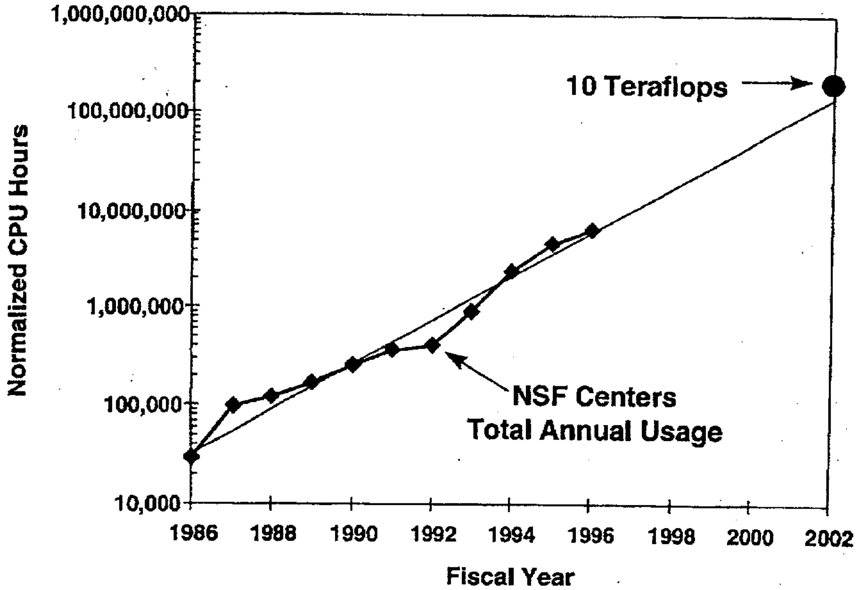

Source: Quantum Research Database

NCSA
University of Illinois at Urbana-Champaign

Organizational Structure

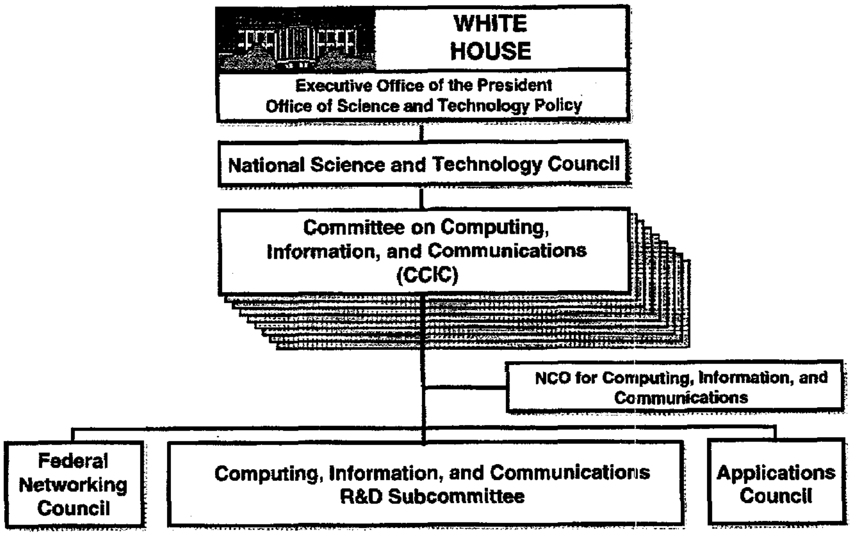

Federal HPCC Program Contributions

- Scalable parallel systems
- Enabling technologies for workstations, distributed systems
- Microkernel operating systems
- Internet networking technology
- Information infrastructure, including early WWW browsers
- Research for Digital Libraries
- Gigabit testbeds
- Supercomputer centers
- Grand Challenge Applications
- National Challenge Applications
- Mission applications: e.g., national security, medicine, environment, and education

History of the Federal HPCC Program

- The Federal HPCC Program:
  - Chartered by Congress FY 1992 through FY 1996 with the High Performance Computing Act of 1991
  - Focused on:
    - High Performance Computing Systems
    - Advanced Software Technology and Algorithms
    - National Research and Educational Network
    - Information Infrastructure Technology and Applications
    - Basic Research and Human Resources
  - Was coordinated through the High Performance Computing, Communications, and Information Technology (HPCCIT) Subcommittee and NCO
  - Administration did not seek reauthorization, but wants to continue a strong program as part of NSTC and agencies' programs

Presidential Advisory Committee

- President Clinton signed an Executive Order in February
- Establishes the Advisory Committee on High Performance Computing and Communications, Information Technology, and the Next Generation Internet:
  - Non-federal members; community representatives from research, education, and libraries; network providers; and critical industries
  - Provides the National Science and Technology Council, through the Director, OSTP, with advice and information on high-performance computing and communications, information technology, and the Next Generation Internet
  - Terminates two years from date of Executive Order, unless extended by the President
- 21 Members have been appointed, including two Co-Chairs
  - Bill Joy, Sun Microsystems
  - Ken Kennedy, Rice University

Presidential Advisory Committee Members

Co-Chairs: Bill Joy, *Sun Microsystems*
Ken Kennedy, *Rice University*

Eric A. Benhamou, *3Com Corporation*

Vinton Cerf, *MCI Communications*

Ching-chih Chen, *Simmons College*

David Cooper, *Lawrence Livermore National Laboratory*

Steven D. Dorfman, *Hughes Telecomunications and Space Company*

Robert Ewald, *Cray Research*

David J. Farber, *University of Pennsylvania*

Sherrilynne S. Fuller, *University of Washington*

Hector Garcia-Molina, *Stanford University*

Susan L. Graham, *University of California, Berkeley*

James N. Gray, *Microsoft Research*

W. Daniel Hillis, *Walt Disney Imagineering*

David C. Nagel, *AT&T Labs*

Raj Reddy, *Carnegie Mellon University*

Edward H. Shortliffe, *Stanford University School of Medicine*

Larry Smarr, *National Center for Supercomputing Applications*

Leslie Vadasz, *Intel Corporation*

Andrew J. Viterbi, *QUALCOMM Incorporated*

Steven J. Wallach, *Hewlett-Packard, Convex Technology Center*

Advisory Committee Update

- Logistics:
  - First meeting held February 27-28, 1997
  - Next meeting June 24-26 1997
    - June 26 will be subcommittee meetings
  - October meeting anticipated
- Subcommittee formation:
  - NGI (Reddy, Gray)
  - High-End (Cooper, Graham)
  - Broad-Based (Nagel, Shortliffe)

CCIC Strategic Planning - Phase II

(February 1997)

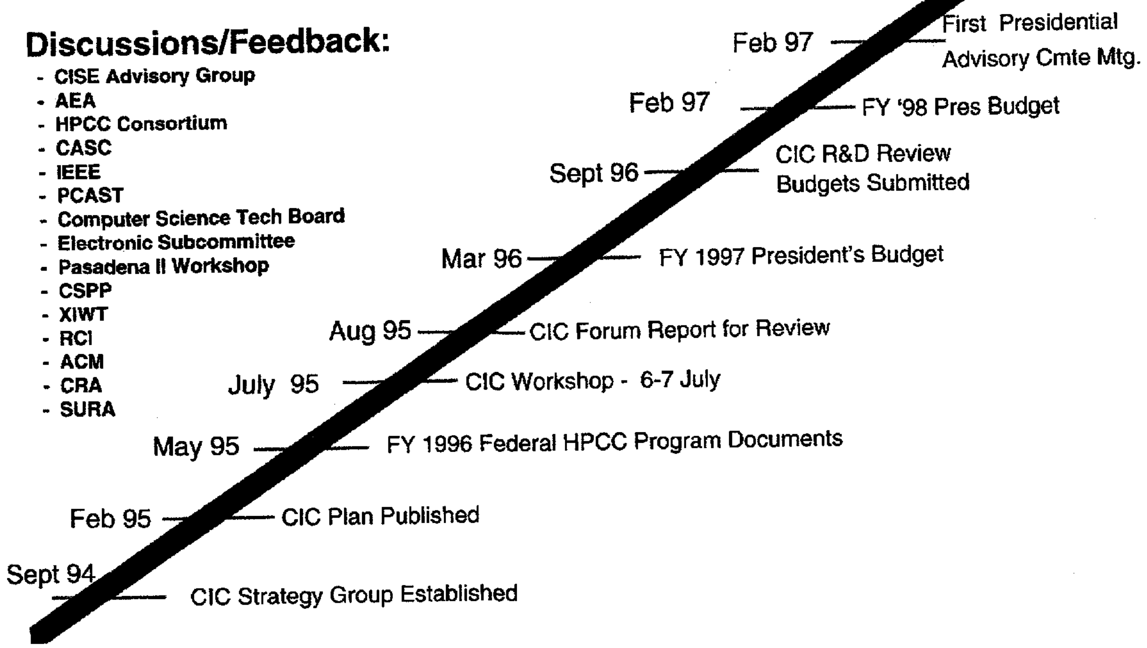

Organizational Structure

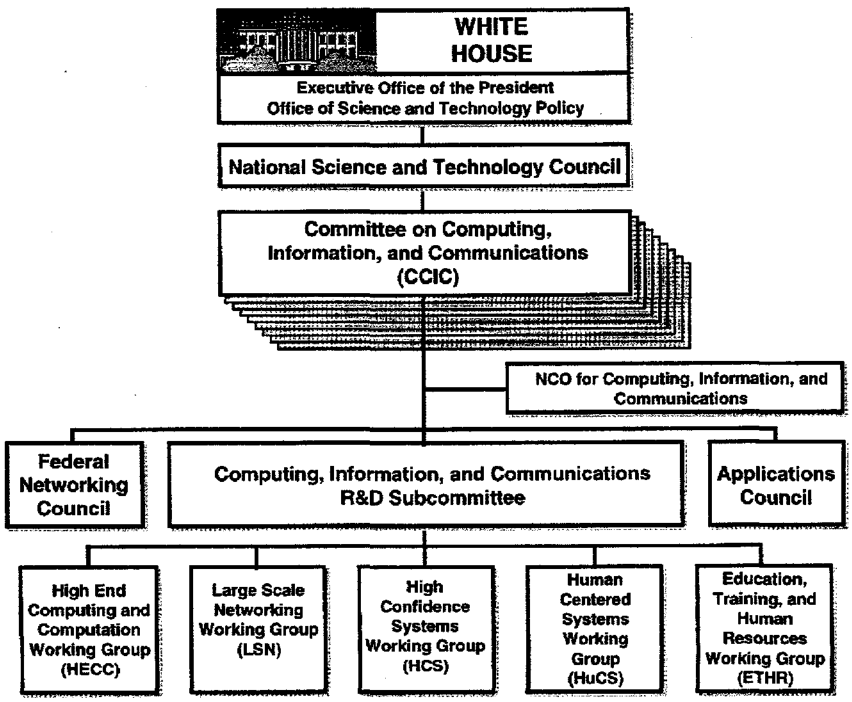

New CCIC R&D Structure

- Broader collaborative R&D investments
- Five Program Component Areas (PCA)
  - High End Computing and Computation
  - Large Scale Networking
  - High Confidence Systems
  - Human Centered Systems
  - Education, Training, and Human Resources
- Each PCA will:
  - Span areas in which multiple agencies have interests/programs
  - Include hardware, software, algorithms, and applications
  - Focus on specific research goals, ensure adequate investments, and maintain necessary budget visibility
- Technology R&D may span multiple PCAs

Participating Agencies

- Defense Advanced Research Projects Agency (DARPA)
- National Science Foundation (NSF)
- Department of Energy (DOE)
- National Aeronautics and Space Administration (NASA)
- National Institutes of Health (NIH)
- National Security Agency (NSA)
- National Institute of Standards and Technology (NIST)
- Department of Education (ED)
- Department of Veterans Affairs (VA)
- National Oceanic and Atmospheric Administration (NOAA)
- Environmental Protection Agency (EPA)
- Agency for Health Care Policy and Research (AHCPR)

Draft FY 1998 HPCC Crosscut Budget Request (Dollars in Millions)

Agency	HECC	LSN	HCS	HuCS	ETHR	TOTAL
DARPA	84.8	89.2	9.4	137.9		321.3
NSF	132.9	79.2	0.9	60.2	21.0	294.2
DOE	90.8	48.8		9.9	3.0	152.5
NASA	90.9	25.1	2.8	4.9	4.7	128.4
NIH	23.7	28.2	4.1	29.3	6.4	91.7
NSA	26.4	2.2	7.2			35.8
NIST	4.0	5.5	3.4	13.6		26.5
VA		7.4	5.4	9.2		22.0
ED					12.0	12.0
NOAA	4.3	2.7		0.5		7.5
EPA	5.4			0.8		6.2
AHCPR				5.5		5.5
TOTAL	463.2	288.3+	33.2	271.8	47.1	1,103.6*

+ The requested FY 1988 LSN budget includes funds for the Next Generation Internet (NGI) Initiative. It also reflects the transition of DARPA's mature technology research from networking development to networking applications. For example, DARPA's FY 1998 allocation for HuCS includes $21M transferred from the FY 1997 networking research budget.

* These totals vary slightly from the President's HPCC Budget. For example, funding for the Department of Transportation, one of the candidate agencies for participation in CIC R&D activities, is not included.

Example Research Agenda for:

High End Computing and Computation

- System software and tools
- Application development environments
- Fast, efficient algorithms for simulation, modeling and visualization
- System architectures
- Device technologies
- Interconnection technologies
- I/O, and multi-level data storage
- Laboratory demonstration prototypes
- Advanced simulation of physical phenomena and other grand challenge applications

Examples of High End Computing and Computation Projects

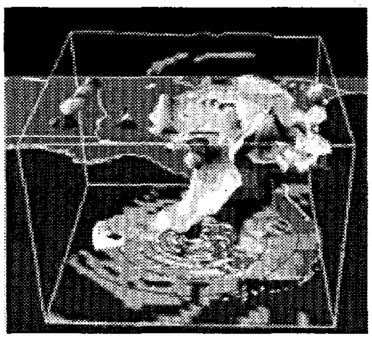

Tropical Storm Gordon, just before becoming a hurricane

GFDL/NOAA Hurricane Model

- Hurricane Prediction System
- 1993 - Run in semi-operational test mode
- 1994 - Part of National Centers for Environmental Prediction operational hurricane forecast suite
- 1995 -Model became fully operational
- 1996 - High resolution models run on highly parallel systems for the Centennial Olympics
- Run for tropical systems at all development stages
- Found to be in top performance group for forecasts out to 36 hours and superior to all other forecast models at 48 and 72 hours

Environmental Modeling

- Parallel logarithms used to model movement of groundwater contaminants
- Numerical simulations to estimate impact of decades of toxic pollution
- Research on nonlinear optimization and control techniques used to minimize groundwater cleanup costs
- Regional Particulate Model (RPM) developed to monitor air quality by region
- Parallel computing used to quantify potential effects of earthquakes on new construction and to assist in creating new safety codes

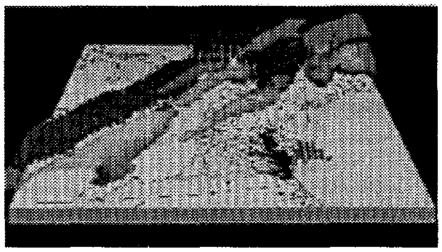

Simulation of nitrogen deposition to the Chesapeake Bay and surrounding areas during a rainstorm.

PetaFLOPS Workshops

- Enabling Technologies for PetaFLOPS Computing (PetaFLOPS I), February 22-24, 1994
- The PetaFLOPS Frontier Workshop, February 6, 1995
- 1995 PetaFLOPS Computing Summer School/Workshop: Applications and Algorithms Challenges for PetaFLOPS Computing, August 14-23, 1995
- PetaFLOPS Architecture Workshop (PAWS'96), April 22-25, 1996
- PetaFLOPS Summer Study on System Software (PetaSoft'96), June 17-21, 1996
- PetaFLOPS Algorithm Workshop (PAL'97) April 14-18, 1997

Example Research Agenda for:

Large Scale Networking

- Fast, efficient routing and enhanced network service
- High performance networking research infrastructure
- Design and scalability issues of large multi-node networks
- Very high performance network components and testbeds
- Network-centric computing including system software and program development environments for scalable distributed systems
- National Challenge-class application drivers such as:
  - Distributed information acquisition tools
  - Digital library technologies
  - Computer-based patient records
  - Electronic commerce

Examples of Large Scale Networking Projects

- I-Way was an extensive networking project introduced at Supercomputing '95
- Experimental, high performance network
- Linked over a dozen high performance computers and advanced visualization machines
- Used as testbed to prototype:
  - Teraflop-class wide area computing
  - An advanced application development resource

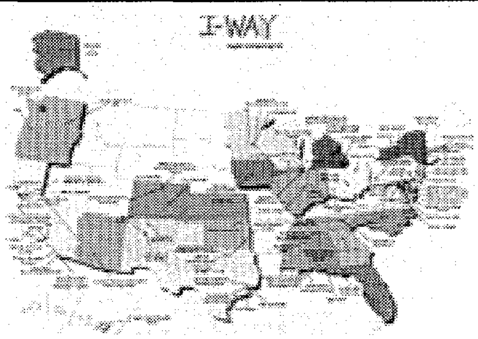

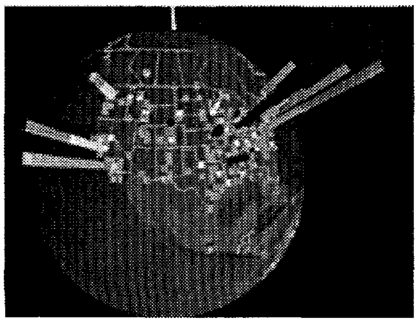

- Begun with development of ARPANET in late 1960s
- Network infrastructure created with NSFNET in 1985
- Development of Mosaic browser in early 1990s
- Community level service turned over to private sector Internet service providers in April 1995
- Unprecedented Internet growth stimulated by HPCC R&D and from public and private investment

The NGI Vision

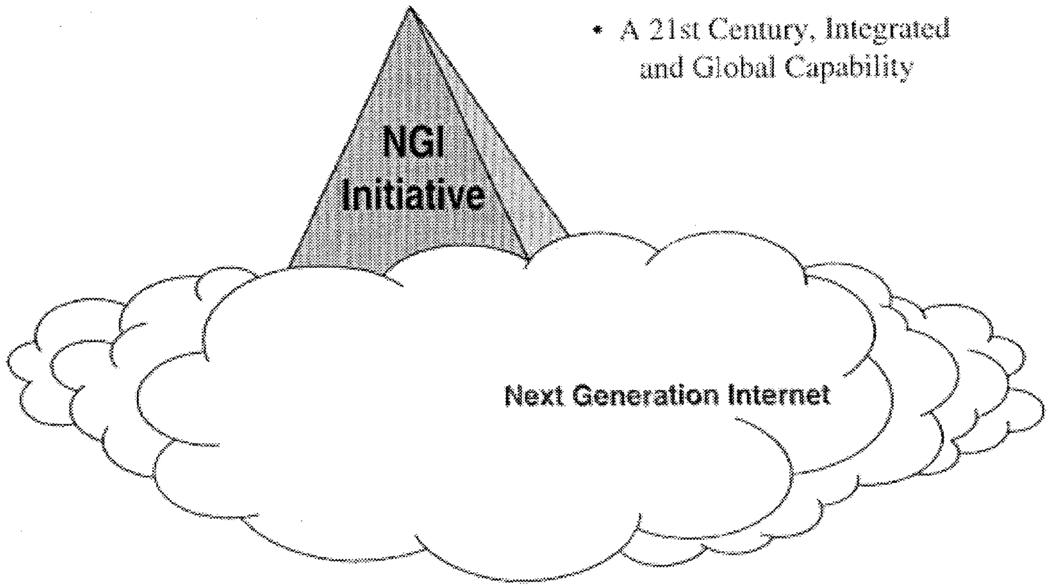

- A 21st Century, Integrated and Global Capability

Next Generation Internet Initiative

Goals	Metrics
1. Develop the next generation network fabric and connect universities and federal research institutions at rates that are 100 to 1000 times faster than today's Internet	end-to-end performance; number of institutions connected
2. Promote experimentation with the next generation of networking technologies	quality of service; adoption of technologies by commercial Internet
3. Demonstrate new applications that meet important national goals and missions	100+ high importance applications; value of applications in testing network technologies

NGI: "The Programs"

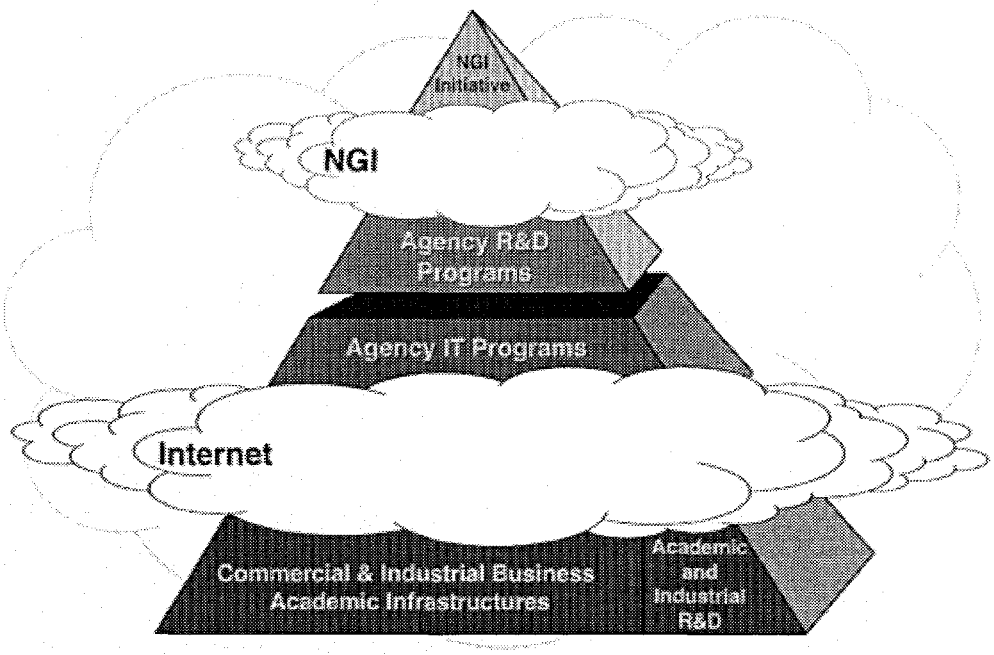

But We Are Really Impacting People!

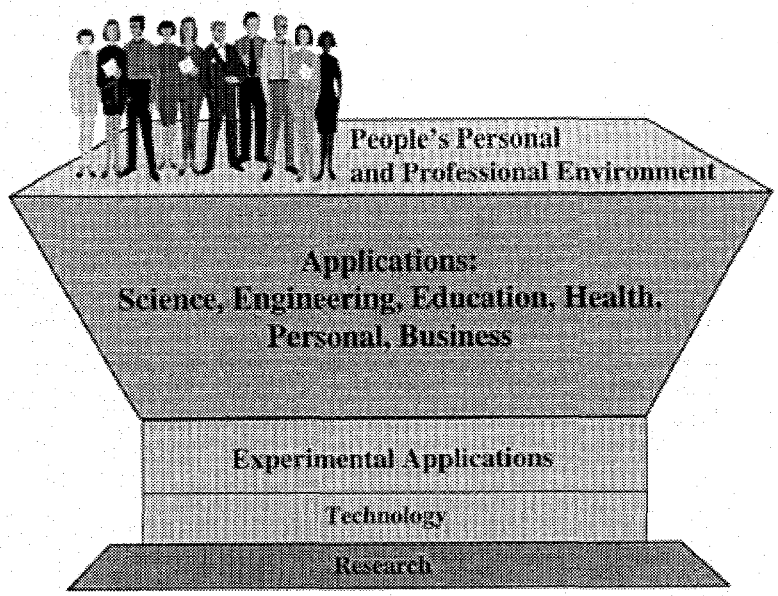

Overview of NGI Activities

Discussions/Feedback:

- Presidential Advisory Committee
- High-level advisory groups
- Special community teams
- Outreach with industry, education, and private industry
- Directly funded agencies
- Briefings

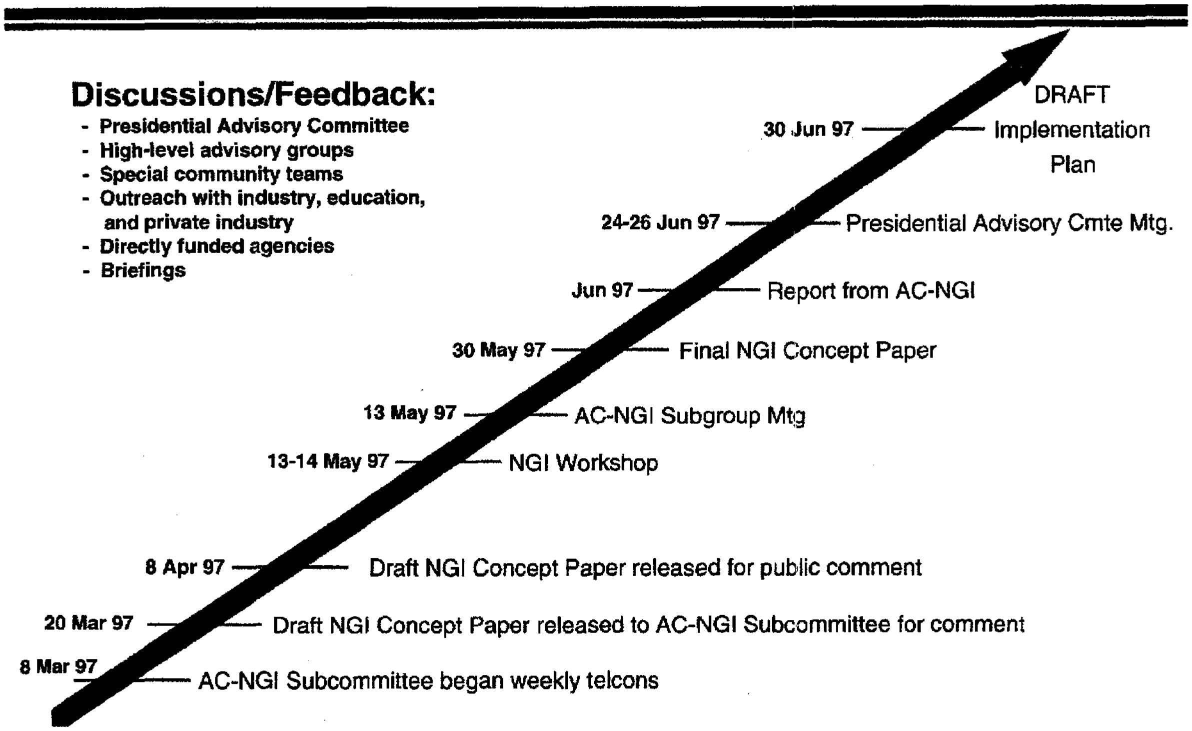

NGI Concept Paper

- Draft for comment now available
  - http://www.hpcc.gov/ngi-concept-08Apr97
  - http://www.ngi.gov (being created)
- Comments:
  - Encouraged to be submitted via WWW or to ngi@hpcc.gov
  - Comments received by 15 May 1997 will be used in preparing a final version of the document shortly thereafter

Agency Roles in the NGI

- DARPA: long-term, general expertise in networking research, general skill in high-end network technology and testbeds, experience in managing networks.
- DOE: long-term experience in managing production and research networks, specialized skills in networking technology, great strength in mission-driven applications and in system integration.
- NASA: experience in network management and in specialized network testbeds, strength in mission-driven applications involving high data rates, great strength in system engineering and integration.
- NSF: special relationships with the academic community, experience in network research and in managing networks, great strength in scientific applications.
- NIST: long experience in standards development, networking research, and in testbeds involving many industrial partners.

NGI FY98 Proposed $100 Million Budget (Dollars in Millions)

	DARPA	NSF	DOE	NASA	NIST
Goal 1: High Speed Connectivity	20	7	25	3	
Goal 2: Technologies	20	2	6	2	2
Goal 3: Applications		1	4	5	3
Total:	40	10	35	10	5

Note:Future versions of this paper are expected to include funding from additional agencies who want to be part of the initiative, for example NIH has expressed interest in joining the NGI.

Example Research Agenda for:

High Confidence Systems

- High Confidence Computing
  - Fault tolerant, secure hardware, and operating systems
  - Secure, trusted agents
- High Confidence Networking
  - Dynamically secure enclaves
  - Robust firewalls
  - Management of massive networks under load, failure, and intrusion
- Encryption, Assurance, and Integration
  - Formal system specifications and verification
  - Advanced algorithms
  - Tools for designing, integrating, and evaluating systems
- Large Scale Systems
  - Analysis and simulation
  - Testing, evaluation, and standards development

Examples of High Confidence Systems Projects

- Merging of computerized patient record systems and telemedicine systems in way that assures integrity and confidentiality of records

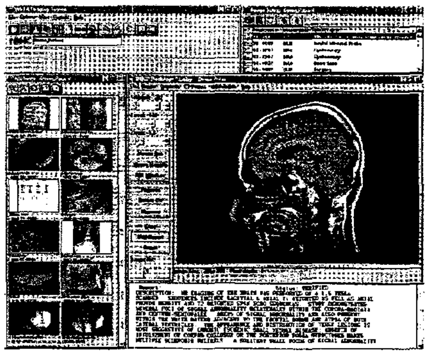

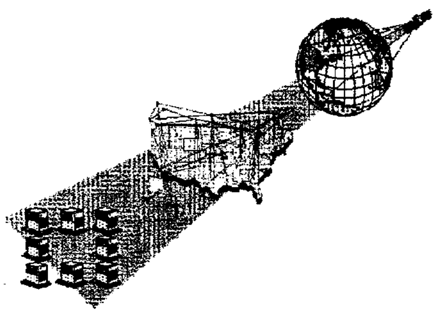

- Critical technologies to ensure global, survivable, secure networks that support high performance distributed computing and information systems with dynamic topologies

Example Research Agenda for:

Human Centered Systems

- Collaboration technologies and applications, including collaboratories, distributed multi-media, middleware, and control of remote instruments
- Multi-modal human-system interactions
- Document understanding and multilingual language technology
- Information agents and environments
- Virtual reality environments and applications
- Knowledge repositories for information access, management, and applications
- Interdisciplinary research in human cognition applied to tools for content recognition and human-system interaction

Examples of Human Centered Systems Projects

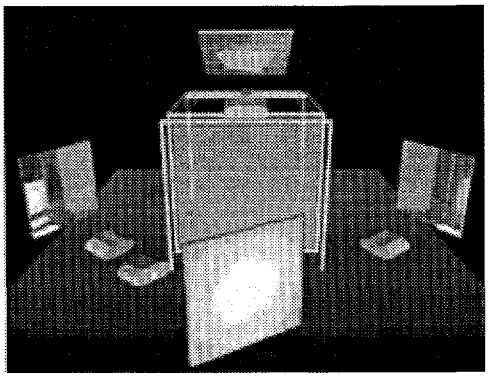

Surround-screen, surround-sound, projection based virtual reality systems (CAVE)

- Visualization, virtual reality, and human-machine interfaces
- Ability to generate, collect, and manipulate vast amounts of data
- Development of a prototype "omnifocus" electronic camera

- Health care and biomedical imaging
- Enhance patient care, improved drug design, and broaden access to medical information

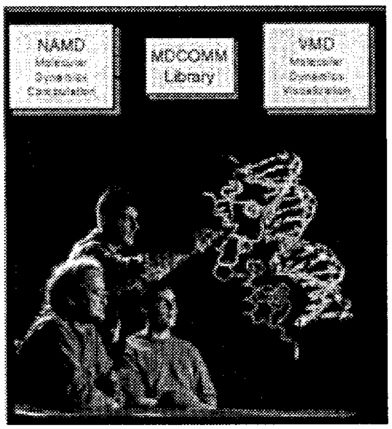

Prototype visual medical chart integrating multi-disciplinary medical images

Example Research Agenda for:

Education, Training, & Human Resources

- Lifelong learning technologies
- Curricula, fellowships, and scholarships for computational, computer, and information sciences and engineering
- Demonstration and development of innovative information infrastructure and technologies for education
- Research supporting distance learning for remote sites

Examples of Education, Training, and Human Resources Projects

- Math and Science Gateway and Gateway for Educators provide web site with access to science and math resources for students and teachers
- Enable design and deployment of multimedia modules used in interactive tutorials
- HorizonNet demonstrates low cost Internet access for a school consortium
- Develop interactive and visualization tools for special educational projects

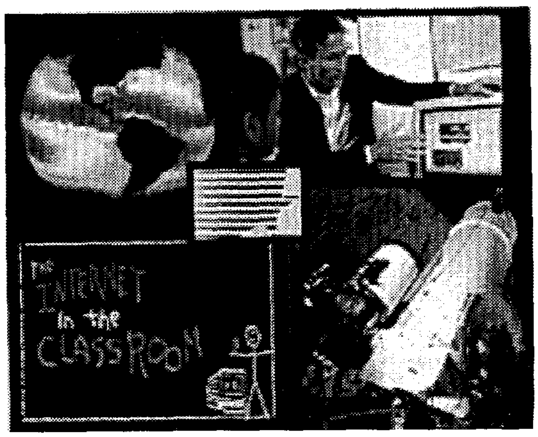

H P C I P T

High End Computing for National Security

Conclusions of an
Integrated Process Team

April 1997

1

H P C I P T

HPC IPT Tasks Addressed by this Report

- What are the alternative future scenarios with respect to the US international lead in the high end of high performance computing?
- What is the threat to national security associated with each scenario?
- What technology advantage is necessary for the US, if any?

2

DDR&E called for an Integrated Process Team (IPT) in April 1996. The IPT was assembled and initiated work in July 1996. The team collected needs, industry views, and reviewed market analyses. Based on the information collected, the IPT developed its conclusions. The IPT consisted of members from DARPA, NSA, DOE, NASA, NRL, Air Force, Army, NCO, OSD, CIA, and HPMO and was chaired by the DARPA member.

HPC IPT

What Is High End Computing?

Manufacturer's Performance: A Moving Target

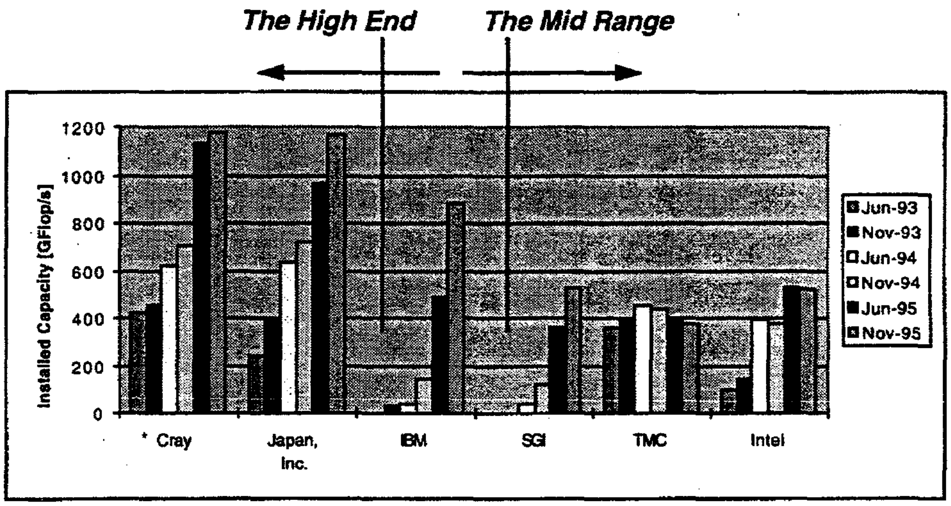

Source: Top 500 Supercomputer Sites compiled at Mannheim University
http://www.netlib.org/benchmark/top500

* Now SGI/Cray

3

The high end is a moving target. Some define it as the capabilities of the top 50 machines in the world. Here is a picture of the installed computational power of the largest facilities in the world. In 1996, TMC and Intel were no longer producing commercial machines.

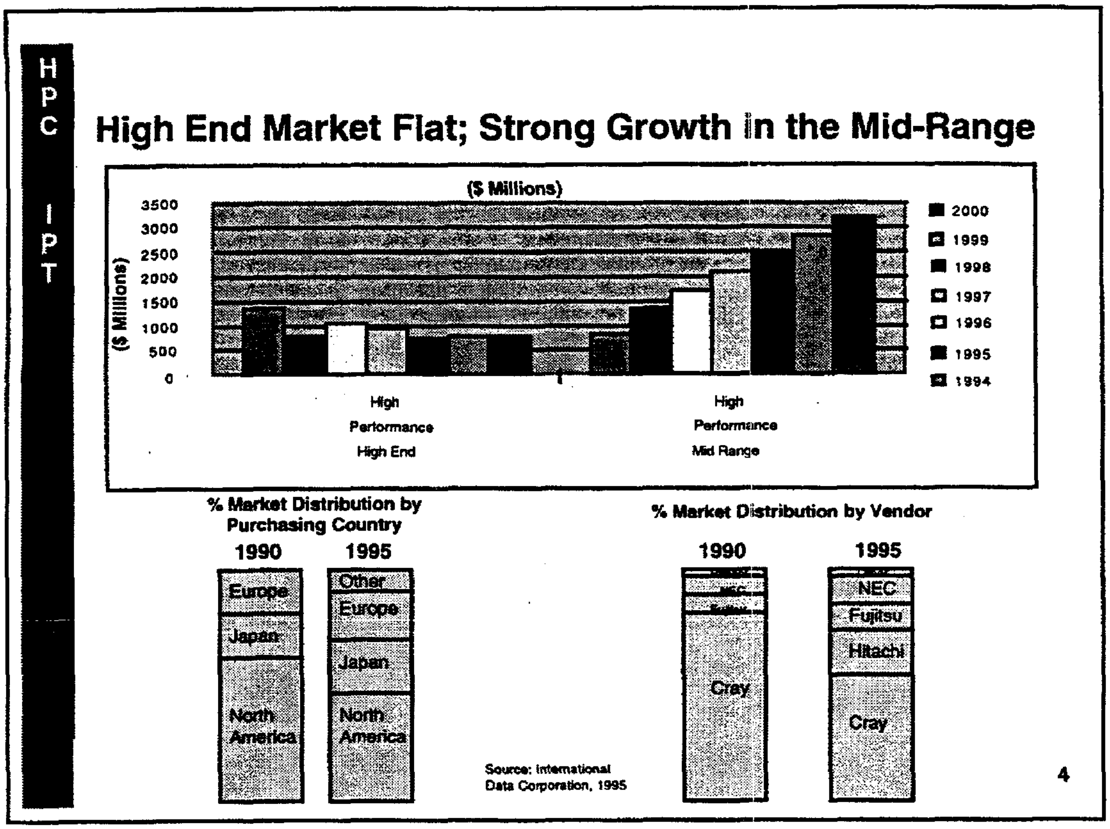

Many US companies abandoned the high end because of flat/declining market with insufficient ROI for future investments, aggressive competition and the potential for higher returns in the mid range.

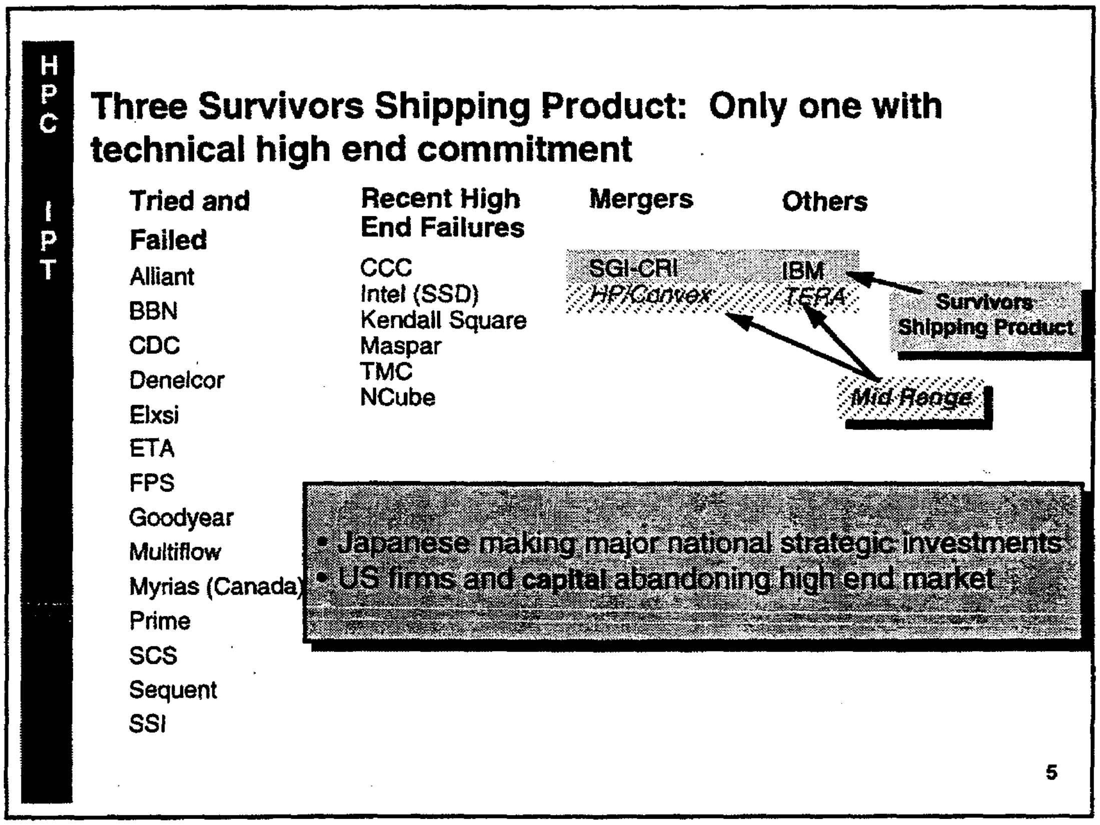

Worldwide revenues are down 30% from 1994. Numerous companies failed in the market with high end losses by the remaining providers. There are only two US companies shipping high end computers today: SGI/Cray which supplies systems with direct applicability to national security problems and IBM which ships principally to the business market but which is providing a system to DOE's ASCI program. HP/Convex produces mid-range systems. Tera is a start-up which has not yet shipped a commercial system. Japanese companies have embarked on an aggressive investment program supported by vigorous Japanese government procurements at a number of universities.

HPC IPT

Federal Investments Fueled Today's US Technology Success

6

HPC IPT

Estimated Overall Federal High Performance Computing and Communications Program Expenditures by Components

	FY1992	FY1993	FY1994	FY1995	FY1996
HPCS	156.80	235.50	171.20	177.90	146.70
ASTA	285.10	325.90	352.10	350.30	371.00
NREN	91.50	123.40	142.10	163.50	185.60
BRHR	124.50	140.40	118.50	106.30	128.10
IITA	0.00	0.00	156.00	240.00	228.20
Total	638.30	795.70	937.90	1,038.00	1,059.60

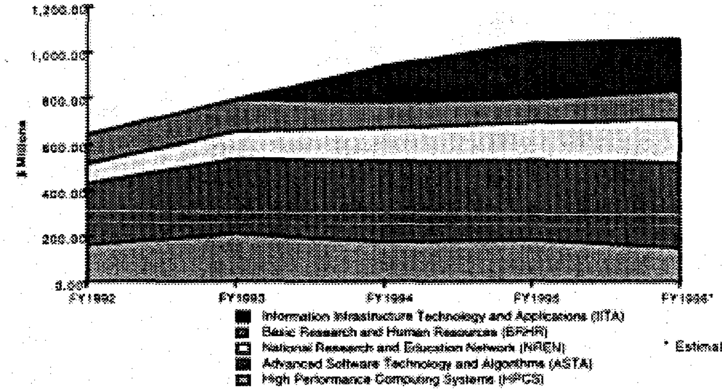

* Estimated

7

The Federal HPCC program was initiated in 1992. It has many activities including human centered systems, information infrastructure technology, networking and high end computing. The IPT estimates that HPCC investments applicable to the high-end of computing have fallen from a peak of approximately $230 million in 1995 to about $150 million in 1997. The IPT estimates that approximately one half of this amount ($75 million) is directly applicable to high end ***national security*** *needs.*

H P C I P T

Federal R&D High End Investments Fueled Today's Technology Successes

- Established scalable parallel processing as the commercial standard for high performance computing
- Enabled the technology base for the $2 billion middle range high performance market which expanded access to high performance computing while reducing costs to the government
- Invented and proved massively parallel systems as effective high end computing devices
- Enabled the near-term computing technology for DOE's ASCI program
- Created the scientific base for High End Computing including trained scientists and engineers, new architectural approaches and next generation technologies

8

The Federal Program produced significant technological successes. It did not address procurement and near-term prototyping issues. The emergence of mid-range processors enabled by Federal technology increased overall availability of powerful machines for the defense and national security, while reducing the cost of these systems to the average defense user. Partially because of the availability of this more attractive option for smaller scale problems, the very high end of the commercial market is now smaller and more specialized. Market growth in the middle proved to meet the needs of most commercial and government customers leaving a principally government high end market.

H P C I P T

National security policy-driven needs

US requires superiority in high end computing for:

- Superiority in weapons design
- Comprehensive test ban treaty
- Critical reaction capability/time for defense
- Battlespace dominance; revolution in military affairs

Superiority in critical defense technologies
requires
Superiority in high end computing technology.

9

It is the belief of the Integrated Process Team that superiority in critical defense technologies requires superiority in high end computing technologies. The US needs assured and timely access to these technologies.

HPC IPT

National Security Applications

- Nuclear weapons stockpile stewardship: Weapon effects simulation to extend stockpile in era of no new design
- Wide area imagery: Near real time analysis of imagery with 3D Resolution, 100,000 sq. mile coverage and high resolution
- Cryptology: Rapid decryption of multiple messages
- Vehicle and weapons design and test
  - High performance aircraft design and test: Full 3D multi-disciplinary (aero, structures, magnetics, propulsion, controls) simulation
  - Weapons systems such as high power RF weapons: End-to-end simulation to predict complex systems response to weapons effects
  - Target discrimination: Combat identification, signature extraction and rapid target insertion for advanced platforms

10

There are many national security applications which require high end capabilities.

HPC IPT

National Security Applications (con't)

- Intelligence data and information extraction
  - Remote sensing exploitation: Parameter analysis of chemical species and spectral band)
  - Chemical detection (FTIR remote sensor): Infrared image generation for radiometric, thermal and emissivity analyses
  - Intelligence data and Information extraction: Reduce time for analysis of complex sensor data from two weeks to real time
- Synthetic theater of war C^3: Mission rehearsal and decision support for Desert Storm size scenarios in faster than real time

11

HPC IPT

Computational Speedups of 1,000-1,000,000 required by 2010

- Nuclear weapons stockpile stewardship: 100,000 - 1,000,000X
- Wide area imagery: 2,000X
- Vehicle and weapons design and test
  - High performance aircraft design and test: 10,000-100,000X
  - Weapons systems such as high power RF weapons: 1,000+X
  - Target discrimination: 3,000
- Intelligence data and information extraction
  - Remote sensing exploitation: 10,000-100,000
  - Chemical detection (FTIR remote sensor): 300,000X
  - Intelligence data / Information extraction: 3,000X
- Synthetic theater of war C^3: 100,000X

12

National Security needs for high end computing exceed projections of commercial mid-range processing by factors of 10 to 10,000. Emerging defense applications will make even greater demands on the technology. According to industry suppliers, national security needs have been 10 to 100 times greater than the greatest needs of their commercial customers with trends continuing this disparity.

H
P
C

I
P
T

Impact of High End Technology Deficit

- Nuclear weapons stockpile stewardship: US unable to extend stockpile life within "zero-yield" nuclear test ban
- Wide area imagery: Unable to provide tactical warfare information superiority
- Cryptology: US loses ability to decrypt industrial-strength message traffic

13

Some of the impacts of the high end technology deficit are listed here. DOE will not be able to have the computational base required to extend stockpile lives without testing. Real-time image analysis and target insertion in battle systems will not be possible. Modern cryptographic techniques will defeat attempts to decrypt traffic. Continued aircraft performance increases may be imperiled because of the inability to model complete vehicles and weapons. We will be unable to predict weapons effects for weapons systems such as high power RF weapons. Real time insertion for target discrimination will remain out of reach. We will be unable to utilize advanced sensor capabilities to predict atmospheric effects and our adversaries' chemical warfare capabilities may go undetected. Data volumes and analysis needs will overwhelm analyst capabilities threatening our information dominance strategy. Course of action decision support for Desert Storm size engagements will not be feasible.

H P C I P T

US Computer Industry Not Producing Next High End Generation to Meet National Security Needs

US computer industry unable or unwilling to invest their own resources to meet high end national security needs

- Insufficient ROI
- Opportunity Cost
- Uncertain Market Future

14

We cannot expect industry to invest their own resources to meet national security needs without a predictable and profitable customer base. Industry needs include stable and predictable procurements, at a level which encourages R&D investment, pricing structures which cover one-time costs of engineering and development and long-term research to pioneer architectures and technologies.

H P C I P T

National Security Threats

- US technology and products may not be available to meet high end National Security needs.
- A potential collapse by withdrawal from market because of insufficient return on investment would lead to dependencies and unpredictable lead times for procurement
- Possible denial of high end systems to US Defense Programs (e.g. Nuclear)
- Expensive special purpose development systems would require years for development.
- Loss of technology base for special purpose development systems with very large costs for replacement
- Long-term technology loss in critical weapons leadership

15

US vendor issues in the high end of computing market pose a number of threats to national security. In addition, it can not be presumed that the current investment strategy of the Japanese companies will be continued if the overall market remains relatively small and flat. Therefore the threat to national security caused by the lack of investment and small number of viable competitors is significant..

Conclusions and Recommendations

- US has critical national security needs for high end technology and technology leadership that are not being met
- Without additional government action, US high end technology will not be available to meet national security leading edge needs
- These concerns require a proactive approach by the defense and national security community to determine an appropriate government response
- The community should implement measures to maintain technology leadership by establishing a joint, high end, national security partnership with industry

16

National security interests require a strong, proactive approach to high end computing. The IPT believes that the national security needs require government to establish strong partnership with industry. This approach should address procurement policy, the costs of prototypes and non-recurring engineering for advanced high end systems, software development, and long-term research. The IPT believes that near-term action is required to reverse the eroding US technology position and without this action risks to national security will increase significantly.

Users of High Performance Computing Technologies

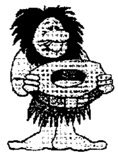

The First 10

- Tolerate anything to find a solution

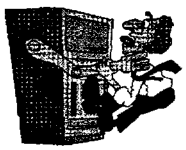

The Next 100

- Presume stability, but with driving needs tolerate some pain

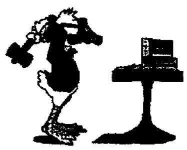

The Next 1,000

- Want stability, test the water, in their own discipline

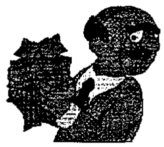

The Next 10,000,000

- Want everything...and for free

What's Important?

