COMPAQ
Inspiration Technology

Alpha 21364 (EV7)

Alpha Microprocessor Roadmap

	EV68C	EV68C	EV7	EV79
Chip Characteristics				
Frequency (GHz)	1	1.25	~1.2	~1.6-1.7
Power (W) max	65	75	155	120
Die Size (mm2)	125	125	400	300
Technology				
Vdd (V)	1.65	1.65	1.65	1.2
CMOS (drawn um)	0.18	0.18	0.18	0.13 - SOI
Packaging	FC/LGA	FC/LGA	FC/LGA	FC/LGA
Pins	675	675	1443	1443
Schedule				
FirstTapeOut	Mar-00	Mar-00	Apr-01	Q1'03
Volume	Apr-01	Dec-01	Q3'02	H1'04

Estimated time for TPC-C

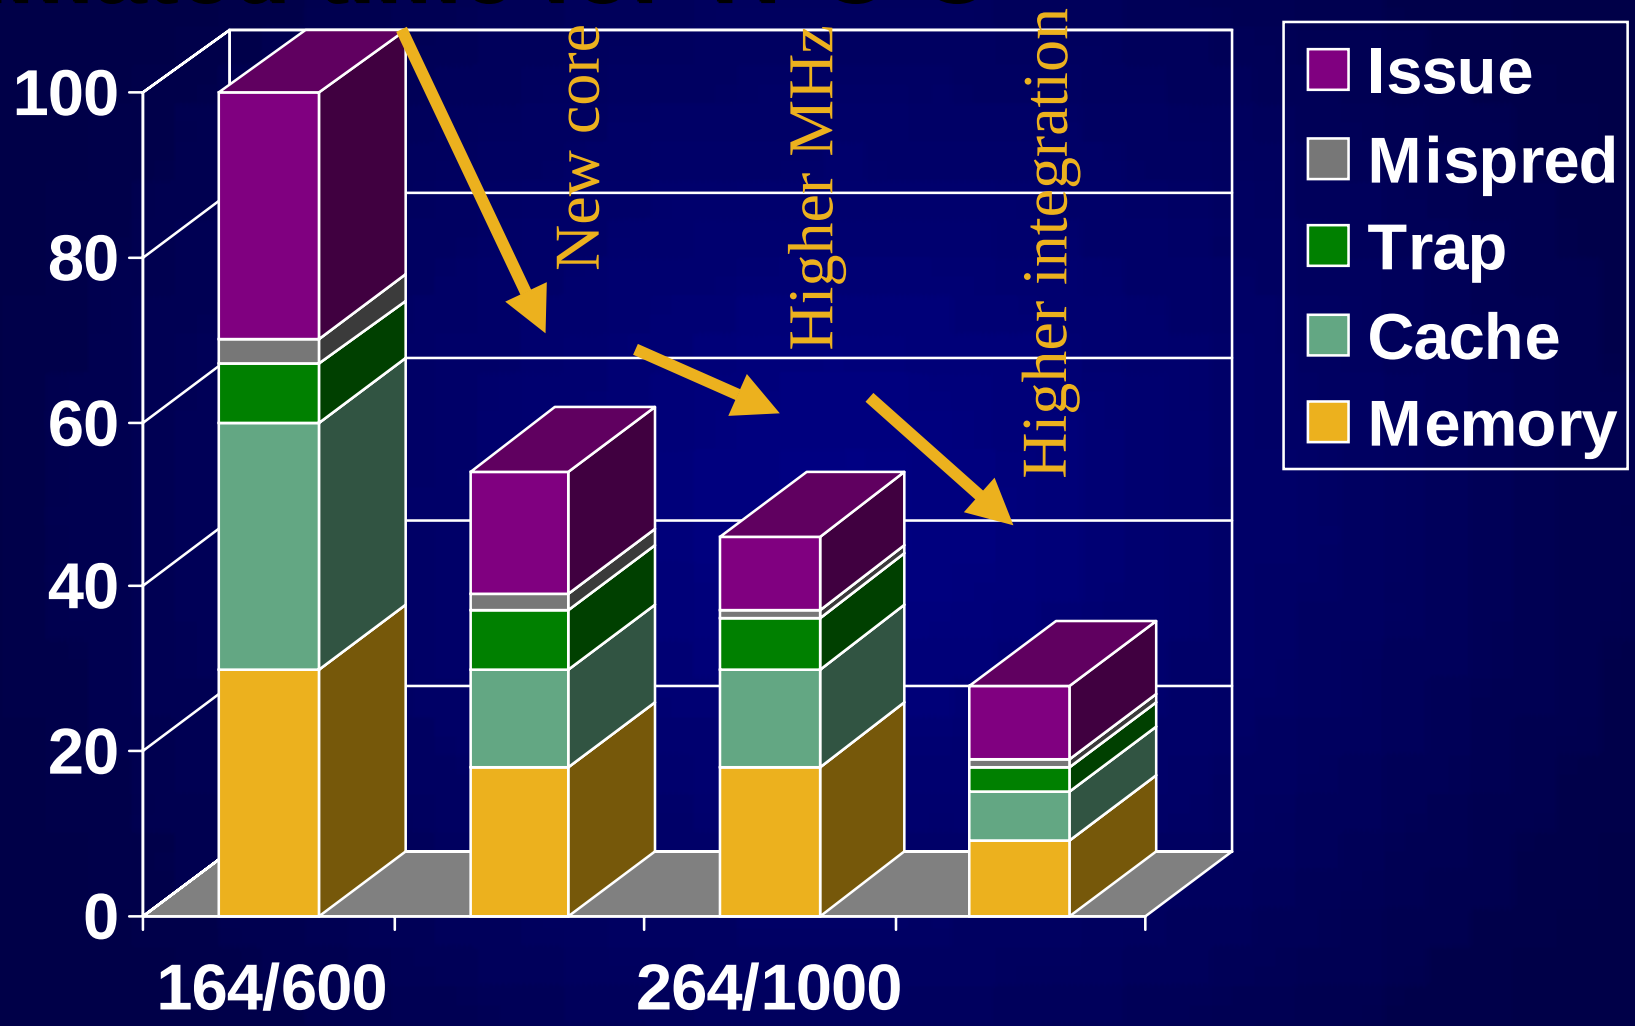

Alpha 21364 Goals

- Improve
  - Single processor performance, operating frequency, and memory system
  - SMP scaling
  - System performance density (computes/ft^3)
  - Reliability and availability
- Decrease
  - System cost
  - System complexity

Alpha 21364 Features

- Alpha 21264 core with enhancements
- Integrated L2 Cache
- Integrated memory controller
- Integrated network interface
- Support for lock-step operation to enable high-availability systems.

Alpha 21364 Technology

- 0.18 μm CMOS
- 1250 MHz
- 135 Watts @ 1.65 volts
- 4 cm^2
- 7 Layer Metal
- 152 million transistors
  - 15 million logic
  - 137 million SRAM

