

# CONCURRENCY: DATA STRUCTURES

Shivaram Venkataraman

CS 537, Spring 2019

# ADMINISTRIVIA

Spring break!

# AGENDA / LEARNING OUTCOMES

Concurrency: How to build concurrent data structures?

Summary of virtualization, concurrency

**RECAP**

# CONCURRENCY OBJECTIVES

**Mutual exclusion** (e.g., A and B don't run at same time)

solved with *locks*

**Ordering** (e.g., B runs after A does something)

solved with *condition variables* and *semaphores*

# ABSTRACTIONS

Objects, Lists, Hashtable

---

Semaphores

Locks, Condition variables

Atomic Primitives

# 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");
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 }
```

**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

# OPERATING SYSTEMS: THREE EASY PIECES

Three conceptual pieces

1. Virtualization

2. Concurrency

3. Persistence

# VIRTUALIZATION

Make each application believe it has each **resource to itself**  
**CPU and Memory**

Abstraction: Process API, Address spaces

Mechanism:

- Limited direct execution, CPU scheduling

- Address translation (segmentation, paging, TLB)

Policy: MLFQ, LRU etc.



# CONCURRENCY

Events occur simultaneously and may interact with one another

Need to

- Hide concurrency from independent processes

- Manage concurrency with interacting processes

Provide abstractions (locks, semaphores, condition variables etc.)

Correctness: mutual exclusion, ordering

Performance: scaling data structures, fairness

Common Bugs!

# NEXT STEPS

Spring break!