

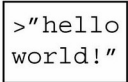


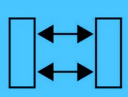
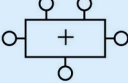
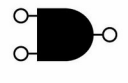
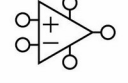
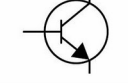

Chapter 7

Digital Design and Computer Architecture, 2nd Edition

David Money Harris and Sarah L. Harris

Chapter 7 :: Topics

- Introduction (done)
- Performance Analysis (done)
- Single-Cycle Processor (done)
- Multicycle Processor (done)
- Pipelined Processor (done)
- Exceptions (done)
- Advanced Microarchitecture (now)

Application Software	
Operating Systems	
Architecture	
Micro-architecture	
Logic	
Digital Circuits	
Analog Circuits	
Devices	
Physics	

Advanced Microarchitecture

- Deep Pipelining
- Branch Prediction
- Superscalar Processors
- Out of Order Processors
- Register Renaming
- SIMD
- Multithreading
- Multiprocessors

Deep Pipelining

- 10-20 stages typical
- Number of stages limited by:
 - Pipeline hazards
 - Sequencing overhead
 - Power
 - Cost

Branch Prediction

- Ideal pipelined processor: $CPI = 1$
- Branch misprediction increases CPI
- **Static branch prediction:**
 - Check direction of branch (forward or backward)
 - If backward, predict taken
 - Else, predict not taken
- **Dynamic branch prediction:**
 - Keep history of last (several hundred) branches in *branch target buffer*, record:
 - Branch destination
 - Whether branch was taken

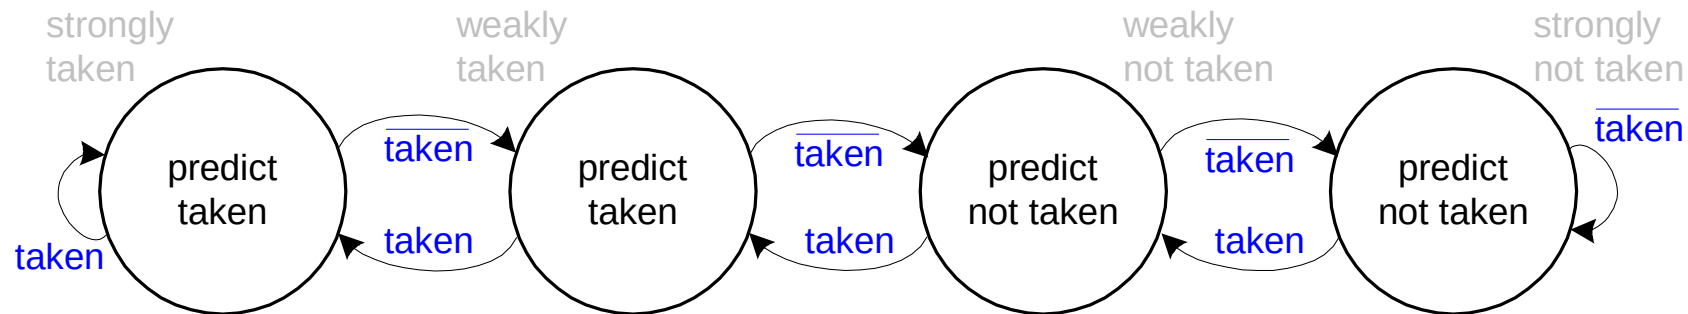
Branch Prediction Example

```
add  $s1, $0, $0      # sum = 0
add  $s0, $0, $0      # i   = 0
addi $t0, $0, 10      # $t0 = 10
for:
    beq $s0, $t0, done # if i == 10, branch
    add $s1, $s1, $s0  # sum = sum + i
    addi $s0, $s0, 1   # increment i
    j   for
done:
```

1-Bit Branch Predictor

- Remembers whether branch was taken the last time and does the same thing
- Mispredicts first and last branch of loop
- If the loop is entered again, the "taken" is incorrect.

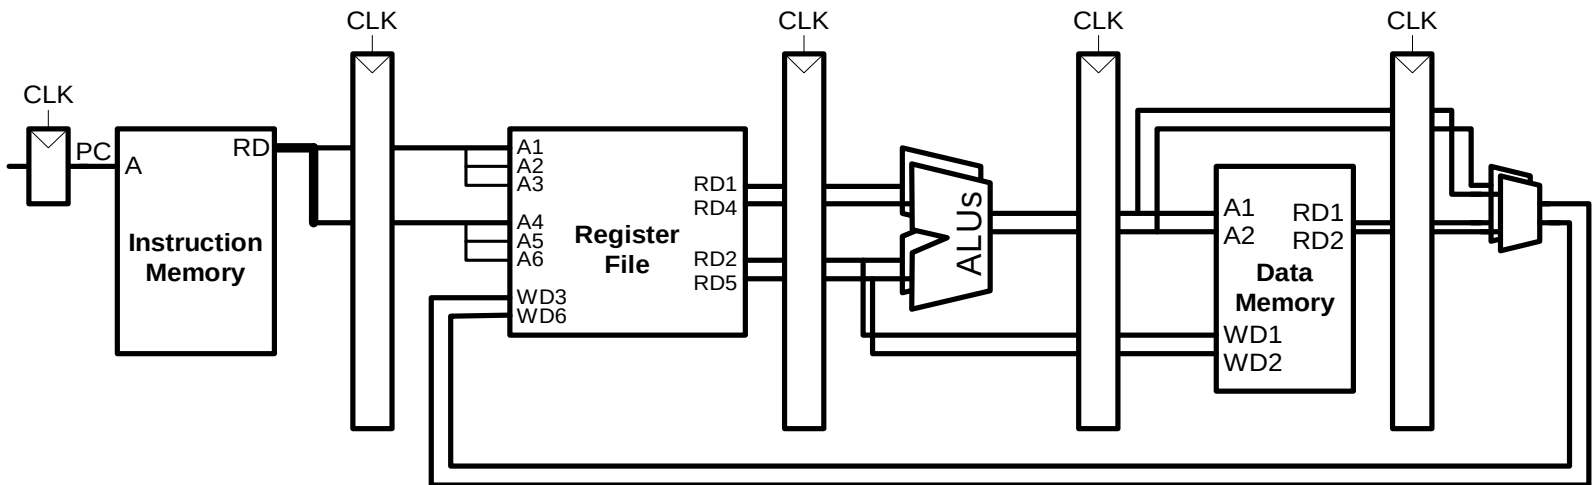
2-Bit Branch Predictor



- When loop is exited, "strongly not taken" state transitions to "weakly not taken".
- If loop is entered again, it correctly predicts "taken" and the state returns to "strongly not taken".
- Only mispredicts last branch of loop

Superscalar

- Multiple copies of datapath execute multiple instructions at once
- Dependencies make it tricky to issue multiple instructions at once



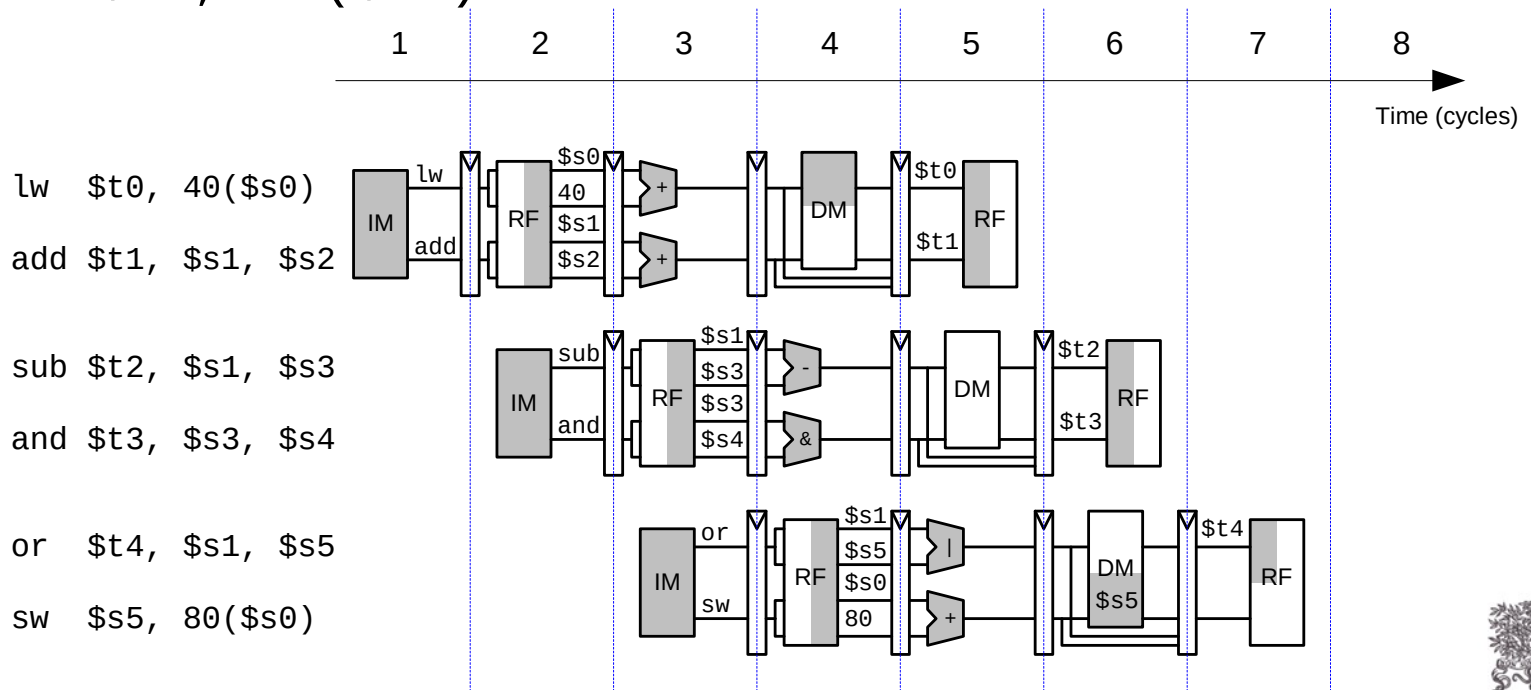
Superscalar Example with No Dependencies

```

lw  $t0, 40($s0)
add $t1, $s1, $s2
sub $t2, $s1, $s3
and $t3, $s3, $s4
or  $t4, $s1, $s5
sw  $s5, 80($s0)
    
```

