

# Recursive read deadlocks and Where to find them

冯博群 Boqun Feng

[boqun.feng@gmail.com](mailto:boqun.feng@gmail.com)



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

# 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 (&amp;A) ;</code>	<code>spin_lock (&amp;B) ;</code>	<code>spin_lock (&amp;C) ;</code>
<code>...</code>	<code>...</code>	<code>...</code>
<code>spin_lock (&amp;B) ;</code>	<code>spin_lock (&amp;C) ;</code>	<code>spin_lock (&amp;A) ;</code>

# Lockdep


- Locks are grouped by classes
- Lock dependency
  - A -> B
- Dependency graph



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

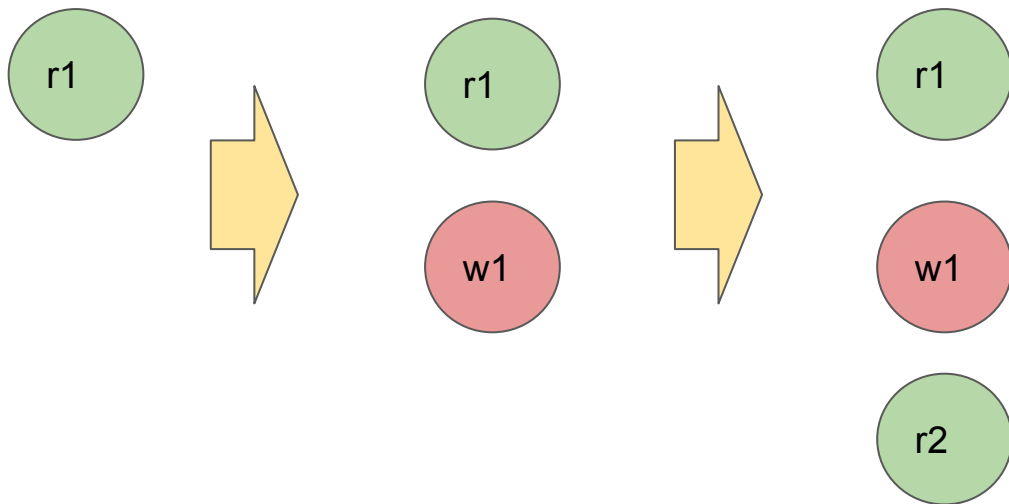
# Lockdep (cont.)

- 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	N	Y
non-recursive(r & w)	Y	Y

# 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	P1
<code>read_lock (&amp;A);</code>	<code>spin_lock (&amp;B);</code>
<code>...</code>	<code>...</code>
<code>spin_lock (&amp;B);</code>	<code>read_lock (&amp;A);</code>



# More deadlock cases

- But this is a deadlock

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




## 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>&lt;in irq handler&gt;</code>		
<code>read_lock(&amp;B);</code>	<code>spin_lock_irq(&amp;A);</code>	
<code>spin_lock(&amp;A);</code>	<code>read_lock(&amp;B);</code>	<code>write_lock_irq(&amp;B);</code>



# More deadlock cases


- Recursive *\*not\** deadlock case



P0	P1	P2
<code>&lt;in irq handler&gt;</code>		
<code>spin_lock(&amp;A);</code>	<code>read_lock(&amp;B);</code>	
<code>read_lock(&amp;B);</code>	<code>spin_lock_irq(&amp;A);</code>	<code>write_lock_irq(&amp;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
  - Informal Proof

# Classification of lock dependencies

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

# Classification of lock dependencies

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



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

# 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) ;
```



# Informal Proof

- We want to prove:
  - A strong dependency circle is equivalent to deadlock possibility
- Necessary condition
  - a strong dependency circle  $\Rightarrow$  deadlock possibility
  - Easy, because a strong dependency circle means we can build a combination of locking sequences that cause deadlock.

# Informal Proof (cont.)

- Sufficient condition
  - deadlock possibility  $\Rightarrow$  a strong dependency circle
  - My trick
    - deadlock possibility  $\Rightarrow$  circular wait (deadlock necessary condition according to wikipedia)
    - circular wait  $\Rightarrow$  a strong dependency circle

# 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](https://git.kernel.org/pub/scm/linux/kernel/git/boqun/linux.git) arr-rfc-wip