

Transactional Memory 4/14/10

Parallel processors

- so far locks

- issues w/ locks (esp. on parallel processors)

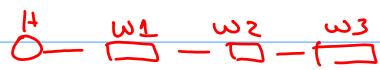
1) overhead of locks

$a = a + 1$

~ 3 instructions

lock $a = a + 1$ unlock

- identified in the database literature early
on.
- * 2) "convoing" : if current holder of a lock gets delayed
(e.g. by page fault or timeslicing) — everybody else who
wants that lock has to wait



- 3) deadlock — static avoidance may be impossible
dynamic detection or avoidance — expensive

- * 4) Unbounded priority inversion
strict priority: the highest-priority ^{ready} runnable processes are assigned to processors.
- Low priority ^{thread} holds a lock.
High priority thread wants the lock. ← priority inversion because LP Thread Runs and HP doesn't.
- Medium priority CPU-bound thread.
- Ad hocery to fix this like priority inheritance — when HP thread waits
LP thread's priority is temporarily increased to that of highest priority waiter.

Almost led to failure of Mars lander mission

Beginning in mid 1980's : lock-free synchronization -
rather than prevent synchronization, check for them before
committing changes.

→ transactional memory

Today's paper Herlihy & Moss ↑ 1993 — hardware-based approach.

How to implement lock on a multiprocessor:

build using Test-and-set TAS Compare-and-Swap CAS — check then
act atomically.

Spinlock using TAS

TAS is defined as

TAS (int *l) {
 tmp = *l; } atomic
 *l = 1
 return tmp

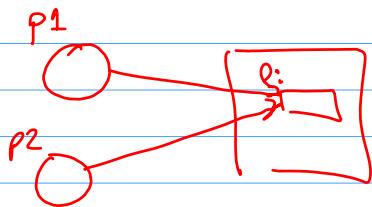
while (TAS(l)) {} ← lock

Assert *l == 1

critical section

*l = 0 ← unlock

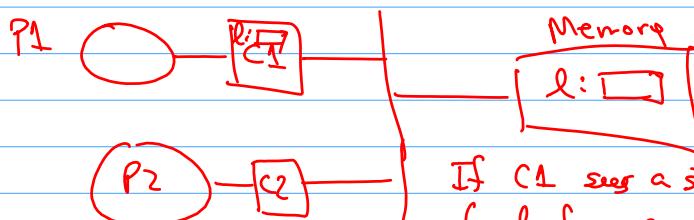
The above spinlock has terrible memory contention



To fix this - test-and-test-and-set TTS

```
TTS (int *l) {  
    tmp = *l  
    if tmp == 0 return TAS(l)  
    return tmp  
}
```

This works because modern multiprocessors have caches betw.
processors and shared memory Bus



Snoopy cache:

Caches monitor the Bus,
snooping

If C1 sees a store sees a fetch and the
for l from C2 cached version is "dirty" -
it drops l from then cache responds
C1 faster than M.

Still hammering the cache.

Suggestion is - back off - by doing some amount of
purely local waiting if you find the lock held.

Design based on optimistic concurrency control using
LL - load-locked } instruction pairs
SC - store-conditional

$\text{tmp} = \text{LL}(l)$

later

$\text{indicator} = \text{SC}(l, v)$ if no change has been made to l
since the LL

$\text{tmp} = *l$
and processor remembers that l is "locked".

$*l = v$

return true

else return false.

```
do {  
    tmp = LL(a)  
    succ = SC(a, tmp + 1) }  
until succ
```

→ very basic transaction mechanism -
succ == True ~ commit
succ == False ~ abort

What this paper does is decouple the commit part of SC from the store part and makes processor keep track of success or failure for multiple locations.

for 1-variable transaction

```
ini(l) {  
    do {  
        ST(l, LTX(l) + 1)  
    } until (commit())  
}
```

load transactional

commit

doubly linked list :