Runtime Verification for Programmable Switches

Apoorv Shukla^(D), Kevin Hudemann, Zsolt Vági, Lily Hügerich^(D), Georgios Smaragdakis^(D),

Artur Hecker, Stefan Schmid[®], and Anja Feldmann

Abstract-We introduce a runtime verification framework for programmable switches that complements static analysis. To evaluate our approach, we design and develop P6, a runtime verification system that automatically detects, localizes, and patches software bugs in P4 programs. Bugs are reported via a violation of pre-specified expected behavior that is captured by P6. P6 is based on machine learning-guided fuzzing that tests P4 switch non-intrusively, i.e., without modifying the P4 program for detecting runtime bugs. This enables an automated and realtime localization and patching of bugs. We used a P6 prototype to detect and patch existing bugs in various publicly available P4 application programs deployed on two different switch platforms, namely, behavioral model (bmv2) and Tofino. Our evaluation shows that P6 significantly outperforms bug detection baselines while generating fewer packets and patches bugs in large P4 programs, e.g., switch.p4 without triggering any regressions.

Index Terms-Programmable networks, P4, verification.

I. INTRODUCTION

THE introduction of programmable networks is a paradigm shift in network design, management, and operation. Network devices on the data plane, e.g., switches, that traditionally have fixed and vendor-specific network functionality and rely on proprietary hardware and software, can now be programmed and customized by network operators. The P4 language [1], [2] was introduced to enable the programmability and customization of data plane functionalities in network devices. P4 is an open-source domain-specific language designed to allow programming of packet forwarding planes, and is now supported by a number of network vendors.

While programmable networks enable to break the tie between vendor-specific hardware and proprietary software, they facilitate an independent evolution of software and

Manuscript received 6 March 2022; revised 4 October 2022; accepted 1 December 2022; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Balasubramanian. This work was supported in part by the European Research Council (ERC) Starting Grant ResolutioNet under Grant ERC-StG-679158; and in part by the Austrian Science Fund (FWF) Project I 4800-N (ADVISE), 2020–2023. (*Corresponding author: Apoorv Shukla.*)

Apoorv Shukla and Artur Hecker are with the Huawei Munich Research Center, 80992 Munich, Germany (e-mail: apoorv.shukla1@huawei.com).

Kevin Hudemann is with SAP, 69190 Walldorf, Germany.

Zsolt Vági is with Swisscom AG, 3048 Bern, Switzerland.

Lily Hügerich is with the Technical University of Berlin (TU Berlin), 10623 Berlin, Germany.

Georgios Smaragdakis is with the Delft University of Technology (TU Delft), 2600 CD Delft, The Netherlands.

Stefan Schmid is with the Technical University of Berlin (TU Berlin), 10623 Berlin, Germany, and also with the Fraunhofer Institute for Secure Information Technology (Fraunhofer SIT), 64295 Darmstadt, Germany.

Anja Feldmann is with the Max Planck Institute for Informatics, 66123 Saarbrücken, Germany.

Digital Object Identifier 10.1109/TNET.2023.3234931

hardware. With the P4 language, one can define in a P4 program, the instructions for processing the packets, e.g., how the received packet should be read, manipulated, and forwarded by a network device, e.g., a switch. Nevertheless, with the new capabilities, new challenges in P4 software verification, i.e., ensuring that the software fully satisfies all the expected requirements, have been unleashed. The P4 switch behavior depends on the *correctness* of the P4 programs running on them. We realize that a bug in a P4 program, i.e., a small fault such as a missing line of code or a fat finger error, or a vendor-specific implementation error, can trigger unexpected and abnormal switch behavior. In the worst case, it can result in a network outage, or even a security compromise [3].

Problem Statement: In this paper, we examine and verify the behavior of P4 switches after the P4 programs are deployed. We pose the question: "Is it possible to automatically detect, localize, and patch software bugs in a P4 program running on P4 switches?". We believe that being able to answer this question, even partially, unlocks full potential of programmable networks, improves their security, and will hence increase their penetration in operational and mission-critical networks.

Recently, a panoply of P4 program verification tools [4], [5], [6], [7], [8], [9] has been proposed. These verification systems, however, fail to repair the P4 program containing bugs. Most of them [4], [5], [6], [7] aim to statically verify user-defined P4 programs which are later, compiled to run on a target switch. They mostly find bugs that violate the memory safety properties, e.g., invalid memory access, buffer overflow, etc. Furthermore, they are prone to false positives and are unable to verify the *runtime* behavior on real packets. In addition, classes of bugs, e.g., checksum-related, ECMP (Equal-Cost Multi-path) hash calculations-related or platform-dependent bugs, cannot be detected by static analysis approaches. Since runtime verification aims to verify the actual behavior against the *expected* behavior of a switch by passing specially-crafted input packets to the switch and observing the behavior, such verification is complementary to static analysis. Note, the detection of bugs causing the abnormal runtime behavior is a complex and challenging task. In particular, the P4 switch does not throw any runtime exceptions. Furthermore, the detection of bugs can be a nightmare if there is no output, i.e., packets are silently dropped instead of being forwarded. Thus, the runtime verification of the switch behavior is crucial.

A useful approach to verify the runtime behavior is fuzz testing or fuzzing [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], a well-known dynamic program testing

1558-2566 © 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

IEEE/ACM TRANSACTIONS ON NETWORKING

technique that generates semi-valid, random inputs which may trigger abnormal program behavior. However, for fuzzing to be efficient, intelligence needs to be added to the input generation, so that the inputs are not rejected by the parser and it maximizes the chances of triggering bugs. This becomes crucial especially in networking, where the input space is huge, e.g., a 32-bit destination IPv4 address field in a packet header has 2^{32} possibilities. With the 5-tuple flows, the input space gets even more complex and large. To make fuzzing more effective, we consider the use of machine learning, to guide the fuzzing process to generate smart inputs that trigger abnormal target behavior. Recently, Shukla et al. [20] have shown that Reinforcement Learning (RL) [21], [22] can be used to train a verification system. We build upon [20] by adding (a) static analysis to the fuzzing process to significantly reduce the input search space, and thus, adding input structure awareness, and (b) support for platform-dependent bug detection.

Even if a bug in a P4 program is detected, the localization of code statements in the P4 code that are responsible for the bug, is non-trivial. The difficulty stems from the fact that practical P4 programs can be large with a dense conditional structure. In addition, the same faulty statements in a P4 program may be executed for both passed as well as failed test cases and thus, it gets hard to pinpoint the actual faulty line/s of code. Tarantula [23], [24], [25] is a dynamic program analysis technique that helps in fault localization by pinpointing the potential faulty lines of code. To localize the software bugs, we tailor Tarantula for generic software to P4 programs by building a localizer called P4Tarantula and integrating it with the bug detection of machine learning-guided fuzzing. In this paper, we combine these two approaches to detect and localize bugs in P4 programs in real-time.

Our P6 System: In P4, the automated program repair [26] is an uncharted territory and becomes increasingly important as the software development lifecycle in programmable networks is short [27] with insufficient testing. In this paper, we show that due to the structure of P4 programs, it is possible to automate patching of platform-independent bugs (P4 program-specific software bugs) in P4 programs, if the patch is available. To this end, we present P6, P4 with runtime Program Patching, a novel runtime P4-switch verification system that (a) detects, (b) localizes, and (c) patches software bugs in a P4 program. P6 improves existing work on machine learning-guided fuzzing [20] in P4 by extending it and augmenting it with: (a) automated localization, and (b) runtime patching. P6 relies on the combination of static analysis of the P4 program and Reinforcement Learning (RL) technique to guide the fuzzing process to verify the P4-switch behavior at runtime.

In a nutshell, in P6, the first step is to capture the expected behavior of a P4 switch, which is achieved using information from three different sources: (i) the control plane configuration, (ii) queries in (§ III-B1), a query language which we leverage to describe expected behavior using conditional statements, and (iii) accepted header layouts, e.g., IPv4, IPv6, etc., learned via static analysis of the P4 program. If the *actual* runtime behavior to the test packets generated via machine-learning guided fuzzing differs from the *expected*

behavior through the violation of the p4q queries, it signals a bug to P6 which then identifies a patch from a library of patches. If the patch is available, P6 modifies the original P4 program to fix the bug signaled by the p4q queries. Then, the patched P4 program is subjected to sanity and regression testing.

We develop a prototype of P6 and evaluate it by testing it on eight P4₁₆ application programs from switch.p4 [28], P4 tutorial solutions [29], and NetPaxos codebase [30] across two P4 switch platforms, namely, behavioral model version 2 (bmv2) [31] and Tofino [32]. Our results show that P6 successfully detects, localizes, and patches diverse bugs in all P4₁₆ programs while significantly outperforming bug detection baselines without introducing any regressions. Part of this work appeared in [20], [33], and [34].

Contributions: Our main contributions are:

• We introduce a runtime verification framework for programmable switches that complements previously proposed static analysis approaches.

• We design, implement, and evaluate P6, the first end-toend P4 runtime verification system that detects, localizes, and patches bugs in P4 programs non-intrusively. (§III)

• We observe that the success of P6 relies on the increased patchability of P4 program from old $(P4_{14})$ to the new version $(P4_{16})$. We confirm this by studying the code of P4 applications in the two different P4 versions. (§II)

• We evaluate our P6 prototype on a P4 switch running eight $P4_{16}$ programs (including switch.p4 with 8,715 LOC) from publicly available sources [28], [29], [30] across two platforms, namely, behavioral model and Tofino. Our results show that P6 non-intrusively detects both platform-dependent and platform-independent bugs, and significantly outperforms state-of-the-art bug detection baselines. (§V)

• We show that in case of platform-independent bugs, P6 can localize bugs and fix the P4 program, when the patch is available, without causing any regressions or introducing new bugs in an automated fashion. (§III, §V)

• To ensure reproducibility and facilitate follow-up work, we release the P6 software and library of ready patches for all existing bugs in the P4 programs to the research community [35].

