
*SP***World**99

IBM Servers: Exploiting Industry Leading Technology

Dr. Joel M. Tendler
IBM Server Group
Austin, TX
jtendler@us.ibm.com

Legal Notices

IBM, PowerPC, PowerPC-AS, RS64, POWER3, POWER4, S/390, RS/6000, and AIX are trademarks of the International Business Machines Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Limited.

This information is provide "as is" without warranty of any kind, including the implied warranties of merchantability and fitness for a particular purpose.


All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your IBM local Branch Office or IBM Authorized Reseller for the full text of a specific Statement of General Direction.

IBM may have patents or pending patent applications covering subject matter in this presentation. The furnishing of this presentation does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

System Performance Components

$$\begin{aligned} \text{Job Time} &= \text{Cycle Time} * \text{Cycles/Instruction} * \text{Instructions/Job} \\ &\quad \downarrow \qquad \qquad \downarrow \qquad \qquad \downarrow \\ &\quad \text{Technology} \quad \text{ISA} \quad \text{Software} \\ &\quad \text{Process} \quad \text{System Structure} \quad \text{Application} \\ &\quad \text{Packaging} \quad \text{Microarchitecture} \quad \text{Middleware} \\ &\quad \quad \text{Memory Hierarchy} \quad \text{Libraries} \\ &\quad \quad \text{I/O Subsystem} \quad \text{Operating System} \\ &\quad \quad \quad \text{Compiler} \\ &= (1/\text{Frequency}) * \text{CPI} * \text{Pathlength} \\ &= \text{Pathlength/MIPS} \end{aligned}$$

The Original Moore's Law Proposal



Increased Integration

Function implemented with:


Many silicon components

Few silicon components

One Chip

- Cost ↓
- Speed ↑
- Function ↑
- RAS ↑

Integrated Circuit Performance Trends

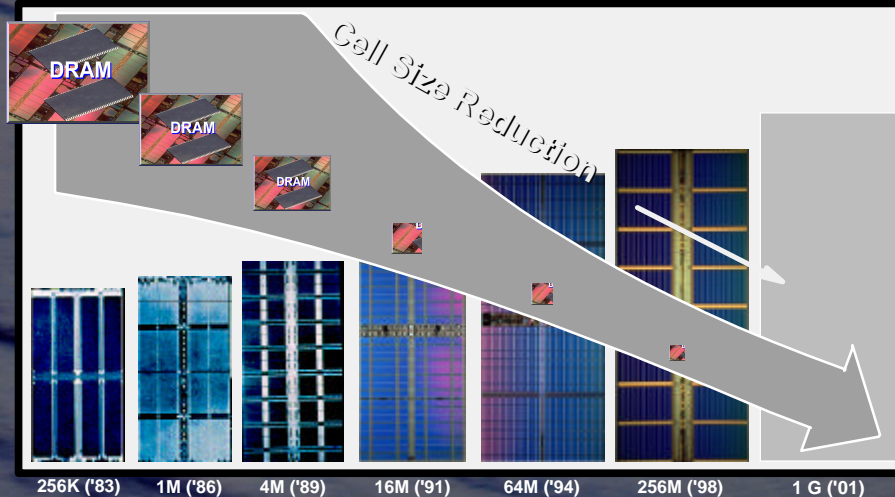


Information Technology Evolution

- Typical luxury car has 20+ microprocessors
 - ▶ More compute power by far than inside the Apollo 11 Lunar Lander
- 2.7 Trillion e-mails worldwide in 1998
 - ▶ Five times more than paper mail
- Typical Laptop is roughly 2X performance of first supercomputer of mid '70s
 - ▶ \$3K compared to several million dollars
- Disk storage: \$100/MB in early '80s to \$0.10/MB today to \$0.02/MB in two years
 - ▶ Similar density improvements would shrink 500 miles of shelf space for 17M books in Library of Congress to ~50 yards

