The Alpha 21264 – Data Stream

Matt Ziegler

Alpha 21264 – Pipeline Stages 4-6 [1]

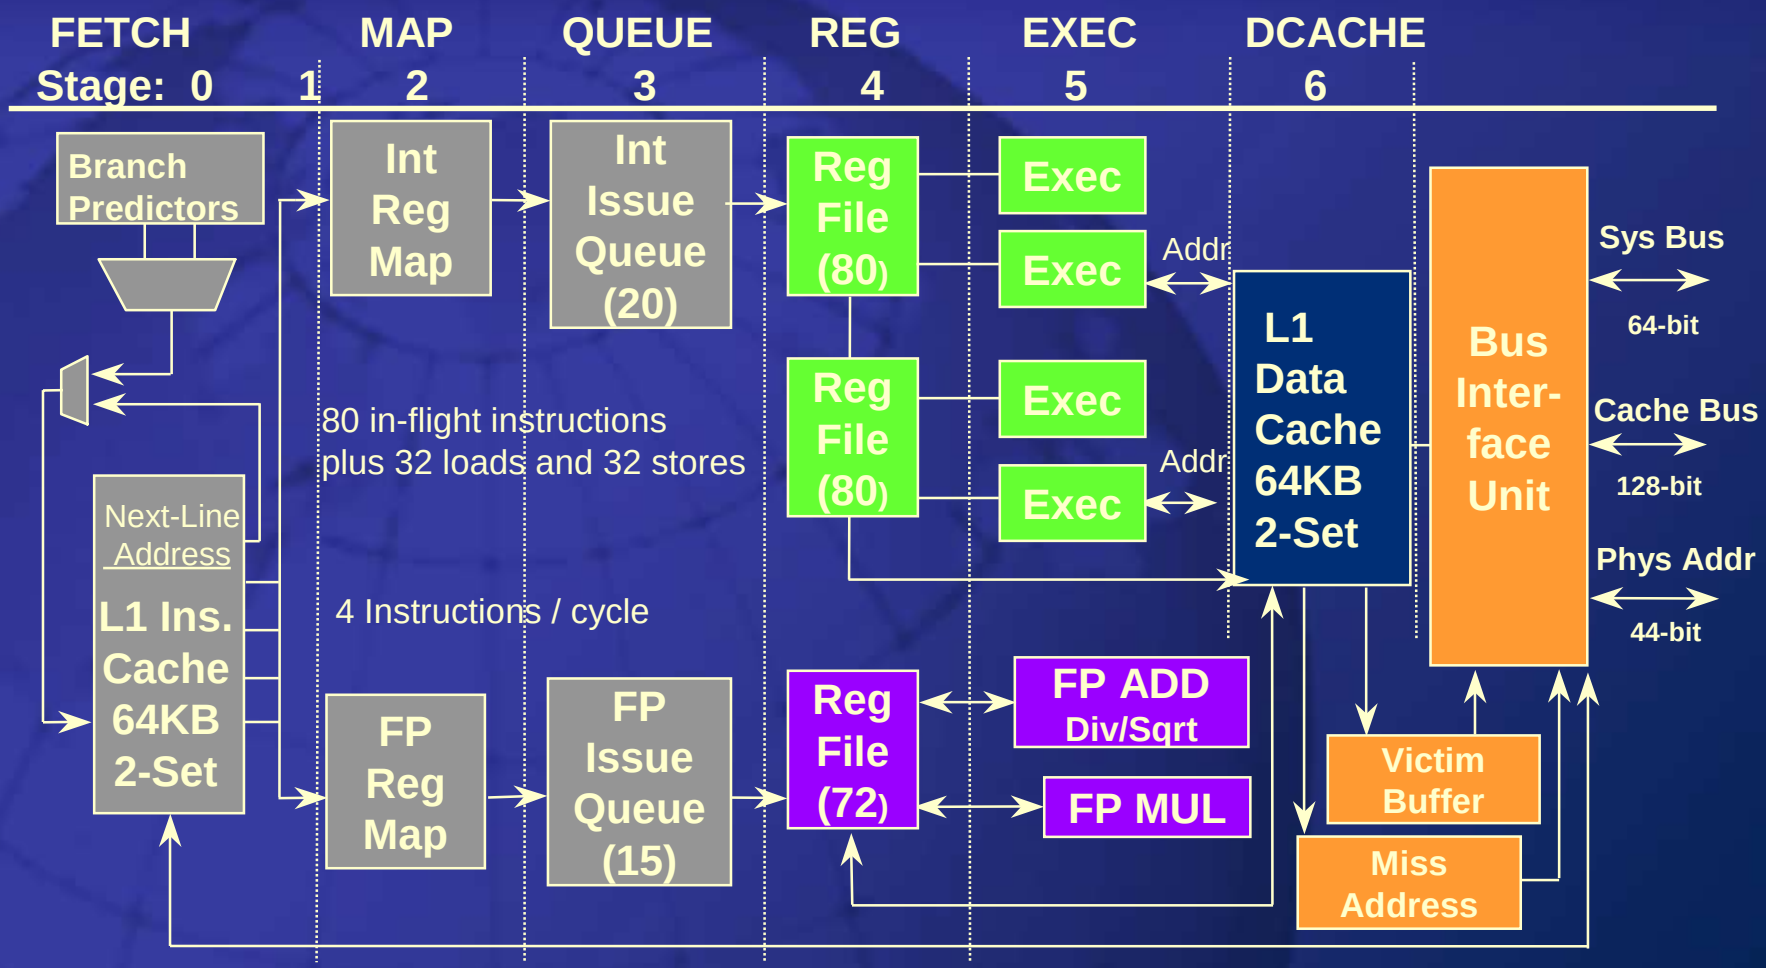

COMPAQ

Alpha 21264 – Block Diagram

