Overview of Intelligent RAM (IRAM)

David Patterson, Krste Asanovic, Aaron Bown,
Ben Gribstad, Richard Fromm, Jason Golbus,
Kimberly Keeton, Christoforos Kozyrakis,
Stelianos Perissakis, Randi Thomas,
Noah Treuhaft, John Wawrzynek,
and Katherine Yelick
patterson@cs.berkeley.edu
http://iram.cs.berkeley.edu/
EECS, University of California
Berkeley, CA 94720-1776

Motivations for IRAM

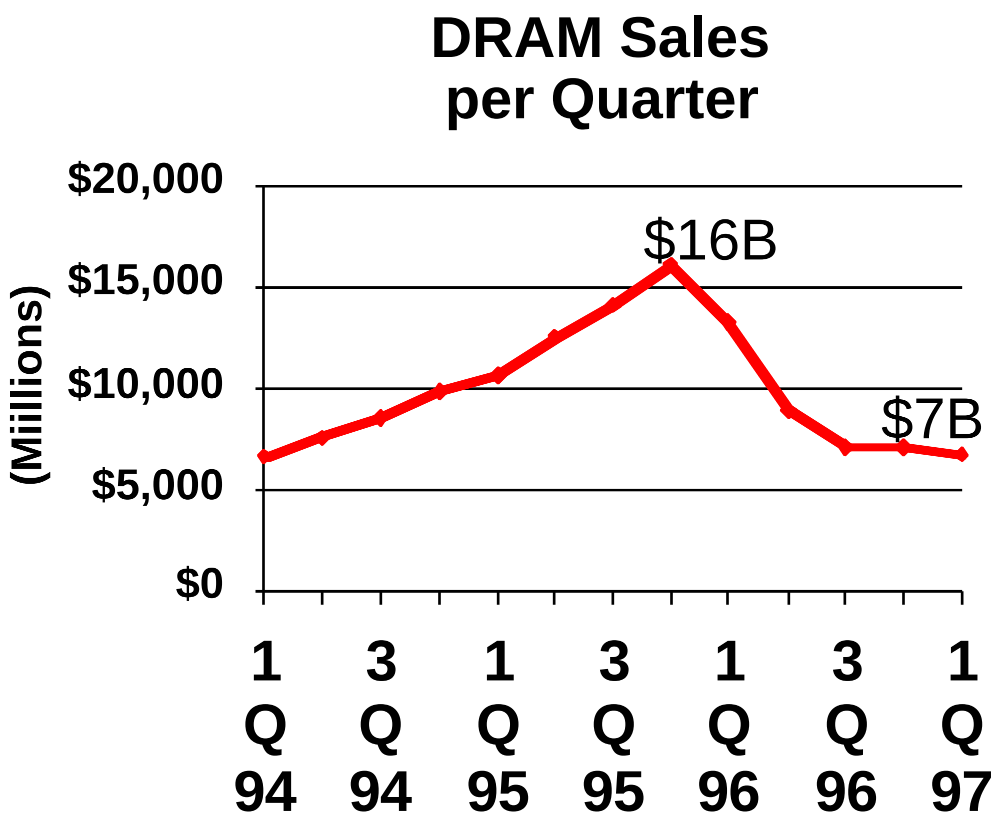

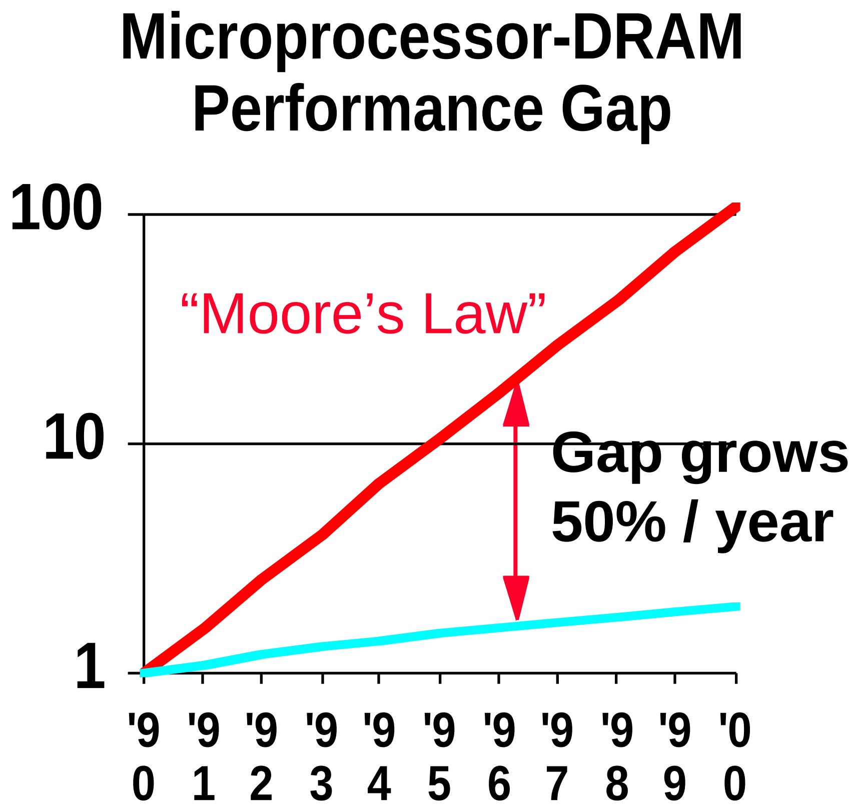

- More profit for DRAM industry?
- MPU close gap?

