The Role of NSF in Computer Architecture Research

A. Yavuz Oruç

Professor
University of Maryland at College Park

Former Director
Computer Systems Architecture Program
2000-2002
National Science Foundation

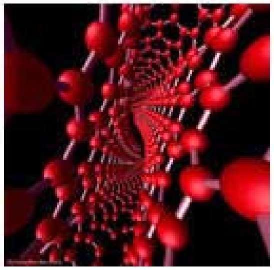

Nanotube Flattened View

THE OUTLINE

OVERVIEW -NSF

OVERVIEW -CCR

OVERVIEW -CSA

PARTING THOUGHTS

BRIEF

THREADED

MULTI-THREADED

NEXT SLIDE PREVIOUS SLIDE

THE OUTLINE

OVERVIEW -NSF ->

ORGANIZATION ->

CISE -> CCR -> CSA ->

OVERVIEW OF CSA ->

PROGRAM HIGHLIGHTS ->

PARTING THOUGHTS

BRIEF

THREADED

MULTI-THREADED

NEXT SLIDE PREVIOUS SLIDE

THE OUTLINE

OVERVIEW-NSF ->
OVERVIEW OF CISE ->
FY2003 BUDGET ->
OVERVIEW OF CCR ->
CISE BUDGET ->
OVERVIEW OF CSA ->
CCR BUDGET ->
PROGRAM HIGHLIGHTS ->
CSA BUDGET ->
PARTING THOUGHTS

BRIEF

THREADED

MULTI-THREADED

NEXT SLIDE PREVIOUS SLIDE

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

"It is our [NSF's] job to keep all fields of science and engineering focused on the furthest frontier, to recognize and nurture emerging fields, to support the work of those with the most insightful reach, and to prepare coming generations of scientific talent."

NSF Director Rita Colwell

"The National Science Foundation at 50"

New York Academy of Sciences, September 30, 1999

MISSION-2000

MISSION-1950

MISSION-ITEMIZED

HISTORY-1

HISTORY-2

HISTORY-3

HISTORY-4

ORGANIZATION

CISE

CCR

NEXT SLIDE

PREVIOUS SLIDE

NATIONAL SCIENCE FOUNDATION
–OVERVIEW –

NSF's mission is set out in the preamble to the National Science Foundation Act of 1950 (Public Law 810507):

To promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense; and for other purposes.

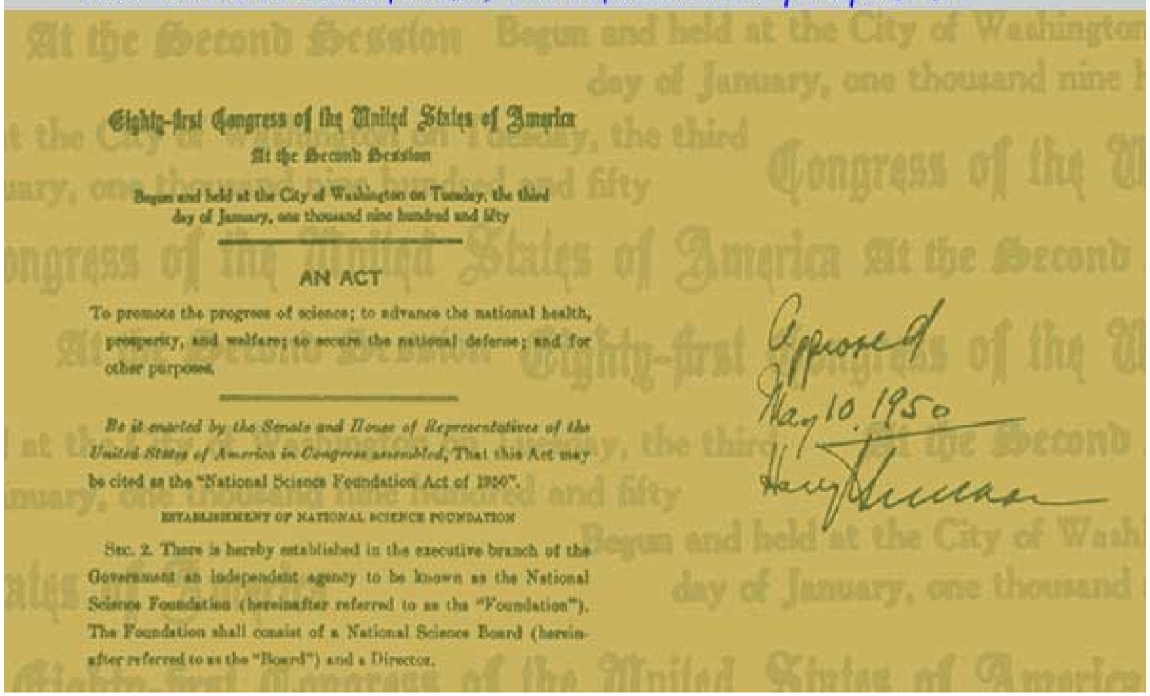

MISSION-2000

MISSION-1950

MISSION-ITEMIZED

HISTORY-1

HISTORY-2

HISTORY-3

HISTORY-4

ORGANIZATION

CISE

CCR

NEXT SLIDE

PREVIOUS SLIDE

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

The Act authorizes and directs NSF to initiate and support:

* basic scientific research and research fundamental to the engineering process,

* programs to strengthen scientific and engineering research potential,

* science and engineering education programs at all levels and in all the various fields of science and engineering,

* an information base for science and engineering appropriate for development of national and international policy.

