

Paul E. McKenney, Facebook

Linux Plumbers Conference: Toolchains & Kernel MC, September 24, 2021



Report From The Standards Committees

What Has Been Happening?

- Concurrency TS 2 (hazard pointers, RCU)
- Lifetime-end pointer zap
- Undefined behavior
- Relaxed guide to relaxed
- `volatile_load` and `volatile_store`
- Address/data dependency ordering
 - Control dependencies covered in next session (“The never-ending saga of control dependencies”)

Concurrency Technical Specification 2

Concurrency Technical Specification 2

- In June 2021, C++ plenary session requested a Concurrency TS 2:
 - Hazard pointers
 - RCU: Adjusted to allow bare-bones implementations
 - Asymmetric fences? (`sys_membarrier()`)
- Maybe into C++26 or C++29

Concurrency TS 2: RCU Adjustments

- Naming (e.g., “rcu_synchronize()”)
- No rcu_head: Instead inherit from rcu_obj_base template class
- “Non-intrusive” retire() (AKA call_rcu() in kernel)
 - Zero storage overhead, similar to single-argument kfree_rcu() in kernel
- Callback invocation from retire()
 - Allows use in constrained environments, as in without softirq or any RCU grace-period kthread
- RAII readers: Automatic rcu_read_unlock() at end of scope
 - There are mechanisms to allow explicit unlock

Lifetime-End Pointer Zap


Lifetime-End Pointer Zap

- Important concurrent algorithms intentionally access pointers to lifetime-ended objects
 - LIFO Push (similar to Treiber's stack ca. 1973)
 - Single-element push and pop-all operations
 - Hazard pointers
 - Variants of sharded-locking methodology
 - Any number of debugging schemes


LIFO Push Algorithm Outline

- Push a single element
 - Allocate and initialize data fields
 - Repeat until `cmpxchg()` succeeds:
 - Initialize `->next` pointer to `top` pointer
 - Use `cmpxchg()` to point `top` to new element
- Pop entire stack
 - Use `xchg()` to pop entire list, setting `top` to `NULL`

