

Abstract Interpretation of Stateful Networks

Kalev Alpernas¹, Roman Manevich², Aurojit Panda³, Mooly Sagiv¹, Scott Shenker⁴, Sharon Shoham¹, and Yaron Velner⁵

¹ Tel Aviv University

² Ben-Gurion University of the Negev

³ NYU

⁴ UC Berkeley

⁵ Hebrew University of Jerusalem

Abstract. Modern networks achieve robustness and scalability by maintaining states on their nodes. These nodes are referred to as middleboxes and are essential for network functionality. However, the presence of middleboxes drastically complicates the task of network verification. Previous work showed that the problem is undecidable in general and EXPSpace-complete when abstracting away the order of packet arrival.

We describe a new algorithm for conservatively checking isolation properties of stateful networks. The asymptotic complexity of the algorithm is polynomial in the size of the network, albeit being exponential in the maximal number of queries of the local state that a middlebox can do, which is often small.

Our algorithm is sound, i.e., it can never miss a violation of safety but may fail to verify some properties. The algorithm performs on-the-fly abstract interpretation by (1) abstracting away the order of packet processing and the number of times each packet arrives, (2) abstracting away correlations between states of different middleboxes and channel contents, and (3) representing middlebox states by their effect on each packet separately, rather than taking into account the entire state space. We show that the abstractions do not lose precision when middleboxes may reset in any state. This is encouraging since many real middleboxes reset, e.g., after some session timeout is reached or due to hardware failure.

1 Introduction

Modern computer networks are extremely complex, leading to many bugs and vulnerabilities that affect our daily life. Therefore, network verification is an increasingly important topic addressed by the programming languages and networking communities [17,5,15,16,14,30,23,12]. Previous network verification tools leverage a simple network forwarding model, which renders the datapath *immutable*. That is, normal packets going through the network do not change its forwarding behaviour, and the control plane explicitly alters the forwarding state at relatively slow time scales.

While the notion of an immutable datapath supported by an assemblage of routers makes verification tractable, it does not reflect reality. *Middleboxes* are widespread in modern enterprise networks [31]. A simple example of a middlebox is a stateful firewall which permits traffic from untrusted hosts only after they have received a packet from a trusted host. Middleboxes, such as firewalls, WAN optimizers, transcoders, proxies,




Fig. 1: A middlebox chain with a buggy topology.

load-balancers and the like, are the most common way to insert new functionality in the network datapath, and are commonly used to improve network performance and security. Middleboxes maintain a state and may change their state and forwarding behavior in response to packet arrivals. While useful, middleboxes are a common source of errors in the network [27].

As a simple example, consider the middlebox chain described in Fig. 1. In this network, a firewall is used to ensure that low security hosts (l_1, \dots, l_m) do not receive packets from the S_h server, and a cache and load balancer are used to improve performance. Unfortunately, the configuration of the network is incorrect since the cache may respond with a stored packet, bypassing the security policy enforced by the firewall. Swapping the order of the cache and the firewall results in a correct configuration.

Safety of Stateful Networks. We address the problem of verifying safety of networks with middleboxes, referred to as *stateful networks*. We target verification of *isolation* properties, namely, that packets sent from one host (or class of hosts) can never reach another host (or class of hosts). Yet, our approach is sound for any safety property. For example, it detects the safety violation described in Fig. 1, and verifies the safety of the correct configuration of this network.

Our focus is on verifying the configuration of stateful networks, i.e., addressing errors that arise from the interactions between middleboxes, and not from the complexity of individual middleboxes. Hence, we follow [35] and use an abstraction of middleboxes as finite-state programs. Previous work [35,32] has shown that many kinds of middleboxes, including proxy, cache proxy, NAT, and various kinds of load-balancers can be modeled in this way, sometimes using non-determinism to over-approximate the behaviour, e.g. to model timers, counters, etc. Since we are interested in safety properties, such an abstraction (overapproximation) is suitable.

As shown in [35], it is undecidable to check safety properties in general and isolation in particular, even for middleboxes with a finite state space, and even when the order of packets pending for each middlebox is abstracted away the complexity is quite high (EXPSpace-complete). Therefore, in this paper we develop additional abstractions for scaling up the verification.

Our approach. This paper makes a first attempt to apply abstract interpretation [7] to automatically prove the safety of stateful networks. Our approach combines sound network-level abstractions and middlebox-level abstractions that, together, make the verification task tractable. Roughly speaking, we apply (i) order abstraction [35], abstracting away the order of packets on channels, (ii) counter abstraction [26], abstracting

away their cardinality, (iii) network-level Cartesian abstraction [7,11,13], abstracting away the correlation between the states of different middleboxes and different channel contents, and (iv) middlebox-level Cartesian abstraction, abstracting away the correlation between states of different packets within each middlebox.

The network-level abstractions, (i)-(iii), lead to a chaotic iteration algorithm that is polynomial in the state space of the individual middleboxes and packets. However, the number of middlebox states can be exponential in the size of the network. For example, a firewall may record the set of trusted hosts and thus its states are subsets of hosts. Therefore, the resulting analysis is exponential in the number of hosts⁶.

The middlebox-level Cartesian abstraction, (iv), is the key to reducing the complexity to polynomial. The crux of this abstraction is the observation that the abstraction of middleboxes as reactive processes that query and update their state in a restricted way (e.g., [35]) allows to represent a middlebox state as a product of loosely-coupled *packet states*, one per potential packet. This lets us define a novel, non-standard, semantics of middlebox programs that we call *packet effect semantics*. The packet effect semantics is equivalent (bisimilar) to the natural semantics. However, while the natural semantics is monolithic, the packet effect semantics decomposes a single middlebox state into the parts that determine the forwarding behavior of different packets, and therefore facilitates the use of Cartesian abstraction to further reduce the complexity.

One of the main challenges for abstract interpretation is evaluating its precision. To address this challenge, we provide sufficient conditions that ensure precision of our analysis. Namely, we show that if the network is safe in the presence of packet re-ordering and middlebox reverts, where a middlebox may revert to its initial state at any moment, then our analysis is guaranteed to be precise, and will never report false alarms. This is, to a great extent, due to the packet effect semantics, which allows to use a middlebox-level Cartesian abstraction without incurring additional precision loss for such networks. Notice that middlebox reverts enable modelling arbitrary hardware failures, which have not been addressed by previous work on stateful network verification (e.g., in [35]). Surprisingly, verification becomes easier under the assumption that middleboxes may reset at any time. (Recall that for arbitrary unordered networks safety checking is EXPSPACE-complete.)

In summary, the main contributions of this paper are

- We introduce the first abstract interpretation algorithm for verifying safety of stateful networks, whose time complexity is polynomial in the size of the network, albeit exponential in the maximal number of queries of the local state that a middlebox can do, which is often small even for complex middleboxes (up to 5 in our examples).
- We develop *packet effect semantics*, a non-standard semantics of middlebox programs that facilitates middlebox-level Cartesian abstraction, reducing the complexity of the abstract interpretation algorithm from exponential in the size of the network to polynomial without incurring any additional precision loss for unordered reverting networks.

