

Transactional Memory

Zichao Yang and Qing Zheng

Applying Transactional Memory to Concurrency Bugs

Haris Volos, Andres Jaan Tack, Michael M. Swift and Shan Lu
ASPLOS 12

Transaction

A block of code declared to be an atomic region, executed by a single processor, and isolated from other regions

Atomicity → execute to completion/not at all

Consistency → correctness (programmer's job)

Isolation → side-effects remain invisible until completion

Durability → not needed

Transactional Memory

An underlying memory system capable of execution transactions

H/W TM → speculative execution (with limitations)

S/W TM → code instrumentation (performance issues)

Hybrid TM → use H/W, fall back to S/W if necessary

Bug-Fix Techniques

Atomic regions → execute atomically with isolation
(speculative execution, locks)

Explicit rollback → abort a partially executed transaction
(opportunities for retries, mimic conditional variables)

Preemptive resources → transaction-friendly
(revertible locks, I/O, system calls)

Atomic/lock serialization → a bridge between atomic region
and traditional locks

Concurrency Bugs

Writing correctly **synchronized** code can be challenge!

Deadlock (**DL**) → use locks with a wrong order

Atomicity Violation (**AL**) → fail to protect critical sections

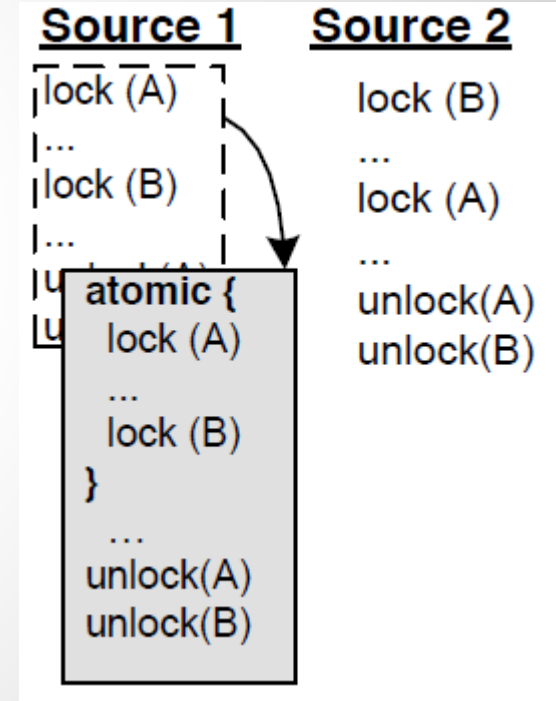
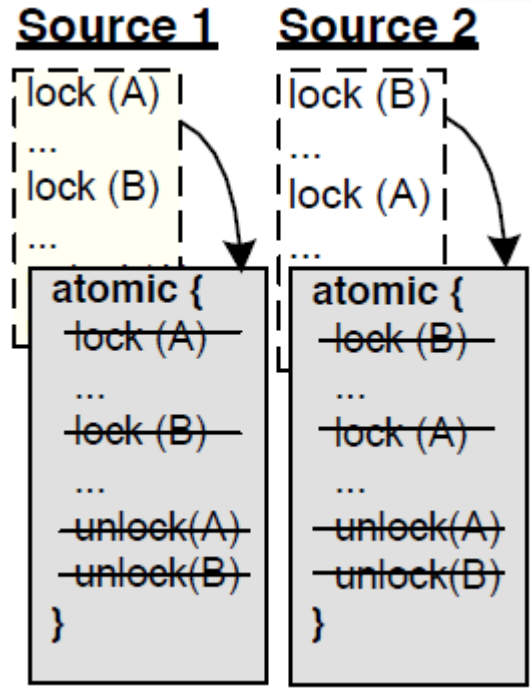
Resolving Deadlock Bugs

Recipe 1: Replacement of Deadlock-prone Locks

- (1) remove the use of multiple locks
- (2) automatically aborting conflicting threads
- (3) preserves concurrency

Recipe 2: Asymmetric Deadlock Preemption

- (1) retry lock acquisition if deadlock happens
- (2) require preemptive resources and deadlock detector
- (3) use TM only for atomicity, not isolation