- Long term research
  - Agency-specific
  - Interagency Initiatives
  - Duality of computational science for mission and science
- Strategies & ideas across the disciplines
- Leadership
  - Communicating the visions carefully
  - Constructively moving forward
  - Being adaptive, critical, and sensitive simultaneously
- People
  - Building the right culture, attitudes
  - Attracting the best and brightest - to government, industry, academia

Encryption Future/Trends

Mike McConnell

Vice President, Booz-Allen & Hamilton

Conference on High Speed Computing

21-24 April 1997

Gleneden Beach, Oregon

Background

- Kahn - The Codebreakers (1970's)
  - WW II
    - German Enigma
    - Japanese AN25
- Symmetrical Encryption - the stuff of Governments
- Asymmetrical (Public Key) Encryption (1970's)
- Banking in the 1970's
  - DES - for Secure Global Banking
- Cold War Arms Control Legislation
  - Export Controls on Cryptography
- The Current Debate -- Future Prediction

Information

Offense	Equities	Defense
Information Warfare	?	Information Assurance
DoD–Programmatics Joint Staff–Operations Services–Development	Citizen Privacy Business Privacy Law Enforcement Signals Intelligence	Federal Classified–NSA Federal Unclassified–NIST Commercial } ? Private } ?
↓	↓	↓
Control/ Degrade/Destroy	Balance	Protect

Policy Dilemma

Need To Reconcile:

- Personal Privacy
- Business Privacy
- Law Enforcement
- National Security

} Interests -- Equities

New Issue: National Interdependencies & Information Assurance

- Information Warfare (Threat? & Vulnerabilities!)

Political and Legal Limitation Issues

- PL 100-235, National Computer Security Act of 1987
  - NSA Protects Classified and National Security Information Only
  - NIST Responsible for All Other Federal Data
  - Does Not Address Protection of Critical *Public or Commercial Systems*
- Political Realities
  - Lingering Mistrust of NSA From Church/Pike Investigation
  - Negative Publicity/Reaction to Clipper Chip

How Vulnerable Are We?

- Massive Networking Makes the U.S. the World's Most Vulnerable Target for Information Warfare
  - Intelligent Exploitation
  - Disruption of Network Infrastructure
- U.S. Has Orders of Magnitude More To Lose to Information Warfare Attacks Than Its Competitors
- Reliance on Unprotected Networks Carries Risk of Military Failure and Catastrophic Economic Loss

What Is At Risk?

- Power and Utility Distribution
- Telephone System/Public Switched Network
- Stock Exchange/Security Trading
- Federal Reserve/IRS/Social Security
- Banking
- Strategically Important Companies
- Research and Development
- Air Traffic Control System

** Degradation of Any of the Above Impacts National Security*

Wall Street Journal
(September 12, 1995)

- CITICORP Raid Details
  - Russian Hacker
  - International Actors and Venues
  - Millions of Dollars Moved
  - Hundreds of Thousands of Dollars Lost
- No Wire Fraud Laws in Russia
- No Extradition Treaty With the U.S.
- Dutch Won't Extradite for Financial Fraud
- Implies Problem Is Widespread

Observations

- Organized Crime Looting U.S. Firms via Computer
  - Russian and Others
  - Drug Cartels
  - Terrorists
  - Over $5 Billion/Year Estimate
- International Legal Structure Inadequate
  - Computer Crime Not Always Illegal
  - Cooperation Is a Problem
  - Extradition Agreements Lacking
- Risks Go Beyond Financial Loss

Where Are We Now?

- Law Enforcement & National Security Concerns
  - Influence World Cryptography
- Public Key Infrastructure (PKI)
  - Cheaper/Faster/More Secure/Greater Confidence
  - Threat to National Security & Law Enforcement
- Equities Issue – Balancing
  1. Personal Privacy
  2. Business Privacy
  3. Law Enforcement
  4. National Security
- Perfect Solutions for 1 & 2 have strong impact on 3 & 4

Current Legislative Debate

- "Pro-Code:" Robust code for Everyman
  - Senators Burns (R-MT), Leahy (D-VT), & Wyden (D-OR)
    - -- "My Lock, My Key"
- "SAFE Act:" Rights of U.S. persons
  - Representative Goodlatte (R-VA)
    - -- Security And Freedom through Encryption (SAFE) Act
- Presidential Policy Decision/Legislative Draft
  - Key Escrow/ Data Recovery
  - Transfer Export Controls from State to Commerce
  - Temporary Relaxation of Export Controls
    - - Tied to Key Recovery

Relaxation of Export Controls

- 56-Bit Key Length DES Without Key Recovery
- Six-Month Renewable Licenses for Only 2 Years
  - One Time Review
- Contingent Upon:
  - Detailed Key Recovery Development Plan
    - Program Plan
    - Promotion Schedule
    - Marketing Plan
  - Six-Month Renewal Depends on Measured Progress

PCCIP Findings

- Threats to US Infrastructure Greater than in Past
  - Growth of Networking
- Vulnerabilities Increased by Major Industry Trends
  - Consolidation/Geographic Concentration
  - Interdependencies
  - Reengineering
  - Alternative Channels (Internet & Intranet)
  - Globalization
- Strong Insight into Minor Threats (Hacking, Crime)
- Less is Known about Leading Edge Offensive IW
  - Closely Guarded Secrets of Nation States

IBM CEO

We'll look back on the Internet fifteen years from now...and realize its impact is analogous to manned flight.

Louis V. Gertsner, CEO, IBM

(Paraphrased)

Current Realities

- Vulnerabilities on the Rise
- Threats Will Increase
- Public Key Infrastucture will grow; if true, will...
  - Raise the Bar for Protection
  - Enable Commerce/Activities not Presently Envisioned
  - Will Negatively Impact LE/Nat'l Security
- Data Recovery (Key Escrow) is a Business Need

Security Terms

Term	Definition/Analogy
Data Integrity	Absolute Verification Data Has Not Been Modified *Detection of a Single Bit Change*
Authentication	Originator Verification *Signature on Check*
Nonrepudiation	Undeniable Proof (Audit) of Participation *Sender/Receiver in Bank Transaction*
Confidentiality	Privacy With Encryption *Scrambled Text*
Availability	Service on Demand Assurance *Guaranteed Dial Tone*

Symmetric Encryption

(Secret Key)

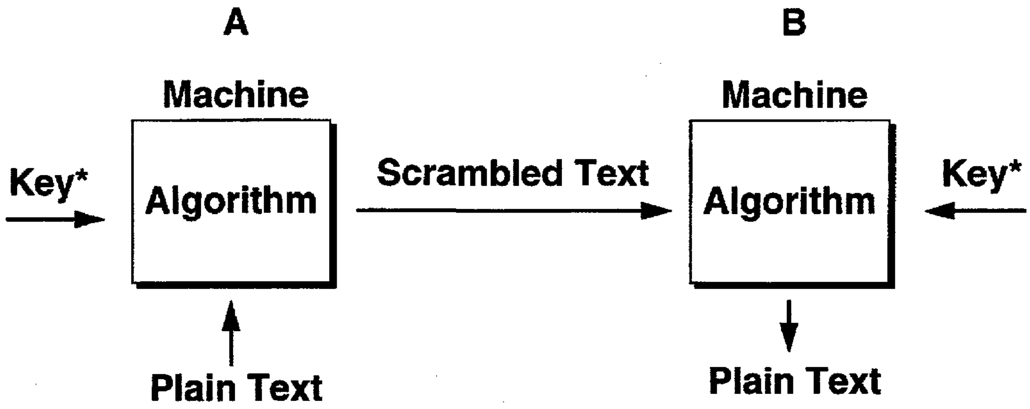

Vulnerabilities/Burden

- **Key Management/Control/Accounting**
- **Expense**
- **One Compromise will Compromise All**

* Keys Are Identical/Millions of Copies

Asymmetric Encryption

(Public Key)

Certification Authority Issues Public and Private Keys

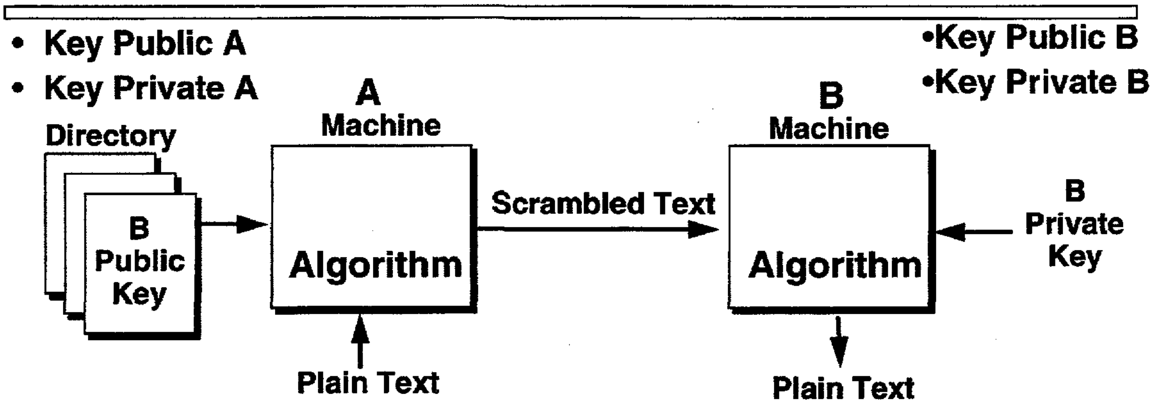

- **A Uses B's Public Key To Encrypt**
- **B Uses B's Private Key To Decrypt (Only B Can Decrypt Text)**
- **Only B Can Decrypt Text**
- **"Signature:" A Signs with A's Private Key & B Verifies by Decypting A's Public Key**

Advantages: Cheaper/Faster/More Secure/Greater Confidence

Cryptography

- Exploitable
- Robust
- Robust With Escrowed Key

Hardware

- Complete

Software

- Cannot yet Bind Escrow and Algorithm?

Cost To Reverse Engineer and Attack

Hardware
- Millions of Dollars
- Man Years

Software
- Thousands of Dollars
- Man Days

Information Security and Protection

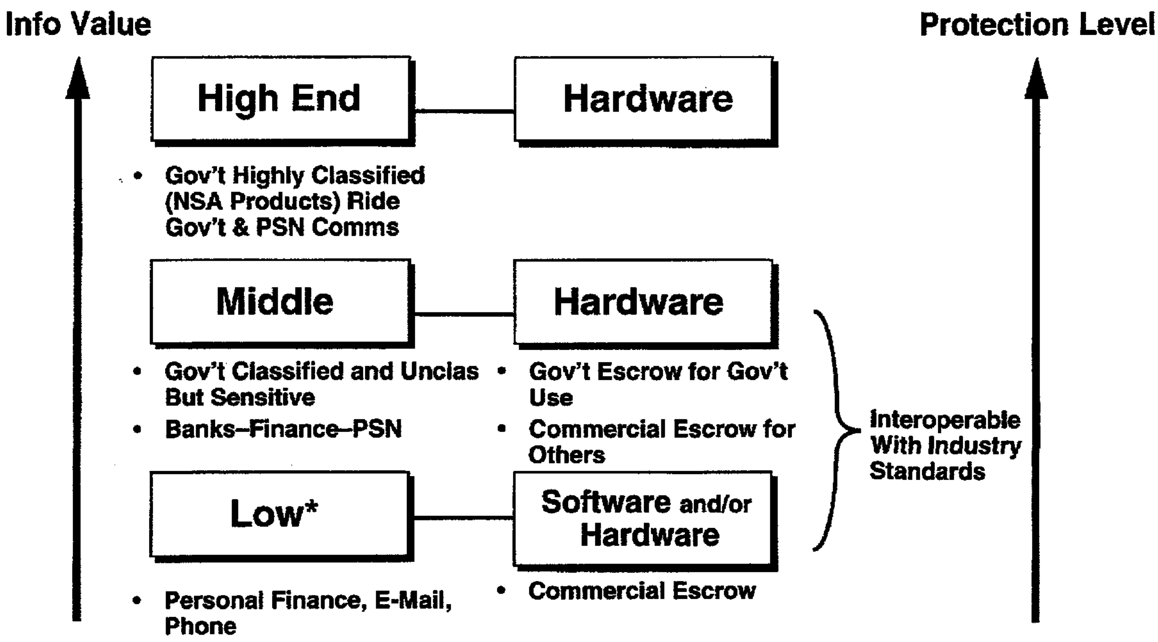

* If Degraded, Destroyed, Does Not Impact Security of Nation at Large

Drivers for Widespread Encryption

- Network Growth
- Ubiquitous PCs
- Need for Electronic Commerce Protection
- Need for Electronic Mail Protection
- Data Transfer Protection
- Theft of Service Protection
- Intellectual Property Rights Protection

Information Security Tiers

1	2	3	4
Physical	**Protection**	**Application**	**Enterprise** ....
Vaults Guards Fences Barriers	Diagnostic Curative Ops/Admin COMSEC & INFOSEC	PKT–Public Key Technologies Applied to Networks	Electronic Commerce

Identity Certificates → Authority Certificates

EBT

EDI

Lower Costs
Increased Speed
and *Confidence*
based on Trusted
Infrastructure

Past → | Present → | Future →

OECD Encryption Guidelines
27 March 1997

1. Crypto should be trustworthy for confidence
2. Right to choose any crypto, <u>subject to applicable law</u>
3. Crypto should be responsive to people, business & govt
4. Standards developed & spread at national & internat'l level
5. Respect fundamental right of individuals to privacy
6. National policies <u>may allow</u> lawful access
7. Liability of crypto providers clearly stated
8. Gov'ts should co-operate to co-ordinate crypto policies and should remove or avoid unjustified trade obstacles

BACKUP SLIDES

- **Following provide Background Information**

OECD Guidelines
(1-4)

1. **Cryptographic methods should be trustworthy in order to generate confidence in the use of information and communications systems.**
2. **Users should have a right to choose any cryptographic method, <u>subject to applicable law</u>.**
3. **Cryptographic methods should be developed in response to the needs, demands and responsibilities of individuals, businesses and governments.**
4. **Technical standards, criteria and protocols for cryptographic methods should be developed and promulgated at the national and international level.**

OECD Guidelines
(5-8)

5. The fundamental rights of individuals to privacy, including secrecy of communications and protection of personal data, should be respected in national cryptography policies and in the implementation and use of cryptographic methods.

6. National cryptography policies may allow lawful access to plaintext, or cryptographic keys, of encrypted data. These policies must respect the other principles contained in the guidelines to the greatest extent possible.

7. Whether established by contract or legislation, the liability of individuals and entities that offer cryptographic services or hold or access cryptographic keys should be clearly stated.

8. Governments should co-operate to co-ordinate cryptography policies. As part of this effort, governments should remove, or avoid creating in the name of cryptography policy, unjustified obstacles to trade.

Export Controls To Commerce

- Industry Desires
  - Dual-Use Technology
  - More Predictable, Timely, Friendly Process
- Executive Order/Presidential Decision Directive
  - Protect Encryption Control
  - Include Justice In Licensing Decisions
  - If New Legislation Fails To Protect Encryption, the Jurisdiction Returns to State

Key Recovery Legislation

To Facilitate Development of Market for Key Recovery Products and KMI

- **Key Recovery Agents:**
  - License
  - Liability Protection
- **Legal Conditions for Release for Law Enforcement**
- **Criminal Charges:**
  - Misuse of Another's Keys
  - Using Encryption in Committing A Crime

Encryption Policy

	Yesterday	Today/Tomorrow
Technology	• Private Key (Symmetric) • Centralized Key Management • Key Used for Long Time Periods • Small User Set • Mainframes • Encryption Use Constrained	• Public Key (Asymmetric) • Distributed Key Management • New Key Every Message • Global Applications • Workstations • Encryption Use Proliferates
Policy	• One-Sided • Classified Algorithms Implemented In Hardware Only • Stringent Export Controls	• Standards Based on Public Debate • Published Algorithms for Hardware and Software • Refocused Export Controls
NSA	• Proprietary Products • Risk Avoidance • Control • SIGINT Access Assured	• COTS Applications • Risk Management • Influence • SIGINT Loss Inevitable
Law Enforcement	• Government Escrow	• Trusted Third Party + Corporate Escrow

Cooperation With Industry

Today

- Export Controls
- Foreign Controls
- Financial Incentives
- Patriotism
- Endorsement

Future

- Relaxed Export Controls
- Foreign Cooperation
- Financial Incentives
- Patriotism
- Value Added
  - Building Codes and Standards
  - Product Differentiation
  - Liability Protection
- Domestic Policy
- Public Key Infrastructure

Secret vs. Public Key

Secret (Symmetric)

- Identical Key Used By Both Communicants
- Requires a Trusted Key Generation Authority
- Requires Distribution By Safe Means of Key To All Communicants
- Geometric Growth: 1,000 People Need ~1,000,000 Keys

Public (Asymmetric)

- Each Communicant Has a Public Key and a Private Key
- To Be Used Widely Requires a Trusted Party To "Certify" the Authenticity of a Public Key
- Sender Uses Public Key of Destination Communicant To Encrypt Communications
- Recipient Uses Private Key To Decrypt
- Linear Growth (Straight Line): 1,000 People Need 1,000 Key Pairs

Higher Trust
(Who Makes The Directory?)

- **Certificate Authority = Source of All Trust**
- **Different Certificate Authorities for Different Roles**

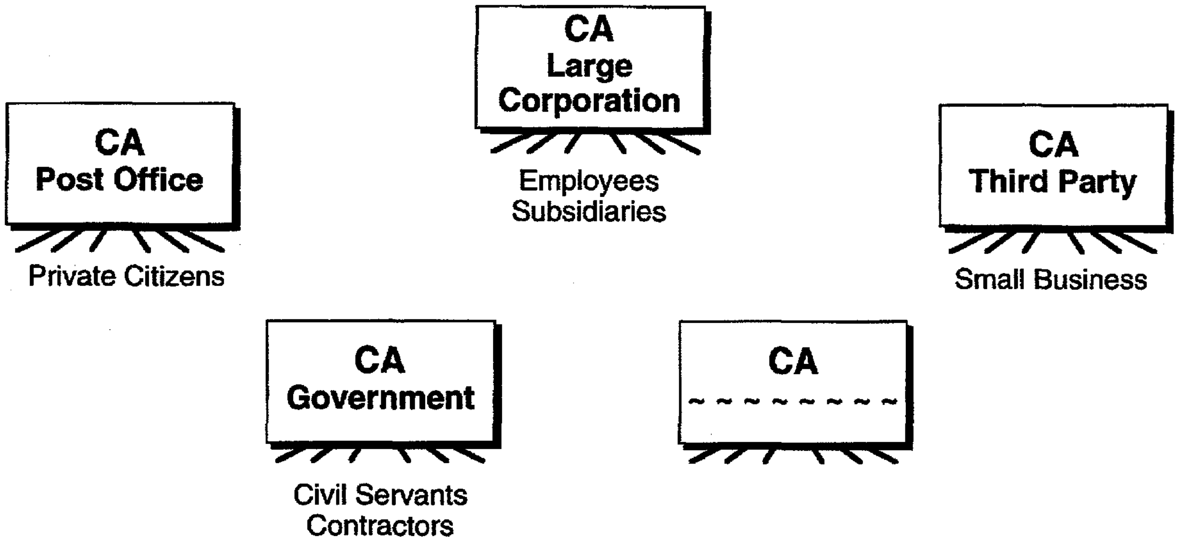

Public Key Encryption
(Confidentiality)

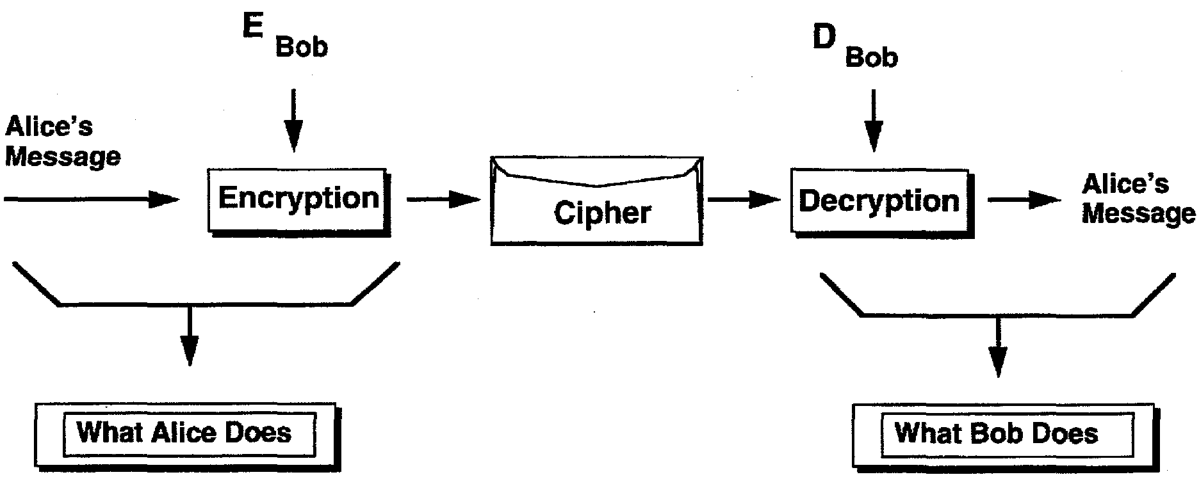

FUNDAMENTAL QUESTION:
Do You Trust That the Public Key Really Is the Person You Think It Is???

Public Key Cryptography

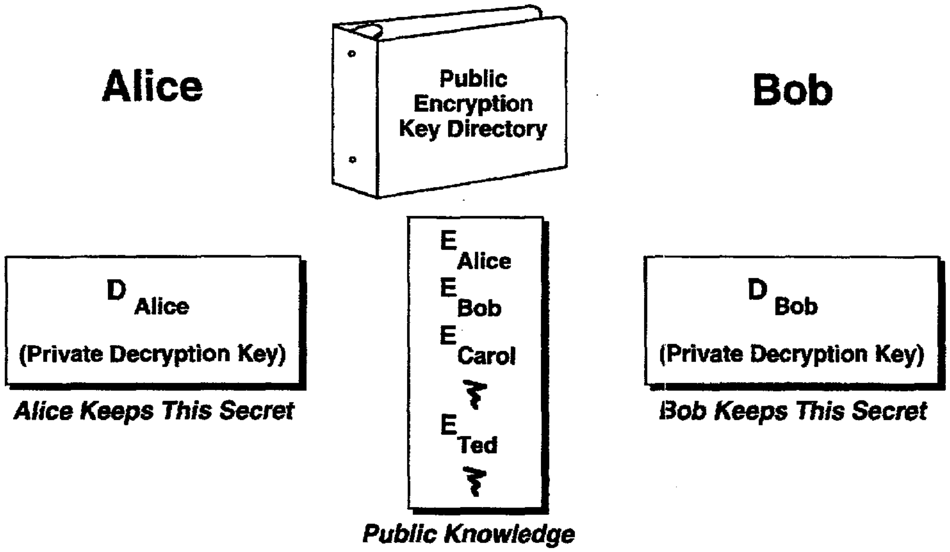

Who Makes The Directory???

Evolution of Crypto Policy Debate

Yesterday

- Crypto Products Produce/Purchased Mostly by Military/Government
- Exports Controlled by All Major Producing Countries
- Crypto Products Were Expensive
- Poor Quality Commercial Solutions
- Little Public Interest/Need in Commercial Solutions

Today

- Internet/Information Superhighway Creates Opportunities for Electronic Commerce and Network Based Services
- Use of the Network Is Unsafe Without Crypto
- Public Key Techniques and Microelectronics
  - Inexpensive
  - Good Quality
  - Easy To Use
- U.S. Industry Sees Worldwide Market Potential
  - Pressures To Eliminate Export Controls

Interests

- U.S. Market Dominance (Vendors)
- U.S. Privacy/Security (Citizens, Business, Government)
- U.S. Public Safety (Law Enforcement)
- U.S. National Security (SIGINT)
- Foreign Policy Safety/National Security (Allies)

Initiatives

- Security Management Infrastructure
  - Strong Information Security
  - Protects Law Enforcement via Escrow in KMI
  - Relax Export Controls With Key Escrow in KMI

Defensive Information Warfare
(aka)
Information Assurance

The Story So Far...

CAPT Bill Gravell, USN
Joint Staff / J6K
Washington, DC 20318-6000

Outline

- When last we met...
- Since then
- Now
- Tomorrow
- Someday

Requirements drivers for supercomputer technology

When Last We Met...

- Vision had emerged
- Vision was being translated into a plan of action
- Aggressive marketing campaign was being conducted

The IW Vision

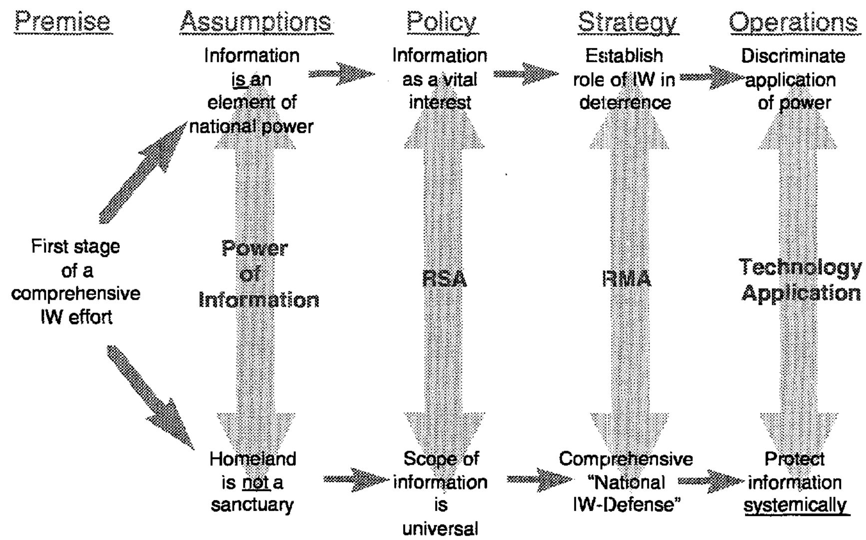

The IW-D System

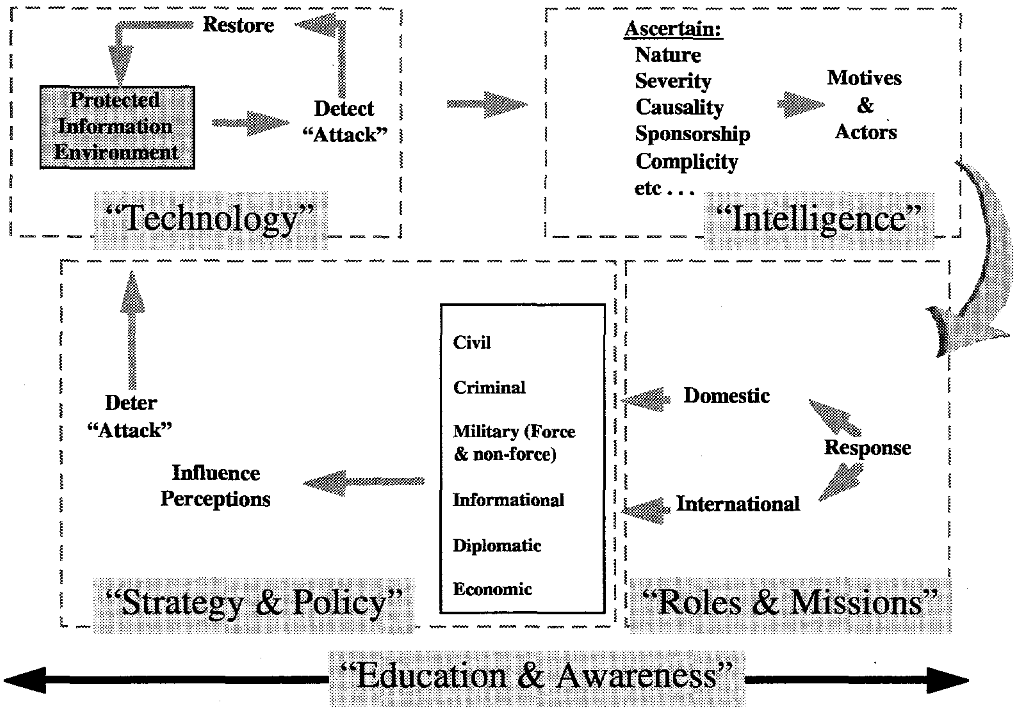

The Beginning

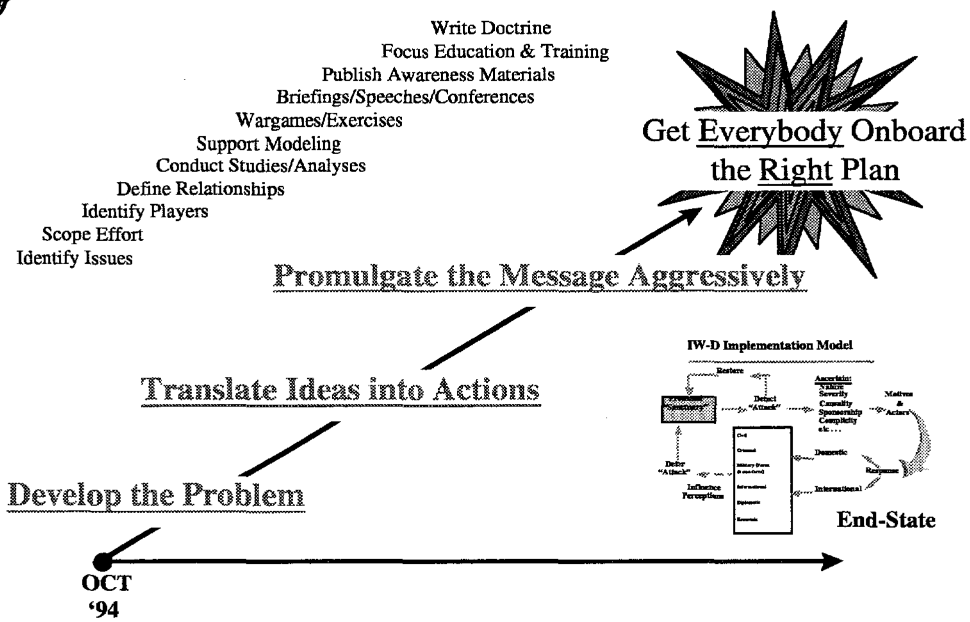

Since Then

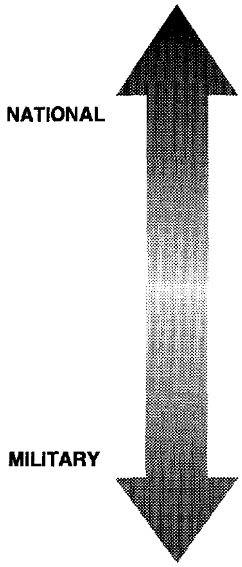

- Growing Congressional Interest/Action
  - Cohen/Nunn Hearings (SASC)
  - Information Superiority Hearings (HNSC)
  - 18/USC 1030 ("NII Protection Act of 1996")
- Gained Presidential-level Attention
  - EO 13010 (15 Jul 1996)
- Defense Science Board Task Force on Defensive IW (25 Nov 1996)
- "Joint Vision 2010"

DSB Task Force IW-D Recommendations

1. Designate an accountable IW focal point (CZAR)
2. Organize for IW-D
3. Increase awareness
4. Assess infrastructure dependencies and vulnerabilities
5. Define threat conditions and responses
6. Assess IW-D readiness
7. "Raise the bar" (with high-payoff, low-cost items)
8. Establish a minimum essential information infrastructure
9. Focus the R&D
10. Staff for success
11. Resolve the legal issues
12. Participate fully in critical infrastructure protection
13. Provide the resources

The Joint Vision: *Emerging Operational Concepts*

- **The Lens of Information Superiority**
  - **Integrates and Amplifies Four New Operational Concepts:**

Enables New Operational Capabilities:

- **Self-Synchronizing Forces**
- **Increased Speed of Command**
- **C4ISR Matched to Combat Power**
- **Decentralized Empowerment**
- **Enable Alternative Command Structures and Procedures**
- **Self Adapting and Learning Organizations**

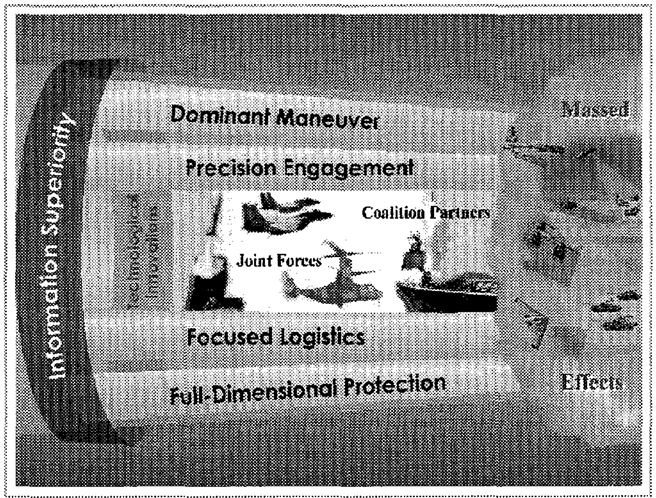

Now

- JV2010 Implementation Efforts
- Quadrennial Defense Review
  - Strategy - Joint Strategy Review
  - Modernization - IA Panel
  - Infrastructure
  - Readiness
  - Force Structure

(Strategy - Joint Strategy Review:)
- Asymmetric Threats
- Terrorism
- Vulnerability of Information Processes
- Dependency on Infrastructure
- "C4ISR"

- President's Commission on Critical Infrastructure Protection
  - J-6 role in CIPWG
- Secretary of Defense William Cohen Onboard

Defense/Information Assurance

- Society
- Industry
- Congress

- Participate in national security debate
- Advocate DOD mission support requirements
- Promote team building
- Seek organizational and technological solutions

Other US Gov't organizations ← **SECDEF**

- Promote collaboration
- Seek jurisdictional agreements

- CINCs
- Services
- Agencies

- Support mission execution
- Harmonize departmental efforts

Tomorrow

- "Information Terrorism"
- Advanced Concept Technology Demonstration (ACTD) for Information Assurance
- Maturation of Legal and Regulatory Regimes (?)

★ Risk Management Study (Summer '97)

- Models and Methodologies
- Commercial Approaches
- DoD Applications

★ Computer Forensics Study (Fall '97)

- Concepts for Info-Attack Source identification
- Technology Shortfalls and Opportunities

Policy: Terrorism (CbT)

- **Informational Considerations Just Now Emerging**
  - Terrorist use of Information Systems for their intended utility
    - Intelligence source
    - Propaganda medium
    - WWW.CRSP.ORG/PERU/SENDERO.HTML
  - Terrorist use of Information Systems to circumvent detection
    - Bulletin boards and cellular phones for C2
    - Encryption for privacy
  - Attacks against Information Systems themselves
    - Injurious effect achieved by disrupting or disabling critical infrastructures

Automated Intrusion Detection Environment ACTD

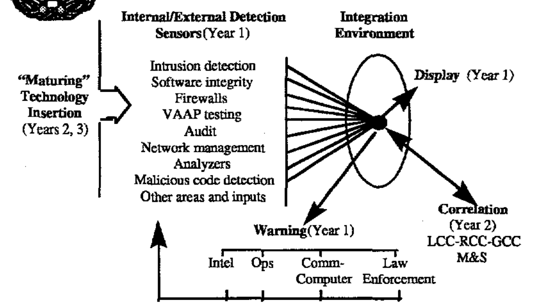

- **Objectives**
  - Answer "Are our systems under attack?"
  - Demonstrate integrated intrusion detection defense in-depth capability

Joint and Interoperable

- **Technology**
  - Firewalls, network management, network analyzers, auditing, encryption, malicious code detection, software integrity and correlation
- **Residuals**
  - Deployed intrusion detection environment at up to 30 sites
  - Links between Intel, Ops, C4I, and LEAs

- **Participants**
  - STRATCOM(CINC Sponsor) and TRANSCOM
  - DISA (Lead Agency)
  - Air Force, Army, Navy, DARPA, DIA, NSA, CIA, DoJ, FBI, AFIWC, Rome Labs and AFOSI
- **Schedule**
  - FY 98 ACTD Candidate
  - FY 98 start, 3 year duration
    - Yrs 1 & 2: Develop integrated environment and correlate sensor outputs
    - Yrs 2 & 3: Insert maturing technology every 6 mos

The Legal Setting

- **Domestic**: **Law follows technology**
  - Expresses evolving need of Society and its Institutions
- **International**: **Market forces shape relationships**
  - Global EC/EDI with digital signature, but not confidentiality

Evolution of Core Social Issues:

- Encryption

  PRIVACY Vs. SECURITY → PRIVACY (ANONYMITY) Vs. ACCOUNTABILITY

- Risk Management

  PERSONAL LOSS → - CONSEQUENTIAL DAMAGE - LIABILITY EXPOSURE

Someday

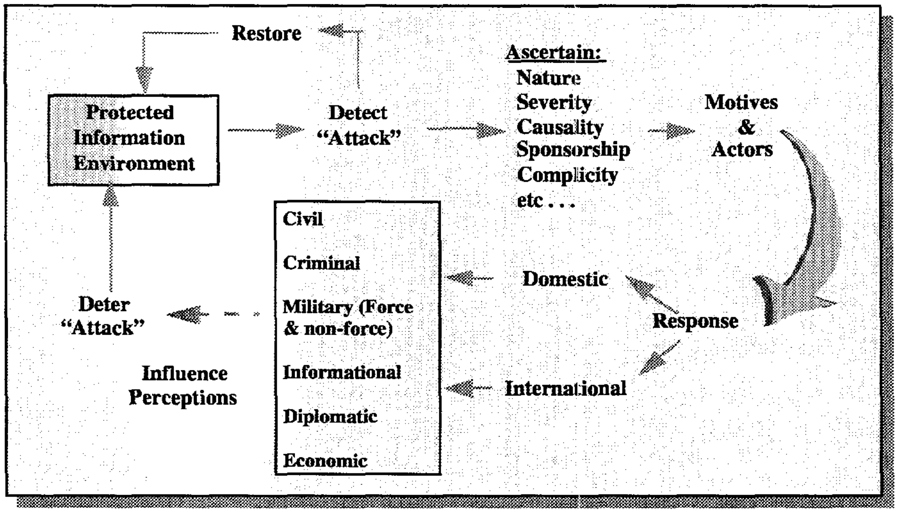

In the full bloom of the Information Age:

- Personal freedom and collective security are in socially-accepted balance
- Mature legal/regulatory environment protects/preserves these relationships
- Pivotal role of info technologies in total, National security fully recognized

Some Requirements Drivers for Leading-edge Computing Technology

- **National-level**
  - Protection Against Information Attack
    - Detection
    - Compilation / Analysis
    - Warning Dissemination
    - Dynamic Response

- **Military Operations**
  - Ballistic Missile Defense (Not just of U.S. Territory)
  - "Brilliant Sensors"
    - Traditional platforms/packages, plus:
    - Small/unmanned
    - Expendable (?)

Crisis Management Using Geographically-based Integrated Decision Support Systems:

Every Town Hall Needs a 4 Gig Drive

Paul S. Fischbeck
Department of Engineering and Public Policy
Department of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh, PA 15213-1414

paul.fischbeck@andrew.cmu.edu
412-268-3240

Outline of Talk

- Overview
- Military origins
- GIS systems and data layers
- Integrating optimization software
- RISES
- IISIS

Crisis Management

"Communities" face significant and/or unexpected risks. Preventing or mitigating these risks is often a complex, multi-attribute, data-intensive task.

- Natural disasters, accidents, crimes
- Contingency planning, real-time evaluation, post-mortem analysis
- Economic, environmental, societal impacts

Information Management

Critical to managing a crisis is managing the necessary information

- Sorting, displaying, and analyzing data
- Integrating different types of data
- Varied sources with different biases and reliability compete for the decision maker's attention
- Overcome cognitive biases with data

Multi-dimensional Data Display

- What are people familiar with?
  - linear file lists
  - two-dimensional "desktops"
    - » maps, graphs, organization charts
  - three-dimensional library of bookshelves
    - » virtual reality
- How should the data be displayed?
  - user defined?
  - pr(individual death/yr), # of deaths/yr/city, etc.

Data Display (continued)

- What should be displayed?
  - level of detail (reduced-form datasets)
  - amplifying information, justification
  - pedigree (prove that it should be trusted)
  - comparability ("bad" compared to what?)
- How should the data be organized?
  - efficiency: minimize steps to find data point
  - clarity: minimize getting lost and training

Data Concerns

- Maintainability
  - updates (how often?)
- Reliability and accuracy
- Security
  - levels of users: "Who gets to see/change what?"
  - "hackers"
  - individual confidentiality
  - commercial businesses

A Military Beginning: USS Carl Vinson (1981-1985)

- Captain Richard Martin
- DARPA and ONR funding
- Technology transfer
  - integrate, demonstrate, and evaluate technology as it is being developed
  - provide rapid transition of proven technology into the Navy
  - train Navy personnel

The Problem

- Need for decisive actions:
  - "He who hesitates is lost..."
- Uncertainty
- High rate of data flow, large database
  - static and dynamic
  - qualitative and quantitative
  - unfiltered and unanalyzed
- Integrate available optimization algorithms
  - search patterns, satellite overflight, data fusion

Research Areas

- Local area networks
- Distributed database management systems
- Human computer interface design
- Real-time decision support systems
- Artifical intelligence

Systems

- SDMS (Spatial Data Management System)
  - tactical database display
- ZOG
  - distributed database
  - human computer interface
- AirPlan
  - AI expert system for air operations
- CAT (Command Action Team)
  - AI expert system for threat analysis

SDMS

- CCA (Computer Corporation of America)
- VAX 11/780
- 300 Mb hard drives, 1.2 Mb RAM
  - fiber optic link to secure tactical database
- Two integrated laser video discs
  - 55,000 photographs
- Interactive graphics display and database
  - "driving" the database with a joy stick

ZOG

- Carnegie Mellon University
- 27 Perq mini computers
  - 24 Mb hard drives, 512 k RAM
  - 10 meg ethernet network
  - three "high speed" laser printers
  - integrated laser disks
- Distributed "hypercard" database
  - 150,000 frames of information

ZOG Applications

- Long-range planning
- On-line operational checklists
- Weapons elevator maintenance and troubleshooting
- Space management
- Virtual ship, virtual organization
- Email

AirPlan

- Carnegie Mellon and ship personnel
- 4 Perq computers using ZOG interface
- Rule-based expert system written in OPS-7
- Real-time analysis of air operations
  - background simulation of aircraft flight profile
  - depth of analysis varied with workload
  - landing priorities, air refueling, divert fields
  - quick access to emergency procedures

CAT

- Shipboard personnel
- 4 Perq computers
- Rule-based expert system written in OPS 83
- Threat analysis

Integration?

- "The James T. Kirk Memorial Bridge"
  - three keyboards
  - four CRTs
  - two mouse pads
- Data communication links
  - VAX-HP-Perq-Wang
  - security
- Hardware reliability
- Training personnel (input and use)

Other "Communities" Are Now Building "Battlefield" Displays

- Hardware/software developments
  - geographic information systems (GIS)
- Availability of geographic information
  - federal sources
    - » census data: geographic and demographic
    - » environmental data
    - » satellite
  - state and local
- Tactical and strategic decisions

Geographic Information Systems: "More than a map"

Geographic objects (points, lines, polygons) are an integral part of a database, each can have it own set of attributes

- For example, a road segment (a line) ...
  - width/length
  - traffic flow peak/off-peak
  - date last paved
  - traffic tickets issued
  - garbage truck route identification

Data Layers

- Geographic information is stored in layers
  - geographic: rivers, roads, elevation contours
  - political: census blocks, voting districts
  - social: 911 phone calls, welfare families
- A "workbook" is created by overlaying these data layers
- Data is accessed with mouse clicks

Mine Subsidence in Pittsburgh

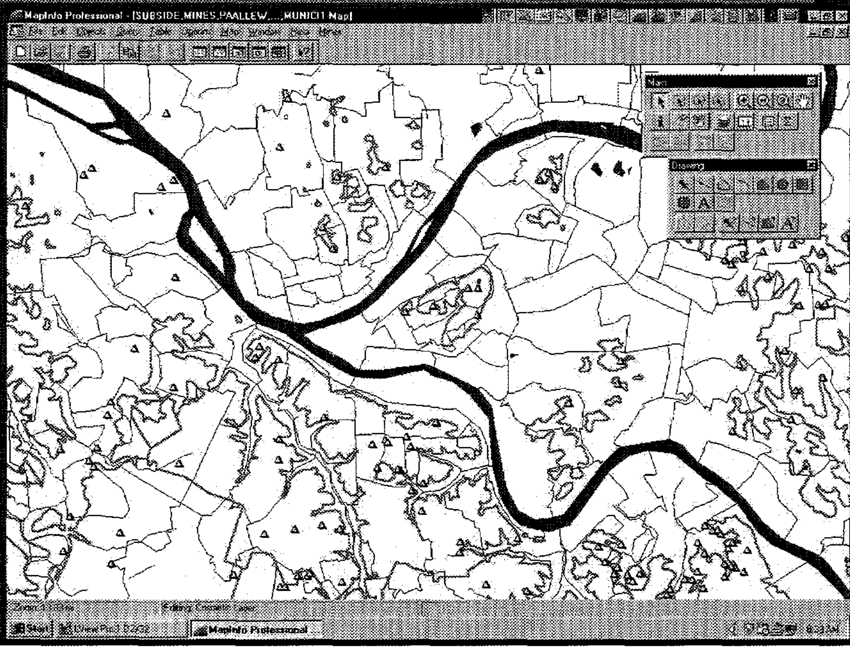

Data Access (Mine subsidence event)

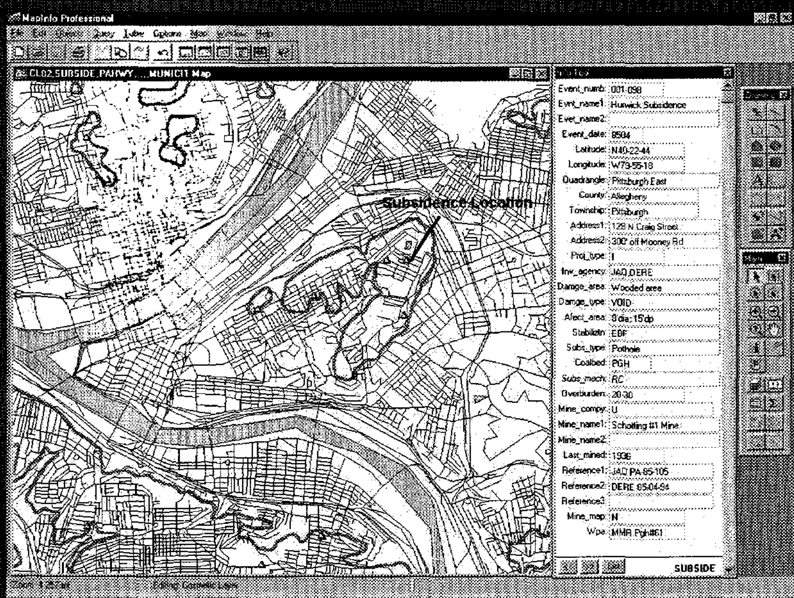

Integrated GIS DSS

- Simplified interface
  - comfortable, understandable
- Tied to other progams
  - calculations/optimzation
- ... and other information sources
  - text
  - web
  - graphics

DSS for Mine Insurance Decision

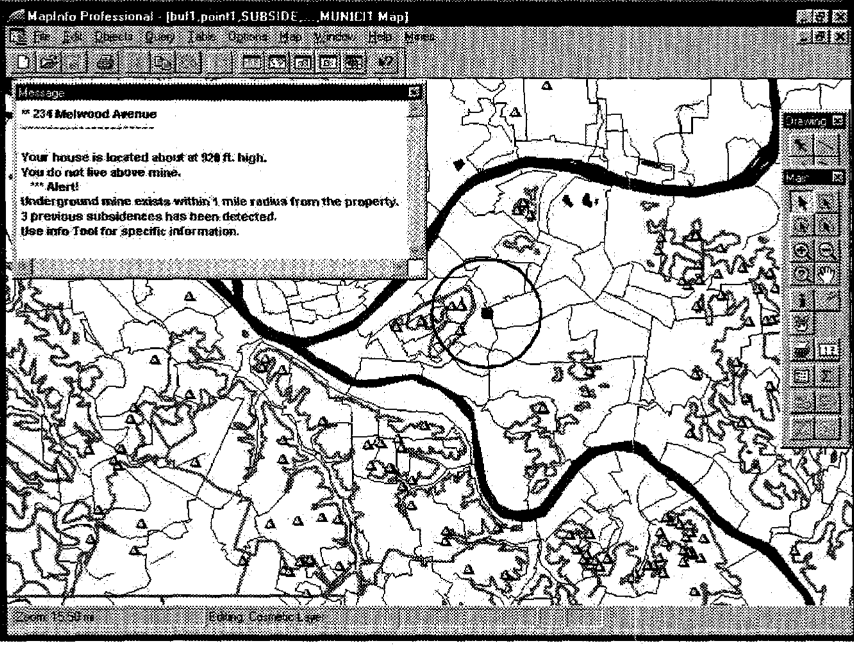

DSS for Mine Insurance Decision

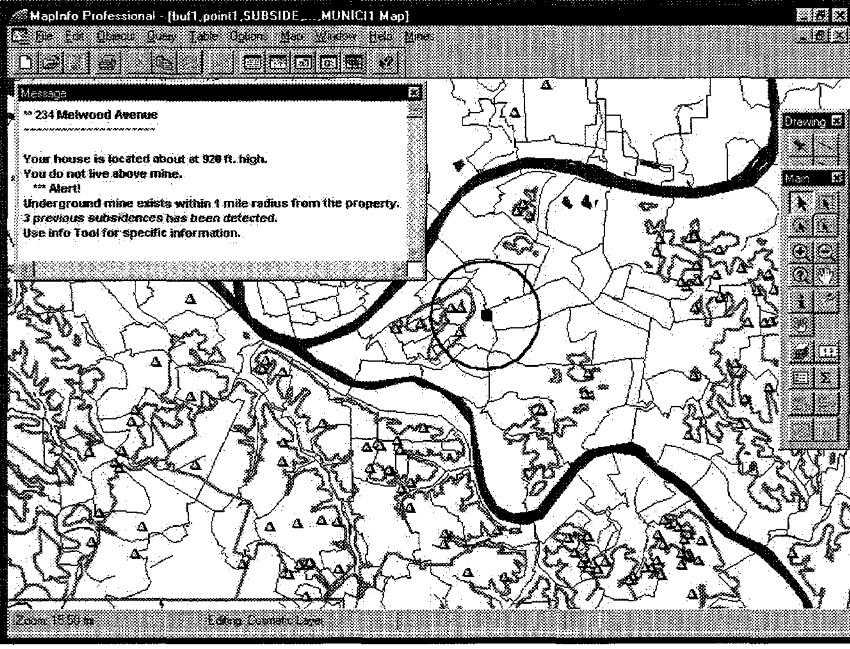

Subsidence Insurance

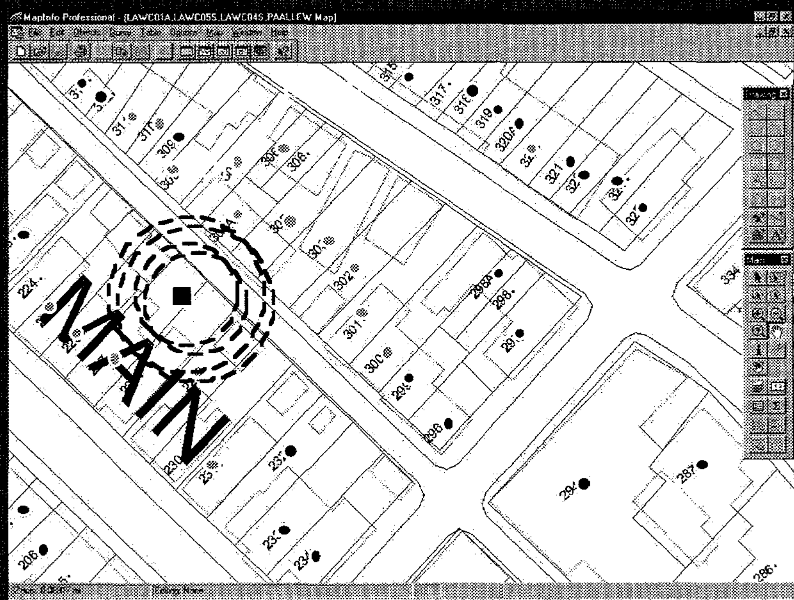

Anticipate Risks

- Use physical and social models to predict problems
  - physical models alone are not sufficient
  - interaction of people with the environment
- Traffic congestion and travel times
- Need for new integrated models

Flood Plain

Location of Emergency Services

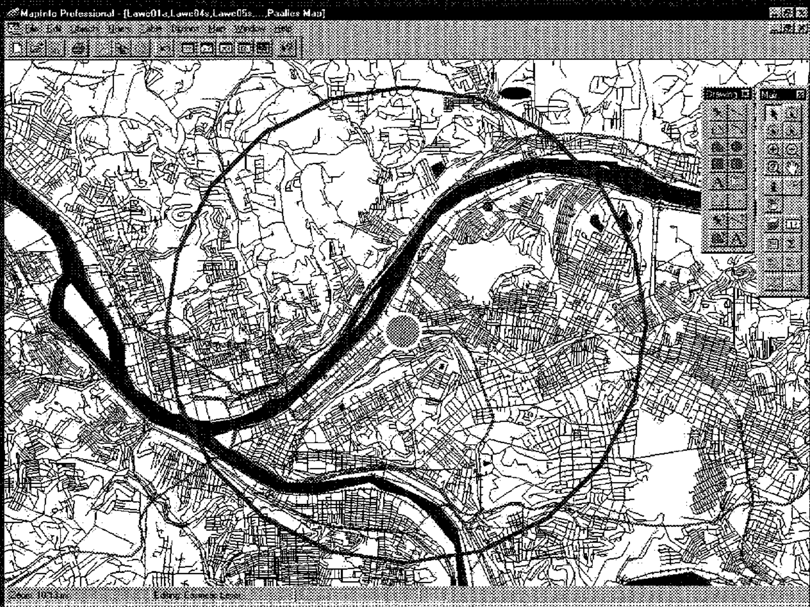

Need for Integrated Optimization

- Shortest path
- Routing vehicles
- GPS
- Monte Carlo simulation of traffic flows
- What-if analysis
- Reduced-form models for real-time analysis
  - neural nets

Example: Network Visualization

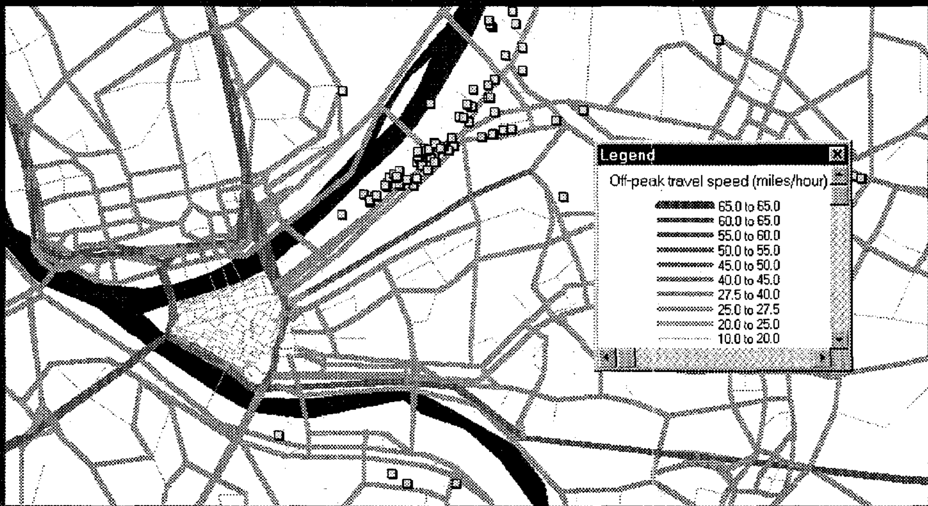

Shortest Path: (with/without ...)

