

Formally Verified Sandboxing for Packet Processing Programs



Srinivas Narayana and Santosh Nagarakatte

<https://people.cs.rutgers.edu/~sn624/verified-sandboxing.html>

Award 2019302

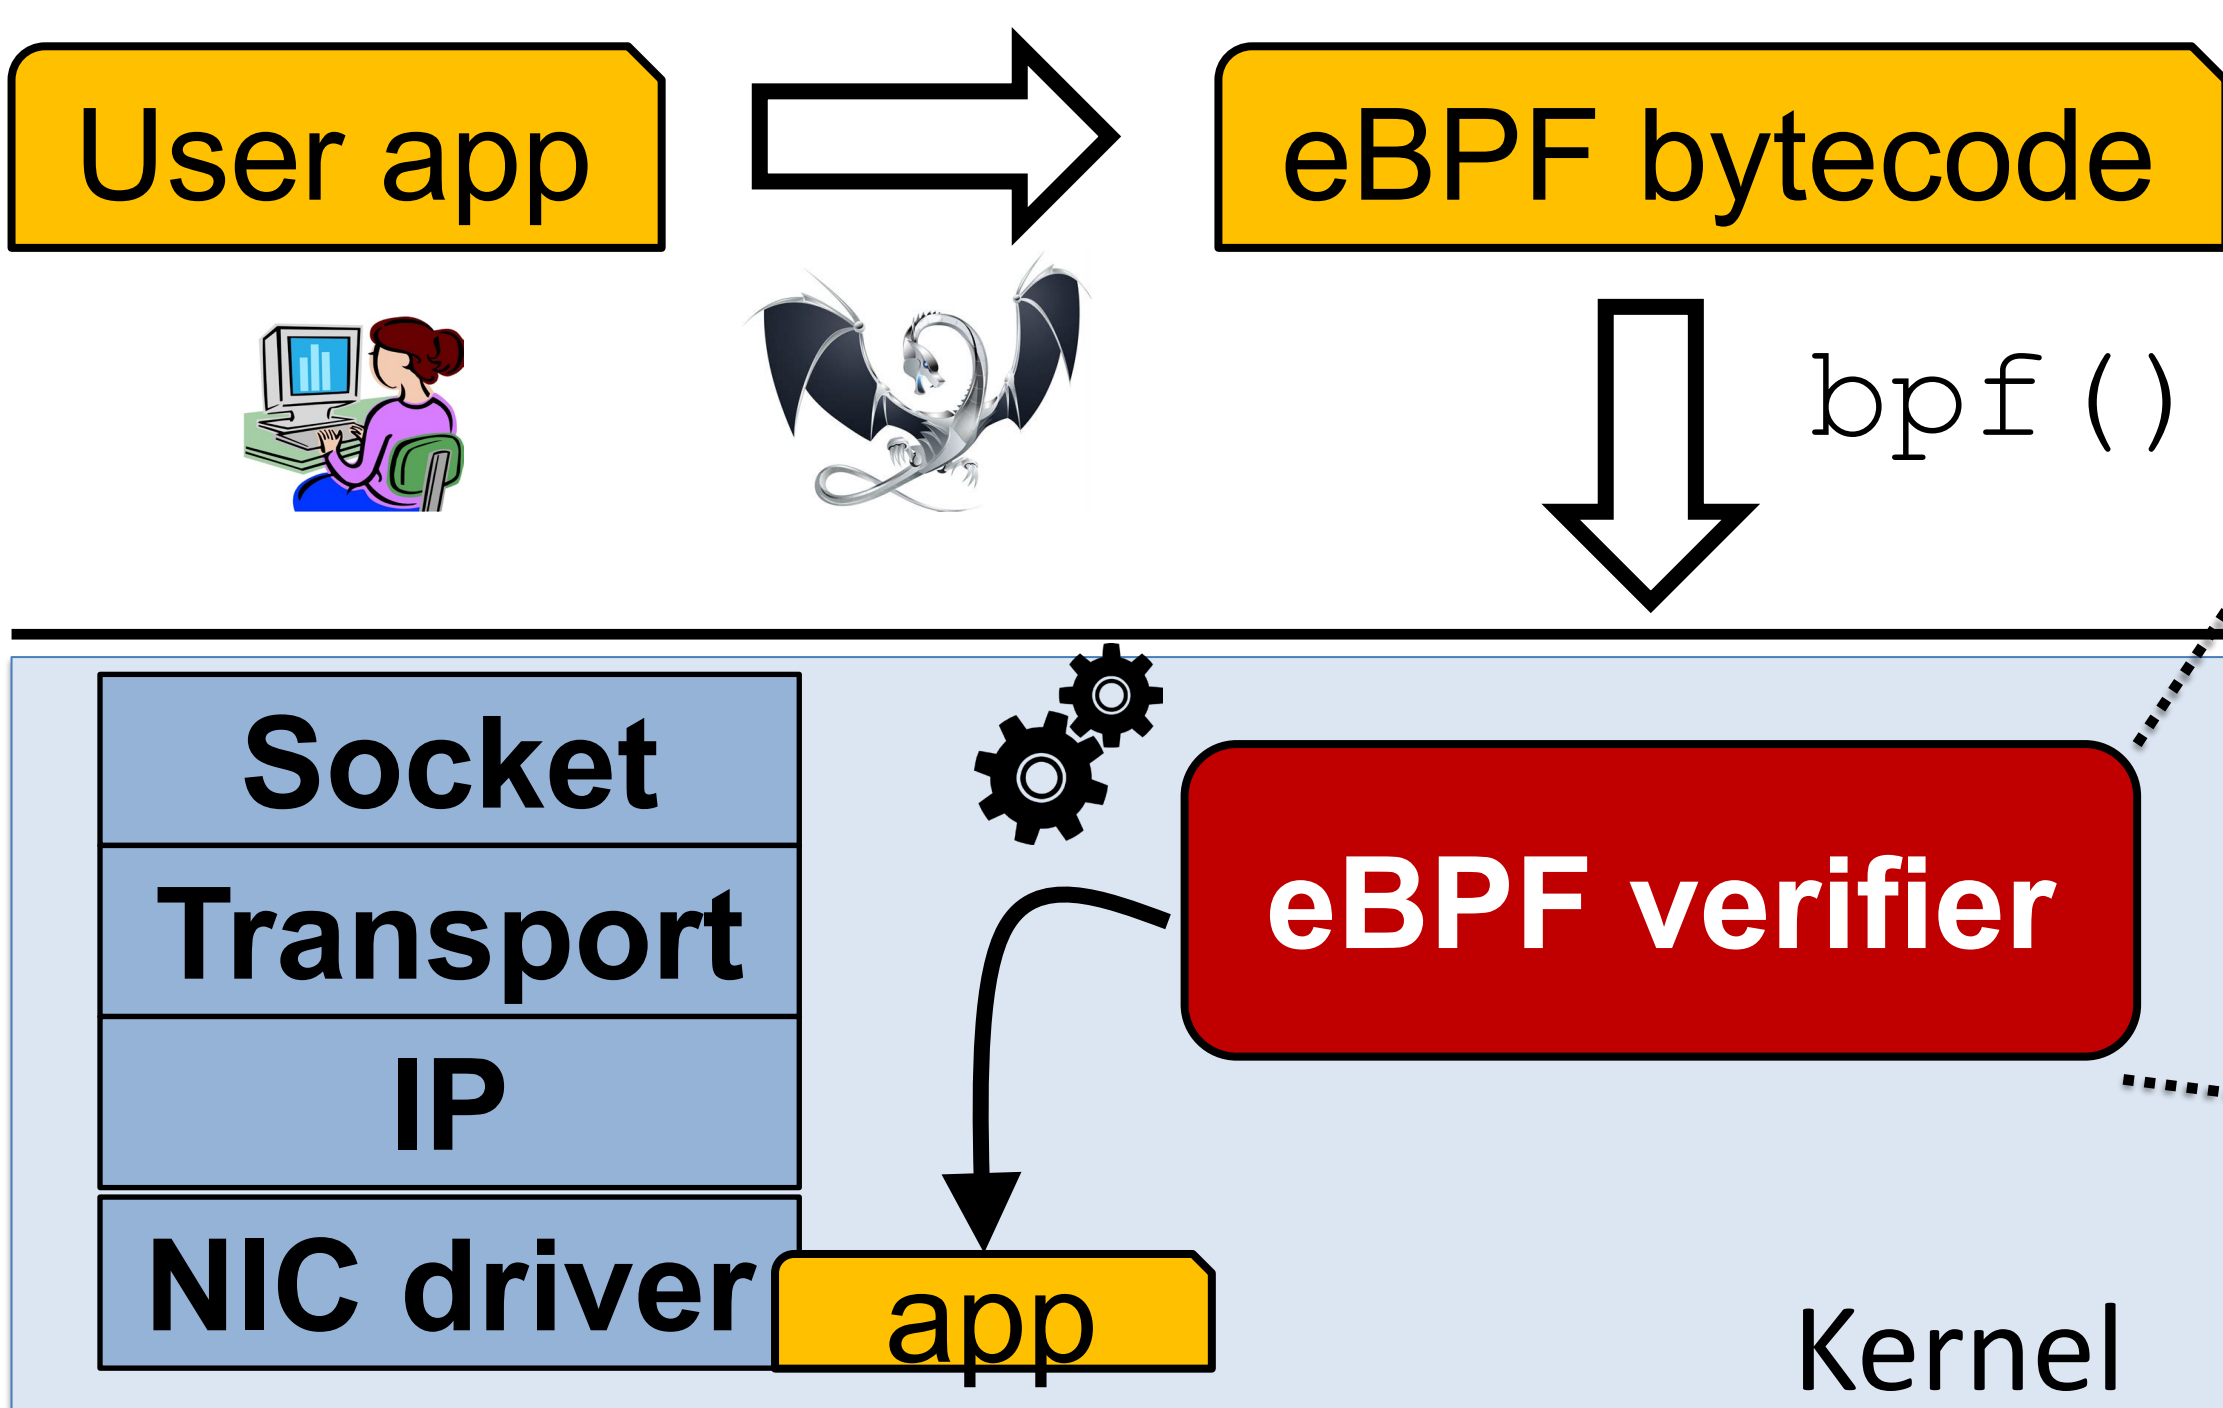
Linux kernel extensions for networking using eBPF are widely deployed.