IRAM Vision Statement

Microprocessor & DRAM on a single chip:

- on-chip memory latency 5-10X, bandwidth 50-100X
- improve energy efficiency 2X-4X (no off-chip bus)
- serial I/O 5-10X v. buses
- smaller board area/volume
- adjustable memory size/width

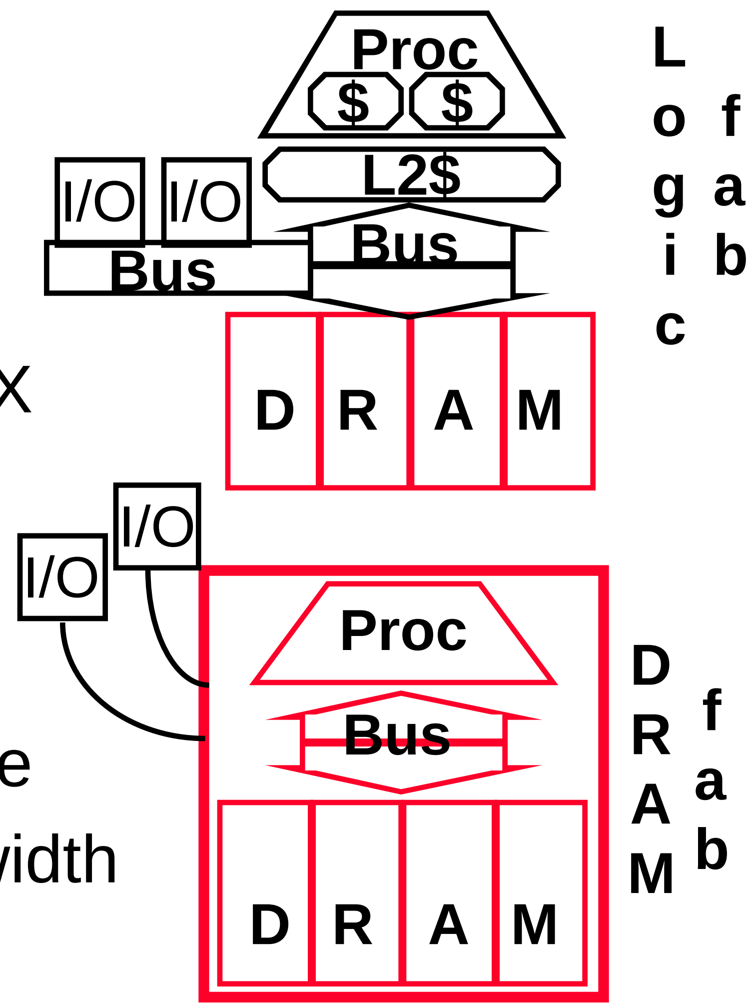

Near-term IRAM Applications

- "Intelligent" Set-top
  - 2.6M Ninetendo 64 (≈ $150) sold in 1st year
  - 4-chip Nintendo ⇒ 1-chip: 3D graphics, sound, fun!
- "Intelligent" Personal Digital Assistant
  - 1.0M PalmPilots (≈ $300) sold in 1st year
  - Super PDA/Smart Phone: speech I/O + "voice" email...

Long-term App: Decision Support?

