

Local Reasoning for RCU

Hongseok Yang (Queen Mary, Univ. of London)

Joint work with

Peter O'Hearn, Matthew Parkinson, Noam Rinetzky, Mooly Sagiv
and Viktor Vafeiadis

Preparing Multicore Era

- Active research by verification communities.
- Various workshops/conferences (E.g. EC^2).
- At least 9 CAV papers on concurrency.
- Concurrency is a hot topic inside separation logic.

Preparing Multicore Era

- Active research by verification communities.
- Various workshops/conferences (E.g. EC^2).
- At least 9 CAV papers on concurrency.
- Concurrency is a hot topic inside separation logic.
- What about Linux hackers?

We have demonstrated that RGSep and shape-value abstraction enable effective automatic linearizability proofs. The examples verified are typical of the research literature 5–10 years ago. The techniques can also cope with more complex al-

From “Shape-value abstraction for verifying linearizability” by Viktor Vafeiadis

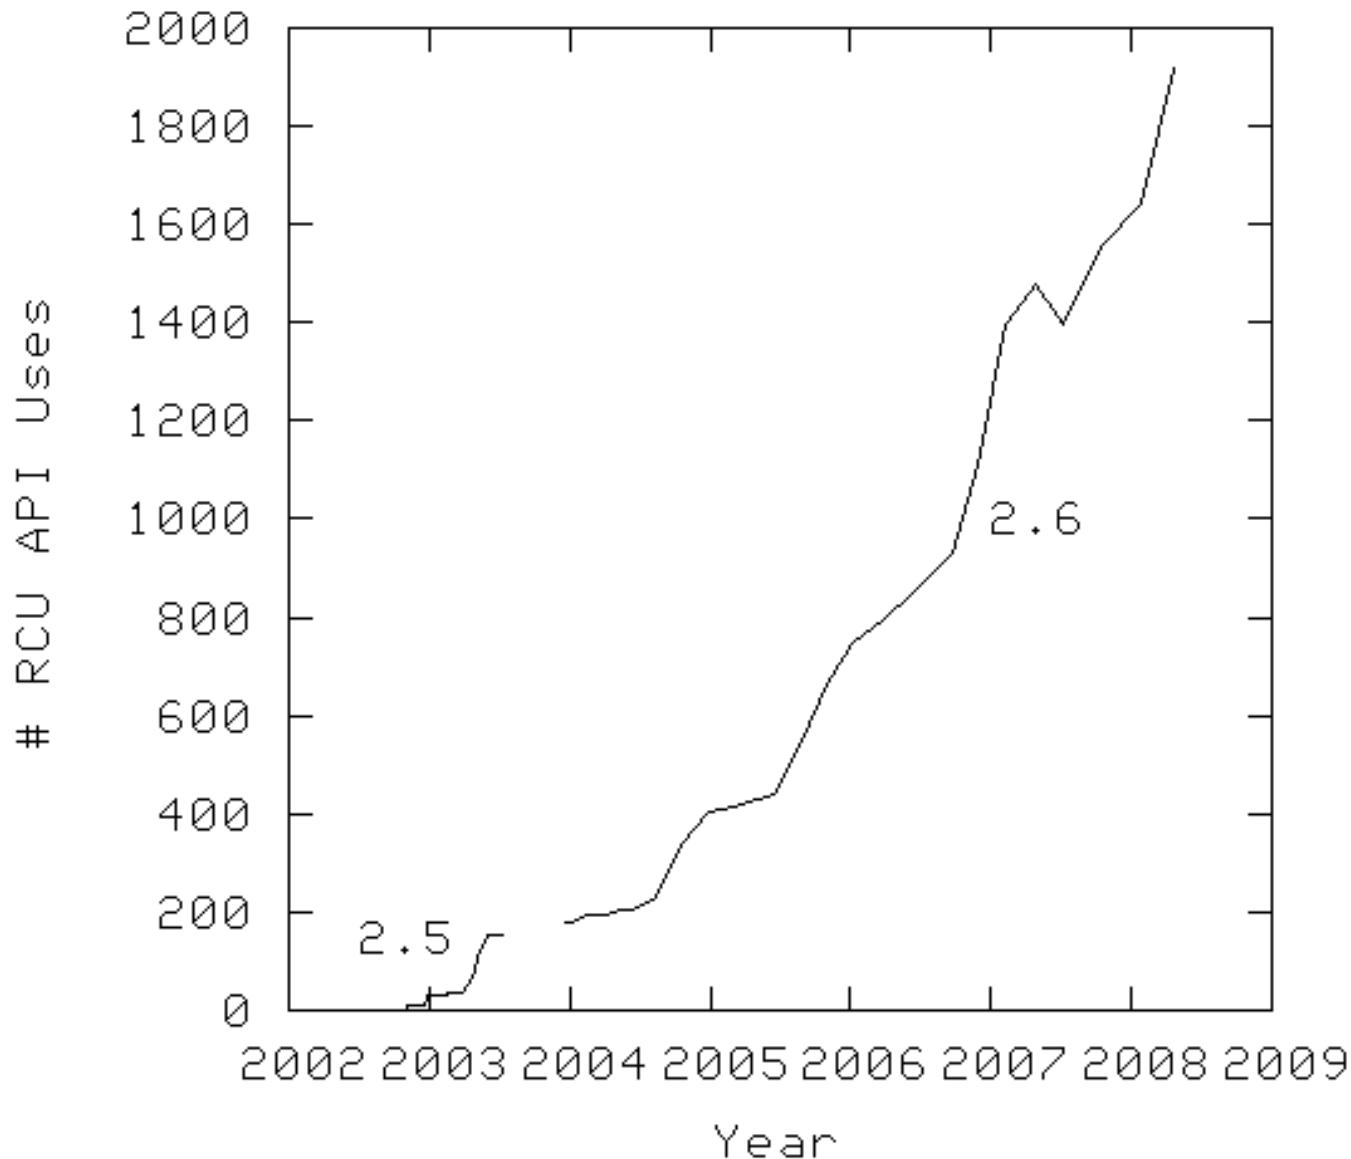
- Active research by verification communities.
- Various workshops/conferences (E.g. EC[^]2).
- At least 9 CAV papers on concurrency.
- Concurrency is a hot topic inside separation logic.
- What about Linux hackers?

We have demonstrated that RGSep and shape-value abstraction enable effective automatic linearizability proofs. The examples verified are typical of the research literature 5–10 years ago. The techniques can also cope with more complex al-

From “Shape-value abstraction for verifying linearizability” by Viktor Vafeiadis

- Active research by verification communities.
- Various workshops/conferences (E.g. EC[^]2).
- At least 9 CAV papers on concurrency.
- Concurrency is a hot topic inside separation logic.
- What about Linux hackers?

RCU Usage in Linux



Data from
Paul McKenney's
RCU Web Page

Objectives

- Explain RCU and research problems related to RCU.
- Present a preliminary result on program logic for RCU.

RCU (Read-Copy-Update)

- Back to the future: old style locking for new style architecture.
- Very approximately, reader/writer locks.
- Allows the concurrent access by multiple readers and a single writer.
- Almost zero concurrency overhead for readers.
- Intended to work well with weak memory.

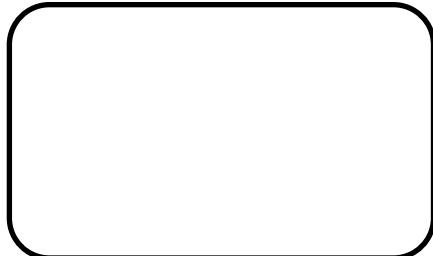
RCU Reader/ Writer

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
.. USE x,y ..
```

reader0



reader1

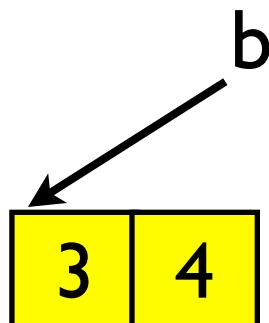


```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    barrier();  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

writer0



writer1



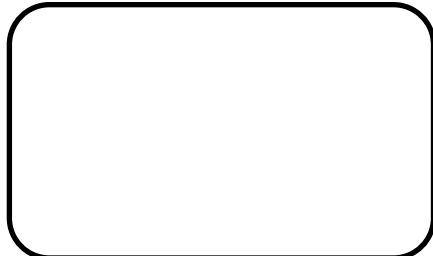
RCU Reader/ Writer

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
.. USE x,y ..
```

reader0



reader1

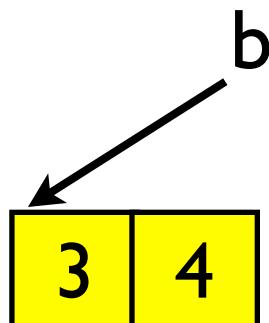


```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    barrier();  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

writer0



writer1

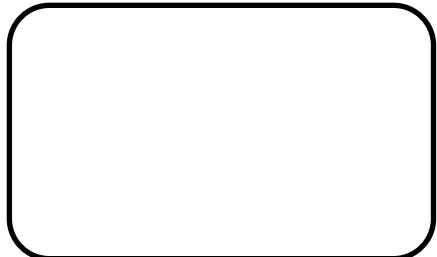


RCU Reader/ Writer

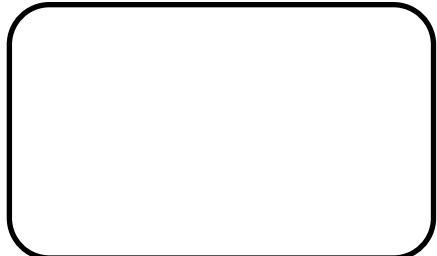
```
b' = b;  
  
x = *b';  
  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0



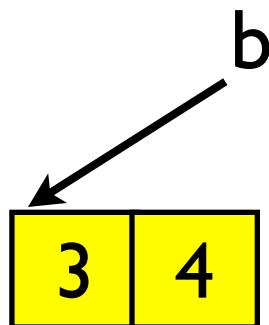
reader1



writer0



writer1



RCU Reader/ Writer

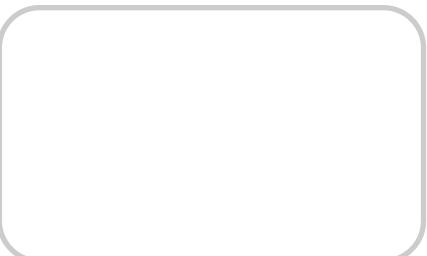
```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
... PRODUCE u,v ...  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0

x: , y:
b'

