

KATch: A Fast Symbolic Verifier for NetKAT*

MARK MOELLER[†], Cornell University, USA

JULES JACOBS[†], Cornell University, USA

OLIVIER SAVARY BELANGER, Galois, USA

DAVID DARAIS, Galois, USA

COLE SCHLESINGER, Galois, USA

STEFFEN SMOLKA, Google, USA

NATE FOSTER, Cornell University, USA

ALEXANDRA SILVA, Cornell University, USA

We develop new data structures and algorithms for checking verification queries in NetKAT, a domain-specific language for specifying the behavior of network data planes. Our results extend the techniques obtained in prior work on symbolic automata and provide a framework for building efficient and scalable verification tools. We present KATch, an implementation of these ideas in Scala, featuring an extended set of NetKAT operators that are useful for expressing network-wide specifications, and a verification engine that constructs a bisimulation or generates a counter-example showing that none exists. We evaluate the performance of our implementation on real-world and synthetic benchmarks, verifying properties such as reachability and slice isolation, typically returning a result in well under a second, which is orders of magnitude faster than previous approaches. Our advancements underscore NetKAT's potential as a practical, declarative language for network specification and verification.

ACM Reference Format:

Mark Moeller, Jules Jacobs, Olivier Savary Belanger, David Darais, Cole Schlesinger, Steffen Smolka, Nate Foster, and Alexandra Silva. 2024. KATch: A Fast Symbolic Verifier for NetKAT. *Proc. ACM Program. Lang.* 8, PLDI, Article 224 (June 2024), 24 pages. <https://doi.org/10.1145/3656454>

1 INTRODUCTION

In the automata-theoretic approach to verification, programs and specifications are each encoded as automata, and verification tasks are reduced to standard questions in formal language theory—e.g., membership, emptiness, containment, etc. [41]. The approach became popular in the mid-1980s, driven by use of temporal logics and model checking in hardware verification, and it remains a

*This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA) and the Office of Naval Research (ONR). The views, opinions, and/or findings contained in this material are those of the authors and should not be interpreted as representing the official views or policies of the Department of the Navy, DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. DISTRIBUTION A. Approved for public release; distribution unlimited.

[†]Equal contribution.

Authors' addresses: Mark Moeller, Cornell University, Ithaca, NY, USA, mdm367@cornell.edu; Jules Jacobs, Cornell University, Ithaca, NY, USA, jj758@cornell.edu; Olivier Savary Belanger, Galois, Portland, OR, USA, olivier@galois.com; David Darais, Galois, Portland, OR, USA, darais@galois.com; Cole Schlesinger, Galois, Portland, OR, USA, coles@galois.com; Steffen Smolka, Google, Mountain View, CA, USA, smolkaj@google.com; Nate Foster, Cornell University, Ithaca, NY, USA, jnfoster@cs.cornell.edu; Alexandra Silva, Cornell University, Ithaca, NY, USA, alexandra.silva@cornell.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

2475-1421/2024/6-ART224

<https://doi.org/10.1145/3656454>

powerful tool today. In particular, automata provide natural models of phenomena like transitive closure, which arises often in programs but is impossible to express in pure first-order logic.

NetKAT, a domain-specific language for specifying and verifying the behavior of network data planes, is a modern success story for the automata-theoretic approach. NetKAT programs denote sets of traces (or “histories”) where a trace is the list of packets “seen” at each hop in the network. The NetKAT language specifies these traces using a regular expression-like syntax, as follows:

$$p, q ::= \perp \mid \top \mid f=v \mid f \neq v \mid f \leftarrow v \mid \text{dup} \mid p + q \mid p \cdot q \mid p^*$$

Unlike ordinary regular expressions, which are stateless, NetKAT programs manipulate state in packets. Accordingly, NetKAT’s atoms are not letters, but *actions* that either drop or forward the current packet (\perp or \top), modify a header field ($f \leftarrow v$), test a header field against a value ($f=v$, $f \neq v$), or append the current packet to the trace (dup). The regular expression operations copy and forward a packet to two different programs ($p + q$), sequence two programs ($p \cdot q$), or loop a program ($p^* \equiv \top + p + p \cdot p + p \cdot p \cdot p + \dots$).

A NetKAT program thus gives a declarative specification of the network’s global behavior in terms of sets of traces. Specifically, we can model the location of a packet in the network using a special header field (usually named sw for “switch”), that we can update ($\text{sw} \leftarrow \ell$) to logically move the packet from its current location to a new location ℓ . Given a declarative description of the intended behavior, the NetKAT compiler generates a set of local forwarding tables for individual switches that together realize the global behavior [19, 34].

Moreover, NetKAT not only plays a role as an implementation language, but also as a language for expressing verification queries, analogous to verification tools powered by SMT solvers [3, 30, 40]. In particular, as NetKAT includes a union operator (“+”), containment can be reduced to program equivalence—i.e., $p \sqsubseteq q$ if and only if $p + q \equiv q$. Hence, unlike other contemporary tools, which rely on bespoke encodings and algorithms for checking network properties like reachability, slice isolation, etc. [18, 25, 26, 43], NetKAT allows a wide range of practical verification questions to be answered using the same foundational mechanism, namely program equivalence. For example,

- *Network Reachability*: Are there any packets that can go from location 121 to 543 in the network specified by the NetKAT program net ? Formally: $(\text{sw} \leftarrow 121) \cdot \text{net}^* \cdot (\text{sw} = 543) \stackrel{?}{\equiv} \perp$
- *Slice Isolation*: Are slices net_1 and net_2 logically disjoint, even though their implementation on shared infrastructure uses the same devices? Formally: $\text{net}_1^* + \text{net}_2^* \stackrel{?}{\equiv} (\text{net}_1 + \text{net}_2)^*$

NetKAT’s dup primitive is a key construct that enables reducing a wide range of network-specific properties to program equivalence. Recall that dup appends the headers of the current packet to the trace. Hence, to verify properties involving the network’s internal behavior, we add intermediate packets to the trace using dup , making them relevant for program equivalence. Conversely, if we only care about the input-output behavior of the entire network, we can omit dup from the specification, so intermediate packets are not considered by the check for program equivalence.

The story so far is appealing, but there is a major fly in the ointment. To decide program equivalence, the automata-theoretic approach relies on the translation from NetKAT programs to automata. But standard NetKAT automata have an enormous space of potential transitions—i.e., the “alphabet” has a “character” for every possible packet. So using textbook algorithms for building automata and checking equivalence would clearly be impractical. In fact, NetKAT equivalence is PSPACE-complete [1]. Instead, what’s needed are symbolic techniques for encoding the state space and transition structure of NetKAT automata that avoid combinatorial blowup in the common case.

Prior work on symbolic automata provides a potential solution to the problems that arise when working with large (or even infinite) alphabets. The idea is to describe the transitions in the automata using logical formulas [15, 16] or Binary Decision Diagrams (BDDs) [33] rather than concrete

characters. Algorithms such as membership, containment, and equivalence can then be reformulated to work with these symbolic representations. Symbolic automata have been successfully applied to practical problems such as building input sanitizers for Unicode, which has tens of thousands of characters [23]. However, NetKAT’s richer semantics precludes the direct adoption of standard notions of symbolic automata. In particular, NetKAT’s transitions describe not only predicates but also transformations on the current packet. As Pous writes, it “seems feasible” to generalize his work to NetKAT, but “not straightforward” [33].

In this paper, we close this gap and develop symbolic techniques for checking program equivalence for NetKAT. In doing so, we address three key challenges:

Challenge 1: Expressive but Compact Symbolic Representations. The state space and transitions of NetKAT automata are very large due to the way the “alphabet” is built from the space of all possible packets. Orthogonally, these characters also encode packet transformations, as reflected in NetKAT’s packet-processing semantics. It is crucial that the symbolic representations for NetKAT programs be compact and admit efficient equivalence checks.

Challenge 2: Extended Logical Operators. In the tradition of regular expressions, NetKAT only includes sequential composition ($e_1 \cdot e_2$) and union operators ($e_1 + e_2$)—the latter computes the union of the sets of traces described by e_1 and e_2 . However, when verifying network-wide properties it is often useful to have operators for intersection ($e_1 \cap e_2$), difference ($e_1 \setminus e_2$), and symmetric difference ($e_1 \oplus e_2$). Having native support for these operators at the level of syntax and in equivalence checking algorithms allows the reduction of all verification queries to an emptiness check—e.g., $A \equiv B$ reduces to $A \oplus B \equiv \perp$, and $A \sqsubseteq B$ reduces to $A \setminus B \equiv \perp$.

Challenge 3: Support for Counter-Examples. When an equivalence check fails, it can be hard to understand the cause of the failure, and which changes might be needed to resolve the problem. It is therefore important to be able to construct *symbolic counter-examples* that precisely capture the input packets and traces that cause equivalence to fail.

In addressing the above challenges, we make the following technical contributions:

Contribution 1: Efficient symbolic representations (Section 3). We design a symbolic data structure called Symbolic Packet Programs (SPPs) for representing both sets of packets and transformations on packets in a symbolic manner. SPPs generalize classic BDDs and are asymptotically more efficient than the representations used in prior work on NetKAT. In particular, SPPs support efficient sequential composition ($e_1 \cdot e_2$) and can be made *canonical* which, with hash consing, allows us to check equivalence of SPPs in constant time.

Contribution 2: Symbolic Brzozowski derivatives (Section 4). We introduce a new form of symbolic Brzozowski derivative to produce automata for NetKAT programs. Our approach naturally and efficiently supports the extended “negative” logical operators mentioned above, in contrast to prior approaches that only support “positive” operators like union.

Contribution 3: Symbolic bisimilarity checking (Section 5). Building on the foundation provided by SPPs and deterministic NetKAT automata, we develop symbolic algorithms for checking bisimilarity. We present a symbolic version of the standard algorithm that searches in the forward direction through the state space of the automata under consideration. We also present a novel backward algorithm that computes symbolic counter-examples to equivalence—i.e., the precise set of input packets for which equivalence fails.

Contribution 4: KATch implementation (Section 6) and evaluation (Section 7). Finally, we present a new verification tool, called KATch, implemented in Scala. KATch implements symbolic

Syntax	Description	Semantics
$p, q ::=$		$\llbracket p \rrbracket(\alpha : \text{Pk}) : \mathcal{P}(\text{Pk}^+) =$
\perp	<i>False</i>	\emptyset
\top	<i>True</i>	$\{\alpha\}$
$f = v$	<i>Test equals</i>	if $\alpha_f = v$ then $\{\alpha\}$ else \emptyset
$f \neq v$	<i>Test not equals</i>	if $\alpha_f \neq v$ then $\{\alpha\}$ else \emptyset
$f \leftarrow v$	<i>Modification</i>	$\{\alpha[f \leftarrow v]\}$
dup	<i>Duplication</i>	$\{\alpha\alpha\}$
$p + q$	<i>Union</i>	$\llbracket p \rrbracket(\alpha) \cup \llbracket q \rrbracket(\alpha)$
$p \cdot q$	<i>Sequencing</i>	$\{\mathbf{ab} \mid \mathbf{a}\beta \in \llbracket p \rrbracket(\alpha), \mathbf{b} \in \llbracket q \rrbracket(\beta)\}$
p^*	<i>Iteration</i>	$\bigcup_{n \geq 0} \llbracket p^n \rrbracket$

Values $v ::= 0 \mid 1 \mid \dots \mid n$ Fields $f ::= f_1 \mid \dots \mid f_k$ Packets $\alpha ::= \{f_1 = v_1, \dots, f_k = v_k\}$

Fig. 1. NetKAT syntax and semantics.

bisimilarity checking, including an extended set of extended logical operators, as well as symbolic counter-examples. We evaluate the performance of KATch on a variety of real-world topologies, and show that it efficiently answers a variety of verification queries and scales to much larger networks than prior work. Due to its use of our efficient data structures, KATch is several orders of magnitude faster than prior implementations of NetKAT on these realistic examples. Moreover, on synthetic combinatorial benchmarks, which require symbolic representations to be tractable, KATch is faster than prior work by arbitrary large factors.

