

高级计算机体系结构

陈文智 浙江大学计算机学院

chenwz@zju.edu.cn

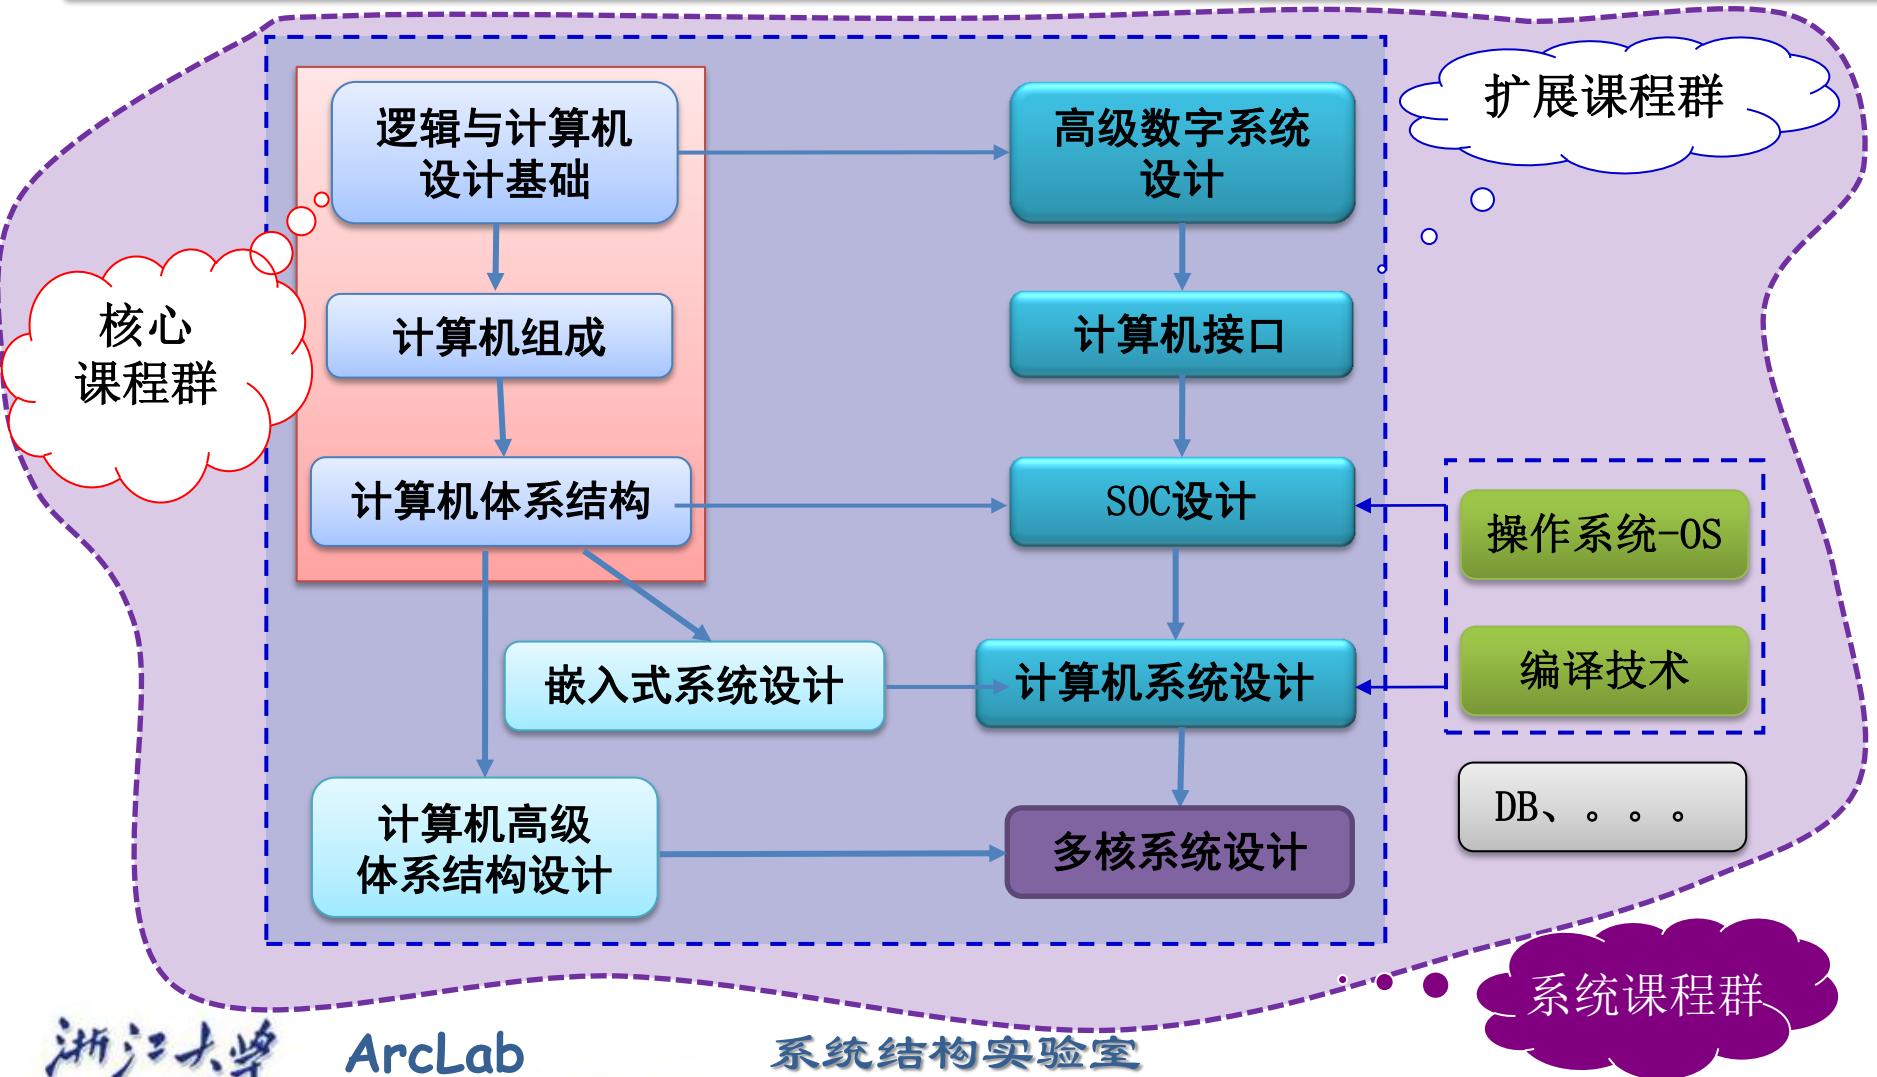
2014年9月

课程现状

- 课程性质
 - 本科/硕士/博士
- 课程衔接
 - 本科体系/相关课程
- 课程目的
 - 掌握计算机系统的设计思想、设计原理、设计技术。能够分析、设计计算机的相关功能部件及处理器，掌握现代计算机处理器中流水线技术、指令级并行性开发技术、多处理器计算机设计技术。