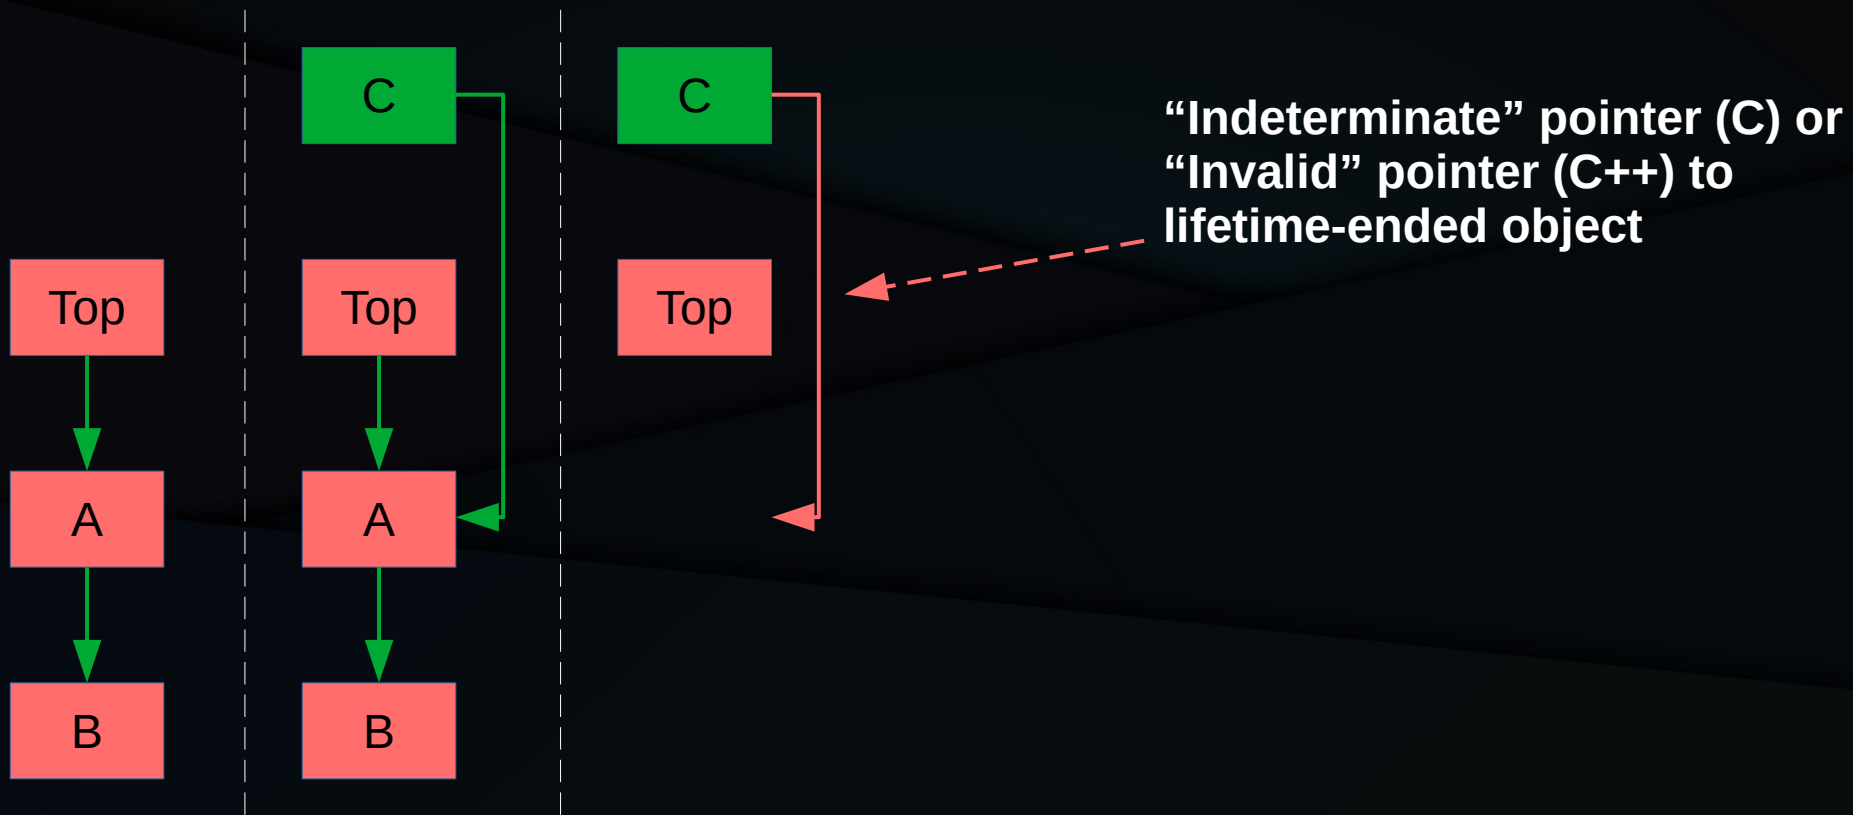
Initial State (Red == Concurrency)