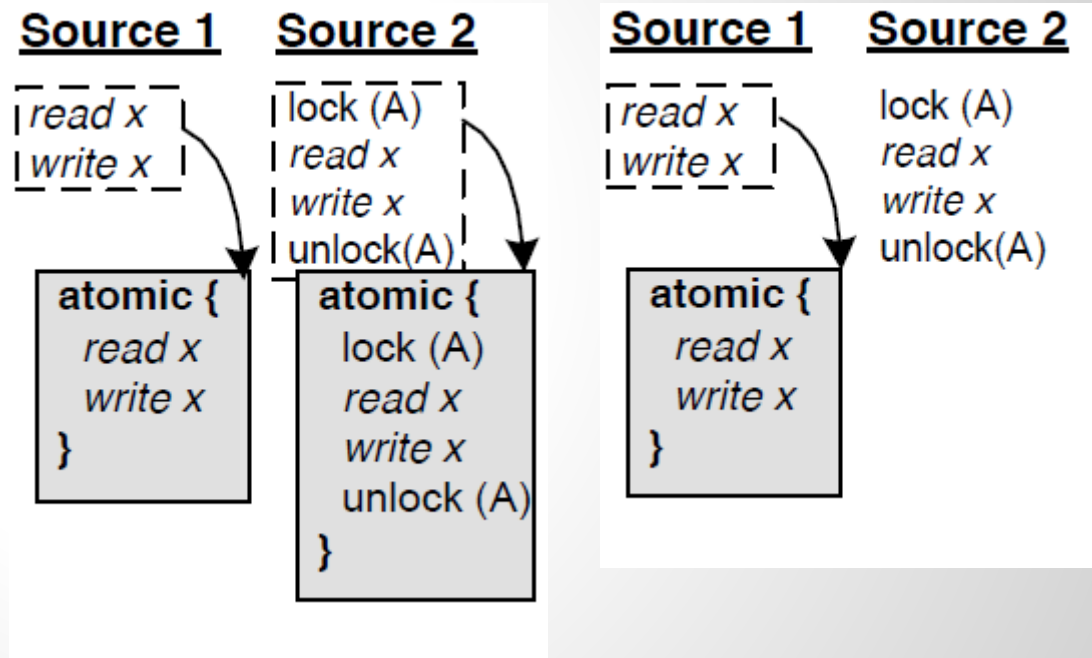
Addressing Atomic Violations

Recipe 1: Wrap All

- (1) use TM to achieve isolation and mutual exclusion
- (2) compatible with existing locks
- (3) no need to introduce new locks

Recipe 2: Warp Unprotected

- (1) reuse existing correct codes and keeps them unchanged
- (2) require atomic/lock serialization
- (3) may lead to performance issues



Evaluation & Case Study

Effectiveness

Bug	App	All	Transactional Memory Fixes				
			Total	R1	R2	R3	R4
DL	Mozilla	13	9	8(2)	-	7(1)	-
	Apache	4	2	1(1)	-	1(1)	-
	MySQL	5	1	0	-	1(1)	-
AV	Mozilla	25	20	-	20(12)	-	8(0)
	Apache	7	5	-	5(3)	-	2(0)
	MySQL	6	6	-	6(2)	-	4(0)
Total		60	43	9(2)	31(17)	9(3)	14(0)

- fix 43 out of 60 bugs: DL 12/22 AV31/38
- R1 and R2 can fix 40 out of 43

Case study

- Mozilla-I: Deadlock

```
1 js_SetProtoOrParent (...)  
2 {  
3  LOCK (rt->setSlotLock);  
4  obj2 = pobj;  
5  while (obj2) {  
6    if (obj2 == obj) {  
7      ...  
8    }  
9    LOCK_SCOPE (obj2);  
10   next_obj2 = OBJ_GET_PROTO(obj2);  
11   UNLOCK_SCOPE (obj2);  
12   obj2 = next_obj2;  
13 }  
14 /* Proceed with setting */  
15 ...  
16 UNLOCK (rt->setSlotLock);  
17 }
```

(a) Buggy code

```
1 js_SetProtoOrParent (...)  
2 {  
3  atomic {  
4    obj2 = pobj;  
5    while (obj2) {  
6      if (obj2 == obj) {  
7        ...  
8      }  
9      atomic {  
10       next_obj2 =  
11         OBJ_GET_PROTO(obj2);  
12       }  
13       obj2 = next_obj2;  
14     }  
15     /* Proceed with setting */  
16     ...  
17   }  
}
```

(b) Fixed with TM

- Deadlock
 - two threads access the two locks in different order

Case study

- Mozilla-I: Deadlock
 - developer fix (hard)
 - force threads to drop ownership before blocking
 - new conditional variables
 - TM fixes
 - recipe 1: replace lock with atomic sections
 - recipe 3: revocable locks
 - Comparison
 - recipe1 solves four other bugs (side effect)
 - performance: recipe 1: 79% worse

Case study

- Apache-I: deadlock

```
1 worker_thread(...)
2 {
3   ...
4   LOCK (timeout);
5   ...
6   UNLOCK (timeout);
7   ...
8   LOCK (idlers)
9   ...
10  SIGNAL (wait_for_idlers)
11  ...
12  UNLOCK (idlers)
13 }
```

(a) Buggy code

```
1 listener_thread (...)
2 {
3   ...
4   LOCK (timeout);
5   ...
6   LOCK (idlers);
7   ...
8   COND_WAIT (wait_for_idler,
9             idlers)
10  UNLOCK (idlers)
11  ...
12  UNLOCK (timeout)
13 }
```

```
1 listener_thread (...)
2 {
3   ...
4   atomic {
5     LOCK (timeout);
6     ...
7     LOCK (idlers);
8     ...
9     if (!COND_TRY_WAIT(...))
10    retry;
11    UNLOCK (idlers)
12  }
13  ...
14  UNLOCK (timeout)
15 }
```