Neal Lane- President's Science Advisor

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

A Political Director?

It was an imposing group that gathered at the White House on December 12, 1950. Though President Truman had not yet arrived, Steelman opened the meeting. Board members elected Conant as chairman and Charles Dollard of Carnegie Corporation of New York as vice chairman. The NSF Act mandated an executive committee, of which Detlev W. Bronk was elected chairman. Bronk, a biologist, was president of Johns Hopkins University and of the National Academy of Sciences.

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

Members at this first meeting had heard rumors that Truman had offered the post of the Foundation's Director to someone they considered less than qualified for the job: Frank P. Graham, a lame-duck U.S. Senator and former history professor. According to later accounts, Truman showed up and asked what they had been talking about. Someone replied that they'd been wondering what qualifications Truman thought were appropriate for the Foundation's Director.

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

Truman answered, "There's only one criterion. He must get along with me." Thus continued the tension of how much the White House would control a Foundation explicitly endowed by Congress with its own independent governing Board.

By law, the President was required to seek the Board's advice before making a formal nomination. Board protests eventually caused Graham to withdraw from consideration. At the Board's fourth meeting in March 1951, a telegram from Truman announced he would nominate Alan T. Waterman, former Yale physicist and chief scientist of the ONR, as Director. Waterman had been on the Board's list of candidates for Director, and his nomination was greeted "with audible relief " by the members.

NATIONAL SCIENCE FOUNDATION
–OVERVIEW –

October 2, 1954

Dear Mr. Lednicer:

Your good letter of September 21 was very much appreciated.

I always knew that the Science Foundation would do a great amount of good for the country and for the world. It took a terrific fight and three years to get it through the Congress, and some smart fellows who thought they knew more than the President of the United States tried to fix it so it would not work.

It is a great pleasure to hear that it is working and I know it will grow into one of our greatest educational foundations.

Sincerely yours,
/s/ Harry S Truman

MISSION-2000

MISSION-1950

MISSION-ITEMIZED

HISTORY-1

HISTORY-2

HISTORY-3

HISTORY-4

ORGANIZATION

CISE

CCR

NEXT SLIDE

PREVIOUS SLIDE

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

National Science Board---Office of Inspector General

NSF Director--------Office of EQOpportunity Prog

| ----------Office of Legislative and Pub. Aff.

| ----------Office of Integrative Activities

| ----------Office of Polar Programs

| ----------Office of The General Counsel

NSF Deputy Director

Biology Directorate

Computer and Information Science
and Engineering Directorate

Education and Human Resources Directorate

Engineering Directorate

Geosciences Directorate

Mathematical and Physical Science Directorate

Social,Behavorial and Economic Sciences Directorate

Office of Budget, Finance and Award Managment

Office of Information and Resource Management

MISSION-2000

MISSION-1950

MISSION-ITEMIZED

HISTORY-1

HISTORY-2

HISTORY-3

HISTORY-4

ORGANIZATION

CISE

CCR

NEXT SLIDE

PREVIOUS SLIDE

NATIONAL SCIENCE FOUNDATION
–OVERVIEW –

CISE Directorate- Peter Freeman-Director

COMPUTER-COMMUNICATIONS RESEARCH DIVISION-
- Kamal Abdali, Act. Div.Director

INFORMATION AND INTELLIGENT SYSTEMS DIVISION -
William S. Bainbridge, Deputy Div. Director

EXPERIMENTAL AND INTEGRATIVE ACTIVITIES DIVISION -
Gary W. Strong, Act. Div.Director

ADVANCED COMPUTATIONAL INFRASTRUCTURE
AND RESEARCH DIVISION, Act. Richard S. Hirsh, Director

ADVANCED NETWORKING INFRASTRUCTURE
AND RESEARCH DIVISION, Aubrey Bush-Director

INFORMATION TECHNOLOGY RESEARCH (ITR) -Taeb Znati

MISSION-2000
MISSION-1950
MISSION-ITEMIZED
HISTORY-1
HISTORY-2
HISTORY-3
HISTORY-4
ORGANIZATION
CISE
CCR
NEXT SLIDE
PREVIOUS SLIDE

NATIONAL SCIENCE FOUNDATION
-OVERVIEW -

COMPUTER COMMUNICATIONS RESEARCH

Kamal Abdali Acting Director

Frank Anger, Acting Deputy Division Director

- COMPUTER SYSTEMS ARCHITECTURE- Peter Varman
- DESIGN AUTOMATION FOR MICRO AND NANO SYSTEMS- Bob Grafton
- EMBEDDED & HYBRID SYSTEMS- Helen Gill
- COMMUNICATIONS- Julia Abrahams
- SIGNAL PROCESSING- John Cozzens
- TRUSTED COMPUTING- Carl Landwehr
- SOFTWARE ENGINEERING & LANGUAGES- Frank Anger
- OPERATING SYSTEMS &COMPILERS- Randy Chao
- THEORY OF COMPUTING- D. Z. Hu
- GGRAPHICS, SYMBOLIC, & GEOMETRIC
- COMPUTATION-John Staudhammer

FISCAL YEAR 2003 BUDGET

NSF Priority Area Investments
(Millions of Dollars)

Priority Area	FY 2001 Actual	FY 2002 Current Plan	FY 2003 Request	Change Amount	Change Percent
Biocomplexity in the Environment	54.88	58.10	79.20	21.10	36.3%
Information Technology Research	216.27	277.52	285.83	8.31	3.0%
Nanoscale Science and Engineering	149.68	198.71	221.25	22.54	11.3%
Learning for the 21st Century Workforce	143.33	144.82	184.69	39.87	27.5%
Mathematical Sciences	0.00	30.00	60.09	30.09	100.3%
Social, Behavioral and Economic Sciences	0.00	0.00	10.00	10.00	N/A
Total, Priority Areas	$564.16	$709.15	$841.06	$131.91	18.6%

Totals may not add due to rounding.

