

# Synchronization

## Part 1: Synchronization - Locks

Dekker's Algorithm and the Bakery Algorithm provide software-only synchronization. Thanks to advancements in hardware, synchronization approaches have been simplified and have significantly improvements in performance. One basis for synchronization is a simple primitive, TestAndSet.

TestAndSet (&x) = Set a 1 in location of x, return old value

- 1) If a location is initialized to 0, what will the first invocation of TestAndSet on that location return?

0

- 2) If 8 threads all call TestAndSet on the same location, how many different values will be seen by the threads?

2 values, 0 and 1

- 3) Would 1 of the 8 threads receive a unique value, i.e. one that no other thread received?

Yes, 1 thread will receive 0

- 4) Given the following code for ReleaseLock(Lock\* L), provide an implementation for AcquireLock(Lock\* L) using TestAndSet:

```
ReleaseLock(Lock* L) :                               AcquireLock(Lock* L) :
    *L = 0                                           While ( TestAndSet(L) == 1 ) :
                                                    // SPIN
```

- 5) Executing TestAndSet requires the M coherence state. If 8 threads all call AcquireLock on the same lock and there are no calls to ReleaseLock, how many GetRdX and GetRd requests have been made after every thread has executed the first line in your AcquireLock code once?

8 GetRdX requests

- 6) Implement AcquireLock, so that (after the first attempt) it only tries to modify the lock if it is available.

```
AcquireLock(Lock* L):  
    While ( TestAndSet(L) == 1 ) :  
        While ( *L == 1 ) :  
            // SPIN
```

- 7) How long a thread should wait before trying to acquire the lock varies depends. Waiting too long results in the wasted cycles, while not long enough results in extra coherence requests. What is the shortest time a thread should wait?

**1 Cycle, which is the shortest possible time until the current holder of the lock may try to release the lock.**

- 8) Operating systems can context switch threads when one thread is blocked. Is a thread waiting on a lock, blocked? If a context switch costs X cycles, how long should a thread wait for a lock before it yields to the operating system?

**A waiting thread is blocked and not producing useful work. If we knew the thread would be waiting for more than  $2X + 1$  cycles, then it would be useful to switch.**

- 9) If we have a range of values (Q7 and Q8), what is an algorithm that will find the correct waiting time with minimal coherence requests, and no more than half of the cycles wasted?

**Try 1, then 2, then 4, ...  
Have a spin counter and double it after each failure.**

- 10) Copy down what the two features are known as:

Q6 - Test and Test-and-Set

Q9 - Exponential Backoff

## Part 2 – Synchronization – Queued Spinlocks

Queued locks provide better guarantees, under certain scenarios than the wait algorithm from Q9. Queued locks rely on a different synchronization primitive, compare-and-swap / compare-exchange (CMPXCHG):

```
int CMPXCHG (int* loc, int oldval, int newval)
{
    ATOMIC();
    int old_req_val = *loc;
    if (old_req_val == oldval)
        *loc = newval;
    END_ATOMIC();
    return old_req_val;
}
```

11) Describe when compare and swap updates the provided location.

Compare and swap updates the location when the provided old value is the current value in the location. (i.e., `old_req_val == oldval`)

12) Provide an implementation of TestAndSet using CMPXCHG.

```
TestAndSetWithXchg(int* loc):
    return CMPXCHG(loc, 0, 1);
```

13) Sketch the execution of two threads calling TestAndSetWithXchg() on the same location.

The execution of two threads will show that one thread will get `old_req_val` of 0 and update the location, while the second thread will not update the location.

This lock is also known as the mcs lock, from the inventors' initials. See - <http://dx.doi.org/10.1145/103727.103729>

Given the Acquire and Release routines for a queued lock:

```
AcquireQLock(*glock, *mlock)                ReleaseQLock(*glock, *mlock)
{
    mlock->next = NULL;
    mlock->state = UNLOCKED;
    ATOMIC();
    prev = glock
    *glock = mlock
    END_ATOMIC();
    if (prev == NULL) return;
    mlock->state = LOCKED;
    prev->next = mlock;
    while (mlock->state == LOCKED)
        ; // SPIN
}

do {
    if (mlock->next == NULL) {
        x = CMPXCHG(glock, mlock, NULL);
        if (x == mlock) return;
    }
    else
    {
        mlock->next->state = UNLOCKED;
        return;
    }
} while (1);
```

14) Knowing that `glock` is the global lock (with initial value of `NULL`) and `mlock` is space allocated by a thread; after one thread calls acquire, what is the state of `glock`?