II. P4 VERIFICATION: CHALLENGES & OPPORTUNITIES

In this section, we describe the P4 packet processing pipeline. Then, we outline the challenges in discovering bugs in P4 programs or switches, and we motivate the need for runtime verification. We conclude this section by characterizing the evolution and structure of P4 programs, that, to our surprise provides opportunities for automated P4 program patching.

A. Packet Processing Pipeline of P4

P4 [1], [2] is a domain-specific language comprising of packet-processing abstractions, e.g., headers, parsers, tables, actions, and controls. The P4 packet processing pipeline evolved from [36] to its current form P4₁₆ [2], see Figure 1. In P4₁₆ packet processing pipeline, there are

SHUKLA et al.: RUNTIME VERIFICATION FOR PROGRAMMABLE SWITCHES



Fig. 1. An example of a platform-independent bug in P4₁₆ packet processing pipeline.



Fig. 2. An example of a platform-dependent bug in $P4_{16}$ packet processing pipeline.

six programmable blocks that are platform-independent, namely, ingress parser, ingress match-action, ingress deparser, egress parser, egress match-action, and egress deparser. The programmable blocks are annotated with a solid line in Figures 1 and 2.

The ingress parser transforms the packet from bits into headers according to a parser specification provided by the programmer. After parsing, an ingress match-action (also called ingress control function) decides how the packet will be processed. Then, the packet is queued for egress processing in the ingress deparser. Upon dequeuing, the packet is processed by an egress match-action (also called egress control function). The egress deparser specification dictates how packets are deparsed from separate headers into a bit representation on output, and finally, the packet leaves the switch. Note that both ingress and egress match-actions (control functions) direct the packet through any number of match-action tables.

In the P4₁₆ packet processing pipeline, there are also two platform-dependent blocks (annotated with dashed lines in Figures 1 and 2), that rely on proprietary implementations of the hardware vendors and are non-programmable. These blocks are the packet replication engine (PRE) and the buffer queuing engine (BQE).

B. Challenges: Runtime Bugs in P4

Bugs or errors can occur at any stage in the P4 pipeline. If a bug occurs in any of the programmable blocks, then the bug is platform-independent and software patching can solve the problem. If the bug appears in the non-programmable or platform-dependent blocks, namely, the PRE or BQE, then the vendor has to be informed to fix the issue as the implementation is vendor-specific. P4 program verification systems [4], [5], [6], [7] are able to detect bugs using static analysis. Unfortunately, static analysis is (i) prone to false positives, (ii) cannot detect platform-dependent bugs, and (iii) cannot detect runtime bugs that require to actively send real packets.

As an example, consider the scenario in Figure 1 (solid line blocks) that illustrates part of the implementation of Layer-3 (L3) switch, provided in the P4 tutorial solutions [29]. Here, the parser does not check if the IPv4 header contains IPv4 options or not, i.e., if the IPv4 ihl field is equal to 5 or not. When updating the IPv4 checksum of the packets during egress processing, IPv4 options are not taken into account, hence for those IPv4 packets with options, the resulting checksum is wrong causing such packets to be forwarded and *incorrectly* dropped at the next hop. This leads to anomalies in network behavior. Other bugs that fall in this category are those related to IPv4/6 checksum in the packet (see Figure 5 later). Such bugs are *platform-independent*, as they only result from programming errors.

Figure 4 illustrates a simple scenario where due to a bug or fault, the packet reaching a P4 switch has a time-to-live (TTL) field in the IP header with value as 0. The expected behavior is that the packet gets dropped, but currently, there is no such check in the P4 code. Thus, the switch forwards the packet and decreases the value of the TTL field, causing it to be increased to 255 incorrectly. This happens due to wraparound caused by underflow in an 8-bit field. Such an anomaly can be nontrivial to detect and localize. In addition, it can be responsible for the abnormal behavior of a P4 switch.

As illustrated in Figure 5, the problem lies in the fact that the P4 program fails to specify that the IPv4 checksum inside the packet needs to be verified before forwarding. An adversary can easily intercept the packets and modify them (e.g., as a Man-in-the-Middle attack (MitM)) and does not even need to recalculate the checksum. Therefore, additional information can be inserted even into encrypted packets. When such a malicious or malformed packet arrives at the P4 switch, it selects the corresponding action based on the match-action tables and forwards the packet without verifying the checksum. Such checksum-related bugs may inflict serious damages to critical servers and can be a nightmare to debug. We must clarify that P6 signals a bug whenever the actual behavior differs from the expected behavior. This may or may not be a vulnerability of security compromise. In Figure 4, accepting the packets with abnormal TTL, i.e., TTL < 2 signals a violation of expected behavior where actual behavior differs from the specification. Therefore, it is a bug. For Figure 5, packet with the wrong checksum is accepted but it should be dropped in the first place. Since the actual





4 program from $P4_{14}$ to $P4_{16}$ in bmv2 platform [31], i.e., it does not include i.e., includes explicit ingress deparser and egress explicit ingress deparser and egress parser.

Fig. 3. Evolution of the P4 program structure from P414 to P416 version.



Fig. 4. P4 switch running the P4 program does not check if the time-to-live (TTL) value in the packet is 0. Blue arrows show the expected, red dashed arrows show the actual path.



Fig. 5. P4 switch running the P4 program does not check the faulty IPv4 checksum in the packet. Blue arrows show the expected path, red dashed arrows show the actual path.

behavior differs from expected behavior, it is designated as a bug by P6. It is noteworthy that in all of these unexpected behaviors, the malicious payload in the processed packet can be leveraged to compromise the target system via algorithmic complexity attacks (ACAs) [39], [40] or via a volumetric DoS attack.

For a *platform-dependent bug*, consider the scenario shown in Figure 2 (dashed line blocks). Here, we assume a P4 program implements at least two match-action tables. Any table except the last one could be a longest prefix match (LPM) table, offering unicast, clone and drop actions (ingress match-action block). The last match-action table implements an access control list (ACL). So, the packets can either be dropped or forwarded according to the chosen actions by the previous tables. In this case, it is possible that conflicting forwarding decisions are made. Consider packets are matched by the first table (Table I) and a clone decision is made, later, those are dropped by the ACL table (Table n). In such a case, the forwarding behavior depends on the implementation of the PRE, which is platform-dependent. PRE implementation for the SimpleSwitch target in the behavioral model (bmv2) is illustrated in Figure 2. It would drop the original packet, however, forward the cloned packer copy. Similar bugs can occur, if instead of the clone action, the resubmit action is



(c) The patchable code percentage for bmv2 switch (b) The particulate code percentage for bin2 switch platform [31], i.e., it does not include explicit ingress deparser and egress parser in all different $P4_{16}$ programs from [29], [30].



(d) The patchable code percentage for generic PSA switch platforms [37], [38] like Tofino [32], i.e., includes explicit ingress deparser and egress parser in all different $P4_{16}$ programs from [29], [30].

chosen (blue). Another bug can be found when implementing multicast (green).

The above motivates us to turn our attention to runtime detection of bugs. Runtime verification is a useful and complementary tool in the P4 verification repertoire that detects both *platform-independent bugs* resulting from programming errors as well as *platform-dependent bugs*.

C. Opportunities for Patching: P4 Program Structure

In the evolution of P4, there are two recent versions: $P4_{14}$ [41] and $P4_{16}$ [2]. $P4_{16}$ allows programmers to use definitions of a target switch-specific architecture, PSA (Portable Switch Architecture) [37], [38]. P4₁₆ is an upgraded version of P414. In particular, a large number of language features have been eliminated from $P4_{14}$ and moved into libraries including counters, checksum units, meters, etc., in $P4_{16}$. $P4_{14}$ allowed the programmer to explicitly program three blocks: ingress parser (including header definitions of accepted header layouts), ingress control and egress control functions. Recall that P416 allows to explicitly program six programmable blocks (see Figure 1).

By analyzing programs in the $P4_{14}$ and $P4_{16}$ versions, we realize that as more blocks of the P4 program get programmable, there is more onus on the programmer to write a program that behaves as expected (when it gets compiled and deployed on the P4 switch). Missing checks or fat finger errors can cause havoc in the network. However, this is a blessing in disguise as the more programmable the code is, the more patchable it is. Thus, programming errors can be fixed. Figures 3a and 3b illustrate that the potentially patchable code percentage increases from $P4_{14}$ to P4₁₆ in all applications (excluding calculator) from P4 tutorial solutions [29] and NetPaxos codebase [30] in behavioral model (bmv2) switch platform [31] and other generic PSA switch platforms [37], [38], e.g., Tofino [32] respectively. Figures 3c and 3d illustrate the patchable code percentage in the latest $P4_{16}$ version. The patchable code percentage comes from the six programmable blocks in $P4_{16}$. Roughly, whatever is programmable, is patchable. In principle, around 40-45% of a P4 program is patchable in $P4_{16}$ programs for behavioral model (bmv2) switch platform [31] (Figure 3c). This increases to 50-55% if the ingress deparser and egress

SHUKLA et al.: RUNTIME VERIFICATION FOR PROGRAMMABLE SWITCHES



Fig. 6. P6 in Action: depicting the automated detection, localization and patching of a bug in a L3 switch P4 program [29].



Fig. 7. P6 Workflow. Modules of P6 (in solid green boxes).

parser are programmable for other target switch platforms, e.g., Tofino [32] (Figure 3d).

In particular, the parser and header definitions account for 20-40% of the total patchable code. If there is no bug in parser or header, packets with incorrect header get dropped. However, the bug still can be either in the non-patchable platform-dependent block or in the application code logic or deparser which is patchable as it is platform-independent.

Observation: From $P4_{14}$ to $P4_{16}$, P4 program possesses twice as many programmable blocks increasing the chances for patchability. Bugs detected in the platform-independent part can be localized and patched; a platform-dependent bug may not be patchable if it is hardware-related.