21364 System Block Diagram

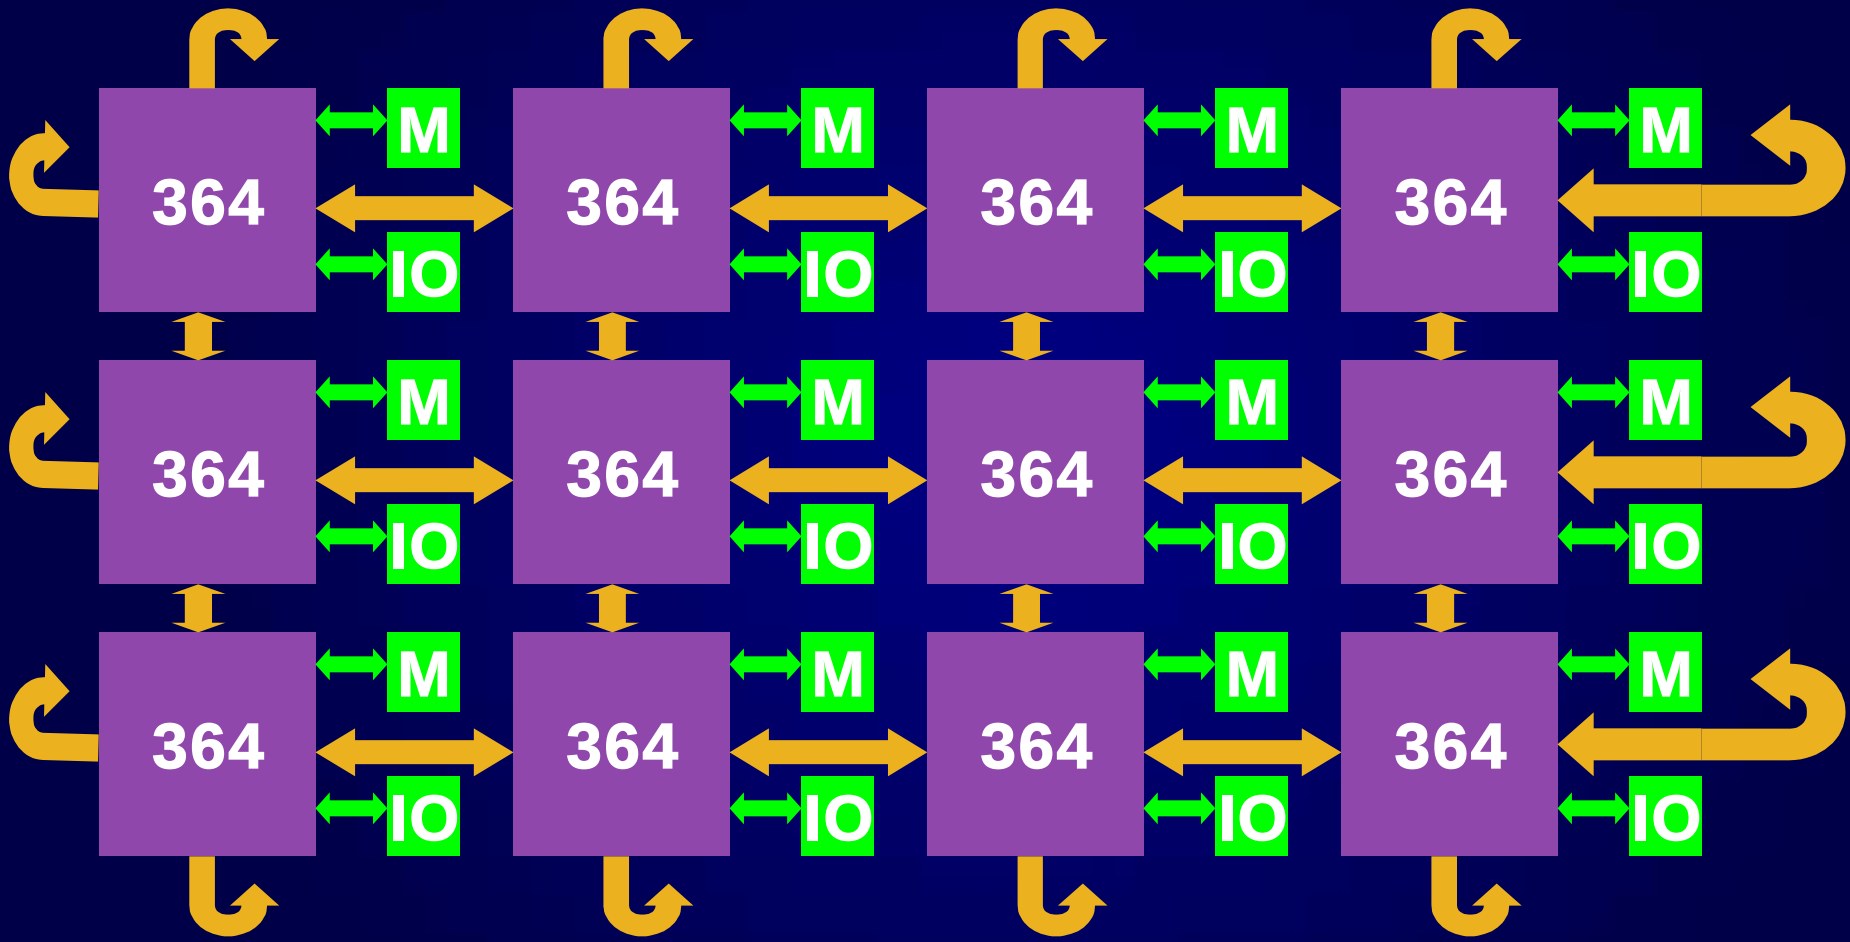

Dual Processor Building Block Module

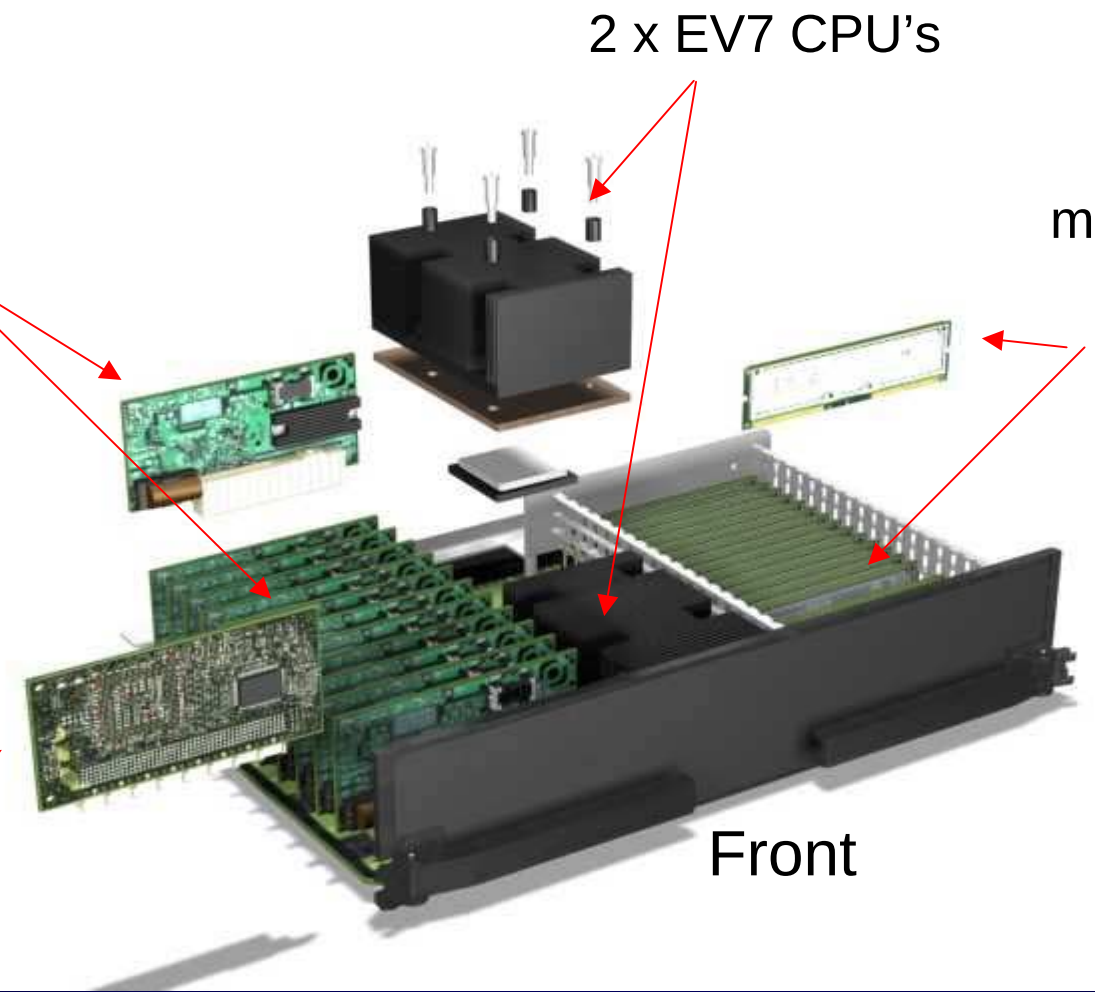

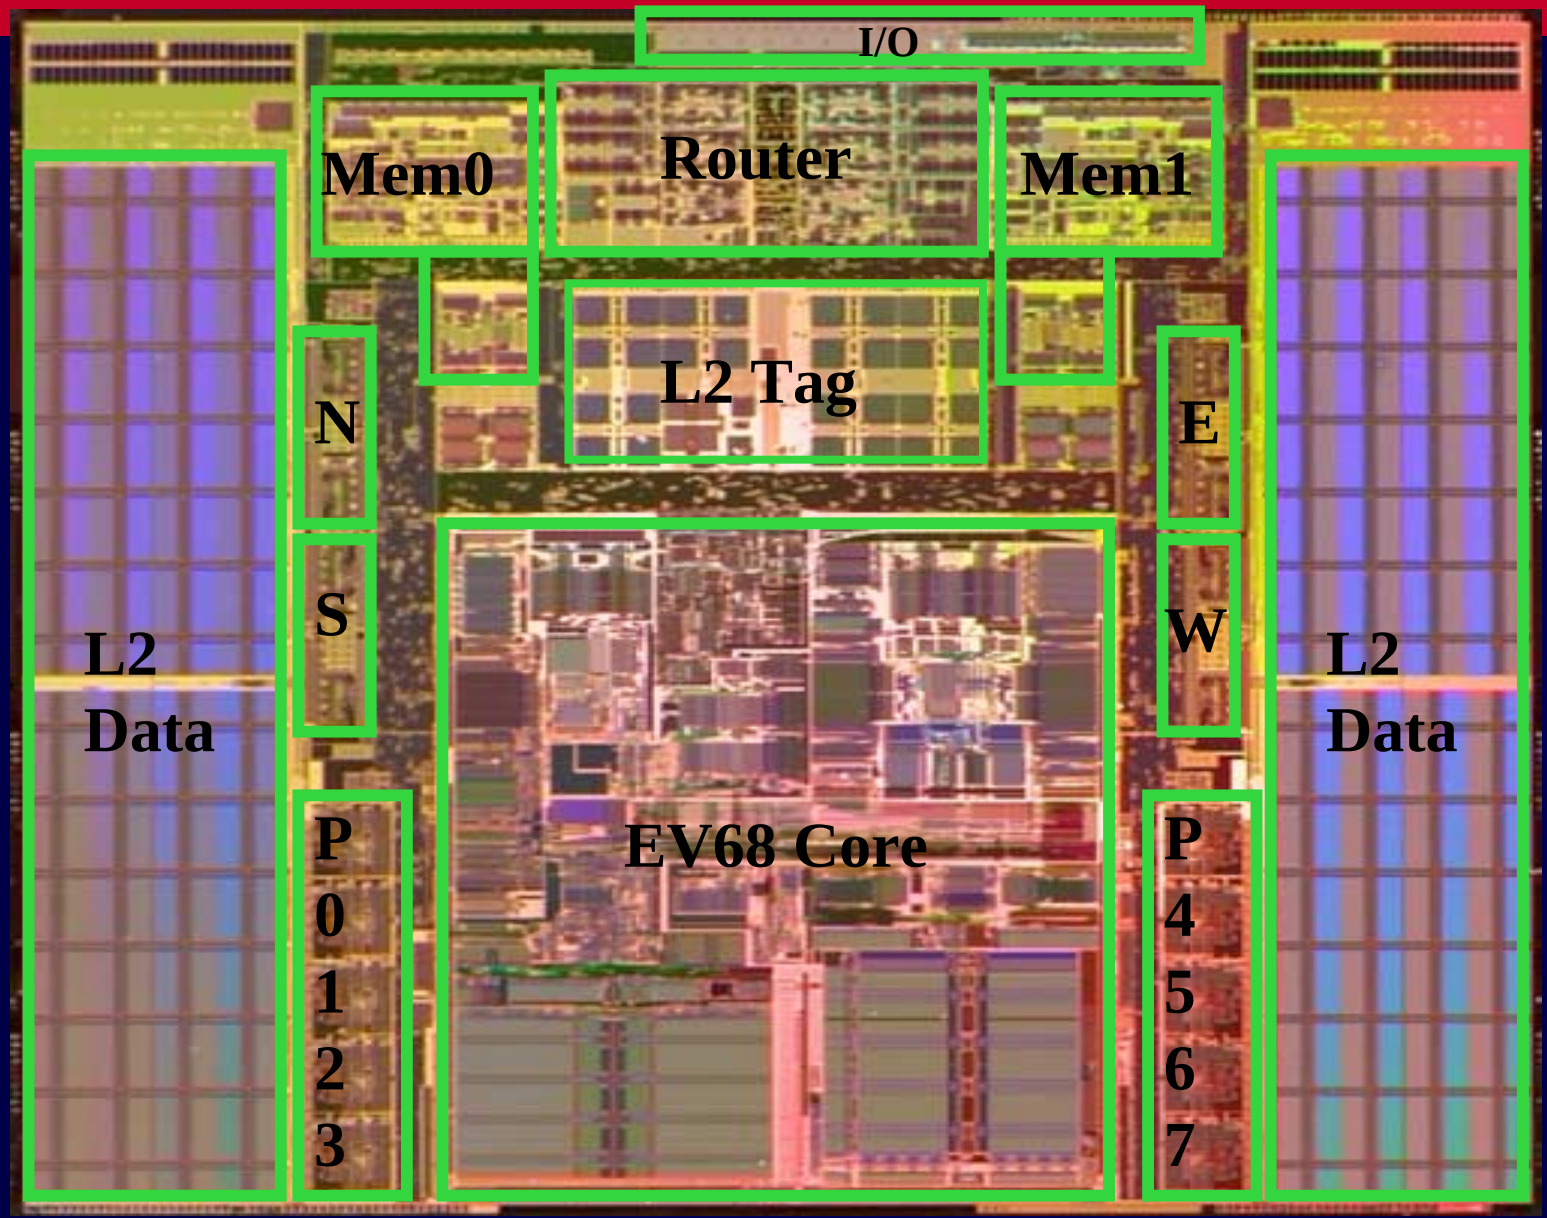
I/O
Mem0
Router
Mem1
L2 Tag
N
E
S
W
L2
Data
L2
Data
EV68 Core
P
0
1
2
3
P
4
5
6
7

21364 Core

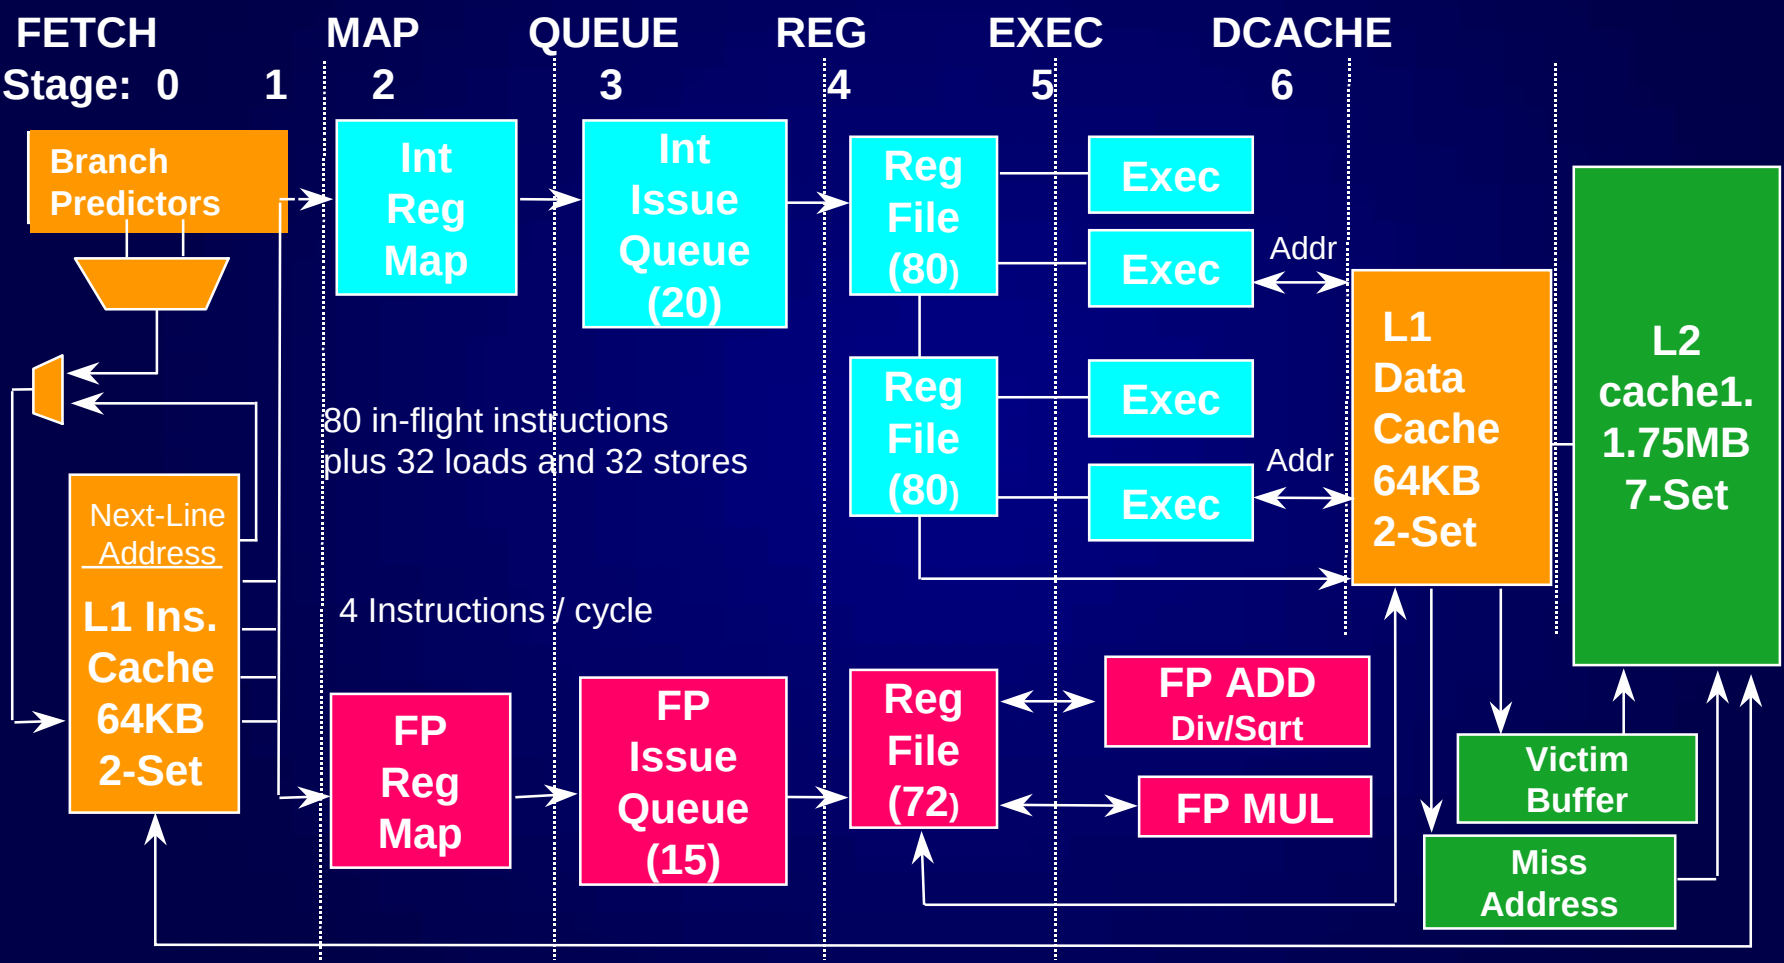
FETCH
Stage: 0
1
MAP
2
QUEUE
3
REG
4
EXEC
5
DCACHE
6
Branch Predictors
Int Reg Map
Int Issue Queue (20)
Reg File (80)
Exec
Exec
Addr
Reg File (80)
Exec
Exec
Addr
L1 Data Cache 64KB 2-Set
L2 cache1. 1.75MB 7-Set
80 in-flight instructions
plus 32 loads and 32 stores
Next-Line Address
L1 Ins. Cache 64KB 2-Set
4 Instructions / cycle
FP Reg Map
FP Issue Queue (15)
Reg File (72)
FP ADD Div/Sqrt
FP MUL
Victim Buffer
Miss Address

