

Sep 25, 24 9:47

spinlock-mutex.txt

Page 3/3

```

95 3. Mutex implementation
96
97      The intent of a mutex is to avoid busy waiting: if the lock is not
98      available, the locking thread is put to sleep, and tracked by a
99      queue in the mutex. The next page has an implementation.
100
101

```

Sep 25, 24 9:47

fair-mutex.c

Page 1/1

```

1 #include <sys/queue.h>
2
3 typedef struct thread {
4     // ... Entries elided.
5     STAILQ_ENTRY(thread_t) qlink; // Tail queue entry.
6 } thread_t;
7
8 struct Mutex {
9     // Current owner, or 0 when mutex is not held.
10    thread_t *owner;
11
12    // List of threads waiting on mutex
13    STAILQ(thread_t) waiters;
14
15    // A lock protecting the internals of the mutex.
16    Spinlock splock; // as in item 1, above
17 };
18
19 void mutex_acquire(struct Mutex *m) {
20
21     acquire(&m->splock);
22
23     // Check if the mutex is held; if not, current thread gets mutex and returns
24     if (m->owner == 0) {
25         m->owner = id_of_this_thread;
26         release(&m->splock);
27     } else {
28         // Add thread to waiters.
29         STAILQ_INSERT_TAIL(&m->waiters, id_of_this_thread, qlink);
30
31         // Tell the scheduler to add current thread to the list
32         // of blocked threads. The scheduler needs to be careful
33         // when a corresponding sched_wakeup call is executed to
34         // make sure that it treats running threads correctly.
35         sched_mark_blocked(&id_of_this_thread);
36
37         // Unlock spinlock.
38         release(&m->splock);
39
40         // Stop executing until woken.
41         sched_swtch();
42
43         // When we get to this line, we are guaranteed to hold the mutex. This
44         // is because we can get here only if context-switched-TO, which itself
45         // can happen only if this thread is removed from the waiting queue,
46         // marked "unblocked", and set to be the owner (in mutex_release()
47         // below). However, we might have held the mutex in lines 39-42
48         // (if we were context-switched out after the spinlock release()),
49         // followed by being run as a result of another thread's release of the
50         // mutex). But if that happens, it just means that we are
51         // context-switched out an "extra" time before proceeding.
52     }
53 }
54
55 void mutex_release(struct Mutex *m) {
56     // Acquire the spinlock in order to make changes.
57     acquire(&m->splock);
58
59     // Assert that the current thread actually owns the mutex
60     assert(m->owner == id_of_this_thread);
61
62     // Check if anyone is waiting.
63     m->owner = STAILQ_GET_HEAD(&m->waiters);
64
65     // If so, wake them up.
66     if (m->owner) {
67         sched_wakeone(&m->owner);
68         STAILQ_REMOVE_HEAD(&m->waiters, qlink);
69     }
70
71     // Release the internal spinlock
72     release(&m->splock);
73 }

```

Sep 29, 24 21:39

handout07.txt

Page 1/3

```

1 CS 202, Fall 2024
2 Handout 7 (Class 8)
3
4 1. Simple deadlock example
5
6 T1:
7     acquire(mutexA);
8     acquire(mutexB);
9
10    // do some stuff
11
12    release(mutexB);
13    release(mutexA);
14
15 T2:
16    acquire(mutexB);
17    acquire(mutexA);
18
19    // do some stuff
20
21    release(mutexA);
22    release(mutexB);
23

```

Sep 29, 24 21:39

handout07.txt

Page 2/3

```

24 2. More subtle deadlock example
25
26     Let M be a monitor (shared object with methods protected by mutex)
27     Let N be another monitor
28
29 class M {
30     private:
31         Mutex mutex_m;
32
33     // instance of monitor N
34     N another_monitor;
35
36     // Assumption: no other objects in the system hold a pointer
37     // to our "another_monitor"
38
39     public:
40         M();
41         ~M();
42         void methodA();
43         void methodB();
44     };
45
46 class N {
47     private:
48         Mutex mutex_n;
49         Cond cond_n;
50         int available;
51
52     public:
53         N();
54         ~N();
55         void* alloc(int nwanted);
56         void free(void*);
57     }
58
59     int
60     N::alloc(int nwanted) {
61         acquire(&mutex_n);
62         while (available < nwanted) {
63             wait(&cond_n, &mutex_n);
64         }
65
66         // peel off the memory
67
68         available -= nwanted;
69         release(&mutex_n);
70     }
71
72     void
73     N::free(void* returning_mem) {
74
75         acquire(&mutex_n);
76
77         // put the memory back
78
79         available += returning_mem;
80
81         broadcast(&cond_n, &mutex_n);
82
83         release(&mutex_n);
84     }
85