2 NETKAT EXPRESSIONS AND AUTOMATA

This section reviews basic definitions for NetKAT to set the stage for the new contributions presented in the subsequent sections. NetKAT is a language for specifying the packet-forwarding behavior of network data planes [1, 21, 34].¹ The syntax and semantics of NetKAT is presented in Figure 1, and is based on Kozen's Kleene Algebra with Tests (KAT) [28].

To a first approximation, NetKAT can be thought of as a simple, imperative language that operates over packets, where a *packet* is a finite record assigning values to fields. The basic primitives in NetKAT are packet tests ($f = v$, $f \neq v$) and packet modifications ($f \leftarrow v$). Program expressions are then compositionally built from tests and packet modifications, using union (+), sequencing (\cdot), and iteration (\star). Conditionals and loops can be encoded in the standard way:

$$\mathbf{if } b \mathbf{ then } p \mathbf{ else } q \triangleq b \cdot p + \neg b \cdot q \qquad \mathbf{while } b \mathbf{ do } p \triangleq (b \cdot p)^\star \cdot \neg b.$$

In a network, conditionals can be used to model the behavior of the forwarding tables on individual switches while iteration can be used to model the iterated processing performed by the network as a whole; the original paper on NetKAT provides further details [1]. Note that assignments and tests in NetKAT are always against constant values. This is a key restriction that makes equivalence decidable and aligns with the capabilities of data plane hardware. The dup primitive makes a copy of the current packet and appends it to the trace, which only ever grows as the packet goes through

¹Networks have a control plane, which computes paths through the topology using distributed routing protocols (or a software-defined networking controller), and a data plane, which implements paths using high-speed hardware and software pipelines. NetKAT models the behavior of the latter.

the network. This primitive is crucial for expressing network-wide properties, as it allows the semantics to “see” the intermediate packets processed on internal switches.

NetKAT’s formal semantics, given in [Figure 1](#), defines the action of a program on an input packet α , producing a set of output traces. The semantics for \top gives a single trace containing the single input packet α , whereas \perp produces no traces. Tests ($f = v$ and $f \neq v$) produce a singleton trace or no traces, depending on whether the test succeeds or fails. Modifications ($f \leftarrow v$) produce a singleton trace with the modified packet. Duplication (dup) produces a single trace with two copies of the input packet. Union ($p + q$) produces the union of the traces produced by p and q . Sequential composition ($p \cdot q$) produces the concatenation of the traces produced by p and q , where the last packet in output traces of p is used as the input to q . Finally, iteration (p^*) produces the union of the traces produced by p iterated zero or more times. Consider the following example:

$$(x=0 \cdot x \leftarrow 1 \cdot \text{dup} + x=1 \cdot x \leftarrow 0 \cdot \text{dup})^*$$

This program repeatedly flips the value of x between 0 and 1, and traces the packet at each step. On input packet $x=0$, it produces the following traces:

$$\{[x=0], [x=0, x=1], [x=0, x=1, x=0], [x=0, x=1, x=0, x=1], \dots\}$$

Consider what happens if we change the program to $(x=0 \cdot x \leftarrow 1 \cdot \text{dup} + x \leftarrow 0 \cdot \text{dup})^*$. Unlike the previous example, where the tests were disjoint, this program can generate multiple outputs for a given input packet, as the right branch of the union can always be taken, leading to sequences $x=0, x=0, \dots$. On the other hand, if the $x \leftarrow 1$ assignment is performed, then the left branch cannot be taken the next time, because the test $x=0$ will fail. Hence, the program produces traces with sequences of $x=0$ packets, with singular $x=1$ packets between them.

As these simple examples show, despite the fact that assignments and tests in NetKAT programs are always against constant values, the behavior of NetKAT programs can be complicated, particularly when conditionals, iteration, and multiple packet fields are involved.

Traditional presentations of NetKAT distinguish a syntactic test fragment that includes negation:

$$\begin{aligned} t, r &::= \perp \mid \top \mid f=v \mid t \vee r \mid t \wedge r \mid \neg t && \text{(Test fragment)} \\ p, q &::= t \mid f \leftarrow v \mid p + q \mid p \cdot q \mid \text{dup} \mid p^* && \text{(General expressions)} \end{aligned}$$

This distinction is necessary because negation $\neg p$ only makes sense on the test fragment. Our presentation simplifies and unifies the syntax and semantics, replacing general negation $\neg t$ with negated field tests $f \neq v$ only. This comes at no loss of generality, as one can use DeMorgan’s laws to push negations to the leaves of an expression. See [?] for the details of the translation.

Conditional dup, equivalence, and subsumption. An interesting feature of NetKAT is that dup need not appear the same number of times on all paths through the expression. Consider the program $(f \leftarrow v + f \leftarrow w + \text{dup})^*$, which either sets field f to v or to w , or logs the packet to the trace, and then repeats. Any number of writes can happen between two dups, and only the effect of the last write before the dup gets logged to the trace. This is allowed and fully supported by KATch.

The number of dups controls the type of equivalence that is checked. If we insert a dup after every packet field write, then we obtain trace equivalence. If we insert no dups whatsoever, we obtain input-output equivalence. By inserting dups in some places but not others, we can control the equivalence that is checked in a fine-grained manner.

In addition to checking equivalence $p \equiv q$, we may want to check inclusion $p \sqsubseteq q$. This can be expressed as $p + q \equiv q$ or equivalently as $p \setminus q \equiv \perp$, provided the language is extended with difference. The number of dups controls the type of inclusion that is checked: from trace inclusion to input-output inclusion, or something in between.

Input: A pair of NetKAT automata $(S, s_0, \delta, \epsilon)$, $(S', s'_0, \delta', \epsilon')$.
Returns: A boolean indicating whether the automata are equivalent.

```

 $W \leftarrow \{(s_0, s'_0, \alpha) \mid \alpha \in \text{Pk}\};$ 
while  $W$  changes do
  for  $(s, s', \alpha) \in W$  do
    if  $\epsilon(s, \alpha) \neq \epsilon'(s', \alpha)$  then return false;
     $W \leftarrow W \cup \{(\delta(s, \alpha)(\alpha'), \delta'(s', \alpha)(\alpha'), \alpha') \mid \alpha' \in \text{Pk}\}$ 
return true;

```

Fig. 2. Bisimulation algorithm à la Hopcroft & Karp [24]

2.1 NetKAT Automata

NetKAT's semantics induce an equivalence $(p \equiv q) \triangleq (\llbracket p \rrbracket = \llbracket q \rrbracket)$ on syntactic programs. This equivalence is decidable, and can be computed by converting programs to automata and checking for automata equivalence. To convert NetKAT programs to automata, the standard approach is to use Antimirov derivatives. We do not give the details here, but the interested reader can find them in [21], and for our symbolic automata, in Section 4. We continue with a brief overview of NetKAT automata and how to check their equivalence. A NetKAT automaton $(S, s_0, \langle \epsilon, \delta \rangle)$ consists of a set of states S , a start state s_0 , and a pair of functions that depend on an input packet:

$$\epsilon : S \times \text{Pk} \rightarrow 2^{\text{Pk}} \qquad \delta : S \times \text{Pk} \rightarrow S^{\text{Pk}}$$

Intuitively, the observation function ϵ is analogous to the notion of a final state, and models the output packets produced from an input packet at a given state. The transition function δ models the state transitions that can occur when processing an input packet at a given state. Because transitions can modify the input packet, the transition structure $\delta(s, \alpha)$ is itself a function of the modified packet, and tells us to transition to state $\delta(s, \alpha)(\alpha')$ when input packet α is modified to α' .

In terms of the semantics, a transition in the automaton corresponds to executing to the next dup in a NetKAT program, and appends the current packet to the trace. It should also be noted that states in the automaton do not necessarily correspond to network devices; because the location of the packet is modeled as just a field in the packet, there can be states in the automaton that handle packets from multiple devices, and the packets of a device can be handled by multiple states.

Carry-on packets. Unlike traditional regular expressions and automata, NetKAT is stateful, and the output packet of a transition is carried on to the next state. This makes NetKAT fundamentally different from traditional regular expressions and automata and results in challenges in the semantics (which is not a straightforward trace or language semantics) and in designing equivalence procedures, as they have to account for the carry-on packet. Technically it would be possible to fold the packet into the state of the automaton and approach the semantics and problem of checking equivalence using classical automata. Because the number of possible packets is exponential in the number of fields and values, however, such a conversion results in enormous state blowup and is therefore impractical.

2.2 Bisimulation of NetKAT Automata

NetKAT automata can be used to decide $p \equiv q$. Intuitively, while traces may be infinite, transitions only depend on the current packet, which has a finite number of distinct possible values.

Given two automata $(S, s_0, \langle \epsilon, \delta \rangle)$ and $(S', s'_0, \langle \epsilon', \delta' \rangle)$, we check their equivalence by considering all possible input packets separately. We run the two automata in parallel, starting from an input




Fig. 3. Examples of SPs, where $p \triangleq (a=3 \cdot b=4 + b \neq 5 \cdot c=5)$ and $q \triangleq (b=3 \cdot c=4 + a=5 \cdot c \neq 5)$.

packet α at their respective start states. We first check that the immediate outputs $\epsilon(s_0, \alpha)$ and $\epsilon'(s'_0, \alpha)$ are the same. If so, we check that the transitions $\delta(s_0, \alpha)$ and $\delta'(s'_0, \alpha)$ are equivalent, by recursively checking the equivalence of the states $\delta(s_0, \alpha)(\alpha')$ and $\delta'(s'_0, \alpha)(\alpha')$ for all α' .

We can implement this strategy using a work list algorithm. The work list contains triples (s, s', α) , where s and s' are states in the two automata, and α is an input packet. Initially, the work list contains (s_0, s'_0, α) for all packets α . We then repeatedly take triples (s, s', α) from the work list, and check that $\epsilon(s, \alpha) = \epsilon'(s', \alpha)$ and add $(\delta(s, \alpha)(\alpha'), \delta'(s', \alpha)(\alpha'), \alpha')$ to the work list for all α' . If at any point we find that $\epsilon(s, \alpha) \neq \epsilon'(s', \alpha)$, then the automata are not equivalent. If the work list stabilizes, the automata are equivalent. This algorithm is shown in Figure 2.

Of course, the issue with this algorithm is that the space of packets is huge. The rest of this paper develops techniques for working with symbolic representations of packets and automata, allowing us to represent and manipulate large sets of packets and to efficiently check equivalence of automata without explicitly enumerating the space of packets.

3 SYMBOLIC NETKAT REPRESENTATIONS

In this section, we develop symbolic techniques to trace an entire set of packets at once, vastly reducing the number of iterations required to compute a bisimulation in practice. These techniques will allow us to revise the approach of Figure 2 with a much more efficient version in Section 5.