III. P6: SYSTEM DESIGN

A. P6: Overview

The goal of P6 (see Figure 7) is to detect, localize and patch the software bugs in a P4 program at runtime with minimal human effort. This is achieved by verifying the *actual* runtime behavior against the *expected* behavior of a P4 switch running a pre-compiled P4 program to the incoming packets.

The P6 contains three main modules:

(1) **Fuzzer** : Generates test packets using RL-guided fuzzing, static analysis, and p4q queries (§III-B.1) to the P4 switch running the pre-compiled P4 program. (§III-B)

(2) Localizer : P4Tarantula is the Localizer which pinpoints faulty lines of code causing bugs in the P4 program. (§III-C)
(3) Patcher : Automates patching of the bugs localized by P4Tarantula Localizer, if patchable. Then, Patcher compiles and loads the patched P4 program on the P4 switch. (§III-D)

P6 Workflow. P6 is a closed-loop control system. Through a pre-generated dictionary from control plane configuration, p4q queries, and static analysis of a P4 program, the expected runtime behavior of the P4 switch is captured and sent as an input to the Fuzzer containing the RL Agent and the Reward System (§III-B). As shown in Figure 7, the Fuzzer selects appropriate mutation actions such as add/delete/modify bytes in a packet to generate test packets towards the P4 switch running the pre-compiled P4 program ①. If the actual runtime behavior towards the packets defies the expected behavior through the violation of the p4q queries, it signals a bug in the form of a reward as a feedback to the Reward System which is then, exploited by the RL Agent to improve the training process by selecting better mutation actions on the packet 2. After the bug detection, the Fuzzer automatically triggers Localizer (§III-C), P4Tarantula (only for platformindependent bugs; for platform-dependent bugs, the vendor is informed) which pinpoints the faulty line of code 3 to trigger the Patcher (§III-D) which searches for the appropriate patch from a library of patches for the corresponding P4 program ④. If the patch is available, Patcher modifies the original P4 program, compiles and loads it on the P4 switch and checks if the bug is no longer triggered by p4q queries by repeating the whole-cycle and executing sanity and regression testing 5. Note, P6 is non-intrusive and thus, requires no modification to the P4 program for testing before patching.

P6 in Action. Before we dive into the details of Fuzzer, Localizer and Patcher, we demonstrate the operation of P6. Figure 6 illustrates how P6 detects, localizes, and patches an existing bug in a layer-3 (L3) switch P4 source code (program) from [29] in an automated fashion. The left part of Figure 6 shows the P4 program containing a platform-independent bug in the parser code, i.e., no header field validation is implemented, hence all IPv4 packets are *incorrectly* accepted by the parser. After the P4 program is deployed on the P4 switch, P6 is triggered. Initially, the Fuzzer detects the bug violating the corresponding p4q query based on the feedback (reward) received from the P4 switch. Then, it triggers the P4Tarantula for localization (shown in the center of Figure 6) where it pinpoints the problematic part of the code (highlighted). Afterwards, the Patcher is triggered



Fig. 8. Fuzzer. Reward System (in yellow) and Agent (in pink).

automatically, patching the necessary problematic parts of the code, i.e., adding header field verification statements (highlighted in right), after checking if the patch was indeed missing from the P4 program. Finally, Patcher automatically compiles [42] and deploys the patched P4 program on the P4 switch, and triggers P6 to ensure that the patches caused no regressions and fixed the detected bug.

B. Fuzzer: RL-Guided Fuzzing

The goal of Fuzzer is to detect the runtime bugs discussed in §II-B. We improve [20] by augmenting Fuzzer with the static analysis of a P4 program which makes the Fuzzer aware of the input structure or accepted header layouts, e.g., IPv4, IPv6, etc., and thus, it significantly reduces the input search space. Indeed, techniques to further reduce the input search space within the accepted headers are discussed in [43], which can be augmented to static analysis. We guide the mutationbased white-box fuzzing [12] via RL [21], [22]. The feedback in the form of *rewards* is received from the switch based on the evaluation of actual against expected runtime behavior. Note, the expected behavior is determined using the static analysis, the control plane configuration, i.e., forwarding rules and p4q queries (§III-B.1). p4q queries are conditional queries (if-then-else) where each query has multiple conditions and each condition acts as a test case. A violation of a test case represents a bug detection.

Reinforcement Learning (RL). Reinforcement learning [21], [22] is a machine learning technique that aims at enabling an Agent to learn how to interact with an environment, based on a series of reinforcements, i.e., rewards or punishments received from the target environment, in our case, a switch. The Agent observes the switch and chooses an action to be executed. After the action is executed, the Agent receives a reward or punishment from the switch. While the goal of learning is to maximize the rewards, we argue it is equally crucial to design a machine learning model which is general enough for any kind of target environment. To detect the bugs triggered by fuzzing, one can observe the output of the target switch in response to the input packets. Thus, reinforcement learning allows developing a Reward System where feedback in the form of rewards from the switch trains the Agent and thus, guides the fuzzing process.

In our RL-based model, we define states, actions, and rewards as follows:

States: The sequence of bytes forming the packet header. **Actions:** The set of mutation actions for each individual packet header field, e.g., add, modify or delete bytes at a given position in the packet header. Note, the add and modify actions either use random bytes or bytes from a pre-generated dictionary (explained below).

Rewards: The *Agent* can immediately receive the reward, after a mutated packet was sent to the target switch and the results of the execution are evaluated. It is likely to experience sparse rewards when most of the sent packets do not trigger any bug. Thus, the reward is defined as 0, if the packet did not trigger a bug and 1, if the packet successfully triggered a bug.

The input to the Fuzzer is a dictionary (hereafter, referred to as dict) that comprises information extracted from static analysis, the control plane configuration, and the queries defined with p4q (§III-B.1). The static analysis is used to derive the input structure awareness such as accepted header layouts and available header fields in the P4 program. The control plane configuration comprises the forwarding table contents and the platform-dependent configuration. Boundary values for the header fields may be extracted from the p4q queries, i.e., when queries explicitly compare packet header fields with values, e.g., TTL > 0.

Figure 8 depicts the Fuzzer workflow. In step 0 (initialization), the Reward System receives the dict as an input. Then, the Agent observes the current state or the current packet header (see the initialization in §III-B.2). The observed state is the input for the neural networks of the Agent (§III-B.2), which outputs the appropriate mutation action. The selected action is applied for the given packet, and the packet is sent to the P4 switch. After the packet is processed by the switch, the behavior is evaluated, the reward of 1 is generated when the p4q query specifying the expected behavior is violated and returned to the Agent. In particular, the packet which was sent to the P4 switch is saved together with a final verdict (pass or fail). A packet's verdict is considered either passed: if the generated reward is equal to 0, i.e., *actual* runtime behavior matches *expected* behavior when the p4q query is not violated or failed: if the generated reward is equal to 1, i.e., actual runtime behavior does not match expected behavior when the p4q query is violated. Then, the Agent (§III-B.2) uses the received reward to improve the action selection in subsequent executions (exploitation).

1) p4q Query Language: Before diving into the details of the Agent training, we explain the query language, p4q [20], used for specifying the *expected* switch behavior. To achieve the goal of an automated runtime verification system, P6 system must query the actual runtime behavior of a P4 switch against a specification defining the *expected* behavior. To extend the query repertoire of p4q from [20], we augment it with platform-dependent queries. In a nutshell, p4q queries are used to compare *expected* against *actual* switch behavior. Currently, our p4q queries act as specification of Layer 3 IPv4 protocol. In particular, we mostly focus on 20 bytes of Layer 3 IPv4 header space (from IPv4 Version field till Destination IP Address). Our P6 system is trained with such queries on a P4 program and thus, a trained model can find errors in other P4 programs quickly by choosing mutation actions within that IP space that lead to rewards, i.e., bugs or unexpected behavior in our case. However, we believe that

SHUKLA et al.: RUNTIME VERIFICATION FOR PROGRAMMABLE SWITCHES



Fig. 9. p4q Queries. Queries 1-6 represent platform-independent, and Query 7-8 represent platform-dependent queries respectively.



by training the P6 Agent for other packet header spaces such as L4 and L2, we can increase the bug scope of P6 to a larger header space.

p4q queries. In a p4q query, the behavior is expressed using if-then-else statements in the form of tuples. The programmer can specify conditions for packets to fulfill at ingress of the switch (if), with corresponding conditions to fulfill at egress (then). In addition, the programmer can describe alternative conditions (else), e.g., if the condition of the then branch is not fulfilled at egress. *To automate the usage of* P6, *an option to execute all the queries of* p4q *with a single command is provided* (see §IV). To define these conditions, the p4q syntax and grammar are used.

p4q Grammar. Figure 10 depicts the grammar and constructs defined in p4q. The p4q grammar allows common boolean expressions and relational operators as they can be found in many programming languages like C, Java or Python, to ease the work for the programmer. The boolean expressions and relational operators have the same semantics as common logical operators and expressions. Variables can either be integers, header fields, header field values, or the evaluation result of the primitive methods, e.g., calcChksum() and table_val(). Each header has a prefix (ing. or egr.)

indicating if it is the packet arriving at ingress or exiting the switch at egress.

Figure 9 illustrates an example of how the packet processing behavior of an IPv4 layer 3 (L3) switch, written in P4, can be queried easily using p4q. Query 1 (lines 1-3), defines that incoming packets with a wrong IPv4 checksum are expected to be dropped. Similarly, the following four queries (lines 4-13) express the validation of the IPv4 version field, the IPv4 header length, the packet length and the IPv4 time-to-live (TTL) field for packets at ingress of the switch respectively. However, there are also conditions for packets at the egress of the switch. These conditions are described by Query 6 (lines 15-20). Namely, changing source and destination MAC addresses to the correct values, decrementing the TTL value by 1, recalculating the IPv4 checksum and emitting the packet on the correct port as instructed by the control plane configuration (forwarding rules). Query 7 (lines 22-27) corresponds to the platform-dependent part of the switch (PRE) and defines conditions for packets that are cloned by the switch. Such packets need to fulfill the same conditions as per Query 6, but the egress port should correspond to the clone session configuration of the target switch. Similarly, Query 8 (lines 29-34) expresses the conditions for multicast packets that need to fulfill the same conditions as per Query 7 but the egress ports should correspond to the configured multicast group configuration of the target switch.

