



**Universiteit
Leiden**
The Netherlands

Bachelor Computer Science

Practical Equivalence Checking of P4 Packet Parsers

Jort van Leenen

J.P.van.Leenen@umail.leidenuniv.nl

Supervisors:

Dr. T.W.J. Kappé and Dr. J.J.M. Martens

BACHELOR THESIS

Leiden Institute of Advanced Computer Science (LIACS)

<https://liacs.leidenuniv.nl>

08/07/2025

Abstract

In computer networks, packets function as structured containers for transmitting information. Packet parsers validate these packets in devices like packet switches and firewalls. P4 is a domain-specific programming language that can specify such parsers in software-defined networks. As with any software system, bugs may be present, or improvements are to be made. These subsequent modifications may introduce discrepancies relative to the intended specification. To address this, we introduce OCTOPUS, a tool for the practical equivalence checking of packet parsers written in P4. Here, practical signifies a runtime suitable for execution on a typical, modern personal computer. The tool additionally produces a certificate that enables manual verification of the result's correctness.

Contents

1. Introduction	1
1.1. Thesis Overview	2
2. Background and Related Work	2
3. Research Question	4
4. Methodology	5
5. Overview	6
6. Frontend	6
7. Naive Bisimulation	8
7.1. Bisimulation	9
8. Symbolic Bisimulation	10
9. Evaluation	11
9.1. Utility	12
9.2. Applicability	14
9.3. Extendability	15
9.4. Octopus Variants	16
10. Conclusion	17
11. Future Work	17
A. Supported P4 Grammar	22
B. OCTOPUS P4 Template	25
C. Benchmarks in P4	25
C.1. State Rearrangement	25
C.2. Variable-Length Formatting	26
C.3. Header Initialisation	33
C.4. Speculative Extraction	34
C.5. Relational Verification and External Filtering	35
C.6. Edge and Translation Validation	37
C.7. Service Provider	41
C.8. Datacenter	42

C.9. Enterprise 45

1. Introduction

The Internet can be considered the largest engineered system by humankind [1, Ch. 1]. By 2023, over two-thirds of the global population was expected to have Internet access, with networked devices projected to exceed three times the global population [2]. Predictions indicate this trend will persist [2]. As the Internet continues to grow more interconnected and diverse, it remains a key focus of active research.

Like any other extensive system, the Internet relies on many subcomponents for its functioning. One of these subareas, including its ongoing research, concerns packet parsers. In a *packet-switched network*, such as the Internet, hosts (*i.e.* devices) communicate by sending packets through a network of communication links and packet switches. A *packet* is a structured container of information comprising two components: a header and a payload [1, Ch. 1]. The *packet switch* is responsible for forwarding packets to their destinations, indicated in the header field(s).

As an example, consider the illustration of a UDP packet with its header and payload components in Figure 1. It depicts how a UDP packet's header consists of four fields, including a destination (port). The data field represents the payload of a UDP packet. Due to the layered organisation of protocols [1, Sec. 1.5], *i.e.* the Internet's *protocol stack*, an example of a payload in a UDP packet could very well be another packet.

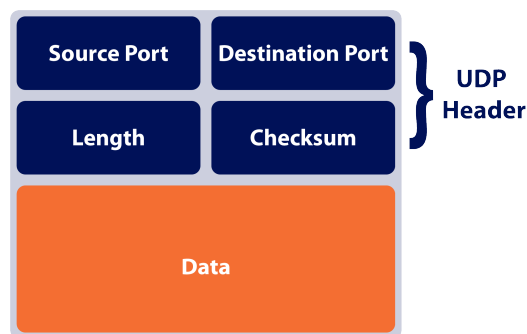


Figure 1: An illustration of the structured information as contained by a UDP packet. The top four blocks represent the fields comprising the packet's header. The bottom block represents the packet's payload.

A packet switch receives a packet as a stream of bits on its incoming link(s). It then has to decide whether the packet is valid, and, if so, to which outgoing link(s) to forward it. To make this decision, the stream of bits must be parsed into the packet's components and respective field(s). This is performed by the packet switch's *packet parser*.

Packet parsers are not limited to use in packet switches; they are also essential in another class of devices known as *middleboxes*. These devices operate on the data path between source and destination hosts and perform functions beyond standard packet forwarding [1, Sec. 4.5]. A subcategory of middleboxes performs security services. For instance, consider a *firewall*, which may block traffic based on specific values of header fields. Similarly, an *intrusion detection system* (IDS) could inspect header fields and payload content, possibly up to a particular layer or "depth". By their very nature, these devices also depend heavily on packet parsing.

Given the central role of packet parsers in both enabling communication and enforcing security, it is critical that they behave according to their functional specifications. Nonetheless, implementation bugs do still occur, and these can lead to incorrect behaviour or security vulnerabilities. Searching for the keyword "packet" in the CVE database [3], a widely used reference for publicly disclosed cybersecurity issues, returns various results. Consider CVE-2020-1350, a vulnerability where, due to an integer overflow in a function responsible for parsing DNS response types, remote code execution (RCE) became possible in a wide range of WINDOWS SERVER editions. Another example is *Ripple20* (CVE-2020-11896–CVE-2020-11914), a set of vulnerabilities discovered in the IP stack implementation

found in a lot of embedded systems. The consequences ranged from RCE to denial of service (DoS) in a variety of devices, including medical applications. Both examples underline the importance of applying formal methods to ensure software correctness [4].

P4 is a domain-specific programming language used to define packet processors, including the behaviour of their packet parsers. As with any software system, P4 programs are susceptible to bugs. They may also undergo modifications over time. These changes can introduce unintended deviations from the intended specification.

To address this problem of unintended deviations, we present OCTOPUS, a tool developed for practical equivalence checking of packet parsers written in P4. The tool aims to verify, within reasonable computational bounds, whether two P4 parsers are behaviourally equivalent. This is relevant as it allows one to compare different revisions of the same parser, representing the possible subsequent modifications. In addition, OCTOPUS generates certificates that facilitate verification of the correctness of its findings. As an example, suppose that a version of a P4 packet parser is deemed correct concerning the intended specification. Now, a programmer has rewritten (parts of) the code in an attempt to improve its performance. OCTOPUS can be used to check whether the two implementations are behaviourally equivalent, *i.e.* no bugs or other deviations from the specification were introduced.

1.1. Thesis Overview

This thesis investigates the problem of practical equivalence checking for packet parsers defined in P4. The primary question we aim to answer is: given an intentional modification to a parser, is its behaviour modulo this deliberate change preserved, *i.e.* equivalent? Beyond constructing a tool for equivalence checking, our goal is to provide guarantees on the correctness of the results produced by this tool. We will attempt to do so by generating certificates alongside the equivalence checks that back the tool's findings.

We begin by presenting the necessary background, including relevant foundational and theoretical concepts, as well as a review of related work, to position our contribution within the existing research landscape. Building on this foundation, we present our methodology for achieving practical equivalence checking. We then describe the design and architecture of the tool, along with its usage and supported syntax. Finally, we evaluate the performance of the tool in comparison with related work, draw conclusions from our findings, and outline directions for future research.


2. Background and Related Work

Software-Defined Networking Within the Internet's protocol stack, the *network layer* is responsible for enabling communication between hosts; at this layer, packets are transferred from a sending to a receiving host. To achieve this, the network layer can be viewed as consisting of two interacting parts: the data plane and the control plane. Within what is considered the *data plane*, it is determined whether and how to forward a packet from an incoming to an outgoing link. Forwarding decisions are made using a *forwarding table*, which maps one or more packet header fields to an output link. These decisions are local to each router. In contrast, the *control plane* embodies network-wide logic that coordinates these local actions to establish end-to-end delivery paths.

Software-defined networking (SDN) is an emergent paradigm within the field of computer networking. Its main promise is better manageability [5]. In a traditional computer network, the control and data planes are vertically integrated, monolithically on each device, and configurable through an interface, often a vendor-specific CLI or GUI. See Figure 2a for an illustration. On the other hand, in SDN, the control and data planes are separated. A logically centralised controller implements the control plane. The devices in the data plane, such as routers and switches, become simple forwarding devices. This is also where SDN finds its name. Namely, the control plane, which forms the heart of the network, is defined by a software-based controller. In practice, this is often achieved through a distributed control system to maximise performance, scalability, and reliability [6], [7]. See Figure 2b for an illustration.

OPENFLOW and P4 Recall that forwarding tables are used in the data plane of a traditional network. In SDN, the software-based controller enables a more universal paradigm called *generalised forwarding* [1, Sec. 4.4]. This approach employs *match-plus-action* tables, which enable more expressive packet matching across multiple fields, protocols, and layers, and which support advanced actions such as load balancing and firewalling.

In practice, the OPENFLOW protocol [8] pioneered the match-plus-action table abstraction. It remains a widely used protocol for defining communication between the controller and the data plane [5]. The protocol is used between an SDN controller and SDN-enabled devices; essentially, the tables defined by OPENFLOW act as an API through which the behaviour of a packet switch can be programmed. While expressive, this model offers only limited programmability. For example, OPENFLOW specifies a fixed set of protocols, restricting extensibility. One could envision a more expressive model that incorporates constructs familiar from general-purpose programming languages, such as variables, conditional logic, and functions. To this end, a domain-specific programming language for *programming protocol-independent packet processors* (P4) [9] was introduced. P4 enables the specification of packet parsing, processing, and forwarding logic in supported devices, without relying on fixed rules, such as those imposed by OPENFLOW.



(a) A network where the control and data planes are vertically integrated [1, Ch. 4]. (b) A network where the control and data planes are separated [1, Ch. 4].

Figure 2: A comparison of two network types. In (a), a network is depicted in which the control and data planes are vertically integrated, *i.e.* monolithically on each device. On the other hand, (b) shows a network where the control and data planes are separated. It is an example of a software-defined network.

Verification and Formal Semantics of P4 As conceptualised in Section 1, software may deviate from its intended specification, potentially leading to security vulnerabilities or other issues. Programs written in P4 are, by definition, also vulnerable to these problems. Formal methods were proposed as one means to mitigate such risks. Applying formal methods to P4 is particularly promising due to the language's expressiveness and the inherent complexity of networked systems, which can allow subtle bugs to go unnoticed and/or untested by developers [4].

To assist developers in verifying the functional correctness of their P4 programs, several push-button verifiers have been developed. These tools, given an annotated P4 program, are capable of checking user-specified properties with minimal manual intervention. Examples include `p4v` [10], `ASSERT-P4` [11], and `AQUILA` [12]. Each allows the specification of functional properties through a manner of annotation. The annotated programs are then translated into intermediate representations (IRs), such as guarded command language (GCL) [13] and C, and verified using formal techniques: verification

condition generation in the cases of P4v and AQUILA, and symbolic execution in the case of ASSERT-P4.

Another approach to checking the correctness of a P4 program is employed by VERA [14]. It operates fully automatically and can detect a wide range of practical bugs. However, it is limited to analysing concrete table snapshots and does not reason over the whole space of possible table entries.

A separate class of tools focuses on verifying the correctness of the translation to concrete network configurations. This area of research is often referred to as *toolstack verification*. These tools have goals different to those of OCTOPUS, which is concerned with parser equivalence rather than end-to-end translation correctness. An example is P4PKTGEN [15], which uses symbolic execution to generate test packets along with their expected processing outcomes. These packets are then put into a software switch (BMV2 [16]), and it is observed whether the actual behaviour aligns with the predictions.

Each of the previous tools relies on ad hoc models of semantics, or utilises an existing implementation, such as P4C¹, the reference compiler of P4. Effort has been made to provide formal semantics for P4, for example, in P4K [17], which specifies these in the K framework [18]. The downside of P4K is that it uses the P4₁₄ language specification, which has a significantly different syntax and semantics compared to the modern P4₁₆ [19]. A more recent development is PETR4 [20], which defines formal semantics for a subset of P4₁₆. Major contributors to the P4 language (specification) have worked on this research. The goal is to extend this approach to the full language and to write a verified compiler for P4. By ensuring correctness upstream in the compiler, the tools that rely on it can also be considered more trustworthy, as a consequence.

Leapfrog From the discussed tools, OCTOPUS is most similar to the functional verifiers. However, as mentioned, we are interested in the relational property of behavioural equivalence between two parsers. All the tools we have seen so far work on a single program instead. In our direction, only a single work has conducted research for P4, namely LEAPFROG [21]. As a result, it is the most relevant.

LEAPFROG, implemented in RocQ, a proof assistant, is a behavioural equivalence checker of P4 packet parsers. The tool accepts a custom syntax representing a subset of P4 functionality. The semantics of this syntax enable a written program to be modelled as a deterministic finite automaton (DFA). Then, by writing packet parsers in this syntax, equivalence checking becomes a matter of proving the equivalence of two DFAs. By formulating equivalence as a (backwards) bisimulation [22] between these underlying DFAs, and using RocQ together with an SMT solver (see Section 8) for formal verification, equivalence can be formally proven. While this approach is effective, it requires a substantial amount of time and computational resources. For example, the implementation of a real-world parser required more than 500 GB of memory and had a runtime of nearly a day, until a timeout occurred due to the machine running out of memory. This is the result of generating a proof object on the fly, which also contains a lot of redundant information. What if such a proof is not constructed, but the equivalence is merely checked? This could significantly improve the algorithm's runtime, possibly allowing it to run on regular PC equipment instead of requiring high-memory servers.

3. Research Question

The principal research question (PRQ) addressed in this thesis is as follows.