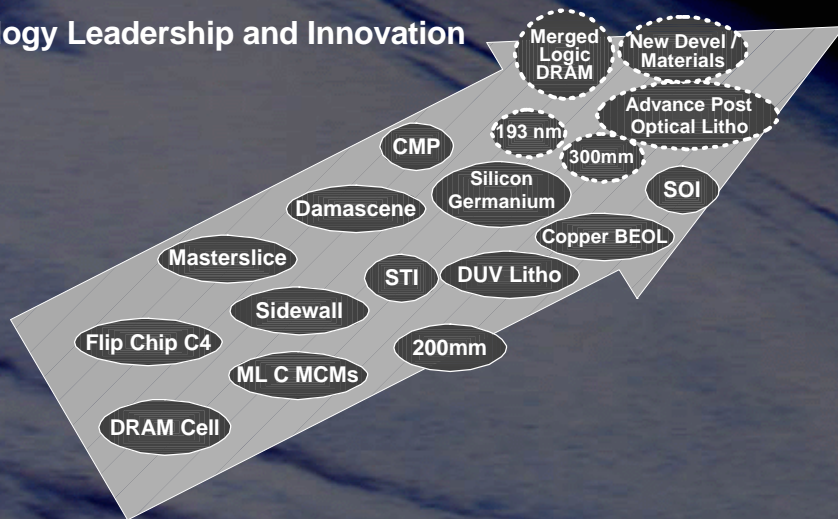
DRAM Evolution

Bits/Chip Increase --- 4X/Gen. (~3 Years)




Continuing Our Heritage:

Technology Leadership and Innovation



S/390 Uniprocessor Performance




S/390: Comparison of Bipolar and CMOS

| | ES9000 9X2 | S/390 G6 |
|---------------------|------------|----------|
| Technology | Bipolar | CMOS |
| Total Chips | 5000 | 31 |
| Total Parts | 6659 | 92 |
| Weight (lbs) | 31,145 | 2057 |
| Power Req (kva) | 153 | 5.5 |
| Chips/processor | 390 | 1 |
| Total Processors | 10 | 12 |
| Maximum Memory (GB) | 10 | 32 |
| Space (sq ft) | 671.6 | 51.9 |

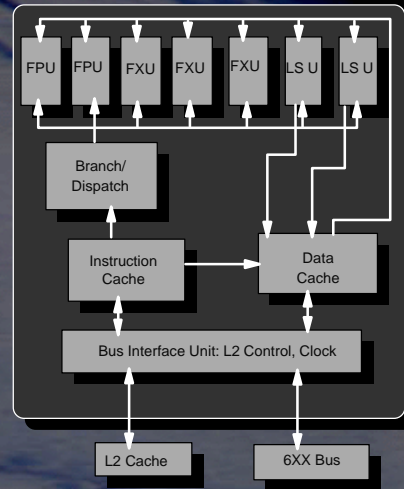
POWER3: 64 bit Technical / Scientific

- Size and power:
 - 15 million transistors
 - Die: 270 mm²
 - Power: 46 W worst case @ 2.5V & 225 MHz
- Third generation superscalar design
 - Up to 8 instructions per cycle
 - Up to 2 operations per f-p unit
- 64-bit SMP
- .25 micron hybrid lithography
- 1088 pin ceramic packaging
- 748 I/O Signals
- Hardware memory prefetch
- 8.0 GB/s Aggregate memory bandwidth

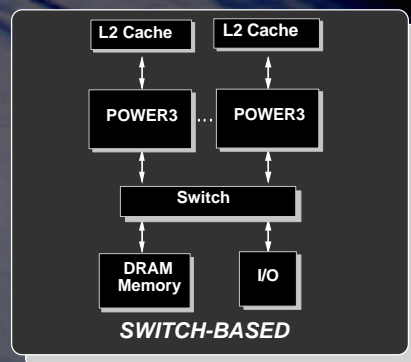
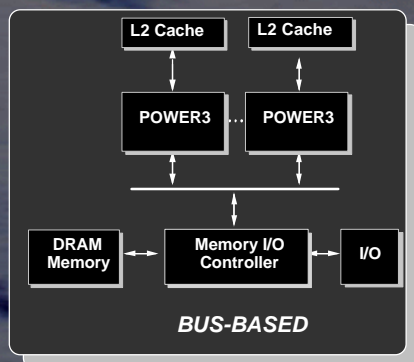


POWER3 - 64 bit Technical/Scientific

- POWER3 Features
 - ▶ Floating point applications, high perf WS, SVRs
 - ▶ FP performance enhanced
 - ▶ CPU clock speeds: 200 to 270 MHz
- Caches
 - ▶ Integrated L1
 - 32KB I, 128W, 128 byte line
 - 64KB D, 128W, 128 byte line
 - ▶ Integrated L2 controller
 - Support for 1 MB to 16 MB 1W
 - 32 bytes / beat
 - 256-bit data, 6.4 GB/s @200 MHz
 - Bus ratios: 1:1, 2:1, and 3:1
- PowerPC 6XX bus
 - ▶ 128-bit data, 1.6 GB/s @ 100 MHz
 - ▶ Bus ratios: 2:1, 3:1, 4:1
- 2.5 V output levels, 3.3 V tolerant on 6XX bus
- Performance @ 200/200/100 MHz
 - ▶ 13.1 SPECint95
 - ▶ 30.1 SPECfp95