```

Sep 29, 24 21:39

handout07.txt

Page 3/3

```

86     void
87     M::methodA() {
88
89         acquire(&mutex_m);
90
91         void* new_mem = another_monitor.alloc(int nbytes);
92
93         // do a bunch of stuff using this nice
94         // chunk of memory n allocated for us
95
96         release(&mutex_m);
97     }
98
99     void
100    M::methodB() {
101
102        acquire(&mutex_m);
103
104        // do a bunch of stuff
105
106        another_monitor.free(some_pointer);
107
108        release(&mutex_m);
109    }
110
111 QUESTION: What's the problem?

```

Sep 29, 24 21:38

filemap.txt

Page 1/2

```

1  2. Locking brings a performance vs. complexity trade-off
2
3  /*
4   *      linux/mm/filemap.c
5   *
6   * Copyright (C) 1994-1999 Linus Torvalds
7   */
8
9  /*
10  * This file handles the generic file mmap semantics used by
11  * most "normal" filesystems (but you don't /have/ to use this:
12  * the NFS filesystem used to do this differently, for example)
13  */
14 #include <linux/export.h>
15 #include <linux/compiler.h>
16 #include <linux/dax.h>
17 #include <linux/fs.h>
18 #include <linux/sched/signal.h>
19 #include <linux/uaccess.h>
20 #include <linux/capability.h>
21 #include <linux/kernel_stat.h>
22 #include <linux/gfp.h>
23 #include <linux/mm.h>
24 #include <linux/swap.h>
25 #include <linux/mman.h>
26 #include <linux/pagemap.h>
27 #include <linux/file.h>
28 #include <linux/uio.h>
29 #include <linux/hash.h>
30 #include <linux/writeback.h>
31 #include <linux/backing-dev.h>
32 #include <linux/pagevec.h>
33 #include <linux/blkdev.h>
34 #include <linux/security.h>
35 #include <linux/cpuset.h>
36 #include <linux/hugetlb.h>
37 #include <linux/memcontrol.h>
38 #include <linux/cleancache.h>
39 #include <linux/shmem_fs.h>
40 #include <linux/rmap.h>
41 #include "internal.h"
42
43 #define CREATE_TRACE_POINTS
44 #include <trace/events/filemap.h>
45
46 /*
47  * FIXME: remove all knowledge of the buffer layer from the core VM
48  */
49 #include <linux/buffer_head.h> /* for try_to_free_buffers */
50
51 #include <asm/mman.h>
52
53 /*
54  * Shared mappings implemented 30.11.1994. It's not fully working yet,
55  * though.
56  *
57  * Shared mappings now work. 15.8.1995 Bruno.
58  *
59  * finished 'unifying' the page and buffer cache and SMP-threaded the
60  * page-cache, 21.05.1999, Ingo Molnar <mingo@redhat.com>
61  *
62  * SMP-threaded pagemap-LRU 1999, Andrea Arcangeli <andrea@suse.de>
63  */
64
65 /*
66  * Lock ordering:
67  *
68  * ->i_mmap_rwsem          (truncate_pagecache)
69  *   ->private_lock          (_free_pte->_set_page_dirty_buffers)
70  *   ->swap_lock              (exclusive_swap_page, others)
71  *     ->i_pages lock
72  *
73  * ->i_mutex