(PRQ) What is the impact of not constructing a proof in parallel when checking the behavioural equivalence of P4 packet parsers?

In our context, impact has two equally valid and practical interpretations, leading to the formulation of two sub-questions (RQ1, RQ2) in support of the PRQ.

(RQ1) What effect does avoiding proof construction have on the runtime of P4 packet parser behavioural equivalence checking?

¹<https://github.com/p4lang/p4c>

(RQ2) To what degree can an implementation that does not construct a proof produce a certificate that allows an external program to verify its correctness?

4. Methodology

To investigate a LEAPFROG implementation that does not construct a proof in parallel, we must reimplement their approach. This should be done in a general-purpose programming language that does not construct a proof object like RocQ builds during computation. A lot of work with P4 has been done in C++. For example, the reference compiler of P4 programs, P4C, is written in C++. However, instead, we will use PYTHON. The reason for this is twofold:

- PYTHON is widely used and well-supported, with readily available packages on PyPI for the SMT solvers required in our work, such as Z3 [23] and cvc5 [24], as well as a library package that serves as an abstraction over using different solvers (see Section 8). In contrast, the C++ ecosystem lacks a unified source for all required solvers. For instance, vcpkg does not provide a package for cvc5 at the time of writing.
- We do not expect a significant performance loss from using PYTHON instead of C++. A substantial portion of the runtime will be spent within external SMT solvers, which are implemented as efficient C or C++ libraries and accessed via PYTHON bindings. As such, the performance-critical components remain largely unaffected by the choice of host language.

Besides using a different programming language, it is also worth noting that, while our goal is to reimplement the approach proposed by LEAPFROG, our approach to bisimulation differs in a key aspect. Rather than using a backwards bisimulation strategy as in the original work, we adopt a forward bisimulation approach. Further details are provided in Section 7.

We need to ensure that we obtain results that can be compared to those obtained by LEAPFROG. As a result, our evaluation methodology consists of two separate experiment runs. In the first run, we execute the benchmarks of LEAPFROG in an identical manner as its authors did. Then, after we have written our implementation, we run the same benchmarks with our tool. We perform this task on the same machine under the same conditions, using the same utilities. Ultimately, we obtain findings that can be compared. We will use the outcome of the testing methodology to make a statement about the runtime of our implementation, answering RQ1.

The correctness of our implementation of LEAPFROG will not be proved. Additionally, our algorithm does not construct a proof that verifies the correctness of its findings. Thus, without modifying the algorithm, one is limited to trusting the implementation and its results. A solution would be to extend the algorithm and have it generate a certificate alongside its decision on equality (RQ2). This *certificate* could consist of an encoding of a bisimulation when it deems two parser programs to be equivalent, and a counterexample trace otherwise.

Such a certificate can then be used by an external program to verify the correctness of the certificate's accompanying result without having to fully rerun the analysis. A limitation of this solution is that it cannot guarantee that our implementation produces a correct result for every input. However, we do not consider this an issue, as the individual statements regarding the equivalence of parsers can be verified. This is what is essential for practical applications.

After addressing both sub-questions, we can formulate an answer to the main research question (PRQ) of this work. In doing so, we assign equal weight to the outcomes of both sub-questions. This reflects the view that, for our tool to be practical, it must be both easy to run and capable of producing verifiable results that can be trusted.

5. Overview

In this and the upcoming sections, we will look at our implementation of OCTOPUS. It is available as an open-source project under the MIT license and can be found on [GITHUB](https://github.com/jortvanleenen/Octopus)². The README file in the repository provides a more detailed overview of how to use the tool in practice. To aid practical usage, a DOCKER image has also been pushed to [DOCKER HUB](https://hub.docker.com/repository/docker/jortvanleenen/octopus)³. It is ready-to-use, comes with two SMT solvers installed (Z3 and cvc5), and includes the benchmarks that are part of our testing methodology, facilitating easy verification of our findings.

We divide the description of the tool into three components, each corresponding to a distinct stage in the processing pipeline. These are examined in the following order:

1. The *frontend* (Section 6), which is responsible for obtaining the intermediate representation (IR) of a parser block with its associated types from a P4 program. Afterwards, it parses this IR into a custom object.
2. The *midend* (Section 7), consisting of a DFA implementation, which serves as an abstraction over the custom object produced by the frontend.
3. The *backend* (Section 7, Section 8), ultimately responsible for performing bisimulation, either on top of the DFA abstraction (naive bisimulation, Section 7) or directly over the parser block representations built by the frontend (symbolic bisimulation, Section 8).

6. Frontend

OCTOPUS is to be provided with two file paths: either to P4 programs, or to their IR in JSON format, as generated by P4C, the P4 reference compiler. If provided IR JSONs, it attempts to parse these into PYTHON objects directly. When supplied two P4 programs, two calls are made to the P4C-GRAPHS command-line tool. This CLI application is a backend of P4C which outputs the IR in JSON format that can then be parsed into a PYTHON object on success.

Once PYTHON objects for both P4 programs are obtained, they are parsed into ParserProgram objects. These represent solely the parser block within a P4 program with its types annotated. This conversion is performed by a hand-rolled parser, implemented as a class hierarchy. An illustration of the full pipeline from P4 program to ParserProgram can be seen in Figure 3.




Figure 3: The preprocessing pipeline as used by OCTOPUS. It is run twice, once for each input. Note that entry to the pipeline can occur at either “P4 Program”, or “IR JSON”. Once two ParserProgram objects are obtained, we can move on to checking for their equivalence.

P4C has support for the full P4 language, including both P4₁₄, as well as P4₁₆ [19], the most recent revision of the P4 language. This is not the case for OCTOPUS, which supports only a subset of P4₁₆. It should be noted, though, that P4C can be used to translate P4₁₄ into P4₁₆. So, as long as the P4₁₆ syntax to which the P4₁₄ source is translated is supported, OCTOPUS could accept such programs as well. Whether a language feature is included or excluded from OCTOPUS’s subset comes down to three criteria:

1. Is the feature relevant for the parser block in a P4 program? If not, then it is irrelevant to OCTOPUS and thus ignored.

²<https://github.com/jortvanleenen/Octopus>

³hub.docker.com/repository/docker/jortvanleenen/octopus

2. Is the feature translatable into a DFA-centric approach? If not, then it falls outside the scope of this research. An example of this includes `varbit`, a variable-length bitstring type. Its behaviour resembles that of a stack-like data structure, which may require modelling with a pushdown automaton (PDA) rather than a DFA. As a result, features like `varbit` are excluded from this work.
3. Is the feature essential to the implementation? Certain features of P4 serve primarily as syntactic conveniences. For example, the `enum` type can be represented equivalently using fixed-size bit-vectors within the context of a packet parser. Due to time constraints, such non-essential features have been omitted from this work.

Based on the criteria above, we define a formal grammar that captures the syntax accepted by OCTOPUS; see Listing 3. This grammar is written to closely mirror the formal grammar specified in [19], ensuring comparability. Our grammar, and thus accepted syntax, fully covers the syntax and semantics supported by LEAPFROG, making it at least equally expressive. In addition, it extends upon their language by introducing two extra operators: bit-vector AND and bit-vector shift-right. These extensions demonstrate the ease with which OCTOPUS’s language subset can be expanded, owing to its modular and hierarchically structured implementation. Further details are provided in Section 9.3.

Features that fall under the categories described by points 2 and 3, namely those that are either non-translatable into a DFA-centric approach or non-essential to the implementation, will trigger a warning from OCTOPUS at runtime. The tool will attempt to recover gracefully, but will raise an exception if recovery fails. This behaviour contrasts with that for features falling under point 1. By default, OCTOPUS will quietly ignore them, as some of which are considered essential for the compilability of the P4 program. Such features are safe to ignore as they are not relevant to the semantics of parser blocks. This distinction reflects the broader observation that parser blocks represent only one component of a complete P4 program.

To support this separation of concerns between the syntax required for the actual parser analysis and the syntax just required for program compilation, we define a minimal template, as shown in Listing 4. This template serves as a lightweight framework within which arbitrary parser blocks can be embedded. It was used for all benchmarks in this work, demonstrating its general applicability. On a similar note, the formal grammar defined earlier (Listing 3) deliberately excludes language constructs unrelated to parser semantics.

More informally, we restrict the P4₁₆ language according to the following conventions:

- OCTOPUS expects headers to consist solely of fixed-width bit-vectors. These headers must eventually be grouped into a `struct` that represents the parsed packet upon completion of the parser.
- A parser block must take exactly two parameters: one variable of type `packet_in`, representing the input stream, and one out-directed variable of the aforementioned `struct` type.
- Within a parser block, local variables and instantiations are disallowed. Only assignment statements are permitted, possibly using a method call to extract bits from the stream. Each parser state must include at least one extraction; states consisting solely of transitions are not allowed. We also disallow nested blocks and direct application.
- Slice expressions are not permitted on the left-hand side of assignments (*i.e.* as `lvalues`).
- Set types are not supported. Consequently, the mask and range operators are disallowed in select expressions. List expressions, used as syntactic shorthand, are also excluded because they are not essential to our implementation.
- Only a limited subset of operators is supported: slicing, concatenation, bit-vector shift-right, and bit-vector AND.

We adhere to the semantics of P4 as defined in [19], and formalised in [20]. As such, we do not redefine it here.

7. Naive Bisimulation

As illustrated in Figure 3, we obtain two `ParserProgram` objects representing the two parser blocks with their associated types. This section illustrates how one can transition from this implementation to a DFA, which is defined as follows [25]. It is a 5-tuple $(Q, \Sigma, \delta, q_0, F)$, where:

1. Q is a finite set called the *states*,
2. Σ is a finite set called the *alphabet*,
3. $\delta : Q \times \Sigma \rightarrow Q$ is the *transition function*,
4. $q_0 \in Q$ is the *initial state*,
5. $F \subseteq Q$, is the set of *accepting states*.

A DFA works as follows. We begin in an initial state and process an input string symbol by symbol from a fixed alphabet. At each step, the current state and the input symbol determine a unique transition to the next state. The machine continues until all input symbols are consumed. If it halts in an accepting state, the input is accepted; otherwise, due to the absence of a valid transition or ending in a non-accepting state, it is rejected.

We will directly base our translation from `ParserProgram` to DFA on P4A as defined in [21]. As a result, most of the formulas that we will describe are (near-)equivalent. However, we will not specify the full semantics as they do. These are already defined for P4 in, e.g., [20] as indicated before. Instead, we will state the essential formulas and notations required to understand the eventual naive bisimulation algorithm and its implementation.

Every packet parser works on an input stream of bits. More formally, when we consider a packet parser as a state machine, its alphabet can be regarded as the set of symbols for which it needs to determine whether to accept or reject, so $\Sigma = \{0, 1\}$ for all P4 packet parsers.

A P4 packet parser always has a finite number of states per the language definition, which we will consider Q , and a finite set of associated types, which we will consider as a set of *header names* H for simplicity's sake.

Each *header* $h \in H$ has an associated size $sz(h) \in \mathbb{N}^+$. When we write $|bv|$ for the length of $bv \in \{0, 1\}^*$, we can define S as the set of finite functions: $s : H \rightarrow \{0, 1\}^*$, where $|s(h)| = sz(h)$.

In P4, each state $q \in Q$ can be considered to consist of two components: an operation block ($op(q)$) and a transition block ($tz(q)$). The *operation block* is responsible for reading in new bits and updating the *store* S , the structured representation of what has been read. The *transition block* is responsible for determining to which state to transition next. An operation block will *buffer*, ensuring that it will only execute once the total number of bits it will read is available on the input stream/buffer. To know how many bits it will read, we need to consider the summation of all $sz(h)$ for which a $extract(h)$ method call exists in $op(q)$. Note that $sz(h)$ is known for all h because we have parsed the associated types with each parser block in a `ParserProgram` object.

A P4 packet parser, represented as a state machine, can be seen in Figure 4. A state named “start” always represents the initial state. It transitions to other states based on what has been read and what is currently in the buffer. A packet parser terminates once the buffer is empty and a terminal state has been reached. In P4, each packet parser has two such states, namely “reject” and “accept”.

As in [21], we can now consider a *configuration*, which represents the state of the DFA representation of a `ParserProgram`. A *configuration* is a triple:

$$\langle q, s, w \rangle \in (Q \cup \{\text{accept}, \text{reject}\}) \times S \times \{0, 1\}^*.$$

It can be interpreted as being in a state in Q , with the contents of a store defined by S , and with w in the buffer. We define C as the finite set of configurations, and F as the set of accepting configurations. Building up the buffer can be defined as a bit-by-bit operation, representing our transition function




Figure 4: An illustration [19] of a parser defined in P4. Its initial state is always called “start” and is part of the definition of the parser block. It always has precisely two terminal states, “accept”, an accepting state, and “reject”, a rejecting state. These are implicit, namely, they are not part of the specification as supplied by the programmer.

δ , where we append the buffer by bit b if $|wb| < sz(op(q))$, and else we execute the operation block $op(q)$ and transition according to the transition block $tz(q)$.

As in P4, any further input on the buffer when the parser is in a terminal configuration will always result in a transition to $\langle reject, s, \epsilon \rangle$.

Finally, as DFA representation, we can consider the triple: $\langle C, \delta, F \rangle$. Note how the alphabet and the initial state are implicit. That is okay, as these are equivalent for every P4 packet parser.

7.1. Bisimulation

To determine whether two packet parsers are behaviourally equivalent, we construct a bisimulation. Given its central role in this thesis, we illustrate the concept using an accessible analogy inspired by [22].

