The eBPF Runtime in the Linux Kernel

Bolaji Gbadamosi
Karlstad University

Luigi Leonardi
University of Pisa

Tobias Pulls
Karlstad University

Toke Høiland-Jørgensen
Red Hat

Simone Ferlin-Reiter
Red Hat

Simo Sorce
Red Hat

Anna Brunström
Karlstad University

Abstract

Extended Berkeley Packet Filter (eBPF) is a runtime that enables users to load programs into the operating system (OS) kernel, like Linux or Windows, and execute them safely and efficiently at designated kernel hooks. Each program passes through a verifier that reasons about the safety guarantees for execution. Hosting a safe virtual machine runtime within the kernel makes it dynamically programmable. Unlike the popular approach of bypassing or completely replacing the kernel, eBPF gives users the flexibility to modify the kernel on the fly, rapidly experiment and iterate, and deploy solutions to achieve their workload-specific needs, while working in concert with the kernel.

In this paper, we present the first comprehensive description of the design and implementation of the eBPF runtime in the Linux kernel. We argue that eBPF today provides a mature and safe programming environment for the kernel. It has seen wide adoption since its inception and is increasingly being used not just to extend, but program entire components of the kernel, while preserving its runtime integrity. We outline the compelling advantages it offers for real-world production usage, and illustrate current use cases. Finally, we identify its key challenges, and discuss possible future directions.

1 Introduction

Contemporary monolithic operating systems like Linux are designed to be general-purpose, and cater to a wide variety of users. They define the necessary abstractions to multiplex and share hardware resources safely and efficiently. As such, their design choices play a key role in influencing an application's performance, scalability, and security.

Evidently, the monolithic design choice of the Linux kernel does come with a cost such as increased complexity and maintenance challenges due to tightly coupled components, higher security risks from a larger attack surface [32], reduced scalability and flexibility in resource-constrained environments.

As a result of these challenges, the approach of bypassing or replacing the OS kernel has gained momentum. Kernel bypass solutions [53, 81, 89] and library OS [77] allow specializing the entire OS stack for a specific workload, giving significant performance improvements. However, these solutions may inhibit resource multiplexing by taking complete control of the hardware. They may also require application logic to be rewritten, and give up traditional benefits of an OS like the security and isolation model.

Users managing a large fleet of machines care deeply about extracting maximum performance and utilization out of their hardware, while running their production workloads on a battle-tested OS kernel like Linux to reduce their maintenance burden. Ideally, they wish to have comparable performance to prior approaches without abandoning their well understood tooling for performance monitoring, administration, and management.

To this end, another direction is to explore tailoring the kernel mechanisms and policies for specific workloads, which can yield drastic improvements. However, this is far from trivial on Linux. It involves changing the kernel code, and deploying workload specific kernel changes to a large set of machines, which is impractical at scale, due to the huge variety of target applications, the complexities of frequent kernel redeployments. There has been various research projects for safer ways to extend the kernel. For instance, VINO (Virtual Integrated Network Operating system) [84] and SPIN [34] allow users to customize their kernels through user-defined functions, SPIN is written in Modula-3 language to make sure that dynamically loaded modules are well-secured and efficient. On the other hand, VINO primarily focus on how extensions can be put together safely via fault isolation techniques

1

within them. However, the TockOS (Tock Operating System) [44] is a more recent initiative that uses Rust to improve security and dependability in its microkernel design for embedded systems. Unlike these systems which often require new operating environments or adjustments, eBPF (Extended Berkeley Packet Filter) allows dynamic programming capabilities directly into monolithic kernels like Linux and Windows. As opposed to programs such as VINO, SPIN and TockOS, this integration enables eBPF develop secure enhancements without requiring a new kernel structure. Introduced in the kernel 3.18 (released on December, 2014) **eBPF** functions as a **safe** programmable virtual machine hosted on top of a **performant** in-kernel runtime. It allows users to write programs and load them into the kernel, and attach them at designated hooks to begin execution. To ensure safety, every program is statically checked by a **verifier** when being loaded. To ensure performance, all programs are **Just-in-time (JIT) compiled** to native machine instructions.

In essence, eBPF makes the Linux kernel *dynamically programmable at runtime*, while ensuring its runtime integrity remains intact. However, thus far, since its initial release until version 6.7 (released on January, 2024), which this paper is based on, there exists no complete description of the design and implementation of eBPF within the Linux kernel. Hence, through this paper, we make the following contributions:

- A comprehensive description of the design and implementation of the eBPF runtime in the Linux kernel up until version 6.7.

- An exhaustive characterisation of eBPF's safety properties (§5).

- Identification of limitations and key challenges concerning eBPF's current design (§11).

The rest of this paper is organized as follows. §2 presents background, and §3 is about the general overview of eBPF program. In §4 we illustrated the high-level programming and execution model of eBPF. We discuss the four major passes involved during the verification of an eBPF program in §6,§7,§8,§9. We then present some use cases in §10, challenges in §11 and conclude the paper in §12.

For brevity, the rest of the paper uses 'eBPF' to refer to the Linux runtime, as is common in the Linux kernel community.

2 Background and Design Principles

In recent years, Linux kernel customization has become a prevalent need across different sectors, with regard to performance, security, and observability goals. The last two decades have seen Linux customization evolve from being only a peripheral consideration to a critical necessity. in order to align the kernel behavior with various operational requirement.

2.1 Challenges with Kernel Customization

Developers often encounter significant hurdles when they need to make direct changes to the Linux kernel. Below are some of the key challenges.

Changing the Kernel Linux provides a huge number of configuration knobs both during compilation and runtime [11, 19]. However, these knobs do not fundamentally address performance bottlenecks or provide insight into the kernel's behavior. They are appropriate for tuning parameters of existing kernel policies, but not expressive enough for encoding new ones.

In such cases when customizing the kernel is warranted, developers would need to make changes to the kernel code, and/or write kernel modules. However, this makes testing and debugging changes much harder. Making involved changes to the kernel requires deep familiarity with its complex codebase, incurring a huge maintenance burden. Unless such changes are accepted by the kernel community, they need to be forward ported on kernel upgrades due to unstable kernel APIs. Any uncaught bugs in the code could easily lead to *system crashes* and take down production servers, which directly translates to downtime and lost revenue.

Deploying Kernel Changes Deploying a change to the kernel, however small, across a fleet of servers is a lengthy process [52]. Replacing the running kernel with a new one involves disruption of workloads hosted on the machine, as they must be stopped. The machine must then boot into the new kernel, and reinitialize all services again. This cycle is particularly harmful for workloads which incur cold starts and have a significant ramp-up time. After all these steps, an extensive testing phase needs to be carried out to test and qualify these changes, while fixing any detected regressions. This process has to be repeated for every server which receives a kernel update, making the process expensive and taking the total time to rollout a new kernel to the fleet in the order of weeks or months.

2.2 Alternative Approaches and Limitations

As mentioned in §2.1 developers face complexities when making changes to the kernel. We discuss alter-

native methods that promise performance benefits but also pose limitations and maintenance challenges.

Completely Bypassing or Replacing the Kernel. Although kernel bypass frameworks and library operating systems specialized for a particular workload are appealing in terms of performance, they come with their own downsides. Kernel bypass solutions require taking complete control of the hardware, and waste CPU cycles in busy polling. Library operating systems are tailored for specific workloads, and require code modifications to adapt to different requirements. In addition, both of these approaches make it difficult or impossible for two workloads to coexist and share hardware resources, which is not acceptable for some users. Finally, they require rewriting application logic, which is a significant maintenance burden.

eBPF achieves these requirements by allowing efficient customization of the Linux kernel safely and dynamically at runtime, reducing the risk of introducing kernel bugs when making changes to the kernel, and accelerating development and deployment velocity. Prior to eBPF's introduction, the kernel had multiple domain-specific register-based virtual machines serving dedicated use cases in the networking subsystem [36], but none of them were designed to be general-purpose[1].

2.3 Design Principles

We now outline the key design principles behind eBPF's continued development and evolution.

Safe and Dynamic OS Customization Hosting a safe and efficient virtual machine runtime within the kernel allows customizing the kernel dynamically, achieving performance similar to existing in-kernel code, while ensuring that the kernel's integrity is not undermined under any circumstances. Programming against a safe virtual machine with a constrained environment allows lowering the barrier to making modifications to the kernel. The guarantee of safety instills greater confidence when such programs are deployed to the kernel, as opposed to direct changes to the kernel or kernel modules reducing the risk of introducing kernel bugs that could lead to service outages.

Rapid Deployment and Upgrades Changing the kernel and deploying it across a large fleet of servers is significantly expensive, since the reboot of a machine incurs the cost of killing and reinitializing the

workload. Loading programs into the kernel dynamically at runtime would allow for rapid deployment across a fleet without any service disruptions. At the same time, fixing bugs and reiteration incorporating feedback based on collected telemetry becomes much faster by unloading and reloading the program at runtime. This significantly simplifies deployment, accelerates the feedback loop of testing and qualifying kernel changes, and speeds up development.

Integration with the Kernel While programs attach to the kernel to introduce alternative behavior, they should be able to interact safely with existing kernel state, and manipulate it if needed. At the same time, programs should have the option to fallback to the existing kernel processing in case they have nothing of interest to do. This presents a flexible model to the user, where the existing kernel implementation can be made use of, if needed, without having to reimplement similar code within the program.

3 Overview

eBPF is defined as an abstract virtual machine that supports the eBPF instruction set [73]. The virtual machine has a set of 11 registers and a fixed size stack. An eBPF program operates within a restricted virtual machine environment provided by the Linux kernel. The eBPF instruction set is a small but versatile collection of 64-bit instructions. These instructions provide a wide range of functionalities, enabling eBPF programs to efficiently perform various tasks within the kernel space [82]. The instruction set supports arithmetic operations such as addition, subtraction, multiplication, and bit operations (e.g, `AND`, `OR`, `XOR`). Additionally, it supports load and store instruction and jump instructions encompassing both conditional and unconditional jumps to alter the program flow, as well as function calls and exits. The instruction set also features atomic operations designed for safe memory access and modification. These atomic operations ensure that concurrent access does not lead to inconsistencies or corruption.

3.1 The eBPF Runtime

An eBPF runtime is the set of necessary components required to map the abstract virtual machine and related entities onto another OS or hardware platform. For the Linux kernel, the eBPF subsystem within the kernel implements the eBPF runtime and defines the system call interface to interact with it. Below, we present the key components that makes up the eBPF ecosystem as shown in Figure 1.