Ideal IPC: 2

Actual IPC: 2



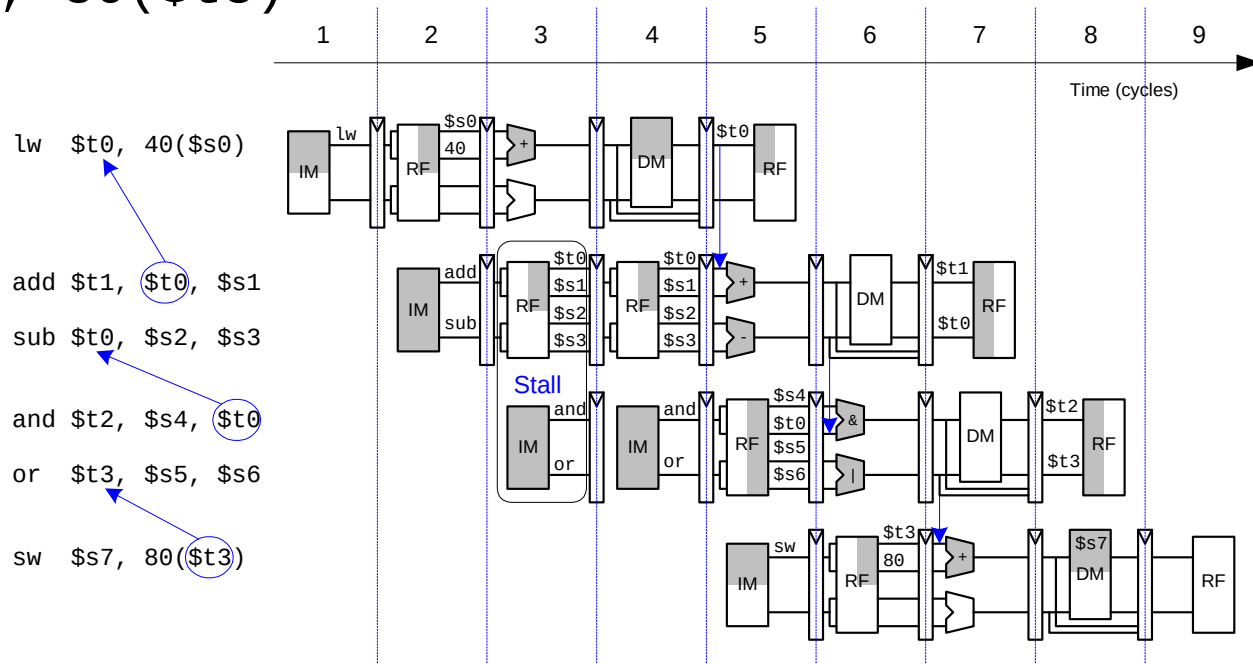
Superscalar Example with Dependencies

```

lw   $t0, 40($s0)
add  $t1, $t0, $s1
sub  $t0, $s2, $s3
and  $t2, $s4, $t0
or   $t3, $s5, $s6
sw   $s7, 80($t3)
    
```

Ideal IPC: 2

Actual IPC: 6/5 = 1.17



Out of Order Processor

- Looks ahead across multiple instructions
- Issues as many instructions as possible at once
- Issues instructions out of order (as long as no dependencies)
- **Dependencies:**
 - **RAW** (read after write): one instruction writes, later instruction reads a register
 - **WAR** (write after read): one instruction reads, later instruction writes a register
 - **WAW** (write after write): one instruction writes, later instruction writes a register

Out of Order Processor

- **Instruction level parallelism (ILP):** number of instruction that can be issued simultaneously (average < 3)
- **Scoreboard:** table that keeps track of:
 - Instructions waiting to issue
 - Available functional units
 - Dependencies

Out of Order Processor Example

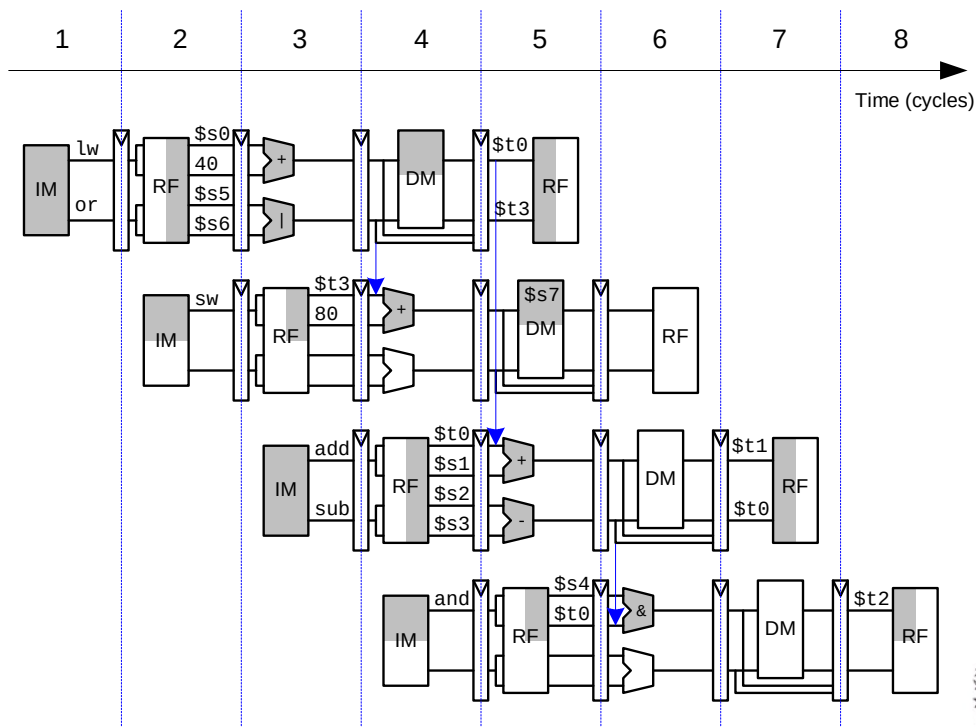
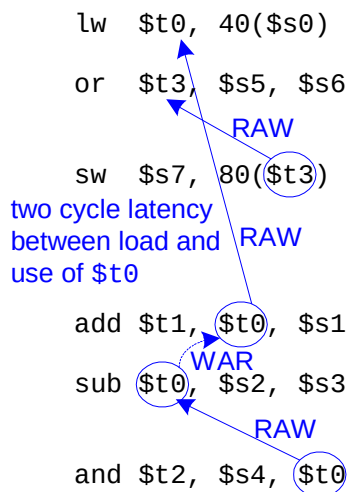
```

lw  $t0, 40($s0)
add $t1, $t0, $s1
sub $t0, $s2, $s3
and $t2, $s4, $t0
or  $t3, $s5, $s6
sw  $s7, 80($t3)

```

Ideal IPC: 2

Actual IPC: 6/4 = 1.5



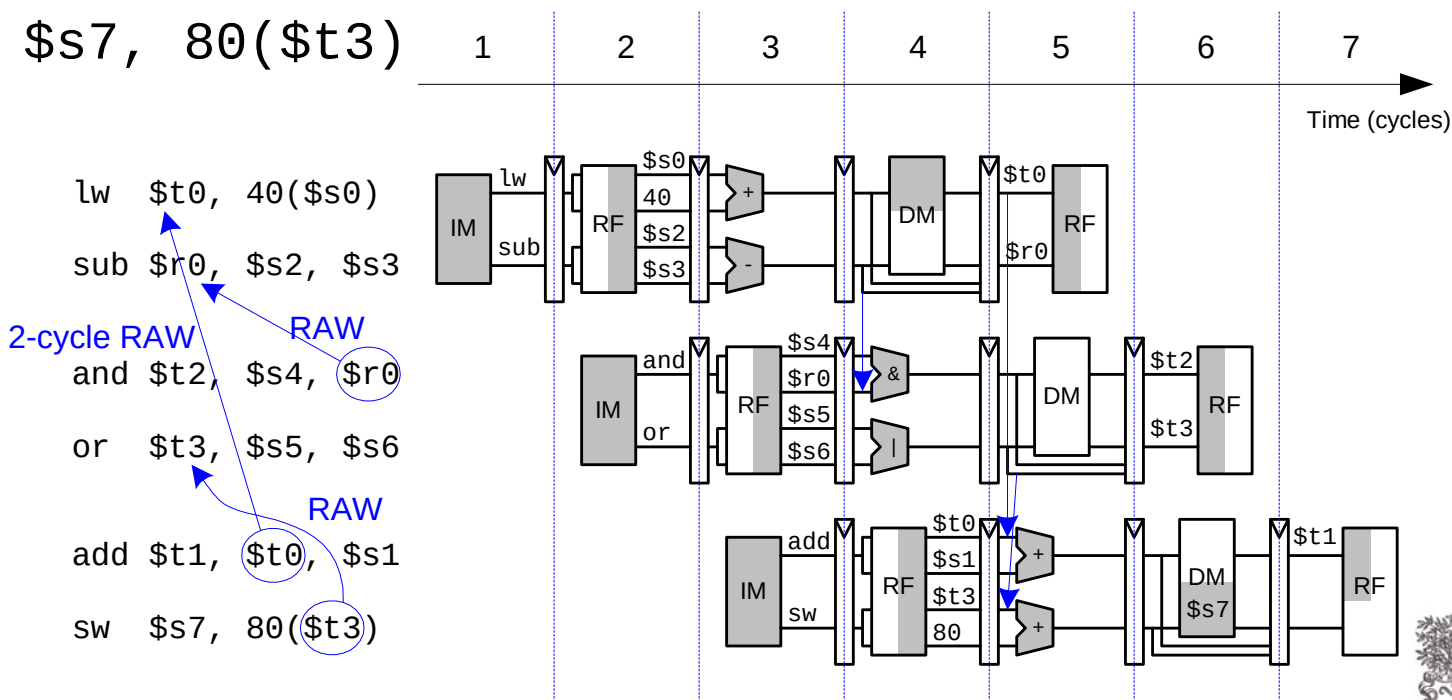
Register Renaming

```

lw $t0, 40($s0)
add $t1, $t0, $s1
sub $t0, $s2, $s3
and $t2, $s4, $t0
or $t3, $s5, $s6
sw $s7, 80($t3)
    
```