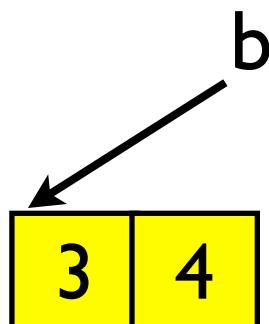
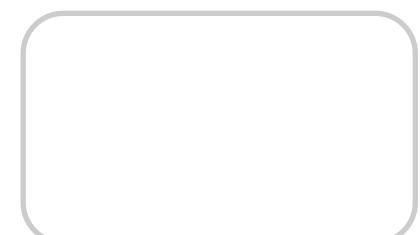
reader1



writer0



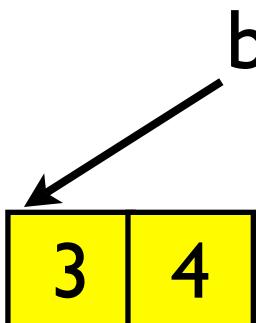
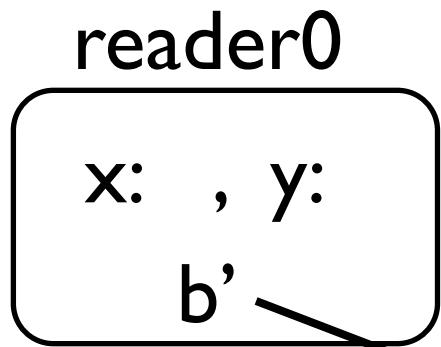
writer1



RCU Reader/ Writer

```
b' = b;  
  
x = *b';  
  
y = *(b'+1);  
  
... USE x,y ...
```

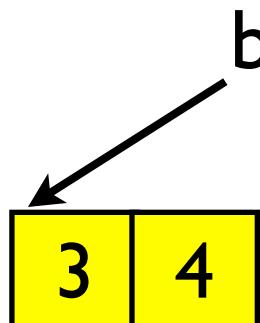
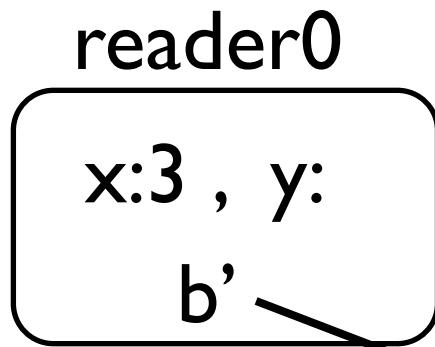
```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
.. USE x,y ..
```

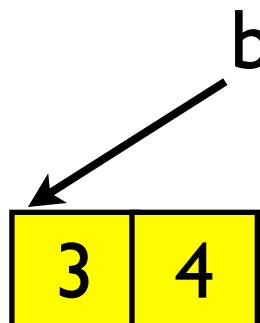
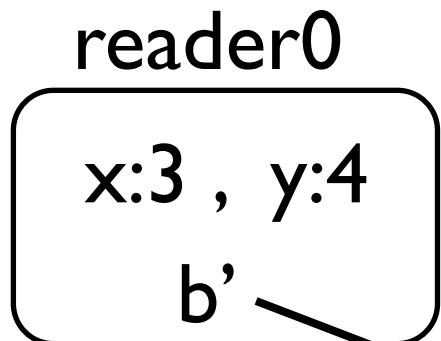
```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```



RCU Reader/ Writer

```
b' = b;  
  
x = *b';  
  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);
```

... USE x,y ...

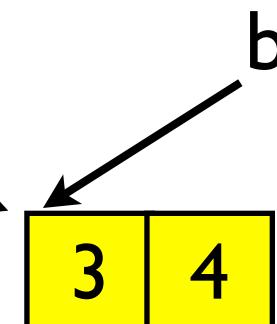
```
.. PRODUCE u,v ..  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0

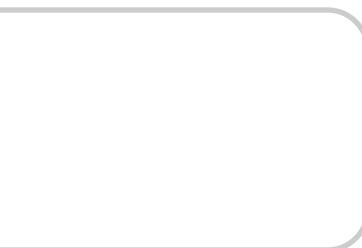
x:3 , y:4

b'

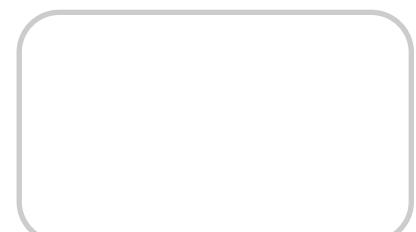
reader1



writer0



writer1



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
... PRODUCE u,v ...  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0

x:3 , y:4

b'

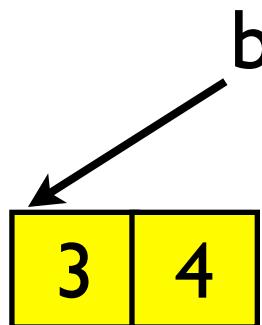
reader1

writer0

u:11, v:12

ob nb

writer1



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

.. PRODUCE u,v ..

nb = malloc(2);

*nb = u; *(nb+1) = v;

ob = b; b = nb;

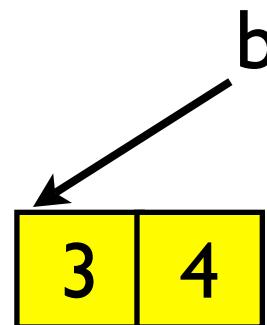
free(ob); free(ob+1);

reader0

x:3 , y:4

b'

reader1



writer0

u:11, v:12

ob nb

writer1



RCU Reader/ Writer

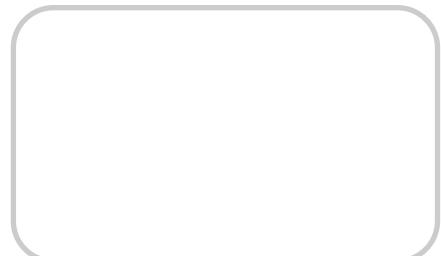
```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0

x:3 , y:4
b'

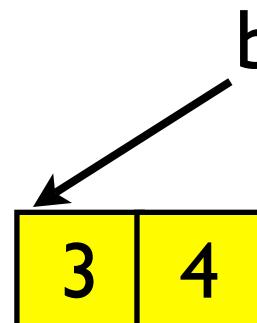
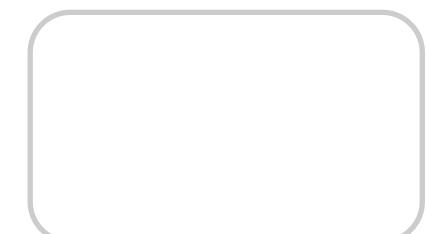
reader1



writer0

u:11, v:12
ob nb

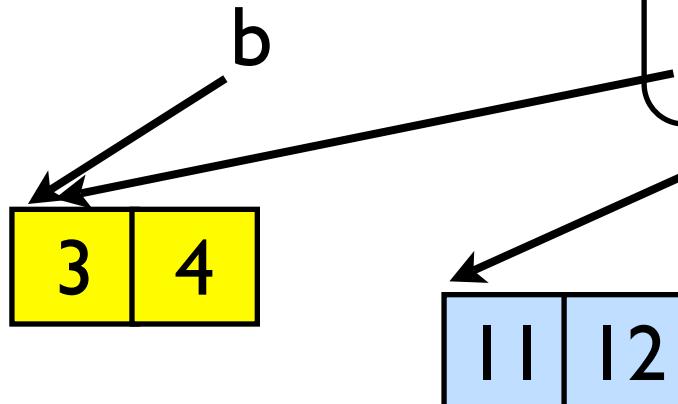
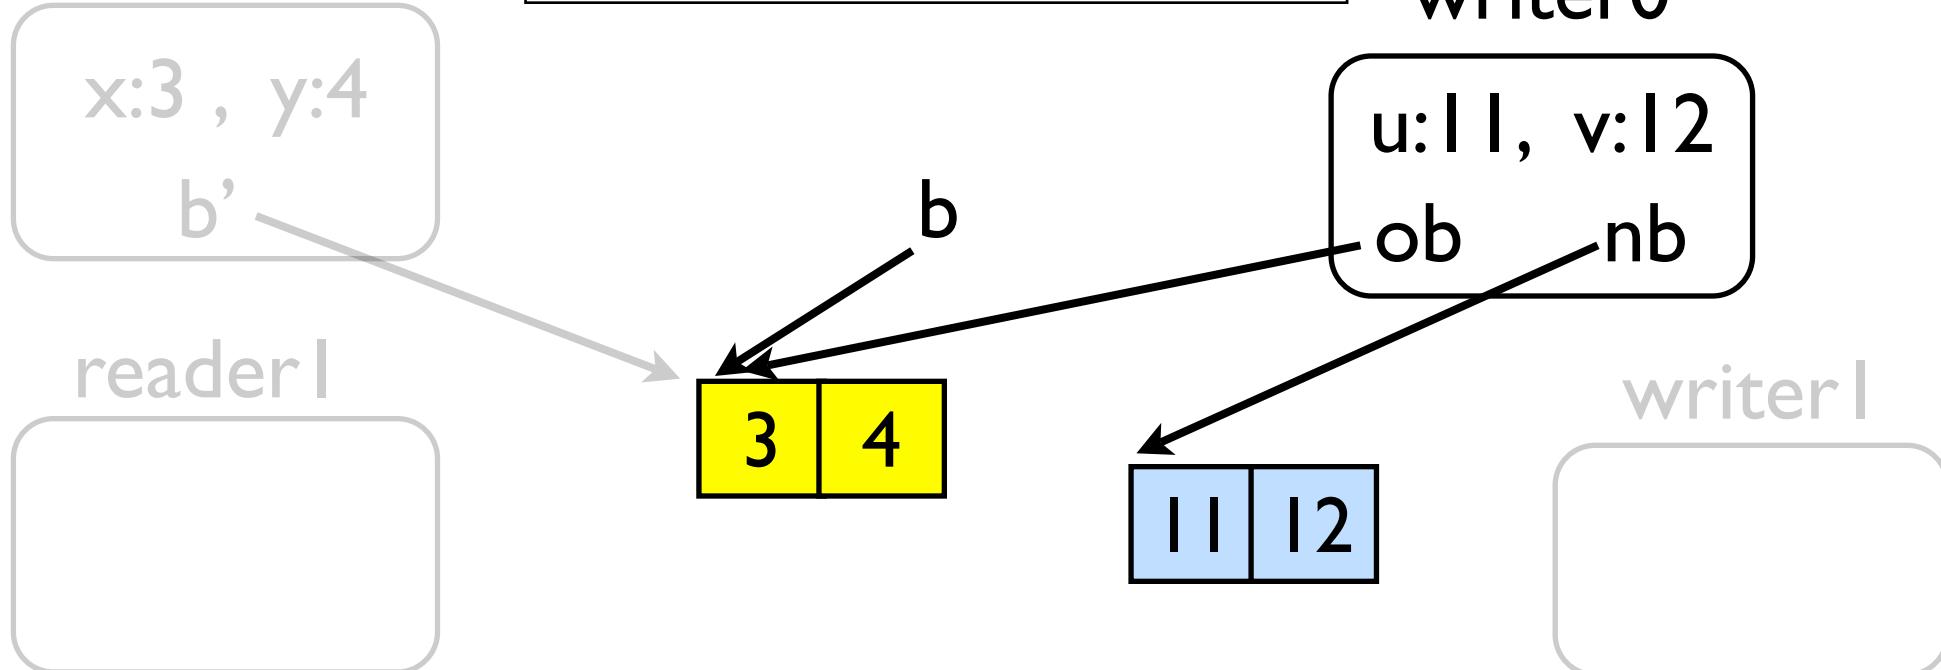
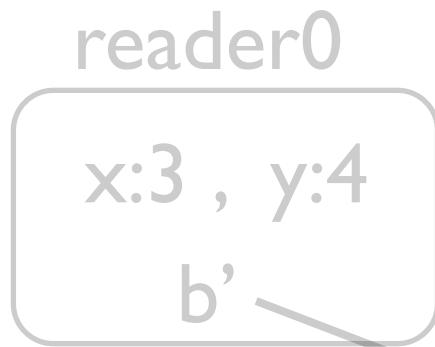
writer1



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0

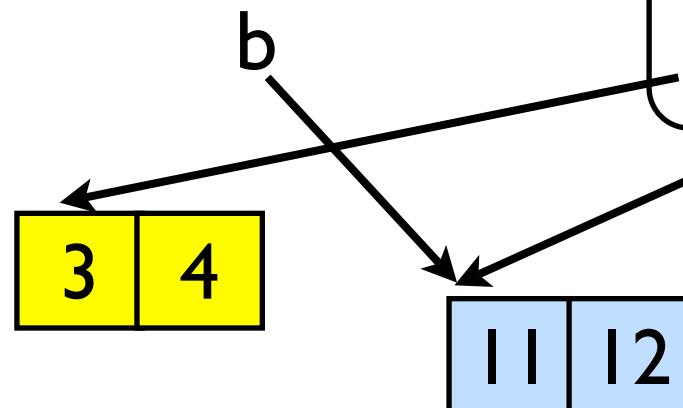
x:3 , y:4
b'

reader1

writer0

u:11, v:12
ob nb

writer1



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0

x:3 , y:4
b'

reader1

writer0

u:l1, v:l2
ob nb

writer1



b



RCU Reader/ Writer

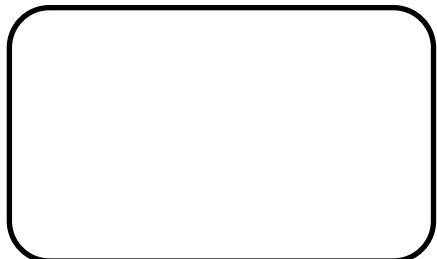
```
b' = b;  
  