---

[1]Unlike the widely accepted misconception, eBPF's design was *not* influenced by Classic BPF [7, 65], and the name was only chosen for familiarity [22].
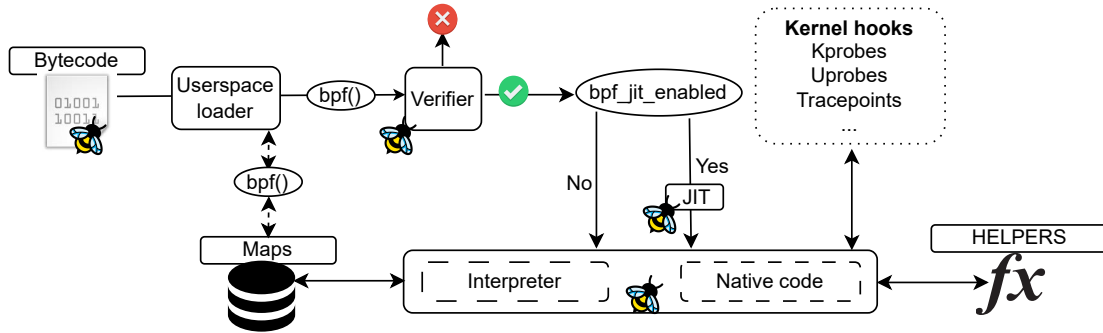
Figure 1: An overview of the eBPF key components and their correlation based on [24].

eBPF Bytecode   eBPF bytecode is defined as a finite sequence of eBPF instructions. The eBPF virtual machine executes `eBPF programs`, which are encoded using eBPF bytecode. Each program is composed of one or more `subprograms` (or subprogs in short). These are simply self-contained units of bytecode analogous to functions. Execution of a program begins at the `main` subprog.

eBPF Userspace Loader   eBPF has a vibrant ecosystem of user space loaders such as BCC [55], Bptrace [56] and libbpf [15] that loads eBPF bytecode into the kernel using the `BPF_PROG_LOAD` system call and attaches the program to relevant hooks while managing any corresponding maps [25]. A file descriptor is returned, which represents the program that has been loaded, which can then be used to attach the program to specific kernel hooks. For the purposes of this paper, we keep the discussion and expositions centered around the LLVM toolchain, its C frontend, and libbpf, which together remain the most popular and featureful reference implementations[2]

eBPF Verifier   The verifier, a crucial component in eBPF systems, inspects the bytecode before it is accepted into the kernel, ensuring safety properties associated with eBPF and that the loading of programs does not negatively impact the kernel's integrity and safety under any circumstances.

eBPF Just-In-Time Compiler and Interpreter
Upon completion of the verification process, the program is compiled into native machine instructions using the Just-In-Time (JIT) compiler. The kernel then

assigns a file descriptor to the loaded program, simplifying its attachment to various execution hooks within the kernel. In cases where JIT is disabled or unsupported, the eBPF interpreter takes on the responsibility of program execution, dynamically decoding and executing the bytecode in real time.

eBPF Hooks   The eBPF program's flow is determined by events, which are executed when the kernel or an application encounters specific hook points. These predefined hook points are placed in various locations within the kernel and cover a wide range of events, including system calls, function entry and exit, network sockets, tracepoints, and more. Based on its attach and/or program type, the program is attached to the hook where it is supposed to run. If a predefined hook does not meet the requirement, developers can create custom hook points called kernel probes (kprobes) or user probes (uprobes). These probes enable the attachment of eBPF programs to almost any position inside the kernel or user applications.

eBPF Program Types   The `BPF_PROG_TYPE` is the categorization of an eBPF program that determines its function, input parameters, acceptable actions, and attach points in the kernel. Each program type has characteristics that define its behavior and interaction with the system. For example, the socket filter program type `BPF_PROG_TYPE_SOCKET_FILTER` is designed to examine and manage network packets at the socket level. Developers can design customized logic within these programs to analyze and modify incoming and outgoing packets as needed. Conversely, tracing program types `BPF_PROG_TYPE_TRACING` are capable of monitoring kernel events and providing valuable information about system operation. These types of programs and their attach types are defined in the kernel codebase and serve as blueprints for creating eBPF

---

[2]These tools are the standard tools that have gained popularity in modern development environments, especially in systems programming and applications that emphasize performance.
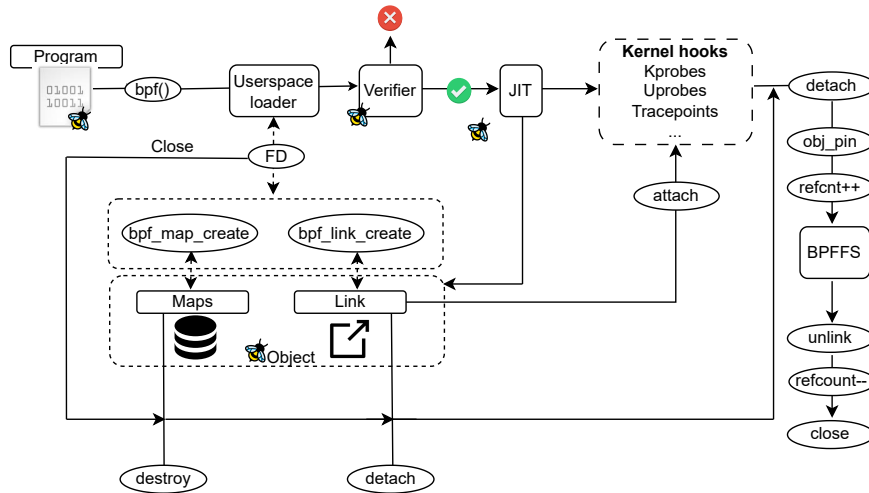
Figure 2: This diagram depicts the lifecycle of eBPF objects within the kernel, with the in-kernel representation, interaction with file descriptors, and the role of pinning in the *bpffs* file system [30].

programs to meet different requirements.

eBPF Helpers eBPF helpers are specialized functions accessible to eBPF programs, enabling interaction with the system and their execution context. These helpers facilitate tasks such as printing debugging messages, retrieving system boot time, manipulating network packets, and interacting with eBPF maps. Each eBPF program type can access a specific subset of these helpers, tailored to its context and requirements. For details about bpf helpers, see the documentation provided by the kernel [3].

eBPF Maps An eBPF map is an abstract data structure of a certain type, such as an array or hash map that facilitates data exchange between the user space and the kernel. The programs running within the eBPF virtual machine may obtain access to one or more maps through platform-specific load instructions.

3.2   eBPF Objects and their Lifecycle

Each eBPF object has an in-kernel representation for the management of eBPF program within the kernel, and is exposed through a file descriptor to user space. The lifecycle of the eBPF object is tied to the lifecycle of the file descriptor [62]. Once the final file descriptor corresponding to the eBPF object is released, its state inside the kernel is also released. To enable persistence beyond a process's lifetime, the kernel allows *pinning* these file descriptors on a special pseudo-file system called *bpffs*. Each pinning operation takes a a reference on the eBPF object, thus extending its lifetime.

eBPF Programs These objects represent the actual program being loaded into the kernel. A file descriptor representing the program is returned from the `BPF_PROG_LOAD` command of the `bpf` system call, after the eBPF verifier verifies and JITs the program and creates its in-kernel representation. Once loaded, the program is ready to be attached to a designated kernel hook.

eBPF Maps When eBPF programs are created, maps are defined using `BPF_MAP_CREATE` command of the `bpf` system call, which returns a file descriptor. This descriptor is used in pseudo load instructions within the eBPF program to reference the map. The verifier resolves the file descriptor to the actual map object in the kernel during program verification, and treats the destination register of the pseudo load instruction as a eBPF map pointer in the program.

eBPF Links eBPF links play a crucial role in ensuring that eBPF probes outlive the lifecycle of the application triggering them. eBPF links are created using the `BPF_LINK_CREATE` command of the `bpf` system call. The link enables developers to indirectly attach eBPF programs to kernel hooks providing a more flexible and sustainable method than regular direct attachment methods [74]. Instead of attaching directly to a hook, creating an eBPF link ties the lifetime of a program's attachment to a file descriptor, thus simplifying the management of program references and maintaining the probe even if the application loading it terminates unexpectedly [28]. The file descriptor associated with an eBPF link controls its lifecycle. When

the last file descriptor is closed, the link detaches its program from the kernel hook, allowing for resource cleanup. Only the link owner can detach or update it, ensuring system integrity and preventing unauthorized modifications [27].

BTF BTF objects represent the BTF type information for a eBPF program or map which has been submitted into the kernel from user space, to allow the verifier to tie this type information to the program or map during its verification procedure. In the case of the kernel and its modules, these objects are automatically created during boot and when any of the kernel modules are loaded. We elaborate on how BTF has been instrumental in enabling several other use cases in §3.4.

3.3 eBPF Instruction Set

The eBPF instruction set is defined in terms of the eBPF virtual machine. It supports two types of instruction encoding (64-bit and 128-bit), general-purpose instructions (such as arithmetic, jump, call, load, and store), and addressing the set of 11 64-bit registers (r0-r10) with a well-defined calling convention, where r10 is read-only and points to the top of the stack. We defer to the eBPF Instruction Set Specification [73] for the complete formal description of the instruction set and the calling convention.

A key principle used throughout the instruction set's design is maintaining near-equivalence with actual hardware instruction set architectures (ISA), which simplifies the implementation of interpreters and JITs. Moreover, this near-equivalence allows for optimizing compiler backends to emit eBPF assembly whose performance is close to natively compiled programs. This is because JITs can translate eBPF instructions to native machine instructions mostly with a one-on-one mapping, without introducing any extra instrumentation to handle the translation.

One of the strengths of decoupling the instruction set from the eBPF runtime in the Linux kernel is the ability to use it outside of the operating system. The eBPF ISA can either be supported directly by hardware [80], or translated to the target architecture of the hardware. This also has immense potential for computational hardware, as the eBPF runtime can control and decide whether offloading should be performed, while programs for a certain hook may be written in an oblivious fashion. Kicinski *et al.* [59, 60] have demonstrated the offloading of eBPF programs for the XDP hook [50] to programmable NICs, while Lukken *et al.* [63] explored its use for computational storage devices.

3.4 BPF Type Format

The BPF Type Format, or BTF for short, is a debug information format designed specifically for use with eBPF. It is produced by the compiler when the kernel or an eBPF program is compiled. In addition to information about C types, it contains function prototype information, custom annotations for types and declarations, which carry context-sensitive meaning for the eBPF verifier, and source information for better introspection and debugging. We defer to the BTF documentation [6] for the formal description of the metadata format.

The benefits of a new debug information format are twofold. Firstly, the existing debug information format used by the kernel, Debugging with arbitrary record formats (DWARF) [68], has a large overhead in terms of memory consumption if embedded in the kernel image. This means that always shipping kernels with debug information which could be used by the eBPF verifier to enrich its static analysis was infeasible. Secondly, for better introspection and analysis of eBPF programs and maps, they would need to supply their own debug information which could be inspected by the verifier. This meant introducing complex code to parse DWARF debug information into the kernel, which was undesirable for maintenance and security reasons.

BTF addresses all of this concern. Due to its compact representation, there is an order of magnitude difference [69] in the memory consumption between DWARF and BTF debug information for the same kernel image produced by the compiler. This is primarily due to the aggressive deduplication algorithm devised by Nakryiko *et al.* [69] to reduce BTF's memory footprint.

In turn, this allows the BTF to always be shipped with the kernel and eBPF programs, which the verifier heavily relies on to perform its static analysis. Due to its simpler representation, BTF is also faster to process, which is critical as an in-kernel representation is created at runtime for the kernel, any loaded kernel modules, and all eBPF programs and maps supplying their BTF information.