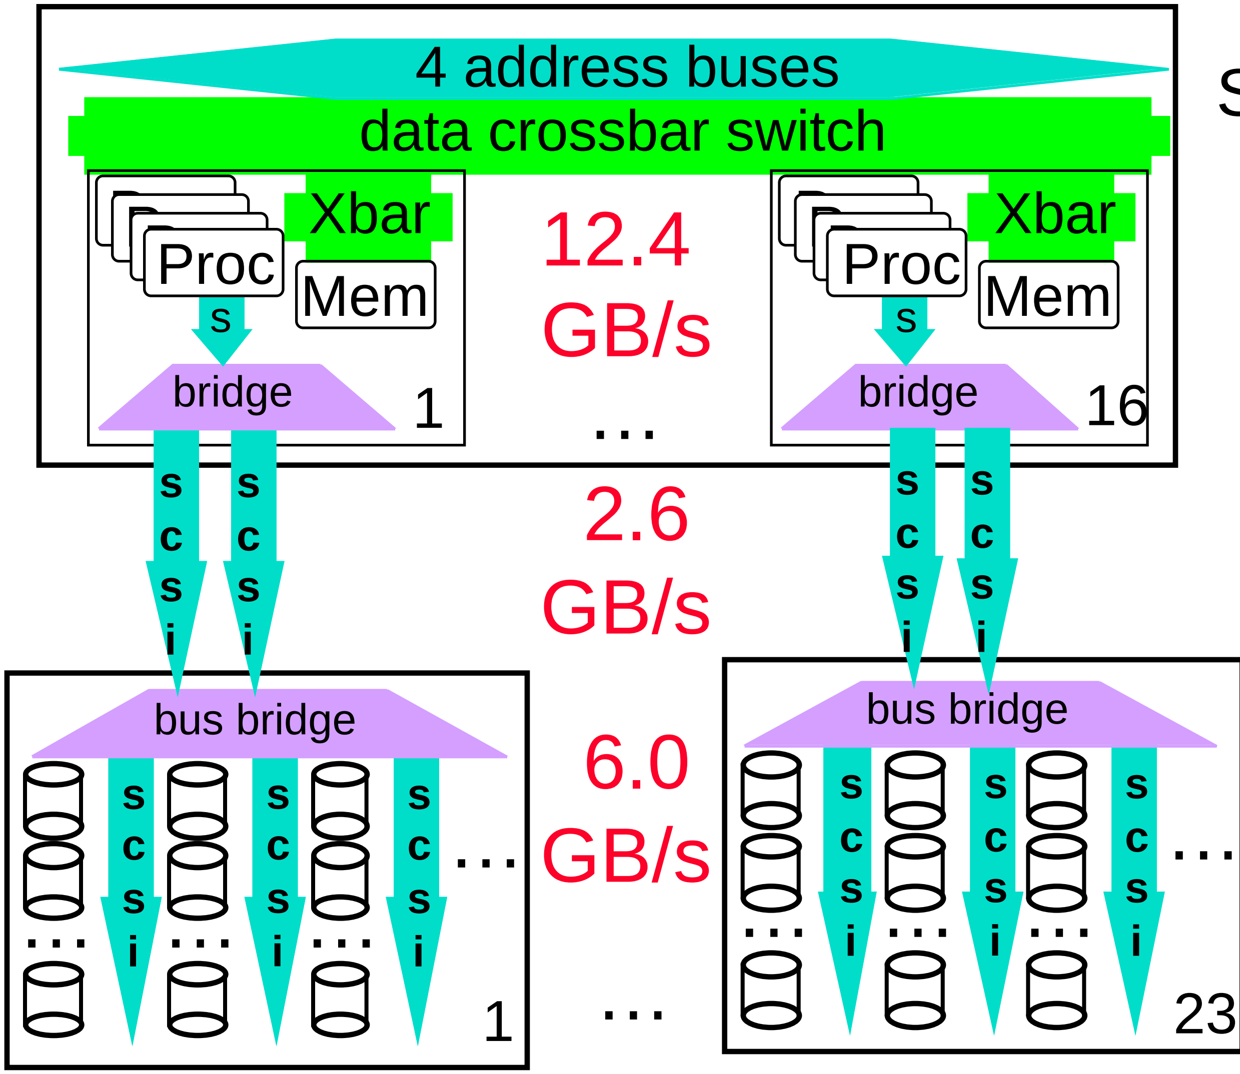

Sun 10000 (Oracle 8):

- TPC-D (1TB) leader
- SMP 64 CPUs, 64GB dram, 603 disks

Disks,encl.	$2,348k
DRAM	$2,328k
Boards,encl.	$983k
CPUs	$912k
Cables,I/O	$139k
Misc	$65k
HW total	$6,775k

IRAM Application Inspiration: Database Demand vs. Processor/DRAM speed

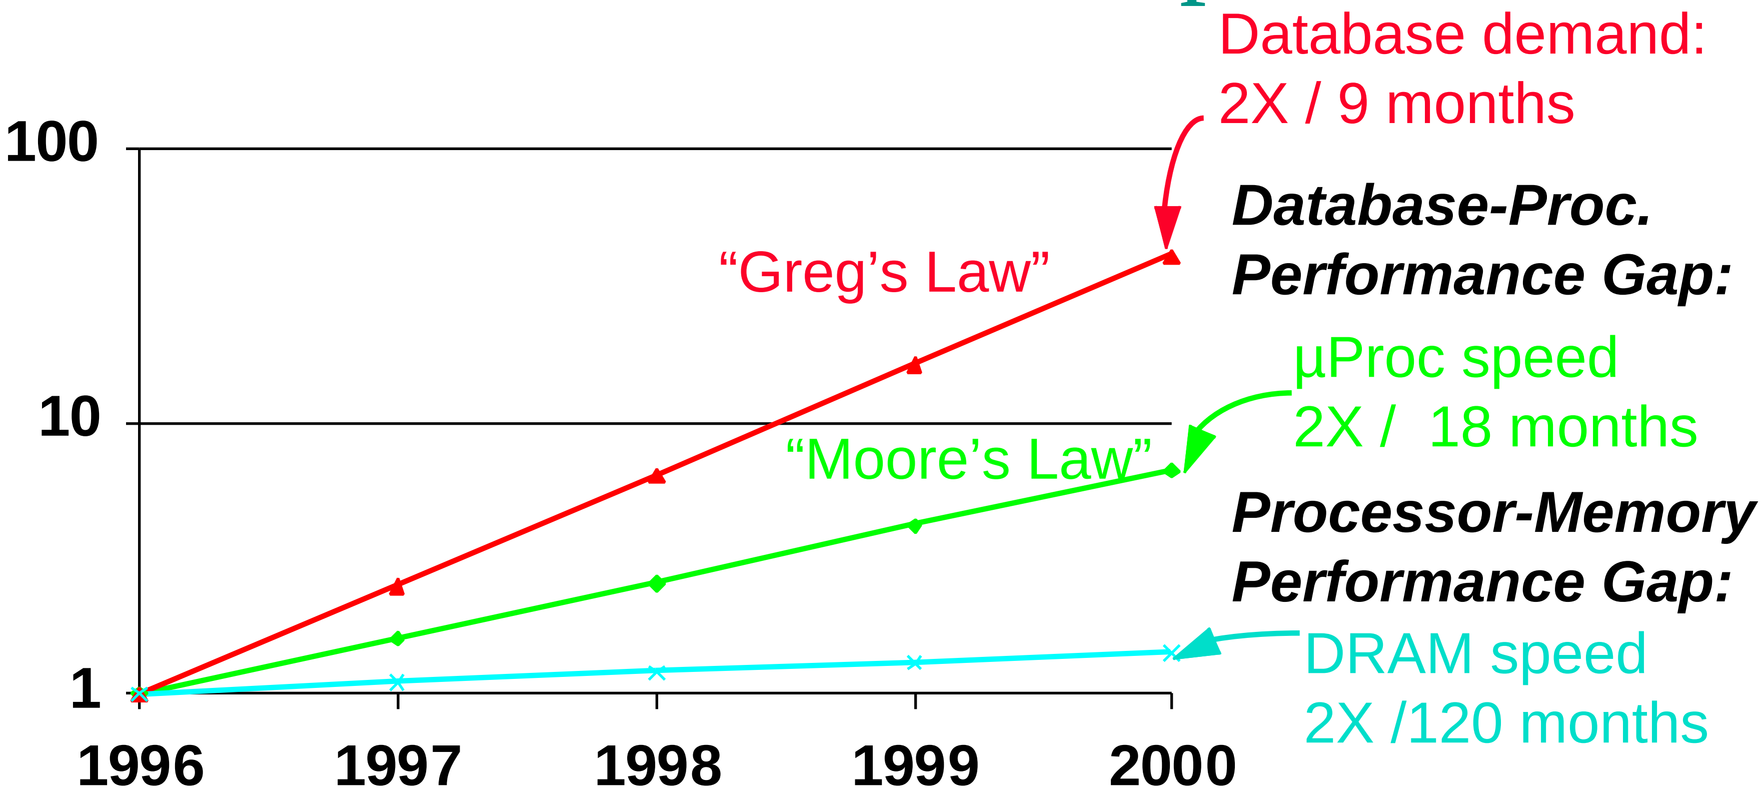

"Intelligent Disk": Scalable Decision Support?