⁶ Unfortunately, if the set of hosts is not fixed, the safety problem becomes undecidable (even under the unordered abstraction) [1]. This means that, in general, it is not possible to alleviate the dependency of the complexity on the hosts.

- We provide sufficient conditions for precision of the analysis that have a natural interpretation in the domain of stateful networks: ignoring the order of packet processing and letting middleboxes revert to their initial states at any time.
- We prove lower bounds on the complexity of safety verification in the presence of packet reordering and/or middlebox reverts, showing that our algorithm is essentially optimal.
- We implement our analysis and show that it scales well with the number of hosts and middleboxes in the network.

We defer proofs of key claims to the extended version of this paper [1].

2 Expressing Middlebox Effects

This section defines our programming language for modeling the abstract behavior of middleboxes in the network. Our modeling language is independent of the particular network topology, which is defined in Sec. 3. The proposed language, AMDL (Abstract Middlebox Definition Language), is a restricted form of OCCAM [29], similar to the languages of [35,32].

We first define the syntax and informal semantics of AMDL (Sec. 2.1); we then define a formal “standard” *relation effect semantics* (Sec. 2.2); we continue by defining an alternative *packet effect semantics* (Sec. 2.3), which is bisimilar to the relation effect semantics (Sec. 2.4); and finally we present a localized version of the packet effect semantics (Sec. 2.5), which is suitable for Cartesian abstraction.

Packets. Middlebox behavior in our model is defined with respect to packets that consist of a fixed, finite, number of packet fields, ranging over finite domains. As such, a packet $p \in P$ in our formalism is a tuple of packet fields over predefined finite sorts. In our examples, a packet is a tuple $\langle s, d, t \rangle$, where s, d are the source and destination hosts, respectively, taken from a finite set of hosts H , and t is a packet tag (or type) that ranges over a finite domain T . In this case, $|P|$ is polynomial in $|H|$. (Our approach is also applicable when additional fields are added, e.g., for modeling the packet’s payload via an abstract finite domain.)

2.1 Syntax and Informal Semantics

Fig. 3 describes the syntax of the AMDL language⁷. Middleboxes are implemented as reactive processes, with events triggered by the arrival of packets. If multiple packets are pending, the AMDL process non-deterministically reads a packet from one of the incoming channels of the process. The packet processing code is a loop-free block of guarded-commands, which may update relations and forward potentially modified packets to some of the output ports. AMDL uses *relations* over finite domains to store the middlebox state. These are the only data structures allowed in AMDL. The only relation operations allowed are inserting a value to a relation, removing a value from a relation, and *membership queries* — checking whether a value is in a relation. For a

⁷ In the code examples, we write p for the triple $(src, dst, type)$ and use access path notation to refer to the fields, e.g., $p.src$.

```

sfirewall = do
  internal_port ? p =>
    if
      p.dst in trusted => external_port ! p
    □
    p.type = 0 => // request packet
      external_port ! p;
      requested(p.dst) := true
    fi
  □
  external_port ? p =>
    if
      p.src in trusted => internal_port ! p
    □
    p.type = 1 and p.src in requested =>
      // response packet with a request
      trusted(p.src) := true
    fi
  od

```

Fig. 2: AMDL code for session firewall.

membership query of the form \bar{a} in r , we denote the relation, r , used in the query by $rel(q)$ and denote the tuple of atoms \bar{a} by $atoms(q)$. For example, the code for a session firewall is depicted in Fig. 2.

Middleboxes may enforce safety properties using the **abort** command. For example, an isolation middlebox would abort when a forbidden packet is received.

2.2 Middlebox Relation Effect Semantics

We now sketch the semantics of AMDL. The definitions below supply a part of the full network semantics, which is given in Sec. 3.

Middlebox States. Each middlebox $m \in M$ maintains its own local state as a set of relations. The domain of a relation r defined over sorts $s_{1..k}$ is $D(r) \stackrel{\text{def}}{=} D(s_1) \times \dots \times D(s_k)$, where $D(s_i)$ is the domain of sort s_i . We use $rels(m)$ to denote the set of relations in m , and $D(m)$ to denote the union of $D(r)$ over $r \in rels(m)$.

The *middlebox state* of m is then a function $s \in \Sigma^R[m] \stackrel{\text{def}}{=} rels(m) \rightarrow \wp(D(m))$, mapping each $r \in rels(m)$ to $v \subseteq D(r)$. In addition, we introduce a unique *error* middlebox state, denoted err . We assume that $err \in \Sigma^R[m]$ for every middlebox m .

Middlebox Transitions. Middlebox transitions have the form

$$\frac{(p,c)/(p_i,c_i)_{i=1..k}}{\rightarrow_{R} \subseteq \Sigma^R[m] \times \Sigma^R[m]}$$

where (p,c) denotes packet-channel at the input, and $(p_i,c_i)_{i=1..k}$ is the sequence of packet-channel pairs that the middlebox outputs.

$$\begin{aligned}
\langle mbox \rangle & ::= m = \mathbf{do} \langle pblock \rangle [\square \langle pblock \rangle]^* \mathbf{od} \\
\langle pblock \rangle & ::= c ? \overline{pfld} \Rightarrow \langle gc \rangle \\
\langle gc \rangle & ::= \langle cond \rangle \Rightarrow \langle action \rangle \mid \mathbf{if} \langle gc \rangle [\square \langle gc \rangle]^* \mathbf{fi} \\
\langle action \rangle & ::= \langle action \rangle ; \langle action \rangle \mid c ! \overline{\langle atom \rangle} \mid r(\overline{\langle atom \rangle}) := \langle cond \rangle \mid \mathbf{abort} \\
\langle cond \rangle & ::= \mathbf{true} \mid \langle cond \rangle \mathbf{and} \langle cond \rangle \mid \mathbf{not} \langle cond \rangle \mid \langle atom \rangle = \langle atom \rangle \mid \overline{\langle atom \rangle} \mathbf{in} r \\
\langle atom \rangle & ::= \overline{pfld} \mid \mathit{const}
\end{aligned}$$

Fig. 3: AMDL syntax. \bar{e} denotes a comma-separated list of elements drawn from the domain e . **abort** imposes a safety condition. $c ? p$ reads p from a channel c and $c ! p$ writes p into c . We write m for a middlebox name, r for a relation name, and c for a channel name. We write const for a constant symbol and \overline{pfld} for identifiers used to match fields in packets, e.g., `src`. Non-deterministic choice is denoted by \square .

