

# Linux Plumbers Conference

Dublin, Ireland **September 12-14, 2022**

# Status Report: Broken Dependency Orderings in the Linux Kernel

Paul Heidekrüger\*, Marco Elver\*\*

\* *Technical University of Munich (TUM)*


\*\* *Google*

Pramod Bhatotia (TUM)

Soham Chakraborty (TU Delft)

Martin Fink (TUM)

Charalampos Mainas (TUM)



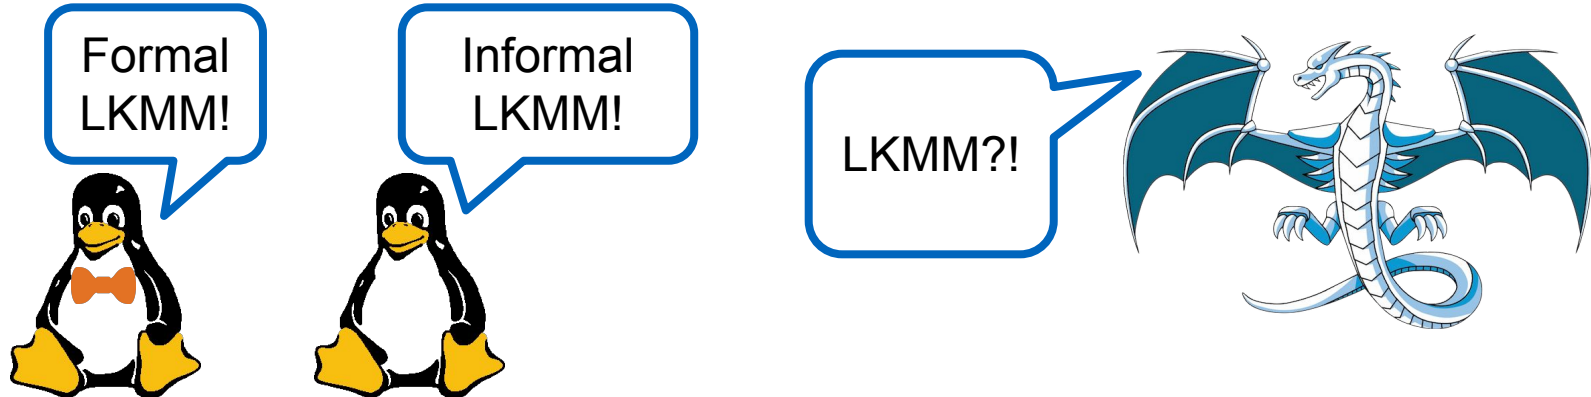
*Uhrenturm der TUM*

# The “Fear” of Broken Dependencies

- **Linux kernel** uses **non-standard C**.
- **Linux-kernel Memory Consistency Model (LKMM)** differs from C11 memory model.
- **Compilers unaware of LKMM**  $\Rightarrow$  potential for miscompilations (?)

# The “Fear” of Broken Dependencies

- **Linux kernel** uses **non-standard C**.
- **Linux-kernel Memory Consistency Model (LKMM)** differs from C11 memory model.
- **Compilers unaware of LKMM**  $\Rightarrow$  potential for miscompilations (?)



## Previously on ... Broken Dependency Orderings in the Linux Kernel




“[...] but dammit, I want to see an actual real example arguing for why it would be relevant and why the compiler would need our help.”

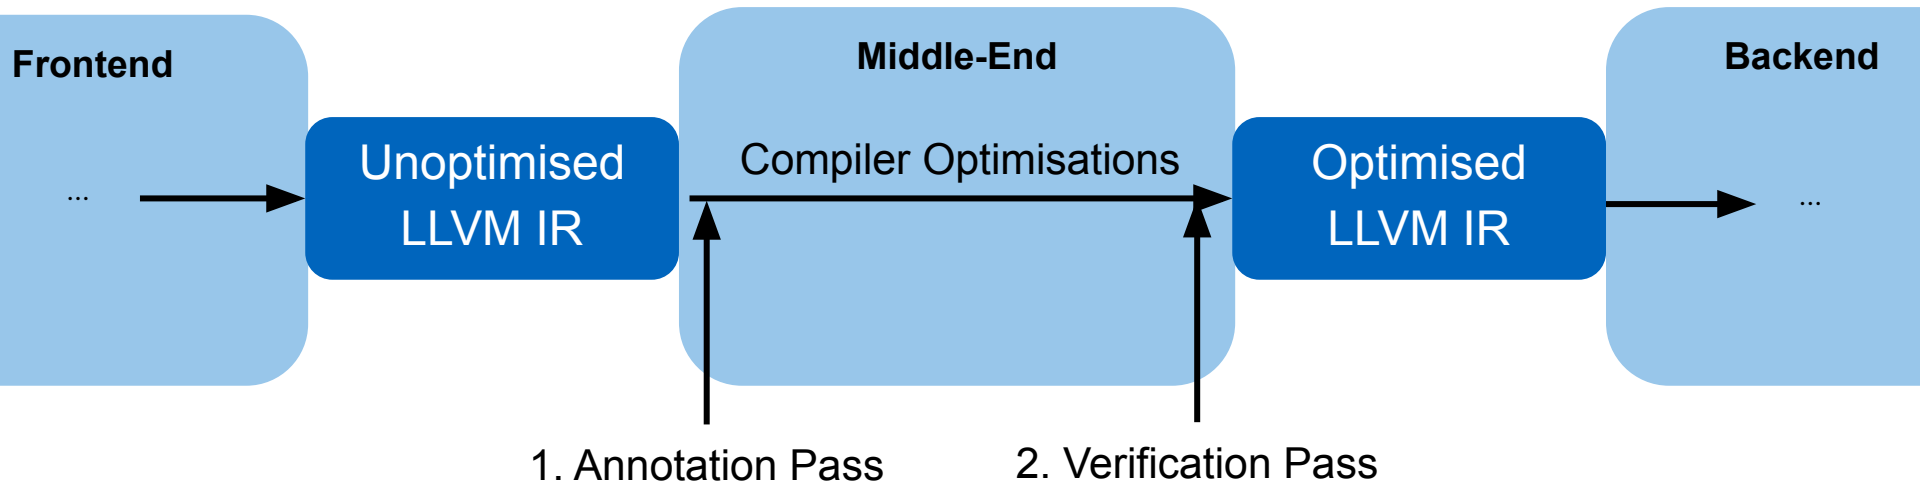
- Linus Torvalds -

## LKMM Dependency Checker

- **Annotates and verifies syntactic addr and ctrl dependencies**
- **Interprocedural analysis up to a given depth**
- **ClangBuiltLinux**