First, we introduce a representation for symbolic packets, which represent sets of packets, or equivalently, the fragment of NetKAT where all atoms are tests. This representation is essentially a natural n-ary variant of binary decision diagrams (BDDs) [12]. Second, we introduce Symbolic Packet Programs (SPPs), a new representation for symbolic transitions in NetKAT automata, representing the dup-free fragment of NetKAT—i.e., the fragment in which all atoms are tests or assignments. The primary challenge is to design this representation so that it is efficiently closed under the NetKAT operations. Furthermore, we need to be able to efficiently compute the transition of a symbolic packet over a SPP, both forward and backward.

3.1 Symbolic Packets

We begin by choosing a representation for symbolic packets. A symbolic packet $p \subseteq \text{Pk}$ is a set of concrete packets, represented compactly as a decision diagram. Syntactically, symbolic packets are

NetKAT expressions with atoms restricted to tests, represented in the following canonical form:

$$p \in \text{SP} ::= \perp \mid \top \mid \underbrace{\text{SP}(f, \{\dots, v_i \mapsto q_i, \dots\}, q)}_{\equiv \sum_i f=v_i \cdot q_i + (\prod_i f \neq v_i) \cdot q}$$

That is, symbolic packets form an n-ary tree, where each child q_i is labeled with a test of the current field f , and the default case q is labeled with the negation of the other tests. Hence, a given concrete packet has a unique path through the tree.

The following conditions need to be satisfied for a symbolic packet to be in canonical form, inspired by the analogous properties of BDDs:

Reduced If a child q_i is equal to the default case q , it is removed. If only the default case remains, the symbolic packet is reduced to the default case itself.

Ordered A path down the tree always follows the same order of fields, and the children q_i are ordered by the value v_i .

These conditions ensure the representation of a symbolic packet is unique, in the sense that two symbolic packets are semantically equal if and only if they are syntactically equal.

Representation. As with BDDs, we share nodes in the tree, making the representation of a symbolic packet a directed acyclic graph, as shown in Figure 3. Vertices are labeled with the packet field they test. Solid arrows labeled with a number encode the test value, and dashed arrows represent a default case. The sinks of the graph are labeled with \top or \perp , indicating membership in the set. On the left, we have a symbolic packet representing all packets where $a = 3$ and $b = 4$, or $b \neq 5$ and $c = 5$. In the middle, we have a symbolic packet representing all packets where $a = 3$ and $c = 4$, or $a = 5$ and $c \neq 5$. On the right, we have the union of these symbolic packets.

Operations. Symbolic packets are closed under the NetKAT operations:

$$\hat{+}, \hat{\wedge}, \hat{\vee}, \hat{\ominus}, \hat{\ominus} : \text{SP} \times \text{SP} \rightarrow \text{SP} \quad \hat{\star} : \text{SP} \rightarrow \text{SP} \quad f=v, f \neq v : \text{SP}$$

As defined in Figure 4, operations on symbolic packets are computed by traversing the two trees in parallel, recursively taking the operation of the children, making sure to maintain canonical form using the smart constructor. Repetition $p^{\star} \triangleq \top$ is trivial for symbolic packets. All other operations are defined in [?].

Specifications. Each of these operations is uniquely specified by two correctness conditions:

- (1) They semantically match their counterpart: $\llbracket p \hat{+} q \rrbracket = \llbracket p + q \rrbracket$ for all $p, q \in \text{SP}$.
- (2) They maintain canonical form: $p \hat{+} q$ is reduced and ordered if p, q are.

For further formal treatment of SPs, see [?].

3.2 Symbolic Transitions

We now introduce a representation for symbolic transitions in NetKAT automata. Whereas symbolic packets represent the fragment of NetKAT where all atoms are tests, symbolic transitions correspond to the larger dup-free fragment, where all atoms are tests or assignments. This introduces additional challenges for a canonical representation, as well as for the operations. For instance, sequential composition is no longer equivalent to intersection, and the star operator is no longer trivial.

$$p \in \text{SPP} ::= \perp \mid \top \mid \underbrace{\text{SPP}(f, \{\dots, v_i \mapsto \{\dots, w_{ij} \mapsto q_{ij}, \dots\}, \dots\}, \{\dots, w_i \mapsto q_i, \dots\}, q)}_{\equiv \sum_i f=v_i \cdot \sum_j f \leftarrow w_{ij} \cdot q_{ij} + (\prod_i f \neq v_i) \cdot (\sum_i f \leftarrow w_i + (\prod_i f \neq w_i) \cdot q)}$$

Like SPs, SPPs have two base cases, \top and \perp . Also like SPs, SPPs test a field f of the input packet against a series of values v_0, \dots, v_n , with a default case $f \neq v_0 \cdots f \neq v_n$. However, instead of

Primitive tests:

$$(f=v) \triangleq \text{SP}(f, \{v \mapsto \top\}, \perp) \quad (f \neq v) \triangleq \text{SP}(f, \{v \mapsto \perp\}, \top)$$

Operations, base cases:

$$\begin{aligned} \top \hat{+} \top &\triangleq \top & \top \hat{+} \perp &\triangleq \top & \top \hat{\wedge} \top &\triangleq \top & \top \hat{\oplus} \top &\triangleq \perp & \top \hat{\wedge} \perp &\triangleq \perp \\ \top \hat{+} \perp &\triangleq \top & \top \hat{+} \perp &\triangleq \perp & \top \hat{\wedge} \perp &\triangleq \perp & \top \hat{\oplus} \perp &\triangleq \top & \top \hat{\wedge} \perp &\triangleq \top \\ \perp \hat{+} \top &\triangleq \top & \perp \hat{+} \perp &\triangleq \perp & \perp \hat{\wedge} \top &\triangleq \perp & \perp \hat{\oplus} \top &\triangleq \top & \perp \hat{\wedge} \perp &\triangleq \perp \\ \perp \hat{+} \perp &\triangleq \perp & \perp \hat{+} \perp &\triangleq \perp & \perp \hat{\wedge} \perp &\triangleq \perp & \perp \hat{\oplus} \perp &\triangleq \perp & \perp \hat{\wedge} \perp &\triangleq \perp \end{aligned}$$

Operations, inductive case:

$$\begin{aligned} \text{SP}(f, b_p, d_p) \hat{+} \text{SP}(f, b_q, d_q) &\triangleq \text{sp}(f, b', d_p \hat{+} d_q) \\ \text{where } b' &= \{v \mapsto b_p(v; d_p) \hat{+} b_q(v; d_q) \mid v \in \text{dom}(b_p \cup b_q)\} \end{aligned}$$

Expansion:

$$p \equiv \text{SP}(f, \emptyset, p) \quad \text{if } p \in \{\top, \perp, \text{SP}(f', b, d)\} \text{ where } f \sqsubset f'$$

Smart constructor:

$$\begin{aligned} \text{sp}(f, b, d) &\triangleq \text{if } b' = \emptyset \text{ then } d \text{ else } \text{SP}(f, b', d) \\ \text{where } b' &\triangleq \{v \mapsto p \mid v \mapsto p \in b, p \neq d\} \end{aligned}$$

Fig. 4. Definition of the SP operations. The inductive case is identical for all operations (indicated by $\hat{+}$), and applies when both SPs test the same field. Expansion inserts a trivial SP node to reduce the remaining cases to the inductive case. The notation $b(v; d)$ means the child $b(v)$ if $v \in \text{dom}(b)$, or the default case d otherwise.

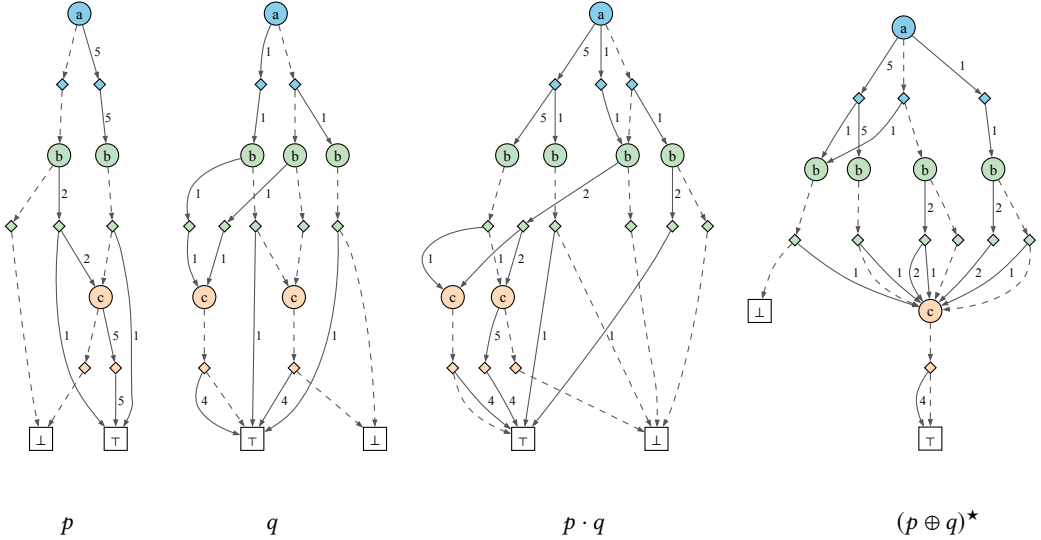


Fig. 5. Examples of SPPs, where $p \triangleq (a=5 + b=2) \cdot (b \leftarrow 1 + c=5)$, and $q \triangleq (b=1 + c \leftarrow 4 + a \leftarrow 1 \cdot b \leftarrow 1)$.

continuing recursively after the test, SPPs non-deterministically assign a value w_{ij} to the field f of the output packet, and continue recursively with the corresponding child q_{ij} . This way, SPPs can output more than one packet for a given input packet, and can also output packets with different values for the same field. The default case is further split into two cases, one where the field f is non-deterministically assigned a new value (like in the other cases), and an identity case where the field f keeps the same value as the input packet. However, the packet for the latter case is not

produced when the input packet's field f had a value that could also be produced by the explicit assignments in the default case. This ensures that for a given input packet, a given output packet is produced by a unique path through the SPP.

The following conditions need to be satisfied for a SPP to be in canonical form:

Reduced If a child q_{ij} is equal to \perp , it is removed (with the exception of the default-identity case, which is always kept). If one of the default-assignment cases for a value w is \perp , it is removed, but to keep the behavior equivalent, an additional test for the same value w is added, with an empty sequence of assignments (if a test for the value w was already present, nothing is added). Further, each of the non-default cases is analyzed, and if we determine that it behaves semantically like the default case for that input value, then it is removed. If only the default case remains, the SPP is reduced to the default case itself.

Ordered A path down the tree always follows the same order of fields, and both the tests and the assignments are ordered by the value v_i or w_{ij} .

These conditions ensure the representation of a SPP is unique, in the sense that two SPPs are semantically equal if and only if they are syntactically equal.

Representation. Like symbolic packets, we represent SPPs as a directed acyclic graph, with duplicate nodes shared. Examples of SPPs are shown in Figure 5. Vertices are labeled with the packet field that they test. Solid arrows labeled with a number represent the tests, and dashed arrows represent the default case. Each test is followed by a non-deterministic assignment of a new value to the field, indicated by the small diamonds. In the default case, the non-deterministic assignment also has an identity case (keeping the value of the field unchanged), indicated by a dashed arrow emanating from the diamond.