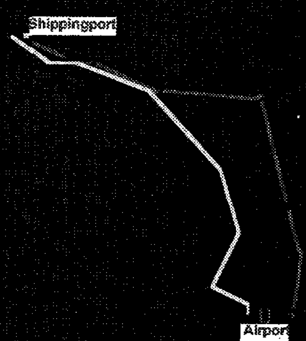

Travel Times: (Peak and Off-peak)

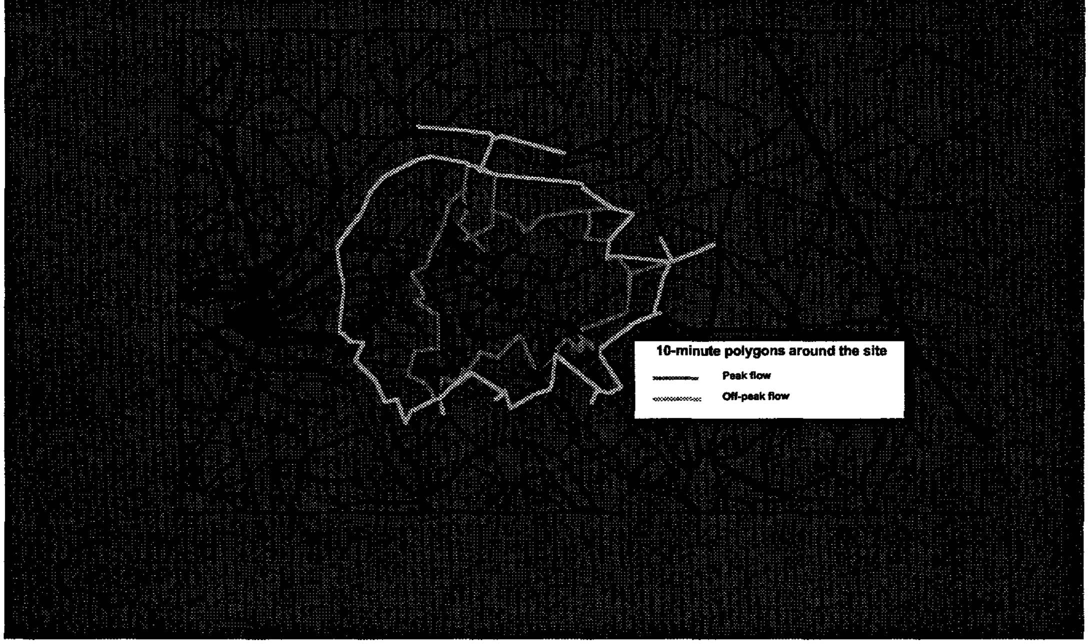

Shortest Path Analysis:
Time Between Crime and Arrest

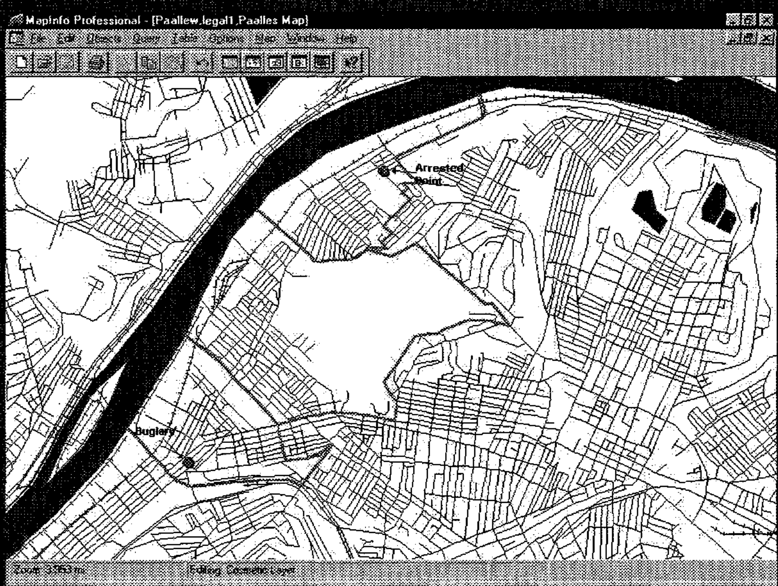

Routing of City Services:
Shortest Path for Garbage Pick-up

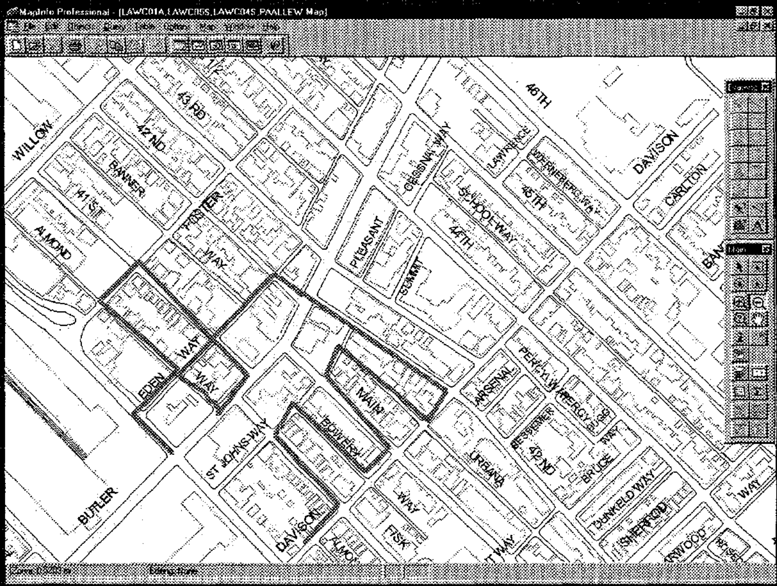

Block Watch Effectiveness

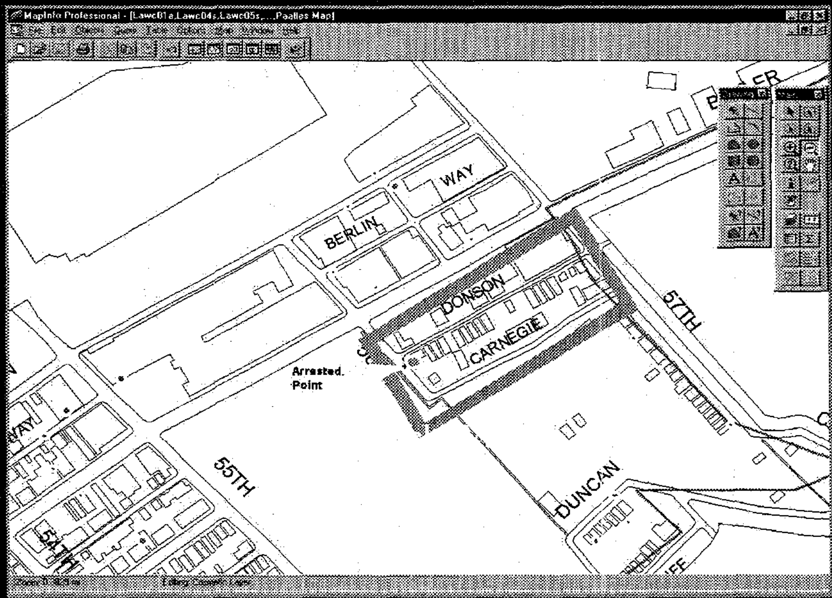

District Analysis (i.e., Zoning, Voting)

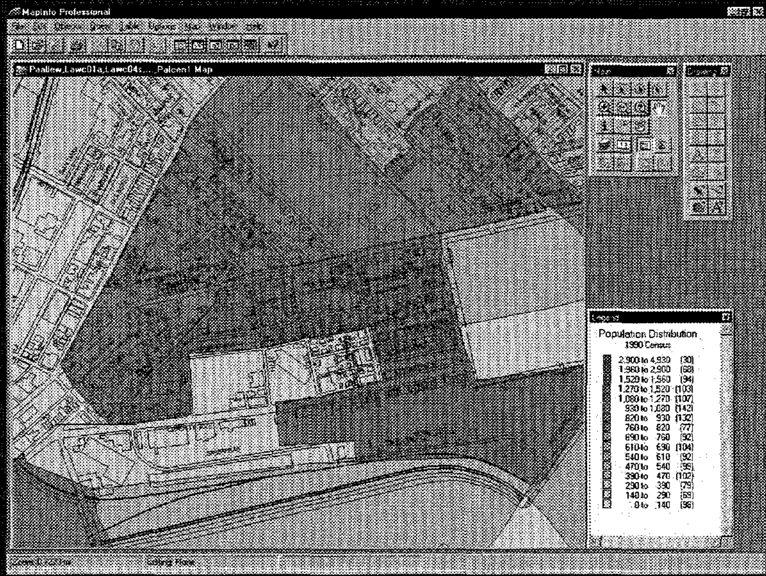

Neighborhood Analysis

Building Footprints:
100 megabyte file

Understanding Local Risks

- Risk equity issues
  - effect on property values
- Risk perception
  - community assessment
  - risk communication
  - legal problems
- Selection of mitigation strategies
  - which objective function?

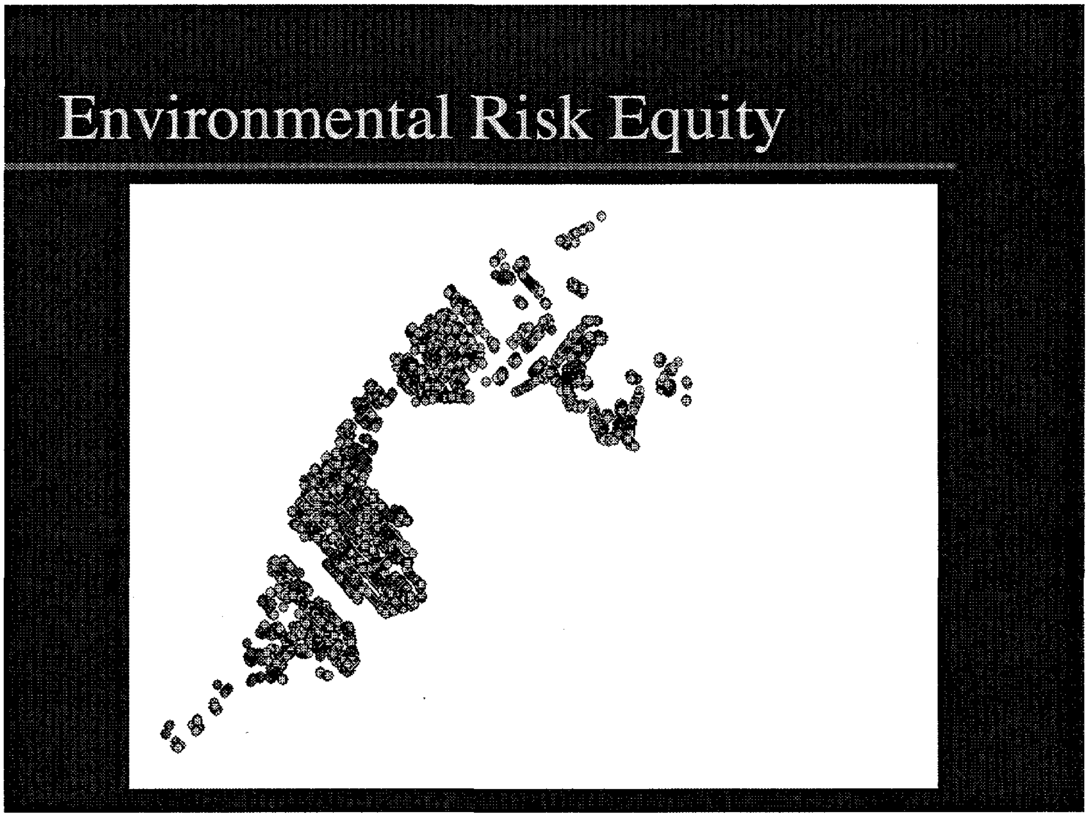
Environmental Risk Equity

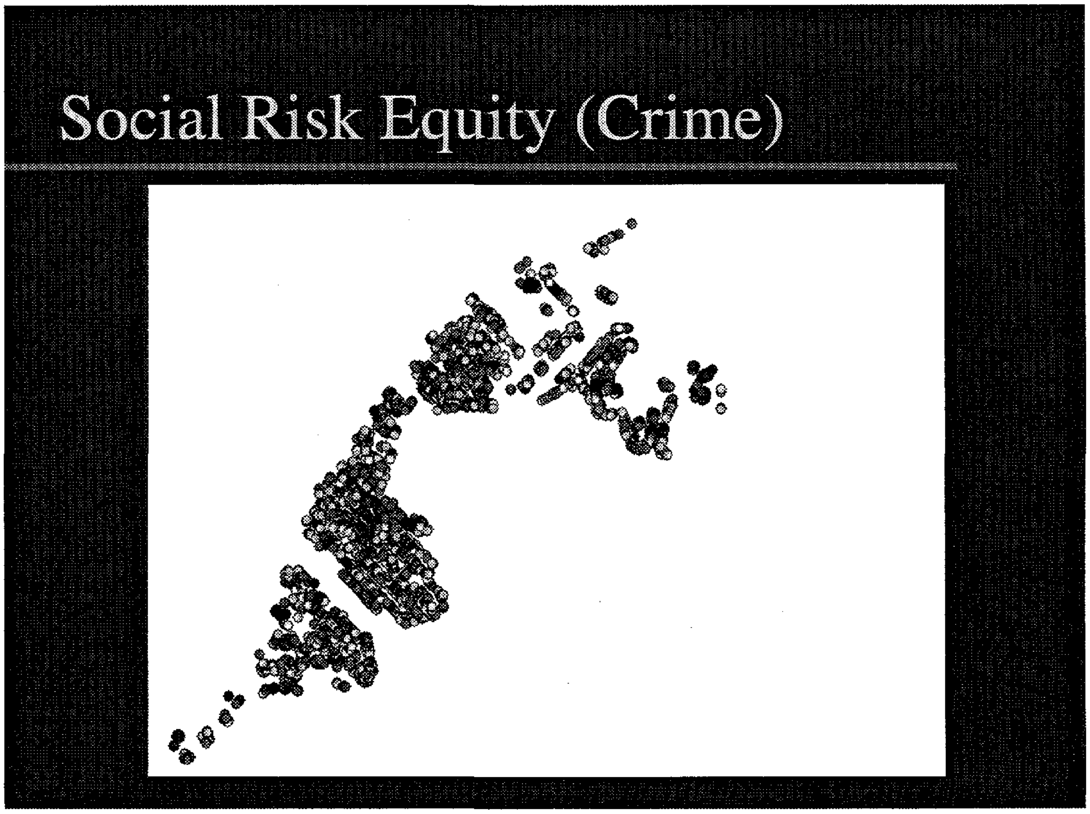
Social Risk Equity (Crime)

RISES:
Regional Industrial Site Evaluation Systems

- Integrate DSS/GIS
- Analysis of industrial (and other) sites
  - demographic
  - environmental
  - economic
  - transportation
- Internet presence
  - Access to other relevant data

Demographic & Economic Analysis

- Description of population in the vicinity of a site
- Based on 1990 Census data
  - Age, education, work force, education, etc.
- Graphical display of the information
- Statistical comparison across sites
- Market/labor force characteristics
- Availability of production inputs/outputs
  - SIC
  - Warehouses/transportation hubs
- Impact of a successful siting
- Value of infrastructure improvements

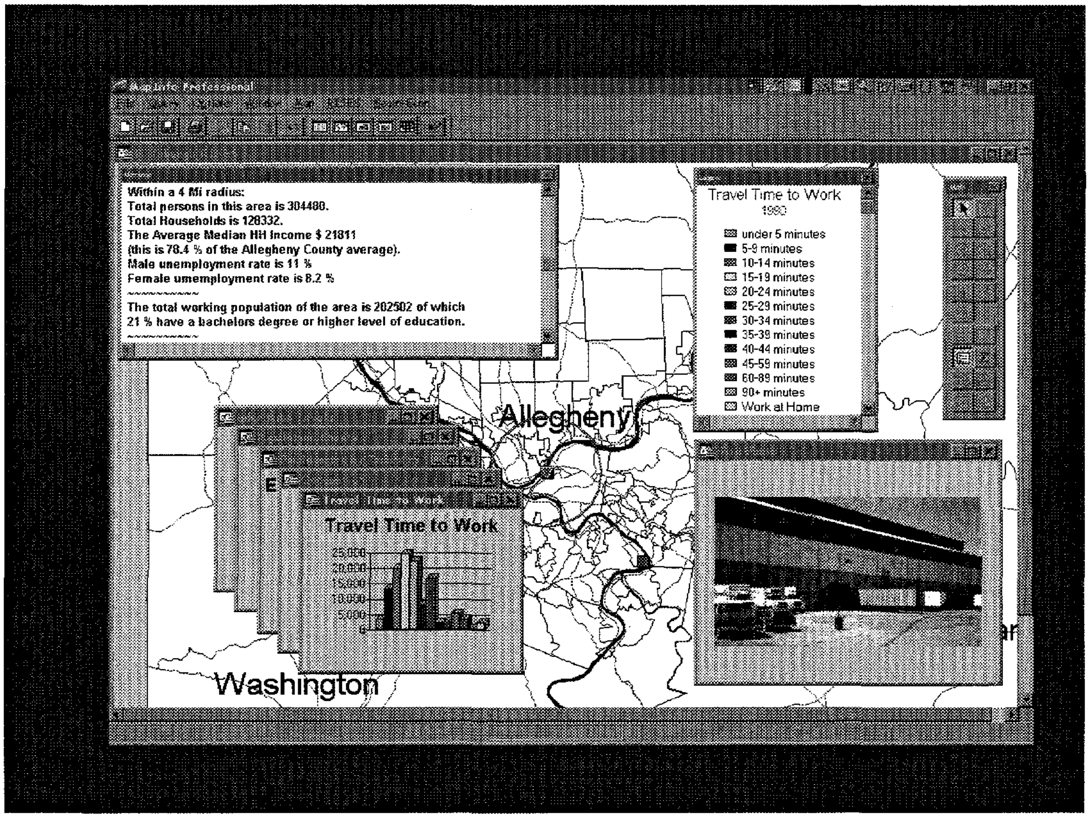

Environmental Analyses

- Preliminary Phase I site assessment
  - ASTM standards
  - RCRA, CERCLIS, UST, LUST, TRI, etc.
- Transaction Screen Process
- Historic environmental
  - Sanborn
  - Site specific data: previous use by SIC

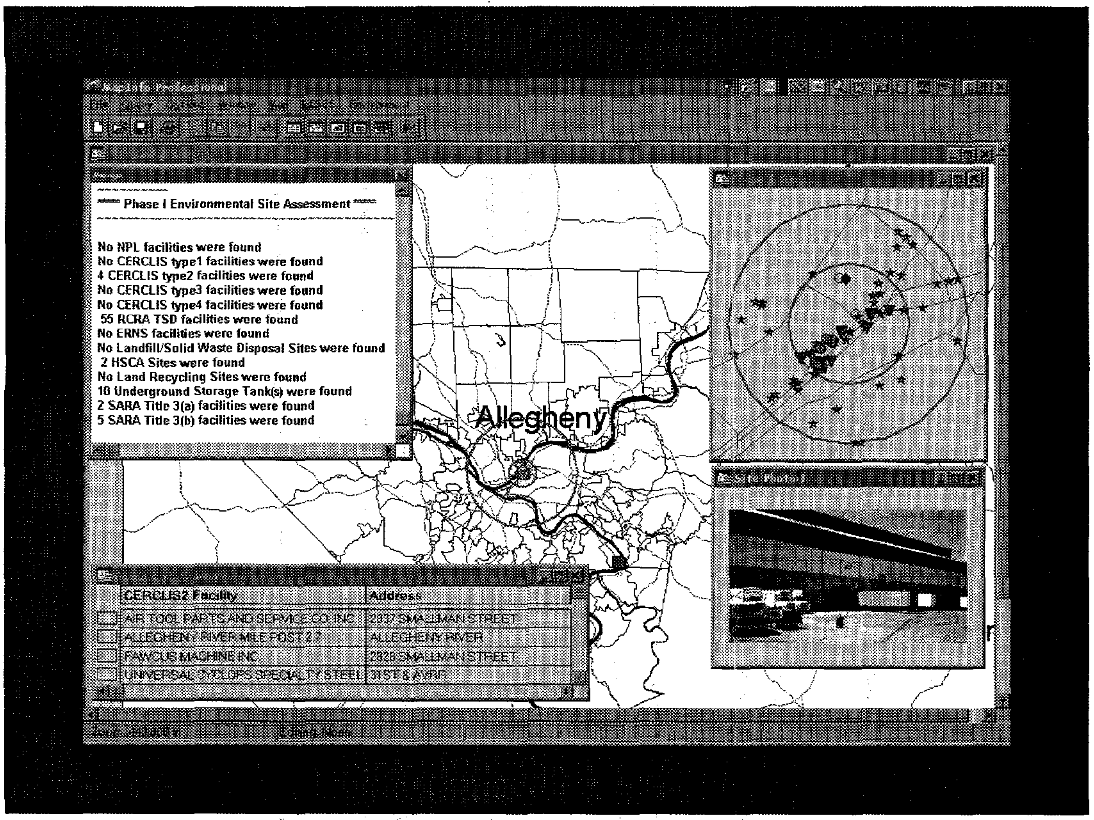

Production Inputs/Outputs

- List and display suppliers, customers, competitors
- Search by 4-digit Standard Industrial Classification
- Identify industry clusters, geographically target strong markets

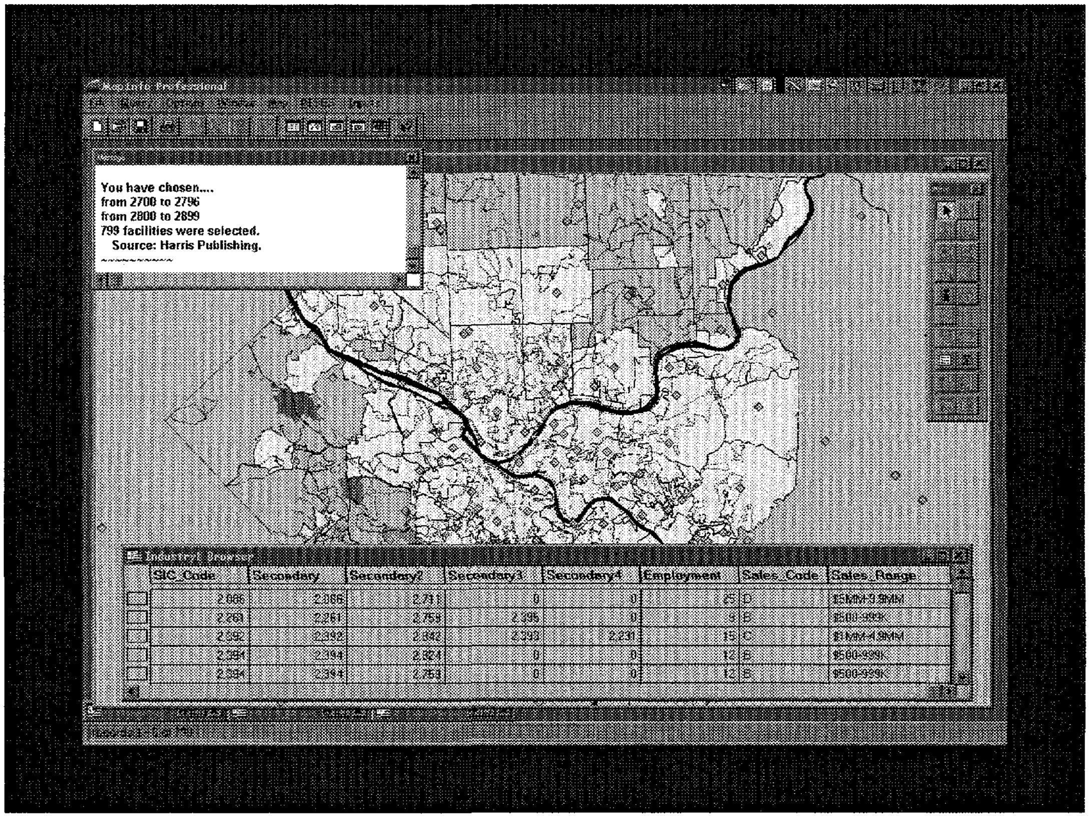

IISIS

- Interactive, intelligent, spatial information system for disaster management
- University of Pittsburgh
  - Louise K. Comfort, Principal Investigator
  - Mark W. Dunn, System Administrator
  - Bruce G. Buchanan, Consultant
- Combining GIS displays with expert systems and communication over the Web

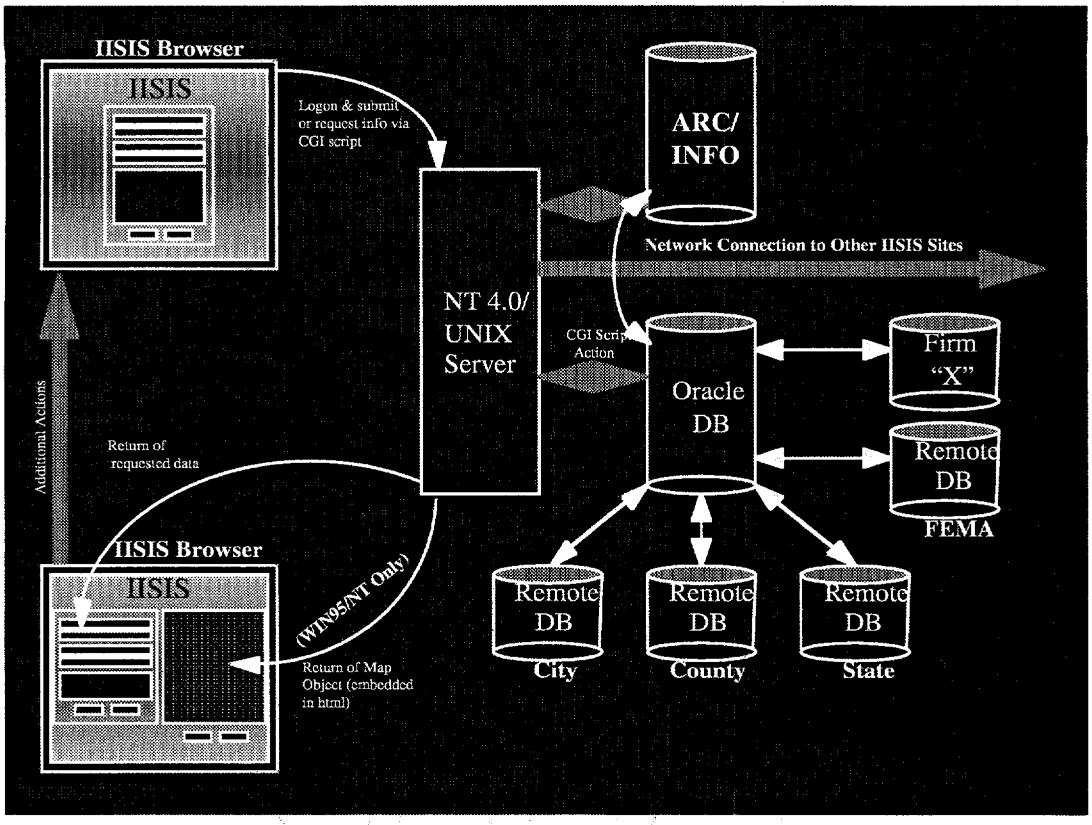

Platform Issues

PROs

- NT 4.0 is rapidly being adapted as a standard in many organizations
- NT WWW Server is easy to administer and scripts for proto will be easier to develop (less programming)
- MapObjects, necessary for multiple GIS displays, operates on Win95/NT 4.0 platforms only
- Cost for NT 4.0 Oracle Server is affordable

CONs

- At present, project has no hardware which will meet performance demands
- MapObjects will cost $500-800 to implement (aside from training costs and time for training)
- Independent of platform issues, Oracle Database will be necessary to develop working knowledge base

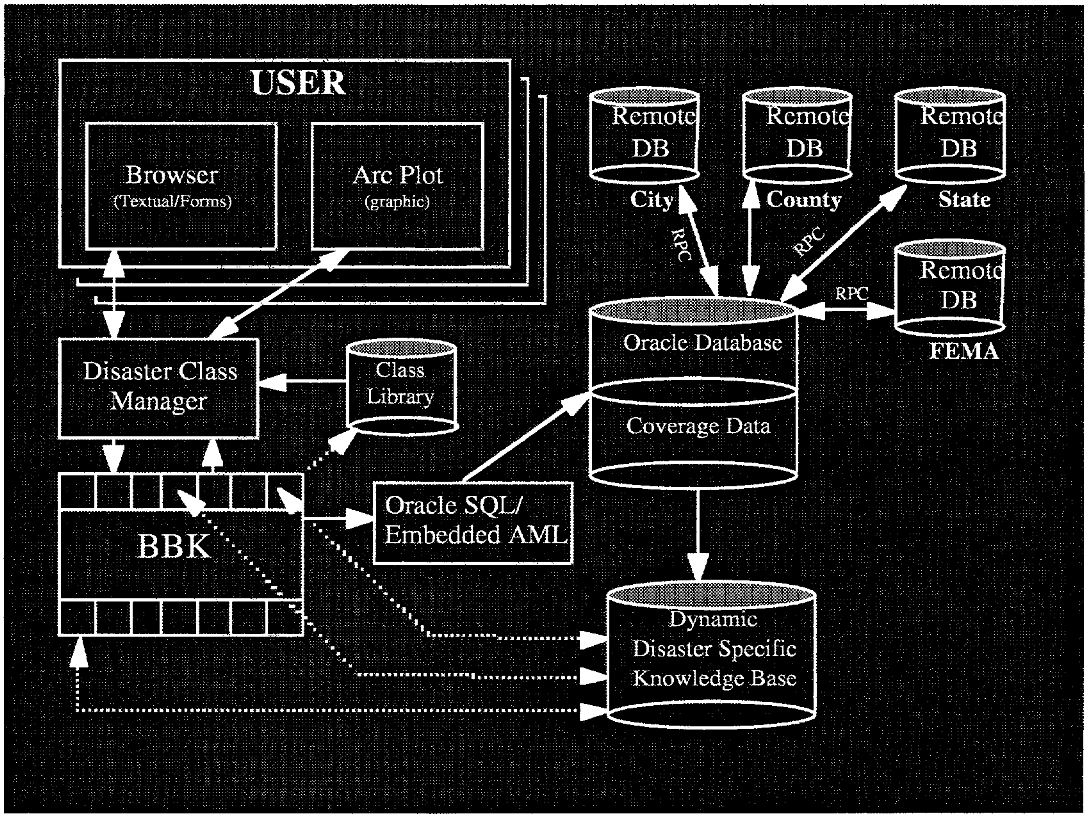
USER
Browser
(Textual/Forms)
Arc Plot
(graphic)
Remote DB
City
Remote DB
County
Remote DB
State
Remote DB
FEMA
RPC
RPC
RPC
Oracle Database
Coverage Data
Disaster Class Manager
Class Library
BBK
Oracle SQL/ Embedded AML
Dynamic Disaster Specific Knowledge Base

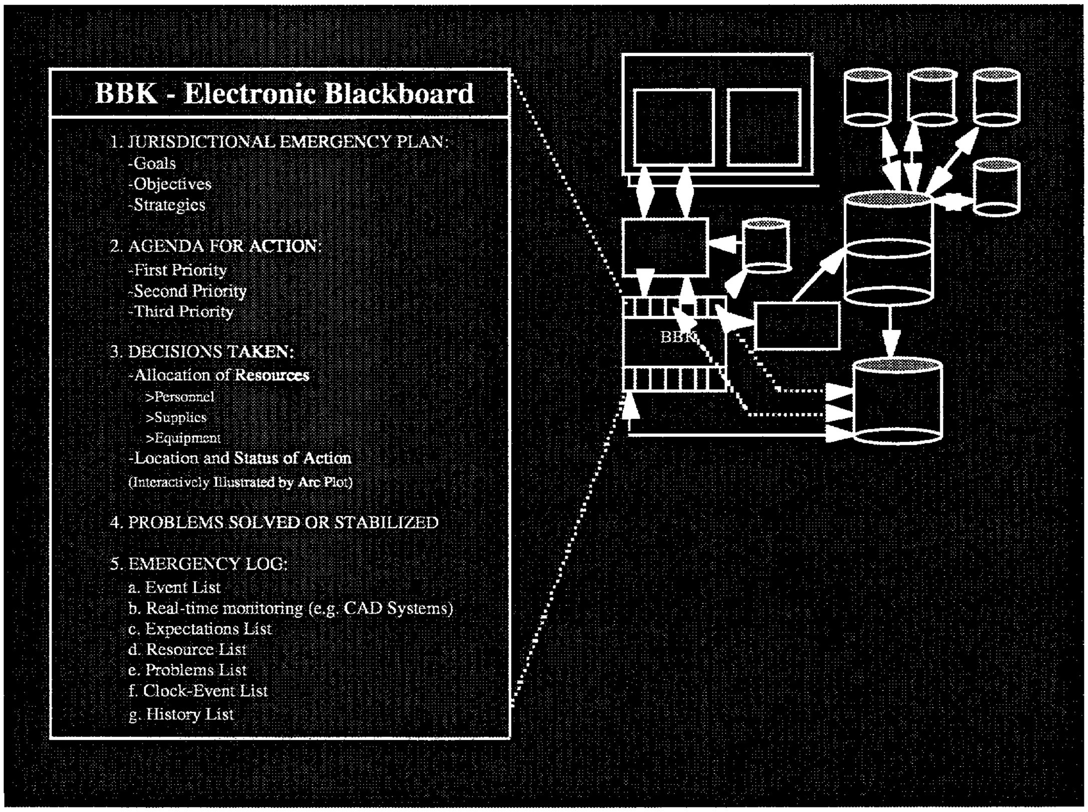
BBK - Electronic Blackboard
1. JURISDICTIONAL EMERGENCY PLAN:
-Goals
-Objectives
-Strategies
2. AGENDA FOR ACTION:
-First Priority
-Second Priority
-Third Priority
3. DECISIONS TAKEN:
-Allocation of Resources
>Personnel
>Supplies
>Equipment
-Location and Status of Action
(Interactively Illustrated by Arc Plot)
4. PROBLEMS SOLVED OR STABILIZED
5. EMERGENCY LOG:
a. Event List
b. Real-time monitoring (e.g. CAD Systems)
c. Expectations List
d. Resource List
e. Problems List
f. Clock-Event List
g. History List
BBK

Responding to an Oil Spill

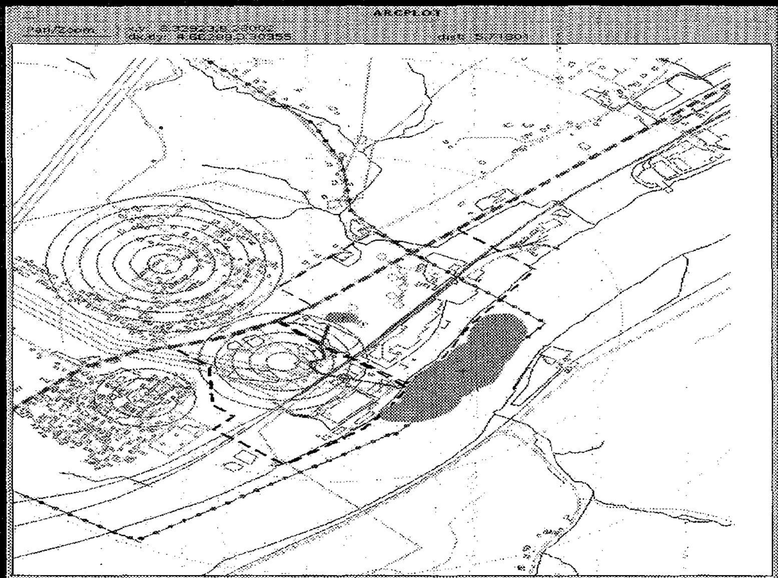

Conclusions

- Technology transfer from the military
- Need for research and tool development
  - data capture and storage
  - database design
  - display of information
  - communication
  - expert system development
  - reduced-form models and datasets
  - uncertainty analysis

Conclusions (continued)

- Cost of system is critical
  - acquisition, maintenance, training
  - hardware, software, and data
- Standardize the process so it can be replicated
- Develop reduced-form models and datasets
- Build on existing (free) datasets
  - socializing the pixel

Crisis Forecasting

Andy White
Salishan, OR
April 22, 1997

Advanced Computing Laboratory

How did all of this start?

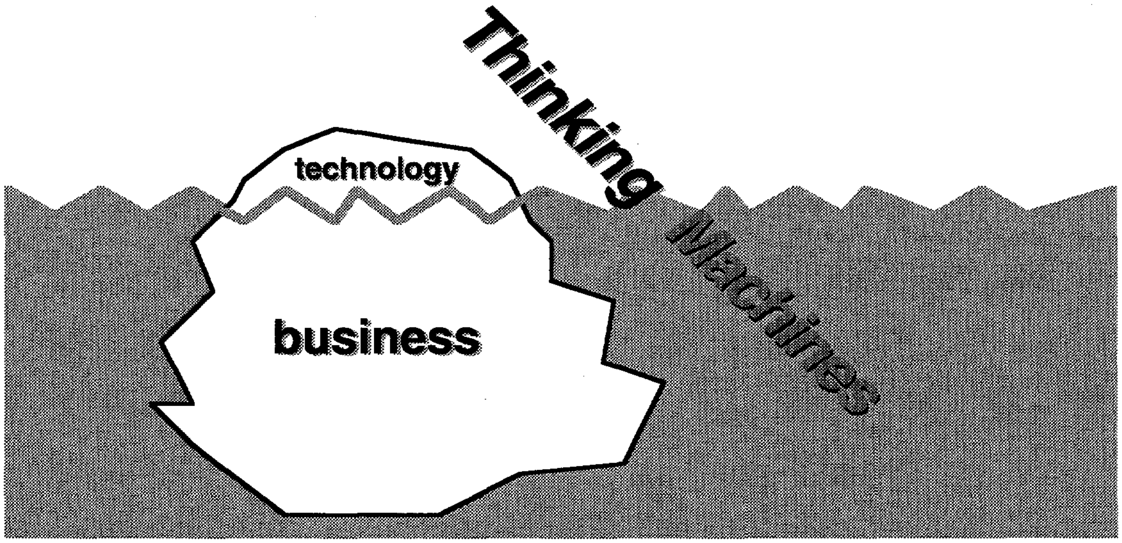

Advanced Computing Laboratory

Algorithms • Architecture • Language

1994: "There is good news and bad news."

- Availability of advanced computing resources has increased a thousand fold
- Dramatic improvements in fidelity of physical models has been demonstrated (e.g. ocean models)
- New techniques are available (e.g. AMR, e^convergent schemes, adaptive systems, wavelets)
- Marketplace and vendor population have declined
- Urgency of solving traditional problems has eroded

Advanced Computing Laboratory

1994 Marketspace

- No room for new (e.g. TMC) or established (e.g. CRI) companies exclusively at the high end

	1993	1998
Supercomputers	$2.0B	$4.0B
High-end systems	$0.8B	$0.5B

19
$3

- Commercial success must be achieved within a single hardware generation
- Focus has been on computing *fast*, not solving problems
- Virtually no interest from the ISVs

Advanced Computing Laboratory

Technical Computing Market*

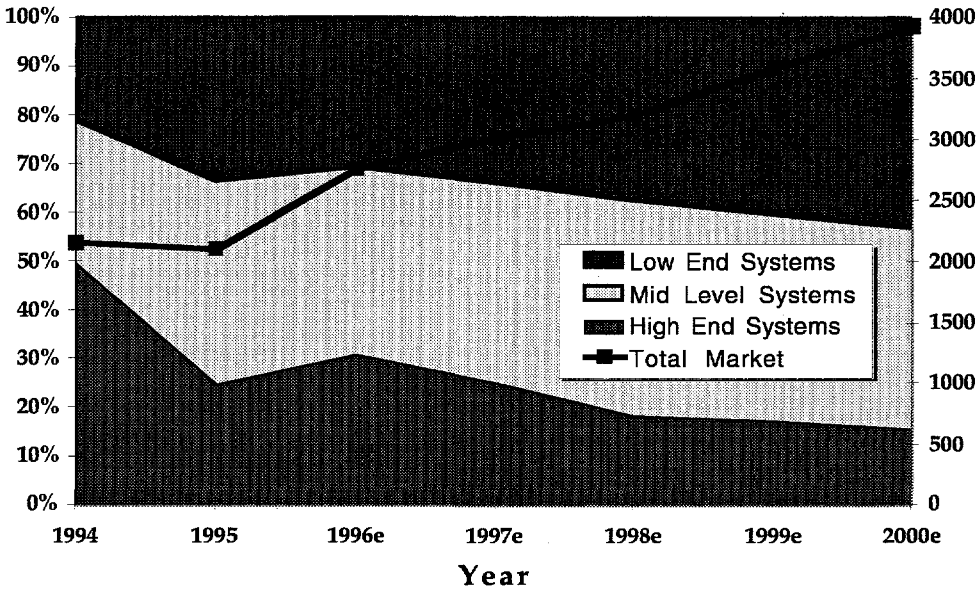

* International Data Corporation

Advanced Computing Laboratory

Recent Events in Marketplace

- Cray Computer Corporation bankrupt
- Convex acquired by HP
- nCUBE targets video-on-demand market
- Thinking Machines targets software
- MasPar suspends hardware development
- Intel eliminates scientific computing business sector
- Cray Research acquired by Silicon Graphics
- ECMRWF chooses Fujitsu, NCAR chooses NEC

Advanced Computing Laboratory

1994 Appspace

- Cold war competition is over
- The arcane and obscure have been our focus in past
- Predictive modeling and simulation unknown
  - Industry interested in capability, but not willing to take risks on hardware time scales
- Software machine slow in emerging
- Machines scale down not up
- Half-ass-good solutions, cost one tenth as much

Advanced Computing Laboratory

Conclusions & Future Directions

- The marketplace must be independent and yet we must have a reliable, domestic supply of high performance computing resources
- We should focus on the solution of complex, time-critical problems of national importance requiring predictive modeling and simulation
- The applications development cycle must not be slaved to the computer system development cycle

Advanced Computing Laboratory

Conclusions & Future Directions

- The marketplace must be independent and yet we must ... of high performance computing resources
- We should focus on the solution of complex, time-c... ...ance requiring predictive modeling and simulation
- The applications development cycle must not be sla... ...development cycle

- **ASCI & Nirvana Blue**
- **Crisis Forecasting**
- **POOMA, PAWS, ...**

Advanced Computing Laboratory

Why Crisis Forecasting?

- **Most challenging modeling and simulation problems are time-critical, complex, extrapolatory**
- **Affects lives and property**
- **Requires High*est* Performance Technology**
  - **- data assimilation (interactive, multi-modal input)**
  - **- urgency (speed)**
  - **- fidelity (methods, data verisimilitude, memory)**
  - **- information delivery (context agile output)**
- **Provides a demonstrable, challenging, visceral proof of predictive modeling and simulation**

Advanced Computing Laboratory

"For a variety of reasons, ... science and scientists have fallen into the habit of separateness ... This separation, which reflects both isolation and autonomy, nullifies a great many contributions that scientists and engineers can make to our larger societal goals."

Neal Lane, Director, NSF

"If we are to reverse the substantial momentum that lies behind proposed cuts in science and technology, we need to engage many different communities: the R&D community, the industrial community, the education community -- ultimately the entire American populace, because they are the ones who will suffer the most from these changes."

Jack Gibbons, Director, OSTP

A Spectrum of Possibilities

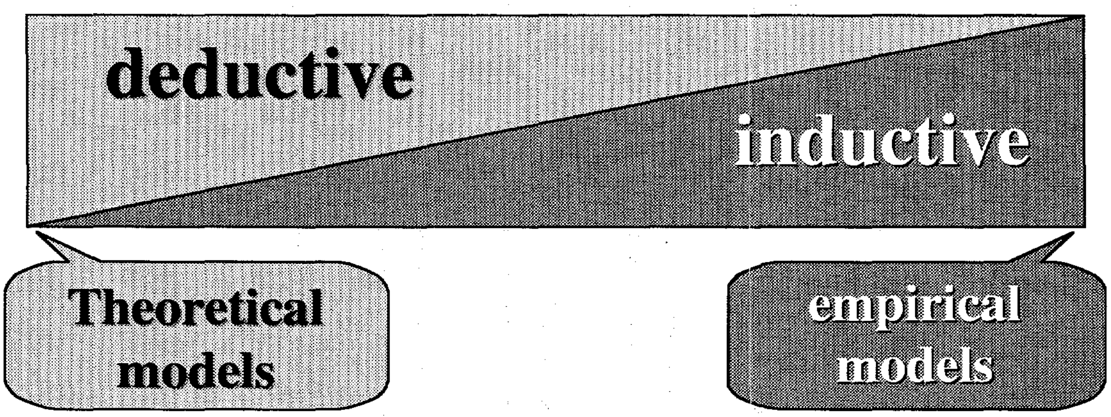

Advanced Computing Laboratory

Pick one from column A and ...

