

---

---

# Topic 22: Multi-Processor Parallelism

COS / ELE 375

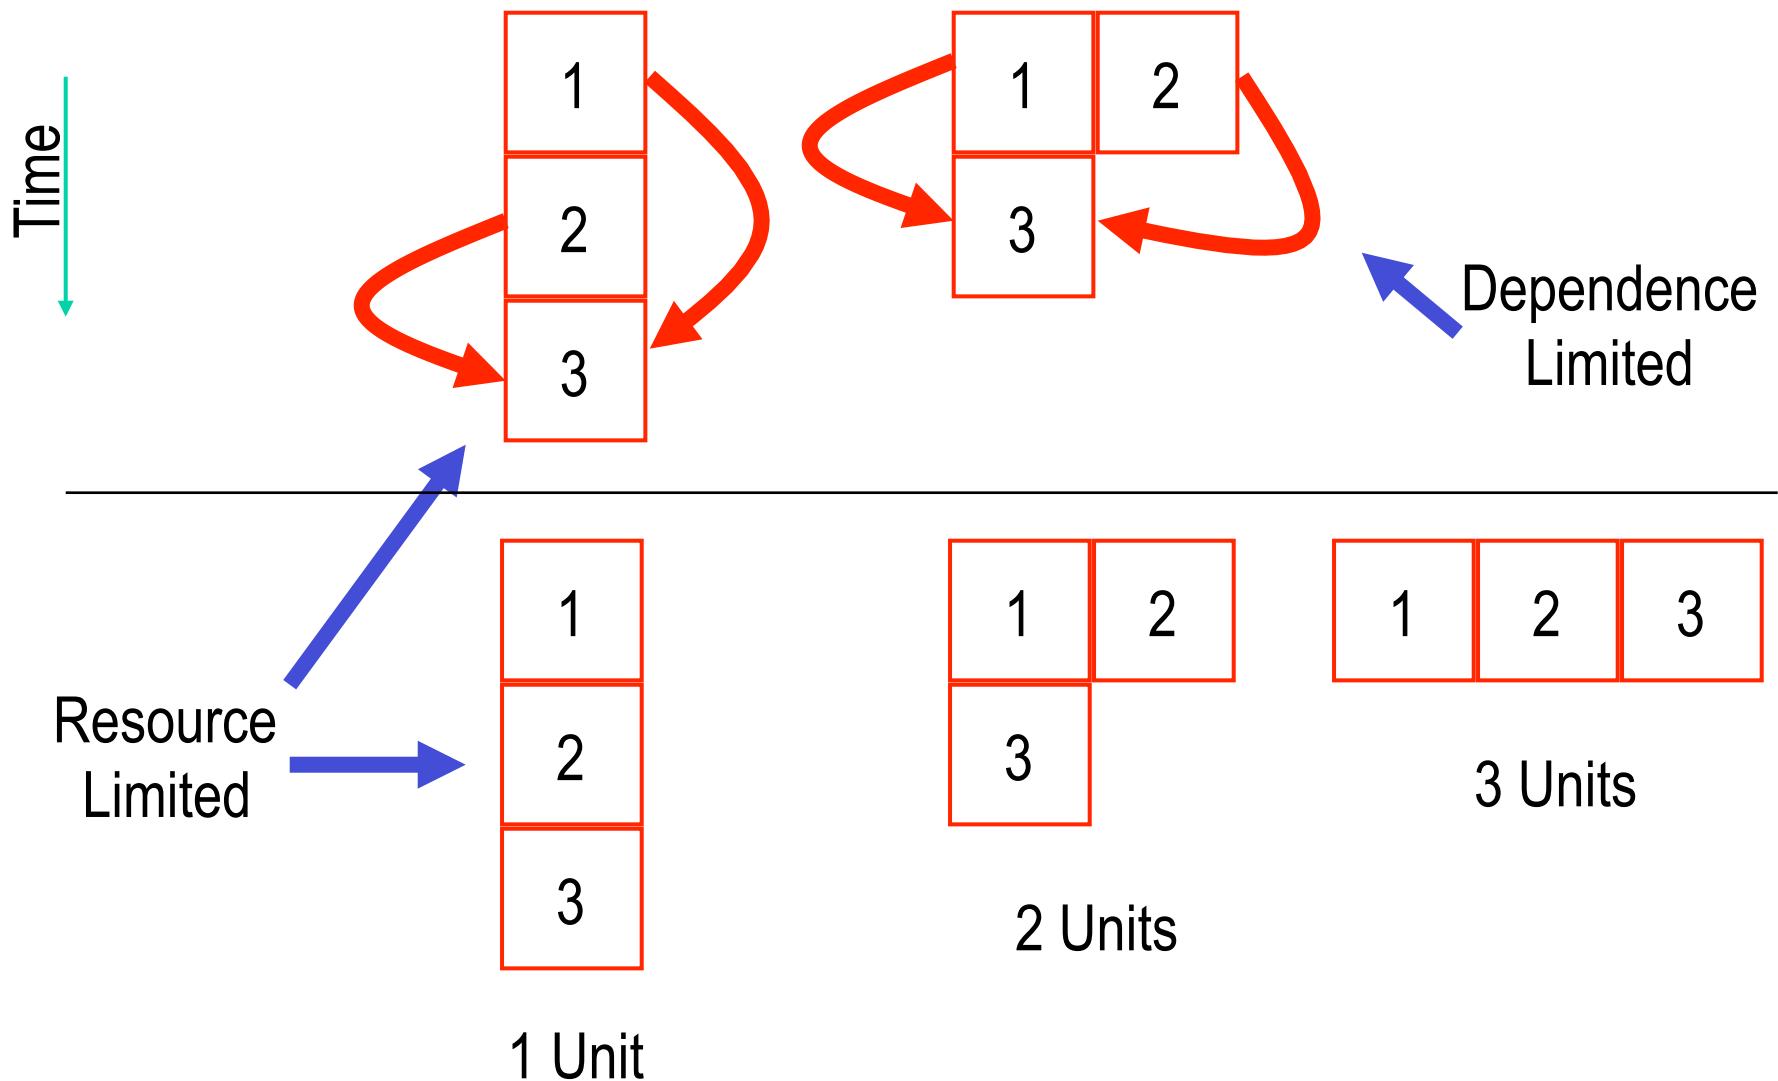
Computer Architecture and Organization

Princeton University  
Fall 2015

Prof. David August

## Review: Parallelism

Independent units of work can execute concurrently if sufficient resources exist



# Review: Where to Find Parallelism?

---

Parallelism can be found/exists at different granularities

- Instruction Level
  - Ex: add instruction executes with multiply instruction
  - Compiler and hardware good at finding this
- Thread Level
  - Ex: screen redraw function executes with recalculate in spreadsheet
  - Programmers OK at finding this
- Process Level
  - Ex: Simulation job runs on same machines as spreadsheet
  - Users good at creating this

## Thread Level Parallelism

---

Programmer generally makes TLP explicit

Compilers can extract threads in regular programs

```
for (i = 0; i < 200; i++)
    for(j = 1; j < 20000; j++)
        val[i,j] = val[i,j-1] + 1;
```

```
forall(i = 0; i < 200; i++)
    for(j = 1; j < 20000; j++)
        val[i,j] = val[i,j-1] + 1;
```

# Thread Level Parallelism

---

## Synchronization

- Unlike in ILP, flow of data/dependences must be explicit

```
while(ptr = ptr->next)
    sum += ptr->val;
```

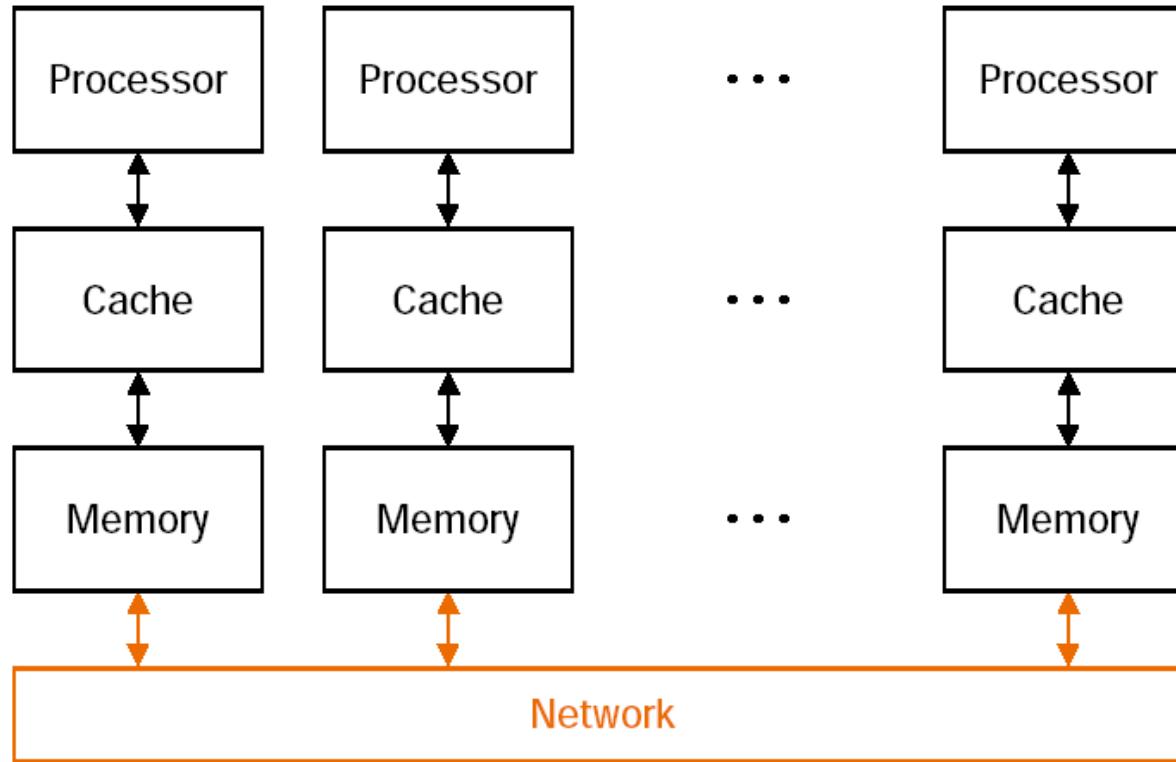
```
while(ptr = ptr->next)
    produce(ptr) ;
produce(NULL) ;
```

```
while(ptr = consume(ptr))
    sum += ptr->val;
```

Communication and Synchronization... (order and flow)

# Multiple Processor Organization

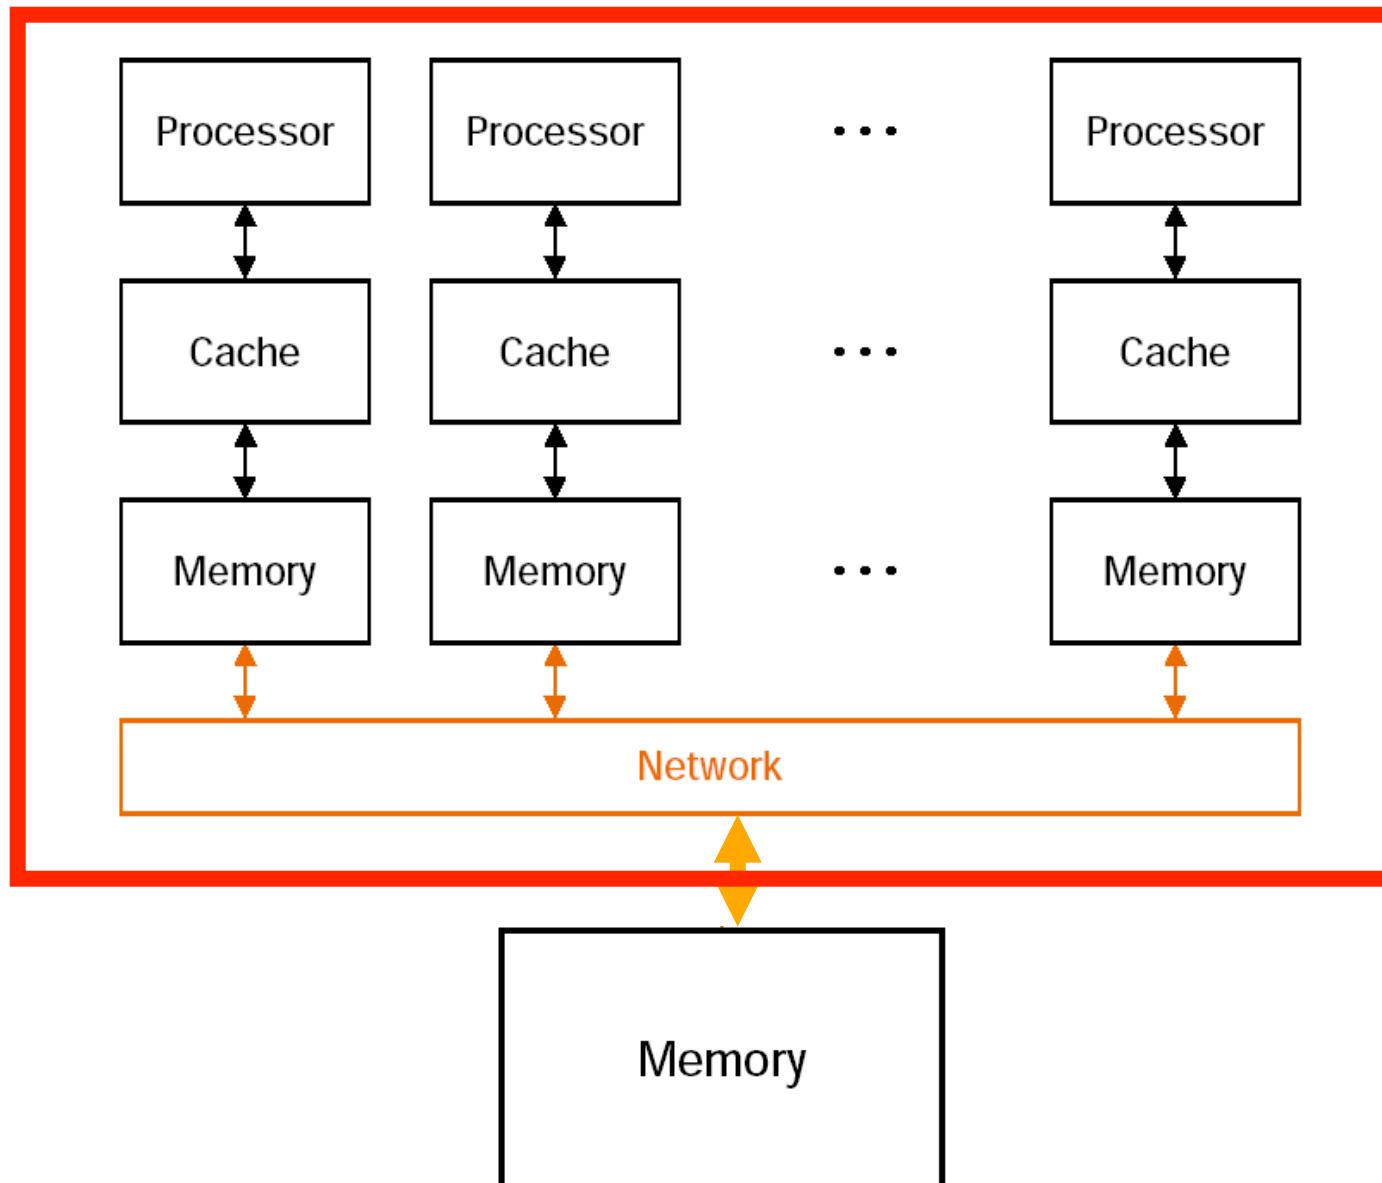
## Message Passing/Private Memory



- Threads communicate directly (send, receive)
- Scales relatively well
- No memory coherence problem (for the hardware at least)

# Multiple Processor Organization

May exist on single chip



## Thread Level Parallelism

---

Programmer generally makes TLP explicit

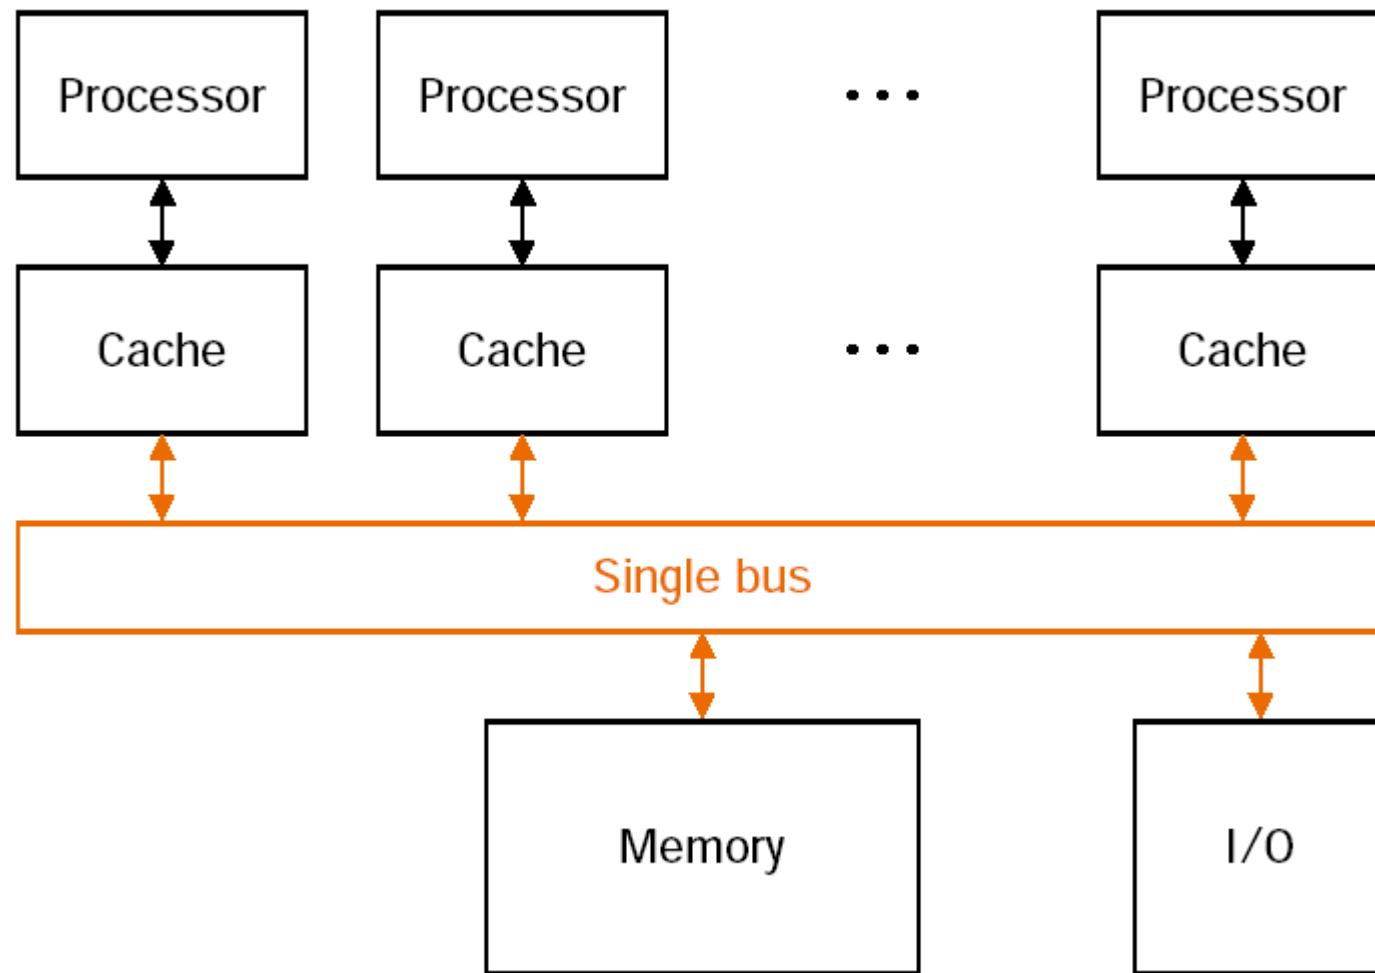
Compilers can extract threads in regular programs

```
for (i = 0; i < 200; i++)
    for(j = 1; j < 20000; j++)
        val[i,j] = val[i,j-1] + 1;
```

```
forall(i = 0; i < 200; i++)
    for(j = 1; j < 20000; j++)
        val[i,j] = val[i,j-1] + 1;
```