Note, queries written in p4q can be extended, reused and provided in the form of libraries. More importantly, the p4q queries help in the availability of a library of pre-defined patches for the corresponding violations. Note, an easily extensible interface is provided to augment p4q further with user queries as per the deployment scenario to allow detection of more bugs.

2) Agent : The Agent houses the RL algorithm (Algorithm 1), inspired by Double Deep Q Network (Double DQN) [44] (improved Deep Q Networks (DQN) [45]).

Double Deep Q Network (DDQN). DDQN algorithm [44] is a recently-developed algorithm based on Q-learning [22], hence a model-free reinforcement learning algorithm. Modelfree means, the Agent does not need to learn a model of the dynamics of an environment and how different actions affect it. This is beneficial, as it can be difficult to retrieve accurate models of the environment. At the same time, the goal is to provide sample efficient learning, i.e. reduce the number of packets sent to the target switch, makes the DDQN a suitable choice. The basic concept of the algorithm is to use the current state (packet header) as an input to a neural network, which predicts the action the Agent shall select to maximize future rewards. In addition, Double DQN algorithm splits action selection in a certain state from the evaluation of that action. To achieve, it uses two neural networks: (i) the online network responsible for action selection, and (ii) the target network evaluating the selected action. This improves the learning process of the Agent, as overoptimism of the future reward when selecting a certain action, is reduced and thus, helps to avoid overfitting.

Prioritized Experience Replay. Experience replay [46] was introduced to eliminate problems of oscillation or divergence

IEEE/ACM TRANSACTIONS ON NETWORKING

Algorithm 1 Agent Training	Algorithm 2 P4Tarantula (Localizer)
Input : Empty prioritized experience replay memory M ,	Input: P4 source code (SC) , sent packets (Ps) and
uninitialized online and target network	corresponding verdicts (V)
Output: Trained online and target network models	Output: S[j] - suspiciousness score for the corresponding
1 Initialize online network with random weights	line j
2 Initialize target network with copy of online network	// V[p] represents the <i>verdict</i> about
parameters	packet p (pass or fail)
3 for $i = 1$ to $num_episodes$ do	<pre>// SC[j] represents line j of the</pre>
4 Initialize byte sequence b_1	source code
5 Preprocess b_1 to get the initial state	// Initialization
$s_1 = preprocess(b_1)$	$1 \ totalFailed = 0, totalPassed = 0$
6 for $step = 1$ to max_ep_len do	2 foreach p in Ps do
7 Select action a_{step} randomly with probability ϵ	3 if $V[p] == pass$ then
(exploration) or use online network to predict a_{step}	4 $totalPassed+=1$
(exploitation)	5 else
8 Execute action a_{step} , observe reward r_{step} and	$6 \qquad totalFailed + = 1$
byte sequence b_{step+1}	7 end
9 Set $s_{step+1} = preprocess(b_{step+1})$	8 follow p through SC :
• Save the transition	9 foreach executed line j in SC do
$(s_{step}, a_{step}, r_{step}, s_{step+1}, terminal)$ in M	10 if $V[p] == pass$ then
Sample batch of transitions (s_j, a_j, r_j, s_{j+1}) from	11 $SC[j].pass+=1$
M	12 else
$\mathbf{z} \mid \mathbf{y}_j =$	13 $SC[j].fail+=1$
$\int r_i$ if terminal	14 end
$\int_{r_i}^{r_i} + \gamma * Q(s_{i+1} \max_{\sigma} Q(s_{i+1} a_i; \Theta) \Theta')$	15 $S[j] = \frac{SC[j].fail/totalFailed}{SC[i].rass/totalPassed+SC[i].fail/totalFailed}$
$\left\{\begin{array}{c} y_{j+1} + y_{j+2} \in (j_{j+1}, \max_{l} \in (j_{j+1}, \max_{l} j_{l}, j_{l}, j_{l}, j_{l}), 0 \end{array}\right\}$	16 end
	17 end
Bertorin stochastic gradient descent using	18 call Patcher
categorical cross entropy loss function	
-	

of parameters, resulting from correlated data. To overcome this problem, the experiences of the *Agent*, i.e., a tuple comprising the current state, predicted action, reward received, and resulting state are saved in the memory of *Agent*. To enable learning by experience replay, the neural network model is updated using random samples from past experiences. To counter the scenario of sparse rewards, a simple form of prioritized experience replay, inspired by [47], is applied. The memory is sorted by absolute reward and each experience is prioritized by a configurable factor and the index.

Agent Training Algorithm. Algorithm 1 presents the training algorithm of the Agent in the P6 system. For our algorithm, we rely on the use of Multi-Layer Perceptron (MLP) [48]. In the initialization phase, the weights of online and target neural networks are initialized (Lines 1-2). For each execution, the current state is reinitialized by randomly choosing a packet header in byte representation from a pre-generated set of packet headers. The corresponding bytes are then converted to a sequence of float representations (Lines 4-5). An ϵ greedy policy is applied to determine the action to be executed (Line 7). Applying an ϵ -greedy policy means that during training of the Agent, an action is selected randomly by the Agent with probability ϵ to ensure sufficient exploration. As the training progresses, probability ϵ is decreased linearly until a lower bound is reached. This helps in reducing overfitting as well, since the Agent never stops exploring the

effects of other actions on the environment during training. The determined action will be executed, the result is observed and saved in the experience memory (Lines 8-10). As a last step, a sample out of the experience memory is selected to calculate y_j which is used to calculate the categorical crossentropy loss and perform the stochastic gradient descent step to update the network weights (Lines 11-13).

C. Localizer: P4Tarantula

P4Tarantula is the Localizer or the bug localization module of P6. P4Tarantula is based on a dynamic program analysis technique for generic software, Tarantula [23], [25]. In case a bug is discovered by Fuzzer, it automatically notifies P4Tarantula. *Note, P4Tarantula will not be notified in case of platform-dependent bugs, as they are neither localizable nor patchable in the P4 program.* As an input, P4Tarantula uses the P4 program or source code, the packets that were sent by Fuzzer as per the p4q query (test cases) to trigger the bug and the pass (if p4q query is not violated) or fail verdict (if p4q query is violated) corresponding to those sent packets. Recall, a *verdict* corresponds to a condition of the p4q query which acts as a test case.

Algorithm 2 presents the localization algorithm used by P4Tarantula. First, P4Tarantula initializes two counters, measuring the number of passed or failed *verdicts* corresponding to the sent packets (Line 1). In the next step, P4Tarantula

increments the counters according to the verdicts made for the given packet (Lines 3-7). Now, the P4 source code needs to be traversed line-by-line (similar to symbolic execution but with actual packet header values to avoid all possible header values), to find the code execution path for the given packet (Line 8). For each line of the P4 source code executed for the given packet, counters for the corresponding verdicts are incremented (Lines 10-14). For the executed lines of the P4 source code, a suspiciousness score [25] is calculated (Line 15). The suspiciousness score is between 0 and 1 as the same line/s can be executed for passed and failed verdicts corresponding to packets. This score corresponds to the likelihood of a line of code causing a potential bug. The closer it is to 1, the more likely it is that the corresponding line of code is problematic. Finally, the P4 program lines are ordered as per their suspiciousness score to localize the bug. Then, Patcher is notified.

D. Patcher

Patcher is the novel automated patching module of the P6 system. If a bug is localized by P4Tarantula, it notifies Patcher. The input for Patcher is the P4 source code, the results of static analysis of the P4 source code, the localization results of P4Tarantula, and the violated p4q query. Patcher compares the localized problematic parts of the code with appropriate available patches. Note, Patcher comes with a library of patches for P4 programs, i.e., those which violate p4q queries. Nevertheless, it can be easily extended when, previously unseen bugs, e.g., bugs in application code logic, are detected.

From the results of the static analysis, Patcher extracts the required parser state names, header names, header field names, metadata names and metadata field names for the patches in the current version of the library of patches. In P4, metadata is used to pass information from one of the programmable or non-programmable blocks to another.

Note, in most P4 programs (including the publicly available programs from [28], [29], and [30]) no variables apart from user-defined names for parser states or header/metadata fields are present. Thus, with the gathered knowledge about user-defined names Patcher can compare through, e.g., regex or string comparison, if the patch (correct code) is already present in the P4 source code or if missing, the patch needs to be applied. Note, if the patches in the patch library require the analysis of custom variables or stateful components, e.g., registers and meters, the comparison if the patch is present or not requires further code analysis.

In case no appropriate patch is available, the programmer is informed by the Patcher. After Patcher finishes the execution, it calls the P4 compiler (p4c) to re-compile the patched version of the P4 program and triggers the re-deployment of the code on the P4 switch. In addition, the Fuzzer is notified automatically by Patcher to test the patched program again, to confirm the patches and ensure no regressions were caused by the patches by testing via the p4q queries and executing regression testing.

Algorithm 3 Patcher

- **Input**: P4 source code (SC), static analysis results (Sr), localization results (Lr) and violated p4q query (q)
- **Output**: A patched version of the source-code (*PSC*)
- // The patcher offers a patch only for those lines where the suspiciousness score ≥ 0.5
- 1 Import & process user-defined parser state names, header and header field names, metadata and metadata field names from Sr required for patches in the patch-library
 2 for lines in Lr do
- 3 **if** Suspiciousness score ≥ 0.5 then
- check corresponding line/s of code pinpointed by P4Tarantula
 if the patch is missing and violating q then
 apply the preferred patch
- else
- 8 inform the programmer
- 9 end
- 10 Goto next line

11 else

12 Goto next line

13 end 14 end

7

15 Compile & re-deploy the patched P4 program (PSC)

and notify Fuzzer for testing the patches and regressions

A patch has the following properties: (a) preferably, few lines of code, e.g., missing checks in parser, (b) makes the P4 program conform to the expected behavior, (c) passes the sanity testing or checks for basic functionality, (d) does not cause regressions breaking existing functionality, and (e) should not only fix the error under the assumption that the match action tables will be populated as needed. The patch should be configuration-agnostic or the proprieties of the patch should be equally upheld for all potential types of incoming packets.