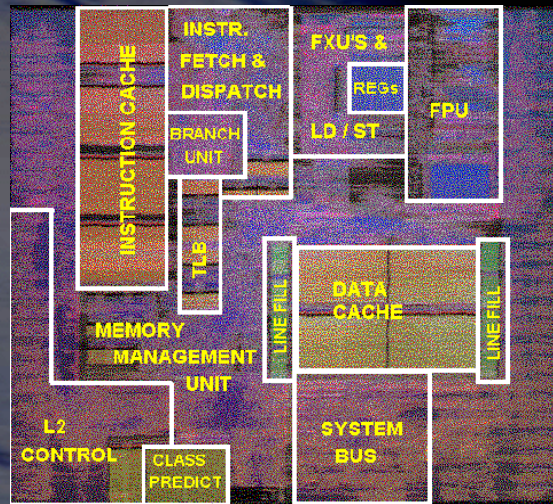
POWER3 System Implementation



• Supports bus- and switch-based MP bus memory configurations

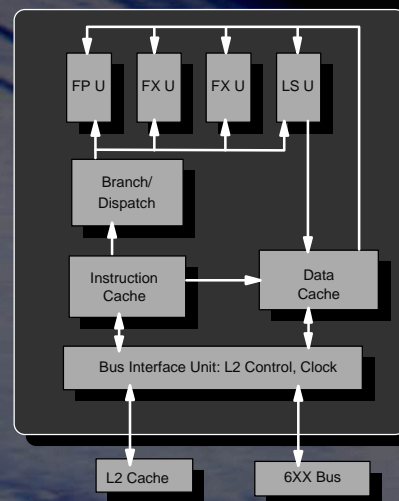
Northstar: 64 bit Commercial

- Size & power
 - 12 million transistors
 - Die: 162 mm²
 - 30W @ 250 MHz
- Technology
 - CMOS 6S2
 - 0.25 micron hybrid lithography
 - Five levels of metal
 - 2.5 volts
- Module
 - 1657 pin ceramic package
 - 985 signal I/O
- Clock distribution
 - Grid distribution for Clock
 - < 100 psec skew at latches




Northstar - 64 bit Commercial

- NorthStar Features
 - Commercial UNIX and transaction processing, workgroup and file servers
- Features
 - Superscalar
 - 4 execution units
 - 4 instructions issued per cycle
 - 64KB instruction + 64KB Data Cache
 - L2 Cache Control
 - 1/4/8 MB using industry standard SRAM's
 - 4 Way Set Associative
- CPU Clock Speeds: 260 - 340 MHz
- Performance @ 340 MHz
 - 16.0 SPECint95
 - 21.2 SPECfp95
- PowerPC 6XX Bus Architecture
 - CPU bus ratio: 2:1, 3:1, 4:1
- Package - 42 MM CCGA
- Technology - 0.25 Im hybrid CMOS6S2




CMOS 7 - World Leadership!

- Advanced Chip Technology
 - New semiconductor manufacturing process
 - First in the industry
- Copper Interconnect for Silicon Wafers
 - Aluminum used during the past 30 years
 - Copper is a superior electrical conductor (40%)
- Higher Performing Microprocessors, SRAM's, ASIC's
 - Up to 200 million transistors
 - Lower voltages (power)




Copper Technology: Enables Better Chip Performance

- Electrical Resistance Increases as Dimensions Get Narrow
- Copper Metallurgy Lowers Resistance
- Allows Designs to Stay on Performance Curve Set By Transistor Scaling / Technology



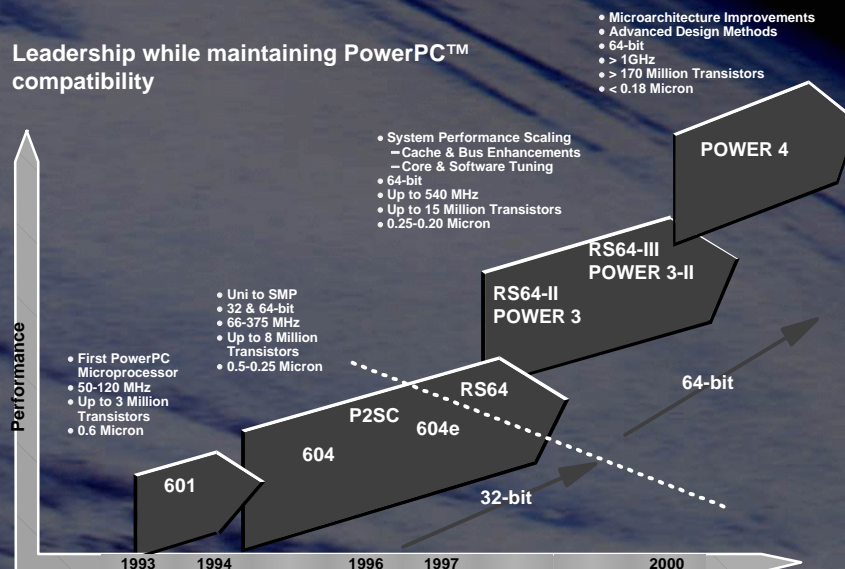
Silicon on Insulator: World Leadership!