# Our Approach



# **Challenge:** Unambiguous and Implementable Dependency Definitions for Analysis of Real Kernel Code



# Part 1:

## Address Dependencies (and How to Break Them)

## But What Is an Address Dependency?

“A read event and another memory access event  
are linked by an address dependency  
if the value obtained by the read affects the  
location accessed by the other event.”

- tools/memory-model/Documentation/explanation.txt -

## Here's an Address Dependency

```
x = READ_ONCE(foo);
```

```
bar = &x[42];
```

```
y = READ_ONCE(*bar);
```

## Is This an Address Dependency?

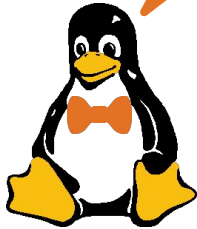
```
int arr[1];
```

```
x = READ_ONCE(foo);  
y = READ_ONCE(arr[x]);
```

## Is This an Address Dependency?

```
int arr[1];
```

```
x = READ_ONCE(foo);  
y = READ_ONCE(arr[x]);
```

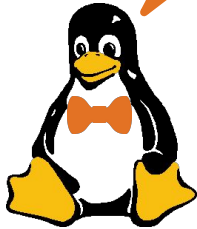


No.  
It's syntactic.

## Is This an Address Dependency?

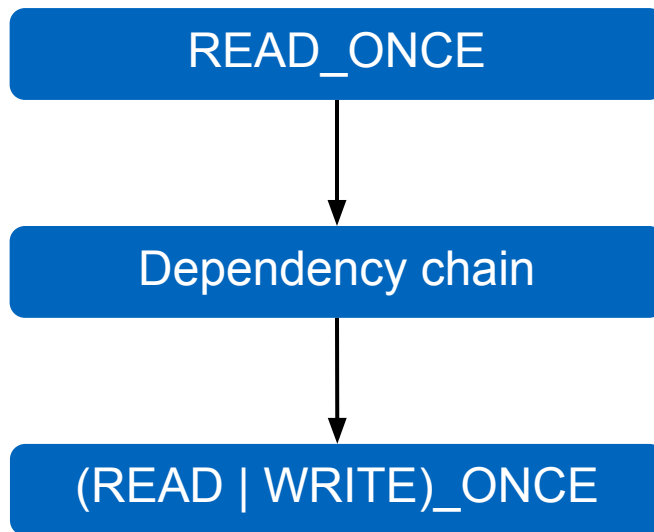
```
int arr[1];
```

```
x = READ_ONCE(foo);  
y = READ_ONCE(arr[0]);
```

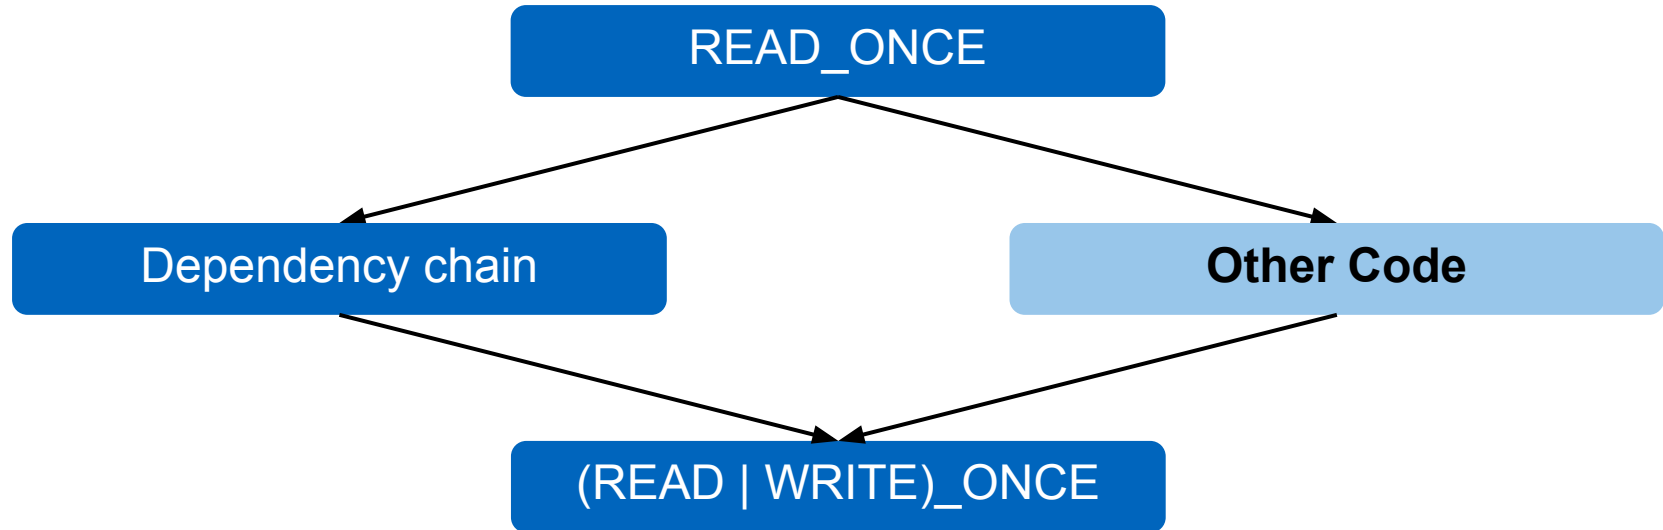


No.  
It's syntactic.