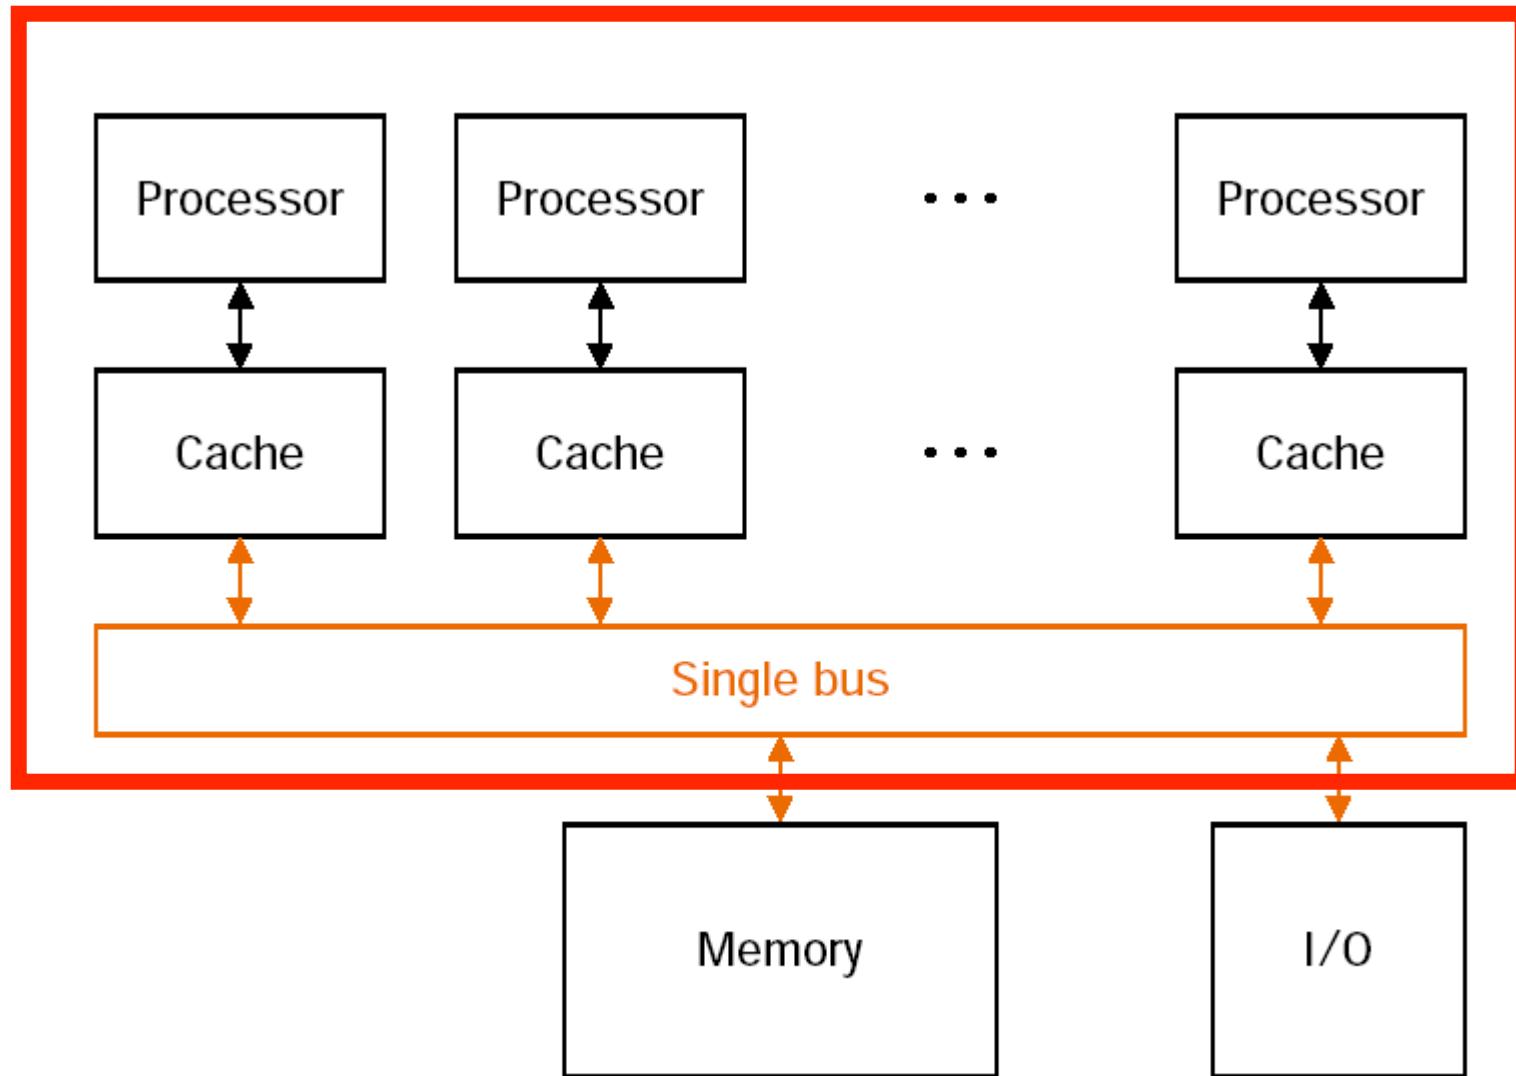
# Multiple Processor Organizations

## Shared Memory/Shared Bus



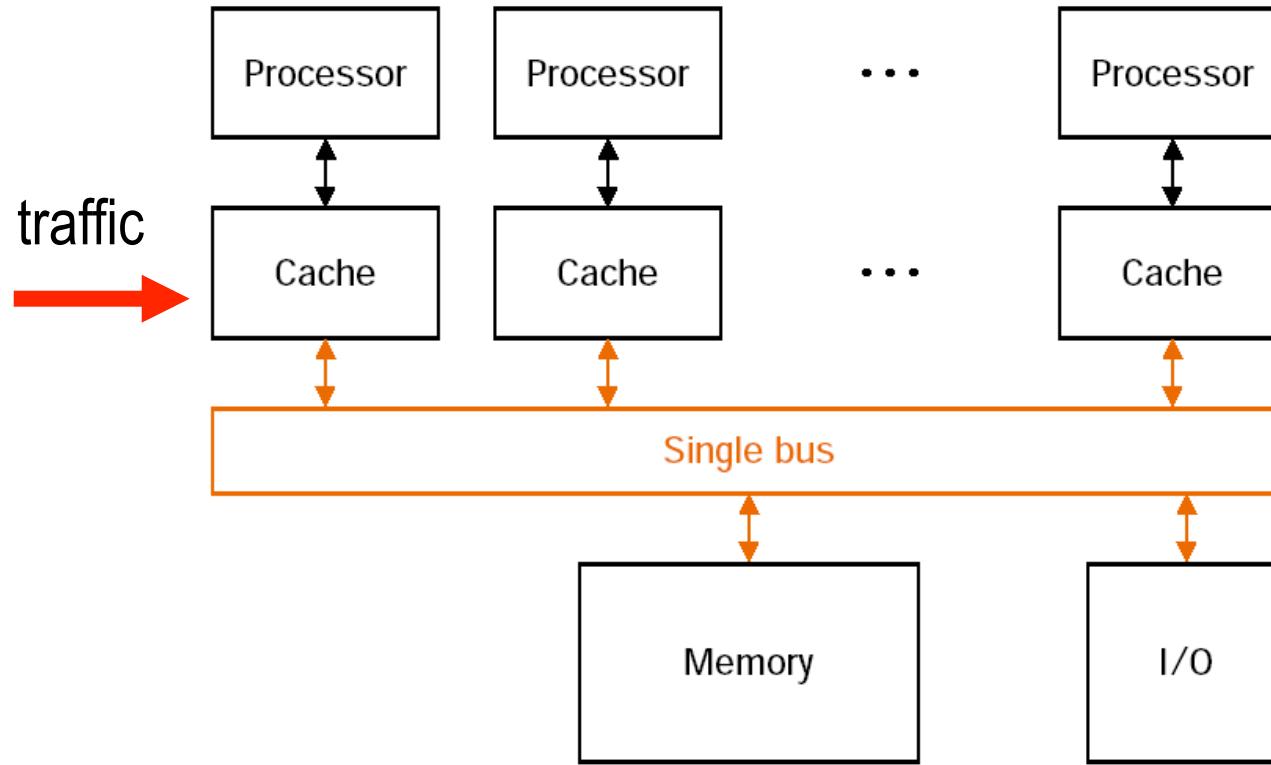
# Multiple Processor Organizations

May be on single chip!



# Multiple Processor Organizations

## Shared Bus

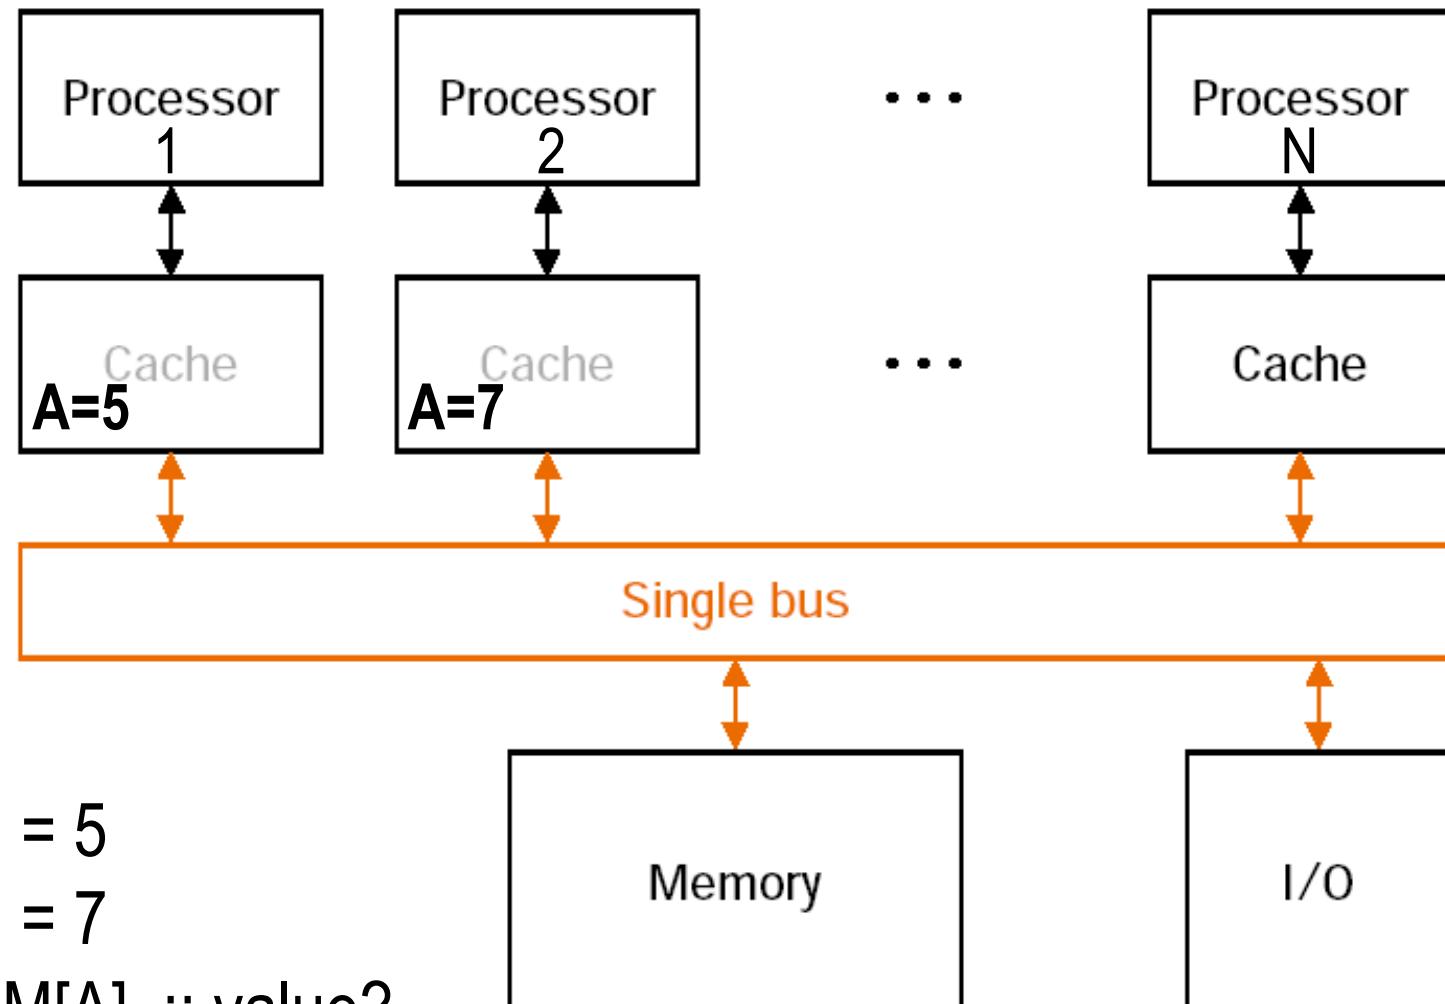


Synchronization and Communication through memory

The cache coherency problem

# Multiple Processor Organizations

## Shared Bus



P1:  $M[A] = 5$

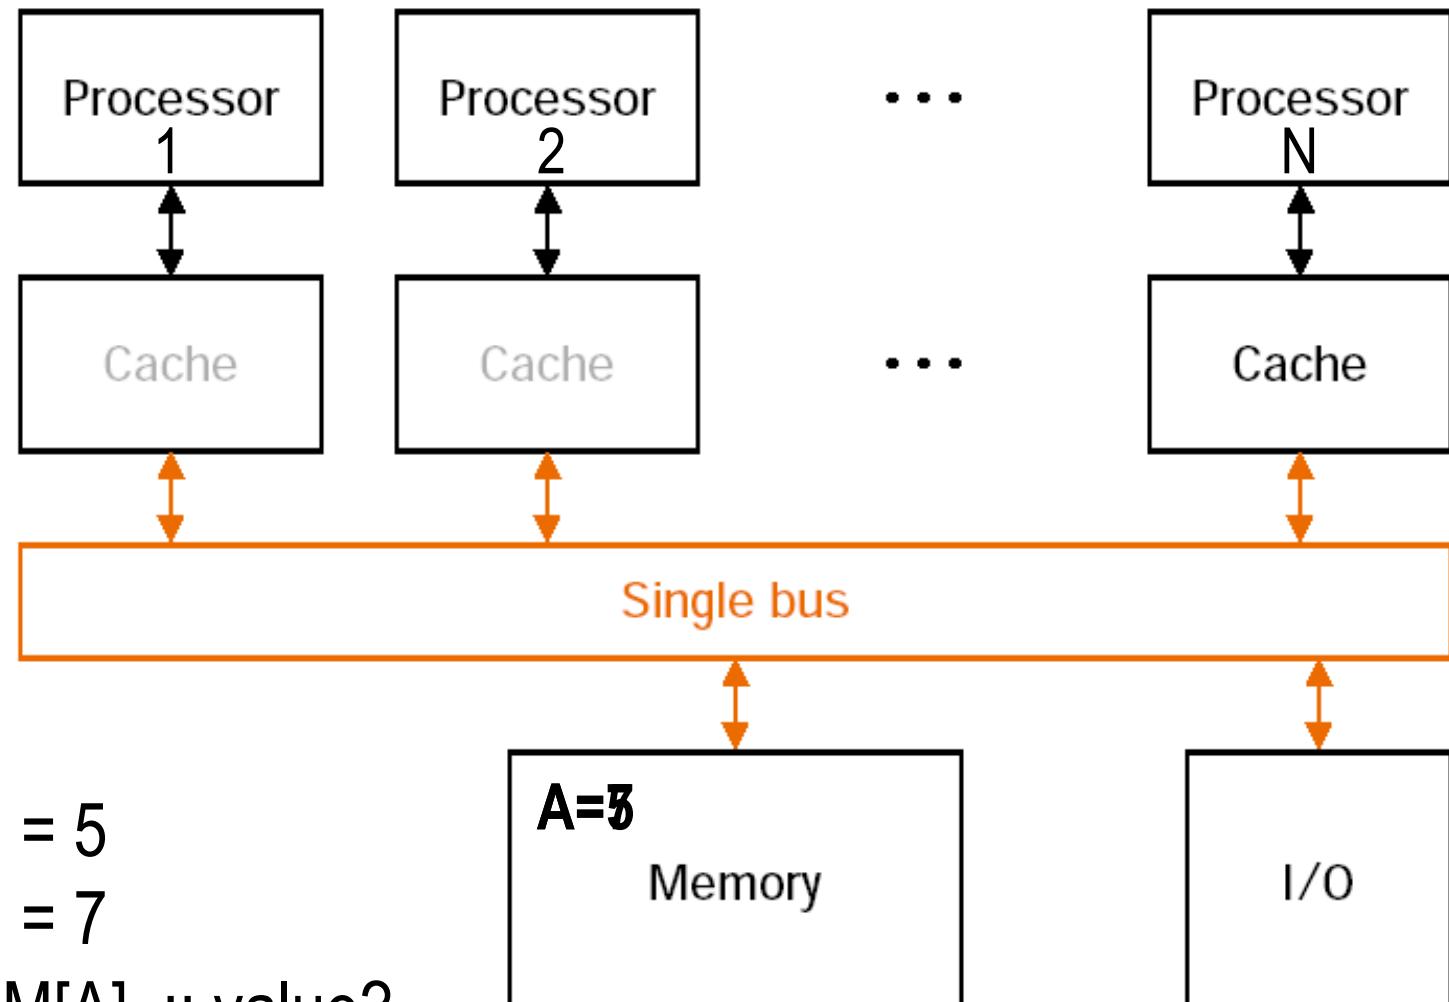
P2:  $M[A] = 7$

P1:  $r1 = M[A]$  ;; value?