Consider two vending machines. Each has a payment terminal, a button to request hot chocolate, a button to request coffee, and a dispensing slot. By interacting with the machine, *e.g.*, by pressing buttons and observing the output, we can assess its behaviour. Most importantly, we are not concerned with the machine’s internal mechanisms, physical design, or implementation details. What matters is how it behaves from an external perspective: what inputs it accepts and what outputs it produces in response to those inputs.

This mirrors our treatment of packet parsers. Their internal structure is irrelevant for our purposes. Instead, we focus on their observable behaviour, specifically, whether they accept or reject the same input. If two parsers respond identically to all possible inputs, we consider them behaviourally equivalent.

Formally, a bisimulation [22] is defined as being a relation $R \subseteq Q_0 \times Q_1$ for which holds:

1. $q_0 \in F_0 \iff q_1 \in F_1$
2. $\forall_{b \in \{0,1\}} \delta_0(q_0, b) R \delta_1(q_1, b)$

We have seen that we can model each parser as a DFA defined by the tuple $\langle C, \delta, F \rangle$. Here, the set of states C corresponds to the set Q as used in the formal definition of bisimulation. Thus, we can check for bisimilarity of ParserProgram objects by representing them as DFAs. In practice, each ParserProgram is passed to a DFA class, which defines DFA operations like those specified above over the semantics of P4. The resulting automata can then be checked for bisimilarity. An overview of this architecture is provided in Figure 5.

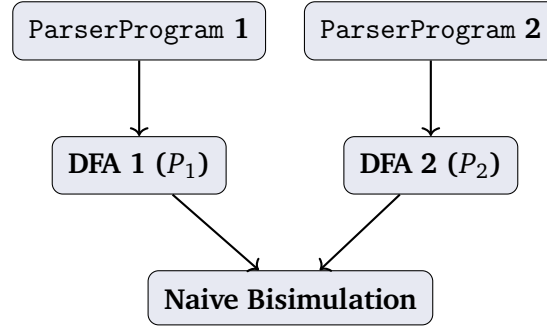


Figure 5: An overview of how naive bisimulation is implemented.

Using pseudocode, we will now describe the operation of the transition block, a component we have so far only briefly touched upon. See Listing 1.

The transition block begins by evaluating the expression used in the select statement. It then iterates over each specified keyset in the order they appear. For each keyset, its value is evaluated and compared against the evaluated key. If the key matches the current keyset, we transition to the respective state. If no match is found, a runtime error is triggered. In our implementation, we model this behaviour by treating it as an implicit transition to the reject state. If a transition block is not explicitly defined for a state, the behaviour defaults to an implicit transition reject.

Listing 1: Pseudocode [19] for the transition block of an arbitrary state in a P4 parser block.

```

key = eval(e);
for (int i=0; i < n; i++) {
    keyset = eval(ks[i]);
    if (keyset.contains(key)) return s[i];
}
verify(false, error.NoMatch);
  
```

Due to the state space explosion associated with naive bisimulation, we will not be investigating this approach further. However, it was very helpful to verify the correctness of our parsing logic, as its straightforward translation allowed for easy debugging. To get an idea of the performance that can be expected when using naive bisimulation, it was made available as an option for OCTOPUS. Additionally, on the author’s computer, 16 bits took around 7 seconds to run, where each bit scales this number by a factor of two. It is clear to see why naive bisimulation is practically impossible for parsers with say more than 24 bits.

8. Symbolic Bisimulation

Just as with naive bisimulation, we also perform a translation for symbolic bisimulation. Instead of using DFAs on top of the parsers, we work with formulas. See Figure 6 for an illustration.




Figure 6: An overview of how symbolic bisimulation is implemented.

The formulas we consider are referred to as template-guarded. They are five-tuples: $(q^<, q^>, n^<, n^>, \varphi)$. The first two elements indicate in which state the left and right parsers are. The third and fourth elements

indicate the length of the buffers in both the left and right parsers. Finally, φ represents the current knowledge. It can be considered a statement of what the group of states we symbolically consider has in common.

Key to symbolic bisimulation are the functions for the strongest postcondition (SP) for operations, and symbolic transition (ST) for symbolically representing the transition block of a state.

To understand the strongest postcondition, consider: $\{p\} S \{q\}$. Here, p represents a precondition, which are assertions that we assume to hold before the execution of program S . q represents the postcondition, which are assertions that hold after S assuming p . It is referred to as a Hoare-triple. This q can vary in ‘specificity’, where the official term for this is strength. The more ‘specific’ q is, the stronger it is considered. Consider that p stands for $Cow(x)$ and that S is a no-operation (NOP); in other words, it does not modify the precondition. A valid postcondition q could be $Animal(x)$, as this is implied by $Cow(x)$. However, the strongest postcondition would be $Cow(x)$. Another way to look at this is from the perspective of the Liskov Substitution Principle in object-oriented programming. Cow can be placed at positions where $Animal$ is used, but not the other way around ($Cow \subset animal$, but $Animal \not\subset Cow$).

Similarly to the above illustration for programs S , we can define the strongest postcondition for operations that can occur within the operation block of the packer parser. See Figure 7.

$$\begin{aligned}
& SP(\varphi, \text{prog}) \\
& SP^{\leq} \{ \varphi, \text{extract}(hdr) \} := \exists y_i \in \{0, 1\}^{sz(hdr.field_i)} \quad \varphi \left[\frac{y_i}{st^{\leq}.hdr.field_i}, \frac{st^{\leq}.hdr.field_i ++ buf^{\leq}}{buf^{\leq}} \right] \\
& SP^{\leq}(\varphi, \text{hdr}.x = e) := \exists y \in \{0, 1\}^{sz(x)} \quad \varphi \left[\frac{y}{st^{\leq}.x} \right] \wedge st^{\leq}.x = e \\
& SP^{\leq}(\varphi, \text{prog}_1; \text{prog}_2) := SP^{\leq}(SP^{\leq}(\varphi, \text{prog}_1), \text{prog}_2)
\end{aligned}$$

Figure 7: The strongest postcondition of components of the operation block, mathematically defined.

Remember from Section 7 that a pair was in the bisimulation if it agreed on acceptance, and all the reachable pairs from it were in the bisimulation as well. In symbolic bisimulation, we do not explicitly represent pairs of states. Instead, we represent groups of equivalent states using template-guarded formulas. In order to answer questions about these formulas, such as two transitions, left and right, are satisfiable, *i.e.* doable based on our current knowledge φ , we need to ask questions to SMT solvers. In our implementation, we work with a library that abstracts away the individual APIs of SMT solvers. Instead, it defines a general interface through which we can interact with supported solvers. The package is called PySMT. For an overview, see Figure 8.

The symbolic transition function will translate transition blocks into their symbolic equivalent. See Algorithm 1 for pseudocode.

9. Evaluation

Having defined and implemented a tool for checking the equivalence of packet parsers written in P4, we now evaluate the impact of our approach. Specifically, we compare its runtime performance against that of LEAPFROG and demonstrate the generation of a certificate in the presence of a counterexample. Together, these evaluations allow us to assess the practical applicability of OCTOPUS.

To enable a direct comparison, we adopt the same benchmark set used in LEAPFROG’s evaluation. As LEAPFROG does not accept P4 syntax, we converted their benchmarks to P4. Each converted benchmark can be found in Appendix C. An overview of the benchmarks and their associated properties is presented in Table 1. Following the methodology outlined in [21], we distinguish between two classes of case

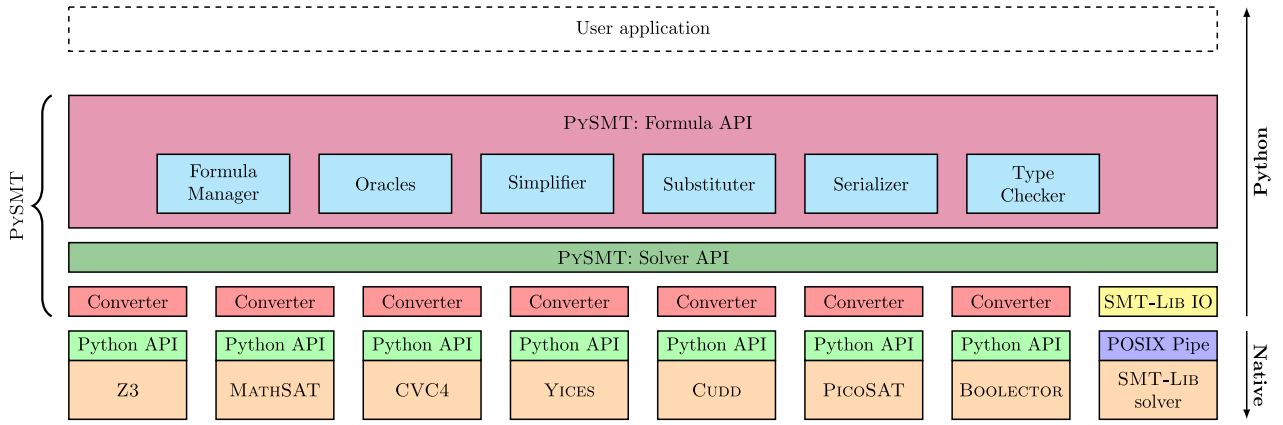


Figure 8: A graphical representation [26] of how PySMT is structured.

studies: utility and applicability. Because detailed definitions for all benchmarks are available in [21], we provide only summaries in the following sections for completeness.

As described in Section 4, we reran the LEAPFROG benchmarks on a high-memory machine at LEIDEN UNIVERSITY. With assistance from Dr. Tobias Kappé, we resolved several dependency issues and successfully executed the benchmarks. All series of experiments were conducted on LATINUM, a server within the REL COMPUTE lab at the LIACS RESEARCH AND EDUCATION LABORATORY (REL). This machine was selected based on the findings reported in [21], which demonstrated the necessity of a high-memory environment, an essential requirement met by LATINUM. Namely, it is equipped with ~ 1.5 TB of system memory. All experiments were executed using the respective DOCKER containers provided by each tool. these were run on an INTEL XEON E5-2630v3. The results can be found in Table 2. All the OCTOPUS benchmark results displayed are averages over four runs, where the first run is ignored to avoid cold start skewing the results.

9.1. Utility

The *utility* class of case studies consists of problems that are designed to evaluate the usefulness of equivalence checking in the networking domain.

State Rearrangement Compilers may merge or split parser states to optimise for hardware constraints. In this case study, we compare two versions of a parser. In one version, fields are extracted in separate states. In the alternate version, we extract both a field and a common prefix simultaneously, thereby merging two states into one. After which, the first field is checked to determine whether additional parsing is needed. This represents a typical example of state merging. OCTOPUS is used to show that the variants accept the same packets, *i.e.* they are behaviourally equivalent.

Variable-Length Formats Many network protocols contain fields whose lengths are not fixed but are instead encoded within the packet itself. In this case study, we examine two scenarios: one in which up to two such fields may appear, and another with up to three. The permissible lengths for each field range from 0 to 6 bytes. We use OCTOPUS to check for equivalence in both cases. Notably, this case study goes beyond what was demonstrated in LEAPFROG. Due to performance limitations, LEAPFROG was unable to verify equivalence when three variable-length fields were involved. In contrast, the improved runtime performance of OCTOPUS on the two-field case encouraged us to attempt the three-field variant as well, which we eventually successfully verified.

Header Initialisation A common mistake when implementing packet parsers is failing to initialise a header field along all execution paths. This may lead to accessing uninitialised data in certain branches.

Algorithm 1: Conversion of a transition block to symbolic transitions.

Input : cases $((\text{expr}, \dots) \rightarrow \text{state})$, selectors $((\text{expr}, \dots))$

Output: symbolic_transitions $\{(\varphi \rightarrow \text{state}), \dots\}$

```
1 if #selectors = 0 then
2   | return  $\{(\top, \text{cases}[\text{DontCare()}])\}$ 
3 fi

4 symbolic_transitions  $\leftarrow \emptyset$ 
5 seen  $\leftarrow \emptyset$ 

6 foreach (for_exprs, state)  $\in$  cases do
7   formula  $\leftarrow \top$ 
8   for  $i \leftarrow 0$  to #for_exprs do
9     if for_exprs[i] is not an instance of DontCare then
10      expression  $\leftarrow \text{deepcopy}(\text{for\_exprs}[i])$ 
11      expression.to_formula()
12      selector  $\leftarrow \text{deepcopy}(\text{selectors}[i])$ 
13      selector.to_formula()
14      formula  $\leftarrow \text{formula} \wedge (\text{expression} = \text{selector})$ 
15    fi
16  od

17  appended_formula  $\leftarrow \text{formula}$ 
18  foreach  $\varphi \in \text{seen}$  do
19    | appended_formula  $\leftarrow \text{appended\_formula} \wedge \neg \varphi$ 
20  end
21  seen  $\leftarrow \text{seen} \cup \{\text{formula}\}$ 
22  symbolic_transitions  $\leftarrow \text{symbolic\_transitions} \cup \{(\text{appended\_formula}, \text{state})\}$ 
23 end
24 return symbolic_transitions
```

To evaluate whether OCTOPUS can detect such issues, we constructed two versions of a parser. In the first version, a header field is initialised in one state, but not for all possible paths. Using equivalence checking, OCTOPUS can confirm that the program’s behaviour does not depend on the uninitialised version of the field. This result is reported in Table 2. In the second version, the initialisation is removed entirely, so even for the execution flow that depends on the field. In this case, OCTOPUS raises an exception, correctly identifying the access to an uninitialised header. This demonstrates the tool’s applicability to both scenarios.

