

Verifying the Verifier: eBPF Range Analysis Verification

Harishankar Vishwanathan

Matan Shachnai

Srinivas Narayana

Santosh Nagarakatte

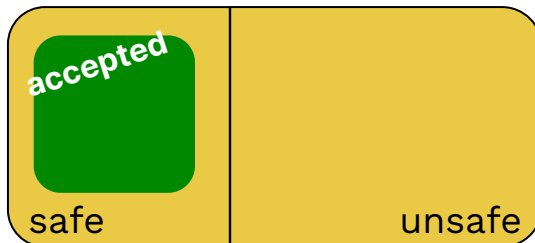
One Page Summary

- The eBPF Verifier is crucial.
- Static Analysis is a crucial part of the eBPF verifier.
 - Writing correct static analysis is hard. Formal verification can be useful!
- Contributions: A tool, Agni:
 - Automatically checks the correctness of (part of) the static analysis in the eBPF verifier, on every commit.
- Results
 - Analyses in kernels starting from 5.13 — 5.19 are correct.
 - 4.14 — 5.12: Agni reports bugs.
 - Agni can generate proof-of-concept eBPF programs that manifest bugs.
- Please give your feedback

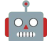
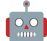
eBPF Verification Must be Sound

- **Soundness:** Unsafe programs should be rejected
- **Safety:**
 - Termination
 - Illegal operations
 - Memory access

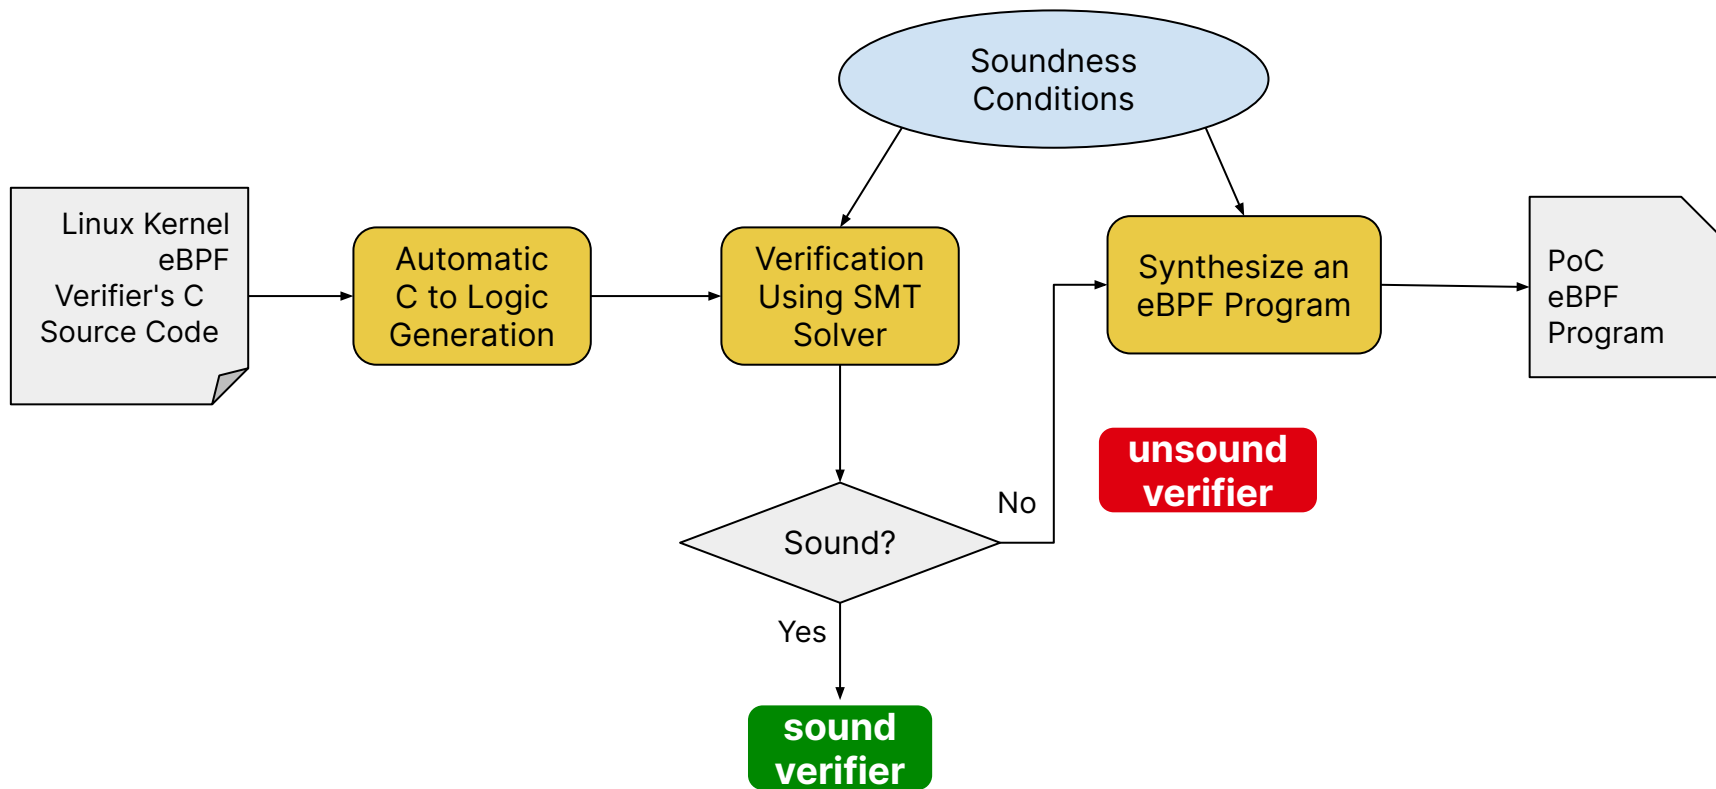
Can we formally verify the soundness of the static analysis in the eBPF verifier?