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY
IT-RESEARCH
NANOTECHNOLOGY
21ST CENTURY LEARNING
MATHEMATICAL SCIENCES
FUNDING LEVELS

FISCAL YEAR 2003 BUDGET

Proposed funding for the Biocomplexity in the Environment priority area is as follows:

(Millions of Dollars)

	FY 2002 Current Plan	FY 2003 Request	Change Amount	Change Percent
Biological Sciences	16.90	35.86	18.96	112.2%
Computer and Information Science and Engineering	6.10	7.36	1.26	20.7%
Engineering	3.69	6.00	2.31	62.6%
Geosciences	23.00	22.22	-0.78	-3.4%
Mathematical and Physical Sciences	5.35	4.70	-0.65	-12.1%
Social, Behavioral and Economic Sciences	1.65	1.65	0.00	0.0%
Office of Polar Programs	1.41	1.41	0.00	0.0%
Total, Biocomplexity in the Environment	$58.10	$79.20	$21.10	36.3%

Totals may not add due to rounding.

Long-term funding for the Biocomplexity in the Environment priority area is as follows:

(Millions of Dollars)

FY 2000 Actual	FY 2001 Actual	FY 2002 Current Plan	FY 2003 Request	FY 2004	FY 2005
50.00	54.88	58.10	79.20	87.76	92.24

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY
IT-RESEARCH
NANOTECHNOLOGY
21ST CENTURY LEARNING
MATHEMATICAL SCIENCES
FUNDING LEVELS

FISCAL YEAR 2003 BUDGET

Proposed funding for the Information Technology Research priority area is as follows:

(Millions of Dollars)

	FY 2002 Current Plan	FY 2003 Request	Change Amount	Change Percent
Biological Sciences	6.08	6.80	0.72	11.8%
Computer and Information Science and Engineering	173.51	190.67	17.16	9.9%
Engineering	10.23	11.17	0.94	9.2%
Geosciences	12.16	13.21	1.05	8.6%
Mathematical and Physical Sciences	33.06	35.52	2.46	7.4%
Social, Behavioral and Economic Sciences	4.26	4.65	0.39	9.2%
Office of Polar Programs	1.22	1.33	0.11	9.0%
Subtotal, Research and Related Activities	240.52	263.35	22.83	9.5%
Education and Human Resources	2.00	2.48	0.48	24.0%
Subtotal, R&RA and Education and Human Resources	242.52	265.83	23.31	9.6%
Major Research Equipment and Facilities Construction	35.00	20.00	-15.00	-42.9%
Total, Information Technology Research	$277.52	$285.83	$8.31	3.0%

Totals may not add due to rounding.

Long-term funding for the Information Technology Research priority area is as follows:

(Millions of Dollars)

FY 2000 Actual	FY 2001 Actual	FY 2002 Current Plan	FY 2003 Request	FY 2004
126.00	216.27	277.52	285.83	291.21

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY
IT-RESEARCH
NANOTECHNOLOGY
21ST CENTURY LEARNING
MATHEMATICAL SCIENCES
FUNDING LEVELS

FISCAL YEAR 2003 BUDGET

Proposed funding for the Nanoscale Science and Engineering priority area is as follows:

(Millions of Dollars)

	FY 2002 Current Plan	FY 2003 Request	Change Amount	Change Percent
Biological Sciences	2.33	2.98	0.65	27.9%
Computer and Information Science and Engineering	10.20	11.14	0.94	9.2%
Engineering	86.30	94.35	8.05	9.3%
Geosciences	6.80	7.53	0.73	10.7%
Mathematical and Physical Sciences	93.08	103.92	10.84	11.6%
Social, Behavorial and Economic Sciences	0.00	1.11	1.11	N/A
Subtotal, Research and Related Activities	198.71	221.03	22.32	11.2%
Education and Human Resources	0.00	0.22	0.22	N/A
Total, Nanoscale Science and Engineering	$198.71	$221.25	$22.54	11.3%

Totals may not add due to rounding.

Long-term funding for the Nanoscale Science and Engineering priority area is as follows:

(Millions of Dollars)

FY 2001 Actual	FY 2002 Current Plan	FY 2003 Request	FY 2004	FY 2005
149.68	198.71	221.25	251.25	266.25

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY
IT-RESEARCH
NANOTECHNOLOGY
21ST CENTURY LEARNING
MATHEMATICAL SCIENCES
FUNDING LEVELS

FISCAL YEAR 2003 BUDGET

Proposed funding for Learning for the 21st Century Workforce priority area is as follows:

(Millions of Dollars)

	FY 2002 Current Plan	FY 2003 Request	Change Amount	Change Percent
Biological Sciences	1.70	1.93	0.23	13.5%
Computer and Information Science and Engineering	1.15	1.20	0.05	4.3%
Engineering	3.40	4.87	1.47	43.2%
Geosciences	3.90	4.23	0.33	8.5%
Mathematical and Physical Sciences	5.00	5.97	0.97	19.4%
Social, Behavioral and Economic Sciences	5.40	5.46	0.06	1.1%
Office of Polar Programs	1.10	1.12	0.02	1.8%
Integrative Activities	0.00	20.00	20.00	N/A
Subtotal, Research and Related Activities	21.65	44.78	23.13	106.8%
Education and Human Resources	123.17	139.91	16.74	13.6%
Total, Learning for the 21st Century Workforce	$144.82	$184.69	$39.87	27.5%

Totals may not add due to rounding.