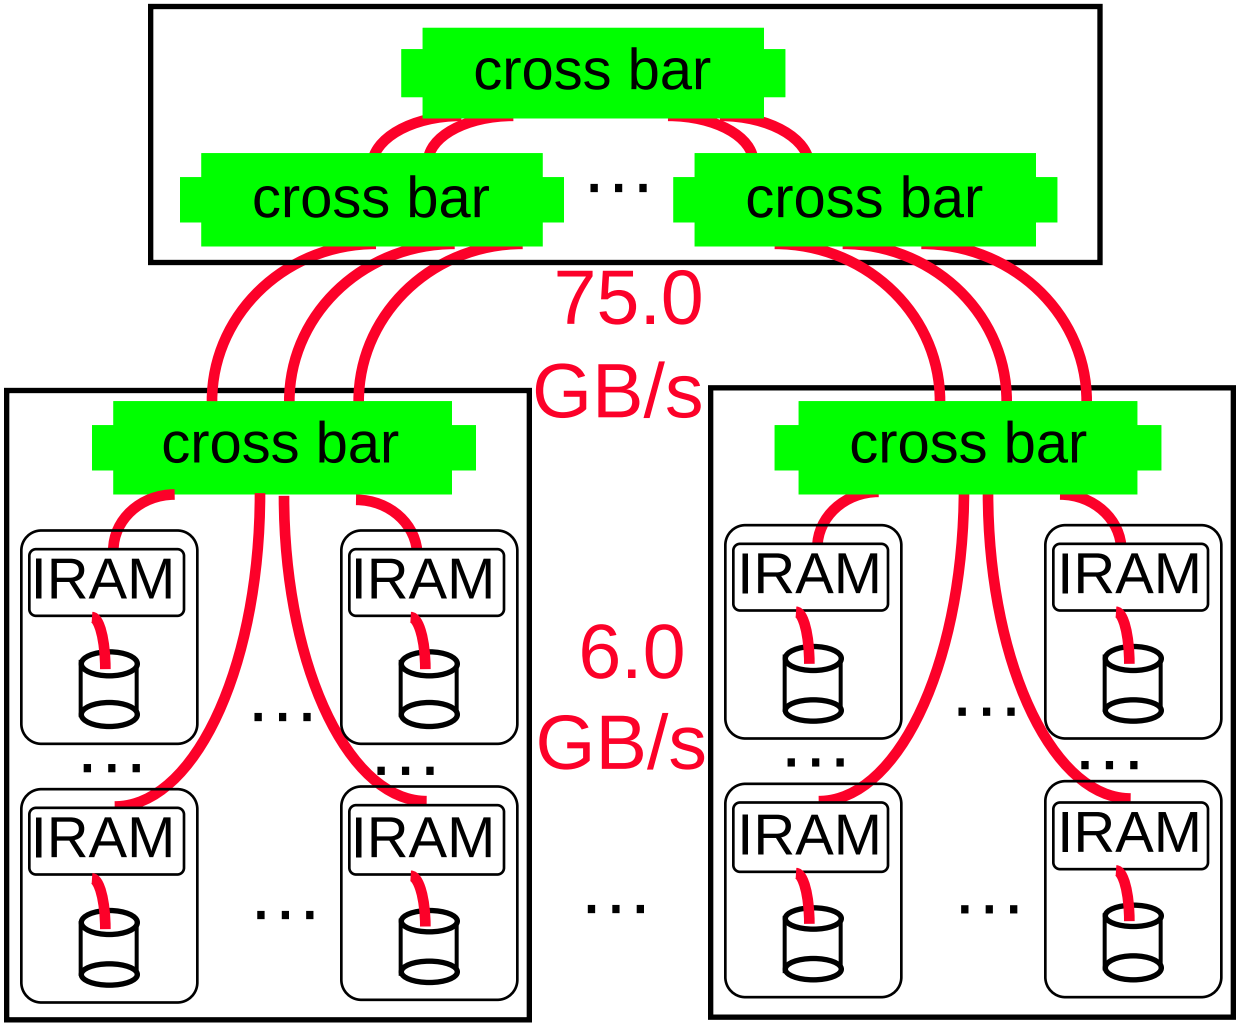

1 IRAM/disk + shared nothing database

– 603 CPUs,
  14GB dram, 603 disks

Disks (market)	$840k
IRAM (@$150)	$90k
Disk encl., racks	$150k
Switches/cables	$150k
Misc	$60k
Subtotal	$1,300k
Markup 2X?	≈ $2,600k

≈1/3 price, 2X-5X perf

New Architecture Directions

- "...wires are not keeping pace with scaling of other features. … In fact, for CMOS processes below 0.25 micron ... *an unacceptably small percentage of the die will be reachable during a single clock cycle*."
- "Architectures that require long-distance, rapid interaction will not scale well ..."
  - "Will Physical Scalability Sabotage Performance Gains?" Matzke, *IEEE Computer* (9/97)

New Architecture Directions

- "…media processing will become the dominant force in computer arch. & microprocessor design."
- "... new media-rich applications... involve signficant real-time processing of continuous media streams, and make heavy use of vectors of packed 8-, 16-, and 32-bit integer and Fl. Pt."
- Needs include high memory BW, high network BW, continuous media data types, real-time response, fine grain parallelism
  - "How Multimedia Workloads Will Change Processor Design", Difendorff & Dubey, *IEEE Computer* (9/97)

Revive Vector Architecture!