x = *b';  
  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

reader0



reader1



writer0



writer1



RCU Reader/ Writer

```
b' = b;  
  
x = *b';  
  
y = *(b'+1);  
  
... USE x,y ...
```

reader0

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
*nb = u; *(nb+1) = v;  
  
ob = b; b = nb;  
  
free(ob); free(ob+1);
```

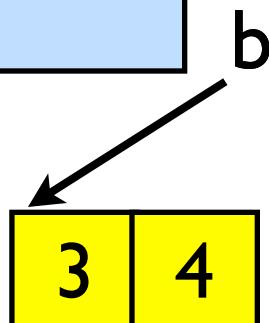
writer0

Pb1: Racy writers.

reader1



writer1



RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

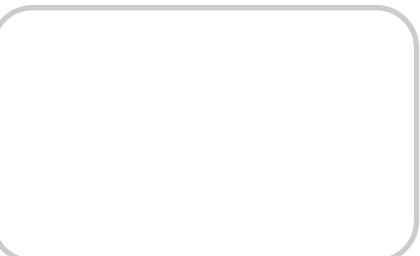
```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
    free(ob); free(ob+1); }
```

reader0

writer0

Pbl: Racy writers.
Sol: Use a standard lock.

reader1



b



writer1



RCU Reader/ Writer

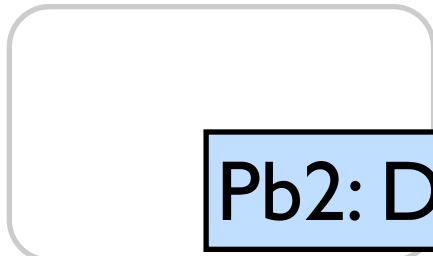
```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
    free(ob); free(ob+1); }
```

reader0



reader1

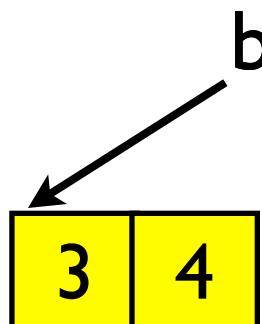


writer0



writer1

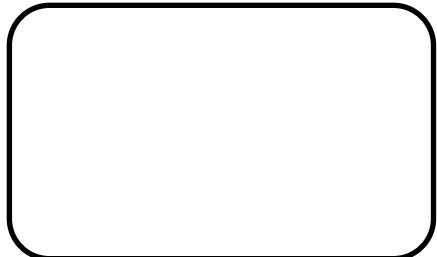
Pb2: Dangling pointer dereference by readers.



RCU Reader/ Writer

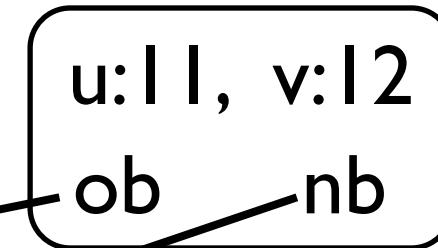
```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

reader0

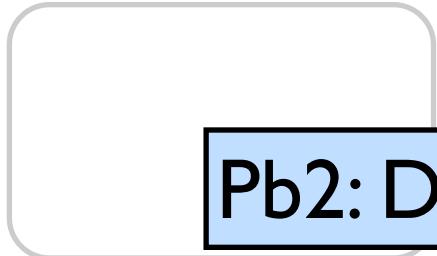


```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
    free(ob); free(ob+1); }
```

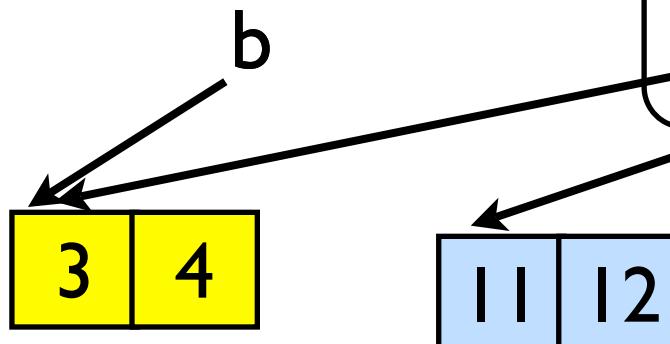
writer0



reader1



writer1



Pb2: Dangling pointer dereference by readers.

RCU Reader/ Writer

```
b' = b;  
x = *b,  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
    free(ob); free(ob+1); }
```

reader0

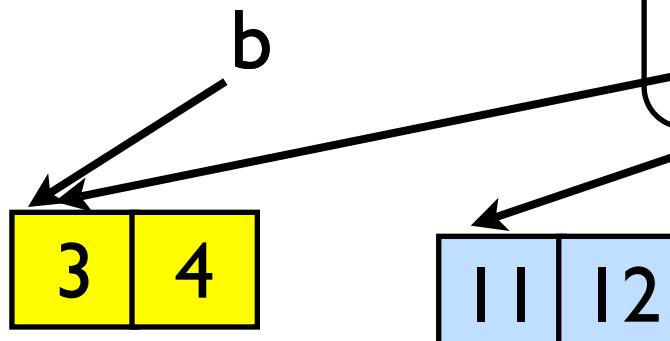
x: , y:
b'

reader1

writer0

u:l1, v:l2
ob nb

writer1



Pb2: Dangling pointer dereference by readers.

RCU Reader/ Writer

```
b' = b;  
x = *b,  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
    free(ob); free(ob+1); }
```

reader0

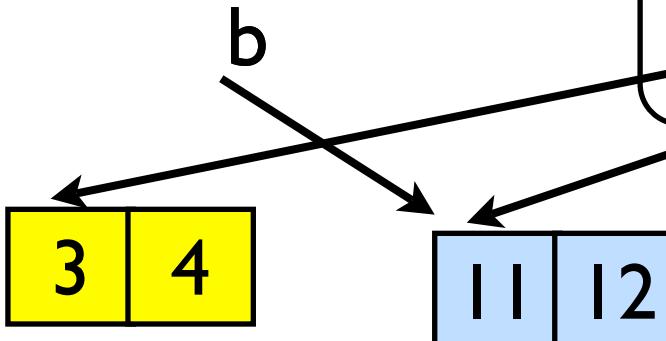
x: , y:
b'

writer0

u:11, v:12
ob nb

reader1

writer1



Pb2: Dangling pointer dereference by readers.