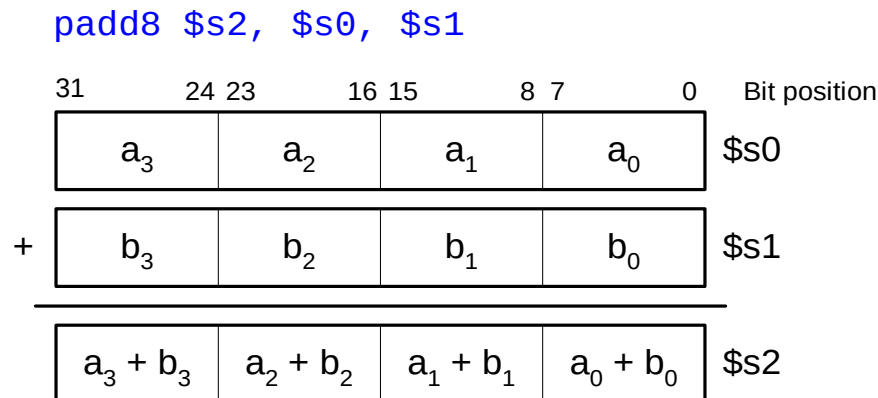
Ideal IPC: 2

Actual IPC: $6/3 = 2$



SIMD

- Single Instruction Multiple Data (SIMD)
 - Single instruction acts on multiple pieces of data at once
 - Common application: graphics
 - Perform short arithmetic operations (also called *packed arithmetic*)
- For example, add four 8-bit elements



Advanced Architecture Techniques

- **Multithreading**
 - Wordprocessor: thread for typing, spell checking, printing
- **Multiprocessors**
 - Multiple processors (cores) on a single chip

Threading: Definitions

- **Process:** program running on a computer
 - Multiple processes can run at once: e.g., surfing Web, playing music, writing a paper
- **Thread:** part of a program
 - Each process has multiple threads: e.g., a word processor may have threads for typing (perhaps for multiple collaborators), spell checking, printing

Threads in Conventional Processor

- One thread runs at once
- When one thread stalls (for example, waiting for memory):
 - Architectural state of that thread stored
 - Architectural state of waiting thread loaded into processor and it runs
 - Called **context switching**
- Appears to user like all threads running simultaneously

Multithreading

- Multiple copies of architectural state
- Multiple threads **active** at once:
 - When one thread stalls, another runs immediately
 - If one thread can't keep all execution units busy, another thread can use them
- Does not increase instruction-level parallelism (ILP) of single thread, but increases throughput

Intel calls this “hypertreading”

Multiprocessors

- Multiple processors (cores) with a method of communication between them
- Types:
 - **Homogeneous:** multiple cores with shared memory
 - **Heterogeneous:** separate cores for different tasks (for example, DSP, GPU, and CPU in cell phone)
 - **Clusters:** each core has own memory system



Chapter 7

Multicores,
Multiprocessors, and
Clusters

Introduction

- Goal: connecting multiple computers to get higher performance
 - Multiprocessors
 - Scalability, availability, power efficiency
- Job-level (process-level) parallelism
 - High throughput for independent jobs
- Parallel processing program
 - Single program run on multiple processors
- Multicore microprocessors
 - Chips with multiple processors (cores)

Hardware and Software

- Hardware
 - Serial: e.g., Pentium 4
 - Parallel: e.g., 16 cores on an Intel i7-10700 CPU
- Software
 - Sequential: e.g., matrix multiplication
 - Concurrent: e.g., operating system
- Sequential/concurrent software can run on serial/parallel hardware
 - Challenge: making effective use of parallel hardware, aka "load balancing"

What We've Already Covered

- §2.11: Parallelism and Instructions
 - Synchronization
- §3.6: Parallelism and Computer Arithmetic
 - Associativity
- §4.10: Parallelism and Advanced Instruction-Level Parallelism
- §5.8: Parallelism and Memory Hierarchies
 - Cache Coherence
- §6.9: Parallelism and I/O:
 - Redundant Arrays of Inexpensive Disks

Parallel Programming

- Parallel software is the problem
- Need to get significant performance improvement
 - Otherwise, just use a faster uniprocessor, since it's easier!
- Difficulties
 - Partitioning
 - Coordination
 - Communications overhead

Amdahl's Law (Review)

- Sequential part can limit speedup
- $T = T_p + T_s$
("p": parallelizeable, "s": sequential)
- $T_{new} = T_p / N + T_s$ for N processors
- let $f = T_p / T_{old}$, the fraction of time in the original program that can be parallelized
- then $T_p = f T$ and $T_s = (1 - f) T$
- we derive

$$speedup \equiv \frac{T}{T_{new}} = \frac{T}{(1 - f) T + \frac{f}{N} T} = \frac{1}{(1 - f) + \frac{f}{N}}$$

Amdahl's Law (Review)

- Sequential part can limit speedup
- Example: $N = 100$ processors, want $90\times$ speedup?
- We solve:

$$\frac{1}{(1-f) + \frac{f}{N}} = 90$$

- Yields $f \approx 0.999$
- Need sequential part to be 0.1% of original time

Scaling Example

- Workload: (a) sum of 10 scalars, and (b) 10×10 matrix sum
 - Speed up from 10 to 100 processors
- Single processor: Time = $(10 + 100) \times t_{\text{add}}$
- 10 processors
 - Time = $10 \times t_{\text{add}} + 100/10 \times t_{\text{add}} = 20 \times t_{\text{add}}$
 - Speedup = $110/20 = 5.5$ (55% of potential)
- 100 processors
 - Time = $10 \times t_{\text{add}} + 100/100 \times t_{\text{add}} = 11 \times t_{\text{add}}$
 - Speedup = $110/11 = 10$ (10% of potential)
- Assumes load can be balanced across processors

Scaling Example (cont)

- What if matrix size is 100×100 ?
- Single processor: Time = $(10 + 10000) \times t_{\text{add}}$
- 10 processors
 - Time = $10 \times t_{\text{add}} + 10000/10 \times t_{\text{add}} = 1010 \times t_{\text{add}}$
 - Speedup = $10010/1010 = 9.9$ (99% of potential)
- 100 processors
 - Time = $10 \times t_{\text{add}} + 10000/100 \times t_{\text{add}} = 110 \times t_{\text{add}}$
 - Speedup = $10010/110 = 91$ (91% of potential)
- Assuming load balanced

Strong vs Weak Scaling

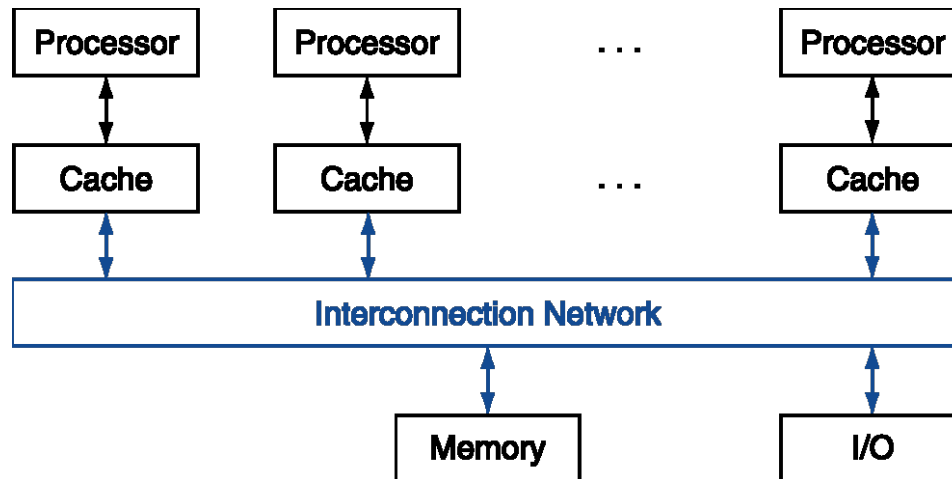
- Strong scaling: problem size fixed (speedup is easy)
 - As in example
- Weak scaling: problem size proportional to number of processors (speedup is hard)
 - 10 processors, 10×10 matrix
 - Time = $20 \times t_{\text{add}}$
 - 100 processors, 32×32 matrix
 - Time = $10 \times t_{\text{add}} + 1000/100 \times t_{\text{add}} = 20 \times t_{\text{add}}$
 - Constant performance in this example

Strong vs Weak Scaling

- Which to use?
 - depends on the application:
What are you trying to prove to whom?
 - given problem size M and P processors
- Strong scaling (solve existing problems faster)
 - memory/processor $\sim M/P$
- Weak scaling (take on bigger problems)
 - memory/processor $\sim M$
 - may be affected by memory bottleneck

Shared Memory

- SMP: shared memory multiprocessor
 - Hardware provides single physical address space for all processors
 - Synchronize shared variables using locks
 - Memory access time
 - UMA (uniform) vs. NUMA (nonuniform)



Example: Sum Reduction

- Sum 100,000 numbers on 100 processor UMA
 - Each processor has ID: $0 \leq P_n \leq 99$
 - Partition 1000 numbers per processor
 - Initial summation on each processor

```
sum[Pn] = 0;  
  for (i = 1000*Pn;  
       i < 1000*(Pn+1); i = i + 1)  
    sum[Pn] = sum[Pn] + A[i];
```

- Now need to add these partial sums
 - Reduction: divide and conquer
 - Half the processors add pairs, then quarter, ...
 - Need to synchronize between reduction steps

Example: Sum Reduction

```
half = 100; /* private */  
repeat
```

```
  synch();
```

```
  if (half % 2 != 0 && Pn == 0)
```

```
    sum[0] = sum[0] + sum[half-1];
```

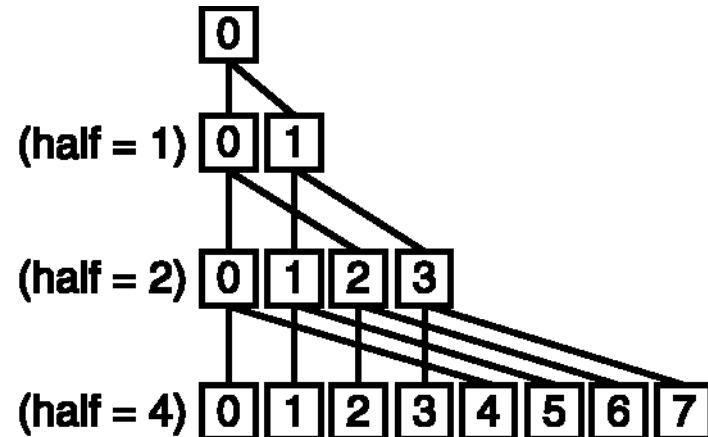
```
    /* Conditional sum needed when half is odd;
```

```
       Processor0 gets missing element */
```

```
  half = half / 2; /* dividing line on who sums */
```

