

## Lecture 9: Avoiding Locks

CSC 469H1F / CSC 2208H1F  
Fall 2007  
Angela Demke Brown

(with thanks to Paul McKenney)



## Locking: A necessary evil?

- Locks are an easy to understand solution to critical section problem
  - Protect shared data from corruption due to simultaneous updates
  - Protect against inconsistent views of intermediate states
- But locks have lots of problems
  - Deadlock
  - Priority inversion
  - Not fault tolerant
  - Convoying
  - Expensive, even when uncontended
- *Not* easy to use correctly!


CSC469



## Deadlock

- Textbook definition: Set of threads blocked waiting for event that can only be caused by another thread in the same set

- Classic example:



- Self-deadlock also a big issue
  - Thread holds lock on shared data structure and is interrupted
  - Interrupt handler needs same lock!
    - Solutions exist (e.g., disable interrupts while holding lock), but add complexity

CSC469



## Priority Inversion

- Lower priority thread gets spinlock
  - Higher priority thread becomes runnable and preempts it
    - needs lock, starts spinning
    - Lock holder can't run and release lock
      - May get to run on another CPU
- 

CSC469



## Not fault tolerant

- Lock holder crashes, or suffers indefinite delay, no one makes progress




CSC469



## Convoying

- Suppose we have set of threads, similar work per thread, but started at different times, occasionally accessing shared data
  - E.g. multi-threaded web server
- Expect access to shared objects (and hence times when locks are needed) to be spread out over time
  - Delay of lock holder allows other threads to catch up
  - Lock becomes contended and tends to stay that way



CSC469



## Expensive, even when uncontended


Operation	Nanoseconds
Instruction	0.24
Clock Cycle	0.69
Atomic Increment	42.09
Cmpxchg Blind Cache Transfer	56.80
Cmpxchg Cache Transfer and Invalidate	59.10
SMP Memory Barrier (eieio)	75.53
Full Memory Barrier (sync)	92.16
CPU-Local Lock	243.10

McKenney, 2005 - 8-CPU 1.45 GHz PPC

CSC469



## Causes: Deeper Memory Hierarchy



- Memory speeds have not kept up with CPU speeds
  - 1984: no caches needed, since instructions slower than memory accesses
  - 2005: 3-4 level cache hierarchies, since instructions orders of magnitude faster than memory accesses
- Synch. ops typically execute at memory speed

CSC469





## NBS Basics

- Make change optimistically, roll back and retry if conflict detected

```
atomic_inc(int *counter) {
    int value;
    do {
        value = *counter;
    } while (!CAS(counter, value, value+1));
}
```

- Complex updates (e.g. modifying multiple values in a structure) are hidden behind a single commit point using atomic instructions

CSC469



## Example: Stack Data Structure

- Lock-based synchronization:

```
typedef struct node_s {
    int val;
    struct node_s *next;
} node_t;

typedef struct stack_s {
    node_t *top;
    lock_t *stack_lock;
} stack_t;

void push(stack_t *s,
          node_t *n) {
    lock(s->stack_lock);
    n->next = s->top; s->top=n;
    unlock(s->stack_lock);
}

node_t* pop(stack_t *s){
    node_t *rnode = NULL;
    lock(s->stack_lock);
    if (s->top != NULL) {
        rnode = s->top;
        s->top = s->top->next;
    }
    unlock(s->stack_lock);
    return rnode;
}
```

CSC469



## Non-blocking stack (take 1)

```
typedef struct node_s {
    int val;
    struct node_s *next;
} node_t;

typedef node_t *stack_t;

void push(stack_t *s, node_t
*n) {
    node_t *first;
    do {
        first = *s;
        n->next = first;
    } while (!CAS(s,first,n));
}
```

```
node_t* pop(stack_t *s) {
    node_t *first, *second;
    do {
        first = *s;
        if (first != NULL) {
            second = first->next;
        } else return NULL;
    } while
    (!CAS(s,first,second));
    return first;
}
```

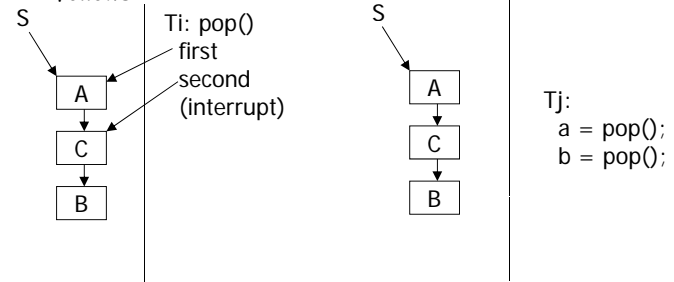
What's wrong?