We now illustrate BTF's advantages in the context of its primary use cases.

Verification The main consumer of BTF is the eBPF verifier. It uses the kernel's BTF information to enforce *type safety* (§5.1) in eBPF programs for kernel pointers they gain access to. The verifier will introspect the type information using BTF to ascertain the size of the

object and introspect members of a structure type.

Annotations BTF can carry custom annotations for types and declarations of functions and variables. These are used to attach context-sensitive meaning [79], to types used in the kernel or the program, to aid verification [87].

Debugging The use of BTF in eBPF programs and maps allows for better introspection and debuggability. When analyzing a program, the verifier can print the source and line information to its log, which comes in handy whenever a program is rejected for user output. It also allows annotating the eBPF bytecode and JIT compiled code with source information. For eBPF maps, dumping of their data can be made structural by recognizing the type of data from their BTF.

Compile Once Run Everywhere (CO-RE) CO-RE is the collective name of a set of relocations for eBPF programs. These allow compiled programs to be more portable by encoding symbolic references for memory accesses to members of structure types, named enumerator constant values and named kernel configuration options. All of these relocations are resolved either by the verifier or `libbpf` when loading the program. This dynamic resolution ensures that eBPF programs can adapt to different kernel versions and architectures without the need for recompilation.

4 Workflow

In this section, we illustrate the high-level programming and execution model of eBPF. Figure 3 illustrates the entire sequence involved in the process of writing and executing an eBPF program for our chosen example from start to finish.

A user typically begins at step ⑤₁ by authoring an eBPF program in a high-level programming language. For our example, we consider the C program in Listing 1 written for the XDP hook [50], which invokes the program for network packets at the network device driver layer before they are processed by the networking stack. The `ctx` argument of type `struct xdp_md` represents the raw network packet the program gets access to. The `data` and `data_end` pointer variables point to start and one past the end of the network packet's data area. Conditional branches comparing `data` and `data_end` are used to ensure that enough room is available, to ensure that memory accesses of the packet's data will be safe.

```
SEC("xdp")
int bpf_program(struct xdp_md *ctx)
{
  void *data_end = (void *)(long)ctx->
      data_end;
  void *data = (void *)(long)ctx->data;
  struct ethhdr *eth = data;

  if (eth + 1 < data_end) {
    if (eth->h_proto == bpf_htons(ETH_P_IP))
        {
      struct iphdr *iph = (void *)(eth + 1);
      if (iph + 1 < data_end && iph->
          protocol == IPPROTO_UDP)
        return XDP_DROP;
    }
  }
  return XDP_PASS;
}
```

Listing 1: An example BPF program for the XDP hook, which drops all incoming IPv4 UDP traffic.

The next step ⑤₂ involves the compilation of this C program using the LLVM toolchain's clang compiler into an object file. The target for compilation is chosen as `bpf`, which instructs the compiler to use the eBPF backend to emit binary code for the produced object file.

Step ⑤₃ is concerned with the processing of the produced object file and submitting the program encoded within it to the kernel through the `bpf(2)` [58] system call for loading it. For our example, we use the `bpftool` user space tool which in turn uses the `libbpf` library to perform the loading of the program. Once the object file has been processed and the program has been extracted from it, step ⑤₄ submits it to the kernel using the `bpf(2)` system call's `BPF_PROG_LOAD` command, which invokes the eBPF verifier.

The eBPF verifier then performs verification of the program to decide whether it will be safe for execution within the kernel. If the verifier fails to determine the program's safety, it rejects the program and returns an error to user space. Otherwise, a successfully verified program is JIT compiled, and a file descriptor corresponding to the eBPF program is returned to user space.

Once user space has the file descriptor for the program, it can now attach to a network device's XDP hook. For our example, the network device is eth0, and step ⑤₅ involves invoking the `BPF_LINK_CREATE` command of the `bpf(2)` system call to perform the attachment of the program to the network device. If all parameters for the command were valid, the kernel returns a file descriptor corresponding to the eBPF link is returned to user space.
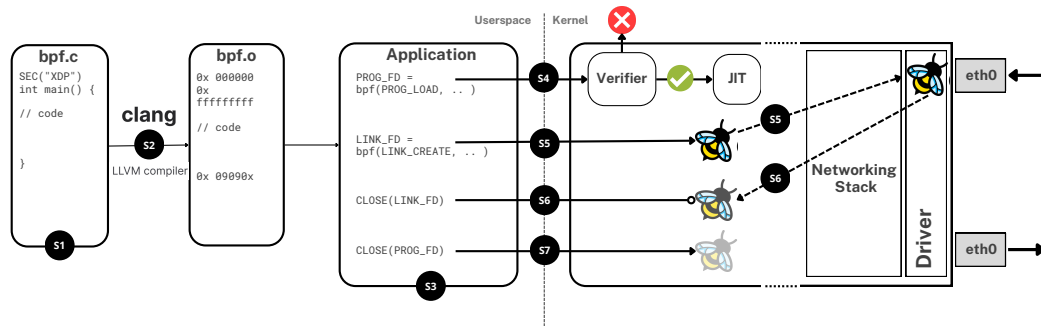
Figure 3: Workflow diagram of an eBPF program based on [21, 30].

At this point, the eBPF program is attached to the `eth0` network interface and rejects all incoming IPv4 UDP traffic for it. It is invoked for every raw network packet received by the kernel's network device driver, and performs its processing on it. The rest of the traffic passes through to the kernel's networking stack.

In step **S6**, once the user space application closes the eBPF link file descriptor, the program is detached from the network interface. Once the eBPF program file descriptor is closed in step **S7**, the kernel will free the resources it occupies such as memory for its code.

5 Safety of eBPF Program

Program safety is a critical aspect of eBPF programs, ensuring they execute correctly and securely without compromising the stability and security of the Linux kernel. In the context of eBPF program, program safety refers to a set of properties that must be satisfied to protect the runtime integrity of the kernel and uphold the invariants of the kernel context in which the eBPF program will execute. Programs that violate any of these safety properties should be rejected when being loaded into the kernel.

5.1 Safety Properties

We now classify and discuss each safety property as detailed in the BPF verifier code and related documentation [75, 76].

- **Memory Safety:** For all memory regions accessible to the program, the verifier endeavors to guarantee that there can be no out-of-bounds accesses, invalid or arbitrary memory accesses, or use-after-free errors. The verifier performs precise bounds tracking for all memory regions accessible to the program and checks that all accesses lie within bounds. Any allocated memory region cannot be accessed once it has been freed. Arbitrary values cannot be used as pointers for memory accesses.

- **Type Safety:** The verifier has precise knowledge of the type of each register that points to a memory region accessible to the program. It keeps track of the type of objects on the stack. This process helps avoid errors caused by type confusion and ensures that different types of programs can utilize various utilities without corrupting kernel memory. The verifier also leverages BTF (§3.4) to capture essential information about kernel and eBPF program types and code. BTF helps identify eBPF kernel data structures and ensures that aggregate types, such as structs, are not accessed beyond their allowable limits.

- **Resource Safety:** The verifier checks that the program leaves no lingering resources when it exits. This implies that the program must release all allocated memory, acquired locks, and reference counts to kernel objects using an appropriate helper function.

- **Information Leak Safety:** The eBPF verifier outlaws any kernel information leaks by analyzing pointers that could potentially reference kernel memory, aiming to prevent any leaks into user-accessible memory regions. It performs thorough escape analysis [78] for all pointers that may point to kernel memory, ensuring they do not escape into memory regions accessible by user space. Additionally, it rejects any attempts to read uninitialized regions of the stack.

- **Data Race Freedom:** The verifier aims to ensure that the program's accesses to kernel state are free from data races. It enforces that any manipulation of kernel state occurs through helpers implementing appropriate synchronization. However, the verifier does not diagnose data races for accesses to memory owned by the program itself (*e.g.* values of eBPF maps), because it has no bearing on the kernel's runtime integrity.

8

- **Termination:** The verifier places a limit on the maximum number of instructions it will explore across all paths of a program, known as the `instruction complexity limit`, aiming to ensure that all programs terminate. If the program does not demonstrate termination for all paths within this limit, the verifier rejects it. This could be due to infinite loops, unbounded loops with unproven termination, or simply because of programs that are too large.

- **Deadlock Freedom:** The verifier aims at ensuring that the program is free of deadlocks. By definition, a deadlock requires two programs executing in parallel to hold two locks in the opposite order. To achieve deadlock freedom, the verifier simply disallows holding more than one lock at a time at any point in the program.

- **Upholding Execution Context Invariants:** The verifier checks to ensure that the program upholds all invariants of the execution context within the kernel. In other words, a program's execution may not violate the invariants and assumptions of existing kernel code. This information is encoded in the verifier every time support for a new kernel hook is introduced.

By delineating these safety properties, we have defined the state of the art for what the verifier will enforce for eBPF programs. However, it is crucial to recognize that these standards are not static; rather an ongoing work to protect the integrity and security of eBPF programs within the kernel, and will be periodically updated as the ecosystem changes and new approaches are explored to reflect the dynamic nature of this field.

5.2 The eBPF Verifier

To enforce the safety properties (§5.1), the eBPF verifier plays a major role. The eBPF verifier is a static analyzer for eBPF programs. It is invoked when a program is loaded into the kernel, and is tasked with ensuring that the program is *safe* to execute within the kernel's context. Once this has been determined, the verifier submits the bytecode for JIT compilation, where it is converted into native machine instructions. The program can then be attached to one of the many available kernel hooks to begin its execution.

The verifier performs static analysis of the program at the level of its eBPF bytecode. It has access to the program's BTF (§3.4) debug information which was produced during its compilation from a high-level language to eBPF bytecode. This design choice of operating over eBPF bytecode and BTF allows the eBPF
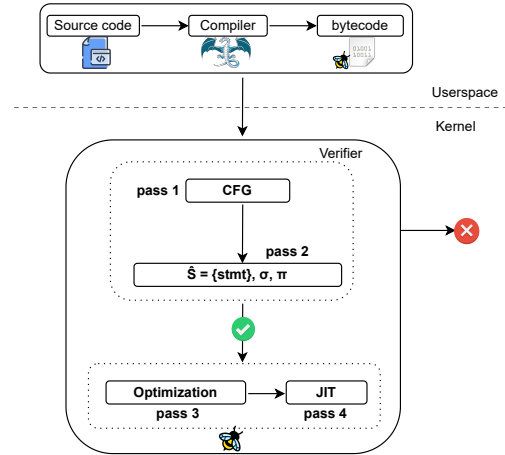

Figure 4: The four major passes of the eBPF verification process

verifier to be usable for programs written in multiple higher-level languages.

There are four major passes involved during the verification of a program, as shown in Figure 4. The first pass validates the control-flow graph of the program (§6). The second pass performs exhaustive symbolic execution of the program to ensure that it is *safe* (§7). Thereafter, the third pass performs optimizations and transformations (§8) for the program before it is submitted to the Just-In-Time compiler (§9) in the final pass. The following sections describe the verifier passes, explaining how the verifier enforces the safety properties.