Consider the action of the first SPP in Figure 5 on the concrete packet $a=5 \cdot b=3 \cdot c=5$. The first test $a=5$ succeeds, and the field a is assigned the value 5. The b field is then tested, but the b node only has a default case. The default case does have a non-deterministic assignment, which can set the b field to the new value 1, or it can keep the old value 3. In case the new value of b is 1, the packet is immediately accepted by the \top node, so the packet $a=5 \cdot b=1 \cdot c=5$ is produced. In case the old value of b is kept, the c field is tested, and in our case the value of the c field is 5. In this case, the value of the field is unchanged, and the packet $a=5 \cdot b=3 \cdot c=5$ is produced. In summary, for input packet $a=5 \cdot b=3 \cdot c=5$, the SPP produces packets $a=5 \cdot b=1 \cdot c=5$ and $a=5 \cdot b=3 \cdot c=5$.

When we sequentially compose the two SPPs on the left, we get the third SPP shown. In other words, if we take a concrete packet, and first apply the first SPP, and then apply the second SPP to all of the resulting packets, then we get the same result as if we apply the SPP on the right directly to the concrete packet. Sequential composition is a relatively complex operation, but algorithmically it is a key strength of our representation and reduced/ordered invariant. To understand why, consider taking a concrete packet, and applying the first SPP to it. Once we have applied the a -layer of the SPP to the packet, we already know what the value of the a field in the output packet will be. We can therefore immediately continue with the a layer of the second SPP, without having to consider the other layers of the first SPP. This is in contrast to a naive algorithm, which would have to consider the entire first SPP before being able to apply the second SPP. This property makes it possible to compute the sequential composition of two SPPs efficiently in practice, and is a key contributor to the performance and scalability of our system.

Operations. SPPs are closed not just under sequential composition, but under all of the NetKAT operations. These operations are largely mechanical, but more complex than for SPs, as we need to take assignments into account while respecting the conditions that ensure uniqueness. In particular, sequential composition is no longer equivalent to intersection, as the assignments in the first SPP

Primitive tests and mutation:

$$(f = v) \triangleq \text{SPP}(f, \{v \mapsto \{v \mapsto \top\}\}, \emptyset, \perp) \quad (f \leftarrow v) \triangleq \text{SPP}(f, \emptyset, \{v \mapsto \top\}, \perp)$$

$$(f \neq v) \triangleq \text{SPP}(f, \{v \mapsto \{v \mapsto \perp\}\}, \emptyset, \top)$$

Operations, base cases:

Identical to those in Figure 4.

Operations, inductive cases:

$$\text{SPP}(f, b_p, m_p, d_p) \hat{\pm} \text{SPP}(f, b_q, m_q, d_q) \triangleq \text{spp}(f, b', m_p \bar{\pm} m_q, d_p \hat{\pm} d_q)$$

where $b' = \{v \mapsto p(v) \bar{\pm} q(v) \mid v \in \text{dom}(b_p \cup b_q \cup m_p \cup m_q)\}$

$$\text{SPP}(f, b_p, m_p, d_p) \hat{\wedge} \text{SPP}(f, b_q, m_q, d_q) \triangleq \text{spp}(f, b', m_A \bar{\wedge} m_B, d_p \hat{\wedge} d_q)$$

where $b' = \{v \mapsto \sum_{v' \mapsto p'(v)} \{w' \mapsto p' \hat{\wedge} q' \mid w' \mapsto q' \in q(v')\} \mid v \in \text{dom}(b_p \cup b_q \cup m_p \cup m_q \cup m_A)\}$

$$m_A = \sum_{v' \mapsto p' \in m_p} \{w' \mapsto p' \hat{\wedge} q' \mid w' \mapsto q' \in q(v')\}, \quad m_B = \{w' \mapsto d_p \hat{\wedge} q' \mid w' \mapsto q' \in m_q\}$$

$$m_1 \bar{\pm} m_2 \triangleq \{v \mapsto m_1(v; \perp) \hat{\pm} m_2(v; \perp) \mid v \in \text{dom}(m_1 \cup m_2)\}, \quad \bar{\Sigma} \triangleq n\text{-ary sum w.r.t. } \bar{\wedge}$$

$$p(v) \triangleq b_p(v; \text{if } v \in m_p \vee d_p = \perp \text{ then } m_p \text{ else } m_p \cup \{v \mapsto d_p\}) \quad (\text{similarly for } q(v))$$

Expansion:

$$p \equiv \text{SPP}(f, \emptyset, \emptyset, p) \quad \text{if} \quad p \in \{\top, \perp, \text{SPP}(f', b, m, d)\} \text{ where } f \sqsubset f'$$

Repetition:

$$p^* \triangleq \hat{\Sigma}_i p^i = \top \hat{\wedge} p \hat{\wedge} p \hat{\wedge} p \hat{\wedge} p \hat{\wedge} p \hat{\wedge} p \hat{\wedge} \dots \quad (\text{sums in SPP converge in a finite number of steps})$$

Fig. 6. Definition of the SPP operations. The inductive case is identical for all operations (indicated by $\hat{\pm}$), except $\hat{\wedge}$, which is given separately. Expansion inserts a trivial SPP node to reduce the remaining cases to the inductive case. The notation $b(v; d)$ means the child $b(v)$ if $v \in \text{dom}(b)$, or by default d otherwise.

clearly affect the tests in the second SPP. Furthermore, the star operator is no longer trivial, as the assignments in the SPP can affect the tests in the SPP itself, and a fixed point needs to be taken. We have the following operations on SPPs:

$$\hat{\wedge}, \hat{\wedge}, \hat{\wedge}, \hat{\wedge}, \hat{\wedge} : \text{SPP} \times \text{SPP} \rightarrow \text{SPP} \quad \hat{\star} : \text{SPP} \rightarrow \text{SPP} \quad f = v, f \neq v, f \leftarrow v : \text{SPP}$$

Push and pull. In addition to these operations for combining SPPs, we also have the following operations, which “push” and “pull” a symbolic packet through a SPP:

$$\text{push} : \text{SP} \times \text{SPP} \rightarrow \text{SP} \quad \text{pull} : \text{SPP} \times \text{SP} \rightarrow \text{SP}$$

The push operation computes the effect of a SPP on a symbolic packet, giving a symbolic packet as a result. The new symbolic packet contains all of the packets that are produced by the SPP when applied to the packets in the input symbolic packet.

The pull operation simulates the effect of a SPP in reverse. This operation is used when computing the backward transition of a symbolic packet over a symbolic transition, for counter-example generation. The pull operation answers this question: given a set of output packets (represented symbolically), what are the possible input packets (also represented symbolically) that could have produced them? In other words, a concrete packet α is an element of $\text{pull}(p, q)$ if and only if running the SPP p on α produces an output packet in q . In particular, it is okay if running the SPP on α produces multiple output packets, as long as at least one of them is in q .

Specifications. Each of these operations is uniquely specified by two correctness conditions:

- (1) They semantically match their counterpart (e.g. $\llbracket p \hat{\wedge} q \rrbracket \equiv \llbracket p + q \rrbracket$ for all $p, q \in \text{SPP}$)
- (2) They maintain the reduced & ordered invariant (e.g., $p \hat{\wedge} q$ is reduced & ordered if p, q are)

For push and pull, the correctness condition is as follows:

$$\beta \in \llbracket \text{push } p \ s \rrbracket \iff \exists \alpha \in \llbracket p \rrbracket. \beta \in \llbracket s \rrbracket(\alpha) \quad \text{and} \quad \beta \in \llbracket \text{pull } s \ p \rrbracket \iff \exists \alpha \in \llbracket p \rrbracket. \alpha \in \llbracket s \rrbracket(\beta)$$

For further details, see [?].

4 SYMBOLIC NETKAT AUTOMATA VIA BRZOWSKI DERIVATIVES

With the symbolic packet and symbolic transition representations in place, we can now define symbolic NetKAT automata. The construction of NetKAT automata shares some similarities with the construction of automata for regular expressions. In prior work [21, 34], NetKAT automata were constructed via Antimirov derivatives [2], which can be extended from regular expressions to NetKAT. The Antimirov derivative constructs a non-deterministic automaton, and as such, is well suited for handling NetKAT’s union operator by inserting transitions for the two sub-terms. However, the Antimirov derivative is not well suited for our extended set of logical operators, as we cannot simply insert (non-deterministic) transitions for the operands of intersection, difference, and symmetric difference. Instead, we extend the Brzowski derivative [13] to NetKAT, which constructs a deterministic automaton directly, and is better suited for the logical operators.

In the rest of this section, we will first describe what symbolic NetKAT automata are, and then describe how to construct them via the Brzowski derivative.

4.1 Symbolic NetKAT Automata

An automaton for a regular expression consists of a set of states, and transitions labeled with symbols from the alphabet. Additionally, one of the states is designated as the initial state, and a subset of the states are designated as accepting states. This way, an automaton for a regular expression models the set of strings that are accepted by the regular expression.

An automaton for a NetKAT program is similar, but instead of modeling a set of strings, it models the traces that are produced for a given input packet. When a packet travels through the automaton, every state that it traverses acts as a dup operation, appending a copy of the packet to the packet’s trace. Therefore, the transition between two states is labeled with a dup-free NetKAT program, represented symbolically as a SPP. That is, each edge in the automaton is not labeled with a single letter, as it would be in a standard DFA or NFA, but with a potentially large SPP. This is necessary because an edge in the automaton intuitively represents “what happens between two dups”. This is precisely a dup-free NetKAT program, which can itself have rich behavior, represented canonically as an SPP. This is needed to support fine-grained control over equivalence and inclusion as discussed in Section 2—e.g., via conditionally executed dups or assignments between dups.

Secondly, instead of having a boolean at each state that determines whether it is accepting or not, NetKAT automata have an additional SPP at each state that determines the set of packets that are produced as output of the automaton when a packet reaches that state. Therefore, a symbolic NetKAT automaton is a tuple $\mathcal{A} = \langle Q, q_0, \delta, \epsilon \rangle$ where

- Q is a finite set of states,
- $q_0 \in Q$ is the initial state,
- $\delta : Q \times Q \rightarrow \text{SPP}$ is the transition function, and
- $\epsilon : Q \rightarrow \text{SPP}$ is the output function.

Deterministic NetKAT automata. The notion of a deterministic NetKAT automaton is more subtle than for classic finite automata. For classic automata, a deterministic automaton is one where for every state q and symbol a , there is at most one transition from q labeled with a . For a NetKAT automaton, we need to take into account the fact that the transitions are labeled with SPPs, which may produce multiple packets for a given input packet. Therefore, we define a deterministic

NetKAT automaton as one where for every state q and packet α , the set of packets produced by $\text{push}(\alpha, \delta(q, q'))$ are disjoint for all $q' \in Q$. This is equivalent to saying that $\delta(q, q'_1) \hat{\cap} \delta(q, q'_2) = \perp$ for all $q'_1, q'_2 \in Q$ ($q'_1 \neq q'_2$). Note that an input packet α may produce multiple different packets at each successor state, but these packet sets at different successor states are disjoint.