改革中的课程体系

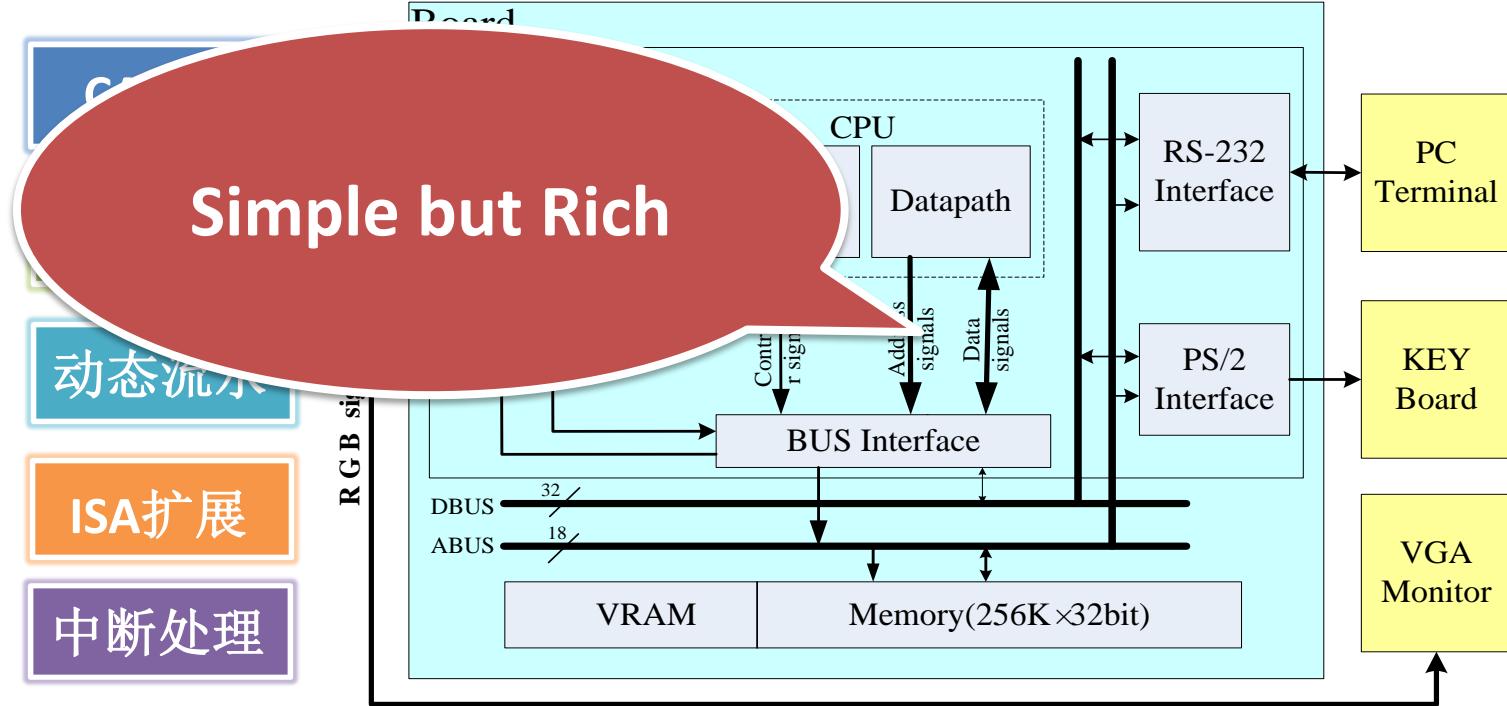
以系统设计为统一视图，建立层次化、循序递进、开放式课程群的课程体系



改革中的实验体系

自主研发统一的实验平台，建立实验成果库，构建系列化、递进式的实验体系

渐近的计算机硬件系统



数字逻辑 → 计算机组成 (CPU) → 体系结构 (流水CPU) → 简易计算机系统

高级数字系统设计 → 计算机接口 → SOC设计 → 计算机系统设计 → 多核系统设计
→ 嵌入式系统设计

操作系统 + 编译系统 + 应用软件

课程内容

教学内容	备注
☆ 绪论: 计算机分类,计算机体系结构定义,实现技术的发展趋势,集成电路功耗的发展趋势,成本、价格的走势,可靠性,性能测量与分析报告,计算机量化设计原则.	第1、A章 复习
☆ 基本流水线技术: 流水线竞争,流水线是如何实现的,流水线实现难点,多周期操作的指令流水线,MIPS R4000流水线	
☆ 指令级并行性及其开发技术: 指令级并行概念和挑战,数据相关和数据冲突,控制相关,开发ILP的基本编译技术,动态转移预测技术.	第A、2章 复习
★ 指令级并行性及其开发技术: 动态调度技术,动态调度实例及算法.	第2章
★ 指令级并行性及其开发技术: 基于硬件的投机技术, 多发射和静态调度技术, 用动态调度、多发射和投机技术开发.	第2章
★ ILP的局限性: 硬件和软件投机技术, 多线程技术, 多发射处理器的性能。	第3章
★ 多处理器和线程级并行性: 多处理器计算机结构分类, 共享存储器结构及其Cache监听协议, 分布存储结构及其基于目录的协议	第4章
☆ 多处理器和线程级并行性: 多处理器同步技术, 存储器连贯性模型, 实例: Sun T1 多处理器计算机	第4章
☆ 向量处理器技术	
☆ VLIW和EPIC的硬件和软件技术	第F、G章

教材

英文（主要）：

《Computer Architecture—A Quantitative Approach》，4th，机械工业出版社原版影印，**¥78。**

中文（参考）：

《计算机体系结构——一种定量的方法》，清华出版社，**¥72**

教材作者介绍

- **John L. Hennessy**

Standford University校长、**IEEE**和**ACM**会员，美国国家工程研究院院士及美国科学艺术研究院院士，2001获**Eckert-Mauchly奖(RISC)**，2001年**Seymour Cray**计算机工程奖得主，2000年获**IEEE John von Neumann奖**。

- **David A. Patterson**

University of California, Berkeley教授，美国国家工程研究院院士，**IEEE**和**ACM**会员，获**IEEE James H. Muligan教育奖章**，1995获**IEEE技术成就奖(RISC)**，1999年**IEEE Reynold Johnson信息存储奖**，2000年获**IEEE John von Neumann奖**。