PN:  $r1 = M[A]$  ;; value?

# Cache Coherency

## “Solution”



P1:  $M[A] = 5$

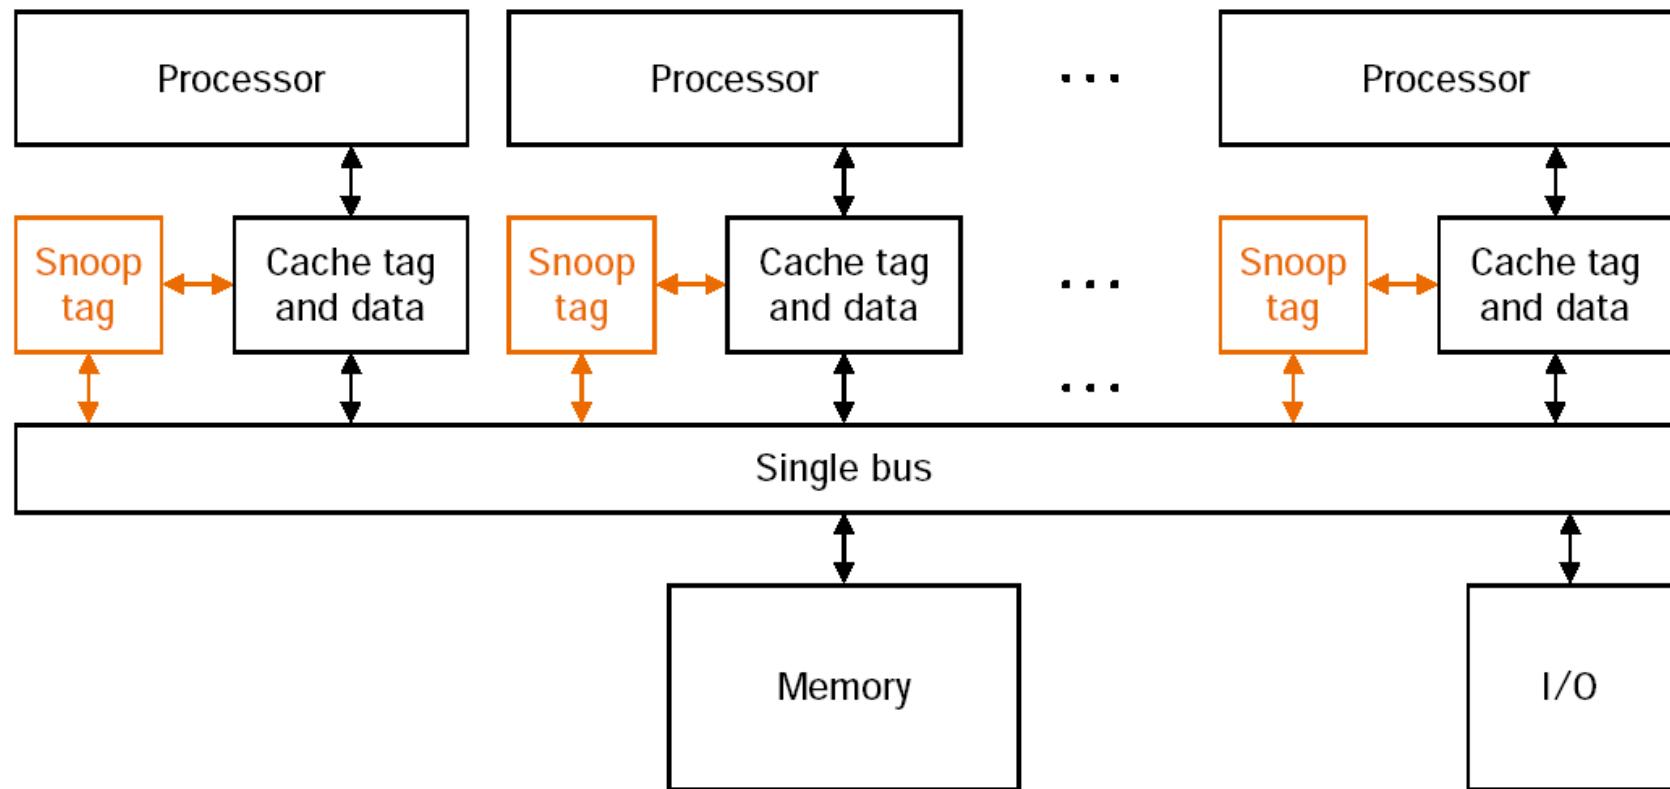
P2:  $M[A] = 7$

P1:  $r1 = M[A]$  ;; value?

PN:  $r1 = M[A]$  ;; value?

# Cache Coherency

## Solution: Snoopy Bus



# Snooping Protocols

---

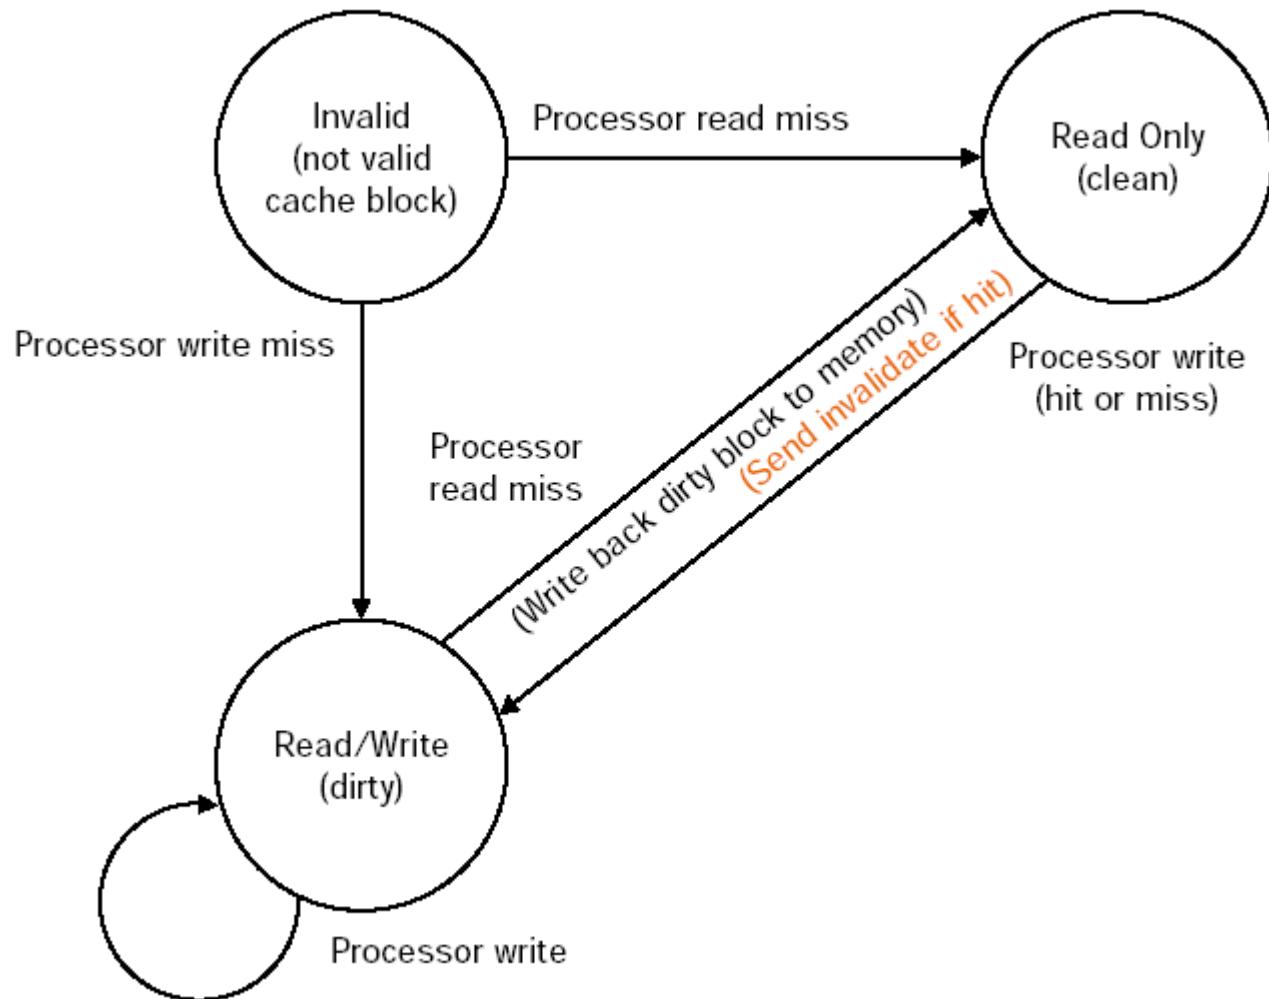
- Variety of protocols minimize traffic for different situations
- Generally many states including:
  - invalid, dirty read/write, clean/read-only

Reads: just work

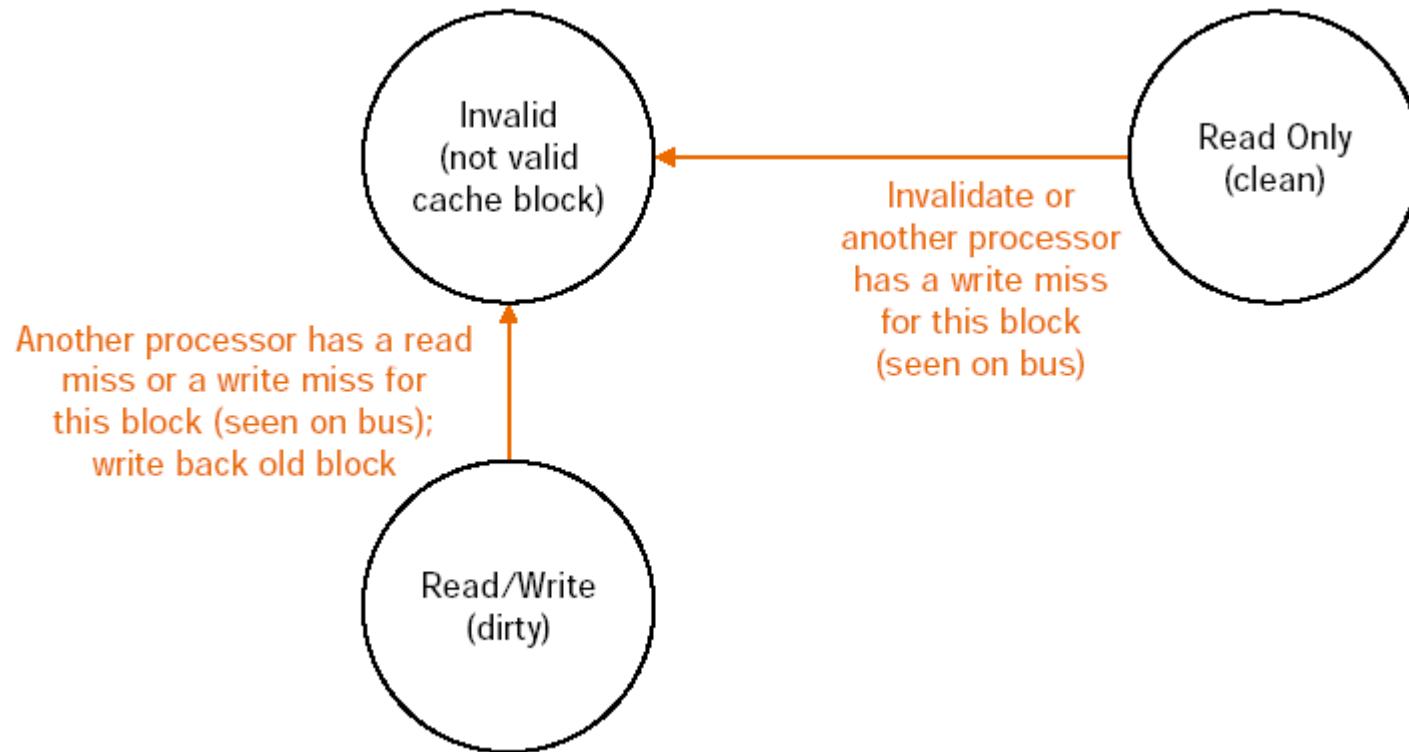
Writes:

- Write-Invalidate - other caches with address invalidate line (block) - only first write generates traffic
- Write-Update - other caches with address update the values in the line (block) - like write through

# Sample Protocol Signals From Processor



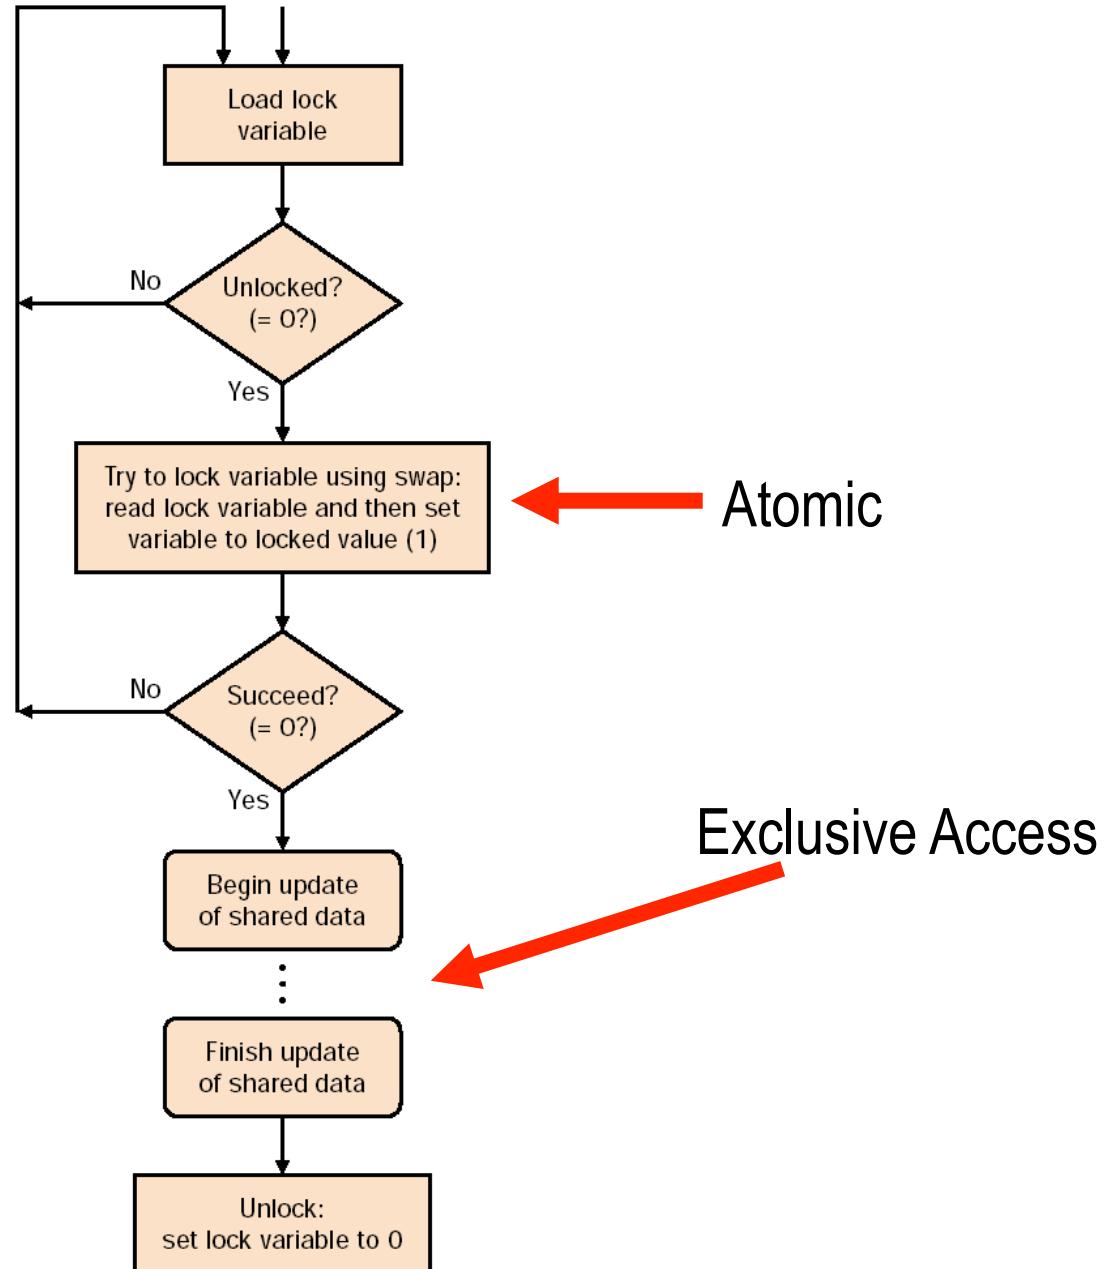
# Sample Protocol Signals From Bus



Other protocols are MESler

# Synchronization/Semaphores

Spin lock



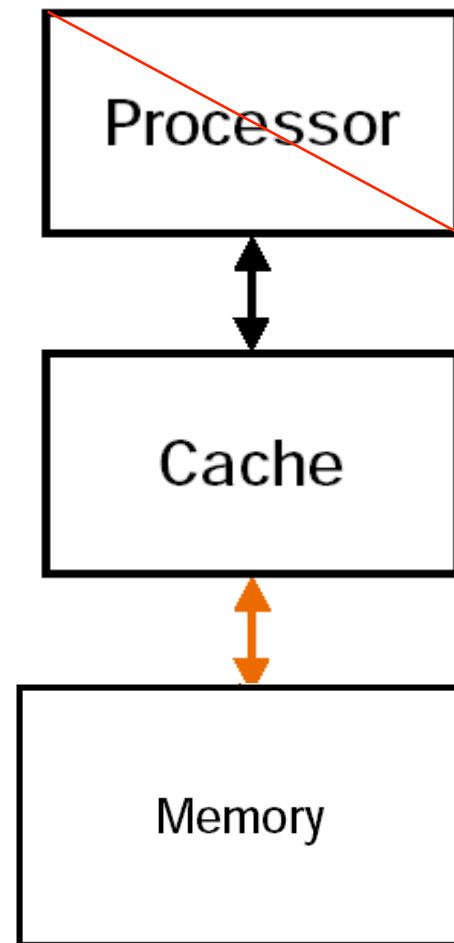
# Multiple Processor Organizations

## Simultaneous Multithreading (“Hyperthreading”)

- Multiple threads in single core
- Helps when single thread ILP is low

Like ILP processor, but

- Multiple PCs, one per thread
- Instructions are tagged with thread ID
- Architectural register file per thread
- Threads share execution resources
- Cross thread synchronization and communication through memory/ cache