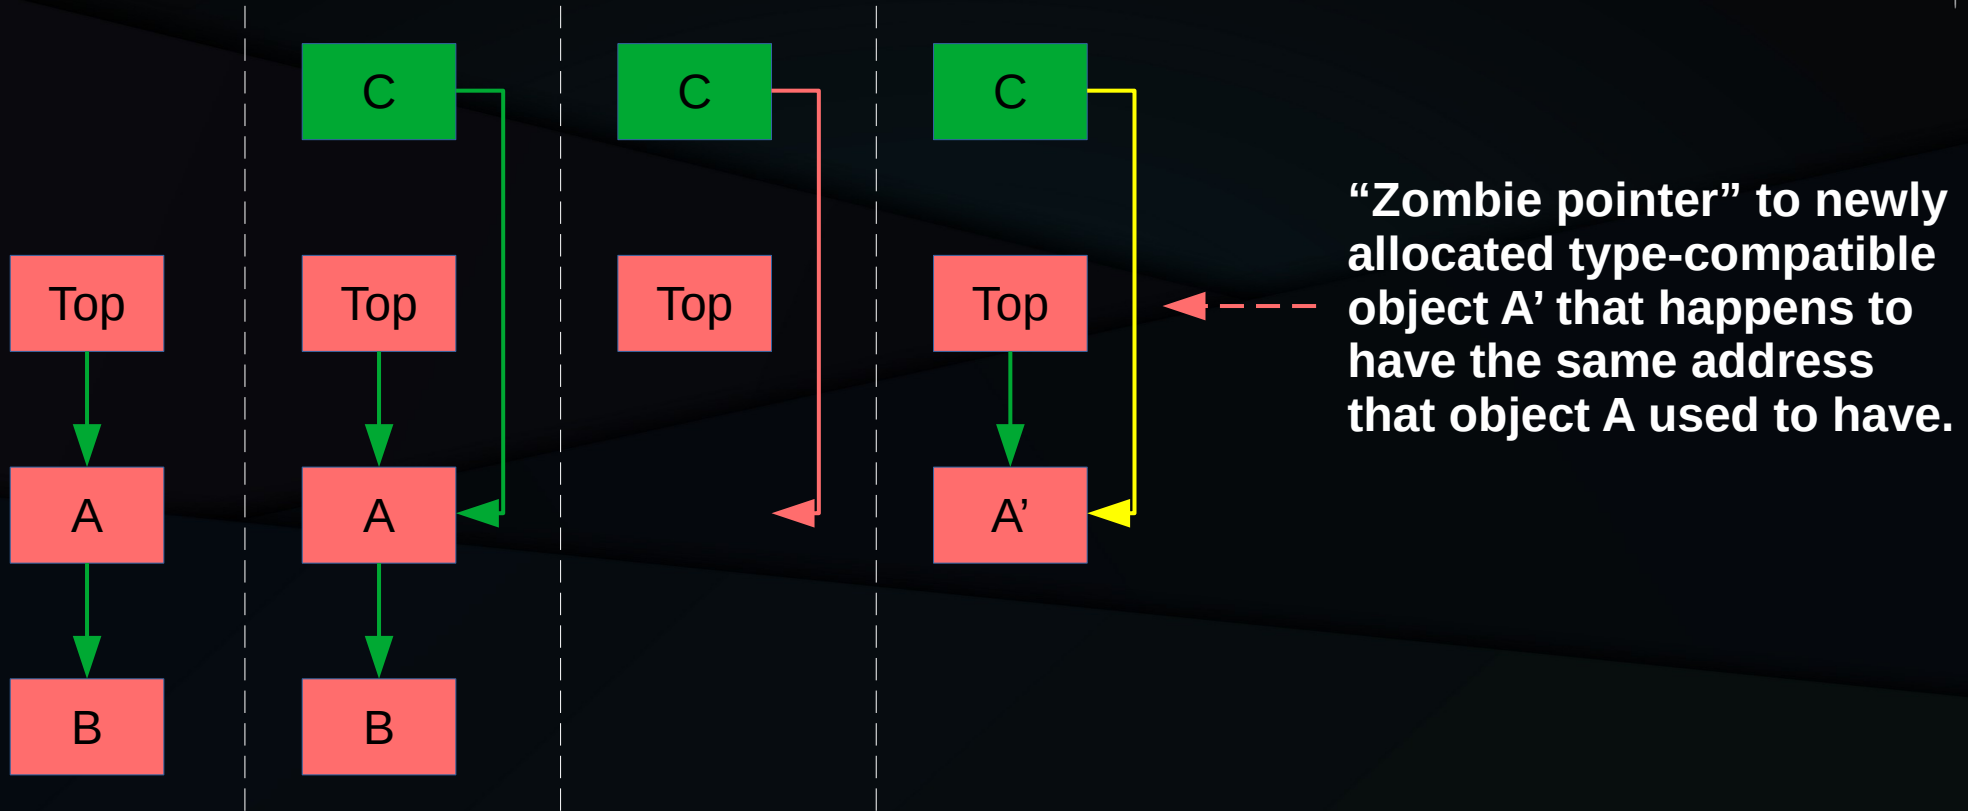
Push of C Begin: Allocate & Initialize