Eckert-Mauchly Award

- | | |
|--|---|
| 2010 <u>Bill Dally</u> | 1994 <u>Thornton, James E.</u> |
| 2009 <u>Joel Emer</u> | 1993 <u>Kuck, David J</u> |
| 2008 <u>David Patterson</u> | 1992 <u>Flynn, Michael J.</u> |
| 2007 <u>Valero, Mateo</u> | 1991 <u>Smith, Burton J.</u> |
| 2006 <u>Pomerene, James H</u> | 1990 <u>Batcher, Kenneth E.</u> |
| 2005 <u>Colwell, Robert P.</u> | 1989 <u>Cray, Seymour</u> |
| 2004 <u>Brooks, Frederick P.</u> | 1988 <u>Siewiorek, Daniel P.</u> |
| 2003 <u>Fisher, Joseph A. (Josh)</u> | 1987 <u>Amdahl, Gene M.</u> |
| 2002 <u>Rau, B. Ramakrishna (Bob)</u> | 1986 <u>Cragon, Harvey G</u> |
| 2001 <u>Hennessy, John</u> | 1985 <u>Cocke, John</u> |
| 2000 <u>Davidson, Edward</u> | 1984 <u>Dennis, Jack B.</u> |
| 1999 <u>Smith, James E.</u> | 1983 <u>Kilburn, Tom</u> |
| 1998 <u>Watanabe, T.</u> | 1982 <u>Bell, C. Gordon</u> |
| 1997 <u>Tomasulo, Robert</u> | 1981 <u>Clark, Wesley A.</u> |
| 1996 <u>Patt, Yale</u> | 1980 <u>Wilkes, Maurice V.</u> |
| 1995 <u>Crawford, John</u> | 1979 <u>Barton, Robert S.</u> |

Big Men in Architecture(1)

- **2007 Mateo Valero**

<http://personals.ac.upc.edu/mateo/>

For important contributions to instruction level parallelism and superscalar processor design.



