


## → divergence

eg: while (true); // infinite computation.

→ looks stuck, but it's not waiting for anything / anyone.

→ FLAW in a sequential program

→ could be valid in a concurrent program → but don't use this! wastes CPU for no reason.



## THREAD TERMINATION.

→ Safety issues exist with ASYNCHRONOUS termination.



→ we don't really know the state of T2

→ what locks it holds?

→ what function/private state it's in?

eg: In early versions of Java, 64-bit integer operations were done by function calls.

→ In general, it's too dangerous since you can't be too sure exactly what's going on.

→ instead set a flag.

eg: pthreads → cancellation (polling a flag).

→ POSIX lets you call stop which sets a flag. SYSCALLS check the flag.

→ this is called CANCELLATION.

→ when you cancel a thread it sets the flag

→ flag is checked frequently, how frequently? ⇒ Nearly every BLOCKING syscall.

- Error Interrupt code: EINTR ⇒ every syscall that could return this checks the flag.

- Doesn't check on mutex-locks. (purely for performance) but there's an API to check yourself.

COMP 409  
OCT 21, 2010

PAUL HUBBARD

LAST TIME ⇒ finished looking at most of the core synchronization primitives

⇒ Deadlock & Termination ⇒ Coffman's Conditions → Must satisfy all 4 properties for resource deadlock. Avoid deadlock ⇒ other resources

⇒ divergence as a possibility.

⇒ Termination → dangerous to stop/destroy for subtle reasons

⇒ Cancellation → polling method in pthreads.

## Cancellation Cont'd

pthread\_cancel (&t);

→ sets the flag TRUE

→ flag will be checked in various places, but there's also a way to check at a specific time.

→ CANCELLATION HANDLES.

→ it's a stack → you can push/pop handles.

→ handlers get run on exit.

eg: m.lock();

→ push new handler to unlock m if it gets cancelled.

→ pop handler

m.unlock(); ← from C.

→ Very much like atexit, except for Threads

→ in pthreads, you can turn cancellation on or off at any point, WITHOUT affecting flag stacks.

→ PTHREAD\_CANCEL\_DISABLE; PTHREAD\_CANCEL\_ENABLE.

→ Cancellation is not recommended, but if you want → PTHREAD\_CANCEL\_ASYNCHRONOUS literally kills the thread even if it's in the middle of something.


by default → PTHREAD\_CANCEL\_DEFERRED: safely terminates a thread.

JAVA → doesn't have cancellation, but DOES have interrupts.

→ try { ... } catch (InterruptedException ie) { ... }

↑ eg: Thread.sleep();

Thread.interrupt sets a flag in the target thread which gets periodically checked, and if it's interrupted, it puts you into the catch clause ⇒ do whatever you want.





NOTE: NOT EVERY BLOCKING CALL CAN BE INTERRUPTED/CANCELLED

→ there's no guarantee with this method that it stops right away (or at all)

→ Interrupted IO Exception

Concurrency Problems (the classic ones)

→ Producer Consumer (as we've seen) → used very often eg: buffers etc.

→ Dining Philosophers (as we've seen) → deadlocks

→ Oriental Gardens (assignment 1)

→ great-big garden



→ Readers & Writers (as we've seen)

→ One-lane bridge

- Similar to reader/writer

- can allow multiple east-bound at once,

and can have multiple west-bound at once.

⇒ "Readers & Other Readers"




→ Unisex bathroom

→ one door, can't have men + women together

→ fairness → don't want to keep people waiting too long

→ Rollercoaster Problem

→ let threads be the people riding the rollercoaster.



→ say car holds 4 people.

→ wait for 4 people to sit

→ only then can it START

→ once moving no rider does anything ⇒ threads SLEEP

→ at the end, they must be woken up

→ wait for them all to leave

→ then repeat

→ Cigarette smoker's Problem

→ 3 smokers sitting around some sort of table.


→ they sit there making or smoking cigarettes.

→ ingredients: tobacco, papers, match

→ each smoker only has ONE of the 3 ingredients (but an infinite supply).

→ a mysterious agent deposits two ingredients on the table.

→ originally this was proposed as something you needed condition synchronization for (to avoid an inherent deadlock situation), but this was not true.