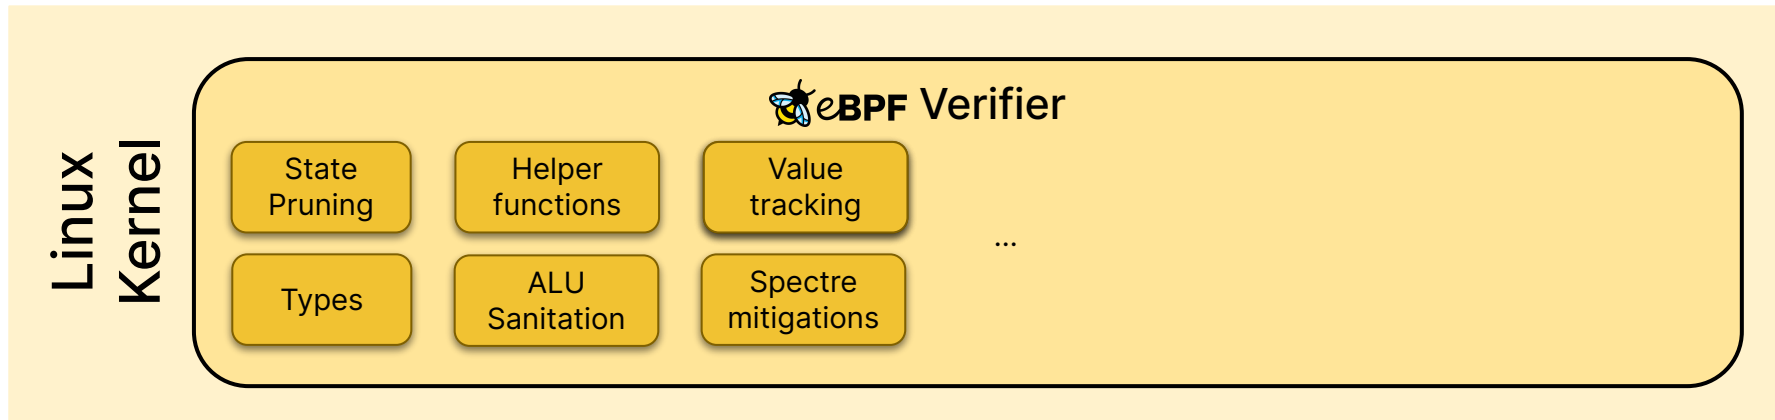
Goals

- Goal: Verify the soundness on every commit
 - Verifier is a large code base
 - Constantly changing
- Options:
 - Manually verify
 - Manually write the kernel code in Logic (SMT)
 - Repeat on every commit
 - Tedious and error prone
 -  Automate 