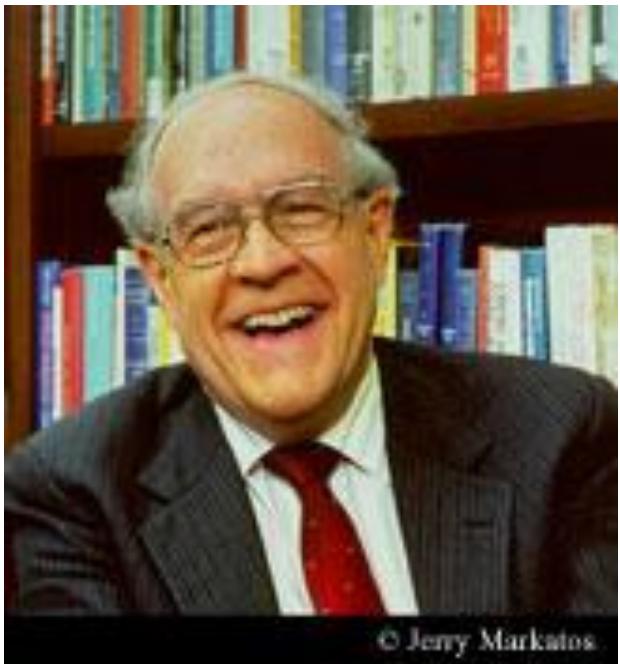
Big Men in Architecture(2)



- **2001 Hennessy, John**

For being the founder and chief architect of the MIPS Computer Systems and contributing to the development of the landmark MIPS R2000 microprocessor.

Big Men in Architecture(3)



© Jerry Markatos

Frederick P. Brooks

<http://www.cs.unc.edu/~brooks/>

- **1999 ACM Turing Award**

landmark contributions to computer architecture, operating systems, and software engineering."

2004 Eckert-Mauchly Award

"For the definition of computer architecture and contributions to the concept of computer families and to the principles of instruction set design; for seminal contributions in instruction sequencing, including interrupt systems and execute instructions; and for contributions to the IBM 360 instruction set architecture."

Big Men in Architecture(4)

1997 Robert Tomasulo

For the ingenious Tomasulo's algorithm, which enabled out-of-order execution processors to be implemented.

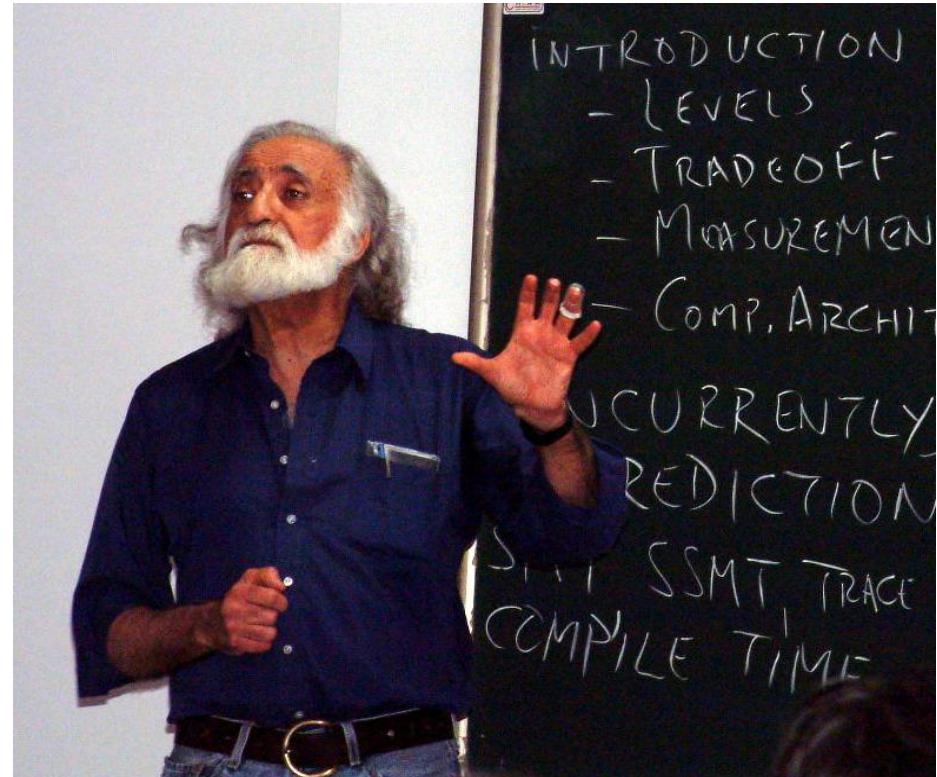
January 30 2008, Tomasulo spoke at the College of Engineering at the University of Michigan about his career and the history and development of out-of-order execution

<http://inst-tech.engin.umich.edu/media/index.php?sk=tomasulo>

Big Men in Architecture(5)

- **1996 Yale Patt**

For important contributions to instruction level parallelism and superscalar processor design.



