18-447 Lecture 25: Synchronization

James C. Hoe Department of ECE Carnegie Mellon University

18-447-S23-L25-S1, James C. Hoe, CMU/ECE/CALCM, ©2023

Housekeeping

- Your goal today
 - be introduced to synchronization concepts
 - see hardware support for synchronization
- Notices
 - HW5, due Friday 4/28 midnight
 - Lab 4, due this week
 - Final Exam, May 4 Thu, 8:30am-11:30am
- Readings
 - P&H Ch2.11, Ch6
 - Synthesis Lecture: Shared-Memory
 Synchronization, 2013 (advanced optional)

This is not 18-447





[Wikimedia Creative Commons]

18-447-

What is 18-447

- Lab 1~3: knowledge and skill
 - anyone with a wrench can take apart a car
 - Google Lens can tell you what each part is

[Wikimedia Common]

- trained person can put back a working car
- Lab 4: analyze and optimize
- R&D
 - what design decisions make for a car that is fast vs good mileage?
 - how to decide how fast or efficient to make it?



Think, Ask, Invent: what is the "right"

future for personal transport?



18-447-S23-L25-S4, James C. Hoe, CMU/ECE/CALCM, ©2023

Computer Architecture is Engineering

- An applied discipline of finding and optimizing solutions under the joint constraints of demand, technology, economics, and ethics
 Thus, instances of what we recent to the provision of the second seco
- Thus, instances of what we practice evolve continuously
- Need to learn the principles that govern how to develop solutions to meet constraints
- Don't memorize instances; understand why it is that way



What exponential really looks like



18-447-523-L25-S6, James C. Hoe, CMU/ECE/CALCM, ©2023 You just have to stand far back enough to see it

Returning to normally scheduled programming

A simple example: producer-consumer

- Consumer waiting for result from producer in shared-memory variable Data
- Producer uses another shared-memory variable
 Ready to indicate readiness (R=0 initially)

(upper-case for shared-mem **V**ariables)



 Straightforward if SC; if WC, need memory fences to order operations on R and D

18-447-S23-L25-S8, James C. Hoe, CMU/ECE/CALCM, ©2023

Data Races

• E.g., threads T1 and T2 increment a sharedmemory variable V initially 0 (assume SC)

T1:	T2:
t=V	t=V
t=t+1	t=t+1
V=t	V=t

Both threads both read and write V

- What happens depends on what T2 does in between T1's read and write to V (and vice versa)
- Correctness depends on T2 not reading or writing V between T1's read and write ("critical section")

Mutual Exclusion: General Strategy

• **Goal:** allow only either **T1** or **T2** to execute their respective critical sections at one time

No overlapping of critical sections!

- Idea: use a shared-memory variable Lock to indicate whether a thread is already in critical section and the other thread should wait
- Conceptual Primitives:
 - wait-on: to check and block if L is already set
 - acquire: to set L before a thread enters critical sect
 - release: to clear L when a thread leaves critical sect

Mutual Exclusion: 1st Try

• Assume **L=0** initially



But now have same problem with data race on L

Mutual Exclusion: Dekker's

 Using 3 shared-memory variables: Clear1=1, Clear2=1, Turn=1 or 2 initially (assumes SC)

```
C1=0;
                                 C2=0;
while (C2==0)
                                 while(C1==0)
                                                                the tie breaker)
      if (T==2) {
                                       if (T==1) {
            C1=1;
                                             C2=1;
                                             while (T==1);
            while (T == 2);
            C1=0;
                                             C2=0;
      }
                                        }
                                                                hint: T is
{ . . . Critical Section . . . }
                                 { . . . Critical Section . . . }
T=2;
                                 T=1;
C1=1;
                                 C2=1;
```

• Can you decipher this? Extend to 3-way?

