Hello!

CONCURRENCY: LOCKS

Shivaram Venkataraman

CS 537, Spring 2020

ADMINISTRIVIA

- Project 2b is done!
- Project 3 is out! → last project before Spring break
- Project 1b grades out today?
- Project 2a grades this week

Discussion section Today 5:30 pm 105 Psych

No class on Thursday

AGENDA / LEARNING OUTCOMES

Concurrency

- What are some of the challenges in concurrent execution?
- How do we design locks to address this?

RECAP

MOTIVATION FOR CONCURRENCY

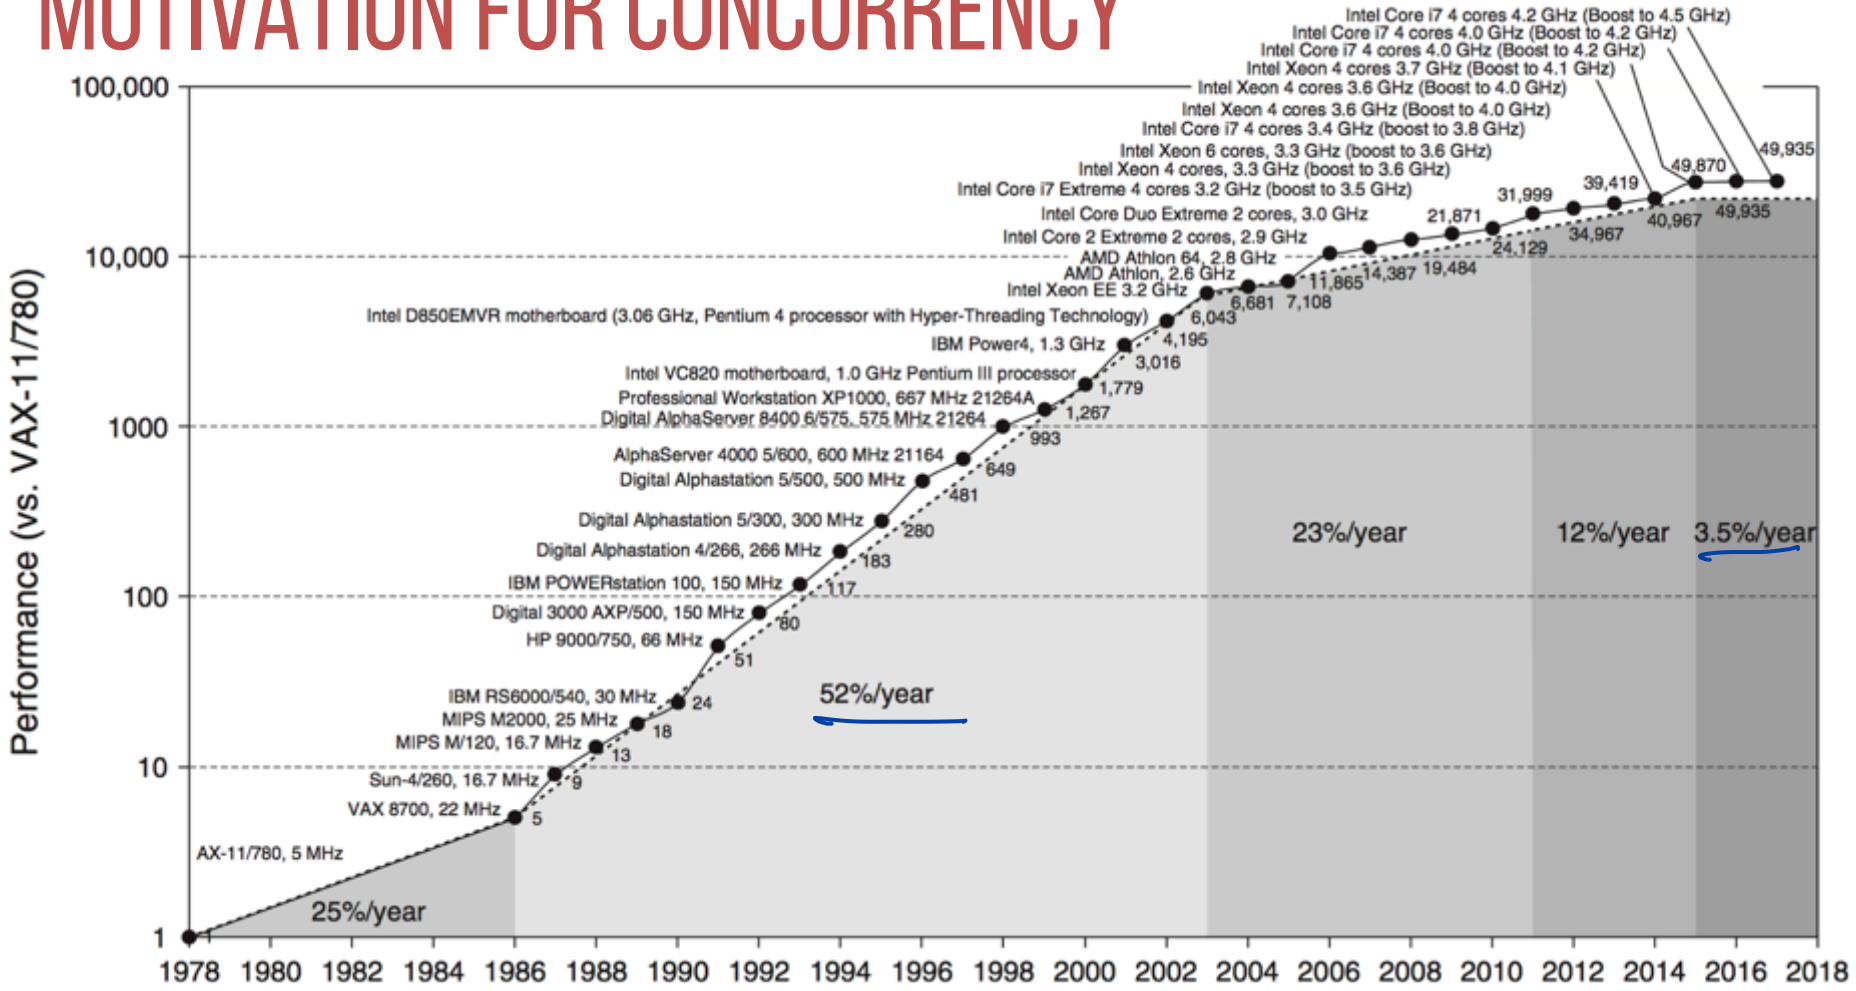

TIMELINE VIEW

Thread 1	**Thread 2**
mov 0x123, %eax	
	mov 0x123, %eax
	add %0x2, %eax
	mov %eax, 0x123
add %0x1, %eax	
mov %eax, 0x123	

WHAT DO WE WANT?

Want 3 instructions to execute as an uninterruptable group

That is, we want them to be atomic

```
mov 0x123, %eax
add %0x1, %eax
mov %eax, 0x123
```

→ reading
→ adding
→ storing

Atomically

More general: Need mutual exclusion for critical sections

if thread A is in critical section C, thread B isn't

(okay if other threads do unrelated work)

LOCK IMPLEMENTATION GOALS

Correctness

- *Mutual exclusion*
  Only one thread in critical section at a time
- *Progress* (deadlock-free)
  If several simultaneous requests, must allow one to proceed
- *Bounded* (starvation-free)
  Must eventually allow each waiting thread to enter

Fairness: Each thread waits for same amount of time

Performance: CPU is not used unnecessarily

```
lock() {
  return → grab
}

unlock () {
  // Thread is done
}
```

RACE CONDITION WITH LOAD AND STORE

shared global

```
*lock == 0 initially

Thread 1                          Thread 2
while(*lock == 1)
                                  while(*lock == 1)
                                  *lock = 1
*lock = 1
```

Both threads grab lock! → violates mutual exclusion

Problem: Testing lock and setting lock are not atomic

XCHG: ATOMIC EXCHANGE OR TEST-AND-SET

How do we solve this ? Get help from the hardware!