Overview

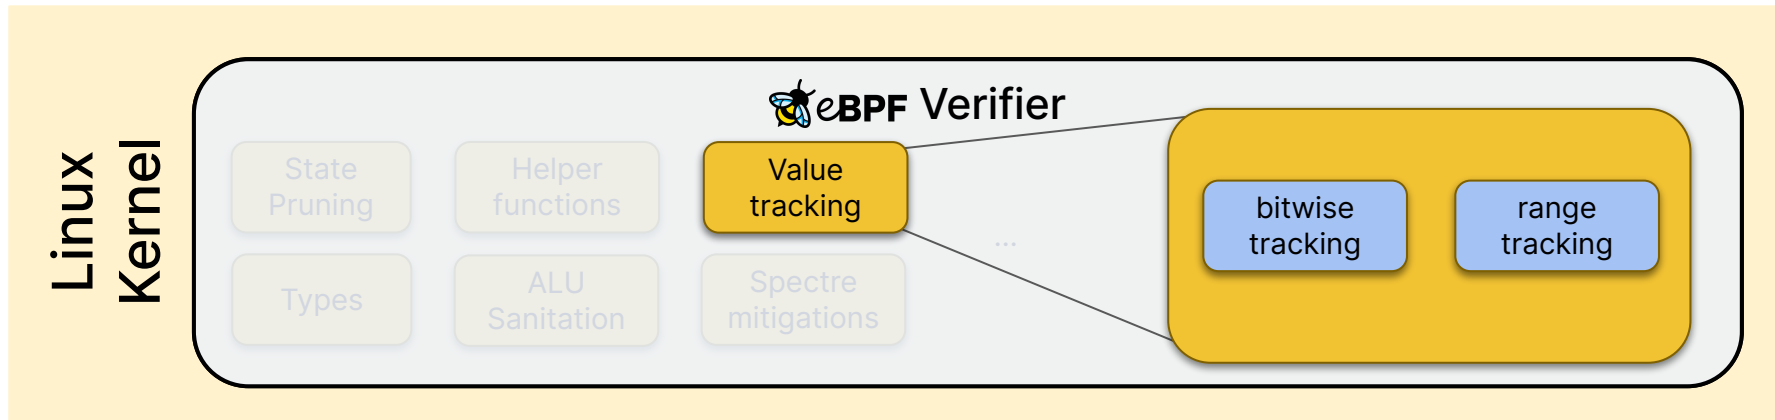


Static Analyses in the eBPF verifier



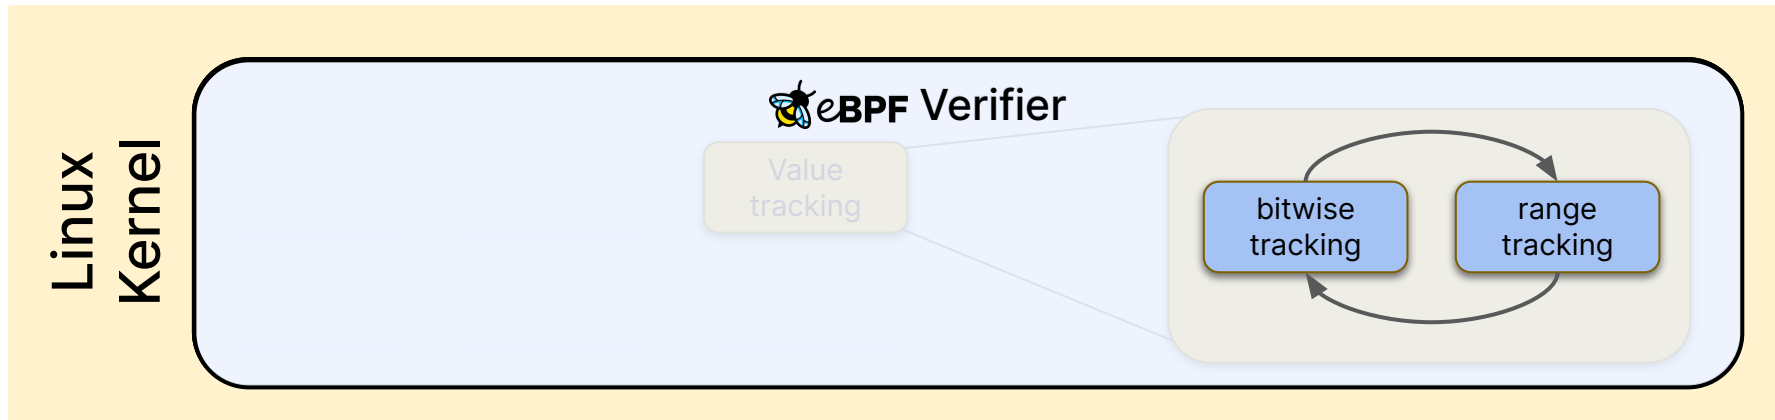
- Tracking the **values** of program variables across **all** executions of the program
- Our work: Reasoning about the soundness and precision of the range analysis + bitwise tracking + their combination

Static Analyses in the eBPF verifier



- Tracking the **values** of program variables across **all** executions of the program
- Our work: Reasoning about the soundness and precision of the range analysis + bitwise tracking + their combination