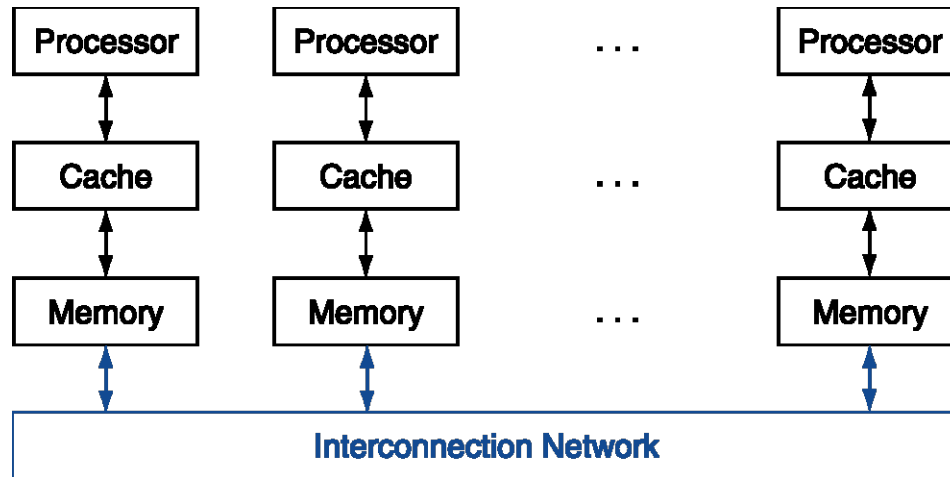
```
  if (Pn < half) sum[Pn] = sum[Pn] + sum[Pn+half];
```

```
until (half == 1);
```



Message Passing

- Each processor has private physical address space
- Hardware sends/receives messages between processors



Loosely Coupled Clusters

- Network of independent computers
 - Each has private memory and OS
 - Connected using I/O system
 - E.g., Ethernet/switch, Internet
- Suitable for applications with independent tasks
 - Web servers, databases, simulations, ...
- High availability, scalable, affordable
- Problems
 - Administration cost (prefer virtual machines)
 - Low interconnect bandwidth
 - c.f. processor/memory bandwidth on an SMP

Sum Reduction (Again)

- Sum 100,000 on 100 processors
- First distribute 1000 numbers to each
 - Then do partial sums

```
sum = 0;
```

```
for (i = 0; i < 1000; i = i + 1)
```

```
    sum = sum + AN[i];
```

- Reduction
 - Half the processors send, other half receive and add
 - The quarter send, quarter receive and add, ...

Sum Reduction (Again)

- Given send() and receive() operations

```
limit = 100; half = 100; /* 100 processors */
repeat
    half = (half+1)/2; /* send vs. receive
                        dividing line */
    if (Pn >= half && Pn < limit)
        send(Pn - half, sum);
    if (Pn < (limit/2))
        sum = sum + receive();
    limit = half; /* upper limit of senders */
until (half == 1); /* exit with final sum */
```

- Send/receive also provide synchronization
- Assumes send/receive take similar time to addition

Grid Computing

- Separate computers interconnected by long-haul networks
 - E.g., Internet connections
 - Work units farmed out, results sent back
- Can make use of idle time on PCs
 - E.g., SETI@home, World Community Grid

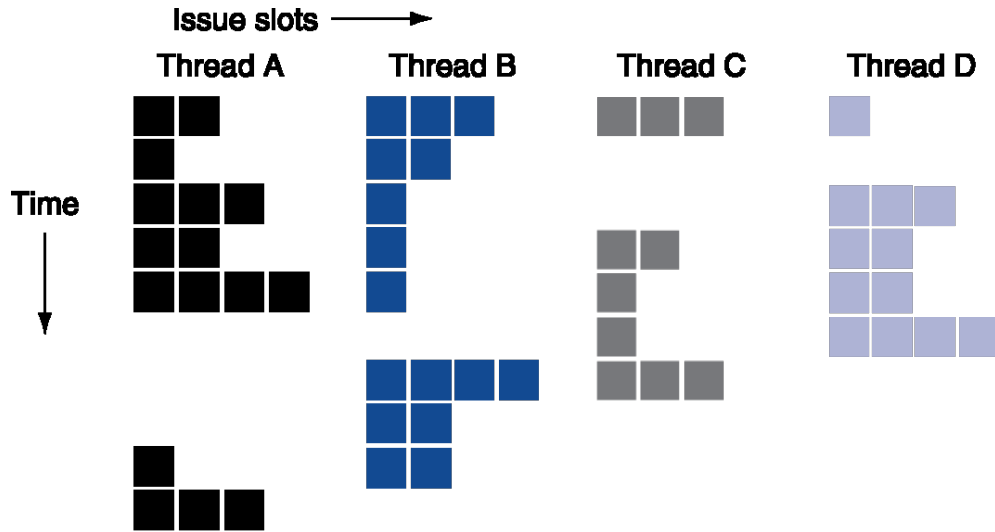
Multithreading

- Performing multiple threads of execution in parallel
 - Replicate registers, PC, etc.
 - Fast switching between threads
- Fine-grain multithreading
 - Switch threads after each cycle
 - Interleave instruction execution
 - If one thread stalls, others are executed
- Coarse-grain multithreading
 - Only switch on long stall (e.g., L2-cache miss)
 - Simplifies hardware, but doesn't hide short stalls (eg, data hazards)

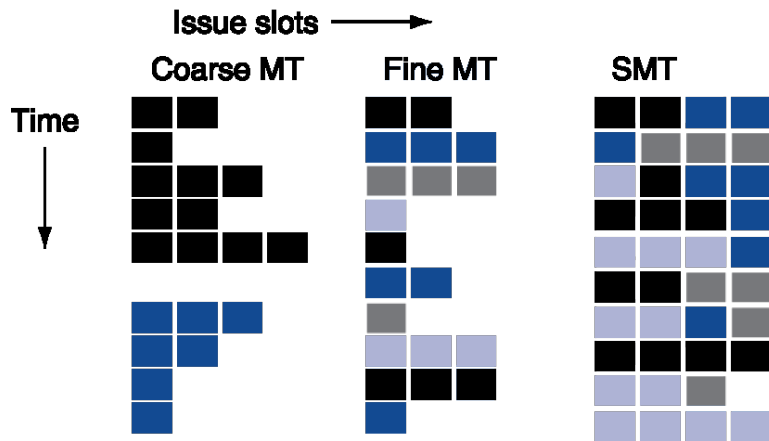
Simultaneous Multithreading

- In multiple-issue dynamically scheduled processor
 - Schedule instructions from multiple threads
 - Instructions from independent threads execute when function units are available
 - Within threads, dependencies handled by scheduling and register renaming
- Example: Intel Pentium-4 HT
 - Two threads: duplicated registers, shared function units and caches

Multithreading Example



superscalar, but
no multithreading



superscalar with
various kinds of
multithreading

Future of Multithreading

- Will it survive? In what form?
- Power considerations \Rightarrow simplified microarchitectures
 - Simpler forms of multithreading
- Tolerating cache-miss latency
 - Thread switch may be most effective
- Multiple simple cores might share resources more effectively

Instruction and Data Streams

- An alternate classification

		Data Streams	
		Single	Multiple
Instruction Streams	Single	SISD: Intel Pentium 4	SIMD: SSE instructions of x86
	Multiple	MISD: No examples today	MIMD: Intel Xeon e5345

- SPMD: Single Program Multiple Data
 - A parallel program on a MIMD computer
 - Conditional code for different processors

SIMD

- Operate elementwise on vectors of data
 - E.g., MMX and SSE instructions in x86
 - Multiple data elements in 128-bit wide registers
- All processors execute the same instruction at the same time
 - Each with different data address, etc.
- Simplifies synchronization
- Reduced instruction control hardware
- Works best for highly data-parallel applications

Vector Processors

- Highly pipelined function units
- Stream data from/to vector registers to units
 - Data collected from memory into registers
 - Results stored from registers to memory
- Example: Vector extension to MIPS
 - 32×64 -element registers (64-bit elements)
 - Vector instructions
 - lv, sv: load/store vector
 - addv.d: add vectors of double
 - addvs.d: add scalar to each element of vector of double
- Significantly reduces instruction-fetch bandwidth

Example: DAXPY ($Y = a \times X + Y$)

■ Conventional MIPS code

```
l.d    $f0, a($sp)      ;load scalar a
      addiu $r4, $s0, #512 ;upper bound of what to load
loop: l.d    $f2, 0($s0)  ;load x(i)
      mul.d  $f2, $f2, $f0 ;a × x(i)
      l.d    $f4, 0($s1)  ;load y(i)
      add.d  $f4, $f4, $f2 ;a × x(i) + y(i)
      s.d    $f4, 0($s1)  ;store into y(i)
      addiu $s0, $s0, #8   ;increment index to x
      addiu $s1, $s1, #8   ;increment index to y
      subu  $t0, $r4, $s0  ;compute bound
      bne   $t0, $zero, loop ;check if done
```

■ Vector MIPS code

```
l.d    $f0, a($sp)      ;load scalar a
lv     $v1, 0($s0)      ;load vector x
mulvs.d $v2, $v1, $f0   ;vector-scalar multiply
lv     $v3, 0($s1)      ;load vector y
addv.d $v4, $v2, $v3    ;add y to product
sv     $v4, 0($s1)      ;store the result
```