→ Sleeping Barber (most complex)

→ Multiple Client-Server Interactions

→ considers rendezvous situations → eventually two threads end up at the same place to communicate.

→ Barber shop has an entrance & an exit.

→ customers show up, take a chair when available, notifies a barber, sleeps until done, he wakes us up, then sleeps himself until someone else notifies/wakes.

→ sequence of steps

→ barber sleeps

→ customer comes in

→ wakes up barber

→ sits in chair (sleeps)

→ barber cuts hair

→ wakes up customer when done

→ customer leaves

→ barber goes back to sleep.

→ if someone is already in chair, customer sits in a wait chair until barber notifies the customer that a chair is free.



→ CHECKING ATOMICITY → "should be fairly important to a lot of us"

→ BE VERY CAREFUL TO AVOID RACE CONDITIONS!!!

- we can have failures of atomicity with OR without race conditions, eg:  $\left[ \begin{array}{c} T1 \\ x++ \end{array} \quad \begin{array}{c} T2 \\ x++ \end{array} \right]$

- there is a technique to CHECK ATOMICITY.

→ starting point for checking atomicity in general, but these processes are VERY EXPENSIVE.

→ Lipton 1975 → "left/right movers"

→ reduction technique → makes larger atomic chunks, & thus have a simple program.

→ take a statement S. We want to make it atomic.  $S \rightarrow \llbracket S \rrbracket$ .

→ we can't naively do this everywhere or it could change the semantics/meaning of the program.

→ places to consider.

→ whether the program terminates or not.

→ equivalence in terms of halting.

eg:  $a = b = 1$ .

$\left[ \begin{array}{c} P(a) \\ P(b) \\ V(a) \\ V(b) \end{array} \right]$   $\left[ \begin{array}{c} P(b) \\ P(a) \\ V(b) \\ V(a) \end{array} \right]$

↑ marking - this atomic changes behavior (deadlock).

→ with no atomicity, deadlock is possible in different ways, so we can halt.

→ if its atomic, it is infinite → doesn't halt.

→ So: can't arbitrarily make chunks of code atomic.

→ So when can we make stuff atomic.

→ set of rules:

Rule 1: If statement S is ever entered, it must be possible to leave (eventually)

- this alone is not yet enough

eg:  $x = y = 0$ .

$x = 0 \rightarrow \left[ \begin{array}{c} x = 1 \\ y = x \\ P(y) \end{array} \right]$

is this valid as an atomic statement?  
y would always be 1 at this point.

→ NOW IT NEVER HALTS

(if  $y = 0$ , this halts)

↳ changed behavior. (BAD)

COMP 409  
OCT 26, 2010

\* Assignment #2 ⇒ as written → no dependencies  
(Problem) fix: loop body

$a[i] = f[i]$

if ( $i > 1$ )

$b[i] = b[i-1] * c[i-2]$

else

$b[i] = g[i]$

$c[i] = a[i+3] + b[i]$

Last time: atomicity checking  $S \Rightarrow \llbracket S \rrbracket$ , behavior differences, different data output, halting behavior  
- continuing set of rules. (R1 is not enough).

Rule 2: The effect of statement  $\llbracket S \rrbracket$  must be the same as S.

- Let f, g & h be statements inside a program, and let  $\alpha$  be some sequence of statements from start.

- f is a RIGHT MOVER if for any computable  $\alpha$  f h, when f, h are executed in different threads, then  $\alpha$  h f is also a computation, and the result (ie effect on all vars) of  $\alpha$  f h and  $\alpha$  h f are the same.

- g is a LEFT MOVER if for any computation  $\alpha$  h g where g & h are executed in different threads, the  $\alpha$  g h is a computation and the result of  $\alpha$  g h and  $\alpha$  h g is the same.



so we will use this to decide what is atomic and what is not.  
 left movers: release, unlock, V() operations  $\rightarrow \alpha P()h \Leftrightarrow \alpha h P()$   
 right movers: acquire, lock, P() operations  $\rightarrow \alpha h V() \Leftrightarrow \alpha V()h$

D-Reduction


$S_1, S_2, S_3, \dots, S_n \Rightarrow [[S_1, S_2, S_3, \dots, S_n]]$   
 sequence of statements  $\Rightarrow$  atomic version of same sequence.  
 This can be made with D-reduction.