For example, for $s \stackrel{\text{def}}{=} [\text{requested} \mapsto \emptyset, \text{trusted} \mapsto \emptyset]$, the guarded command corresponding to the internal port of the firewall middlebox (Fig. 2) induces a transition $s \xrightarrow{((h_1, h_2, 0), \overline{c_{in}}) / ((h_1, h_2, 0), \overline{c_{out}})}_R s'$ where $s' \stackrel{\text{def}}{=} [\text{requested} \mapsto \{h_2\}, \text{trusted} \mapsto \emptyset]$.

abort commands induce transitions to the *err* state.

The formal definition of the middlebox transitions appears in the extended version of this paper [1].

2.3 Middlebox Packet Effect Semantics

We now present a semantics that is equivalent to the relation effect semantics. The semantics is based on an alternative (yet isomorphic) representation of middlebox states that reveals a loose coupling between the parts of the state that are relevant for different packets. This loose coupling then facilitates a Cartesian abstraction that abstracts away correlations between packets in the same state.

Packet Effect Representation of Middlebox State Recall that in Sec. 2.1 we restrict the values that can be used in a middlebox program to either constants or the values of fields of the currently processed packet. We do not allow extracting tuples from the relation (e.g., by having a `get` command, or by iterating over the contents of the relation). Instead, we limit the interaction with the relation to checking whether a tuple (that consists of packet fields or constants) exists in the relation. Consequently, instead of storing the contents of all relations, the state of the middlebox can be represented by mapping all potential packets in the network to their effect on the middlebox. Specifically, we map each packet and membership query in the program to whether that membership query will be evaluated to **True** when the program is executed on that packet.

For every middlebox m , we denote by $Q(m)$ the set of membership queries in m 's program. (We need not distinguish between different instances of the same query.) For example, in Fig. 2, $Q(fw) = \{\text{p.dst in trusted}, \text{p.src in trusted}, \text{p.src in requested}\}$.

The *packet effect state* of a middlebox m is a function $s \in \Sigma^P[m] \stackrel{\text{def}}{=} P \rightarrow Q(m) \rightarrow \{\text{True}, \text{False}\}$, mapping each packet $p \in P$ to the evaluation of all queries of m when p is the input packet, thus capturing the way in which p traverses m 's program. We refer to $s(p) \in Q(m) \rightarrow \{\text{True}, \text{False}\}$ as the *packet state* of packet p in middlebox state s . We extend $\Sigma^P[m]$ with an error state $\lambda p \in P. \text{err}$, which is also denoted *err*.

Middlebox Transition Relation in the Packet Space The semantics of middlebox m in the packet space is defined via a transition relation $\frac{(p,c)/(p_i,c_i)_{i=1..k}}{\rightarrow_{P,m} \subseteq \Sigma^P[m] \times \Sigma^P[m]}$. When m is clear, we omit it from the notation. A transition $\tilde{s} \xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}} \tilde{s}'$ exists if (one of) the sequence of operations applied on \tilde{s} when packet p arrives on channel c outputs $(p_i, c_i)_{i=1..k}$ and leads to \tilde{s}' .

The semantics of operations is defined similarly to the “standard” relation effect semantics. The semantics of error and output actions (that do not change the middlebox state) is straightforward. Next, we explain the semantics of the operations that depend on or change the middlebox state — membership queries and relation updates.

Consider a membership query q . Let \tilde{s} be the middlebox state before evaluating q , i.e., \tilde{s} is the state that results from executing all previous relation updates, and let p be the packet that invoked the middlebox transition. Then q is evaluated to $\tilde{s}(p)(q)$.

Next, consider a relation update. A relation update $r(\bar{a}) := \text{cond}$ updates the packet states of all packets that are affected by the operation. This is done as follows. As before, let \tilde{s} be the intermediate state of m right before executing the operation, and let p be the packet that the middlebox program is operating on. Consider the case where cond evaluates to **True** in \tilde{s} , corresponding to addition of a value. (Removal of a value is symmetric.) We denote by $\bar{a}(p)$ the result of substituting each field name in \bar{a} by its value in p . That is, $\bar{a}(p) \in D(r)$ is the value being added to r . This addition may affect the value of membership queries $q \in Q(m)$ with $\text{rel}(q) = r$ (querying the same relation r) for other packets \tilde{p} as well, in case that $\text{atoms}(q)(\tilde{p})$, i.e., the value being queried on \tilde{p} , is the same as the value $\bar{a}(p)$ being added to r . Therefore, the intermediate state obtained after the relation update operation has been applied is

$$\tilde{s}' = \lambda \tilde{p} \in P. \lambda q \in Q(m). \begin{cases} \text{True}, & \text{if } \text{rel}(q) = r \wedge \text{atoms}(q)(\tilde{p}) = \bar{a}(p). \\ \tilde{s}(\tilde{p})(q), & \text{otherwise.} \end{cases}$$

Namely, the operation updates to **True** the value of queries that coincide with the tuple of elements inserted to the relation.

Example 1. Consider the packet effect state $\tilde{s} \stackrel{\text{def}}{=} \lambda p. \lambda q. \text{False} \in \Sigma^P[\text{fw}]$ of the firewall (Fig. 2), where q ranges over the three membership queries in the code. Upon reading the packet $(h_1, h_2, 0)$ from an internal port, the middlebox performs a sequence of internal transitions which includes evaluating the expression “ $p.\text{type}=0$ ” to **True**, outputting the packet $(h_1, h_2, 0)$ to the output port, and executing the command `requested(p.dst) := true`, which results in updating the state to:

$$\tilde{s}' \stackrel{\text{def}}{=} \lambda \tilde{p}. \lambda q. \begin{cases} \text{True}, & \text{if } \text{rel}(q) = \text{requested} \wedge \text{atoms}(q)(\tilde{p}) = h_2 \\ \text{False}, & \text{otherwise.} \end{cases}$$

That is, $\tilde{s}'((h_2, *, *))(\text{p.src in requested}) = \text{True}$ and all the other values in \tilde{s}' remain **False** as before. Therefore, $\tilde{s} \xrightarrow{((h_1, h_2, 0), c_{in}) / ((h_1, h_2, 0), c_{out})} \tilde{s}'$. \square

2.4 Bisimulation of Packet Effect Semantics and Relation Effect Semantics

We continue by showing that the transition systems defining the semantics of middleboxes in the packet effect and in the relation effect representations are bisimilar.

To do so, we first define a mapping $ps: \Sigma^R[m] \rightarrow \Sigma^P[m]$ from the relation state representation to the packet effect state representation. Recall that the relation state representation of middlebox states is $s \in \Sigma^R[m] \stackrel{\text{def}}{=} \text{rels}(m) \rightarrow \wp(D(m))$. Given a state $s \in \Sigma^R[m]$, ps maps it to the packet effect state s^P defined as follows:

$$s^P \stackrel{\text{def}}{=} \lambda \tilde{p} \in P. \lambda q \in Q(m). \text{atoms}(q)(\tilde{p}) \in s(\text{rel}(q)).$$

That is, for every input packet \tilde{p} , the value in s^P of the query $q \in Q(m)$ is equal to the evaluation of the same query in s based on an input packet \tilde{p} .