(Millions of Dollars)

FY 2001 Actual	FY 2002 Current Plan	FY 2003 Request	FY 2004	FY 2005
143.33	144.82	184.69	191.97	197.00

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY

IT-RESEARCH

NANOTECHNOLOGY

21ST CENTURY LEARNING

MATHEMATICAL SCIENCES

FUNDING LEVELS

FISCAL YEAR 2003 BUDGET

Proposed funding for the Mathematical Sciences priority area is as follows:

(Millions of Dollars)

	FY 2002 Current Plan	FY 2003 Request	Change Amount	Change Percent
Biological Sciences	0.00	0.91	0.91	N/A
Computer and Information Science and Engineering	0.00	2.29	2.29	N/A
Engineering	0.00	0.91	0.91	N/A
Geosciences	0.00	4.57	4.57	N/A
Mathematical and Physical Sciences	30.00	47.39	17.39	58.0%
Social, Behavioral and Economic Sciences	0.00	1.10	1.10	N/A
Office of Polar Programs	0.00	0.18	0.18	N/A
Subtotal, Research and Related Activities	$30.00	$57.35	$27.35	91.2%
Education and Human Resources	$0.00	$2.74	2.74	N/A
Total, Mathematical Sciences	$30.00	$60.09	$30.09	100.3%

Totals may not add due to rounding.

Long-term funding for the Mathematical Sciences priority area is as follows:

(Millions of Dollars)

FY 2002 Current Plan	FY 2003 Request	FY 2004	FY 2005	FY 2006	FY 2007
30.00	60.09	72.10	86.50	99.50	109.50

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY
IT-RESEARCH
NANOTECHNOLOGY
21ST CENTURY LEARNING
MATHEMATICAL SCIENCES
FUNDING LEVELS

FISCAL YEAR 2003 BUDGET

NSF FUNDING PROFILE	FY 2001 Actual	FY 2002 Estimate	FY 2003 Estimate
Number of Requests for Funding	43,515	44,550	45,900
Dollars Requested (in millions)	$28,784	$28,910	$30,000
Total Number of Awards	20,923	21,590	22,050
Statistics for Competitive Awards			
Number	9,925	10,430	10,630
Funding Rate	31%	32%	32%
Statistics for Research Grants			
Number of Research Grants	6,220	6,390	6,580
Median Annualized Award Size	$84,612	$86,000	$87,400
Average Annualized Award Size	$113,601	$113,000	$125,000
Average Duration (yrs.)	2.9	3.0	3.0

Research Grant excludes infrastructure + equipment + travel + workshops

PRIORITY AREAS

PRIORITY DOLLARS BY PROGRAM

BIOCOMPLEXITY
IT-RESEARCH
NANOTECHNOLOGY
21ST CENTURY LEARNING
MATHEMATICAL SCIENCES
FUNDING LEVELS

FISCAL YEAR 2003 CISE BUDGET
REGULAR DIVISION+PROGRAM+ITR FUNDS

CISE FUNDING LEVEL WAS RAISED BY 2.3% IN FY2003 BUDGET REQUEST

COMPUTER COMMUNICATIONS RESEARCH FUNDING LEVEL WAS RAISED BY 0.5%

INFORMATION TECHNOLOGY RESEARCH FUNDING LEVEL WAS RAISED BY 9.9%

ALL OTHER PROGRAMS IN CISE SAW A DROP IN THEIR FUNDING LEVELS.

HIGHLIGHTS

DOLLARS BY CISE DIVISION

DOLLARS BY CCR PROGRAM

NEXT SLIDE

PREVIOUS SLIDE

FISCAL YEAR 2003 CISE BUDGET
REGULAR DIVISION+PROGRAM+ITR FUNDS

	FY02	FY03	Change
COMPUTER-COMMUNICATIONS RESEARCH	$69,810	$70,170	0.5%
INFORMATION AND INTELLIGENT SYSTEMS	$52,060	$50,610	-2.8%
EXPERIMENTAL AND INTEGRATIVE ACTIVITIES	$62,670	$62,160	-0.8%
ADVANCED COMPUTATIONAL INFRASTRUCTURE AND RESEARCH ($80M INFRASTRUCTURE)	$86,970	$85,420	-1.8%
ADVANCED NETWORKING INFRASTRUCTURE AND RESEARCH ($47M INFRASTRUCTURE)	$69,860	$67,910	-2.8%
ITR	$173,510	$190,670	9.9%
CISE	$514,880	$526,940	2.3%