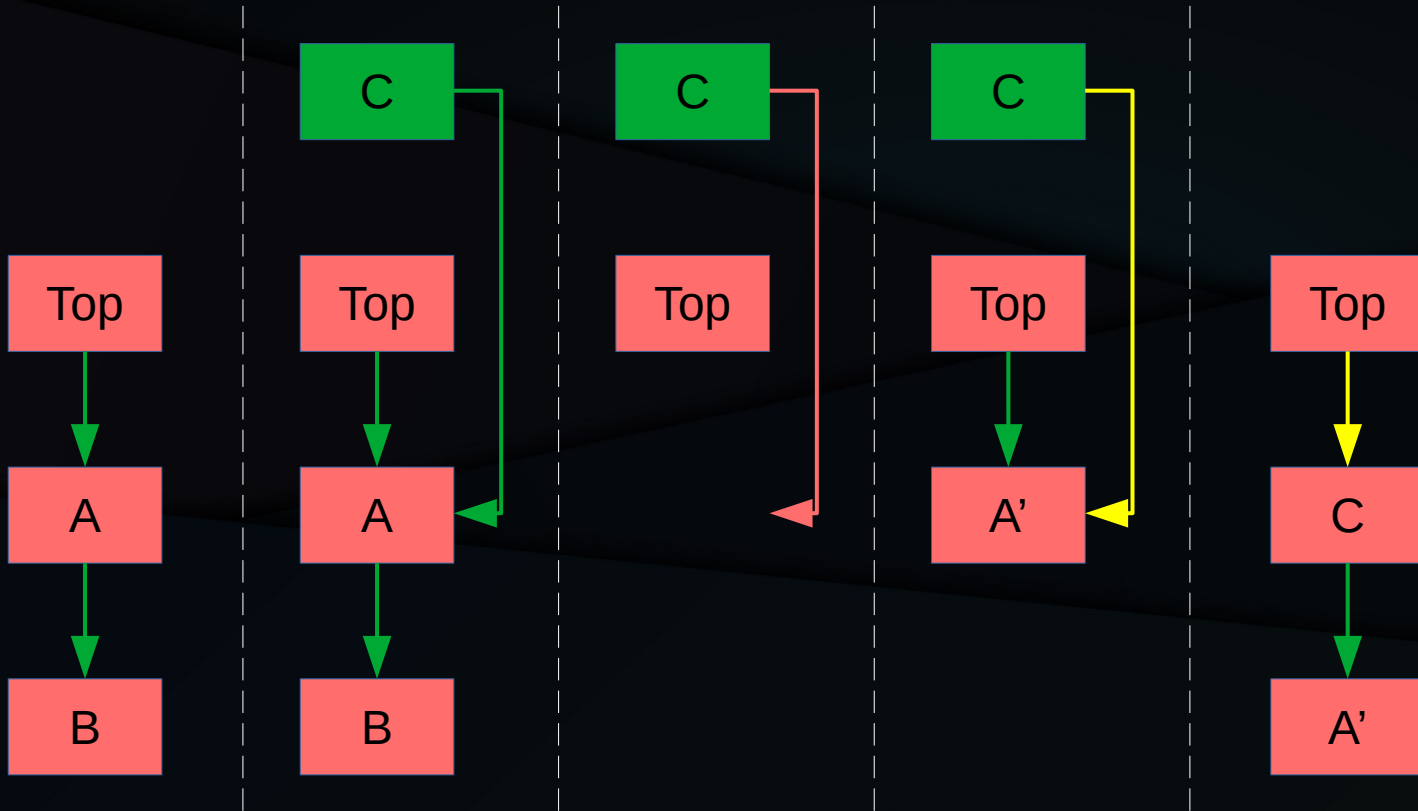
Intervening Pop-All Operation!!!