# Breaking Address Dependencies



# Breaking Address Dependencies - Making the Dep Chain Conditional






## Breaking Address Dependencies - Making the Dep Chain Conditional

```
x = READ_ONCE(foo);
```

```
bar = &x[42];
```

```
y = READ_ONCE(*bar);
```



```
x = READ_ONCE(foo);
```

```
if (x == baz) /* Oh no! */
```

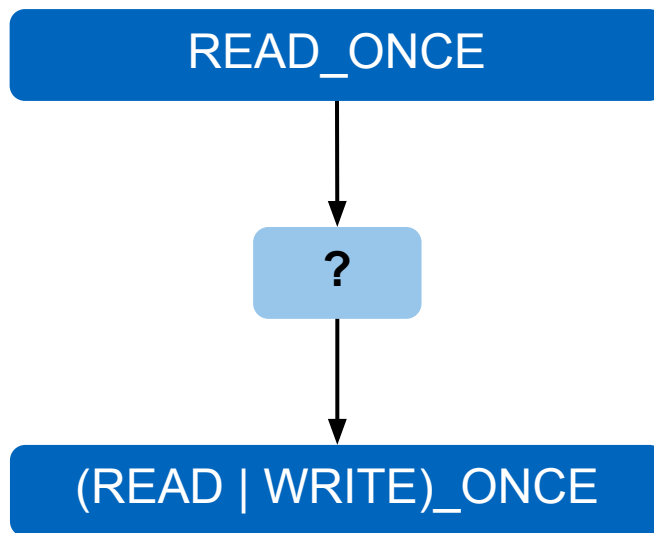
```
    bar = &baz[42];
```

```
else
```

```
    bar = &x[42];
```

```
y = READ_ONCE(*bar);
```


## Breaking Address Dependencies - No Dep Chain



## Broken Dependency Chain (mm/ksm.c::2032)

```
stable_node = page_stable_node(page); /* READ_ONCE(...) */  
if (stable_node)  
    /*More code*/  
    stable_node->head = &migrate_nodes;  
    list_add(&stable_node->list, stable_node->head); /* WRITE_ONCE() */
```

```
stable_node = page_stable_node(page); /* READ_ONCE(...) */
if (stable_node)
    /*More code*/
    stable_node->head = &migrate_nodes;
    list_add(&stable_node->list, stable_node->head); /* WRITE_ONCE() */
```



```
stable_node = page_stable_node(page); /* READ_ONCE(...) */
if (stable_node)
    /*More code*/
    /* Oh no! */
    list_add(&stable_node->list, &migrate_nodes); /* WRITE_ONCE() */
```

## Transformation to Conditional (fs/nfs/delegation.c::617 - 622)

```
if (place_holder)
    delegation = rcu_dereference([...]); /* READ_ONCE() */
if (!delegation || delegation != place_holder_deleg)
    delegation = list_entry_rcu([...]);
for(..., delegation = list_entry_rcu(delegation, [...])) /* READ_ONCE() */
```

```
if(place_holder == NULL) {
    delegation = list_entry_rcu([...]);
else {
    cmp = rcu_dereference(place_holder->delegation); /* READ_ONCE(...) */
    if(cmp != NULL) {
        if(cmp == place_holder_deleg) /* Oh no! */
            delegation = place_holder_deleg; /* Oh no! (cont.) */
        else
            delegation = list_entry_rcu([...]);
    } else {
        delegation = list_entry_rcu([...]);
    }
}

for([...], delegation = list_entry_rcu(delegation, [...]) /* READ_ONCE() */
```

# **Challenge:** Unambiguous and Implementable Dependency Definitions for Analysis of Real Kernel Code

# Part 2:

## Control Dependencies (and How to Break Them)



## But What Is a Control Dependency?

“Finally, a read event  $X$  and a write event  $Y$  are linked by a control dependency if  $Y$  syntactically lies within an arm of an if statement and  $X$  affects the evaluation of the if condition via a data or address dependency (or similarly for a switch statement)”

- tools/memory-model/Documentation/explanation.txt (in a few weeks) -


## Here's a Control Dependency

```
x = READ_ONCE(foo);
```


```
if(x == 42)
```

```
    WRITE_ONCE(bar, 42);
```

```
return true;
```

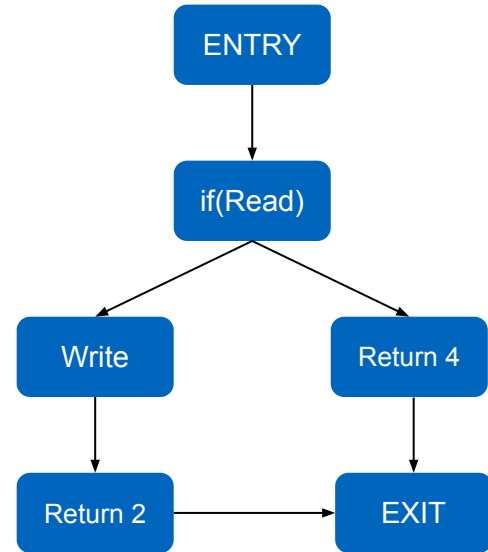


Agreed.