- High cost: ≈ $1M / processor?
- Low latency, high BW memory system?
- Compilers?
- Performance?
- Limited to scientific applications?
- Real-time?

- Single-chip CMOS microprocessor/IRAM
- IRAM = low latency, high bandwidth memory
- For sale, mature (>20 years)
- Easy to scale speed with technology (e.g, hides latency)
- Multimedia apps vectorizable too: N*64b,2N*32b,4N*16b,8N*8b
- No caches, no speculation ⇒ repeatable speed as vary input

V-IRAM-2: 0.13 µm, Fast Logic, 1GHz
16 GFLOPS(64b) / 128 GOPS(8b) / 96MB

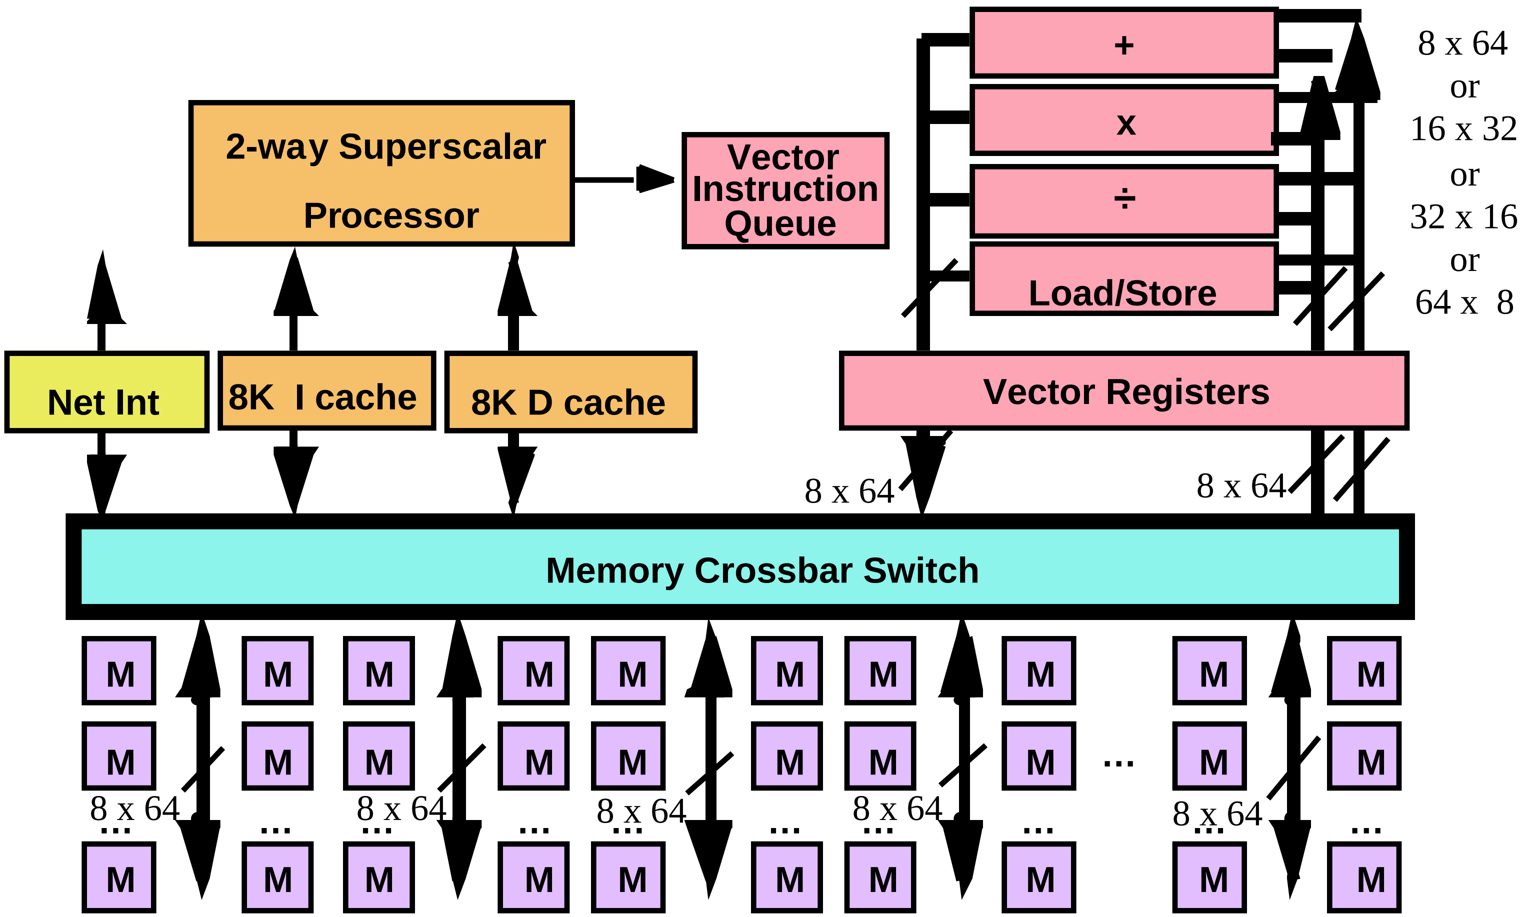

V-IRAM-2 Floorplan

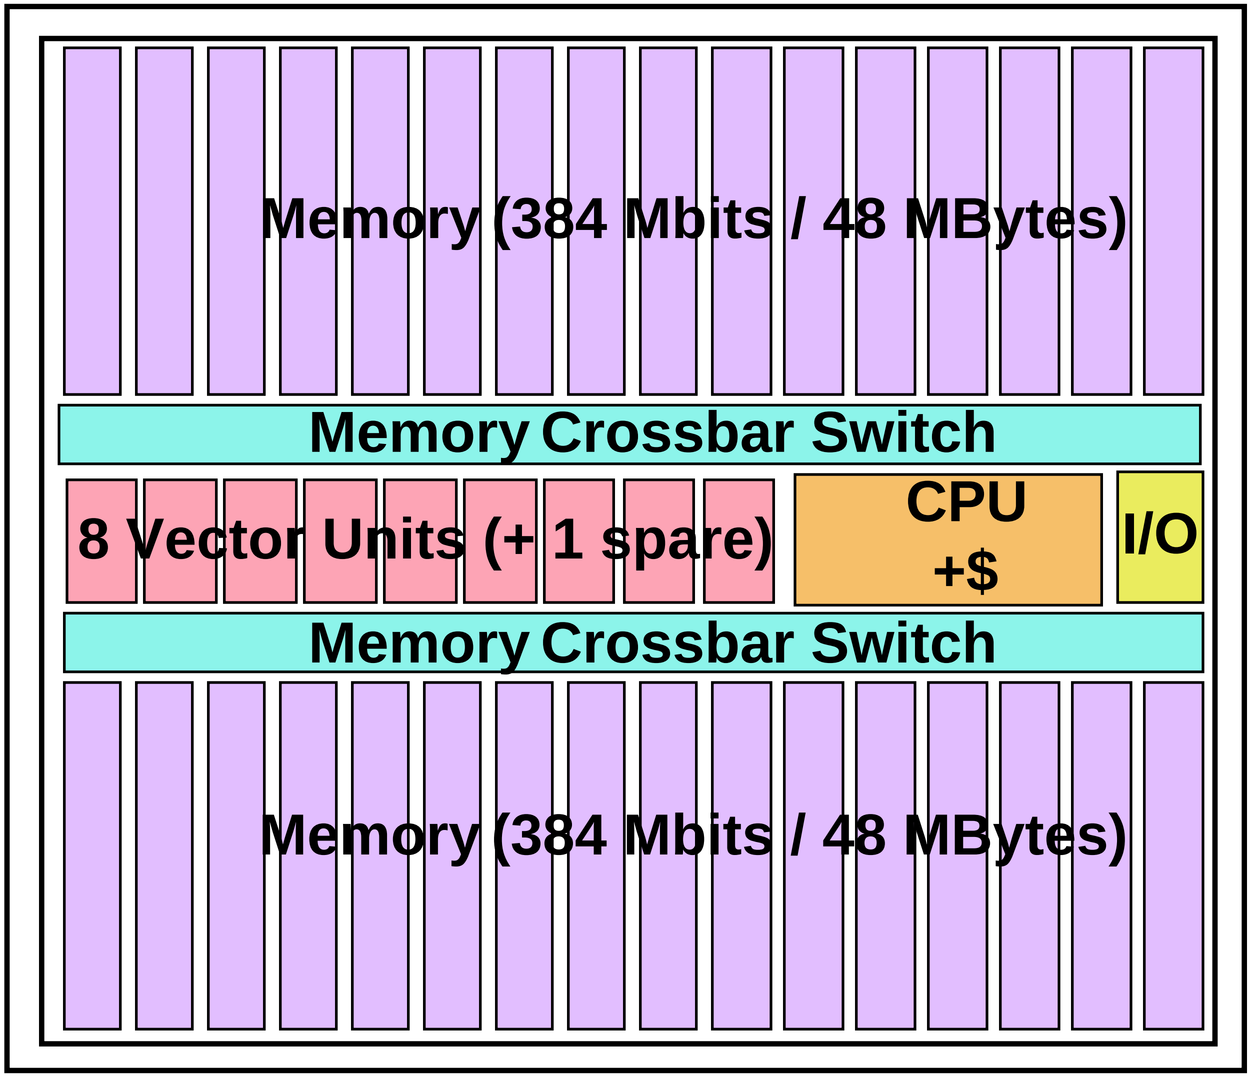

- 0.13 µm, 1 Gbit DRAM
- 1B Xtors: 90% Memory, Xbar, Vector ⇒ regular design
- Spare VU & Memory ⇒ 90% die repairable
- Short signal distance ⇒ speed scales <0.2 µm

IRAM Conclusion

- IRAM potential in bandwidth (memory and I/O), latency, energy, capacity, board area; challenges in power/performance, testing, yield
- V-IRAM can show potential (+compilers,+testing)
- 10X-100X improvements based on technology shipping for 20 years (not JJ, photons, MEMS, ...)
- Potential upheaval in database server industry?
- Potential shift in balance of power in DRAM/ microprocessor industry in 5-7 years?
  - Who ships the most memory?
