

Spinlocks

Samy Al Bahra Devon H. O'Dell

Message Systems, Inc.

April 8, 2011

A mutex is an object which implements acquire and relinquish operations such that the execution following an acquire operation and up to the relinquish operation is executed in a mutually exclusive manner relative to the object implementing a mutex.


Locks are an implementation of a mutex.

- Sleep lock
 - Any mutex type which deactivates processes that attempt to acquire a mutex that has already been acquired by another process until a relinquish operation on the mutex activates one or more of them.
- Spinlock
 - Any mutex type which forces callers of an acquire operation to spend an unbounded number of processor cycles re-evaluating the availability of the mutex until it has been acquired. The process that invokes acquire is never deactivated before the completion of the acquire operation.
- Spinlocks are preferred to sleep mutexes when the waiting time for a resource is less than the time for the scheduling overhead of process activation/deactivation or when scheduling simply is not possible.

```
void
lock(uint32_t *mutex)
{
    while (ck_pr_fas_32(mutex, true) != false)
        ck_pr_stall();

    return;
}


void
unlock(uint32_t *mutex)
{
    *mutex = false;
    return;
}
```



Spinlocks

- └ Non-Arbitrating Spinlocks


- └ Cache Coherency




Spinlocks

- └ Non-Arbitrating Spinlocks

- └ Cache Coherency






Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

- Non-Arbitrating Spinlocks


- Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



Spinlocks

└ Non-Arbitrating Spinlocks


└ Cache Coherency



```
void
lock(uint32_t *mutex)
{
    while (ck_pr_fas_32(mutex, true) != false) {
        while (ck_pr_load_32(mutex) == true)
            ck_pr_stall();
    }

    return;
}

void
unlock(uint32_t *mutex)
{
    *mutex = false;
    return;
}
```



```
void
lock(uint32_t *mutex)
{


    ck_backoff_t backoff = CK_BACKOFF_INITIALIZER;


    while (ck_pr_fas_32(mutex, true) != false)
        ck_backoff_eb(&backoff);

    return;
}

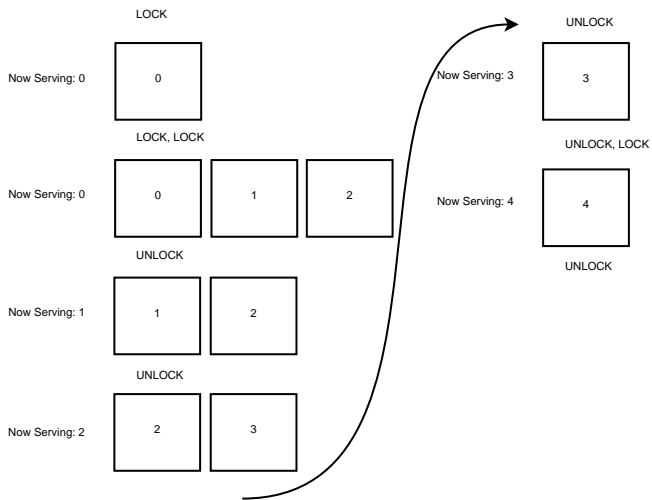
void
unlock(uint32_t *mutex)
{


    *mutex = false;
    return;
}
```

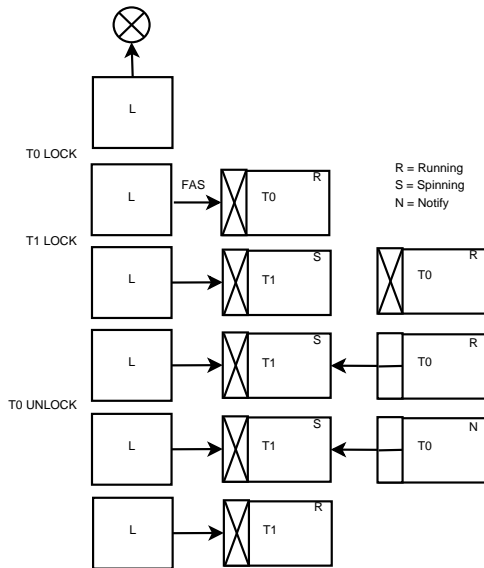



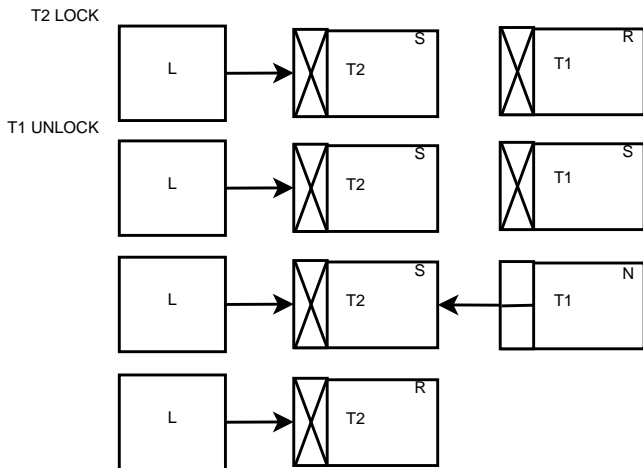


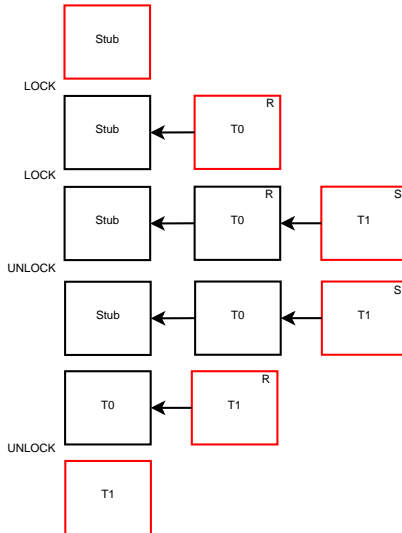
- Non-arbitrating spinlocks do not provide fairness (or starvation-freedom) guarantees.











Fast path latency

See <http://concurrencykit.org/doc/appendixZ.html>

- Mutexes in general are not composable.
- Subtle ordering issues can lead to hard-to-detect deadlock conditions.
- Blocking synchronization is sensitive to preemption.

Questions?