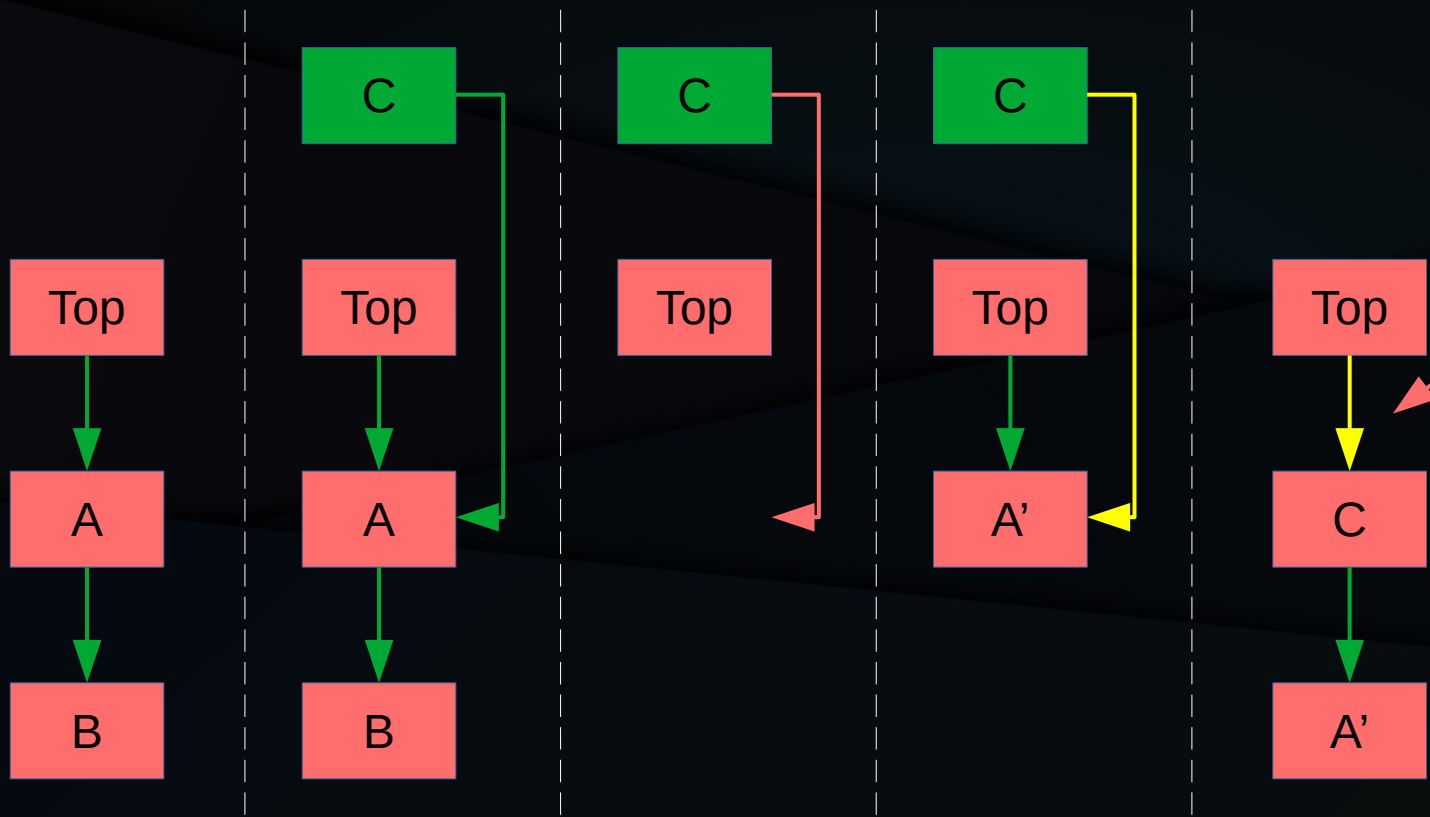
Push of A' (Reuses Memory of A)



Push of C Finally Completes



Compilers Hate Zombie Pointers!!!



The list's bits are just fine, but the compiler hates this zombie pointer!!!

Undefined Behavior

Undefined Behavior (UB)

- UB can back-propagate
- UB anywhere? Undefined everywhere!

```
int a[5];
```

```
int i = 3;
```

```
a[i] = 5;
```


Array-Out-Of-Bounds UB

- UB can back-propagate
- UB anywhere? Undefined everywhere!

```
int a[5];  
int i = 3;
```

```
a[i] = 5; ←----- UB here might set "i" to 5...
```

Back-Propagated UB

- UB can back-propagate
- UB anywhere? Undefined everywhere!

```
int a[5];  
int i = 3;
```

```
a[i] = 5; ←----- UB here might set "i" to 5...
```

Back-propagate...