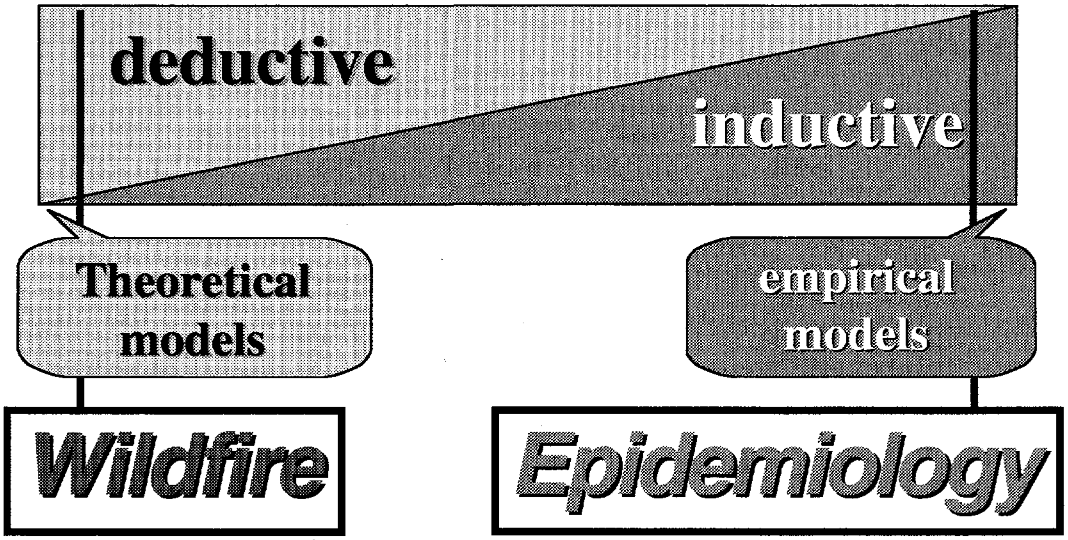

Wildfire

Modeling of wildland fires including weather, fuel state, combustion, topography, and surveillance

- ***Examples:***
  controlled burns: Kennedy Space Center
  wildlands: Glenwood Springs, CO
  wildland-urban interface: Calabassas Canyon, CA
- ***Partners:***
  Forest Service Fire Labs, Los Angeles County Fire Department, Fish and Wildlife, Dynamac, NASA, Air Force

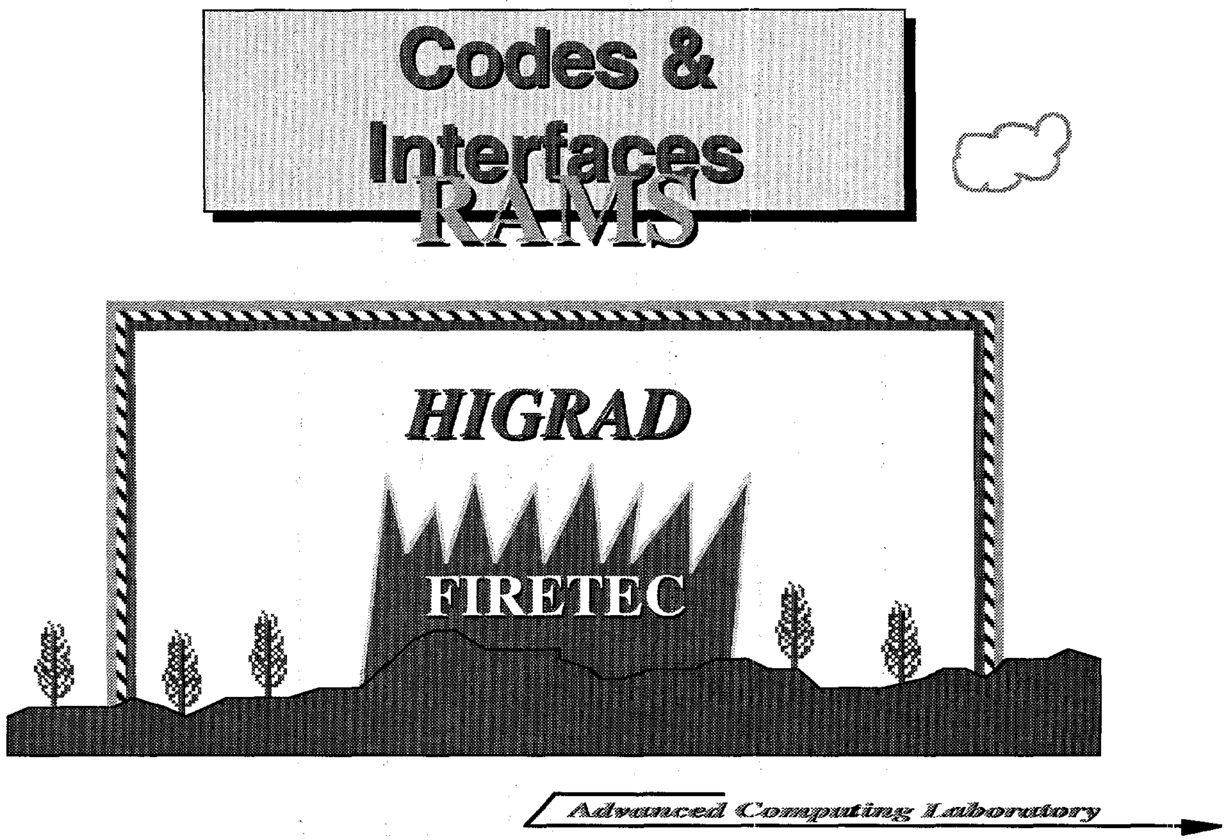
Codes &
Interfaces
RAMS
HIGRAD
FIRETEC
Advanced Computing Laboratory

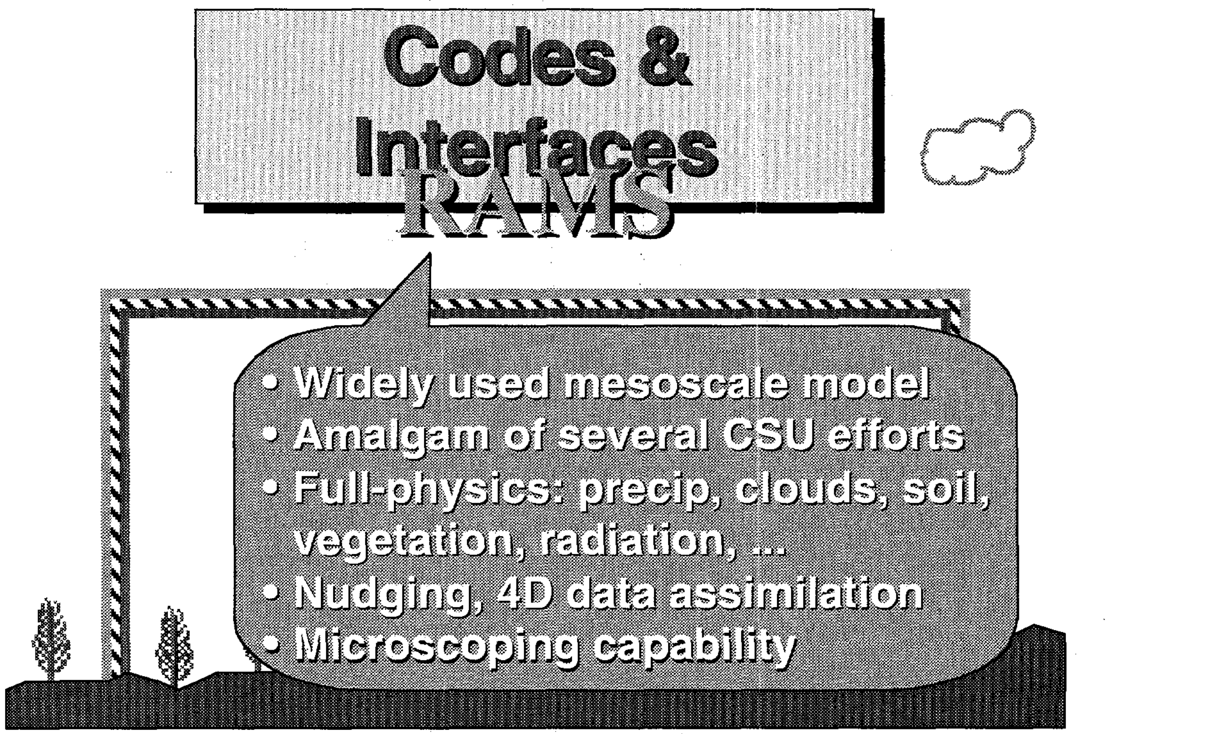
Codes &
Interfaces
RAMS
• Widely used mesoscale model
• Amalgam of several CSU efforts
• Full-physics: precip, clouds, soil, vegetation, radiation, ...
• Nudging, 4D data assimilation
• Microscoping capability
Advanced Computing Laboratory

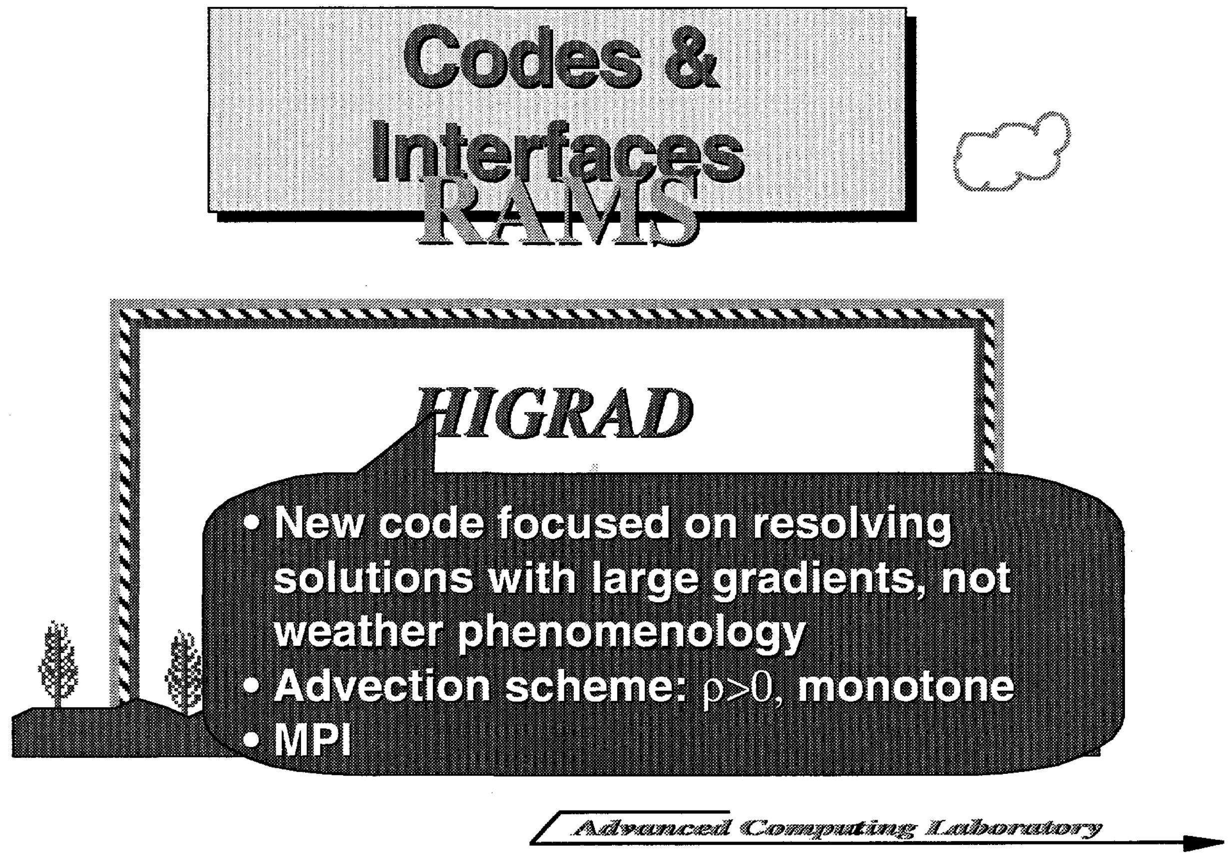
Codes &
Interfaces
RAMS
HIGRAD
• New code focused on resolving solutions with large gradients, not weather phenomenology
• Advection scheme: $\rho>0$, monotone
• MPI
Advanced Computing Laboratory

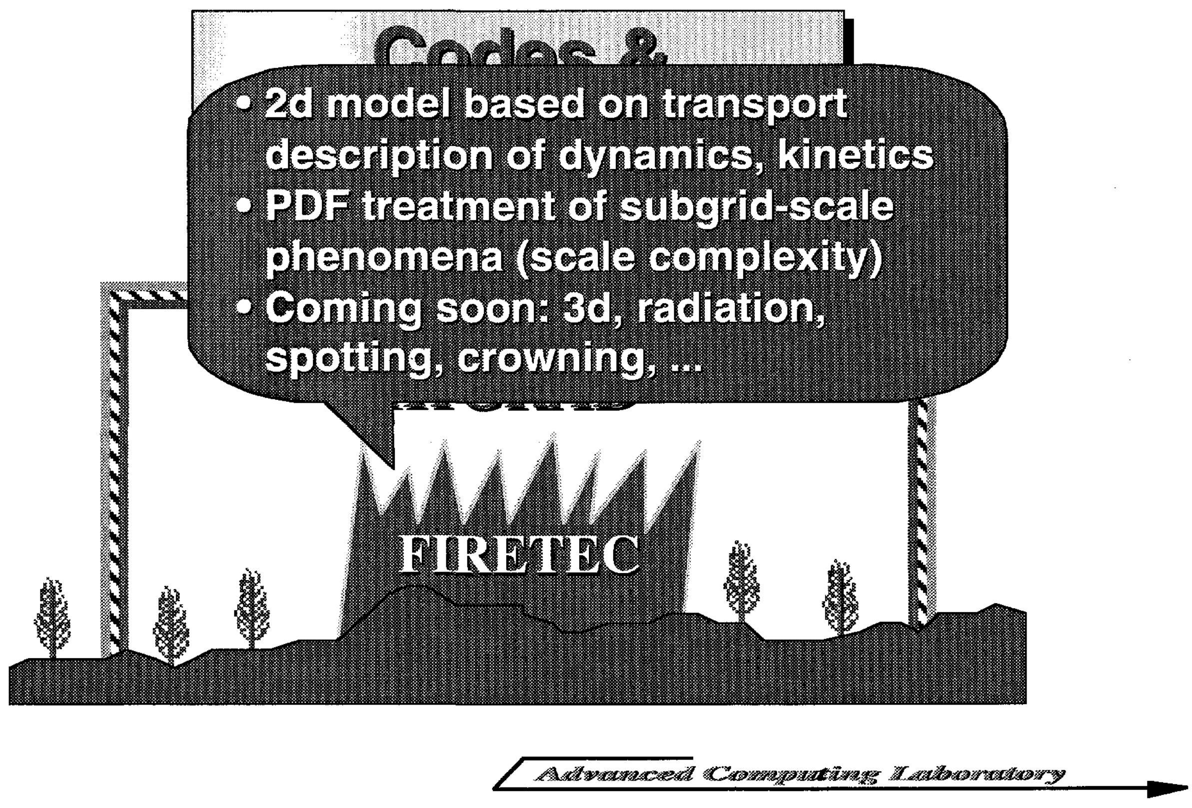
• 2d model based on transport description of dynamics, kinetics
• PDF treatment of subgrid-scale phenomena (scale complexity)
• Coming soon: 3d, radiation, spotting, crowning, ...
FIRETEC
Advanced Computing Laboratory

Another view

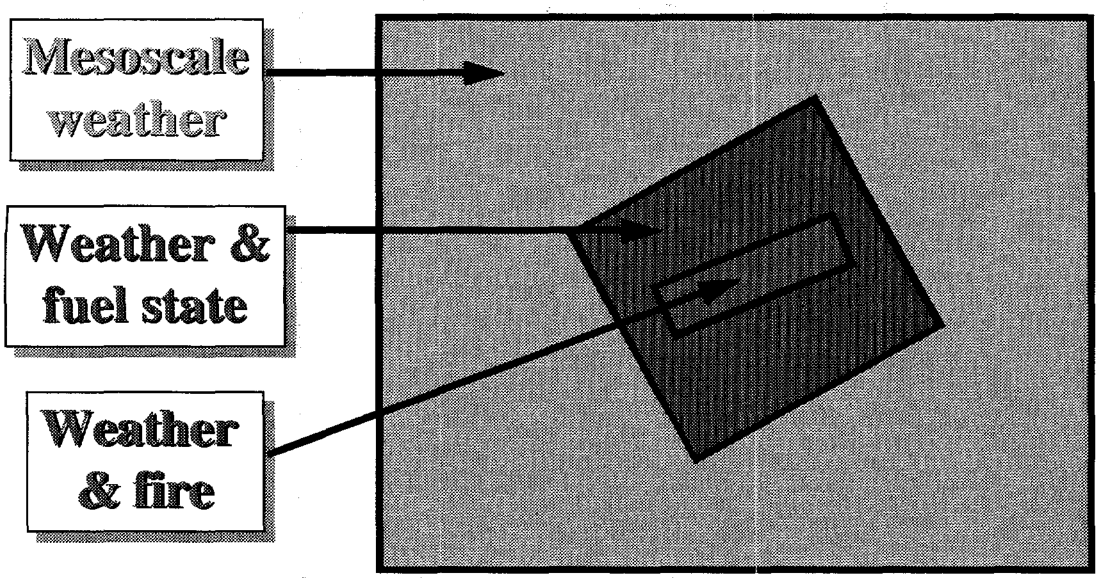

Operational Requirements

Speed

- ***Merely*** **predicting the course of a wildfire: 12x Δt^{real}**
- **Scenarios for decision making: 720x Δt^{real}**
- **Evaluation of objective function: 8640x Δt^{real}**

Fidelity

- **The required resolution is dependent on the accuracy of the initial, boundary, and in-flight data**
- **Missoula Fire Lab tests, Kennedy Space Center and Los Angeles County controlled burns**

Epidemiology, Influenza

Predict the molecular changes in the influenza virus & and the course of pandemic outbreaks

- ***Data***
  100 sentinel stations world-wide; 150 direct US reports; 300 human, 200 avian sequences
- ***Annual Effects***
  $5B in medical costs; $12B in lost productivity; 20,000 excess deaths per year
- ***Extraordinary biological threats***
  natural pandemics, weaponization

Influenza A

- **Characterized by its coat proteins, hemagglutinin (HA) and neuraminidase (NA)**
- **Focus is on hemagglutinin**
  - **mediates binding and fusion of virus to cell**
  - **principal target of the human immune system**
  - **HA has 550 amino acids**
- **Sequence drift causes 'annual' epidemics in spite of vaccines and previous immune responses**
- **Sequence reassortment causes periodic shifts**

Recent History of Influenza

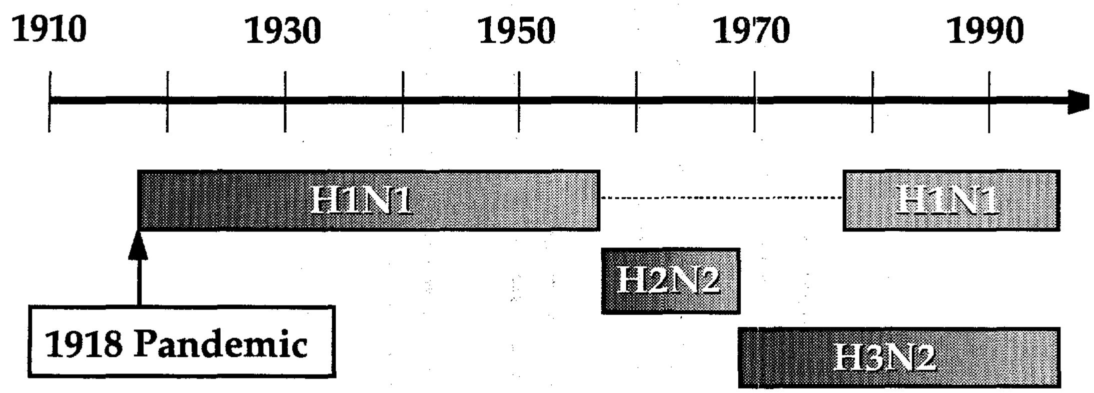

Advanced Computing Laboratory

Ecology of Influenza

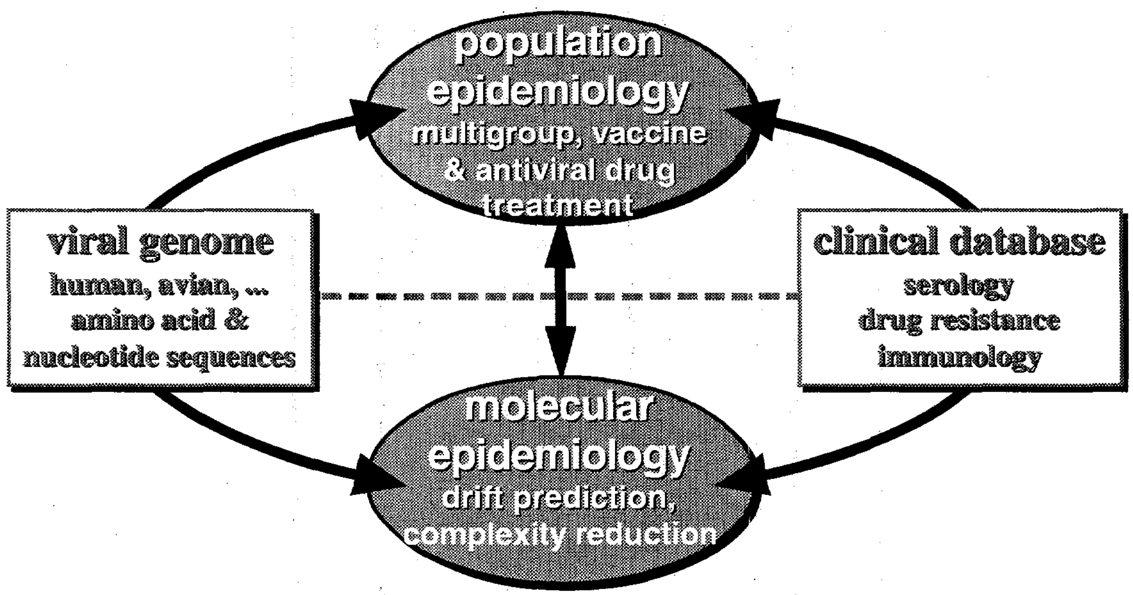

Advanced Computing Laboratory

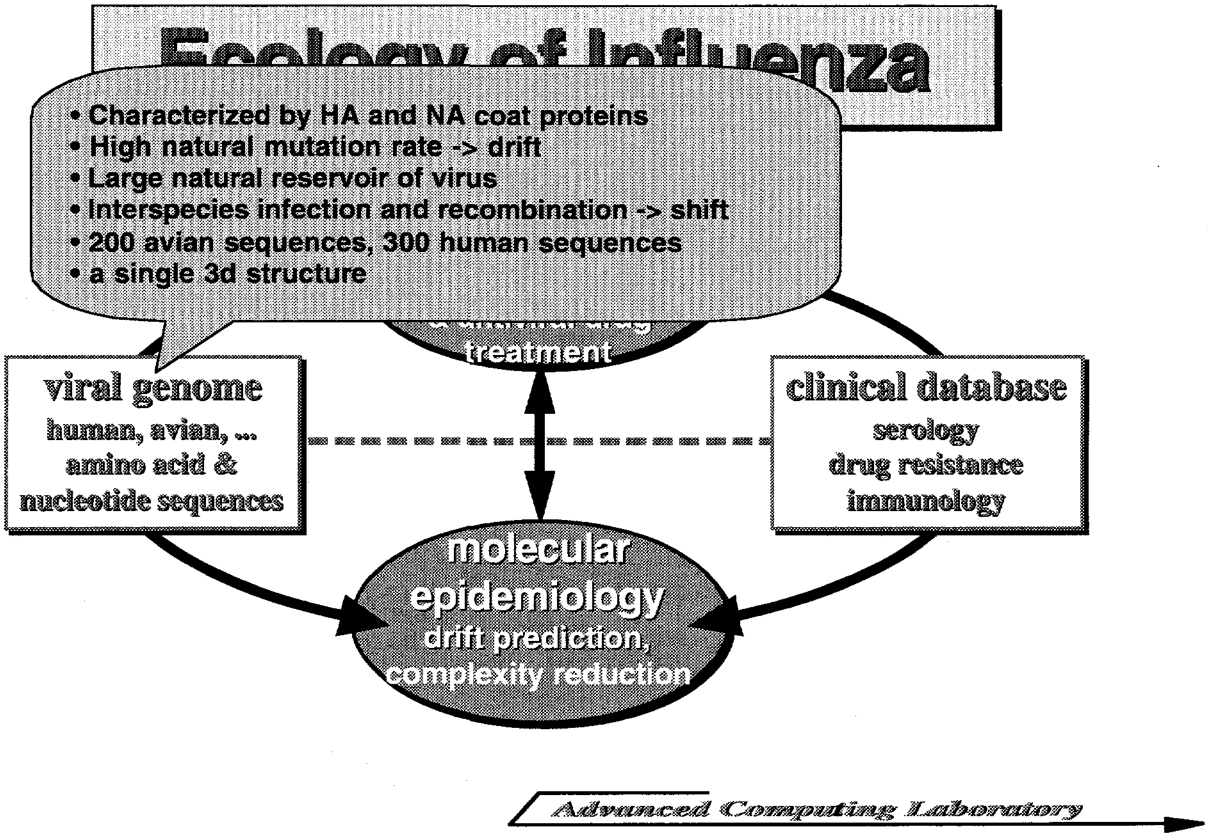
Ecology of Influenza
• Characterized by HA and NA coat proteins
• High natural mutation rate -> drift
• Large natural reservoir of virus
• Interspecies infection and recombination -> shift
• 200 avian sequences, 300 human sequences
• a single 3d structure
treatment
viral genome
human, avian, ...
amino acid &
nucleotide sequences
clinical database
serology
drug resistance
immunology
molecular
epidemiology
drift prediction,
complexity reduction
Advanced Computing Laboratory

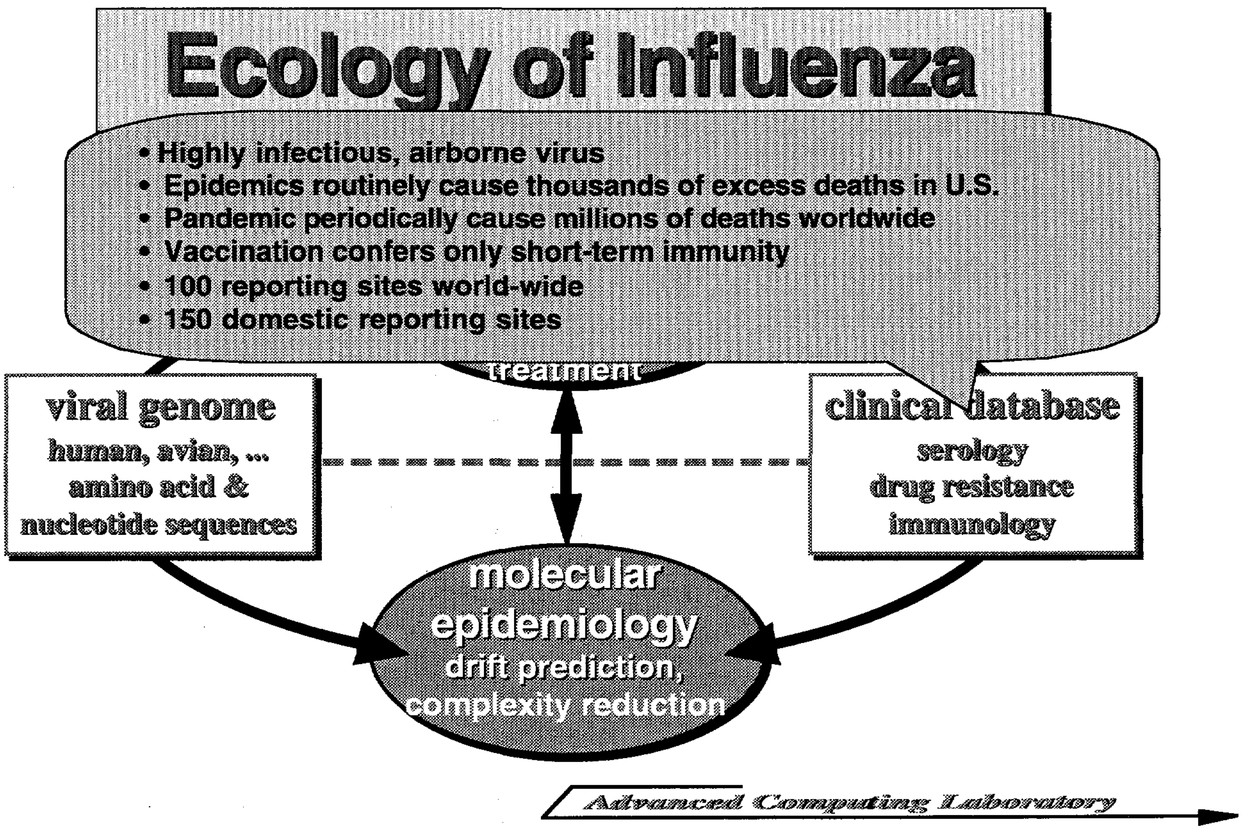
Ecology of Influenza
• Highly infectious, airborne virus
• Epidemics routinely cause thousands of excess deaths in U.S.
• Pandemic periodically cause millions of deaths worldwide
• Vaccination confers only short-term immunity
• 100 reporting sites world-wide
• 150 domestic reporting sites
treatment
viral genome
human, avian, ...
amino acid &
nucleotide sequences
clinical database
serology
drug resistance
immunology
molecular
epidemiology
drift prediction,
complexity reduction
Advanced Computing Laboratory

Ecology of Influenza

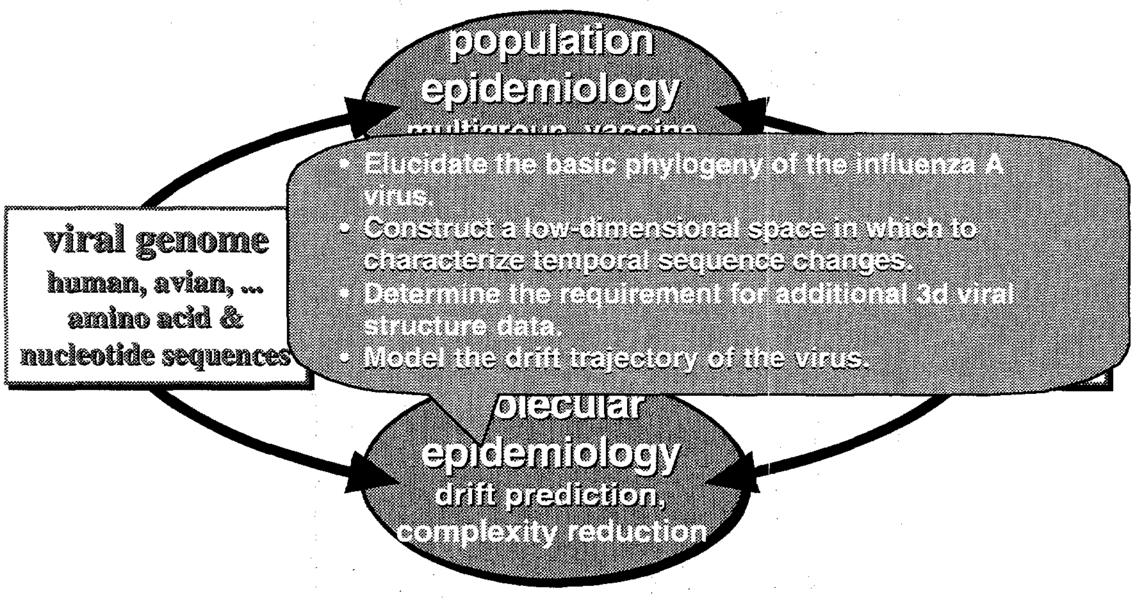

Advanced Computing Laboratory

Ecology of Influenza

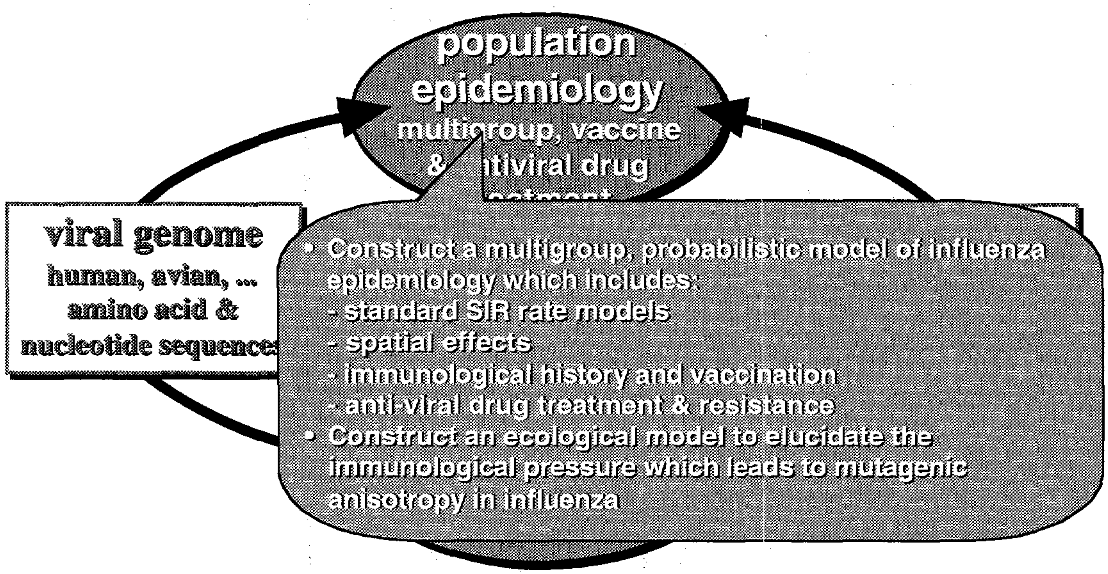

Advanced Computing Laboratory

Taxonomy for Predictive Modeling & Simulation

DESIGN: Able to control solution space and thus maximize stability & use of interpolatory simulations

PROCESS: Unable to control solution space ("take what you get") and thus may be in unstable regions or outside envelope of experiments

TRANSITION: By definition in an unstable region of problem space, likely with little or no prior data

Advanced Computing Laboratory

Computer Crimes - 1997

Computer Crimes & Internet Security

4/29/97 Computer Crimes Division, NASA OIG 1

Overview

- Computer intrusions via the Internet
- Tools and techniques used by offenders
- Internet security basics
- Legal, Policy, other initiatives
- Future challenges

The NASA OIG

- So why are we here?
- We're users/defenders of the technologies that you develop
- We're natural allies on this issue
- Technology is public treasure . . . merits protection

What Is The OIG Mission?

- Ensure programmatic & systemic integrity of NASA
  - » Authority: Public Law 95-452, as amended
- Mission is executed through a staff of
  - » Auditors *(ounce of prevention function)*
  - » Inspectors *(ounce of prevention function)*
  - » Investigators *(ounce of prevention and pound of cure function)*
- Oversight (forward & backward looking)

The Computer Crimes Division

- Investigates criminal intrusions into NASA's networks
  - » Types: Internet (packet based)
  - » CBX (circuit switched)
  - » Spacecraft command & control
- Promotes awareness
- Supports criminal investigations of other matters

Computer Crimes Personnel

- Posted at each NASA Center
- Trained and equipped to deal with cyber crimes
- We can "Email" you a roster

Computers & Crimes

- Computers can be used as ...

- Instruments in the commission of crimes
  - » Books & records analogy

- Storage devices
  - » Electronic dead drops analogy

- Weapons
  - » Tools to attack electronics sites

We'll Focus On Weapons

- Using computers to attack networks
  - » Overview of intrusions issue
  - » Types of systems usually exploited

- Types of crimes involved . . .
  - » Industrial Espionage
  - » Theft
  - » Extortion
  - » Denial of service

Case Examples

- External attacks on NASA . . .
  - Mayhem via the Internet . . .
    - » Denial of service attacks against NASA
    - » Attacking corporate America from NASA
    - » Competitors using a NASA domain to attack each other
- Example: extortion . . . $500K via the Internet
    - » Disabling networks and sending threatening messages to NASA

Case Examples

- Insider Schemes . . .
  - Combined schemes via CBX's & Internet
    - Trade secret fraud
    - Voice mail fraud
    - Attacking external sites from NASA
- Traditional contract fraud via the Internet...
  - Bribery, trade secret, and contract fraud by the Internet: $500 Million

Part I.

Computer Intrusions and the Internet

Internet "Axioms"

- What you "touch electronically, touches you"
- The Internet is "non-contextual"
  - Who is behind the terminal?
  - Understanding issues
  - Example given...contextual versus non-contextual

Analogies

- These are general crimes in an electronic environment . . .
  - » Analogies: murder, rape, robbery, fraud
- Murder or Mayhem= a denial of service attack

Analogies

- "Theft"=industrial espionage (example cited)
- "Fraud"=use of the Internet to execute a fraudulent scheme
  - » Examples cited

Nature of attacks

- Automated
- Varied
- Often exploit default configurations of operating systems
- Attack signatures (description provided)

Types of intruders ...

- Sophisticated intruder...
  - » Works alone or with trusted ring
  - » Develops hacking tools
  - » Uses them
  - » Distributes them *(floods Internet...deception tactic)*
- "Lamer" or cookbook intruder...
  - » Gathers automated tools from "Net"
  - » Uses them

Part II:

Some Tools and Techniques

Examples: Tools

- Scanner software
  - » Looks for vulnerabilities on networks
  - » SATAN (is an example of this)
- Web daemon exploits
  - » Exploits default or common configurations

Examples: Tools

- Internet Packet Filters aka "sniffers"
  - Passive tools to capture network traffic
- Password crackers
  - Tool to decrypt password files
- Spoofers
  - Tools to masquerade as other users with

Examples: Tools

- Nukers
  - Tools to destroy system log trails with *(evidence)*
- Anonymous Remailers
  - Can cause business problems
  - Masks attacks
- Steganography
  - Graphics encryption...electronic dead drops
  - Industrial espionage issue

A Standard Attack Pattern

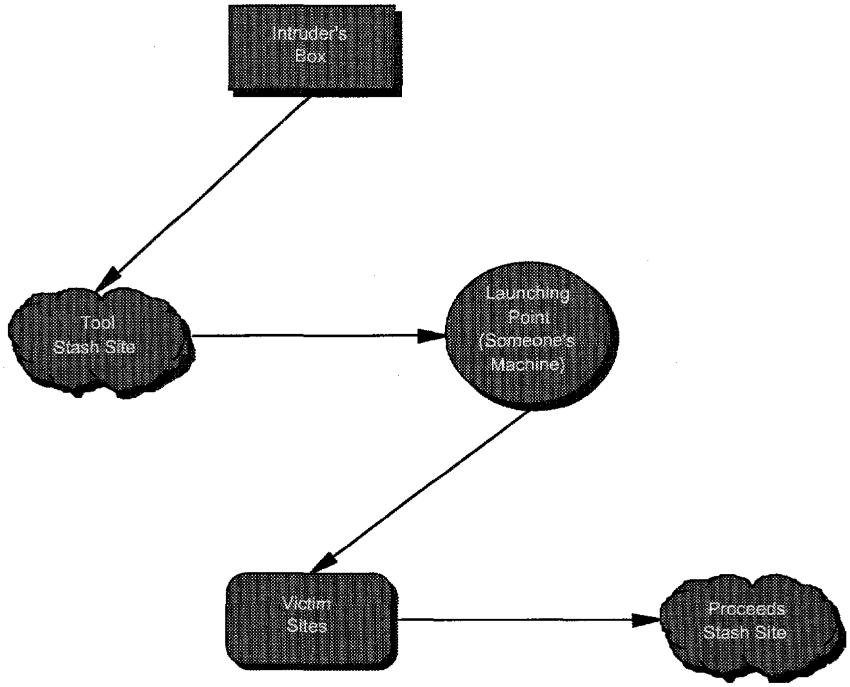

Sample Attack Sequence

- Scans for vulnerabilities
- Executes a tool to get user or root access
- Plants "trojan" tools to permit hidden, future access
- Obscures access trails
- Attack or collect on other targets

A Denial Of Service Attack

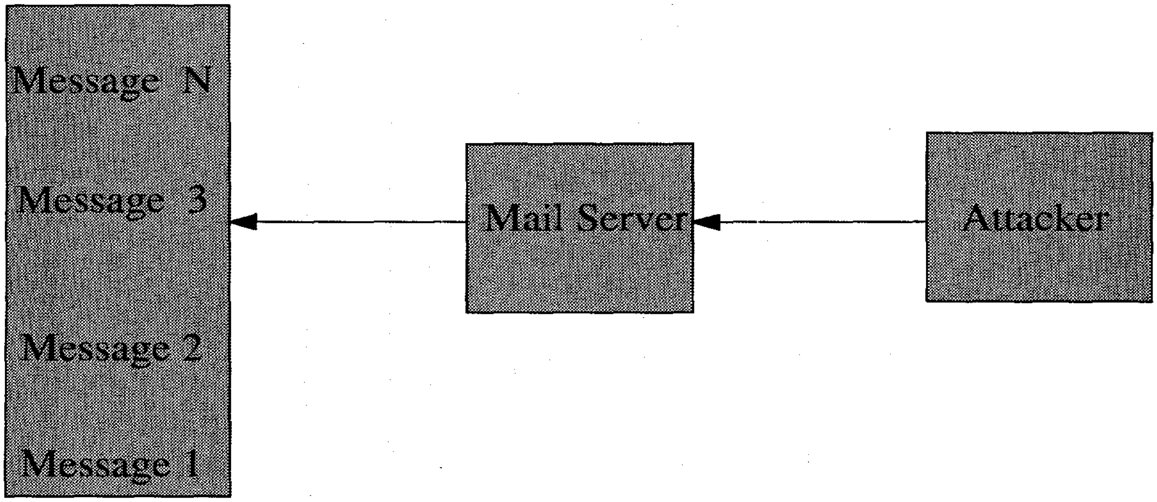

Denial of Service Attacks

- Some take the form of "choking" victim machines as depicted on the previous slide
  - » Sending endless "pings" to a server
  - » Scripting routines that constantly access large files
- SYN Flood Attacks
  - » "Triple handshake" exploitation

Part III:

Internet Security Basics

Why Does This Happen?

- Challenges posed by Internet
- Struggle between "openness versus security"
- TCP/IP network systems are designed to promote openness

What's The Point?

- Internet security work is "never done"
- It is a process of adding and managing "layers" of security
- Many folks think of this area from one or two points of view . . .
  - » Firewalls
  - » Controlling users
  - » Controlling modem deployment

There's Much More Involved

- Training
  - » Awareness training for users
    - re threats/exploitations
  - » Security training for system administrators
    - What to do and how to do it
  - » Legal and policy training for all employees
    - Includes all persons using NASA funded systems
- Detection and response mechanisms
  - » Automated incident detection tools
  - » Coordinated incident response mechanisms

There's More Involved . . .

- Integrated policies
  - » What should and shouldn't be out on Internet servers?
  - » Need for banners that address 4th amendment issue
- Integrated software/hardware solutions
  - » In terms of network monitoring tools

Part IV:

Legal Issues

Policy and other initiatives

To Grapple With All Of This ...

- Government had a choice . . .
- Go back through entire US federal code & modify it . . .
- Or . . . try to capture all concerns in one or more new federal laws

First Problem: Legal Definition Of "Computer"

- Broad enough to include your . . .
  - Wristwatch . . .
  - The chip in your microwave oven . . .

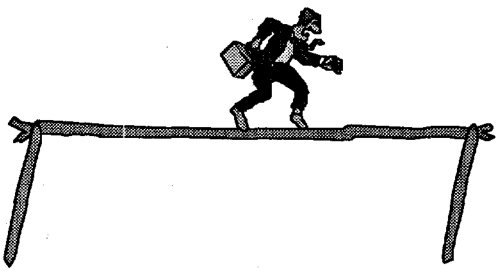

- The Justice Department worked with Congress on "reasonable" language for statutes

Federal Laws

- 18 USC 1029
  - » "Fraud and related activity in connection with access devices"
    - Stolen password file counts as access device
- 18 USC 1030
  - » "Fraud and related activity in connection with computers"

Federal Laws

- 18 USC 1831-39
  - » The Economic Espionage Act of 1996
- 42 USC 2000
  - » The Privacy Protection Act
- 18 USC 2511
  - » The Federal Wiretapping Statute
    - Exempts system administrators from prosecution if they monitor their systems for "auditing or security" purposes

Federal Laws

- 18 USC 2701-10
  - » "Stored wire and electronic communications and transactional records access"
    - Search warrant needed for unopened Email less than 180 days old
    - Opened Email (stored)is no longer Email
- Civil liability
  - » Built into most of these laws
  - » Court orders required for many things they are not required for under "normal" law

Example:

The Computer Fraud and Abuse Act

(18 USC 1030)

Example: 18 USC 1030

- 1030 (a) (1) =
  Makes "computer espionage" a felony

- 1030 (a) (2) =
  Says accessing without authority or *exceeding* authority on a computer
  - » Is a felony
  - » "protected computer" vs. "federal interest computer"
    - Former is now operative language...covers financial institutions, etc.
  - » Use in interstate or foreign communications

Example: 18 USC 1030

- 1030 (a) (3) =
  Is the trepass provision . . .
  - » "non-public" vs. "public" access of a machine
  - » Web server example given

- 1030 (a) (4) =
  The fraud provision . . .
  - » Use of computer time: felony

Example: 18 USC 1030

- 1030 (a) (5) =
  The damages provision . . .
  - » Targets breakins by outsiders
  - » Felony
  - » Keep track of costs to fix system…Damages!
- 1030 (a) (6) =
  Exceeding authorized access provision
  - » Targets inside offenders

Example: 18 USC 1030

- 1030 (a) (7) =
  The extortion provision . . .
  - » Designed to protect system administrators & owners
  - » Felony

Economic Espionage Act of 1996

- 18 USC 1831 =
  - » Covers foreign power involvement/benefit
- 18 USC 1832 =
  - » Covers interstate, foreign commerce, or benefit for someone other than the owner

Economic Espionage Act of 1996

- What's a trade secret?
  - » Something "not known"...doesn't have to be "new"
  - » Patents generally = "new"
- Have to show harm to the owner . . .
  - » Knowledge it will cause harm by the offender
  - » Must truly be a secret
  - » Financial harm must be demonstrated

Economic Espionage Act of 1996

- Designed to protect corporate America
  - » law not written with idea that Government could be a victim
- NASA could be an exception . . .
  - » Holding the trade secrets of companies (on servers !(?))
  - » "Novating" of NASA inventions into trade secrets by firms
  - » Partnerships, alliances, and teaming issues

Solution To Legal Issues

- International agreements

Part V:

Future Challenges

Encryption

- Public-Private Key Cryptography
  - » Almost absolute privacy possible for the contents of stored electronic communications (18 USC 2703)
  - » Never envisioned by founding fathers
  - » Legal issues raised . . .4th Amendment
- Encryption infrastructure issues
  - » Transnational storage issues (Examples given)
  - » National key management infrastructure
  - » Sheathed key proposal (Example given)
  - » Data recovery will drive issue (business example cited)

Encryption . . .

- Some legal issues with encryption . . .
  - » How you "lock up" a document has nothing to do with it's discoverability under the law
- "Act of production immunity"
  - » Needed WHEN the ACT of production is BOTH testimonial AND incriminating
    - Compelling passwords/keys
- Steganography. . . issues with

Melding Of Technologies

- Integration will present challenges
- E.G. . . . Internet type capabilities with . . .
  - » Command, Control, and Communications with air/ spacecraft
  - » Cable TV
  - » Banking
  - » Audio/visual telephony
  - » Personel satellite operations

Other Issues

- Truly "remote" perpetration of crimes
  - » Examples cited
- Some contend that next release of IP will have an "unbreakable" encrypted stack . . .
  - » All problems cited earlier will "go away", (!) (?)
- History has shown us that the "secrets" of technological advances are the most ephemeral of all

"Washington" Initiatives

- Critical Infrastructure Protection Commission
- Infrastructure Protection Task Force
- INFOTECH Training Working Group
- Internet Guidelines Working Group
- International efforts...multiple agencies

Other Initiatives

- Departmental/Agency Programs
  - » NASA OIG
  - » Air Force OSI
  - » FBI
  - » IRS CID/Inspections
  - » Secret Service

Want More Information?