4.2 Constructing Automata via Brzowski Derivatives

We now describe how to construct a symbolic NetKAT automaton for a NetKAT program via the Brzowski derivative. We take the set of states to be the set of NetKAT expressions, and the initial state to be the NetKAT program itself. Because we want to construct a deterministic automaton, we need a way to represent a non-intersecting outgoing symbolic transition structure (STS). We represent such a transition structure as a NetKAT expression in the following form:

$$r \in \text{STS} ::= p_1 \cdot \text{dup} \cdot q_1 + \dots + p_n \cdot \text{dup} \cdot q_n$$

where $p_i \in \text{SPP}$, and q_i are NetKAT expressions, and the p_i are disjoint ($p_i \hat{\cap} p_j = \perp$ for $i \neq j$).

Operations. Dup is an STS, by taking $r = \top \cdot \text{dup} \cdot \top$. We extend the logical operators to STSs:

$$\tilde{\top}, \tilde{\cap}, \tilde{\oplus}, \tilde{\leftarrow} : \text{STS} \times \text{STS} \rightarrow \text{STS}$$

These operations need to be defined such that the resulting STS is deterministic. For $r_1 \cap r_2$:

$$(p_1 \cdot \text{dup} \cdot q_1 + \dots + p_n \cdot \text{dup} \cdot q_n) \tilde{\cap} (p'_1 \cdot \text{dup} \cdot q'_1 + \dots + p'_n \cdot \text{dup} \cdot q'_n)$$

To bring this in STS form, we distribute the intersection over the union, and combine the terms:

$$(p_1 \hat{\cap} p'_1) \cdot \text{dup} \cdot (q_1 \cap q'_1) + (p_1 \hat{\cap} p'_2) \cdot \text{dup} \cdot (q_1 \cap q'_2) + \dots + (p_n \hat{\cap} p'_n) \cdot \text{dup} \cdot (q_n \cap q'_n)$$

This is not yet in STS form, as the $q_i \cap q'_i$ terms may not all be different, so we need to collect the terms with the same $q_i \cap q'_i$, and union their SPPs. The other operators need a similar (albeit slightly more complicated) treatment in order to maintain determinism.

Second, we extend sequential composition to STSs, in two forms (denoted with the same symbol). We can compose an STS with a SPP on the left, or with a NetKAT expression on the right:

$$\tilde{\cdot} : \text{SPP} \times \text{STS} \rightarrow \text{STS} \quad \tilde{\cdot} : \text{STS} \times \text{Exp} \rightarrow \text{STS}$$

Like the other operators, these need to be defined such that the resulting STS is deterministic.

All of the STS operations are defined in terms of SPP operations. Therefore, an efficient SPP implementation is a key component of our system, not only for the operations on symbolic packets, but also for the operations on STSs that are used to construct the NetKAT automaton.

Specifications. Each of these operations is uniquely specified by two correctness conditions:

- (1) They semantically match their counterpart (e.g. $\llbracket p \tilde{\top} q \rrbracket \equiv \llbracket p + q \rrbracket$ for all $p, q \in \text{STS}$)
- (2) They maintain the invariant that the transitions are pairwise disjoint.

Brzowski derivative. With these operations in place, it is straightforward to define the Brzowski derivative for NetKAT expressions:

$$\begin{array}{llll} \epsilon(p + q) \triangleq \epsilon(p) \hat{\top} \epsilon(q) & \epsilon(\text{dup}) \triangleq \perp & \delta(p + q) \triangleq \delta(p) \tilde{\top} \delta(q) & \delta(\text{dup}) \triangleq \text{dup} \\ \epsilon(p \cap q) \triangleq \epsilon(p) \hat{\cap} \epsilon(q) & \epsilon(f = v) \triangleq f = v & \delta(p \cap q) \triangleq \delta(p) \tilde{\cap} \delta(q) & \delta(f = v) \triangleq \perp \\ \epsilon(p \oplus q) \triangleq \epsilon(p) \hat{\oplus} \epsilon(q) & \epsilon(f \neq v) \triangleq f \neq v & \delta(p \oplus q) \triangleq \delta(p) \tilde{\oplus} \delta(q) & \delta(f \neq v) \triangleq \perp \\ \epsilon(p - q) \triangleq \epsilon(p) \hat{\leftarrow} \epsilon(q) & \epsilon(f \leftarrow v) \triangleq f \leftarrow v & \delta(p - q) \triangleq \delta(p) \tilde{\leftarrow} \delta(q) & \delta(f \leftarrow v) \triangleq \perp \\ \epsilon(p \cdot q) \triangleq \epsilon(p) \hat{\cdot} \epsilon(q) & \epsilon(\top) \triangleq \top & \delta(p \cdot q) \triangleq \delta(p) \tilde{\cdot} q \tilde{\cdot} \epsilon(p) \tilde{\cdot} \delta(q) & \delta(\top) \triangleq \perp \\ \epsilon(p^\star) \triangleq \epsilon(p)^\star & \epsilon(\perp) \triangleq \perp & \delta(p^\star) \triangleq \epsilon(p)^\star \tilde{\cdot} \delta(p) \tilde{\cdot} p^\star & \delta(\perp) \triangleq \perp \end{array}$$

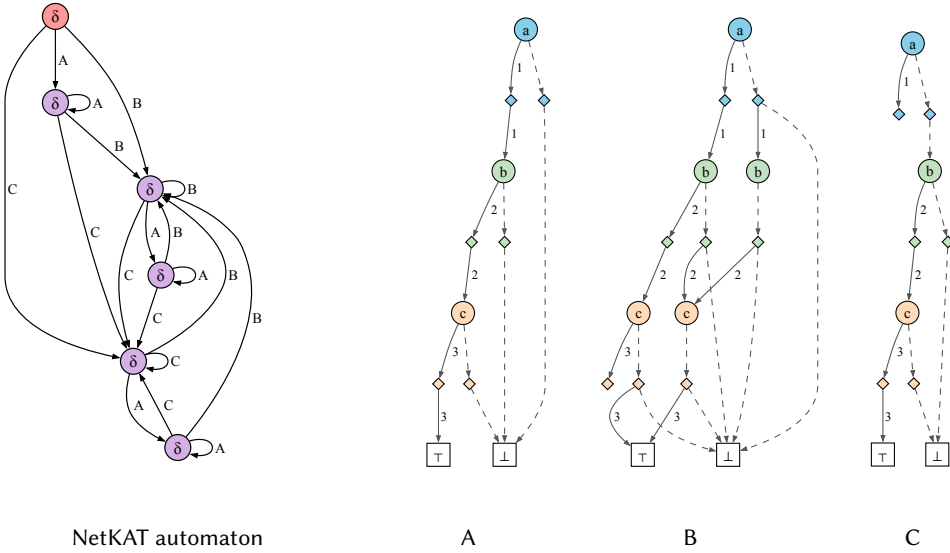


Fig. 7. Symbolic NetKAT automaton for $p \triangleq ((a \leftarrow 1 \cdot b \leftarrow 2 \cdot c \leftarrow 3 \cdot \text{dup})^* + (b = 2 \cdot c = 3 \cdot \text{dup})^*)^*$

We construct a deterministic symbolic NetKAT automaton for a NetKAT program p as follows:

States $Q \triangleq \text{Exp}$.

Initial state $q_0 \triangleq p$.

Transitions take $\delta(q, q')$ to be the SPP of q' in the Brzozowski derivative of $\delta(q)$.

Output $\epsilon : Q \rightarrow \text{SPP}$, defined above.

The Brzozowski derivative is guaranteed to transitively reach only finitely many essentially different NetKAT expressions from a given start state. Therefore, in the actual implementation, we do not use the infinite set Exp for the states, but instead use only the finitely many essentially different NetKAT terms reached from the start state.

The NetKAT automaton for an example program is shown in Figure 7. The edges are labeled with the SPPs that represent the transitions. The output SPP of every state is \top , i.e., the automaton accepts all packets that reach a state. Note that the SPPs in the figure feature several diamonds without outgoing edges. These diamonds produce no output packet, i.e., the packet is dropped.

For more details, see [?].

5 BISIMILARITY AND COUNTER-EXAMPLE GENERATION

After converting the NetKAT programs to automata, we can check equivalence of two NetKAT programs by checking whether their automata are bisimilar. Bisimilarity is a well-known notion of equivalence for automata. For NetKAT automata, two states are bisimilar for a given packet α if (1) the states produce the same output packets (2) for each possible modified packet α' , the two states transition to states that are bisimilar for α' .

Historically, methods for computing bisimulations of automata [8, 17] have been based on either on Moore's algorithm for minimization [32] or Hopcroft and Karp's algorithm [24]. Indeed, the naive approach shown in Figure 2 follows the basic structure of Hopcroft-Karp, in the sense that it relates the start states and proceeds to follow transitions forward in the two automata.

Our situation is different, because we have already added the symmetric difference operator to NetKAT. Therefore, we can reduce equivalence checks $A \equiv B$ to emptiness checks $A \oplus B \equiv \perp$.

<pre> done(q) ← ⊥ for q ∈ Q; todo(q) ← ⊥ for q ∈ Q \ {q₀}; todo(q₀) ← ⊤; while ∃q. todo(q) ≠ ⊥ do set p := todo(q) $\hat{=}$ done(q); todo(q) ← ⊥; done(q) $\hat{=}$ p; for q' ∈ Q do todo(q') $\hat{=}$ push(p, δ(q, q')); return ∑_{q∈Q} push(done(q), ε(q)); </pre> <p style="text-align: center;">(i) Forward algorithm</p>	<pre> done(q) ← ⊥ for q ∈ Q; todo(q) ← pull(ε(q), ⊤) for q ∈ Q; while ∃q. todo(q) ≠ ⊥ do set p := todo(q) $\hat{=}$ done(q); todo(q) ← ⊥; done(q) $\hat{=}$ p; for q' ∈ Q do todo(q') $\hat{=}$ pull(δ(q', q), p); return done(q₀); </pre> <p style="text-align: center;">(ii) Backward algorithm</p>
---	--

Fig. 8. Forward and backward algorithms for NetKAT automata

Subtlety of symmetric difference. The reader should note that the situation is a bit more subtle than it may seem at first sight. In particular, consider the query $A \equiv B$, which asks whether two NetKAT programs A and B are equivalent. If they are inequivalent, then there must be some input packet that causes A to produce a different output than B . Output in this sense does not just mean the final output packet, but also the trace of the packet. It can be the case that A and B produce the same final output packets, but differ in the trace of the packets. Therefore, the symmetric difference $A \oplus B$ takes the symmetric difference of the traces, not just of the output packets, i.e., the symmetric difference $A \oplus B$ contains precisely the traces of the packets that are produced by A or B , but not by both. In other words, the symmetric difference $A \oplus B$ contains precisely the traces that are counter-examples to the equivalence of A and B .

Subtlety of the emptiness of an automaton. A second subtlety is the emptiness of an automaton. Whereas for regular expressions, it is easy to check whether a DFA is empty (just check whether there are any reachable accepting states), this is not the case for NetKAT automata. Because NetKAT automata manipulate and test the fields of packets, it is possible that a packet travels through the automaton, and even causes multiple packets to be produced (e.g., due to the presence of union in the original NetKAT expression), but nevertheless, it is possible that all of these packets are eventually dropped by the automaton, before producing any output packets.

The goal of this section is to develop a symbolic algorithm for this check, which is more efficient than the naive algorithm that checks all concrete input packets separately.