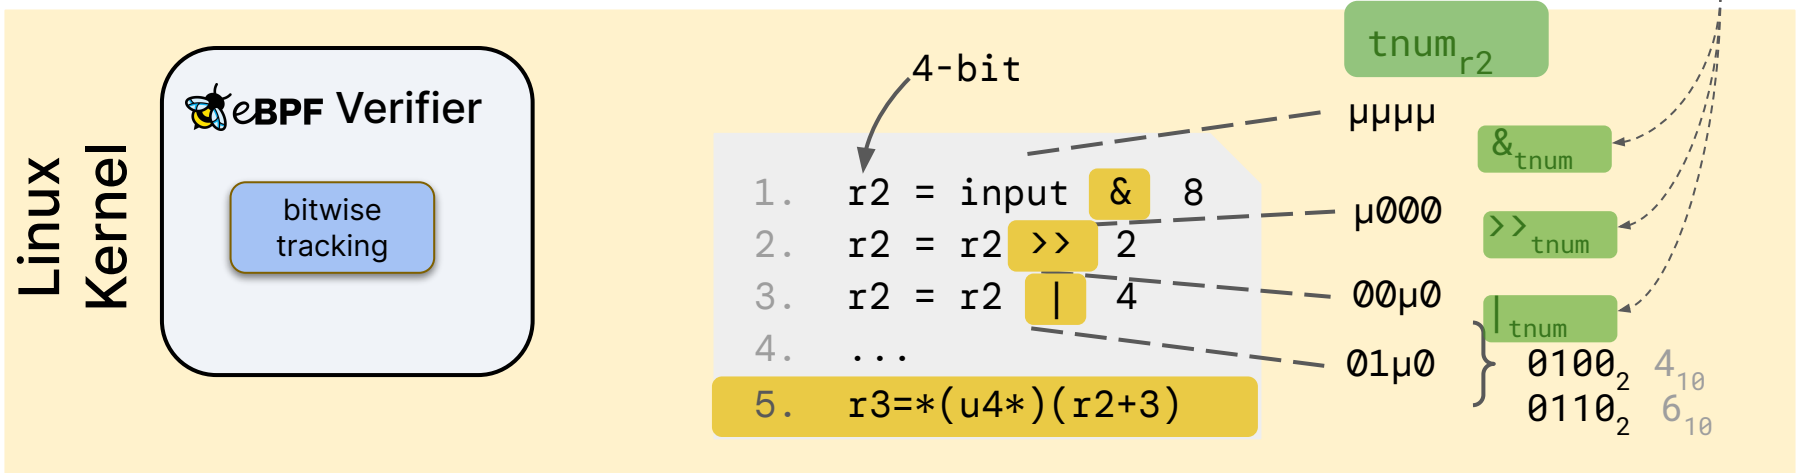
Static Analyses in the eBPF verifier



- Tracking the **values** of program variables across **all** executions of the program
- Our work: Reasoning about the soundness of the range analysis + bitwise tracking + their combination

Bitwise Tracking

- Task: track the **values** of program variables across **all** executions
 - Using abstract values from an abstract domain
- Bitwise domain: track **individual bits** of a program variable .
 - Kernel term: tristate numbers (tnums) $\{0, 1, \mu\}$



Bitwise Tracking

- Task: track the **values** of program variables across **all** executions
 - Using abstract values from an abstract domain
- Bitwise domain: track **individual bits** of a program variable .
 - Kernel term: tristate numbers (tnums) $\{0, 1, \mu\}$

```
1. struct tnum tnum_and(struct tnum a, struct tnum b)
2. {
3.     u6
4.     al
5.     be
6.     v = a.value & b.value;
7.     return TNUM(v, alpha & beta & ~v);
8. }
```

Is tnum_and sound?

&
tnum

>>
tnum

tnum

0100_2 4_{10}
 0110_2 6_{10}

Constraints for a Sound Tnum Operator

$P, Q \in \text{tnums}$

$x, y \in \text{integers}_{64}$:

$\text{member}(x, P) \quad \wedge$

$\text{member}(y, Q) \quad \wedge$

$z = x \ \& \ y \quad \wedge$

$R = P \ \&_{\text{tnum}} \ Q \quad \wedge$

$\neg \text{member}(z, R) \quad ?$

$01\mu 0 \quad \left. \begin{array}{l} 0100_2 \quad 4_{10} \\ 0110_2 \quad 6_{10} \end{array} \right\}$

$\text{member}(1, 01\mu 0) = \text{false}$

$\text{member}(4, 01\mu 0) = \text{true}$

Constraints for a Sound Tnum Operator

$P, Q \in \text{tnums}$

$x, y \in \text{integers}_{64}$:


$\text{member}(x, P) \quad \wedge$

$\text{member}(y, Q) \quad \wedge$


$z = x \ \& \ y \quad \wedge$

$R = P \ \&_{\text{tnum}} \ Q \quad \wedge$

$\neg \text{member}(z, R) \quad ?$




Range Analysis



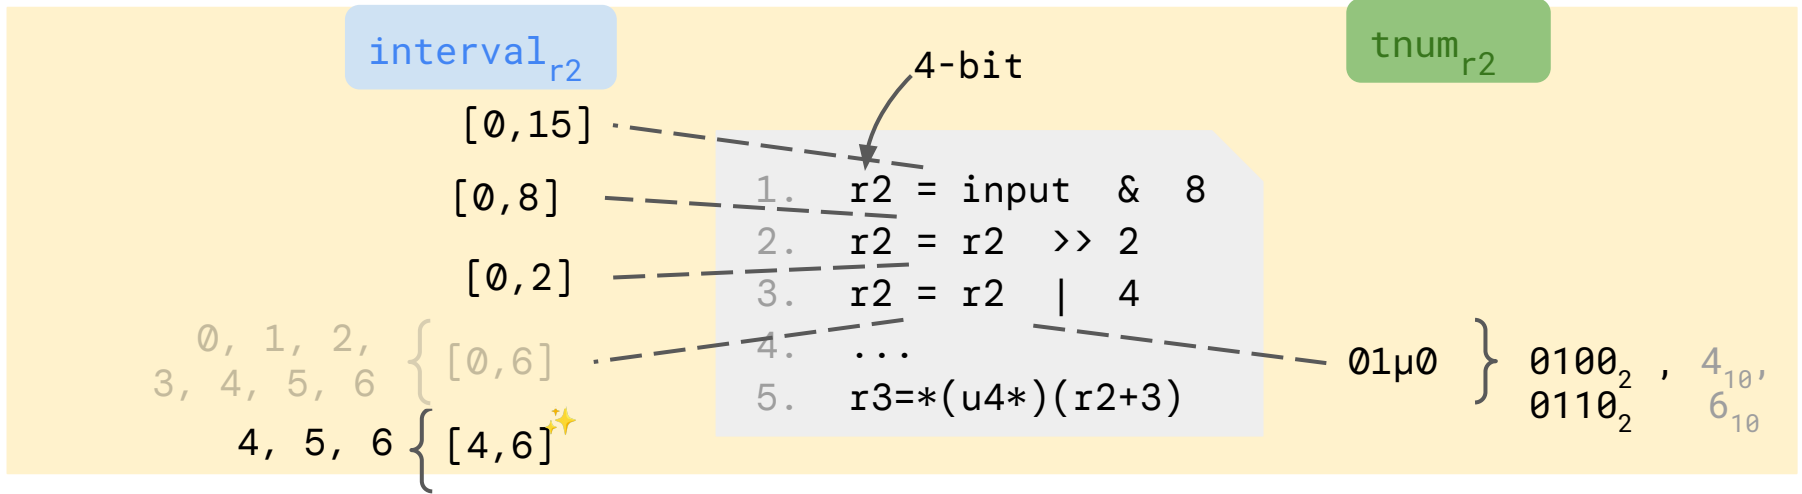
- Range Analysis: tracks range of possible values – [min, max]
 - Interval domain