Definition 1 (Bisimulation Relation). For a middlebox m , we define the relation $\sim_m \subseteq \Sigma^R[m] \times \Sigma^P[m]$ as the set of all pairs (s, s^P) such that $s = s^P = \text{err}$ or $ps(s) = s^P$.

Lemma 1. Let $s \in \Sigma^R[m]$ and $\tilde{s} \in \Sigma^P[m]$ and $s \sim_m \tilde{s}$. Then the following holds:

- For every state $s' \in \Sigma^R[m]$, if $s \xrightarrow{(p,c)/o} s'$ then there exists a state $\tilde{s}' \in \Sigma^P[m]$ s.t. $\tilde{s} \xrightarrow{(p,c)/o} \tilde{s}'$ and $s' \sim_m \tilde{s}'$, and
- For every state $\tilde{s}' \in \Sigma^P[m]$ if $\tilde{s} \xrightarrow{(p,c)/o} \tilde{s}'$ then there exists a state $s' \in \Sigma^R[m]$ s.t. $s \xrightarrow{(p,c)/o} s'$ and $s' \sim_m \tilde{s}'$.

2.5 Locality of Packet-Effect Middlebox Transitions

In this section we present a locality property of the packet effect semantics that will allow us to efficiently compute an abstract transformer when applying a Cartesian abstraction. Namely, we observe that an execution of an operation $r(\bar{a}) := \text{cond}$, in the context of processing an input packet p , potentially updates the packet states of all packets. However, for each packet \tilde{p} , the updated packet state $\tilde{s}'(\tilde{p})$ depends only on its pre-state $\tilde{s}(\tilde{p})$, the input channel c , the input packet p , and $\tilde{s}(p)$, which determines the value of queries; it is completely independent of the packet states of all other packets. Since, in addition, the execution path of the middlebox when processing input packet p depends only on the packet state of p , this form of *locality*, which we formalize next, extends to entire middlebox programs.

Definition 2 (Substate). Let $\tilde{s} \in P \rightarrow Q(m) \rightarrow \{\text{True}, \text{False}\}$ be a packet effect state. We denote by $\tilde{s}|_{\{p, \tilde{p}\}} \in \{p, \tilde{p}\} \rightarrow Q(m) \rightarrow \{\text{True}, \text{False}\}$ the substate obtained from \tilde{s} by dropping all packet states other than those of p and \tilde{p} . Let $\Sigma^P[m, p, \tilde{p}] \stackrel{\text{def}}{=} \{p, \tilde{p}\} \rightarrow Q(m) \rightarrow \{\text{True}, \text{False}\}$ denote the set of substates for p and \tilde{p} .

Definition 3 (Substate transition relation). We define the substate transition relation $\xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}}_{P[p,\tilde{p}]}: \Sigma^P[m, p, \tilde{p}] \times \Sigma^P[m, p, \tilde{p}]$ as follows. A substate transition $\tilde{s}[p, \tilde{p}] \xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}}_{P[p,\tilde{p}]} \tilde{s}[p, \tilde{p}]'$ holds if there exist \tilde{s} and \tilde{s}' such that $\tilde{s}|_{[p,\tilde{p}]} = \tilde{s}[p, \tilde{p}]$, $\tilde{s}'|_{[p,\tilde{p}]} = \tilde{s}[p, \tilde{p}]'$ and $\tilde{s} \xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}}_P \tilde{s}'$.

The locality of AMDL programs manifests itself in the ability to compute the substate transition relation, $\xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}}_{P[p,\tilde{p}]}$, directly from the code (without first computing the transition relation and then using projection). This property will be important later to efficiently compute a network-level abstract transformer (Sec. 4.1):

Lemma 2 (2-Locality). Given $\tilde{s}[p, \tilde{p}]$ and $\tilde{s}[p, \tilde{p}]'$, checking whether

$$\tilde{s}[p, \tilde{p}] \xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}}_{P[p,\tilde{p}]} \tilde{s}[p, \tilde{p}]'$$

can be done in time linear in the size of the middlebox program.

3 Network Semantics

This section defines the semantics of stateful networks by defining the semantics of packet traversal over communication channels in the network, and the transitions between network configurations. We first define a concrete semantics, followed by two relaxations: unordered semantics and reverting semantics. These relaxations provide sufficient conditions for completeness of the abstract interpretation performed in Sec. 4. **Network Topology.** A network N is a finite bidirected⁸ graph of *hosts* and *middleboxes*, equipped with a *packet domain*. Formally, $N = (H \cup M, E, P)$, where:

- P is a set of packets.
- H is a finite set of *hosts*. A *host* $h \in H$ consists of a unique identifier and a set of packets $P_h \subseteq P$ that it can send.
- M is a finite set of *middleboxes*. A *middlebox* $m \in M$ is associated with a set of communication channels C_m .
- $E \subseteq \{\langle h, c_m, m \rangle, \langle m, c_m, h \rangle \mid h \in H, m \in M, c_m \in C_m\} \cup \{\langle m_1, c_{m_1}, c_{m_2}, m_2 \rangle \mid m_1, m_2 \in M, c_{m_1} \in C_{m_1}, c_{m_2} \in C_{m_2}\}$ is the set of directed communication channels in the network, each connecting a communication channel $c_{m_1} \in C_{m_1}$ of middlebox m_1 either to a host, or to a communication channel $c_{m_2} \in C_{m_2}$ of middlebox m_2 . For e of the form $\langle m, c_m, h \rangle$ or $\langle m, c_m, c_{m_2}, m_2 \rangle$, we say that e is an *egress* channel of middlebox m connected to channel c_m and an *ingress* channel of host h , respectively middlebox m_2 , connected to channel c_{m_2} .

The network semantics is parametric in the middlebox semantics. It considers the semantics of a middlebox $m \in M$ to be a transition system with a finite set of states $\Sigma[m]$, an initial state $\sigma_I(m) \in \Sigma[m]$ and a set of transitions $\xrightarrow{(p,c)/(p_i,c_i)_{i=1..k}}_{\subseteq} \Sigma[m] \times \Sigma[m]$. This can be realized with either the relation effect semantics or the packet effect semantics defined in Sec. 2.2 and Sec. 2.3, respectively.

⁸ A *bidirected graph* is a directed graph in which every edge has a matching edge in the opposite direction. i.e., $(u, v) \in E \iff (v, u) \in E$.