Speculative Extraction Another performance optimisation that compilers tend to make is to speculatively extract bits from the buffer in an attempt to effectively vectorise this operation. We have implemented two parsers: one employing speculative extraction and one that does not. We again use OCTOPUS to show their behavioural equivalence.

External Filtering and Relational Verification A parser may be strict or lenient. A *strict* parser explicitly checks for values, whereas a *lenient* parser may test for one value x and, if the test fails, silently assume the alternative y . This is useful for tolerating malformed packets. We consider two parsers, one “sloppy” (lenient) and one strict.

LEAPFROG can express *external filters*. These initial relations constrain the final stores and are capable of verifying richer relational properties (e.g., if both parsers accept, their stores must agree on a certain field). OCTOPUS, by contrast, currently only supports behavioural equivalence checking of the parser

Table 1: The benchmark set [21] as used for evaluation. In OCTOPUS, the benchmarks are recognisable by their **Name**, whereas in LEAPFROG they can be recognised by their **File**. **States** represents the total number of states in both parser programs. **Branched** indicates the considered number of bits in all `transition select` statements. **Total** reflects the number of bits over all variables. An explicit state space would contain 2^T states. An optimal verification algorithm would need to represent 2^B states [21].

	Name	File ^a	States	Branched (b)	Total (b)
Utility	State rearrangement	IPFilter ^b	5	8	136
	Variable-length form. 2	IPOptions2	30	64	632
	Variable-length form. 3	IPOptions3	45	80	672
	Header initialisation	SelfComparison	10	10	320
	Speculative extraction	MPLSVectorized	5	2	160
	Relational verification	SloppyStrictStores	6	64	1056
	External filtering	SloppyStrictFilter	6	64	1056
Applicability	Edge	EdgeSelf	28	52	3184
	Service provider	ServiceproviderSelf	22	50	2536
	Datacenter	DataCenterSelf	30	242	2944
	Enterprise	EnterpriseSelf	22	176	2144
	Translation validation	EdgeTrans	30	56	3148

^a The full filenames are: <name-in-table>Proof.v, except SloppyStrictStores/Filter.

^b The GITHUB repository of LEAPFROG denotes EthernetProof.v as the benchmark file of “state rearrangement”, this is incorrect.

blocks themselves; it does not support external filters. We therefore use the sloppy-strict pair solely as a sanity check. As is clear from their definition, these parsers are not equivalent. Comparing them using OCTOPUS indeed results in a counterexample trace; see Listing 2 for a truncated example. It clearly shows how, from both start states, a contradiction can be achieved. These traces have also been proven very helpful during the debugging of OCTOPUS.

Listing 2: A truncated counterexample trace obtained by checking the sloppy-strict parser pair from the external filtering and relational verification case studies.

--- Counterexample ---

Step 0:

Left state: start, Right state: start
Buffer lengths: left=0, right=0

Step 1:

Left state: parse_ipv4, Right state: reject
Buffer lengths: left=0, right=0

Step 2:

Left state: accept, Right state: reject
Buffer lengths: left=0, right=128

9.2. Applicability

Under the applicability class, we consider real-world parsers. This helps us estimate the performance and thus practical applicability of OCTOPUS. To this end, we use the PARSER-GEN framework [27]. As in LEAPFROG, we examine four representative parsers:

- *edge*: modelling an edge router,
- *service provider*: modelling a core router,
- *datacenter*: modelling a top-of-rack switch, and
- *enterprise*: modelling a router for a typical campus or mid-sized enterprise.

As in LEAPFROG, we omit the *big-union* parser, as it serves primarily to estimate hardware resource requirements and is not representative of practical, real-world parser configurations [21].

We based our implementations of these parsers on the representations written by the authors of LEAPFROG. We did have to modify the enterprise parser with respect to the one found in LEAPFROG’s GITHUB repository⁴. Namely, an indexing error occurred because the size of the IPv4 header was incorrectly specified. We were able to adjust this to the correct value by looking at the original code at the PARSE-GEN repository⁵, where the size was correctly defined (as $4 \times 4 \times 8 = 128$).

We used OCTOPUS to perform a self-equivalence check for each parser, verifying that each is equivalent to itself. This serves both as a sanity check and as a means of assessing the tool’s real-world performance. Additionally, the PARSE-GEN framework includes a compiler that produces a hardware-optimised version of a parser. In [21], this was used to compile the edge benchmark. We replicate this experiment (translation validation) by using OCTOPUS to check the plain and compiled versions of the edge parser in P4 for behavioural equivalence.

Name	LEAPFROG		OCTOPUS (Z3, cvc5)	
	Runtime (m)	Memory (GiB)	Runtime (m)	Memory (GiB)
State rearrangement	1.13	1.07	0.08	0.04
Variable-length form. 2	3825.70	391.91	1.91	0.33
Variable-length form. 3	–	–	46.22	4.36
Header initialisation	21.56	13.48	0.01	0.04
Speculative extraction	5.45	3.26	0.01	0.04
Relational verification	1.59	1.65	–	–
External filtering	2.37	2.03	–	–
Edge	730.12	243.35	0.19	0.08
Service provider	7167.03	823.70	0.08	0.05
Datacenter	1975.92	382.17	4.66	0.76
Enterprise	817.12	66.35	2.77	0.39
Translation validation	1035.38	349.15	0.38	0.12

Table 2: Runtime and memory usage for each case study under both LEAPFROG and OCTOPUS. Runtime, given in minutes, refers to the total wall-clock time required for the complete execution of the equivalence checking process. Memory usage, reported in GiB, indicates the peak resident memory consumed during this process.

9.3. Extendability

As discussed in Section 6, we extended the syntax as supported by OCTOPUS to include bit-vector AND and bit-vector shift-right operations. This demonstrated the ease with which the accepted language subset can be expanded. To also show how to validate the correctness of these additions easily, we introduced a small benchmark titled *extended syntax*. It builds on the plain parser from the speculative

⁴<https://github.com/verified-network-toolchain/leapfrog/blob/main/lib/Benchmarks/Enterprise.v>

⁵<https://github.com/grg/parser-gen/blob/master/examples/headers-enterprise.txt>

extraction case study, but redefines the alternate version. Rather than performing speculative extraction, it modifies the `select` statement.

Specifically, the original version extracts the 9th bit using a slice: `field[8:8]`, while the alternate version uses the extended syntax: `field > 8 & 0x1`. Using OCTOPUS, we confirmed that both parsers are behaviourally equivalent.

9.4. Octopus Variants

In addition to comparing OCTOPUS against LEAPFROG, we also conducted benchmarking across different configurations of OCTOPUS itself. Specifically, we evaluated two different solver configurations: one using only Z3, and the other using only cvc5. The results are summarised in Table 3.

As expected, memory usage remained largely consistent across both configurations. This is reasonable, as both solvers should result in OCTOPUS examining the same formulas in its work queue, and, ultimately, in generating a certificate with equivalent content, albeit possibly in a different order. The more notable difference lies in runtime: across all benchmarks, the configuration using cvc5 consistently achieved faster execution. A Wilcoxon signed-rank test, applied to the paired average runtimes with a significance threshold of $p = 0.05$, confirmed that this performance difference is statistically significant. In other words, we can state with high confidence that cvc5 outperforms Z3 for our benchmark suite. It is also noteworthy that both configurations outperform the parallel solver portfolio, probably due to some overhead in parser creation and calling. Nonetheless, given the relatively small number of benchmarks, we retain a parallel solver configuration as the default OCTOPUS portfolio to fence against instances that do not follow this trend.

We also evaluated the impact of disabling the leaps optimisation in OCTOPUS. As shown in Table 3, disabling leaps results in a significant increase in both runtime and memory consumption, even for the smallest benchmarks. This increase can be attributed to the larger state space that must be explored without making leaps, resulting in higher computational overhead and memory requirements. These results underline the importance of leaps in ensuring the practical scalability of OCTOPUS as we have observed.

Name	OCTOPUS (Z3)		OCTOPUS (cvc5)		OCTOPUS (NO LEAPS)	
	Run (m)	Mem (GiB)	Run (m)	Mem (GiB)	Run (m)	Mem (GiB)
State rearrangement	0.01	0.05	0.01	0.04	0.35	0.24
Variable-length form. 2	2.11	0.33	1.84	0.31	–	–
Variable-length form. 3	42.45	4.36	40.58	4.34	–	–
Header initialisation	0.02	0.05	0.01	0.04	5.64	3.13
Speculative extraction	0.02	0.05	0.01	0.04	0.48	0.30
Relational verification	–	–	–	–	–	–
External filtering	–	–	–	–	–	–
Edge	0.31	0.09	0.17	0.07	–	–
Service provider	0.14	0.07	0.08	0.05	149.18	49.68
Datacenter	5.68	0.77	4.16	0.74	–	–
Enterprise	3.29	0.39	2.54	0.37	–	–
Translation validation	0.54	0.13	0.35	0.11	–	–

Table 3: A tabular overview of the performance of different OCTOPUS configurations. Note that “runtime” and “memory” have been abbreviated to “run” and “mem” respectively. The first column group details the performance of OCTOPUS using only the Z3 SMT solver in its solver portfolio. Likewise, the configuration depicted in the second column group employs just the cvc5 solver. The final configuration in the third, right-most column group uses the default solver portfolio, but without leaps enabled.

10. Conclusion

It has been demonstrated that we were able to implement forward bisimulation and symbolic execution with the goal of equivalence checking of P4 packet parsers, inspired by LEAPFROG. We have done so in a manner that accepts a subset of the P4 syntax and accurately follows its semantics. Furthermore, due to the modular structure of the codebase, it has been shown to be easily extendable to accept a larger subset of P4₁₆.

We successfully replicated all findings reported by LEAPFROG, with the exception of those that rely on external filtering, which is not supported by OCTOPUS. Conversely, OCTOPUS was able to verify the variable-length format benchmark with up to three fields, whereas LEAPFROG was limited to two.

By not constructing a proof object on the fly during computation, we have achieved a more practical runtime performance. Each of the equivalence checks over parser blocks that we have considered could have been executed on a modern PC. We therefore conclude that OCTOPUS exhibits practical runtime performance, thereby answering **RQ1**.

Additionally, we have shown that our method of generating certificates is helpful and sufficient for practical usage in both positive and negative cases of behavioural equivalence. Especially the counterexample trace has already proven helpful, both during debugging, as mentioned before, but also during the conversion of LEAPFROG benchmarks to P4, as it allowed us to pinpoint issues that would have been hard to spot otherwise. We therefore consider OCTOPUS sufficiently capable of producing certificates suitable for practical and trusted use, thereby addressing **RQ2**.

In conclusion, obtaining favourable outcomes for both **RQ1** and **RQ2**, we have described and implemented a tool that can be used for the practical equivalence checking of P4 packet parsers, answering **PRQ**.

11. Future Work

While OCTOPUS demonstrates practical runtime performance, there remains potential for further optimisation. Currently, the implementation is fully synchronous and does not leverage multiprogramming. Although we utilise some asynchronous features of PySMT, such as portfolio solving with solvers running in parallel, the validity and satisfiability queries are invoked synchronously from within our codebase. A promising direction for future improvement would be to parallelise the processing of the work queue maintained during symbolic bisimulation. Given that all pure formula objects should be unique copies, this could be a straightforward modification.

While we have experimented with the Z3 and cvc5 SMT solvers, PySMT supports a wider range of solvers. Built-in support exists, for example, for BOOLECTOR 3.0 [28], MATHSAT 5 [29], and YICES 2.2 [30], all of which support the bit-vector SMT theory. PySMT also supports any SMT solver that is SMT-LIB 2 compliant. It could be interesting to explore a different or more extensive combination of SMT solvers for the portfolio that symbolic bisimulation uses. This is especially interesting for solvers for which we do not have directly comparable results available, such as cvc5, which performed well in SMT-COMP 2024. Another strong contender in SMT-COMP 2024 was BITWUZLA [31], which is SMT-LIB 2 compliant and won most subcategories in the QF_BITVEC (Single Query Track) of SMT-COMP 2024.

We chose to support a subset of the P4₁₆ language specification, as shown in Listing 3, which matches the supported syntax by LEAPFROG. As we have shown, this subset can be easily expanded due to the structure of the OCTOPUS codebase. In further work, one could expand this subset to support other language constructs. For example, one could allow the use of slices in lvalues, which currently have to be emulated by concatenating slices on the right-hand side of the assignment operator. Another example could be implementing more operators, allowing for more expressive expressions to be made.

Finally, as mentioned, we were unable to perform the external filtering and relational verification benchmarks as required for the proper comparison of the sloppy and strict P4 programs. The codebase of OCTOPUS could be expanded to support this behaviour, rather than, for example, at the moment

having to parse the certificate and possibly performing external filtering there, before returning true or false for equivalence.