EV7 Addressing - 4TB

Linear Addressing:

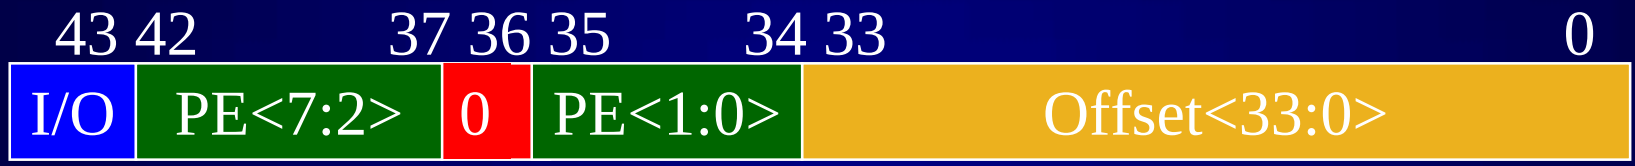

Quad CPU Interleaving:

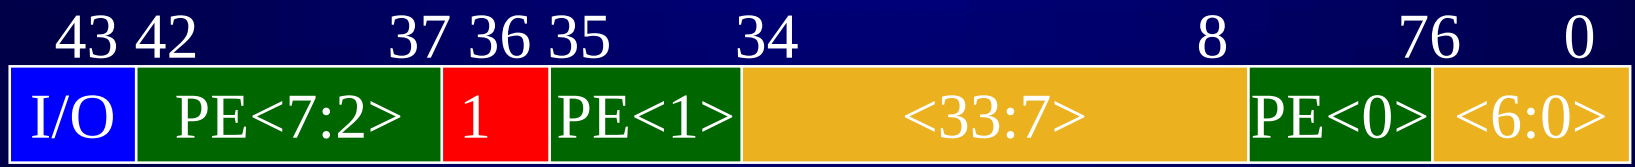

Virtual Page Size

- Current virtual page size
  - 8K
  - 64K
  - 512K
  - 4M
- New virtual page size (boot time selection)
  - 64K
  - 2M
  - 64M
  - 512M

Integrated L2 Cache

- 1.75 MB, 7-way set associative, with ECC
- 20 GB/s total read/write bandwidth
- 16 Victim buffers for L1 -> L2
- 16 Victim buffers for L2 -> Memory
- 9.6ns load to use latency
- Tag access start every cycle
- Data access in 4 cycle blocks
- Couple Tag/Data access to minimize latency
- Decoupled Tag access to minimize resource use.

Two Integrated Memory Controllers

- RDRAM memory
  - Directly connect to the processor
  - High data capacity per pin
  - 800 Mb/s operation
- 75ns load to use latency
- 12.8 GB/sec peak bandwidth
- 6 GB/sec read or write bandwidth
- 2048 open pages
- 64 entry directory based cache coherence engine
- ECC SECDED
- Optional 4+1 parity in memory