CSC469



## ABA Problem

- $T_i, T_j$  both doing pops and pushes, interleaved as follows:

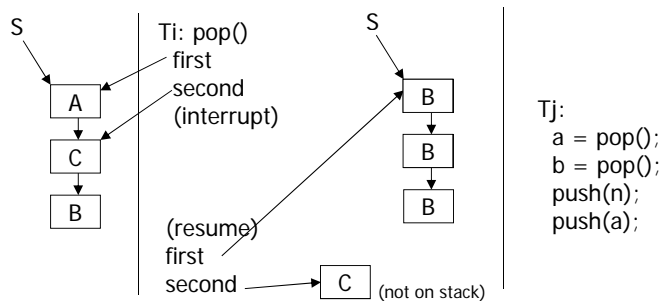


CSC469



## ABA Problem

- $CAS(x, y, z)$  succeeds if value stored at  $x$  matches  $y$



CSC469



## One Solution

- Include a version number with every pointer
  - $pointer\_t = \langle pointer, version \rangle$
  - Increment version number (atomically) every time you modify pointer
  - Change to version number guarantees CAS will fail if pointer has changed
  - Requires double-word CAS operation (not every architecture provides this)
  - May restrict reuse of memory

CSC469



## Using NBS

- Good for simple data structures, update heavy
- When you need NBS constraints/guarantees
  - Progress in face of failure
  - Linearizability
    - Everyone agrees on all intermediate states
- Both constraints are often irrelevant!

CSC469



## Constraints Irrelevant?

- Real systems don't fail the way theoretical ones do
  - Software bugs are not always fail-stop
  - Preemption/interrupt is not a failure
    - And can be controlled by system programmer or scheduler-conscious synchronization
  - Page fault is not a failure
    - Over-provision memory... if shared data really is paged out, it will have to be brought into memory before progress is made anyway
- Don't always need intermediate states, just final
  - Linearizability implies dependency  $\rightarrow$  limits parallelism
  - If events are unrelated, asynchronous, does it matter which happened first?

CSC469



## Read-Copy Update (RCU)


- What is RCU?
  - Paul McKenney's PhD thesis, a key part of the Linux scalability effort, and one of the key technologies in the SCO lawsuit against IBM.
- Ok, what is it really?
  - Reader-writer synchronization mechanism
    - Readers use no locks; best for read-mostly data structures
    - Writers create new versions atomically (typically by locking out other writers)
    - Readers can continue to access old versions
      - Old versions must be deleted at some point ("poor man's garbage collection")

CSC469



## RCU Example

- T1 traversing linked list, T2 removes an element:




CSC469



## RCU Example (2)

- After removal - T1 continues to use N and later nodes in the list



When is it ok to delete N (and reuse the memory for something else)?

CSC469



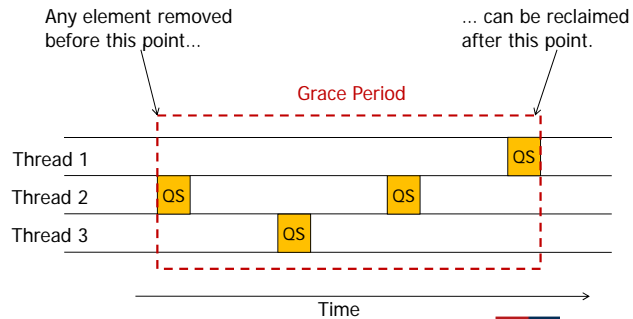
## Handling read-reclaim races

- RCU uses *quiescent state based reclamation* (QSBR)
- **Defn:** A *quiescent state* for a thread T is a state in which T holds no references to shared data
- **Defn:** A *grace period* is any interval in which every thread has passed through at least one quiescent state
- **Basic Idea:** elements removed from a data structure can be reclaimed after a grace period, since no thread can still be holding a reference to the old element at that point

CSC469



## Illustration



CSC469



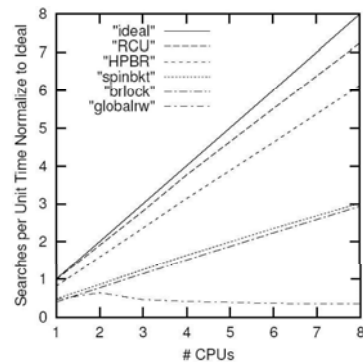
## How to define Quiescent States?

- Application dependent!
- For OS kernels, some natural ones exist
  - E.g. a context switch in a non-preemptive kernel
- RCU primitives
  - `rcu_read_lock()` and `rcu_read_unlock`
    - Surround read-side critical sections
    - No overhead (`#define'd as nothing`) in non-preemptive kernels
    - Quiescent state may not occur inside read-side critical section
  - `synchronize_rcu()`
    - Wait until all pre-existing RCU read-side critical sections complete

CSC469



## PPC Hash Table with RCU



CSC469



## When to use which tool

- Read-mostly situations
    - RCU (if algorithm can tolerate concurrent reads and updates)
  - Update-heavy situations
    - Simple data structures and algorithms: NBS
    - Complex data structures and algorithms: Locking
- "When the only tool you have is a hammer, everything looks like a nail."*
- It's good to have lots of tools in your toolbox

CSC469