Algorithm 3 shows the Patcher algorithm. First, Patcher imports the needed header or metadata field names, as well as parser state names for the currently available patches in the library of patches. Then, for each line in the localization results, Patcher checks if the suspiciousness score is greater than the defined threshold of 0.5 (threshold is configurable as per deployment scenario) (Line 2), as it is highly likely that the corresponding line of code is responsible for triggering the detected bug. In case the suspiciousness score is above the defined threshold, Patcher will check the corresponding line of code. The Patcher, then, checks if the patch is available, e.g., through string comparison with the appropriate patch to be applied for the violated p4q query. If the patch is indeed missing, then the problematic line of code is patched, else the programmer is informed as the appropriate patch is not available (Lines 3-8). Once, all the localization results are processed (Lines 9-12), the patched P4 program is compiled by triggering the compiler (p4c) to be re-deployed on the P4

IEEE/ACM TRANSACTIONS ON NETWORKING

TABLE I

BUGS (WITH BUG IDS) DETECTED BY THE P6 PROTOTYPE THROUGH THE VIOLATION OF THE CORRESPONDING p4q QUERIES (IN FIGURE 9). NOTE, PI AND PD REFER TO PLATFORM-INDEPENDENT AND -DEPENDENT RESPECTIVELY

	D	
Bug IDs	Bugs	Queries (Figure 9)
1	Accepted wrong checksum (PI)	Query 1
2	Generated wrong checksum (PI)	Query 6 (Line 19)
3	Incorrect IP version (PI)	Query 2
4	IP IHL value out of bounds (PI)	Query 3
5	IP TotalLen value is too small (PI)	Query 4
6	TTL 0 or 1 is accepted (PI)	Query 5
7	TTL not decremented (PI)	Query 6 (Line 18)
8	Clone not dropped (PD)	Query 7 (Line 27)
9	Resubmitted packet not dropped (PD)	Query 6 (Line 20)
10	Multicast packet not dropped (PD)	Query 8 (Line 34)

switch and the Fuzzer is triggered to re-test the patched code (Line 14).

IV. P6 PROTOTYPE

We develop a P6 prototype using Python version 3.6 with \approx 3, 100 lines of code (LOC); Fuzzer with \approx 2, 200 LOC, P4Tarantula with \approx 490 LOC and Patcher with \approx 430 LOC. Fuzzer is implemented using Keras [49] library with Tensorflow [50] backend and Scapy [51] for packet generation and monitoring. Currently, P6 only supports programs written in P4₁₆ [2] as the P4 compiler (p4c [42]) supports the translation of programs written in P4₁₄ [41] to P4₁₆. The *Agent* was trained separately, for each condition of each query written in p4q. The training process as well as the later execution using the trained Agents, however, can be parallelized. For queries described in Figure 9, the trained model of the *Agent* can be reused for testing different P4 programs that implement IPv4 packet processing.

In addition to the modules described in §III, we implement a control plane module using P4Runtime [52] and Python. For the P4 switch, we rely on software switches supporting P4₁₆, namely behavioral model (bmv2) [31] with SimpleSwitchGrpc target (Version 1.12.0), and Barefoot Tofino Model [32] (Version 8.3.0). To simplify and automate the usage of P6, a default option is provided where all queries of p4q are executed by: P6 'p4/source_code_location' -default

V. EVALUATION

In this section, we evaluate the P6 verification capabilities.

A. Baselines

We compare P6 against three baseline fuzzing approaches:

(1) Advanced Agent. The first baseline is an Advanced Agent only relying on random fuzz action selection, i.e., without prioritized experience replay. Thus, Advanced Agent can execute the same mutation actions as P6, but cannot learn which actions lead to rewards. It represents the intelligent baseline.

(2) **IPv4-based fuzzer.** The second baseline is an IPv4based fuzzer, which is aware of the IPv4 header layout and randomizes the different available header fields, except IP options fields and the destination IP as it prevents the packets from being dropped by the forwarding rules of the P4 switch. The actual behavior is evaluated using the queries of p4q.

(3) Naïve fuzzer. The third baseline is a simple naïve fuzzer, which is not aware of any packet header layouts. It generates and sends Ethernet frames from purely random mutation of bytes. The actual behavior is evaluated using the p4q queries.

B. Bugs

Table I provides an overview of *existing* bug types (with bug IDs) detected in the publicly available P4 programs from [28], [29], and [30] by the P6 prototype. These bugs are detected as they violate the corresponding p4q queries (from Figure 9). In total, P6 prototype can detect 10 distinct bugs in the P4 programs. Out of these 10 bugs, 7 are patchable platform-independent (bugs 1 - 7), and 3 are platform-dependent bugs (bug 8 - 10).

Platform-independent bugs. The two detected bugs with bug ID 1 and 2, are related to wrong IPv4 checksum computation and missing checksum validation. P6 is able to detect, localize and patch these bugs. The four bugs with IDs 3-6 are missing or wrong IPv4 packet header validation. Specifically, missing validation of IP version (bug 3), IPv4 header length (bug 4), IPv4 total length (bug 5) and IPv4 time-to-live (TTL) (bug 6). P6 can detect, localize and patch these bugs. While the current approaches *may* be able to detect the bugs, they still lack localization and patching of the bugs. The last of the platform-independent bugs is faulty TTL decrement bug (bug 7). In this case, the P4 program accepts packets with IPv4 TTL 0, the TTL decrement is still executed, causing an incorrect increment of TTL to 255. Note, these 7 platformindependent bugs already exist in the publicly available P4 programs of [29] and [30]. In switch.p4 program [28], bugs with IDs 1,2,4 and 5 exist.

Platform-dependent bugs. In addition to the aforementioned platform-independent bugs, P6 is able to detect three platformdependent bugs. The first bug (Bug ID 8) is described in Figure 2 and occurs when using ingress-to-egress clone action. It violates Query 7 in Figure 9 leading to incorrect forwarding of cloned packets when they are supposed to be dropped. The second bug (Bug ID 9) involves the resubmit operation. Packets with the resubmit metadata field set which are marked to be dropped in a later stage, will be incorrectly resubmitted, i.e., the packet will be processed again, starting at the ingress parser. Thus, packets that are not expected to be resubmitted will be processed again. In the worst case, this can lead to packets being resubmitted over and over again or other unexpected behavior. It violates Query 6 in Figure 9. The third bug (Bug ID 10) involves the multicast operation. If a packet is marked to be dropped and later the mcast_grp metadata field is set, then the multicasted copies of the packet are incorrectly forwarded and do not get dropped. It violates Query 8 in Figure 9. Note, we found all platform-dependent bugs specifically, in basic.p4 program [29] on bmv2 platform [31]. Furthermore, these bugs *cannot* be detected by current approaches that are based on static analysis.

SHUKLA et al.: RUNTIME VERIFICATION FOR PROGRAMMABLE SWITCHES



Fig. 11. Bug detection, localization and patching times of different P4 programs in bmv2 and Tofino. Each plot represents a median over 10 runs.



Fig. 12. P6 vs Baselines. Each plot represents a median over 10 runs.

TABLE II P4_{16} Programs' LOC in BMV2 and Tofino

P4 ₁₆ Applications	bmv2 (LOC)	Tofino (LOC)	% increment
basic.p4 [29]	181	257	41.9
basic_tunnel.p4 [29]	219	302	37.9
advanced_tunnel.p4 [29]	242	316	30.6
mri.p4 [29]	277	372	34.3
netpaxos-acceptor.p4 [30]	270	327	21.1
netpaxos-leader.p4 [30]	259	323	24.7
netpaxos-learner.p4 [30]	288	355	23.2
switch.p4 [28]	8715	N/A for Tofino	-



Fig. 13. Experiment topology.

C. Experiment Strategy

For conducting our experiments and to evaluate P6 prototype, we ran P6 together with the P4 switch and

control plane module in a Vagrant [53] environment with VirtualBox [54]. We emulate the network shown in Figure 13. For each program, separate Vagrant machines, each with 10 CPU cores and 7.5 GiB RAM, are used. The Vagrant machines ran on a server running Debian 9 OS (Version 4.9.110), with Intel Xeon CPU and 256 GiB RAM. Each experiment was executed ten times on each of the eight P4₁₆ programs shown in Table II from P4 tutorials [29], NetPaxos [30] and switch.p4 [28] repository. For each of the ten runs, 9 test-cases were executed, where each test-case corresponds to one condition of the queries 1-6 illustrated in Figure 9. Note, basic.p4 program has 10 test-cases as it is also tested using query 7 for the platform-dependent bug. Furthermore, we currently trained the Agent separately for each test-case and sequentially execute the test cases. This can, however, be parallelized easily. We observed that for only two conditions of the p4q queries, no bugs were detected.

1) Experiment Topology: Figure 13 shows the topology used for the experiments as part of the P6 system evaluation. In total, five virtual machines are used. One of the machines is running the P4 switch, which is connected to all other machines. The controller is connected to the P4 switch, in order to deploy the P4 program and fill the forwarding

12

tables. The P6 system is connected to one port of the switch for sending the packets generated by Fuzzer. Two virtual machines act as the receiver for these packets. The feedback of the receivers is used by the P6 system to evaluate the actual behavior of the P4 switch. In addition, the machine running P6 is connected to the controller to trigger re-deployment of the patched P4 program, in case a bug is detected.