Lower Capacitance: More Performance / Less Power




IBM Server & Workstation Microprocessor Roadmap

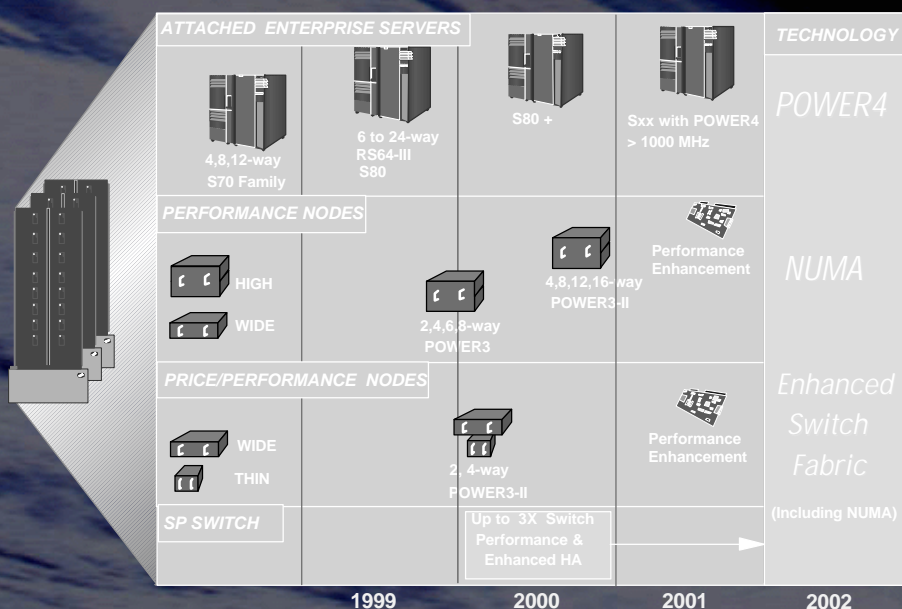
Leadership while maintaining PowerPC™ compatibility



Enterprise Server Roadmap



Large Scale SP Roadmap

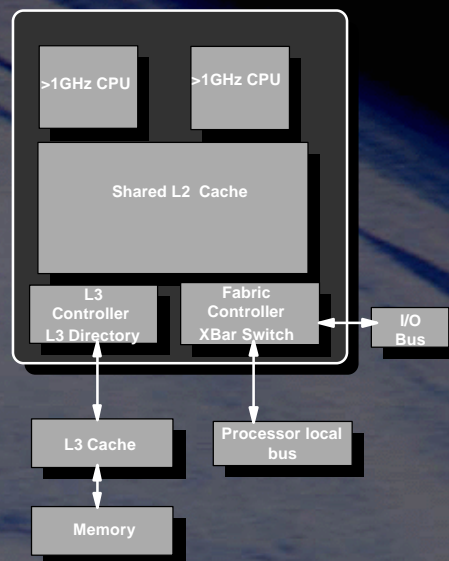


POWER4 Design Objectives

- Very High Frequency Design
 - Leverage IBM's leadership semiconductor technology
 - Full custom circuit design and integration
 - Extended pipeline and layered machine organization
- Commercial Workloads and Technical Workloads
 - Leadership system performance
 - Low latency memory access and very high bandwidth
 - Usable for AS/400, RS/6000 Commercial and RS/6000 Technical Servers
- Binary Compatibility
 - Applications, Middleware, and Operating Systems (OS/400 and AIX)
- SMP Optimization
- Leverage IBM's leadership packaging technology
 - Incorporate coherency optimizations (shared caches, special cache states, etc)
- Scalability and upgradability
- Full System Design Approach
 - Partnerships with compiler, O/S, semiconductor, architecture, and I/O teams
 - Early involvement of Brands