So for some  $i$  we have  
 $S_1 \dots S_{i-1}$  as all right-movers (NB:  $S_i$  is unconstrained)  
 $S_{i+1} \dots S_n$  as all left-movers

\*  $\rightarrow$  and each of  $S_2 \dots S_n$  is executed  
 Then:  $[[S]]$  is a D-reduction of  $S$ .

- If  $S'$  is a D-Reduction of  $S$ , then  $S$  halts IFF  $S'$  halts.

Example: Launching 3 threads:



we consider particular sequences:

$P(c) \xrightarrow{\text{right}} V(b) \xleftarrow{\text{left}} V(c) \xleftarrow{\text{left}}$  and decide that this sequence can be atomic

$[[P(c) V(b) V(c)]]$  which makes the problem much smaller.

- we can do this again, and again... we notice that with  $[[P(c) V(b) V(c)]]$  the P(c) and V(c) happens no matter what and since it's atomic, it's similar/equivalent to V(b) on its own.
- Then after, we see this;  $\begin{matrix} P(a) & P(b) & P(c) \\ \cancel{V(a)} & \cancel{V(b)} & \cancel{V(c)} \end{matrix} \Rightarrow$  this program does nothing.
- This is actually used in practice (research practice) as an atomicity checker, but it slows down the code by a factor of 10.
- This was done in (2003: Flanagan, Qader).
- This is NOT EASY!
- And they actually did find several atomicity bugs in JAVA Runtime Library, which were later corrected.

Scheduling & Priority

- there are different mappings
- $\rightarrow$  The OS used to schedule PROCESSES. Not threads.
- $\rightarrow$  Now the OS schedules KERNEL THREADS, which can then create USER THREADS.
- $\rightarrow$  A LIGHT-WEIGHT-PROCESS (LWP) is a KERNEL thread partitioned differently.

Models: One-to-One

- User Threads ARE Kernel threads (some thing).
- Advantage  $\rightarrow$  Real Concurrency with multiple CPUs.
- Disadvantage  $\rightarrow$  everytime we make a thread, there's MORE OVERHEAD through the OS.

Many-to-One (like green threads)

- many user threads on ONE KERNEL thread.
- Advantage  $\rightarrow$  Requires no limited OS support, LESS overhead (NB: avoid getting stuck in blocking calls)
- Disadvantage  $\rightarrow$  No Real Concurrency with several CPUs.



Many-to-Many (arbitrary mapping, most flexible model  $\rightarrow$  most general, gives OS more freedom)

- $\rightarrow$  most flexible,
- $\rightarrow$  processor affinity  $\Rightarrow$  assigning certain things to specific CPUs.

## PRIORITIES

$\rightarrow$  **DON'T USE THESE!** (most often not necessary).

- $\rightarrow$  With RealTime priority  $\rightarrow$  easy to lock yourself out of the system.
- $\rightarrow$  there's different priority in different contexts.

Java has 1-10, where 10 is high priority.

$\rightarrow$  windows NT had 7 priority levels, so Java's 1-10 gets squashed somewhere in the 1-7.

$\rightarrow$  Linux: 128 levels  $\rightarrow$  is priority 10 in Java = 128 in Linux?  $\neq$  10? we don't know.

$\rightarrow$  Pthreads:  $\geq 32$  levels of priority.  $\rightarrow$  sometimes it gets mapped down to fit, or mapped up & spread.

$\rightarrow$  in Pthreads priority levels are separate for SCHED-RR and SCHED-FIFO.

$\rightarrow$  how these mix?  $\rightarrow$  SOMEHOW, (we don't know).

## Priority Inversion

$\rightarrow$  3 threads with 3 different priorities: Low, Med, High.

If Low gets a lock before the others start and a High wants the lock, it's blocked.

- So at (\*HERE) in time (see graph  $\rightarrow$ ), High is effectively blocked by both threads that are lower priority!

Note: Med blocks L from unlocking, which blocks High.

$\rightarrow$  This actually happened

$\rightarrow$  Mars Pathfinder


- $\rightarrow$  info-bus with lock to protect it.
- $\rightarrow$  bus manager thread is HIGH priority.
- $\rightarrow$  meteorological data is LOW priority.
- $\rightarrow$  communication thread is MED priority.