6 Validation of the Control Flow Graph

The verifier begins by analyzing the Control-Flow Graph (CFG) of the program. Instructions in the CFG are represented as nodes, and the different forms of control flow represent edges. All, but the first instruction, will have at least one incoming control-flow edge. Except `EXIT`, all instructions also have one outgoing `fallthrough` control-flow edge, which points to the next instruction. For unconditional `JUMP` instructions, the fallthrough edge points to the instruction they jump to. In case of `CALL` and conditional `JUMP` instructions, there is an extra outgoing *branch* control-flow edge to the branch target. For `CALL` instructions, the branch control-flow edge points to the callee's first instruction. For conditional `JUMP` instructions, it points to their branch target instruction (when the condition is true).

The verifier walks the CFG (Figure 6) of the program using a depth-first search algorithm. This allows it to label all reachable instructions and the control-flow

edges followed to visit them. Simply completing the traversal allows it to detect any unreachable instructions in the program by checking if they were unvisited. Additionally, while walking the CFG, it tracks `state pruning` points.

The verifier checks for the following properties, and rejects programs that do not adhere to it:

- Ensures no infinite loops or loops with complex termination conditions that cannot be statically determined.

- Ensure that there exist no unreachable instructions in the program.

- Ensure that all subprogs end with `EXIT` or `JUMP` instruction, *i.e.* no automatic fallthrough to the next subprog.

State Pruning. The verifier performs precise symbolic execution of the eBPF program (described in §7). However, this approach can be very expensive when multiple paths need to be explored due to the presence of multiple branch conditions. The distinct state constraints, symbolic and concrete values for registers and the stack have to be maintained in separate verifier states. Thus, the path explosion problem also translates into a state explosion problem. Since the verification algorithm's complexity limit concerns instructions explored in *all* possible paths, larger and more complicated programs would eagerly hit this limit.

To alleviate such scalability issues, the eBPF verifier implements the `state pruning` approach, borrowing ideas from the RWSet analysis technique by Boonstoppel *et al.* [38]. This involves pruning of redundant path walks while exploring the program by detecting variable *liveness*. The idea is to establish equivalence in terms of program side-effects between the current program state and an already explored program state which has passed through the same point in the program. During this state equivalence check, the verifier skip variables which are not used in subsequent basic blocks. These pruning points are simply instructions in the program for which the states will be stored and compared for equivalence every time the verifier encounters them while exploring a path.

The program is annotated with `pruning points` where the verifier's current state will be saved. Later, when the verifier arrives at the same pruning point while exploring another path, it traverses through the checkpointed states which have been fully explored and compares the current state to them.

This is meant to establish `state equivalence`, a property which takes into consideration whether given the old checkpointed state, which has already been verified to be safe, it can be considered equivalent to the current state from the point of view of program safety. If this property is true, then the verifier need not continue exploration of the current path, as the program from this point has already been verified to be safe for an equivalent verifier state. Otherwise, the verifier keeps exploring the current path until the next pruning point, where it repeats this step, or it encounters the `BPF_EXIT` instruction.

7 Symbolic Execution

The verifier's second pass symbolically executes the eBPF bytecode. To ensure safety, as defined previously in §5.1, the verifier must exhaustively explore all feasible paths through the program and precisely track state of the stack (at byte-level granularity) and the program registers through every point in the the control flow. Note that pass one (§6) already ensured that all points in the CFG are reachable.

State tracking in *symbolic execution* is defined in the form of a three-tuple $(\hat{S} = \{stmt\}, \sigma, \pi)$ for every path explored. $\{stmt\}$ denotes the list of instructions visited while executing the current path symbolically. σ is a symbolic store that maps *program variables*, here eBPF registers and stack, to symbolic or concrete values and tracks associated type information. Finally, π represents the path state, which includes additional state information maintained during path exploration. In the context of eBPF, these conditions maintain dataflow information collected during path exploration. For example, π maintains *precision* markings to know which variables need to be precisely tracked (§6), lock-held sections of the program. In addition to tracking this three-tuple over all paths, the verifier also caches the state (σ, π), by checkpointing them at *prune points* that we computed in pass one (§6). We call this as the *verifier state* in the context of the eBPF verifier's symbolic execution pass.

The verifier begins path exploration from the first instruction of a program, by setting up the *symbolic store*. In the initial state, only register $r1$ has some value. It holds a pointer to the context object which every program receives when being invoked from the kernel; all other registers are empty and the stack is uninitialized. The internal structure of the object pointed to by the context pointer is dependent on the program type. Thereafter, the verifier begins symbolically executing the program, instruction by instruction, and continues to update the *verifier state*.

When the verifier encounters a conditional `JUMP` instruction, it first tries to predict which branch will be taken based on the values of operand registers in the symbolic store σ. If it cannot predict the evaluation of the branch condition, the verifier splits exploration by forking the current *verifier state* into two. It then continues exploration of one path of the branch, and enqueues the other path into a worklist for later exploration. The *verifier state* in each path exploration will have additional information about the program state based on the truth assignment of the branch condition. The symbolic value of the operand register will be updated to have more precision.

Unfortunately, with branching programs such as loops, the number of paths to explore increases exponentially, triggering path explosion. To keep the verification task tractable, the verifier employs an approach called *state pruning* which reduces the number of paths it needs to explore. As we discussed earlier in §6, this ties back to the *prune* points marked by the verifier during its CFG validation pass. We now discuss in detail the pieces needed for symbolic execution.

7.1 Symbolic State

The symbolic store σ in the *verifier state* tracks the following information about registers and stack state:

- It tracks whether the value in a register, or data in the stack (at byte-granularity), is a *scalar* or a *pointer*.

- For *scalars*, the verifier tracks precise integer bounds. This means that the verifier does not over-approximate the range based on visiting a *stmt*.

- For *pointers*, the verifier tracks the type of object it points to, the `offset` of the pointer in an accessible memory region, and the `len` in which de-referencing the pointer is valid.

The path state π in the verifier tracks the following information at every *prune* point:

- **Liveness Tracking:** The verifier performs, in-place, live variable analysis to backpropagate which registers are live.

- **Scalar Precision:** In addition to liveness, the verifier also tracks scalar values that are used later in contexts where precision tracking is necessary, such as `JUMP` targets and helper calls. Registers that hold scalars that are live but need not be precise also assist in pruning redundant path exploration.

- **Resource Tracking:** π also keeps track of resources like allocated memory regions, ref-counts and spinlock held in a certain section of the program to ensure their timely release before program termination. It enables the management of various mechanisms including lock tracking and synchronization, kernel object reference management, and memory allocation and deallocation.

- **Pointer Alignment:** The verifier inspects the alignment of register types and pointers to prevent any access to memory with incorrect alignments due to variable offsets. It scrutinizes the amount of data being read from or written to memory in a single operation. Additionally, it enforces alignment checks for stack pointers, as it depends on tracking stack spills. Any misaligned stack access can lead to the corruption of spill registers, thereby posing potential threats of exploitation.

7.2 Instruction Simulation

The eBPF instruction set defines all instruction types, their expected operands, and expected behavior. All instructions are encoded into the verifier as transfer functions that take an initial state $\hat{S}_0$ and returns a new state $\hat{S}_1$. If any simulated transition causes an error, verification fails.

We now discuss what the modeling of each instruction-class looks like in the verifier.

Arithmetic Operations   In cases where the opcode is valid, the function goes one step further and checks the validity of the source and destination operands by ensuring that the register used as the source operand is readable and that the register used as the destination operand can be written to. If the source and destination registers are valid, arithmetic operations with pointers and scalars are processed and new signed and unsigned bounds can be calculated for all *32/64-bit* ALU operations, excluding `BPF_END`, `BPF_NEG` and `BPF_MOV` operations as detailed in Appendix A.

Load and Store Instructions   Store instructions have two modes of operation. The `BPF_STX` class takes a register operand whose value will be stored, while `BPF_ST` works with an immediate as the source operand. The destination operand is always a register with a pointer type of a memory region. While the handling of each pointer type is distinct, overall all of them share a few high level properties. First of all, the actual offset where the store is done is formed by accumulating the offset specified as part of the store instruction and the

symbolic or concrete offset value associated with the destination register.

Value Tracking   Value tracking ensures the safety and accuracy of memory operations in eBPF programs by analyzing register and stack slot values and enforcing constraints. Each register's state is carefully observed and assigned a specific type based on its content. For instance, registers may have type `NOT_INIT` if they haven't been written to yet, or `SCALAR_VALUE` if they hold scalar values not used as pointers. Pointers, when present, are categorized into various types like `PTR_TO_CTX` or `PTR_TO_MAP_VALUE`. Additionally, pointers can undergo arithmetic operations, resulting in fixed or variable offsets. The verifier tracks these offsets, adjusting their minimum and maximum values accordingly. This insight of register types and pointer offsets ensures that memory accesses within eBPF programs adhere to designated areas and comply with safety constraints.

Numerical Abstract Domain   Tnum is an abstract range that approximates the set of possible values that can be stored by a variable. It is used by the eBPF verifier to check the validity of memory usage during read or write operations. Tnum is a 64-bit value with a mask that indicates a range of possible values. Each bit in the binary representation of a tnum can be set to 1, 0, or unknown [88], representing a precise value, a known range, or an unknown range, respectively. The verifier also tracks register constants using an interval abstraction to determine the minimum and maximum possible values of each register at any given time. By analyzing this data, the verifier can detect potential out-of-bounds memory accesses and prevent them from occurring during program execution. For the complete formal description of the verifier's numerical abstract domain, including the proofs for its soundness and optimality, we defer to Vishwanathan *et al.* [88].

Conditional and Unconditional Jumps   In the eBPF architecture, jumps are classified into unconditional and conditional. Unconditional jumps move the program counter to a new instruction by adding a fixed offset to the current instruction's position. This offset can be positive or negative, enabling jumps forward or backward, as long as they stay within the program's boundaries and avoid loops.

Conditional jumps, however, depend on a runtime condition involving the evaluation of certain variables or registers. If the condition evaluates to true, the program counter is adjusted by a specific offset to reach a new instruction. If the condition is false, the execution continues to the next sequential instruction. Thus, unconditional jumps follow a predetermined path based on a fixed offset, while conditional jumps adjust the path according to whether the evaluated condition is met or not [42].

Function Calls   Different eBPF program types have access to a different set of functions, reflecting their specific use cases. Function calls in eBPF programs extend their functionality. To ensure safety, the verifier enforces strict checks to guarantee that functions are called with valid arguments. This process involves checking that the registers used for function arguments such as (r1 - r5) match the expected types, only calls to known functions are allowed and function calls and dynamic linking are prohibited [26].

Registers used temporarily within functions, known as caller-saved or volatile registers, need not retain their values after function calls. Therefore, it is the caller's responsibility to save these values if they are needed later. In contrast, callee-saved registers r6 - r9 must be preserved across function calls to maintain their values. The verifier's symbolic execution pass ensures these rules are followed, enhancing function call safety.