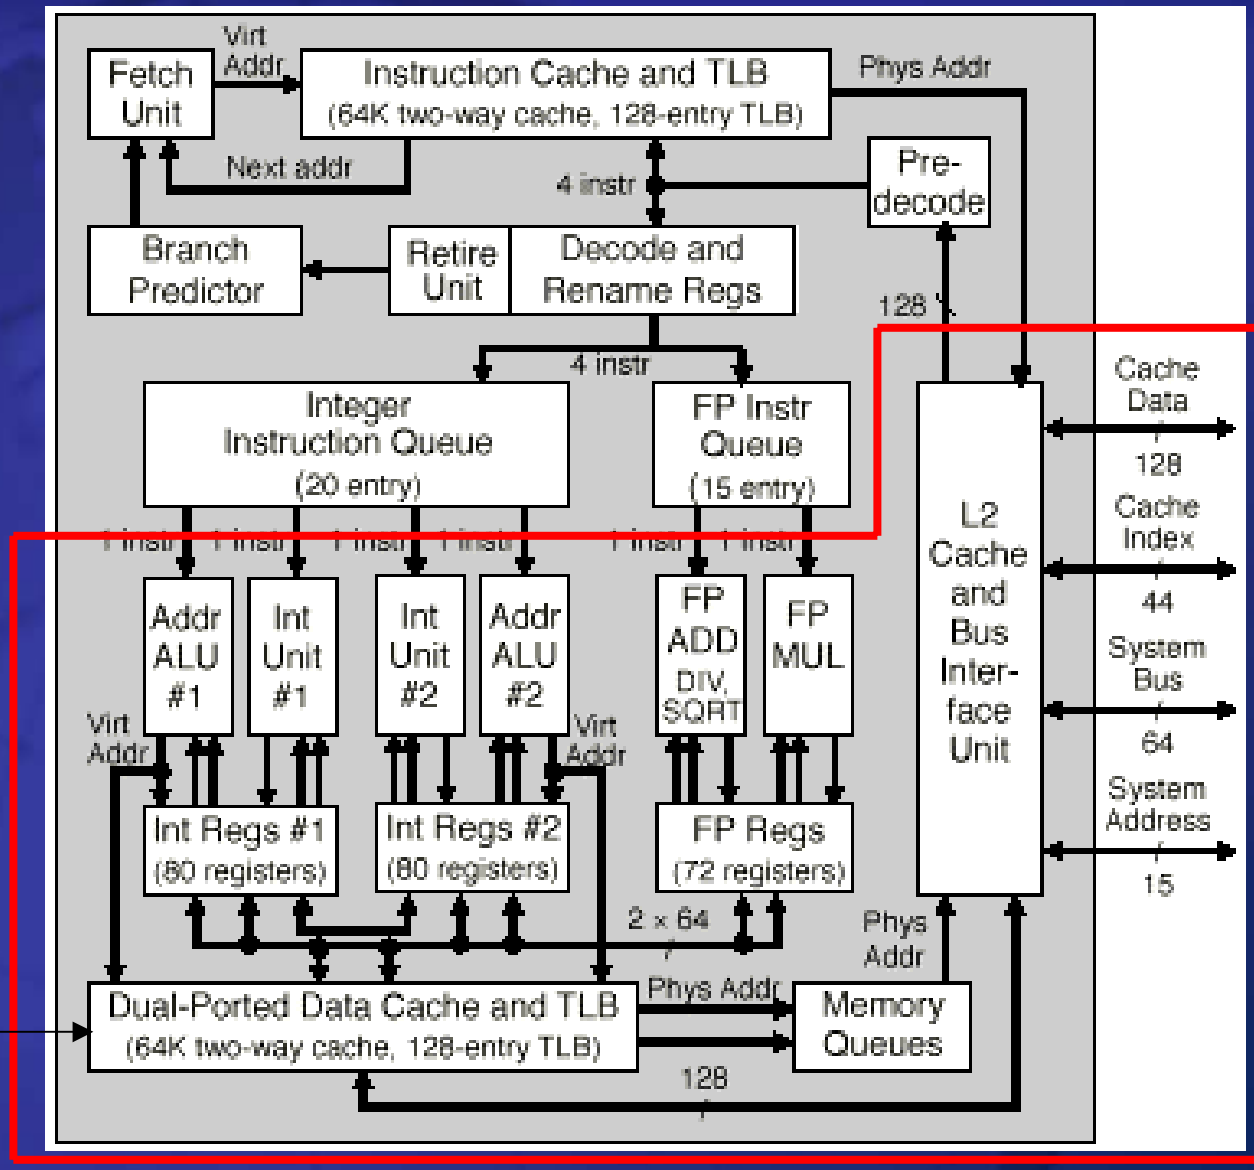

Data Stream Overview [1]

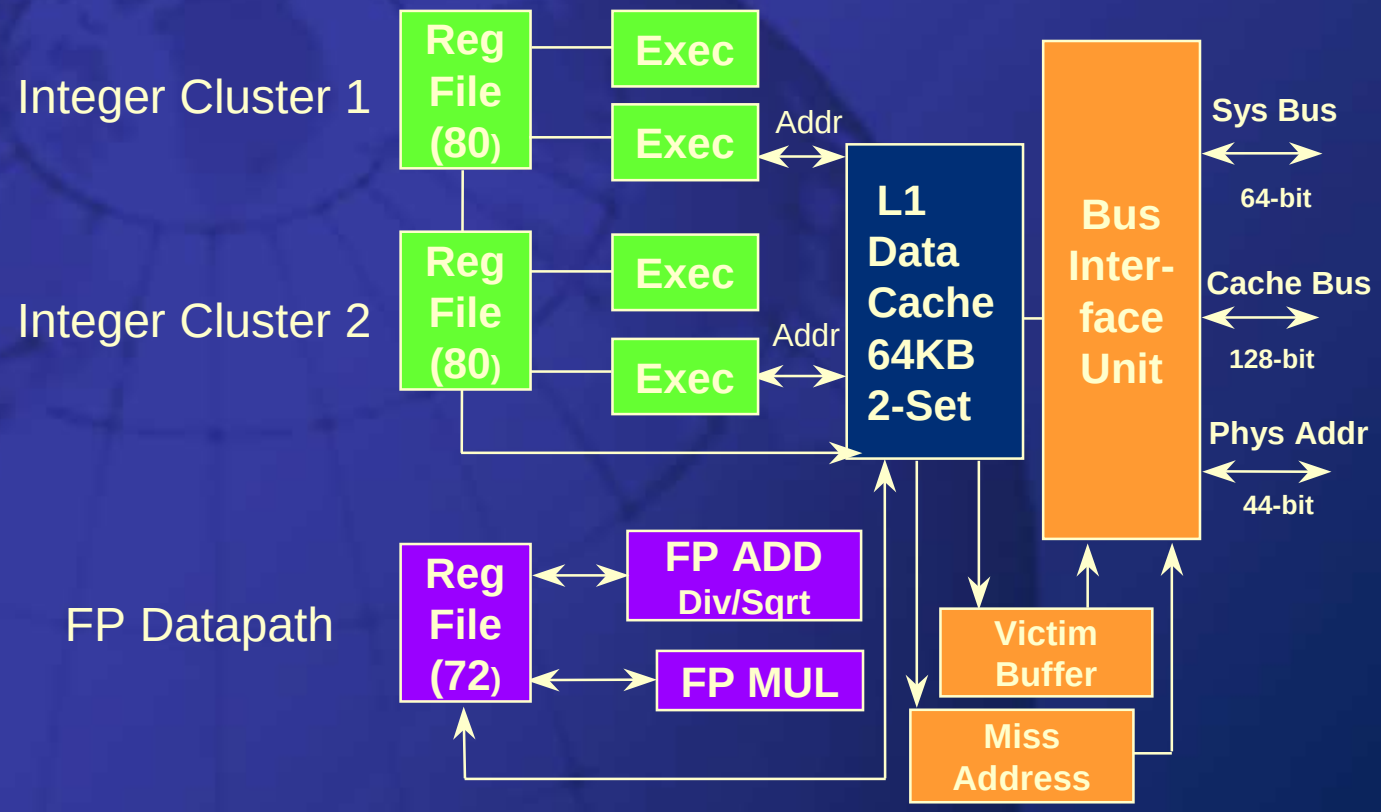

COMPAQ

Register and Execute Stages [1]