(b) Fixed with TM

- Deadlock
 - listener first hold timeout then waiting for idle worker thread
 - worker thread first get timeout lock then signal

Case study

- Apache-I: deadlock
 - developer fix (hard)
 - release time out before wait
 - three failed attempts
 - TM fixes
 - recipe 3: abort transaction if wait
 - Comparison
 - simpler to fix
 - 28% worse in performance

Case study

- Apache-II: missing synchronization

```
1 void ap_buffered_log_writer (...)  
2 {  
3     ...  
4     s = &buffer[buf->outputCount];  
5     memcpy (s, str, len);  
6     temp = buf->outputCount + len;  
7     buf->outputCount = temp;  
8     apr_file_write(buf->handle);  
9     ...  
10 }
```

(a) Buggy code

```
1 void ap_buffered_log_writer (...)  
2 {  
3     ...  
4     atomic {  
5         s = &buffer[buf->outputCount];  
6         memcpy (s, str, len);  
7         temp = buf->outputCount + len;  
8         buf->outputCount = temp;  
9         apr_file_write(buf->handle);  
10    }  
11    ...  
12 }
```

(b) Fixed with TM

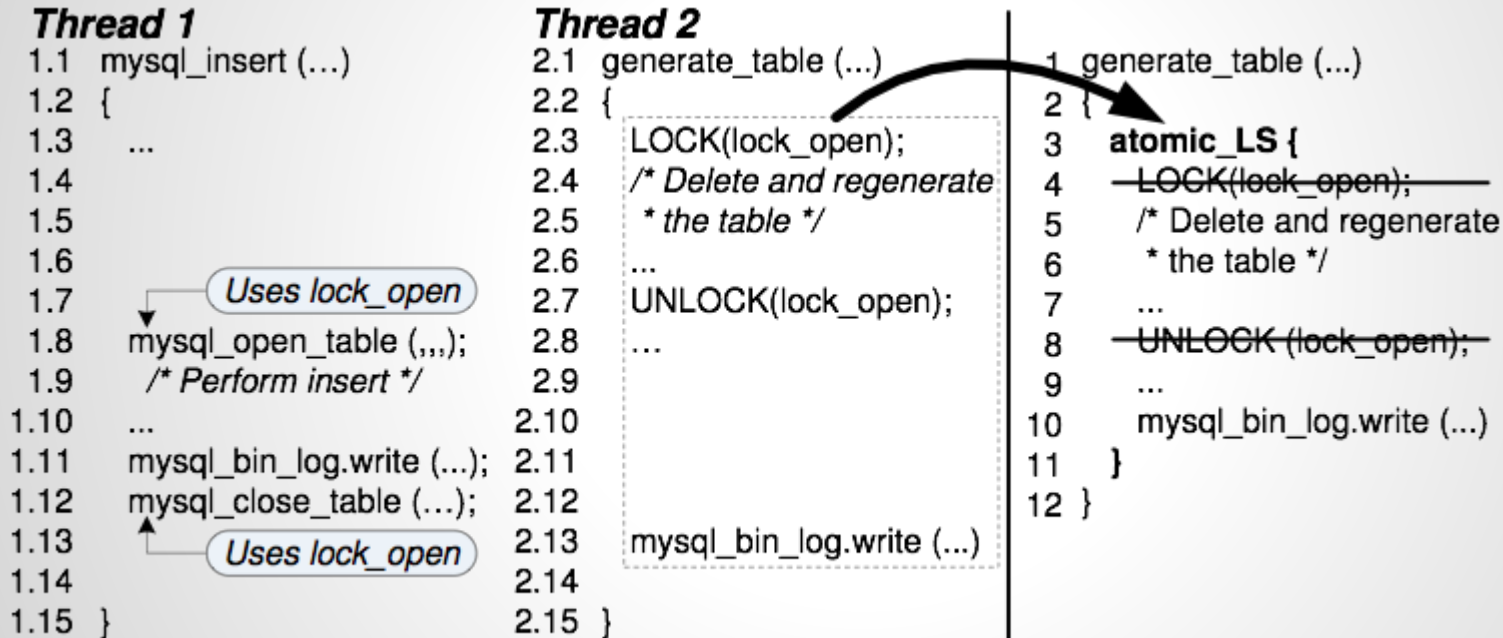
- race condition
 - two threads compete over the buffer

Case study

- Apache-II: missing synchronization
 - developer fix (medium difficulty)
 - use lock for each log device
 - TM fixes
 - insert single atomic block
 - easy
 - Comparison
 - TM simpler to fix
 - 4% slower

Case study

- MySQL-I: missing synchronization



(a) an incorrect interleaving

(b) Fixed with TM

- Deadlock
 - delete release lock too early

Case study

- Apache-II: missing synchronization
 - developer fix (hard)
 - extent lock_open to the end
 - performance implications
 - TM fixes
 - insert single atomic block
 - easy
 - Comparison
 - simple, expressive and non-invasive
 - same performance

Conclusion

- Very simple to use TM to fix bugs
- performance remains to be improved
- deal with cases can't tackle now