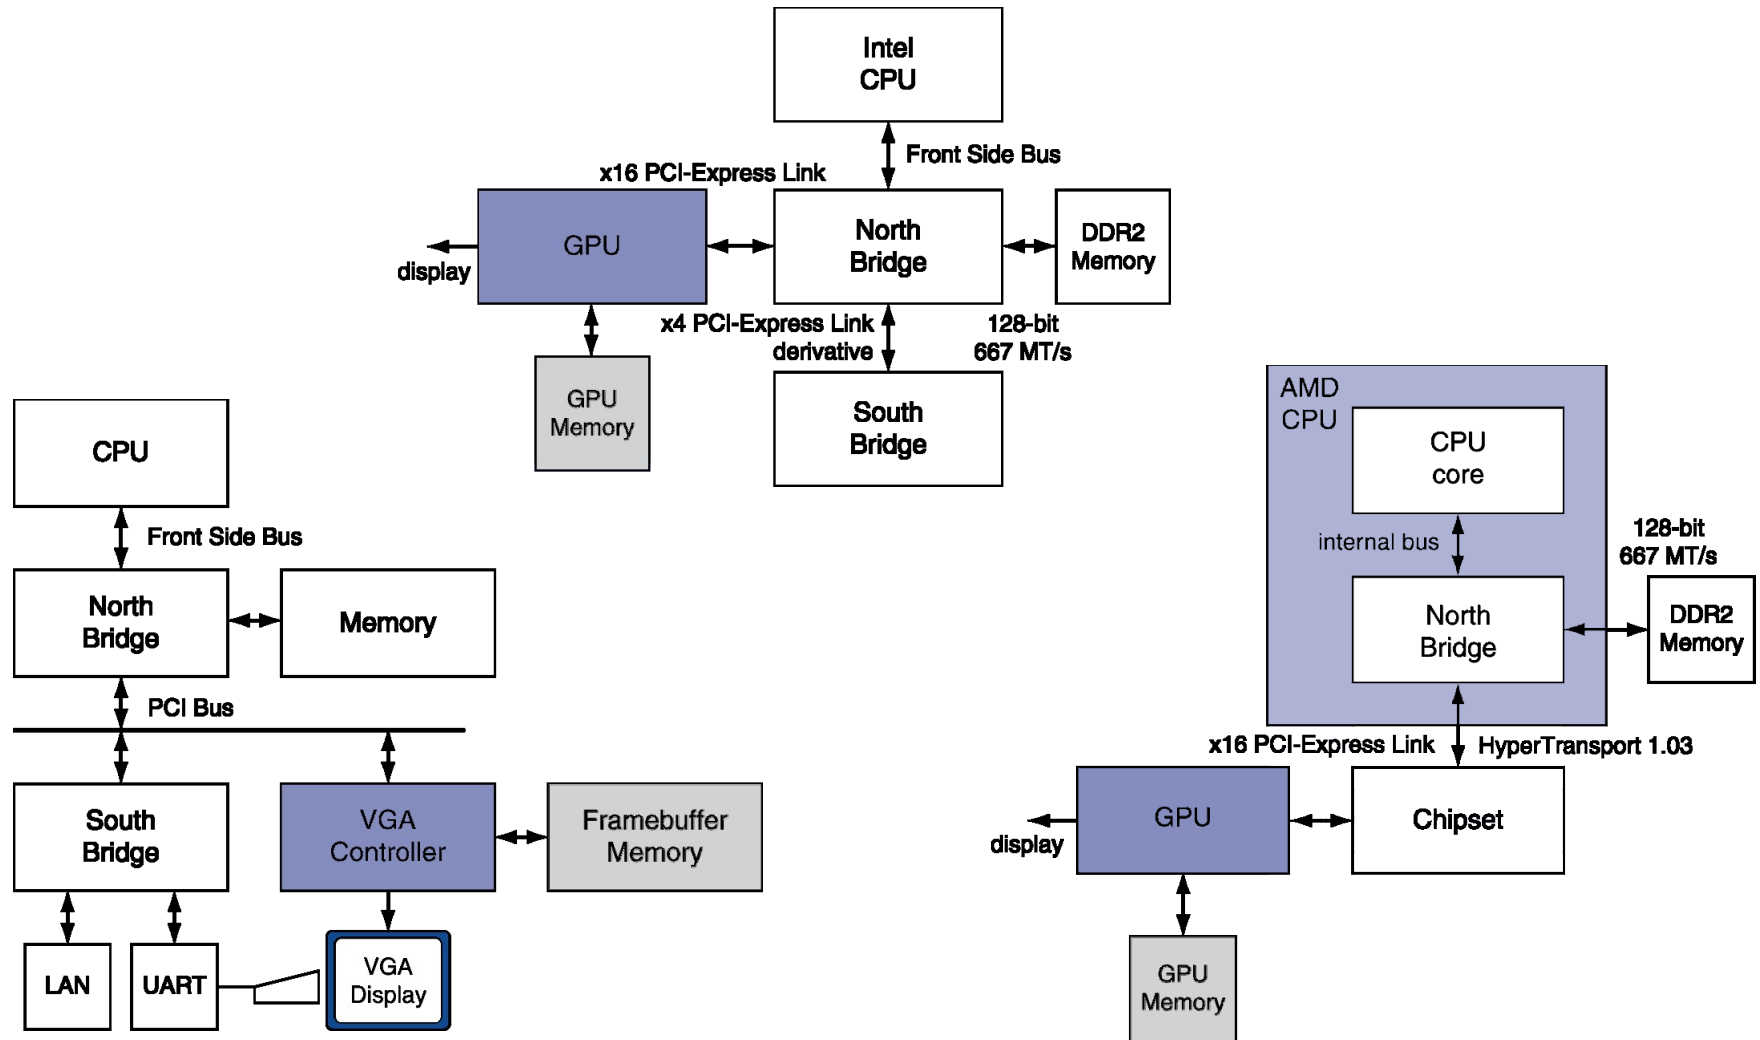

Vector vs. Scalar

- Vector architectures and compilers
 - Simplify data-parallel programming
 - Explicit statement of absence of loop-carried dependences
 - Reduced checking in hardware
 - Regular access patterns benefit from interleaved and burst memory
 - Avoid control hazards by avoiding loops
- More general than ad-hoc media extensions (such as MMX, SSE)
 - Better match with compiler technology

History of GPUs

- Early video cards
 - Frame buffer memory with address generation for video output
- 3D graphics processing
 - Originally high-end computers (e.g., SGI)
 - Moore's Law ▲ lower cost, higher density
 - 3D graphics cards for PCs and game consoles
- Graphics Processing Units
 - Processors oriented to 3D graphics tasks
 - Vertex/pixel processing, shading, texture mapping, rasterization