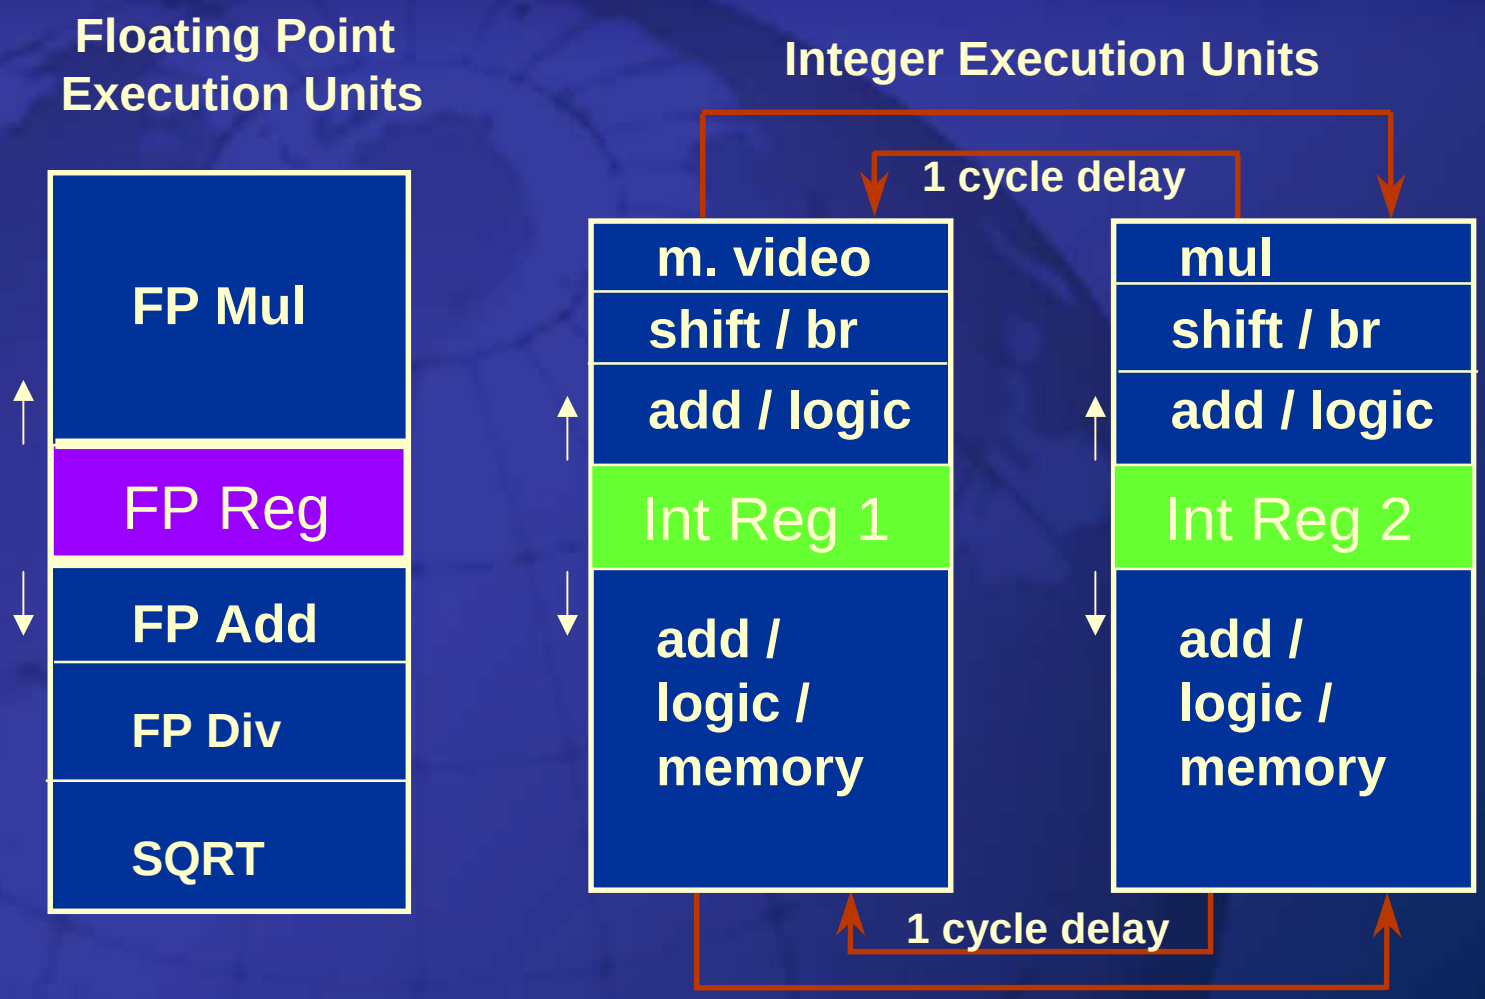

COMPAQ

Alpha 21264 Floor Plan – The 6 Datapaths [1]

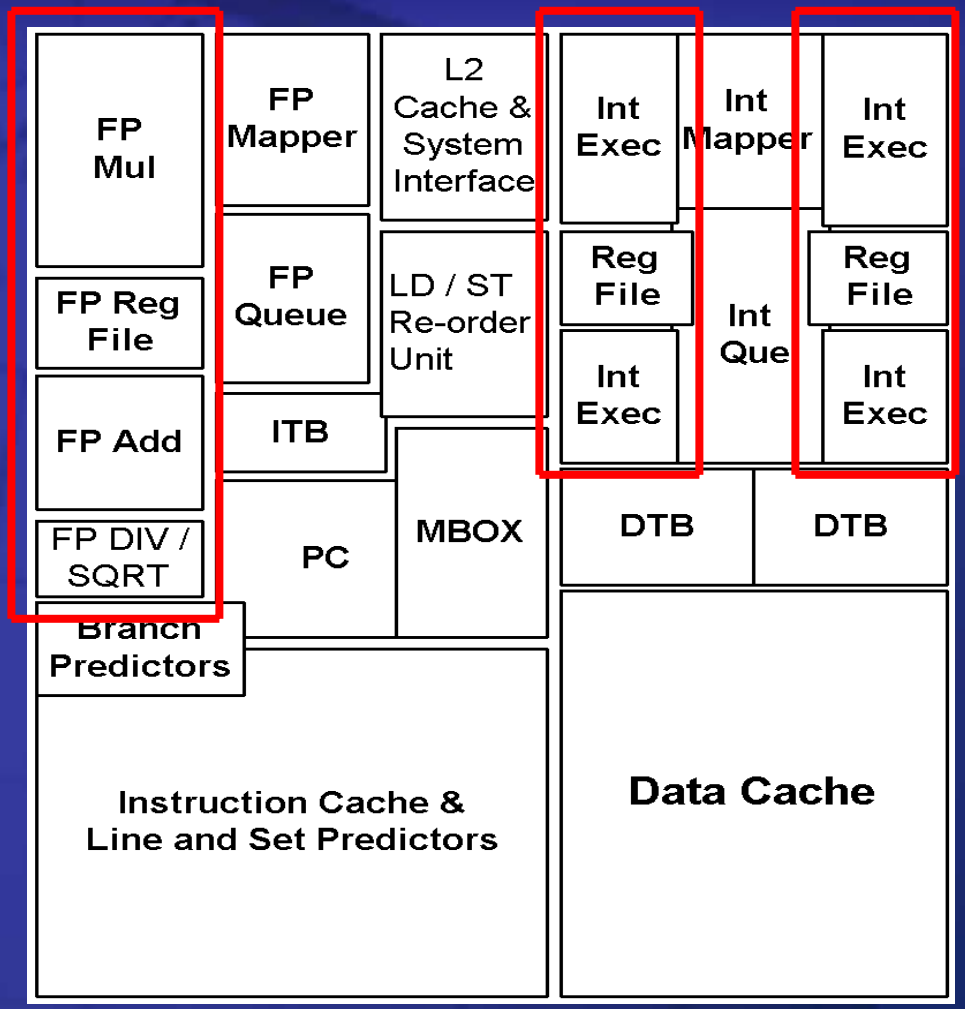

COMPAQ

Integer Datapath

- 2 Integer Clusters, 2 Pipes per Cluster → 4 Pipes
  - Each Cluster has a copy of the Register File

- Cluster 1
  - Upper Pipe
    - MVI/PLZ
    - Shifter/Branch
    - Add/Logic
  - Lower Pipe
    - Add/Logical
    - Load/Store

- Cluster 2
  - Upper Pipe
    - Integer Multiplier
    - Shifter/Branch
    - Add/Logic
  - Lower Pipe
    - Add/Logical
    - Load/Store

COMPAQ

Floating Point Datapath

- Hardware support for the IEEE FP standard
  - NaN, Infinity processing, Denormals, etc.
- 2 Pipes
  - Upper Pipe
    - FP Multiply
  - Lower Pipe
    - FP Add
    - FP Divide
    - FP Square Root

COMPAQ

Instruction Latencies

- Simple Integer Ops 1
- MVI / PLZ 3
- Int Multiply 7
- Int Load 3
- FP Load 4
- FP Add 4
- FP Multiply 4
- FP Divide 12 s-p, 15 d-p
- FP Square Root 15 s-p, 30 d-p