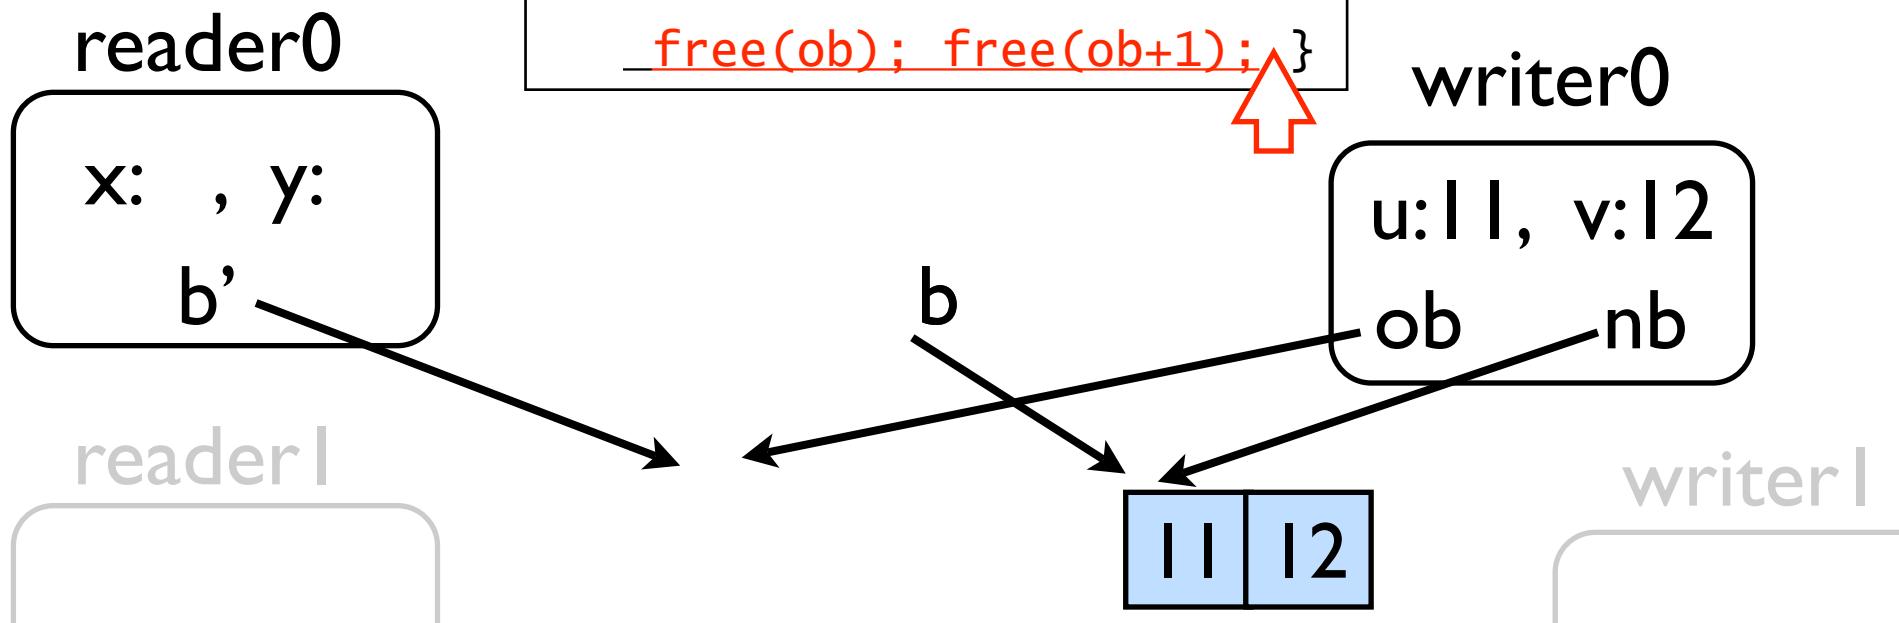
RCU Reader/ Writer

```
b' = b;  
x = *b,  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;
```

```
ob = b; b = nb;
```

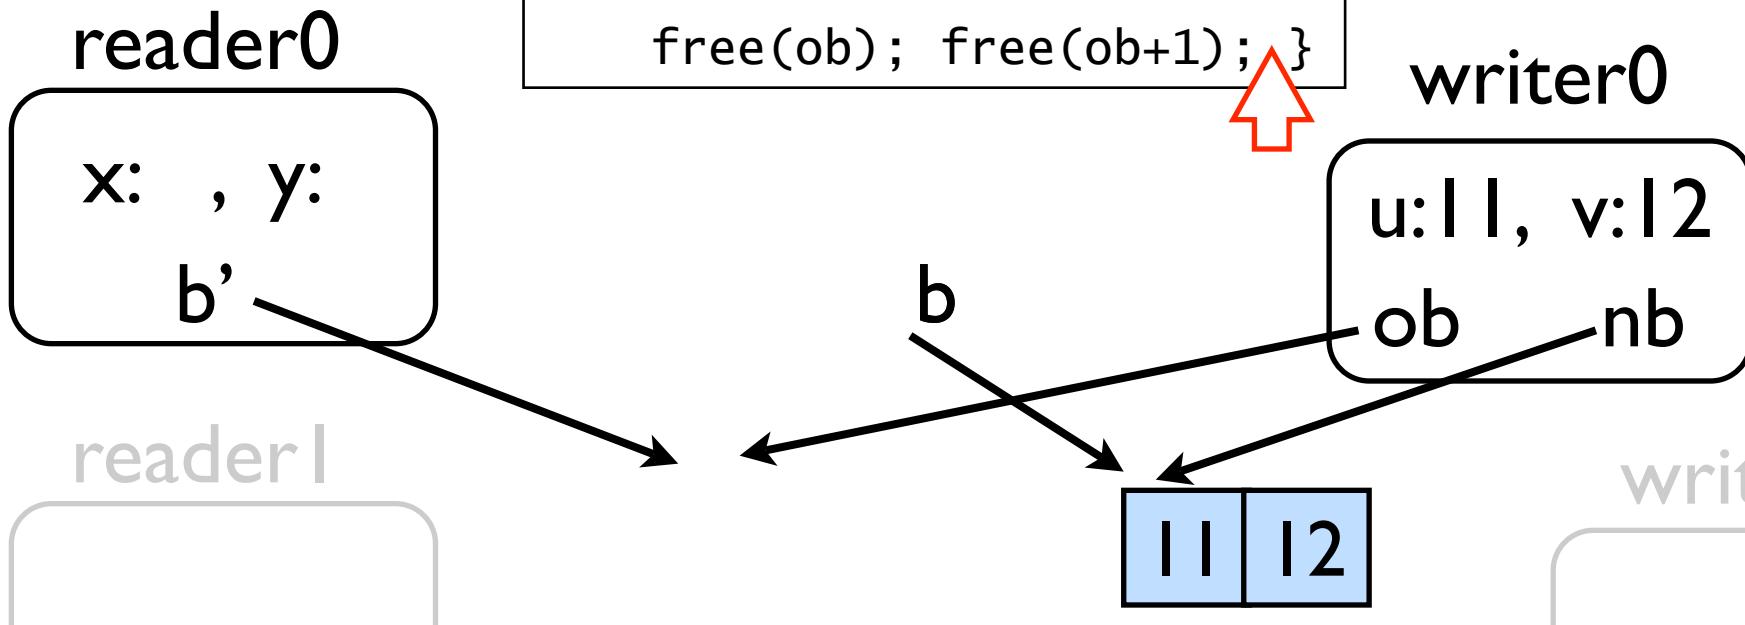
```
free(ob); free(ob+1); }
```



Pb2: Dangling pointer dereference by readers.

RCU Reader/ Writer

```
b' = b;  
x = *b';  
y = *(b'+1);  
... PRODUCE u,v ...  
nb = malloc(2);  
{  
    b+1) = v;  
... USE x,y ...  
ob = b; b = nb;  
free(ob); free(ob+1); }
```



Pb2: Dangling pointer dereference by readers.

RCU Reader/ Writer

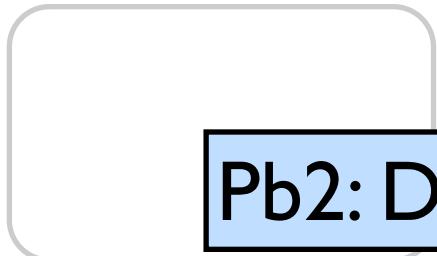
```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
    free(ob); free(ob+1); }
```

reader0



reader1

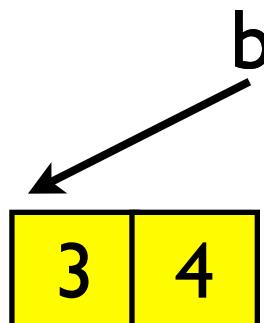


writer0



writer1

Pb2: Dangling pointer dereference by readers.



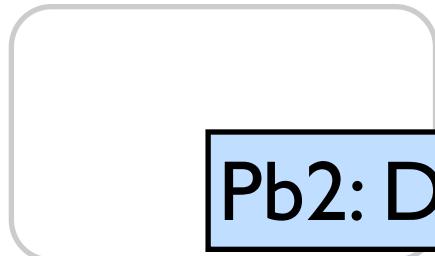
RCU Reader/

```
b' = b;  
x = *b';  
y = *(b'+1);  
  
... USE x,y ...
```