ZBox Block Diagram

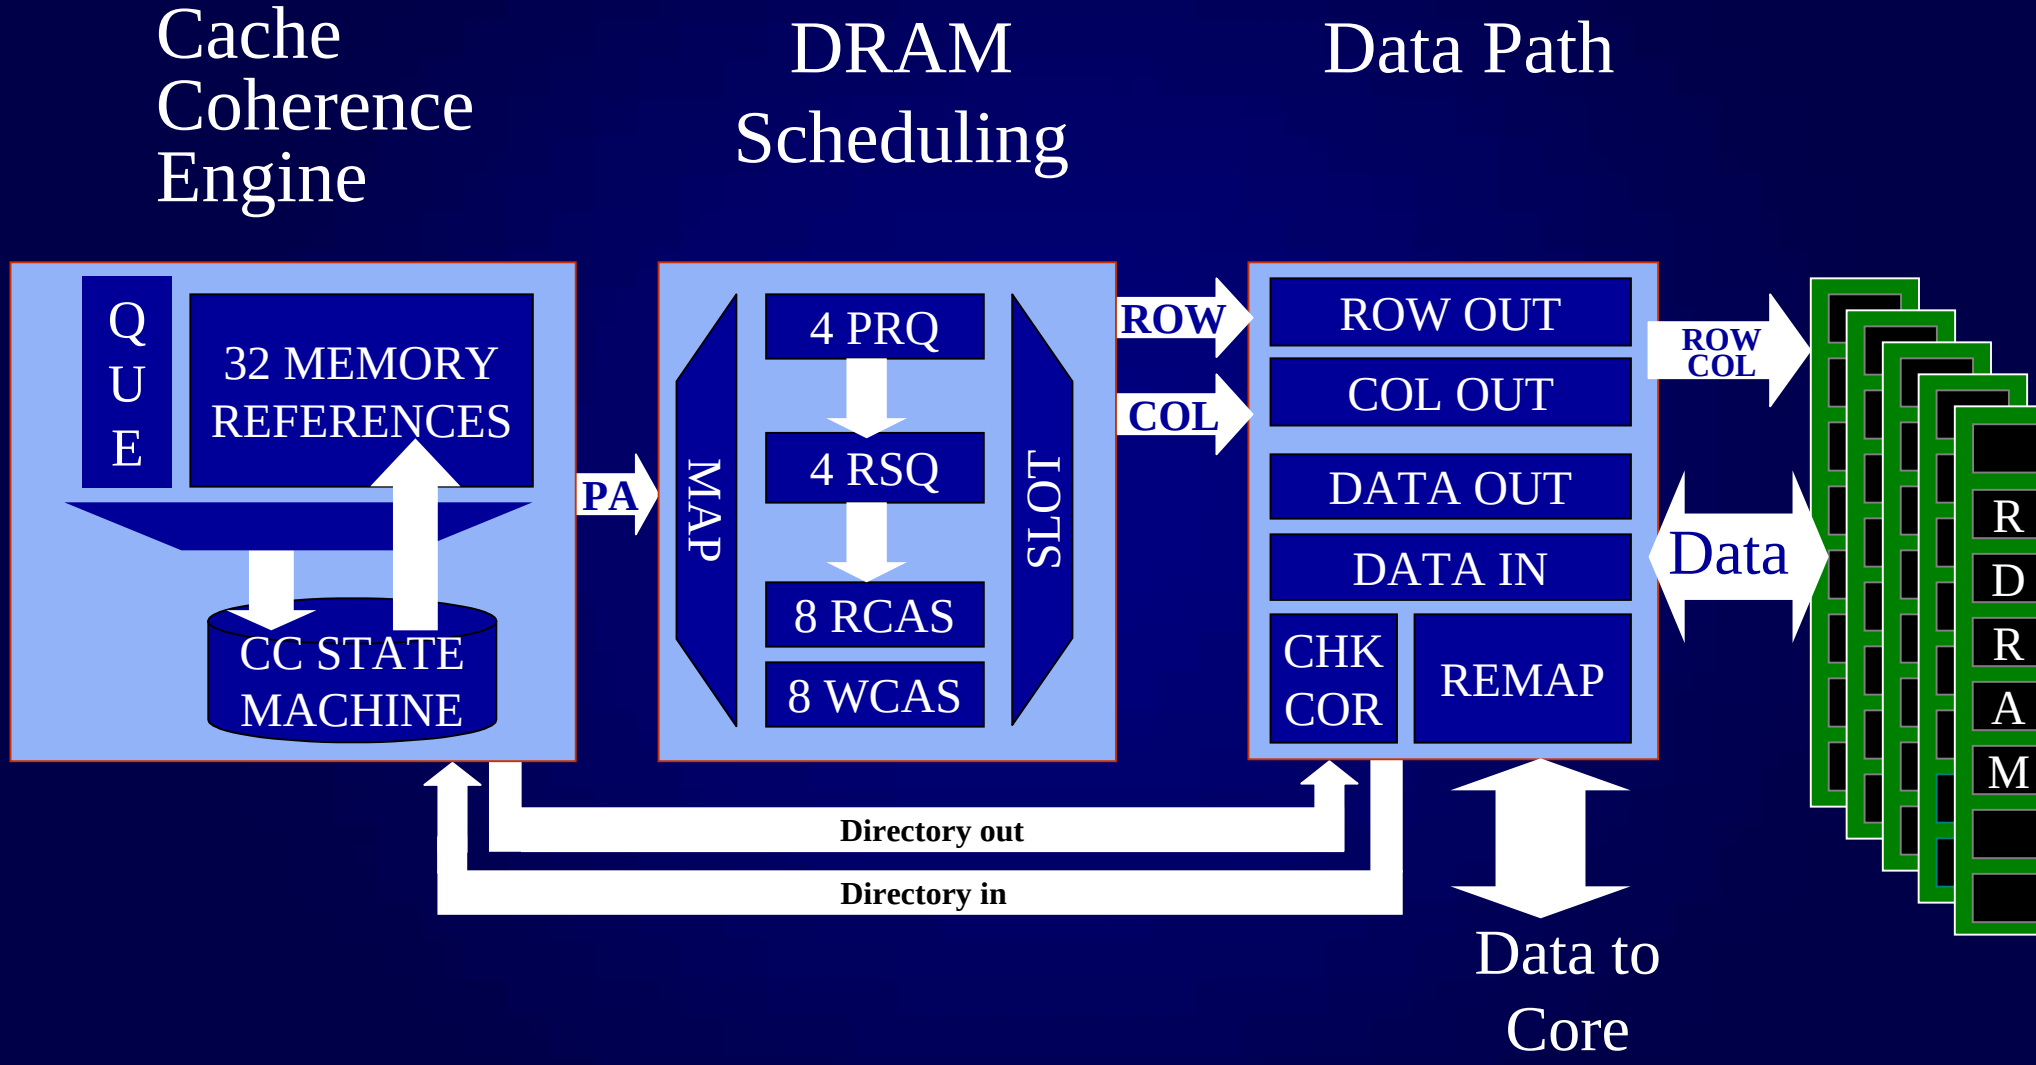

RDRAM Memory Interface

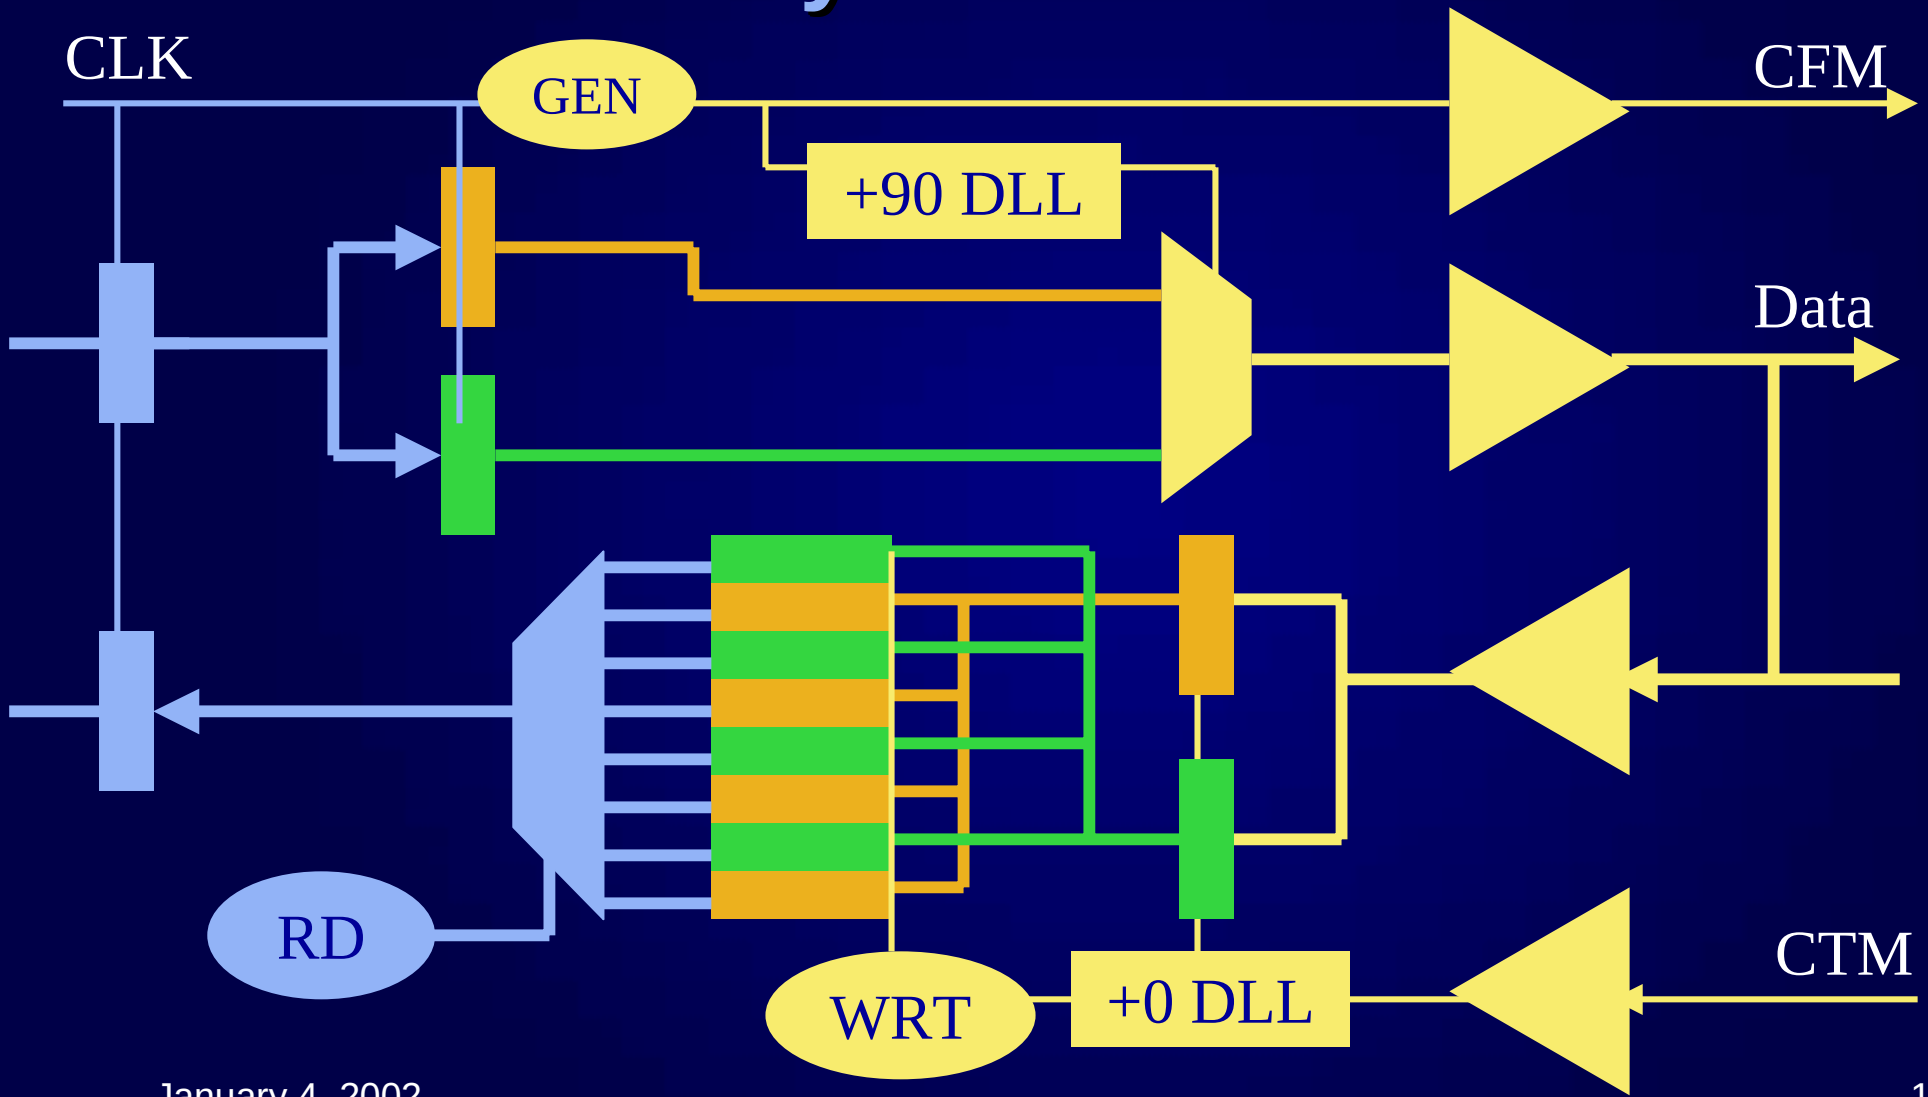

Integrated Network Interface

- Direct processor-to-processor interconnect
- 4 links 6.4 GB/second per link
  - 32 bits + ECC at 800 Mb/s each direction
- 18ns processor-to-processor latency
- ECC, single error correct, double error detect, per hop
- Out-of-order network with adaptive routing
- Asynchronous clocking between processors
- 3 GB/second I/O interface per processor