They are statically checked for safety before running in the kernel.

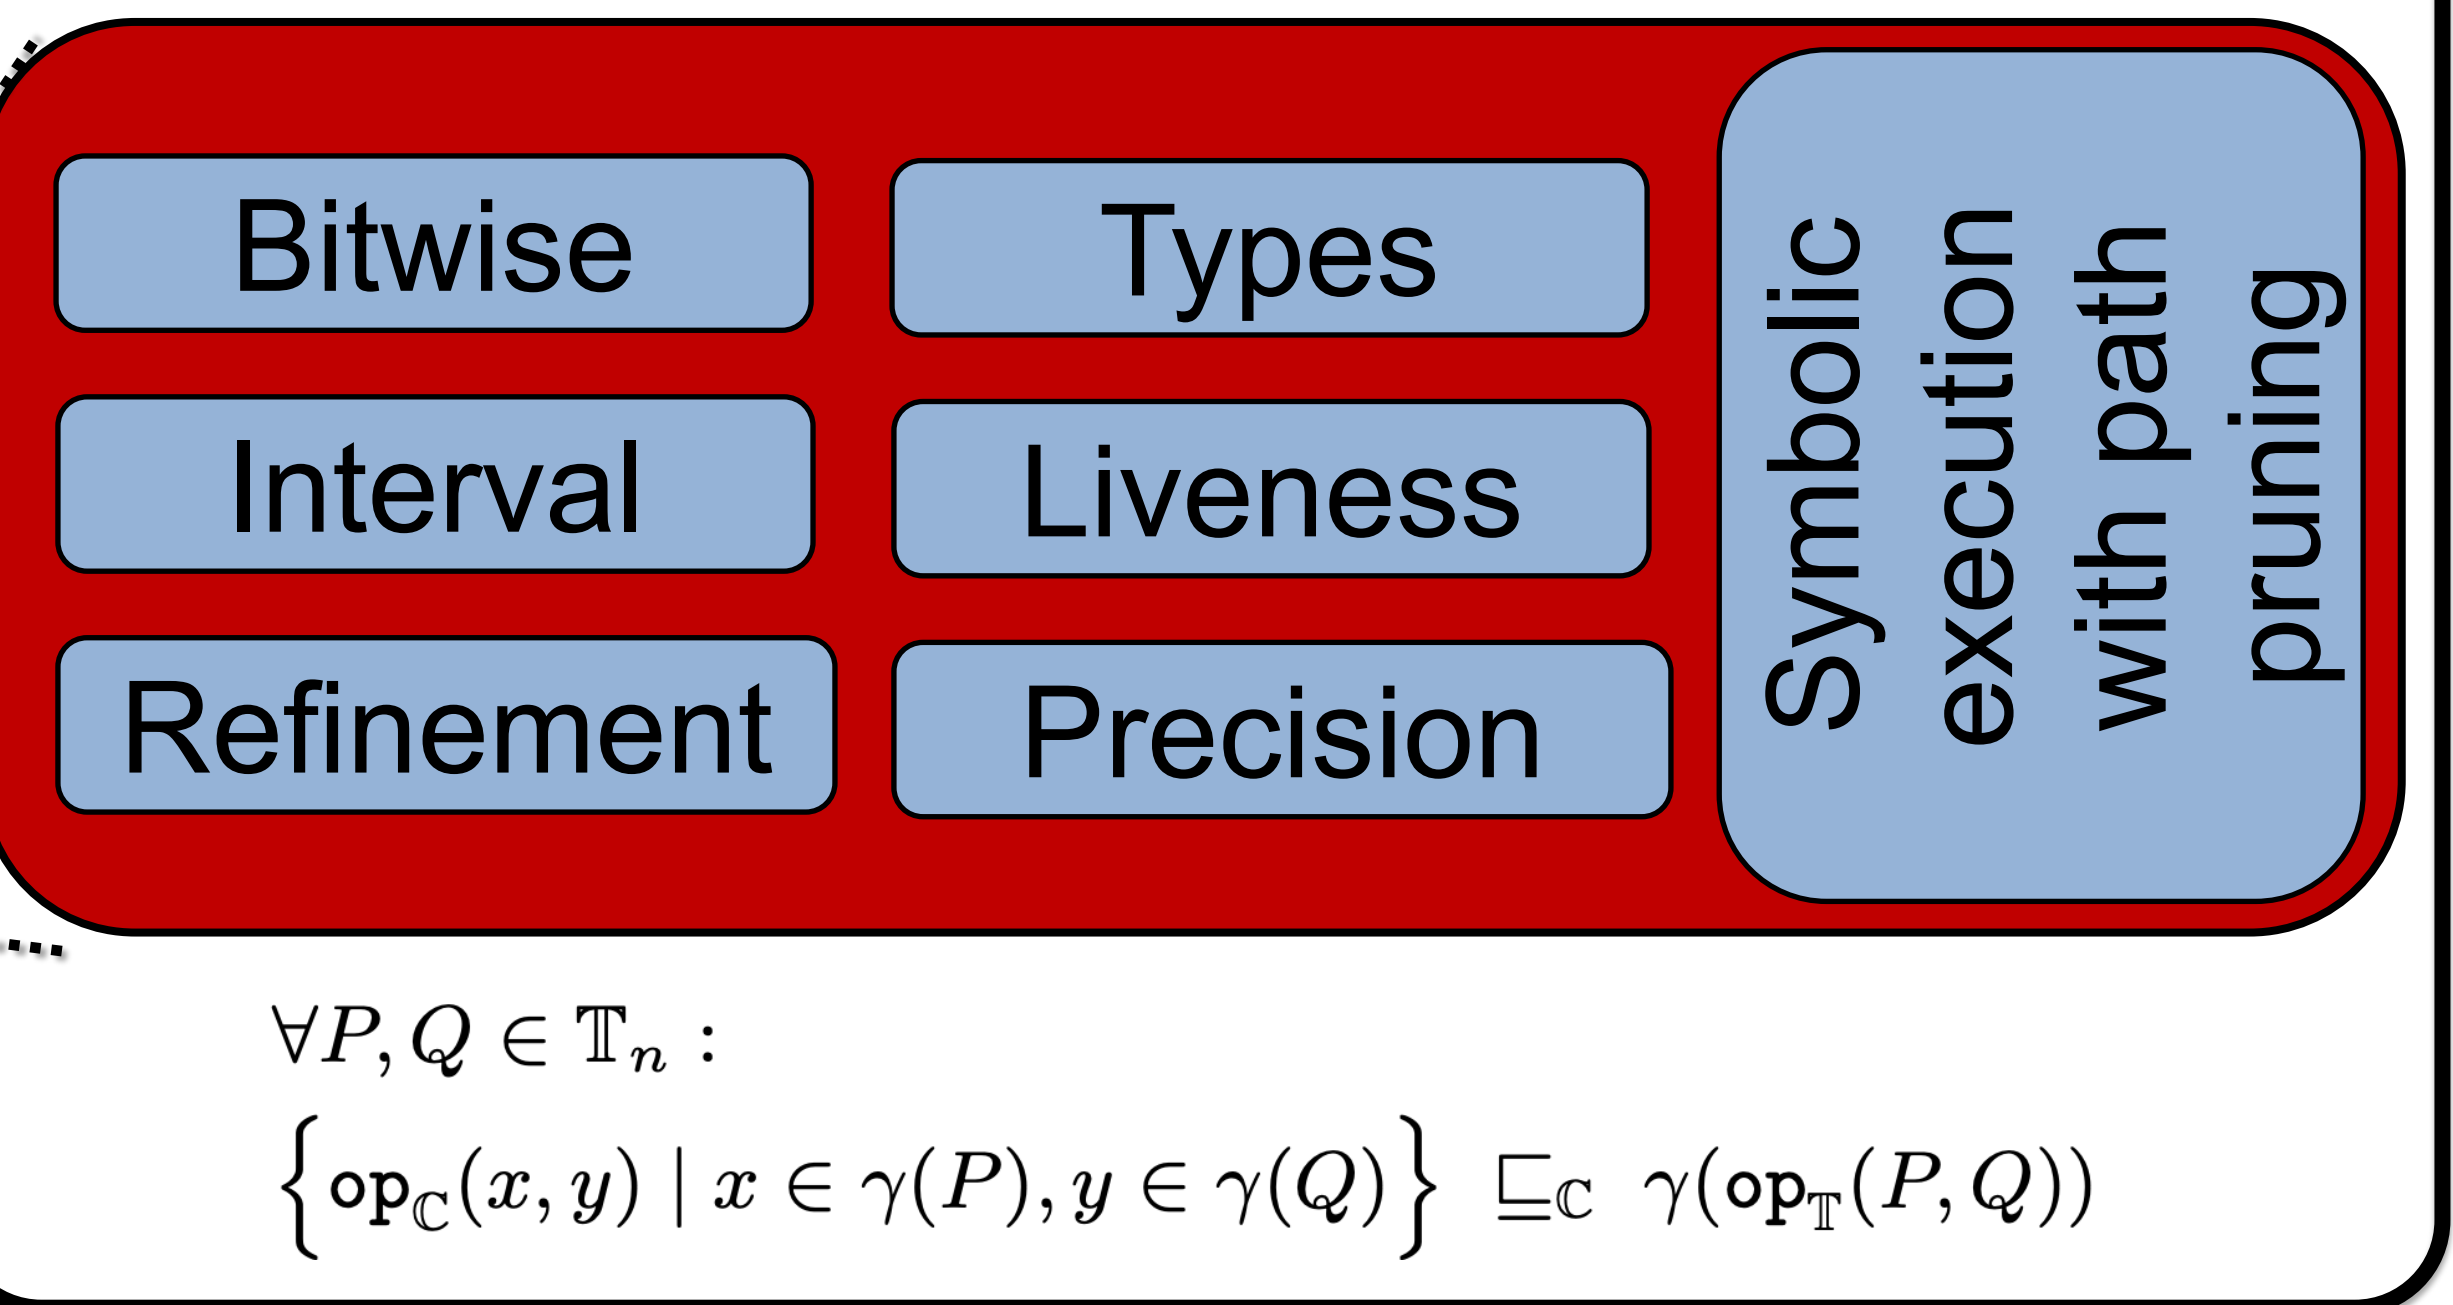
However, the in-kernel verifier has bugs, opening the kernel up to attack.