5.1 Forward Algorithm

To check whether a NetKAT automaton drops all input packets, we use an algorithm that works in the forward direction. The forward algorithm computes *all* output packets that the automaton can produce, across all possible input packets. More formally: given a NetKAT automaton for a NetKAT program p , the forward algorithm computes the set of final packets occurring in the traces $\llbracket p \rrbracket(\alpha)$ across all input packets α . The forward algorithm therefore starts with the complete symbolic packet \top at the input state, and repeatedly applies all outgoing transitions δ to it. In this way, the algorithm iteratively accumulates a symbolic packet at every state, which represents the set of packets that can reach that state from the start state. Once we know the set of packets that can reach a state, we can determine the set of output packets by applying the output function ϵ to the symbolic packet of each state, and taking the union of the results.

The forward algorithm is shown in [Figure 8](#). To determine the set of packets that a given NetKAT expression p can produce, we convert it to a NetKAT automaton, and then use the forward algorithm to determine the symbolic output packet. In our implementation, we also have a way to stop the algorithm early, if the user is only interested in a “yes” or “no” answer for the query $p \equiv \perp$. In this case, we can stop the algorithm as soon as any output packet is produced.

5.2 Backward Algorithm

The forward algorithm gives us the set of output packets that a NetKAT automaton can produce, but we are also interested in which input packets cause this output to be produced. For a given check $A \equiv B$, we want to know *which* input packets cause A and B to produce different sets of output traces. More generally, for a NetKAT automaton, we want to know which input packets cause the automaton to produce a non-empty set of output packets. More formally: given a NetKAT automaton for a NetKAT program p , the backward algorithm computes the set of initial packets α for which $\llbracket p \rrbracket(\alpha)$ is nonempty.

To answer this question, we developed a *backward algorithm*, shown in [Figure 8](#). The backward algorithm computes *all* input packets that can cause the automaton to produce a non-empty set of output packets. The backward algorithm therefore starts with the complete symbolic packet \top at every state, and pulls it backwards through all output transitions ϵ . The algorithm then iteratively accumulates a symbolic packet at every state, which represents the set of packets that will cause the automaton to produce a non-empty set of output packets, when starting from that state. The accumulation is done by pulling the symbolic packet backwards through all transitions δ , until fixpoint is reached. Once fixpoint is reached, we simply return the symbolic packet at the start state, which represents the set of input packets that will cause the automaton to produce a non-empty set of output packets when starting from the start state.

6 IMPLEMENTATION

We have implemented the algorithms in a new system, KATch, comprising 2500 lines of Scala. In this section we explain the system’s interface, and in the next section we evaluate its performance. The implementation provides a surface syntax for expressing queries, which extends the core NetKAT syntax from [Figure 1](#). The extended syntax is shown in [Figure 9](#). The language has statements and expressions, which we describe below.

Statements. The check $e_1 \equiv e_2$ command runs the bisimulation algorithm. The system reports success if the expressions are equivalent (when \equiv is used) or inequivalent (when \neq) is used, or reports failure otherwise. The other statements behave as expected: The print statement invokes the pretty printer, $x = e$ let-binds names to expressions, and for runs a statement in a loop.

Expressions. The forward e expression computes, in forward-flowing mode, the set of output packets resulting from running the symbolic packet \top through the given expression e (see [Figure 8](#)). Conversely backward e computes the set of input packets which generate some output packet when run on the given expression ([Figure 8](#)). These are generally used in conjunction with the \oplus , $-$, \cap operators to express the desired query. The user may combine these expressions with the $\hat{\exists}$ and $\hat{\forall}$ operators to reason about symbolic packets, and the print operator to pretty print the symbolic packets, or the check operator to assert (in)equivalence of two expressions.

Statements

```
check  $e_1 \equiv e_2$ 
check  $e_1 \neq e_2$ 
print  $e$ 
 $x = e$ 
for  $i \in n_1..n_2$  do  $c$ 
```

Expressions

```
forward  $e$ , backward  $e$ 
 $e_1 \cap e_2$ ,  $e_1 \oplus e_2$ ,  $e_1 - e_2$ 
 $\hat{\exists} f e$ ,  $\hat{\forall} f e$ 
(+ NetKAT, see Figure 1)
```

Fig. 9. NKPL syntax.

We note that the generality of the language allows us to express some queries in different, equivalent ways. For example, the two checks: $e_1 \equiv e_2$ and $e_1 \oplus e_2 \equiv \perp$ are equivalent. However, the expression on the right lends itself to inspection of counterexample input packets: print (backward $e_1 \oplus e_2$). This statement pretty prints a symbolic packet having different behavior on e_1 and e_2 (i.e., packets that result in some valid history in one expression and not the other). In other words, it prints the set of all counter-example input packets for $e_1 \equiv e_2$.

Topologies and routing tables. While NetKAT can be used to specify routing policies declaratively, it is also possible to import topologies and routing tables into NetKAT. For example, a simple way to define routing tables and topologies in NetKAT is as follows:

$$\begin{array}{l|l} R \triangleq \text{sw}=5 \cdot (\text{dst}=6 \cdot \text{pt} \leftarrow 3 + \text{dst}=8 \cdot \text{pt} \leftarrow 4) & T \triangleq \text{sw}=5 \cdot (\text{pt} = 3 \cdot \text{sw} \leftarrow 6 + \text{pt} = 4 \cdot \text{sw} \leftarrow 8) \\ \quad \quad \quad + \cdots + & \quad \quad \quad + \cdots + \\ \text{sw}=7 \cdot (\text{dst}=8 \cdot \text{pt} \leftarrow 2 + \text{dst}=9 \cdot \text{pt} \leftarrow 1) & \text{sw}=7 \cdot (\text{pt} = 2 \cdot \text{sw} \leftarrow 8 + \text{pt} = 1 \cdot \text{sw} \leftarrow 9) \end{array}$$

The routing R tests the switch field (where the packet currently is), and the destination field (where the packet is supposed to end up), and then sets the port over which the packet should be sent out. The topology T tests the switch and port fields, and then transports the packet to the next switch. We define the action of the network by composing the route and topology:

$$\text{net} \triangleq R \cdot T \cdot \text{dup}$$

We include a dup to extend the trace of the packet at every hop.

Example 6.1 (All-Pairs Reachability Queries). A naive way to check reachability of all pairs of hosts in a network is to run the following command for each pair of end hosts n_i, n_j :

$$\text{check } (\text{sw}=n_i) \cdot \text{net}^* \cdot (\text{sw}=n_j) \neq \perp$$

Of course, this requires a number of queries which is quadratic in the number of hosts—quickly becoming prohibitive. One might think that we could reduce the number of queries to n by running:

$$\text{for } i \in 1..n \text{ do check (forward } (\text{sw}=i \cdot \text{net}^*)) \equiv (\text{sw} \in 1..n)$$

This query does work if sw is the only field of our packets, because the left hand side contains the packets that can be reached from i via the network, and the right hand side contains packets where the sw field is any value in $1..n$. Unfortunately, this does not quite work if the network also operates on other packet fields (say, fields f_1 and f_2), as the packets on the left hand side will have those fields, whereas the packets on the right hand side will only have a sw field. The $\hat{\exists}$ and $\hat{\forall}$ operators allow us to manipulate symbolic packets and express all pairs reachability in n queries:

$$\text{for } i \in 1..n \text{ do check } (\hat{\exists} f_1 (\hat{\exists} f_2 (\text{forward } (\text{sw} = i \cdot \text{net}^*)))) \equiv (\text{sw} \in 1..n)$$

The operators $\hat{\exists}$ and $\hat{\forall}$ give the programmer the ability to reason about symbolic packets that may be computed, for instance, by forward or backward. The specifications are:

$$p \in \hat{\exists} f e \iff \exists v \in V: p[f \leftarrow v] \in \llbracket e \rrbracket \quad \text{and} \quad p \in \hat{\forall} f e \iff \forall v \in V: p[f \leftarrow v] \in \llbracket e \rrbracket$$

The implementation does not iterate over V (indeed, we need not even know V). Rather, $\hat{\exists}$ and $\hat{\forall}$ are implemented directly as operations on symbolic packets. In fact, the implementation does not need to fix the set of fields or the set of values up-front at all. Instead, it operates on a conceptually infinite set of fields and values. This works because SPPs' default cases handle all remaining fields and all values.




Fig. 10. Full reachability queries on Topology Zoo. KATch and APKEEP results are averages of 100 runs, preceded by 10 runs of JIT warmup. KATch-naive is without JIT warmup and uses $O(n^2)$ 1-to-1 queries.

Correctness and testing methodology. We validated our implementation with the following property-based fuzz testing methodology:

- (1) We implemented the semantics $\llbracket p \rrbracket(\alpha)$ in Scala, as described in Figure 1.
- (2) We repeatedly pick two SPs/SPPs p and q and a packet α . We enumerated small SPPs up to a bound and generated larger SPPs randomly. We generated packets exhaustively.
- (3) We check the soundness (e.g., $\llbracket p \hat{+} q \rrbracket(\alpha) = \llbracket p \rrbracket(\alpha) \cup \llbracket q \rrbracket(\alpha)$) and canonicity (e.g., $\llbracket p \rrbracket(\alpha) = \llbracket q \rrbracket(\alpha)$ for all α if and only if $p = q$) of the operations, as well as additional algebraic laws derived from the NetKAT axioms (e.g., $p \hat{+} q = q \hat{+} p$).

In addition to this property-based testing for SPs and SPPs, we also generated hundreds of thousands of pairs of random NetKAT expressions and checked that KATch’s bisimilarity check matches the output of FRENETIC. These efforts revealed a subtle bug in Frenetic, which we reported and is now fixed.

7 EVALUATION

To evaluate KATch, we conducted experiments in which we used it to solve a variety of verification tasks for a range of topologies and routes, as well as challenging combinatorial NetKAT terms. The goal of our evaluation is to answer the following three questions:

- (1) How does KATch perform compared to the state of the art NetKAT verifier, FRENETIC?
- (2) How does KATch perform compared to the state of the art specialized network verification tool, APKEEP?
- (3) How well does KATch scale with the size of topology?
- (4) When does KATch perform asymptotically better than prior work?

7.1 Topology Zoo

To begin to answer the first three questions, we conducted our experiments using *The Internet Topology Zoo* [27] dataset, a publicly available set of 261 network topologies, ranging in size from just 4 nodes (the original ARPANet) to 754 nodes (KDL, the Kentucky Data Link ISP topology). For each topology, we generated a destination-based routing policy using an all-pairs shortest path scheme that connects every pair of routers to each other.

Name (Topology Zoo)	Size (atoms)	1-to-1 reachability		Slicing		Full reachability	
		KATch	Frenetic	KATch	Frenetic	KATch	APKEEP
Layer42	135	0.00	0.04	0.00	0.07	0.00	0.00
Compuserv	539	0.00	0.35	0.01	0.81	0.01	0.01
Airtel	785	0.01	0.82	0.02	1.98	0.01	0.01
Belnet	1388	0.01	3.19	0.03	7.99	0.02	0.01
Shentel	1865	0.01	3.93	0.03	9.76	0.02	0.03
Arpa	1964	0.01	4.23	0.04	10.42	0.03	0.03
Sanet	4100	0.03	23.42	0.07	62.21	0.05	0.07
Unet	5456	0.04	80.85	0.11	203.80	0.07	0.07
Missouri	9680	0.06	166.87	0.21	441.72	0.13	0.19
Telcove	10720	0.07	441.47	0.21	1121.30	0.12	0.16
Deltacom	27092	0.18	2087.34	0.48	5098.57	0.36	0.63
Cogentco	79682	0.53	18910.82	1.38	54247.73	0.98	2.67
Kdl	1144691	9.87	out of memory	23.67	out of memory	19.44	110.92

Fig. 11. Running time (in seconds) of KATch, FRENETIC, and APKEEP on Topology Zoo queries. KATch and APKEEP results are averages of 100 runs, preceded by 10 runs of JIT warmup.