COMPAQ

Integer and FP Register Files

- Int Reg File – 31 Visible, 80 total
- Two Integer Processing Clusters
  - Two pipes in each cluster
  - Each cluster has its own copy of Int Reg File
  - Reduces the number of access ports from 8 read / 6 write to 4 read / 6 write
- FP Reg File – 31 Visible, 72 total

COMPAQ

Alpha 21264 Loads and Stores [1]

- L1 Data Cache
  - Two loads / stores per cycle (any combination)
    - 2x clock Frequency: Phase pipelined - no bank conflict
    - 16 Byte read / write per cycle
- Loads and stores issue out-of-order
  - 32-entry load and store reorder buffers
  - Memory references check buffers to enforce ordering
  - Uncommitted stores forward data to loads

COMPAQ

Alpha 21264 Floor Plan – Data Cache [1]

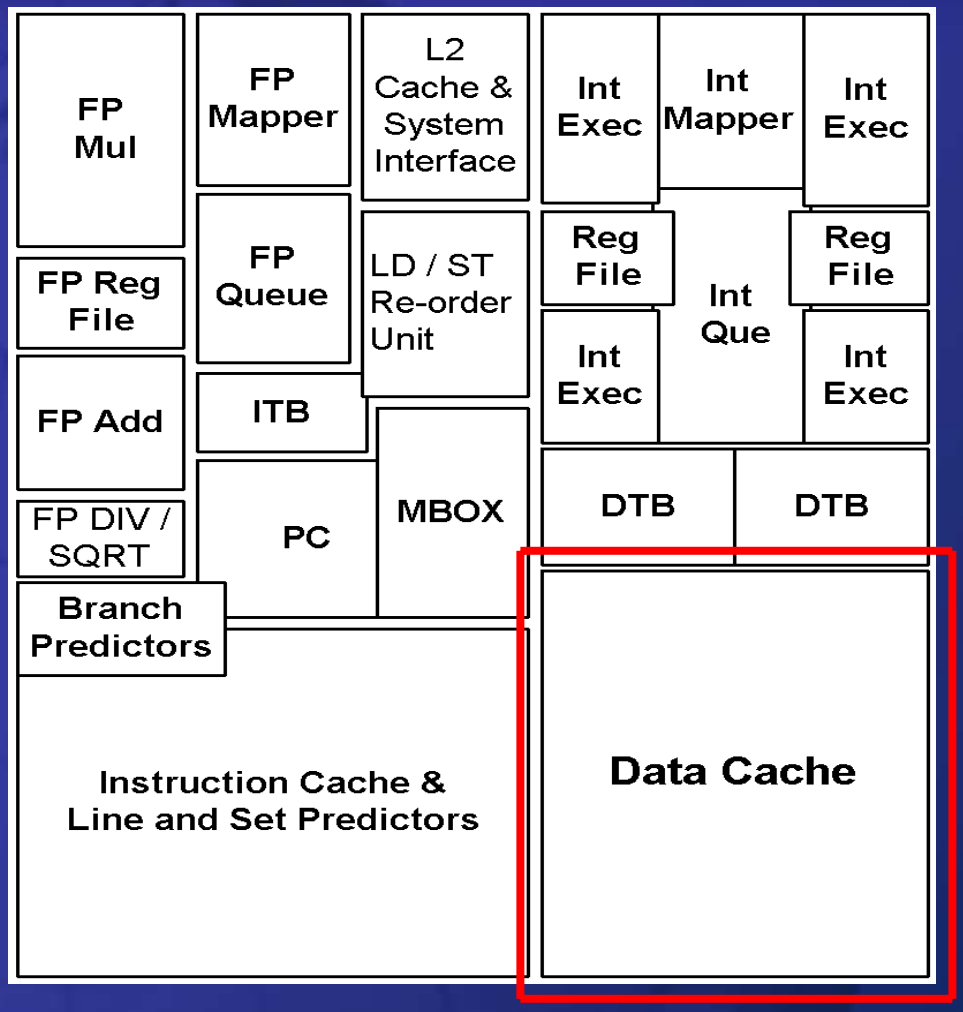

COMPAQ

Data Cache Overview

- 64 KB L1 On-Chip Data Cache
- 2 Way Set Associative
- 64 Byte Blocks
- Write-Back, Read/Write Allocate
- Virtually Indexed / Physically Tagged
- 128 Entry Fully Associative TLB
- 8 Entry Victim Buffer