3.1 Concrete (Ordered) Network Configurations

All variants of the network semantics defined in this section are defined over the same set of configurations. Let $\Sigma[M] \stackrel{\text{def}}{=} \bigcup_{m \in M} \Sigma[m]$ denote the set of middlebox states of all middleboxes in a network. An *ordered network configuration* $(\sigma, \pi) \in \Sigma = (M \rightarrow \Sigma[M]) \times (E \rightarrow P^*)$ assigns middleboxes to their (local) middlebox states and communication channels to sequences of packets. The sequence of packets on each channel represents all packets sent from the source and not yet processed by the destination.

Initial Configuration. We denote the ordered initial configuration by $(\sigma_I, \lambda e \in E . \epsilon)$, where $\sigma_I: M \rightarrow \Sigma[M]$ denotes the initial state of all middleboxes.

Error Configurations. We say that a configuration is an *error configuration* if any of its middleboxes is in the error state. We denote all error configurations by *err*.

3.2 Concrete (FIFO) Network Semantics

We first consider the First-In-First-Out (FIFO) network semantics, under which communication channels retain the order in which packets were sent.

Ordered Network Transitions. The network semantics is defined via *middlebox transitions* and *host transitions*.

A middlebox transition is $(\sigma, \pi) \xrightarrow[p, e, m]{}_o (\sigma', \pi')$ where the following holds: (i) p is the *first* packet on the channel $e \in E$, (ii) the channel e is an ingress channel of middlebox m connected to channel $c \in C_m$, (iii) $\sigma(m) \xrightarrow[(p, c) / (p_i, c_i)_{i=1..k}]{} \sigma'(m)$, meaning that $\sigma'(m)$ is the result of updating $\sigma(m)$ according to the middlebox semantics, (iv) the channels e_i are egress channels of middlebox m connected to the channels $c_i \in C_m$, (v) π' is the result of removing packet p from (the head of) channel e and appending p_i to the tails of the appropriate channels e_i , and (vi) the states of all other middleboxes equal their states in σ .

A host transition is $(\sigma, \pi) \xrightarrow[h, e, p]{}_o (\sigma, \pi')$ where one of the following holds:

Packet Production (i) the channel e is an egress channel of host h , (ii) $p \in P_h$ is a packet sent by h , and (iii) π' is the result of appending p to the tail of e ; or

Packet Consumption (i) the channel e is an ingress channel of host h , (ii) p is the first packet on the channel e , and (iii) π' is the result of removing p from the head of e .

We denote the ordered transition relation obtained by the union of all middlebox and host transitions by \Rightarrow_o . It is naturally lifted to a concrete transformer $\mathcal{T}^o: \wp(\Sigma) \rightarrow \wp(\Sigma)$ defined as:

$$\mathcal{T}^o(X) \stackrel{\text{def}}{=} \{(\sigma', \pi') \mid (\sigma, \pi) \in X \wedge (\sigma, \pi) \Rightarrow_o (\sigma', \pi')\} .$$

Collecting Semantics. The ordered collecting semantics of a network N is the set of configurations reachable from the initial configuration.

$$\llbracket N \rrbracket^o \stackrel{\text{def}}{=} \text{LeastFixpoint}(\mathcal{T}^o)(\sigma_I, \lambda e \in E . \epsilon) = \bigcup_{i=1}^{\infty} (\mathcal{T}^o)^i(\sigma_I, \lambda e \in E . \epsilon) .$$

Definition 4 (Safety Verification Problem). *For a network N and initial state σ_I for the middleboxes, the safety verification problem is to determine whether an error configuration is reachable from the initial configuration. That is, whether $err \in \llbracket N \rrbracket^o$.*

Theorem 1. [35] *The safety verification problem for ordered networks is undecidable.*

In this work, we tackle the undecidability of verification by developing a sound abstract interpretation that can be used to check the safety of networks. Before doing so, we present two relaxed network semantics that motivate the abstractions we employ, and also provide sufficient conditions for their completeness.

3.3 Unordered and Reverting Network Semantics

The “unordered” semantics allows channels to not preserve the packet transmission order. Namely, packets in the same channel may be processed in a different order than the order in which they were received. The “reverting” semantics allows middleboxes to revert to their initial state after every transition. Formally, these relaxed semantics extend the set of network transitions (and consequently, the transformer and the collecting semantics) with reordering transitions and reverting transitions, respectively.

A *reordering transition* has the form $(\sigma, \pi) \xrightarrow{e} (\sigma, \pi')$ where for the channel $e \in E$, $\pi'(e)$ is a permutation of $\pi(e)$ and for all other channels $e' \neq e$, $\pi'(e') = \pi(e')$.

A *reverting transition* has the form $(\sigma, \pi) \xrightarrow{m} (\sigma', \pi)$ where for the middlebox $m \in M$, $\sigma'(m) = \sigma_I(m)$ and for all other middleboxes $m' \neq m$, $\sigma'(m') = \sigma(m')$.

The *unordered network transitions* consist of the ordered transitions as well as the reordering transitions; the *ordered reverting transitions* consist of the ordered transitions and the reverting transitions; and the *unordered reverting transitions* consist of all of the above. We denote the corresponding collecting semantics by $\llbracket N \rrbracket^u$, $\llbracket N \rrbracket^{or}$ and $\llbracket N \rrbracket^{ur}$, respectively. Clearly,

$$\llbracket N \rrbracket^o \subseteq \llbracket N \rrbracket^u \subseteq \llbracket N \rrbracket^{ur} \quad \text{and} \quad \llbracket N \rrbracket^o \subseteq \llbracket N \rrbracket^{or} \subseteq \llbracket N \rrbracket^{ur}$$

By plugging-in the two representations of middleboxes in the definition of the network semantics, we obtain two variants of the network semantics for each of the four variants considered so far. In the sequel, we use a *pa* subscript to refer to the packet effect semantics, and no subscript to refer to the relation effect semantics. The bisimulation between middlebox representations is lifted to a bisimulation between each relation state network semantics and the corresponding packet state network semantics. Therefore, the following holds:

Lemma 3. *For every semantic identifier $i \in \{o, u, or, ur\}$, $err \in \llbracket N \rrbracket^i$ iff $err \in \llbracket N \rrbracket_{pa}^i$.*

The safety verification problem is adapted for the different variants of the network semantics. The following theorem summarizes the complexity of the obtained problems. (We do not distinguish the packet effect semantics from the relation effect semantics, since due to Lem. 3 they induce the same safety verification problem.)

Theorem 2. *The safety verification problem is*

- (i) *EXPSACE-complete for unordered networks [35].*
- (ii) *undecidable for ordered reverting networks.*
- (iii) *coNP-hard for unordered reverting networks.*

Thm. 2(ii) justifies the need for the unordered abstraction even in reverting networks. Thm. 2(iii) implies that our abstract interpretation algorithm, presented in Sec. 4, which is both sound and complete for the unordered reverting semantics, is essentially optimal since it essentially meets the lower bound stated in the theorem (it is exponential in the number of state queries of any middlebox and polynomial in the number of middleboxes, hosts and packets).