  - Who ships the most microprocessors?

Interested in Participating?

- Looking for ideas of IRAM enabled apps
- Contact us if you're interested:
  http://iram.cs.berkeley.edu/
  email: patterson@cs.berkeley.edu
- Thanks for advice/support: DARPA, Intel, LG Semiconductor, Neomagic, Samsung, SGI/Cray, Sun Microsystems

Backup Slides

(The following slides are used to help answer questions)

Characterzing IRAM Cost/Performance

- Cost ≈ embedded processor + memory
- Small memory on-chip (25 - 100 MB)
- High vector performance (2 -16 GFLOPS)
- High multimedia performance (4 - 64 GOPS)
- Low latency main memory (15 - 30ns)
- High BW main memory (50 - 200 GB/sec)
- High BW I/O (0.5 - 2 GB/sec via N serial lines)
  - Integrated CPU/cache/memory with high memory BW ideal for fast serial I/O

IRAM Challenges

- Chip
  - Good performance and reasonable power?
  - Speed, area, power, yield, cost in DRAM process?
  - BW/Latency oriented DRAM tradeoffs?
  - Testing time of IRAM vs DRAM vs microprocessor?
  - Reconfigurable logic to make IRAM more generic?
- Architecture
  - How to turn high memory bandwidth into performance for real applications?
  - Extensible IRAM: Large program/data solution? (e.g., external DRAM, clusters, CC-NUMA, ...)