```

Sep 29, 24 21:38

filemap.txt

Page 2/2

```

74 *      ->i_mmap_rwsem          (truncate->unmap_mapping_range)
75 *
76 *      ->mmap_sem
77 *      ->j_mmap_rwsem
78 *          ->page_table_lock or pte_lock  (various, mainly in memory.c)
79 *          ->i_pages lock           (arch-dependent flush_dcache_mmap_lock)
80 *
81 *      ->mmap_sem
82 *          ->lock_page          (access_process_vm)
83 *
84 *      ->i_mutex
85 *          ->mmap_sem
86 *
87 *      bdi->wb.list_lock
88 *          sb_lock              (fs/fs-writeback.c)
89 *          ->i_pages lock       (_sync_single_inode)
90 *
91 *      ->i_mmap_rwsem
92 *          ->anon_vma.lock      (vma_adjust)
93 *
94 *      ->anon_vma.lock
95 *          ->page_table_lock or pte_lock   (anon_vma_prepare and various)
96 *
97 *      ->page_table_lock or pte_lock
98 *          ->swap_lock            (try_to_unmap_one)
99 *          ->private_lock         (try_to_unmap_one)
100 *          ->i_pages lock        (try_to_unmap_one)
101 *          ->zone_lru_lock(zone)  (follow_page->mark_page_accessed)
102 *          ->zone_lru_lock(zone)  (check_pte_range->isolate_lru_page)
103 *          ->private_lock         (page_remove_rmap->set_page_dirty)
104 *          ->i_pages lock        (page_remove_rmap->set_page_dirty)
105 *          bdi.wb->list_lock     (page_remove_rmap->set_page_dirty)
106 *          ->inode->i_lock       (page_remove_rmap->set_page_dirty)
107 *          ->memcg->move_lock    (page_remove_rmap->lock_page_memcg)
108 *          bdi.wb->list_lock     (zap_pte_range->set_page_dirty)
109 *          ->inode->i_lock       (zap_pte_range->set_page_dirty)
110 *          ->private_lock         (zap_pte_range->_set_page_dirty_buffers)
111 *
112 *      ->i_mmap_rwsem
113 *          ->tasklist_lock        (memory_failure, collect_procs_ao)
114 */
115
116 static int page_cache_tree_insert(struct address_space *mapping,
117                                   struct page *page, void **shadowp)
118 {
119     struct radix_tree_node *node;
120     .....
121 [the point is: fine-grained locking leads to complexity.]
```

Sep 29, 24 21:38

dclp.txt

Page 1/2

```

1 3. Cautionary tale
2
3 Consider the code below:
4
5  struct foo {
6      int abc;
7      int def;
8  };
9  static int ready = 0;
10 static mutex_t mutex;
11 static struct foo* ptr = 0;
12
13 void
14 doublecheck_alloc()
15 {
16     if (!ready) { /* <-- accesses shared variable w/out holding mutex */
17         mutex_acquire(&mutex);
18         if (!ready) {
19             ptr = alloc_foo(); /* <-- sets ptr to be non-zero */
20             ready = 1;
21         }
22     }
23     mutex_release(&mutex);
24
25 }
26
27 return;
28 }

29 This is an example of the so-called "double-checked locking pattern."
30 The programmer's intent is to avoid a mutex acquisition in the common
31 case that 'ptr' is already initialized. So the programmer checks a flag
32 called 'ready' before deciding whether to acquire the mutex and
33 initialize 'ptr'. The intended use of doublecheck_alloc() is something
34 like this:
35
36 void f() {
37     doublecheck_alloc();
38     ptr->abc = 5;
39 }
40
41 void g() {
42     doublecheck_alloc();
43     ptr->def = 6;
44 }
45
46 We assume here that mutex_acquire() and mutex_release() are implemented
47 correctly (each contains memory barriers internally, etc.). Furthermore,
48 we assume that the compiler does not reorder instructions.
49
50 NEVERTHELESS, on multi-CPU machines that do not offer sequential
51 consistency, doublecheck_alloc() is broken. What is the bug?
52
53 -----
54
55 Unfortunately, double-checked initialization (or double-checked locking
56 as it's sometimes known) is a common coding pattern. Even some
57 references on threads suggest it! Still, it's broken.
58
59 While you can fix it (in C) by adding barriers (exercise:
60 where?), this is not recommended, as the code is tricky to reason about.
61 One of the points of this example is to show you why it's so important
62 to protect global data with a mutex, even if "all" one is doing is
63 reading memory, and even if the shortcut looks harmless.
64
65
```