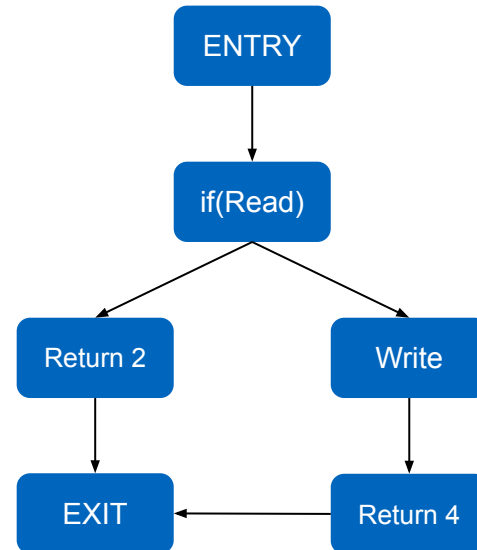
## Is this a Control Dependency?

```
if(READ_ONCE(x))  
    return 4;  
WRITE_ONCE(y, 21);  
return 2;
```



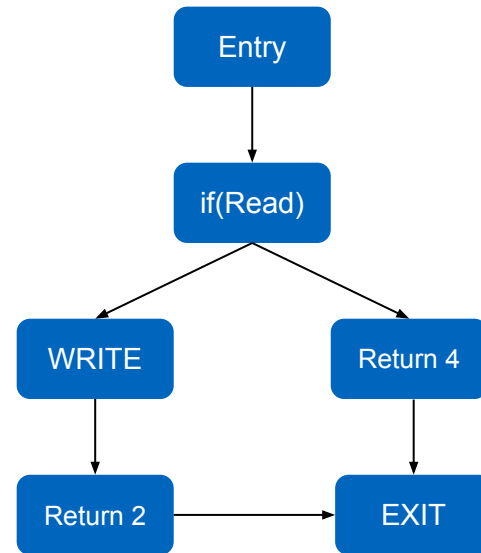

## Is this a Control Dependency?

```
if(!READ_ONCE(x)) {  
    WRITE_ONCE(y, 42);  
    return 2;  
}  
return 4;
```




# Is this a Control Dependency?

```
if(READ_ONCE(x))  
    return 42;  
WRITE_ONCE(y, 42); /* Uhm? */  
return 0;
```




# Is this a Control Dependency?

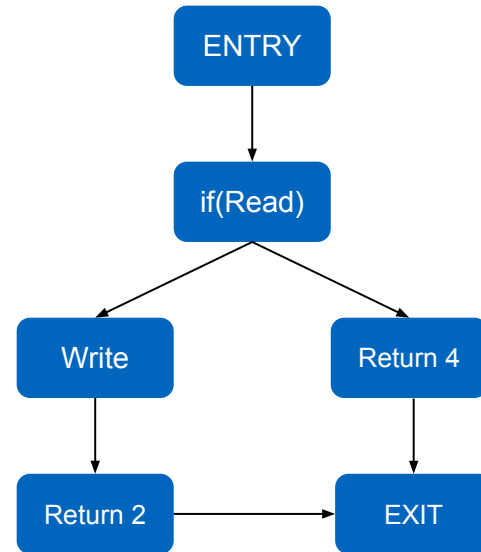
```
if(READ_ONCE(x))  
    return 42;  
WRITE_ONCE(y, 42);  
  
/* The answer has to be Yes! */  
return 0;
```



No.



Yes!



## But What Is a Control Dependency?

“Let  $G$  be a control flow graph. Let  $X$  and  $Y$  be nodes in  $G$ .

$Y$  is control dependent on  $X$  iff


- (1) there exists a directed path  $P$  from  $X$  to  $Y$  with any  $Z$  in  $P$   
(excluding  $X$  and  $Y$ ) post-dominated by  $Y$  and
- (2)  $X$  is not post-dominated by  $Y$ .”

## But What Is Post Dominance?

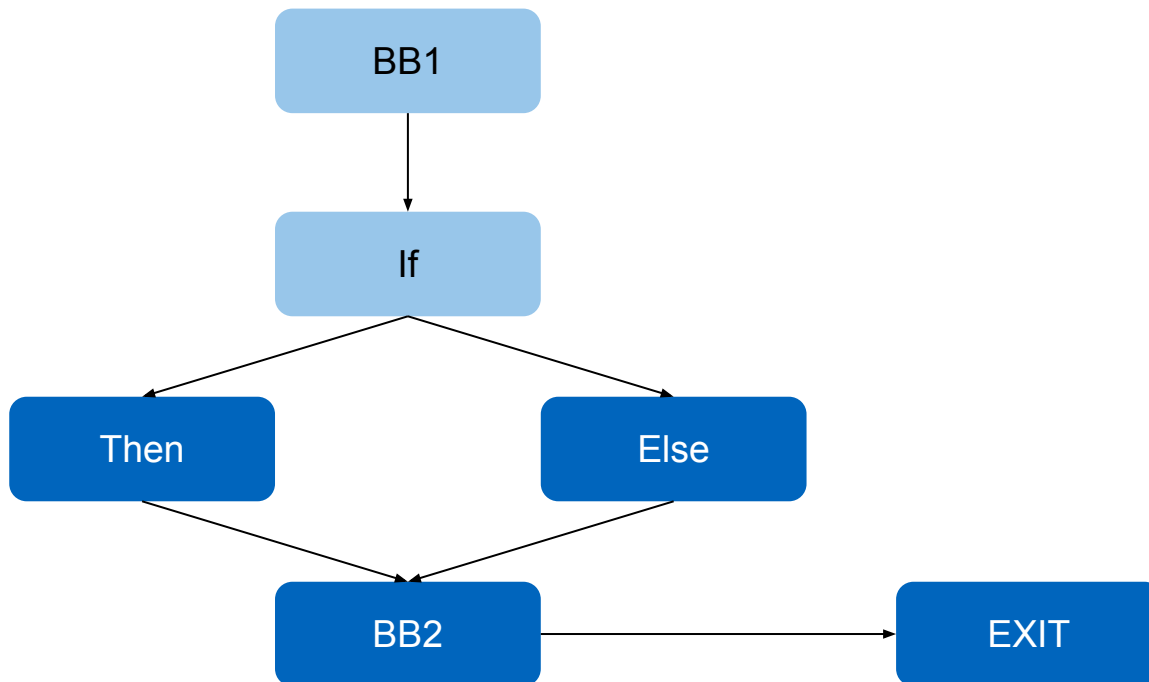
“A node  $V$  is post-dominated by a node  $W$  in  $G$   
if every directed path  
from  $V$  to  $STOP$  (not including  $V$ )  
contains  $W$ .”