To demonstrate KATch’s scalability, we first ran full (i.e., $O(n^2)$) reachability queries for every topology in the zoo using KATch, FRENETIC², and APKEEP [44]. The results are shown in Figure 10 in a log-log plot, so straight lines correspond to polynomials with exponents related to their slope.

The figure shows two different configurations of KATch: one that verifies full reachability using a quadratic number of point-to-point queries and does not use JIT warmup (“katch-naive”), and one that verifies full reachability using a linear number of queries and does use JIT warmup (“katch”).

Because of the size of the dataset, we set a timeout of 5 minutes per topology. Under these conditions, FRENETIC was unable to complete for all but the smallest topologies. KATch-naive handles most of the topologies in well under a second, and all but the largest in under 2 minutes. KATch-naive exceeds the timeout on Kentucky Data Link—using a quadratic number of queries to check full reachability produces over 500k individual point-to-point queries in a network with 754 nodes! (However, KATch-naive is able to finish it in just under 20 minutes.)

To avoid combinatorial blowup in the verification query itself, we also used KATch’s high-level verification interface, to check full reachability using a *linear* number of queries, as discussed in Example 6.1. FRENETIC does not have an analogous linear mode. APKEEP does have an analogous mode, as it has specialized support for full reachability queries. APKEEP is faster than KATch for the smaller topologies, but slower for the larger ones. Indeed, the slope of the APKEEP line is steeper than the KATch line, indicating that KATch scales slightly better with the size of the network on these queries. It is important to note that APKEEP has support for prefix-matching, ACLs, NAT, incrementality, and other features for which KATch does not have specialized support and which are not tested in this comparison. One should therefore not draw strong conclusions from this comparison; we include it as it is encouraging that reachability via NetKAT equivalence can be competitive with a state-of-the-art specialized tool.

We also randomly sampled a subset of topologies from the Topology Zoo of varying size and generated point-to-point (1-to-1) reachability and slicing queries to be checked against the routing configurations with no timeout. For each query, we ran KATch and also generated an equivalent query in the syntax of FRENETIC, and ran FRENETIC’s bisimulation verifier on those queries. We present a full table of results of these experiments in Figure 11. The table shows that KATch’s relative

²<https://github.com/frenetic-lang/frenetic>




Fig. 12. Results of running KATch and FRENETIC on combinatorial benchmarks

speedup over FRENETIC is considerable, and increases problem size. This is encouraging, because it shows that KATch is more scalable. FRENETIC was unable to complete KDL within a 200GB memory limit (for comparison, KATch is able to complete the same query with well under 1GB). We also include a selection of full-reachability results from Figure 10 in the table. The reader may wonder why full reachability takes only twice as long as 1-to-1 reachability. This is because a significant fraction of the time is spent constructing automata (which are similar for both queries), and much of the work in the full reachability queries is shared due to memoization of SPP operations.

7.2 Combinatorial Examples

Finally, we ran experiments to test the hypothesis that SPPs have an asymptotic advantage for certain types of queries. FRENETIC uses forwarding decision diagrams (FDDs), which is a different (and non-canonical) representation for automaton transitions. A key advantage of SPPs over FDDs is that SPPs keep the updates to each field next to the tests of the same field. On the other hand, FDDs keep all updates at the leaves, which can result in combinatorial explosion during sequencing. To test this, we generated the following NetKAT programs:

Inc: Treating the input packet’s n boolean fields as a binary number, increment it by one.

Flip: Sequentially flip the value of each of the n boolean fields.

Nondet: Set each field of the packet to a range of values from 0 to n .

For Inc, we tested that repeatedly incrementing (using the \star operator) can turn packet $00 \cdots 0$ into $11 \cdots 1$. For Flip, we tested that flipping all bits twice returns the original packet. For Nondet, we tested that setting the fields non-deterministically twice is the same as doing it once. The results of this experiment are shown in Figure 12. Because the available fields are hardcoded in FRENETIC, we only ran the Inc and Flip experiments up to $n = 10$. We ran the non-determinism test up to $n = 15$. KATch finishes all three queries up to $n = 100$ in under a one minute, demonstrating its asymptotic advantage for these queries, while Frenetic shows combinatorial blowup on larger packets.

Source of speedup. The speedup of KATch over FRENETIC comes from several sources: (1) the use of SPPs instead of FDDs, which can be exponentially more compact and support more efficient sequential composition (2) the symbolic bisimulation algorithm operating on SPs, which can handle exponentially large sets of packets at once (3) a quadratic-doubling implementation of star, which can handle exponentially long traces quickly. These differences show up most strongly in the combinatorially adversarial examples above, but note that the speedup is also considerable for the Topology Zoo benchmarks, which are based on real-world topologies and not designed to be combinatorially adversarial.

Field order. The field order of SPs and SPPs can affect performance, just as for BDDs. In practice, field orders that keep related fields close together are beneficial. Our implementation allows the user to control the field order, but we did not use this for our benchmarks. By default, the fields are in the order in which they first occur in the input file, which turned out to work well enough.

8 RELATED WORK

This section discusses the most closely related prior work to this paper, focusing on three areas: NetKAT, network verification, and automata-theoretic approaches to verification.

NetKAT. NetKAT was originally proposed as a semantic foundation for SDN data planes [1]. Indeed, being based on KAT [28], the language provides a sound and complete algebraic reasoning system. Later work on NetKAT developed an automata-theoretic (or coalgebraic) account of the language [21], including a decision procedure based on Brzozowski derivatives and bisimulation, implemented in FRENETIC. However, the performance of this approach turns out to be poor, as shown in our experiments, due to the use of ad hoc data structures (“bases”) to encode packets and automata. NetKAT’s compiler uses a variant of BDDs, called Forwarding Decision Diagrams (FDDs), as well as an algorithm for converting programs to automata using Antimirov derivatives [34], improving on the earlier representation of transitions as sets of bases [21]. The SPPs proposed in this paper improve on FDDs by ensuring uniqueness and supporting efficient sequential composition in the common case. In addition, the deterministic automata used in KATch support additional “negative” operators that are useful for verification. Other papers based on NetKAT have explored use of the language in other settings such as distributed control planes [7], probabilistic networks [20, 36, 37], and Kleene algebra over composable theories with unbounded state [22]. It would be interesting to extend the techniques developed in this paper to these richer settings. Another interesting direction for future work is to build a symbolic verifier for the guarded fragment of NetKAT [35].

Network Verification. Early work by Xie et al. [42] proposed a unifying mathematical model for Internet routers and developed algorithms for analyzing network-wide reachability. Although the paper did not discuss an implementation, its elegant formal model has been extremely influential in the community and has served as the foundation for many follow-on efforts, including this work. The emergence of software-defined networking (SDN) led to a surge of interest in static data plane verification, including systems such as Header Space Analysis (HSA) [25], Anteater [31], VeriFlow [26], Atomic Predicates (AP) [43], APKeep [44]. These systems all follow a common approach: they build a model of the network-wide forwarding behavior and then check whether given properties hold. However, they vary in the data structures and algorithms used to represent and analyze the network model. For instance, Anteater relies on first-order logic and SAT solvers, while VeriFlow uses prefix trees and custom graph-based algorithms. HSA, AP and APKeep are arguably the most related to our work as they use symbolic representations and BDDs respectively.

The primary difference between KATch and these systems is that the latter implement specialized algorithms for network-wide analysis queries, while KATch (and FRENETIC) solve the more general NetKAT equivalence problem, into which network-wide analysis queries can be encoded. This generic approach is based on principled automata-theoretic foundations, but one might expect it to be less efficient than specialized algorithms. FRENETIC was found to be comparable to or faster than HSA [21], so by transitivity we expect KATch to be faster than HSA. However, the current state of the art is APKEEP, which is significantly faster than FRENETIC and HSA. We include a comparison with APKEEP in Section 7 on reachability queries. The results show that KATch is competitive with APKEEP on these benchmarks, despite solving general NetKAT queries. It is important to note that APKEEP has support for prefix-matching, ACLs, NAT, incrementality, and other features for which KATch does not have specialized support and which are not tested in this comparison. One should

therefore not draw strong conclusions from this comparison; we include it as an indication that reachability via NetKAT equivalence can be competitive with a state-of-the-art specialized tool.

Another line of work has explored how to lift verification from the data plane to the control plane. Batfish [10, 18] proposed using symbolic simulation to analyze distributed control planes—i.e., generating all possible data planes that might be produced starting from a given control-plane configuration. Like KATch, Batfish uses a BDD-based representation for data plane analysis. MineSweeper [4] improves on Batfish using an SMT encoding of the converged states of the control plane that avoids having to explicitly simulate the underlying routing protocols. Recent work has focused on using techniques like modular reasoning [38, 39] abstract interpretation [6], and a form of symmetry reduction [5] to further improve the scalability of control-plane verification.

Automata-Theoretic Approach and Symbolic Automata. Our work on KATch builds on the large body of work on the automata-theoretic approach to verification and symbolic automata. The automata-theoretic approach was pioneered in the 1980s, with applications of temporal logics and model checking to hardware verification. BDDs, originally proposed by Lee [29], were further developed by Bryant [11], and used by McMillan for symbolic model checking [14].

An influential line of work by D’Antoni and Veaus developed techniques for representing and transforming finite automata where the transitions are not labeled with individual characters but with elements of a so-called effective Boolean algebra [15, 16]. Shifting from a concrete to a symbolic representation requires generalizing classical algorithms, such as minimization and equivalence, but facilitates building automata that work over huge alphabets, such as Unicode.

Pous [33] developed symbolic techniques for checking the equivalence of automata where transitions are specified using BDDs. The methods developed in this paper take Pous’s work as a starting point but develop a non-trivial extension for NetKAT. In particular, SPPs provide a compact representation for “carry-on” packets, which is a critical and unique aspect of NetKAT’s semantics. Bonchi and Pous [8] explored the use of up-to techniques for checking equivalence of automata. In principle, one can view the invariants enforced by KATch’s term representations as a kind of up-to technique, but a full investigation of this idea requires more work.

Our backward algorithm for computing bisimulations can be seen of a variant of Moore’s classic algorithm for computing the greatest bisimulation to NetKAT automata. Doenges et al. [17] proposed an analogous approach for checking equivalence of automata that model the behavior of P4 packet parsers [9]. However, Leapfrog’s model is simpler than NetKAT’s, being based on classic finite automata, and it achieves scalability primarily due to a novel up-to technique that “leaps” over internal buffering transitions rather than symbolic representations.

ACKNOWLEDGEMENTS

We wish to thank the PLDI reviewers for many helpful comments and suggestions, which improved the paper significantly. We are also grateful to Aaron Gember-Jacobson and Peng Zhang for making AP Keep available as open-source software, to Caleb Koch for early discussions on compact data structures for data plane verification, and to the Cornell PLDG for helpful comments on early drafts. This work was supported in part by ONR grant award N68335-22-C-0411 and DARPA grant W912CG-23-C-0032.