1. Stores newval into addr
2. returns what used to be in addr

```
// xchg(int *addr, int newval)
// return what was pointed to by addr
// at the same time, store newval into addr
int xchg(int *addr, int newval) {
  int old = *addr;
  *addr = newval;
  return old;
}
```

Atomically

```
movl 4(%esp), %edx
movl 8(%esp), %eax
xchgl (%edx), %eax
ret
```

addr
newval

return value

LOCK IMPLEMENTATION WITH XCHG

0 - unlocked
1 - locked

T1 T2

flag = 1 xchg(&flag, 1) → 1

flag = 0 xchg(&flag, 1) → 0

```
typedef struct __lock_t {
    int flag;
} lock_t;

void init(lock_t *lock) {
    lock->flag = 0;
}

void acquire(lock_t *lock) {
    ????;
    // spin-wait (do nothing)
}

void release(lock_t *lock) {
    lock->flag = 0;
}
```

```
int xchg(int *addr, int newval)
```

Wait until lock is held
else grab the lock

```
while ( xchg(&lock->flag, 1) == 1) {
    //spin
}
```

Somebody else has lock

OTHER ATOMIC HW INSTRUCTIONS

CAS

```
int CompareAndSwap(int *addr, int expected, int new) {
  int actual = *addr;
  if (actual == expected)
    *addr = new;
  return actual;
}
```

old value

if compare:
swap
return old value

```
void acquire(lock_t *lock) {
    while(CompareAndSwap(&lock->flag, 0 , 1) == 1) ;
    // spin-wait (do nothing)
}
```

we want to swap if unlocked

we want to lock if unlocked

if other thread has lock spin.

0 → unlock
1 → lock

BASIC SPINLOCKS ARE UNFAIR

Thread B keeps waiting for the lock

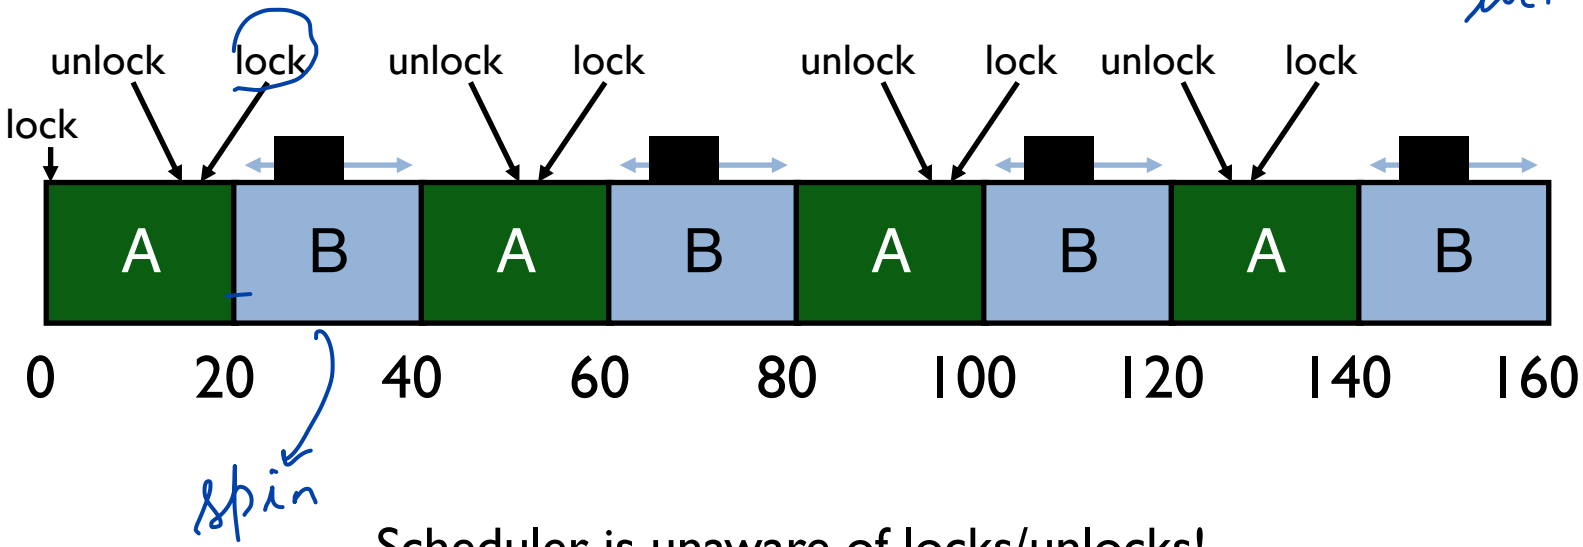

Scheduler is unaware of locks/unlocks!