- Computer Crimes Division
  NASA Office of Inspector General
  300 E Street, Southwest
  Washington, DC 20546

- Tom Talleur: (202)-358-2587
  - » Email: thomas.talleur@hq.nasa.gov
- Steve Nesbitt:(202)-358-2576
  - » Email: steve.nesbitt@hq.nasa.gov

DREAMS TO MACHINES AND BACK

DAVID URIE

THE CONFERENCE ON HIGH SPEED COMPUTING BANQUET

APRIL 22, 1997

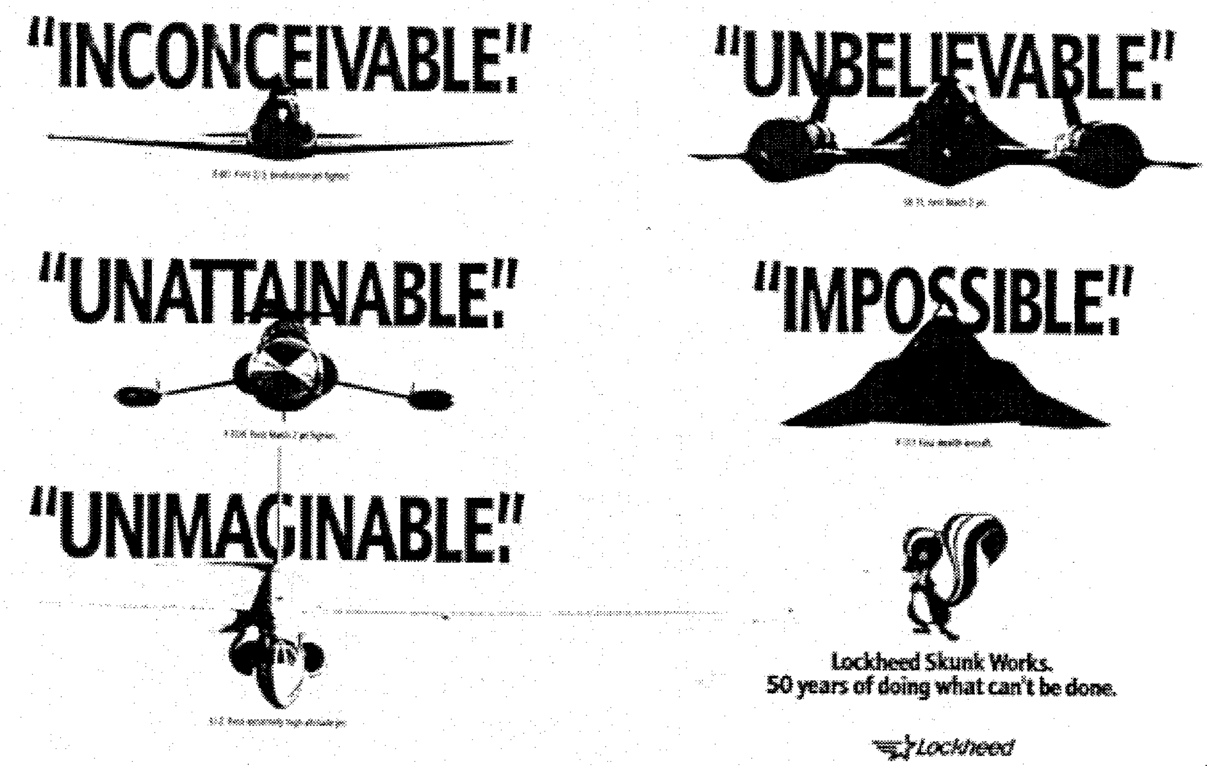

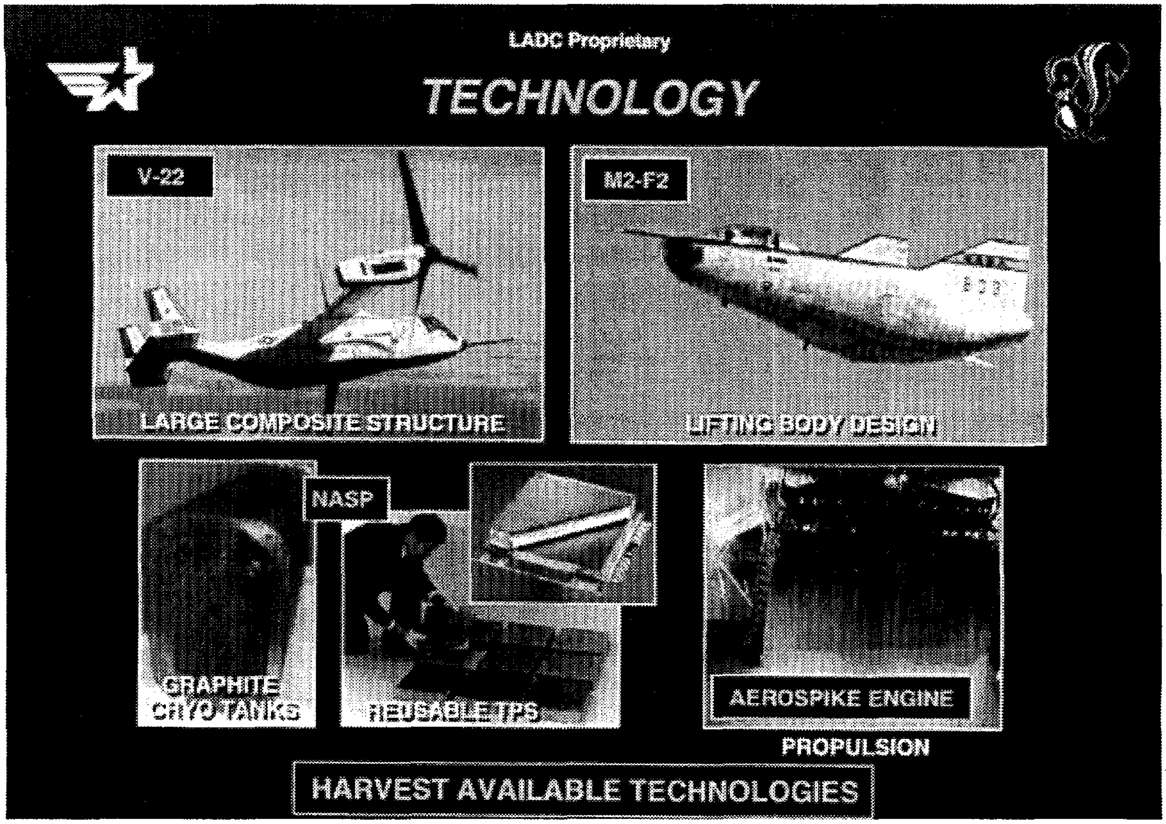
LADC Proprietary
TECHNOLOGY
V-22
LARGE COMPOSITE STRUCTURE
M2-F2
LIFTING BODY DESIGN
NASP
GRAPHITE CRYO TANKS
REUSABLE TPS
AEROSPIKE ENGINE
PROPULSION
HARVEST AVAILABLE TECHNOLOGIES

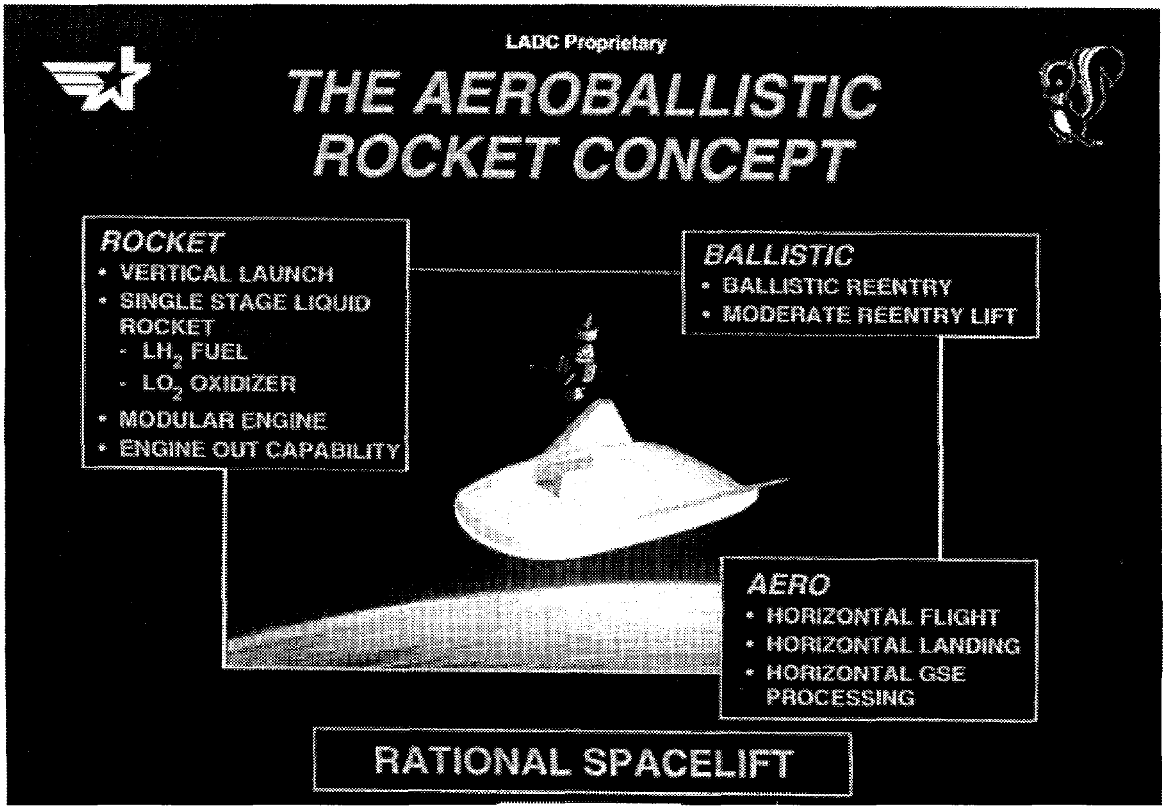
LADC Proprietary
THE AEROBALLISTIC ROCKET CONCEPT
ROCKET
• VERTICAL LAUNCH
• SINGLE STAGE LIQUID ROCKET
- LH_2 FUEL
- LO_2 OXIDIZER
• MODULAR ENGINE
• ENGINE OUT CAPABILITY
BALLISTIC
• BALLISTIC REENTRY
• MODERATE REENTRY LIFT
AERO
• HORIZONTAL FLIGHT
• HORIZONTAL LANDING
• HORIZONTAL GSE PROCESSING
RATIONAL SPACELIFT

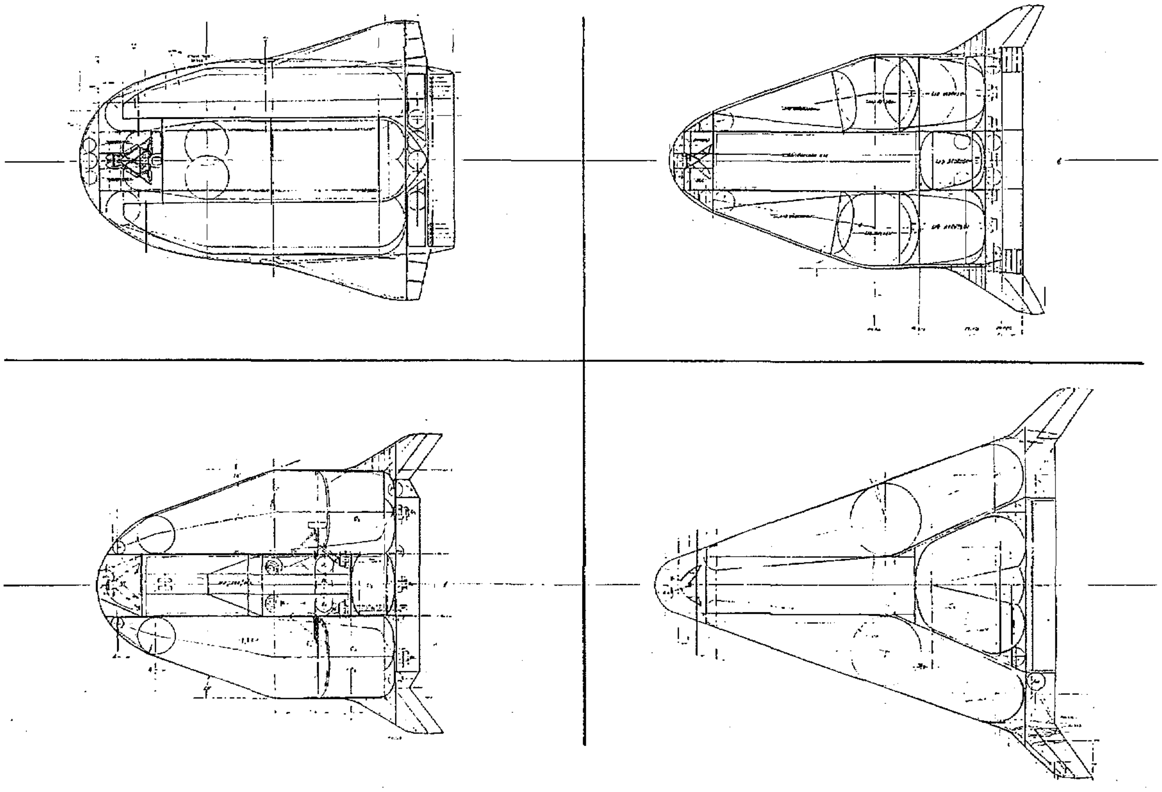

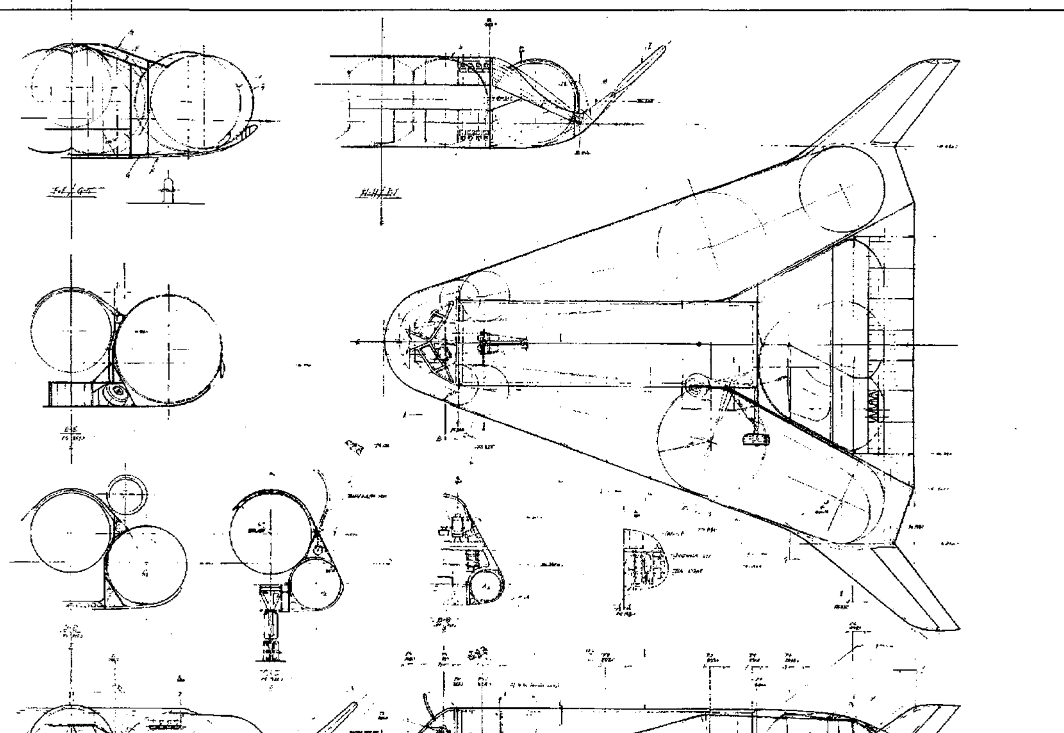


Affordable
Access
to Space
. . . A Skunk Works Initiative
June 1993

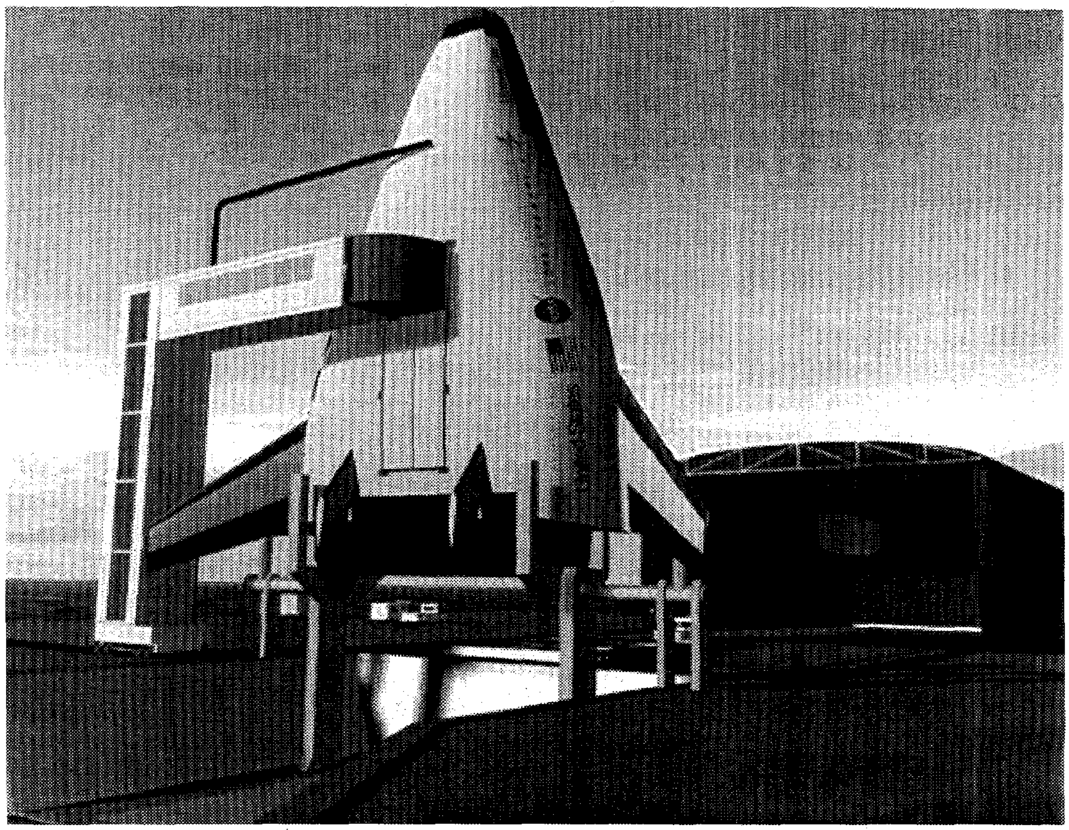

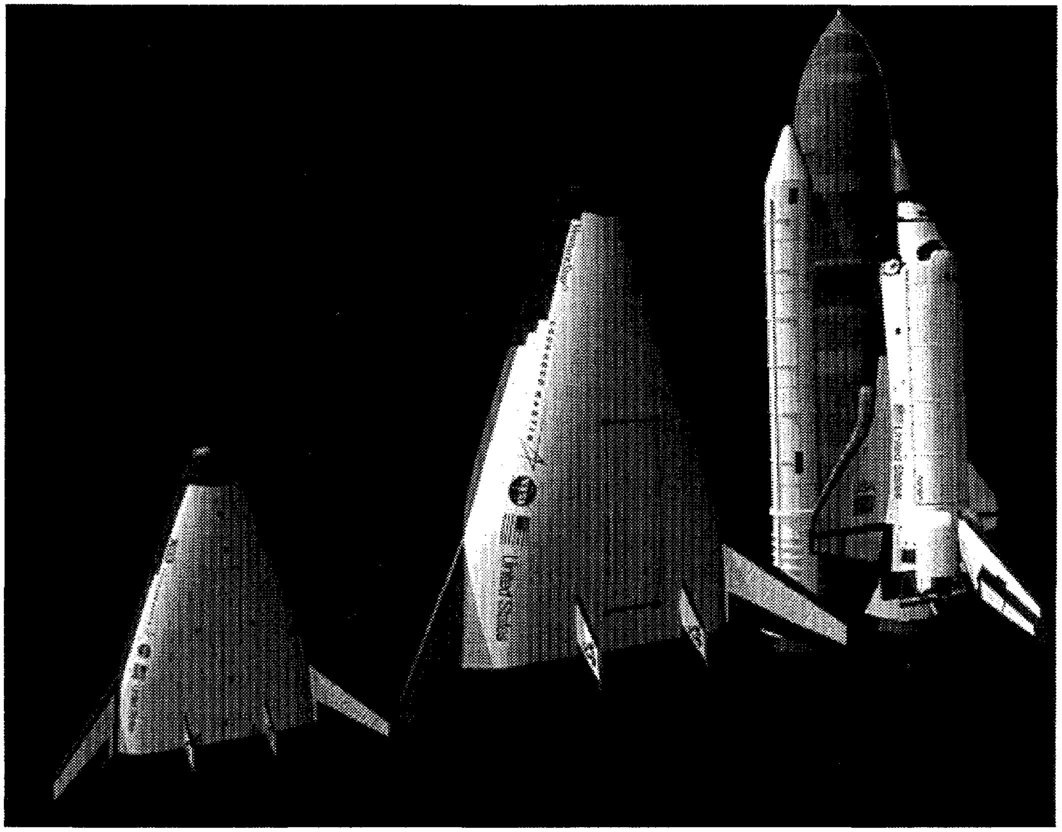

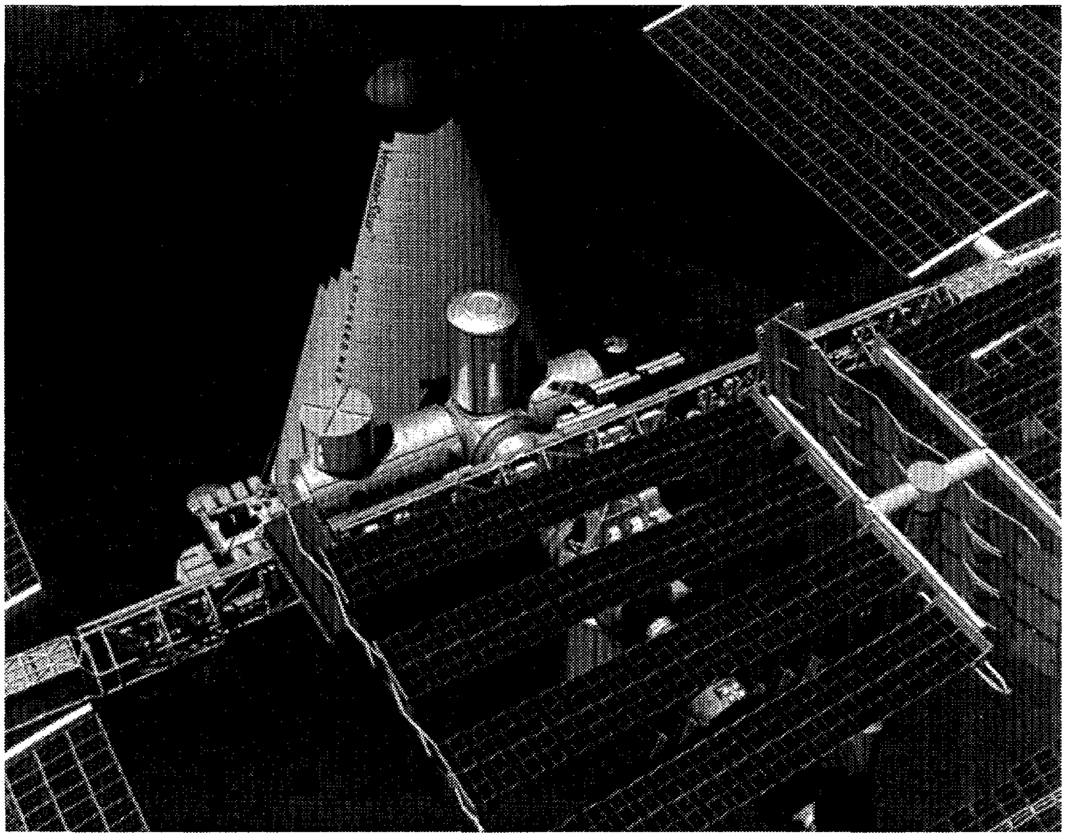

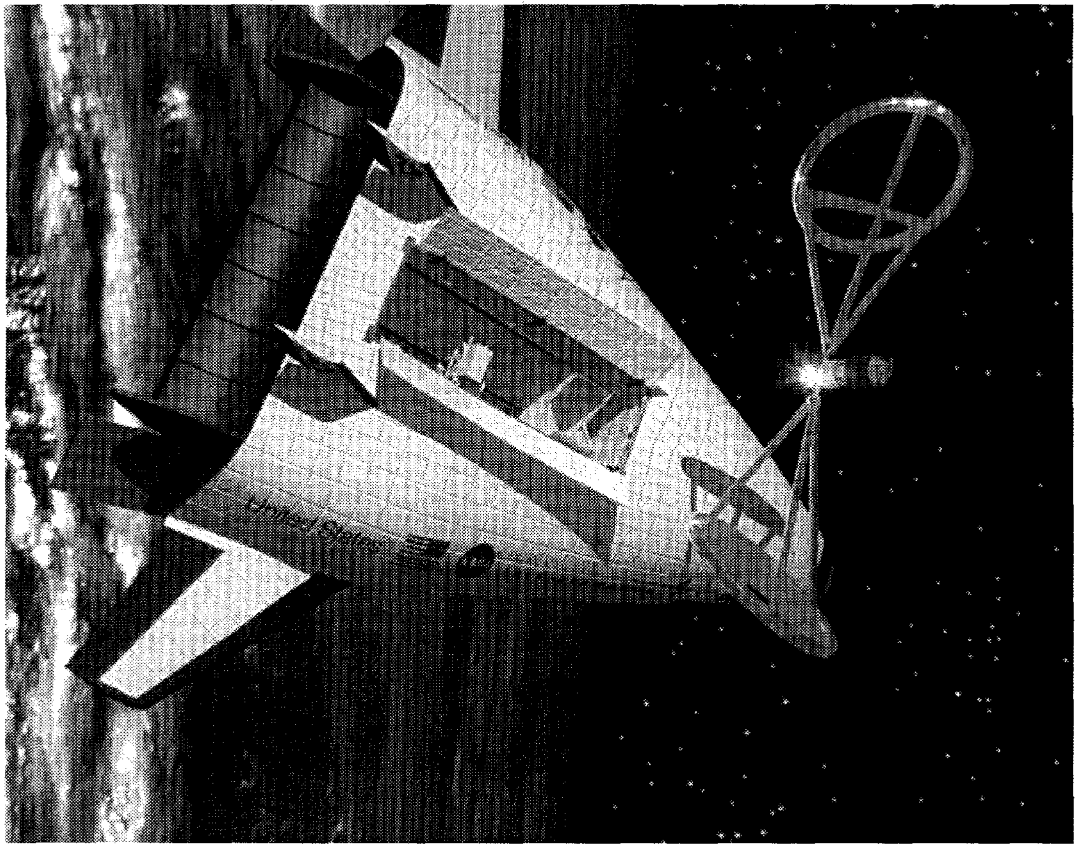

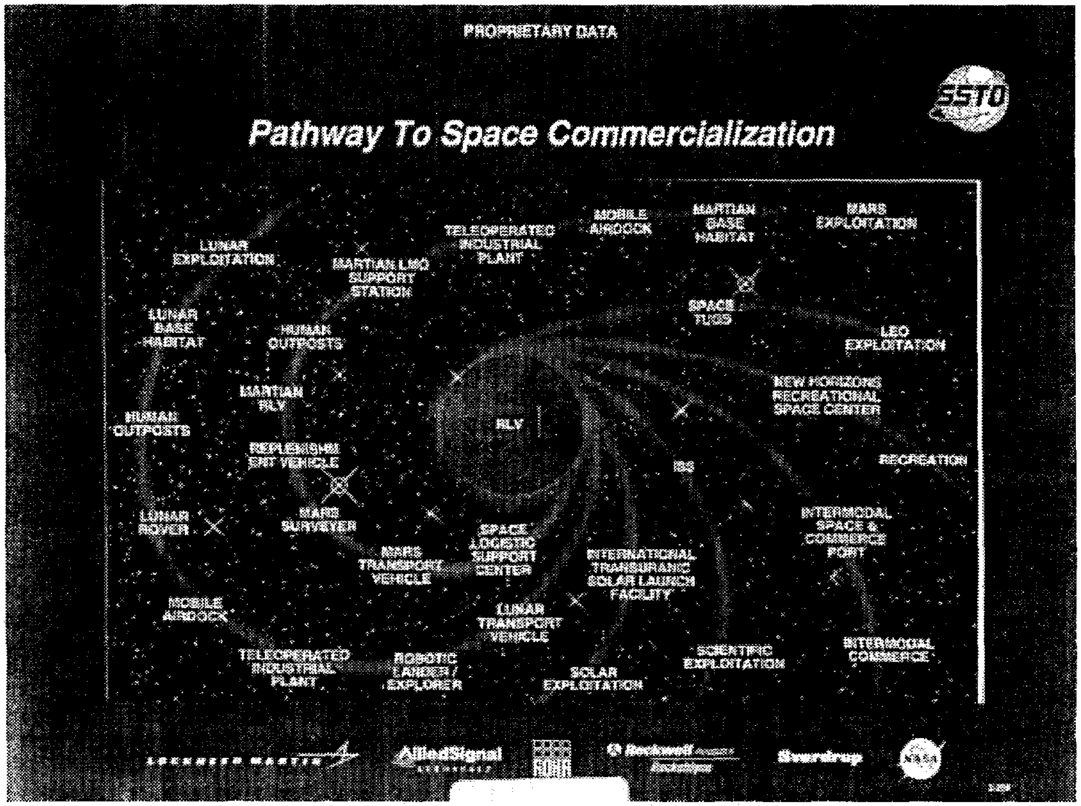
PROPRIETARY DATA
Pathway To Space Commercialization
LUNAR EXPLOITATION
MARTIAN LMO SUPPORT STATION
TELEOPERATED INDUSTRIAL PLANT
MOBILE AIRDOCK
MARTIAN BASE HABITAT
MARS EXPLOITATION
LUNAR BASE HABITAT
HUMAN OUTPOSTS
SPACE TUGS
LEO EXPLOITATION
MARTIAN RLV
NEW HORIZONS RECREATIONAL SPACE CENTER
HUMAN OUTPOSTS
RLV
REPLENISHMENT VEHICLE
ISS
RECREATION
LUNAR ROVER
MARS SURVEYER
SPACE LOGISTIC SUPPORT CENTER
MARS TRANSPORT VEHICLE
INTERMODAL SPACE & COMMERCE PORT
INTERNATIONAL TRANSURANIC SOLAR LAUNCH FACILITY
MOBILE AIRDOCK
LUNAR TRANSPORT VEHICLE
TELEOPERATED INDUSTRIAL PLANT
ROBOTIC LANDER / EXPLORER
SOLAR EXPLOITATION
SCIENTIFIC EXPLOITATION
INTERMODAL COMMERCE
AlliedSignal
Sverdrup

Trade-Offs in High Performance Pipelined Processor Design

Apoorv Srivastava and Alvin M. Despain

apoorv@isi.edu, despain@isi.edu

Advanced Computer Architecture Laboratory
http://www.isi.edu/acal

Information Sciences Institute
University of Southern California

Pipelined Processor Performance

$$Perf_{cpu} = Freq_{clock} \cdot \frac{1}{\sum_{\forall (i \in I)} \left(DynamicInstructionCount_i \cdot \frac{1}{IPCeffective_i} \right)}$$

$Perf_{cpu}$ **= CPU Performance**
$Freq_{clock}$ **= Clock Frequency**
Dynamic Instruction Count$_i$ **= Dynamic Instruction Count for Instruction *i*.**
IPCeffective$_i$ **= Effective Instructions Per Clock for Instruction *i*.**
I **= The set of all instructions.**

- **Goal of High-Performance Design:**
  - Increase $IPC_{eff.}$ ⇧ (*Superpipe*: $IPC_{eff.} = 1$ Superscalar: $IPC_{eff.} > 1$)
  - Increase Frequency ⇧

Limitations to Speedup

- Instruction Dependencies (Decreases IPC)

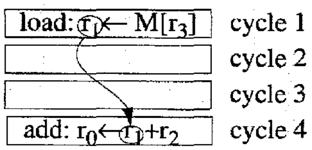

Load Latency = 3 cycles

- Pipelining Overhead (Limits Frequency)

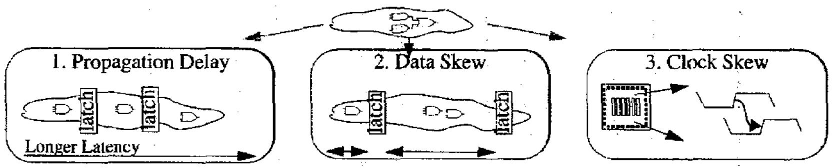

- Control Path Limitations (Limits Frequency)

Minimum amount of control logic for each pipeline stage

Performance Space of Pipelined Processors

A Qualitative Performance Analysis

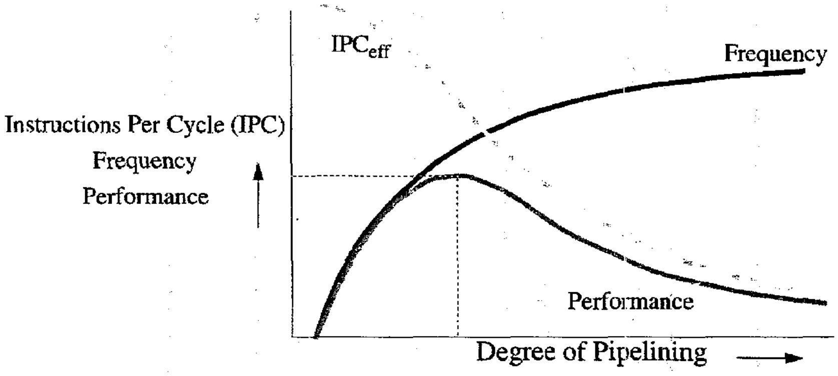

$$Perf_{CPU} = (IPC_{eff} * Frequency)/(InstructionCount_{dynamic})$$

What's the optimal degree of pipelining for symbolic benchmarks?

Challenges to High-Performance Pipelined Processors

- Design of very High-Speed Processors (Increasing Frequency)
- Increasing the $IPC_{effective}$

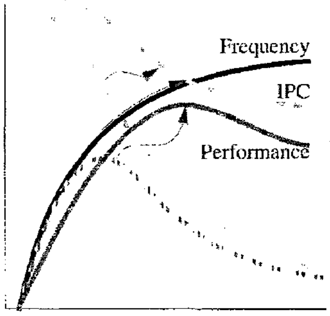

Proposed Approach

- ⇧ Freq.: **Deep Pipelining** and **Innovative Circuits**
- ⇧ IPC_{eff}: Exploiting **Instruction Level Parallelism (ILP)**; i.e. finding independent instructions

Selected References (VLSI)

MHz-GHz

- The Road to High-Speed VLSI Design
  - Reducing propagation delay (Pipelining): **Single-Phase Clock Dynamic Logic**

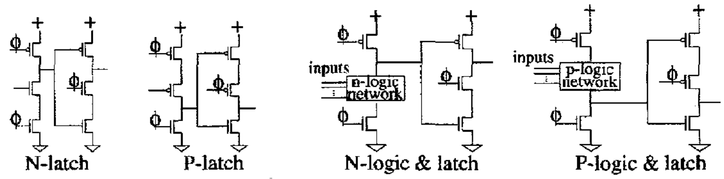

  - Reducing Clock Skew

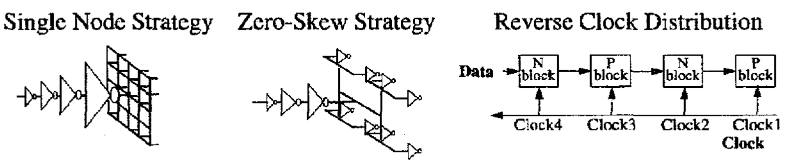

(Dobberpuhl '92, Tsay '91, Yuan & Svensson '89, Afgahi & Svensson '90)

Trends in Pipelined Processor Design

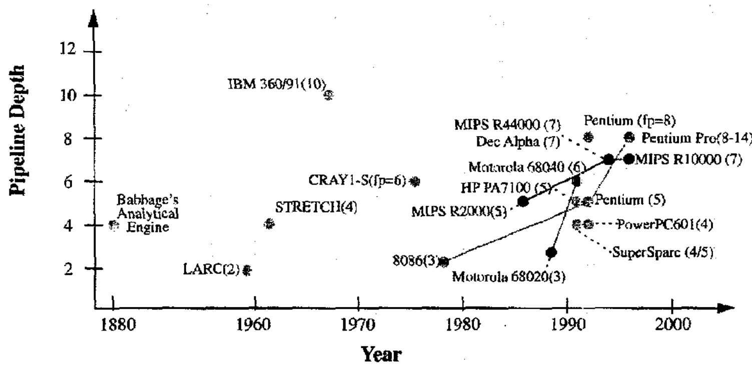

Charting the Performance Space

- Pipelined Models

IF	ID	EX	MEM	WB

5-Stage Pipeline (7 logic levels) est. freq.: 130 MHz

IF1	IF2	ID	EX	M1	M2	TC	WB

8-Stage Pipeline (5 logic level) estimated frequency: 200 MHz

IF1	IF2	IF3	IF4	ID1	ID2	BP	EX1	EX2	EX3	M1	M2	M3	M4	TC1	WB1	WB2

17-Stage Pipeline (2 logic levels) estimated frequency: 380 MHz

IF1	IF2	IF3	IF4	IF5	IF6	ID1	ID2	ID3	BP	EX1	EX2	EX3	EX4	EX5	M1	M2	M3	M4	M5	M6	TC1	WB1	WB2	WB3

25-Stage Pipeline (1 logic level) estimated frequency: 500 MHz

IF	Instruction Cache Access
BP	Bypass Data
ID	Instruction Decode/Register Read
EX	Execute
M	Memory Cache Access
TC	Tag Check
WB	Retire Instrctuion/Writeback Result in Regfile

Key Features
- 0.8 μm CMOS
- On-chip Caches (64k Data & 64k Instruction)
- Bypassing

- Note: 64 k byte caches are used for simulation only. VLSI timing is based on smaller cache sizes (4k).

8-Stage Pipeline

IF1	IF2	ID	EX	M1	M2	TC	WB

8-Stage Pipeline *200MHz*

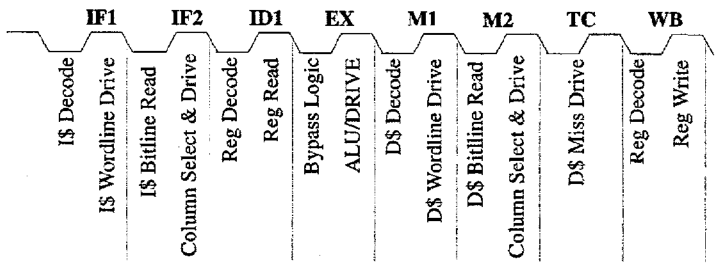

Delay Slots:

Branch	= 4
Alu-Any	= 0
Load-Any	= 2
Store-load	= 1

17-Stage Pipeline

IF1	IF2	IF3	IF4	ID1	ID2	BP	EX1	EX2	EX3	M1	M2	M3	M4	TC1	WB1	WB2

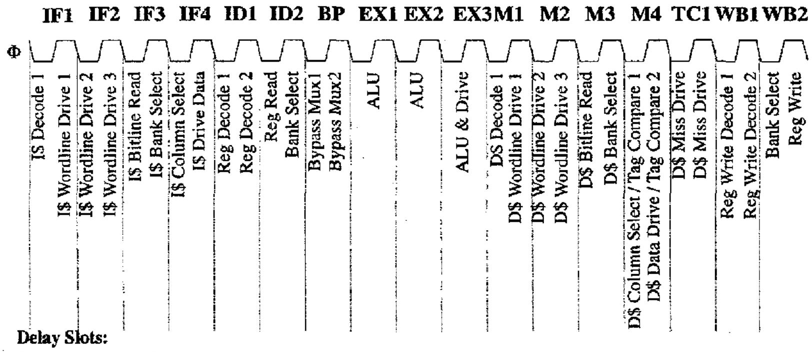

Delay Slots:

Branch	= 9
Alu-Any	= 3
Load-Any	= 7
Store-load	= 1

Simulation Environment

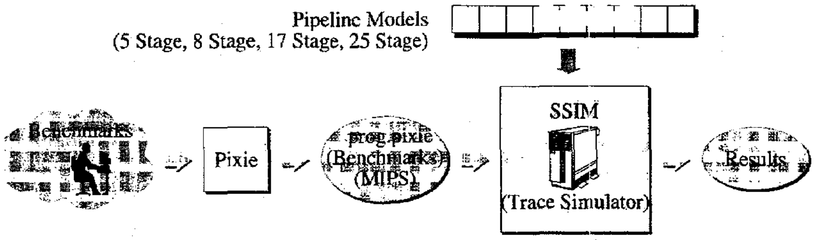

⇨ 64 k I & D Cache
⇨ 2048 Branch Target Buffer
⇨ 12 Cycle Cache Miss Penalty

Benchmark - SpecInt '92

Benchmark	Description	No. of Instructions (millions)	Input Set
compress	data compression algorithm	6.88	short input
espresso	boolean function minimization	0.79	short input *dc1*
eqntott	translate logical eqn. to truth table	0.47	short input *ex10*

Simulation Result (Compress)

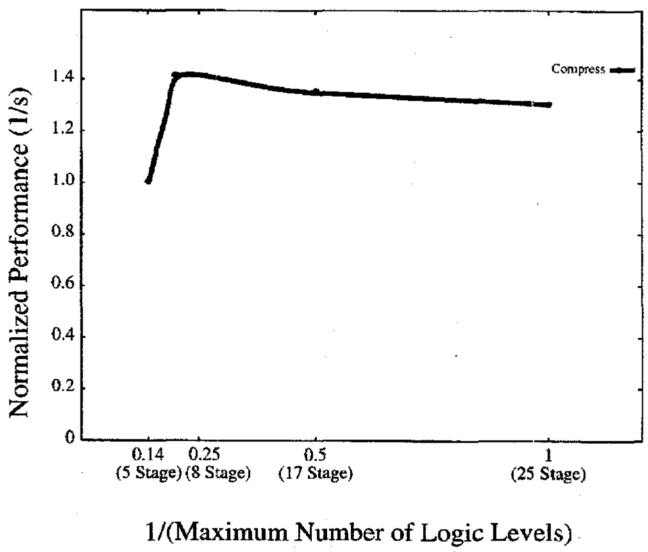

Simulation Result (Espresso)

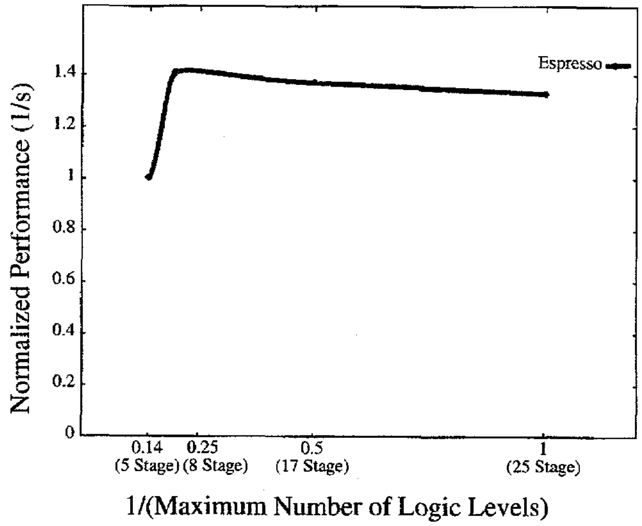

USC INFORMATION SCIENCES INSTITUTE

ISI Advanced Computer Architecture Laboratory

Simulation Result (Eqntott)

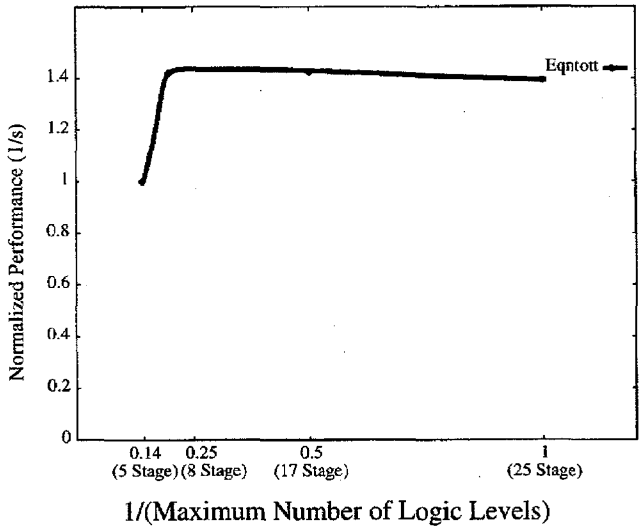

USC INFORMATION SCIENCES INSTITUTE

ISI Advanced Computer Architecture Laboratory

Conclusion

- Comparison of 4 Pipeline Models (for 0.8 μm CMOS):
  - ⇨ 5 Stage Pipeline - maximum 7 logic levels/phase est. freq. 130 MHz
  - ⇨ 8 Stage Pipeline - maximum 5 logic levels/phase est. freq. 200 MHz
  - ⇨ 17 Stage Pipeline - maximum 2 logic levels/phase est. freq. 380 MHz
  - ⇨ 25 Stage Pipeline - maximum 1 logic level/phase est. freq. 500 MHz
- Benchmarks (compress, espresso, eqntott)
  - ⇨ Optimal performance = 8 Stage Pipeline (5 logic levels/phase)
- Branch prediction is very low:
  - ⇨ *Compress* 83%
  - ⇨ *Eqntott* 69%
  - ⇨ *Espresso* 68%
- Results would favor higher degree of pipelining, if $IPC_{effective}$ is higher.

Future Work

- Study other benchmarks (floating point)
- Study pipeline models with a pipeline simulator running optimized code.
- Study impact of incomplete bypassing.
- Investigate the use of carry-save adders.
- Investigate impact of speculative execution
  (software support and hardware cost).

Parallel Programs from Constraint Specifications

Ajita John

Department of Computer Sciences

University of Texas at Austin

Advisor: Prof. J.C. Browne

ajohn@cs.utexas.edu
http://www.cs.utexas.edu/users/ajohn

Application Oriented High-level Specification of Parallel Programs

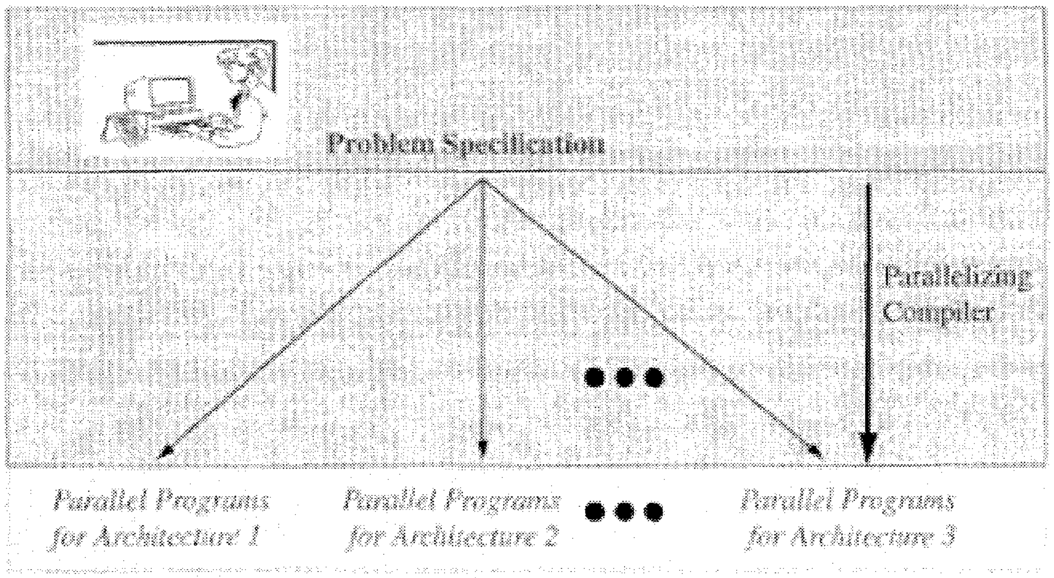

Goal

High-level Specifications ---> High Performance

A representation for problem specification which

- places minimum burden on user
  - no explicit parallelism in specification
- allows extraction of maximal parallelism
- allows reuse of components

Constraints:

attractive choice for a representation for problem specification

Why ?

Goal

Extraction of efficient parallel structures from constraints

Contributions

1. Programming system based on constraints for matrix computation.
    - *Hierarchical Type System*
    - *Programming Representation*
    - *Modular Structures for Constraint Systems*
    - *Basic Compilation Algorithm*
    - *Extension for Cyclic Constraints*
    - *Extraction of Task and Data Parallelism*
2. Translation to multiple parallel architectures.
3. Development of applications yielding high performance.

A constraint is a relationship between a set of variables.

E.g. 1

C == (F-32) * 5/9

E.g. 2

F == (9/5) * C + 32

== is equality as opposed to assignment

1 ≡ 2

Constraint Programming vs. Imperative Programming

Constraint Programming	Imperative Programming
C == (F-32) * 5/9	Program 1 C = (F-32) * 5/9
	Program 2 F = 32 + (9/5) * C

Related Work (Representative Examples)

Constraint Programming

- *Thinglab - Alan Borning*
- *Consul - Doug Baldwin*
- *CCP, Oz - Vijay Saraswat, Gert Smolka et. al.*

Parallel Programming

Imperative Programming

- *HPF - Ken Kennedy et. al.*

Equational Programming

- *Unity - Chandy, Misra*

Logic Programming

- *PCN - Chandy, Taylor*
- *Strand - Foster, Taylor*

Approach

Constraint program (constraint specification & input set)

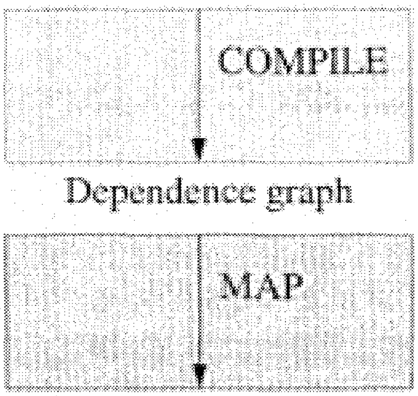

Sequential and parallel C programs generated for

CRAY J90, SPARCcenter 2000, Sequent,

PVM,

MPI (under development)

CODE - Browne, Newton et. al.

Hierarchical Type System

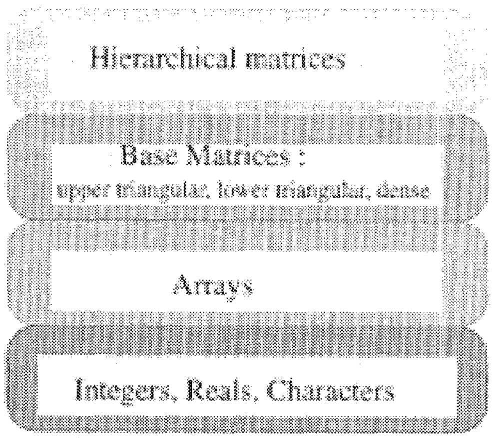

- Constructs new types at higher levels from existing types at lower levels.
- Provides a tool for controlling granularity.
- Each high level type may incorporate domain information in its definition.
- Different efficient algorithms can implement operations depending on the type.
- We have a translator for specifications over a rich matrix type set to parallel programs.

Operator Expressions

- the building blocks for constraints

Expressions involving:

1. Scalar operators, pure function invocations

   E.g. $a + b * sqr(c)$

2. Matrix operators

   E.g. $M_1 * M_2 - M_3$

3. Indexed operators: <op> FOR (<index> <b1> <b2>) <expression>

   E.g. + *FOR (i 1 5)* *A*[i]

   Equivalent to $A[1] + A[2] + A[3] + A[4] + A[5]$

Constraint Representation

Simple Constraints: relational operators (<, <=, >, >=, ==, !=) over operator expressions

E.g. $a + b > c$

$M_1 * M_2 == M_3$ | only == for matrices

Propositional Operators: NOT, AND, OR, on constraints:

E.g. $a + b > c$ *AND* $a == 5$

Constraints over indexed sets:

AND/OR FOR (<index> <b1> <b2>) { $C_1, C_2, \ldots, C_n$ } C_i is a constraint

E.g. *AND FOR (i 1 5) { A[i] == A[i-1]+1, B[i] == A[i] }*

- represents 10 constraints

Constraint Modules: encapsulation of constraints over parameters

E.g. *EquivalentTemp(F,C)*

Quadratic Equation Solver

$$ax^2 + bx + c == 0, \quad 2ax = -b \pm \sqrt{b^2 - 4ac}$$

```
/* Main */
    a == 0  AND  r1 == r2  AND  b*r1 + c == 0

    OR

    a != 0 AND DefinedRoots(a, b, c, r1, r2)
- - - - - - - - - - - - - - - - - - - - - - - - - -
/* Constraint module */

 DefinedRoots(a, b, c, r1, r2)
    t == sqr(b)- 4*a*c
    AND t>=0
    AND 2*a*r1==-b+sqrt(abs(t))
    AND 2*a*r2==-(b+sqrt(abs((t))))
```

Compilation

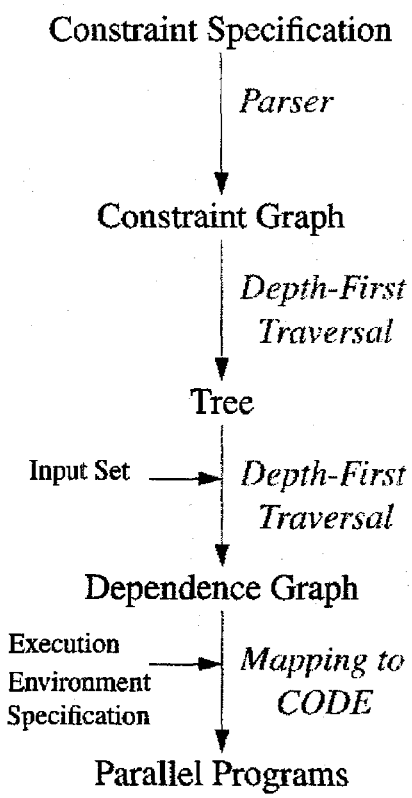

Tree from Depth-First Traversal (1)

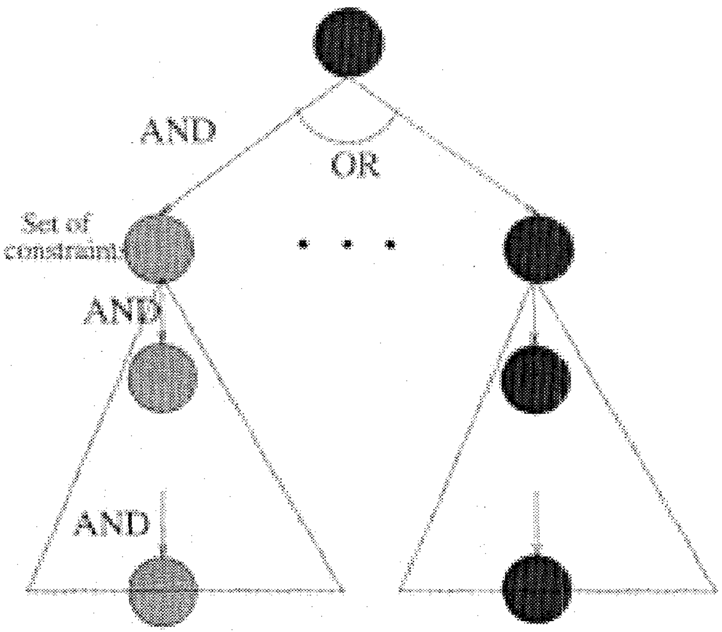

Each path may yield a solution.