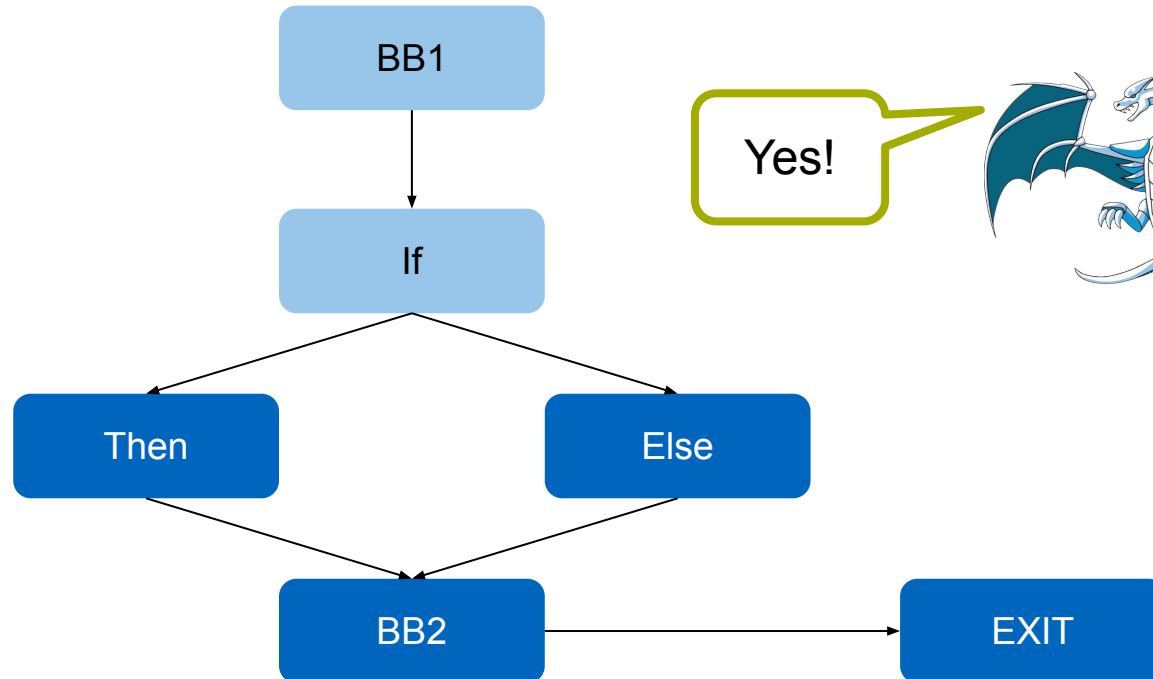
# Illustrating Post-Dominance




## Does “If” Post-Dominate “BB1”?



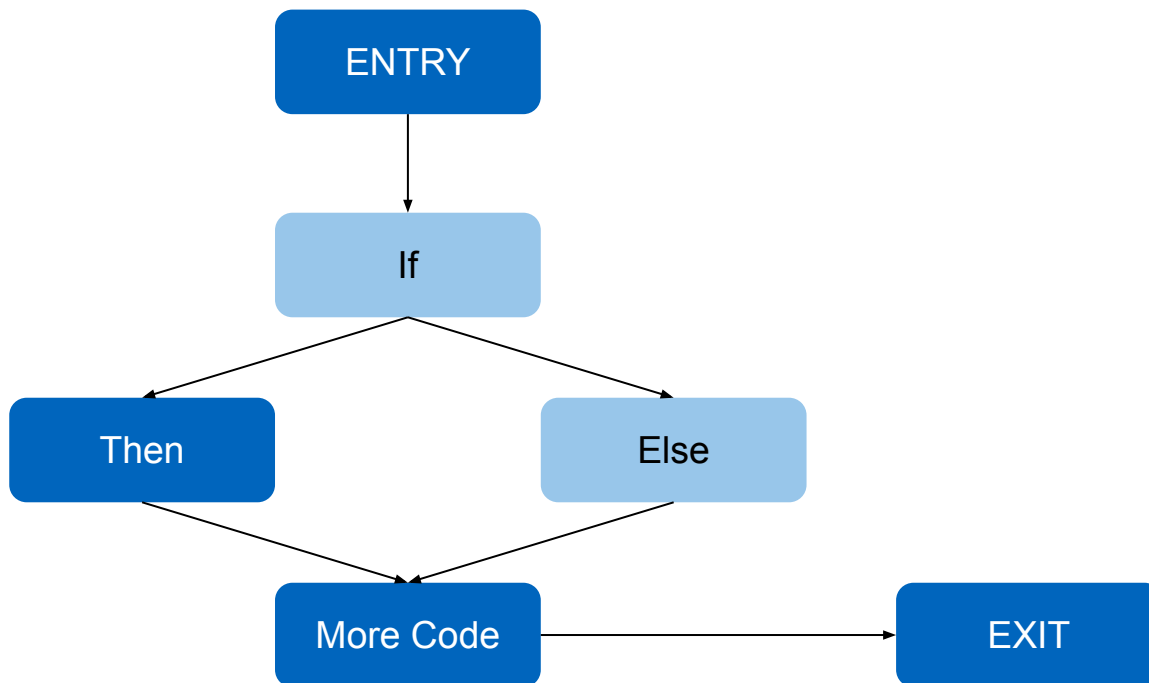
# Does “If” Post-Dominate “BB1”?