FAIRNESS: TICKET LOCKS

Idea: reserve each thread's turn to use a lock.

Each thread spins until their turn.

Use new atomic primitive, fetch-and-add

example C code

but is atomic instr.

```
int FetchAndAdd(int *ptr) {
   int old = *ptr;
   *ptr = old + 1;
   return old;
}
```

Acquire: Grab ticket; Spin while not thread's ticket != turn

Release: Advance to next turn

TICKET LOCK EXAMPLE

RR sched

A lock(): ticket = 0 → grabs the lock
B lock(): ticket = 1 → spin
C lock(): ticket = 2 → spin

A unlock(): inc turn value

A lock(): ticket = 3. turn = 1

B unlock(): B gets lock, turn = 2

C unlock(): C gets lock turn = 3

A unlock(): A gets lock turn = 4

Avoided unfair

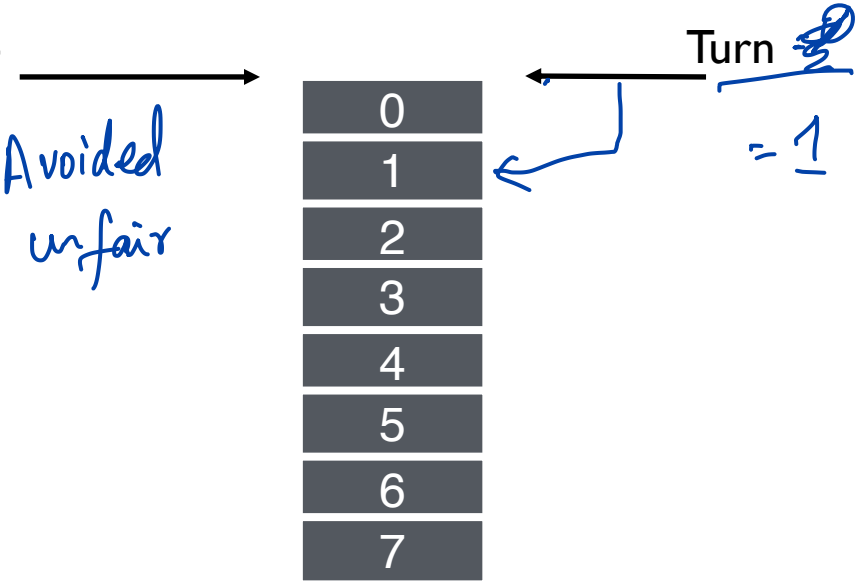

TICKET LOCK IMPLEMENTATION

```
typedef struct __lock_t {
    int ticket;
    int turn;
}

void lock_init(lock_t *lock) {
    lock->ticket = 0;
    lock->turn = 0;
}
```

Grab a ticket

increment to 1

```
void acquire(lock_t *lock) {
    int myturn = FAA(&lock->ticket);
    // spin
    while (lock->turn != myturn);
}

void release(lock_t *lock) {
    FAA(&lock->turn);
}
```

0

indicate that next ticket holder can get the lock.