HIGHLIGHTS

DOLLARS BY CISE DIVISION

DOLLARS BY CCR PROGRAM

NEXT SLIDE

PREVIOUS SLIDE

FISCAL YEAR 2003 CISE BUDGET
REGULAR DIVISION+PROGRAM+ITR FUNDS

	FY02	FY03
THEORY OF COMPUTING	$6,108	N/A
GRAPHICS,SYMBOLIC,GEOM. COMPUTING	$8,492	N/A
OPERATING SYSTEMS AND COMPILERS	$6,493	N/A
SOFTWARE ENGINEERING AND LANGUAGES	$6,191	N/A
COMMUNICATIONS	$5,526	N/A
DESIGN AUTOMATION	$6,526	N/A
COMPUTER SYSTEMS ARCHITECTURE	$4,891	N/A
SIGNAL PROCESSING SYSTEMS	$5,628	N/A
TRUSTED COMPUTING	$5,420	N/A
EMBEDDED & HYBRID SYSTEMS	$4,397	N/A

HIGHLIGHTS

DOLLARS BY CISE DIVISION

DOLLARS BY CCR PROGRAM

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM

The CSA's main mission is
to foster basic research
on all aspects of computing
systems that have the potential
to lead to broad use and
applications

MISSION

SCOPE-SHORT TERM

SCOPE-LONG TERM

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM

NEAR TERM SCOPE

-*Metrics and parallelism:* performance evaluation of single processor and multiprocessor architectures.

-*Systems:* system latency, fault-tolerance , reliability, reconfiguration, quality of service, interprocessor communication, interconnection networks.

-*Memory:* memory architectures, bandwidth, latency, data prefetching and forwarding techniques, cache coherence and synchronization, processor-in-memory, active memory, and memory management problems.

-*Microarchitecture:* compiler-architecture interaction, out-of-order execution, VLIW, instruction and data prefetching and prediction, speculation, superscalar processing, multithreading, embedded processor design, lower power architectures.

-*Input/Output:* Disk organization, scheduling, data stream management, low-overhead protection, latency reduction, active disk, and high performance I/O design.

MISSION

SCOPE-SHORT TERM

SCOPE-LONG TERM

NEXT SLIDE

PREVIOUS SLIDE

LONGER TERM

Nanoscale Systems and
Molecular Architectures

About 10 hydrogen atoms standing side-by-side = 1 nanometer

A DNA Molecule = 2.5 nanometer

A Carbon nanotube = 1.5 nanometer

"Nanotechnology has given us the tools... to play with the ultimate toy box of nature- atoms and molecules. Everything is made from it... Possibilities to creat new things appear limitless."

Horst Stormer, Nobel Prize Winner-Physics
Columbia University

MISSION

SCOPE-SHORT TERM

SCOPE-LONG TERM

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM IN NUMBERS

FY2001 Operating Budget*:	$5,846,939
Continuing Commitments:	$955,061
#New Regular Proposals:	$3,858,826
#Career Proposals:	$968,318
#REUs:	$64,734
FY2002 Operating Budget*:	$4,677,413
Continuing Commitments:	$1,434,905
#New Regular Proposals:	$2,264,027
#Career Proposals:	$881,481
#REUs:	$47,000
ITR Commitments	$50,000

*Excludes IPA + Workshops

BUDGET

FUNDING Regular

FUNDING Career

FUNDING Total

FY02+ COMMITTMENTS

BUDGET TREND

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM IN NUMBERS

FY2001

Regular Proposals Received: 36

Regular Proposals Awarded : 17 (%48)

Regular Proposals Declined : 19 (%52)

Award Range: $189K - $422K

Average Award/year: $87K

FY2002

Regular Proposals Received: 50

Regular Proposals Awarded : 13 (%26)

Regular Proposals Declined : 37 (%74)

Award Range: $144K - $331K

Average Award/year: $85K

BUDGET

FUNDING Regular

FUNDING Career

FUNDING Total

FY02+ COMMITTMENTS

BUDGET TREND

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM IN NUMBERS

FY2001

Career Proposals Received: 24

Career Proposals Awarded : 8 (%33)

Career Proposals Declined : 16 (%67)

Award Range: $268K - $300K

Average Award/year: ~$60K

FY2002

Career Proposals Received: 20

Career Proposals Awarded : 7 (%35)

Career Proposals Declined : 13 (%65)

Award Range: $349K - $375K

Average Award: ~$75K

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM IN NUMBERS

FY2001

Proposals Received: 60

Proposals Awarded : 25 (%42)

Proposals Declined : 35 (%58)

Award Range: $189K - $422K

Average Award/year: ~$78K

FY2002

Proposals Received: 70

Proposals Awarded : 20 (%28)

Proposals Declined : 50 (%72)

Award Range: $144K - $375K

Average Award: ~$81K

BUDGET

FUNDING

Regular

FUNDING

Career

FUNDING

Total

FY02+ COMMITTMENTS

BUDGET TREND

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM IN NUMBERS

FY2003: $1,173,770

FY2004: $1,301,517

FY2005: $ 760,308

BUDGET

FUNDING

Regular

FUNDING

Career

FUNDING

Total

FY02+ COMMITTMENTS

BUDGET TREND

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM IN NUMBERS

FY1997	$ 5.6M
FY1998	$ 6.5M
FY1999	$ 5,972,535
FY2000	$ 5,733,751
FY2001	$ 5,846,939
FY2002	$ 4,677,413
Total:	~$34M