# CFGs, PCs, and Cross-Iteration\_deps

---

```
1. r1 = 10
```

```
1. r1 = r1 + 1
```

```
2. r2 = MEM[r1]
```

```
3. r2 = r2 + 1
```

```
4. MEM[r1] = r2
```

```
5. Branch r1 < 1000
```

No register live outs

# Loop-Level Parallelization: DOALL

---

```
1. r1 = 10
```

```
1. r1 = r1 + 1
```

```
2. r2 = MEM[r1]
```

```
3. r2 = r2 + 1
```

```
4. MEM[r1] = r2
```

```
5. Branch r1 < 1000
```

```
1. r1 = 9
```

```
1. r1 = r1 + 2
```

```
2. r2 = MEM[r1]
```

```
3. r2 = r2 + 1
```

```
4. MEM[r1] = r2
```

```
5. Branch r1 < 999
```

```
1. r1 = 10
```

```
1. r1 = r1 + 2
```

```
2. r2 = MEM[r1]
```

```
3. r2 = r2 + 1
```

```
4. MEM[r1] = r2
```

```
5. Branch r1 < 1000
```

No register live outs

# Another Example

---

1.  $r1 = 10$

1.  $r1 = r1 + 1$

2.  $r2 = \text{MEM}[r1]$

3.  $r2 = r2 + 1$

4.  $\text{MEM}[r1] = r2$

5. Branch  $r2 == 10$

No register live outs

# Another Example

---

1.  $r1 = 10$

1.  $r1 = r1 + 1$

2.  $r2 = \text{MEM}[r1]$

3.  $r2 = r2 + 1$

4.  $\text{MEM}[r1] = r2$

5. Branch  $r2 == 10$

1.  $r1 = 9$

1.  $r1 = r1 + 2$

2.  $r2 = \text{MEM}[r1]$

3.  $r2 = r2 + 1$

4.  $\text{MEM}[r1] = r2$

5. Branch  $r2 == 10$

1.  $r1 = 10$

1.  $r1 = r1 + 2$

2.  $r2 = \text{MEM}[r1]$

3.  $r2 = r2 + 1$

4.  $\text{MEM}[r1] = r2$

5. Branch  $r2 == 10$

No register live outs

# Speculation

---

1.  $r1 = 9$

1.  $r1 = 10$

1.  $r1 = r1 + 2$

2.  $r2 = \text{MEM}[r1]$

3.  $r2 = r2 + 1$

4.  $\text{MEM}[r1] = r2$

5. Branch  $r2 == 10$

1.  $r1 = r1 + 2$

2.  $r2 = \text{MEM}[r1]$

3.  $r2 = r2 + 1$

4.  $\text{MEM}[r1] = r2$

5. Branch  $r2 == 10$

No register live outs

# Speculation, Commit, and Recovery

---

```
1. r1 = 9
```

```
1. r1 = r1 + 2  
2. r2 = MEM[r1]  
3. r2 = r2 + 1  
4. Send{1} r2  
5. Jump
```

```
1. r2 = Receive{1}
```

```
2. Branch r2 != 10
```

```
3. MEM[r1] = r2
```

```
4. r2 = Receive{2}
```

```
5. Branch r2 != 10
```

```
6. MEM[r1] = r2
```

```
7. Jump
```

```
1. r1 = 10
```

```
1. r1 = r1 + 2  
2. r2 = MEM[r1]  
3. r2 = r2 + 1  
4. MEM[r1] = r2  
5. Jump
```

No register live outs

```
1. Kill and Continue
```

# Difficult Dependencies

---

1.  $r1 = \text{Head}$

1.  $r1 = \text{MEM}[r1]$

2. Branch  $r1 == 0$

3.  $r2 = \text{MEM}[r1 + 4]$

4.  $r3 = \text{Work}(r2)$

5. Print(  $r3$  )

6. Jump

No register live outs

# DOACROSS

---

1.  $r1 = \text{Head}$

1.  $r1 = \text{MEM}[r1]$

2. Branch  $r1 == 0$

3.  $r2 = \text{MEM}[r1 + 4]$

4.  $r3 = \text{Work}(r2)$

5. Print(  $r3$  )

6. Jump

No register live outs

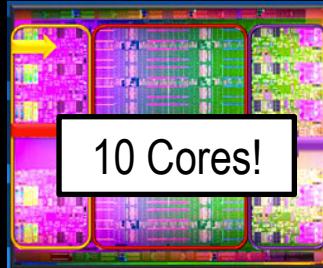
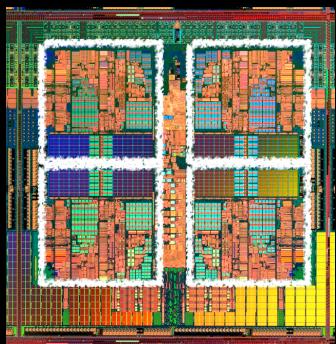
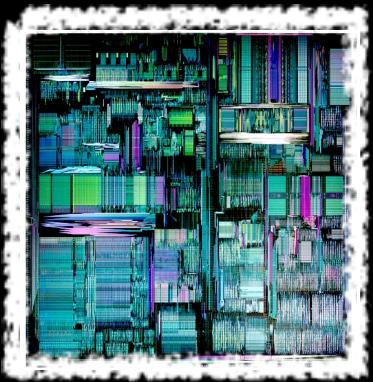
# PS-DSWP

---

1.  $r1 = \text{Head}$

1.  $r1 = \text{MEM}[r1]$
2. Branch  $r1 == 0$
3.  $r2 = \text{MEM}[r1 + 4]$
4.  $r3 = \text{Work}(r2)$
5. Print(  $r3$  )
6. Jump

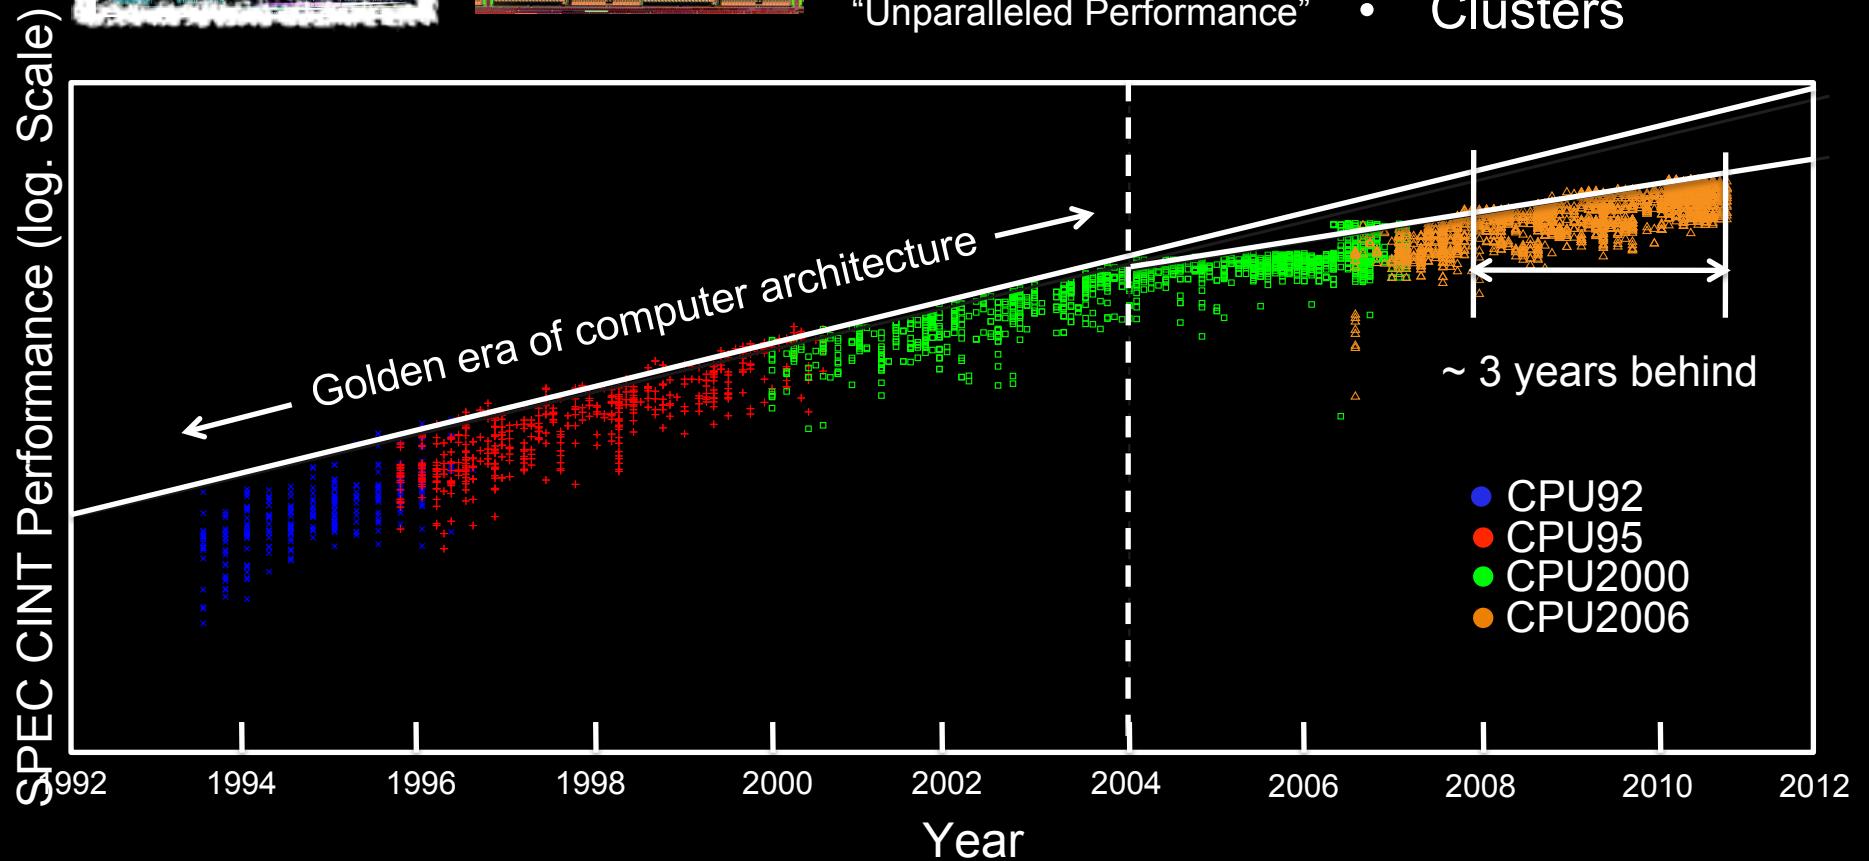
No register live outs



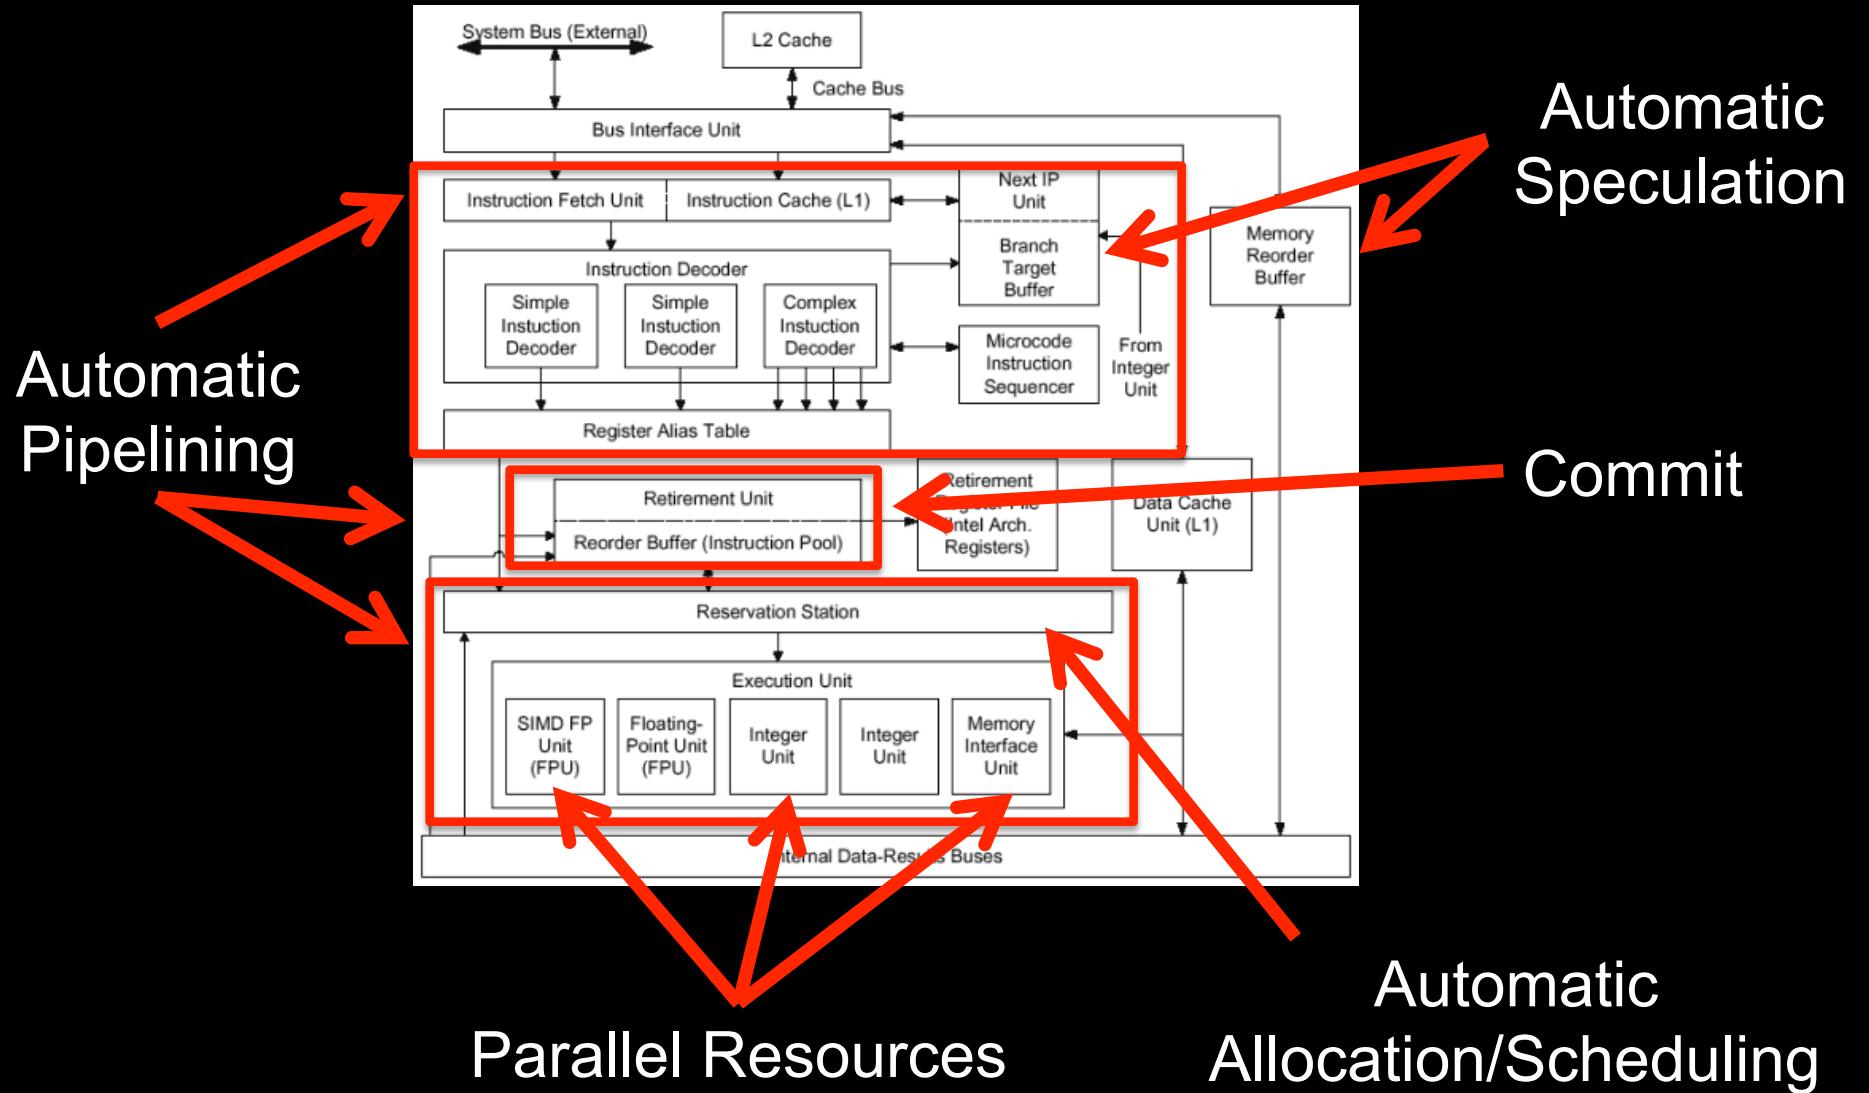
10-Core Intel Xeon  
“Unparalleled Performance”