Big Men in Architecture(6)

- 1992 **Michael J. Flynn**

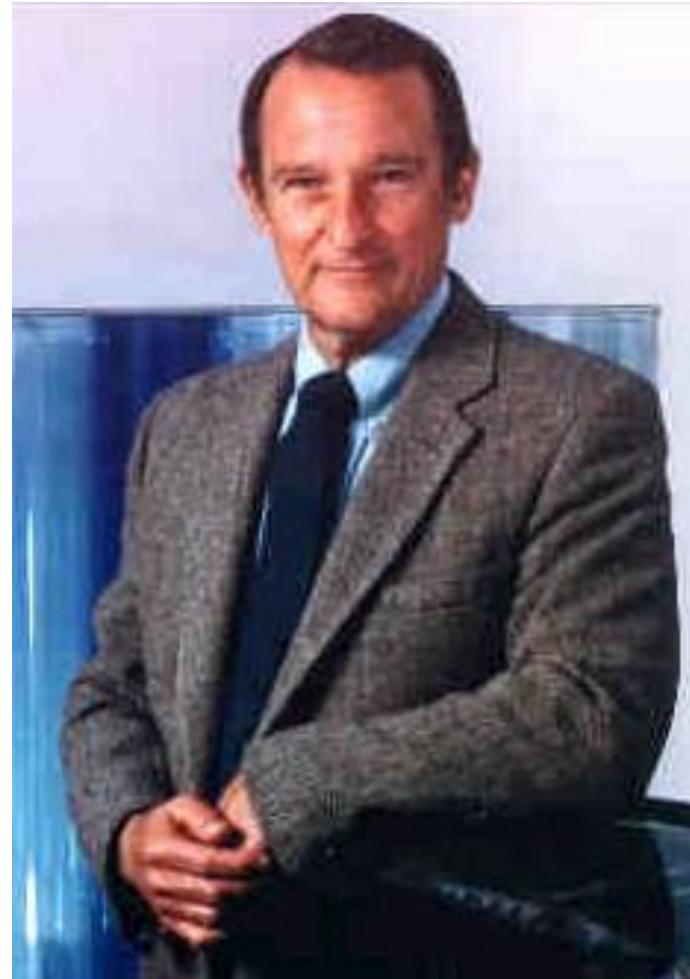
<http://www.cpe.calpoly.edu/IAB/flynn.html>

- For his important and seminal contributions to processor organization and classification, computer arithmetic and performance evaluation.

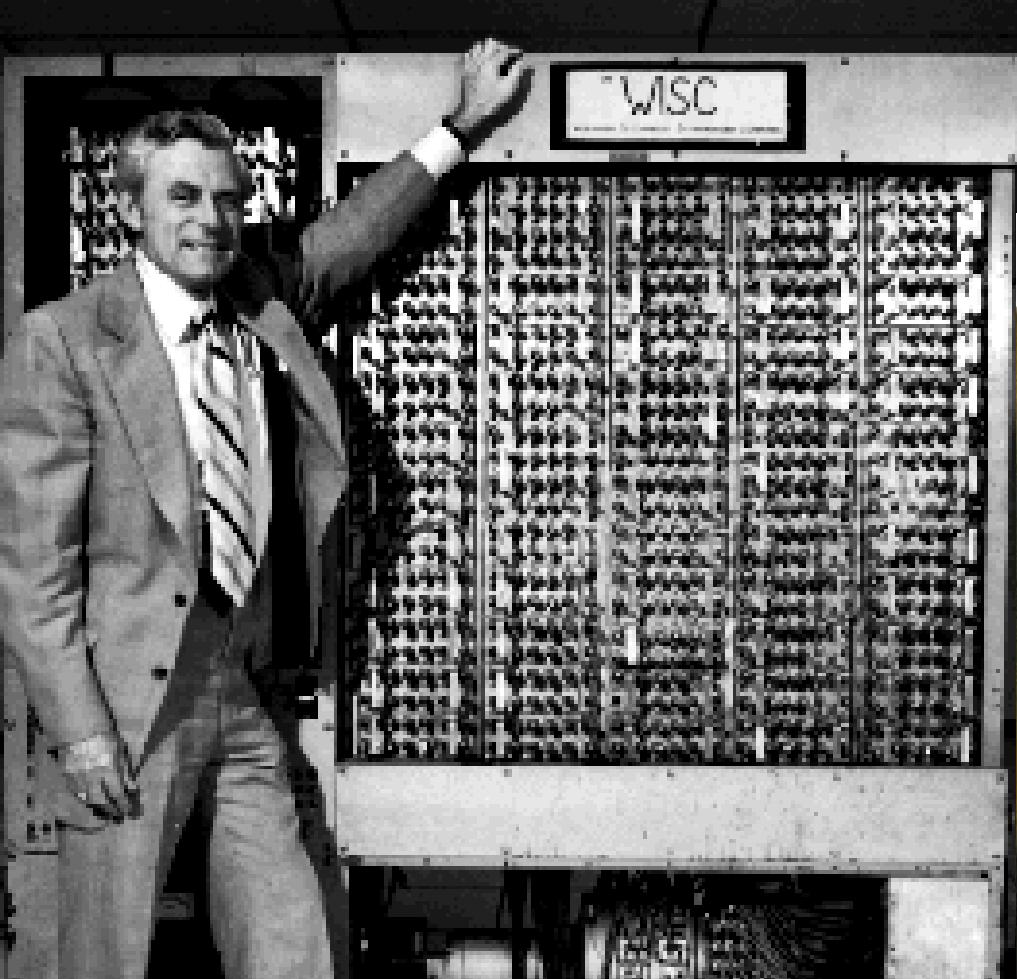


Big Men in Architecture(7)

- 1989 Cray, Seymour
- For a career of achievements that have advanced supercomputing design.