D. Metrics

In particular, we ask the following questions:

Q1. How much time does P6 *take to detect, localize, and patch all bugs?* (§V-D.1)

Q2. How does P6 perform against the baselines? (§V-D.2)

Q3. How many rewards does P6 generate against the baseline of an Advanced Agent for Agent training? (§V-D.3)

Q4. How many packets does P6 generate to detect bugs against the baselines? (\S V-D.4)

Q5. What is the accuracy of P6 ? (§V-D.5)

1) Performance (Speed) of P6 : To evaluate the performance of P6, we execute the detection, localization, and patching on 8 publicly available P4 programs from the P4 tutorials [29], NetPaxos [30] and switch.p4 [28] repository with minimal manual efforts.

Figure 11a and 11d show the median bug detection time of P6 over ten runs for the different programs using bmv2 SimpleSwitchGrpc and Barefoot Tofino Model, respectively. Note, switch.p4 program is only available for bmv2 and was not tested using Tofino. In all runs on bmv2 except for switch.p4 program, P6 was able to detect all bugs in less than two seconds. In switch.p4, P6 was able to detect all bugs in less than ten seconds. The detection time is higher for switch.p4 as compared to the other tested programs since more packets get dropped making bug detection more difficult. On Tofino, the median detection time was slightly higher for four out of seven programs. The reason for the increased bug detection time with the NetPaxos programs [30] can be due to the instrumentation of these programs by us to make them run on CPU intensive Tofino Model.

Figures 11b and 11e illustrate the median bug localization time of P6 for the different programs using bmv2 SimpleSwitchGrpc and Barefoot Tofino Model. Overall, all bugs for 7 of the programs were localized by P6 in just above 0.12 seconds on bmv2 and Tofino. To our surprise, the bug localization time for switch.p4 program running on bmv2 is only increased by a factor of $4\times$, even though the program has about $30\times$ more lines of code compared to the other tested programs (see Table II). The median time of patching the code is shown in Figures 11c and 11f for bmv2 and Tofino respectively. P6 is able to patch the P4 programs with millisecond scale performance (max. 98 milliseconds).

2) P6 Vs Baselines: Detection Time: We compare P6 against the three baseline approaches in terms of bug detection time. We observe that the Advanced Agent baseline, see Figure 12a (with quartiles), was able to detect all the bugs present in the tested programs, which is due to the similarity with the P6 Agent. Advanced Agent, however, cannot learn from the rewards, hence generates more packets and thus,

IEEE/ACM TRANSACTIONS ON NETWORKING

TABLE III P6 Vs Baselines. Median #packets Sent per Run Over 10 Runs

P416 Applications	P6	Advanced Agent	IPv4-based	Naïve
basic.p4 [29]	13	59	8,035	16,000
basic_tunnel.p4 [29]	11	59	6,044	14,000
advanced_tunnel.p4 [29]	12	57	6,038	14,000
mri.p4 [29]	10	61	6,058	14,000
netpaxos-acceptor.p4 [30]	11	52	6,021	14,000
netpaxos-leader.p4 [30]	9	49	6,024	14,000
netpaxos-learner.p4 [30]	12	44	6,026	14,000
switch.p4 [28]	28	113	2,132	14,000

takes more time to detect the bugs than P6 Agent. IPv4based fuzzer was only able to detect 4 out of 10 bugs in the seven programs from [29] and [30]. For switch.p4 program [28], IPv4-based fuzzer was able to detect 3 out of 4 bugs which were IPv4-based. In Figure 12b (with quartiles), the speedup is defined as infinite for the test-cases where IPv4based fuzzer could not detect the bug. Accordingly, the bars representing these test-cases range until the top of the figure. Note, Naïve fuzzer was not able to detect bugs at all, even though generating 16k packets.

Figure 12a shows the speedup (Advanced Agent/P6 Agent) for all bugs detected in the seven tested programs from [29] and [30]. The results show that P6 Agent can detect bugs up to $10.96 \times$ faster than the Advanced Agent baseline. Only bug 7 was detected faster by the Advanced Agent in 3 of the 7 P4₁₆ applications tested as the Advanced Agent needs less time for random action selection than P6 Agent for intelligent action selection, based on its neural networks. In addition, Advanced Agent can make use of the same mutation actions and the pre-generated dict, hence when triggering the bug, the overall execution time will be slightly lower than that of P6 Agent. In 94% of the test-cases, Advanced Agent required more time and packets to detect the bugs than the P6 Agent. For switch.p4 program [28], the results show that P6 Agent is able to detect bugs up to $30 \times$ faster than the Advanced Agent baseline.

Figure 12b shows the speedup (IPv4-based fuzzer/P6 *Agent*) for all bugs detected in the seven tested programs from [29] and [30]. For the test-cases where IPv4-based fuzzer was able to detect the bug, we observe that in 89% of the test-cases P6 *Agent* is able to detect the bugs faster while sending significantly fewer packets. P6 *Agent* outperforms IPv4-based fuzzer by up to $8.88 \times$ even though IPv4-based fuzzer sends packets at a higher rate. For switch.p4 program, the results show that P6 *Agent* is able to detect bugs up to $30 \times$ faster than IPv4-based fuzzer.

3) P6 Vs Advanced Agent Training: To verify that P6 Agent is able to effectively learn to detect bugs, we compare P6 Agent against an Advanced Agent, that only relies on random action selection. This makes Advanced Agent similar, but not as intelligent as P6 Agent. Advanced Agent can still execute the same mutation actions but is not able to reason about which actions lead to maximized rewards. Figure 12c shows a comparison of the mean cumulative reward (MCR) of the training process of both agents for bug ID 4 of Table I. We observe that the P6 Agent is able to outperform the baseline by a factor of $3.56 \times$ for the mentioned case. Especially, the prioritized experience replay helps the P6



Fig. 14. Training: P6 vs Advanced Agent (random action selection) in the case of bug ID 2 in Table I.



Fig. 15. Training: P6 vs Advanced Agent (random action selection) in the case of bug ID 5 in Table I.



Fig. 16. Training: P6 vs Advanced Agent (random action selection) in the case of bug ID 6 in Table I.



Fig. 17. Training: P6 vs Advanced Agent (random action selection) in the case of bug ID 7 in Table I.

Agent to quickly learn about which actions lead to reward, hence trigger bugs in the program. Since the P6 Agent is

trained only using experiences which are *valuable* for the training.

P6 Agent is trained for each condition of each query described by Figure 9 using the same set of hyper-parameters. **More results.** Figure 14 shows the training comparison results for the bug ID 2 (Generated wrong checksum) in Table I. Also, in this case, the P6 Agent is able to outperform the Advanced Agent baseline by a factor of $2.62 \times$.

The training comparison results for the bug ID 5 (IP TotalLen Value is too small) can be seen in Figure 15, showing that P6 Agent is able to outperform Advanced Agent baseline by a factor of $2.06 \times$.

In the case of the bug ID 6 (TTL 0 or 1 is accepted), P6 *Agent* is able to outperform Advanced Agent baseline by a factor of $2.11 \times$, as illustrated by Figure 16.

Figure 17 shows that P6 Agent outperforms Advanced Agent for the bug ID 7 (TTL not decremented) by a factor of $2.03\times$. These results show the clear advantage of P6 Agent over the Advanced Agent baseline.

Above results show P6 Agent consistently outperforms Advanced Agent in MCR for other queries. Overall, this shows the clear advantage of P6 Agent over the baselines and the ability to detect bugs with fewer packets.

4) P6 Vs Baselines: Dataplane Overhead: Table III illustrates the number of packets sent by P6 and the baselines. This shows the usefulness of the P6 Agent which generates less packets by learning about rewards, and generates packets that trigger bugs. In this case, the Advanced Agent is almost similar. IPv4-based fuzzer can detect 4 out of 10 bugs, but generates around 6k packets per run. For each test-case, naïve fuzzer sends around 2k packets (in total between 12k and 16k) but it was not able to trigger any bug.

5) P6 Accuracy: During the experiments, we could not observe any false positives in bug detection. Note, false positives can occur if the p4q queries are not correct and complete. In the case of bug localization by P4Tarantula, false positives are not observed. We did not observe any false positives in patching by the Patcher as it only fixes the code if the correct code is missing. Note, if localization results by P4Tarantula contain false positives, then Patcher is prone to false positives as it makes the patching decisions based on the localization results.

Summary. Our results show that P6 due to RL-guided fuzzing significantly outperforms the baselines across two different platforms: bmv2 and Tofino in terms of detecting runtime bugs (including platform-dependent bugs) with minimal dataplane overhead non-intrusively. We observe that most of the platform-independent bugs existed in the parser or header part, otherwise packets with invalid headers get rejected. P6 accurately and swiftly localizes and patches (millisecond scale) the bugs due to the P4 program structure in an automated fashion.

VI. RELATED WORK

Verification of programmable networks has been in a constant state of flux. Approaches like [57], [58], and [59] perform modeling of the network from the control plane to check the reachability, loop freedom, and slice isolation.

TABLE IV

Related Work in P4 Verification. PD Corresponds to the platform-dependent Bugs. Note, \checkmark Denotes the Capability, (\checkmark) Denotes a Part of Full Capability, and \times Denotes the Missing Capability

Related work in P4	Runtime Verification	Detection	Localization	Patching	Detection of PD bugs
Cocoon [55]	X	 ✓ 	√	×	X
Vera [4]	X	√	×	×	×
p4v [5]	X	√	×	×	X
ASSERT-P4 [7], [6]	X	√	×	×	X
P4NOD [56]	X	√	×	×	X
p4pktgen [8]	X	 ✓ 	X	×	X
P4CONSIST [9]	√	~	X	×	X
P4rl [20]	√	√	×	×	X
P 6	√	√	 ✓ 	√	√

ATPG [60] generates test packets based on control plane configuration using [57] for functional and performance verification in traditional networks and SDNs (Softwaredefined Networks). All of the aforementioned tools [57], [58], [59], [60], however, assume that the control plane has a consistent or correct view of the data plane in traditional IP-based networks or SDNs only. P6 does not assume the correctness of the control plane and observes the runtime behavior to detect, localize and patch the software bugs in P4 switches. References [61], [62], and [63] use different machine learning approaches for finding security vulnerabilities or compiler specific-bugs which cause crashes, however, they are insufficient for network-related verification. P6 executes switch verification to identify the bugs in a P4 switch.