Range Analysis



- Range Analysis: tracks range of possible values – [min, max]
 - Interval domain

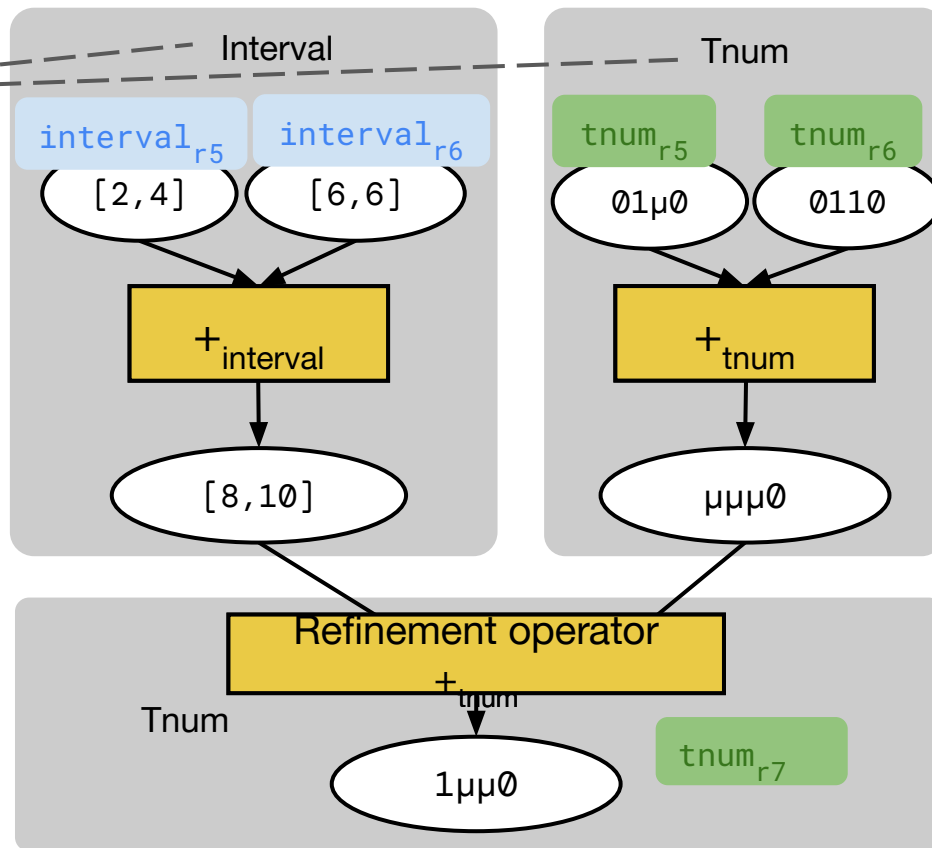
Range Analysis Refinement



- Range Analysis: tracks range of possible values – [min, max]
 - Interval domain
- Refinement: Abstract values in one domain can be used to refine abstract values in another domain

Modular Reasoning: The Usual Setting

```
42. ...  
43. r7 = r5 + r6  
44. ...
```