References

- [1] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach*, 8th (Global Edition). Harlow, United Kingdom: Pearson Education Limited, 2022, ISBN: 978-1-292-40546-9.
- [2] T. Barnett, S. Jain, U. Andra, and T. Khurana, *Cisco Annual Internet Report (2018–2023)*, en, 2018. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html> (visited on 05/27/2025).
- [3] *CVE: Common Vulnerabilities and Exposures*, May 2025. [Online]. Available: <https://www.cve.org/> (visited on 06/01/2025).
- [4] L. Sassaman, M. L. Patterson, S. Bratus, and M. E. Locasto, “Security applications of formal language theory,” *IEEE Systems Journal*, vol. 7, no. 3, pp. 489–500, 2013, Publisher: IEEE. DOI: [10.1109/JSYST.2012.2222000](https://doi.org/10.1109/JSYST.2012.2222000).
- [5] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, “Software-defined networking: A comprehensive survey,” *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2014, Publisher: Ieee. DOI: [10.1109/JPROC.2014.2371999](https://doi.org/10.1109/JPROC.2014.2371999).
- [6] T. Koponen, M. Casado, N. Gude, *et al.*, “Onix: A distributed control platform for large-scale production networks,” in *9th USENIX symposium on operating systems design and implementation (OSDI 10)*, 2010. DOI: [10.5555/1924943.1924968](https://doi.org/10.5555/1924943.1924968).
- [7] S. Jain, A. Kumar, S. Mandal, *et al.*, “B4: Experience with a globally-deployed software defined wan,” en, *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 3–14, Sep. 2013, ISSN: 0146-4833. DOI: [10.1145/2534169.2486019](https://doi.org/10.1145/2534169.2486019).
- [8] N. McKeown, T. Anderson, H. Balakrishnan, *et al.*, “OpenFlow: Enabling innovation in campus networks,” en, *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, Mar. 2008, ISSN: 0146-4833. DOI: [10.1145/1355734.1355746](https://doi.org/10.1145/1355734.1355746).
- [9] P. Bosshart, D. Daly, G. Gibb, *et al.*, “P4: Programming protocol-independent packet processors,” en, *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 87–95, Jul. 2014, ISSN: 0146-4833. DOI: [10.1145/2656877.2656890](https://doi.org/10.1145/2656877.2656890).
- [10] J. Liu, W. Hallahan, C. Schlesinger, *et al.*, “P4v: Practical verification for programmable data planes,” en, in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, Budapest Hungary: ACM, Aug. 2018, pp. 490–503, ISBN: 978-1-4503-5567-4. DOI: [10.1145/3230543.3230582](https://doi.org/10.1145/3230543.3230582).
- [11] M. Neves, L. Freire, A. Schaeffer-Filho, and M. Barcellos, “Verification of P4 programs in feasible time using assertions,” en, in *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*, Heraklion Greece: ACM, Dec. 2018, pp. 73–85, ISBN: 978-1-4503-6080-7. DOI: [10.1145/3281411.3281421](https://doi.org/10.1145/3281411.3281421).
- [12] B. Tian, J. Gao, M. Liu, *et al.*, “Aquila: A practically usable verification system for production-scale programmable data planes,” en, in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, Virtual Event USA: ACM, Aug. 2021, pp. 17–32, ISBN: 978-1-4503-8383-7. DOI: [10.1145/3452296.3472937](https://doi.org/10.1145/3452296.3472937).
- [13] E. W. Dijkstra, “Guarded commands, nondeterminacy and formal derivation of programs,” en, *Communications of the ACM*, vol. 18, no. 8, pp. 453–457, Aug. 1975, ISSN: 0001-0782, 1557-7317. DOI: [10.1145/360933.360975](https://doi.org/10.1145/360933.360975).
- [14] R. Stoenescu, D. Dumitrescu, M. Popovici, L. Negreanu, and C. Raiciu, “Debugging P4 programs with vera,” en, in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, Budapest Hungary: ACM, Aug. 2018, pp. 518–532, ISBN: 978-1-4503-5567-4. DOI: [10.1145/3230543.3230548](https://doi.org/10.1145/3230543.3230548).

- [15] A. Nötzli, J. Khan, A. Fingerhut, C. Barrett, and P. Athanas, “P4pktgen: Automated Test Case Generation for P4 Programs,” en, in *Proceedings of the Symposium on SDN Research*, Los Angeles CA USA: ACM, Mar. 2018, pp. 1–7, ISBN: 978-1-4503-5664-0. DOI: [10.1145/3185467.3185497](https://doi.org/10.1145/3185467.3185497).
- [16] H. Stubbe, “P4 compiler & interpreter: A survey,” *Future Internet (FI) and Innovative Internet Technologies and Mobile Communication (IITM)*, vol. 47, 2017.
- [17] A. Kheradmand and G. Rosu, *P4K: A Formal Semantics of P4 and Applications*, arXiv:1804.01468 [cs], Apr. 2018. DOI: [10.48550/arXiv.1804.01468](https://doi.org/10.48550/arXiv.1804.01468).
- [18] G. Roşu and T. F. Şerbănuţă, “An overview of the K semantic framework,” *The Journal of Logic and Algebraic Programming*, vol. 79, no. 6, pp. 397–434, 2010, Publisher: Elsevier. DOI: [10.1016/j.jlap.2010.03.012](https://doi.org/10.1016/j.jlap.2010.03.012).
- [19] The P4 Language Consortium, *P4_16 Language Specification - version 1.2.5*, Oct. 2024. [Online]. Available: <https://p4.org/wp-content/uploads/2024/10/P4-16-spec-v1.2.5.html> (visited on 04/24/2025).
- [20] R. Doenges, M. T. Arashloo, S. Bautista, et al., “Petr4: Formal foundations for p4 data planes,” en, *Proceedings of the ACM on Programming Languages*, vol. 5, no. POPL, pp. 1–32, Jan. 2021, ISSN: 2475-1421. DOI: [10.1145/3434322](https://doi.org/10.1145/3434322).
- [21] R. Doenges, T. Kappé, J. Sarracino, N. Foster, and G. Morrisett, “Leapfrog: Certified equivalence for protocol parsers,” en, in *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, San Diego CA USA: ACM, Jun. 2022, pp. 950–965, ISBN: 978-1-4503-9265-5. DOI: [10.1145/3519939.3523715](https://doi.org/10.1145/3519939.3523715).
- [22] D. Sangiorgi, *Introduction to bisimulation and coinduction*. Cambridge University Press, 2011, ISBN: 978-0-511-77711-0. DOI: [10.1017/CBO9780511777110](https://doi.org/10.1017/CBO9780511777110).
- [23] L. De Moura and N. Bjørner, “Z3: An Efficient SMT Solver,” in *Tools and Algorithms for the Construction and Analysis of Systems*, D. Hutchison, T. Kanade, J. Kittler, et al., Eds., vol. 4963, Series Title: Lecture Notes in Computer Science, Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 337–340, ISBN: 978-3-540-78799-0. DOI: [10.1007/978-3-540-78800-3_24](https://doi.org/10.1007/978-3-540-78800-3_24).
- [24] H. Barbosa, C. Barrett, M. Brain, et al., “Cvc5: A Versatile and Industrial-Strength SMT Solver,” en, in *Tools and Algorithms for the Construction and Analysis of Systems*, D. Fisman and G. Rosu, Eds., vol. 13243, Series Title: Lecture Notes in Computer Science, Cham: Springer International Publishing, 2022, pp. 415–442, ISBN: 978-3-030-99524-9. DOI: [10.1007/978-3-030-99524-9_24](https://doi.org/10.1007/978-3-030-99524-9_24).
- [25] M. Sipser, *Introduction to the Theory of Computation*, en. Cengage Learning, Jun. 2012, ISBN: 978-1-133-18779-0.
- [26] M. Gario and A. Micheli, “PySMT: A solver-agnostic library for fast prototyping of SMT-based algorithms,” in *SMT workshop*, vol. 2015, 2015.
- [27] G. Gibb, G. Varghese, M. Horowitz, and N. McKeown, “Design principles for packet parsers,” in *Architectures for Networking and Communications Systems*, IEEE, 2013, pp. 13–24. DOI: [10.1109/ANCS.2013.6665172](https://doi.org/10.1109/ANCS.2013.6665172).
- [28] A. Niemetz, M. Preiner, C. Wolf, and A. Biere, “Btor2 , BtorMC and Boolector 3.0,” in *Computer Aided Verification*, H. Chockler and G. Weissenbacher, Eds., vol. 10981, Series Title: Lecture Notes in Computer Science, Cham: Springer International Publishing, 2018, pp. 587–595, ISBN: 978-3-319-96144-6. DOI: [10.1007/978-3-319-96145-3_32](https://doi.org/10.1007/978-3-319-96145-3_32).
- [29] A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani, “The MathSAT5 SMT Solver,” en, in *Tools and Algorithms for the Construction and Analysis of Systems*, D. Hutchison, T. Kanade, J. Kittler, et al., Eds., vol. 7795, Series Title: Lecture Notes in Computer Science, Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 93–107, ISBN: 978-3-642-36741-0. DOI: [10.1007/978-3-642-36742-7_7](https://doi.org/10.1007/978-3-642-36742-7_7).

- [30] B. Dutertre, “Yices 2.2,” in *Computer Aided Verification*, D. Hutchison, T. Kanade, J. Kittler, *et al.*, Eds., vol. 8559, Series Title: Lecture Notes in Computer Science, Cham: Springer International Publishing, 2014, pp. 737–744, ISBN: 978-3-319-08866-2. DOI: [10.1007/978-3-319-08867-9_49](https://doi.org/10.1007/978-3-319-08867-9_49).
- [31] A. Niemetz and M. Preiner, “Bitwuzla,” en, in *Computer Aided Verification*, C. Enea and A. Lal, Eds., vol. 13965, Series Title: Lecture Notes in Computer Science, Cham: Springer Nature Switzerland, 2023, pp. 3–17, ISBN: 978-3-031-37702-0. DOI: [10.1007/978-3-031-37703-7_1](https://doi.org/10.1007/978-3-031-37703-7_1).

A. Supported P4 Grammar

Listing 3: The P4 grammar as accepted by OCTOPUS. It represents the subset of the P4 language that is both valid under P4C, as well as parseable by OCTOPUS for the purpose of equivalence checking. Grammar rules required by P4C for full program compilation but irrelevant to parser block definitions are omitted. For further details, see Section 6. The grammar is deliberately written to closely follow the YACC/Bison style used in the official P4₁₆ language specification [19], which may lead to some verbose rules. For simplicity, lexical tokens with a fixed string representation, such as HEADER, are written as string literals (in violet). Lexical constructs without a fixed string representation appear in full capitals (in blue). Both should be understood as lexer-level constructs handling related syntactical details.

```
1  /** PROGRAM **/  
2  
3  p4program  
4      : typeDeclarations parserDeclaration  
5      ;  
6  
7  /** TYPES **/  
8  
9  typeDeclarations  
10     : typeDeclaration  
11     | typeDeclarations typeDeclaration  
12     ;  
13  
14  typeDeclaration  
15     : headerTypeDeclaration  
16     | structTypeDeclaration  
17     ;  
18  
19  headerTypeDeclaration  
20     : "header" TYPE "{" typeFieldList "  
21     ;  
22  
23  structTypeDeclaration  
24     : "struct" TYPE "{" typeFieldList "  
25     ;  
26  
27  typeFieldList  
28     : typeField  
29     | typeFieldList typeField  
30     ;  
31  
32  typeField  
33     : typeRef IDENTIFIER ";"  
34     ;  
35  
36  typeRef  
37     : baseType  
38     | typeName  
39     ;  
40  
41  baseType  
42     : "bit" "<" INTEGER ">"  
43     ;  
44  
45  typeName  
46     : prefixedType  
47     ;  
48  
49  prefixedType
```

```

50      : TYPE
51      | "." TYPE
52      ;
53
54  /** PARSER **/
55
56  parserDeclaration
57      : "parser" IDENTIFIER "(" parameterList ")" "{" parserStates "}"
58      ;
59
60  parameterList
61      : parameter "," parameter
62      ;
63
64  parameter
65      : direction typeName IDENTIFIER
66      ;
67
68  direction
69      : /* empty */
70      | "in"
71      | "out"
72      | "inout"
73      ;
74
75  parserStates
76      : parserState
77      | parserStates parserState
78      ;
79
80  parserState
81      : "state" STATE "{" parserStatements transitionStatement "}"
82      ;
83
84  /** PARSER STATEMENTS **/
85
86  parserStatements
87      : parserStatement
88      | parserStatements parserStatement
89      ;
90
91  parserStatement
92      : assignmentOrMethodCallStatement
93      ;
94
95  assignmentOrMethodCallStatement
96      : lvalue "(" expression ")" ";"
97      | lvalue "=" expression ";"
98      ;
99
100  lvalue
101      : prefixedNonTypeName
102      | lvalue "." member
103
104  prefixedNonTypeName
105      : nonTypeName
106      | "." nonTypeName
107      ;
108
109  nonTypeName
110      : IDENTIFIER
111      ;
112

```

```

113 member
114     : IDENTIFIER
115     ;
116
117 /** TRANSITION STATEMENTS */
118
119 transitionStatement
120     : "transition" stateExpression
121     ;
122
123 stateExpression
124     : STATE ";"
125     | selectExpression
126     ;
127
128 selectExpression
129     : "select" "(" expressionList ")" "{" selectCaseList "}"
130     ;
131
132 expressionList
133     : expression
134     | expressionList "," expression
135     ;
136
137 selectCaseList
138     : selectCase
139     | selectCaseList selectCase
140     ;
141
142 selectCase
143     : keysetExpression ":" STATE ";"
144     ;
145
146 keysetExpression
147     : simpleKeysetExpression
148     | "(" simpleKeysetExpression "," simpleExpressionList ")"
149     ;
150
151 simpleExpressionList
152     : simpleKeysetExpression
153     | simpleExpressionList "," simpleKeysetExpression
154     ;
155
156 simpleKeysetExpression
157     : expression
158     | "default"
159     | "_"
160     ;
161
162 /** EXPRESSIONS */
163
164 expression
165     : INTEGER
166     | lvalue
167     | expression "[" INTEGER ":" INTEGER "]"
168     | expression "++" expression
169     | expression ">>" expression
170     | expression "&" expression
171     | "(" expression ")"
172     ;

```

B. OCTOPUS P4 Template

Listing 4: This template defines the minimal structure required for a P4 program to be parseable by OCTOPUS. It imports the standard library, declares a parser block, and instantiates a basic architecture. While all these components are necessary for the program to be compilable, only the parser block and its associated type definitions are relevant to the equivalence checking process. Consequently, all other elements have been abstracted into this template to isolate the parts of interest.

```
1  #include <core.p4>
2
3  // header and struct definitions (TO FILL IN)
4
5  parser Parser(packet_in pkt, out headers_t hdr) {
6
7      // parser states (TO FILL IN)
8
9  }
10
11  parser Parser_t(packet_in pkt, out headers_t hdr);
12  package Package(Parser_t p);
13
14  Package(Parser()) main;
```

C. Benchmarks in P4

This section lists the P4 programs on which OCTOPUS was evaluated. They are directly based on the benchmarks of [21], but have been converted to P4 for use with OCTOPUS. All programs use the template as defined in Listing 4; however, the headers and footers have been omitted for increased readability.