### Era of DIY:

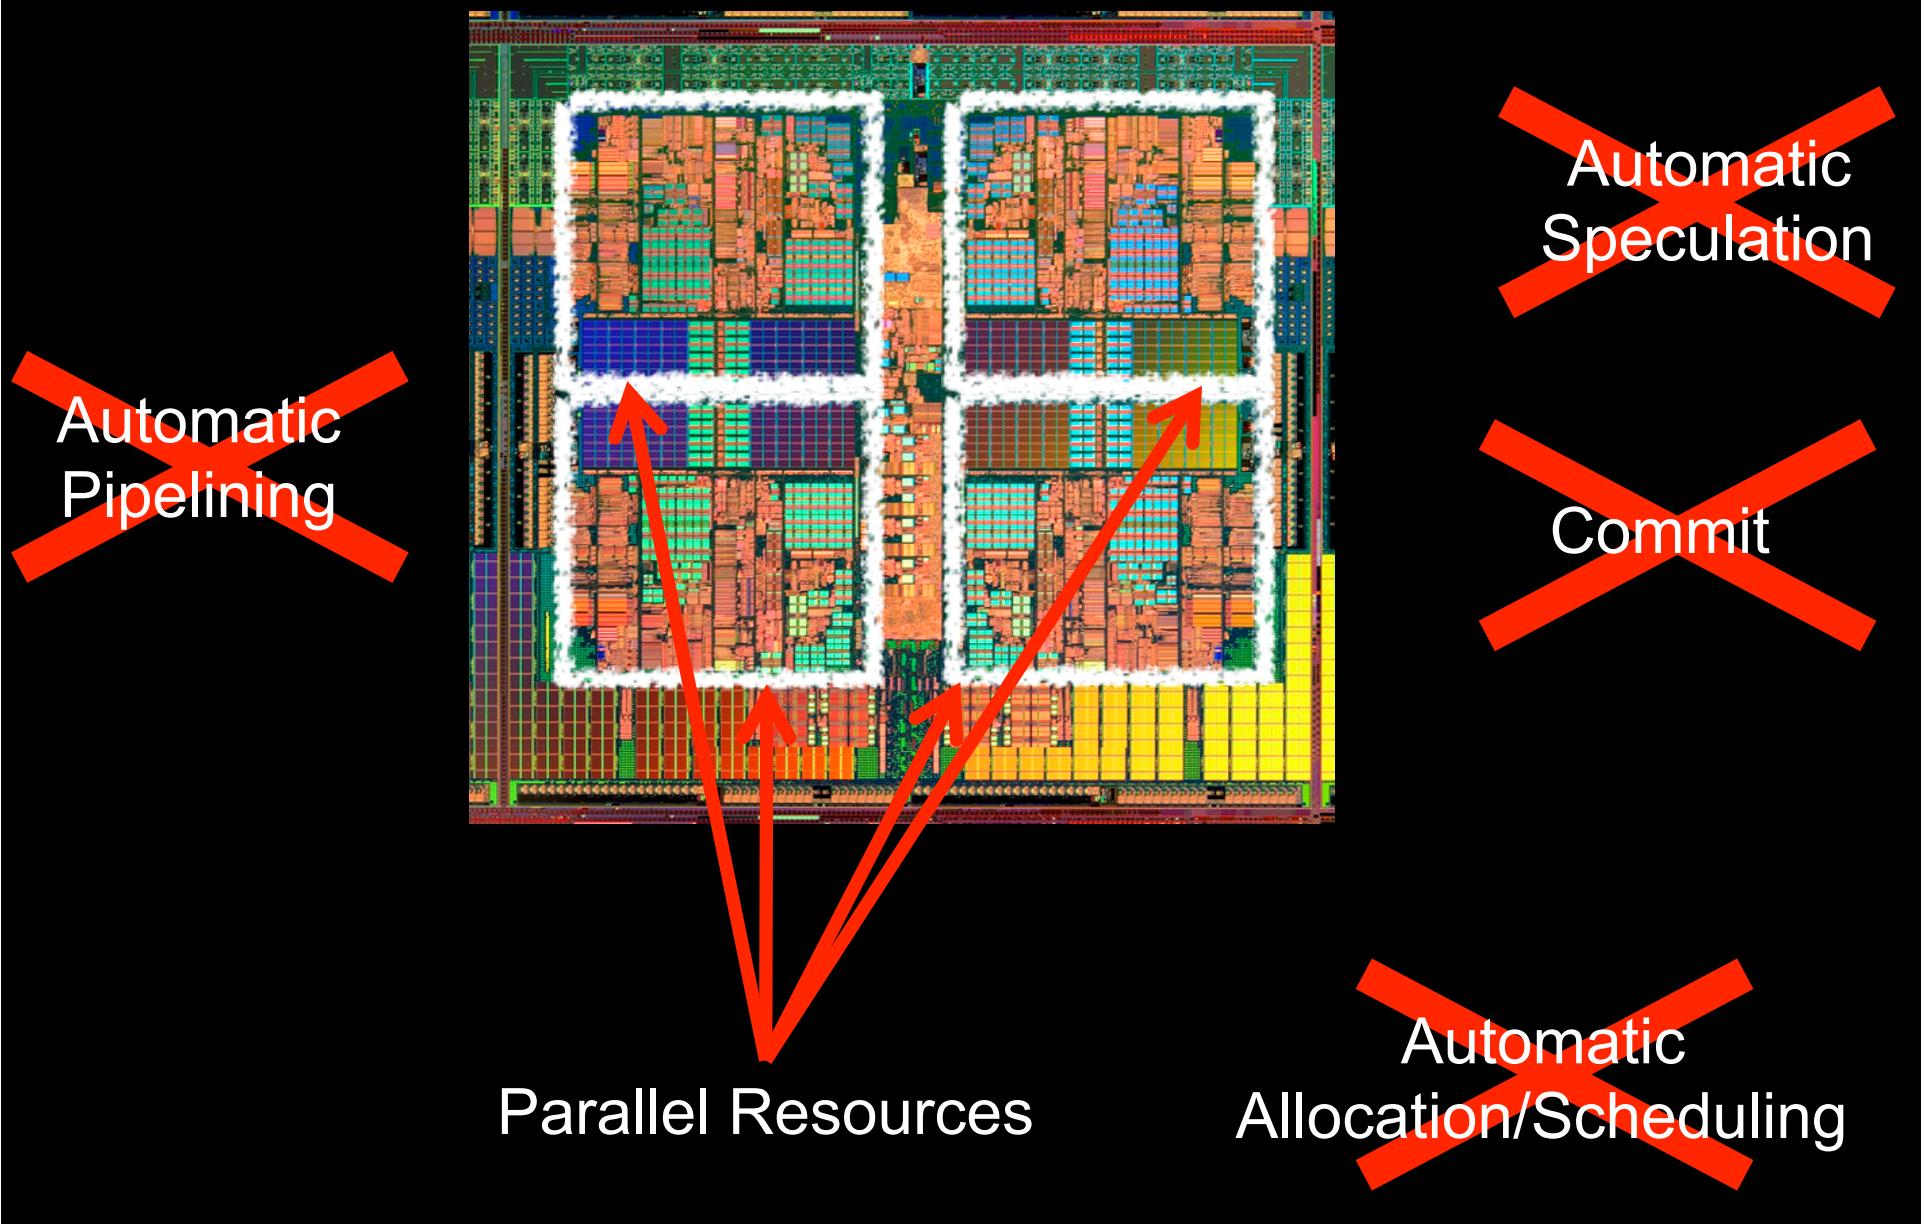
- Multicore
- Reconfigurable
- GPUs
- Clusters



# P6 SUPERSCALAR ARCHITECTURE (CIRCA 1994)

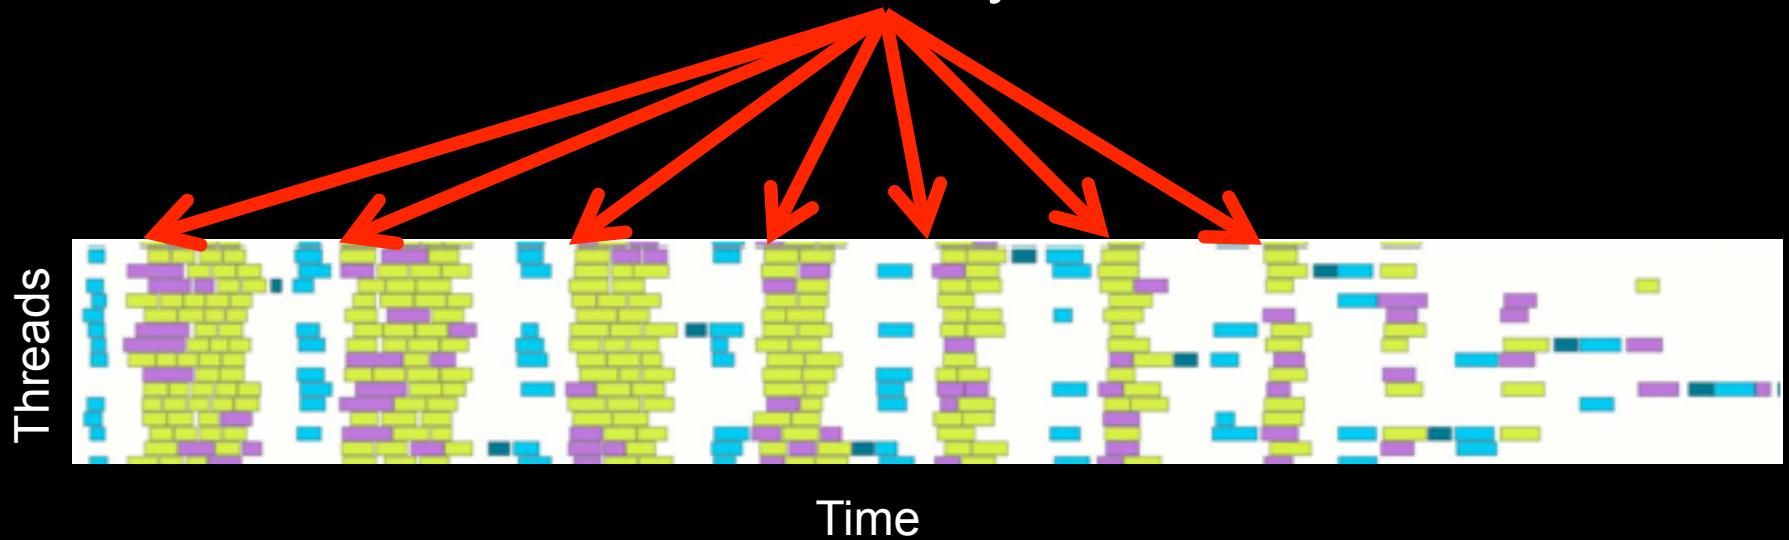


## MULTICORE ARCHITECTURE (CIRCA 2010)

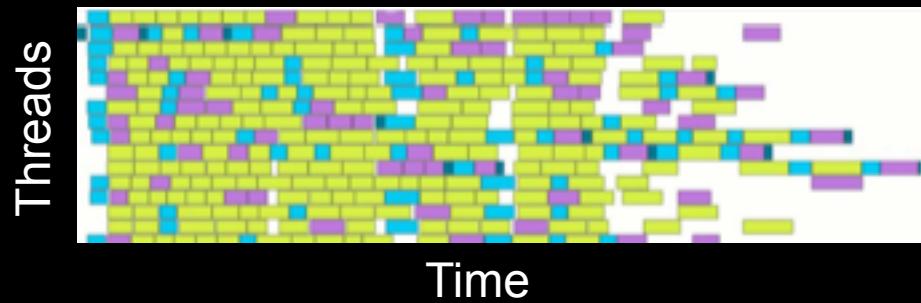


ABCPL	CORRELATE	GLU	Mentat	Paraphrase2	
ACE	CPS	GUARD	Legion	Paralation	pC++
ACT++	CRL	HAsL	Meta Chaos	Parallel-C++	SCHEDULE
Active messages	CSP	Haskell	Midway	Parallaxis	SciTL
Adl	Cthreads	HPC++	Millipede	ParC	POET
Adsmith	CUMULVS	JAVAR.	CparPar	ParLib++	SDDA.
ADDAP	DAGGER	HORUS	Mirage	ParLin	SHMEM
AFAPI	DAPPLE	HPC	MpC	Parmacs	SIMPLE
ALWAN	Data Parallel C	IMPACT	MOSIX	Parti	Sina
AM	DC++	ISIS.	Modula-P	pC	SISAL.
AMDC	DCE++	JAVAR	Modula-2*	pC++	distributed smalltalk
AppLeS	DDD	JADE	Multipol	PCN	SMI.
Amoeba	DICE.	Java RMI	MPI	PCP:	SONiC
ARTS	DIPC	javaPG	MPC++	PH	Split-C.
Athapascan-Ob	DOLIB	JavaSpace	Munin	PEACE	SR
Aurora	DOME	JIDL	Nano-Threads	PCU	SThreads
Automap	DOSMOS.	Joyce	NESL	PET	Strand.
bb_threads	DRL	Khoros	NetClasses++	PETSc	SUIF.
Blaze	DSM-Threads	Karma	Nexus	PENNY	Synergy
BSP	Ease.	KOAN/Fortran-S	Nimrod	Phosphorus	Telephros
BlockComm	ECO	LAM	NOW	POET.	SuperPascal
C*	Eiffel	Lilac	Objective Linda	Polaris	TCGMSG.
"C* in C	Eilean	Linda	Occam	POOMA	Threads.h++.
C**	Emerald	JADA	Omega	POOL-T	TreadMarks
CarlOS	EPL	WWWinda	OpenMP	PRESTO	TRAPPER
Cashmere	Excalibur	ISETL-Linda	Orca	P-RIO	uC++
C4	Express	ParLin	OOF90	Prospero	UNITY
CC++	Falcon	Eilean	P++	Proteus	UC
Chu	Filaments	P4-Linda	P3L	QPC++	V
Charlotte	FM	Glenda	p4-Linda	PVM	ViC*
Charm	FLASH	POSYBL	Pablo	PSI	Visifold V-NUS
Charm++	The FORCE	Objective-Linda	PADE	PSDM	VPE
Cid	Fork	LiPS	PADRE	Quake	Win32 threads
Cilk	Fortran-M	Locust	Panda	Quark	WinPar
CM-Fortran	FX	Lparx	Papers	Quick Threads	WWWinda
Converse	GA	Lucid	AFAPI	Sage++	XENOOPS
Code	GAMMA	Maisie	Para++	SCANDAL	XPC
COOL	Glenda	Manifold	Paradigm	SAM	Zounds
					ZPL