SPINLOCK PERFORMANCE

Fast when…

- many CPUs
- locks held a short time
- advantage: avoid context switch

Slow when…

- one CPU
- locks held a long time
- disadvantage: spinning is wasteful

CPU SCHEDULER IS IGNORANT

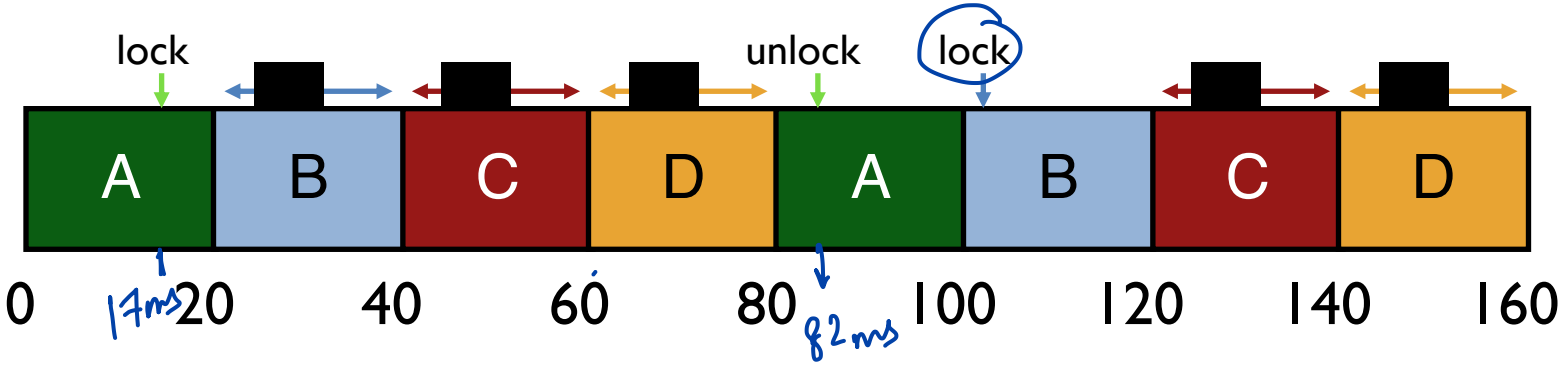

CPU scheduler may run **B, C, D** instead of **A**
even though **B, C, D** are waiting for **A**

TICKET LOCK WITH YIELD

```
typedef struct __lock_t {
    int ticket;
    int turn;
}

void lock_init(lock_t *lock) {
    lock->ticket = 0;
    lock->turn = 0;
}
```

```
void acquire(lock_t *lock) {
    int myturn = FAA(&lock->ticket);
    while (lock->turn != myturn)
        yield();
}

void release(lock_t *lock) {
    FAA(&lock->turn);
}
```

YIELD INSTEAD OF SPIN

A gets a second turn at 35ms with yield vs.