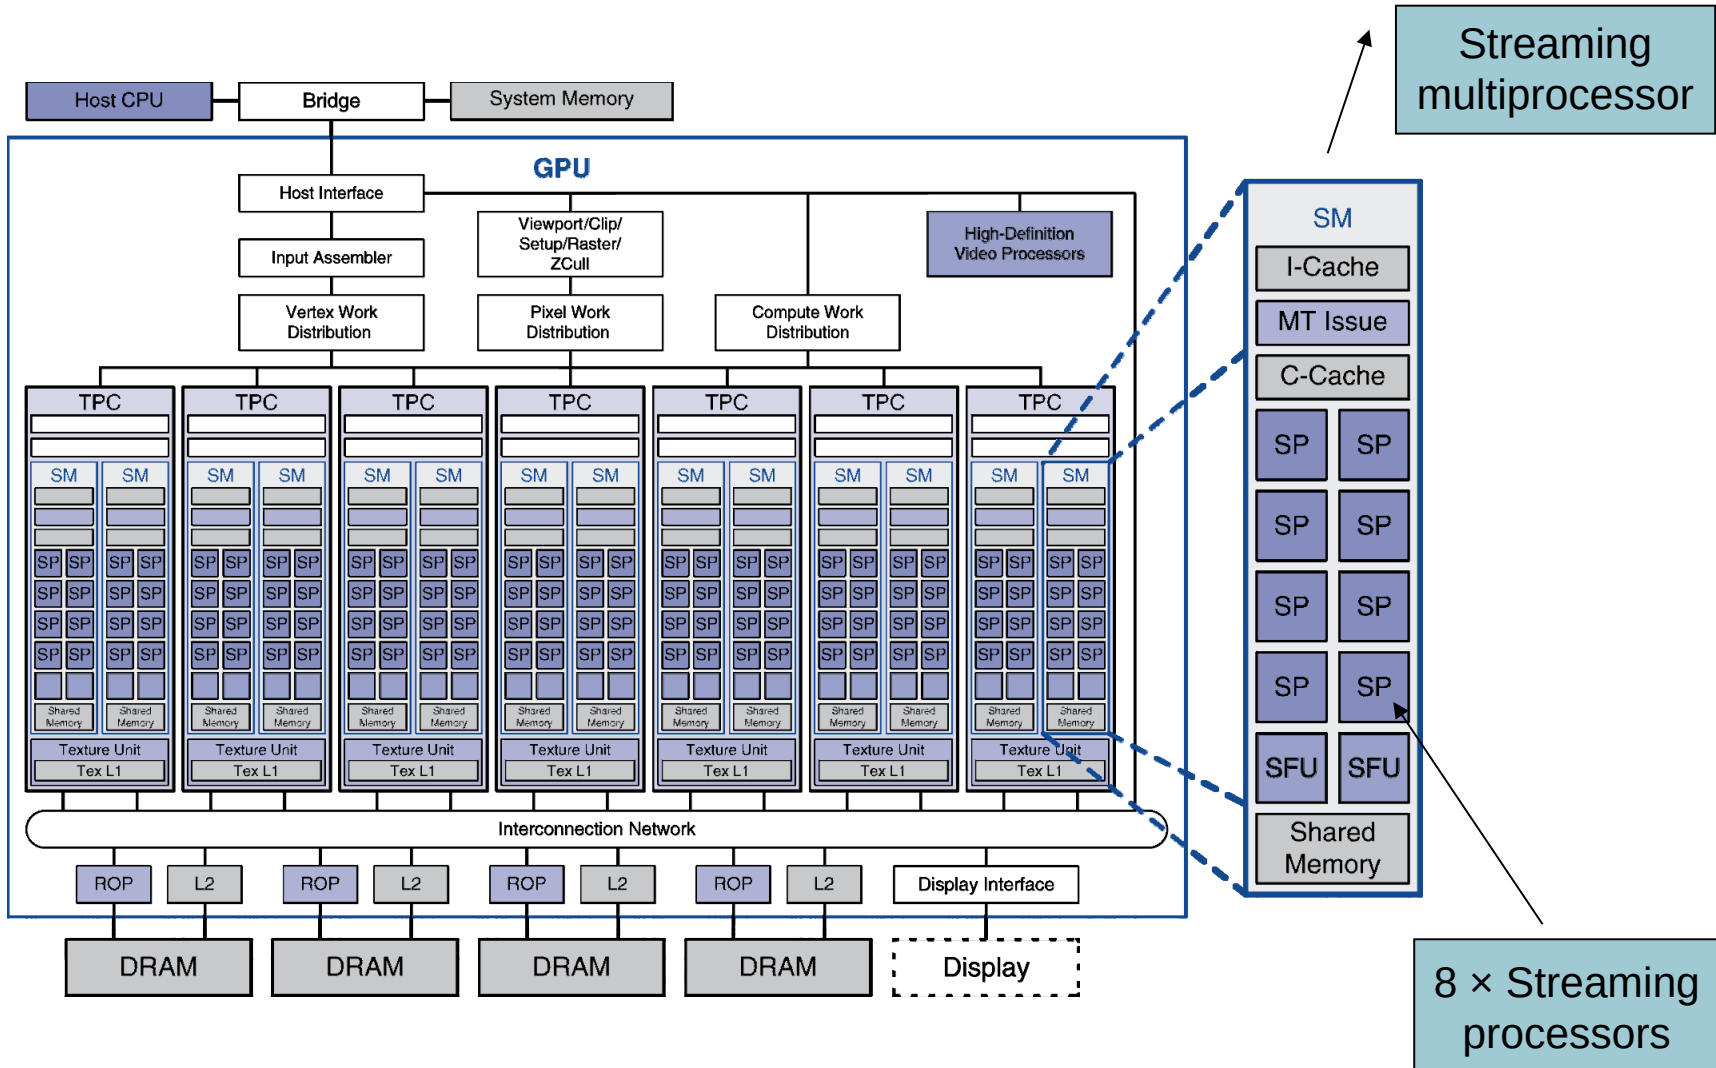
Graphics in the System



GPU Architectures

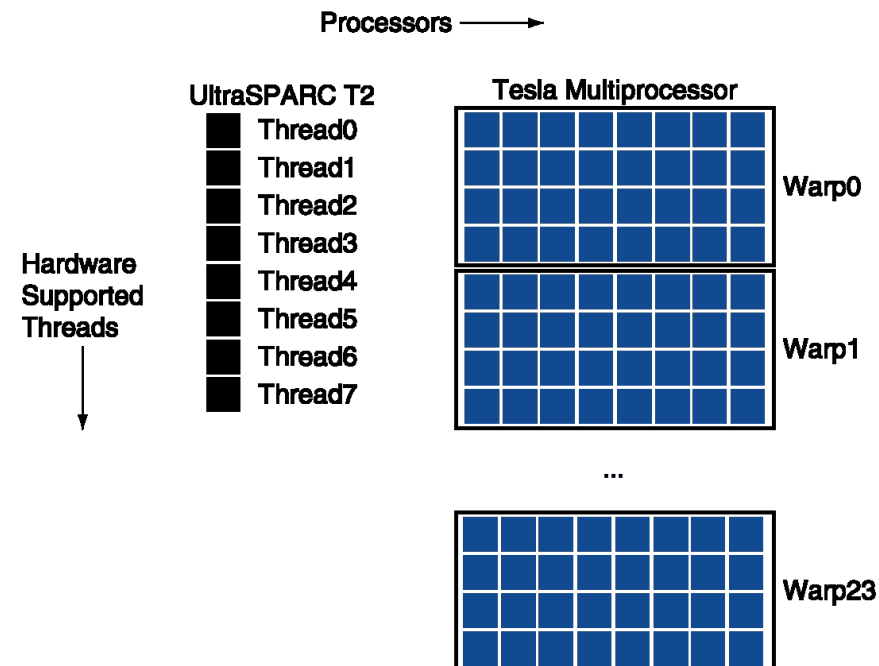
- Processing is highly data-parallel
 - GPUs are highly multithreaded
 - Use thread switching to hide memory latency
 - Less reliance on multi-level caches
 - Graphics memory is wide and high-bandwidth
- Trend toward general purpose GPUs
 - Heterogeneous CPU/GPU systems
 - CPU for sequential code, GPU for parallel code
- Programming languages/APIs
 - DirectX, OpenGL
 - C for Graphics (Cg), High Level Shader Language (HLSL), OpenGL Shading Language (GLSL)
 - Compute Unified Device Architecture (CUDA), OpenCL

Example: NVIDIA Tesla



Example: NVIDIA Tesla

- Streaming Processors
 - Single-precision FP and integer units
 - Each SP is fine-grained multithreaded
- Warp: group of 32 threads
 - Executed in parallel, SIMD style
 - 8 SPs
× 4 clock cycles
 - Hardware contexts for 24 warps
 - Registers, PCs, ...



Classifying GPUs

- Don't fit nicely into SIMD/MIMD model
 - Conditional execution in a thread allows an illusion of MIMD
 - But with performance degradation
 - Need to write general purpose code with care

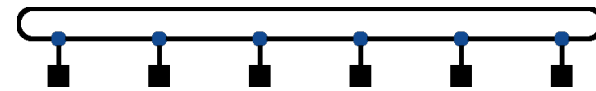
	Static: Discovered at Compile Time	Dynamic: Discovered at Runtime
Instruction-Level Parallelism	VLIW	Superscalar
Data-Level Parallelism	SIMD or Vector	Tesla Multiprocessor

Interconnection Networks

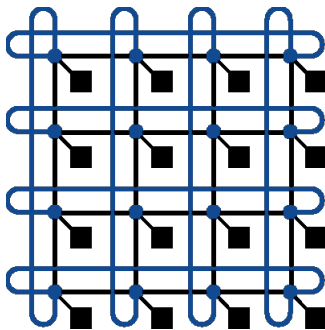
- Network topologies
 - Arrangements of processors, switches, and links



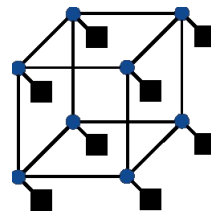
Bus



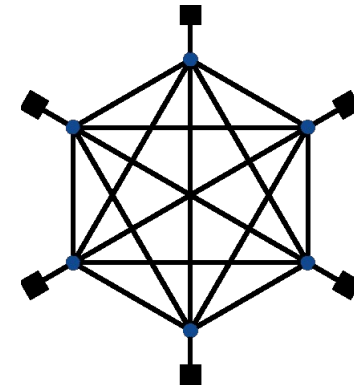
Ring



2D Mesh

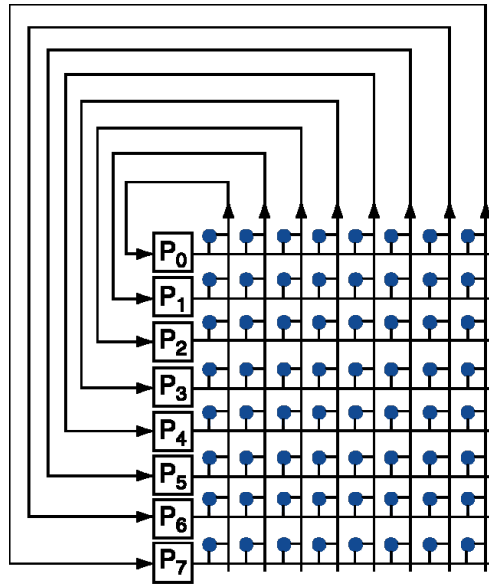


N-cube (N = 3)

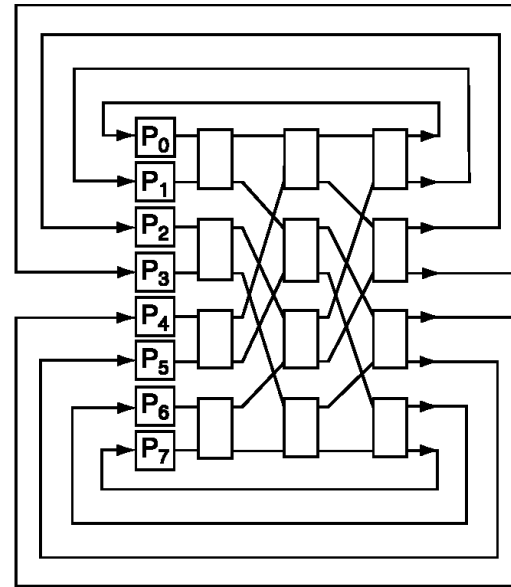


Fully connected

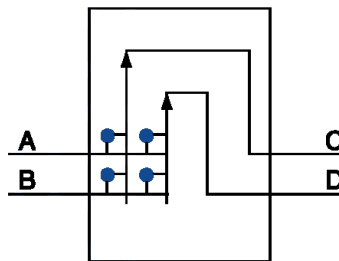
Multistage Networks



a. Crossbar



b. Omega network



c. Omega network switch box

Network Characteristics

- Performance
 - Latency per message (unloaded network)
 - Throughput
 - Link bandwidth
 - Total network bandwidth
 - Bisection bandwidth
 - Congestion delays (depending on traffic)
- Cost
- Power
- Routability in silicon

Parallel Benchmarks

- Linpack: matrix linear algebra
- SPECrate: parallel run of SPEC CPU programs
 - Job-level parallelism
- SPLASH: Stanford Parallel Applications for Shared Memory
 - Mix of kernels and applications, strong scaling
- NAS (NASA Advanced Supercomputing) suite
 - computational fluid dynamics kernels
- PARSEC (Princeton Application Repository for Shared Memory Computers) suite
 - Multithreaded applications using Pthreads and OpenMP

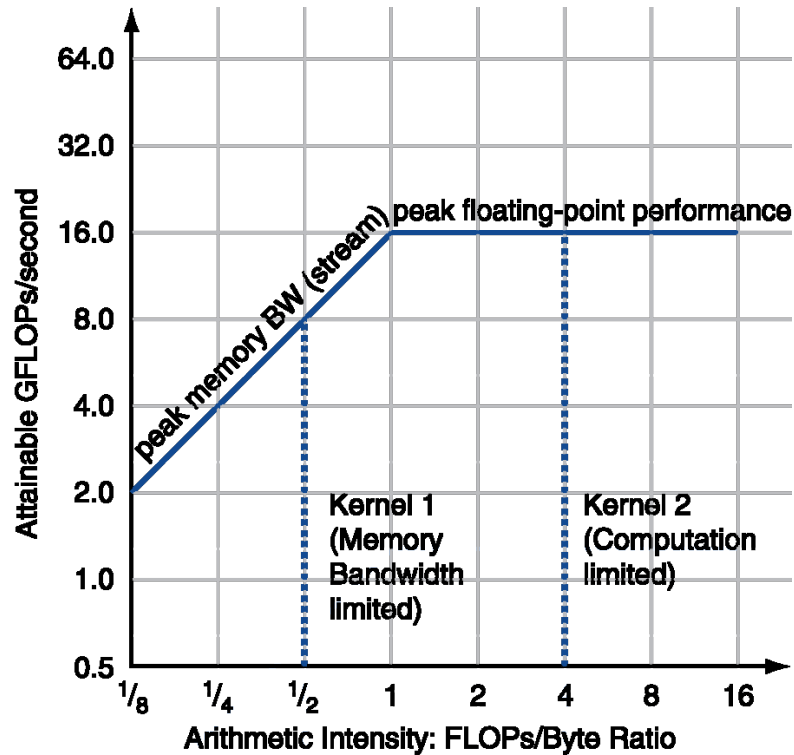
Code or Applications?

