

Design Objectives of the 0.35 μ m Alpha 21164 Microprocessor

(A 500MHz Quad Issue RISC Microprocessor)

Gregg Bouchard

Digital Semiconductor

Digital Equipment Corporation

Hudson, MA




Outline


- 0.35 μm Alpha 21164 Overview
- Design Goals
- Technology Issues
- 21164 Internal Architecture Review
- Architectural Enhancements
- System Performance Enhancements
- Alpha Processor RoadMap
- Summary

0.35 μ m ALPHA 21164

- Technology shrink of 0.5 μ m design
 - + Architectural Enhancements
 - + System Performance Enhancements



0.5 μ m process



0.35 μ m process

Transistor Count	9.3 Million	9.66 Million
Die Size	16.5 mm x 18.1 mm	14.4 mm x 14.5 mm
Power Supply	3.3V external 3.3V internal	3.3V external 2.5V internal
WC Power Dissipation	50W @ 300 MHz	37W @ 433 MHz
Target Cycle Time	300 MHz	433 MHz


Design Goals

- **Reduced Cost**
 - Die size reduction
 - Remove processing steps
- **Higher Performance**
 - 433 MHz design target
 - Architectural enhancements
- **Reduced Power**
 - Lower Core Operating Voltage (2.0v-2.5v)
- **Time to Market**
 - Significant leverage from previous design


Die Size Analysis

	<u>X-dim</u>	<u>Y-dim</u>	<u>NormA</u>
Original	16.5	18.1	1.00
25% shrink	- 4.1	- 4.5	
	12.4	13.6	0.56
Conversion	+0.0	+0.8	
	12.6	14.5	0.62
LI Removal	+1.8	+0.0	
	14.4	14.5	0.70
<hr/>			
Actual	14.4	14.5	0.70
<hr/>			
Ideal 30%	11.5	12.7	0.49


Original Die
16.5mm x 18.1mm
298 mm²



Actual shrink
14.4mm x 14.5mm
209 mm²




Ideal 30% shrink
11.5mm x 12.7mm
146 mm²




Layout Conversion Strategy

- **Full 30% linear shrink was not possible**
- **Solution:**
 - 25% linear shrink
 - Semi-automated conversion of design rules
 - Polysilicon mask layer pushed to full 30% shrink dimensions for performance
- **Redesign of caches**
 - Local Interconnect removed


Layout Conversion Example



Original



After Script



Final

Speed


- **Single-wire, two phase clocking scheme**
- **14 gates per cycle including latches**
- **Single global clock grid**
 - Global clock skew < 90ps
 - Local clock skew < 25ps
- **Clock statistics (0.5 μm design)**
 - Clock load = 3.75 nF
 - Size of final clock inverter = 58 cm
 - Edge rate = 0.5 ns
 - Clocking consumes 40% of chip power
 - $di/dt = 50 \text{ A}$

Speed Verification

- Test circuits were used to determine the speed scaling of different circuit configurations
 - Predicted average process speed up
 - Identified “slow” circuit types
- New circuits were evaluated in SPICE
- Chip sections with major modifications were completely re-verified

Power


- Significant Power Reduction
- $V_{ddi} = 2.2\text{v}$ (to 2.5v)
- 3.3V only interface to external devices



Pre-CLK Drivers

Main CLK Drivers

Pad Ring (I/O)