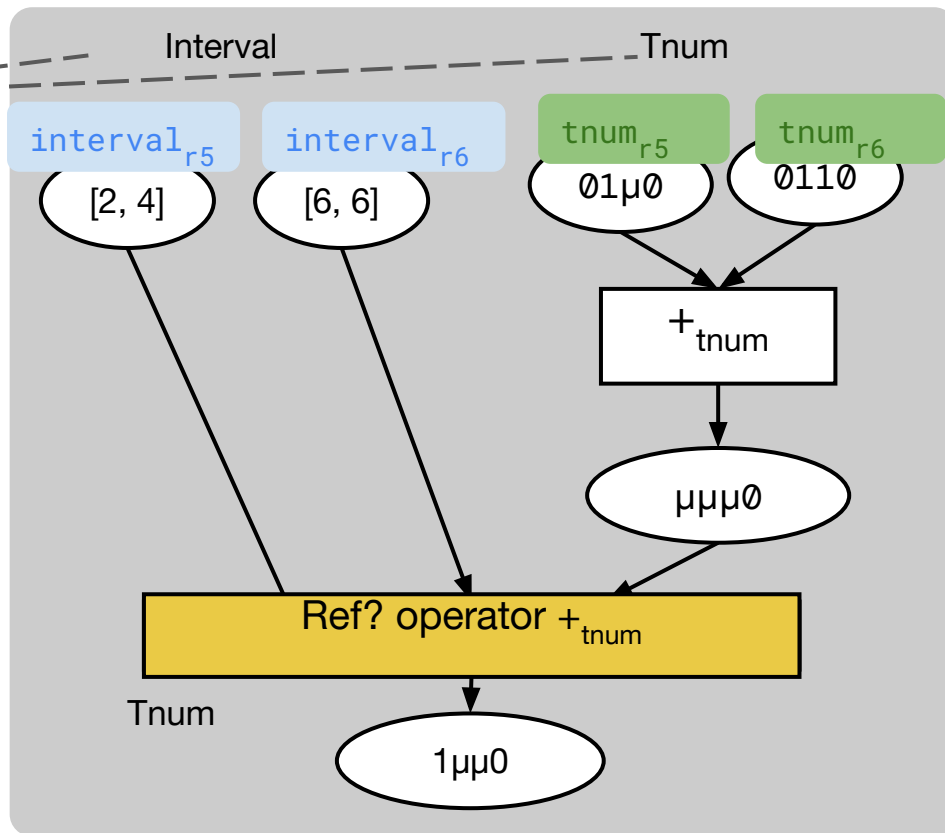


Soundness?

Modular reasoning

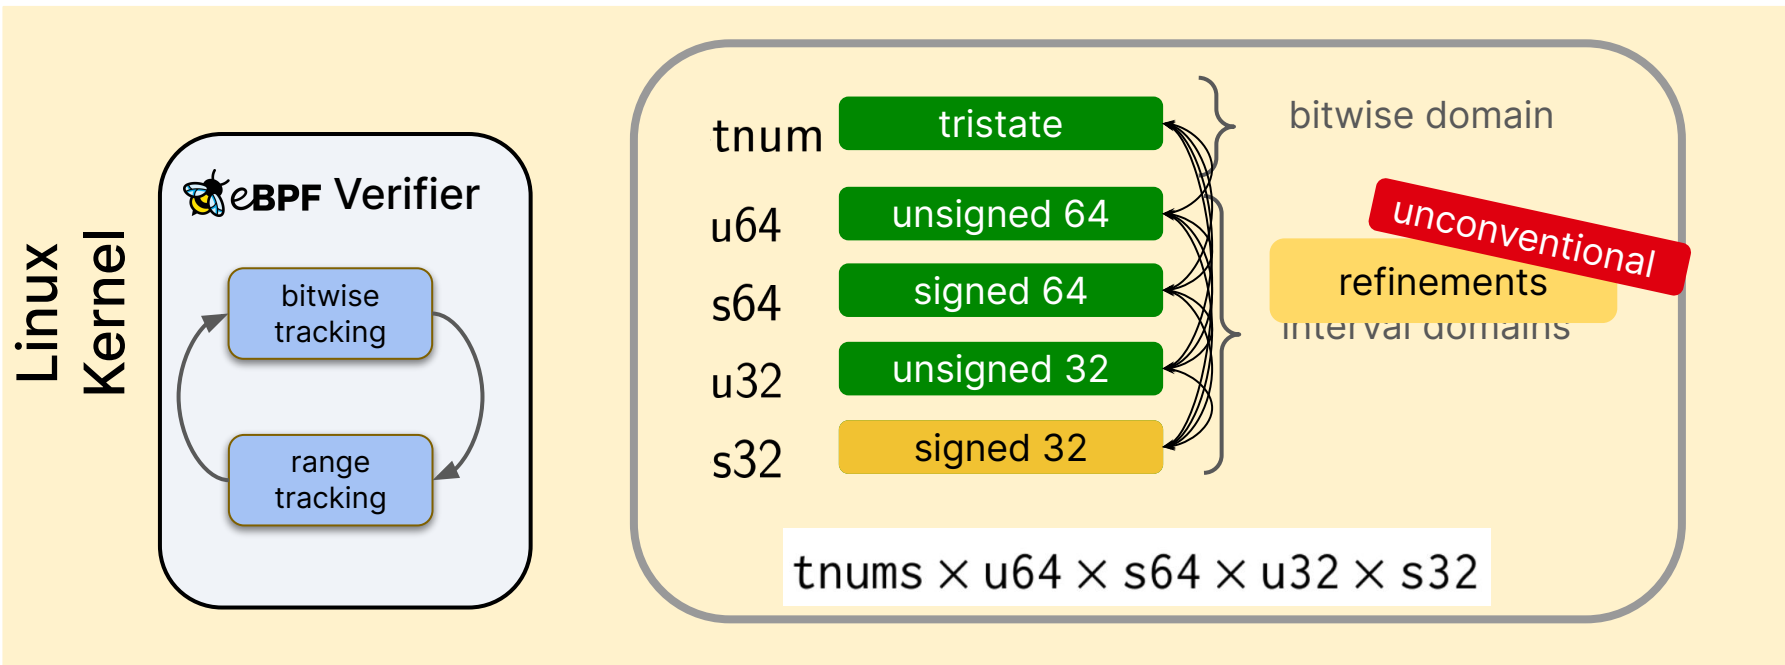
Non-Modular Reasoning: The eBPF Verifier

```
42. ...  
43. r7 = r5 + r6  
44. ...
```