C.1. State Rearrangement

Listing 5: Parser block with combined states

```
1  header ip_t    { bit<64> data; }
2  header data_t  { bit<32> data; }
3
4  struct headers_t {
5      ip_t    ip;
6      data_t  pref;
7      data_t  suff;
8  }
9
10 parser Parser(packet_in pkt, out headers_t hdr) {
11     state start {
12         pkt.extract(hdr.ip);
13         pkt.extract(hdr.pref);
14         transition select(hdr.ip.data[23:20]) {
15             0: parse_suff;
16             1: accept;
17         }
18     }
19
20     state parse_suff {
21         pkt.extract(hdr.suff);
22         transition accept;
23     }
24 }
```

Listing 6: Parser block with separate states

```

1  header ip_t    { bit<64> data; }
2  header udp_t   { bit<32> data; }
3  header tcp_t   { bit<64> data; }
4
5  struct headers_t {
6      ip_t    ip;
7      udp_t   udp;
8      tcp_t   tcp;
9  }
10
11 parser Parser(packet_in pkt, out headers_t hdr) {
12     state start {
13         pkt.extract(hdr.ip);
14         transition select(hdr.ip.data[23:20]) {
15             0: parse_tcp;
16             1: parse_udp;
17         }
18     }
19
20     state parse_udp {
21         pkt.extract(hdr.udp);
22         transition accept;
23     }
24
25     state parse_tcp {
26         pkt.extract(hdr.tcp);
27         transition accept;
28     }
29 }

```

C.2. Variable-Length Formatting

Listing 7: The IPOptions parser block. For variable-length formatting 3, we consider the parser block as given. For variable-length formatting 2, we consider state `parse_1` as the start state instead. Both these versions have been shown in a single listing due to their equivalence otherwise.

```

1  header scratch8_t  { bit<8> data; }
2  header scratch16_t { bit<16> data; }
3  header scratch24_t { bit<24> data; }
4  header scratch32_t { bit<32> data; }
5  header scratch40_t { bit<40> data; }
6
7  header t0_t { bit<8> data; }
8  header l0_t { bit<8> data; }
9  header v0_t { bit<48> data; }
10
11 header t1_t { bit<8> data; }
12 header l1_t { bit<8> data; }
13 header v1_t { bit<48> data; }
14
15 header t2_t { bit<8> data; }
16 header l2_t { bit<8> data; }
17 header v2_t { bit<48> data; }
18
19 struct headers_t {
20     t0_t t0;
21     l0_t l0;
22     v0_t v0;
23     t1_t t1;
24     l1_t l1;

```

```

25     v1_t v1;
26     t2_t t2;
27     l2_t l2;
28     v2_t v2;
29
30     scratch8_t scratch8;
31     scratch16_t scratch16;
32     scratch24_t scratch24;
33     scratch32_t scratch32;
34     scratch40_t scratch40;
35 }
36
37 parser Parser(packet_in pkt, out headers_t hdr) {
38     state start {
39         pkt.extract(hdr.t0);
40         pkt.extract(hdr.l0);
41         transition select(hdr.t0.data, hdr.l0.data) {
42             (0, 0): accept;
43             (1, 0): accept;
44             (_, 1): parse_01;
45             (_, 2): parse_02;
46             (_, 3): parse_03;
47             (_, 4): parse_04;
48             (_, 5): parse_05;
49             (_, 6): parse_06;
50             default: reject;
51         }
52     }
53
54     state parse_01 {
55         pkt.extract(hdr.scratch8);
56         hdr.v0.data = hdr.scratch8.data ++ hdr.v0.data[47:8];
57         transition parse_1;
58     }
59
60     state parse_02 {
61         pkt.extract(hdr.scratch16);
62         hdr.v0.data = hdr.scratch16.data ++ hdr.v0.data[47:16];
63         transition parse_1;
64     }
65
66     state parse_03 {
67         pkt.extract(hdr.scratch24);
68         hdr.v0.data = hdr.scratch24.data ++ hdr.v0.data[47:24];
69         transition parse_1;
70     }
71
72     state parse_04 {
73         pkt.extract(hdr.scratch32);
74         hdr.v0.data = hdr.scratch32.data ++ hdr.v0.data[47:32];
75         transition parse_1;
76     }
77
78     state parse_05 {
79         pkt.extract(hdr.scratch40);
80         hdr.v0.data = hdr.scratch40.data ++ hdr.v0.data[47:40];
81         transition parse_1;
82     }
83
84     state parse_06 {
85         pkt.extract(hdr.v0);
86         transition parse_1;
87     }

```

```

88
89     state parse_1 {
90         pkt.extract(hdr.t1);
91         pkt.extract(hdr.l1);
92         transition select(hdr.t1.data, hdr.l1.data) {
93             (0, 0): accept;
94             (1, 0): accept;
95             (_, 1): parse_11;
96             (_, 2): parse_12;
97             (_, 3): parse_13;
98             (_, 4): parse_14;
99             (_, 5): parse_15;
100            (_, 6): parse_16;
101            default: reject;
102        }
103    }
104
105    state parse_11 {
106        pkt.extract(hdr.scratch8);
107        hdr.v1.data = hdr.scratch8.data ++ hdr.v1.data[47:8];
108        transition parse_2;
109    }
110
111    state parse_12 {
112        pkt.extract(hdr.scratch16);
113        hdr.v1.data = hdr.scratch16.data ++ hdr.v1.data[47:16];
114        transition parse_2;
115    }
116
117    state parse_13 {
118        pkt.extract(hdr.scratch24);
119        hdr.v1.data = hdr.scratch24.data ++ hdr.v1.data[47:24];
120        transition parse_2;
121    }
122
123    state parse_14 {
124        pkt.extract(hdr.scratch32);
125        hdr.v1.data = hdr.scratch32.data ++ hdr.v1.data[47:32];
126        transition parse_2;
127    }
128
129    state parse_15 {
130        pkt.extract(hdr.scratch40);
131        hdr.v1.data = hdr.scratch40.data ++ hdr.v1.data[47:40];
132        transition parse_2;
133    }
134
135    state parse_16 {
136        pkt.extract(hdr.v1);
137        transition parse_2;
138    }
139
140    state parse_2 {
141        pkt.extract(hdr.t2);
142        pkt.extract(hdr.l2);
143        transition select(hdr.t2.data, hdr.l2.data) {
144            (0, 0): accept;
145            (1, 0): accept;
146            (_, 1): parse_21;
147            (_, 2): parse_22;
148            (_, 3): parse_23;
149            (_, 4): parse_24;
150            (_, 5): parse_25;

```

```

151         (_, 6): parse_26;
152         default: reject;
153     }
154 }
155
156 state parse_21 {
157     pkt.extract(hdr.scratch8);
158     hdr.v2.data = hdr.scratch8.data ++ hdr.v2.data[47:8];
159     transition accept;
160 }
161
162 state parse_22 {
163     pkt.extract(hdr.scratch16);
164     hdr.v2.data = hdr.scratch16.data ++ hdr.v2.data[47:16];
165     transition accept;
166 }
167
168 state parse_23 {
169     pkt.extract(hdr.scratch24);
170     hdr.v2.data = hdr.scratch24.data ++ hdr.v2.data[47:24];
171     transition accept;
172 }
173
174 state parse_24 {
175     pkt.extract(hdr.scratch32);
176     hdr.v2.data = hdr.scratch32.data ++ hdr.v2.data[47:32];
177     transition accept;
178 }
179
180 state parse_25 {
181     pkt.extract(hdr.scratch40);
182     hdr.v2.data = hdr.scratch40.data ++ hdr.v2.data[47:40];
183     transition accept;
184 }
185
186 state parse_26 {
187     pkt.extract(hdr.v2);
188     transition accept;
189 }
190 }

```

Listing 8: The Timestamp parser block. For variable-length formatting 3, we consider the parser block as given. For variable-length formatting 2, we consider state parse_1 as the start state instead. Both these versions have been shown in a single listing due to their equivalence otherwise.

```

1 header scratch8_t { bit<8> data; }
2 header scratch16_t { bit<16> data; }
3 header scratch24_t { bit<24> data; }
4 header scratch32_t { bit<32> data; }
5 header scratch40_t { bit<40> data; }
6
7 header t0_t { bit<8> data; }
8 header l0_t { bit<8> data; }
9 header v0_t { bit<48> data; }
10
11 header t1_t { bit<8> data; }
12 header l1_t { bit<8> data; }
13 header v1_t { bit<48> data; }
14
15 header t2_t { bit<8> data; }
16 header l2_t { bit<8> data; }
17 header v2_t { bit<48> data; }
18

```



```

19 header pointer_t    { bit<8> data; }
20 header overflow_t   { bit<4> data; }
21 header flag_t       { bit<4> data; }
22 header timestamp_t  { bit<32> data; }
23
24 struct headers_t {
25     scratch8_t scratch8;
26     scratch16_t scratch16;
27     scratch24_t scratch24;
28     scratch32_t scratch32;
29     scratch40_t scratch40;
30
31     t0_t t0;
32     l0_t l0;
33     v0_t v0;
34     t1_t t1;
35     l1_t l1;
36     v1_t v1;
37     t2_t t2;
38     l2_t l2;
39     v2_t v2;
40
41     pointer_t pointer;
42     overflow_t overflow;
43     flag_t flag;
44     timestamp_t timestamp;
45 }
46
47 parser Parser(packet_in pkt, out headers_t hdr) {
48     state start {
49         pkt.extract(hdr.t0);
50         pkt.extract(hdr.l0);
51         transition select(hdr.t0.data, hdr.l0.data) {
52             (0x44, 6): parse_0s;
53             (0, 0): accept;
54             (1, 0): accept;
55             (_, 1): parse_01;
56             (_, 2): parse_02;
57             (_, 3): parse_03;
58             (_, 4): parse_04;
59             (_, 5): parse_05;
60             (_, 6): parse_06;
61             default: reject;
62         }
63     }
64
65     state parse_0s {
66         pkt.extract(hdr.pointer);
67         pkt.extract(hdr.overflow);
68         pkt.extract(hdr.flag);
69         pkt.extract(hdr.timestamp);
70         transition parse_1;
71     }
72
73     state parse_01 {
74         pkt.extract(hdr.scratch8);
75         hdr.v0.data = hdr.scratch8.data ++ hdr.v0.data[47:8];
76         transition parse_1;
77     }
78
79     state parse_02 {
80         pkt.extract(hdr.scratch16);
81         hdr.v0.data = hdr.scratch16.data ++ hdr.v0.data[47:16];

```

```

82         transition parse_1;
83     }
84
85     state parse_03 {
86         pkt.extract(hdr.scratch24);
87         hdr.v0.data = hdr.scratch24.data ++ hdr.v0.data[47:24];
88         transition parse_1;
89     }
90
91     state parse_04 {
92         pkt.extract(hdr.scratch32);
93         hdr.v0.data = hdr.scratch32.data ++ hdr.v0.data[47:32];
94         transition parse_1;
95     }
96
97     state parse_05 {
98         pkt.extract(hdr.scratch40);
99         hdr.v0.data = hdr.scratch40.data ++ hdr.v0.data[47:40];
100        transition parse_1;
101    }
102
103    state parse_06 {
104        pkt.extract(hdr.v0);
105        transition parse_1;
106    }
107
108    state parse_1 {
109        pkt.extract(hdr.t1);
110        pkt.extract(hdr.l1);
111        transition select(hdr.t1.data, hdr.l1.data) {
112            (0x44, 6): parse_1s;
113            (0, 0): accept;
114            (1, 0): accept;
115            (_, 1): parse_11;
116            (_, 2): parse_12;
117            (_, 3): parse_13;
118            (_, 4): parse_14;
119            (_, 5): parse_15;
120            (_, 6): parse_16;
121            default: reject;
122        }
123    }
124
125    state parse_1s {
126        pkt.extract(hdr.pointer);
127        pkt.extract(hdr.overflow);
128        pkt.extract(hdr.flag);
129        pkt.extract(hdr.timestamp);
130        transition parse_2;
131    }
132
133    state parse_11 {
134        pkt.extract(hdr.scratch8);
135        hdr.v1.data = hdr.scratch8.data ++ hdr.v1.data[47:8];
136        transition parse_2;
137    }
138
139    state parse_12 {
140        pkt.extract(hdr.scratch16);
141        hdr.v1.data = hdr.scratch16.data ++ hdr.v1.data[47:16];
142        transition parse_2;
143    }
144