BUDGET

FUNDING

Regular

FUNDING

Career

FUNDING

Total

FY02+ COMMITTMENTS

BUDGET TREND

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM
BY FUNDED TOPICS

FY2001

Energy-Exposed Instruction Sets

New Directions in Speculative Execution-Microarchitecture

A Program of Research and Education in Storage Systems Design

Energy Efficient Architectures and Their Interaction with Software: A Java Perspective

rePlay: A Microarchitecture to support Dynamic Program Optimization

Cooperative Redundant Threads

Computer Arithmetic Algorithms and Scalable Hardware designs for Cryptographic Applications

Closing the Memory Gap for Unstructured Applications

CAREER-01

CAREER-02

REGULAR-01

REGULAR-02

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM BY FUNDED TOPICS

FY2002

Systematic Design Space Exploration

Memory Controller Interconnect and Policy

Interconnect Design for Programmable Computation

Semantic Decomposition of Instruction Sets

The Evaluation and Design of a Scalable, High Performance and Energy Efficient Microprocessor Architecture

Control-Theoretic Techniques and Thermal/Power Modeling For Dynamically Managing Temperature and Power in Microprocessors

Soft-Instruction Set Computing

CAREER-01

CAREER-02

REGULAR-01

REGULAR-02

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM
BY FUNDED TOPICS

FY2001

High Performance Internet Router Architectures

Exploiting Critical Path in the Design and Performance Analysis of Modern Processors

Some Coding Techniques for VLSI and Computer Systems

New Prediction Paradigms for Parallel and Distributed Systems

Trace-Driven Evaluations of the Memory Behavior of Large Commercial Applications

ADORE - A Framework for Adaptive Object Code Reoptimization

Incorporating Fault-tolerance at the Application Level

Small Scale Dynamic Recofigurability for Large-Scale Benefits

Instruction Set Architecture for Pervasive Security

Efficient Fine Grained Synchronization Support Using Full/Empty Tagged Shared Memory Cache Coherence

Fault Tolerance in System Architectures Implementing the Compression, Transmission and Expansion of Data

Optimizing Integrated Memory-Hierarchy Design

Critical Path Computing

Performance Evaluation of Disk Arrays

Architectural Support for Scalable High-Speed Routers

High Performance Parallel I/O

Compiler-Inserted Control Independence Information for Latency Hiding and Reduced Branch Cost

CAREER-01

CAREER-02

REGULAR-01

REGULAR-02

NEXT SLIDE

PREVIOUS SLIDE

THE COMPUTER SYSTEMS ARCHITECTURE PROGRAM BY FUNDED TOPICS

FY2002

- Using Simultaneous Multithreaded Processors for Soft Real-Time Applications
- High-Performance Asynchronous Computer and SoC Architecture
- Next Generation Load Value Predictors
- Signals Approaches to Computer Architecture Prediction Mechanisms
- Compiler-driven Design Space Exploration for Heterogeneous System-on-a-Chip
- Information Encoding for Energy Efficient Processor Design
- Main Memory Power Management
- Dynamic Superpipelining: Shaping Microarchitecture for Variable Frequency
- Sequential Architectures for Quantum Computation
- QoS Provisioning in InfiniBand Architecture for System Area Networks
- Theoretical Support for Efficient Network Discovery and Reconfiguration Techniques
- Operating System and Architectural Implications of Programmable Network Interfaces
- WDM Optical Interconnect Architectures for Parallel and Distributed Computing and Communications