Sticky Properties. Unordered reverting networks have a useful property of *sticky packets*, meaning that if a packet is pending for a middlebox in some run of the network then any run has an extension in which the packet is pending again with multiplicity $> n$, for any $n \in \mathbb{N}$. This property implies a stronger property:

Lemma 4 (Sticky Packet States Property). *For every channel e , packets p, \tilde{p} , middlebox m and packet state \tilde{v} of \tilde{p} in m : If, in some reachable configuration, channel e contains p and in some (possibly other) reachable configuration the packet state of \tilde{p} in m is \tilde{v} , then there exists a reachable configuration where simultaneously e contains p and the packet state of \tilde{p} in m is \tilde{v} .*

Intuitively, Lem. 4 follows from the fact that all middleboxes can revert to their initial state and the unordered semantics enables a scenario where the particular state and packets are reconstructed. It ensures that ignoring the correlation between the packet states of a middlebox for different packets, the packet states across different middleboxes, and the occurrence (and cardinality) of packets on channels does not incur any precision loss w.r.t. safety. This makes the network-level abstraction defined in Sec. 4, which treats channels as sets of packets and ignores correlations between packet states and channels, precise.

4 Abstract Interpretation for Stateful Networks

In this section, we present our algorithm for safety verification of stateful networks based on abstract interpretation of the semantics $\llbracket \mathbb{N} \rrbracket_{pa}^o$, and discuss its guarantees.

4.1 Abstract Interpretation for Packet Space

We apply sound abstractions to different components of the concrete packet state network domain. Due to space constraints, we do not describe the intermediate steps in the construction of the abstract domain, and only present the final domain used by the analysis. Roughly speaking, the obtained domain abstracts away (i) the order and cardinality of packets on channels; (ii) the correlation between the states of different middleboxes and different channel contents; and (iii) the correlation between states of different packets within each middlebox.

Cartesian Packet Effect Abstract Domain. Let $Q \rightarrow \{T, F\}$ denote the union of $Q(m) \rightarrow \{T, F\}$ over all middleboxes $m \in M$, including the error state *err*. The

Cartesian abstract domain of the packet state of the network is given by the lattice $\mathcal{A} \stackrel{\text{def}}{=} (A, \perp, \sqsubseteq, \sqcup)$, where $A \stackrel{\text{def}}{=} (M \rightarrow P \rightarrow \wp(Q \rightarrow \{T, F\})) \times (E \rightarrow \wp(P))$. That is, an abstract element maps each packet in each middlebox to a set of possible valuations for the queries, and each channel to a set of packets. The bottom element is $\perp \stackrel{\text{def}}{=} (\lambda m. \lambda p. \emptyset, \lambda e. \emptyset)$, the partial order $a_1 \sqsubseteq a_2$ is defined by pointwise set inclusions per middlebox and channel, and join is defined by pointwise unions $(\omega_1, \omega_2) \sqcup (\omega'_1, \omega'_2) \stackrel{\text{def}}{=} (\lambda m. \lambda p. \omega_1(m)(p) \cup \omega'_1(m)(p), \lambda e. \omega_2(e) \cup \omega'_2(e))$.

Let $\mathcal{C} \stackrel{\text{def}}{=} (\wp(\Sigma^P), \subseteq)$ be the concrete network domain. We define the Galois connection $(\mathcal{C}, \gamma, \alpha, \mathcal{A})$ as follows. The abstraction function $\alpha : \wp(\Sigma^P) \rightarrow A$ for a set of packet state configurations $X \subseteq \Sigma^P$ is defined as $\alpha(X) = (\omega_{mboxes}, \omega_{chans})$ where

$$\omega_{mboxes} = \lambda m. \lambda p. \{ \sigma(m)(p) \mid (\sigma, \pi) \in X \} \text{ and } \omega_{chans} = \lambda e. \bigcup_{(\sigma, \pi) \in X} \pi(e) .$$

The concretization function $\gamma : A \rightarrow \wp(\Sigma^P)$ is induced by α and \sqsubseteq . We denote the initial abstract element as $a_I = \alpha(\{(\sigma_I, \lambda e \in E . \emptyset)\})$.

Abstract Transformer. Next, we define the abstract transformer $\mathcal{T}^\sharp : A \rightarrow A$, which soundly abstracts the concrete transformer \mathcal{T}° and show that it is efficient, due to the locality property of middlebox transitions. We use the predicate $in(c, e, m)$ to denote that the network channel e is an ingress channel of middlebox m , connected to its c channel. Similarly, $out(c, e, m)$ means that e is an egress channel of m connected to its c channel. Further, let $[x_1 \mapsto y_1, \dots, x_n \mapsto y_n]$ denote a mapping from each x_i to y_i for $i = 1..n$ and $f[x \mapsto y]$ denote the function f updated by (re-)mapping x to y .

Definition 5. Let $(\omega_1, \omega_2) \in (M \rightarrow P \rightarrow \wp(Q \rightarrow \{T, F\})) \times (E \rightarrow \wp(P))$ be an abstract element. Then $\mathcal{T}^\sharp(\omega_1, \omega_2) \stackrel{\text{def}}{=}$

$$\bigsqcup \left\{ \begin{array}{l} (\omega_1[m \mapsto \tilde{p}\tilde{s}], \\ \omega_2[e_i \mapsto \omega_2(e_i) \cup \{p_i\}]) \end{array} \left| \begin{array}{l} (1) m \in M, \\ (2) p \in \omega_2(e), in(c, e, m), \\ (3) \tilde{s} \in \omega_1(m), \tilde{p} \in P, \\ \tilde{s}[p, \tilde{p}] = [p \mapsto \tilde{s}(p), \tilde{p} \mapsto \tilde{s}(\tilde{p})], \\ (4) \tilde{s}[p, \tilde{p}] \xrightarrow{(p, c)/(p_i, c_i)_{i=1..k}} P_{[p, \tilde{p}]} \tilde{s}[p, \tilde{p}]', \\ (5) \tilde{p}s = \tilde{s}[\tilde{p} \mapsto \{ \tilde{s}[p, \tilde{p}]'(\tilde{p}) \}], \\ (6) out(c_i, e_i, m), i = 1..k \end{array} \right. \right\} .$$

Intuitively, the transformer updates the abstract state by joining the individual effects obtained by: (1) considering each middlebox, (2) considering each input packet to the middlebox, (3) considering every possible substate for the input packet p and every other packet \tilde{p} , (4) considering every possible substate transition, (5) adding the new packet state for \tilde{p} to the relevant set, and (6) adding each output packet to the corresponding edge.

Proposition 1. The running time of \mathcal{T}^\sharp is $O((|M| + |E|) \cdot |P|^2 \cdot 2^{2|Q_{max}|})$, where Q_{max} denotes the maximal set of queries $Q(m)$ over all middleboxes $m \in M$.