reader0



reader1



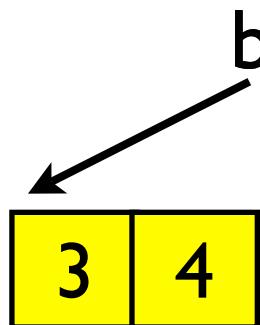
```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

Wait until no readers are accessing the old buffer.

writer0

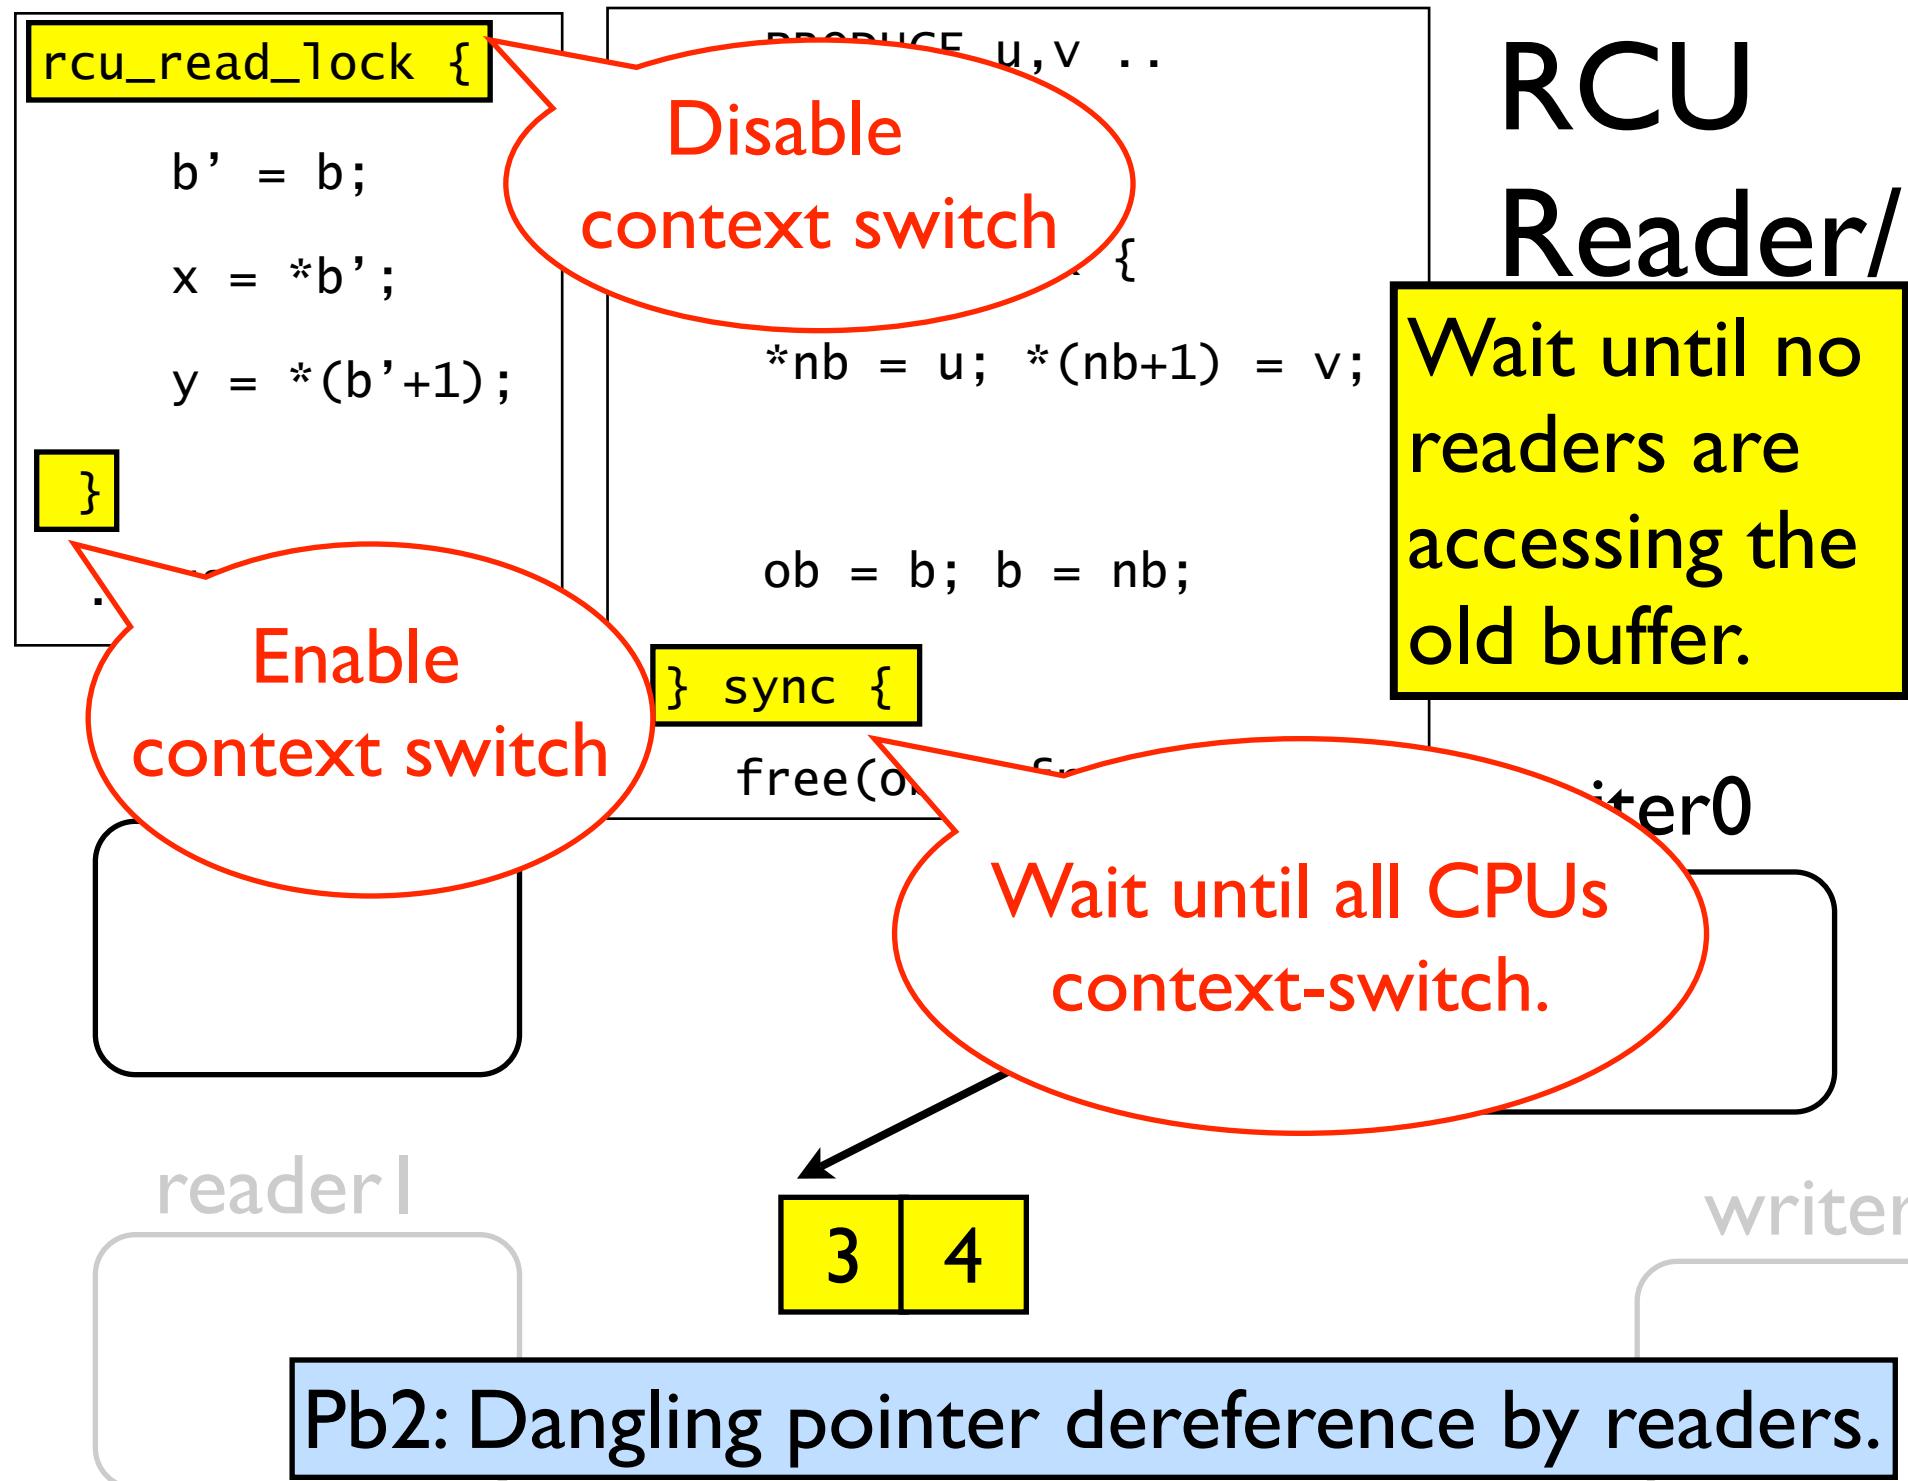


writer1



Pb2: Dangling pointer dereference by readers.

RCU Reader/



RCU Reader/

Wait until no readers are accessing the old buffer.

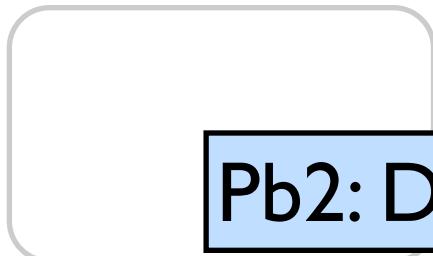
```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
.. USE x,y ..
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

reader0



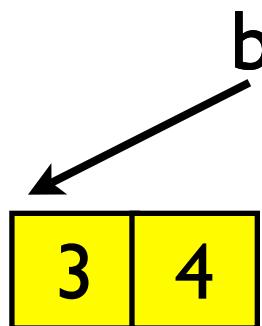
reader1



writer0



writer1



Pb2: Dangling pointer dereference by readers.

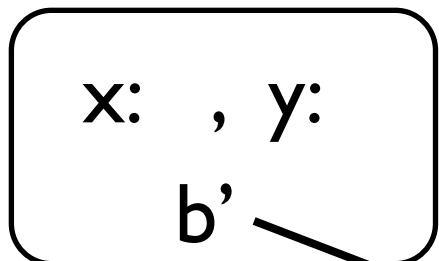
RCU Reader/

Wait until no readers are accessing the old buffer.