PROPOSAL REVIEW AND AWARD/DECLINATION PROCESS

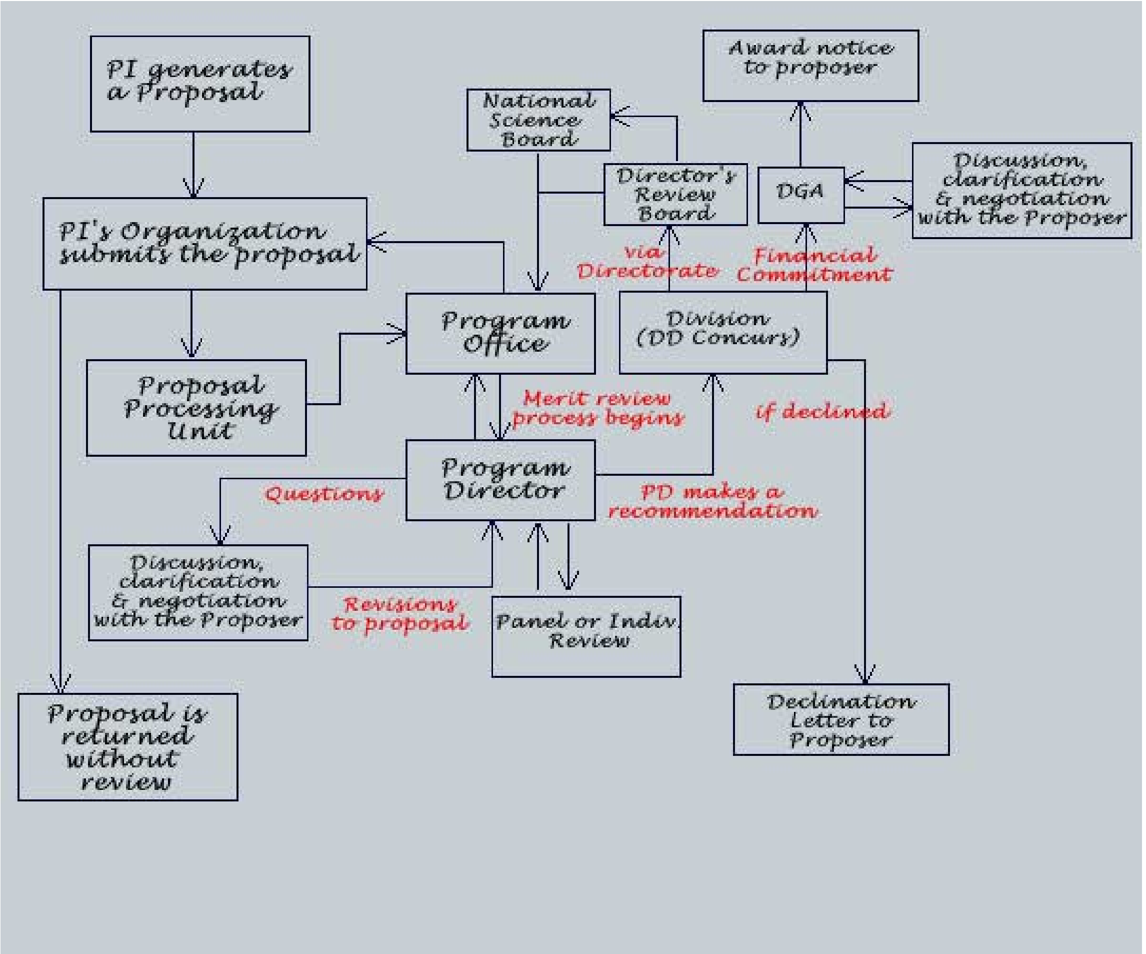

PROCESS

FORM -7

NEXT SLIDE

PREVIOUS SLIDE

PROPOSAL REVIEW AND AWARD/DECLINATION PROCESS

Continuation of Form 7

Computer Systems Architecture Program

DATE:XXXXX
PROPOSAL NUMBER: XXXXX
INSTITUTION: XXXXX
PI: XXXXX
TITLE: XXXXX

This CAREER proposal was evaluated in the C-CR Computer Systems Architecture (CSA) Program, and reviewed by a Special Emphasis Panel. The meeting was devoted to the review and evaluation of those CAREER proposals in the general area of Computer architecture that had been assigned to the CSA program this fiscal year. The review process for the Special Emphasis Panel is described in the panel minutes.

The panel reviewed a total of XX proposals. Results are tabulated in the appendix. Shown there are the individual reviewers' overall ratings and the panel ratings, of Fund (FND), Fund if Possible (FIP) and Do Not Fund (DNF). The results are:

- XX proposals ranked in the FND category;
- XX proposals grouped in the FIP category, and
ranked relative to each other within the category;
- XX proposals grouped in DNF category.

This proposal was initially ranked in the (FIP) category and later moved to the FND category by the panel. It has individual ratings of XEs, YV, and ZV/Gs which are good ratings within this particular panel group. The panel was quite impressed with this proposal, and it suggested that it has a high potential impact. On the less positive side, the panel thought that the goals of the proposal are very ambitious even though the PI is very capable. One panelist stated that the PI is perhaps the best researcher .... Another panelist praised the proposal and PI's background and his publication record. Other panelists... were generally supportive of the proposal...

The main thrust of this proposal is to develop....

I agree with the panel that this is an ambitious task. At the same time, the potential impact and benefits of ... that will be developed during the project for computer architecture research can be huge. The PI's extremely well-qualified, and the work on... is already underway. Given all these facts, I have little reservation, if any, that the proposal may not succeed. Therefore, I recommend that it be funded subject to the reduction of the budget to ....

A. Yavuz Oruc

PROCESS

FORM -7

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/patterson.htm

Top Six ("greatest potential for fruitful research in the next five years"):

-Dynamically scheduled uniprocessors
-Intelligent DRAM/Processor In Memory
-Building systems from systems
-High-performance Networks and Storage
-Benchmarking and Performance Evaluation using Commercial Software
-Fast compilers and novel instruction sets for network applets

Bottom Four ("least potential for fruitful research in the next five years"):

-Multiple processors on a chip
-Massively Parallel Architectures
-Cache studies
-Branch prediction studies

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/anant.htm

Finally, I guess I should say something about the topics suggested by our moderator. I believe the following topics have the least research potential:

uniprocessor issues, ILP, performance evaluation. Most research in this area has been producing minuscule performance improvements.

Multiprocessors and related topics are interesting, but industry is fast catching up. Performance-driven research should be given up in favor of issues of protection, fault tolerance, security, and support for new application domains.

I think application-driven architectures and new workloads hold the most promise but alas industry seems to be way ahead of academic research.

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/mudge.htm

Research Areas That Have The Least Potential For Fruitful Research

-Massively Parallel Processors
-Caches, Instruction Level Parallelism, Cache Coherence
-Instruction Set Design

Research Areas That Still Have Potential For Fruitful Research

-Parallel Processing (Excluding MPP)
- Small Scale Parallel Processing
-Performance evaluation

Research Areas That Have The Greatest Potential For Fruitful Research

-Latency Tolerance Memory System Design
-Architectures for New Computing Situations
-Architecture Implications for Low power
-Design Validation or Correctness

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/smith.htm

... networks of workstations, in my opinion, do not seem to be the solution to large scale computational problems. They seem to be useful primarily for throughput-oriented computation. Many of the technical problems of the first MPPs, e.g. high latency, low bandwidth, difficulties in program development, also seem to be present in networks of workstations -- only much worse.

