

CONCURRENCY: DATA STRUCTURES

Shivaram Venkataraman

CS 537, Fall 2024

ADMINISTRIVIA

Midterm I - !?!

Project 4 out soon?

Office hours: Friday 3-4pm (this week only)

AGENDA / LEARNING OUTCOMES

Concurrency

How do we design locks?

How to build concurrent data structures?

RECAP

LOCK IMPLEMENTATION WITH XCHG

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

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

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

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

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

TICKET LOCK IMPLEMENTATION

```
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);
    // spin
    while (lock->turn != myturn);
}

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

LOCK IMPLEMENTATION: BLOCK WHEN WAITING

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

spin-lock around flag and queue ops

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

time spent spinning is limited

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

thread woken up, as if it is returns from `park`

(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;
    }
}

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;

BLOCK WHEN WAITING: FINAL CORRECT LOCK

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

setpark() fixes race condition

```
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;
}
```

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 ($t \ll C$)

Lock released slowly → Block ($t \geq C$)

Quick and slow are relative to context-switch cost, C

QUIZ 10

<https://tinyurl.com/cs537-fa24-q10>



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

A =

B =

C =

D =

```
int xchg(int *addr, int newval) {
    int old = *addr;
    *addr = newval;
    return old;
}

int CAS(int *addr, int ex, int n) {
    int actual = *addr;
    if (actual == ex)
        *addr = n;
    return actual;
}
```

Assuming round-robin scheduling, 10ms time slice with processes A, B, C, D, E, F, G, H in a single CPU system, with processes C - H long-running jobs.

Timeline

A: lock() ... compute ... unlock()

B: lock() ... compute ... unlock()

C: lock()

A's compute is 20ms, starting at $t = 0$, when does B acquire the lock with spin waiting locks?

B's compute is 30ms long, when does C acquire the lock with spin waiting locks? (in ms)

CONCURRENT DATA STRUCTURES

CONCURRENT DATA STRUCTURES

Counters

Lists

Hashtable

Queues

Start with a correct solution

Make it perform better!

WHAT IS SCALABILITY

N times as much work on N cores as done on 1 core

Strong scaling

Fix input size, increase number of cores

Weak scaling

Increase input size with number of cores

COUNTERS

```
1 typedef struct __counter_t {
2     int value;
3 } counter_t;
4
5 void init(counter_t *c) {
6     c->value = 0;
7 }
8 void increment(counter_t *c) {
9     c->value++;
10 }
11 int get(counter_t *c) {
12     return c->value;
13 }
```


THREAD SAFE COUNTER

```
1 typedef struct __counter_t {
2     int value;
3     pthread_mutex_t lock;
4 } counter_t;
5
6 ...
7
8 10
9
11 void increment(counter_t *c) {
12     Pthread_mutex_lock(&c->lock);
13     c->value++;
14     Pthread_mutex_unlock(&c->lock);
15 }
```

COUNTER SCALABILITY DEMO

UNDERLYING PROBLEM?


```
void spin_lock(spinlock_t *lock)
{
    t = atomic_inc(lock->next_ticket);
    while (t != lock->current_ticket)
        ; /* Spin */
}
```

```
void spin_unlock(spinlock_t *lock)
{
    lock->current_ticket++;
}
```

```
struct spinlock_t {
    int current_ticket;
    int next_ticket;
}
```

An Analysis of Linux Scalability to Many Cores

Boyd-Wickizer et. al
OSDI 2010



APPROXIMATE COUNTERS

Maintain a counter per-core, global counter

Global counter lock

Per-core locks if more than 1 thread per-core?

Increment:

- update local counters

- at threshold update global

Read:

- global counter (maybe inaccurate?)

DEMO

CONCURRENT LINKED LIST

```
18 void List_Insert(list_t *L, int key) {
19     pthread_mutex_lock(&L->lock);
20     node_t *new = malloc(sizeof(node_t));
21     if (new == NULL) {
22         perror("malloc");
23         pthread_mutex_unlock(&L->lock);
24         return; // fail
25     }
26     new->key = key;
27     new->next = L->head;
28     L->head = new;
29     pthread_mutex_unlock(&L->lock);
30     return; // success
31 }
```

BETTER CONCURRENT LINKED LIST?

```
18 void List_Insert(list_t *L, int key) {
19     node_t *new = malloc(sizeof(node_t));
21     if (new == NULL) {
22         perror("malloc");
24         return; // fail
25     }

26     pthread_mutex_lock(&L->lock);
27     new->key = key;
28     new->next = L->head;
29     L->head = new;
30     pthread_mutex_unlock(&L->lock);
31     return; // success
32 }
```

DEMO

HASH TABLE FROM LIST

```
1 #define BUCKETS (101)
2 typedef struct __hash_t {
3     list_t lists[BUCKETS];
4 } hash_t;
5
6 int Hash_Insert(hash_t *H, int key) {
7     int bucket = key % BUCKETS;
8     return List_Insert(&H->lists[bucket], key);
9 }
10
```

DEMO

```
21 void Queue_Enqueue(queue_t *q, int value) {
22     node_t *tmp = malloc(sizeof(node_t));
23     assert(tmp != NULL);
24     tmp->value = value;
25     tmp->next = NULL;
26
27     pthread_mutex_lock(&q->tailLock);
28     q->tail->next = tmp;
29     q->tail = tmp;
30     pthread_mutex_unlock(&q->tailLock);
31 }
32
33 int Queue_Dequeue(queue_t *q, int *value) {
34     pthread_mutex_lock(&q->headLock);
35     node_t *tmp = q->head;
36     node_t *newHead = tmp->next;
37     if (newHead == NULL) {
38         pthread_mutex_unlock(&q->headLock);
39         return -1; // queue was empty
40     }
41     *value = newHead->value;
42     q->head = newHead;
43     pthread_mutex_unlock(&q->headLock);
44     free(tmp);
45     return 0;
46 }
```

CONCURRENT DATA STRUCTURES

Simple approach: Add a lock to each method?!

Check for scalability – weak scaling, strong scaling

Avoid cross-thread, cross-core traffic

- Per-core counter

- Buckets in hashtable

- Keep critical sections small!

NEXT STEPS

Condition Variables