$\rightarrow$  Problem: Priority Inversion.

Scenario: Meteorological Thread (Low) gets lock, bus manager makes up, starts executing but can't get the lock, then communication starts running (takes a LONG time)

There's a separate "watch-dog" process.


- if anything looks weird, it reboots the system.
- it spontaneously reboots, and the reason was hard to find, but this is why



## TWO SOLUTIONS TO PRIORITY INVERSION


### 1) Inheritance

- Holder of a lock gets HIGHEST priority of anyone else trying to acquire it.



### 2) Ceilings

- $\rightarrow$  give the lock a priority.
- a thread holding a lock has the PRIORITY OF THE LOCK.



COMP 409  
OCT 28 2010

PAUL NUSSMAN

Last time  $\Rightarrow$  checking atomicity using Left-movers & Right-movers. (basis of a realistic technique).

- scheduling
- priorities  $\rightarrow$  problems: priority inversion

## Miscellaneous Java/Pthreads

$\rightarrow$  ONE-TIME Initializers (for convenience):

- often we need to do something once after we are already multi-threaded.
- global variable of type pthread\_once\_t initialized to PTHREAD\_ONCE\_INIT.
- now threads can call pthread\_once(&once, &function).  $\Rightarrow$  many threads call it, but it only happens once.



THREAD SPECIFIC DATA

- errno (global) records the error value whenever you make a system call.

```

T1
x = read()
(errno = BADF)
if (x < 0)
    printf("%d", errno);

T2
x = read()
(errno = 0)
    
```

← Problem: errno could be wrong depending on order of execution.

- Solution: T, S, D. (Thread-Specific-Data) in pthreads, or Thread-Local-Storage in Java.

→ Idea: we have a single global variable name, but every thread has its own value.

In Java → allocate an object  
Thread Local Storage tls = new (...)

```


T1
tls.set(____)
o = tls.get()

T2
tls.set(____)
q = tls.get()
    
```

} distinct values stored in the same place for each thread,

How to do this

Java {  
- have a hash table mapping thread ids to their corresponding values.  
- (easy conceptually, but hard for efficiency) ⇒ was really slow!  
- there are some clever schemes, but we won't go into that.



pthreads {  
- works differently,  
- key, value pairs  
→ allocate a key (ie global handle for each thread).  
→ pthreads maintains a table  
- allocating a thread adds a row  
- allocating a key adds a column.

	key 1	key 2	key 3
Thread 1			
Thread 2			
⋮			
⋮			

pthreads <

→ pthread\_create\_key (&key, &fn)  
- when a thread terminates → destructor  
- calls destructor function on each key.  
- you receive the actual TSD value → can do anything with it.  
- in the destructor, what if you create a new key?  
- once all the destructors are called for a thread, pthreads will do another pass to call destructor for any new keys to some limit.  
limit = PTHREAD\_DESTRUCTOR\_TERMINATION  
→ can delete keys explicitly that doesn't call destructor  
pthread\_set\_specific (key, value)  
pthread\_get\_specific (key)  
→ nb! for errno as a special case, all the syntax for this is hidden.

MEMORY MODELS & CONSISTENCY

example	T1	T2	x=1	y=2	x=1	... (etc),
	x=1	y=2	b=y	a=x	x=1	
	b=y	a=x	y=2	x=1	b=y	
			a=x	b=y	a=x	
			(1,0)	(0,2)	(1,2)	

- possible final values of (a, b)  
→ interleavings.  
→ notice (0,0) is not possible to achieve.

- Sequential Consistency ⇒ a multithreaded program is S.C. (Sequentially Consistent) if any execution is the same as if the ops of all processes were executed in some sequential order, and ops of each processor respect program order.

nb: as if means it doesn't actually need to execute in that order but it must appear to.