- Traditional benchmarks
 - Fixed code and data sets
- Parallel programming is evolving
 - Should algorithms, programming languages, and tools be part of the system?
 - Compare systems, provided they implement a given application
 - E.g., Linpack, Berkeley Design Patterns
- Would foster innovation in approaches to parallelism

Modeling Performance

- Assume performance metric of interest is achievable GFLOPs/sec
 - Measured using computational kernels from Berkeley Design Patterns
- Arithmetic intensity of a kernel
 - FLOPs per byte of memory accessed
- For a given computer, determine
 - Peak GFLOPS (from data sheet)
 - Peak memory bytes/sec (using Stream benchmark)

Roofline Diagram

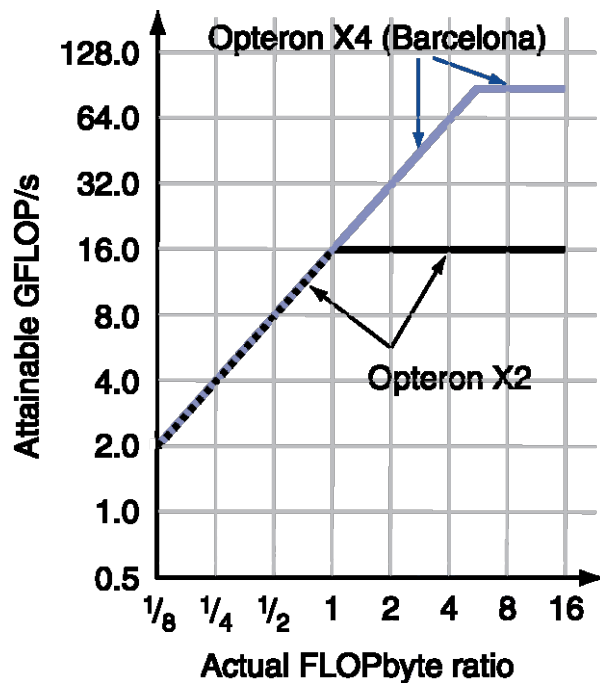


Attainable GPLOPs/sec

= Max (Peak Memory BW × Arithmetic Intensity, Peak FP Performance)

Comparing Systems

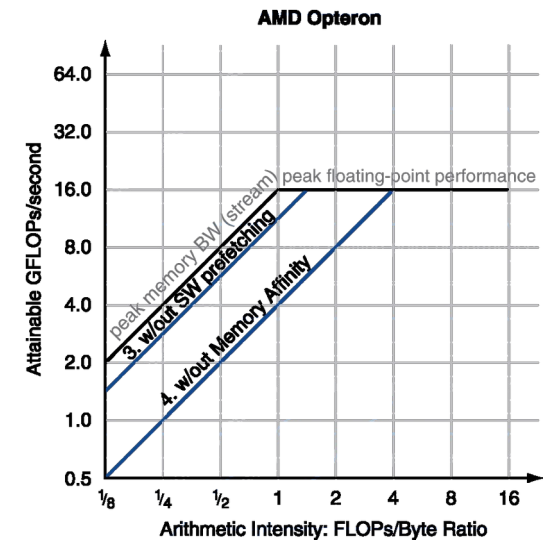
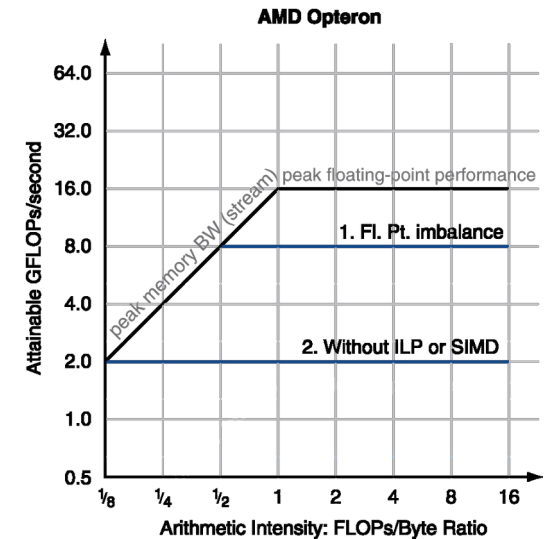
- Example: Opteron X2 vs. Opteron X4
 - 2-core vs. 4-core, 2× FP performance/core, 2.2GHz vs. 2.3GHz
 - Same memory system



- To get higher performance on X4 than X2
 - Need high arithmetic intensity
 - Or working set must fit in X4's 2MB L-3 cache

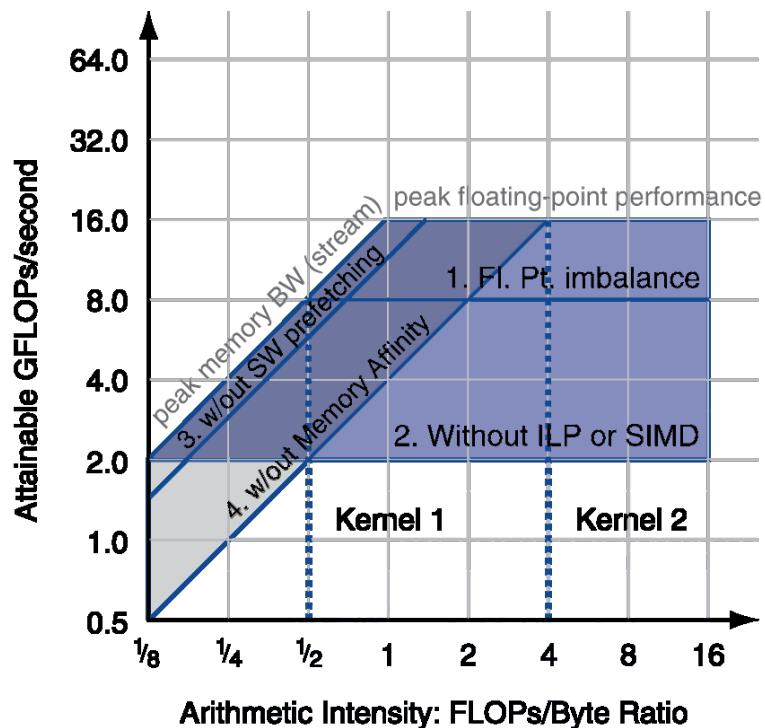
Optimizing Performance

- Optimize FP performance
 - Balance adds & multiplies
 - Improve superscalar ILP and use of SIMD instructions
- Optimize memory usage
 - Software prefetch
 - Avoid load stalls
 - Memory affinity
 - Avoid non-local data accesses



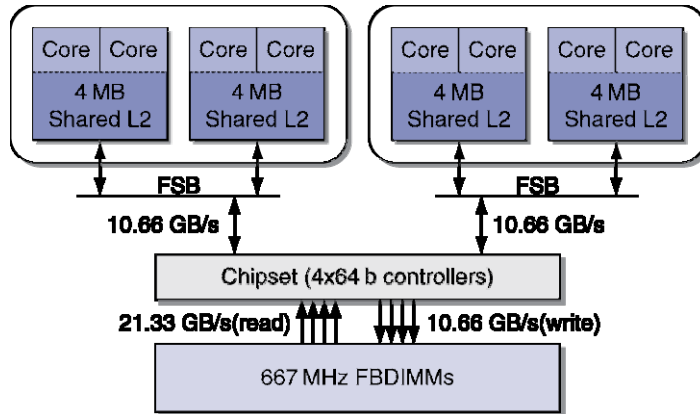
Optimizing Performance

- Choice of optimization depends on arithmetic intensity of code

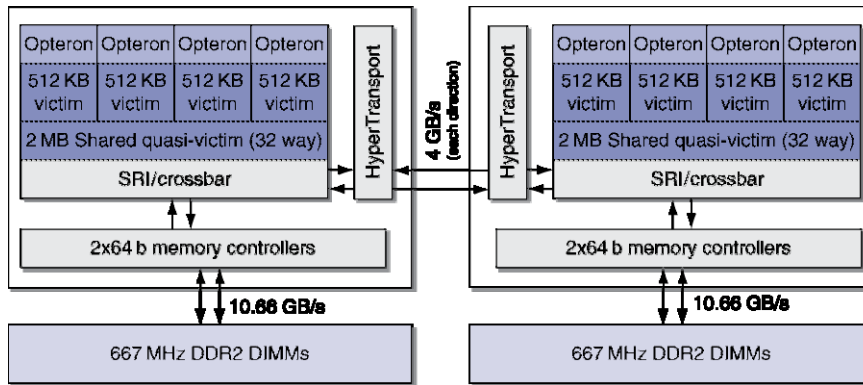


- Arithmetic intensity is not always fixed
 - May scale with problem size
 - Caching reduces memory accesses
 - Increases arithmetic intensity

Four Example Systems

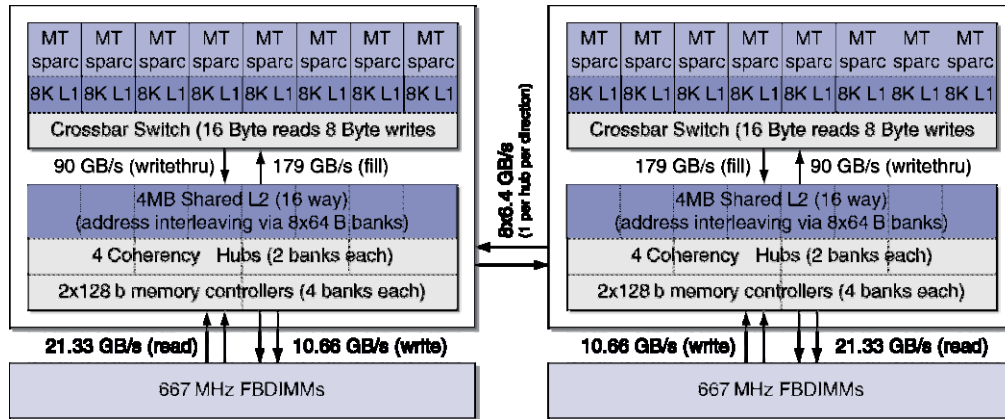


2 × quad-core
Intel Xeon e5345
(Clovertown)