Our algorithm for safety verification computes $\mu^\sharp \stackrel{\text{def}}{=} \text{LeastFixpoint}(\mathcal{T}^\sharp)(a_I) = \bigsqcup_{i=1}^{\infty} \mathcal{T}^{\sharp^i}(a_I)$ and checks whether $err \in \mu^\sharp$.

Complexity of Least Fixpoint Computation. The height of the abstract domain lattice is determined by the number of packets that can be added to the channels of the network— $(|P| \cdot |E|)$, multiplied by the number of state changes that can occur in any of the middleboxes— $O(|M| \cdot |P| \cdot 2^{|Q|})$. The time complexity of the abstract interpretation is bounded by the height of the abstract domain lattice multiplied by the time complexity of the abstract transformer:

$$O(|P|^4 \cdot |E| \cdot |M| \cdot 2^{3|Q_{max}|} \cdot (|M| + |E|)) .$$

4.2 Soundness and Completeness

Our algorithm is sound in the sense that it never misses an error state. This follows from the use of a sound abstract interpretation:

Theorem 3 (Soundness). $\llbracket \mathbb{N} \rrbracket_{pa}^o \subseteq \llbracket \mathbb{N} \rrbracket_{pa}^{ur} \subseteq \gamma(\mu^\sharp)$.

Our algorithm is also complete relative to the reverting unordered semantics.

Theorem 4 (Completeness). $\mu^\sharp \sqsubseteq \alpha(\llbracket \mathbb{N} \rrbracket_{pa}^{ur})$.

The proof of Thm. 4 relies on the sticky property formalized by Lem. 4. The theorem states that for reverting unordered networks μ^\sharp is at least as precise as applying the abstraction function on the concrete packet state network semantics. In particular, this implies that if μ^\sharp is an abstract error element then $err \in \llbracket \mathbb{N} \rrbracket_{pa}^{ur}$. As a result, for such networks our algorithm is a decision procedure. For other networks it may produce false alarms, if safety is not maintained by an unordered reverting abstraction.

Properties. Recall that we express safety properties via middleboxes in the network. Therefore, in unordered reverting networks, the possibility to revert applies to the safety property as well, and may introduce false alarms due to addition of behaviors leading to error. However, for safety properties such as isolation which are suffix-closed (i.e., all the suffixes of a safe run are themselves safe runs), this cannot happen[1].

5 Implementation and Initial Evaluation

In this section, we describe our implementation of the analysis described in Sec. 4, and report our initial experience running the algorithm on a few example networks.

Implementation. We have developed a compiler, `amd1c`, which takes as input a network topology and its initial state (given in `json` format) and AMDL programs for the middleboxes that appear in the topology. The compiler outputs a Datalog program, which can then be efficiently solved by a Datalog solver. Specifically, we use `LogicBlox` [3].

The generated Datalog programs include three relations: (i) `packetsSeen`, which stores the packets sent over the network channels; (ii) `middleboxState`, which stores the packet state of individual packets in each middlebox (i.e., the possible valuation of each middlebox program’s queries for each individual packet); and (iii) `abort`, which stores the middleboxes that have reached an `err` state.




Fig. 4: Topology of the datacenter example.

We encode the packets that hosts can send to their neighboring middleboxes and the initial state of the middleboxes as Datalog *facts* (edb), and the effects of the middlebox programs, i.e. relation update actions and packet output actions, as Datalog *rules* (idb).

We then use the datalog engine to compute the fixed point of the datalog program. That fixed point is exactly the least fixed point $\mu^\# \stackrel{\text{def}}{=} \text{LeastFixpoint}(\mathcal{T}^\#)(a_I) = \bigsqcup_{i=1}^{\infty} \mathcal{T}^{\#i}(a_I)$

Evaluation. The main challenge in acquiring realistic benchmarks is that middlebox configuration and network topology are considered security sensitive, and as a result enterprises and network operators do not release this information to the public. Consequently, we benchmarked our tool using the synthetic topologies and configurations described by [24].

Our benchmarks focus on datacenter networks and enterprise networks. The set of middleboxes we used in our datacenter benchmarks is based on information provided in [27], and on conversations with datacenter providers. We ran both a simple case where each tenant machine is protected by firewalls and an IPS (Intrusion Prevention System); and a more complex case where we use redundant servers and distribute traffic across them using a load balancer. Our enterprise topology is based on the standard topology used in a variety of university departments including UIUC (reported in [18]), UC Berkeley, Stanford, etc. which employ firewalls and an IP gateway.

We ran two scaling experiments, measuring how well our system scales when the number of hosts or the number of middleboxes in the network increases. The experiments were run on Amazon EC2 r4.16 instances with 64-core CPUs and 488GiB RAM.

Multi Tenant Datacenter Network. Fig. 4 illustrates the topology of a multi tenant datacenter. Each rack hosts a different tenant, and the safety property we wish to verify is isolation between the hosts of the two racks. In this example the network also employs an IPS to prevent malicious traffic from reaching the datacenter. Actual IPS code is too complex to be accurately modeled in AMDL; instead we over-approximate the behaviour of an IPS by modeling it as a process that non-deterministically drops incoming packets.

Enterprise Network. Fig. 5a illustrates the topology of an enterprise network. The enterprise network consists of three subnets, each with a different security policy. The *public* subnet is allowed unrestricted access with the outside network. The *quarantined* subnet is not allowed any communication with the outside network. The *private* subnet

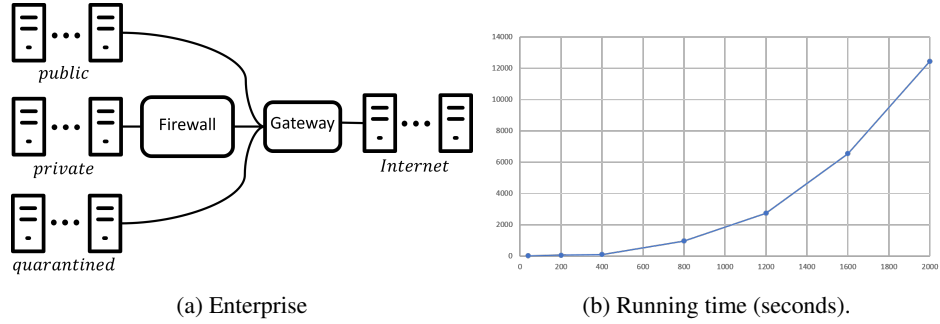


Fig. 5: Topology and running times of the host scalability test.

can initiate communication with a host in the outside network, but hosts in the outside network cannot initiate communication with the hosts in the *private* subnet.

To evaluate the feasibility of our solution, we ran the analysis of Fig. 5a on networks with varying numbers of hosts ranging from 20 to 2,000. Our implementation successfully verified a network with 2,000 hosts in under four hours, suggesting that the implementation could be used to verify realistic networks. Fig. 5b shows the times of the analysis on an enterprise network with 20–2,000 hosts.