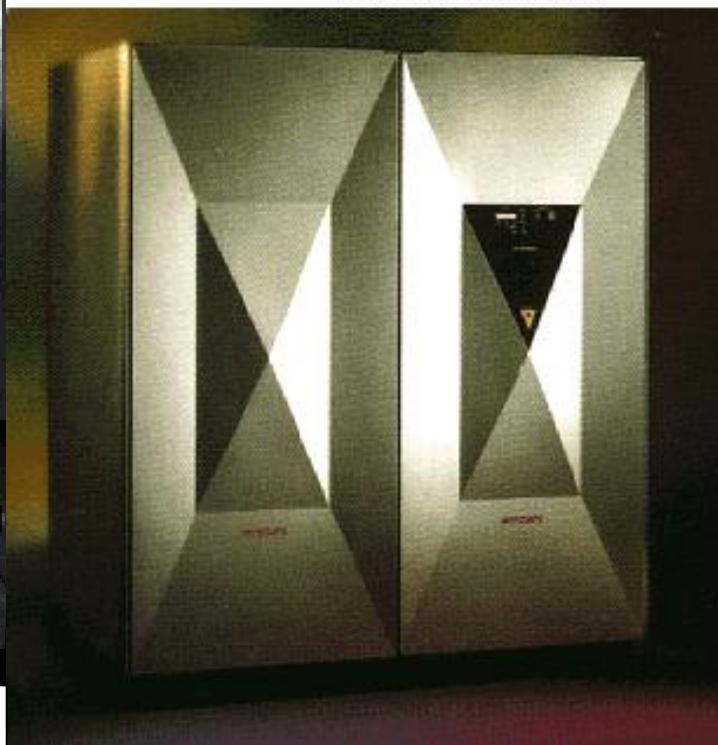


Architecture(8)



In 1975, Dr. Amdahl stands beside the Wisconsin Integrally Synchronized Computer (WISC), which he designed in 1950. It was built in 1952. (Image courtesy of Dr. Gene M. Amdahl.)

From Computer Desktop Encyclopedia
Reproduced with permission.
© 1997 Amdahl Corporation



David A. Patterson (UC Berkeley)

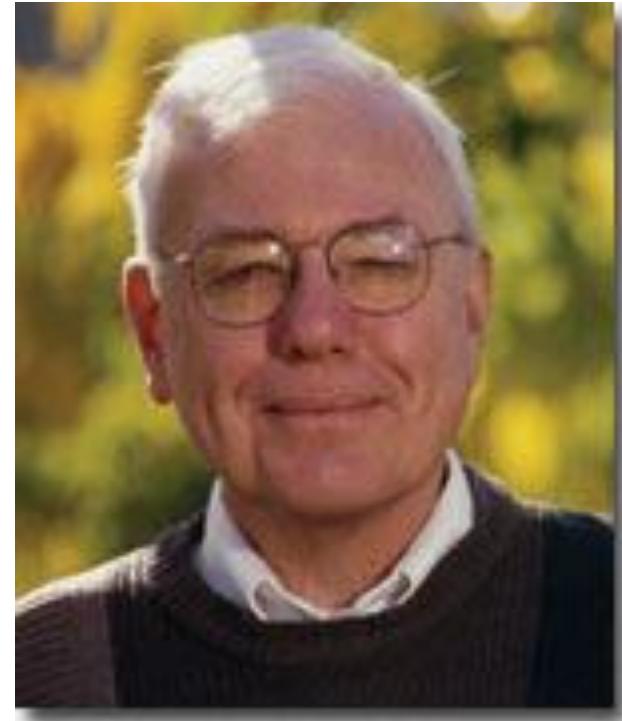
- He led the design and implementation of **RISC I** (the foundation of the **SPARC** architecture)
- Leader of **RAID**
- involved in the Network of Workstations (**NOW**) project
- Research Accelerator for Multiple Processors (**RAMP**)





A. M. Turing Award-2009

- **Charles P Thacker**
- For the pioneering design and realization of the first modern personal computer -- the Alto at Xerox PARC -- and seminal inventions and contributions to local area networks (including the Ethernet), multiprocessor workstations, snooping cache coherence protocols, and tablet personal computers.



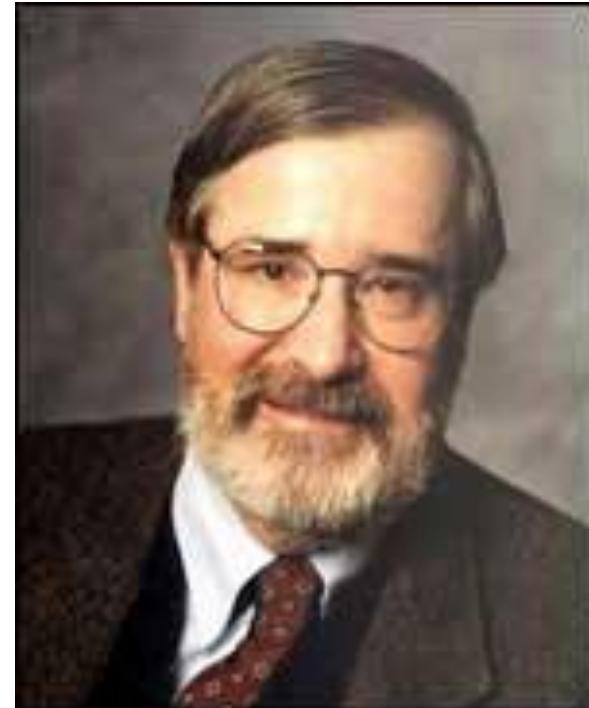
A. M. Turing Award-2008

- **Liskov, Barbara**
- For contributions to practical and theoretical foundations of programming language and system design, especially related to data abstraction, fault tolerance, and distributed computing.