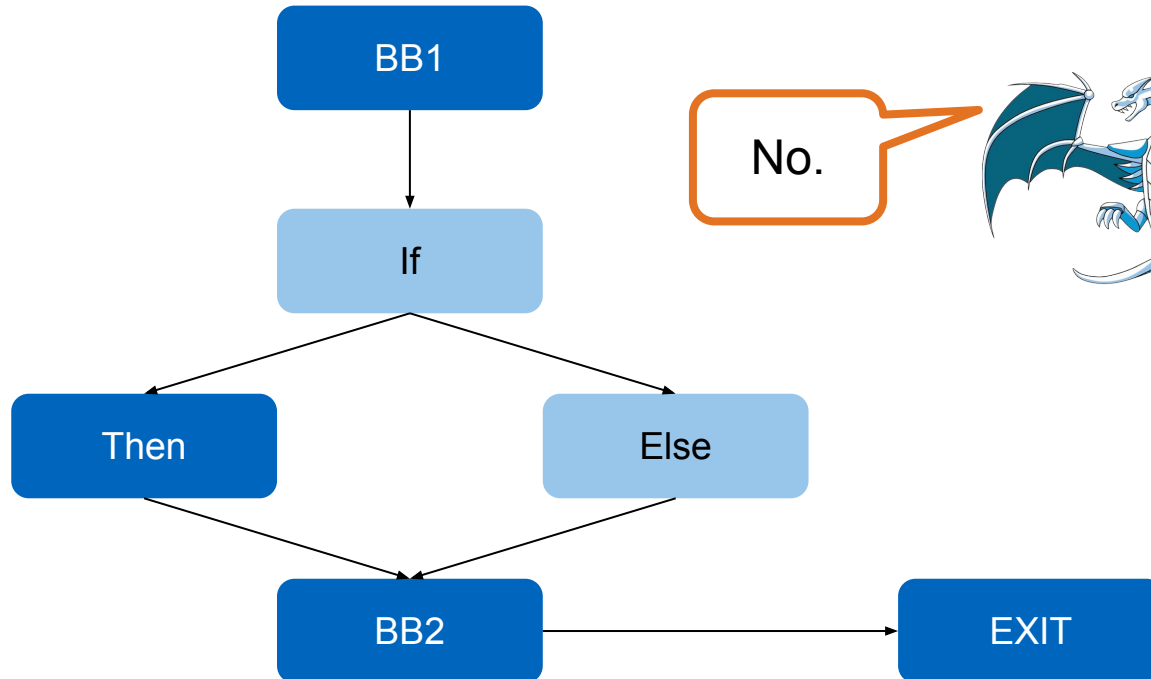
Yes!



# Does “Then” Post-Dominate “If”?




# Does “Then” Post-Dominate “If”?



# Here's a Control Dependency

```
x = READ_ONCE(foo);  
if(x == 42)  
    WRITE_ONCE(bar, 24);  
return true;
```

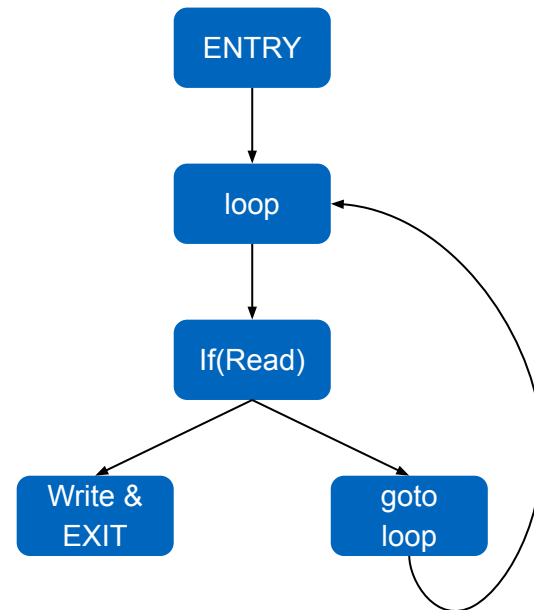


# Is This a Control Dependency?

```
loop:
```

```
if(READ_ONCE(x)) {  
    WRITE_ONCE(y, 42);  
    return 0;  
}
```

```
goto loop;
```





## Is This a Control Dependency?

```
loop:
```

```
if(READ_ONCE(x)) {  
    WRITE_ONCE(y, 42);  
    return 0;  
}
```

```
goto loop;
```

No.






# Is this a Control Dependency?

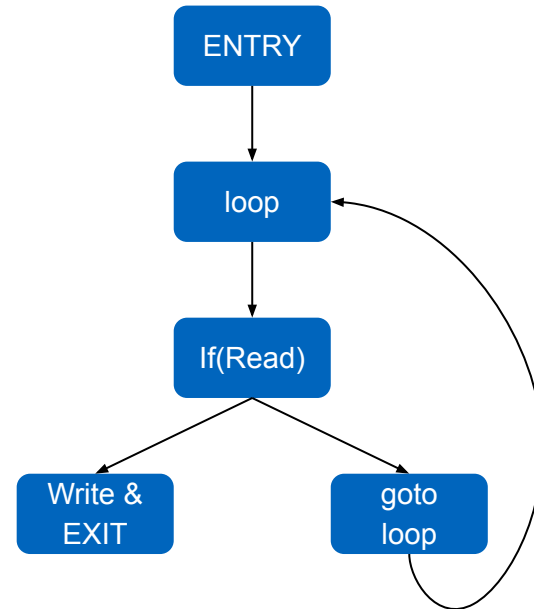
```
loop:
```

```
if(READ_ONCE(x)) {  
    WRITE_ONCE(y, 42);  
    return 0;  
}
```

```
goto loop;
```




Yes!




# Is this a Control Dependency?

loop:

```
if(READ_ONCE(x)) {
    WRITE_ONCE(y, 42); /* Yes! */
    return 0;
}
```




goto loop;