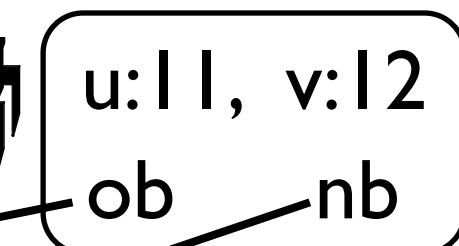
```
rcu_read_lock {  
  
    b' = b;  
    x = *b',  
    y = *(b'+1);  
  
}  
  
.. USE x,y ..
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

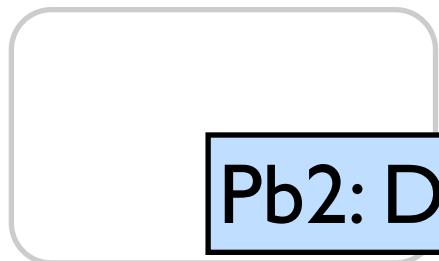
reader0



writer0



reader1



writer1



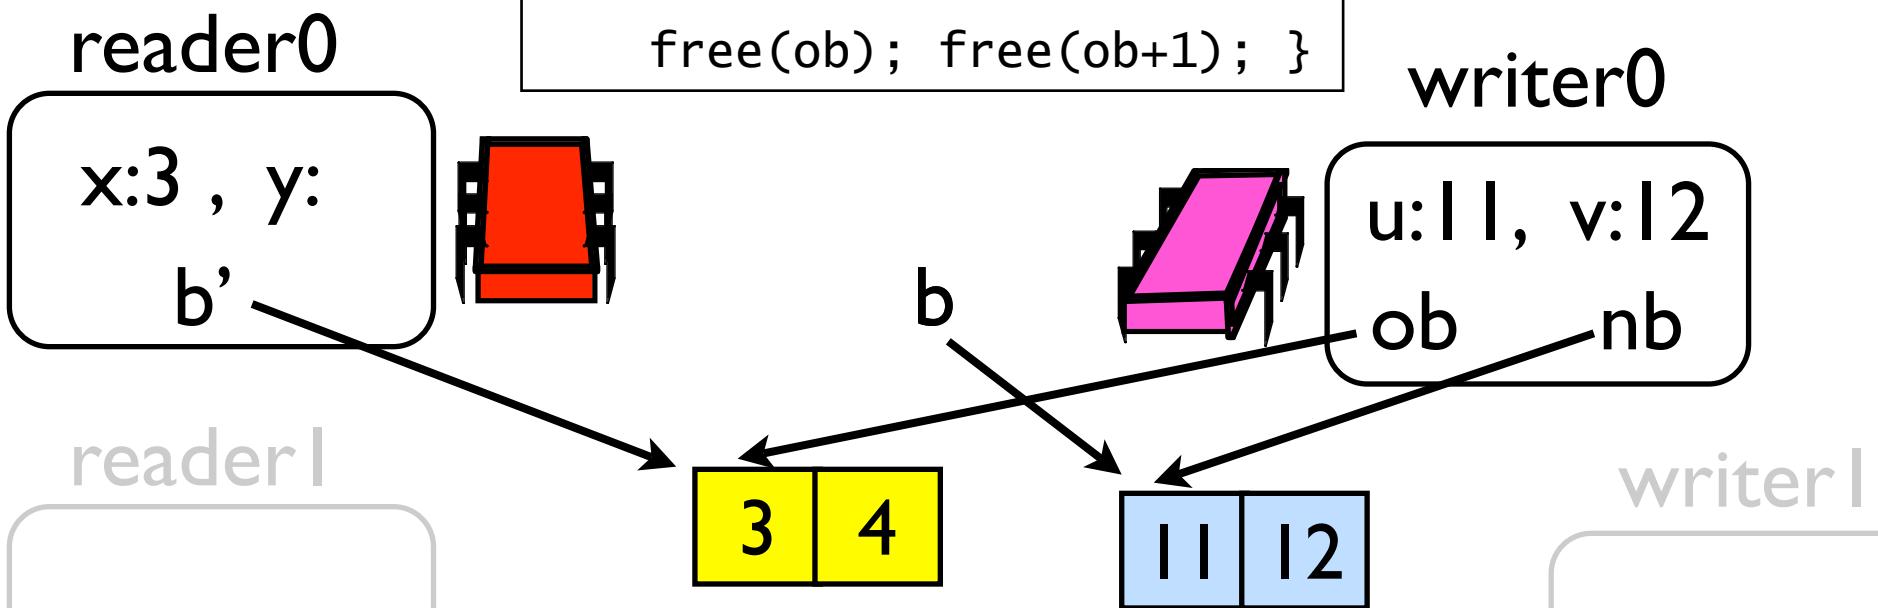
Pb2: Dangling pointer dereference by readers.

RCU Reader/

Wait until no readers are accessing the old buffer.

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b'; // Red arrow points here.  
  
    y = *(b'+1);  
  
}  
  
.. USE x,y ..
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb; // Red arrow points here.  
  
} sync {  
  
    free(ob); free(ob+1); }
```



Pb2: Dangling pointer dereference by readers.

RCU Reader/

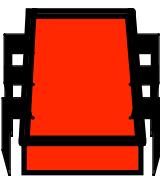
Wait until no readers are accessing the old buffer.

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
}  
  
}  
  
.. USE x,y ..
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

reader0

x:3 , y:4
b'



reader1

writer0

u:11, v:12
ob nb



writer1

Pb2: Dangling pointer dereference by readers.



RCU Reader/

Wait until no readers are accessing the old buffer.

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
... USE x,y ...
```

reader0

x:3 , y:4
b'

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

writer0

u:11, v:12
ob nb

reader1

b



writer1

Pb2: Dangling pointer dereference by readers.

RCU Reader/

Wait until no readers are accessing the old buffer.

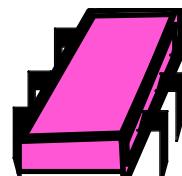
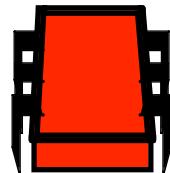
```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
... USE x,y ...
```

```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

writer0

reader0

x:3 , y:4
b'



u:l1, v:l2
ob nb



writer1

Pb2: Dangling pointer dereference by readers.

RCU Reader/

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
... USE x,y ...
```

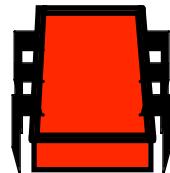
```
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1);}
```

Wait until no readers are accessing the old buffer.

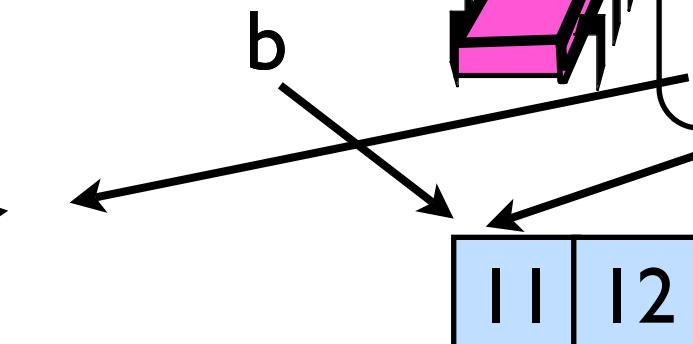
writer0

reader0

x:3 , y:4
b'



u:l1, v:l2
ob nb



writer1

Pb2: Dangling pointer dereference by readers.

RCU Reader/ Writer

```
rcu_read_lock {  
    b' = b;  
    x = *b';  
    y = *(b'+1);  
}  
.. USE x,y ..
```

```
.. PRODUCE u,v ..  
nb = malloc(2);  
rcu_write_lock {  
    *nb = u; *(nb+1) = v;  
    ob = b; b = nb;  
} sync {  
    free(ob); free(ob+1); }
```

RCU Reader/ Writer

```
rcu_read_lock {  
    b' = b;  
    x = *b';  
    y = *(b'+1);  
}  
... USE x,y ...
```

```
... PRODUCE u,v ...  
nb = malloc(2);  
rcu_write_lock {  
    *nb = u; *(nb+1) = v;  
    barrier();  
    ob = b; b = nb;  
} sync {  
    free(ob); free(ob+1); }
```

- 1) Handles weak memory.
- 2) Ensures that a reader never reads values from a non-initialized buffer.

RCU Reader/ Writer

```
rcu_read_lock {  
  
    b' = b;  
  
    x = *b';  
  
    y = *(b'+1);  
  
}  
  
.. USE x,y ..  
  
.. PRODUCE u,v ..  
  
nb = malloc(2);  
  
rcu_write_lock {  
  
    *nb = u; *(nb+1) = v;  
  
    barrier();  
  
    ob = b; b = nb;  
  
} sync {  
  
    free(ob); free(ob+1); }
```

- RCU constructs:

```
rcu_read_lock { .. },  
  
rcu_write_lock { ... } sync { ... }, barrier
```

- Linux API also includes asynchronous sync and

Research Issues

- Full specification of RCU programs.

Research Issues

- Full specification of RCU programs.
- Program logic for RCU.

Research Issues

- Full specification of RCU programs.
- Program logic for RCU.
- Automatic program analysis.
 - Shape analysis, termination analysis.

Research Issues

- Full specification of RCU programs.
- Program logic for RCU.
- Automatic program analysis.
 - Shape analysis, termination analysis.
- Verify the implementations of RCU.
 - Implementations of RCU, SRCU.
 - Prove correctness considering weak memory.

Research Issues

- Full specification of RCU programs.
- Program logic for RCU.
- Automatic program analysis.
 - Shape analysis, termination analysis.
- Verify the implementations of RCU.
 - Implementations of RCU, SRCU.
 - Prove correctness considering weak memory.
- New algorithms based on RCU.

Research Issues

- Full specification of RCU programs.
- Program logic for RCU.
- Automatic program analysis.
 - Shape analysis, termination analysis.
- Verify the implementations of RCU.
 - Implementations of RCU, SRCU.
 - Prove correctness considering weak memory.
- New algorithms based on RCU.

Overview of Logic

- Resource-invariant-based, thread-local reasoning.
- sync, barrier as ownership transfer operations.
- Reasoning about readers should use only those facts preserved by writers' operations.

Proof Sketch of RCU Readers/Writers

rcu_read_lock {

$b' = b;$

$x = *b';$

$y = *(b' + 1);$

}

rcu_write_lock {

$*nb = u; *(nb + 1) = v;$

barrier $_{nb \mapsto _,_};$

$ob = b: b = nb:$

} sync {

free(ob); **free**($ob+1$);

}

Proof Sketch of RCU Readers/Writers

rcu_read_lock {

$b' = b;$

$x = *b';$

$y = *(b' + 1);$

}

rcu_write_lock {

$*nb = u; *(nb + 1) = v;$

barrier $_{nb \mapsto _, _};$

$ob = b: b = nb:$

} sync {

free(ob); free($ob+1$);

}

Invariant-based, thread-local reasoning.

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

$$b' = b;$$

$$x = *b';$$

$$y = *(b' + 1);$$

}

{emp}

{ $nb \mapsto _, __$ }

rcu_write_lock {

$$*nb = u; \ (*nb + 1) = v;$$

barrier $_{nb \mapsto _, __}$;

$$ob = b: \ b = nb:$$

} sync {

free(ob); free($ob+1$);

}

{emp}

Invariant-based, thread-local reasoning.

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

ResInv holds initially for
the shared heap.

}

{emp}

{nb \mapsto -, -}

rcu_write_lock {

{nb \mapsto -, - * [b \mapsto -, -]}

*nb = u; *(nb + 1) = v;

barrier_{nb \mapsto -, -};

ob = b: b = nb:

} sync {

free(ob); free(ob+1);

}

{emp}

ResInv : b \mapsto -, -

Invariant-based, thread-local reasoning.

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

$b' = b;$

$x = *b';$

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * [b \mapsto _, _]$ }

$*nb = u; *(nb + 1) = v;$

ResInv : $b \mapsto _, _$

barrier $_{nb \mapsto _, _};$

$ob = b; b = nb;$

sync {

free(ob); free($ob + 1$);

{emp * [$b \mapsto _, _$]}

}

{emp}

y = $*^{(nb + 1)}$
ResInv holds
right before releasing
the lock.

{e}

Invariant-based, thread-local reasoning.

Proof Sketch of RCU Readers/Writers

{emp}
rcu_read_lock {

$b' = b;$

$x = *b';$

}

{emp}

$y = *(b' + 1);$

{ $nb \mapsto _, _$ }
rcu_write_lock {

{ $nb \mapsto _, _ * [b \mapsto _, _]$ }
 $*nb = u; *(nb + 1) = v;$
{ $nb \mapsto _, _ * [b \mapsto _, _]$ }
barrier $_{nb \mapsto _, _};$

$ob = b: b = nb:$

} sync {

free(ob); free($ob+1$);
{emp * $[b \mapsto _, _]$ }
}
{emp}

Proof Sketch of RCU Readers/Writers