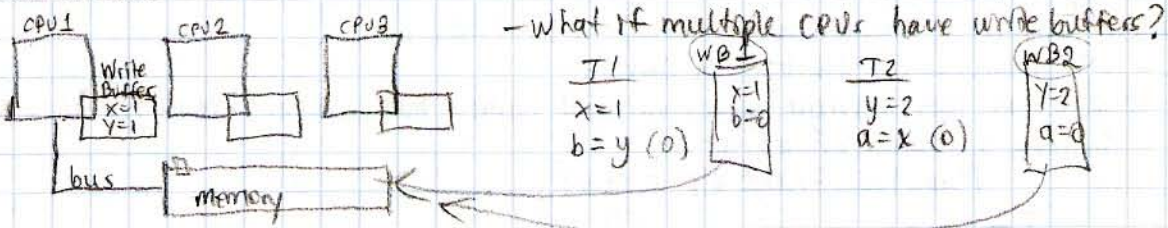
- Strict Consistency ⇒ it MUST actually execute like S.C.



ex: T1                      T2  
 → x=1                      z=3  
 → y=1

- we couldn't really tell if T1's instruction order is changed ⇒ same result no matter what.
- Sequential Consistency gives you the ability to switch things up, as long as the result doesn't change no matter what

ex: Write Buffer ⇒ allows for multiple memory writes at a time as a batch process over a single bus.



- what if multiple CPUs have write buffers?

note: this time both end up as 0 ⇒ (a,b) = (0,0).

- so it's not common to have S.C. because of problems like this.
- No modern hardware guarantees SC for us.
- Intel gets very close to S.C. ⇒ strong model
- PowerPC is more weakly ordered ⇒ less guarantees than Intel.
- DEC Alpha was a very weakly ordered model.

nb: estimate: 80-80% of performance can arise from weak ordering.

→ so we need memory consistency models to tell us what we can actually rely on.

### Memory Models

- spec/model of what is actually allowed, and what is not.

- Coherence (one of the most common).

- weaker than S.C. but quite popular (as weak as you probably want to go for human beings)
- in S.C. we produce one global ordering for all reads & writes.
- in coherence we have an ordering for each variable
- each variable in isolation behaves in a S.C. fashion.

example:

P1  
 x=0  
 ↓  
 x=1  
 y=2

P2  
 =y → can be 2  
 =x → can be 0 ⇒ doesn't assume x must be 1 even though y is already 2.  
 =x → can be 1, but after that it can't be 0.

coherent, but not SC.

- PRAM - Pipelined Ram model

- writes of a single processor will be seen in the same order by every other processor
- still have multiple orderings → one per processor.

P  
 x=0  
 ↓  
 x=1  
 =y (0)  
 =y (1)

Q  
 y=0  
 y=1  
 =x (0)  
 =x (1)

→ as long as x is 0's it can see 0's again, or 1's but once it's a 1, it can't go back to 0. (follows order that x is written in P).

- this is PRAM but not SC.

nb: PRAM & Coherence are NOT comparable. ⇒ one doesn't imply the other.  
 look at above example ⇒ in coherent, but NOT PRAM.

ex:

P  
 x=0  
 =x (1)  
 y=1  
 =y (2)

Q  
 x=1  
 =x (0)  
 y=2  
 =y (1)

} in PRAM, but not coherent.



Release Consistency

→ acquire & release (lock & unlock)

a =  
b =  
acquire → acquire can move up, but not down  
x =  
y =  
z =  
release → release can move down but not up  
c =  
d =

This is called "ROACH MOTEL" semantics  
(can check in but not check out)

⇒ MODERN HARDWARE GIVES YOU NONE OF THESE!

COMP 409  
NOV 2, 2010

Midterm next class. OPEN BOOK, no electronics.

Covers up to but NOT including memory consistency,  
any programming questions (if any) will be pseudocode and short.

Last Time

→ Memory Consistency (Models). simple, basic model (usually assumed): Sequential Consistency.

→ In practice, you never really get Sequential Consistency.

→ There are many → one of the most important: Coherence. (kind of like a lower bound)  
→ another interesting one: PRAM

→ Acquire/Release semantics ⇒ roach motel.

→ Location Consistency

→ Rerun consistency → etc.

→ Hardware has its own models ⇒ whatever is in the CPU.

→ Intel & AMD aren't as bad, you just have to look at the manual for corner cases otherwise is pretty-much close to S.C.

→ all Multiprocessor hardware needs memory models.


→ most languages DON'T. (unless they have explicit concurrency).

→ Java ⇒ interesting example ⇒ has concurrency built in. & runs on a Virtual Machine  
↳ multi-thread.