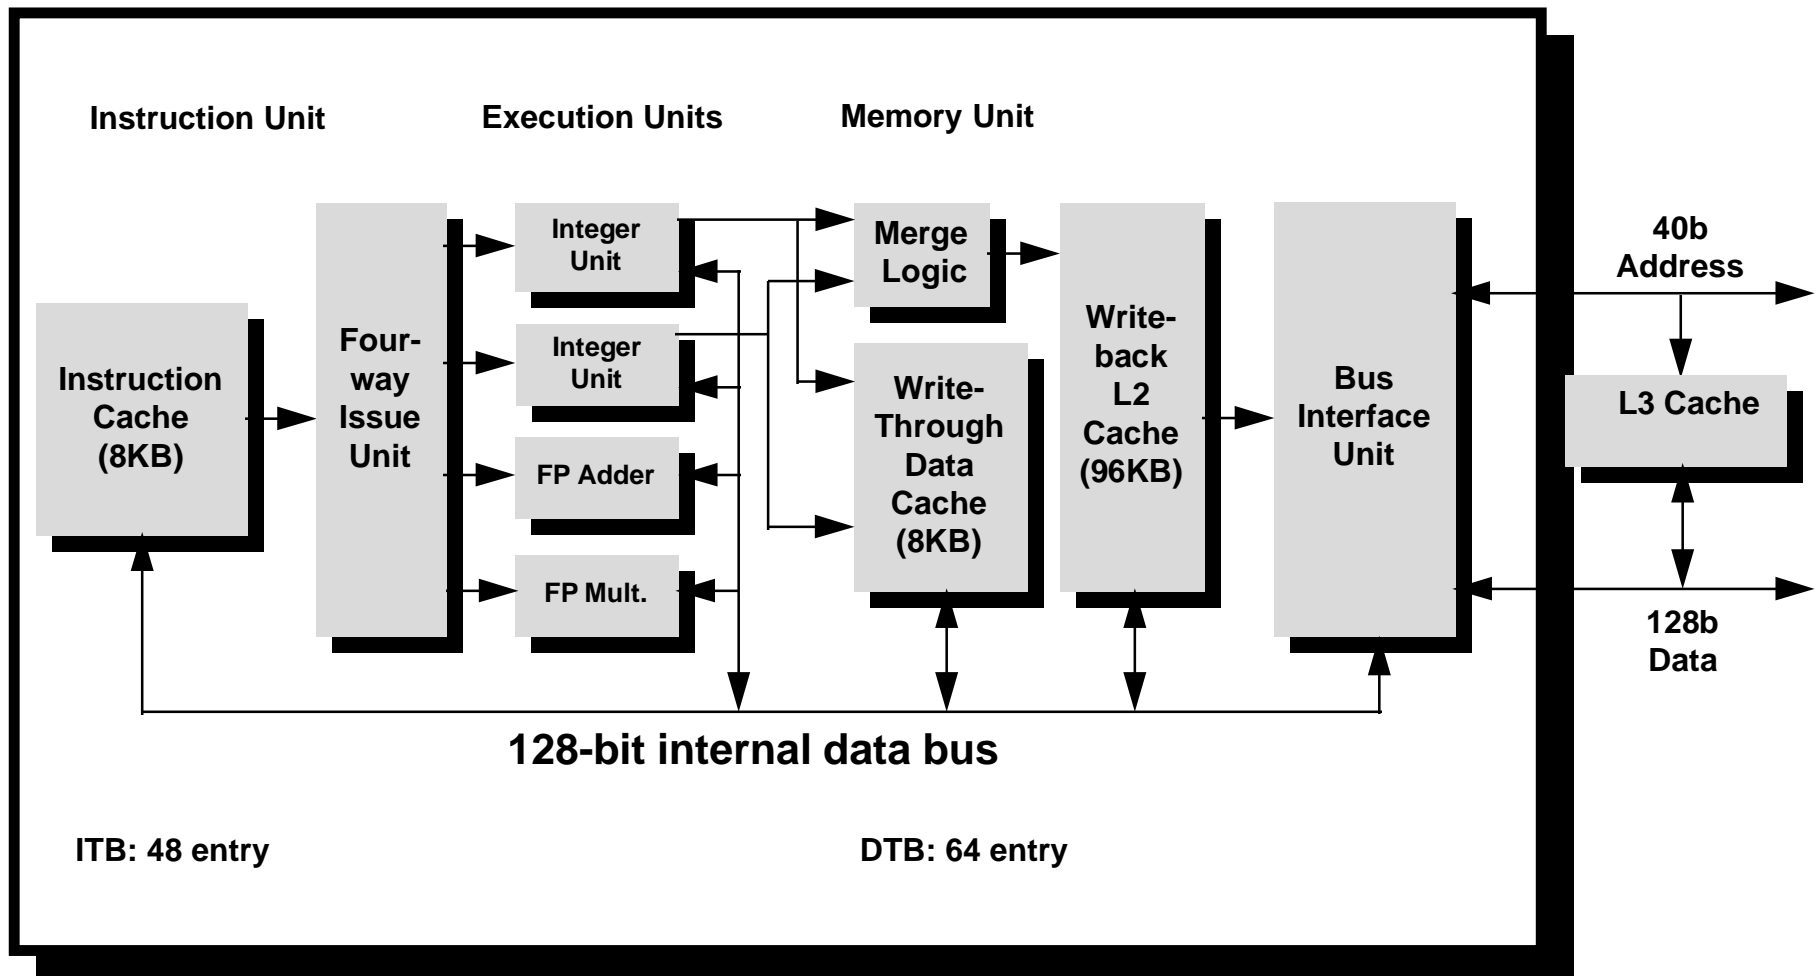


Alpha 21164 Features

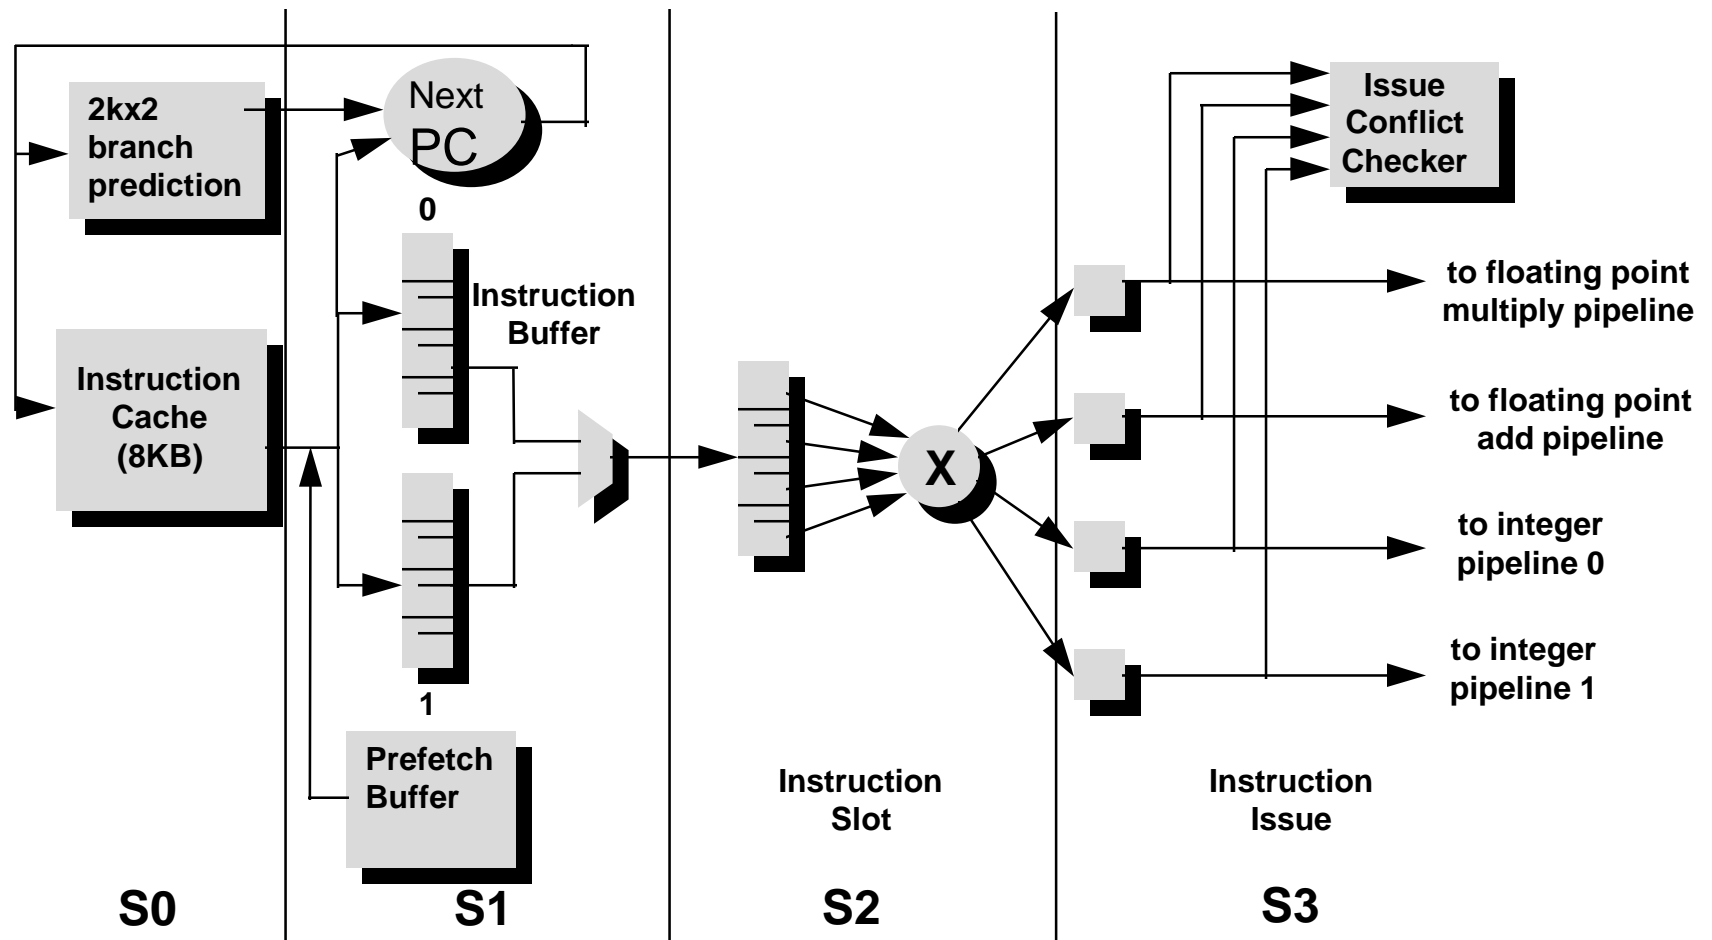
Key Attributes

- **4-way issue superscalar**
 - Up to 2 Integer AND 2 Floating Point instructions issued per CPU cycle.
- **Large on-chip L2 cache**
 - 96KB, writeback, 3-way set associative
- **Fully Pipelined**
 - 7-stage integer pipeline
 - 9-stage floating point pipeline
- **Emphasis on low latency at high clock rate**
- **High-throughput memory subsystem**