```
{emp}  
rcu_read_lock {
```

$b' = b;$

```
{nb $\mapsto$ -, -}  
rcu_write_lock {
```

$\{nb\mapsto_, _ * [b\mapsto_, _]\}$

$*nb = u; *(nb + 1) = v;$

$\{nb\mapsto_, _ * [b\mapsto_, _]\}$

barrier $_{nb\mapsto_, _};$

$\{emp * [nb\mapsto_, _ * b\mapsto_, _]\}$

$ob = b; b = nb;$

} sync {

free(ob); free($ob + 1$);

$\{emp * [b\mapsto_, _]\}$

}

{emp}

ResInv : $b\mapsto_, _$

barrier as ownership
transfer from private to shared

Proof Sketch of RCU Readers/Writers

```
{emp}
rcu_read_lock {
    b' = b;
    x = *b';
    y = *(b' + 1);
}

{emp}
```

```
{nb $\rightarrow$ -, -}
rcu_write_lock {
    {nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    *nb = u; *(nb + 1) = v;
    {nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    barriernb $\rightarrow$ -, -;
    {emp * nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    ob = b; b = nb;
    {emp * b $\rightarrow$ -, - * ob $\rightarrow$ -, -}
} sync {
    free(ob); free(ob+1);
    {emp * b $\rightarrow$ -, -}
}

{emp}
```

ResInv : $b\rightarrow_-, -$

Proof Sketch of RCU Readers/Writers

```
{emp}  
rcu_read_lock {
```

```
b' = b;
```

```
x = *b';
```

```
y = *(b' + 1);
```

```
}
```

```
{nb $\rightarrow$ -, -}  
rcu_write_lock {
```

```
{nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
```

```
*nb = u; *(nb + 1) = v;
```

```
{nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
```

```
barriernb $\rightarrow$ -, -;
```

```
{emp * nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
```

```
ob = b; b = nb;
```

```
{emp * b $\rightarrow$ -, - * ob $\rightarrow$ -, -}
```

```
} sync {
```

```
{ob $\rightarrow$ -, - * b $\rightarrow$ -, -}
```

```
free(ob); free(ob+1);
```

```
{emp * b $\rightarrow$ -, -}
```

```
} {emp}
```

ResInv : $b \rightarrow _, _$

sync as ownership transfer
from shared to private

Proof Sketch of RCU Readers/Writers

```
{emp}
rcu_read_lock {
    b' = b;
    x = *b';
    y = *(b' + 1);
}

{emp}
```

```
{nb $\rightarrow$ -, -}
rcu_write_lock {
    {nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    *nb = u; *(nb + 1) = v;
    {nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    barriernb $\rightarrow$ -, -;
    {emp * nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    ob = b; b = nb;
    {emp * b $\rightarrow$ -, - * ob $\rightarrow$ -, -}
} sync {
    {ob $\rightarrow$ -, - * b $\rightarrow$ -, -}
    free(ob); free(ob+1);
    {emp * b $\rightarrow$ -, -}
}

{emp}
```

ResInv : $b\rightarrow_-, -$

Proof Sketch of RCU Readers/Writers

```
{emp}
rcu_read_lock {
    b' = b;
    x = *b';
    y = *(b' + 1);
}

{emp}
```

```
{nb $\rightarrow$ -, -}
rcu_write_lock {
    {nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    *nb = u; *(nb + 1) = v;
    {nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    barriernb $\rightarrow$ -, -;
    {emp * nb $\rightarrow$ -, - * b $\rightarrow$ -, -}
    ob = b; b = nb;
    {emp * b $\rightarrow$ -, - * ob $\rightarrow$ -, -}
} sync {
    {ob $\rightarrow$ -, - * b $\rightarrow$ -, -}
    free(ob); free(ob+1);
    {emp * b $\rightarrow$ -, -}
}

{emp}
```

ResInv : $b\rightarrow_, _$

AlwaysResInv : $b\rightarrow_, _ * \text{true}$

Proof Sketch of RCU Readers/Writers

```

{emp}
rcu_read_lock {
    b' = b;
    x = *b';
    y = *(b' + 1);
}

{emp}

```

```

{nb ↦ -, -}
rcu_write_lock {
    {nb ↦ -, - * [b ↦ -, -]}
    *nb = u; *(nb + 1) = v;
    {nb ↦ -, - * [b ↦ -, -]}
    barriernb ↦ -, -;
    {emp * [nb ↦ -, - * b ↦ -, -]}
    ob = b; b = nb;
    {emp * [b ↦ -, - * ob ↦ -, -]}
} sync {
    {ob ↦ -, - * [b ↦ -, -]}
    free(ob); free(ob+1);
    {emp * [b ↦ -, -]}
}

{emp}

```

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. ([a \mapsto _, _ * \text{true}] \rightsquigarrow [a \mapsto _, _ * \text{true}])$

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

{emp * $b \mapsto _, _ * \text{true}$ }

$b' = b;$

$x = *b';$

$y = *(b' + 1);$

}

{emp}

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$*nb = u; *(nb + 1) = v;$

* $b \mapsto _, _$

$\vdash nb \mapsto _, _;$

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$ob = b; b = nb;$

{emp * $b \mapsto _, _ * ob \mapsto _, _$ }

} sync {

{ $ob \mapsto _, _ * b \mapsto _, _$ }

free(ob); free($ob + 1$);

{emp * $b \mapsto _, _$ }

}

{emp}

ResInv : $b \mapsto _, _$

Start with
AlwaysResInv.

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. ([a \mapsto _, _ * \text{true}] \rightsquigarrow [a \mapsto _, _ * \text{true}])$

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

{emp * $b \mapsto _, _ * \text{true}$ }

$b' = b;$

$x = *b';$

$y = *(b' + 1);$

{emp * $b' \mapsto _, _ * \text{true}$ }

}

{emp}

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$*nb = u; *(nb + 1) = v;$

{ $nb \mapsto _, _ * b \mapsto _, _$ }

barrier $_{nb \mapsto _, _};$

{emp * $nb \mapsto _, _ * b \mapsto _, _$ }

$ob = b; b = nb;$

{emp * $b \mapsto _, _ * ob \mapsto _, _$ }

$\lambda \text{ sync} \{$

$\mapsto _, _ * b \mapsto _, _$

$(ob); \text{free}(ob + 1);$

$b \mapsto _, _ \}$

ResInv : $b \mapsto _, _$

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. (a \mapsto _, _ * \text{true}) \rightsquigarrow (a \mapsto _, _ * \text{true})$

Can end with
any assertion.

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

{emp * $b \mapsto _, _ * \text{true}$ }

$b' = b;$

{emp * $b' \mapsto _, _ * \text{true}$ }

$x = *b';$

$y = *(b' + 1);$

{emp * $b' \mapsto _, _ * \text{true}$ }

}

{emp}

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$*nb = u; *(nb + 1) = v;$

{ $nb \mapsto _, _ * b \mapsto _, _$ }

barrier $_{nb \mapsto _, _};$

{emp * $nb \mapsto _, _ * b \mapsto _, _$ }

$ob = b; b = nb;$

{emp * $b \mapsto _, _ * ob \mapsto _, _$ }

} sync {

{ $ob \mapsto _, _ * b \mapsto _, _$ }

free(ob); free($ob + 1$);

{emp * $b \mapsto _, _$ }

}

{emp}

ResInv : $b \mapsto _, _$

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. (a \mapsto _, _ * \text{true} \rightsquigarrow a \mapsto _, _ * \text{true})$

Proof Sketch of RCU

```

{emp}
rcu_read_lock {
    {emp *  $b \mapsto \_, \_ * \text{true}$ }  
     $b' = b;$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }  
     $x = *b';$ 
     $y = *(b' + 1);$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }  
}
{emp}

```

Not modify the shared heap.

```

 $*nb = u; *(nb + 1) = v;$ 
{ $nb \mapsto \_, \_ * b \mapsto \_, \_$ }  
barrier $_{nb \mapsto \_, \_};$ 
{emp *  $nb \mapsto \_, \_ * b \mapsto \_, \_$ }  
 $ob = b; b = nb;$ 
{emp *  $b \mapsto \_, \_ * ob \mapsto \_, \_$ }  
} sync {
    { $ob \mapsto \_, \_ * b \mapsto \_, \_$ }  
    free( $ob$ ); free( $ob + 1$ );
    {emp *  $b \mapsto \_, \_$ }  
}
{emp} AlwaysResInv :  $b \mapsto \_, \_ * \text{true}$ 

```

Pres : $\forall a. ([a \mapsto _, _ * \text{true}] \rightsquigarrow [a \mapsto _, _ * \text{true}])$

Proof Sketch of RCU

```

{emp}
rcu_read_lock {
    {emp *  $b \mapsto \_, \_ * \text{true}$ }
     $b' = b;$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }
     $x = *b';$ 
     $y = *(b' + 1);$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }
}
{emp}

```

Not modify the shared heap.