Glock will point to mlock

15) Given Q14, briefly describe how the call to release will, in fact, release the lock.

Since the lock is not contended, `glock` still points to `mlock`, so `CMPXCHG` will put `glock` back to `NULL`.

16) Work through the acquire calls for two threads, after the two calls have been made, what is the "stable" state of the machine? That is, what values do `glock`, and the two `mlocks` contain?

Glock will point to `mlock2`  
`mlock1->next = mlock2`  
`mlock2->state = LOCKED`

17) After the two threads have called Acquire, one has the lock and the other waits. When the holder of the lock calls release queue lock, briefly describe how the other thread knows the lock is available.

The releasing thread will see that its next point is not `NULL`.  
It will then signal the next thread using this pointer and do so via setting the waiting thread's state to `UNLOCKED`.

See also - <http://dx.doi.org/10.1109/MICRO.2010.12>

18) If a thread is in the middle of the acquiring and releasing the lock? (Assume glock and each mlock is a different cache line).

6:	Release path:
1) mlock1	4) mlock1 (by T1)
2) glock	5) mlock2 (by T1)
3) mlock2	6) mlock2 (by T2)

19) What happens if a thread in the middle of the queue is context switched out? That is, if it is signaled before it is context switched back in?

The lock will be held by a thread that has been context switched out. This is always bad. Given that other threads may be waiting on this lock and the lock could have been given to them, a queued lock does worse than other implementations in this scenario.

### Part 3 - Synchronization Barriers

See also - <http://dx.doi.org/10.1145/379539.379566>

Barriers are when the program wants to synchronize a collection of threads together, often allowing the threads to advance to a new "time" or common unit of work.

```
counter = release = 0; //long ago, far away..
BarrierWait():
    if (AtomicIncrement(counter) == N) { // N threads total
        counter = 0; // Reset counter
        release = 1; // Release threads
    } else {
        while (release != 1)
            ; // Spin on release to be 1
    }
}
```

20) Given the simple implementation above for a barrier, discuss how this would perform versus the implementation in Q20.

This approach uses an atomic instruction to handle the updating of the counter (it could be done in a lock). The signal uses a separate global variable (although often the counter and release are part of a barrier struct) to signal to waiting threads that all have arrived.

21) In order to reset the barrier, we need to set release back to 0. Conceptually, when should this be done?

This should be done after the last thread has returned from BarrierWait().

Given that it is difficult to achieve Q21, barriers can instead be written using a toggle. We replace “1” with a toggle variable to signal thread release.

```
counter = release = 0; //long ago, far away...
toggle = !release;
BarrierWait():
    if (AtomicIncrement(counter) == N) { // N threads total
        counter = 0; // Reset counter
        release = !release; // Release threads
        toggle = !toggle; // Reset signal
    } else {
        while (release != toggle)
            ; // Spin on release
    }
```

22) This implementation has a bug in it. Identify why this barrier may not work consistently and what its intended operation is.

There is a narrow window between switching release and toggle when waiting threads can continue. Using a local variable to store the old toggle value allows the switching of the values to be observed.

23) Due to various circumstances, some of the threads may not be running when the barrier is released. In such a scenario, one thread may reach the barrier for time N+1 while other threads are still sleeping on barrier for time N. Does the toggle implementation prevent this thread from executing the barrier for time N+1 early? Why or why not?

The danger remains that if the signaling thread is context switched out before the writes are complete, then the barrier could stay in the signal state. Knowing that `release = !release` happens before `toggle = !toggle`, no thread will leave the barrier until this last write happens. With weak consistency models we would need a barrier / synchronization between these two writes.

24) In the above algorithm, on what is there contention?

counter

25) Conceptually, how might the barrier be modified to reduce contention?

One possibility is a tree combining.