A. M. Turing Award-2007

- **Clarke, Edmund M**
- For his role in developing Model-Checking into a highly effective verification technology, widely adopted in the hardware and software industries.



A. M. Turing Award-2007

- **Emerson, E Allen**
- For his role in developing Model-Checking into a highly effective verification technology, widely adopted in the hardware and software industries.



A. M. Turing Award-2007

- **Sifakis, Joseph**
- For his role in developing Model-Checking into a highly effective verification technology, widely adopted in the hardware and software industries.



课程考核

- 方式：
 - 笔试
 - 报告
 - 随堂测验

My class: Gem5 Academic Program

课程要求

- 1、每个学生查找近三年的国际期刊或国际会议上计算机体系结构相关论文，每篇论文篇幅在6页以上。撰写一篇字数不少于8000字的读书报告。读书报告的参考文献要求在10篇以上，并且引用的参考文献必须在文中加以引用标注。
- 2、Gem5

Top conference in CA

- **ISCA**: The International Symposium on Computer Architecture
 - **MICRO**: Intl Symp on Microarchitecture
 - **ASPLOS**: Architectural Support for Prog Lang and OS
 - **HPCA**: IEEE Symp on High-Perf Comp Architecture
 - **OSDI**: USENIX Symposium on Operating Systems Design and Implementation
 - **SOSP**: The ACM Symposium on Operating Systems Principles
-

Top10 HPC, 2008

Rank	Site	Computer
1	<u>DOE/NNSA/LANL</u> United States	<u>Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz , Voltaire Infiniband</u> IBM
2	<u>Oak Ridge National Laboratory</u> United States	<u>Jaguar - Cray XT5 QC 2.3 GHz</u> Cray Inc.
3	<u>NASA/Ames Research Center/NAS</u> United States	<u>Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0/2.66 GHz</u> SGI
4	<u>DOE/NNSA/LLNL</u> United States	<u>BlueGene/L - eServer Blue Gene Solution</u> IBM
5	<u>Argonne National Laboratory</u> United States	<u>Blue Gene/P Solution</u> IBM
6	<u>Texas Advanced Computing Center/Univ. of Texas</u> United States	<u>Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband</u> Sun Microsystems
7	<u>NERSC/LBNL</u> United States	<u>Franklin - Cray XT4 QuadCore 2.3 GHz</u> Cray Inc.
8	<u>Oak Ridge National Laboratory</u> United States	<u>Jaguar - Cray XT4 QuadCore 2.1 GHz</u> Cray Inc.
9	<u>NNSA/Sandia National Laboratories</u> United States	<u>Red Storm - Sandia/ Cray Red Storm, XT3/4, 2.4/2.2 GHz dual/quad core</u> Cray Inc.
10	<u>Shanghai Supercomputer Center</u> China	<u>Dawning 5000A - Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008</u> Dawning

Fastest computer in China

- 2008 ShuGuang5000A
 - 180.6 TeraFlops
 - 122.88TB RAM, 30720 computing Cell
- 2004 ShuGuang4000A
 - 11 TeraFLOPS
 - rank 10 in top 500 in June,
- 2003 ShenTeng6800
 - 5.324 TeraFLOPS
- 2002 ShenTeng1800
 - 2.04 TeraFLOPS
- 2000 YinHe IV
 - 1024 CPU
 - 1 TeraFLOPS



ShuGuang
6000 ?

TOP 10, 11/2009

1	<u>Jaguar - Cray XT5-HE Opteron Six Core 2.6 GHz</u>
2	<u>Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband</u>
3	<u>Kraken XT5 - Cray XT5-HE Opteron Six Core 2.6 GHz</u>
4	<u>JUGENE - Blue Gene/P Solution</u>
5	<u>Tianhe-1 - NUDT TH-1 Cluster, Xeon E5540/E5450, ATI Radeon HD 4870, Infiniband</u>
6	<u>Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0 GHz/Nehalem EP 2.93 Ghz</u>
7	<u>BlueGene/L - eServer Blue Gene Solution</u>
8	<u>Blue Gene/P Solution</u>
9	<u>Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband</u>
10	<u>Red Sky - Sun Blade x6275, Xeon X55xx 2.93 Ghz, Infiniband</u>

