


- ☑ 1. Last time
- ☑ 2. Deadlock, continued
- ☑ 3. Other progress issues
- ☑ 4. Performance issues
- ☑ 5. Programmability issues
- ☑ 6. Mutexes and interleaving


ONE HANDOUT

2. Deadlock



Happens when all four of these conditions are present:

- i. mutual exclusion
- ii. hold and wait
- iii. no pre-emption
- iv. circular wait



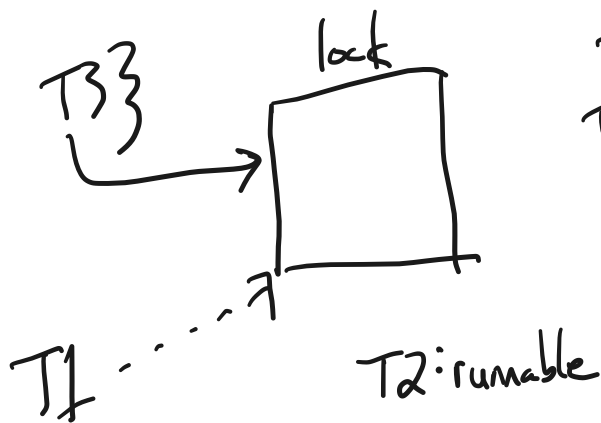
What can we do about deadlock?

- (a) ignore it
- (b) detect + recover
- (c) avoid algorithmically
- (d) negate one of the 4 conditions
- (e) static/dynamic detection tools

3. Other progress issues

- Starvation

- Priority inversion



Assume: highest-prio runnable thread runs.

4. Performance issues and tradeoffs

(a) Invoking spinlocks/mutexes can be expensive

MCS locks

futexes

(b) Coarse locks limit available parallelism...

(c) ... but fine-grained locking leads to complexity and hence bugs

5. Programmability issues

- loss of modularity

```
printf("...");
```

- what's the fundamental problem?



6. mutexes and interleavings

ops

mem barrier

ops

≡ A

release() 
B 

"A stuff" happens before "B stuff"
although, within A and B, the "stuff"
can happen out of order from the perspective
of other cores, assuming they could view the
operations separately.

```

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

```

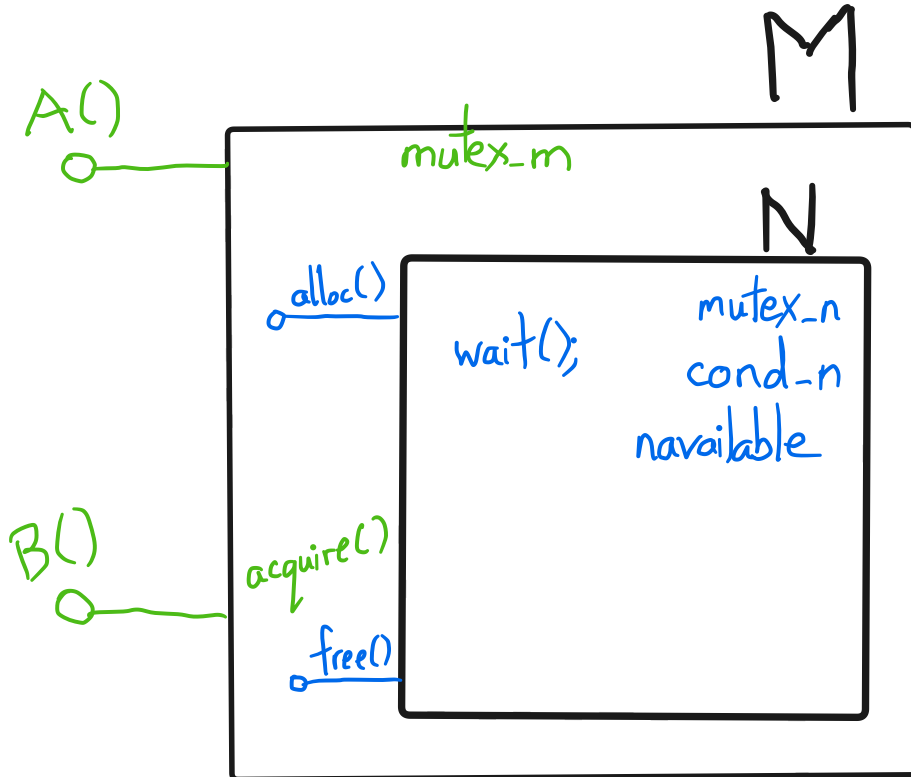
```

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 navailable;
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 (navailable < nwanted) {
63             wait(&cond_n, &mutex_n);
64         }
65
66         // peel off the memory
67
68         navailable -= nwanted;
69         release(&mutex_n);
70     }
71
72     void
73     N::free(void* returning_mem) {
74         acquire(&mutex_n);
75
76         // put the memory back
77
78         navailable += returning_mem;
79
80         broadcast(&cond_n, &mutex_n);
81
82         release(&mutex_n);
83     }
84
85