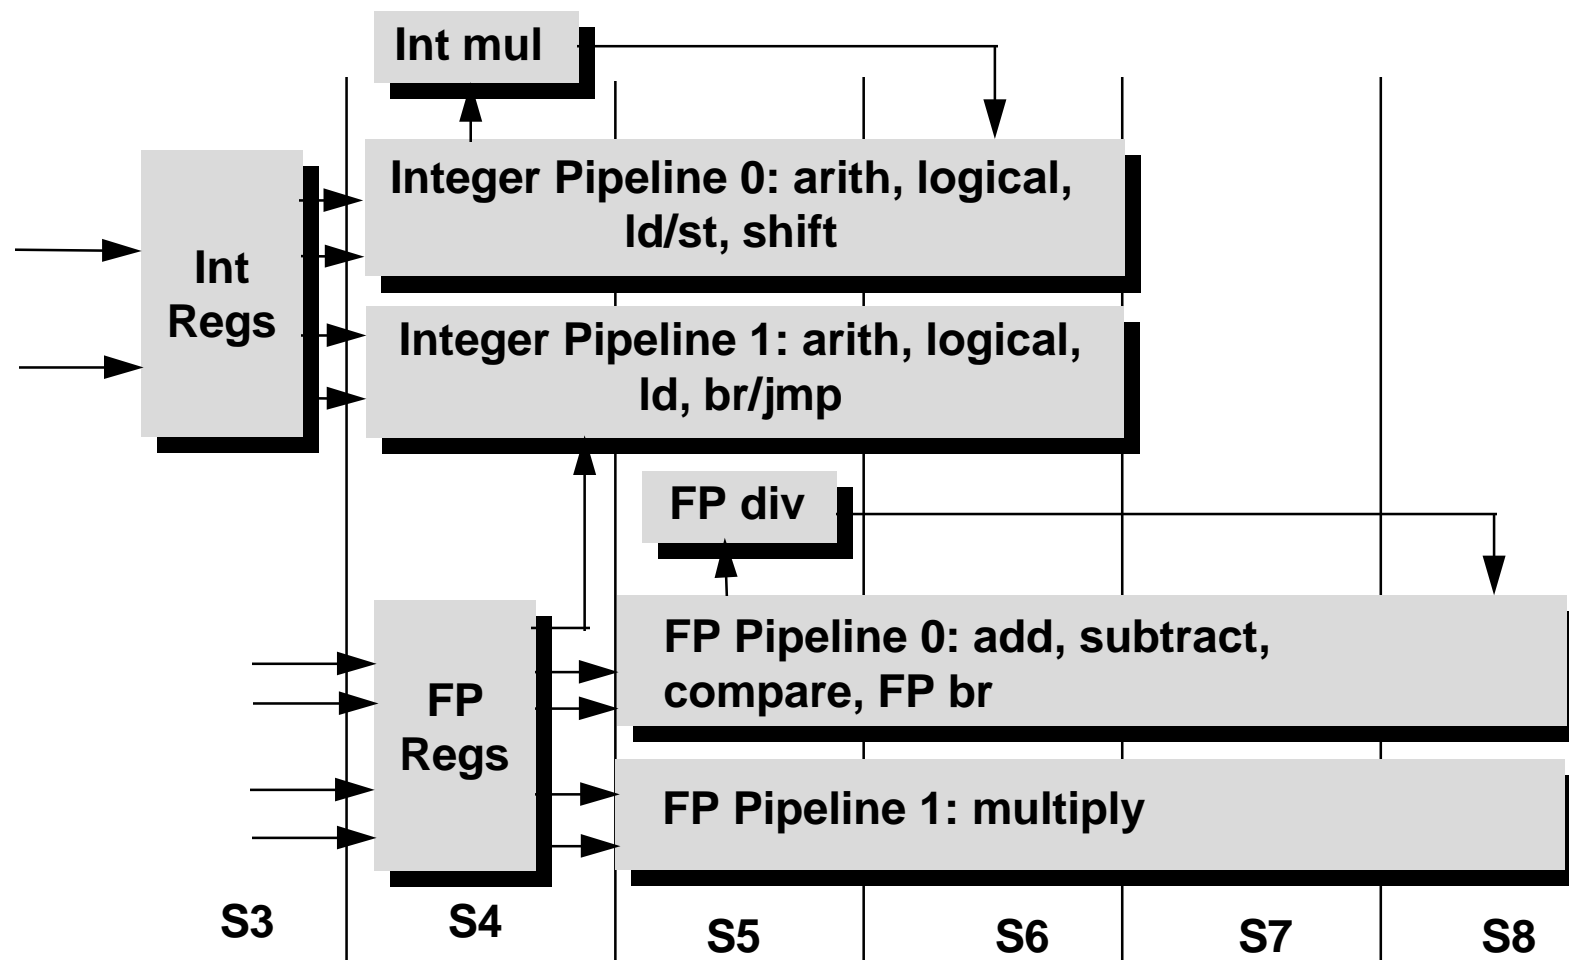
Alpha 21164 Block Diagram



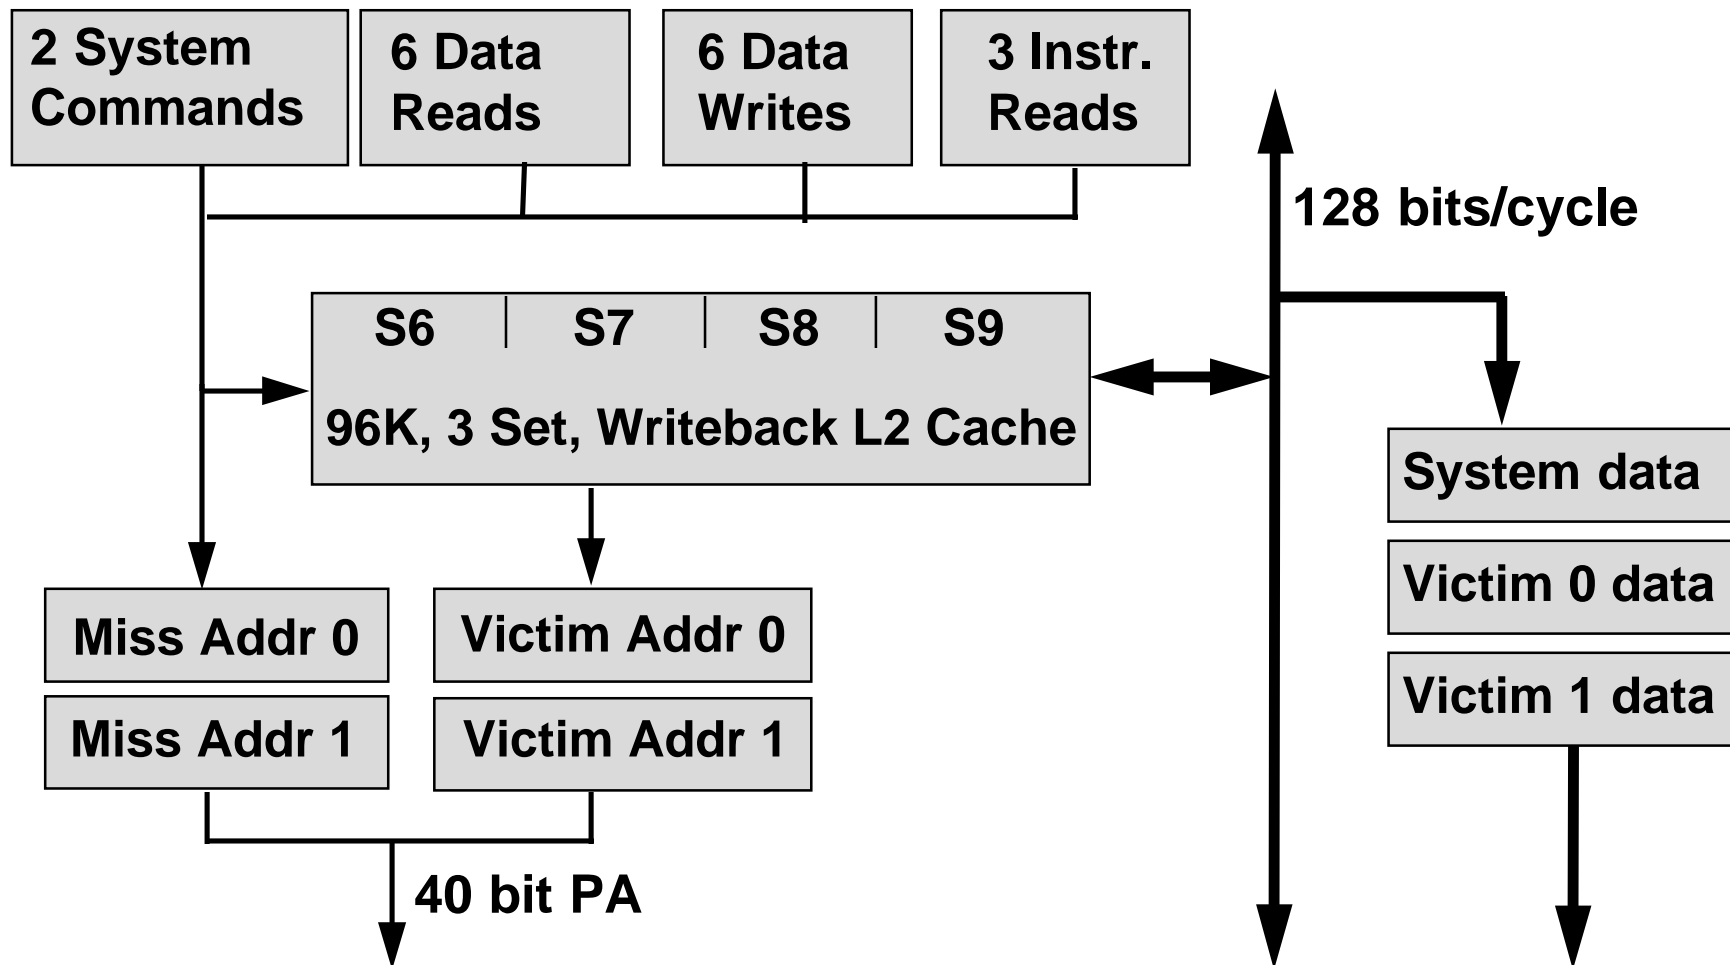
Instruction Issue Pipeline Review



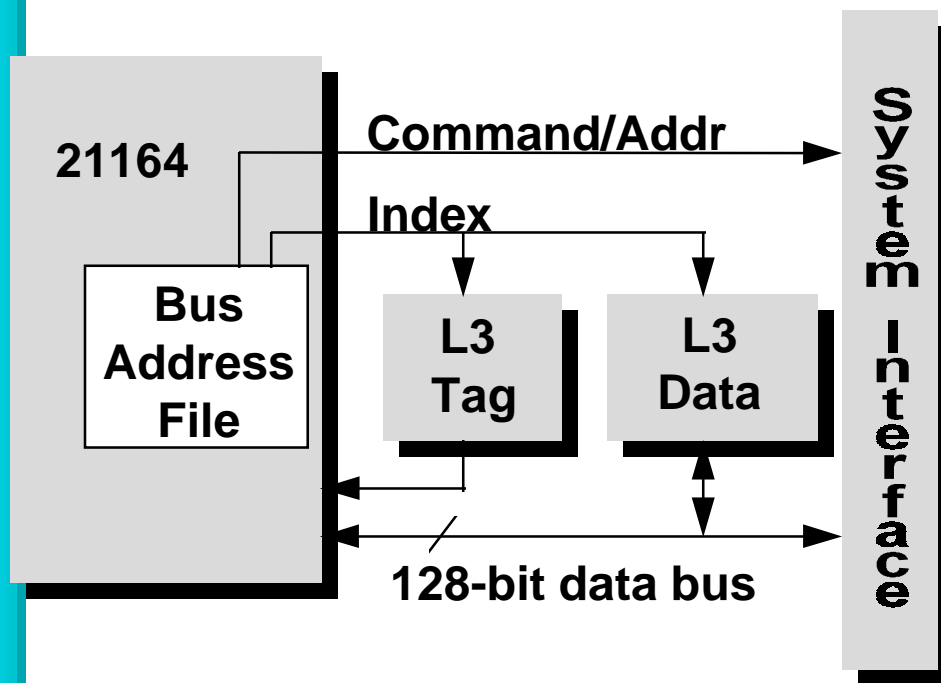
Execution Pipeline Review



On-chip Cache Resources Review



L3 Cache (off-chip)



- L3 cache is a direct-mapped writeback superset of on-chip L2 cache
- Up to 2 reads (or outstanding read commands) in L3 cache
- Programmable wave pipelining for L3 cache
- Support for Synchronous Flow-Thru SRAMs
- L3 cache is optional

Off-Chip L3 Cache Options