COMPAQ

Data Cache – Block Diagram

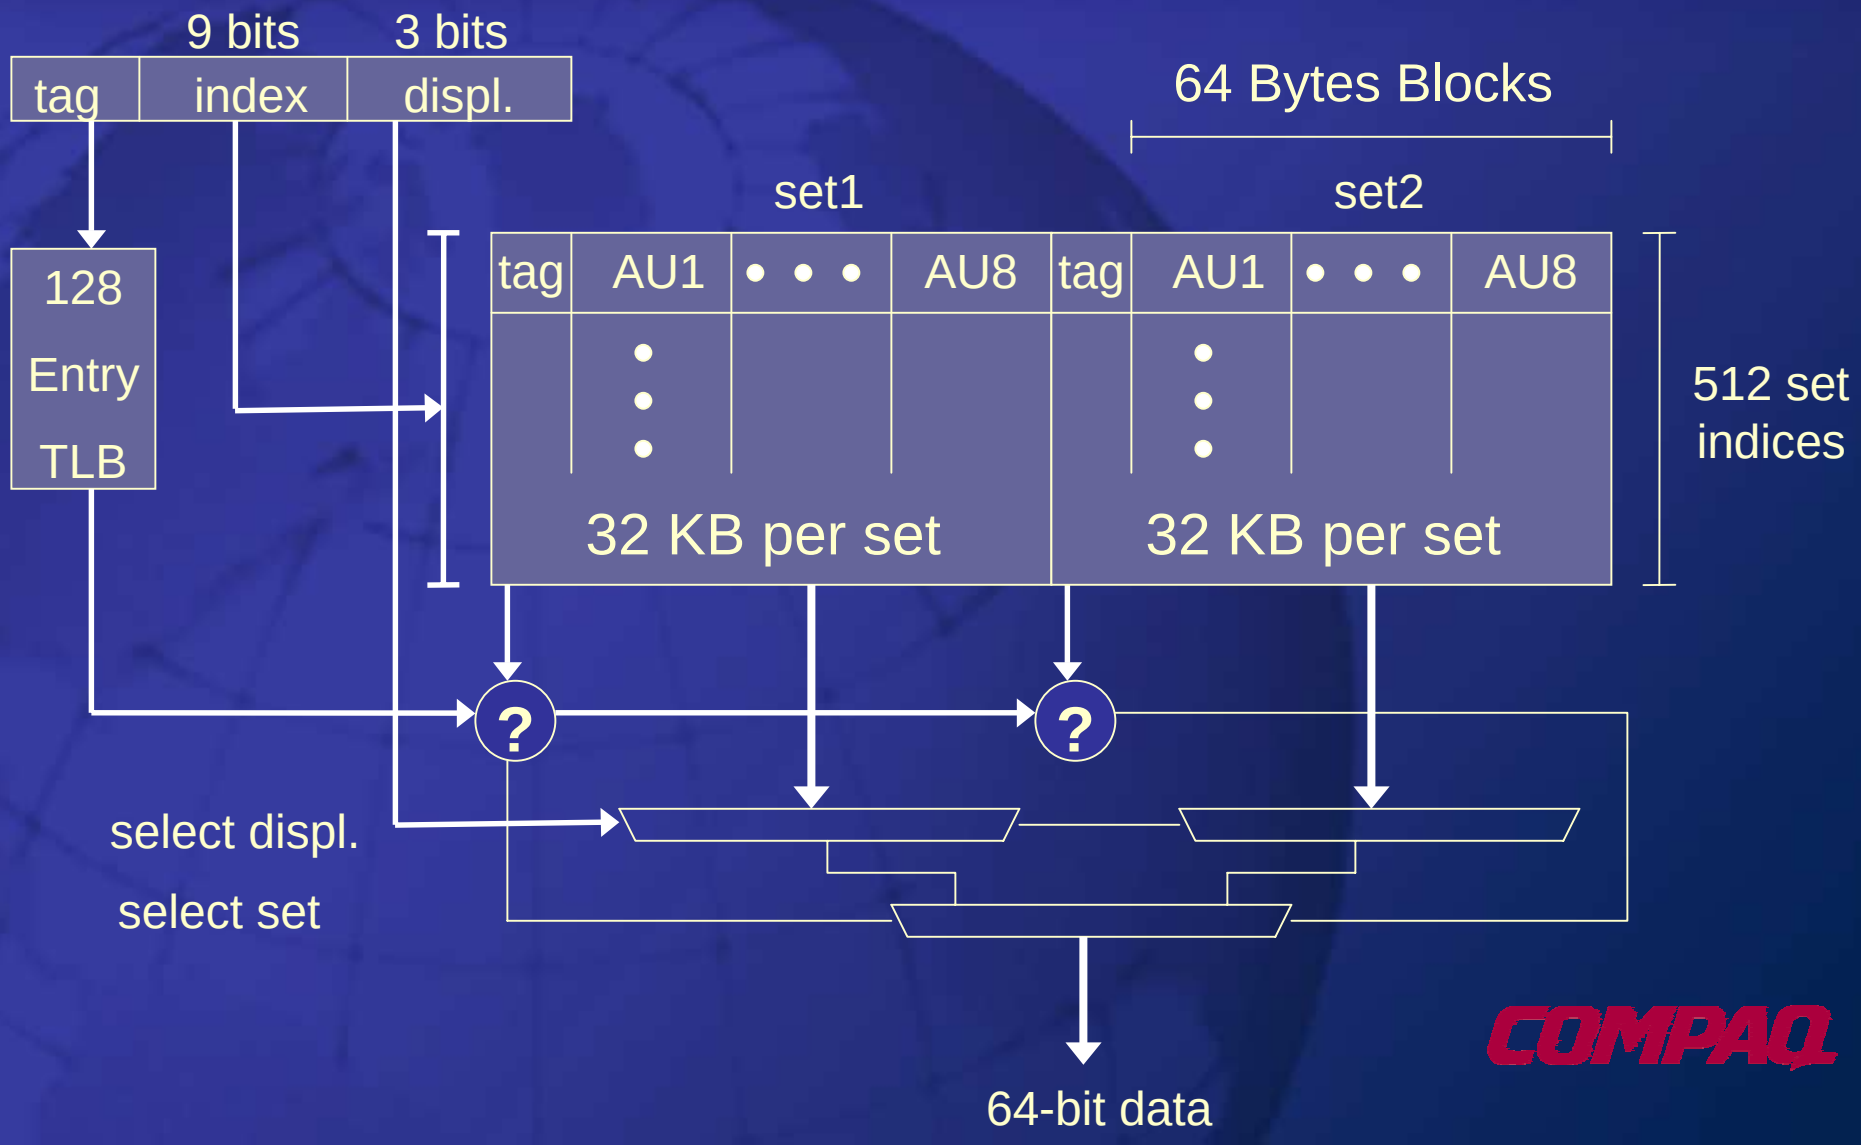

Data Cache – Access Cycle

- 2 memory ops per cycle (loads or stores)
- Single-Ported
  - " Double-pumped" – Accesses data on opposite clock phases
  - Single port reduces area and delay compared to dual ported
- 16 Bytes read / write per cycle
- Pipelined – 2 latency

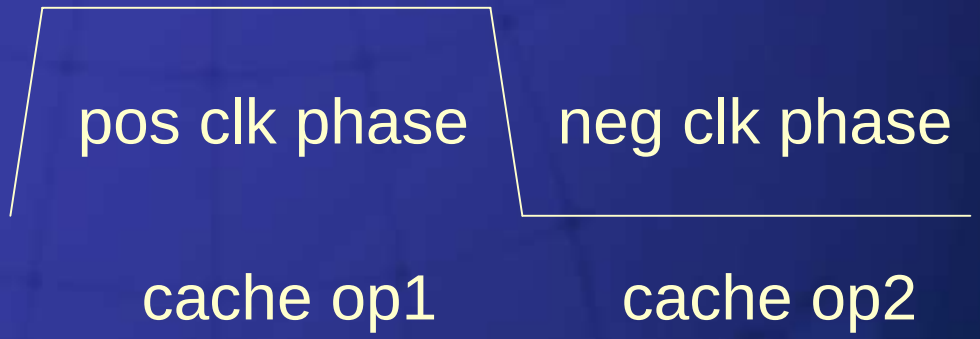

COMPAQ

Data Stream Summary

- 4 Integer pipes, divided among 2 clusters
  - 31 visible Registers, 80 total
- 2 floating point pipes
  - 31 visible Registers, 72 total
- 64 KB D-cache
  - 2 assoc., 64 Byte Blocks
  - Write-Back, Read/Write Allocate
  - Virtually Indexed / Physically Tagged
  - 128 Entry Fully Associative TLB
  - 8 Entry Victim Buffer

COMPAQ

The Alpha 21264 – External L2 and Memory System

Matt Ziegler

Alpha 21264 Performance-Focused Memory System [1]

- L1 Dcache: 9+ GB/s
  - 128b datapath with 3 cycle load-to-use latency
- L2 Cache: 6+ GB/s
  - 128b datapath with 12 cycle load-to-use latency
- System port: 3+ GB/s
  - 64b datapath with 80 cycle load-to-use latency
- 16 64-byte off-chip memory references
  - 8 read-misses and 8 write-backs

COMPAQ

L2 Cache and Memory System Overview

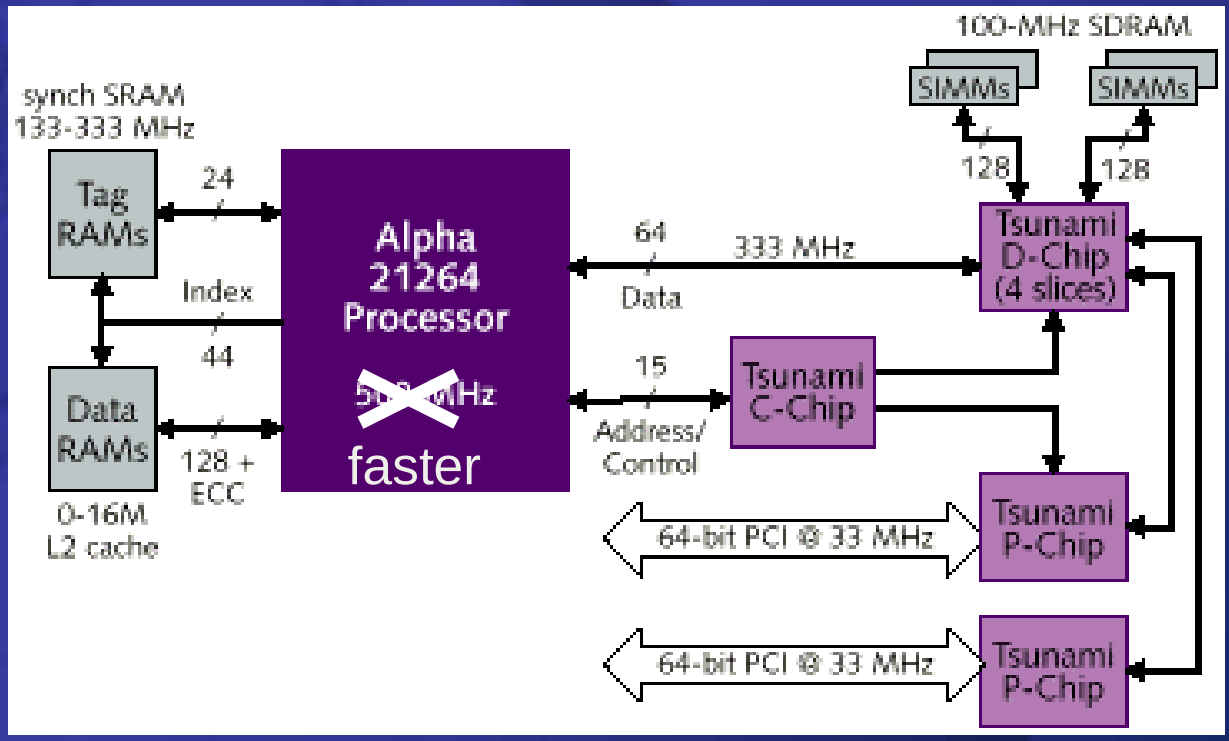

COMPAQ

Alpha 21264 External Interface [1]