```

```

145     state parse_13 {
146         pkt.extract(hdr.scratch24);
147         hdr.v1.data = hdr.scratch24.data ++ hdr.v1.data[47:24];
148         transition parse_2;
149     }
150
151     state parse_14 {
152         pkt.extract(hdr.scratch32);
153         hdr.v1.data = hdr.scratch32.data ++ hdr.v1.data[47:32];
154         transition parse_2;
155     }
156
157     state parse_15 {
158         pkt.extract(hdr.scratch40);
159         hdr.v1.data = hdr.scratch40.data ++ hdr.v1.data[47:40];
160         transition parse_2;
161     }
162
163     state parse_16 {
164         pkt.extract(hdr.v1);
165         transition parse_2;
166     }
167
168     state parse_2 {
169         pkt.extract(hdr.t2);
170         pkt.extract(hdr.l2);
171         transition select(hdr.t2.data, hdr.l2.data) {
172             (0x44, 6): parse_2s;
173             (0, 0): accept;
174             (1, 0): accept;
175             (_, 1): parse_21;
176             (_, 2): parse_22;
177             (_, 3): parse_23;
178             (_, 4): parse_24;
179             (_, 5): parse_25;
180             (_, 6): parse_26;
181             default: reject;
182         }
183     }
184
185     state parse_2s {
186         pkt.extract(hdr.pointer);
187         pkt.extract(hdr.overflow);
188         pkt.extract(hdr.flag);
189         pkt.extract(hdr.timestamp);
190         transition accept;
191     }
192
193     state parse_21 {
194         pkt.extract(hdr.scratch8);
195         hdr.v2.data = hdr.scratch8.data ++ hdr.v2.data[47:8];
196         transition accept;
197     }
198
199     state parse_22 {
200         pkt.extract(hdr.scratch16);
201         hdr.v2.data = hdr.scratch16.data ++ hdr.v2.data[47:16];
202         transition accept;
203     }
204
205     state parse_23 {
206         pkt.extract(hdr.scratch24);
207         hdr.v2.data = hdr.scratch24.data ++ hdr.v2.data[47:24];

```

```

208         transition accept;
209     }
210
211     state parse_24 {
212         pkt.extract(hdr.scratch32);
213         hdr.v2.data = hdr.scratch32.data ++ hdr.v2.data[47:32];
214         transition accept;
215     }
216
217     state parse_25 {
218         pkt.extract(hdr.scratch40);
219         hdr.v2.data = hdr.scratch40.data ++ hdr.v2.data[47:40];
220         transition accept;
221     }
222
223     state parse_26 {
224         pkt.extract(hdr.v2);
225         transition accept;
226     }
227 }

```

C.3. Header Initialisation

Listing 9: The parser block where header initialisation occurs in a single state (which is correct).

```

1  header eth_t  { bit<112> data; }
2  header vlan_t { bit<32> data; }
3  header ip_t   { bit<160> data; }
4  header udp_t  { bit<64> data; }
5
6  struct headers_t {
7      eth_t  eth;
8      vlan_t vlan;
9      ip_t   ip;
10     udp_t  udp;
11 }
12
13 parser Parser(packet_in pkt, out headers_t hdr) {
14     state start {
15         pkt.extract(hdr.eth);
16         transition select(hdr.eth.data[111:111]) {
17             0: default_vlan;
18             1: parse_vlan;
19         }
20     }
21
22     state default_vlan {
23         hdr.vlan.data = 0;
24         pkt.extract(hdr.ip);
25         transition parse_udp;
26     }
27
28     state parse_vlan {
29         pkt.extract(hdr.vlan);
30         transition parse_ip;
31     }
32
33     state parse_ip {
34         pkt.extract(hdr.ip);
35         transition parse_udp;
36     }
37 }

```

```

38     state parse_udp {
39         pkt.extract(hdr.udp);
40         transition select(hdr.vlan.data[31:28]) {
41             0b1111: reject;
42             default: accept;
43         }
44     }
45 }

```

Listing 10: The parser block where header initialisation occurs nowhere (which is incorrect).

```

1  header eth_t  { bit<112> data; }
2  header vlan_t { bit<32> data; }
3  header ip_t   { bit<160> data; }
4  header udp_t  { bit<64> data; }
5
6  struct headers_t {
7      eth_t  eth;
8      vlan_t vlan;
9      ip_t   ip;
10     udp_t  udp;
11 }
12
13 parser Parser(packet_in pkt, out headers_t hdr) {
14     state start {
15         pkt.extract(hdr.eth);
16         transition select(hdr.eth.data[111:111]) {
17             0: default_vlan;
18             1: parse_vlan;
19         }
20     }
21
22     state default_vlan {
23         // BUG: vlan not initialised
24         pkt.extract(hdr.ip);
25         transition parse_udp;
26     }
27
28     state parse_vlan {
29         pkt.extract(hdr.vlan);
30         transition parse_ip;
31     }
32
33     state parse_ip {
34         pkt.extract(hdr.ip);
35         transition parse_udp;
36     }
37
38     state parse_udp {
39         pkt.extract(hdr.udp);
40         transition select(hdr.vlan.data[31:28]) {
41             0b1111: reject;
42             default: accept;
43         }
44     }
45 }

```

C.4. Speculative Extraction

Listing 11: The plain version of the parser block

```

1  header mpls_t { bit<32> label; }

```

```

2  header udp_t    { bit<64> data; }
3
4  struct headers_t {
5      mpls_t mpls;
6      udp_t  udp;
7  }
8
9  parser Parser(packet_in pkt, out headers_t hdr) {
10     state start {
11         pkt.extract(hdr.mpls);
12         transition select(hdr.mpls.label[8:8]) {
13             0: start;
14             1: parse_udp;
15         }
16     }
17
18     state parse_udp {
19         pkt.extract(hdr.udp);
20         transition accept;
21     }
22 }

```

Listing 12: The vectorised version of the parser block

```

1  header udp_t    { bit<64> data; }
2  header tmp_t    { bit<32> field; }
3
4  struct headers_t {
5      udp_t  udp;
6      tmp_t  tmp;
7      tmp_t  old;
8      tmp_t  new;
9  }
10
11 parser Parser(packet_in pkt, out headers_t hdr) {
12     state start {
13         pkt.extract(hdr.old);
14         pkt.extract(hdr.new);
15         transition select(hdr.old.field[8:8], hdr.new.field[8:8]) {
16             (0, 0): start;
17             (0, 1): parse_udp;
18             (1, _): cleanup;
19         }
20     }
21
22     state parse_udp {
23         pkt.extract(hdr.udp);
24         transition accept;
25     }
26
27     state cleanup {
28         pkt.extract(hdr.tmp);
29         hdr.udp.data = hdr.new.field ++ hdr.tmp.field;
30         transition accept;
31     }
32 }

```

C.5. Relational Verification and External Filtering

Listing 13: The sloppy version of the parser block

```

1  header eth_t    { bit<112> data; }

```

```

2  header ipv4_t  { bit<128> data; }
3  header ipv6_t  { bit<288> data; }
4
5  struct headers_t {
6      eth_t    eth;
7      ipv4_t   ipv4;
8      ipv6_t   ipv6;
9  }
10
11 parser Parser(packet_in pkt, out headers_t hdr) {
12     state start {
13         pkt.extract(hdr.eth);
14         transition select(hdr.eth.data[15:0]) {
15             0x86dd: parse_ipv6;
16             default: parse_ipv4;
17         }
18     }
19
20     state parse_ipv4 {
21         pkt.extract(hdr.ipv4);
22         transition accept;
23     }
24
25     state parse_ipv6 {
26         pkt.extract(hdr.ipv6);
27         transition accept;
28     }
29 }

```

Listing 14: The strict version of the parser block

```

1  header eth_t   { bit<112> data; }
2  header ipv4_t  { bit<128> data; }
3  header ipv6_t  { bit<288> data; }
4
5  struct headers_t {
6      eth_t    eth;
7      ipv4_t   ipv4;
8      ipv6_t   ipv6;
9  }
10
11 parser Parser(packet_in pkt, out headers_t hdr) {
12     state start {
13         pkt.extract(hdr.eth);
14         transition select(hdr.eth.data[15:0]) {
15             0x86dd: parse_ipv6;
16             0x8600: parse_ipv4;
17             default: reject;
18         }
19     }
20
21     state parse_ipv4 {
22         pkt.extract(hdr.ipv4);
23         transition accept;
24     }
25
26     state parse_ipv6 {
27         pkt.extract(hdr.ipv6);
28         transition accept;
29     }
30 }

```

C.6. Edge and Translation Validation

Listing 15: The plain version of the parser block

```
1 header eth_t      { bit<112> data; }
2 header mpls_t     { bit<32> data; }
3 header eompls_t   { bit<28> data; }
4 header ipver_t    { bit<4> data; }
5 header ipv4_5_t   { bit<152> data; }
6 header ipv4_6_t   { bit<184> data; }
7 header ipv4_7_t   { bit<216> data; }
8 header ipv4_8_t   { bit<248> data; }
9 header ipv6_t     { bit<316> data; }
10
11 struct headers_t {
12     eth_t      eth0;
13     eth_t      eth1;
14     mpls_t      mpls0;
15     mpls_t      mpls1;
16     eompls_t    eompls;
17     ipver_t     ipver;
18     ipv4_5_t    ipv4_5;
19     ipv4_6_t    ipv4_6;
20     ipv4_7_t    ipv4_7;
21     ipv4_8_t    ipv4_8;
22     ipv6_t      ipv6;
23 }
24
25 parser Parser(packet_in pkt, out headers_t hdr) {
26     state start {
27         pkt.extract(hdr.eth0);
28         transition select(hdr.eth0.data[15:0]) {
29             0x8847: parse_mpls0;
30             0x8848: parse_mpls0;
31             0x0800: ignore_ipver4;
32             0x86dd: ignore_ipver6;
33             default: accept;
34         }
35     }
36
37     state parse_mpls0 {
38         pkt.extract(hdr.mpls0);
39         transition select(hdr.mpls0.data[8:8]) {
40             0: parse_mpls1;
41             1: parse_ipver;
42             default: reject;
43         }
44     }
45
46     state parse_mpls1 {
47         pkt.extract(hdr.mpls1);
48         transition select(hdr.mpls1.data[8:8]) {
49             1: parse_ipver;
50             default: reject;
51         }
52     }
53
54     state parse_ipver {
55         pkt.extract(hdr.ipver);
56         transition select(hdr.ipver.data) {
57             0: parse_eompls;
58             4: parse_ipv4;
59             6: parse_ipv6;
```



```

60         default: reject;
61     }
62 }
63
64 state ignore_ipver4 {
65     pkt.extract(hdr.ipver);
66     transition parse_ipv4;
67 }
68
69 state ignore_ipver6 {
70     pkt.extract(hdr.ipver);
71     transition parse_ipv6;
72 }
73
74 state parse_eompls {
75     pkt.extract(hdr.eompls);
76     transition parse_eth1;
77 }
78
79 state parse_eth1 {
80     pkt.extract(hdr.eth1);
81     transition accept;
82 }
83
84 state parse_ipv4 {
85     pkt.extract(hdr.ipver);
86     transition select(hdr.ipver.data) {
87         5: parse_ipv4_5;
88         6: parse_ipv4_6;
89         7: parse_ipv4_7;
90         8: parse_ipv4_8;
91         default: reject;
92     }
93 }
94
95 state parse_ipv4_5 {
96     pkt.extract(hdr.ipv4_5);
97     transition accept;
98 }
99
100 state parse_ipv4_6 {
101     pkt.extract(hdr.ipv4_6);
102     transition accept;
103 }
104
105 state parse_ipv4_7 {
106     pkt.extract(hdr.ipv4_7);
107     transition accept;
108 }
109
110 state parse_ipv4_8 {
111     pkt.extract(hdr.ipv4_8);
112     transition accept;
113 }
114
115 state parse_ipv6 {
116     pkt.extract(hdr.ipv6);
117     transition accept;
118 }
119 }

```

Listing 16: This is the optimised version of the parser block. Together with its plain version, it is used in the translation validation benchmark.

```

1  header buf_16_t    { bit<16> data; }
2  header buf_32_t    { bit<32> data; }
3  header buf_64_t    { bit<64> data; }
4  header buf_112_t   { bit<112> data; }
5  header buf_128_t   { bit<128> data; }
6  header buf_144_t   { bit<144> data; }
7  header buf_176_t   { bit<176> data; }
8  header buf_208_t   { bit<208> data; }
9  header buf_240_t   { bit<240> data; }
10 header buf_304_t   { bit<304> data; }
11 header buf_320_t   { bit<320> data; }
12
13 struct headers_t {
14     buf_16_t    b16;
15     buf_32_t    b32;
16     buf_64_t    b64;
17     buf_112_t   b112;
18     buf_128_t   b128;
19     buf_144_t   b144;
20     buf_176_t   b176;
21     buf_208_t   b208;
22     buf_240_t   b240;
23     buf_304_t   b304;
24     buf_320_t   b320;
25 }
26
27 parser Parser(packet_in pkt, out headers_t hdr) {
28     state start {
29         pkt.extract(hdr.b112);
30         transition select(hdr.b112.data[15:0]) {
31             0x0800: state_3;
32             0x86dd: state_0_suff_1;
33             0x8847: state_0_suff_2;
34             0x8848: state_0_suff_3;
35             default: accept;
36         }
37     }
38
39     state state_0_suff_1 {
40         pkt.extract(hdr.b320);
41         transition accept;
42     }
43
44     state state_0_suff_2 {
45         pkt.extract(hdr.b16);
46         transition state_4;
47     }
48
49     state state_0_suff_3 {
50         pkt.extract(hdr.b16);
51         transition state_4;
52     }
53
54     state state_1 {
55         pkt.extract(hdr.b16);
56         transition state_1_suff_0;
57     }
58
59     state state_1_suff_0 {
60         pkt.extract(hdr.b128);
61         transition accept;
62     }
63