Self-Justifying UB (Anonymous, 2007)

- UB can back-propagate
- UB anywhere? Undefined everywhere!

```
int a[5];
```

```
int i = 3; ←----- ... thus justifying the UB!!!
```

...back-propagate...

```
a[i] = 5; ←----- UB here might set "i" to 5...
```

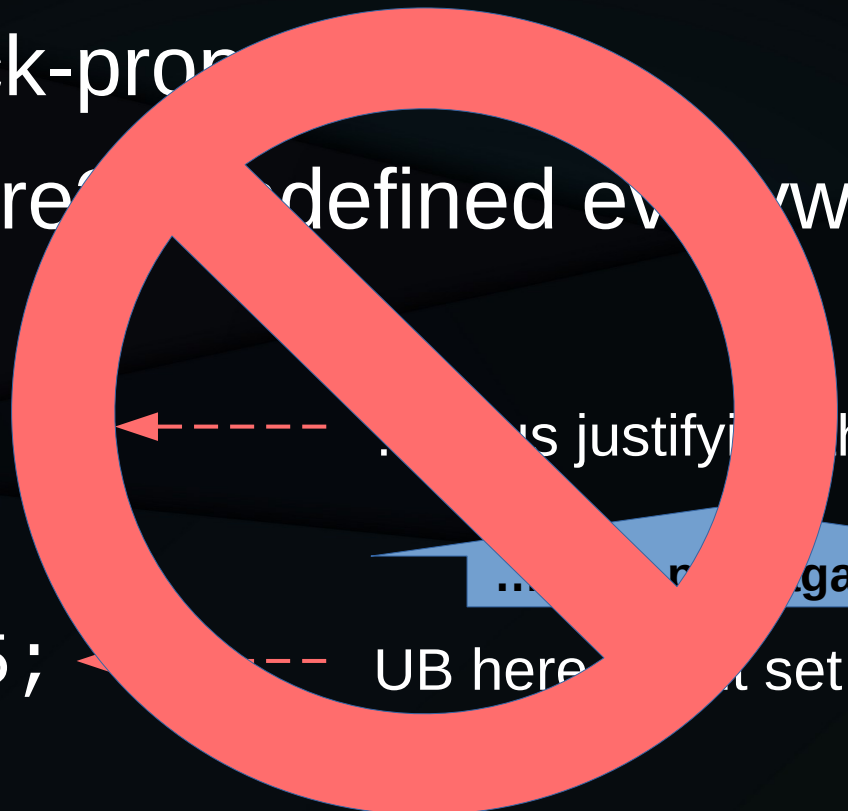
Self-Justifying UB (Anonymous, 2007)

- UB can back-propagate
- UB anywhere can be defined everywhere!

```
int a[5]
```

```
int i =
```

```
a[i] = 5;
```



... is justifying the UB!!!

... propagate...

UB here ... set "i" to 5...

“Cannot happen”, but no formal limitation justifying this.