Depth-first Traversal (2)
Tree ------> Dependence Graph

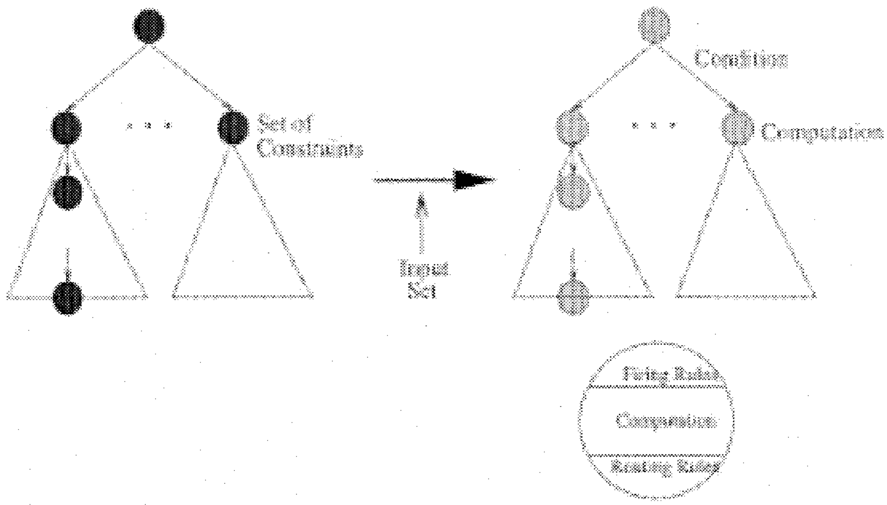

Dependence Graph

Nodes are computational units.

Nodes execute dependent on certain *conditions* and the availability of data.

Resolution of Constraint Modules

Constraint Modules can be compiled to procedures

input and output parameters determined by the input set at the point of call

E.g. Constraint Module *DefinedRoots(a, b, c, r1, r2)* can be compiled to

three procedures:

- *DefinedRoots(a, b, c, r1, r2)*
- *DefinedRoots(a, b, r1, c, r2)*
- *DefinedRoots(a, b, r2, c, r1,)*

Red: Inputs, Blue: Outputs

Resolution of Constraints over Indexed Sets

Constraints over Indexed Sets can be compiled to loops over computations extracted from constraints

```
AND FOR ( i 1 10 ) {
    A[i] + A[i-1]  ==  0,
    A[i-1] - B[i]  ==  2
}

Input A[0]
```

→

```
for ( i=1;  i<=10;  i++) {
    A[i] = -A[i-1];
    B[i] =  A[i-1] - 2;
}
```

Extraction of Parallelism

Types of Parallelism

AND-OR

Task

Data

Parallelism to be exploited:

1. in different paths of dependence graph (OR, Task)
2. between computations at a node (AND, Task)
3. in operations within a computation (Task)

   $\underline{M_1 * M_2} + \underline{M_1 * M_3}$

4. in iterations of loops (Data)

↓ Implemented

5. in primitive operations over base types (Data)

↓ Future

Control over Granularity

Modules

Functions

Hierarchical Types

Block Odd-Even Reduction Algorithm (BOER)

Solution of A * x == d, where

$$A == \begin{bmatrix} B & C & 0 & 0 & \cdots & 0 & 0 & 0 \\ C & B & C & 0 & \cdots & 0 & 0 & 0 \\ 0 & C & B & C & \cdots & 0 & 0 & 0 \\ . & . & . & . & . & . & . & . \\ 0 & 0 & 0 & 0 & & C & B & C \\ 0 & 0 & 0 & 0 & & 0 & C & B \end{bmatrix}$$

B, C are block matrices of order n >= 2

A is of order M*n, where $M = 2^k-1$, k >= 2

Constraint Program for BOER

```
1 → BP[k-1] * x[2^(k-1)] == dP[2^(k-1)][k-1]

    AND
2 → AND  FOR (j 1 k-1)  {
        2 * CP[j-1] * CP[j-1] == BP[j] + BP[j-1] * BP[j-1]

        CP[j-1] * CP[j-1] - CP[j] == 0

        AND FOR (i 0 2^(k-j)-2)  {
            CP[j-1] * ( dP[f1(i,j)][j-1] + dP[f2(i,j)][j-1] )
                == dP[i*2^j][j] + BP[j-1] * dP[i*2^j][j-1]        }}

    AND
3 → AND FOR (j k-1 1)  {
      AND FOR (i 0 2^(k-j)-1)  {
          CP[j-1] * ( x[(i+1)*2^j] + x[i*2^j] )
            == dP[f3(i,j)][j-1] - BP[j-1] * x[(i+1)*2^j-2^(j-1)]   }}
```

Inputs=
{ BP[0],
CP[0],
dP[0][*]
}

Dependence Graph for BOER Algorithm

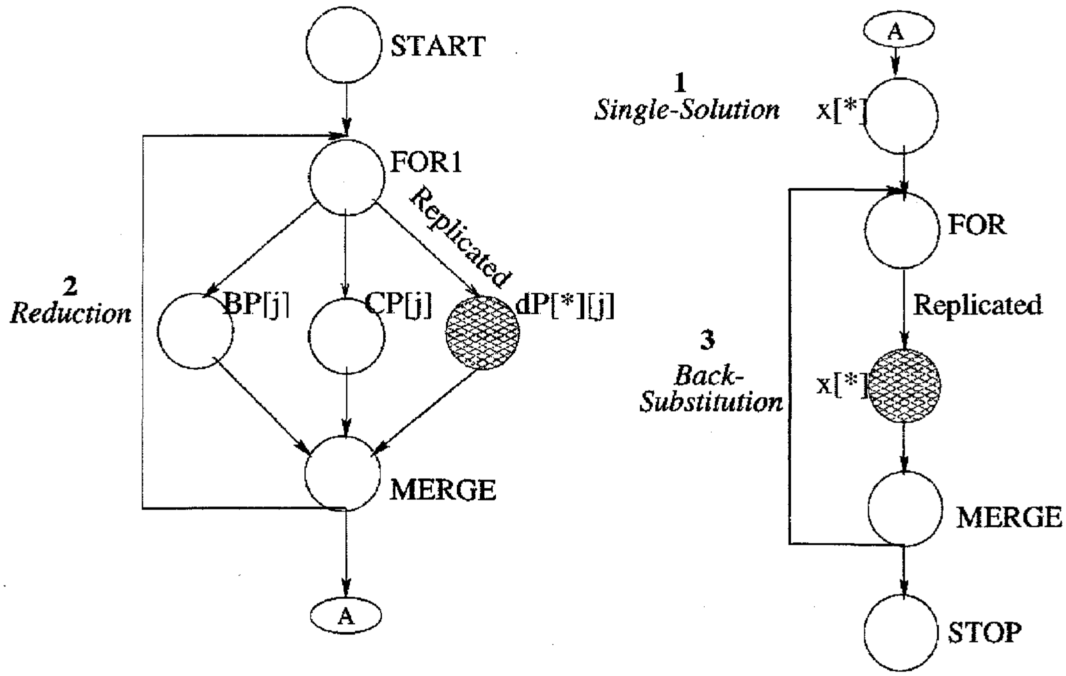

Summary

Constraint systems & appropriate sets of input variables

↓

generalized dependence graphs

↓

efficient parallel programs.

Granularity is controlled through
modules,
functions,
hierarchical types.

Both task and data parallelism can be targeted.

Constraints:
attractive choice for a representation for problem specification

Why ?

Declarative specification

Compiler has great freedom in the generation of efficient parallel structures

Extraction of different programs from same specification

Generation of complete or effective programs

Conclusion

Feasibility of compilation of efficient parallel procedural programs from constraint programs has been established.

Future Work

Short Term

- Incorporate powerful data partitioning mechanisms.
- Complete development of an execution environment specification.
- Relax restrictions on indexed sets

Long Term

- Extraction of algorithms for applications where algorithms specification is difficult.
- How to write "good" specifications ?
- Which path should be selected in effective programs ?
- Investigate other application domains.
- Applications:
  1. Boundary Matching Problems for Decomposed Domains
  2. Data Mediation
  3. Back Tracing for Circuits
  4. Development of Distributed Network Applications

Defining the Rules of the Net: The Visions and Lessons of Self-Governance

Jeffrey B. Ritter
Director, ECLIPS
Ohio Supercomputer Center
www.osc.edu/eclips

Defining the Rules of the Net

- The compelling force of greed
- The architecture of rule-making for the Net
- The visions and lessons of self-governance (so far)
- Recommended strategies

The Compelling Force of Greed-Why we make rules

- Protection of Property Assets
- Protection of Competitive Advantage
- Protection of Communities

Realities of the Net

- Networks, including the Net, are defined by their rules.
  - technical rules
  - process rules
  - dispute resolution

Realities of the Net

- Networks, including the Net, are defined by their rules.
- Information is emerging as a new species of property.

Realities of the Net

- Networks, including the Net, are defined by their rules.
- Information is emerging as a new species of property.

We will create rules for the Net
to protect the value of
our information assets.

How will we make the rules?

- Who is "we"?
- For what space?
- What are rules?

Rules--Controlling Assumptions

- He who makes the rules wins.
- To be free, we must regulate the Net.
- The rules will be driven by the commercial requirement for adequate returns on investment.

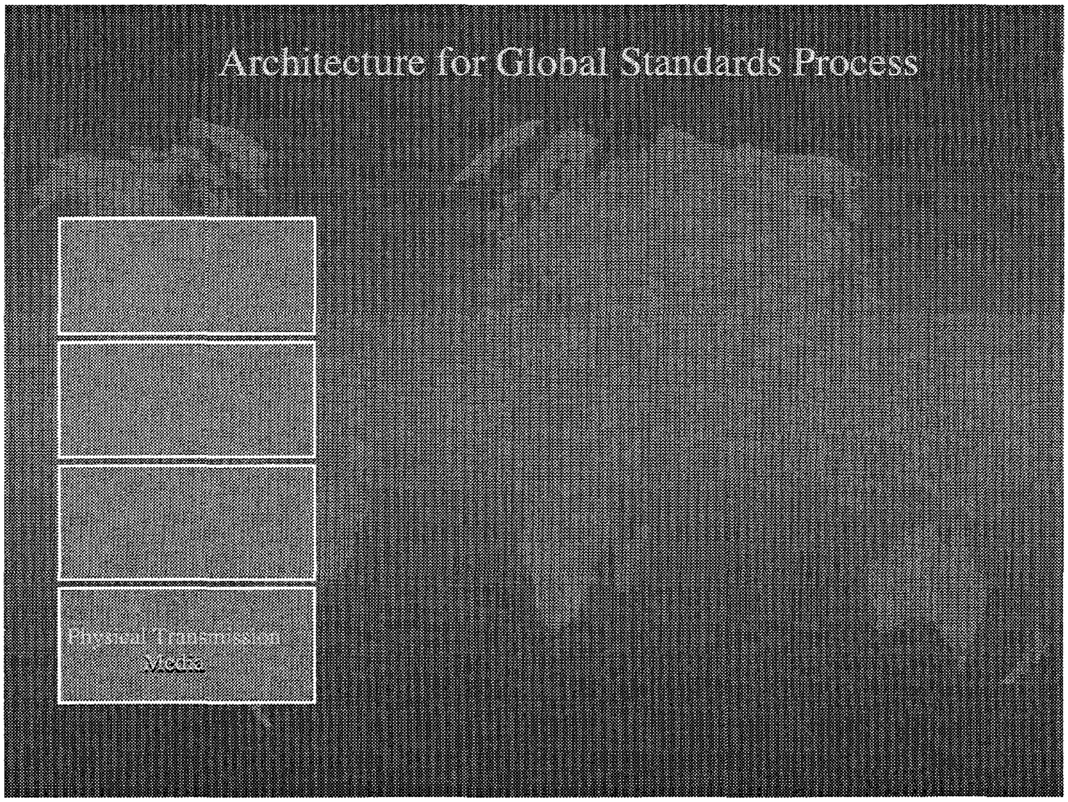
Architecture for Global Standards Process
Physical Transmission Media

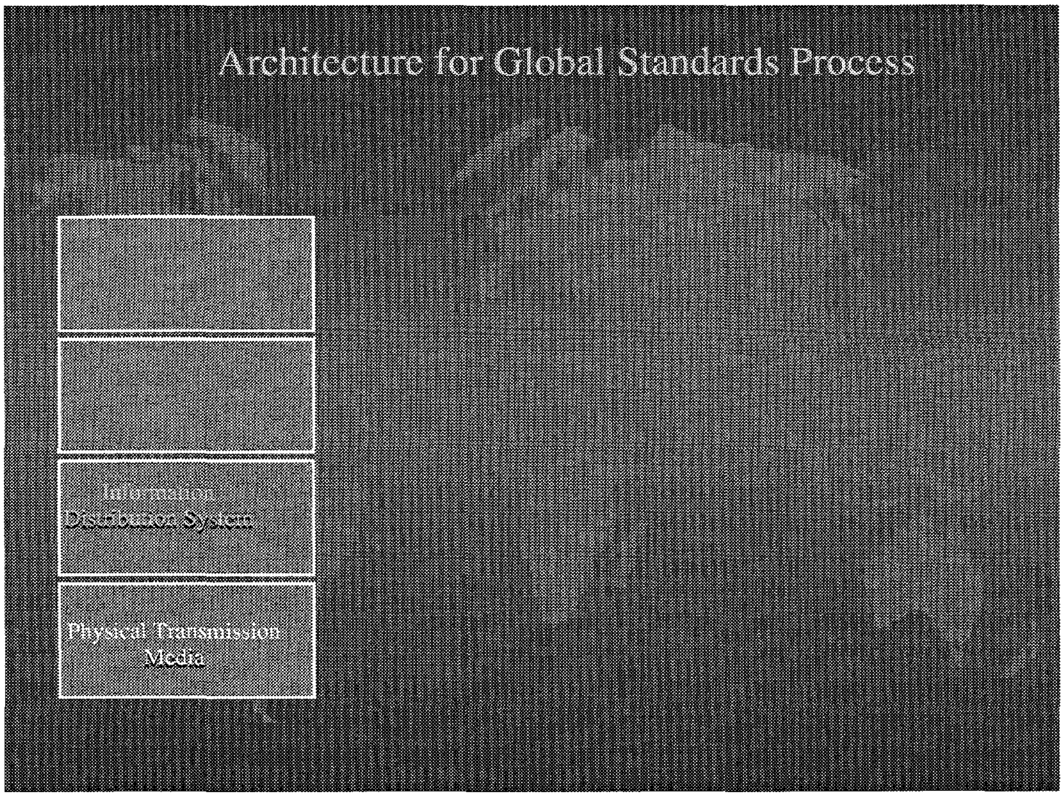
Architecture for Global Standards Process
Information Distribution System
Physical Transmission Media

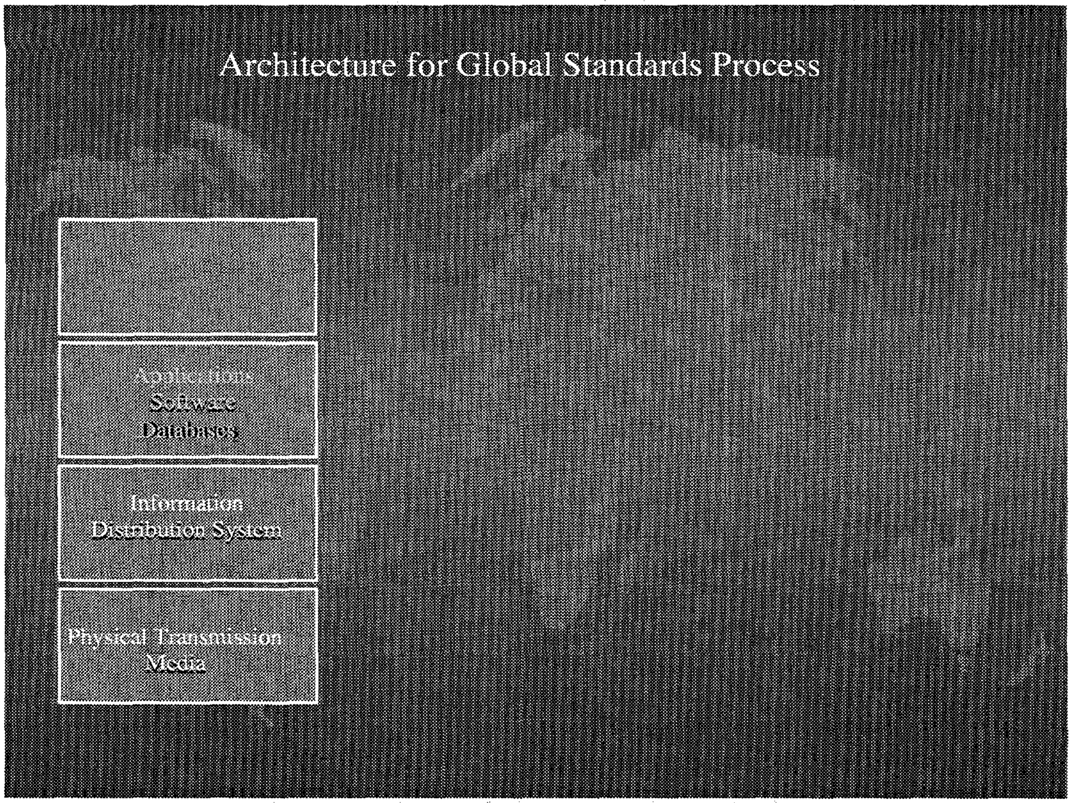
Architecture for Global Standards Process
Applications
Software
Databases
Information
Distribution System
Physical Transmission
Media

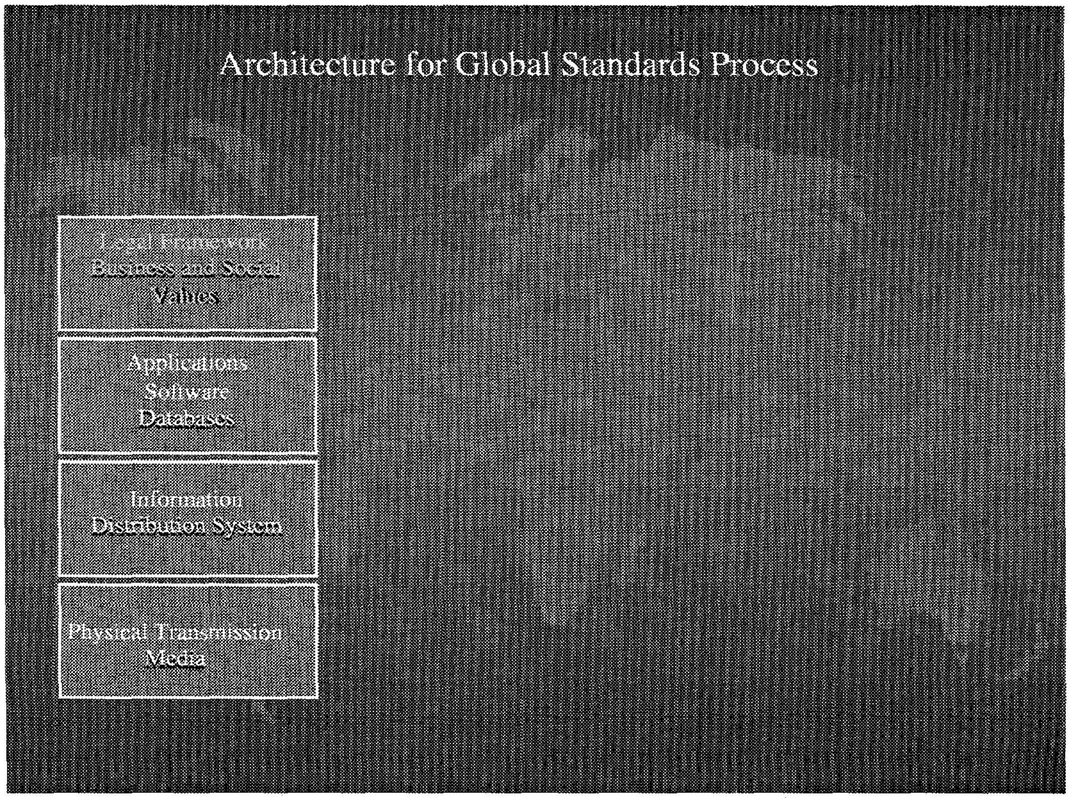
Architecture for Global Standards Process
Legal Framework
Business and Social
Values
Applications
Software
Databases
Information
Distribution System
Physical Transmission
Media

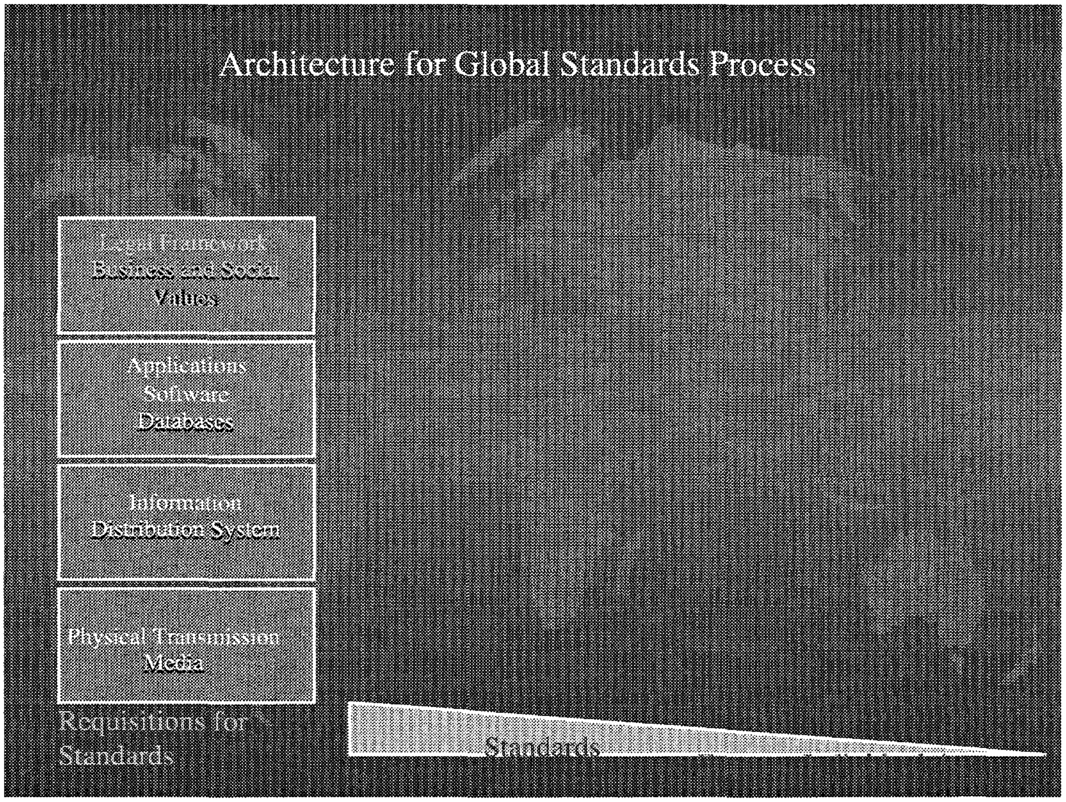
Architecture for Global Standards Process
Legal Framework Business and Social Values
Applications Software Databases
Information Distribution System
Physical Transmission Media
Requisitions for Standards
Standards

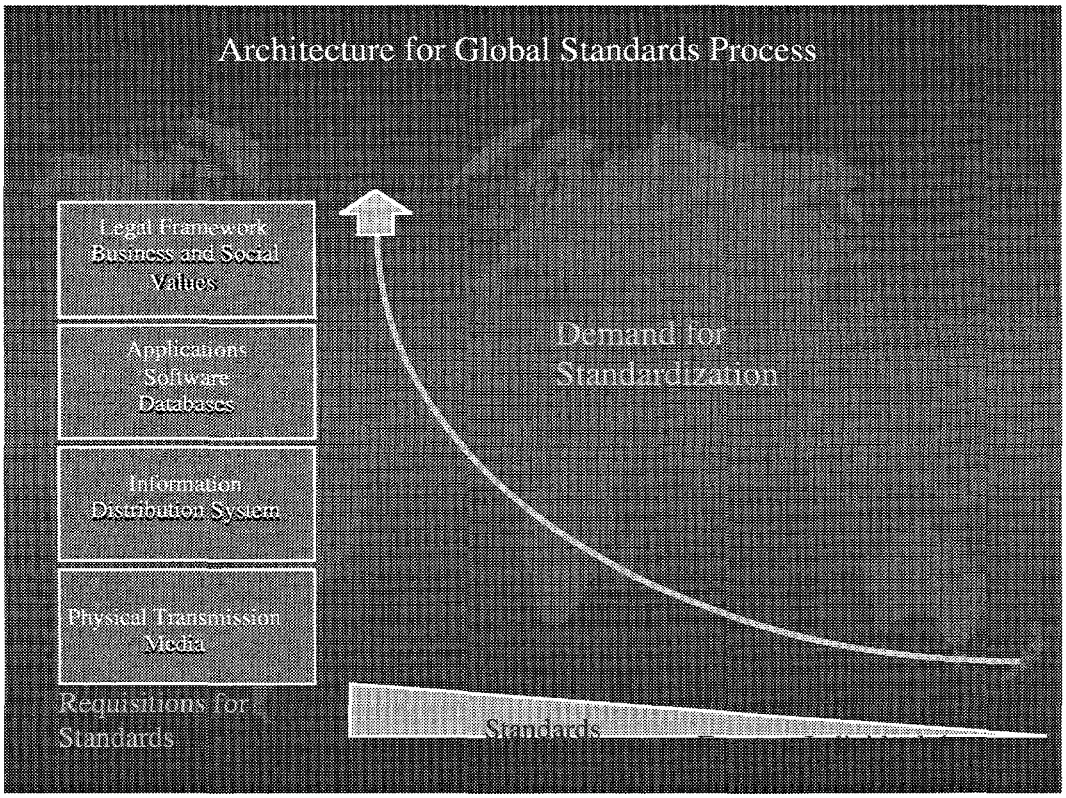
Architecture for Global Standards Process
Legal Framework Business and Social Values
Applications Software Databases
Information Distribution System
Physical Transmission Media
Demand for Standardization
Requisitions for Standards
Standards

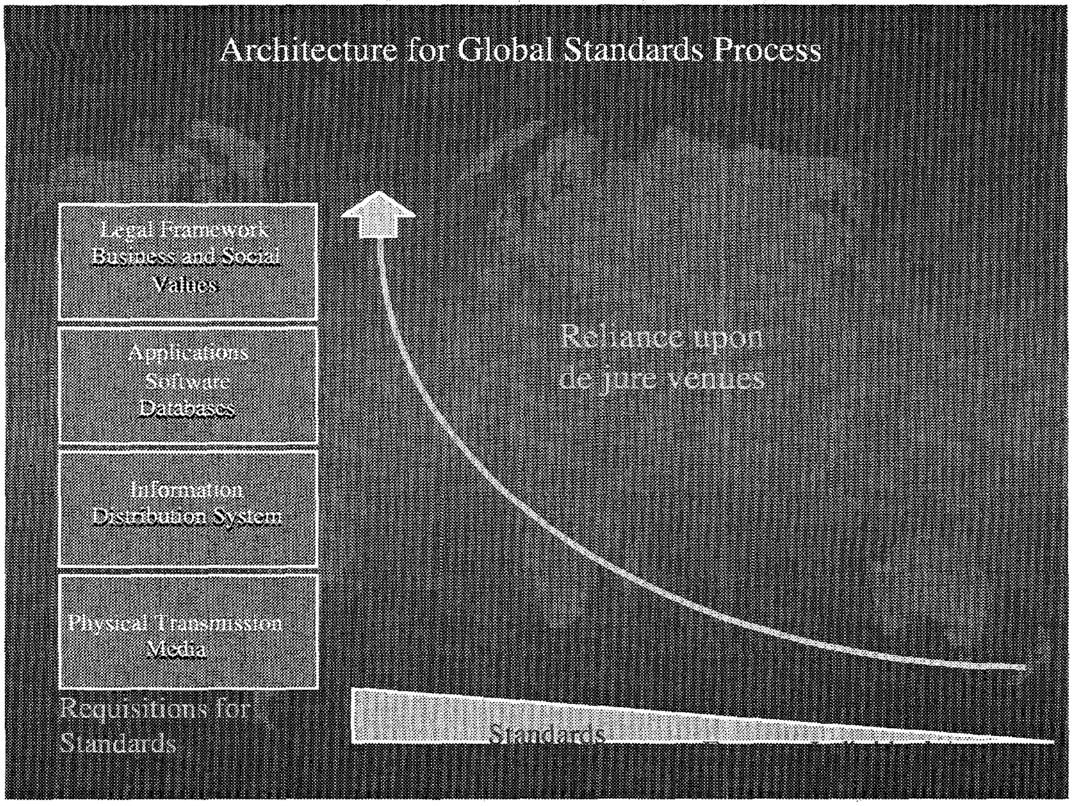
Architecture for Global Standards Process
Legal Framework
Business and Social
Values
Applications
Software
Databases
Information
Distribution System
Physical Transmission
Media
Reliance upon
de jure venues
Requisitions for
Standards
Standards

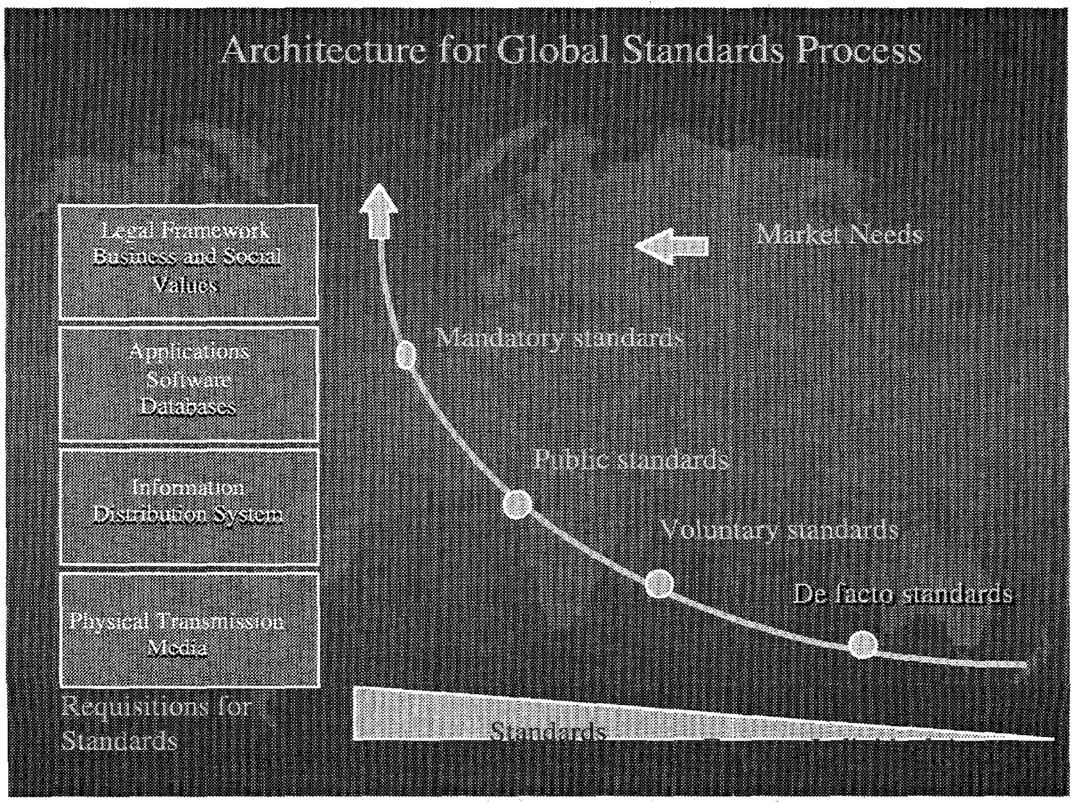
Architecture for Global Standards Process
Legal Framework
Business and Social
Values
Applications
Software
Databases
Information
Distribution System
Physical Transmission
Media
Market Needs
Mandatory standards
Public standards
Voluntary standards
De facto standards
Requisitions for
Standards
Standards

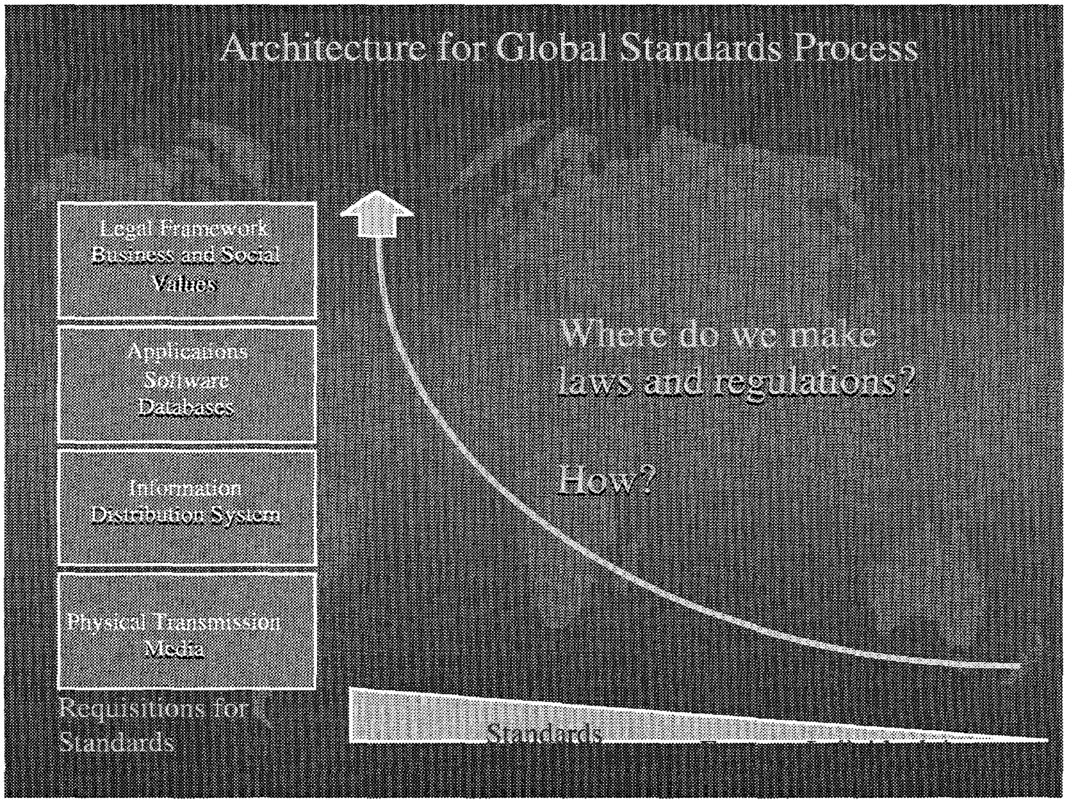

Architecture for Global Standards Process
Legal Framework Business and Social Values
Applications Software Databases
Information Distribution System
Physical Transmission Media
Requisitions for Standards
Standards
Where do we make laws and regulations?
How?

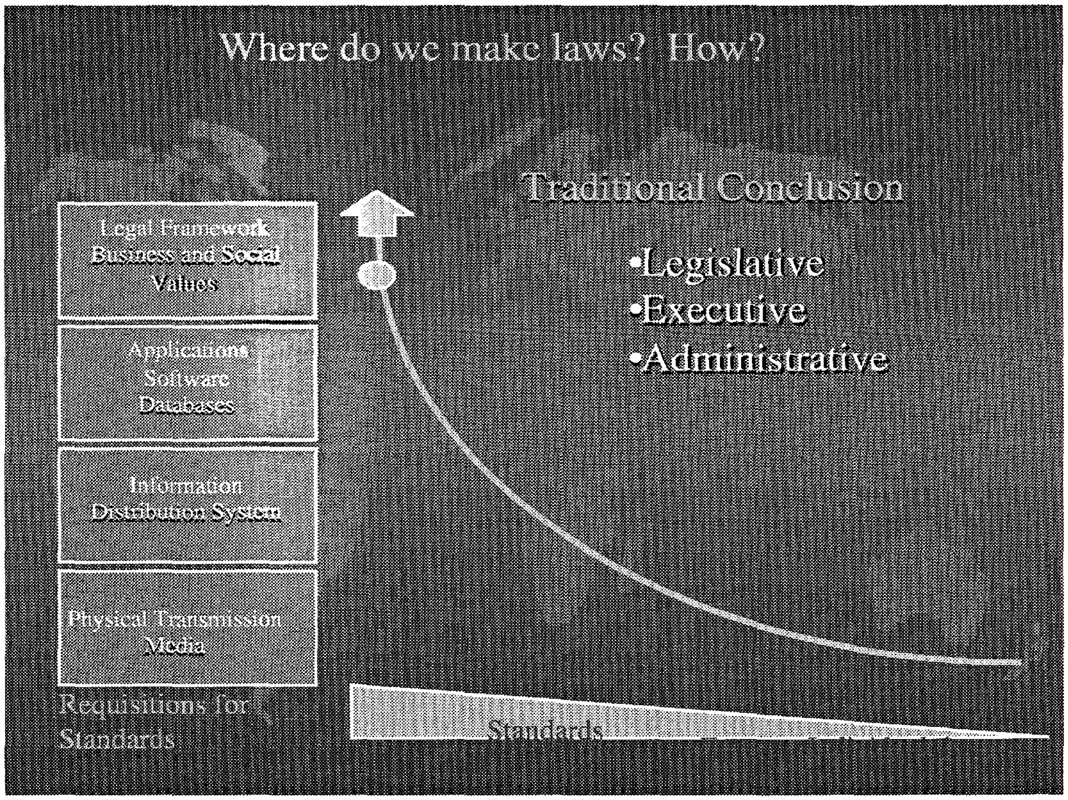

Where do we make laws? How?
Traditional Conclusion
•Legislative
•Executive
•Administrative
Legal Framework Business and Social Values
Applications Software Databases
Information Distribution System
Physical Transmission Media
Requisitions for Standards
Standards

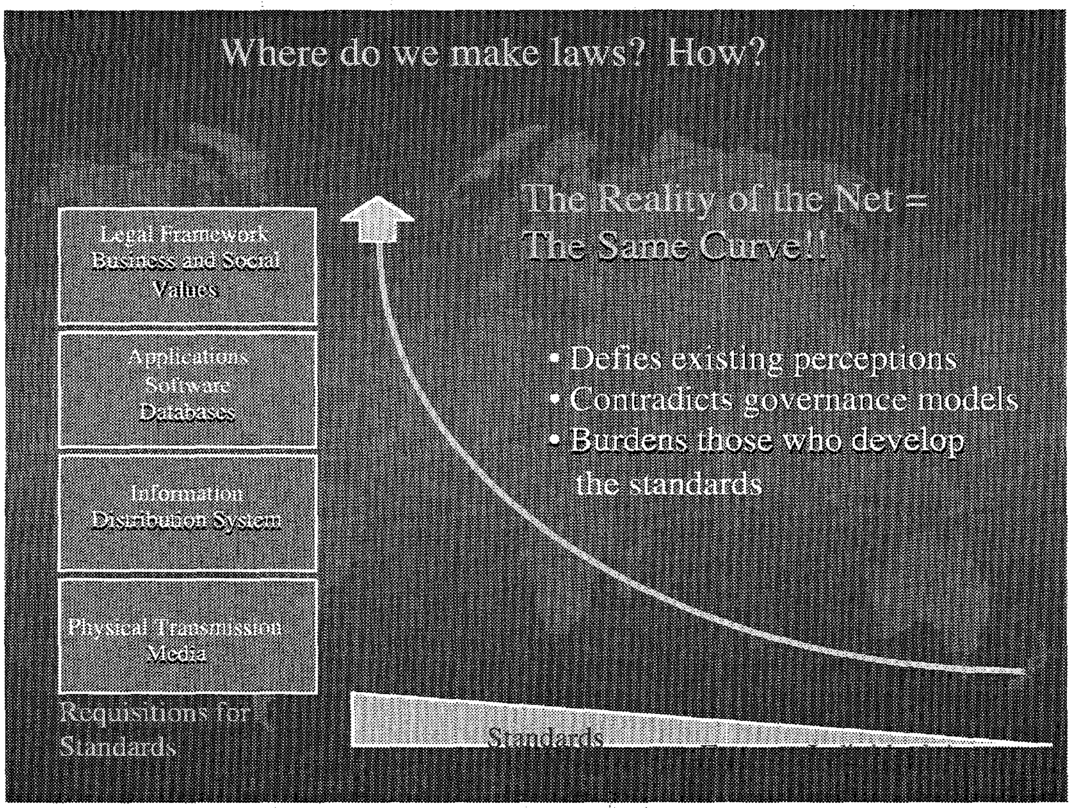
Where do we make laws? How?
The Reality of the Net =
The Same Curve!!
• Defies existing perceptions
• Contradicts governance models
• Burdens those who develop the standards
Legal Framework
Business and Social
Values
Applications
Software
Databases
Information
Distribution System
Physical Transmission
Media
Requisitions for
Standards
Standards

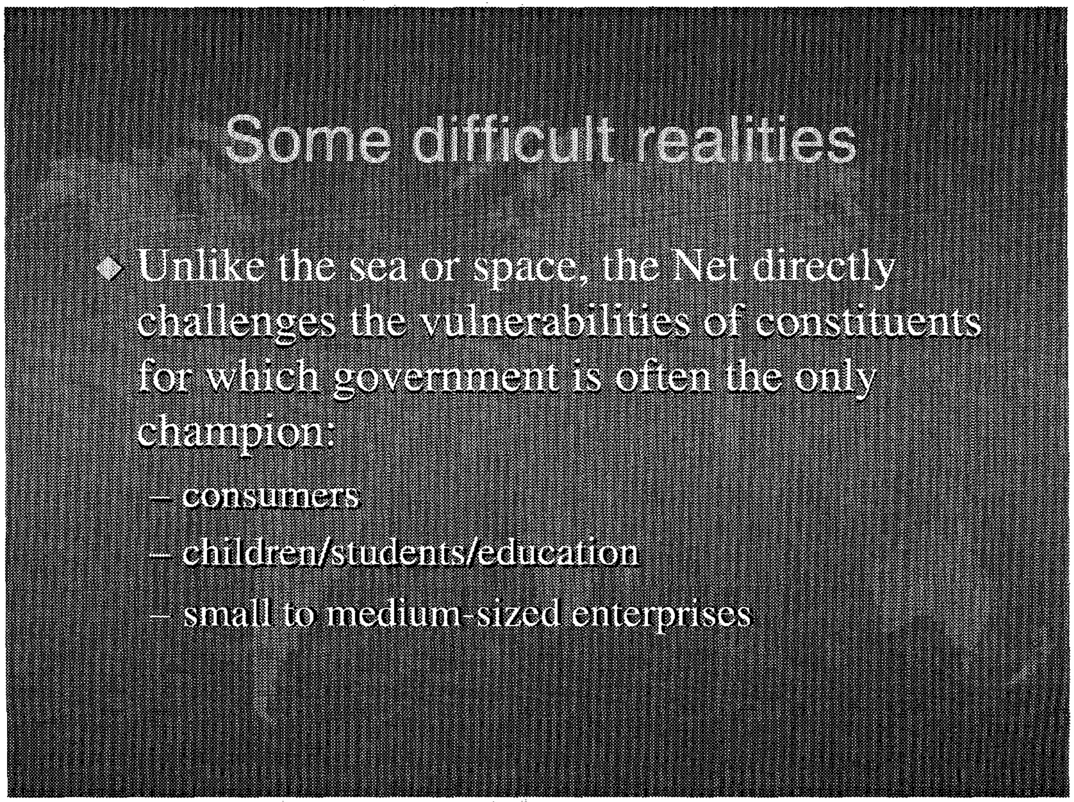
Some difficult realities
Unlike the sea or space, the Net directly challenges the vulnerabilities of constituents for which government is often the only champion:
– consumers
– children/students/education
– small to medium-sized enterprises

Some difficult realities

- Governments (fueled by the press) overreact to "market failure" in standards development
- Governments retain "veto power" over the results of self-regulation
- Standards development (including self-regulation) is capital intensive

The Consequences for Law-Making

- No "safe-harbors" readily provided for the results.
- No predictability for returns on investments.
- Reactive vs. pro-active attitudes among management and developers.

The Consequences for Law-Making

- No "safe-harbors" readily provided for the results.
- No predictability for returns on investments.
- Reactive vs. pro-active attitudes among management and developers.

[illegible]
costs for acquiring proprietary advantages
in the final layer of the standards infrastructure.

Essential Strategies

- Must encourage governments to defer the exercise of inherent jurisdiction
- Must empower standards development to embrace the complete architecture
- Must deliver methodologies for developing accelerated solutions for implementation

Deferral of Jurisdiction

- For electronic commerce, regulatory and executive sectors are prepared to defer:
  - US--Magaziner paper
  - Japan
  - European Union
  - Australia (content regulation)

Embrace the full architecture

- The gap between law and technology will diminish (and is diminishing).
- Law and policy venues are committing to technology-informed processes.
- Private sector is experimenting with self-governance.

Deliver the Methodologies of Governance

- Must develop and disseminate effective models for solutions process:
  - market-driven
  - agile
  - inclusive
  - global

Deliver the Methodologies of Governance

- Must embrace an international view-- government must be considered in the process and brought to the table.
  - technology standards have consequences for national and international policies.

Deliver the Methodologies of Governance

- Must empower the value of informed choices among:
  - model agreements
  - *mniform codes of conduct*
  - statements of best practices
  - icon-based community-entry
- Must promote 'safe harbors' for deferral of government jurisdiction.

Some portending consequences

- Private sector funding of rulemaking will enhance TNO's challenge to the nation-state.
- Bruce Sterling was right--the private sector will design enforcement power into its rules (and the new technologies).
- Diversity will undermine the consensus required for maximum security effectiveness.


Some portending consequences
The U.S. will not be able to control the policy momentum of Europe and Asia, putting at risk:
– commercial interests
– military security and management
– privacy of the individual

Privacy in the Digital Age

Deirdre Mulligan

At its core, the Digital Age represents a dramatic shift in computing and communication power. The decentralized, open nature of the network coupled, with an emphasis on user control over information, are central to achieving the First Amendment potential of the Internet. Through interactive technology, individuals today can enjoy a heretofore unknown ability to exercise First Amendment freedoms. Access to the Internet empowers individuals with an enormous capacity to speak and be heard, and listen and learn. The development of filtering and blocking devices that empower individuals to control the inflow of information gives new meaning to the core First Amendment principle that individuals should determine the ideas and beliefs deserving of expression, consideration and adherence.[1]

However, at this moment the impact of the Digital Age on individual privacy remains an open question. Will the Digital Age be a period in which individuals lose all control over personal information? Or does the Digital Age offer a renewed opportunity for privacy? The development of technologies that empower individuals to control the collection and use of personal information and communications - such as encryption, and anonymous remailers, web browsers and payment systems - are inspiring examples of the privacy-enhancing possibilities of interactive technology. However, we believe that the architecture of the Internet must be designed to advance individual privacy by facilitating individual control over personal information.

The rise of technologies that empower users of interactive communications media to affirmatively express control over personal information can fundamentally shift the balance of power between the individual and those seeking information. CDT believes this technological shift is possible and necessary, and offers us an unprecedented opportunity to advance individual privacy. However, this shift will occur if interactive media is harnessed to advance individual privacy.

Rather than responding to the very real risks posed by new technology with the Luddite-call of "smash the machine," we are calling for a reversal of the technological status quo by demanding that technology be designed to empower people. We should seize the opportunity to vest individuals with the information and tools to express their desire for privacy in clear and effective ways, and by having those desires acknowledged by information users, we can advance privacy. We believe that this post-Luddite approach will reinvigorate individual privacy in the Digital Age.

While strengthening existing laws, such as the Fair Credit Reporting Act and the Right to Financial Privacy Act, and enacting legislation to protect health records, are crucial to protecting individual privacy, individual empowerment technologies offer a powerful method of implementing the core principle of individual control where current gaps and weaknesses leave individual privacy vulnerable. We believe that user controlled technologies that enable individuals to protect the privacy of their communications and personal information, offer an unprecedented opportunity to extend real protections for individual privacy around the world.

[1]See Turner Broadcasting Syste. Inc. v. FCC, 114 S. Ct. 2445, 2458 (1994)

The TFLOPS Era is Here

Art Hale

Sandia National Laboratories

April 24, 1997

TFLOPS Culminates 50 Years of Supercomputing

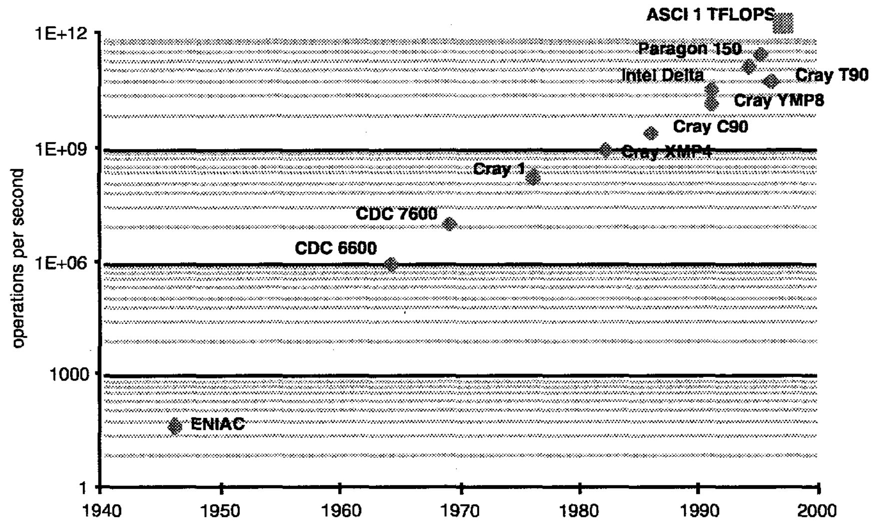

In The Beginning of MPP

"Parallel machines are hard to program and we should make them even harder – to keep the riff-raff off them." *... an expert*

"Conversion of any code to parallel takes a few weeks, perhaps longer." *... a manager*

"Massively Parallel machines are generally considered to be hard to program special-purpose engines, but the nCUBE is different." *... a salesman*

What's the difference between a computer salesman and a used-car salesman?

The used-car salesman knows when he's lying."

... anonymous

Heyday of MPP Competition

On speed:

"The nCUBE 2 supercomputers are the fastest computing systems available today." *nCUBE brochure, 11/91*

"The Touchstone Delta system is the world's fastest production computer." *Intel brochure, 11/91*

On Popularity:

"CM systems are the most popular parallel supercomputers in the world." *TMC document, 11/91*

"The iPSC/860 has become the most widely used parallel supercomputer in the world." *Intel brochure, 11/91*

On the Future:

"The CM-5 incorporates the first parallel supercomputer architecture that scales to TeraFlops." *TMC document, 11/91*

"Intel is the first to offer a proven technology path to take you to TeraFlops." *Intel brochure, 11/91*

There are three rules for programming parallel computers.

We just don't know what they are yet.

Gary Montry

Rules of Parallel Programming?