```

```

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?
    
```



```

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
    
```

```

74 *      ->i_mmap_rwsem          (truncate->unmap_mapping_range)
75 *
76 *      ->mmap_sem
77 *      ->i_mmap_rwsem
78 *      ->page_table_lock or pte_lock  (various, mainly in memory.c)
79 *      ->i_pages lock          (arch-dependent flush_dcachemmap_lock)
80 *
81 *      ->mmap_sem
82 *      ->lock_page            (access_process_vm)
83 *
84 *      ->i_mutex                (generic_perform_write)
85 *      ->mmap_sem              (fault_in_pages_readable->do_page_fault)
86 *
87 *      bdi->wb.list_lock        (fs/fs-writeback.c)
88 *      sb_lock                  (__sync_single_inode)
89 *      ->i_pages lock
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
122 [the point is: fine-grained locking leads to complexity.]

```

```

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
18             mutex_acquire(&mutex);
19             if (!ready) {
20                 ptr = alloc_foo(); /* <-- sets ptr to be non-zero */
21                 ready = 1;
22
23                 CPU1
24
25                 mutex_release(&mutex);
26                 CPU2
27             }
28             return;
29         }
30
31 This is an example of the so-called "double-checked locking pattern."
32 The programmer's intent is to avoid a mutex acquisition in the common
33 case that 'ptr' is already initialized. So the programmer checks a flag
34 called 'ready' before deciding whether to acquire the mutex and
35 initialize 'ptr'. The intended use of doublecheck_alloc() is something
36 like this:
37
38     void f() {
39         doublecheck_alloc();
40         ptr->abc = 5;
41     }
42
43     void g() {
44         doublecheck_alloc();
45         ptr->def = 6;
46     }
47
48 We assume here that mutex_acquire() and mutex_release() are implemented
49 correctly (each contains memory barriers internally, etc.). Furthermore,
50 we assume that the compiler does not reorder instructions.
51
52 NEVERTHELESS, on multi-CPU machines that do not offer sequential
53 consistency, doublecheck_alloc() is broken. What is the bug?
54 -----
55
56 Unfortunately, double-checked initialization (or double-checked locking
57 as it's sometimes known) is a common coding pattern. Even some
58 references on threads suggest it! Still, it's broken.
59
60 While you can fix it (in C) by adding barriers (exercise:
61 where?), this is not recommended, as the code is tricky to reason about.
62 One of the points of this example is to show you why it's so important
63 to protect global data with a mutex, even if "all" one is doing is
64 reading memory, and even if the shortcut looks harmless.
65

```

Sep 29, 24 21:38

dclp.txt

Page 2/2

```
66 Finally, here are some references on this topic:
67
68 --http://www.aristeia.com/Papers/DDJ\_Jul\_Aug\_2004\_revised.pdf
69 explores issues with this pattern in C++
70
71 --The "Double-Checked Locking is Broken" Declaration:
72 http://www.cs.umd.edu/~pugh/java/memoryModel/DoubleCheckedLocking.html
73
74 --C++11 provides a way to implement the pattern correctly and
75 portably (again, using memory barriers):
76 https://preshing.com/20130930/double-checked-locking-is-fixed-in-cpp11/
```