Selectable via on-chip programmable registers

- ◆ **Cache Size** - 1 to 64M Byte
- ◆ **Cache Read/Write Speed** - 4 to 15 cpu cycles
- ◆ **Read to Write Spacing** - 1 to 7 cpu cycles
- ◆ **Write Pulse (Bit Mask)** - Up to 9 cpu cycles
- ◆ **Wave pipelining** - 0 to 3 cpu cycles
- ◆ **Support for Synchronous SRAM's**

Architectural Enhancements

New Instructions

- Scalar support for Byte and Short data types
 - LDBU, LDWU load an unsigned byte or short
 - STB, STW, store an byte or short
 - SEXTB, SEXTW sign extend a byte or short
- Eases porting of device drivers to Alpha
- Improves emulation of Intel code on Alpha
- Implemented in this and all future Alpha microprocessors

Architectural Enhancements

New Instructions (continued)

- IMPLVER
 - returns a small integer indicating the core design
 - used for code scheduling decisions
 - implemented in all Alpha microprocessors
- AMASK
 - clears bits to indicate which features are present
 - implemented in all Alpha microprocessors

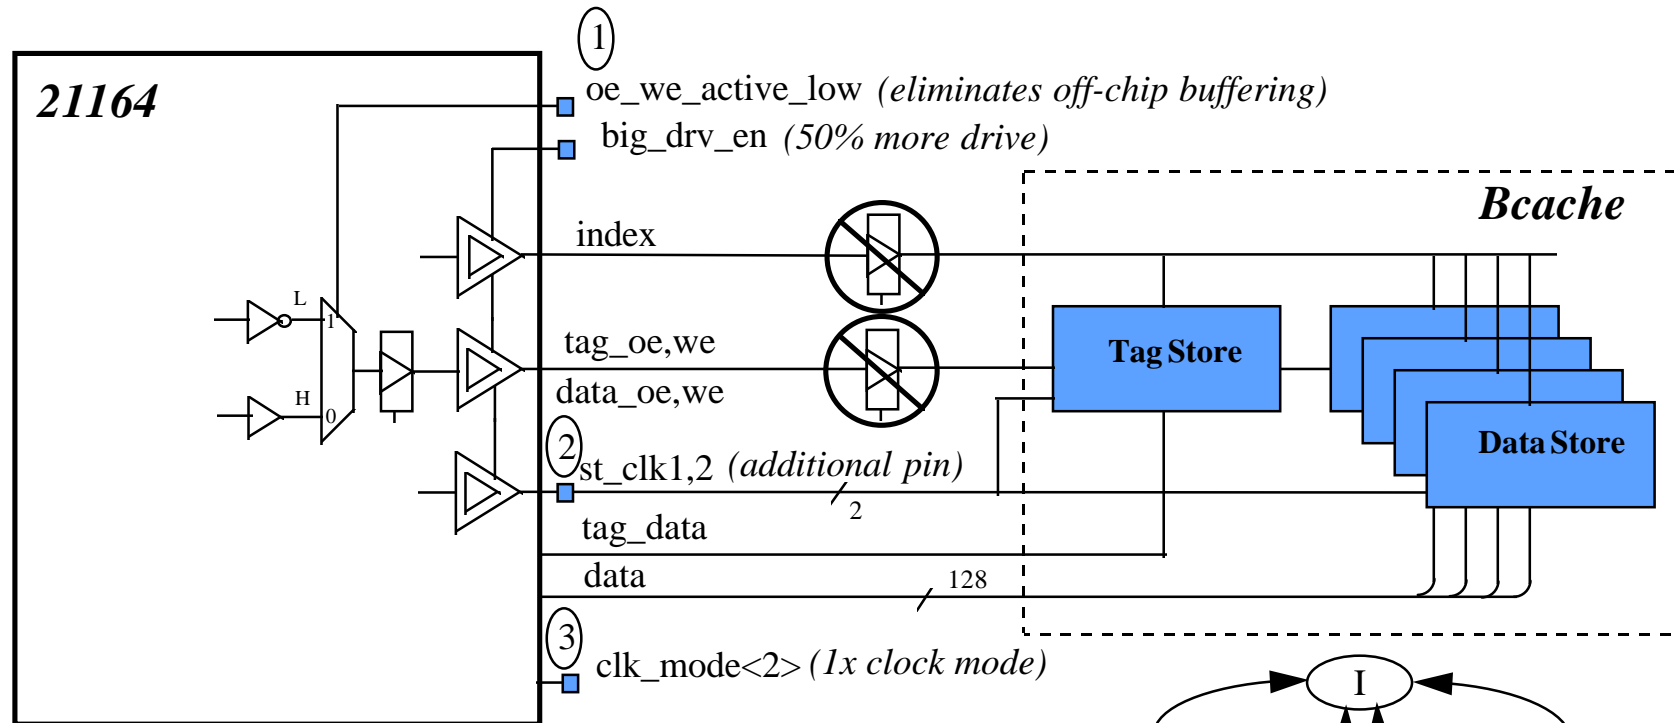
Example:

AMASK #1, R0 ;byte/word present

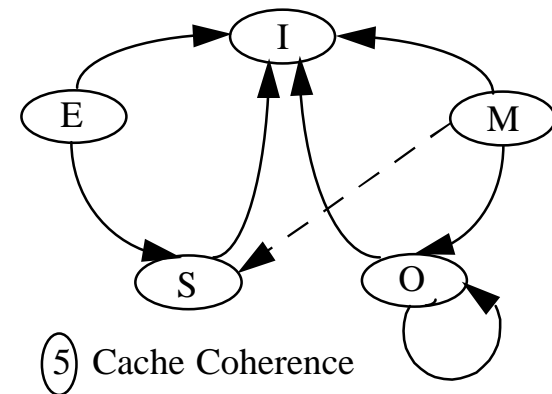

BNE R0, emulate ;if not emulate

...

System Performance Enhancements



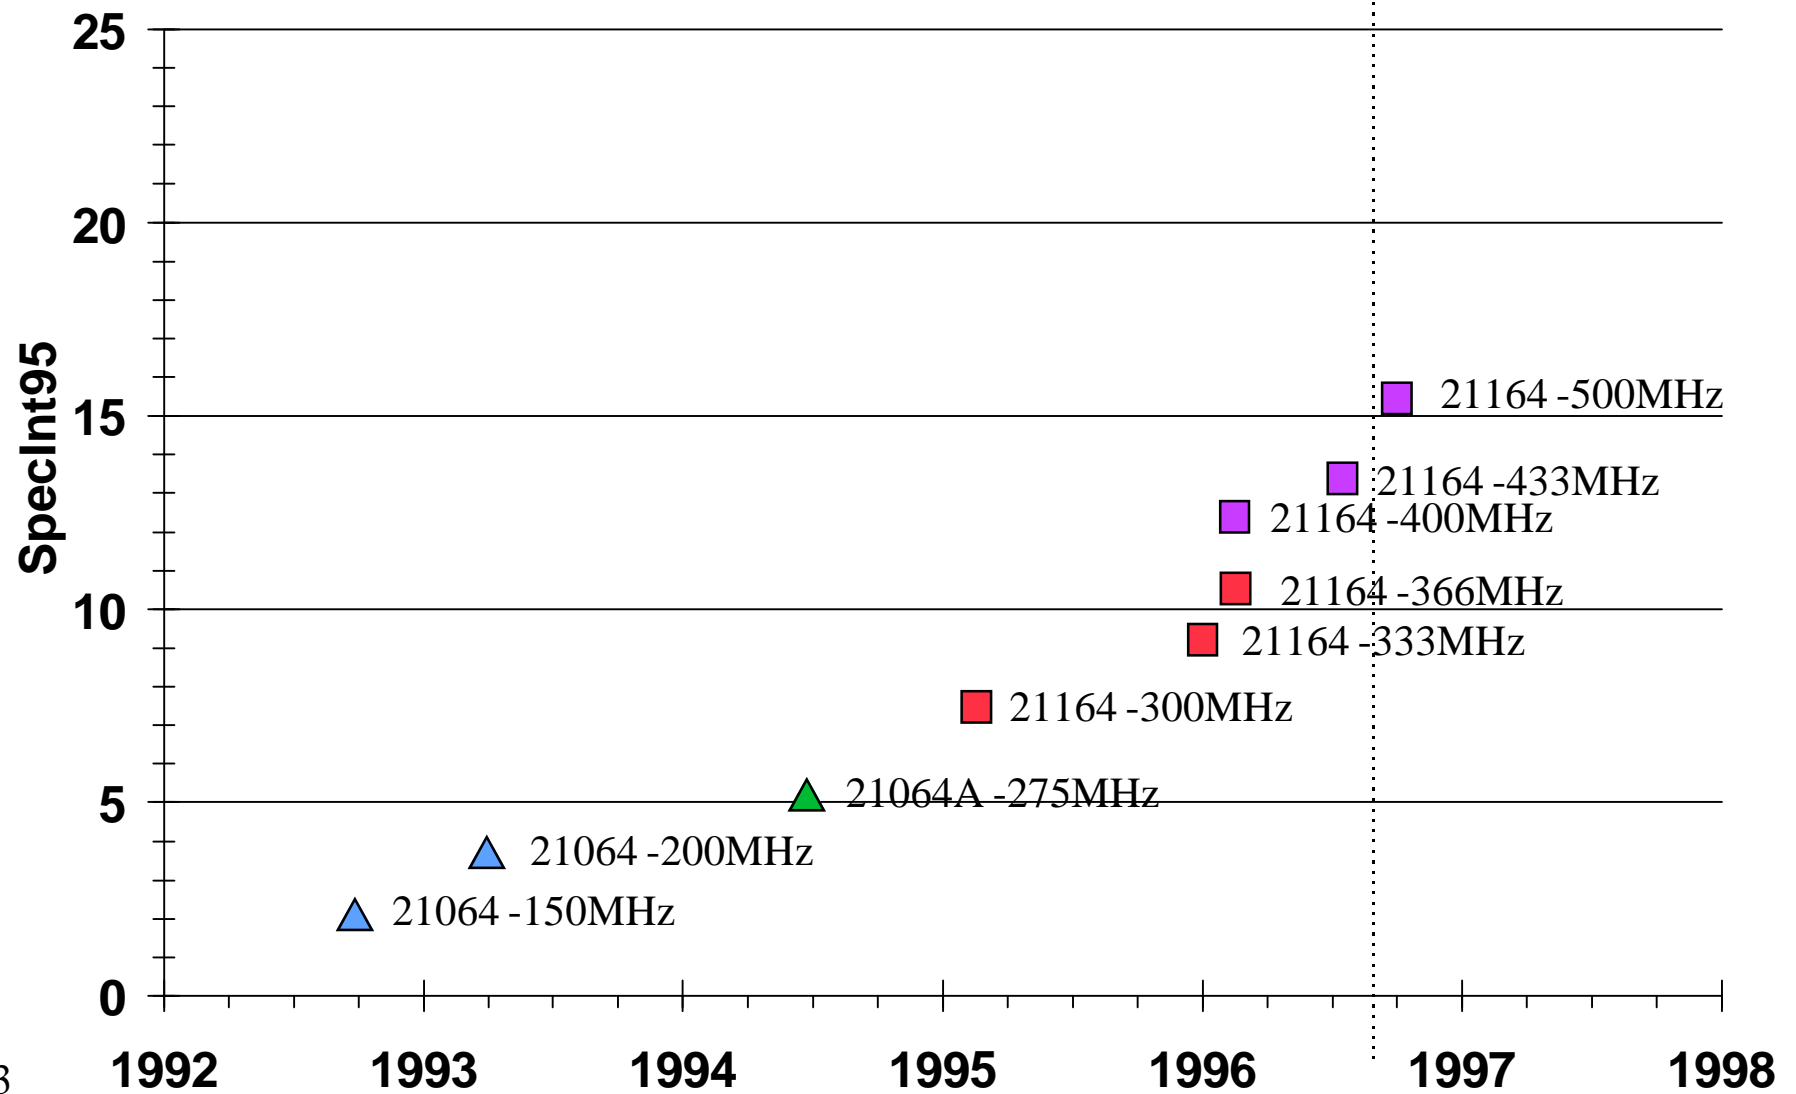
④ Bus utilization improvements (Auto Dack)



Pre-Silicon Logic Verification

- Used extensive test suite developed for original chip as testing baseline
- Random and focused testing
- Coverage analysis to ensure excellent test coverage
- Three simulation systems used:
 - RTL
 - Transistor-level
 - Gate-level

Alpha Microprocessor Road Map



Summary

- FIRST PASS SILICON - October 1995
- Booted first operating system in 2 days!
- Continued Performance Leadership (4+ years)
- Look for next generation 30+ SpecInt95 by Q3'97

Acknowledgments

Peter J. Bannon, Michael S. Bertone, Randel P. Blake Campos, William J. Bowhill, David A. Carlson, Ruben W. Castelino, Dale R. Donchin, Richard M. Fromm, Mary K. Gowan, Paul E. Gronowski, Anil K. Jain, Bruce J. Loughlin, Shekhar Mehta, Jeanne E. Meyer, Robert O. Mueller, Andy Olesin, Tung N. Pham, Ronald P. Preston, Paul I. Rubinfeld