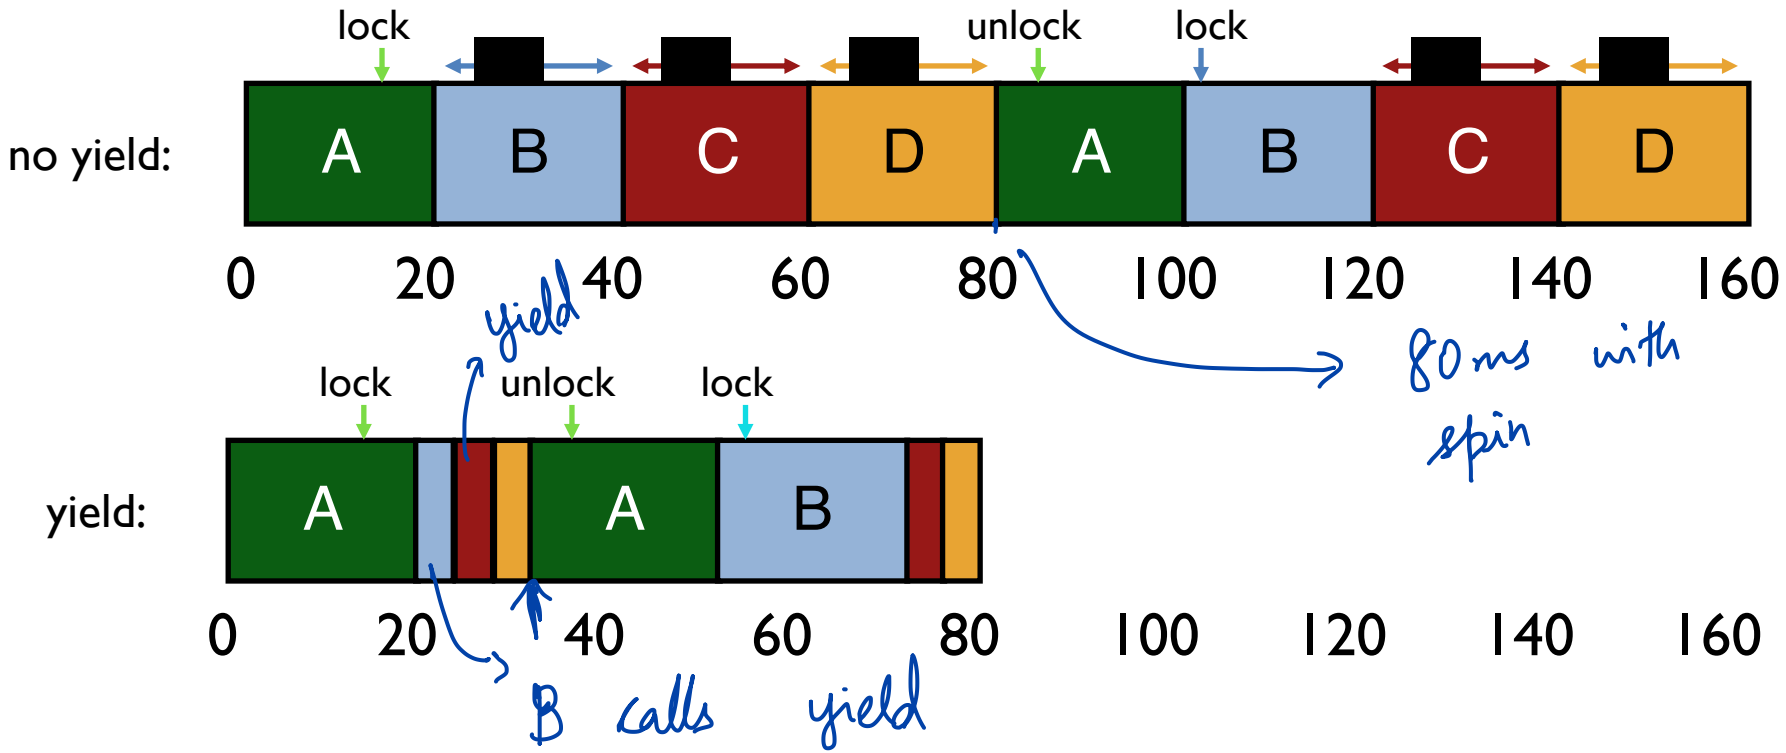

QUIZ 16

https://tinyurl.com/cs537-sp20-quiz16

```
a = 1
int b = xchg(&a, 2)
int c = CAS(&b, 2, 3)
int d = CAS(&b, 1, 3)
```

1

1

Compare success!

Final values

a = 2

b = 3

c = 1

d = 1

Assuming round-robin scheduling, 10ms time slice. ~~Processes~~ Threads A, B, C, D, E, F, G, H in the system

Timeline

A: lock() … compute … unlock() (20ms)

B: lock() … compute … unlock()

C: lock()

(a)

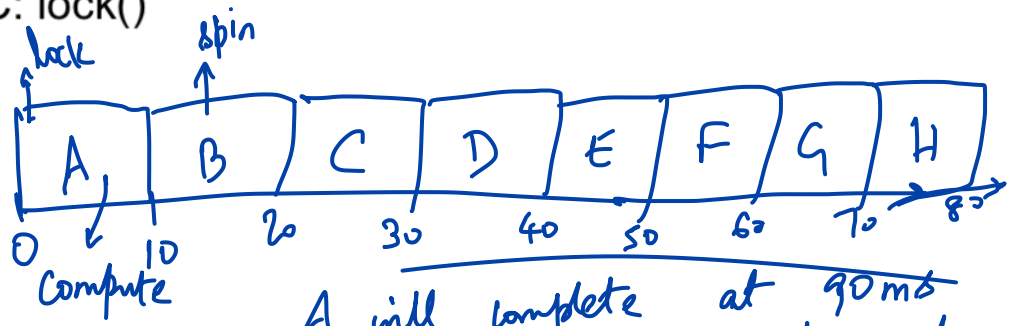

A will complete at 90ms

B will get the lock at 90 ms → SPIN

B has 30 ms of compute, when does C get the lock?