Currently, most of the P4-related verification techniques, use static analysis of P4 programs using symbolic execution [4], [6], [7] or Hoare logic [5]. The static analysis is prone to false positives as it analyzes the P4 program without passing any real inputs, e.g., packets. Therefore, checksumrelated bugs where computations are required on input packets and platform-dependent bugs cannot be detected. Such bugs require P6-like runtime verification. In addition, [5], [6], [7] require a P4 program to be manually annotated by the programmer which is cumbersome and prone to manual errors whereas P6 is non-intrusive as it does not require to modify P4 program for bug detection and localization. p4pktgen [8] focuses on locating errors in the toolchains used to compile and run P4 code, e.g., p4c, and uses symbolic execution to create exemplary packets which can execute a selected path in the program. However, it cannot detect platformdependent bugs or egress pipeline bugs. Such a verification method can complement our solution. P4NOD [56] statically models the network, however, it does not check how the actual P4 switches behave upon receiving the malformed packets e.g., incorrect IPv4 checksum. Cocoon [55] suggests refinement-based programming for network verification. While this approach tries to ensure that programs match their specification, it requires a huge amount of additional and manual user input. For runtime verification, such a formal method is insufficient. Recently, [64] propose data-plane primitive for detecting and localizing bugs: tracking each packet's execution path through the P4 program by augmenting P4 programs. However, this remains an in-progress and intrusive approach as it requires augmenting P4 code whereas P6 does not change P4 program. Furthermore, recently [65] was developed which uses a mix of static verification, code

changes and runtime checks to ensure that the deployed P4 program is bug free. It complements our approach as it also focuses on control plane related bugs. Also, [66] detects packet-processing errors induced by bugs that are not caught (or are difficult to catch) before the P4 program deployment.

In-band network telemetry (INT) [67], [68] enables to collect telemetry data from each switch, however, unlike P6, it cannot localize or patch bugs if packets get dropped. Recently, Shukla et al. proposed P4CONSIST [9], a system that gathers the control and data plane states independently for comparison to verify the control-data plane consistency of P4 SDNs by detecting the path violations for critical flows in P4, however, without localization or patch support. In [20], a machine learning guided-fuzzing system is used to only detect platform-independent bugs in P4 programs. In the context of fuzzing, two approaches like [10] and [17] are worth-considering. [10] is insufficient as it uses program coverage feedback to guide fuzzing without knowing which mutations lead to bugs. Reference [17] transforms the target program to remove sanity checks for fuzzing, however, it is intrusive as the target program requires modification for testing. P4wn [69] is a program profiler that can analyze program behaviors of stateful data plane systems. We note that such an approach can be complementary to P6. Furthermore, P4wn differs from P6 in its use case as P4wn is more securitycentric relying on adverserial testing of stateful P4 programs for vulnerability detection whereas P6 performs an end-to-end network verification of stateless P4 programs which not just focuses on bug detection but also on localization and patching. Aquila [70] focuses on verifying P4 programs for bugs and also, localizes bugs but unlike P6, it is not in its focus to patch the bugs or even report platform-dependent bugs.

Unlike P6, P4-based verification approaches [4], [5], [6], [7], [8], [9], [20], [55], [56], [67], [68] are insufficient in localizing and patching the runtime bugs in P4 programs. Besides, they cannot detect the platform-dependent bugs. Table IV illustrates the capabilities of other P4 verification tools as compared to P6.

VII. DISCUSSION

Traditionally, fuzz testing or fuzzing is known to offer a partial testing solution as it is prone to false negatives. Rice's theorem [71] states that all the non-trivial, semantic properties of a program are undecidable. Semantic properties refer to the behavior of a target program for all inputs. Therefore, if fuzzing does not detect any problems, it does not ensure that there is no problem at all. Statistical techniques like Good-Turing frequency estimations [72], [73], [74] for fuzz testing partially, aid in inferring the probability that the next generated test input leads to the discovery of a previously unseen species.

P6 can only verify functional or binary properties such as whether actual behavior is the same as expected behavior or not. However, we cannot verify properties such as performance-related ones. In P6, static analysis acts as a guiding process for P6 Fuzzer which helps to significantly reduce the input search space and adds input structure awareness. In general, the quality of the input seeds, e.g., the relevance of mutations to the target and program coverage serve as good indicators to assess the quality of a Fuzzer. However, there is a tradeoff between speed and precision of fuzz testing as there is an instrumentation overhead involved in generating a dictionary of meaningful inputs for effective testing and significantly reducing the input search space as compared to random mutations of bits in the inputs.

Machine learning techniques like reinforcement learning [21], [22] helps to some extent with the training of the models based on the feedback from the target. However, *what is a good feedback?* is debatable. Traditionally, feedback depends on the coverage but it all boils down to the program under test. In addition, machine learning is as good as the input training data and thus, offers insufficient guarantees as instead of learning, the model may memorize. A generalized model applicable to *any* kind of training data is highly desirable as it avoids the problems of overfitting and underfitting [75].

Dynamic program analysis technique for fault-localization like Tarantula [23], [24], [25] benefits from utilizing the information from multiple failed test cases which helps it in leveraging the richer information base. Therefore, it helps to have at least one failed test case. In addition, allowing tolerance for passed test cases that occasionally execute faults is essential for an effective fault-localization technique.

Software patching facilitates in fixing the software code errors. The patches, however, *may* cause regressions which reflect in the form of abnormal behavior of the software. Sanity testing and modular code design facilitate in ensuring that basic functionality is not affected by the patch. However, one cannot assure that there are no other problems (false negatives) caused by the fix.

We note that leveraging programmability, future programmable networks will encompass even more possibilities of faults with a mix of vendor-code, reusable libraries, and in-house code. As such the general problem of network verification will persist and we will have to explore how to extend the P6 system to traditional IP-based networks and verify important classes of properties such as performancerelated ones (e.g., mean and tail flow completion times). To address state loss and flow interruption due to the deployment delay of the patched P4 program, the existing flows can be migrated to the relevant servers for processing via approaches like [76].

VIII. CONCLUSION

We presented a new approach to verify programmable switches at runtime. We also present the design, implementation, and evaluation of P6, the first system that enables runtime verification of P4 switches in a non-intrusive fashion. P6 uses static analysis- and machine learning-guided fuzzing to detect multiple runtime bugs which are then, localized and patched on the fly with minimal human effort. Through experiments on existing P4 application programs, we showed that P6 significantly outperforms the baseline bug detection approaches to detect existing platform-independent and -dependent bugs. In the case of platform-independent bugs, P6 takes advantage of the increased programmable blocks to localize them and repair the $P4_{16}$ programs, if and when a patch is available.

We believe P6 is an important foray into self-driving networks [77], which come with stringent requirements on dependability and automation. With P6, developers of P4 programs and operators of P4-enabled devices can improve the security of their products. To make the results of this work available to the public, we release the P6 software and library of ready patches that we, respectively, developed and used in this work. As a part of our future agenda, we plan to apply P6 on commercial-grade P4 programs and networks to report our experience.

REFERENCES

- P. Bosshart et al., "P4: Programming protocol-independent packet processors," ACM SIGCOMM Comput. Commun. Rev., vol. 44, no. 3, pp. 87–95, Jul. 2014.
- P. L. Consortium. (2018). P4₁₆ Language Specs, Version 1.1.0. [Online]. Available: https://p4.org/specs/
- [3] P. Kazemian. (2017). Network Path Not Found?. [Online]. Available: https://bit.ly/2FzpEEZ
- [4] R. Stoenescu, D. Dumitrescu, M. Popovici, L. Negreanu, and C. Raiciu, "Debugging P4 programs with vera," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2018, pp. 518–532.
- [5] J. Liu et al., "P4V: Practical verification for programmable data planes," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2018, pp. 490–503.
- [6] M. Neves, L. Freire, A. Schaeffer-Filho, and M. Barcellos, "Verification of P4 programs in feasible time using assertions," in *Proc. 14th Int. Conf. Emerg. Netw. Exp. Technol.*, Dec. 2018, pp. 73–85.
- [7] L. Freire, M. Neves, L. Leal, K. Levchenko, A. Schaeffer-Filho, and M. Barcellos, "Uncovering bugs in P4 programs with assertion-based verification," in *Proc. Symp. SDN Res.*, Mar. 2018, pp. 1–7.
- [8] A. Nötzli, J. Khan, A. Fingerhut, C. Barrett, and P. Athanas, "P4pktgen: Automated test case generation for P4 programs," in *Proc. Symp. SDN Res.*, Mar. 2018, pp. 1–7.
- [9] A. Shukla, S. Fathalli, T. Zinner, A. Hecker, and S. Schmid, "P4Consist: Toward consistent P4 SDNs," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 7, pp. 1293–1307, Jul. 2020.
- [10] M. Zalewski. (2010). American Fuzzy Lop: A Security-Oriented Fuzzer. Accessed: Apr. 6, 2019. [Online]. Available: http://lcamtuf.coredump.cx/afl/
- [11] Llvm Compiler Infrastructure: Libfuzzer: A Library for Coverage-Guided Fuzz Testing. [Online]. Available: http://llvm.org/docs/LibFuzzer.html
- [12] P. Godefroid, M. Y. Levin, and D. Molnar, "SAGE: Whitebox fuzzing for security testing," *Commun. ACM*, vol. 55, no. 3, pp. 40–44, 2012.
- [13] Peach Fuzzer. [Online]. Available: https://www.peach.tech/
- [14] OpenRCE: Sulley. [Online]. Available: https://github.com/OpenRCE/ sulley
- [15] Radamsa. [Online]. Available: https://gitlab.com/akihe/radamsa
- [16] ZZUF-MULTI-PURPOSE FUZZER. [Online]. Available: http://caca.zoy.org/wiki/zzuf
- [17] H. Peng, Y. Shoshitaishvili, and M. Payer, "T-fuzz: Fuzzing by program transformation," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2018, pp. 697–710.
- [18] N. Stephens et al., "Driller: Augmenting fuzzing through selective symbolic execution," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2016, pp. 1–16.
- [19] A. Shukla, S. J. Saidi, S. Schmid, M. Canini, T. Zinner, and A. Feldmann, "Toward consistent SDNs: A case for network state fuzzing," *IEEE Trans. Netw. Service Manag.*, vol. 17, no. 2, pp. 668–681, Jun. 2020.
- [20] A. Shukla, K. N. Hudemann, A. Hecker, and S. Schmid, "Runtime verification of P4 switches with reinforcement learning," in *Proc. Workshop Netw. Meets AI ML (NetAI)*, 2019, pp. 1–7.
- [21] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [22] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 3rd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2009.
- [23] J. A. Jones, M. J. Harrold, and J. T. Stasko, "Visualization for fault localization," in *Proc. ACM/IEEE ICSE Workshops*, Jan. 2001.