- Java HAS a memory model

- in fact it's on its second version

OLD MODEL



- read-load-process  
- process-export-write.

- lots of little things to worry about

- old model implied coherence. (serious problem!)

example by (Bill Pugh)

→ int i = p.x;

int j = q.x;

int k = p.x;

↳ compiler optimization problem

⇒ could say int k = i; // FASTER

- what if at runtime another thread writes to p or q and p & q are both pointing to the same thing. Say it writes a 1, and it used to be 0.

- You would see the OLD value in k ⇒ violates coherence! ⇒ The optimization would be wrong.



so it turns out that you can't really optimize much this way!  
 (most compiler optimizations are invalid).

Everyone has always tried to improve Java performance, so this is a problem! (most)

A NEW MEMORY MODEL was introduced  $\Rightarrow$  JSR 133 - memory model allowing for compiler optimizations!

Also thinking of the programmer to make things easier without having to follow so many rules  
 The easiest would be Sequential Consistency, but that's very hard to model.

They used something in between.

$\rightarrow$  Diverse Programs into two classes.

(1) Correctly Synchronized Programs.

$\rightarrow$  CONTAIN NO DATA RACES.

$\rightarrow$  In this case sequential consistency is guaranteed. (good for the programmer).

could stop here (as in C++)  $\rightarrow$  incorrectly synchronized programs behave in an undetermined way.

(2) Incorrectly Synchronized Programs

$\rightarrow$  There are many semantics for this

$\rightarrow$  We want defined semantics in this case (unlike C++).

BASIC MODEL  $\rightarrow$  HB "Happens-Before" Consistency

$\rightarrow$  Build a graph from the actions found in execution traces.

$\rightarrow$  nodes are runtime actions.

$\rightarrow$  edges represent "Happens-Before" relations

HB Edges

$\rightarrow$  between each action of a thread  $t$  to the next action in  $t$ . (intra-thread)

$\rightarrow$  From an unlock of monitor  $M$  to all subsequent locks of  $m$

$\rightarrow$  From each write of a variable  $v$  to all subsequent reads

$\rightarrow$  From the initialization of a thread to its FIRST ACTION.

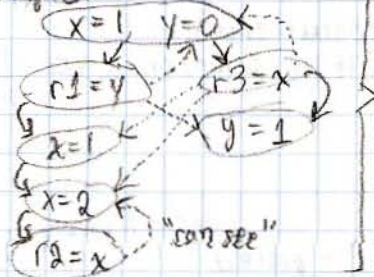
$\rightarrow$  From Thread.start() to the FIRST ACTION.

$\rightarrow$  From the final action in a thread to any subsequent join() or false eval of isAlive().

$\rightarrow$  etc.

So we build this graph

Example:



$\rightarrow$  use the HB edges as a PARTIAL-ORDER.

ie:  $x < y$  iff there is a HB-path from  $x$  to  $y$ .

$\rightarrow$  a read  $r$  of a variable  $x$  is allowed to see a write of  $x$

iff: 1)  $r$  IS NOT ORDERED BEFORE the write

2) THERE IS NO INTERVENING write to  $x$ , such that this intervening write is between the read and the write.

Notice  $r3=x$  can either be  $x=1$  from the top


$x=1$  from the other write in T1

$x=2$  from the 3rd write in T1

- so  $r3$  can be  $=2$  in T1 (due to an optimization) even if it is not yet written in T1.

- so the output can be verified in terms of what we see.

Example:



$\rightarrow$  any value pointed out here is legal according to the memory model. "out-of-thin-air" values.  $\Rightarrow$  considered back! (come from leftover stack/memory data).

(for programmer)

$\rightarrow$  suppose you ran this and got  $x=y=42$ . Would you be able to complain to Sun? No!



- Maybe we should dis-allow these "causal cycles"  
But => are all cycles bad?

```

Example: a=1; b=0;
        i=a      k=b
        j=a      a=k
        if (i==j) {
        } b=2;
  
```

- can it result in i=j=k=2?

- How can we optimise T1's code?

```

// if b=2; is here, it's the same thing.
i=a      k=b
j=i      a=k
b=2; // i=j is always true.
  
```

- so now i=j=k=2