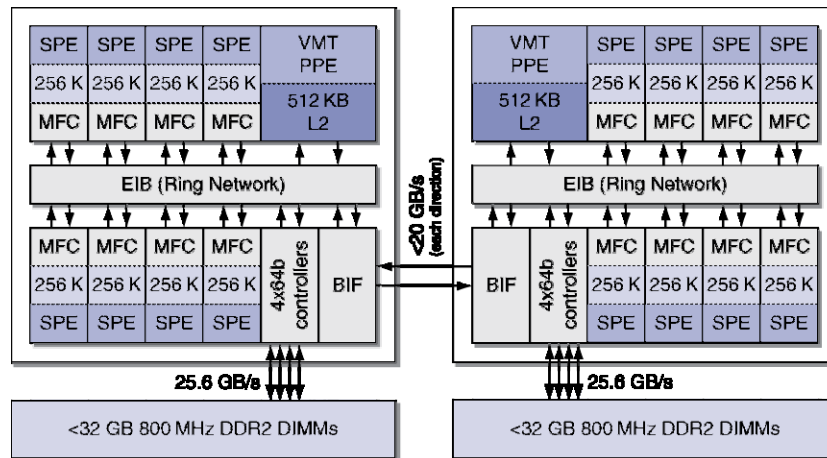


2 × quad-core
AMD Opteron X4 2356
(Barcelona)

Four Example Systems



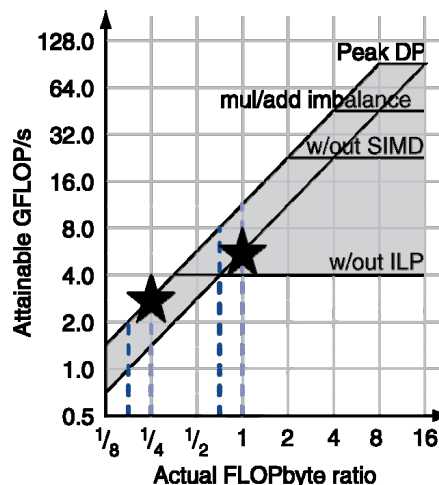
2 × oct-core
Sun UltraSPARC
T2 5140 (Niagara 2)



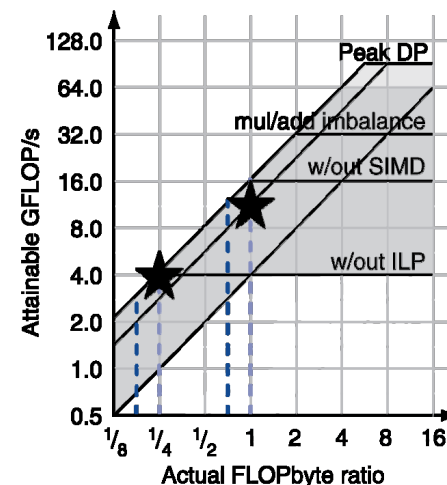
2 × oct-core
IBM Cell QS20

And Their Rooflines

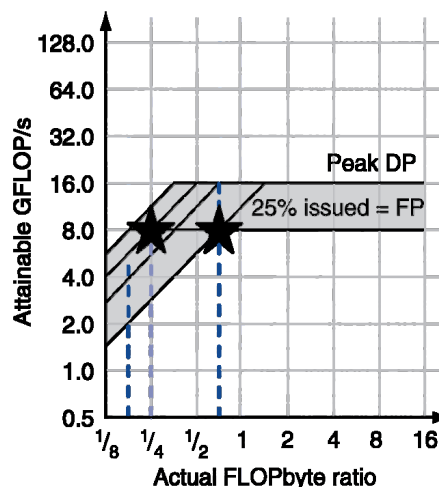
- Kernels
 - SpMV (left)
 - LBHMD (right)
- Some optimizations change arithmetic intensity
- x86 systems have higher peak GFLOPs
 - But harder to achieve, given memory bandwidth



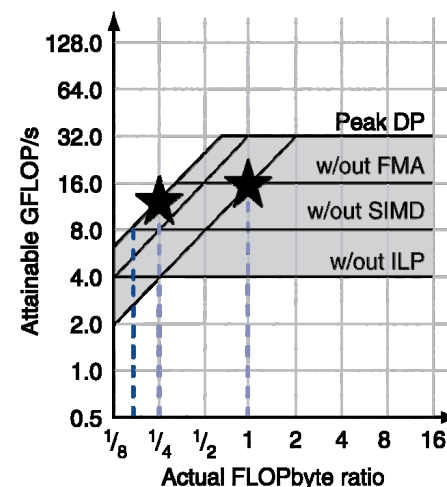
a. Intel Xeon e5345 (Clovertown)



b. AMD Opteron X4 2356 (Barcelona)



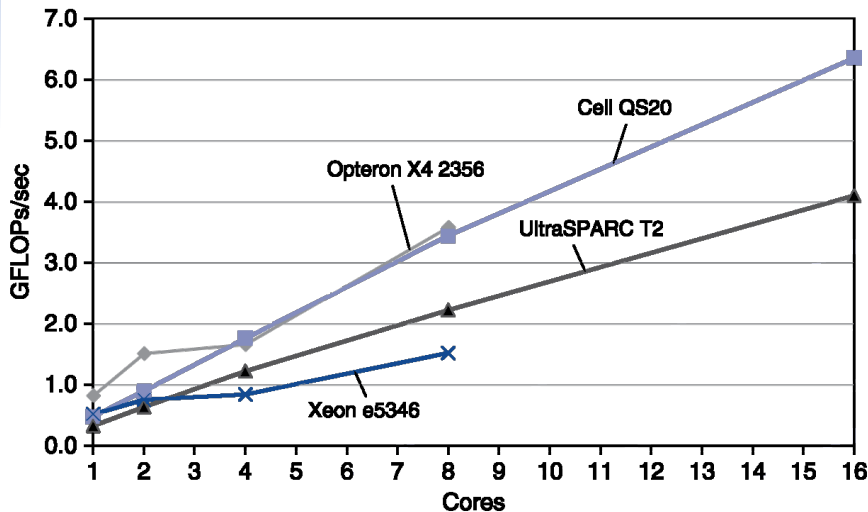
c. Sun UltraSPARC T2 5140 (Niagara 2)



d. IBM Cell QS20

Performance on SpMV

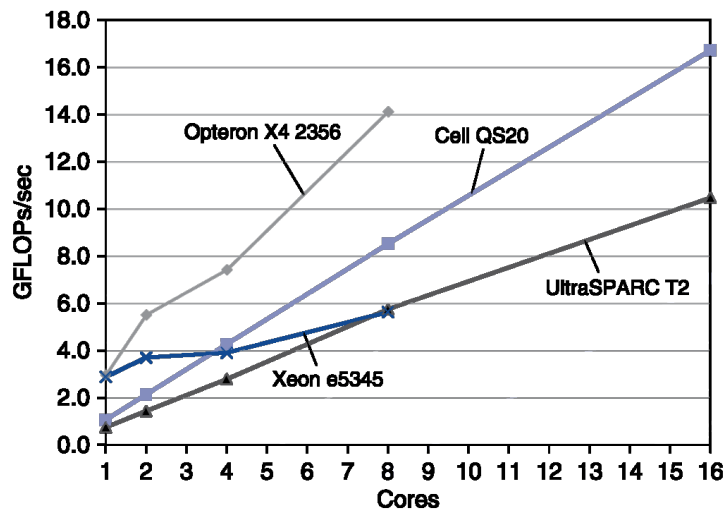
- Sparse matrix/vector multiply
 - Irregular memory accesses, memory bound
- Arithmetic intensity
 - 0.166 before memory optimization, 0.25 after



- Xeon vs. Opteron
 - Similar peak FLOPS
 - Xeon limited by shared FSBs and chipset
- UltraSPARC/Cell vs. x86
 - 20 – 30 vs. 75 peak GFLOPs
 - More cores and memory bandwidth

Performance on LBMHD

- Fluid dynamics: structured grid over time steps
 - Each point: 75 FP read/write, 1300 FP ops
- Arithmetic intensity
 - 0.70 before optimization, 1.07 after



- Opteron vs. UltraSPARC
 - More powerful cores, not limited by memory bandwidth
- Xeon vs. others
 - Still suffers from memory bottlenecks

Achieving Performance

- Compare naïve vs. optimized code
 - If naïve code performs well, it's easier to write high performance code for the system

System	Kernel	Naïve GFLOPs/sec	Optimized GFLOPs/sec	Naïve as % of optimized
Intel Xeon	SpMV	1.0	1.5	64%
	LBMHD	4.6	5.6	82%
AMD Opteron X4	SpMV	1.4	3.6	38%
	LBMHD	7.1	14.1	50%
Sun UltraSPARC T2	SpMV	3.5	4.1	86%
	LBMHD	9.7	10.5	93%
IBM Cell QS20	SpMV	Naïve code not feasible	6.4	0%
	LBMHD	Naïve code not feasible	16.7	0%

Fallacies

- Amdahl's Law doesn't apply to parallel computers
 - Since we can achieve linear speedup
 - But only on applications with weak scaling
- Peak performance tracks observed performance
 - Marketers like this approach!
 - But compare Xeon with others in example
 - Need to be aware of bottlenecks

Pitfalls

- Not developing the software to take account of a multiprocessor architecture
 - Example: using a single lock for a shared composite resource
 - Serializes accesses, even if they could be done in parallel
 - Use finer-granularity locking

Concluding Remarks

- Goal: higher performance by using multiple processors
- Difficulties
 - Developing parallel software
 - Devising appropriate architectures
- Many reasons for optimism
 - Changing software and application environment
 - Chip-level multiprocessors with lower latency, higher bandwidth interconnect
- An ongoing challenge for computer architects!

Other Resources

- Patterson & Hennessy's: *Computer Architecture: A Quantitative Approach*
- Conferences:
 - www.cs.wisc.edu/~arch/www/
 - ISCA (International Symposium on Computer Architecture)
 - HPCA (International Symposium on High Performance Computer Architecture)