The verifier examines whether the program's state has sufficient precision to prove the need to avoid further unrolling. This refers to accurately capturing the critical parts of the state necessary to confirm the program's safety. When precision is achieved, it means the verifier has a stable and detailed understanding of the program state, particularly the aspects that affect safety-sensitive operations. In this case, further actions like unrolling loops are unnecessary. Unrolling is generally used to investigate different possible behaviors of the program, but if the verifier has already obtained sufficient information to ensure safety, it can skip this process.

7.3 Loops

Loops are supported by the verifier in primarily two flavors. The first is loops within the program whose bounds are known to the verifier. The second is when the verifier encounters a loop whose bounds are unknown to it.

Bounded Loops   The process of bounded loop verification is done by unrolling the loop. The verifier continues unrolling until it has exhausted exploring all iterations, or until the instruction complexity limit is hit. Hence, there are no issues stemming from unrolling until a fixed bound which may be less than the actual bound, and the verifier does not have to worry

about lost precision from not exploring the remaining iterations.

Unbounded Loops   Unbounded loop verification in the verifier relies on special helper functions. Typically, this is only supported for `bpf_iter` helpers which initialize an iterator object, get the next item from the iterator, and destroy it. The typical setup is to initialize and destroy the iterator object before and after the loop, and use the return value of the helper returning next item to continue loop iteration. The verifier knows that the helper eventually returns an item which terminates iteration, so it does not have to concern itself with proving the termination of the loop.

At this point, the only relevant property is to ensure that the verifier can precisely track the state of the program after the body of such loops has executed an unknown number of times. Due to the limitation of unrolling until a fixed bound, as premature termination of unrolling may occur, the verifier employs state pruning logic to establish that the current state already possesses sufficient precision, thereby proving that further unrolling is unnecessary. Until the unrolling process converges for such unbounded loops and performs state pruning at the point where the loop condition is checked (to determine whether to break or continue iteration), the verifier continues exploration. If the loop does not converge until the instruction complexity limit is reached, the program is rejected.

Since the verifier relies on state pruning to establish equivalence of verifier states representing different iterations of an unbounded loop's body, it needs to know what registers and stack slots actually need precision to increase the chances of convergence. To achieve this, the verifier first performs a depth-first exploration of the remaining program when encountering the loop condition for the first time. Once it has explored the program past the loop once, all precision and liveness marks have been propagated to the checkpointed states for the relevant pruning points. This then allows the verifier to more aggressively perform state pruning to establish convergence.

Path Explosion   The loop handling in both cases is effective when the loop body is simple. However, the verifier encounters difficulties when the loop body contains branches. For each such branch, when the verifier unrolls the loop body, it has to fork the states to follow both arms of a branch condition. If one of the paths causes loop termination, it will continue exploring the rest of the program for each iteration. The number of states grow exponentially when branch conditions exist within the loop body, and pruning opportunities are typically missed as each state requires precision for certain registers and stack slots which are distinct for either of them. This is a major limitation for the verifier today, where it will simply give up under path explosion once it hits the instruction complexity limit.

7.4   Resource Management

The verifier's resource management logic is concerned with tracking symbolic resources created for a particular verifier state, and ensuring that they have been destroyed by the time the exploration of the path being explored terminates. In simple terms, this means that a eBPF helper can be responsible for acquiring a reference in the verifier state, and another eBPF helper will be responsible for releasing it. This is then mapped to multiple higher-level acquire-release patterns, such as acquiring and releasing reference counts of kernel objects, allocating and freeing memory, obtaining and relinquishing ownership of an object. Mapping verifier level references to values returned from eBPF helpers allows enforcing the invariant that any acquired resources are released before the program's exit. Each reference's unique identifier is attached to the register `r0` after a reference acquiring eBPF helper is invoked. The same pointer value is then passed as an argument to a reference releasing eBPF helper, which releases the reference state corresponding to the identifier. The verifier implements precise tracking of these references and complains whenever there is an unreleased reference state after encountering the `BPF_EXIT` instruction. Among the resources, the verifier keeps track of held spin locks, and ensures deadlock avoidance. Currently, each program can only hold a single spin lock at once, and must release this same lock before the program ends. By ensuring that every program only holds a single eBPF spin lock at a time, it also eliminates the possibility of deadlocks. In the verifier state, the verifier associates a unique identifier with the memory region holding the a given eBPF spin lock, and remembers this identifier. A pointer to the same memory region must be passed to the unlock function to release the spin lock in the verifier state, thus also ensuring that there is no mismatch between the lock and unlock calls.

8   Post-Verification Optimizations

After completing the symbolic execution pass for the eBPF program and deeming it safe for execution, there are a series of post-verification optimizations and fixups that are performed on each program as part of the program optimization and fixup pass. We discuss them in this section.

8.1 Dead Code Elimination

The verifier during symbolic execution does not explore untaken branches of conditional jump instructions when it can determine that a given branch will never be taken by analyzing the program's data flow. Every instruction that has been simulated at least once is labeled as 'seen'. Thus, once the verifier can conclude that an instruction is never reachable at runtime, it can safely eliminate such dead instructions from the program.

Some conditions in the program can be only resolved as late as the verification stage (*e.g.*, configuration option values for the kernel on which the program is being loaded). Thus, the eBPF verifier can eliminate dead code in the program more aggressively than the compiler due to the richer data flow information available to it. This optimization opportunity is readily exploited by the verifier, primarily to enable better portability for eBPF programs across kernel versions, where different code paths are taken for different versions or configurations, and to reduce runtime overhead of always untaken branches and conditional jumps. For programs running on the kernel's latency critical path (*e.g.* XDP), dead code elimination is important for saving precious CPU cycles at runtime due to the nanosecond scale execution requirements [50] while retaining the ability to have conditional fallbacks in the program that are resolved at load time.

8.2 Inlining and Instruction Rewriting

In some cases, the overhead of direct calls starts to add up for programs. The underlying operation involves so little work that the overhead of calling a function to do it would be more than directly executing those instructions. This also applies to timing-sensitive operations, such as the `bpf_jiffies64` helper. However, eBPF helpers also serve as a means to enforce API usage and program context invariants. The verifier cannot permit the program to directly manipulate kernel data structures. In such cases, it will transparently replace the helper call instruction and substitute a set of eBPF assembly instructions implementing equivalent functionality. This technique is used for map accesses. During program verification, the verifier knows which eBPF map objects the program has access to, statically. Thus, whenever the program calls eBPF helpers manipulating those maps (*e.g.*, `bpf_map_lookup_elem`, `bpf_map_update_elem`, and `bpf_map_delete_elem`), the verifier, by knowing the type of the map during verification, can transparently translate indirect calls for the map operation made by these eBPF helpers into direct calls for the underlying map implementation.

9 Just-In-Time Compilation

After after post-verification optimizations, all eBPF programs are Just-In-Time (JIT) compiled to native machine instructions before being enabled for execution. The JIT compiler performs a direct translation of eBPF instructions to the underlying machine instructions, benefiting from the close equivalence between the eBPF ISA and hardware ISAs. The compilation procedure occurs separately for each subprog, and can be decomposed into four major steps as illustrated in Figure 5:

- **Image Allocation:** Estimate the size of the JIT *image* and allocate memory to write native machine instructions to after translation.

- **Emit Prologue:** Emit the code required for the kernel to safely call into the eBPF program. This involves saving the frame pointer and pushing any callee-saved registers to the stack.

- **Emit Body:** Emit the code to translate each eBPF instruction to native machine instructions. Special handling is performed for load instructions which were seen for registers holding untrusted pointer. The JIT compiler prepares an exception table indexed by the instruction, and any page faults on invalid access at runtime are automatically resolved to not crash the kernel by looking up this table, and simulating a load of zeroed memory.

- **Emit Epilogue:** Emit the code to undo the prologue, *i.e.* to restore the frame pointer and pop the callee-saved registers from the stack. The final JIT image is made read-only as a security measure to prevent modifications of the memory later.

When extra hardening is requested [48], the JIT compiler performs *constant blinding* as a mitigation. For every instruction using an immediate as operand, the immediate value is xored with a random constant, and moved into a JIT-specific internal eBPF register (`rAX`). Then, `rAX` is xored again with the same constant [29, 83], and the instruction in question is translated but modified to use register `rAX` as its operand. This changes the immediate value in the produced JIT image after translation, rendering any JIT spraying attempts useless [33].

10 Use Cases

We illustrate the various use cases served by eBPF within and beyond the Linux kernel.
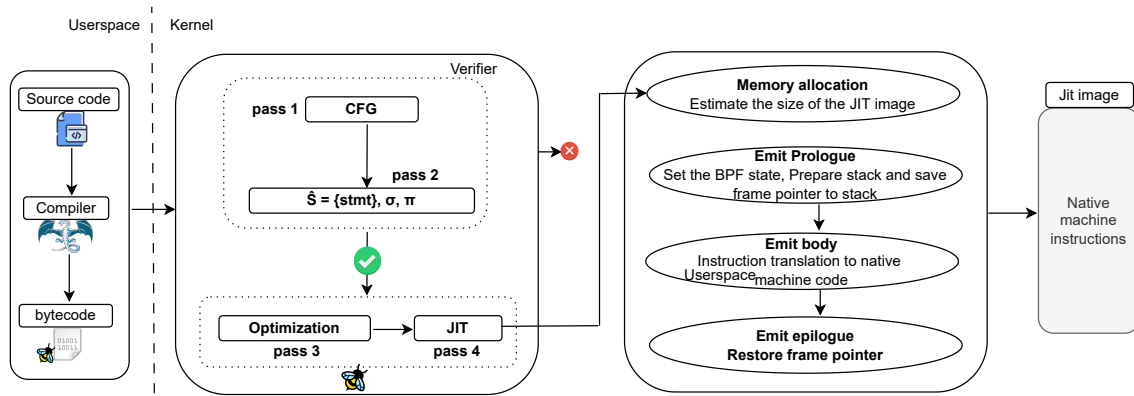
Figure 5: Overview of the process that shows eBPF instructions translation into native machine code

10.1 Networking

We now highlight several key functionalities enabled by eBPF programs in networking, empowering developers to achieve high-performance tasks, ensure flexibility, and enhance efficiency through customizable hooks in the Linux kernel.

- **XDP and TC:** Programs for the XDP [50] and TC hooks [17, 39, 40] perform high-performance network packet processing by bypassing the kernel networking stack, while still having access to its state. This includes dropping packets at high rates for DDoS mitigation [35, 66], load balancing, implementing network functions [67], redirecting flows to different CPUs or network devices, or rewriting and transmitting the packet directly without passing through the kernel.

- **Socket Lookup:** This hook allows programs to select the socket which receives a network packet destined for local delivery [5, 10]. This allows steering traffic from any IP address and port pair to a single socket, without using a socket per pair, which affects socket lookup scalability.

- **Socket Reuseport:** This hook allows programs to choose one from a set of reuse sockets bound to the same IP address and port. Programmability through eBPF allows for more informed selection, such as favoring NUMA locality, connection migrations from one socket to another [61], or based on data within the packet.