C D E F G H

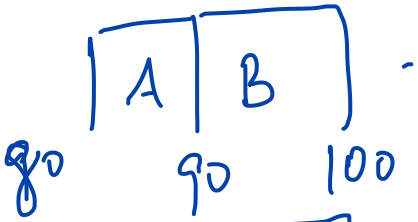

60 ms

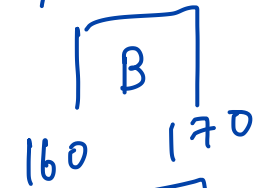

60 ms

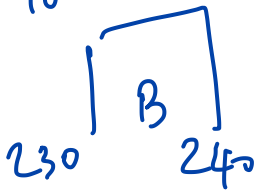

C gets the lock

Context switch = 1ms

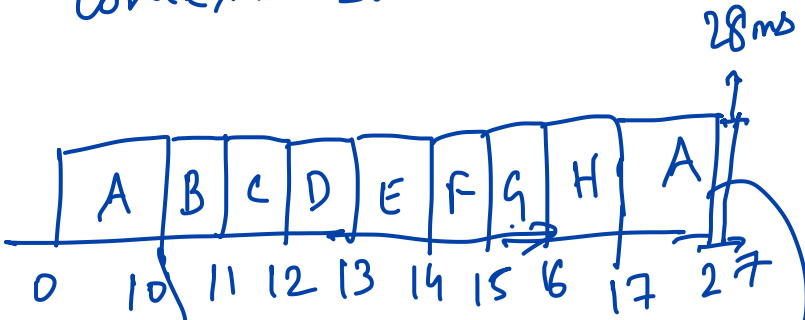

Context switch

B can grab the lock

29??

if we count context switches

SPINLOCK PERFORMANCE

Waste of CPU cycles?

Without yield: O(threads * **time_slice**)

With yield: O(threads * **context_switch**)

Even with yield, spinning is slow with high thread contention

Next improvement: Block and put thread on waiting queue instead of spinning

LOCK IMPLEMENTATION: BLOCK WHEN WAITING

Remove waiting threads from scheduler runnable queue
(e.g., park() and unpark(threadID)) → Solaris SUN Microsystem

Scheduler runs any thread that is **runnable**

A B D contend for lock, C is not contending

A has 60 ms worth of work

20ms is the timeslice

RUNNABLE: A, ~~B~~, C, ~~D~~

RUNNING:

WAITING: B, D

B or D can be moved into Runnable

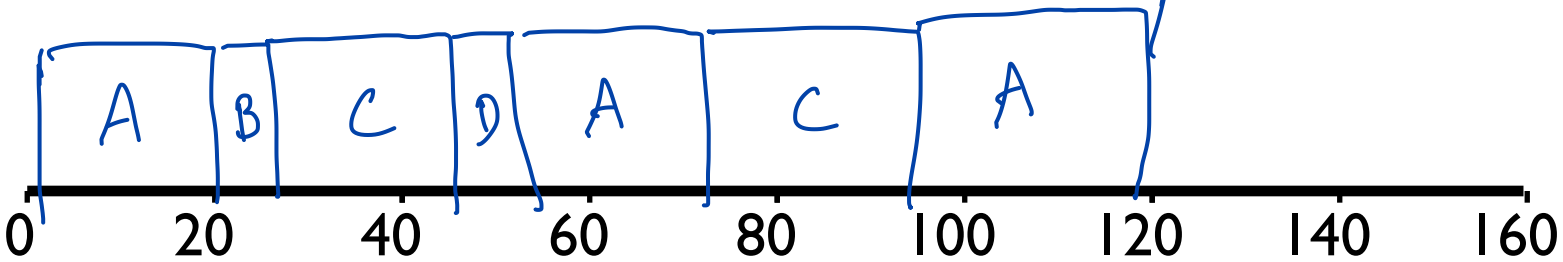

B, D were not considered

LOCK IMPLEMENTATION: BLOCK WHEN WAITING

```
typedef struct {
    bool lock = false;
    bool guard = false;
    queue_t q;
} LockT;
```

queue of threads

spin lock to guard the queue

T2

```
void acquire(LockT *l) {
    while (XCHG(&l->guard, true));
    if (l->lock) {
        qadd(l->q, tid);
        l->guard = false;
        park();        // blocked
    } else {
        l->lock = true;
        l->guard = false;
    }
}
```

