

Feb 19, 15 13:58

handout06.txt

Page 1/7

```

1 CS 202, Spring 2015
2 Handout 6 (Class 7)
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

```

Feb 19, 15 13:58

handout06.txt

Page 2/7

```

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

```

Feb 19, 15 13:58

handout06.txt

Page 3/7

```

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?
112

```

Feb 19, 15 13:58

handout06.txt

Page 4/7

```

113 3. Locking brings a performance vs. complexity trade-off
114
115 /*
116 *      linux/mm/filemap.c
117 *
118 * Copyright (C) 1994-1999 Linus Torvalds
119 */
120
121 /*
122 * This file handles the generic file mmap semantics used by
123 * most "normal" filesystems (but you don't /have/ to use this:
124 * the NFS filesystem used to do this differently, for example)
125 */
126 #include <linux/export.h>
127 #include <linux/compiler.h>
128 #include <linux/fs.h>
129 #include <linux/uaccess.h>
130 #include <linux/aio.h>
131 #include <linux/capability.h>
132 #include <linux/kernel_stat.h>
133 #include <linux/gfp.h>
134 #include <linux/mm.h>
135 #include <linux/swap.h>
136 #include <linux/mman.h>
137 #include <linux/pagemap.h>
138 #include <linux/file.h>
139 #include <linux/uio.h>
140 #include <linux/hash.h>
141 #include <linux/writeback.h>
142 #include <linux/backing-dev.h>
143 #include <linux/pagevec.h>
144 #include <linux/blkdev.h>
145 #include <linux/security.h>
146 #include <linux/cpuset.h>
147 #include <linux/hardirq.h> /* for BUG_ON(!in_atomic()) only */
148 #include <linux/hugetlb.h>
149 #include <linux/memcontrol.h>
150 #include <linux/cleancache.h>
151 #include <linux/rmap.h>
152 #include "internal.h"
153
154 #define CREATE_TRACE_POINTS
155 #include <trace/events/filemap.h>
156
157 /*
158 * FIXME: remove all knowledge of the buffer layer from the core VM
159 */
160 #include <linux/buffer_head.h> /* for try_to_free_buffers */
161
162 #include <asm/mman.h>
163
164 /*
165 * Shared mappings implemented 30.11.1994. It's not fully working yet,
166 * though.
167 *
168 * Shared mappings now work. 15.8.1995 Bruno.
169 *
170 * finished 'unifying' the page and buffer cache and SMP-threaded the
171 * page-cache, 21.05.1999, Ingo Molnar <mingo@redhat.com>
172 *
173 * SMP-threaded pagemap-LRU 1999, Andrea Arcangeli <andrea@suse.de>
174 */
175
176 /*
177 * Lock ordering:
178 *
179 * ->i_mmap_rwsem          (truncate_pagecache)
180 * ->private_lock         (__free_pte->__set_page_dirty_buffers)
181 * ->swap_lock            (exclusive_swap_page, others)
182 * ->mapping->tree_lock
183 *
184 * ->i_mutex
185 * ->i_mmap_rwsem          (truncate->unmap_mapping_range)