It is my opinion that non-numeric applications should be the drivers for future processor architecture research. While numeric applications are important, they tend to be limited more by data path considerations than control.

In general, parallelism tends to be easier to find in numeric applications. And methods that can exploit the irregular parallelism of non-numeric applications can likely be applied to those portions of numeric applications that are more difficult to parallelize.

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/emmer.htm

In summary, the research areas that I feel are most promising are prediction, in the short term; small scale multiprocessing/ multithreading and medium scale multiprocessing in the medium term; and processor-in-memory and more radical ILP schemes in the long term.

The areas that I least expect to see fruitful results are: application-driven architecture or high-performance storage. In each of these cases, I see little need for further refinement or opportunities for new insight from new approaches.

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/colwell.htm

I think research on uniprocessor architectures is likely to suffer the same fate that it has for the last 5 years: industry swamps it with money.

The research results are either obsolete or unusable by real designs. One cure is to look out farther than just the next design.

Do what Patt, Sohi and Uht doing; ask what might become feasible two or three generations out, and set out to find and conquer the limitations that will pop up. "What should one do with (say) 50M transistors?" But even then, I'm skeptical that NSF ought to be funding very much of this kind of work. Yes, it's good prep for the students to enter industry, thank you very much. But a lot of it just plain misses the mark, either too general for justifiable conclusions to be drawn, or so specific that nothing can be extracted for real-world systems.

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

http://www.cise.nsf.gov/evnt/wksp/baskett.htm

1. How can we get LOTS of memory bandwidth into each of our processors?

2. Is it useful to build interconnects that are either programmable or adaptable, "smart interconnects"? ...

3. Consider the possibilities of a processor and memory on the same silicon chip.

4. Are there some innovations in low power design yet to be had?

5. Is there a breakthrough in optical interconnects around the corner? .... Think what an optical backplane might be like. It's easy for the imagination to get ahead of what's currently or soon to be practical. But the promise is alluring.

On the other side of the coin:

1. Can multithreaded architectures ever win any single thread benchmark competitions? Academics have lavished a lot of attention on multithreaded architectures ... Industry says that you can't win any benchmark competitions with multithreaded architectures so why bother? Is there a way to win or is there a way to change the rules?

2. Putting multiple processors on a single chip seems to make the bandwidth problem worse. More processors needing more bandwidth sharing that bandwidth over fewer pins. How can this work?

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

WHERE IS COMPUTER SYSTEMS ARCHITECTURE RESEARCH HEADED?

1. Will it ever possible to build processors using molecules-I mean molecule by molecule? What will be the building blocks? What are the architectural barriers that need to be overcome to get there? How much computation potential does a single molecule hold?

2. Will computers always be built in layers? Can we remove the ISA layer from the equation and "Silicon compile" applications directly over a programmable substrate? What will the architecture of such a system look like?

3. Can we model and predict the behavior of computers more accurately? Is it feasible to build simulators that can be switched across various layers of computers seamlessly like lenses?

4. What are the implications and potential of building "smart" processors? What are the characteristics of such processors? How will the ISA of smart processors differ from conventional processors?

ARCHITECTS' VIEW

DAVE PATTERSON

ANANT AGRAWAL

TREVER MUDGE

JAMES SMITH

JOEL EMMER

ROBERT COLWELL

FOREST BASKETT

NON-ARCHITECT'S VIEW

AKA MY VIEW

NEXT SLIDE

PREVIOUS SLIDE

PARTING REMARKS

Computer architecture research has been a melting pot of theoretical concepts for devising concrete techniques to design and develop computing systems of all sorts eversince the early pioneers embark upon building computers with bulky vacuum tubes.

It appears that, once again, computer architects are presented with a golden opportunity to break new ground, this time, to architect molecules into artifacts of all kinds, not just computer systems as we have known them.

Among the key questions that need to be addressed are (as identified at the 1st Molecular Architecture Workshop, Nov. 2001)

-How can the molecular level interactions be expressed in abstract ways, how can such abstractions be translated into primitive building blocks, and what kinds of models and design tools should be used to emphasize such interactions?

- How does one develop new design strategies for combining such primitive building blocks into larger functional subsystems, and then scale them into even larger molecular systems?

- How could such molecular systems be interfaced with legacy technologies such as silicon?

In my view, confronting and tackling questions such as these on molecular and other nano-scale structures holds the key to building as rich a track-record of computer architecture research in the next two decades as we have seen during the last two decades.

PARTING REMARKS

To All MicroArchitects*

Microarchitects love to
speculate values and such
for more mips and whetstones
but not by so much!

One less clock-tick here,
another saved there
It sure pays to execute
threads everywhere!

Stalls are so painful,
squashes are sheer torture
Yet they are every
microarchitect's overture!

Moore claimed transistors
habitually double in a chip
In the hands of microarchitects
this has become a whip!

Billion transistor circuits
finally appear to be in sight
To speculate them will sure
be every μarchitect's birth right!

*It can be argued that the PCs on our desktops have reached their MFLOPS performance because of the Moore's Law but without the significant discoveries in the microarchitecture research, it is difficult to imagine that the sheer doubling of circuit densities every 18 months would have brought us all of this impressive performance.

CONCLUSIONS

MICROARCHITECT

NEXT SLIDE

PREVIOUS SLIDE