Our goal is to design a sound eBPF verifier for Linux.

Context: eBPF and Verification



Approach: Sound Abstract Interpretation



Bitwise abstractions (CGO'22)

Tri-state domain

$tnum\ P : 10\mu 0\ (1000, 0010)$

$tnum\ Q : 10\mu 1\ (1001, 0010)$

Prove sound+optimal add

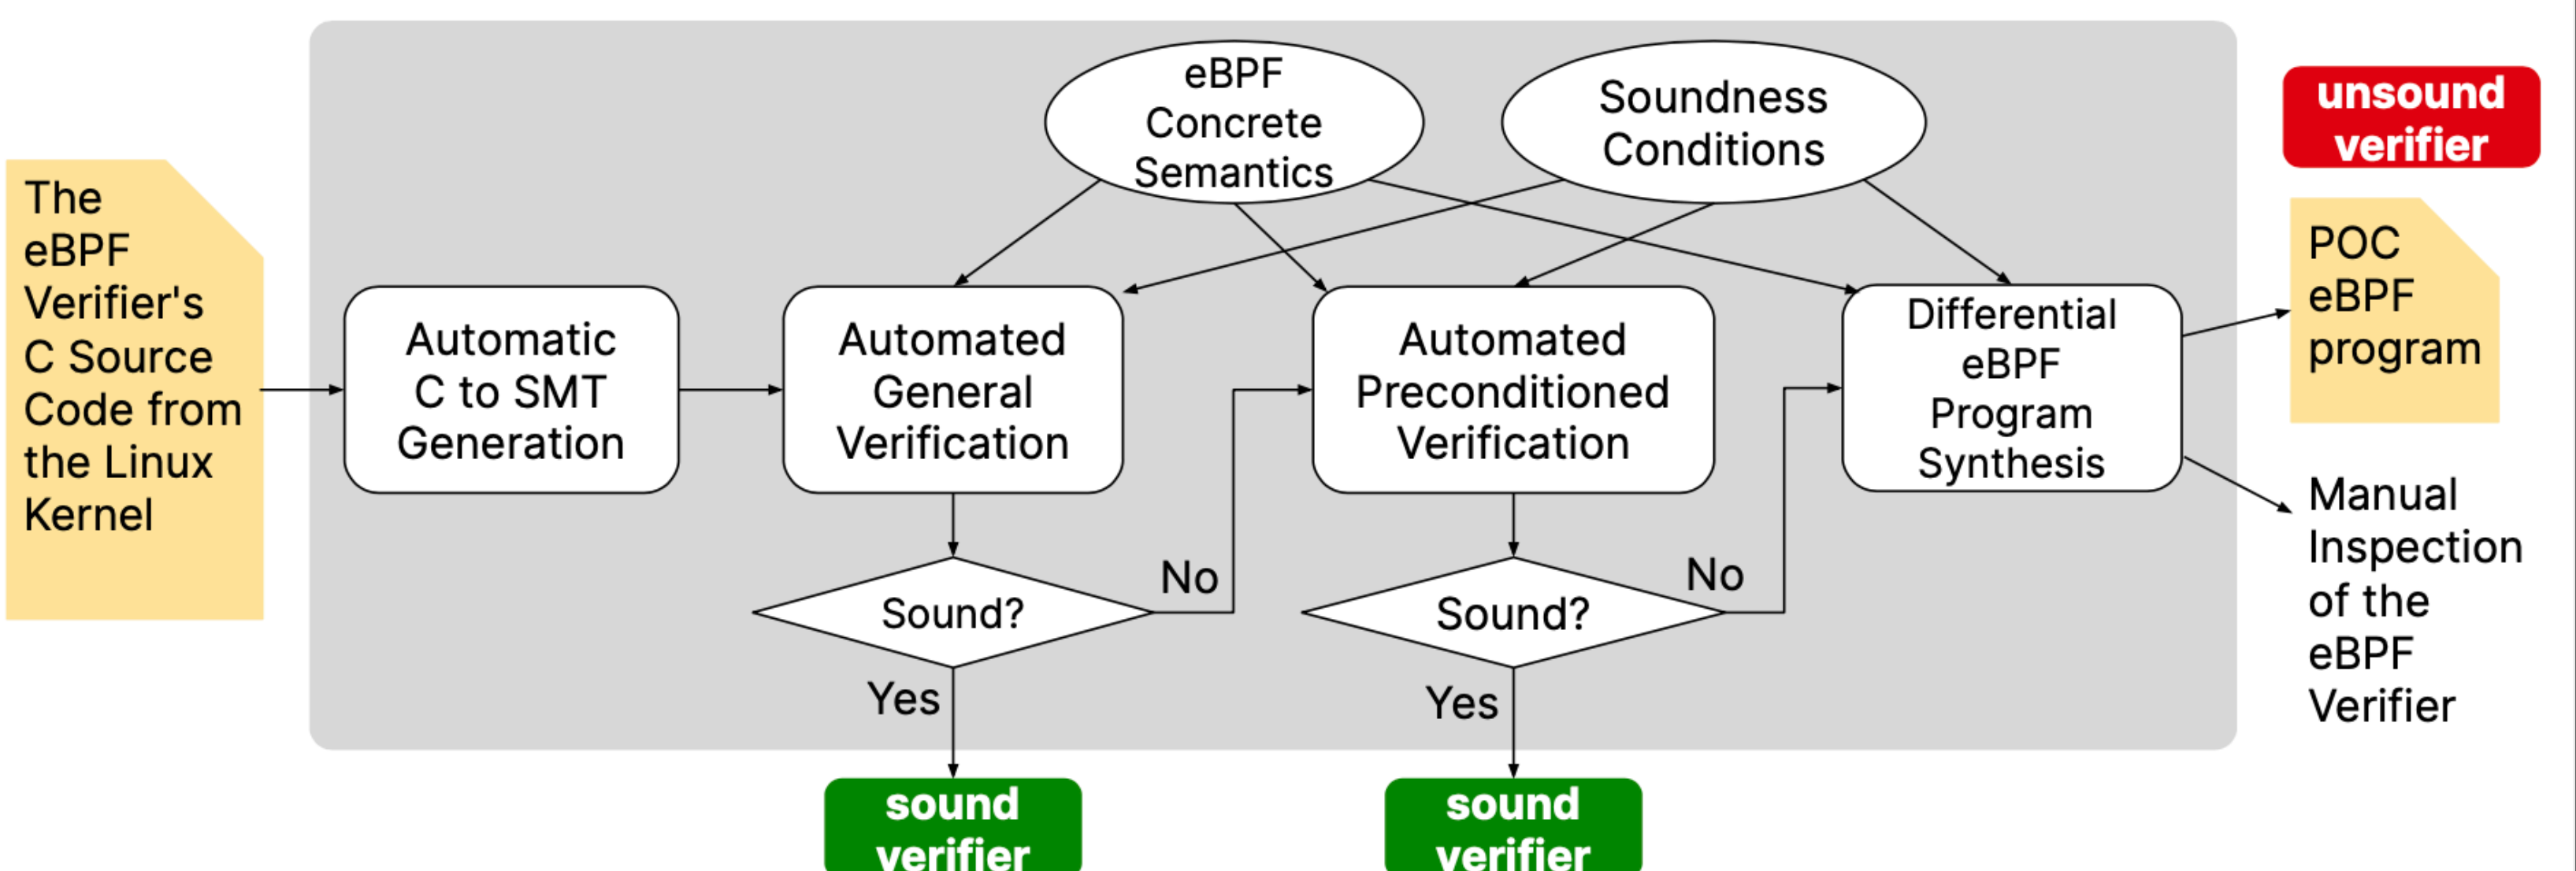
```
def tnum_add(tnum P, tnum Q):
    u64 sv := P.v + Q.v
    u64 sm := P.m + Q.m
    u64 \Sigma := sv + sm
    u64 \chi := \Sigma \oplus sv
    u64 \eta := \chi | P.m | Q.m
    tnum R := tnum(sv & ~\eta, \eta)
    return R

def tnum_mul(tnum P, tnum Q):
    ACC := tnum(P.v * Q.v, 0)
    ACC := tnum(0, 0)
    while P.value or P.mask:
        # LSB of tnum P is a certain 1
        if (P.v[0] == 1) and (P.m[0] == 0):
            ACC := tnum_add(ACC, tnum(Q, Q.m))
        # LSB of tnum P is uncertain
        else if (P.m[0] == 1):
            ACC := tnum_add(ACC, tnum(Q, Q.v|Q.m))
        # Note: no case for LSB is certain 0
        P := tnum_rshift(P, 1)
        Q := tnum_lshift(Q, 1)
    tnum R := tnum_add(ACC, ACC)
    return R
```

New sound multiplication (faster, more precise than kernel's algorithm)

Intervals, multi-domain refinements (CAV'23)

Sound modular operators (SAS'24)



Formalization of multi-domain abstractions in kernel Toolchain (Agni) to extract SMT from eBPF operators. Previously unknown bugs & simpler PoCs for existing

Broader Impacts

Technology transfer: Patches upstreamed to Linux mainline kernel (algorithm, verification fixes)

Outreach: Linux Plumbers Conference '23,'24, eBPF workshop SIGCOMM '23, '24. Actively used as CI by devs

Training: Grads (Hari Vishwanathan, Matan Shachnai) & UGs

```
bpf-next: Avoid goto in regs_refine_cond_op()
In case of GE/GT/SGE/JST instructions, regs_refine_cond_op()
re bpf, tnums: Provably sound, faster, and more precise algorithm for tnum_mul
Th This patch addresses a new class of bugs for multi-branches of tnum_mul
nu bpf: Harden and/or/xor value tracking in verifier
c4 co
```

31 workflow run results			
Event	Status	Branch	Actor
End-to-End Tests	End-to-End Tests #74: Scheduled	main	18 hours ago ...
End-to-End Tests	End-to-End Tests #66: Scheduled	main	2 days ago ...