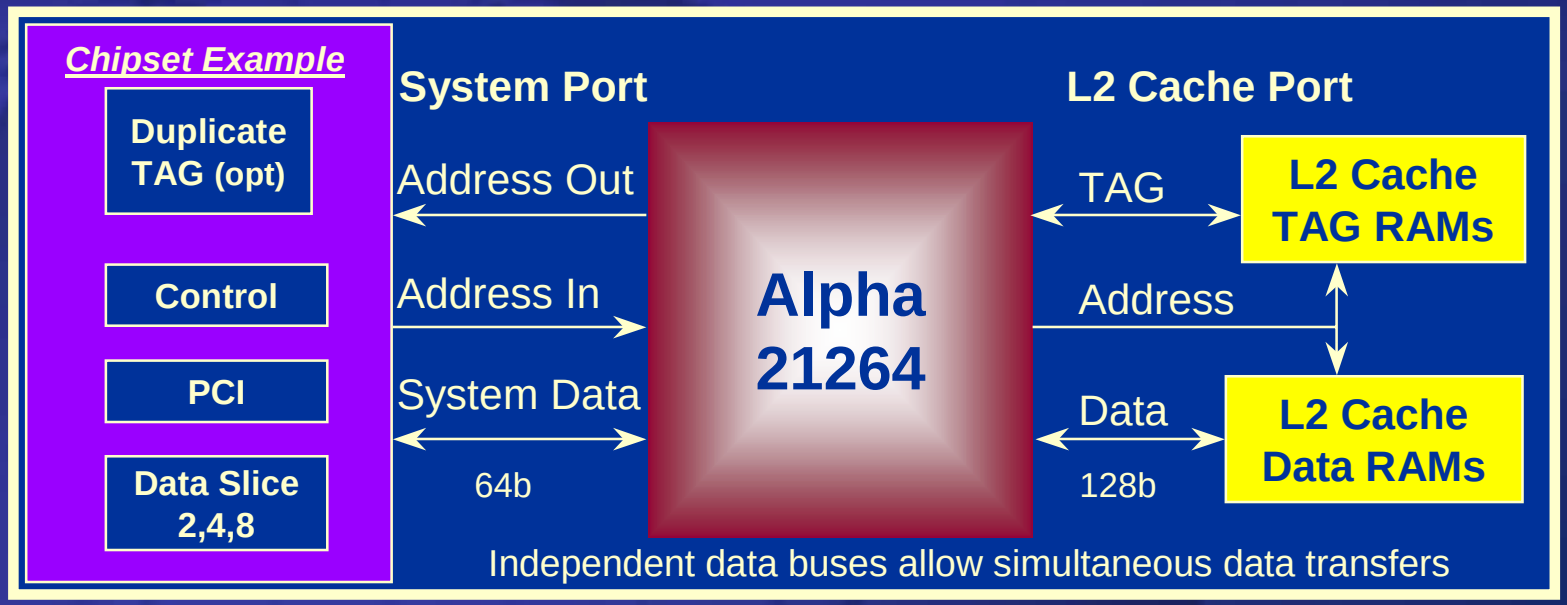

L2 Cache:	0 -16MB	Synch Direct Mapped	**Peak Bandwidth:**
Example 1: Reg-Reg	133	MHz BurstRAM	2.1 GB/sec
Example 2: 'RISC'	200+	MHz Late Write	3.2+ GB/sec
Example 3: Dual Data	400+	MHz	6.4+ GB/sec

COMPAQ

Off-Chip Unified L2 Cache

- 0 – 16 MB
- Physically Indexed
- 128 bit bus connecting L1 to L2
  - 16 bytes every 1.5 cycles
  - 12 cycle latency
- Non-blocking
  - 8 in-flight misses