Add thread to queue

Wakes up

block this thread, not runnable

T1

```
void release(LockT *l) {
    while (XCHG(&l->guard, true));
    if (qempty(l->q)) l->lock=false;
    else unpark(qremove(l->q));
    l->guard = false;
}
```

first TID in Queue

LOCK IMPLEMENTATION: BLOCK WHEN WAITING

(a) Why is **guard** used?

We prevent races in updating the queue

(b) Why okay to **spin** on guard?

Amt of time spin is small

(c) In release(), why not set lock=false when unpark?

l->lock = true is handed over from release to acquire

(d) Is there a race condition?

```
void acquire(LockT *l) {
    while (XCHG(&l->guard, true));
    if (l->lock) {
        qadd(l->q, tid);
        l->guard = false;
        park();      // blocked
    } else {
        l->lock = true;
        l->guard = false;
    }
}
```

l->lock = true

```
void release(LockT *l) {
    while (XCHG(&l->guard, true));
    if (qempty(l->q)) l->lock=false;
    else {unpark(qremove(l->q));}
    l->guard = false;
}
```

RACE CONDITION

Thread 1 (in lock)

```
if (l->lock) {
    qadd(l->q, tid);
    l->guard = false;

    park();     // block
```

Thread 2 (in unlock)

```
while (TAS(&l->guard, true));
if (qempty(l->q)) // false!!
else unpark(qremove(l->q));
l->guard = false;
```

unblock a thread which is not blocked

BLOCK WHEN WAITING: FINAL CORRECT LOCK

```
typedef struct {
    bool lock = false;
    bool guard = false;
    queue_t q;
} LockT;
```

```
void acquire(LockT *l) {
    while (TAS(&l->guard, true));
    if (l->lock) {
        qadd(l->q, tid);
        setpark(); // notify of plan
        l->guard = false;
        park(); // unless unpark()
    } else {
        l->lock = true;
        l->guard = false;
    }
}
void release(LockT *l) {
    while (TAS(&l->guard, true));
    if (qempty(l->q)) l->lock=false;
    else unpark(qremove(l->q));
    l->guard = false;
}
```

setpark() fixes race condition

park that operation will be called next

SPIN-WAITING VS BLOCKING

Each approach is better under different circumstances

Uniprocessor

- Waiting process is scheduled → Process holding lock isn't
- Waiting process should always relinquish processor
- Associate queue of waiters with each lock (as in previous implementation)

Multiprocessor

- Waiting process is scheduled → Process holding lock might be
- Spin or block depends on how long, t, before lock is released
  - Lock released quickly → Spin-wait
  - Lock released slowly → Block
  - Quick and slow are relative to context-switch cost, C

WHEN TO SPIN-WAIT? WHEN TO BLOCK?

If know how long, **t,** before lock released, can determine optimal behavior

How much CPU time is wasted when spin-waiting?

t

How much wasted when blocking?

What is the best action when t<C?

When t>C?

Problem:

Requires knowledge of future; too much overhead to do any special prediction

TWO-PHASE WAITING

Theory: Bound worst-case performance; ratio of actual/optimal

When does worst-possible performance occur?

Spin for very long time t >> C
Ratio: t/C (unbounded)

Algorithm: Spin-wait for C then block → Factor of 2 of optimal

Two cases:

- t < C: optimal spin-waits for t; we spin-wait t too
- t > C: optimal blocks immediately (cost of C);
  we pay spin C then block (cost of 2 C);
  2C / C → 2-competitive algorithm

NEXT STEPS

Project 3: out now!

Discussion: Today at 5.30pm in 105 Psychology

No class on Thursday

Next Tuesday: Condition Variables