## Parallel Library Calls

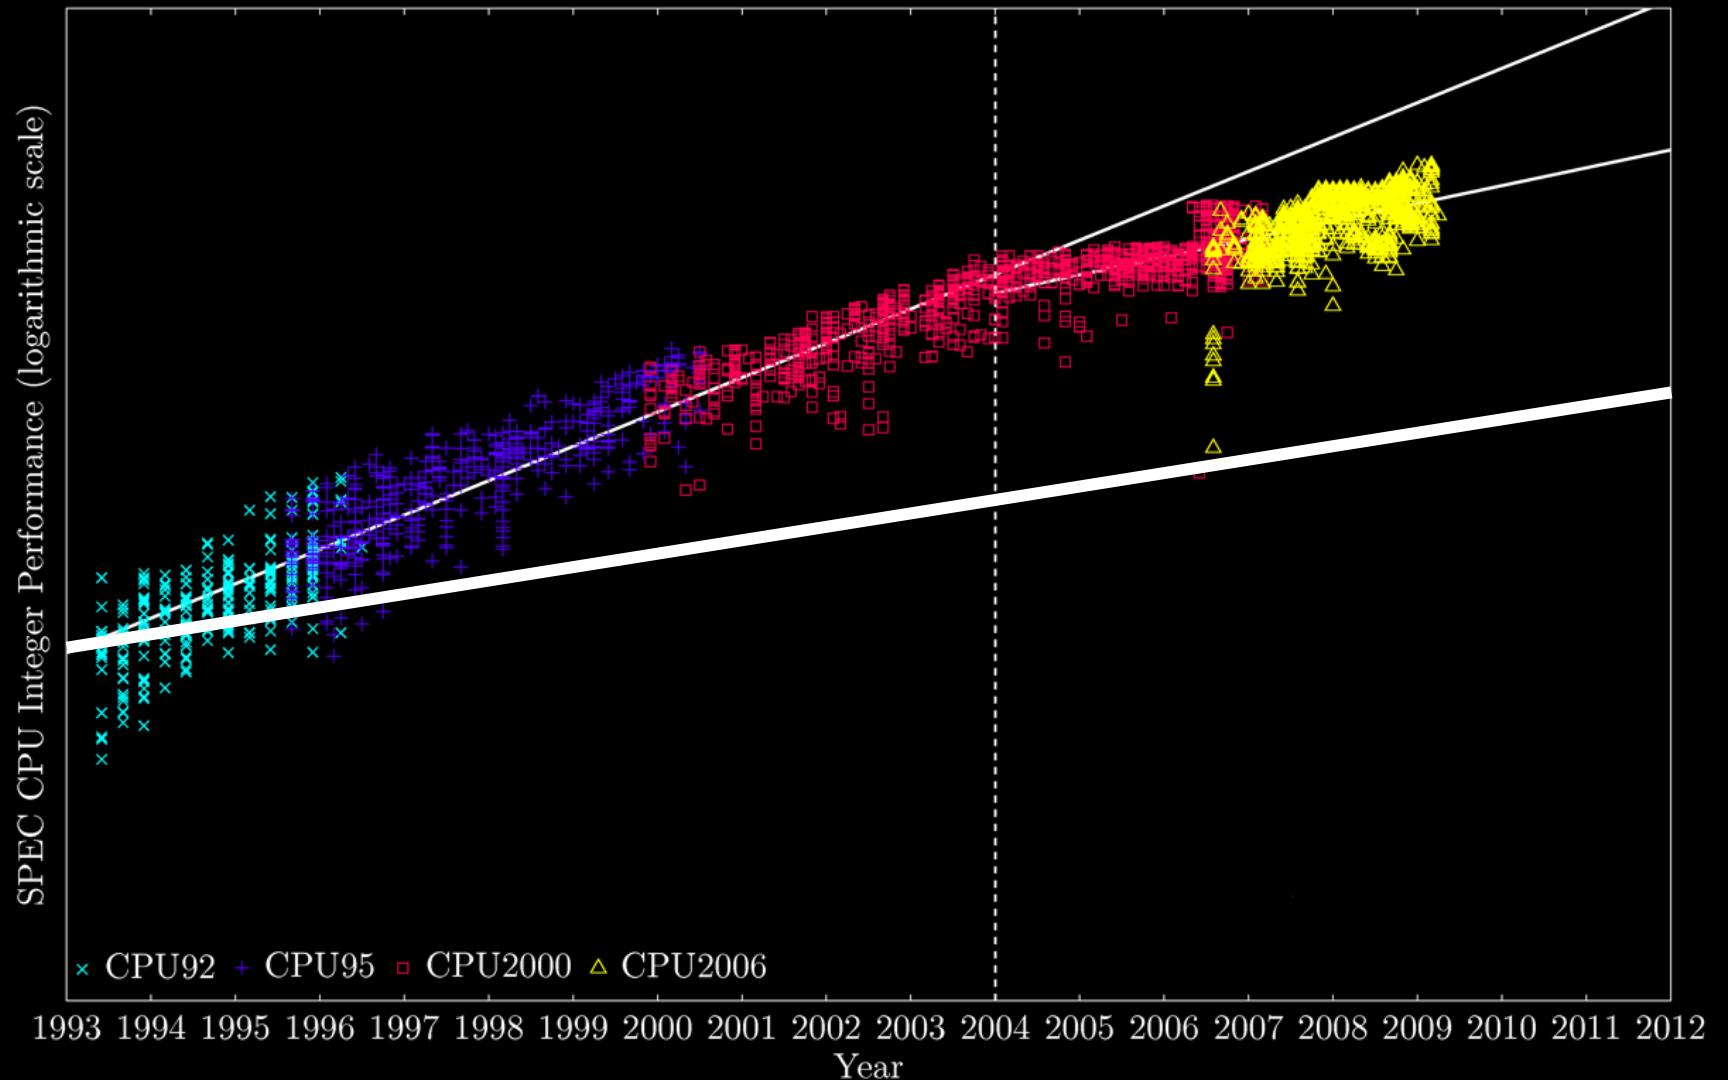


## Realizable parallelism



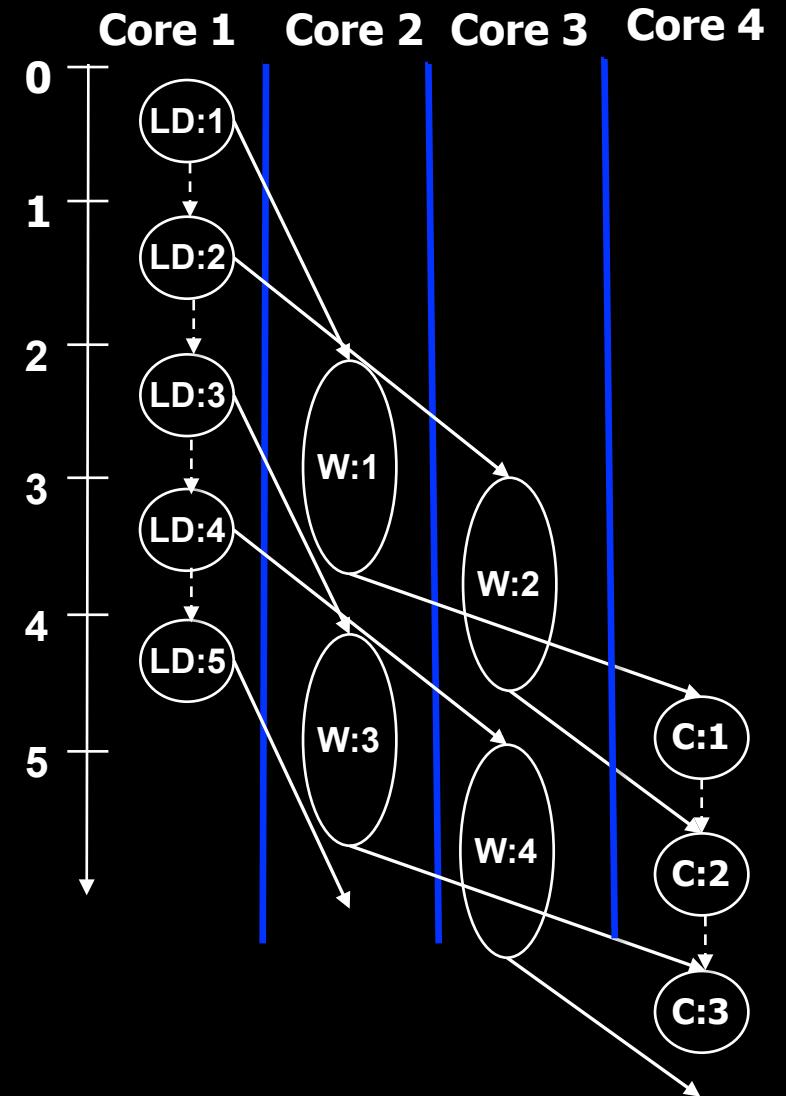
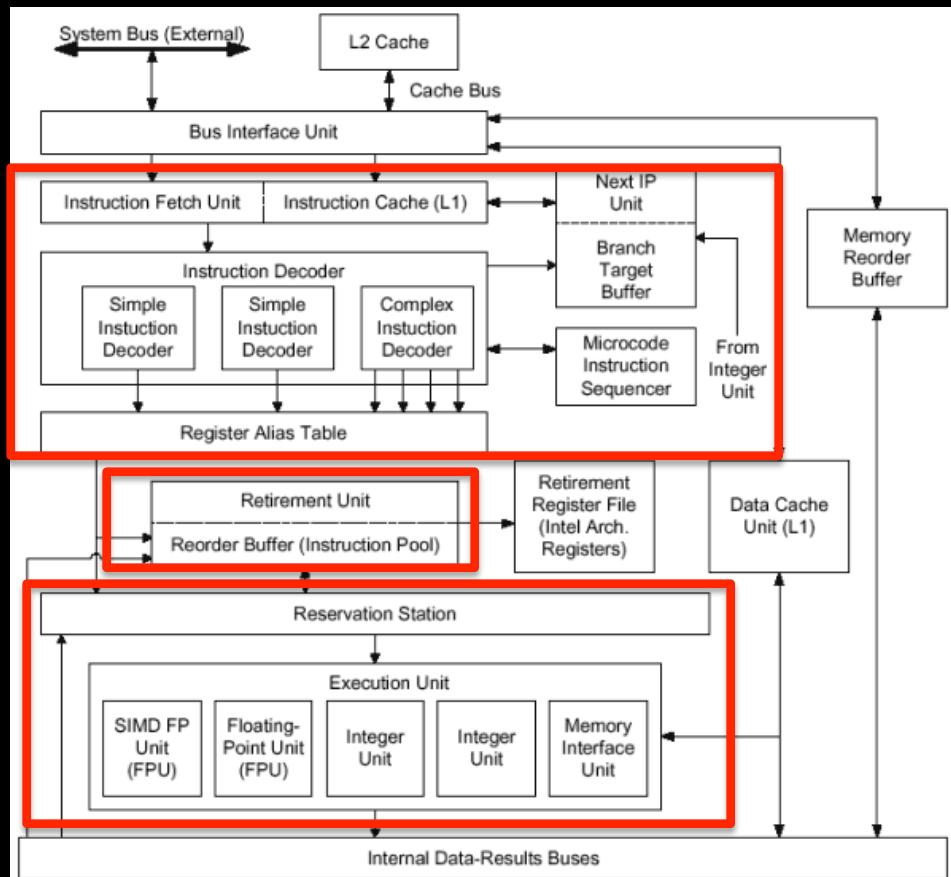
Credit: Jack Dongarra

***“Compiler Advances Double Computing Power Every 18 Years!”***  
– Proebsting’s Law



# Spec-PS-DSWP

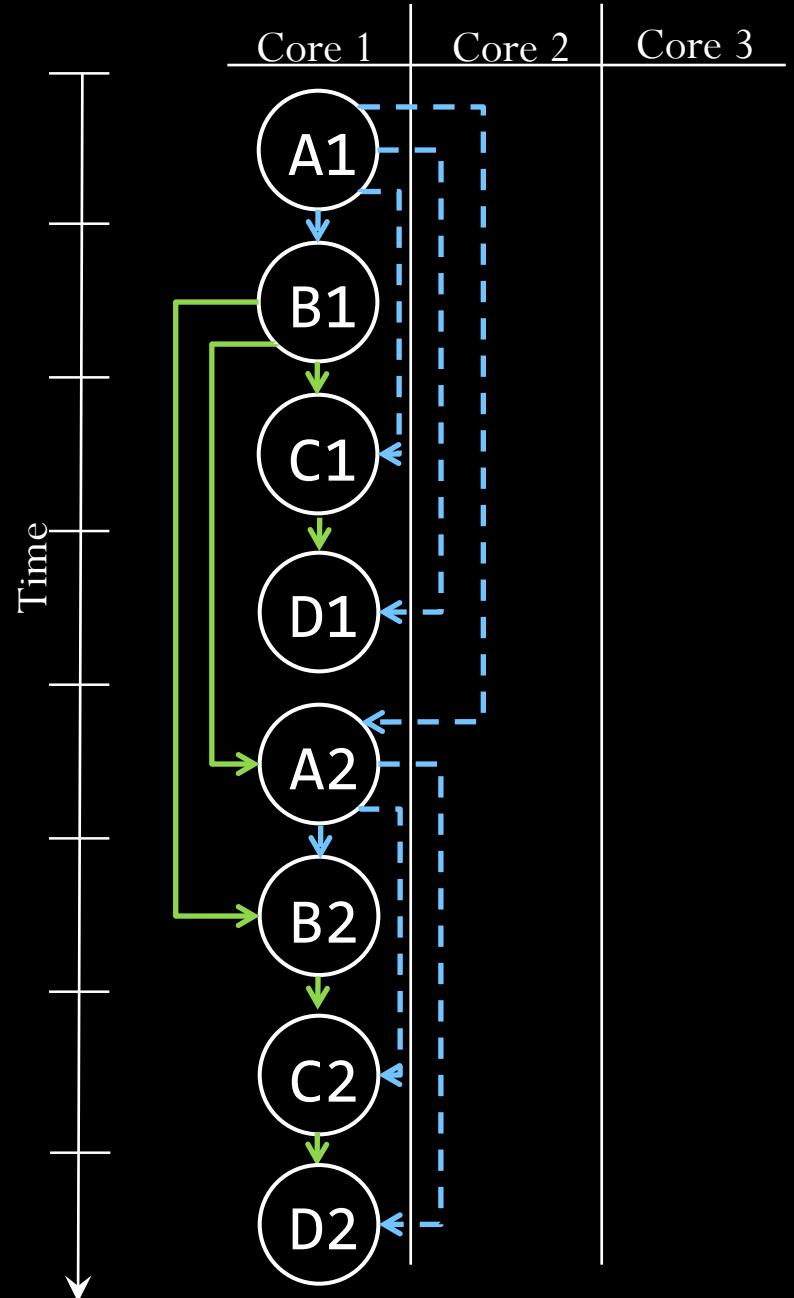
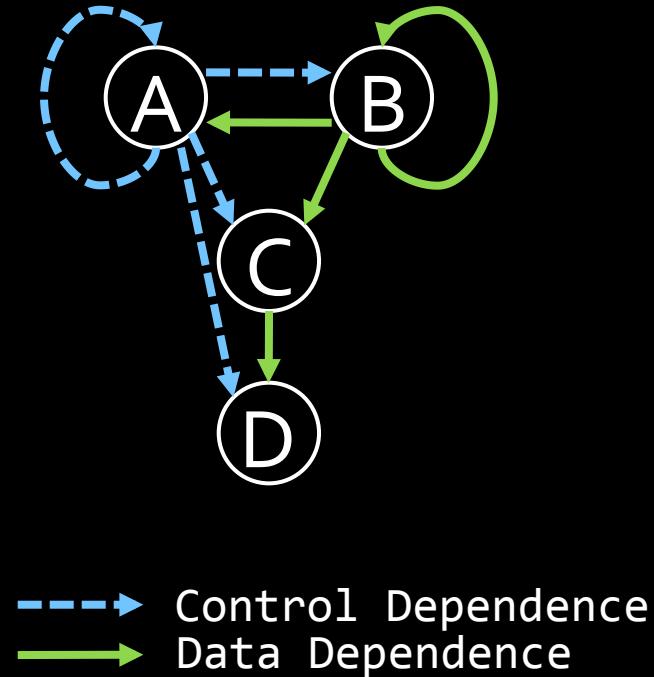
## P6 SUPERSCALAR ARCHITECTURE



### Example

```
A: while (node) {  
B:     node = node->next;  
C:     res = work(node);  
D:     write(res);  
}
```

Program Dependence Graph

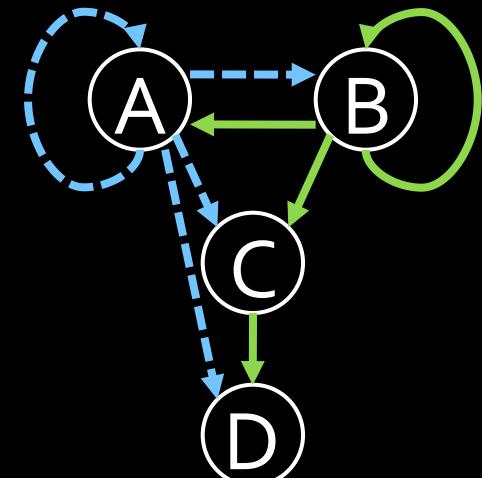


# Spec-DOALL

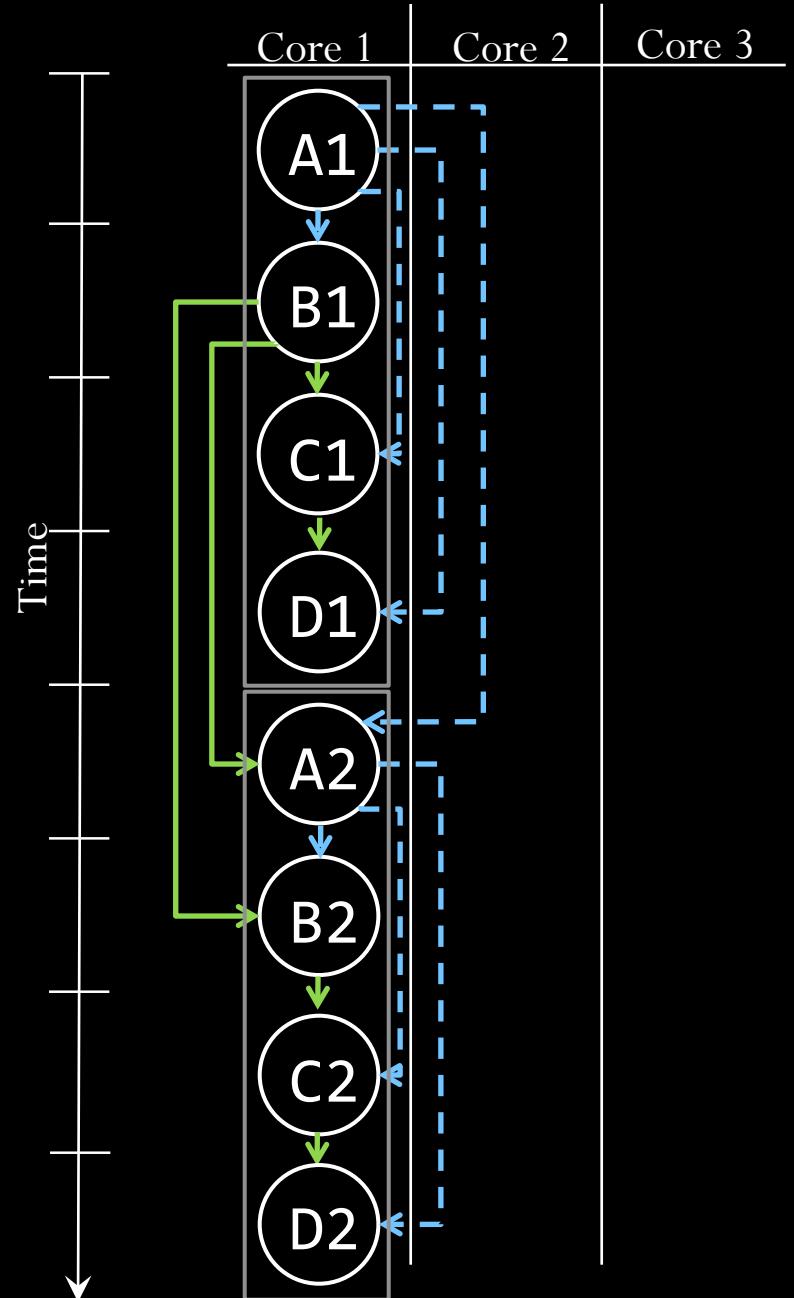
## Example

```
A: while (node) {  
B:     node = node->next;  
C:     res = work(node);  
D:     write(res);  
}
```

Program Dependence Graph



→ Control Dependence  
→ Data Dependence

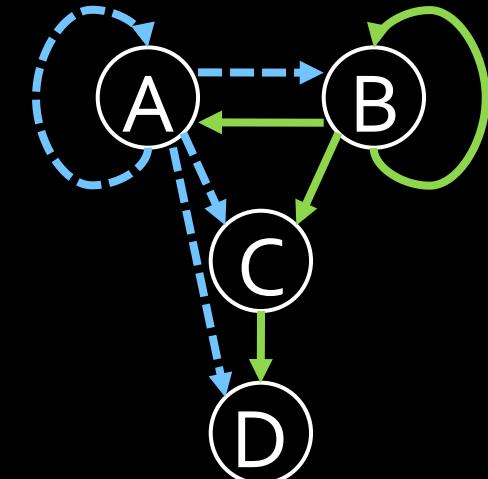


# Spec-DOALL

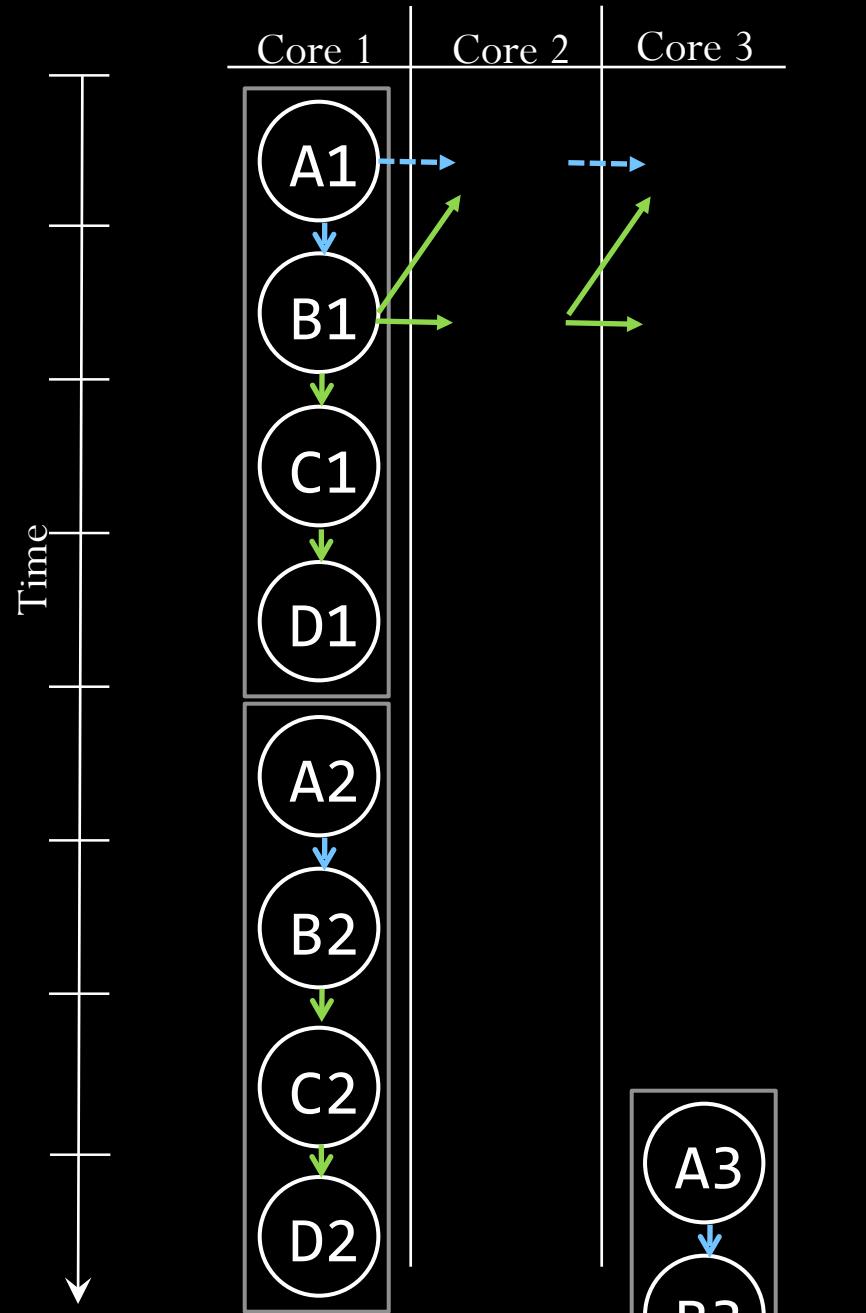
## Example

```
A: while (node) {  
B:     node = node->next;  
C:     res = work(node);  
D:     write(res);  
}
```