- **Control Groups:** Control Group hooks allow programs to be logically tied to a container context. Ingress and egress hooks allow filtering traffic. Egress hook also allows limiting outgoing bandwidth by setting the Earliest Departure Time [18]

for the packet. Other hooks exist for control path operations. Hooks for `bind`, `connect`, `sendmsg`, `getsockname`, `getpeername` system calls allow enforcing policies and manipulating the local or remote address passed in from user space. Hooks for `setsockopt` and `getsockopt` allow changing and setting additional socket options based on the desired policy [8].

- **User-Level Protocol:** The ULP hook (also known as SK_MSG) allows invoking eBPF programs when a `sendmsg` or `sendfile` operation occurs for a socket. These programs are used to enforce policies for the payload of each message being sent [2]. Paired with kTLS [13], it provides transparent enforcement of ULP layer policies even for encrypted traffic [41].

- **Congestion Control:** eBPF programs can be registered as callbacks for TCP congestion control operations [12], allowing faster experimentation, data collection, and iteration of custom congestion control algorithms in production.

- **Socket Operations:** These hooks allow attaching eBPF programs to a TCP socket's state transition events for example listen, connect, active connection vs passive connection establishment, TCP header options [31, 43]. It allows using the peer address to dynamically select socket options and congestion control settings, the retransmission timeout, maximum acknowledgment delay timeout, and manipulating TCP header options.

10.2 Profiling

Applications use eBPF programs attached to `perf` events to collect data from hardware performance counters and capture stack traces for the kernel and user

15

space. Due to its low overhead, eBPF is used in multiple continuous profiling applications [92] without degrading workload performance significantly.

10.3 Tracing

Kernel functions and tracepoints can be traced by attaching eBPF programs which execute at function entry and exit. The programs have access to all arguments of the kernel function or tracepoint. Low overhead instrumentation (due to runtime code modification to emit direct calls to the program) makes runtime tracing practical even for high-performance production workloads. A large set of tools use this facility to perform runtime data collection, performance analysis, and profiling of kernel subsystems [49, 54].

10.4 Security

The LSM BPF subsystem [14, 16] allows eBPF programs to be attached to LSM hooks within the kernel. Programmable LSM hooks allow enforcing security policies and auditing of the system. Object-local storage maps [4] are used to associate policy-specific data with a kernel object (*e.g.* cgroups, inodes, sockets, etc.) acting as a subject of an LSM hook. Additionally, LSM BPF programs can be attached to a cgroup context to constrain the policy's scope to it [47]. These capabilities allow LSM BPF to serve as a basic building block for flexibly building higher-level security frameworks [46].

10.5 Emerging

Listed below are some emerging applications and innovations where eBPF is leveraged to extend functionality beyond traditional use cases within the Linux kernel.

- **Device Drivers:** The HID-BPF framework [9] allows parts of the HID device drivers to be implemented using eBPF programs to filter events, make driver fixes without changing the kernel, and inject additional input events.

- **Scheduling:** The ghOSt scheduler [51] attempts to delegate scheduling to a user space application by using eBPF to communicate scheduler events over shared memory. SCHED-EXT [20, 85] takes a different approach, by completely implementing the scheduling logic within eBPF programs as synchronous callbacks for the Linux scheduler.

- **Storage:** XRP [90, 91] accelerates storage applications by attaching eBPF programs to the NVMe driver layer. The programs then issue read operations directly, bypassing the kernel's storage stack, while still keeping file system state in sync.

11 Challenges

In this section, we focus our attention on various challenges concerning eBPF's current design.

11.1 Usability

Navigating the intricacies of connecting eBPF programs to attach points in the Linux kernel can be challenging. Understanding the various hook points available, discerning their suitability for specific tasks, and seamlessly integrating eBPF programs with them necessitate a profound understanding of kernel intricacies and eBPF methodologies [1].

Furthermore, the dearth of exhaustive documentation and user-friendly development tools exacerbates these usability hurdles. Developers often find themselves grappling to locate pertinent resources and guidance, impeding their progress and efficiency in eBPF development endeavors. Moreover, there are concerns about compatibility and stability across different versions of the kernel. Alterations in kernel APIs or hook placements can disrupt the behavior of eBPF programs, compelling developers to continuously adjust their code for compatibility. This perpetual adjustment introduces unwelcome complexity and overhead [86].

11.2 Scalability of the Verifier

A more scalable verifier directly translates to increased capabilities for eBPF, as a larger set of valid programs can pass through it. The current verifier's static analysis theory is not capable enough to handle extreme pessimistic cases of path explosion. Scalable handling of loops is also critical for increasing expressiveness of programs. Instead of unrolling, techniques such as loop invariant analysis [45] or summarization need further exploration.

11.3 Correctness of the Verifier

The safety guarantees of eBPF hinge on the correctness and soundness of the verifier's implementation. The eBPF verifier serving as the final arbiter deciding whether programs are safe to execute comes with its own set of problems. Firstly, the verifier cannot be too conservative during its analysis, since that leads to the rejection of a large set of valid and safe programs. Secondly, the verification algorithm should ideally terminate within a fixed amount of time, without breaching the complexity limit during the program's symbolic execution.

In practice, the verifier can minimize the impact of these vexing issues to a great extent using the algorithms illustrated before in §6. Any logical bug in the

verification algorithm, failure to account for unsafe program behavior in unexplored paths due to eager redundant path pruning, or failure to capture and enforce kernel-specific invariants correctly directly translate to unsafe or malicious eBPF programs passing through the verifier, undermining eBPF's strong safety guarantees. The complexity and sheer size of the verifier's codebase, paired with the high frequency of changes made every kernel release to meet the needs of ever-increasing use cases make preserving the correctness of the verification logic an increasingly daunting task for eBPF developers.

11.4 Formal Verification

There has been no comprehensive formal investigation of the verifier and whether its safety guarantees are sound. This remains an open research problem, and also a huge undertaking due to large number of features supported by it. The difficulty is further compounded by the high rate of changes made to the verifier's codebase in every kernel release [76].

Nevertheless, some promising attempts have been made thus far. Vishwanathan *et al.* [88] formally specify the verifier's numerical abstract domain (§7.2), providing soundness and optimality proofs. A part of this work has since been adopted in the Linux kernel [23]. Bhat *et al.* [37] create an automated formal verification framework that verifies the correctness of the C implementation of the range analysis logic (§7.2) against a specification describing the correctness invariants. Nelson *et al.* [70] create the Serval framework to produce an automated verifier for the eBPF instruction set. Nelson *et al.* [71] also apply automated proof techniques to verify eBPF's JIT implementations, by constructing a JIT correctness specification.

11.5 Security

The criticality of eBPF bugs affecting the safety, integrity, and security of the kernel, and by extension the rest of the system is self-evident from the set of reported vulnerabilities in the past, which repeatedly abused it as a potent vector for exploitation [64] in unprivileged mode. As a consequence, eBPF developers chose to disable the unprivileged mode by default [72], and even when enabled explicitly, it is extremely restrictive to reduce the kernel's attack surface exposed to untrusted and potentially malicious users. Hence, eBPF today remains largely useful only in the trusted user model.

The introduction of `CAP_BPF` with linux 5.8 aims to separate BPF functionality from the broader `CAP_SYS_ADMIN` capability. The general idea is that a user who has this capability is able to (among other things) create bpf maps or load `SK_REUSEPORT` programs. However, capabilities like `CAP_NET_ADMIN` and `CAP_PERFMON` are still required for loading networking and tracing programs respectively,highlighting the ongoing need for elevated privileges in certain eBPF operations.

Improving the correctness guarantees of the eBPF verifier and JIT implementation might, among other things, allow for the possibility of relaxing the restrictions imposed on the eBPF programs in the unprivileged mode, making eBPF more useful for a larger set of use cases which do not require privileges [57].

11.6 Code Reuse

Code reuse in eBPF programs is a double-edged sword, because on the one hand there is CO-RE, as explained in §3.4, which allows to use the same compiled programs on different Linux versions by fixing memory offsets for data structures at load time. On the other hand, although function calls are allowed in eBPF programs, there is no support for static or dynamic libraries. This means that two or more programs living in separate source files may need to reimplement the same functions.

12 Conclusion

We examine the operational mechanisms of different components within the eBPF subsystem in the Linux kernel, focusing on its objectives to offer safety assurances and analyzing the current implementations of these objectives. We discuss how various components in the eBPF ecosystem come together to enable enhanced capabilities that users leverage to improve system performance, implement observability, security, and monitoring infrastructure, and enable high performance networking applications. We discuss challenges concerning eBPF's current direction and future progress, and the open research questions presenting themselves.

eBPF has seen wide adoption since its introduction into the Linux kernel. It remains under active development. While its design is guided by a strong focus on safety, flexibility, and performance, the primary driving force behind eBPF's continued evolution are its users within and outside the kernel.

More importantly, we illustrate how eBPF empowers users to rethink conventional operating system design. It allows users to perform radical changes to core kernel subsystems with relative ease and confidence, and innovate quickly in a fashion which was impractical

before for a popular OS kernel like Linux.

Acknowledgements

References

[1] 10 Developer Challenges Related to eBPF. https://www.emaallstars.com/ebpf-developer-challenges.html. Accessed: 2024-06-09.

[2] Advanced Network Policy. https://cilium.io/use-cases/network-policy/. Accessed: 2024-09-06.

[3] bpf-helpers(7) — Linux manual page. https://man7.org/linux/man-pages/man7/bpf-helpers.7.html. Accessed: 2024-06-09.

[4] BPF maps — The Linux Kernel documentation. https://www.kernel.org/doc/html/latest/bpf/maps.html. Accessed: 2024-09-06.

[5] BPF sk_lookup program — The Linux Kernel documentation. https://www.kernel.org/doc/html/latest/bpf/prog_sk_lookup.html. Accessed: 2024-09-06.

[6] BPF Type Format (BTF). https://docs.kernel.org/bpf/btf.html. Accessed: 2024-09-06.

[7] Classic BPF vs eBPF. https://docs.kernel.org/bpf/classic_vs_extended.html. Accessed: 2024-09-06.

[8] Control Groups. https://docs.kernel.org/admin-guide/cgroup-v1/cgroups.html. Accessed: 2024-09-06.

[9] HID-BPF — The Linux Kernel documentation. https://docs.kernel.org/hid/hid-bpf.html. Accessed: 2024-09-06.

[10] It's crowded in here! https://blog.cloudflare.com/its-crowded-in-here/. Accessed: 2024-09-06.

[11] Kconfig. https://www.kernel.org/doc/html/latest/kbuild/kconfig-language.html. Accessed: 2024-09-06.

[12] Kernel operations structures in BPF [LWN.net]. https://lwn.net/Articles/811631/. Accessed: 2024-09-06.

[13] Kernel TLS offload. https://docs.kernel.org/networking/tls-offload.html. Accessed: 2024-06-09.

[14] KRSI — the other BPF security module [LWN.net]. https://lwn.net/Articles/808048/. Accessed: 2024-09-06.

[15] libbpf Overview. https://docs.kernel.org/bpf/libbpf/libbpf_overview.html. Accessed: 2024-06-09.