COMPAQ

Alpha 21264 – V LSI Implementation

Matt Ziegler

VLSI Design Strategies

IBM

PowerPC

- Mainly Place & Route
- Low Cost
- Fast Design Cycle

Intel

Pentium

- Some Place & Route
- Some Custom
- Mid Range Cost
- Medium Design Cycle

Compaq

Alpha

- Mostly Custom Design
- High Cost
- High Performance
- Extensive Circuit Design

COMPAQ

Performance vs. Cost

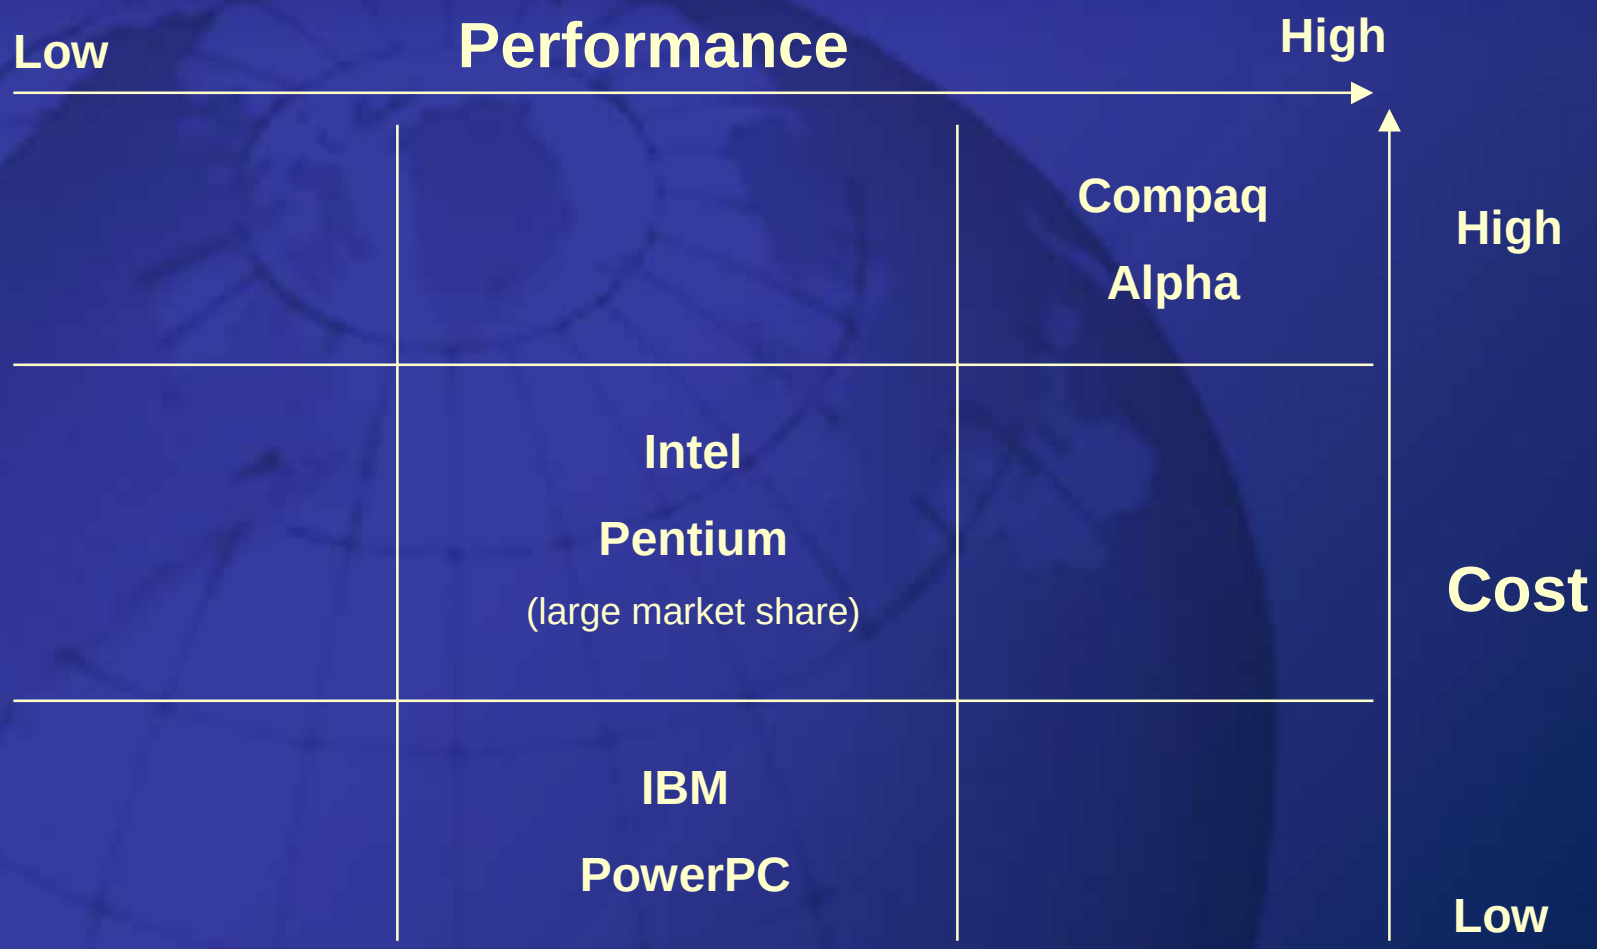

High Performance → High Power

Matt Ziegler

Consider This:

- The Alpha 21264 is like a Ferrari
  - Both are Engineering Masterpieces
  - Both deliver High Performance
  - Both are Gas/Power Guzzlers
- So, if you want High Performance, you have to pay for it!

COMPAQ

Alpha 21264 - Performance vs. Power

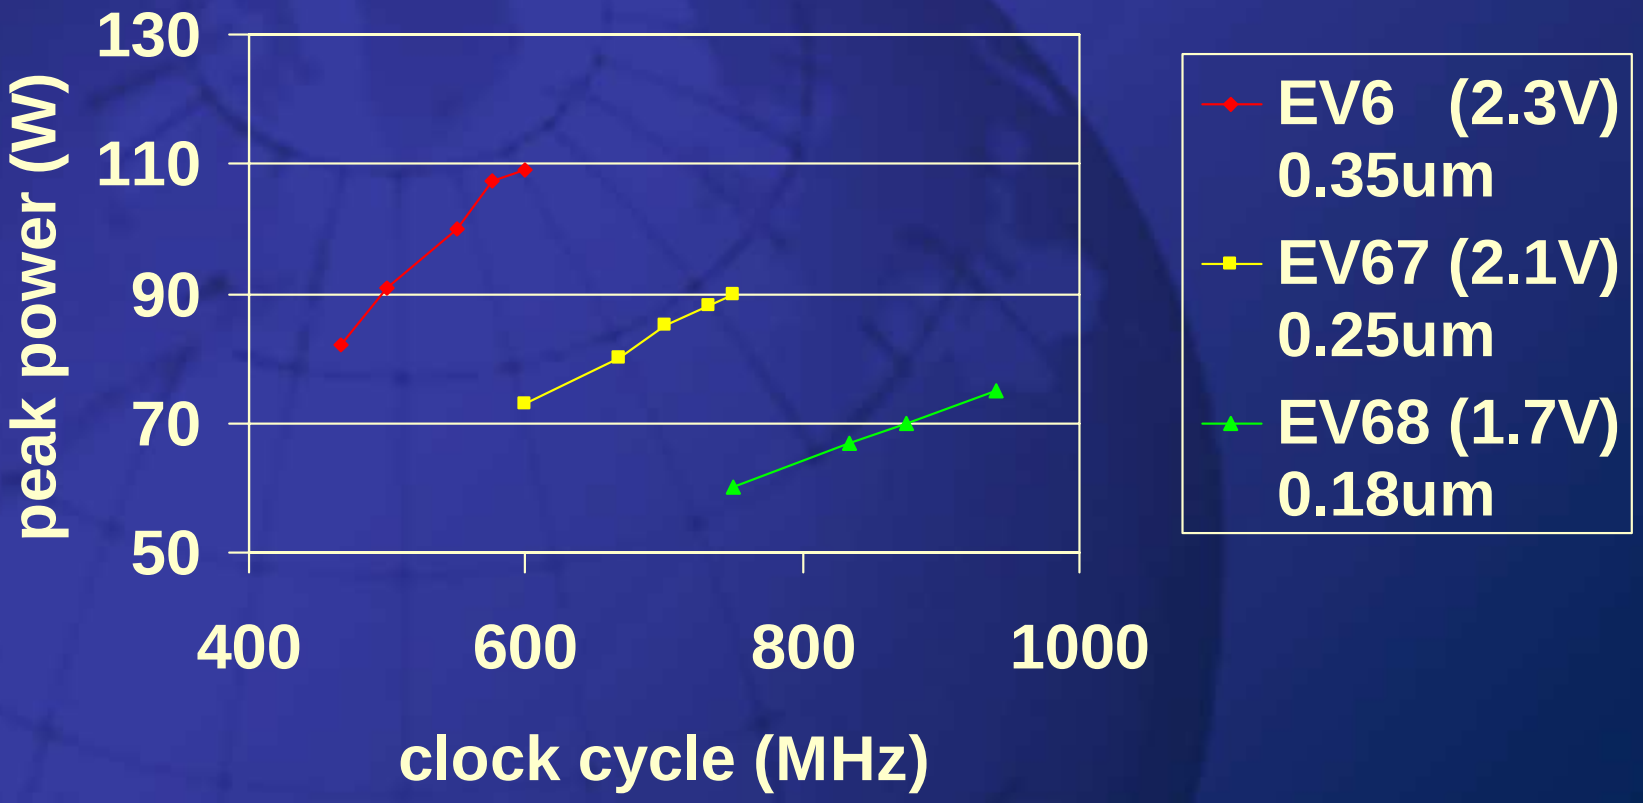

COMPAQ

EV6	2.7V 0.35um		EV67	2.1V 0.25um		EV68	1.7V 0.18um	
MHz	Watts							
466	82							
500	91							
550	100							
575	107.5		MHz	Watts				
600	109		600	73				
			667	80				
			700	85				
			733	88		MHz	Watts	
			750	90		750	60	
						833	67	
						875	70	
						940	75	

Alpha 21264 Power Distribution

Global Clock Network	32%
Inst Issue Units	18%
Caches	15%
FP Exe Units	10%
Integer Exe Units	10%
MMU	8%
I/O	5%
Misc. Logic	2%

COMPAQ

The Alpha Family of Microprocessors the 21264 and Beyond

Matt Ziegler

Alpha Family Overview [1]

- E5 (21164)
  - In-order 4-wide
- EV6 (21264)
  - .35 µm, 600 MHz
  - 4-wide superscalar
  - Out-of-order execution
  - Backside L2 cache port
- EV67
  - .25 µm, ~800 MHz
- EV68
  - .18 µm, >1000 MHz
- EV7 (21364)
  - .18 µm, >1000 MHz
  - L2 cache on-chip
  - RAMBUS
  - Glueless MP
- EV8 (21464)
  - .13 µm, 1400 MHz
  - 8-wide superscalar
  - SMT