1) Focus on Scalability
2) Refer to Rule #1
3) Refer to Rule #2

Integrate Distributed Computing

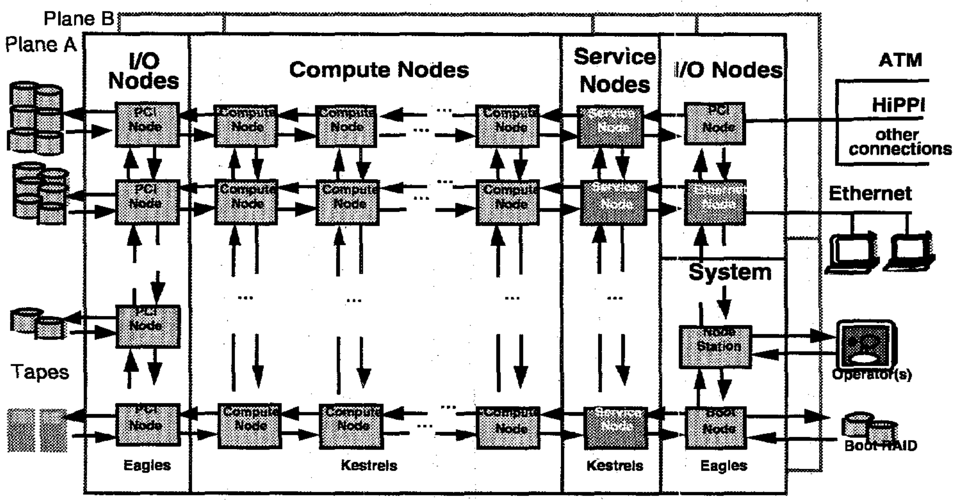

Specialized servers and services
Closely integrated "system image"

Balance Communication/Computation

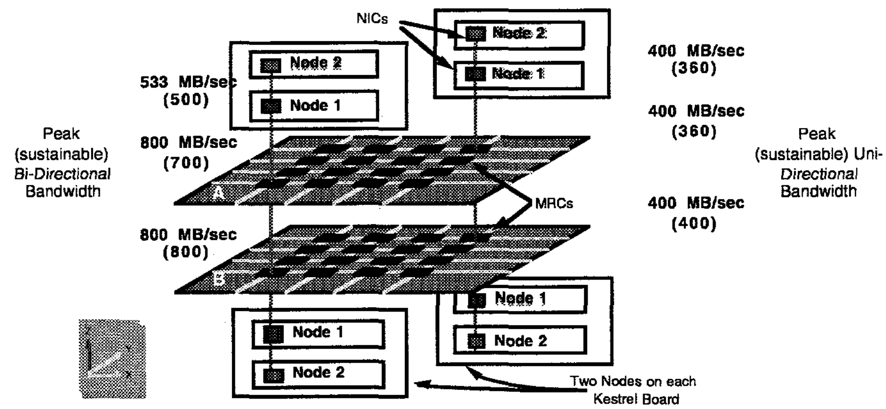

Scalable Operating System

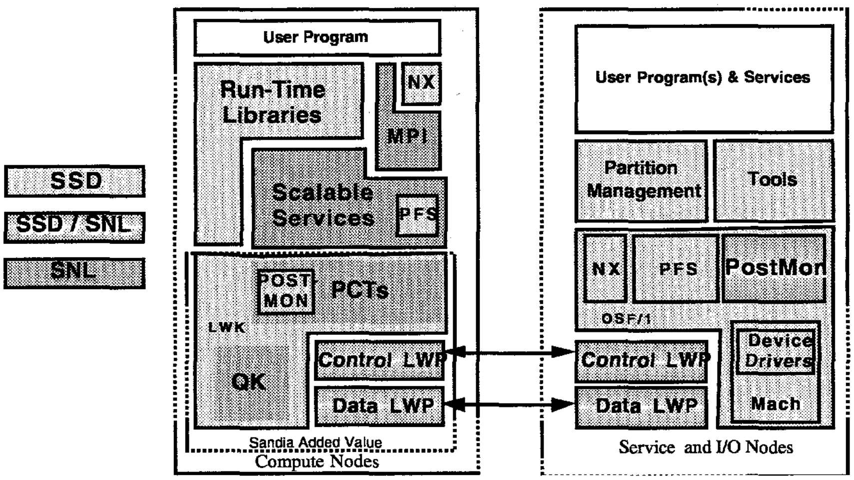

Build In Scalable RAS Features

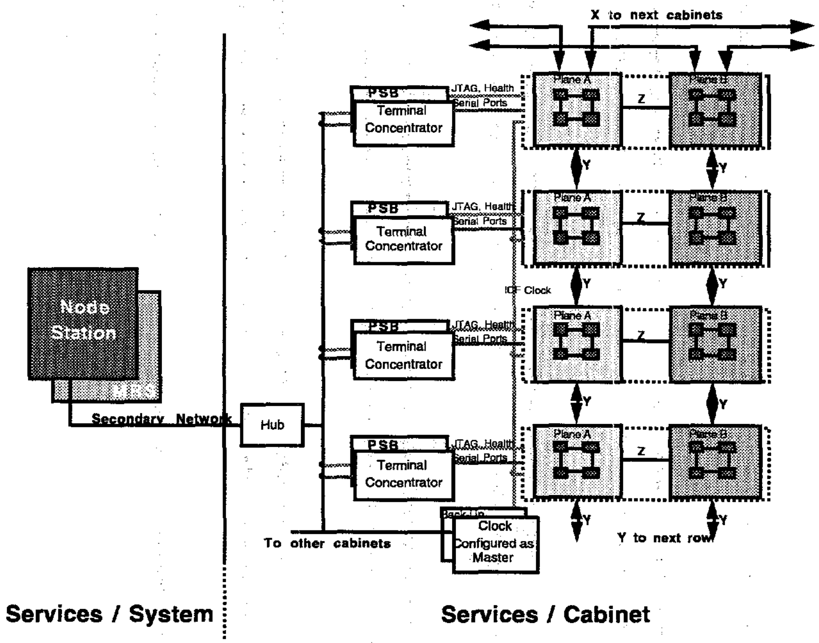

Scalable Administration

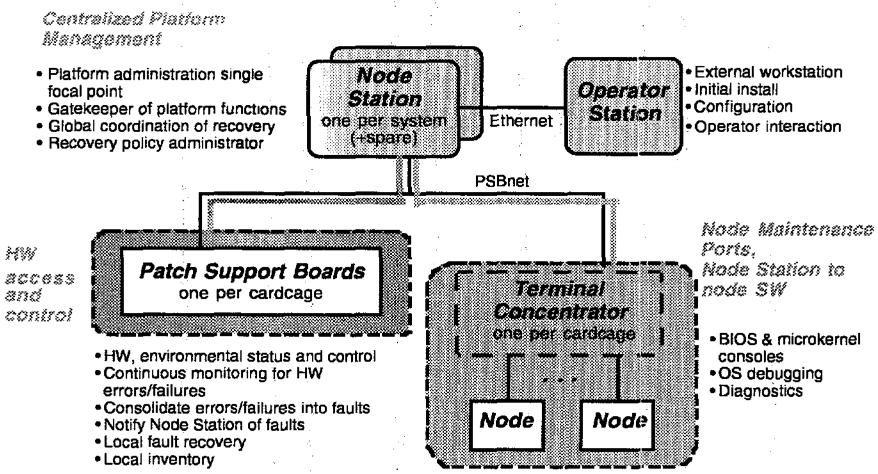

Scale Applications

Materials
- electronic structure (BANDPW)
- quantum chemistry (QUEST)
- MD for molecules (ParBond)
- MD for metals (ParaDyn)
- grain growth evolution (PGG)
- gas separation and adsorption
- GCMD for diffusion
- Quantum Monte Carlo
- Hartree Fock Methods
- electron microscopy (PMC)

Signal and Image Processing
- SAR image simulation (SRIM)
- radar signatures (XPATCH)
- phase gradient autofocus (PGA)
- seismic modeling
- acoustic beam formation for ASW
- magnetic resonance imaging

Engineering Simulations
- shock physics (PCTH)
- low density flows (DSMC)
- chemically reacting flows (SALSA)
- structural mechanics
- integrated shock/structures (ALEGRA)
- combustion
- inverse scattering for semiconductor manufacturing
- ocean modeling
- electromagnetism
- aerodynamics
- integrated thermal/solid mechanics (DELTA)

Information Processing
- cryptanalysis
- tracking and correlating

"When you come to a fork in the road – take it."

Yogi Berra

Some Future Challenges

- **Unbundled building blocks**
  - **nodes, interconnect technologies, system software, programming environment**
- **Load balancing and heterogeneity**
- **Abstract system views**
  - **e.g. application configuration independence**
- **Automated system discovery and configuration (plug & play)**
- **Fault tolerance**
- **Simplifying system integration**

PetaFlops in 2007 is Consistent with Historical Evolution

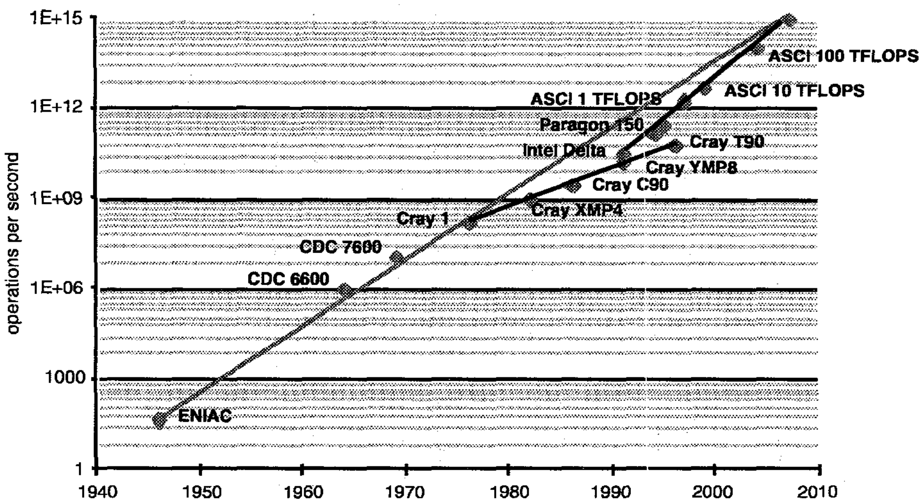

"What is the world's fastest computer and how fast is it?

Currently, it's an HP notebook. It's used on the Space Shuttle to compute orbital position and has been clocked at 17,500 mph."

Robert Hyatt

ASCI Applications Challenges

Ken Koch
ASCI Code Developer
Applied Theoretical & Computational Physics
Los Alamos National Laboratory
krk@lanl.gov

ASCI Applications Overview

What is ASCI?

- Accelerated Strategic Computing Initiative
  - Goal
    - Computational simulation as a tool for addressing safety, reliability, and performance issues in the absence of nuclear testing
  - Structure
    - 1 Program/3 Labs (LANL, LLNL, Sandia,DOE HQ)
    - Applications & Problem Solving Environment
    - Platform strategies
      - Red, Blue Pacific, Blue Mountain, ...(and more)
    - Alliances, Pathforward, ...

Simulation Objectives

- Physics
  - 3D highly resolved multi-material complex geometries with large deformations
  - multiple spatially-collocated coupled physical processes
- Computer Science
  - portability & standards
    - languages, APIs, distributed & shared memory
  - scalable parallel processing
    - 1 million to 1 billion mesh sizes
    - Order(1K to 10K) CPUs

Programming Approaches

- Start with current defacto standards
  - portability & longevity
  - better leveraging through broad user-base
  - lower risk
- Building blocks
  - C++, Fortran90, & C languages
  - MPI & domain decomposition
  - add multiprocessing as 2nd step
    - a wish for good parallelizing compilers
    - within each domain or across domains

Major LANL ASCI Applications Projects

Blanca Project

- Regular structured grids

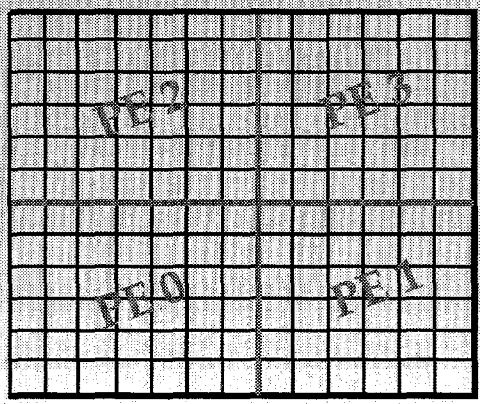

X(imax,jmax)

X(i,j) - X(i+1,j)

Blanca Codes

- Tecolote-2
  - multidimensional Eulerian/ALE hydro
  - POOMA C++ OO framework
    - data distribution, methods, & MPI/threads
  - multiple ordered data blocks per PE
    - domain overloading
    - variable data per cell
- Euler (quick port)
  - Fortran90 & MPI "cshifts"
  - one ordered data block per variable per PE
    - 1 or M data per cell

Crestone Project

- Regular base grid with cell-by-cell adaptive refinement

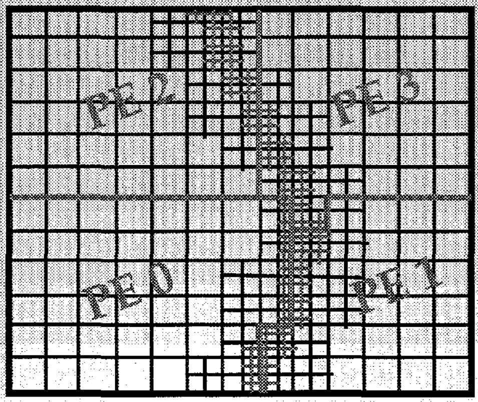

X(ncell)
Face(nface,2)
Index(nface,2,2)

X(Index(i,LO,Xdir)) -
X(Index(i,HI,Xdir))

Crestone Codes

- Rage90
  - multidimensional Eulerian rad-hydro
  - Fortran90 with MPI communications
    - MPI global gather/scatter exchanges
  - oct-tree data stored in 1D arrays with index arrays for connections
    - arbitrary cell ordering within each PE
    - one dynamic array per "physics" variable
    - variable data can be packed per cell
    - data count changes every time step
  - export controlled

Shavano Project

- Unstructured polyhedron grids

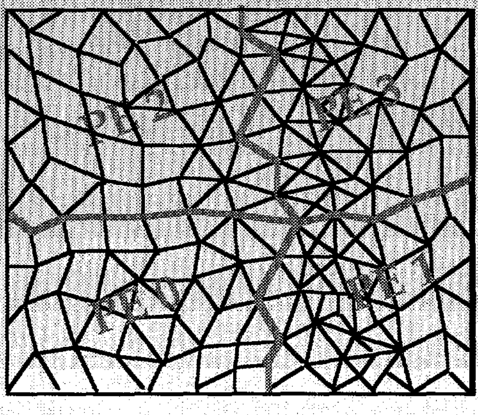

X(imax)
Index(neigh)

X(i) - X(Index(n))

Shavano Codes

- Flag code
  - multidimensional multiple-grid hydro
  - Fortran77 & C & data storage framework
  - polyhedral finite difference meshes
  - 1D data arrays with storage manager
- Chad & Dante codes (alternate track)
  - 3D Eulerian hydro and transport
  - Fortran90 & MPI library for gather/scatter
  - hexahedral non-moving finite elements
  - 1D data arrays with index arrays

Parallel Scalability Issues: Now and in the Future

Load Balancing

- spatial decomposition at least
  - 10^4 to 10^6 cells/PE (think crudely: 20^3 to 100^3)
  - possible additional decompositions over energy, etc., but with more communication
  - possible domain overloading
- dynamic in time
  - even if the grid is static!
    - e.g. mixed materials & turbulent jets
  - everyone is moving toward adaptivity
  - work monitoring & migration necessary
  - more domains aggravate the situation

Load Balancing

- multi-physics timestep
  - inherent imbalances
    - mixed materials, regional physics, particle field vs. grid views, AMR levels
  - timestep example:
    1. global explicit physics
    2. zone-local (zero-D) region-restricted physics
    3. global implicit physics
  - different decompositions per physics?
    - more data motion/copying
    - higher communications

Sparse Matrix Solvers

- currently CG with diagonal scaling
  - easy to implement in message passing
  - it works and is robust for now
- distributed matrix-vector multiply
  - only ~2 FLOPs per communicated element
  - irregular or adaptive grids
    - on-PE memory gathers
    - off-PE communications gathers
    - variable number of coefficients per row

Sparse Matrix Solvers

- millions to billions of unknowns
  - bigger problem sizes on bigger machines
  - convergence problems
    - huge number of degrees of freedom
    - are existing methods robust enough?
    - required iterations likely to increase
  - is parallel irregular multigrid an answer?
- performance
  - time-to-solution **not** parallel speedup
  - latency problems because of limited FLOPs per communicated element

Software Development

- compilers/languages
  - C++ and Fortran90 need to mature more
  - parallelizing compilers
    - support for C++, Fortran90, & C with mixed-language compatibility
    - scalability limit & tradeoffs with MPI
  - Fortran90 + MPI
    - copy-in/copy-out + asynchronous operations
    - user-derived types
    - "choice" arguments & strong typing
    - Fortran95 & MPI-2 are not enough
  - true parallel languages or extensions

Software Development

- debuggers
  - ➢ need solid support for the latest languages and newest features
  - ➢ support for all parallel features
    - – MPI
    - – irregular distributed data
    - – parallelizing compilers
    - – threads
  - ➢ usability on large parallel systems
  - ➢ visualizations instead of prints

Software Development

- parallel I/O
  - ➢ simple methods not scalable
    - – 1 file per MPI process
    - – single PE I/O with MPI gather/scatter
  - ➢ need a proven widely-supported API standard and good implementations
    - – parallel reads/writes to single logical files on high performance parallel filesystems
  - ➢ need for high-level self-describing files, not just fast low-level binary I/O

Better Physics

- better fidelity demands not just higher resolution but better physics as well
- more effort to develop & incorporate new physics since they need to be parallel implementations

The job is never done!

Day-to-Day Programmatic Usage of 100 TFLOP/s Systems Demands Careful Balance in the Overall Computing Environment

ASCI Session at Salishan Conference
April 23, 1997

Dr. Mark K. Seager
Asst. Department Head
Scientific Computing and Communications Department
seager@llnl.gov

Computation Directorate

Salishan Conference April 21 - 25, 1997 ASCI Session 97-515

ASCI Platforms Roadmap

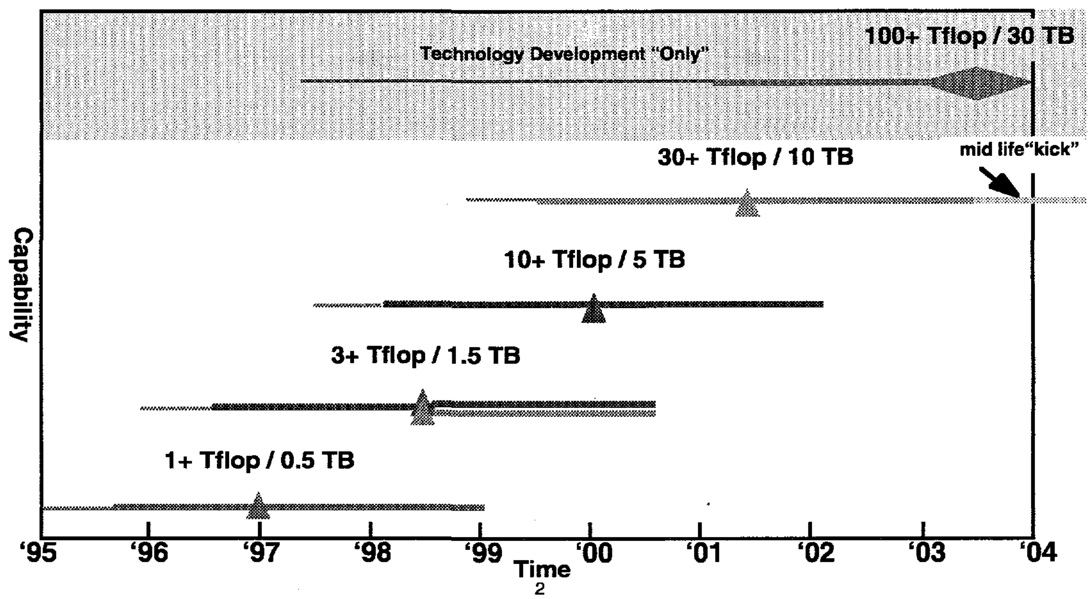

2

Setting for 100 TF Machine

- ★ CY2004
  - » Al Gore starting second term
  - » BART establishing service to SFO
  - » Bullet Train between Sacramento and LA
- ★ Technology Trends
  - » "Processors" achieve ~8 GF performance
  - » 100 TF machine is ~12,500 processors
  - » Power about 5-10 MWatt
  - » Cooling about 1,500-2,300 Tons (3-4 MWatt power)
  - » 5,000-15,000 FT^2 floorspace

3

Balanced System for Science Based Stockpile Stewardship

- ★ *Mission Driven Ratios*
  - » 2.5 TB Memory
  - » 2.5 Tb/s Bi-Sectional B/W
  - » 3.0 TeraFLOP/s Peak
  - » 75 TB Disk
  - » 90 GB/s I/O B/W

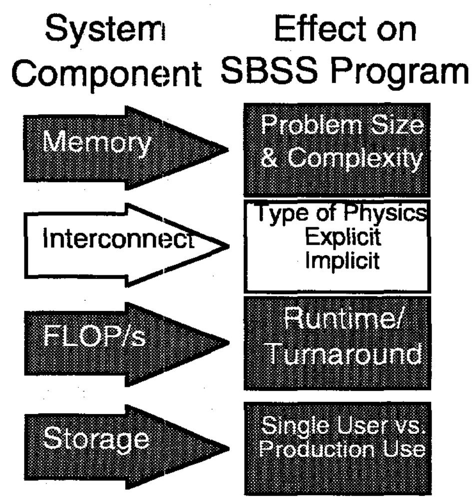

4

Typical Method Comparison Calculation Regime

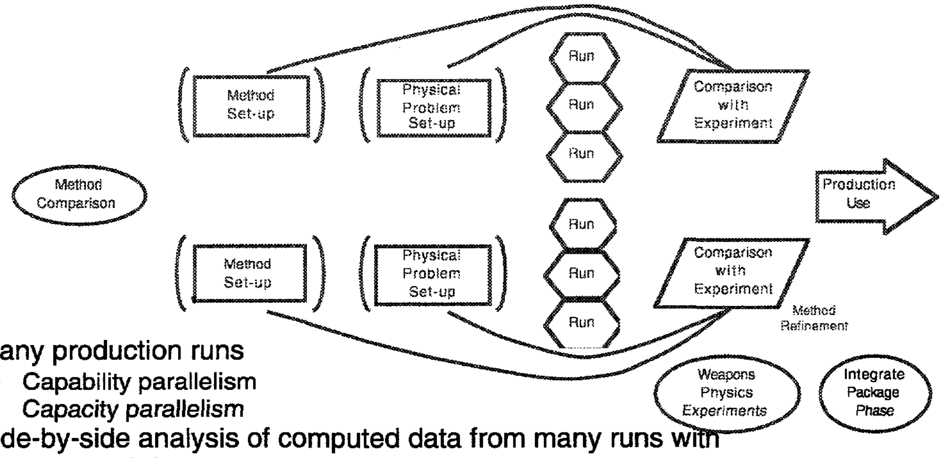

- Many production runs
  - Capability parallelism
  - *Capacity parallelism*
- Side-by-side analysis of computed data from many runs with experimental data
- One, two and three dimensional data comparison; many involving time dependent variables
- Method set-up and physical problem set-up visualization intensive and very time consuming

5

Typical Set-up and Production Utilize Multiple Platforms with Different Levels of Performance

20-30 Users	5-20 Problems 20-40 Runs	50 CPU Hrs 15-150 GB Memory 5-50 GB Disk 10 GB/s I/O 1 Gb/s Network	100-1000 CPU 15-150GB Memory 1.5-15 TB Disk 300 GB/s I/O 100 Gb/s Network	150 CPU Hrs 45-450 GB Memory 4.5-45 TB Disk 900 GB/s I/O 300 Gb/s Network

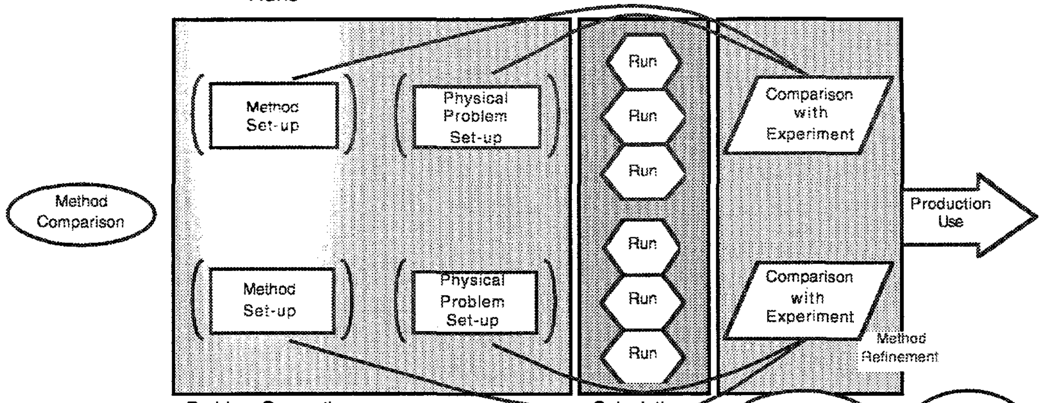

Platform:	Problem Generation Desktop, Local Server	Calculation Setup SMP Cluster	Calculation SMP Cluster or ASCI USCP	
Stress Point:		SMP Memory	CPU, I/O Bandwidth, Disk	SMP Memory Network, I/O, Disk

6

Preperation for "Full-Up Calculation" is Very Labor Intensive

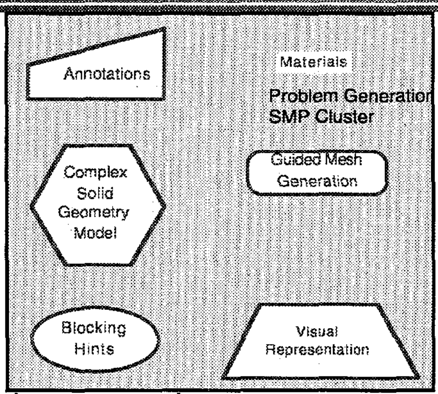

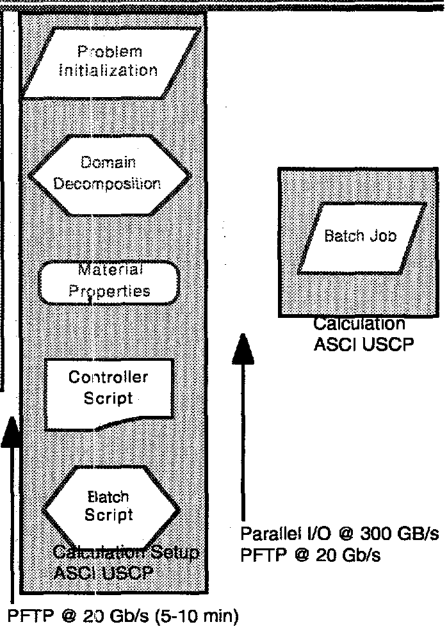

- ★ Very complex geometries
  - » Large number of small details
- ★ Innovative visual techniques needed
- ★ Multiple platforms engaged
- ★ Multi-step process
- ★ Human driven or guided process

7

Secondary "Full-Up Calculation" Scenerio

- **Specific A-Division 3-D Production Example**

CPU
50 TFLOP/s = 80 Hrs runtime
= 1 Week wall = 10^6 CPU Hrs

Memory
30TB Memory (~20 Billion Zones)

Swapping
Hourly, 5% Overhead, Read/Write=>2x
75% total memory on machine
=> 416 GB/s I/O to local disk

Archive
50% Job memory, 15 min intervals
=> 16.7 GB/s I/O to tape
320 restarts => 48 PB

Disk
20 BZ = 2 TB/Frame for vis
15 min intervals @ 5% sync = 40.4 GB/s
320 frames = 640 TB

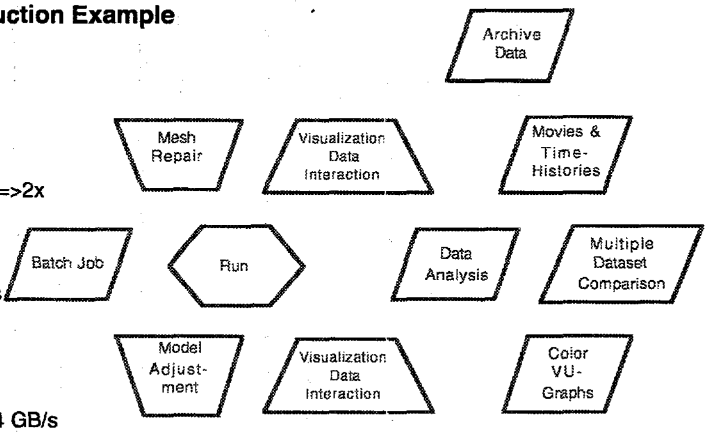

- ★ A calculation is not just a batch job. It is a series of batch jobs
- ★ Most calculations are "steared" in a batch sense.
- ★ Frequent interaction with archived data
- ★ Frequent need to visualize mesh, mesh based variables and materials

8

CY04 Ultra-Scale Computing for Immersive Data Interaction

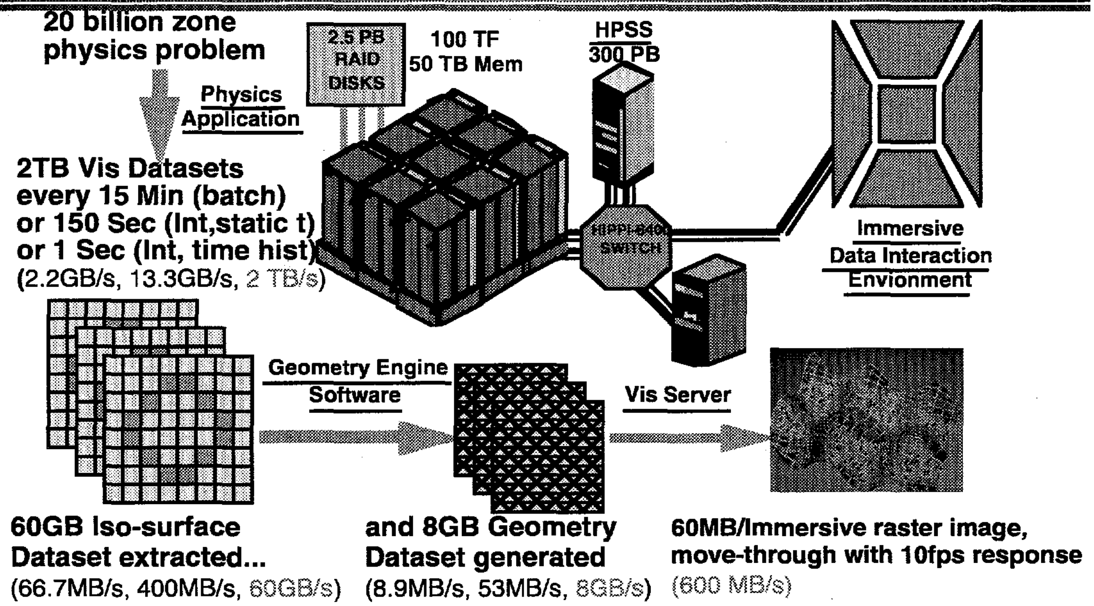

9

The Key to a Usable System is Application Driven Scaling

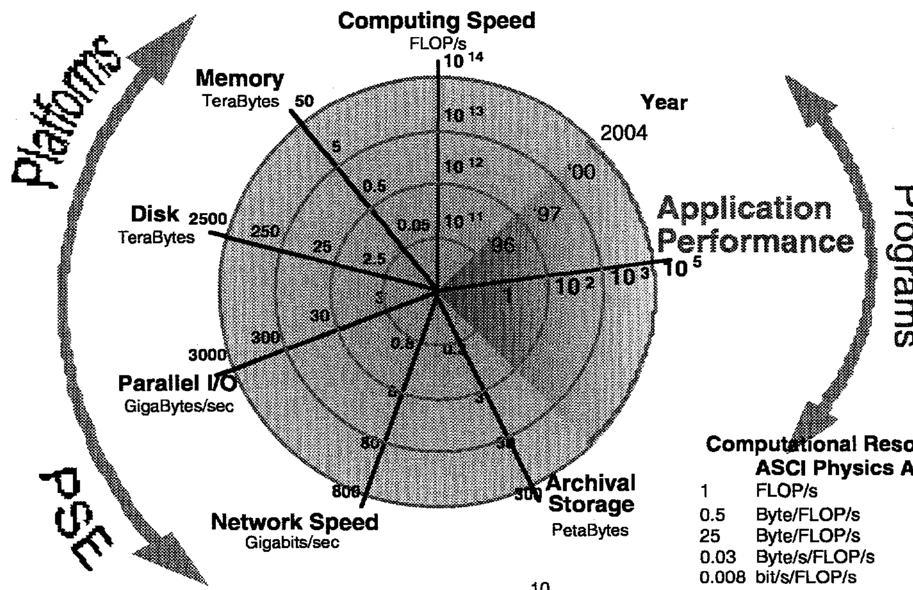

Computational Resource Scaling for ASCI Physics Applications

1	FLOP/s	Peak Compute
0.5	Byte/FLOP/s	Memory
25	Byte/FLOP/s	Disk
0.03	Byte/s/FLOP/s	Peak Parallel I/O
0.008	bit/s/FLOP/s	Peak Network
3000	Byte/FLOP/s	Archive

10

Designer productivity risk due to archival storage limitations

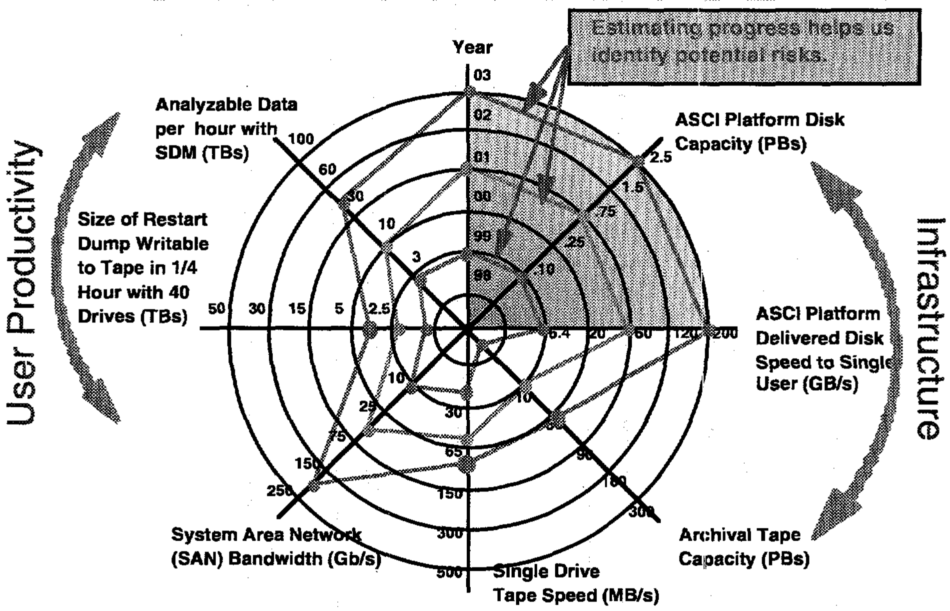

11

Summary

- 100 TF computers will give ASCI users new capability that would otherwise be unattainable.
- The computing environment around these machines must be carefully scaled up as well because this is a huge impact on the day-to-day productivity that ASCI users will experience.
- Multiple classes of machines (full spectrum) are utilized at various stages in the problem set-up, calculation and analysis phases of a series of runs.
- In the 100 TF computing regime, the "big-data" problem becomes the "huge-data" problem. Innovative approaches to storage, networking and visualization must be developed and deployed on a large scale.

12

Examining ASCI Computing Models

by

Karl-Heinz A. Winkler

LANL

Present Situation

- Experienced Radical Changes
- No Nuclear Tests, compute-based approach
- Last Decade: Great Variety of Architectures
- Realigned CRI, Convex Phased out Intel
- RIP CCC, KSR, TMC, ...
- Mass market FLOP/s less demanding
- DOE ASCI 10x every 4 years
- Moore's law 4x every 3 years, last 25 years
- Application development on time critical path

Can we establish a period of stability of several (many) years for the ASCI programming models?

What will be the successful programming models for ASCI class computers?

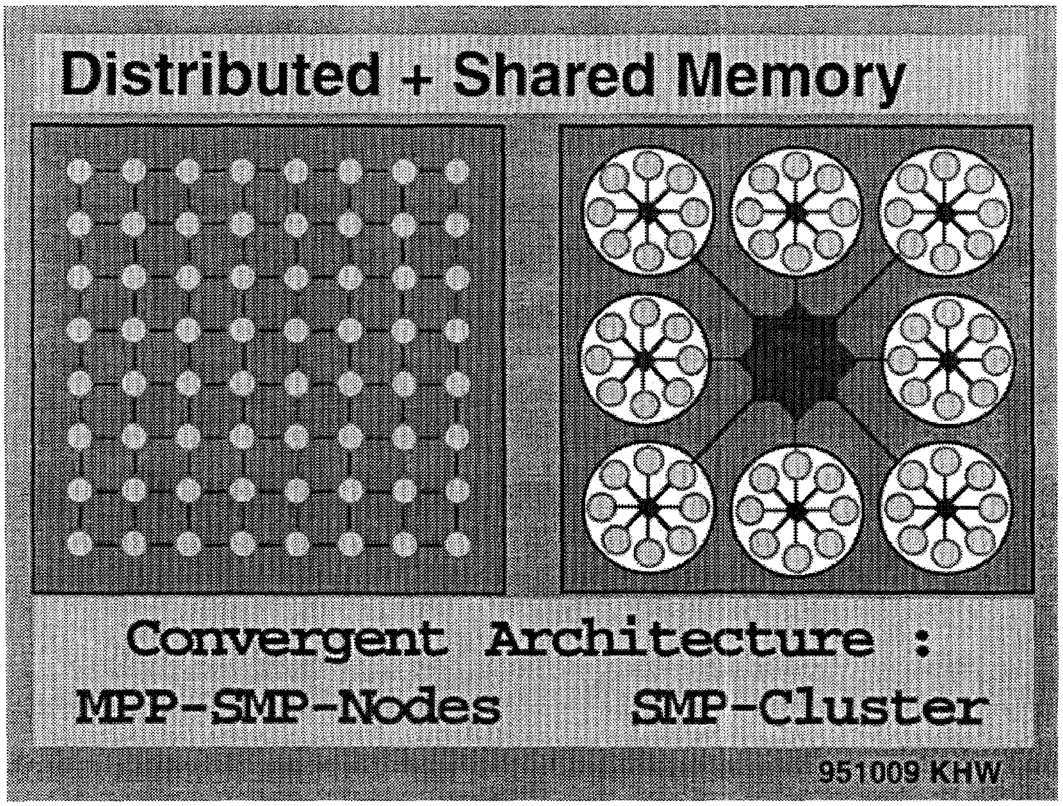

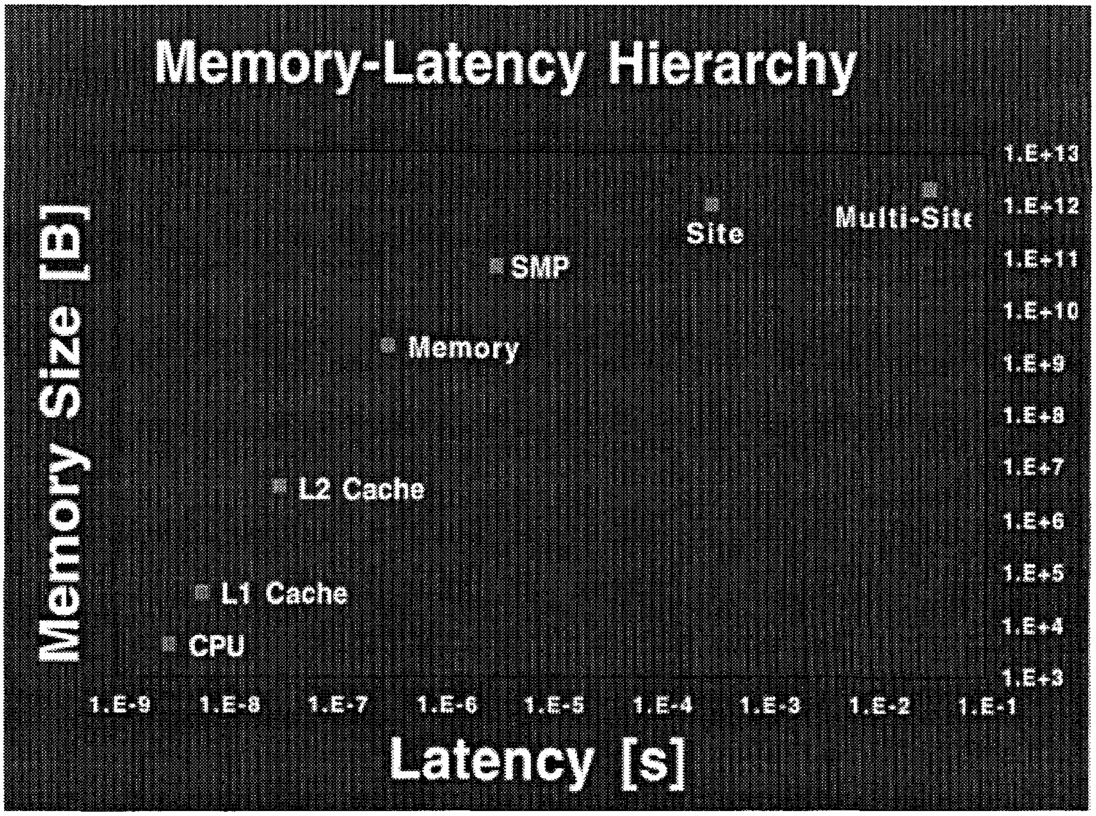

Hierarchical Distributed Shared Memory Programming Model

- **Hierarchical**
  - Take Advantage of Memory-Latency Hierarchy
  - Utilize Memory (Contiguous Data Layout), Cache of SMPs
  - Complete Algorithm Update on Small Subset of Zones
  - No Partial Algorithm Update on all Zones (traditional)
- **Distributed**
  - Explicit Message Passing between SMPs, MPI
  - Logical Shared Memory across Cluster, DSM
  - Exploit Surface to Volume Effect of SMP Boundary Zones
  - Overlap Computation and Communication
- **Shared Memory**
  - Multi-Tasking, POSIX Threads, Java
  - Make Single SMP behave like one Unit in Cluster Context

array d(X,Y, V)

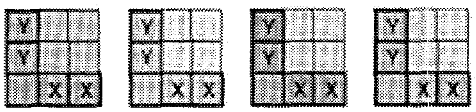

X-sweep

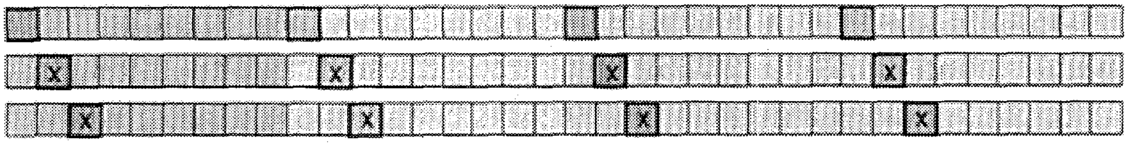

Y-sweep

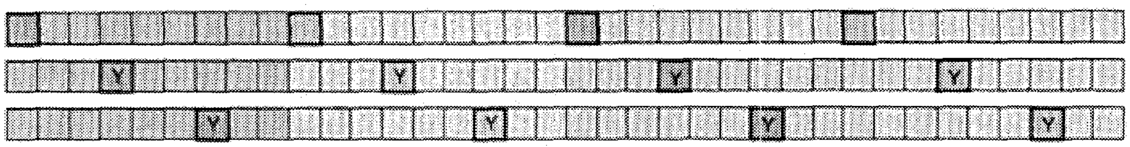

array d(V,X,Y)

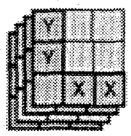

X-sweep

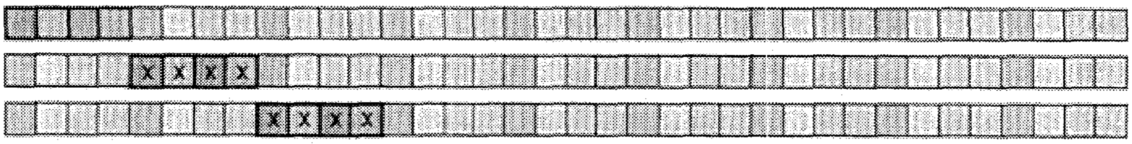

Y-sweep

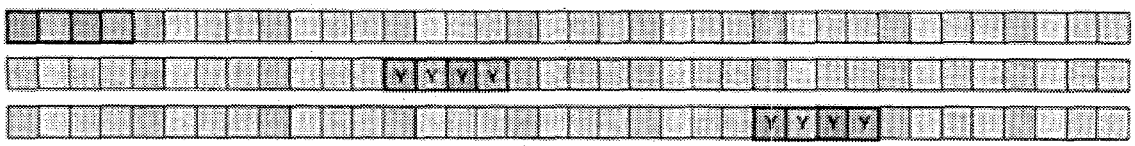

DIGITAL 300 MHz Alpha 21164

$d_i = d_i + d_i * d_i$, vpz=16, 8 Byte Words, f90

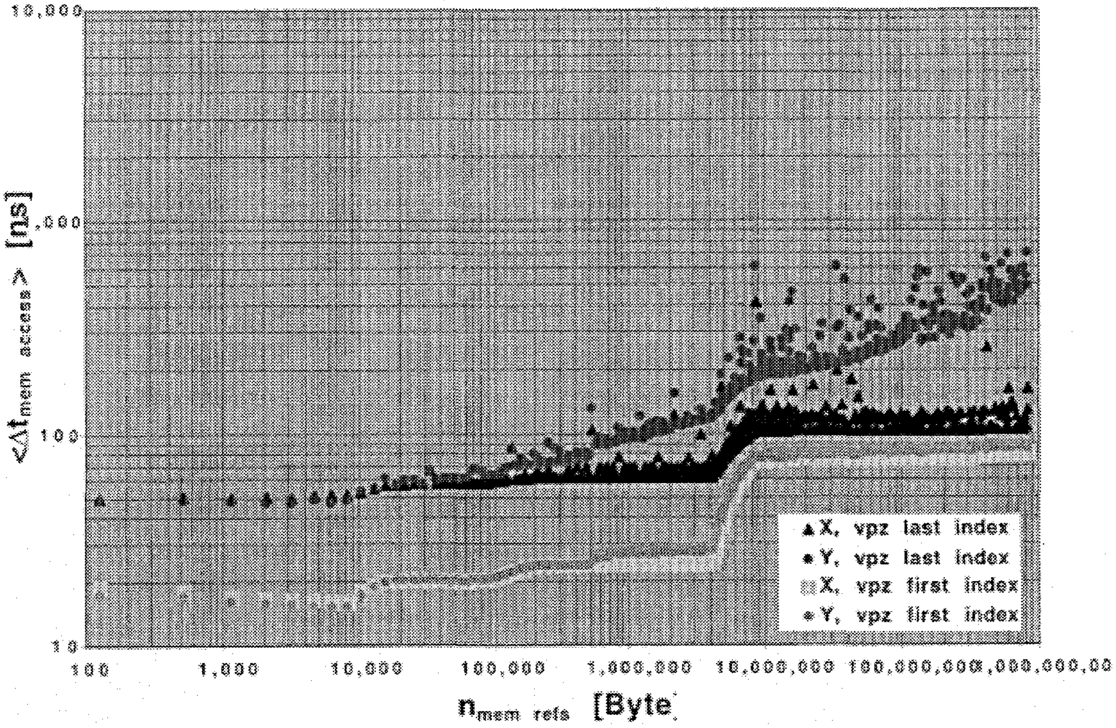

SGI 190 MHz R10k

d_i=1., vpz=4, 8 Byte Words, k77

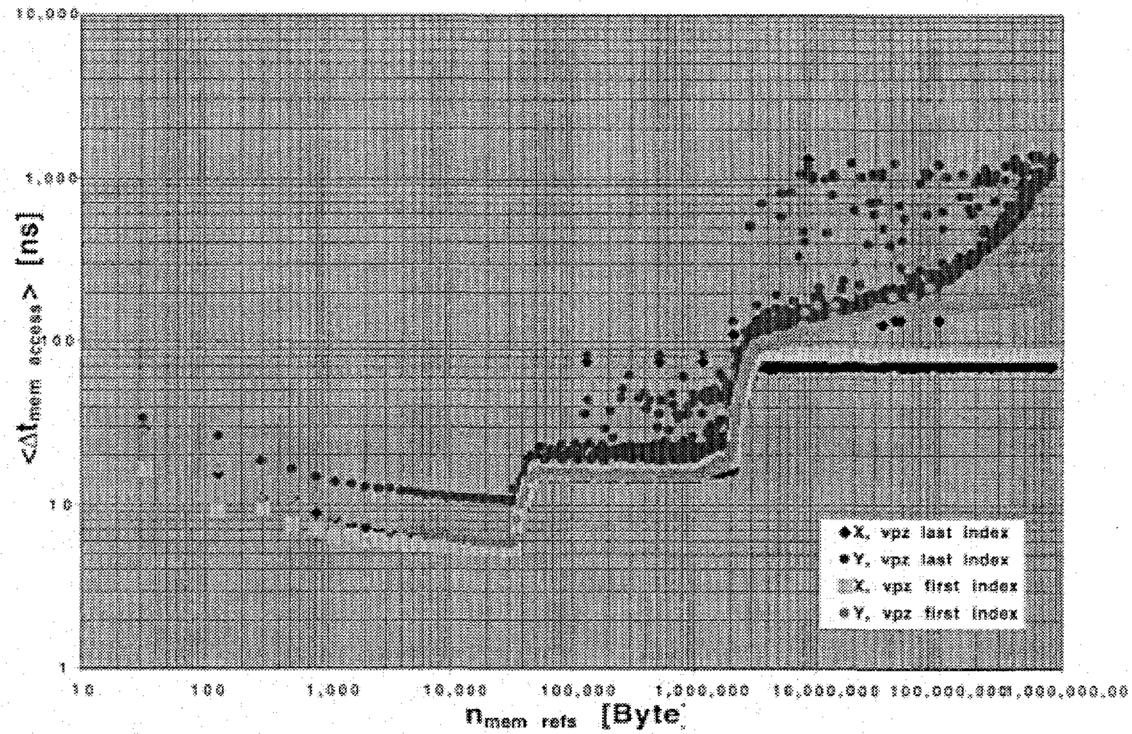

array d(V,X,Y) blocked

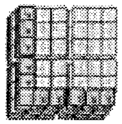

X-sweep

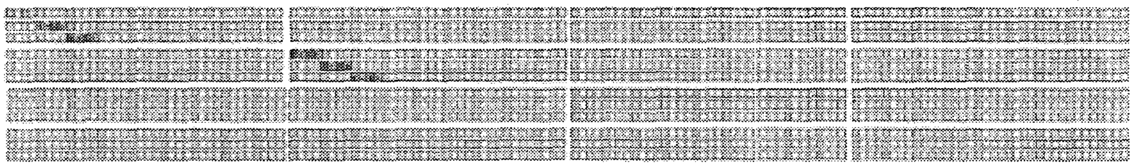

Y-sweep

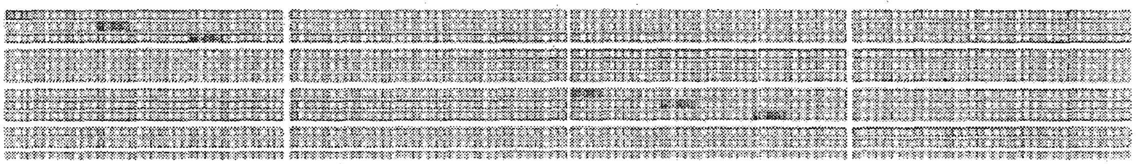

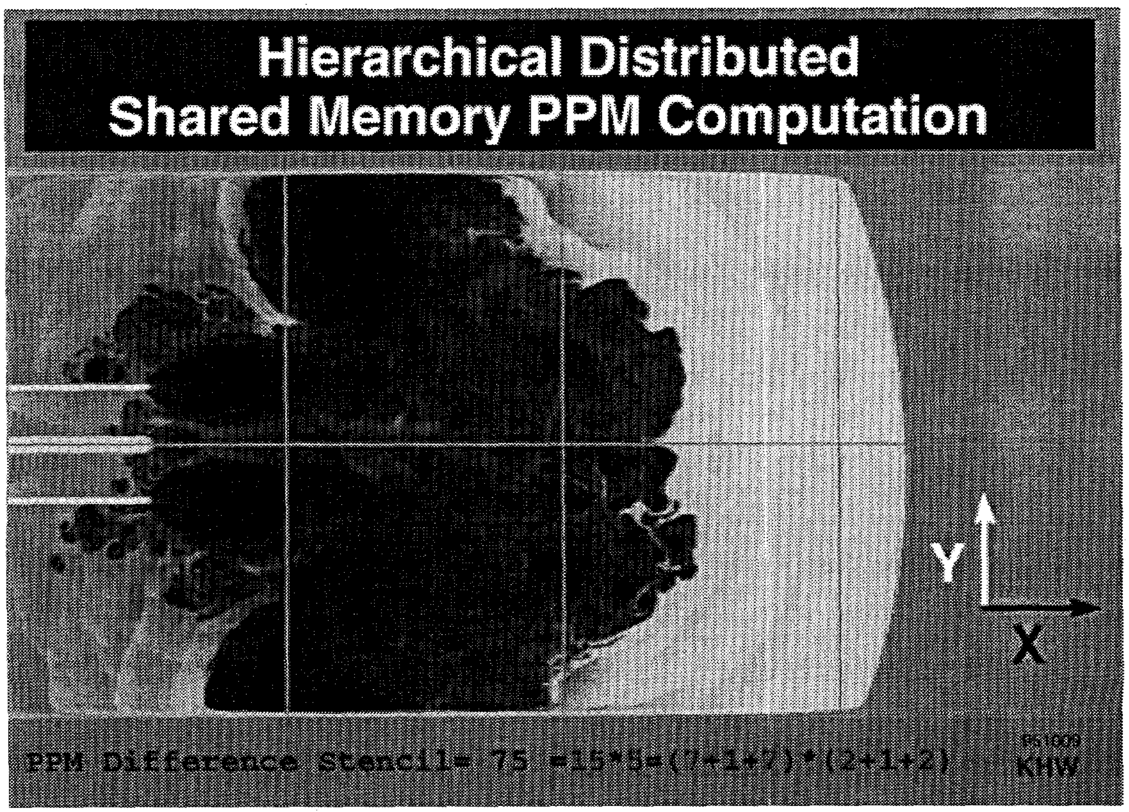

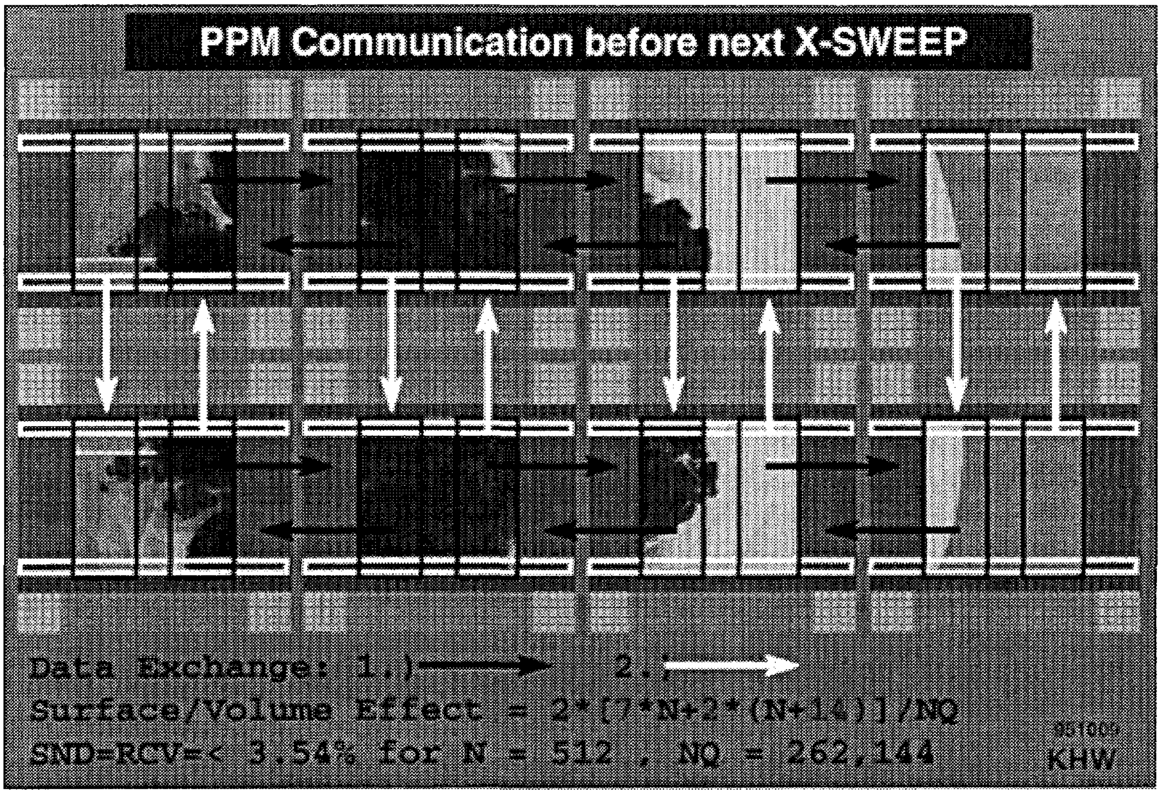
PPM Communication before next X-SWEEP
Data Exchange: 1.) 2.)
Surface/Volume Effect = 2*[7*N+2*(N+14)]/NQ
SND=RCV=< 3.54% for N = 512 , NQ = 262,144
951006
KHW