[16] LSM BPF Programs — The Linux Kernel documentation. https://docs.kernel.org/bpf/prog_lsm.html. Accessed: 2024-09-06.

[17] net, sched: add clsact qdisc [LWN.net]. https://lwn.net/Articles/671458/. Accessed: 2024-09-06.

[18] Replacing HTB with EDT and BPF. https://netdevconf.info//0x14/pub/papers/55/0x14-paper55-talk-paper.pdf. Accessed: 2024-06-09.

[19] sysctl. https://www.kernel.org/doc/Documentation/sysctl/kernel.txt. Accessed: 2024-09-06.

[20] The extensible scheduler class [LWN.net]. https://lwn.net/Articles/922405/. Accessed: 2024-09-06.

[21] Unlocking network performance with xdp and ebpf. https://medium.com/@kcl17/unlocking-network-performance-with-xdp-and-ebpf-67c712128025. Accessed: 2024-06-09.

[22] The untold story of BPF | Kernel Recipes 2022. https://kernel-recipes.org/en/2022/talks/the-untold-story-of-bpf/, 2022. Accessed: 2024-09-06.

[23] bpf-next git commit 05924717ac70. https://git.kernel.org/pub/scm/linux/kernel/git/bpf/bpf-next.git/commit/?id=05924717ac70, 2023. Accessed: 2024-09-06.

[24] ebpf documentation - loader & verification architecture. https://ebpf.io/static/1a1bb6f1e64b1ad5597f57dc17cf1350/6515f/go.png, 2024. Accessed: 2024-06-09.

[25] Linux kernel - ebpf runtime. https://ebpf.io/infrastructure/, 2024. Accessed: 2024-06-09.

[26] ALEXEI, S. bpf: introduce function calls. https://lwn.net/Articles/741773/. Accessed: 2024-06-09.

[27] ANDRII, N. Add support for cgroup bpf_link. https://lore.kernel.org/bpf/CAEf4Bzbt_aVR5HyW_MuYvxT1fMBKD3xZOw3rEuHtj1i_T8NO_g@mail.gmail.com/T/. Accessed: 2024-06-09.

[28] ANDRII, N. Introduce pinnable bpf_link kernel abstraction. https://lore.kernel.org/bpf/87h7z7t620.fsf@toke.dk/T/, 2020. Accessed: 2024-06-09.

[29] ATHANASAKIS, M., ATHANASOPOULOS, E., POLYCHRONAKIS, M., PORTOKALIDIS, G., AND IOANNIDIS, S. The devil is in the constants: Bypassing defenses in browser jit engines. In *NDSS* (2015).

[30] AUTHORS, C. BPF Architecture - Object Pinning. https://docs.cilium.io/en/latest/bpf/architecture/#object-pinning, 2024. Accessed: 2024-06-09.

[31] AUTHORS, C. eBPF Datapath Introduction. https://docs.cilium.io/en/stable/network/ebpf/intro/, 2024. Accessed: 2024-06-09.

[32] BAVENDIEK, S. Attack surface analysis of the linux kernel based on complexity metrics. https://www.ndss-symposium.org/wp-content/uploads/2017/09/03_2_0.pdf, 04 2021. Accessed: 2024-06-09.

[33] BERLAKOVICH, F., NEUGSCHWANDTNER, M., AND BARANY, G. Look ma, no constants: Practical constant blinding in graalvm. In *Proceedings of the 15th European Workshop on Systems Security* (2022), pp. 36–42.

[34] BERSHAD, B. N., SAVAGE, S., PARDYAK, P., SIRER, E. G., FIUCZYNSKI, M. E., BECKER, D., CHAMBERS, C., AND EGGERS, S. Extensibility safety and performance in the spin operating system. In *Proceedings of the fifteenth ACM symposium on Operating systems principles* (1995), pp. 267–283.

[35] BERTIN, G. Xdp in practice: integrating xdp into our ddos mitigation pipeline. In *Technical Conference on Linux Networking, Netdev* (2017), vol. 2, The NetDev Society, pp. 1–5.

[36] BERTRONE, M., MIANO, S., RISSO, F., AND TUMOLO, M. Accelerating linux security with ebpf iptables. In *Proceedings of the ACM SIGCOMM 2018 Conference on Posters and Demos* (New York, NY, USA, 2018), SIGCOMM '18, Association for Computing Machinery.

[37] BHAT, S., AND SHACHAM, H. Formal Verification of the Linux Kernel eBPF Verifier Range Analysis.

[38] BOONSTOPPEL, P., CADAR, C., AND ENGLER, D. RWset: Attacking Path Explosion in Constraint-Based Test Generation. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems* (Berlin, Heidelberg, 2008), TACAS'08/ETAPS'08, Springer-Verlag, p. 351–366.

[39] BORKMANN, D. Advanced programmability and recent updates with tc's cls bpf. *Proc. NetDev 1* (2016). Tokyo, Japan.

[40] BORKMANN, D. On getting tc classifier fully programmable with cls bpf. *Proceedings of netdev 1* (2016). Seville, Spain.

[41] BORKMANN, D., AND FASTABEND, J. Combining kTLS and BPF for Introspection and Policy Enforcement. In *Linux Plumbers Conference* (2018), vol. 18.

[42] CILIUM. Bpf architecture. https://docs.cilium.io/en/latest/bpf/architecture/. Accessed: 2024-06-09.

[43] CORBET, J. Attaching eBPF programs to sockets [LWN.net]. https://lwn.net/Articles/625224/, 2014. Accessed: 2024-06-09.

[44] CULIC, I., VOCHESCU, A., AND RADOVICI, A. A low-latency optimization of a rust-based secure operating system for embedded devices. *Sensors 22*, 22 (2022), 8700.

[45] FASTABEND, J. bpf, bounded loop support work in progress [LWN.net]. https://lwn.net/Articles/756284/. Accessed: 2024-09-06.

[46] FINDLAY, W., SOMAYAJI, A., AND BARRERA, D. Bpfbox: Simple Precise Process Confinement with EBPF. In *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop* (New York, NY, USA, 2020), CCSW'20, Association for Computing Machinery, pp. 91–103.

[47] FOMICHEV, S. bpf: cgroup_sock lsm flavor [LWN.net]. https://lwn.net/Articles/896911/, 2022. Accessed: on 05/21/2024.

[48] FRASSETTO, T., GENS, D., LIEBCHEN, C., AND SADEGHI, A.-R. Jitguard: Hardening just-in-time compilers with sgx. pp. 2405–2419. Dallas, TX, USA.

[49] GREGG, B. A thorough introduction to bpftrace. https://www.brendangregg.com/blog/2019-08-19/bpftrace.html, 2019. Accessed: 2024-09-06.

[50] HØILAND-JØRGENSEN, T., BROUER, J. D., BORKMANN, D., FASTABEND, J., HERBERT, T., AHERN, D., AND MILLER, D. The express data path: Fast programmable packet processing in the operating system kernel. In *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, Association for Computing Machinery, pp. 54–66.

[51] HUMPHRIES, J. T., NATU, N., CHAUGULE, A., WEISSE, O., RHODEN, B., DON, J., RIZZO, L., ROMBAKH, O., TURNER, P., AND KOZYRAKIS, C. ghost: Fast & flexible user-space delegation of linux scheduling. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles* (Germany, 2021), pp. 588–604.

[52] HUO, T. Re: [patchset rfc] sched: Implement bpf extensible scheduler class - tejun heo. https://lore.kernel.org/bpf/Y5pM0ralEr6coT25@slm.duckdns.org/, 2022. Accessed: 2024-09-06.

[53] INTEL. Introduction to the Storage Performance Development Kit (SPDK). https://software.intel.com/content/www/us/en/develop/articles/introduction-to-the-storage-performance-development-kit-spdk.html, 2016. Accessed: 2024-09-06.

[54] IOVISOR. GitHub - bcc reference guide. https://github.com/iovisor/bcc/blob/master/docs/reference_guide.md. Accessed: 2024-06-09.

[55] IOVISOR. GitHub - bpf compiler collection. https://github.com/iovisor/bcc. Accessed: 2024-06-09.

[56] IOVISOR. GitHub - bpftrace. https://github.com/bpftrace/bpftrace. Accessed: 2024-06-09.

[57] KERNEL DEVELOPMENT COMMUNITY, T. Seccomp BPF (SEcure COMPuting with filters). https://docs.kernel.org/userspace-api/seccomp_filter.html. Accessed: 2024-01-12.

[58] KERNEL DEVS, L. Bpf syscall man. https://man7.org/linux/man-pages/man2/bpf.2.html. Accessed: 2024-06-09.

[59] KICINSKI, J., AND VILJOEN, N. eBPF Hardware Offload to SmartNICs: cls bpf and XDP. *Proceedings of netdev 1* (2016). Tokyo, Japan.

[60] KICINSKI, J., AND VILJOEN, N. Comprehensive XDP offload-handling the edge cases. *Proceedings of netdev 1* (2017). Seoul, Korea.

[61] LAMETER, C. An overview of non-uniform memory access. *Commun. ACM 56*, 9 (sep 2013), 59–54.

[62] LINGUAGLOSSA, D. ebpf offensive capabilities – get ready for next-gen malware. https://sysdig.com/blog/ebpf-offensive-capabilities/, 2023. Accessed: 2023-05-08.

[63] LUKKEN, C., FRASCARIA, G., AND TRIVEDI, A. ZCSD: a Computational Storage Device over Zoned Namespaces (ZNS) SSDs. *arXiv preprint arXiv:2112.00142* (2021).

[64] MATAN, L. The good, bad, and compromisable aspects of Linux eBPF. https://pentera.io/blog/the-good-bad-and-compromisable-aspects-of-linux-ebpf/, 2022. Accessed: 2024-06-09.

[65] MCCANNE, S., AND JACOBSON, V. The BSD Packet Filter: A New Architecture for User-level Packet Capture. In *USENIX winter* (1993), vol. 46. San Diego, CA.

[66] MIANO, S., DORIGUZZI-CORIN, R., RISSO, F., SIRACUSA, D., AND SOMMESE, R. Introducing smartnics in server-based data plane processing: The ddos mitigation use case. *IEEE Access 7* (2019), 107161–107170.

[67] MIANO, S., RISSO, F., BERNAL, M. V., BERTRONE, M., AND LU, Y. A framework for ebpf-based network functions in an era of microservices. *IEEE Transactions on Network and Service Management 18*, 1 (2021), 133–151.

[68] MICHAEL J., E. Introduction to the DWARF Debugging Format? https://dwarfstd.org/doc/Debugging%20using%20DWARF-2012.pdf, 2012. [Accessed 14-May-2024].

[69] NAKRYIKO, A. BTF deduplication and Linux kernel BTF. https://nakryiko.com/posts/btf-dedup/, 2018. Accessed: 2024-09-06.

[70] NELSON, L., BORNHOLT, J., GU, R., BAUMANN, A., TORLAK, E., AND WANG, X. Scaling Symbolic Evaluation for Automated Verification of Systems Code with Serval. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles* (New York, NY, USA, 2019), SOSP '19, Association for Computing Machinery, p. 225–242.