Rbox Block Diagram

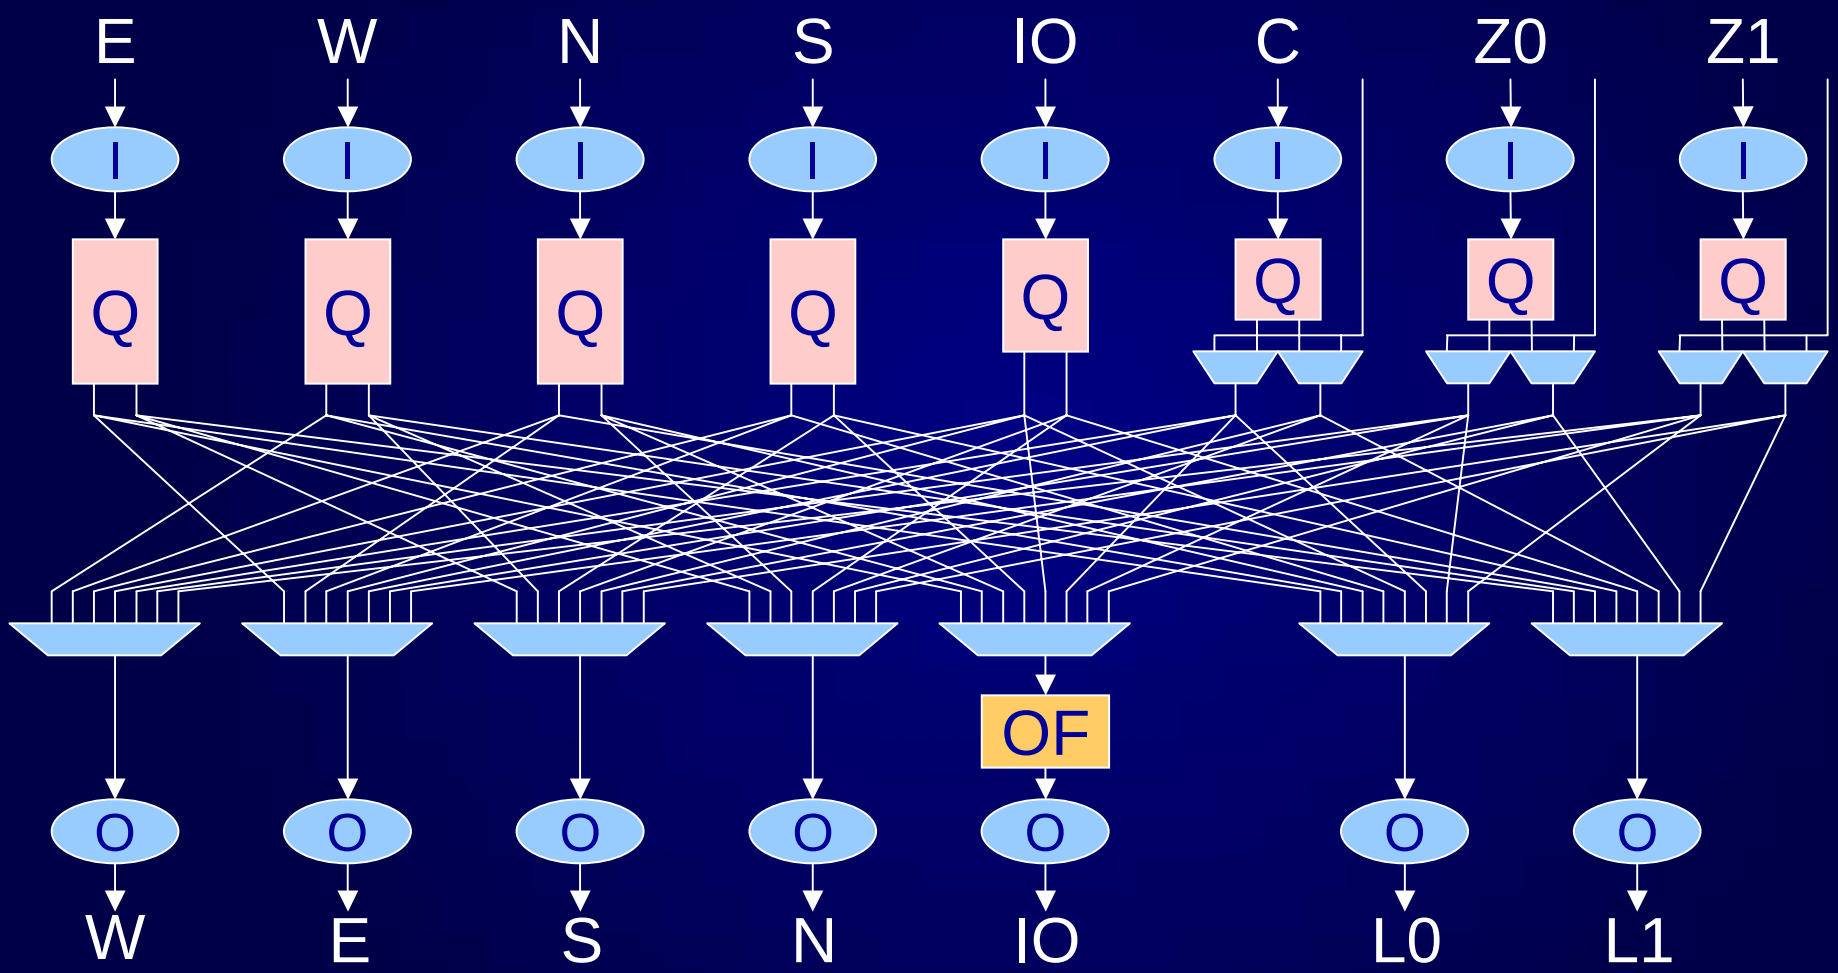

Router Table

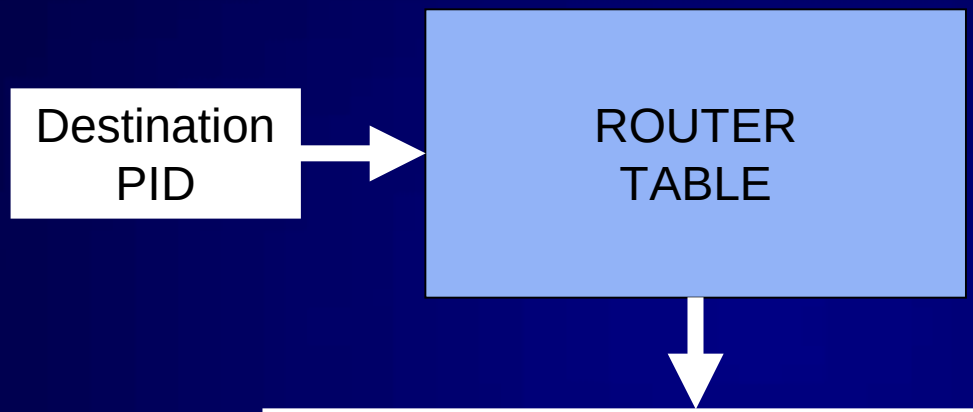

- 128 + 20 entries
- 4 lookups/cycle

- Initial/No initial Hop
- Initial hop direction
- Target EW and NS coordinates
- EW and NS directions
- Can/Can't adapt
- Access priv's

Memory Directory

- 27 bit directory stored with memory data
- Limited pointer based design

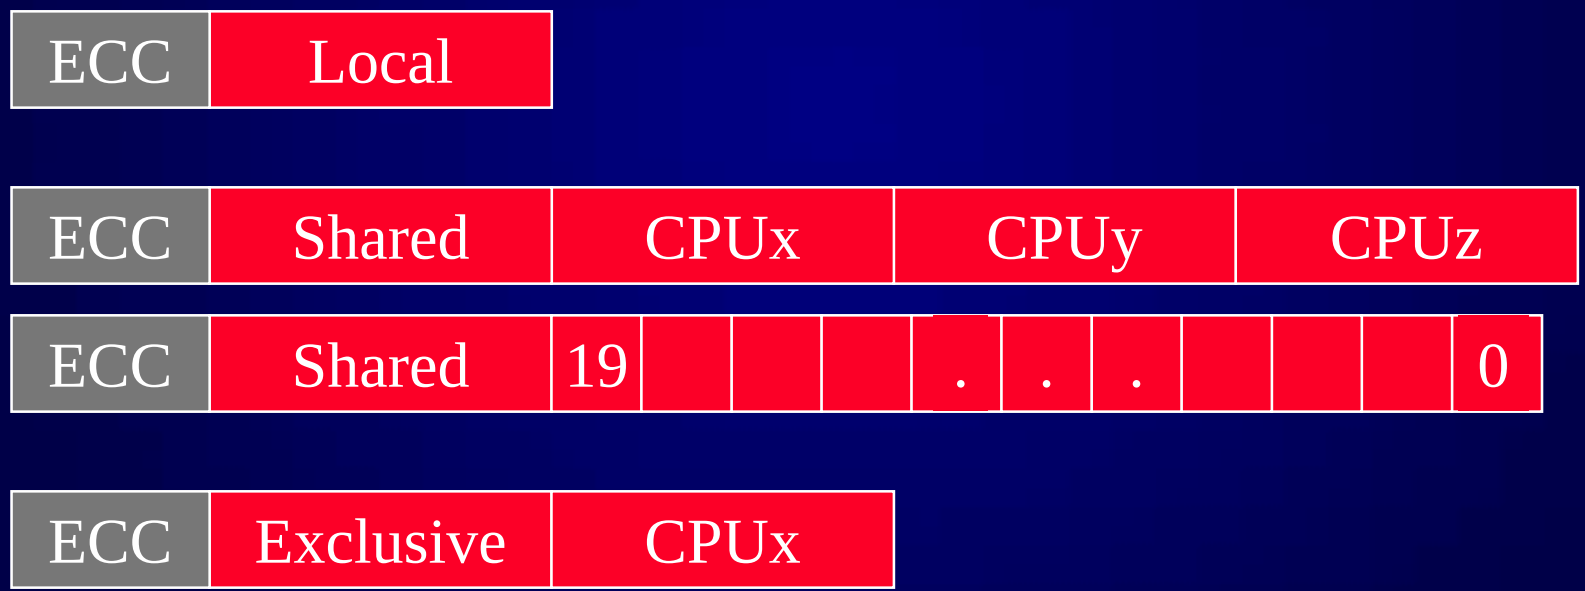

Node Terminology

- Requester (R) - node encountering a read or write miss
- Home (H) - node that contains the memory and directory for the referenced line
- Owner (O) - remote node that contains an exclusive copy of the line in its cache
- Sharer (S) - remote node that contains a shared copy of the line in its cache