"One-shot"
reasoning

Value Tracking Abstract Domains in the Linux Kernel



Constraints for Multi-Domain Soundness

$P, Q \in \text{tnums} \times \text{u64} \times \text{s64} \times \text{u32} \times \text{s32} :$

$x, y \in \text{integers}_{64} :$

$\text{member}(x, P) \quad \wedge$

$\text{member}(y, Q) \quad \wedge$

$z = x \ \& \ y \quad \wedge$

$R = P \ \& \ \text{tnums} \times \text{u64} \times \text{s64} \times \text{u32} \times \text{s32} \ Q \quad \wedge$

$\neg \text{member}(z, R)$

Constraints for Multi-Domain Soundness

$P, Q \in \text{tnums} \times \text{u64}$

$x, y \in \text{integers}_{64}$

$\text{member}(x, P)$

$\text{member}(y, Q)$

$z = x \ \& \ y$

$R = P \ \& \ \text{tnums} \times \text{u64} \times \text{s64} \times \text{u32} \times \text{s32} \ \cdot \ Q$

$\neg \text{member}(z, R)$

```
1. static int adjust_scalar_min_max_vals(...)  
2. {  
3.   switch (opcode) {  
4.  
5.     case BPF_AND:  
6.       scalar_min_max_and(dst_reg, &src_reg);  
7.       tnum_and(dst_reg->var_off, src_reg.var_off);  
8.       break;  
9.  
10.    case BPF_SUB:  
11.      scalar32_min_max_sub(dst_reg, &src_reg);  
12.    ...
```

\wedge

Scope of Automated Verification

- Arithmetic and Logic:
 - add, sub, or, and, lsh,
 - div, mod (Trivially sound)
 - mul **X**
- Jump:
 - ja, jeq, jgt, jge, jlt, jle
- Strict constraints: only report those than

Weaken constraints

true positive

✓ report

false positive

X do not report

```
archive mirror
search help / color / mirror / Atom feed
Sunhao.th@gmail.com>
Starovoitov <ast@kernel.org>,
Daniel Borkmann <daniel@iogearbox.net>,
John Fastabend <john.fastabend@gmail.com>,
Andrii Nakryiko <andrii@kernel.org>,
Martin KaFai Lau <martin.lau@linux.dev>,
Song Liu <song@kernel.org>,
Yonghong Song <yonghong.song@linux.dev>,
KP Singh <kpsingh@kernel.org>,
Stanislav Fomichev <sdf@google.com>, Hao Luo <haoluo@google.com>,
Jiri Olsa <jolsa@kernel.org>
Cc: bpf <bpf@vger.kernel.org>,
Linux Kernel Mailing List <linux-kernel@vger.kernel.org>
Subject: bpf: shift-out-of-bounds in tnum_rshift()
Date: Tue, 24 Oct 2023 14:40:04 +0200 [thread overview]
Message-ID: <CACKBjsY2q1_fUohD7hRmKqV1MV=eP2f6XK8kjkYNw7BaiF8iQ@mail.gmail.com> (raw)

Hi,

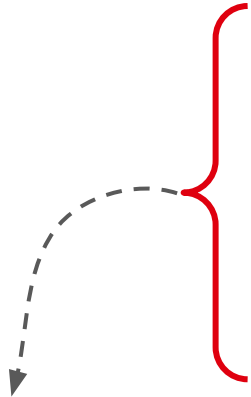
The following program can trigger a shift-out-of-bounds in
tnum_rshift(), called by scalar32_min_max_rsh():

0: (bc) w0 = w1
1: (bf) r2 = r0
2: (18) r3 = 0xd
4: (bc) w4 = w0
5: (bf) r5 = r0
6: (bf) r7 = r3
7: (bf) r8 = r4
8: (2f) r8 *= r5
9: (cf) r5 s>>= r5
10: (a6) if w8 < 0xffffffffb goto pc+10
11: (1f) r7 -= r5
12: (71) r6 = *(u8*)(r1 +17)
13: (5f) r3 &= r8
14: (74) w2 >>= 30
15: (1f) r7 -= r5
16: (5d) if r8 != r6 goto pc+4
17: (c7) r8 s>>= 5
```