COMPAQ

Alpha Family Evolution [1]

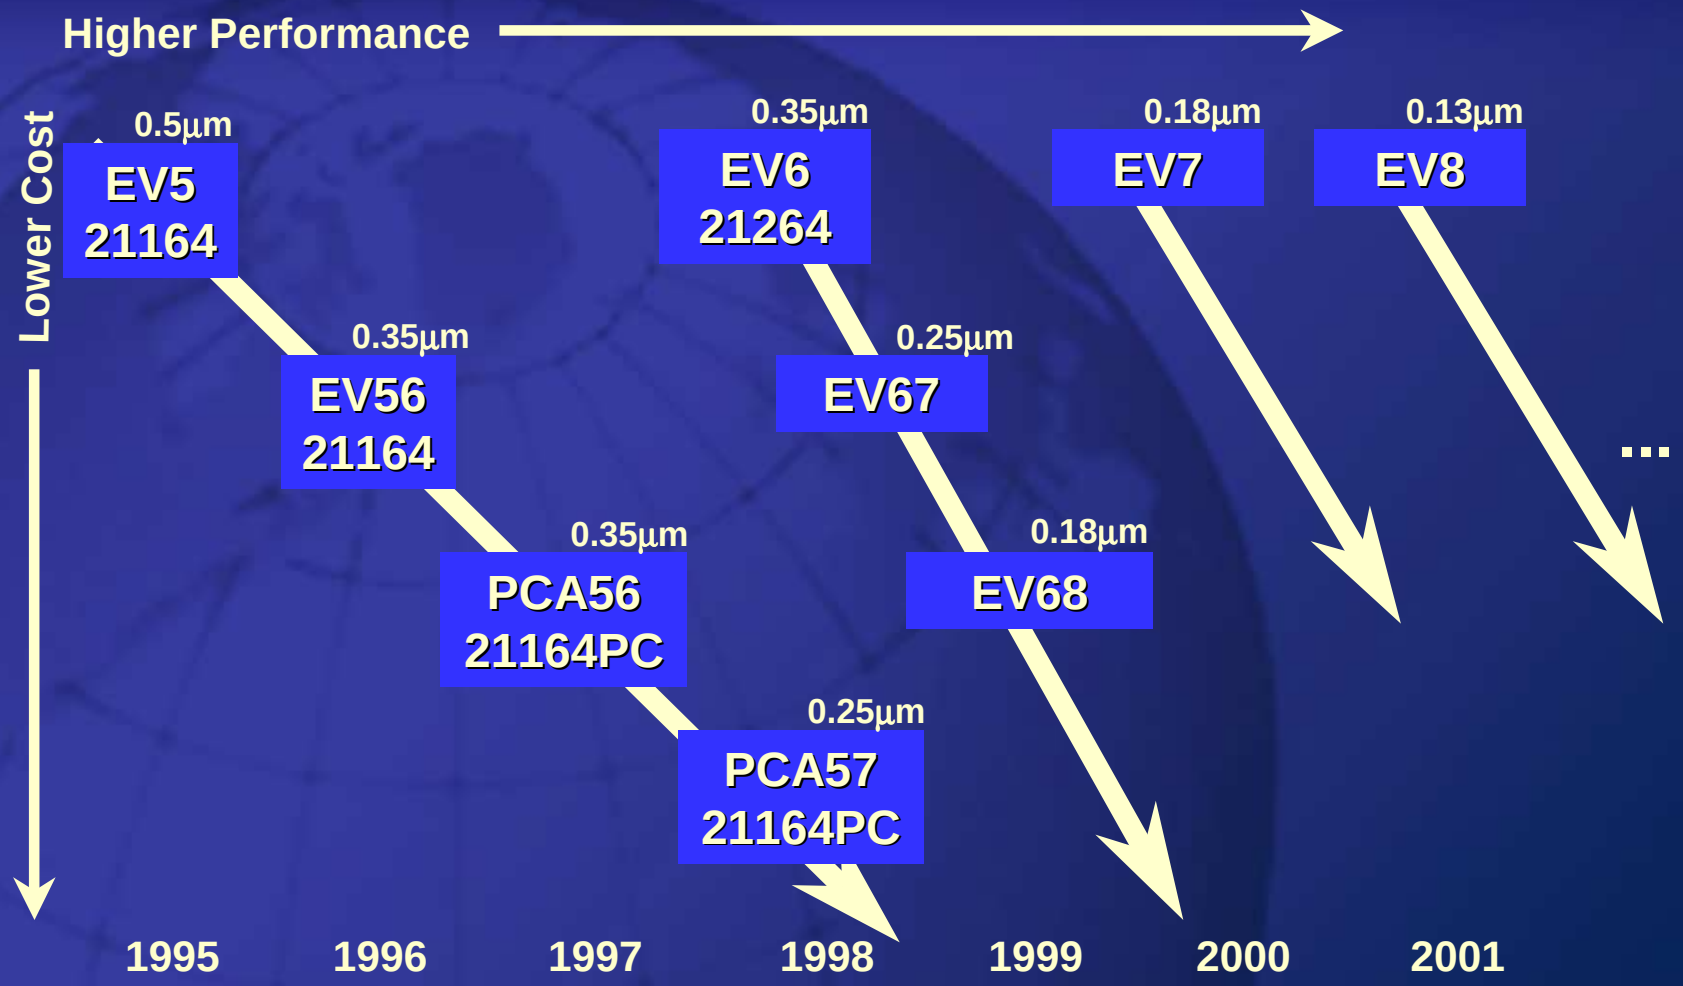

COMPAQ

EV7 Overview [1]

- 21264 core + enhancements:
  - Double the number of read-miss and victims (relative to 21264)
  - Graphics extensions
- On-chip 8-way associative L2 cache, currently 1.5MB
- RAMBUS DRAM memory interface
- Glue-less scalable, reliable system
- 70 - 80 SPECint95, 110-130 SPECfp95
- Sustained memory bandwidth -- 10GB/sec

COMPAQ

EV8 Overview [1]

- Enhanced out-of-order execution
- 8-wide superscalar
- 4-way simultaneous multi-threading (SMT)
- On-chip L2 cache, ≥ 2MB
- RAMBUS interface
- New instruction fetcher and branch predictor
- ~ 200 SPECint95, ~ 300 SPECfp95
- Sustained memory bandwidth -- 10GB/sec

COMPAQ

Alpha Family Roadmap [1]

	EV56	EV6	EV67	EV68	EV7	EV8
Schedule -						
Ship Date	Jun-96	H2 1998	H1 1999	H2 1999	2000	2001
Technology -						
CMOS	.35um	.35um	.28um	.18um	.18um	.13um
Vdd (V)	2.5	2.2	2	1.5	1.5	1.2
Packaging	WB/PGA	WB/PGA	WB/PGA	FC/SCP	FC/SCP	FC/SCP/MCP
Pins	499	587	587	587	~1400	~1800
Chip Characteristics -						
Frequency (MHz)	600	600	~800	>1000	>1000	~1400
Performance (SpecINT95)	18.8	33	~45	~65	~75	~200
Performance (SpecFP95)	29	52	~70	~100	~120	~300
Sustained Memory BW (GB/sec)	0.5	2	2	2	10	10
Sustained Cache BW (GB/sec)	1.3	4	4.5	5.5	16	48
Power (W)	55	95	85	60	100	120
Die Size (mm^2)	210	320	210	150	350	300
Architectural Features		Out-of-order execution, dedicated L2 cache port			RAMBUS, 4-CPU switch, 1.5MB on-chip L2 cache	8-wide superscalar issue, SMT

COMPAQ

DONE