Need an easier, more general solution

Aside: what happens in Dekker's w/o T

 Using shared-memory variables: Clear1=1, Clear2=1 initially (assumes SC)

```
C1=0;
while(C2==0) {
    C1=1;
    some delay;
    C1=0;
}
{...Critical Section...}
C1=1;
C2=0;
while(C1==0) {
    C2=1;
    some delay;
    C2=0;
}
C2=0;
}
...Critical Section...}
```

- Above is safe—if one side in C.S., the other isn't
- Either or both loop forever if pathological timing *Livelock possible*

Aside: Dumb it down more

 Using shared-memory variables: Clear1=1, Clear2=1 initially (assumes SC)

```
C1=0;
while(C2==0) {
    some delay;
}
{...Critical Section...}
C2=0;
while(C1==0) {
    some delay;
}
{...Critical Section...}
```

- Above is still safe—if one side in C.S., the other isn't
- Both loop forever if tried at same time

Deadlock possible

Atomic Read-Modify-Write Instruction

- Special class of memory instructions to facilitate implementations of lock synchronizations
- Effects executed "atomically" (i.e. not interleaved by other reads and writes)
 - reads a memory location
 - performs some simple calculation
 - writes something back to the same location

HW guarantees no intervening read/write by others

E.g., <swap>(addr,reg):
 temp MEM[addr];
 MEM[addr] reg;
 reg temp;

<test&set>(addr,reg): reg←MEM[addr]; if (reg==0) MEM[addr]←1;

Expensive to implement and to execute

18-447-S23-L25-S15, James C. Hoe, CMU/ECE/CALCM, ©2023

Acquire and Release

- Could rewrite earlier examples directly using
 <swap> or <test&set> instead loads and stores
- Better to hide ISA-dependence behind portable
 Acquire () and Release () routines



Note: implicit in Acquire (L) is to wait on L if not free

Acquire and Release

• Using <swap>, L initially 0

```
void Acquire(L) {
   do {
     reg=1;
     <swap>(L,reg);
   } while(reg!=0);
}
```

```
void Release(L) {
  L=0;
}
```

Using <test&set>, L initially 0

```
void Acquire(L) {
    do {
        <test&set>(L, reg);
    }
    while(reg!=0);
}
void Release(L) {
    L=0;
}
```

Many equally powerful variations of atomic

18-447-S23-L25-S17, James C. Hoe, CMU/ECE/CALCM, ©2023

RMW insts can accomplish the same

High Cost of Atomic RMW Instructions

- Literal enforcement of atomicity very early on
- In CC shared-memory multiproc/multicores
 - RMW requires a writeable M/E cache copy
 - lock cacheblock from replacement during RMW
 - expensive when lock contended by many concurrent acquires—a lot of cache misses and cacheblock transfers, just to swap "1" with "1"
- Optimization
 - check lock value using normal load on read-only S copy
 - attempt RMW only when success is possible

```
18-447-S23-L25-S18, James C. Hoe, CMU/ECE/CALCM, ©2023
```

```
do {
    reg=1;
    if (!L) {
        <swap>(L,reg);
    }
} while (reg!=0);
```

RMW without Atomic Instructions

Add per-thread architectural state: *reserved*,
 address and *status* (addr, reg):

<ld-linked>(reg,addr):
 reg MEM[addr];
 reserved 1;
 address addr;

<st-cond>(addr,reg):
 if (reserved &&
 address==addr)
 M[addr] \leftarrow reg;
 status \leftarrow 1;
 else
 status \leftarrow 0;

- <ld-linked> requests S-copy (if not alrdy S or M)
- HW clears *reserved* if cached copy lost due to CC (i.e., store or <st-cond> at another thread)
- If *reserved* stays valid until <st-cond>, request Mcopy (if not already M) and update; can be no other intervening stores to *address* in between!!

Acquire () by Id-linked and st-cond



Resolving Data Race without Lock

• E.g., two threads T1 and T2 increment a sharedmemory variable V initially 0 (assume SC)

- Atomicity not guaranteed, but . . .
- You know if you succeeded; no effect if you don't Just try and try again until you succeed

Barrier Synchronization



18-447-S23-L25-S22, James C. Hoe, CMU/ECE/CALCM, ©2023

(Blocking) Barriers

- Ensure a group of threads have all reached an agreed upon point
 - threads that arrive early have to wait
 - all are released when the last thread enters
- Can build from shared memory on small systems
 e.g., for a simple 1-time-use barrier (B=0 initially)

```
Acquire (L<sub>B</sub>)
B=B+1;
Release (L<sub>B</sub>)
while (B!=NUM_THREADS); wait
```

• Barrier on large systems are expensive, often supported/assisted by dedicated HW