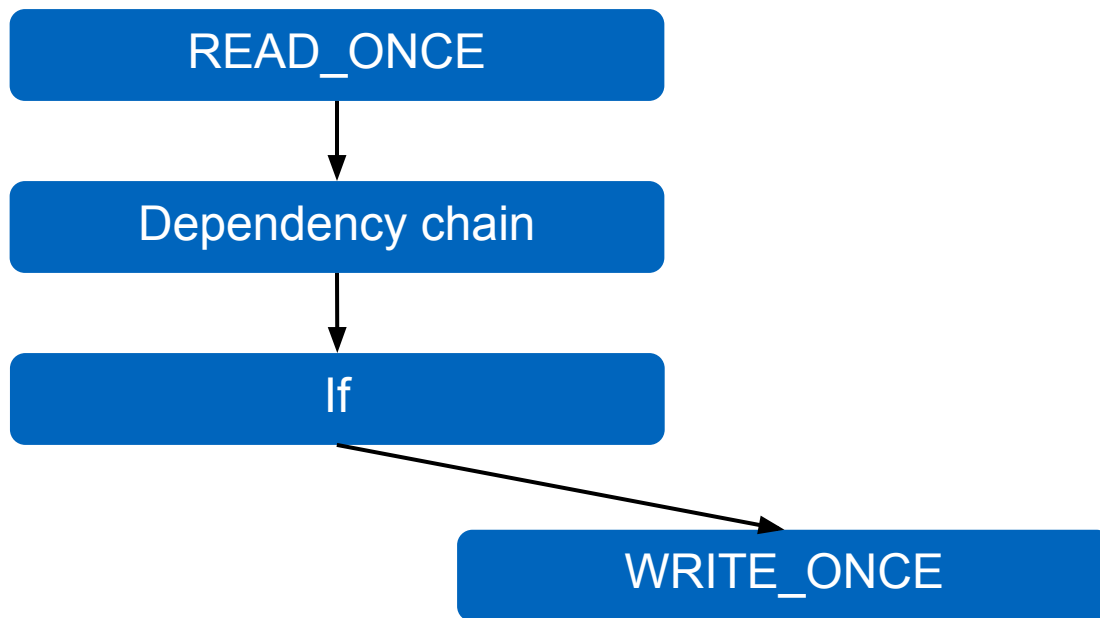
## But What Is a Control Dependency?

There is a control dependency from a marked read A to a marked write B if there is a condition C s.t. B is within the scope of C and C depends on A.

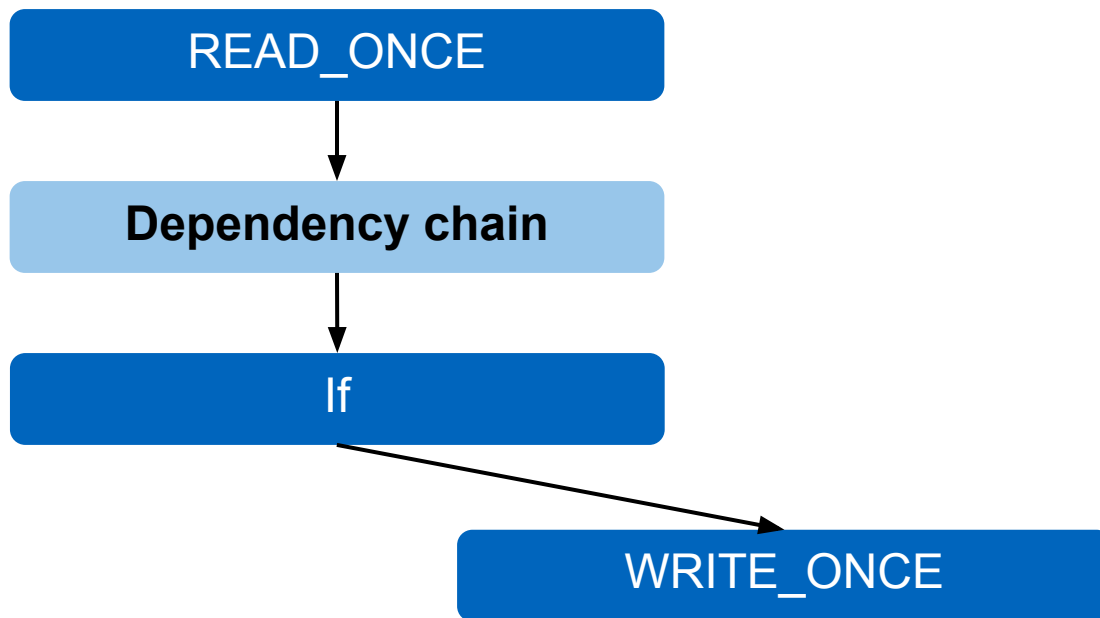
- WIP Control Dependency Definition -



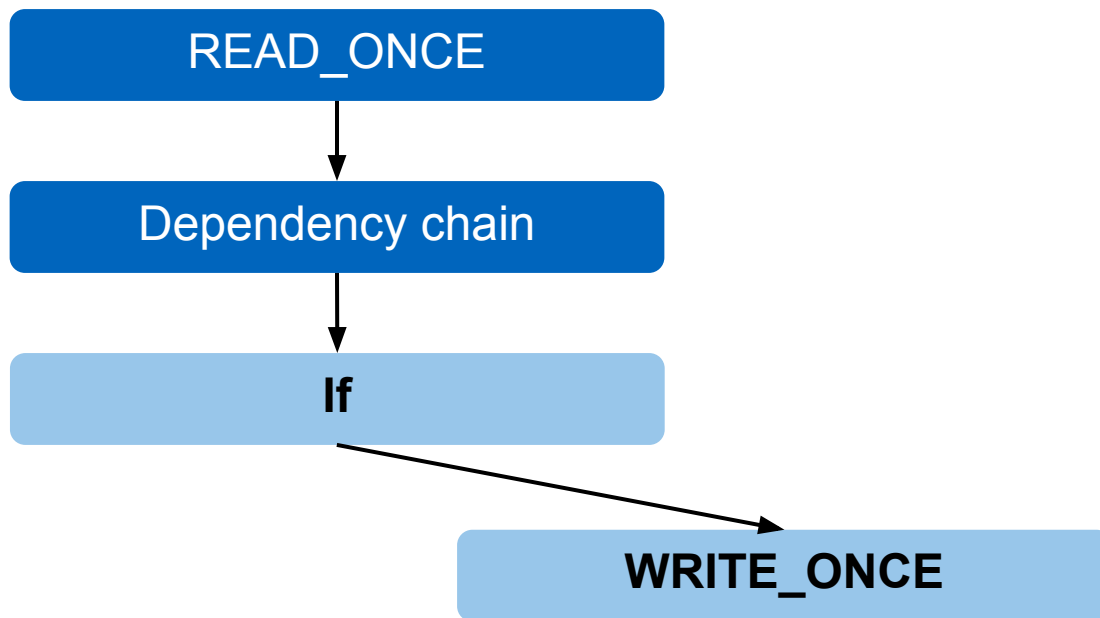
## Breaking Control Dependencies (?)