Example 1: read, local home

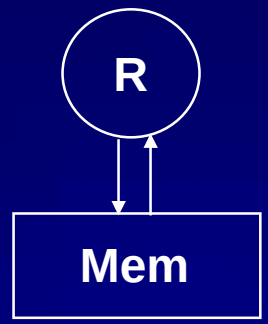

- Conditions:
  - home/memory is at local node
  - directory state is local or shared
- Actions:
  - retrieve data directly from local memory
  - directory is not updated, so very efficient (state of line at home is first determined by cache probe)

Example 2: read, remote home

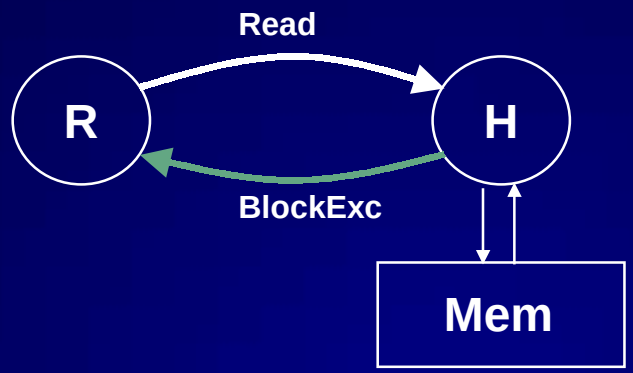

- Conditions:
  - home/memory is remote
  - directory state is shared or local
- Actions:
  - request sent to home
  - home node gets line from cache/memory, updates directory state, and replies

Example 3: read, remote owner

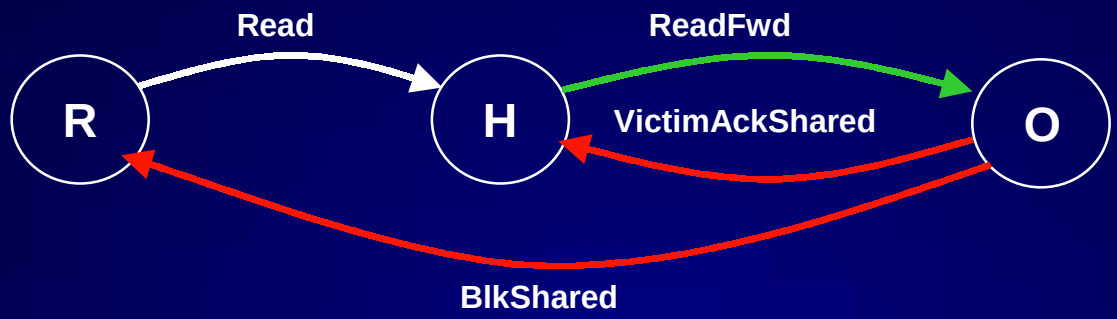

- Conditions:
  - home is remote, directory state is exclusive
- Actions:
  - Read request sent to home
  - home node forwards request to owner, leaves directory entry pending
  - owner sends read reply with data to requester, sharing writeback data to home
  - home makes directory entry not pending when writeback arrives
- Pending state maintains serialization order

Example 4: write, remote owner

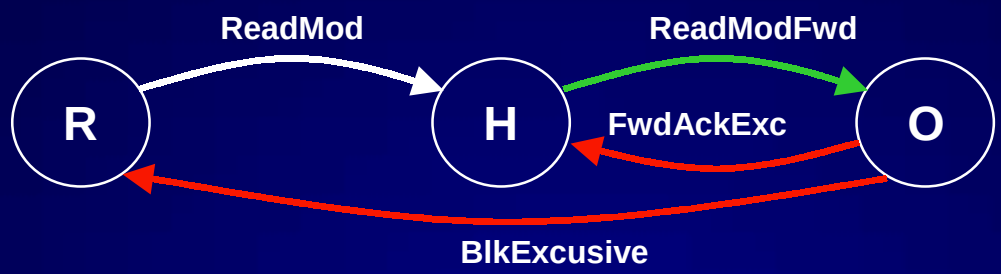

- Conditions:
  - home is remote, directory state is exclusive
- Actions:
  - read modify request sent to home, forwarded to owner
  - directory points to R as new owner
  - owner sends reply with data to requester

Example 5: write, remote sharers

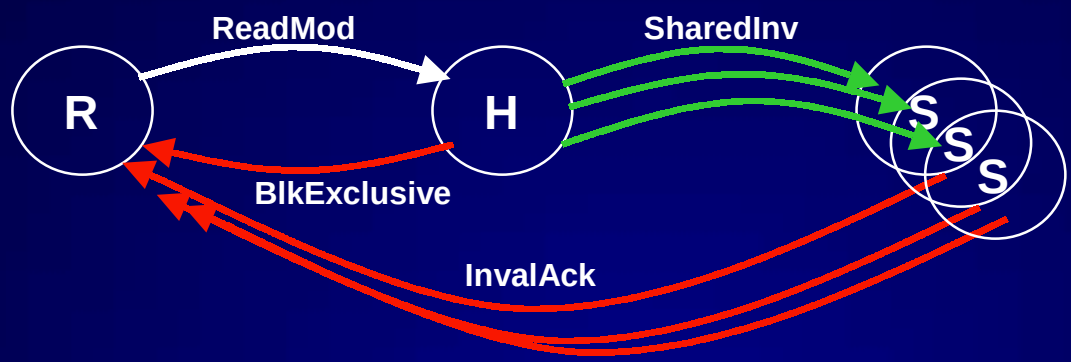

- Conditions:
  - home is remote, directory state is shared
- Actions:
  - read-exclusive request to home
  - home sends invalidation requests to sharers, sends data back to requester with invalidation count (early exclusive reply)
  - sharing nodes reply to *requester* with invalidation acknowledgements
  - requester proceeds when data arrives, but must stall incoming requests and potential writeback of line until all acks are received

Example 6: writebacks

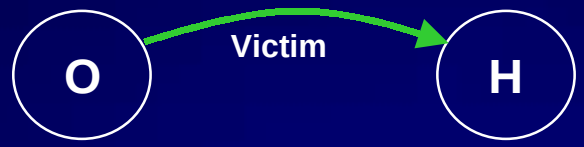

- Conditions:
  - owner has modified line in cache, and must replace line from cache
- Actions:
  - owner sends writeback request with data to home
  - home writes data to memory, changes directory state to local

EV7 Error Correction & Containment

- ECC on cache, memory, IP links, and I/O links
  - Errors corrected at point of detection
  - Uncorrectable errors reported at source and to all consumers.
  - 71% of all pins are covered by ECC
- Optional 4+1 Parity on RDRAM memory covers:
  - Multi-bit errors
  - Control errors
  - Clock errors
  - RDRAM or channel failures
  - 87% of all pins are covered by ECC or RAID
- Partitions