18-447-S23-L25-S23, James C. Hoe, CMU/ECE/CALCM, ©2023

Nonblocking Barriers

- Separate primitives for enter and exit
 - enterBar() is non-blocking and only records that a thread has reached the barrier

```
Acquire(L<sub>B</sub>)
B=B+1;
Release(L<sub>B</sub>)
```

exitBar() blocks until the barrier is complete

while (B!=NUM_THREADS);

- A thread
 - calls enterBar() then go on to independent work
 - calls exitBar() only when no more work that doesn't depend on the barrier

Pass this point not on exams

For more, go read "Synthesis Lecture: Transactional Memory," 2nd Ed., 2010

18-447-S23-L25-S25, James C. Hoe, CMU/ECE/CALCM, ©2023

Transactional Memory

T1 :		T2 :	
	<pre>TxnBegin();</pre>		<pre>TxnBegin();</pre>
	t=V		t=V
	$t=func_1(t, V,)$		$t=func_2(t, V,)$
	V=t		V=t
	TxnEnd();		TxnEnd();

- Acquire (L) /Release (L) say do one at a time
- TxnBegin () /TxnEnd () say "look like" done one at a time

Implementation can allow transactions to overlap and only fixes things if violations observable

Optimistic Execution Strategy

- Allow multiple transaction executions to overlap
- Detect atomicity violations between transactions
- On violation, one of the conflicting transactions is aborted (i.e., restarted from the beginning)
 - TM writes are speculative until reaching TxnEnd
 - speculative TM writes not observable by others
- Effective when actual violation is unlikely, e.g.,
 - multiple threads sharing a large structure/array
 - cannot decide statically which part of structure/array touched by different threads
 - conservative locking adds a cost to every access
 - TM incurs a cost only when data races occur

Detecting Atomicity Violation

- A transaction tracks memory **RdSet** and **WrSet**
- Txn_a appears atomic with respect to Txn_b if
 - $WrSet(Txn_a) \cap (WrSet(Txn_b) \cup RdSet(Txn_b)) = \emptyset$
 - $RdSet(Txn_a) \cap WrSet(Txn_b) = \emptyset$
- Lazy Detection
 - broadcast RdSet and WrSet to other txns at TxnEnd
 - waste time on txns that failed early on
- Eager Detection
 - check violations on-the-fly by monitoring other txns' reads and writes
 - require frequent communications

Oversimplified HW-based TM using CC

- Add RdSet and WrSet status bits to identify cacheblocks accessed since TxnBegin
- Speculative TM writes
 - issue BusRdOwn/Invalidate if starting in I or S
 - issue BusWr(old value) on first write to M block
 - on abort, silently invalidate WrSet cacheblocks
 - on reaching TxnEnd, clear RdSet/WrSet bits
 - Assume **RdSet/WrSet** cacheblocks are never displaced
- Eager Detection
 - snoop for BusRd, BusRdOwn, and Invalidation
 - M→S, M→I or S→I downgrades to RdSet/WrSet indicative of atomicity violation

Which transaction to abort?

Hoe, CMU/ECE/CALCM, ©2023

Why not transaction'ize everything?

```
void *sumParallel
      (void * id) {
  long id=(long) id;
  long i;
  long N=ARRAY SIZE/p;
  TxnBegin();
  for(i=0;i<N;i++) {</pre>
    double v=A[id*N+i];
    if (v>=0)
      SumPos+=v;
    else
      SumNeg+=v;
  TxnEnd();
}
                           }
```

```
void *sumParallel
    (void *_id) {
    long id=(long) _id;
    long i;
    long N=ARRAY_SIZE/p;
```

```
for(i=0;i<N;i++) {
    TxnBegin();
    double v=A[id*N+i];
    if (v>=0)
        SumPos+=v;
    else
        SumNeg+=v;
    TxnEnd();
}
```

Compute separate sums of positive and negative elements of **A** in **SumPos** and **SumNeg**

p=2

```
void *sumParallel
    (void *_id) {
    long id=(long) _id;
    long i;
    long N=ARRAY SIZE/p;
```

```
for (i=0;i<N;i++) {
   double v=A[id*N+i];
   if (v>=0) {
     TxnBegin();
     SumPos+=v;
     TxnEnd();
   } else {
     TxnBegin();
     SumNeg+=v;
     TxnEnd();
   }
}
Better??
```

Overhead vs Likelihood of Succeeding