Why IRAM now? Lower risk than before

- Faster Logic + DRAM available now/soon?
- DRAM manufacturers now willing to listen
  - Before not interested, so early IRAM = SRAM
- Past efforts memory limited ⇒ multiple chips ⇒ 1st solve the unsolved (parallel processing)
  - Gigabit DRAM ⇒ ≈100 MB; OK for many apps?
- Systems headed to 2 chips: CPU + memory
- Embedded apps leverage energy efficiency, adjustable mem. capacity, smaller board area ⇒ OK market v. desktop (55M 32b RISC '96)

IRAM
not a new idea

Stone, ‘70 “Logic-in memory”
Barron, ‘78 “Transputer”
Dally, ‘90 “J-machine”
Patterson, ‘90 panel session
Kogge, ‘94 “Execube”

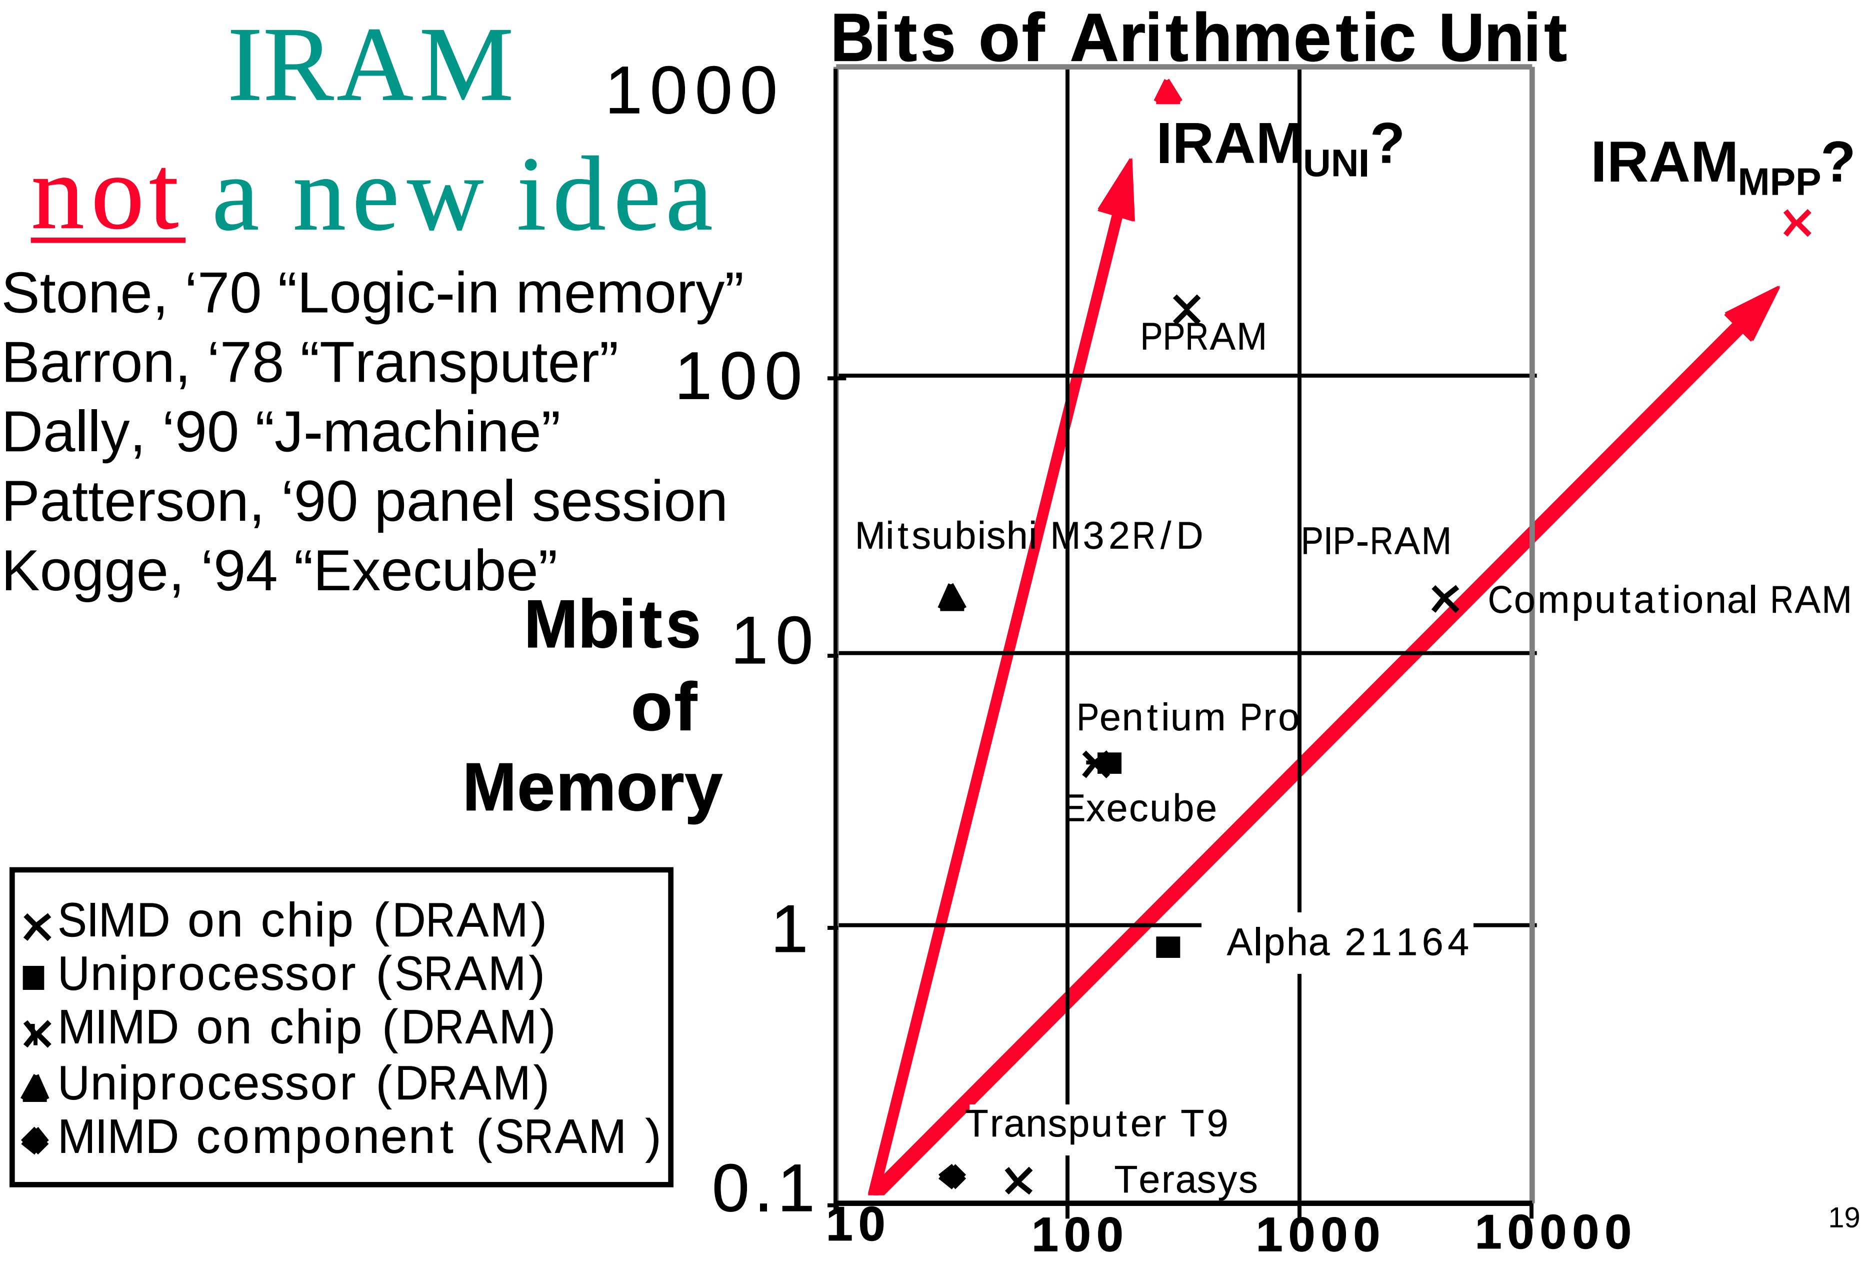

Goal for Vector IRAM Generations

- V-IRAM-1 (≈2000)
- 256 Mbit generation (0.20)
- Die size = 256 Mb DRAM die
- 1.5 - 2.0 v logic, 2-10 watts
- 100 - 500 MHz
- 4 64-bit pipes/lanes
- 1-4 GFLOPS(64b)/6-32G (8b)
- 30 - 50 GB/sec Mem. BW
- 24 MB capacity + DRAM bus
- Several fast serial I/O

- V-IRAM-2 (≈2003)
- 1 Gbit generation (0.13)
- Die size = 1 Gb DRAM die
- 1.0 - 1.5 v logic, 2-10 watts
- 200 - 1000 MHz
- 8 64-bit pipes/lanes
- 2-16 GFLOPS/24-128G