[71] NELSON, L., GEFFEN, J. V., TORLAK, E., AND WANG, X. Specification and verification in the field: Applying formal methods to BPF just-in-time compilers in the Linux kernel. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)* (Nov. 2020), USENIX Association, pp. 41–61.

[72] ORGANIZATION, L. K. bpf: Disallow unprivileged bpf by default. https://lore.kernel.org/bpf/20211027233943. kehyrdbibp2d2u4c@gupta-dev2.localdomain/T/, 2021. Accessed: 2023-05-24.

[73] ORGANIZATION, L. K. eBPF Instruction Set Specification, v1.0. https://docs.kernel.org/bpf/instruction-set. html, 2023. Accessed: 2024-09-06.

[74] ORGANIZATION, L. K. eBPF Syscall. https://docs. kernel.org/userspace-api/ebpf/syscall.html, 2023. Accessed: 2024-06-09.

[75] ORGANIZATION, L. K. eBPF verifier. https://docs. kernel.org/bpf/verifier.html, 2023. Accessed: 2024-06-09.

[76] ORGANIZATION, L. K. verifier.c bpf kernel - kernel/git/-torvalds/linux.git. https://git.kernel.org/pub/scm/ linux/kernel/git/torvalds/linux.git/tree/kernel/ bpf/verifier.c, 2023. Accessed 2023-05-24.

[77] OUSTERHOUT, A., FRIED, J., BEHRENS, J., BELAY, A., AND BALAKRISHNAN, H. Shenango: Achieving high CPU efficiency for latency-sensitive datacenter workloads. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)* (2019), pp. 361–378.

[78] PARK, Y. G., AND GOLDBERG, B. Escape analysis on lists. *SIGPLAN Not. 27*, 7 (jul 1992), 116–127.

[79] PAUL, M. What is RCU, Fundamentally? https://lwn.net/ Articles/773605/, 2017. Accessed: 2024-09-06.

[80] PRINZ, R. GitHub - rprinz08/hBPF: hBPF = eBPF in hardware — github.com. https://github.com/rprinz08/hBPF. Accessed: 2024-09-06.

[81] PROJECTS, T. L. F. DPDK. https://www.dpdk.org/, 2021. Accessed: 2024-09-06.

[82] RICE, L. *Learning eBPF*. " O'Reilly Media, Inc.", 2023.

[83] SAMSUNG. Hardening hostile code in ebpf - analysis on kernel self-protection: Understanding security and performance implication. https://samsung.github.io/kspp-study/ bpf.html, 2020. Accessed: 2024-06-09.

[84] SMALL, C. A., AND SELTZER, M. I. Vino: An integrated platform for operating system and database research. Cambridge, MA 02138.

[85] TEJUN, H. sched: Implement BPF extensible scheduler class. https://lwn.net/Articles/916290/, 2022. Accessed: 2024-09-06.

[86] TRAIAIA, K. Issue 002 | Programming the Kernel with eBPF. https://www.kerno.io/blog/ programming-the-kernel-with-ebpf, 2023. Accessed: 2024-01-12.

[87] VERNET, D. Long-lived kernel pointers in BPF [LWN.net]. https://lwn.net/Articles/900749/, 2022. Accessed: 2023-06-09.

[88] VISHWANATHAN, H., SHACHNAI, M., NARAYANA, S., AND NAGARAKATTE, S. Sound, Precise, and Fast Abstract Interpretation with Tristate Numbers. In *Proceedings of the 20th IEEE/ACM International Symposium on Code Generation and Optimization* (2022), CGO '22.

[89] ZHANG, I., RAYBUCK, A., PATEL, P., OLYNYK, K., NELSON, J., LEIJA, O. S. N., MARTINEZ, A., LIU, J., SIMPSON, A. K., JAYAKAR, S., PENNA, P. H., DEMOULIN, M., CHOUDHURY, P., AND BADAM, A. The Demikernel Datapath OS Architecture for Microsecond-Scale Datacenter Systems. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles* (New York, NY, USA, 2021), SOSP '21, Association for Computing Machinery, p. 195–211.

[90] ZHONG, Y., LI, H., WU, Y. J., ZARKADAS, I., TAO, J., MESTERHAZY, E., MAKRIS, M., YANG, J., TAI, A., STUTSMAN, R., ET AL. XRP: In-Kernel Storage Functions with eBPF. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)* (Carlsbad, CA, USA, 2022), pp. 375–393.

[91] ZHONG, Y., WANG, H., WU, Y. J., CIDON, A., STUTSMAN, R., TAI, A., AND YANG, J. BPF for storage: an exokernel-inspired approach. In *Proceedings of the Workshop on Hot Topics in Operating Systems* (Ann Arbor, MI, USA, 2021), pp. 128–135.

[92] ZOHARI, A. eBPF Profiling: The Key to System Insights. https://www.groundcover.com/ebpf/ebpf-profiling, 2023. Accessed: 2024-01-12.

A eBPF Operations

Table 1 is a list of common eBPF operations along with their descriptions [73].
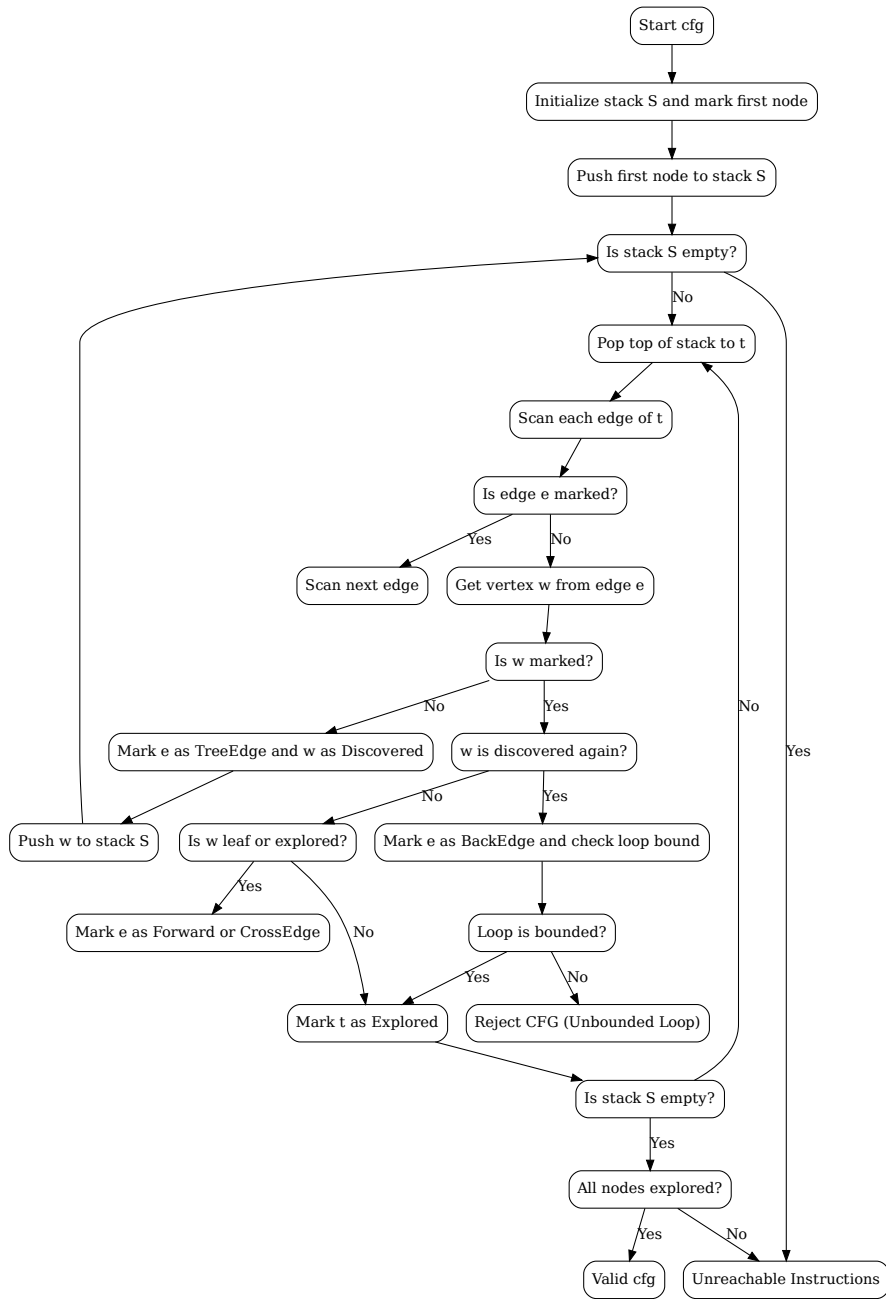
B Control-Flow Graph

Figure 6: The flowchart illustrates the `check_cfg()` [76] function of the eBPF verifier used to verify a Control-Flow Graph (CFG). It begins by initializing a stack (S) and marking the first node. Through depth-first search (DFS), it explores nodes (t) and scans edges (e). Edge classifications such as tree-edge, back-edge, and Forward/cross edge are determined based on traversal states. The function ensures all nodes are explored and checks for unreachable instructions.

Table 1: List of eBPF Operations

Operation	Description
BPF_ALU_ADD	Add two registers
BPF_ALU_SUB	Subtract two registers
BPF_ALU_MUL	Multiply two registers
BPF_ALU_DIV	Divide two registers
BPF_ALU_MOD	Modulus of two registers
BPF_ALU_OR	Bitwise OR of two registers
BPF_ALU_AND	Bitwise AND of two registers
BPF_ALU_LSH	Shift left of register
BPF_ALU_RSH	Shift right of register
BPF_ALU_NEG	Negate value of register
BPF_ALU_XOR	Bitwise XOR of two registers
BPF_ALU_MOV	Move a value from one register to another
BPF_ALU_ARSH	Arithmetic right shift of register
BPF_ALU_END	End marker for ALU operations
BPF_JMP_JEQ	Jump if equal
BPF_JMP_JNE	Jump if not equal
BPF_JMP_JA	Jump always
BPF_JMP_JGT	Jump if greater than
BPF_JMP_JGE	Jump if greater than or equal
BPF_JMP_JLT	Jump if less than
BPF_JMP_JLE	Jump if less than or equal
BPF_JMP_JSET	Jump if bitwise AND with immediate is true
BPF_JMP_CALL	Call function
BPF_JMP_EXIT	Terminate execution
BPF_JMP_ALU64	Perform 64-bit arithmetic and jump
BPF_JMP_X	Reserved for future use
BPF_JMP_ADD	Add offset to register
BPF_JMP_MUL	Multiply register by scale
BPF_JMP_NEG	Negate register
BPF_JMP_AND	Bitwise AND with register
BPF_JMP_OR	Bitwise OR with register
BPF_JMP_XOR	Bitwise XOR with register
BPF_JMP_MOV	Move register to another
BPF_JMP_ARSH	Arithmetic right shift of register
BPF_JMP_END	End marker for JMP operations
BPF_STX	Store into memory
BPF_LDX	Load from memory
BPF_ST	Store value
BPF_LD	Load value