Program Dependence Graph



→ Control Dependence  
→ Data Dependence



# Spec-DOALL

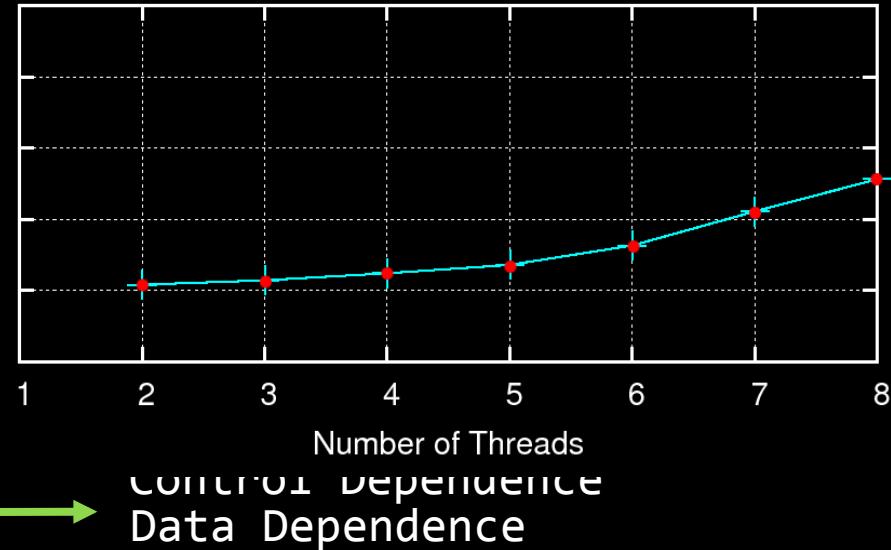
Example

```
A: while (node) {  
B:     node = node->next;  
C:     res = work(node);  
D:     write(res);  
}
```

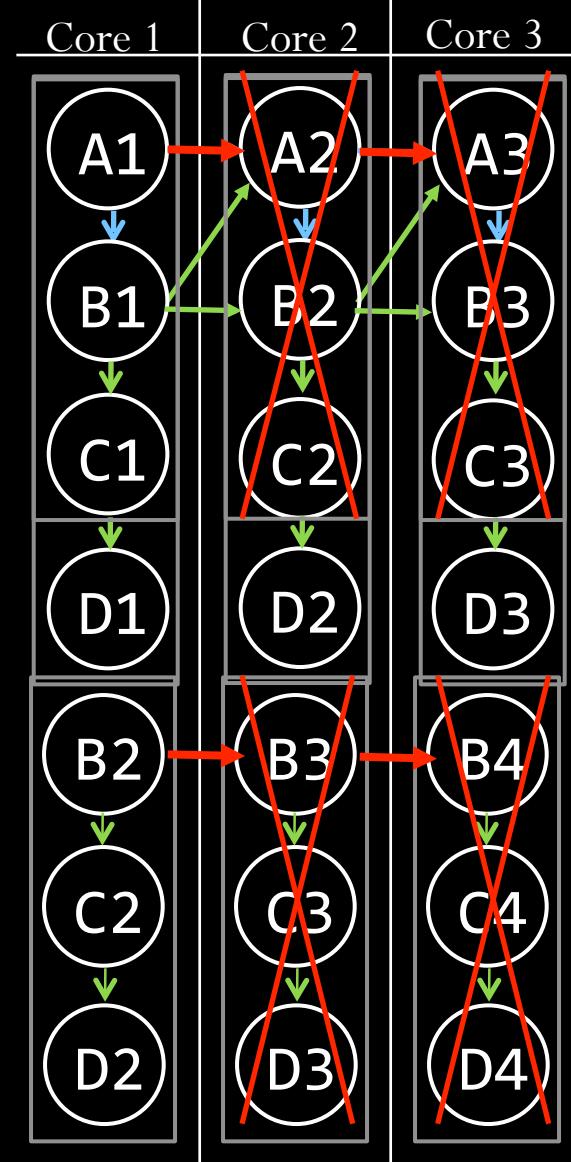
Program Dependence Graph



Slowdown

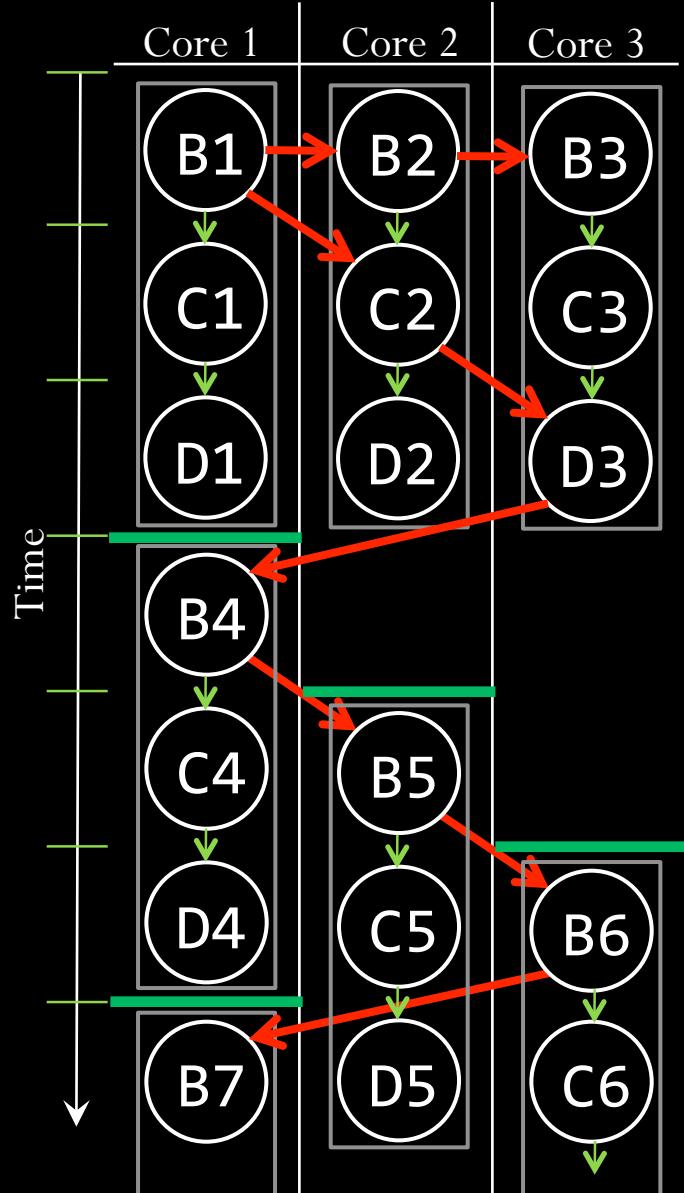


Time ↓



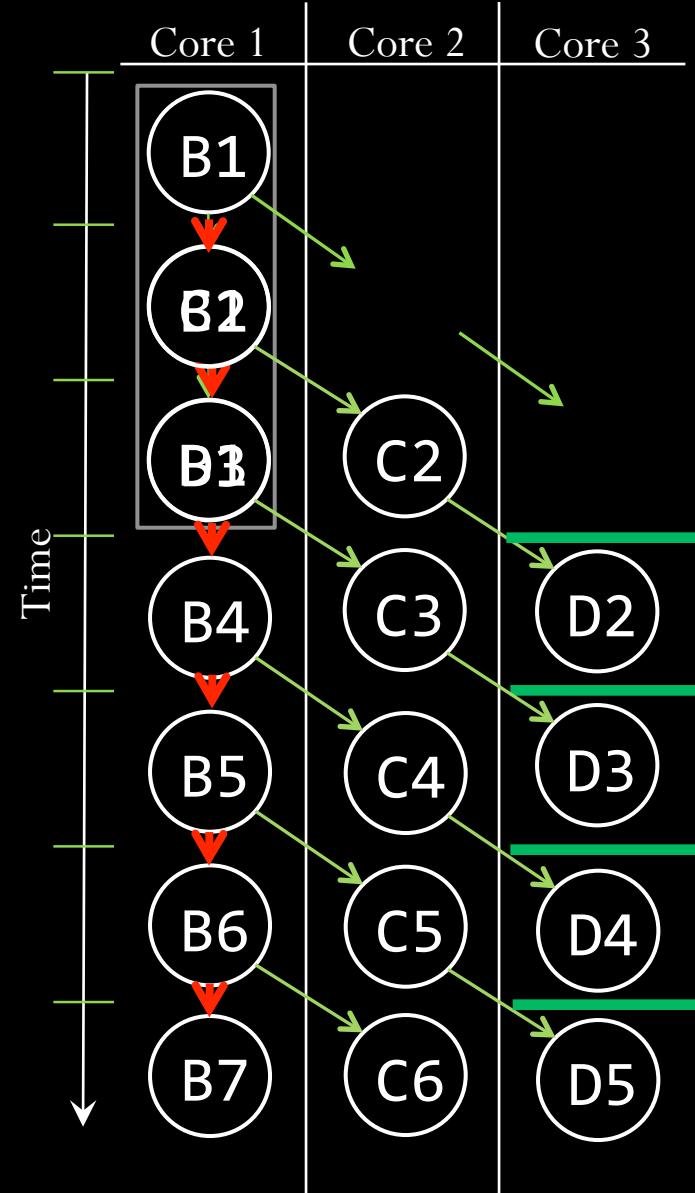
## Spec-DOACROSS

Throughput: 1 iter/cycle



## Spec-DSWP

Throughput: 1 iter/cycle



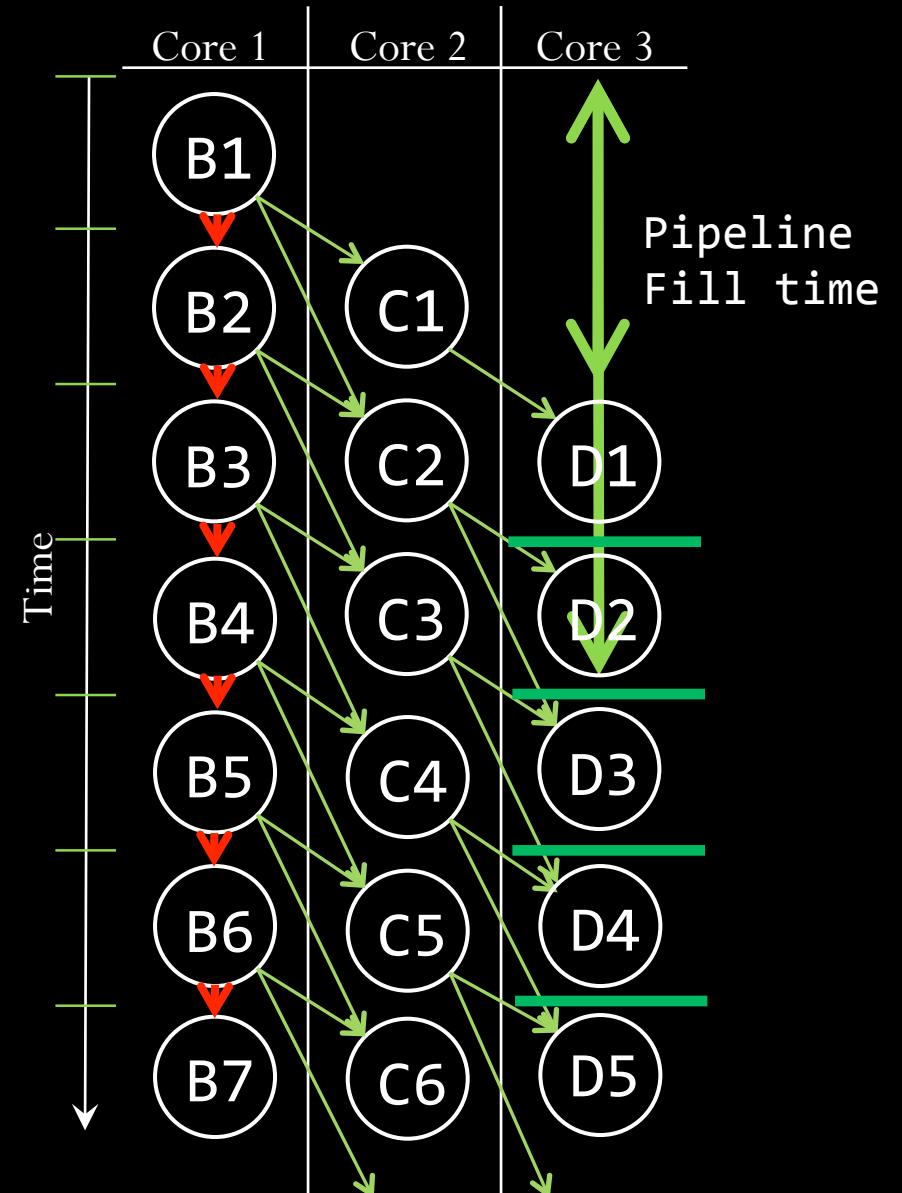
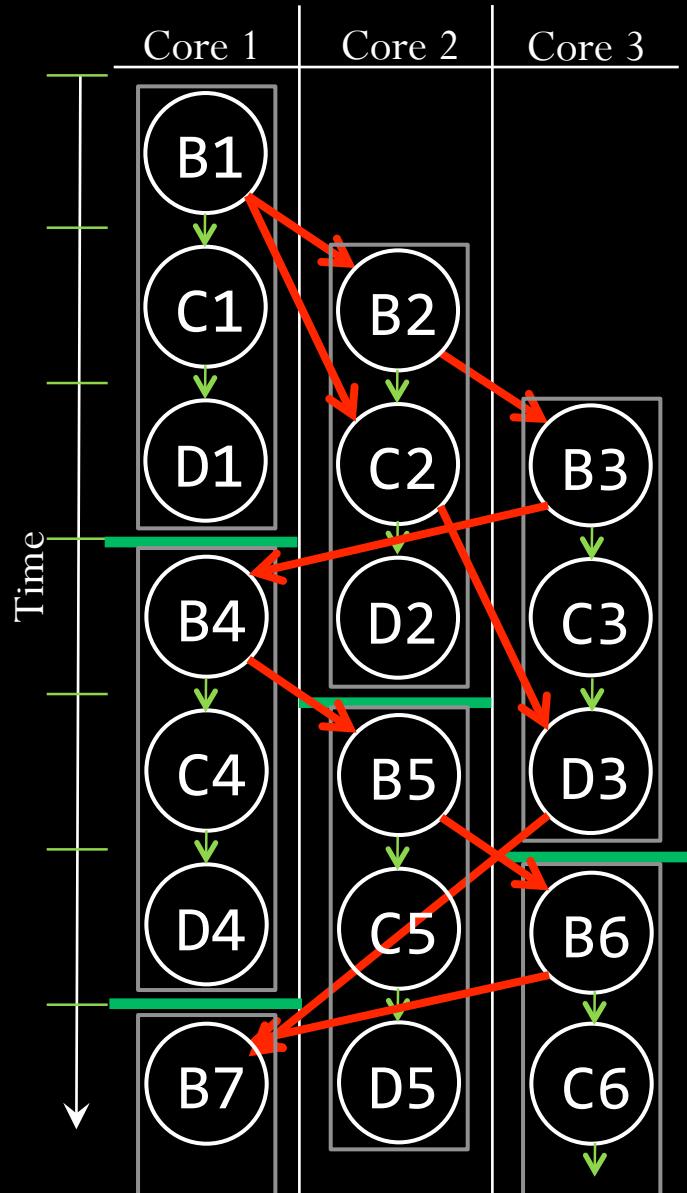
# Comparison: Spec-DOACROSS and Spec-DSWP