```

 $\dots$ 
 $*nb = u; *(nb + 1) = v;$ 
{ $nb \mapsto \_, \_ * [b \mapsto \_, \_]$ }
 $\dots$ 
 $\{ob \mapsto \_, \_ * [b \mapsto \_, \_]\}$ 
 $\text{free}(ob); \text{free}(ob+1);$ 
{ $emp * [b \mapsto \_, \_]$ }
}
{emp}

```

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. ([a \mapsto _, _ * \text{true}] \rightsquigarrow [a \mapsto _, _ * \text{true}])$

Proof Sketch of RCU Readers/Writers

```

{emp}
rcu_read_lock {
    {emp *  $b \mapsto \_, \_ * \text{true}$ } 
     $b' = b;$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }
     $x = *b';$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }
     $y = *(b' + 1);$ 
    {emp *  $b' \mapsto \_, \_ * \text{true}$ }
}
{emp}

```

ResInv : $b \mapsto _, _$

```

{nb  $\mapsto \_, \_$ }
rcu_write_lock {
    {nb  $\mapsto \_, \_ * b \mapsto \_, \_$ }
     $*nb = u; *(nb + 1) = v;$ 
    {nb  $\mapsto \_, \_ * b \mapsto \_, \_$ }
    barrier $_{nb \mapsto \_, \_};$ 
    {emp *  $nb \mapsto \_, \_ * b \mapsto \_, \_$ }
     $ob = b; b = nb;$ 
    {emp *  $b \mapsto \_, \_ * ob \mapsto \_, \_$ }
} sync {
    {ob  $\mapsto \_, \_ * b \mapsto \_, \_$ }
    free( $ob$ ); free( $ob + 1$ );
    {emp *  $b \mapsto \_, \_$ }
}
{emp}

```

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. (a \mapsto _, _ * \text{true}) \rightsquigarrow (a \mapsto _, _ * \text{true})$

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

{emp * $b \mapsto _, _ * \text{true}$ }

$b' = b;$

{emp * $b' \mapsto _, _ * \text{true}$ }

$x = *b';$

{emp * $b' \mapsto _, _ * \text{true}$ }

$y = *(b' + 1);$

{emp * $b' \mapsto _, _ * \text{true}$ }

}

{emp}

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$*nb = u; *(nb + 1) = v;$

{ $nb \mapsto _, _ * b \mapsto _, _$ }

barrier $_{nb \mapsto _, _};$

{emp * $nb \mapsto _, _ * b \mapsto _, _$ }

$ob = b; b = nb;$

{emp * $b \mapsto _, _ * ob \mapsto _, _$ }

} sync {

{ $ob \mapsto _, _ * b \mapsto _, _$ }

free(ob); free($ob + 1$);

{emp * $b \mapsto _, _$ }

}

{emp}

ResInv : $b \mapsto _, _$

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. (\boxed{a \mapsto _, _ * \text{true}} \rightsquigarrow \boxed{a \mapsto _, _ * \text{true}})$

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

{emp * $b \mapsto _, _ * \text{true}$ }

$b' = b;$

{emp * $b' \mapsto _, _ * \text{true}$ }

$x = *b';$

{emp * $b' \mapsto _, _ * \text{true}$ }

$y = *(b' + 1);$

{emp * $b' \mapsto _, _ * \text{true}$ }

}

{emp}

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$*nb = u; *(nb + 1) = v;$

{ $nb \mapsto _, _ * b \mapsto _, _$ }

barrier $_{nb \mapsto _, _};$

{emp * $nb \mapsto _, _ * b \mapsto _, _$ }

$ob = b; b = nb;$

{emp * $b \mapsto _, _ * ob \mapsto _, _$ }

} sync {

{ $ob \mapsto _, _ * b \mapsto _, _$ }

free(ob); free($ob + 1$);

{emp * $b \mapsto _, _$ }

}

{emp}

ResInv : $b \mapsto _, _$

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. (\boxed{a \mapsto _, _ * \text{true}} \rightsquigarrow \boxed{a \mapsto _, _ * \text{true}})$

Proof Sketch of RCU Readers/Writers

{emp}

rcu_read_lock {

{emp * $b \mapsto _, _ * \text{true}$ }

$b' = b;$

{emp * $b' \mapsto _, _ * \text{true}$ }

{ $nb \mapsto _, _$ }

rcu_write_lock {

{ $nb \mapsto _, _ * b \mapsto _, _$ }

$*nb = u; *(nb + 1) = v;$

{ $nb \mapsto _, _ * b \mapsto _, _$ }

ResInv : $b \mapsto _, _$

- 1) Invariant-based, thread-local reasoning.
- 2) barrier, sync as ownership transfer operations.
- 3) Use only preserved facts, when verifying readers.

} {emp}

} sync {

{ $ob \mapsto _, _ * b \mapsto _, _$ }

free(ob); free($ob+1$);

{emp * $b \mapsto _, _$ }

}

{emp}

AlwaysResInv : $b \mapsto _, _ * \text{true}$

Pres : $\forall a. (a \mapsto _, _ * \text{true} \rightsquigarrow a \mapsto _, _ * \text{true})$

Assertion Language

$P ::= \text{true} \mid E \mapsto F \mid \text{emp} \mid P * Q \mid \dots$

$K ::= P \mid \boxed{P} \mid K * L \mid \dots$

- P describes the heap local to the thread.
- \boxed{P} describes the shared heap.
- K describes both.

Judgments

$$(R, I(a)) \vdash^{\text{rd/wr/no}} \{K\}C\{L\}$$

Judgments

Precise
res. invariant.
(E.g. $b \mapsto _, _)$

$(R, I(a)) \vdash^{\text{rd/wr/no}} \{K\}C\{L\}$

Judgments

Precise
res. invariant.
(E.g. $b \mapsto _, _)$

$$(R, I(a)) \vdash^{\text{rd/wr/no}} \{K\} C \{L\}$$

Facts
preserved by writers.
(E.g. $a \mapsto _, _ * \text{true}$)

Judgments

Precise
res. invariant.
(E.g. $b \mapsto _, _)$

$(R, I(a))$

\vdash rd/wr/no

$\{K\}C\{L\}$

Indicates
“lock” held.

Facts
preserved by writers.
(E.g. $a \mapsto _, _ * \text{true}$)

Judgments

Precise
res. invariant.
(E.g. $b \mapsto _, _)$

Indicates
“lock” held.

$$(R, I(a)) \vdash^{\text{rd/wr/no}} \{K\} C \{L\}$$

Facts
preserved by writers.
(E.g. $a \mapsto _, _ * \text{true}$)

Proof Rule for Lock/Sync

$$\frac{(R, I(a)) \vdash^{\text{wr}} \{P * \boxed{R}\} C \{Q' * \boxed{R' * R}\} \\ (R, I(a)) \vdash^{\text{wr}} \{Q' * R' * \boxed{R}\} D \{Q * \boxed{R}\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_write_lock } C \text{ sync } D\{Q\}}$$

Proof Rule for Lock/Sync

Allowed to
access the shared heap
described by R.

$$\frac{(R, I(a)) \vdash^{\text{wr}} \{P * \boxed{R}\} C \{Q' * \boxed{R' * R}\} \quad (R, I(a)) \vdash^{\text{wr}} \{Q' * R' * \boxed{R}\} D \{Q * \boxed{R}\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_write_lock } C \text{ sync } D\{Q\}}$$

Proof Rule for Lock/Sync

$$\frac{(R, I(a)) \vdash^{\text{wr}} \{P * R\} C \{Q' * R' * R\} \quad (R, I(a)) \vdash^{\text{wr}} \{Q' * R' * R\} D \{Q * R\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_write_lock } C \text{ sync } D\{Q\}}$$

sync transfers R' from
shared to private.

Proof Rule for Lock/Sync

Every op in C and D should preserve R*true and I(a).

$$\frac{(R, I(a)) \vdash^{\text{wr}} \{P * \boxed{R}\} C \{Q' * \boxed{R' * R}\} \quad (R, I(a)) \vdash^{\text{wr}} \{Q' * R' * \boxed{R}\} D \{Q * \boxed{R}\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_write_lock } C \text{ sync } D\{Q\}}$$

Proof Rule for Readers

$$\frac{(R, I(a)) \vdash^{\text{rd}} \{P * \boxed{R * \text{true}}\} C \{Q * \boxed{R'}\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_read_lock } C \{Q\}}$$

Proof Rule for Readers

$$\frac{(R, I(a)) \vdash^{\text{rd}} \{P * \boxed{R * \text{true}}\} C \{Q * \boxed{R'}\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_read_lock } C \{Q\}}$$

Start with ResInv * true

Anything

Proof Rule for Readers

Every op in
 C can only read
the shared.

$$\frac{(R, I(a)) \vdash^{\text{rd}} \{P * [R * \text{true}]\} C \{Q * [R']\}}{(R, I(a)) \vdash^{\text{no}} \{P\} \text{rcu_read_lock } C \{Q\}}$$

Proof Rule

$$\frac{}{(R, I(a)) \vdash^{\text{wr}} \{P * Q * \boxed{R}\} \text{barrier}_Q \{P * \boxed{Q * R}\}}$$

Transfer Q from
private to shared.

Q has to
be precise.

Soundness

- No weak memory yet.
- Proved by the compilation into RGSep by Vafeiadis and Parkinson.

Messages

- Sync transfers a heaplet from shared to private.
- Barrier transfers a heaplet from private to shared.

RCU Reader/Writer

```
while (TRUE) {  
    rCU_read_lock {  
        b' = b;  
        x0 = *b';  
        x1 = *(b'+1);  
    }  
    .. USE x0,x1 ..  
}
```

```
while (TRUE) {  
    .. PRODUCE y0,y1 ..  
    nb = malloc(2);  
    rCU_write_lock {  
        *nb = y0; *(nb+1) = y1;  
        barrier();  
        ob = b; b = nb;  
    } sync {  
        free(ob); free(ob+1);  
    } }  
}
```

RCU Reader/Writer Optimized

```
while (TRUE) {  
  
    rcu_read_lock {  
  
        b' = b;  
  
        x0 = *b';  
  
        x1 = *(b'+1);  
  
    }  
  
    .. USE x0,x1 ..  
  
}
```

```
nb = malloc(2);
```

```
while (TRUE) {  
  
    .. PRODUCE y0,y1 ..  
  
    nb = malloc(2);  
  
    rcu_write_lock {  
  
        *nb = y0; *(nb+1) = y1;  
  
        barrier();  
  
        ob = b; b = nb;  
  
    } sync {  
  
        nb = ob;  
free(nb); nb = ob+1;  
  
    } }  
  
}}
```