DATA AVAILABILITY

The current version of the implementation is available at <https://github.com/cornell-netlab/KATch>. A snapshot of the code as of artifact evaluation (with a dependencies-included Docker image) is available at <https://zenodo.org/doi/10.5281/zenodo.10806446>.

REFERENCES

- [1] Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker. 2014. NetKAT: Semantic Foundations for Networks. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/2535838.2535862>
- [2] Valentin Antimirov. 1996. Partial derivatives of regular expressions and finite automaton constructions. *Theoretical Computer Science* (1996). [https://doi.org/10.1016/0304-3975\(95\)00182-4](https://doi.org/10.1016/0304-3975(95)00182-4)
- [3] Michael Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K. Rustan M. Leino. 2005. Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In *Formal Methods for Components and Objects, 4th International Symposium, FMCO 2005, Amsterdam, The Netherlands, November 1-4, 2005, Revised Lectures*. https://doi.org/10.1007/11804192_17
- [4] Ryan Beckett, Aarti Gupta, Ratul Mahajan, and David Walker. 2017. A General Approach to Network Configuration Verification. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. <https://doi.org/10.1145/3098822.3098834>
- [5] Ryan Beckett, Aarti Gupta, Ratul Mahajan, and David Walker. 2018. Control plane compression. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM 2018, Budapest, Hungary, August 20-25, 2018*. <https://doi.org/10.1145/3230543.3230583>
- [6] Ryan Beckett, Aarti Gupta, Ratul Mahajan, and David Walker. 2020. Abstract interpretation of distributed network control planes. *POPL* (2020). <https://doi.org/10.1145/3371110>
- [7] Ryan Beckett, Ratul Mahajan, Todd Millstein, Jitendra Padhye, and David Walker. 2016. Don't Mind the Gap: Bridging Network-Wide Objectives and Device-Level Configurations. In *Proceedings of the 2016 ACM SIGCOMM Conference*. <https://doi.org/10.1145/2934872.2934909>
- [8] Filippo Bonchi and Damien Pous. 2013. Checking NFA Equivalence with Bisimulations up to Congruence. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/2429069.2429124>
- [9] Pat Bosshart, Dan Daly, Glen Gibb, Martin Izzard, Nick McKeown, Jennifer Rexford, Cole Schlesinger, Dan Talayco, Amin Vahdat, George Varghese, and David Walker. 2014. P4: Programming Protocol-Independent Packet Processors. *SIGCOMM Comput. Commun. Rev.* (2014). <https://doi.org/10.1145/2656877.2656890>
- [10] Matt Brown, Ari Fogel, Daniel Halperin, Victor Heorhiadi, Ratul Mahajan, and Todd D. Millstein. 2023. Lessons from the evolution of the Batfish configuration analysis tool. In *Proceedings of the ACM SIGCOMM 2023 Conference, ACM SIGCOMM 2023, New York, NY, USA, 10-14 September 2023*. <https://doi.org/10.1145/3603269.3604866>
- [11] Randal E. Bryant. 1986. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Trans. Computers* (1986). <https://doi.org/10.1109/TC.1986.1676819>
- [12] Randal E. Bryant. 1992. Symbolic Boolean Manipulation with Ordered Binary-Decision Diagrams. *ACM Comput. Surv.* (1992). <https://doi.org/10.1145/136035.136043>
- [13] Janusz A. Brzozowski. 1962. Canonical regular expressions and minimal state graphs for definite events. In *Proc. Symposium of Mathematical Theory of Automata*.
- [14] Jerry R. Burch, Edmund M. Clarke, Kenneth L. McMillan, David L. Dill, and L. J. Hwang. 1990. Symbolic Model Checking: 10²⁰ States and Beyond. In *LICS*. <https://doi.org/10.1109/LICS.1990.113767>
- [15] Loris D'Antoni and Margus Veanes. 2014. Minimization of Symbolic Automata. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/2535838.2535849>
- [16] Loris D'Antoni and Margus Veanes. 2017. Forward Bisimulations for Nondeterministic Symbolic Finite Automata. In *Proceedings, Part I, of the 23rd International Conference on Tools and Algorithms for the Construction and Analysis of Systems - Volume 10205*. https://doi.org/10.1007/978-3-662-54577-5_30
- [17] Ryan Doenges, Tobias Kappé, John Sarracino, Nate Foster, and Greg Morrisett. 2022. Leapfrog: Certified Equivalence for Protocol Parsers. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. <https://doi.org/10.1145/3519939.3523715>
- [18] Ari Fogel, Stanley Fung, Luis Pedrosa, Meg Walraed-Sullivan, Ramesh Govindan, Ratul Mahajan, and Todd Millstein. 2015. A general approach to network configuration analysis. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*. <https://doi.org/10.5555/2789770.2789803>
- [19] Nate Foster, Rob Harrison, Michael J. Freedman, Christopher Monsanto, Jennifer Rexford, Alec Story, and David Walker. 2011. Frenetic: a network programming language. In *ICFP*. <https://doi.org/10.1145/2034773.2034812>
- [20] Nate Foster, Dexter Kozen, Konstantinos Mamouras, Mark Reitblatt, and Alexandra Silva. 2016. Probabilistic NetKAT. In *ESOP*. https://doi.org/10.1007/978-3-662-49498-1_12
- [21] Nate Foster, Dexter Kozen, Mae Milano, Alexandra Silva, and Laure Thompson. 2015. A Coalgebraic Decision Procedure for NetKAT. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/2676726.2677011>

- [22] Michael Greenberg, Ryan Beckett, and Eric Campbell. 2022. Kleene algebra modulo theories: a framework for concrete KATs. In *PLDI*. <https://doi.org/10.1145/3519939.3523722>
- [23] Pieter Hooimeijer, Benjamin Livshits, David Molnar, Prateek Saxena, and Margus Veanes. 2011. Fast and Precise Sanitizer Analysis with BEK. In *USENIX Conference on Security*. <https://doi.org/10.5555/2028067.2028068>
- [24] John E. Hopcroft and Richard M. Karp. 1971. A Linear Algorithm for Testing Equivalence of Finite Automata.
- [25] Peyman Kazemian, George Varghese, and Nick McKeown. 2012. Header Space Analysis: Static Checking for Networks. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*. <https://doi.org/10.5555/2228298.2228311>
- [26] Ahmed Khurshid, Wenxuan Zhou, Matthew Caesar, and P. Brighten Godfrey. 2012. Veriflow: Verifying Network-Wide Invariants in Real Time. *SIGCOMM Comput. Commun. Rev.* (2012). <https://doi.org/10.1145/2377677.2377766>
- [27] Simon Knight, Hung X. Nguyen, Nickolas Falkner, Rhys Bowden, and Matthew Roughan. 2011. The Internet Topology Zoo. *Selected Areas in Communications, IEEE Journal on* (2011). <https://doi.org/10.1109/JSAC.2011.1110002>
- [28] Dexter Kozen. 1996. Kleene Algebra with Tests and Commutativity Conditions. In *Tools and Algorithms for the Construction and Analysis of Systems*. https://doi.org/10.1007/3-540-61042-1_35
- [29] C. Y. Lee. 1959. Representation of switching circuits by binary-decision programs. *The Bell System Technical Journal* (1959). <https://doi.org/10.1002/j.1538-7305.1959.tb01585.x>
- [30] K. Rustan M. Leino and Valentin Wüstholtz. 2014. The Dafny Integrated Development Environment. In *Proceedings 1st Workshop on Formal Integrated Development Environment, F-IDE 2014, Grenoble, France, April 6, 2014*. <https://doi.org/10.4204/EPTCS.149.2>
- [31] Haohui Mai, Ahmed Khurshid, Rachit Agarwal, Matthew Caesar, P. Brighten Godfrey, and Samuel Talmadge King. 2011. Debugging the Data Plane with Anteater. *SIGCOMM Comput. Commun. Rev.* (2011). <https://doi.org/10.1145/2043164.2018470>
- [32] Edward F. Moore. 1956. *Gedanken-Experiments on Sequential Machines*.
- [33] Damien Pous. 2015. Symbolic Algorithms for Language Equivalence and Kleene Algebra with Tests. *SIGPLAN Not.* (2015). <https://doi.org/10.1145/2775051.2677007>
- [34] Steffen Smolka, Spiridon Eliopoulos, Nate Foster, and Arjun Guha. 2015. A Fast Compiler for NetKAT. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*. <https://doi.org/10.1145/2784731.2784761>
- [35] Steffen Smolka, Nate Foster, Justin Hsu, Tobias Kappé, Dexter Kozen, and Alexandra Silva. 2019. Guarded Kleene Algebra with Tests: Verification of Uninterpreted Programs in Nearly Linear Time. *POPL* (2019). <https://doi.org/10.1145/3371129>
- [36] Steffen Smolka, Praveen Kumar, Nate Foster, Justin Hsu, Dexter Kozen, and Alexandra Silva. 2019. Scalable Verification of Probabilistic Networks. In *PLDI*. <https://doi.org/10.1145/3314221.3314639>
- [37] Steffen Smolka, Praveen Kumar, Nate Foster, Dexter Kozen, and Alexandra Silva. 2017. Cantor Meets Scott: Semantic Foundations for Probabilistic Networks. In *POPL*. <https://doi.org/10.1145/3009837.3009843>
- [38] Alan Tang, Ryan Beckett, Steven Benaloh, Karthick Jayaraman, Tejas Patil, Todd D. Millstein, and George Varghese. 2023. Lightyear: Using Modularity to Scale BGP Control Plane Verification. In *Proceedings of the ACM SIGCOMM 2023 Conference, ACM SIGCOMM 2023, New York, NY, USA, 10-14 September 2023*. <https://doi.org/10.1145/3603269.3604842>
- [39] Timothy Alberdingk Thijm, Ryan Beckett, Aarti Gupta, and David Walker. 2023. Modular Control Plane Verification via Temporal Invariants. *PLDI* (2023). <https://doi.org/10.1145/3591222>
- [40] Emina Torlak and Rastislav Bodik. 2013. Growing solver-aided languages with rosette. In *ACM Symposium on New Ideas in Programming and Reflections on Software, Onward! 2013, part of SPLASH '13, Indianapolis, IN, USA, October 26-31, 2013*. <https://doi.org/10.1145/2509578.2509586>
- [41] Moshe Y. Vardi and Pierre Wolper. 1986. An Automata-Theoretic Approach to Automatic Program Verification (Preliminary Report). In *LICS*.
- [42] Geoffry G. Xie, Jibin Zhan, David A. Maltz, Hui Zhang, Albert Greenberg, Gisli Hjalmtýsson, and Jennifer Rexford. 2005. On static reachability analysis of IP networks. In *INFOCOMM*. <https://doi.org/10.1109/INFOCOM.2005.1498492>
- [43] Hongkun Yang and Simon S. Lam. 2016. Real-Time Verification of Network Properties Using Atomic Predicates. *IEEE/ACM Trans. Netw.* (2016). <https://doi.org/10.1109/TNET.2015.2398197>
- [44] Peng Zhang, Xu Liu, Hongkun Yang, Ning Kang, Zhengchang Gu, and Hao Li. 2020. APKeep: Realtime Verification for Real Networks. In *NSDI*. <https://www.usenix.org/conference/nsdi20/presentation/zhang-peng>