EV7 Partition Example

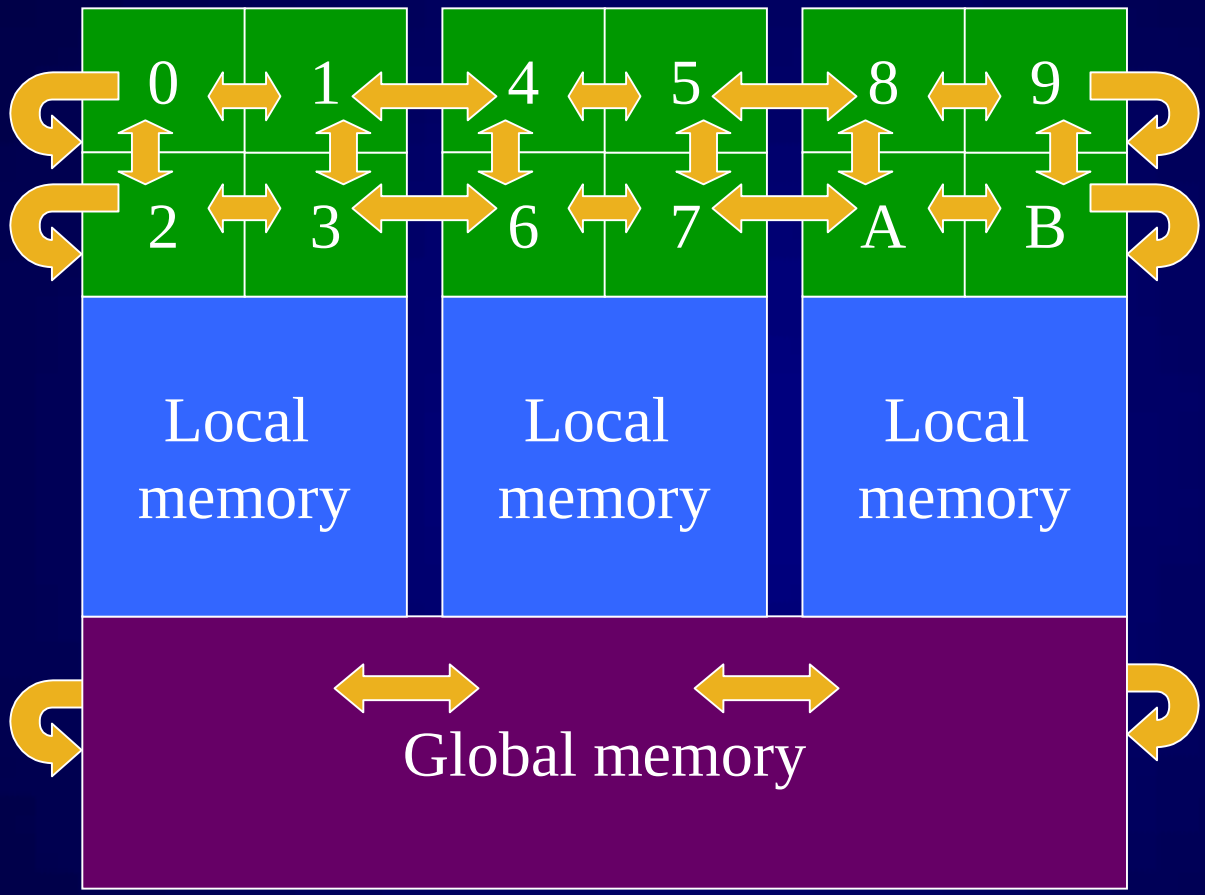

EV7 64P Latency

319	283	247	211	247	283	319	355
283	247	211	175	211	247	283	319
247	211	175	140	175	211	247	283
211	175	140	75	140	175	211	247
247	211	175	140	175	211	247	283
283	247	211	175	211	247	283	319
319	283	247	211	247	283	319	355
355	319	283	247	283	319	355	391

160 GB/s aggregate bandwidth

COMPAQ

Memory Bandwidth (McCalpin Streams)

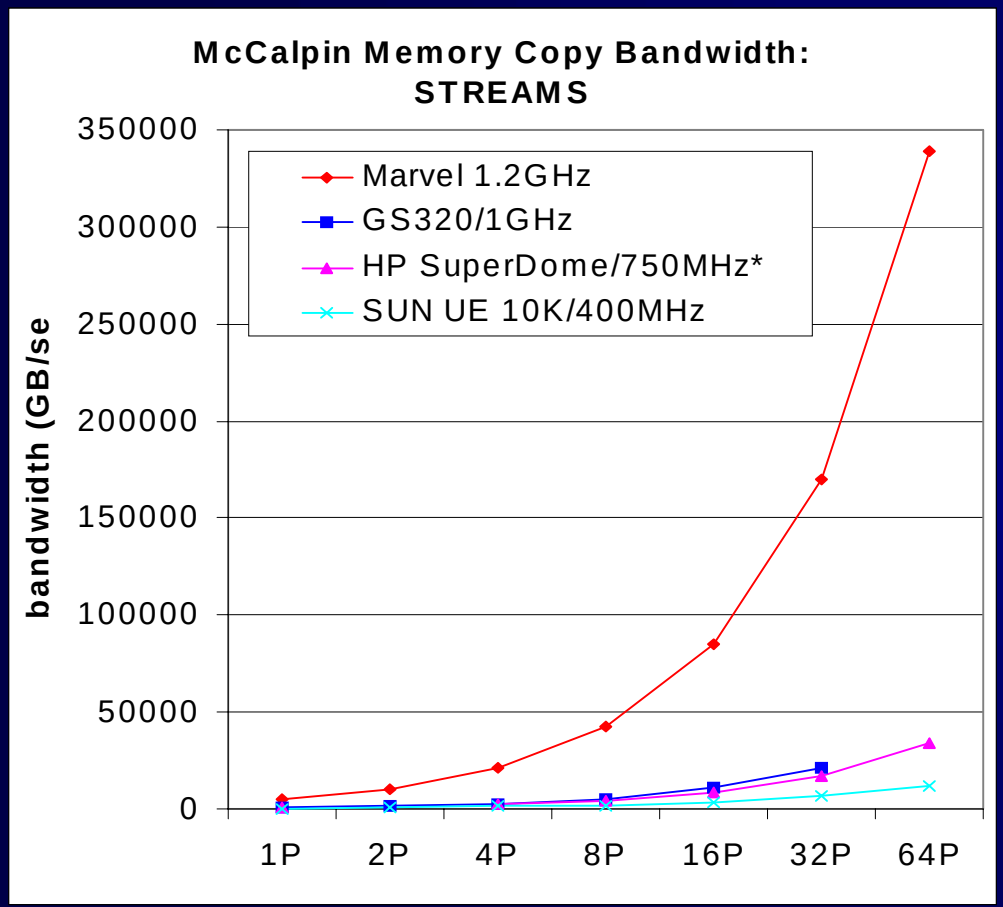

January 4, 2002

COMPAQ

SPEC2000 1-CPU Peak

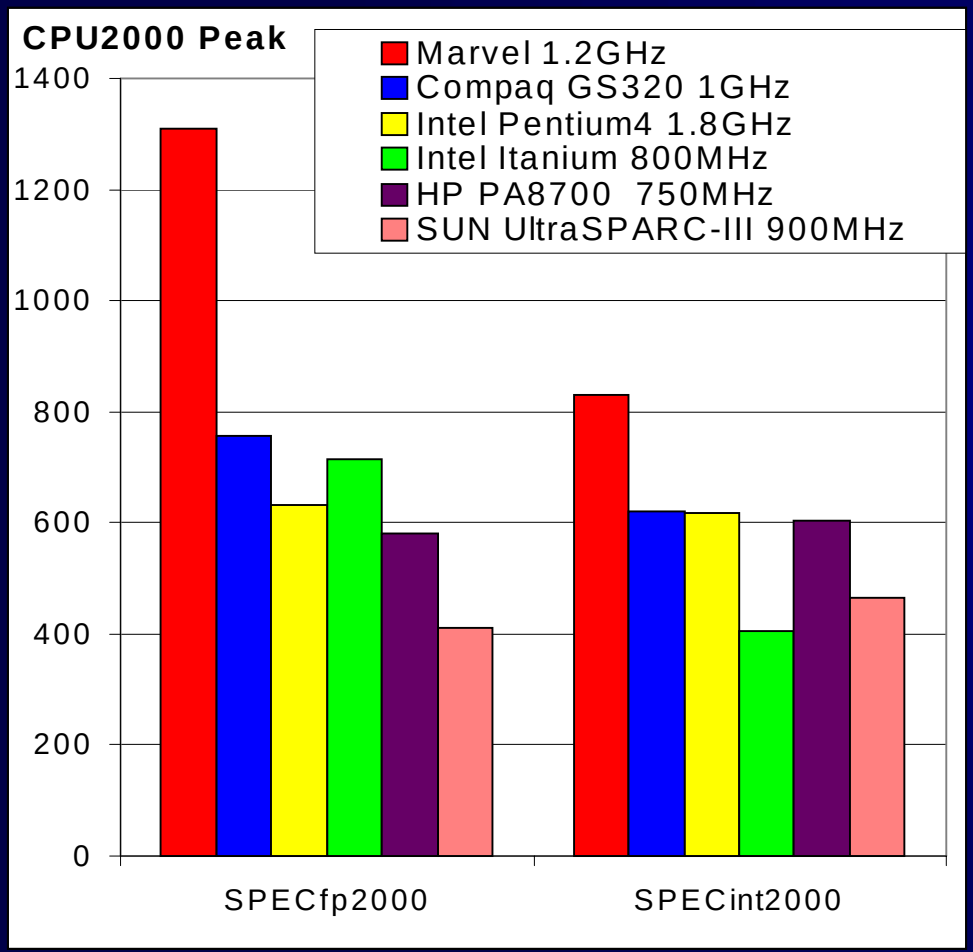

January 4, 2002

SPECint2000 (estimate)

	Itanium	EV7	EV7/Itanium
Frequency	800	1200	1.5
gzip	322	618	1.9
vpr	402	630	1.6
gcc	428	968	2.3
mcf	605	686	1.1
crafty	357	1000	2.8
parser	316	606	1.9
eon	370	982	2.7
perlm bk	320	806	2.5
gap	258	688	2.7
vortex	472	1144	2.4
bzip2	362	786	2.2
twolf	450	946	2.1
SPECint2000	**379**	**804**	**2.1**
int2000/MHz	**0.47**	**0.67**	**1.4**

SPECfp2000 (estimate)

	Itanium	EV7	EV7/Itanium
Frequency	800	1200	1.5
wupwise	469	1050	2.2
swim	1071	3024	2.8
mgrid	871	1110	1.3
applu	542	1378	2.5
mesa	382	1054	2.8
galgel	1377	1722	1.3
art	1638	2418	1.5
equake	565	1316	2.3
facerec	542	1180	2.2
ammp	554	772	1.4
lucas	1020	1612	1.6
fma3d	278	1278	4.6
sixtrack	631	534	0.8
apsi	413	840	2.0
SPECfp2000	**653**	**1253**	**1.9**
fp2000/MHz	**0.82**	**1.04**	**1.3**

Conclusion

- EV68CB upgrade to 1250 MHz
- EV7 will extend the EV6 core with:
  - On chip L2
  - Two memory controllers for directly connected RDRAM memory
  - Glueless SMP
- EV79 will extend the EV7 with
  - Improved clock frequency
  - Increased memory performance with 1066Mb/s RDRAM memory


COMPAQ