POWER4 Features

- High performance commercial servers
- Leading edge technical servers
- Deep Computing nodes
- CPU clock speeds: >1000 MHz
- 2-Way SMP on a chip
 - Superscalar, out-of-order
 - Innovative cache and memory hierarchy
 - High bandwidth scalable interprocessor fabric
- Technology:
 - 0.18 micron lithography
 - CMOS-8S2
 - Copper, SOI
 - 7 Layers of Metal



POWER4 Test Chip

- ★ Operating at 1+ GHz
- 35 Million transistors
- 0.18 micron lithography
- 7 layers of metal, Copper, SOI
- >2200 chip I/Os, >5500 total connections



Powered
On
12/17/98

System Performance Components

Job Time

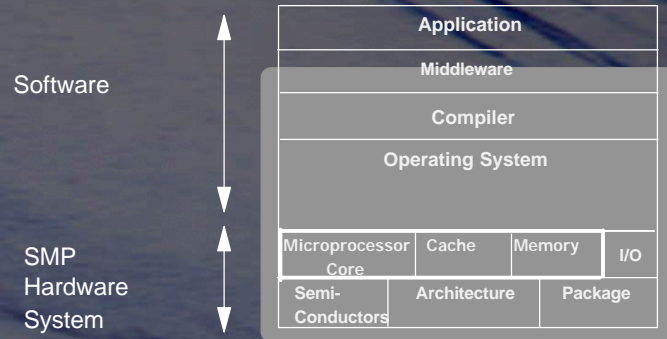
= Cycle Time * Cycles/Instruction * Instructions/Job

↓
Technology
Process
Packaging

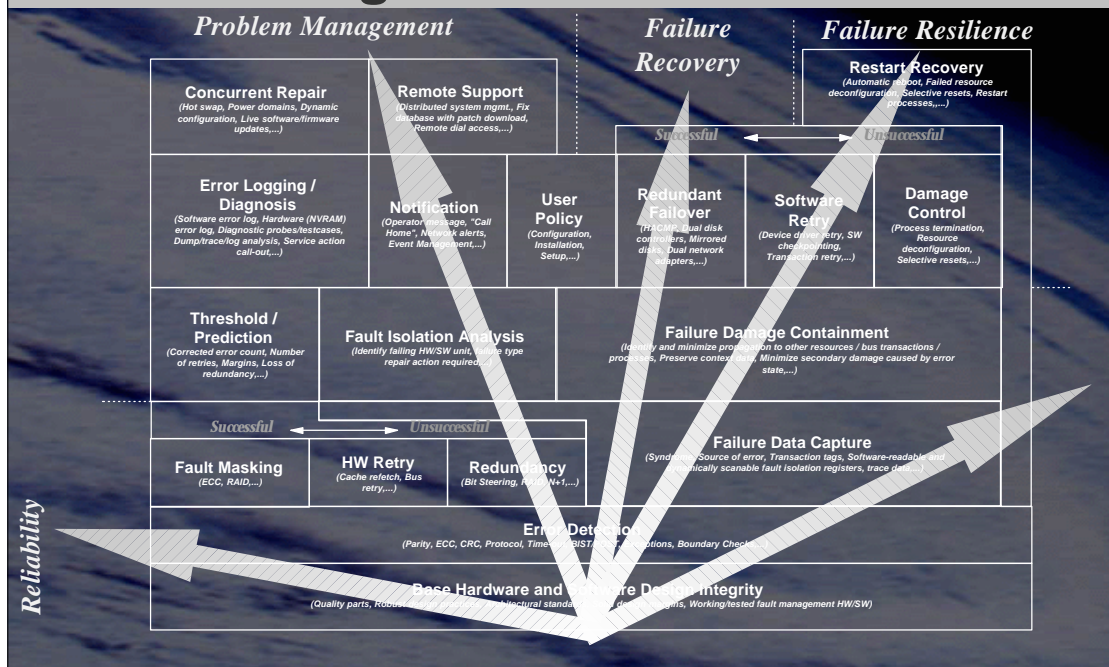
↓
ISA
System Structure
Microarchitecture
Memory Hierarchy
I/O Subsystem

↓
Software
Application
Middleware
Libraries
Operating System
Compiler

System Stack



RAS Building Blocks



IBM Servers --- Engines of e-business

- Technology driving server performance
- IBM continues to be industry leader in introducing and leveraging new technology
 - ▶ Hardware technology
 - ▶ Software technology
 - ▶ System technology
- IBM is only company in industry involved in all aspects of computer systems