- 16
- [24] J. A. Jones, M. J. Harrold, and J. Stasko, "Visualization of test information to assist fault localization," in *Proc. 24th Int. Conf. Softw. Eng. (ICSE)*, 2002, pp. 467–477.
- [25] J. A. Jones and M. J. Harrold, "Empirical evaluation of the tarantula automatic fault-localization technique," in *Proc. 20th IEEE/ACM Int. Conf. Automated Softw. Eng. (ASE)*, Nov. 2005, pp. 273–282.
- [26] C. L. Goues, M. Pradel, and A. Roychoudhury, "Automated program repair," *Commun. ACM*, vol. 62, no. 12, pp. 56–65, 2019.
- [27] Cisco Systems. (2019). daPIPE: DAta Plane Incremental Programming Environment. [Online]. Available: https://p4.org/assets/ P4WS_2019/p4workshop19-final16.pdf
- [28] Switch. P4. [Online]. Available: https://github.com/p4lang/switch
- [29] P4 Tutorial. [Online]. Available: https://github.com/p4lang/tutorials
- [30] NetPaxos. [Online]. Available: https://github.com/usi-systems/p4xospublic
- [31] P4.org. (Oct. 2015). Behavioral Model Repository. Accessed: Dec. 2018. [Online]. Available: https://github.com/p4lang/behavioral-model
- [32] Tofino. [Online]. Available: https://www.barefootnetworks.com/products/ brief-tofino
- [33] A. Shukla et al., "Fix with P6: Verifying programmable switches at runtime," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2021, pp. 1–10.
- [34] A. Shukla, "Towards runtime verification of programmable networks," Doctoral thesis, Technische Univ. Berlin, Berlin, Germany, 2020, doi: 10.14279/DEPOSITONCE-10594.
- [35] A. Shukla. (2022). P6 Software. [Online]. Available: https://github.com/Stevey1234/P6
- [36] P. Bosshart et al., "Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN," ACM SIGCOMM Comput. Commun. Rev., vol. 43, no. 4, pp. 99–110, 2013.
- [37] P416 Portable Switch Architecture (PSA)-Version 1.1: Programmable Blocks. [Online]. Available: https://p4.org/p4-spec/docs/PSAv1.1.0.html#sec-programmable-blocks
- [38] P416 Portable Switch Architecture (PSA)-Version 1.1: Ingress Deparser and Egress parser. [Online]. Available: https://p4.org/p4-spec/docs/PSAv1.1.0.html#appendix-rationale-ingress-deparser-egress-parser
- [39] N. Atre, H. Sadok, E. Chiang, W. Wang, and J. Sherry, "SurgeProtector: Mitigating temporal algorithmic complexity attacks using adversarial scheduling," in *Proc. ACM SIGCOMM Conf.*, Aug. 2022, pp. 723–738.
- [40] S. A. Crosby and D. S. Wallach, "Denial of service via algorithmic complexity attacks," in *Proc. 12th USENIX Secur. Symp.*, 2003.
- [41] The P4 Language Specifications, Version 1.0.5, P4 Language Consortium.
- [42] P. L. Community. (2019). P4c. [Online]. Available: https://github.com/p4lang/p4c
- [43] S. Moon et al., "Alembic: Automated model inference for stateful network functions," in *Proc. NSDI*, 2019, pp. 699–718.
- [44] H. van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI*, 2016, pp. 1–7.
- [45] V. Mnih et al., "Playing Atari with deep reinforcement learning," 2013, arXiv:1312.5602.
- [46] L.-J. Lin, "Reinforcement learning for robots using neural networks," CMU School of Computer Science, Tech. Rep., 1993.
- [47] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," 2015, arXiv:1511.05952.
- [48] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Parallel Distributed Processing, Explorations in the Microstructure of Cognition*, vol. 1, D. E. Rumelhart and J. L. McClelland, Eds., 1986.
- [49] Keras: The Python Deep Learning Library. [Online]. Available: https://keras.io/
- [50] TensorFlow. [Online]. Available: https://www.tensorflow.org/
- [51] Scapy. [Online]. Available: https://scapy.net/
- [52] P4Runtime. [Online]. Available: https://p4.org/p4-runtime/
- [53] Vagrant. [Online]. Available: https://www.vagrantup.com/

- [54] VirtualBox. [Online]. Available: https://www.virtualbox.org/
- [55] L. Ryzhyk et al., "Correct by construction networks using stepwise refinement," in *Proc. NSDI*, 2017, pp. 683–698.
- [56] N. Lopes, N. Bjørner, N. McKeown, A. Rybalchenko, D. Talayco, and G. Varghese, "Automatically verifying reachability and well-formedness in P4 Networks," Microsoft Technical, Tech. Rep., 2016.
- [57] P. Kazemian, G. Varghese, and N. McKeown, "Header space analysis: Static checking for networks," in *Proc. NSDI*, 2012, pp. 113–126.
- [58] P. Kazemian, M. Chang, H. Zeng, G. Varghese, N. Mckeown, and S. Whyte, "Real time network policy checking using header space analysis," in *Proc. NSDI*, 2013, pp. 99–111.
- [59] A. Khurshid, X. Zou, W. Zhou, M. Caesar, and P. B. Godfrey, "VeriFlow: Verifying network-wide invariants in real time," in *Proc. NSDI*, 2013, pp. 49–54.
- [60] H. Zeng, P. Kazemian, G. Varghese, and N. McKeown, "Automatic test packet generation," in *Proc. 8th Int. Conf. Emerg. Netw. Exp. Technol.*, 2012, pp. 241–252.
- [61] M. Rajpal, W. Blum, and R. Singh, "Not all bytes are equal: Neural byte sieve for fuzzing," 2017, arXiv:1711.04596.
- [62] P. Godefroid, H. Peleg, and R. Singh, "Learn&fuzz: Machine learning for input fuzzing," in *Proc. 32nd IEEE/ACM Int. Conf. Automated Softw. Eng. (ASE)*, 2017, pp. 50–59.
- [63] C. Cummins, P. Petoumenos, A. Murray, and H. Leather, "Compiler fuzzing through deep learning," in *Proc. 27th ACM SIGSOFT Int. Symp. Softw. Test. Anal.*, Jul. 2018, pp. 95–105.
- [64] S. Kodeswaran, M. T. Arashloo, P. Tammana, and J. Rexford, "Tracking P4 program execution in the data plane," in *Proc. Symp. SDN Res.*, Mar. 2020, pp. 117–122.
- [65] D. Dumitrescu, R. Stoenescu, L. Negreanu, and C. Raiciu, "BF4: Towards bug-free P4 programs," in *Proc. ACM SIGCOMM*, 2020, pp. 571–585.
- [66] K. S. Kumar et al., "DBVal: Validating P4 data plane runtime behavior," in *Proc. ACM SIGCOMM Symp. SDN Res. (SOSR)*, Oct. 2021, pp. 122–134.
- [67] C. Kim, A. Sivaraman, N. Katta, A. Bas, A. Dixit, and L. J. Wobker, "Inband network telemetry via programmable dataplanes," in *Proc. ACM SIGCOMM*, 2015, pp. 1–2.
- [68] Int Specification. [Online]. Available: https://github.com/p4lang/p4applications/blob/master/docs
- [69] Q. Kang, J. Xing, Y. Qiu, and A. Chen, "Probabilistic profiling of stateful data planes for adversarial testing," in *Proc. 26th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, Apr. 2021, pp. 286–301.
- [70] B. Tian et al., "Aquila: A practically usable verification system for production-scale programmable data planes," in *Proc. ACM SIGCOMM Conf.*, 2021, pp. 17–32.
- [71] H. G. Rice, "Classes of recursively enumerable sets and their decision problems," *Trans. Amer. Math. Soc.*, vol. 74, no. 2, pp. 358–366, Jan. 1953.
- [72] I. J. Good, "The population frequencies of species and the estimation of population parameters," *Biometrika*, vol. 40, nos. 3–4, pp. 237–264, 1953.
- [73] W. A. Gale and G. Sampson, "Good-turing frequency estimation without tears," J. Quant. Linguistics, vol. 2, no. 3, pp. 217–237, 1995.
- [74] M. Bohme, "Assurance in software testing: A roadmap," in Proc. Int. Conf. Softw. Eng., 2019, pp. 5–8.
- [75] K. P. Burnham, D. R. Anderson, and K. P. Huyvaert, "AIC model selection and multimodel inference in behavioral ecology: Some background, observations, and comparisons," *Behav. Ecology Sociobiol.*, vol. 65, no. 1, pp. 23–35, Jan. 2011.
- [76] M. Zhang et al., "Poseidon: Mitigating volumetric DDoS attacks with programmable switches," in *Proc. 27th Netw. Distrib. Syst. Secur. Symp.* (NDSS), 2020.
- [77] N. Feamster and J. Rexford, "Why (and how) networks should run themselves," 2017, arXiv:1710.11583.