- 100 - 200 GB/sec Mem. BW
- 96 MB cap. + DRAM bus
- Many fast serial I/O

Simple v. Complex Case Study

MIPS MPUs	R5000	R10000	10k/5k
▪ Clock Rate	200 MHz	195 MHz	1.0x
▪ On-Chip Caches	32K/32K	32K/32K	1.0x
▪ Instructions/Cycle	1(+ FP)	4	4.0x
▪ Pipe stages	5	5-7	1.2x
▪ Model	In-order	Out-of-order	---
▪ SPECint_base95	5.7	8.8	1.6x
▪ Die Size (mm^2)	84	298	3.5x
– without cache, TLB	32	205	6.3x
▪ Development (man yr.)	60	300	5.0x

Processor-Memory Performance Gap “Tax”

Processor	% Area (≈*cost*)	%Transistors (≈*power*)
■ Alpha 21164	37%	77%
■ StrongArm SA110	61%	94%
■ Pentium Pro	64%	88%

– 2 dies per package: Proc/I$/D$ + L2$

- Caches have no inherent value, only try to close performance gap

How to get Low Power, High Clock rate IRAM?

- Digital Strong ARM 110 (1996): 2.1M Xtors
  - 160 MHz @ 1.5 v = 184 “MIPS” < 0.5 W
  - 215 MHz @ 2.0 v = 245 “MIPS” < 1.0 W
- Start with Alpha 21064 @ 3.5v, 26 W
  - Vdd reduction ⇒ 5.3X ⇒ 4.9 W
  - Reduce functions ⇒ 3.0X ⇒ 1.6 W
  - Scale process ⇒ 2.0X ⇒ 0.8 W
  - Clock load ⇒ 1.3X ⇒ 0.6 W
  - Clock rate ⇒ 1.2X ⇒ 0.5 W
- 6/97: 233 MHz, 268 MIPS, 0.36W typ., $49