Results

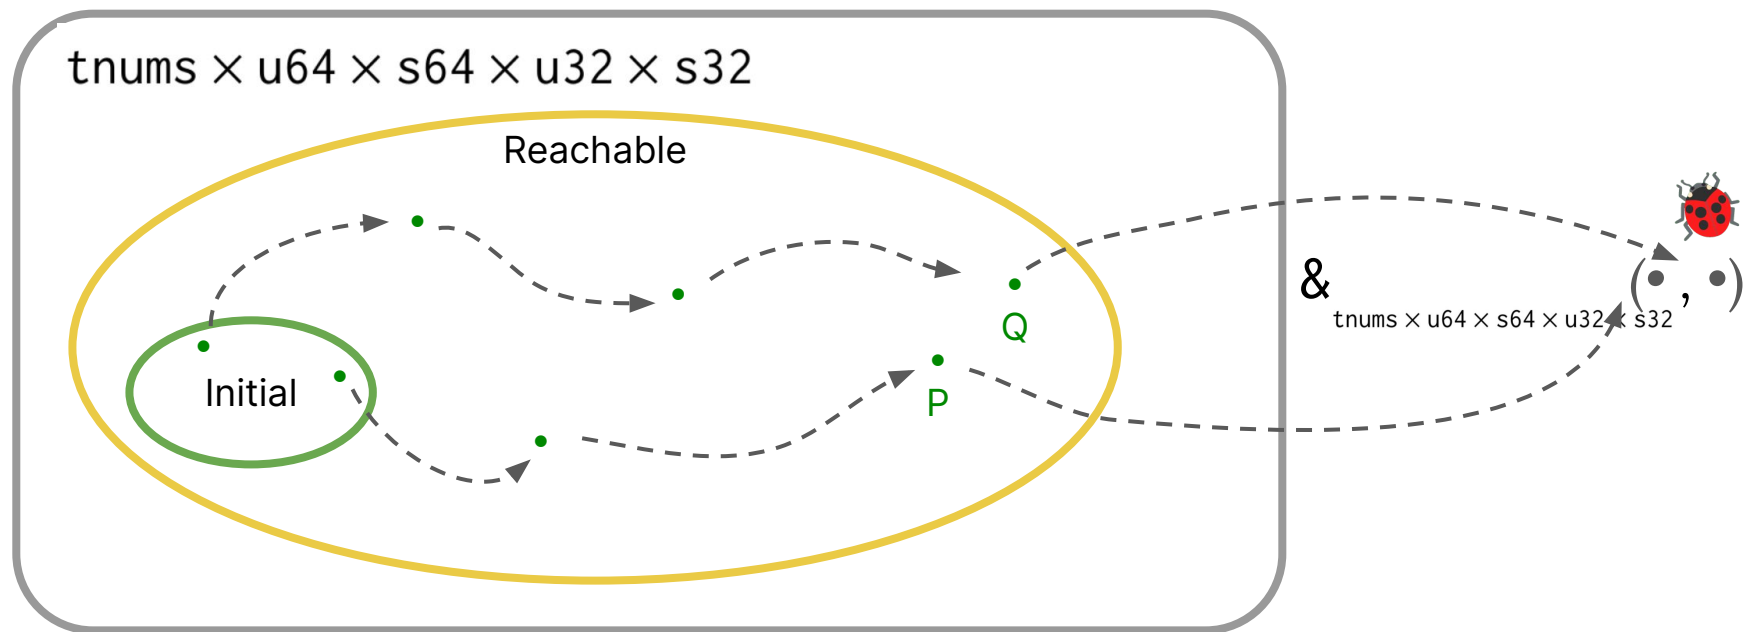
- Proved that all abstract operators in kernels starting from v5.13 are sound
- What can we do about unsound versions?
 - False positive?
 - Are these bugs possible via real eBPF programs?

Generate *actual* eBPF programs!






Kernel Version	Sound?
v4.14	✗
v5.5	✗
v5.7	✗
...	✗
v5.12	✗
v5.13	✓
v5.14	✓
...	✓
v5.19	✓




Challenges in Synthesizing eBPF Programs that Manifest Bugs



Features of Agni's eBPF Program Synthesis

- Technique: A combination of an **enumerator** and a **solver**
- Supported instructions:
 - Arithmetic and Logic: add, sub, or, and, lsh, rsh, xor, arsh
 - **Jump** ja, jeq, jgt, jge, jlt, jle, jset, jne, jsgt, jsge, jslt, jsle

```
1. r1 =   
2. r2 =   
3. r3 =   
4. if r1 > r2 goto +1  
5. exit  
6. r3 = r1 & r2
```


```
1. r1 =   
2. r2 =   
3. r3 = r1 + r2  
4. r4 = r3 & 
```


Demo

Results


Agni: A verification tool for eBPF Range Analysis

Kernel Versions Tested



Proved versions
5.13 through 5.19
sound

Soundness Violations Across Kernel Versions



Part of 8
documented
CVEs

eBPF Programs Synthesized

- Produced a POC for ~97% of soundness violations across kernel versions

Next Steps

- At a high-level: making Agni as push-button as possible
- Exploring the possibility of using Agni in Linux eBPF CI
- Reducing Verification Time
 - Initial work to parallelize verification of each instruction
- Reliable Code Extraction
- Lot's more to do!
- Possibilities for future verification:
 - Symbolic state pruning
 - eBPF helper function interface
- Ask me for a demo!

Hari harishankar.vishwanathan@rutgers.edu

Matan mys35@cs.rutgers.edu

Srinivas srinivas.narayana@rutgers.edu

Santosh santosh.nagarakatte@cs.rutgers.edu

<https://github.com/bpfverif/agni>