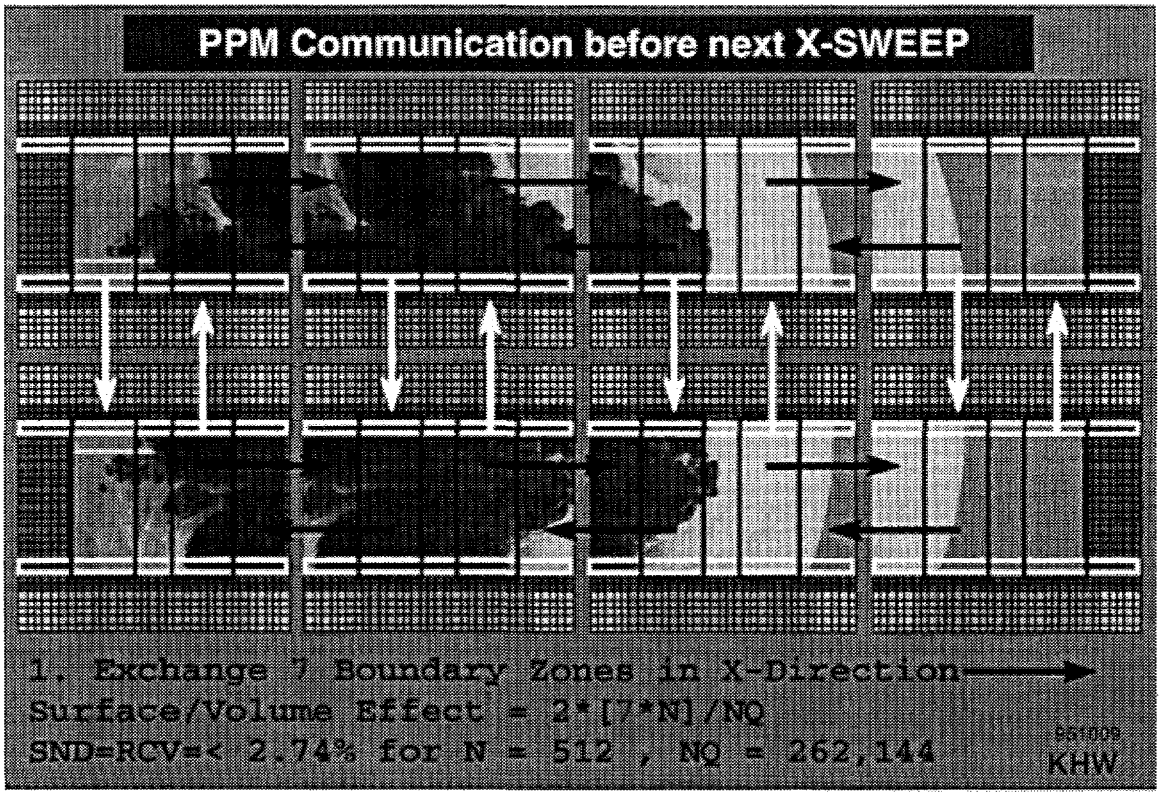
PPM Communication before next X-SWEEP
1. Exchange 7 Boundary Zones in X-Direction
Surface/Volume Effect = 2*[7*N]/NQ
SND=RCV=< 2.74% for N = 512 , NQ = 262,144
951006
KHW

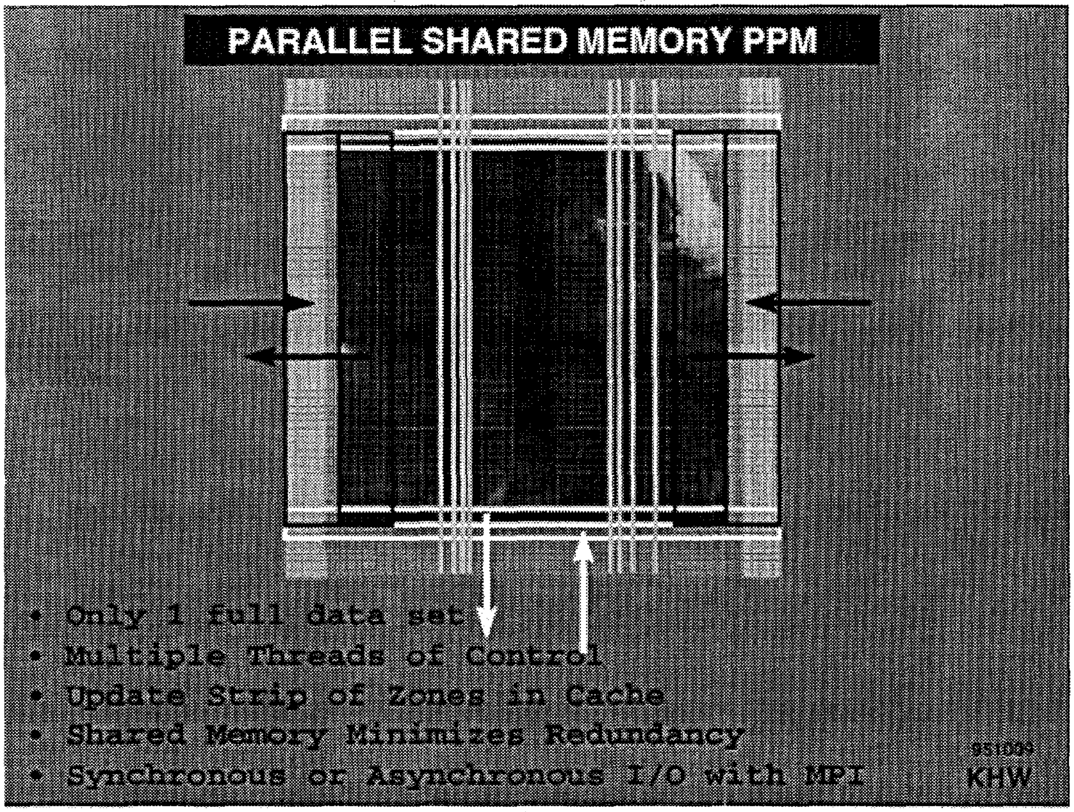

Visualization Done The Same Way

- **Up to 8 Distributed Workstations Sustain 1 Gigabit/s Communication Bandwidth and Display a Single High-Resolution Digital Movie at a Rate of up to 60 Frames per second.**

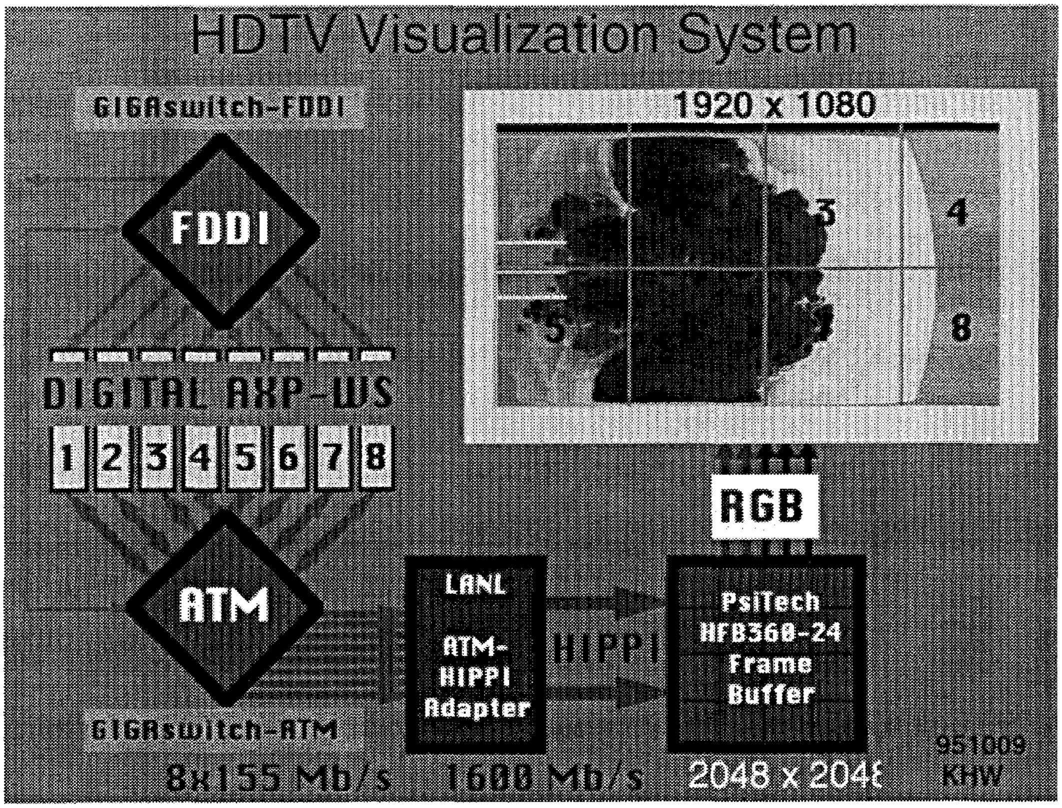
HDTV Visualization System
GIGAswitch-FDDI
FDDI
DIGITAL AXP-WS
1 2 3 4 5 6 7 8
ATM
GIGAswitch-ATM
1920 x 1080
3
4
5
8
RGB
LANL
ATM-
HIPPI
Adapter
HIPPI
PsiTech
HFB360-24
Frame
Buffer
8x155 Mb/s
1600 Mb/s
951009
KHW

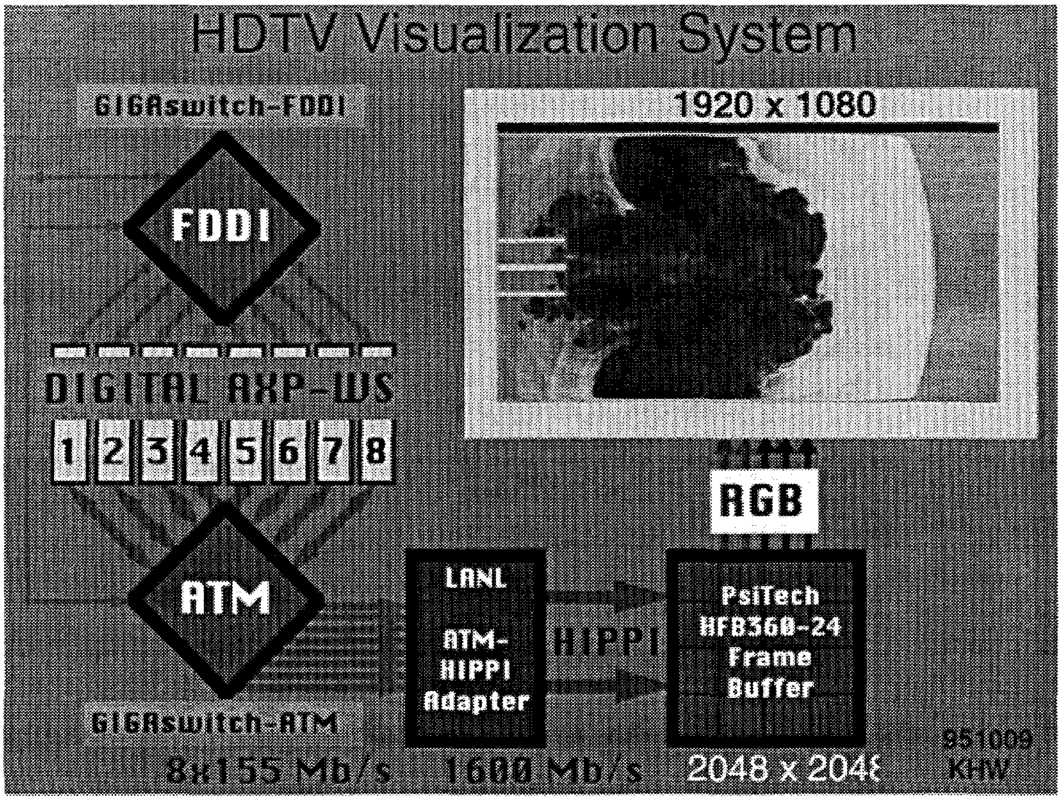
HDTV Visualization System
GIGAswitch-FDDI
FDDI
DIGITAL AXP-WS
1 2 3 4 5 6 7 8
ATM
GIGAswitch-ATM
1920 x 1080
RGB
LANL
ATM-
HIPPI
Adapter
HIPPI
PsiTech
HFB360-24
Frame
Buffer
8x155 Mb/s
1600 Mb/s
951009
KHW

Does the federal government need to support the development and engineering of ASCI scaling and integrating technologies?

Yes!

If yes, how much money is needed?

Depends!

Level	Effective Latency (CPU cycles)	Bandwidth (Random read/write)	Size
On-chip cache*, L1	2-3 •	16-32 B/cycle •	[illegible] B/flop • ▲
Off-chip cache**, L2 (SRAM)	5-6 •	16 B/cycle •	[illegible] B/flop • ▲
Local main memory (DRAM)	30-80 (15-30) ▼	2-8 B/flop pk (2-8 B/flop sustained) ▼	1 B/flop • ▲
"nearby nodes"	300-500 (30-50) ▼	1-8 B/flop (8 B/flop) ▼	1 B/flop •
"far away nodes"	1000 (100-200) ▼	1 B/flop (1 B/flop) ▼	[illegible] B/flop •
I/O (memory disk)	10 ms •	[illegible] B/flop •	10-100 B/flop •
Archive (disk-tape)	Seconds •	[illegible] B/flop (0.001-0.01 B/flop) ▼	[illegible] B/flop [illegible] B/flop ▼
User access	1/10 s (1/60 s)	[illegible]	100 users •
Multiple sites	1/10 s •	•	•

Compute engine

Interconnect

* Equivalent integer and floating-point data calculation rates are required.

** Cacheless systems with equivalent performance are fully acceptable.

Primary investment priority

Secondary investment priority

1996-1998 Situation

(1998-2000 Requirements)

Industry Trend

▲ Industry gets better at meeting requirements

▼ Industry gets worse at meeting requirements

• Industry continues to meet requirements

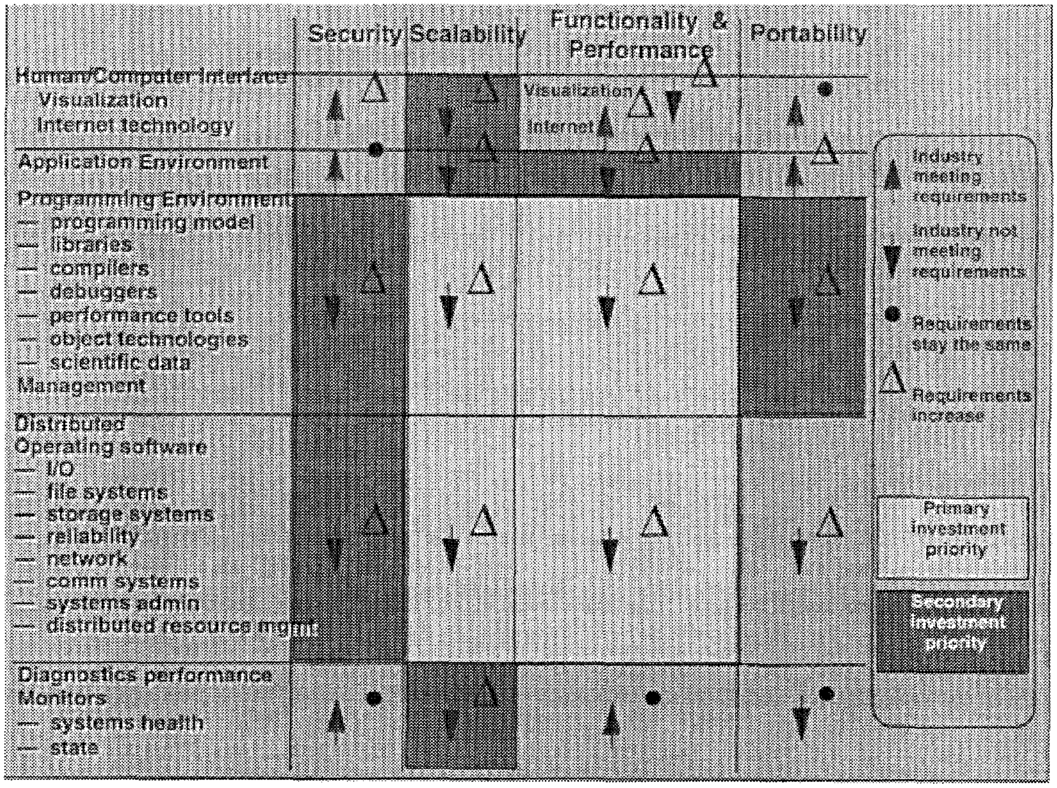

ASCI Coordinated Projects Collaboration

The purpose of this collaboration is to develop and engineer a system area interconnect which will meet the ASCI program's (and other high-end users) needs for computing. The interconnect will have to provide a low-latency and high-performance physical fabric which can be efficiently utilized to interconnect SMP-computers, visualization engines, secondary and tertiary storage devices, and external networks. It is expected that efficient use of the interconnect will require hardware support for global addressability and for cache management. It is not a requirement that the interconnect directly implement a global, cache-coherent, shared memory.

Interoperability:

In order to meet the ASCI program's need for complete, balanced systems it will be critical for the interconnect fabric to provide an open interface to a wide variety of devices and microprocessors. It is desirable that it be possible to directly connect the fabric to both the I/O channels and the memory bus interfaces of most common CMOS microprocessors, such as Intel's Merced, IBM's PowerPC, DEC's Alpha, SGI's Mips, SUN's Sparc, Tera's MTA. Furthermore the fabric should connect to disk and tape controllers, network gateways, and network attached peripherals.

Scalability:

The interconnect will need to support thousands of network interface ports. It must have the link capacity and routing logic to minimize contention and congestion due to a mix of large and small messages being inserted into the network from thousands of nodes at once. A modular design is a must which facilitates the tailoring of the network to traffic requirements and which allows scalable growth. Since the topology of the overall system will probably change over time it is important that the switch fabric support different topologies, such as interconnections of 8way, 16way, 32way switches. It is expected that a modular design will also facilitate the physical distribution of switch resources and the economical production of components.

Performance:

Individual network links should approximate in performance the memory bandwidths of microprocessors of the 1999 - 2001 time frame.

This implies a top performance of between 2 - 10 GBytes/s per link over distances of up to 100 yards. Desired latency through the switch should be of order 100s of nanoseconds, certainly less than 1 microsecond.

Low-level Protocol:

A standardized virtual memory interface, supporting both global push and pull, has to be developed and implemented on the various microprocessor platforms. Virtual memory interface libraries and Hippi6400's scheduled transfer protocol may give an indication of the areas of work where consensus between the collaboration members is needed. It will be critical to support routing, packetization, virtual circuits, health monitoring, performance monitoring, and debugging with a judicious balance of network protocol support (hardware and firmware) and system software support (drivers, OSs, and libraries).

Software:

An efficient system area network requires close collaboration between network resources and system software. It is expected that development will be required in at least the following areas: switch operation, low-level drivers, data movement libraries, distributed resource management, performance monitoring, and debugging.

In the area of switch operation the system software will have to be able to create routes (e.g. by discovering/employing routing tables or by interacting with network self-routing mechanisms), recognize and adjust for failing components and/or replacement components, and perhaps prioritize traffic.

The system software must as always form the bridge between hardware and users. For a system area network this means both providing low-level drivers within all target processor-nodes and peripheral devices and providing standard libraries for message passing and data movement.

The system envisioned will have thousands of network ports and a multiplicity of device types. It will be essential that system software support resource management in a distributed, robust manner. Resource management should include the ability to monitor the health of portions of the network, measure and report the performance of operations throughout the network, and provide debugging information to system designers, system administrators, and application designers/developers.

What are the roles of proprietary and non-proprietary technologies?

Commercial Mass Market: proprietary

Commercial Mass Market x 1000: non-proprietary community effort for development and engineering of base technologies

Licensing:

Development of this fabric most likely necessitates the development of ASICs (Application-Specific Integrated Circuits) for the switch itself, for the symmetric optical interconnect links, and for the adaptation to the specifics of the various memory bus architectures. These base technologies, developed on behalf of and paid for by ASCI, will be licensed or sold as chips to anybody without restrictions in intended use. In this way a larger market for the ASIC chips is created from the start which also supports a competitive business model for derivative products such as the switch itself or memory and I/O interfaces.

ASCI Operating systems of choice: UNIX, Windows NT, others?

- **UNIX:** for now.
- **Windows NT:** many issues like security, performance, scalability, 32 bit, MS-DOS remnants, entirely mass market focused.
- **Java:** virtual machine, threads, interoperable, garbage collection.

What are the business models for a self-sustaining high-end computing industry supporting ASCI capability computing systems?

How would this high-end industry relate to the extremely successful mass market computer industry?

Should ASCI forge relationships with vertically or horizontally integrated virtual companies?

intel

Applications of High Speed Computing in Manufacturing: Godzilla Meets King Kong

Eugene S. Meieran
Intel Fellow
Intel Corp.
Chandler, AZ

intel

Build a Better Mousetrap and...!

Traditional Belief

- The world will beat a path to your door!

Reality

- We will populate the world with smarter and more aggressive mice
- We will have more mouse trap competitors
- The mice will become resistant to mouse-catching procedures

intel®

The Manufacturer's Dilemma: A New Environment

- More expensive products
- More complex processes
- More customization
- A global economy
- More stuff to dispose of
- More information management
- More customers
- More talented workforce
- More automation
- More difficult manufacturing
- More new technology
- With shorter life cycles
- With less tolerance for defects
- But at no higher cost
- With more cultural barriers
- With more responsibility for it
- Generating information overload
- More competitors
- Less stable workforce
- Leading to lots of unskilled workers
- But at a lower cost
- Haven't mastered the old yet!

Gene Meieran Page 3 24 April, 1997

intel®

The Manufacturing Environment: Unstoppable Forces

- Ever more complex products and processes
- More physical automation and information automation
- Continuously improving product, process quality
- Globalization of facilities, suppliers, customers
- Increasingly sophisticated competition
- Increasing environmental concerns and regulations
- Better and more expensive facilities
- Faster development time and time-to-market
- Increasing amount of data and information generation
- More electronic commerce; Inter-and Intranet

Gene Meieran Page 4 24 April, 1997

intel®

Common Issue

All these forces strongly depend on applying more and more sophisticated software and hardware computational and communication tools to improve manufacturing performance.

Godzilla emerges from the depths.

Gene Meieran | Page 5 | 24 April, 1997

intel®

The Manufacturing Environment: Immovable Objects

- Reluctance to monkey with "family jewels" (source of $)
- Huge obsolete software, hardware, equipment baggage
- "Tower of Babel" heterogeneity of SW, HW platforms
- Fear of further productivity increases (loss of jobs)
- Smart but ignorant work force; fear of new skill needs
- Antiquated management styles; rigid organizational boundaries
- Poor ROI metrics
- Operational inertia
- Hard to choose between too many good alternatives
- Too many bad results; negative press and skepticism

Gene Meieran | Page 6 | 24 April, 1997

intel®

Common Issue

All these generate strong resistance to the introduction of new technology into the manufacturing arena.

King Kong emerges from the forest

Gene Meieran Page 7 24 April, 1997

intel®

When Godzilla Meets King Kong

- Long, tedious decision-making process; decision paralysis
- Pocket vetoes and stalemates
- Frustration; loss of motivation, indeed sabotage
- Poorly justified decisions to NOT EVER AGAIN engage in new technology ventures!
- Loss of competitive advantage

Gene Meieran Page 8 24 April, 1997

intel®

Moral

Creating and introducing more advanced technology does NOT fix the manufacturer's dilemma; plenty already exists.

Indeed, advocating introduction of more new technology may be part of the problem!

The problem is to bring new technology into manufacturing facilities when this is regarded as an avoidable and unnecessary threat.

Gene Meieran | Page 9 | 24 April, 1997

---

intel®

Applications for high speed computing

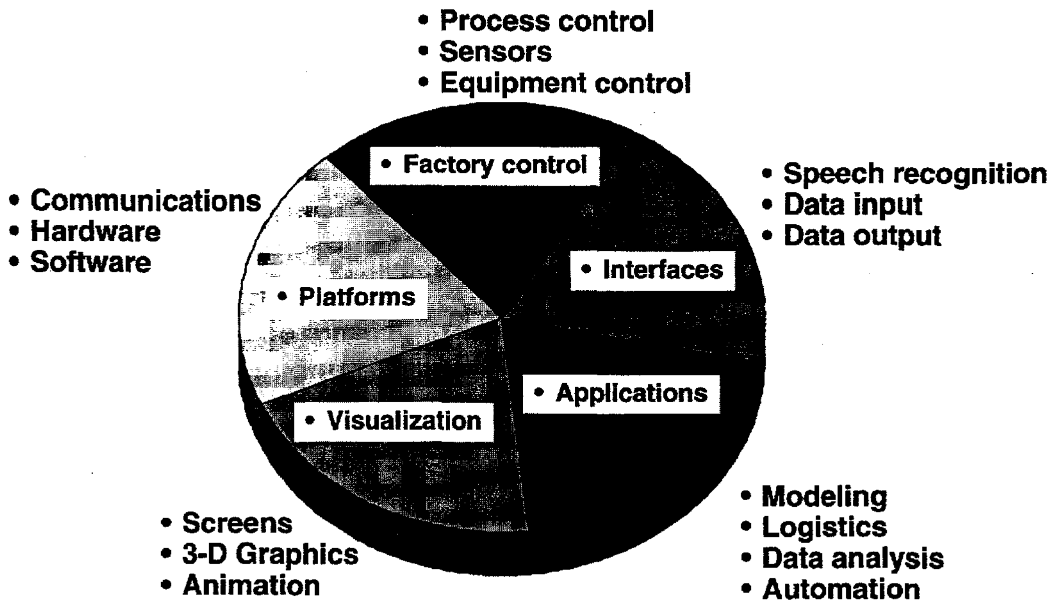

Gene Meieran | Page 10 | 24 April, 1997

intel®

High Speed Computing Needs

- Operational modeling and simulation
- Factory, equipment and process control
- Data analyses and decision-assist systems
- Enterprise integration
- User interface improvements

intel®

Operational Modeling and Simulation

- NEEDS:
  - Cheaper decisions (experiments are costly)
  - Faster decisions (experiments are time consuming)
  - Better decisions (wrong decisions are EXPENSIVE!)
- ISSUES:
  - Experts do not have time to create good models
  - Computer whizzes don't understand manufacturing
  - Manufacturing is NOT analytical; cannot simply apply equations
  - Integration of models at different hierarchical levels are very difficult
  - Model validation and testing
  - Discrete event simulation can ALWAYS be faster!

intel®

Factory, Process and Equipment Control

- NEEDS:
  - Need to deliver higher quality products
  - Inter-relationship of very complex, widely separated process steps
  - Faster detection of trends to prevent problems
  - Improved factory environments (waste, resource usage)
- ISSUES:
  - Equipment manufacturers are often small businesses without sufficient expertise to handle topic
  - APC seems antithetical to SPC; "chasing one's tail"
  - Stability of systems

Gene Meieran Page 13 24 April, 1997

intel®

Decision-Assist Systems

- NEEDS:
  - Faster decisions
    - Avoid reinvention of the wheel
    - Based on past knowledge and experience
  - Better decisions
    - Traceable to source
    - Easy to use
  - Remembered decisions
    - Based on historical record
    - Easy to find and apply
- ISSUES:
  - This is the Holy Grail, so far unsuccessfully pursued
  - These need to be integrated into the factory flow
  - There are too few applicable metrics

Gene Meieran Page 14 24 April, 1997

intel®

Enterprise Integration

- NEEDS:
  - Faster product and process development
    - 24 hour collaborative development
    - 100% electronic information transfer
  - Faster time to customer
    - Customized products delivered anywhere
    - Electronic data-to-product
  - Better utilization of enterprise resources
    - Minimal number of qualified vendors
    - Dynamic network of suppliers, manufacturers and customers
  - Lower overall cost
    - Customized products at mass produced prices
    - Virtual factories, companies, enterprises,
- ISSUES:
  - Multi-nation, multi-culture, multi-language
  - Long term stability
  - Change control

Gene Meieran Page 15 24 April, 1997

intel®

User Interface Development

- NEEDS:
  - Fewer errors on the factory floor
  - Lower resistance to using computer tools
  - Ability to handle more data and information
  - Ability to do wireless communication
  - Ability to integrate different applications
  - Space saving
- ISSUES:
  - Too few computer techies KNOW what the factory really looks like or needs
  - Ability to handle non-keyboard, perhaps non-verbal inputs

Gene Meieran Page 16 24 April, 1997

intel®

Common Needs

• Better hardware platforms	Faster computing Big, reliable servers Parallel processing
• Improved software platforms	Robust code Fast problem resolution Parallel processing
• Expanding application suites	Integration
• Improved user interfaces	Easier, more varied
• Better visualization	Larger, flatter screens Higher resolution screens Personalized screens
• Communications	High speed, high bandwidth Global and wireless Secure

Gene Meieran Page 17 24 April, 1997

intel®

What Must the Computer Technocrat Do?

- Understand the customer from the CUSTOMER'S perspective!
- Identify and work with key stakeholders
- Benchmark within and outside the industry
- Generate mutually agreeable expectations
- Start small; build incrementally
- Continuously validate benefits and issues
- Remember that most of the costs will be in maintaining the system for a LONG time after the implementation!

Gene Meieran Page 18 24 April, 1997

intel®

What Must the Computer Technocrat ~~Do~~?

- Assume that the best technology automatically SHOULD be adopted (even if you're right!)
- Assume that "proven" significant benefits WILL outweigh even moderate risks in the eyes of the stakeholders
- Assume that the main issue is technology introduction
- Assume the main issue is technology in the first place

intel®

The Bottom Line

- <u>The issue is NOT technology:</u>
  Sufficient technology exists to do almost everything a manufacturer wants for the next decade
- <u>The issue is wisdom;</u>
  Understand and be responsive to the wishes, desires, fears and needs of the manufacturing customer (who could care less about technology).
- <u>Prognosis:</u>
  High speed computing WILL benefit manufacturing IF the technologists (Godzilla) assumes the appearance of the manufacturing executive (King Kong) and thinks, acts and is motivated accordingly
- <u>Otherwise..... stalemate!</u>

List of Attendees

Thomas Adams
Los Alamos National Laboratory
NWT-PO, MS F629
Los Alamos, NM 87545

Frank Akers
Lockheed Martin Energy Systems
PO Box 2009
Oak Ridge, TN 37831

Steve Ashby
Lawrence Livermore National Lab
L-561
Livermore, CA 94551

Greg Astfalk
Hewlett-Packard
3000 Waterview Parkway
Richardson, TX 75083-3851

Robert G. Babb II
University of Denver
Dept. of Math & Comp. Science
2360 S. Gaylord Street
Denver, CO 80208

Kimberly Barnes
Oak Ridge National Laboratory
Bldg 6010, MS 6414
Oak Ridge, TN 37831-6414

Doc' Bedard
NSA
505 Hermleigh Road
Silver Springs, MD 20902

Tor Bloch
Danish Comp. Ctr. for Res. in Ed.
UNI-C, DTU, Bldg 304
DK-2800, LYNGBY, DENMARK

Dorothy Bollman
Department of Mathematics
Univ. of Puerto Rico, Managuez
Managuez, Puerto Rico 00680

Bob Borchers
National Science Foundation
4201 Wilson Blvd., Room 1122
Arlington, VA 22230

Joe Brandt
Lawrence Livermore National Lab
P. O. Box 808, L-66
Livermore, CA 94551

Forrest Brown
KAPL
PO Box 9049
Niskayuna, NY 12309

John Brown
Silicon Graphics
2011 N. Shoreline Blvd., MS 580
Mountain View, CA 94043-1389

Susan Brown
Univ. of MI, Ctr. for Human Mtr. Res.
Division of Kinesiology
401 Washtenaw Ave.
Ann Arbor, MI 48109-2214

T. Daniel Butler
Los Alamos National Laboratory
T-3, MS B216
Los Alamos, NM 87545

L. Jean Camp
Sandia National Laboratories
MS 9053
Livermore, CA 94550

Bill Carlson
Inst. for Defense Analysis
17100 Science Drive
Bowie, MD 20715-4300

Ralph Carlson
Lawrence Livermore National Lab
P.O. Box 808; L-711
Livermore, CA 94551

Bill Carothers
Hewlett Packard
5725 W. Las Positas Blvd.
Pleasanton, CA 94588

John Cerutti
Los Alamos National Laboratory
MS F663
Los Alamos, NM 87545

Sam Coleman
Lawrence Livermore National Lab
P. O. Box 808, L-60
Livermore, CA 94551

David Cooke
Dean, Faculty of Health and Sciences
Elborn College
University of Western Ontario
London, Ontario, Canada N6G 1H1

David M. Cooper
Lawrence Livermore National Lab
P. O. Box 808; L-66
Livermore, CA 94550

George Cotter
National Security Agency
9800 Savage Road
Fort Meade, MD 20755-6217

Steve Crago
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292-6714

Art Dana
Los Alaomos National Laboratory
X-NH, MS F664
Los Alamos, NM 87545

Tom DeBoni
Lawrence Livermore National Lab
P. O. Box 808, L-416
Livermore, CA 94551

David J. DiLeonardo
Bettis Atomic Power Laboratory
P. O. Box 79
West Mifflin, PA 15122-0079

Jerry Dobson
Oak Ridge National Laboratory
PO Box 2008, MS 6237, 4500-N
Oak Ridge, TN 37831-6237

Jack Dongarra
University of Tennessee
Computer Science Department
104 Ayres Hall
Knoxville, TN 37996-1301

List of Attendees continued

Pete Eltgroth
Lawrence Livermore National Lab
P.O. Box 808, L-045
Livermore, CA 94551

Beth Erbert
Lawrence Livermore National Lab
L-66
Livermore, CA 94551

John Feo
Tera Computer Company
2815 Eastlake Ave. East
Seattle, WA 98155

Paul Fischbeck
Social and Decision Sciences
Carnegie-Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213

Hugo de Garis
Brain Bldr. Grp, Evolutionary Syst. Dept.
ATR Human Info. Proc. Labs
2-2 HIKARI DAI, Seika-cho Soraku-gun
Kansai Science City, Kyoto 619-02 Japan

Don Geiser
Hewlwtt Packard
24 Inverness Place East
Englewood, CO 80112

Alan Geller
SGI/CRI
3130 Crow Canyon Place; Suite 400
San Ramon, CA 94583

Morven Gentleman
Institute for Information Technology
National Research Council
Montreal Road, Bldg M-50
Ottawa, Ontario, Canada K1A OR6

Mike Gittings
Los Alamos National Laboratory
X-TA, MS F663
Los Alamos, NM 87545

Dick Glassbrook
RCI, Ltd.
17 Turnham Court
Gaithersburg, MD 20878

Maya Gokhale
David Sarnoff Research Center
201 Washington Road, CN 5300
Princeton, NJ 08543-5300

Debra Goldfarb
IDC
5 Speen Street
Framingham, MA 01701

CAPT William Gravell
US Navy
Chief, Information Assurance Div.
The Joint Staff (J6K)
Washington, DC 20318-6000

John Gustafson
Ames Laboratory, Iowa State Univ.
236 Wilhelm
Ames, IA 50011

Art Hale
Sandia National Laboratory
Mail Stop 1109
PO Box 5800
Albuquerque, NM 87185

Susan Hagness
Northwestern Univ.
EECS Department
2145 Sheridan Road
Evanston, IL 60208-3118

Ann Hayes
Los Alamos National Laboratory
ACI, MS B287
Los Alamos, NM 87545

Andrew Heller
Heller Associates
4501 Hennig Drive
Austin, TX 78738

Don Heller
Ames Laboratory
Iowa State University
327 Wilhelm Hall
Ames, Iowa 50011

Michael J. Henesey
IBM
P.O. Box 100
Somers, New York

Rod Hibner
Hewlett Packard

Kathy Hirons
Los Alamos National Laboratory
DDCIC, MS B260
Los Alamos, NM 87545

Tad Hogg
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

Karin Hollerbach
Lawrence Livermore National Lab
ISCR, L-416
Livermore, CA 94550

Don Hossink
IBM

Alan Huang
Stanford Univ.
682 Sixteenth Ave.
Menlo Park, CA 94025

George Ivy
Logicon RDA
P. O. Box 92500
Los Angeles, CA 90009

Ajita John
Dept. of Computer Sciences
TAY 2.124
Univ. of Texas at Austin
Austin, TX 78712

Harry Jordan
University of Colorado at Boulder
Dept. of Electrical & Comp. Engr.
Campus Box 425
Boulder, CO 80309-0425

David Kahaner
ATIP
6-15-21 Roppongi, Harks Bldg 1F
Minato-ku, Tokyo, Japan

List of Attendees continued

Mal Kalos
Cornell Theory Center
Cornell University
524 Rhodes Hall
Ithaca, NY 14853-3801

Dan Katz
NASA, JPL
4800 Oak Grove Drive
MS 168-522
Pasadena, CA 91109

Tom Keyser
University of Southern Colorado
2200 Bonforte Blvd.
Pueblo, CO 81001-4901

Connie Kirchner
Lawrence Livermore Natinal Lab
L-66
Livermore, CA 94550

Tom Kitchens
U. S. Dept. of Energy
Math., Information & Comp. Sci. Div.
19901 Germantown Road
Germantown, MD 20874-1290

Ken Kliewer
Oak Ridge National Laboratory
P.O. Box 2008
Oak Ridge, TN 37831-6203

Robert Kluttz
Naval Security Group Command
9800 Savage Road, Suite 6685
Operations Bldg #3
Fort Meade, MD 20755

Kenneth Koch
Los Alamos National Laboratory
CIC-19 MS B256
Los Alamos, NM 87545

Martin Lades
Lawrence Livermore National Lab
L-416, ISCR,
Livermore, CA 94551

Dale Land
Los Alamos National Laboratory
CIC-2, MS D445
Los Alamos, NM 87545

Alex Larzelere
U.S. Dept. of Energy DP-14
1000 Independence Ave., SW
Washington, DC 20585

John Lekashman
NASA
PO Box 390045
Mt. View, CA 94039

Robert Leland
Sandia National Laboratory
Albuquerque, NM 87185-0441

Bill Lennon
Lawrence Livermore National Lab
L-541
Livermore, CA 94551

John Levesque
Applied Parallel Research, Inc.
1723 Professional Drive
Sacramento, CA 95825

Michael Levine
Pittsburgh Supercomputing Center
4400 Fifth Ave.
Pittsburgh, PA 15213

Debbie Martinez
Los Alamos National Laboratory
CIC-12, MS B295
Los Alamos, NM 87545

Ron Matlock
Fujitsu America
2500 City West Blvd., Suite 300
Houston, TX 77042

Mike McConnell
Booz, Allen & Hamilton
8283 Greensboro Drive
MacLean, VA 22102-3838

Bill McCurdy
Lawrence Berkeley National Lab
MS 50B - 4230
Berkeley, CA 94720

Charlie McMillan
Lawrence Livermore National Lab
L-170
Livermore, CA 94551

Gene Meieran
Intel Corporation CH2-23
5000 West Chandler Blvd.
Chandler, AZ 85226

Bob Meisner
U.S. Department of Energy, DP-52
1000 Independence Ave. SW
Washington, DC 20585

Paul Messina
California Institute of Technology
1200 E. California Blvd. MC 158-79
Pasadena, CA 91125

George Michael
Lawrence Livermore National Lab
P.O. Box 808; L-306
Livermore, CA 94550

George Michaels
George Mason Univ.
Inst. for Comp. Sciences, MS 5C3
Fairfax, VA 22030-4444

Christie Michelsen
Los Alamos National Laboratory
CIC-8, MS B272
Los Alamos, NM 87545

Ken Miura
Fujitsu America, Inc., MS 1-5
3055 Orchard Drive
San Jose, CA 95134

Gloria Montoya
Los Alamos National Laboratory
DDCIC, MS B260
Los Alamos, NM 87545

Gary Montry
Southwest Software
11812 Persimmon NE
Albuquerque, NM 87111

List of Attendees continued

Jim Morel
Los Alamos National Laboratory
CIC-19, MS B265
Los Alamos, NM 87545

Robin Morel
Los Alamos National Laboratory
P.O. Box 1663, B265
Los Alamos, NM 87545

John Morrison
Los Alamos National Laboratory
DDCIC, MS B260
Los Alamos, NM 87545

Fred Mortensen
Los Alamos National Laboratory
X-TA MS B220
Los Alamos, NM 87545

Deirdre Mulligan
Center for Democracy & Technology
1634 I Street NW
Washington, DC 20006

Nathan Myhrvold
Microsoft Corporation
1 Microsoft Way
Redmond, WA 98052-6399

Steve Nesbitt
NASA, Office of the Inspector General
300 E Street SW, Code W
Washington, DC 20546

Steven Oberlin
SGI/Cray
P.O. BOx 4000
Chippewa Falls, WI 54729

Rod Oldehoeft
U. S. Dept. of Energy, ER-31
19901 Germantown Road
Germantown, MD 20874

C. Edward Oliver
Oak Ridge National Laboratory
P. O. Box 2008,
Bldg. 4500N, Room K-244
Oak Ridge, TN 37831-6259

Mel Paisley
Hewlett Packard
2101 Gaither Road
Rockville, MD 20850

Cherri M. Pancake
Oregon State University
Dept. of Computer Science
303 Dearborn Hall
Corvallis, OR 97331

Steve Poole
Performance Group
2830 Kenross
Houston, TX 77043

John Riganati
David Sarnoff Research Center
201 Washington Road, CN 5300
Princeton, NJ 08543-5300

Jeffrey Ritter
ECLIPS
Ohio Supercomputer Center
1224 Kinnear Road
Columbus, OH 43212-1163

Mark Roschke
Los Alamos National Laboratory
CIC-11, MS B269
Los Alamos, NM 87545

Paul Rupert
Lawrence Livermore National Lab
7900 Tesla Road
Livermore, CA 94550

Subhash Saini
NASA Ames Research Center
MS T27A-1
Moffett Field, CA 94035-1000

Jeff Saltzman
Los Alamos National Laboratory
CIC-19, MS B256
Los Alamos, NM 87545

Mark Seager
Lawrence Livermore National Lab
P. O. Box 808, L-060
Livermore, CA 94550

Ron Selva
Silicon Graphics, Inc.
600 Waterfront Drive
Pittsburgh, PA 15222

Bill Sembrat
SRC Computers, Inc.
395 Rider Ridge Road
Santa Cruz, CA 95065

Margaret Simmons
Nat. Coordination Office for HPCC
4201 Wilson Blvd., Suite 665
Washington, DC 22230

Horst Simon
Lawrence Berkeley National Lab
1 Cyclotron Road, MS 50B-4230
Berkeley, CA 94720

Richard Sincovec
Oak Ridge National Laboratory
Bldg. 6025, MS-6359
Oak Ridge, TN 37831-6359

Charles Slocomb
Los Alamos National Laboratory
DDCIC, MS B260
Los Alamos, NM 87545

Burton Smith
Tera Computer Co.
Tera Computer Company
2815 Eastlake Ave. East
Seattle, WA 98102-1318

Paul Smith
U. S. Dept. of Energy, DP-50
1000 Independence Ave. SW
Washington, DC 20585

Bill Spangenberg
Los Alamos National Laboratory
CIC-12 MS B295
Los Alamos, NM 87545

Apoorv Srivastava
University of Southern California
4676 Admiralty Way, #943
Marina Del Rey, CA 90292-6695

List of Attendees continued

Will Stackhouse
Xybernaut
1123 Stanford Ave.
Redondo Beach, CA 90278-4039

Tom Sterling
NASA JPL/Caltech
4800 Oak Grove Drive
Mail Stop 168-522
Pasadena, CA 91109

Gene Stokes
Defense Spec. Weapons Agency
6801 Telegraph Road
Alexandria, VA 22310-3398

Jack Strauss
Spectral Systems, Inc.
5200 Springfield Pike, Suite 320
Dayton, OH 45431

Steve Stringer
SGI/CRI
Los Alamos, NM 87544

Thomas M. Sutton
Knolls Atomic Power Lab
E6-204, P. O. Box 1072
Schenectady, NY 12301-1072

Allen Taflove
Northwestern Univ.
EECS Department
2145 Sheridan Road
Evanston, IL 60208

Tom Talleur
NASA, Office of the Inspector General
300 E Street SW, Code W
Washington, DC 20546

Glenn Tenney
Fantasia Systems, Inc.
2111 Ensenada Way
San Mateo, CA 94403

John Toole
National Coordination Office for CIC
4201 Wilson Blvd., Suite 665
Arlington, VA 22230

Steve Turpin
U.S. Dept. of Energy
DP-50, Forres 4A-045
1000 Independence Ave. SW
Washington, DC 20585

David Urie
Lockheed (Ret.)
3634 Woodcliff Road
Sherman Oaks, CA 91403

George VanDegrift
SGI/CRI
6565 Americas Parkway NE, ST. 660
Albuquerque, NM 87110

Manuel Vigil
Los Alamos National Laboratory
P.O. Box 1663, MS B294
Los Alamos, NM 87545

Mike Vildibill
San Diego Supercomputing Ctr.
P.O. Box 85608
San Diego, CA 92186-5608

Cheryl Wampler
Los Alamos National Laboratory
MS B294
Los Alamos, NM 87545

Dick Watson
Lawrence Livermore National Lab
P. O. Box 808, L-66
Livermore, CA 94551

David Wehrly
ForeFronts Comp. Tech., Inc
PO Box 766, 20101 US Hwy 14
Sundance, WY 82729-0766

Andrew B. White
Los Alamos National Laboratory
ACL, MS B287
Los Alamos, NM 87545

Frank Wilson
UCSF
4096 Sugar Maple Drive
Danville, CA 94506-4611

Karl-Heinz Winkler
Los Alamos National Laboratory
DDCIC, MS B260
Los Alamos, NM 87545

Paul Woodward
Univ. of Minnesota
Astronomy Department
116 Church Street S.E.
Minneapolis, MN 55455