

Recursive read deadlocks and Where to find them

冯博群 Boqun Feng

Agenda

- Deadlock cases
- Lockdep
- Flavors of read/write locks
- More deadlock cases
- (Recursive) read deadlock detection

Deadlock cases

- self deadlock

```
P0  
  
spin_lock (&A) ;  
...  
spin_lock (&A) ;
```

- ABBA deadlock

P0	P1
<code>spin_lock (&A) ;</code>	<code>spin_lock (&B) ;</code>
<code>...</code>	<code>...</code>
<code>spin_lock (&B) ;</code>	<code>spin_lock (&A) ;</code>

Deadlock cases (cont.)

- IRQ safe->unsafe deadlocks
 - IRQs bring more "code combinations"

P0

```
<irq enabled>
spin_lock(&A);
...
<in irq handler>
  spin_lock(&A);
```

P0

```
<irq enabled>
spin_lock(&A);
...
<in irq handler>
  spin_lock(&B);
```

P1

```
<irq disabled>
spin_lock(&B);
...
spin_lock(&A);
```

Deadlock cases (cont.)

- ABCCA deadlocks
 - or more

P0	P1	P2
<code>spin_lock (&A) ;</code>	<code>spin_lock (&B) ;</code>	<code>spin_lock (&C) ;</code>
<code>...</code>	<code>...</code>	<code>...</code>
<code>spin_lock (&B) ;</code>	<code>spin_lock (&C) ;</code>	<code>spin_lock (&A) ;</code>

Lockdep

- Detect deadlock possibility
 - Assume the worst case: all the combinations of code sequences can happen
- Lock dependency
 - A -> B
 - Assume we can catch all dependencies




- Dependency graph

P0

```
spin_lock(&A);  
...  
spin_lock(&B);
```


Lockdep (cont.)

- Deadlock detection
 - A closed path (circle) in the dependency graph




Flavors of read/write locks

- Recursive/unfair rwlocks
 - readers are preferable



Flavors of read/write locks (cont.)

- Non-recursive/fair rwlocks



Flavors of read/write locks (cont.)

flavors	multiple readers	recursive c.s	a reader blocks another reader
recursive	Y	Y	N
non-recursive	Y	N	Y* (via a waiting writer)

Flavors of read/write locks (cont.)

- Block condition

- Recursive readers can get blocked by writers
- Non-recursive readers can get blocked by non-recursive readers (via a waiting writer) or writers

	reader(recursive or not)	writer
recursive reader	No	Yes
non-recursive(r & w)	Yes	Yes

More deadlock cases

- For non-recursive read/write locks
 - Same as spinlocks, since readers can block each other via a waiting writer

P0

`read_lock (&A) ;`

`...`

`spin_lock (&B) ;`

P1

`spin_lock (&B) ;`

`...`

`read_lock (&A) ;`

P2

`write_lock (&A) ;`

More deadlock cases

- For recursive locks, things get interesting:
 - This is not a deadlock

P0

```
read_lock (&A) ;
```

```
...
```

```
spin_lock (&B) ;
```

P1

```
spin_lock (&B) ;
```

```
...
```

```
read_lock (&A) ;
```



More deadlock cases

- But this is a deadlock

P0

```
read_lock(&A);
```

```
...
```

```
spin_lock(&B);
```

P1

```
spin_lock(&B);
```

```
...
```

```
write_lock(&A);
```




More deadlock cases

- Things get complicated when we mixed recursive and non-recursive read locks
- queued rwlock
 - non-recursive read lock in process context
 - recursive read lock in irq context

More deadlock cases

- Recursive deadlock case



P0	P1	P2
<code><in irq handler></code>		
<code>read_lock(&B);</code>	<code>spin_lock_irq(&A);</code>	
		<code>write_lock_irq(&B);</code>
<code>spin_lock(&A);</code>	<code>read_lock(&B);</code>	

More deadlock cases


- Recursive *not* deadlock case



P0	P1	P2
<code><in irq handler></code>		
<code>spin_lock(&A);</code>	<code>read_lock(&B);</code>	
<code>read_lock(&B);</code>	<code>spin_lock_irq(&A);</code>	<code>write_lock_irq(&B);</code>

Recursive read deadlock detection

- Limitation of current lockdep
 - circles mean deadlocks
 - while not all the circles mean deadlocks if we consider recursive readers.



Recursive read deadlock detection

- Goals

- Compatible with original lockdep detection.
- Handle qrwlock semantics.
- No false positive.

Recursive read deadlock detection

- Overview
 - Classification for lock dependencies
 - Definition of "strong" dependencies
 - Deadlock Condition

Classification of lock dependencies


- We used to treat all lock dependencies as the same
- but they are really not.
- $\{R \text{ reader, reader, writer}\} \rightarrow \{R \text{ reader, reader, writer}\} : 9$ combinations

Classification of locks

- (S)hared locks: reader (recursive or not)
- (E)clusive locks: writer (or plain spinlocks)
- (R)ecursive readers
- (N)on-recursive (readers and writers)

blocked by	reader(recursive or not)	writer
recursive reader	No	Yes
non-recursive(r & w)	Yes	Yes

blocked by	S	E
R	No	Yes
N	Yes	Yes



Classification of lock dependencies

- Groups things into 4
 - $-(SN)-> : \{R \text{ reader, reader}\} \rightarrow \{\text{reader, writer}\}$
 - $-(SR)-> : \{R \text{ reader, reader}\} \rightarrow \{R \text{ reader}\}$
 - $-(EN)-> : \{\text{writer}\} \rightarrow \{\text{reader, writer}\}$
 - $-(ER)-> : \{\text{writer}\} \rightarrow \{R \text{ reader}\}$
- Why? Because for a dependency $A \rightarrow B$, we care:
 - Whether A can block anyone
 - Whether B can get blocked by anyone



blocked by	S	E
R	No	Yes
N	Yes	Yes

P0

```
read_lock (&A) ;  
...  
spin_lock (&B) ;
```

P1

```
spin_lock (&B) ;  
...  
write_lock (&A) ;
```



Definition of "strong" dependencies

- Chaining lock dependencies via block conditions
- For dependencies $A \rightarrow B$ and $B \rightarrow C$
 - $A \rightarrow B \rightarrow C$ is a "strong" dependency path iff
 - $A \rightarrow B : -(R^*) \rightarrow$ and $B \rightarrow C : -(E^*) \rightarrow$, or
 - $A \rightarrow B : -(N^*) \rightarrow$ and $B \rightarrow C : -(S^*) \rightarrow$, or
 - $A \rightarrow B : -(N^*) \rightarrow$ and $B \rightarrow C : -(E^*) \rightarrow$
 - IOW, $-(R^*) \rightarrow -(S^*) \rightarrow$ will break the dependency
- works for " $A \rightarrow B, B \rightarrow C$ and $C \rightarrow D$ " case, and so on

blocked by	S	E
R	No	Yes
N	Yes	Yes

Deadlock condition

- A strong dependency chain/path forms a circle



P0

```
spin_lock(&A);  
...  
write_lock(&B);
```

P1

```
read_lock(&B);  
...  
write_lock(&C);
```

P2

```
read_lock(&C);  
...  
spin_lock(&A);
```

Implementation

- Extend `__bfs()` to walk on strong dependency path
- Make `LOCK*_STATE*` part of the chainkeys
- Add test cases
 - also unleash `irq_read_recursion2`
- Enable this for `srcu`
- Code
 - git.kernel.org/pub/scm/linux/kernel/git/boqun/linux.git arr-rfc-wip