## Breaking Control Dependencies (?)



## Breaking Control Dependencies (?)



# Is this a Control Dependency?


```
#define MAX 1
```

```
x = READ_ONCE(*foo);
```

```
if (x % MAX == 0)
```

```
    WRITE_ONCE(*bar, 1);
```

```
Return true;
```



## Is this a Control Dependency?


```
#define MAX 1
```

```
x = READ_ONCE(foo);
```

```
/* No branch?! */
```

```
WRITE_ONCE(*bar, 1);
```

```
Return true;
```





# Is this a Control Dependency?


```
#define MAX 1
```

```
x = READ_ONCE(*foo);
```


```
if (x % MAX == 0)
```

```
    WRITE_ONCE(*bar, 1);
```

No.  
It's syntactic.



It depends.



## Is this a Control Dependency?

```
x = READ_ONCE(*bar);  
if (x == 42)  
    WRITE_ONCE(*baz, 1);  
else  
    WRITE_ONCE(*baz, 2);  
WRITE_ONCE(*y, 42);
```

## Is this a Control Dependency?

```
x = READ_ONCE(*bar);
```

```
if (x == 42)
```


```
    /* arm64 compiler can make this a conditional select */
```

```
    WRITE_ONCE(*baz, 1);
```

```
else
```


```
    WRITE_ONCE(*baz, 2);
```

```
WRITE_ONCE(*y, 42);
```



Plural - Control Dependencies!  
There is more than one.


It depends.



## Where Does this Leave Us?

“To What Extent Are Compilers Undermining the Linux Kernel Memory Model?”

- **LKMM addr dependency checker has found broken dependencies**
- LKMM addr dependency checker **LLVM RFC** hopefully very soon
- WIP LKMM **ctrl dependency checker**

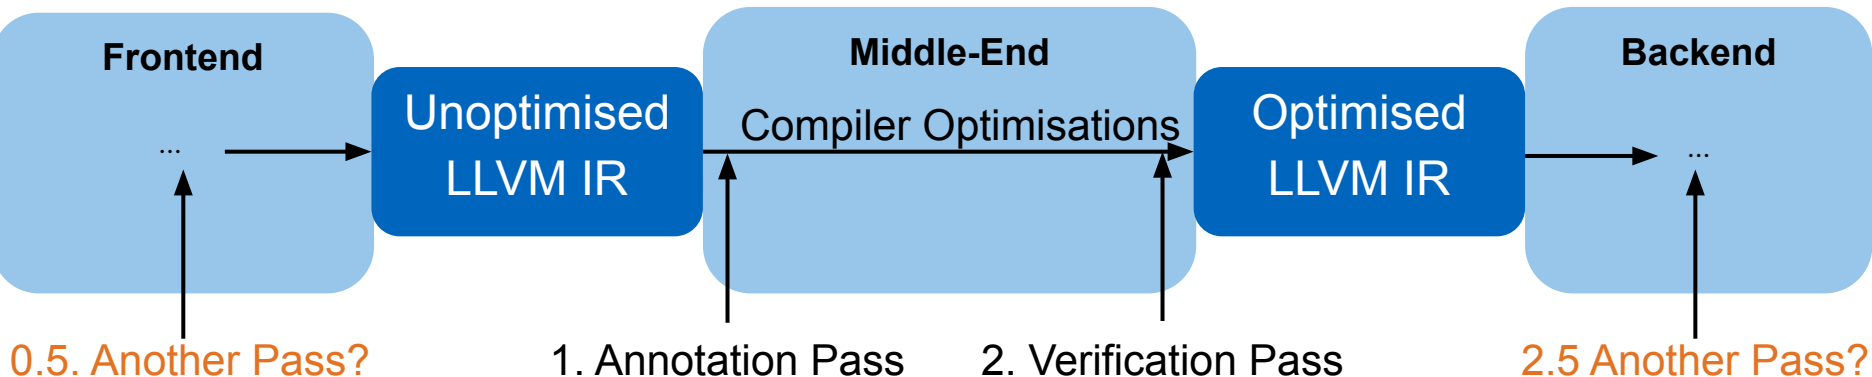


**Broken dependencies are real!** 😱

*The LKMM Dependency Checker found broken dependencies!*

# Limitations


- **Syntactic dependency** !=> **semantic dependency**
- **Computational feasibility** (interprocedural analysis, variadic functions ...)
- Our analysis avoids **backedges** (right now)
- **Frontend** and **backend** optimisations not accounted for



# The Hunt for Broken Dependencies - Future Plans

“To What Extent Are Compilers Undermining the Linux Kernel Memory Model?”

- **AArch64 defconfigs** for various Linux kernel flavours (incl. LTO)
- **Random testing** with random, but relevant, Linux kernel configs
- **LLVM RFCs**
- Find **ctrl to addr** dependency **transformations**
- Extend our analysis to the **backend**
- Get the ball rolling on **short-term** to **long-term fixes** with sufficient evidence



Get in touch!



[twitter.com/pbhdk](https://twitter.com/pbhdk)

[Paul.Heidekrueger@in.tum.de](mailto:Paul.Heidekrueger@in.tum.de)




[linkedin.com/in/pbhdk](https://linkedin.com/in/pbhdk)

# The Hunt for Broken Dependencies - Discussion

“To What Extent Are Compilers Undermining the Linux Kernel Memory Model?”

- **Annotating dependency chains?**
- **-memory-model=1kmm** option for compilers?
- **Combating false positives:**
  - **“safe regions”** via **pragmas** or **function attributes**?
  - **Exclusion lists**?
- **Suggesting fixes** to the user? **Automatically insert barriers**?



Get in touch!



[twitter.com/pbhdk](https://twitter.com/pbhdk)

[Paul.Heidekrueger@in.tum.de](mailto:Paul.Heidekrueger@in.tum.de)



[linkedin.com/in/pbhdk](https://linkedin.com/in/pbhdk)