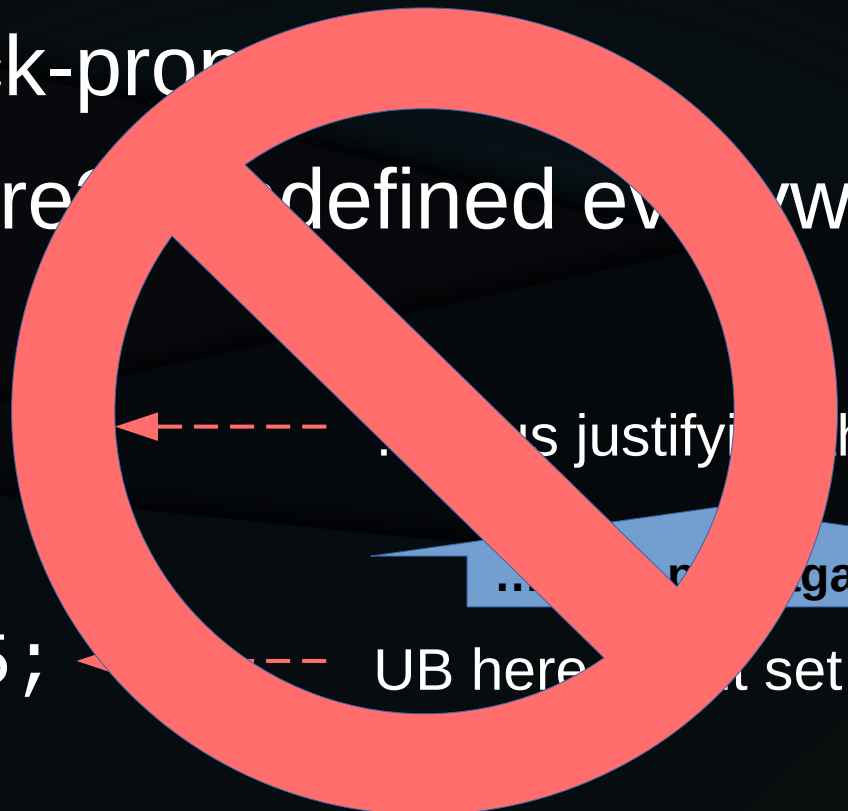
Self-Justifying UB (Anonymous, 2007)

- UB can back-propagate
- UB anywhere can be defined everywhere!

```
int a[5]
```

```
int i =
```

```
a[i] = 5;
```



... cannot justify the UB!!!

... propagate...

UB here ... set "i" to 5...

“Cannot happen”... Unless memory_ordered_relaxed concurrency is in play.

Relaxed Guide to Relaxed

Relaxed Guide to Relaxed

- C/C++ memory model allows OOTA values
 - T1: `x.store(y.load, mo_relaxed), mo_relaxed);`
 - T2: `y.store(x.load, mo_relaxed), mo_relaxed);`
 - Even if `x` & `y` initially `0`, could have `x==y==42`
- Why? C & C++ don't respect dependencies!

Relaxed Guide to Relaxed

- C/C++ memory model allows OOTA values
 - T1: `x.store(y.load, mo_relaxed), mo_relaxed);`
 - T2: `y.store(x.load, mo_relaxed), mo_relaxed);`
 - Even if `x` & `y` initially 0, could have `x==y==42`
- Why? C & C++ don't respect dependencies!
 - Random `mo_relaxed` programs problematic...

Relaxed Guide to Relaxed

- Random `memory_order_relaxed` programs problematic?
 - Don't ever use `memory_order_relaxed!!!`

Relaxed Guide to Relaxed

- Random `memory_order_relaxed` programs problematic?
 - Don't ever use `memory_order_relaxed!!!`
 - Pity about the poor performance...

Relaxed Guide to Relaxed

- Random `memory_order_relaxed` programs problematic?
 - ~~Don't ever use `memory_order_relaxed`!!!~~
 - Pity about the poor performance...
 - Instead, don't randomly generate programs involving `memory_order_relaxed` accesses!!!

Relaxed Guide to Relaxed

- Random `memory_order_relaxed` programs problematic?
 - ~~Don't ever use `memory_order_relaxed`!!!~~
 - Pity about the poor performance...
 - Instead, don't randomly generate programs involving `memory_order_relaxed` accesses!!!
- Design programs using known-good patterns

`volatile_load<T> & volatile_store<T>`

`volatile_load<T> & volatile_store<T>`

- Maybe C++'s answer to `READ_ONCE()` and `WRITE_ONCE()`

Address/Data Dependency Ordering

Address/Data Dependency Ordering

- Lots of electrons burned on this one...
 - <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0371r1.html>
 - <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0098r1.pdf>
 - <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2017/p0462r1.pdf>
 - <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2017/p0190r4.pdf>
 - <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2018/p0750r1.html>
 - <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2019/p0735r1.html>
- Next step: Implementation (GSoC prototype)

Summary

Summary

- C11/C++11 got the concurrency ball rolling
- But these cannot be the final word
- The historical separation of the C/C++ and concurrency communities has bitten us extremely hard!