Top10 HPC, June, 2010

Rank	Site	Computer
1	<u>Oak Ridge National Laboratory</u> United States	<u>Jaguar - Cray XT5-HE Opteron Six Core 2.6 GHz</u> Cray Inc.
2	<u>National Supercomputing Centre in Shenzhen (NSCS)</u> China	<u>Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU</u> Dawning
3	<u>DOE/NNSA/LANL</u> United States	<u>Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband</u> IBM
4	<u>National Institute for Computational Sciences/University of Tennessee</u> United States	<u>Kraken XT5 - Cray XT5-HE Opteron Six Core 2.6 GHz</u> Cray Inc.
5	<u>Forschungszentrum Juelich (FZJ)</u> Germany	<u>JUGENE - Blue Gene/P Solution</u> IBM
6	<u>NASA/Ames Research Center/NAS</u> United States	<u>Pleiades - SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon Westmere 2.93 Ghz, Infiniband</u> SGI
7	<u>National SuperComputer Center in Tianjin/NUDT</u> China	<u>Tianhe-1 - NUDT TH-1 Cluster, Xeon E5540/E5450, ATI Radeon HD 4870 2, Infiniband</u> NUDT
8	<u>DOE/NNSA/LLNL</u> United States	<u>BlueGene/L - eServer Blue Gene Solution</u> IBM
9	<u>Argonne National Laboratory</u> United States	<u>Intrepid - Blue Gene/P Solution</u> IBM
10	<u>Sandia National Laboratories / National Renewable Energy Laboratory</u> United States	<u>Red Sky - Sun Blade x6275, Xeon X55xx 2.93 Ghz, Infiniband</u> Sun Microsystems

TOP 10, 11/2010

1	<u>Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C</u>
2	<u>Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz</u>
3	<u>Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU</u>
4	<u>TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows</u>
5	<u>Hopper - Cray XE6 12-core 2.1 GHz</u>
6	<u>Tera-100 - Bull bullex super-node S6010/S6030</u>
7	<u>Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband</u>
8	<u>Kraken XT5 - Cray XT5-HE Opteron 6-core 2.6 GHz</u>
9	<u>JUGENE - Blue Gene/P Solution</u>
10	<u>Cielo - Cray XE6 8-core 2.4 GHz</u>

TOP 10 Sites for June 2011

Rank	Site	Computer
1	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu
2	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C NUDT
3	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.
4	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, Nvidia Tesla C2050 GPU Dawning
5	GSIC Center, Tokyo Institute of Technology Japan	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows NEC/HP
6	DOE/NNSA/LANL/SNL United States	Cielo - Cray XE6 8-core 2.4 GHz Cray Inc.
7	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon 5570/5670 2.93 Ghz, Infiniband SGI
8	DOE/SC/LBNL/NERSC United States	Hopper - Cray XE6 12-core 2.1 GHz Cray Inc.
9	Commissariat à l'Energie Atomique (CEA) France	Tera-100 - Bull bullex super-node S6010/S6030 Bull SA
10	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband IBM

TOP 10 Sites for November 2011

Rank	Site	Computer
1	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu
2	National Supercomputing Center in Tianjin China	NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 NUDT
3	DOE/SC/Oak Ridge National Laboratory United States	Cray XT5-HE Opteron 6-core 2.6 GHz Cray Inc.
4	National Supercomputing Centre in Shenzhen (NSCS) China	Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 Dawning
5	GSIC Center, Tokyo Institute of Technology Japan	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows NEC/HP
6	DOE/NNSA/LANL/SNL United States	Cray XE6, Opteron 6136 8C 2.40GHz, Custom Cray Inc.
7	NASA/Ames Research Center/NAS United States	SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon 5570/5670 2.93 Ghz, Infiniband SGI
8	DOE/SC/LBNL/NERSC United States	Cray XE6, Opteron 6172 12C 2.10GHz, Custom Cray Inc.
9	Commissariat à l'Energie Atomique (CEA) France	Bull bullx super-node S6010/S6030 Bull
10	DOE/NNSA/LANL United States	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband IBM

TOP 10 Sites for June 2012

Rank	Site	Computer
1	SA/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM
2	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu
3	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM
4	Leibniz Rechenzentrum Germany	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM
5	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 NUDT
6	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XK6, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA 2090 Cray Inc.
7	CINECA Italy	Fermi - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM
8	Forschungszentrum Juelich (FZJ) Germany	JuQUEEN - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM
9	CEA/TGCC-GENCI France	Curie thin nodes - Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR Bull
10	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 Dawning

名称	国家和地区	场所	安装年份	供应商	处理器核心数	Rmax (Tflops)	Rpeak (Tflops)	功率(千瓦)
天河二号	中国	国家超级计算机中心	2013	中国人民解放军国防科学技术大学	32,000	33860.00	54900.00	24000.00
泰坦	美国	橡树岭国家实验室	2012	克雷公司	18,688	17590.00	27113.00	8200.00
红杉	美国	劳伦斯利福摩尔国家实验室	2011	IBM	1572864	16324.75	20132.66	7890.0
京	日本	理化学研究所	2011	富士通	705024	10510.00	11280.38	12659.9
蓝色基因/L	美国	阿贡国家实验室	2012	IBM	786432	8162.38	10066.33	3945.0
美洲虎	美国	德州高级计算中心	2012	戴尔	462462	5168.1	8520.1	4510
JUQUEEN	德国	Forschungszentrum Jülich	2012	IBM	458752	5008.9	5872	2301
天河一號	美国	DOE NNSA LLNL	2012	IBM	393216	4293.30	5033.20	1972.00
普朗克	德国	莱布尼茨超级计算中心	2012	戴尔	147456	2897	3185.1	3422.67
天河一號	中国	国家超级计算天津中心	2010	中国人民解放军国防科学技术大学	186368	2566.00	4701.00	4040.00

再见，
谢谢！