```

```

64     state state_2 {
65         pkt.extract(hdr.b16);
66         transition state_2_suff_0;
67     }
68
69     state state_2_suff_0 {
70         pkt.extract(hdr.b304);
71         transition accept;
72     }
73
74     state state_3 {
75         pkt.extract(hdr.b16);
76         transition select(hdr.b16.data[11:8]) {
77             5: state_3_suff_0;
78             6: state_3_suff_1;
79             7: state_3_suff_2;
80             8: state_3_suff_3;
81             default: reject;
82         }
83     }
84
85     state state_3_suff_0 {
86         pkt.extract(hdr.b144);
87         transition accept;
88     }
89
90     state state_3_suff_1 {
91         pkt.extract(hdr.b176);
92         transition accept;
93     }
94
95     state state_3_suff_2 {
96         pkt.extract(hdr.b208);
97         transition accept;
98     }
99
100    state state_3_suff_3 {
101        pkt.extract(hdr.b240);
102        transition accept;
103    }
104
105    state state_4 {
106        pkt.extract(hdr.b16);
107        transition select(hdr.b16.data[8:8]) {
108            0: state_4_skip;
109            1: state_4_trailer;
110            default: reject;
111        }
112    }
113
114    state state_4_skip {
115        pkt.extract(hdr.b32);
116        transition select(hdr.b32.data[8:8]) {
117            1: state_4_trailer;
118            default: reject;
119        }
120    }
121
122    state state_4_trailer {
123        pkt.extract(hdr.b16);
124        transition select(hdr.b16.data[15:12], hdr.b16.data[11:8]) {
125            (0, _): state_1_suff_0;
126            (6, _): state_2_suff_0;

```

```

127         (4, 5):      state_3_suff_0;
128         (4, 6):      state_3_suff_1;
129         (4, 7):      state_3_suff_2;
130         (4, 8):      state_3_suff_3;
131         default:      reject;
132     }
133 }
134 }

```

C.7. Service Provider

```

1  header eth_t      { bit<112> data; }
2  header mpls_t     { bit<32> data; }
3  header ip_ver_t   { bit<4> data; }
4  header ihl_t      { bit<4> data; }
5  header ipv4_5_t   { bit<152> data; }
6  header ipv4_6_t   { bit<184> data; }
7  header ipv4_7_t   { bit<216> data; }
8  header ipv4_8_t   { bit<248> data; }
9  header ipv6_t     { bit<316> data; }
10
11  struct headers_t {
12      eth_t      eth;
13      mpls_t     mpls;
14      ip_ver_t   ip_ver;
15      ihl_t      ihl;
16      ipv4_5_t   ipv4_5;
17      ipv4_6_t   ipv4_6;
18      ipv4_7_t   ipv4_7;
19      ipv4_8_t   ipv4_8;
20      ipv6_t     ipv6;
21  }
22
23  parser Parser(packet_in pkt, out headers_t hdr) {
24      state start {
25          pkt.extract(hdr.eth);
26          transition select(hdr.eth.data[111:96]) {
27              0x8847: parse_mpls;
28              0x8848: parse_mpls;
29              0x0800: parse_ethv4;
30              0x86dd: parse_ethv6;
31              default: reject;
32          }
33      }
34
35      state parse_ethv4 {
36          pkt.extract(hdr.ip_ver);
37          transition parse_ipv4;
38      }
39
40      state parse_ethv6 {
41          pkt.extract(hdr.ip_ver);
42          transition parse_ipv6;
43      }
44
45      state parse_mpls {
46          pkt.extract(hdr.mpls);
47          transition select(hdr.mpls.data[24:24]) {
48              0: parse_mpls;
49              1: parse_ip_ver;
50              default: reject;
51          }

```

```

52     }
53
54     state parse_ip_ver {
55         pkt.extract(hdr.ip_ver);
56         transition select(hdr.ip_ver.data) {
57             4: parse_ipv4;
58             6: parse_ipv6;
59             default: reject;
60         }
61     }
62
63     state parse_ipv4 {
64         pkt.extract(hdr.ihl);
65         transition select(hdr.ihl.data) {
66             5: parse_ipv4_5;
67             6: parse_ipv4_6;
68             7: parse_ipv4_7;
69             8: parse_ipv4_8;
70             default: reject;
71         }
72     }
73
74     state parse_ipv4_5 {
75         pkt.extract(hdr.ipv4_5);
76         transition accept;
77     }
78
79     state parse_ipv4_6 {
80         pkt.extract(hdr.ipv4_6);
81         transition accept;
82     }
83
84     state parse_ipv4_7 {
85         pkt.extract(hdr.ipv4_7);
86         transition accept;
87     }
88
89     state parse_ipv4_8 {
90         pkt.extract(hdr.ipv4_8);
91         transition accept;
92     }
93
94     state parse_ipv6 {
95         pkt.extract(hdr.ipv6);
96         transition accept;
97     }
98 }

```

C.8. Datacenter

```

1  header eth_t      { bit<112> data; }
2  header vlan_t     { bit<160> data; }
3  header ipv4_t      { bit<160> data; }
4  header icmp_t      { bit<32> data; }
5  header tcp_t       { bit<160> data; }
6  header udp_t       { bit<160> data; }
7  header gre_t       { bit<32> data; }
8  header nvgre_t     { bit<32> data; }
9  header vxlan_t     { bit<64> data; }
10 header arp_t       { bit<64> data; }
11 header arp_ip_t    { bit<160> data; }
12

```

```

13 struct headers_t {
14     eth_t      eth0;
15     eth_t      eth1;
16     vlan_t     vlan0;
17     vlan_t     vlan1;
18     ipv4_t     ipv4;
19     icmp_t     icmp;
20     tcp_t      tcp;
21     udp_t      udp;
22     gre_t      gre0;
23     gre_t      gre1;
24     gre_t      gre2;
25     nvgre_t    nvgre;
26     vxlan_t    vxlan;
27     arp_t      arp;
28     arp_ip_t   arp_ip;
29 }
30
31 parser Parser(packet_in pkt, out headers_t hdr) {
32     state start {
33         pkt.extract(hdr.eth0);
34         transition select(hdr.eth0.data[15:0]) {
35             0x8100: parse_vlan0;
36             0x9100: parse_vlan0;
37             0x9200: parse_vlan0;
38             0x9300: parse_vlan0;
39             0x0800: parse_ipv4;
40             0x0806: parse_arp;
41             0x8035: parse_arp;
42             default: reject;
43         }
44     }
45
46     state parse_vlan0 {
47         pkt.extract(hdr.vlan0);
48         transition select(hdr.vlan0.data[15:0]) {
49             0x8100: parse_vlan1;
50             0x9100: parse_vlan1;
51             0x9200: parse_vlan1;
52             0x9300: parse_vlan1;
53             0x0800: parse_ipv4;
54             0x0806: parse_arp;
55             0x8035: parse_arp;
56             default: reject;
57         }
58     }
59
60     state parse_vlan1 {
61         pkt.extract(hdr.vlan1);
62         transition select(hdr.vlan1.data[15:0]) {
63             0x0800: parse_ipv4;
64             0x0806: parse_arp;
65             0x8035: parse_arp;
66             default: reject;
67         }
68     }
69
70     state parse_ipv4 {
71         pkt.extract(hdr.ipv4);
72         transition select(hdr.ipv4.data[87:80]) {
73             6: parse_tcp;
74             17: parse_udp;
75             47: parse_gre0;

```

```

76         default: accept;
77     }
78 }
79
80 state parse_tcp {
81     pkt.extract(hdr.tcp);
82     transition accept;
83 }
84
85 state parse_udp {
86     pkt.extract(hdr.udp);
87     transition select(hdr.udp.data[143:128]) {
88         0xFFFF: parse_vxlan;
89         default: accept;
90     }
91 }
92
93 state parse_icmp {
94     pkt.extract(hdr.icmp);
95     transition accept;
96 }
97
98 state parse_gre0 {
99     pkt.extract(hdr.gre0);
100    transition select(hdr.gre0.data[29:29], hdr.gre0.data[15:0]) {
101        (1, 0x6558): parse_nvgre;
102        (1, 0x6559): parse_gre1;
103        default: accept;
104    }
105 }
106
107 state parse_gre1 {
108     pkt.extract(hdr.gre1);
109     transition select(hdr.gre1.data[15:0]) {
110         0x16558: parse_nvgre;
111         0x16559: parse_gre2;
112         default: accept;
113     }
114 }
115
116 state parse_gre2 {
117     pkt.extract(hdr.gre2);
118     transition select(hdr.gre2.data[15:0]) {
119         0x16558: parse_nvgre;
120         0x16559: reject;
121         default: accept;
122     }
123 }
124
125 state parse_nvgre {
126     pkt.extract(hdr.nvgre);
127     transition parse_eth1;
128 }
129
130 state parse_vxlan {
131     pkt.extract(hdr.vxlan);
132     transition parse_eth1;
133 }
134
135 state parse_eth1 {
136     pkt.extract(hdr.eth1);
137     transition accept;
138 }

```

```

139
140     state parse_arp {
141         pkt.extract(hdr.arp);
142         transition select(hdr.arp.data[47:32]) {
143             0x0800: parse_arp_ip;
144             default: accept;
145         }
146     }
147
148     state parse_arp_ip {
149         pkt.extract(hdr.arp_ip);
150         transition accept;
151     }
152 }

```

C.9. Enterprise

```

1  header eth_t      { bit<112> data; }
2  header vlan_t     { bit<160> data; }
3  header ipv4_t     { bit<128> data; }
4  header ipv6_t     { bit<64>  data; }
5  header tcp_t      { bit<160> data; }
6  header udp_t      { bit<160> data; }
7  header icmp_t     { bit<32>  data; }
8  header icmp_v6_t  { bit<32>  data; }
9  header arp_t      { bit<64>  data; }
10 header arp_ip_t   { bit<64>  data; }
11
12 struct headers_t {
13     eth_t      eth;
14     vlan_t     vlan0;
15     vlan_t     vlan1;
16     ipv4_t     ipv4;
17     ipv6_t     ipv6;
18     tcp_t      tcp;
19     udp_t      udp;
20     icmp_t     icmp;
21     icmp_v6_t  icmp_v6;
22     arp_t      arp;
23     arp_ip_t   arp_ip;
24 }
25
26 parser Parser(packet_in pkt, out headers_t hdr) {
27     state start {
28         pkt.extract(hdr.eth);
29         transition select(hdr.eth.data[111:96]) {
30             0x8100: parse_vlan0;
31             0x9100: parse_vlan0;
32             0x9200: parse_vlan0;
33             0x9300: parse_vlan0;
34             0x0800: parse_ipv4;
35             0x86dd: parse_ipv6;
36             0x0806: parse_arp;
37             0x8035: parse_arp;
38             default: reject;
39         }
40     }
41
42     state parse_vlan0 {
43         pkt.extract(hdr.vlan0);
44         transition select(hdr.vlan0.data[159:144]) {
45             0x8100: parse_vlan1;

```



```

46     0x9100: parse_vlan1;
47     0x9200: parse_vlan1;
48     0x9300: parse_vlan1;
49     0x0800: parse_ipv4;
50     0x86dd: parse_ipv6;
51     0x0806: parse_arp;
52     0x8035: parse_arp;
53     default: reject;
54 }
55 }
56
57 state parse_vlan1 {
58     pkt.extract(hdr.vlan1);
59     transition select(hdr.vlan1.data[159:144]) {
60         0x0800: parse_ipv4;
61         0x86dd: parse_ipv6;
62         0x0806: parse_arp;
63         0x8035: parse_arp;
64         default: reject;
65     }
66 }
67
68 state parse_ipv4 {
69     pkt.extract(hdr.ipv4);
70     transition select(hdr.ipv4.data[79:72]) {
71         1: parse_icmp;
72         6: parse_tcp;
73         11: parse_udp;
74         default: accept;
75     }
76 }
77
78 state parse_ipv6 {
79     pkt.extract(hdr.ipv6);
80     transition select(hdr.ipv6.data[55:48]) {
81         1: parse_icmp_v6;
82         6: parse_tcp;
83         11: parse_udp;
84         default: accept;
85     }
86 }
87
88 state parse_tcp {
89     pkt.extract(hdr.tcp);
90     transition accept;
91 }
92
93 state parse_udp {
94     pkt.extract(hdr.udp);
95     transition accept;
96 }
97
98 state parse_icmp {
99     pkt.extract(hdr.icmp);
100    transition accept;
101 }
102
103 state parse_icmp_v6 {
104     pkt.extract(hdr.icmp_v6);
105     transition accept;
106 }
107
108 state parse_arp {

```

```
109     pkt.extract(hdr.arp);
110     transition select(hdr.arp.data[31:16]) {
111         0x0800: parse_arp_ip;
112         default: accept;
113     }
114 }
115
116 state parse_arp_ip {
117     pkt.extract(hdr.arp_ip);
118     transition accept;
119 }
120 }
```