```

Feb 19, 15 13:58

handout06.txt

Page 5/7

```

186 *
187 * -> mmap_sem
188 *   -> i_mmap_rwsem
189 *     -> page_table_lock or pte_lock (various, mainly in memory.c)
190 *     -> mapping->tree_lock (arch-dependent flush_dcache_mmap_lock)
191 *
192 * -> mmap_sem
193 *   -> lock_page (access_process_vm)
194 *
195 * -> i_mutex (generic_perform_write)
196 *   -> mmap_sem (fault_in_pages_readable->do_page_fault)
197 *
198 * bdi->wb.list_lock
199 *   sb_lock (fs/fs-writeback.c)
200 *   -> mapping->tree_lock (__sync_single_inode)
201 *
202 * -> i_mmap_rwsem
203 *   -> anon_vma.lock (vma_adjust)
204 *
205 * -> anon_vma.lock
206 *   -> page_table_lock or pte_lock (anon_vma_prepare and various)
207 *
208 * -> page_table_lock or pte_lock
209 *   -> swap_lock (try_to_unmap_one)
210 *   -> private_lock (try_to_unmap_one)
211 *   -> tree_lock (try_to_unmap_one)
212 *   -> zone.lru_lock (follow_page->mark_page_accessed)
213 *   -> zone.lru_lock (check_pte_range->isolate_lru_page)
214 *   -> private_lock (page_remove_rmap->set_page_dirty)
215 *   -> tree_lock (page_remove_rmap->set_page_dirty)
216 *   bdi.wb->list_lock (page_remove_rmap->set_page_dirty)
217 *   -> inode->i_lock (page_remove_rmap->set_page_dirty)
218 *   bdi.wb->list_lock (zap_pte_range->set_page_dirty)
219 *   -> inode->i_lock (zap_pte_range->set_page_dirty)
220 *   -> private_lock (zap_pte_range->__set_page_dirty_buffers)
221 *
222 * -> i_mmap_rwsem
223 *   -> tasklist_lock (memory_failure, collect_procs_ao)
224 */
225
226 static void page_cache_tree_delete(struct address_space *mapping,
227                                   struct page *page, void *shadow)
228 {
229     struct radix_tree_node *node;
230     ....
231
232
233 [the point is: fine-grained locking leads to complexity.]
234

```

Feb 19, 15 13:58

handout06.txt

Page 6/7

```

235 4. Cautionary tale
236
237 Consider the code below:
238
239     struct foo {
240         int abc;
241         int def;
242     };
243     static int ready = 0;
244     static mutex_t mutex;
245     static struct foo* ptr = 0;
246
247     void
248     doublecheck_alloc()
249     {
250         if (!ready) { /* <-- accesses shared variable w/out holding mutex */
251
252             mutex_acquire(&mutex);
253             if (!ready) {
254                 ptr = alloc_foo(); /* <-- sets ptr to be non-zero */
255                 ready = 1;
256             }
257
258             mutex_release(&mutex);
259
260         }
261         return;
262     }
263
264 This is an example of the so-called "double-checked locking pattern."
265 The programmer's intent is to avoid a mutex acquisition in the common
266 case that 'ptr' is already initialized. So the programmer checks a flag
267 called 'ready' before deciding whether to acquire the mutex and
268 initialize 'ptr'. The intended use of doublecheck_alloc() is something
269 like this:
270
271     void f() {
272         doublecheck_alloc();
273         ptr->abc = 5;
274     };
275
276     void g() {
277         doublecheck_alloc();
278         ptr->def = 6;
279     }
280
281 We assume here that mutex_acquire() and mutex_release() are implemented
282 correctly (each contains memory barriers internally, etc.). Furthermore,
283 we assume that the compiler does not reorder instructions.
284
285 NEVERTHELESS, on multi-CPU machines that do not offer sequential
286 consistency, doublecheck_alloc() is broken. What is the bug?
287
288 -----
289
290 Unfortunately, double-checked initialization (or double-checked locking
291 as it's sometimes known) is a common coding pattern. Even some
292 references on threads suggest it! Still, it's broken.
293
294 While you can fix it (in C) by adding another barrier (exercise:
295 where?), this is not recommended, as the code is tricky to reason about.
296 One of the points of this example is to show you why it's so important
297 to protect global data with a mutex, even if "all" one is doing is
298 reading memory, and even if the shortcut looks harmless.
299

```

Feb 19, 15 13:58

handout06.txt

Page 7/7

```
300 Finally, here are some references on this topic:
301
302 --http://www.aristeia.com/Papers/DDJ\_Jul\_Aug\_2004\_revised.pdf
303 explores issues with this pattern in C++
304
305 --The "Double-Checked Locking is Broken" Declaration:
306 [http://www.cs.umd.edu/~pugh/java/memoryModel/DoubleCheckedLocking.html]
307
308 --C++11 provides a way to implement the pattern correctly and
309 portably (again, using memory barriers):
310 http://preshing.com/20130930/double-checked-locking-is-fixed-in-cpp11/
```