Comm.Latency = 1: 1 iter/cycle

Comm.Latency = 2: **0.5 iter/cycle**

Comm.Latency = 1: 1 iter/cycle

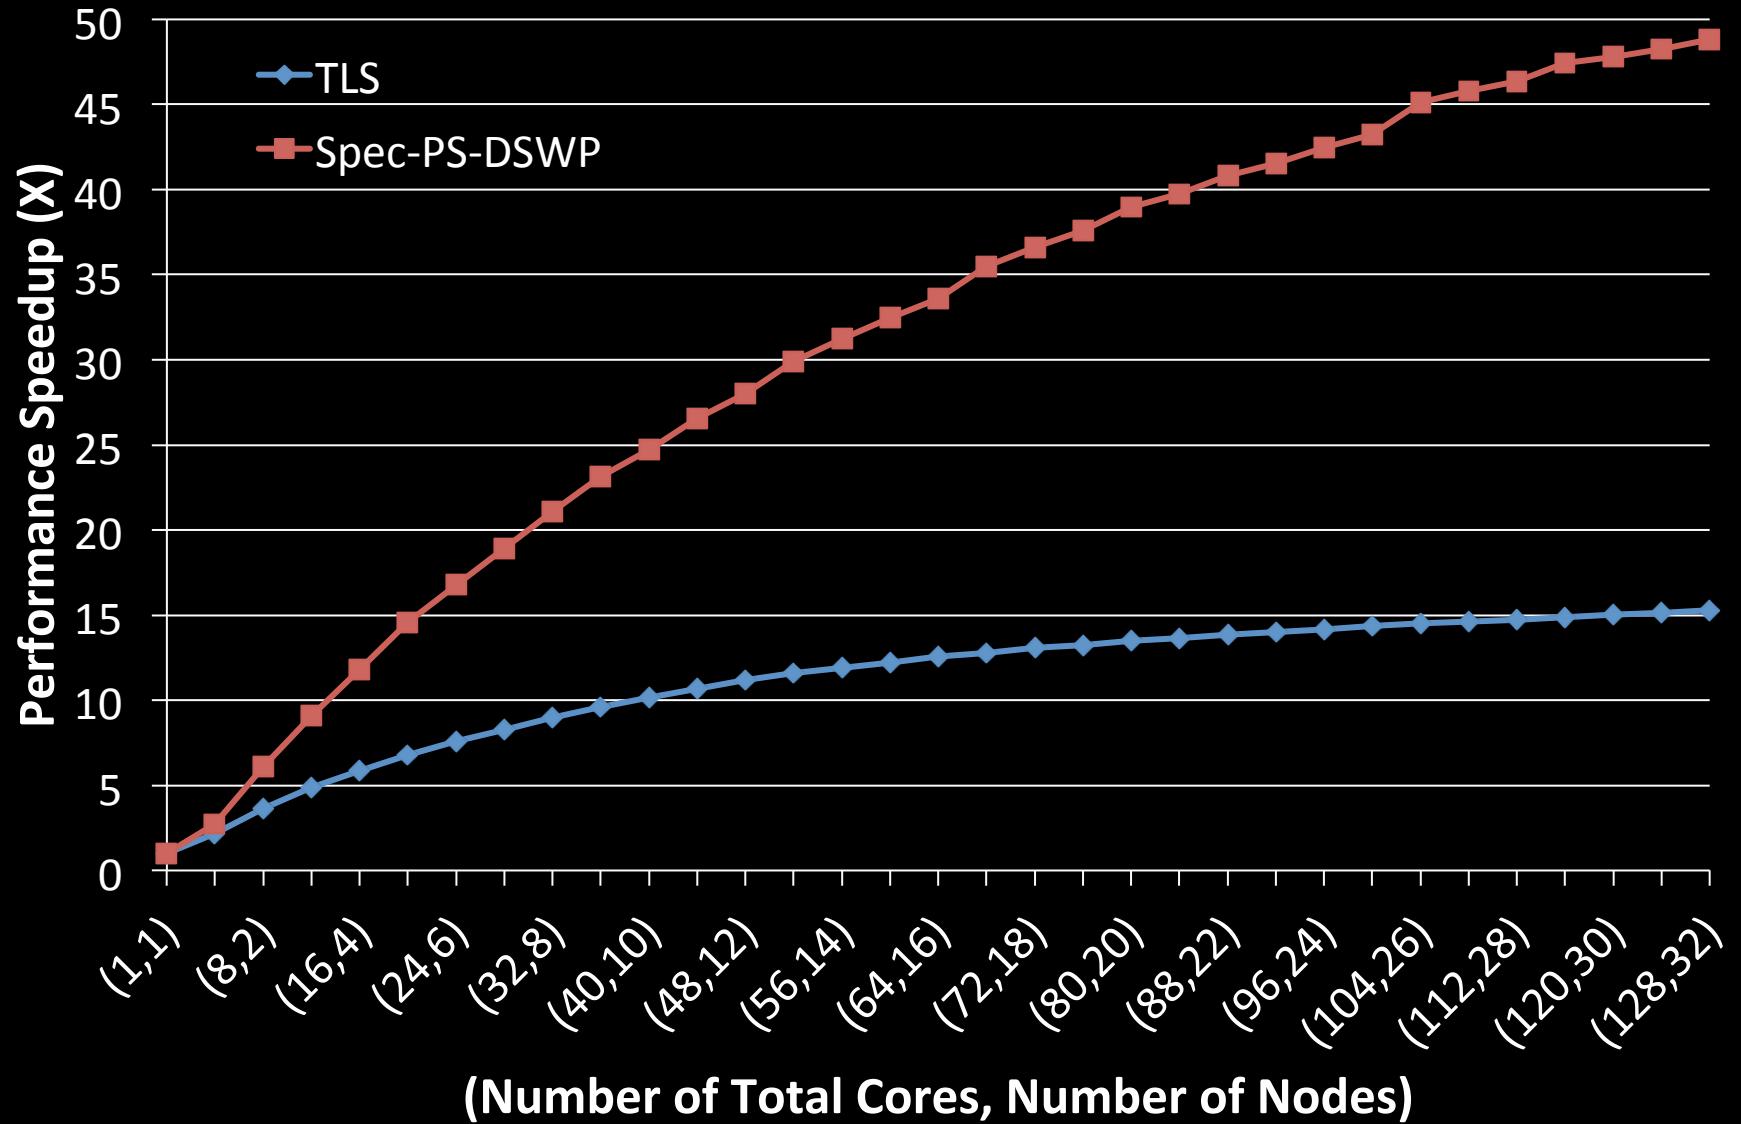
Comm.Latency = 2: **1 iter/cycle**



# Spec-DOACROSS vs. Spec-DSWP

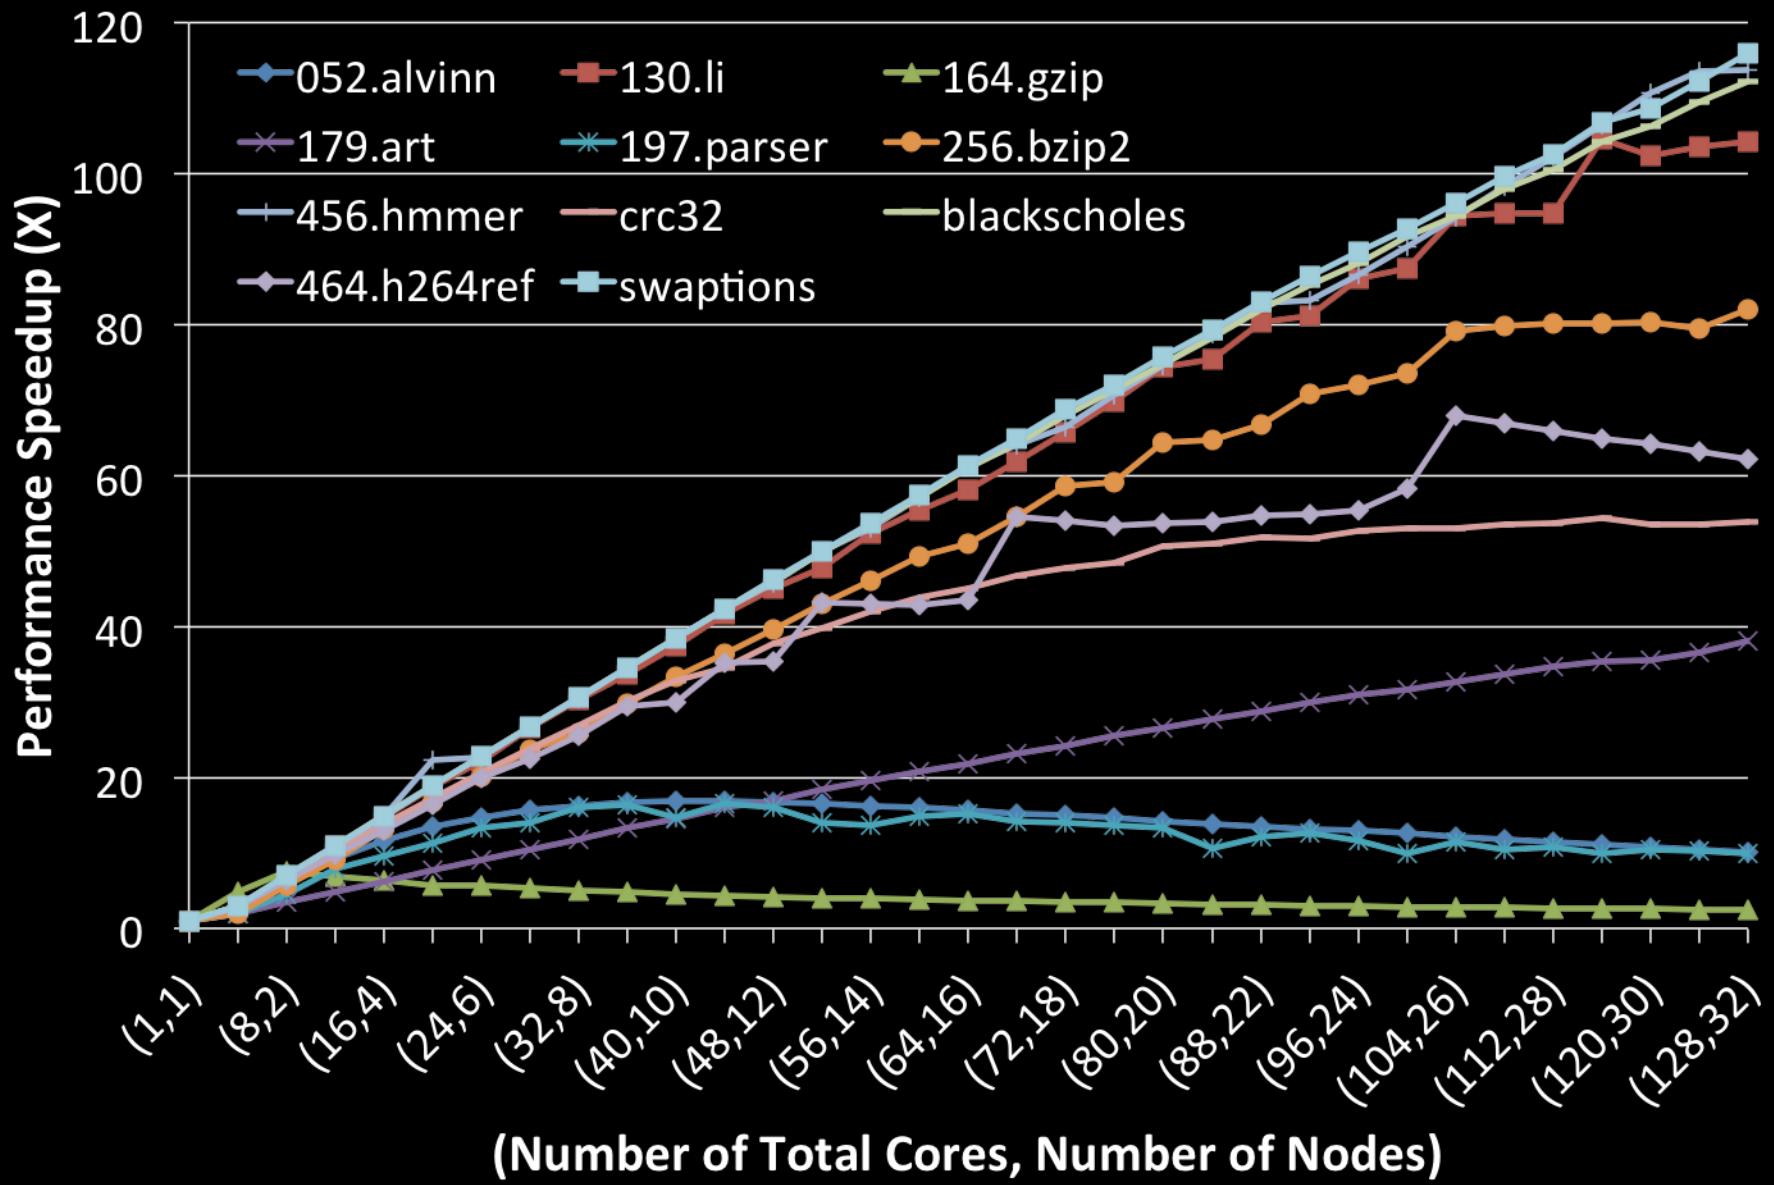
[MICRO 2010]

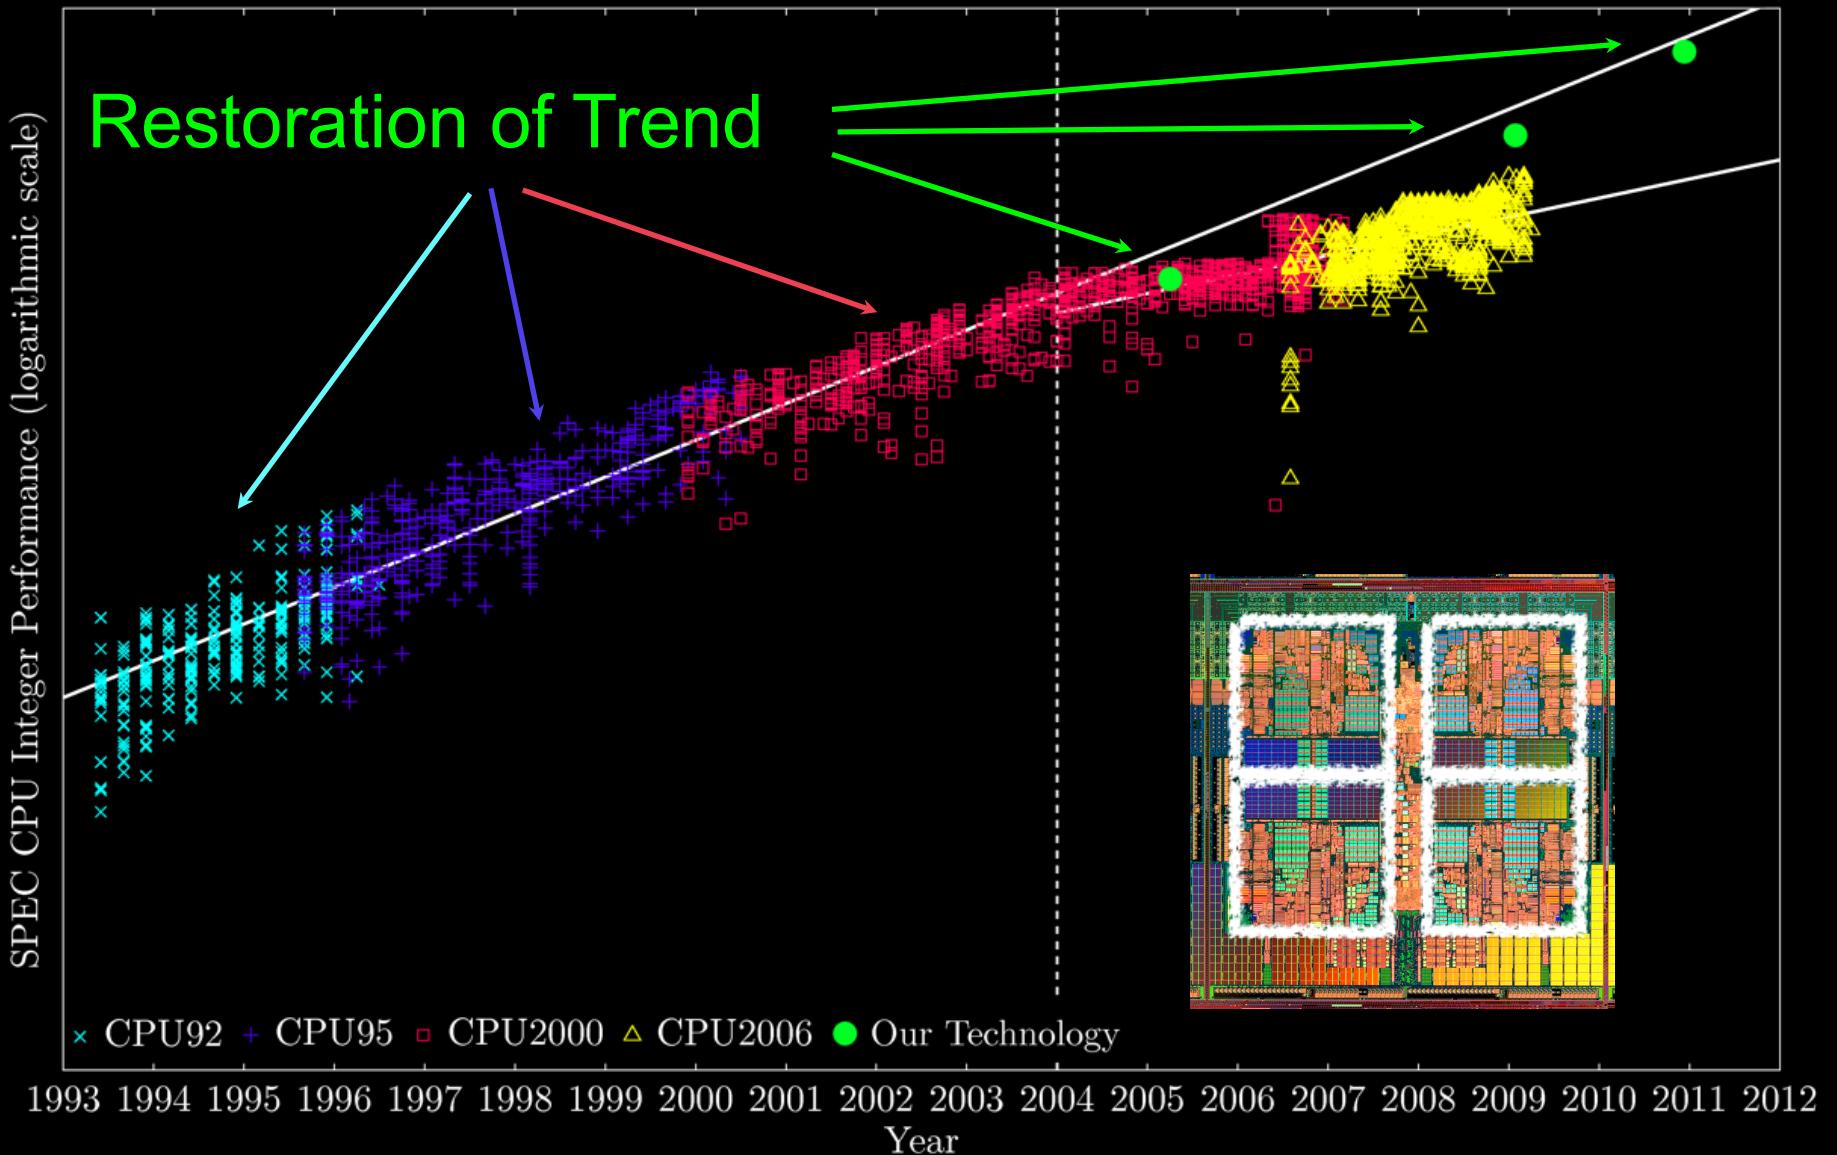
Geomean of 11 benchmarks on the same cluster



# Performance relative to Best Sequential

## 128 Cores in 32 Nodes with Intel Xeon Processors [MICRO 2010]





***“Compiler Advances Double Computing Power Every 18 Years!”***  
– Probsting’s Law