Datacenter Middlebox Pipeline. Fig. 6a describes a datacenter topology with a pipeline of middleboxes connecting servers to the Internet. The topology contains multiple middlebox pipelines for load-balancing purposes and to ensure resiliency. We use this topology to test the scalability of our approach w.r.t the size of the network, by adding additional middlebox pipelines and keeping the number of hosts constant.

Fig. 6b shows the running times of the analysis of a datacenter with 3–189 middleboxes (1–32 middlebox chains). All topologies contained 1000 hosts.

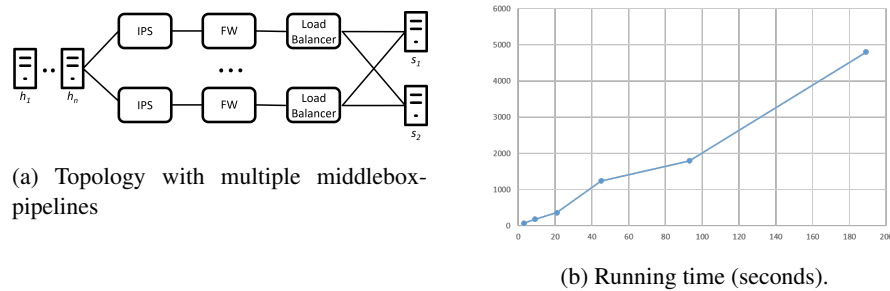


Fig. 6: Topology and running times of the network topology scalability test.

6 Concluding Remarks and Related Work

In this paper, we applied abstract interpretation for efficient verification of networks with stateful nodes. We now briefly survey closely related works in this area.

Topology Independent Network Verification. Early work in network verification focused on proving correctness of network protocols [6,28]. Subsequent work in the context of software define networking (SDN) including Flowlog [23] and VeriCon [4] looked at verifying the correctness of network applications (implemented as middle-boxes or in network controllers) independent of the topology and configuration of the network where these were used. However, since this problem is undecidable, these methods use bounded model checking or user provided inductive invariants, which are hard to specify even in simple network topologies.

Verifying Immutable Network Configurations. Verifying networks with immutable states is an active line of research [18,14,16,5,15,33,30,2,12]. In the future, we hope to combine our abstraction with the techniques used in these papers. We hope to use similar techniques to Veriflow [16] to handle switches more efficiently, and leverage compact header representation described in NetKat [12].

Stateful Network Verification. Previous works provide useful tools for detecting errors in firewalls [20,19,22]. Buzz [9] and SymNet [34] have looked at how to use symbolic execution and packet generation for testing and verifying the behavior of stateful networks. These works implement testing techniques rather than verifying network behavior and are hence complementary to our approach.

Velner et al. [35] show that checking safety in stateful networks is undecidable, necessitating the use of overapproximations. They provide a general algorithm for checking safety using Petri nets. This algorithm has high complexity and scales poorly. They also provide an efficient algorithm for checking safety in a limited class of networks.

Exploring Network Symmetry. Recent work explored the use of bisimulation to leverage the extensive symmetry found in real network topologies [21] to accelerate stateless [25] and stateful [24] network verification. Both approaches are not automatic. We are encouraged by the fact that our automatic approach achieves performance comparable to VMN [24] on the same examples without requiring human intervention. We attribute this improvement to modularity and to the use of packet state representation.

Extensible Semantics. Previous works have explored ideas similar to the reverting semantics, to obtain complexity and decidability results in different settings.

In [8] the authors analyze the complexity of verifying asynchronous shared-memory systems. They use *copycat* processes that mirror the behaviour of another process to show that executions are extensible, similarly to how our work uses the sticky packet states property (Lem. 4). In their model, when the processes are finite state machines, they obtain coNP-complete complexity for verification.

In [10] the authors explore a more general setting of well-structured transition system, and present the *home-state idea*, which allows the system to return to its initial state (essentially, revert). They obtain decidability results for well-structured transition systems with a home-state, but do not show any tighter complexity results.

Acknowledgments We thank our anonymous shepherd, and anonymous referees for insightful comments which improved this paper. We thank LogicBlox for providing us with an academic license for their software, and Todd J. Green and Martin Bravenboer for providing technical support and helping with optimization. This publication is part of projects that have received funding from the European Research Council (ERC) under the European Union’s Seventh Framework Program (FP7/2007–2013) / ERC grant agreement no. [321174-VSSC], and Horizon 2020 research and innovation programme (grant agreement No [759102-SVIS]). The research was supported in part by Len Blavatnik and the Blavatnik Family foundation, the Blavatnik Interdisciplinary Cyber Research Center, Tel Aviv University, and the Pazy Foundation. This material is based upon work supported by the United States-Israel Binational Science Foundation (BSF) grants No. 2016260 and 2012259. This research was also supported in part by NSF grants 1704941 and 1420064, and funding provided by Intel Corporation.

References

1. K. Alpernas, R. Manevich, A. Panda, M. Sagiv, S. Shenker, S. Shoham, and Y. Velner. Abstract interpretation of stateful networks. *arXiv preprint arXiv:1708.05904*, 2018.
2. C. J. Anderson, N. Foster, A. Guha, J.-B. Jeannin, D. Kozen, C. Schlesinger, and D. Walker. NetKAT: Semantic foundations for networks. In *POPL*, 2014.
3. M. Aref, B. ten Cate, T. J. Green, B. Kimelfeld, D. Olteanu, E. Pasalic, T. L. Veldhuizen, and G. Washburn. Design and implementation of the logicblox system. In *ACM SIGMOD International Conference on Management of Data*, pages 1371–1382, 2015.
4. T. Ball, N. Bjørner, A. Gember, S. Itzhaky, A. Karbyshev, M. Sagiv, M. Schapira, and A. Valadarsky. Vericon: towards verifying controller programs in software-defined networks. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI*, page 31, 2014.
5. M. Canini, D. Venzano, P. Peres, D. Kostic, and J. Rexford. A nice way to test openflow applications. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI’12)*, 2012.
6. E. M. Clarke, S. Jha, and W. R. Marrero. Using state space exploration and a natural deduction style message derivation engine to verify security protocols. In *Programming Concepts and Methods, IFIP TC2/WG2.2.2.3 International Conference on Programming Concepts and Methods (PROCOMET ’98) 8-12 June 1998, Shelter Island, New York, USA*, pages 87–106, 1998.
7. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *Proceedings of the 6th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 269–282. ACM, 1979.
8. J. Esparza, P. Ganty, and R. Majumdar. Parameterized verification of asynchronous shared-memory systems. In *International Conference on Computer Aided Verification*, pages 124–140. Springer, 2013.
9. S. K. Fayaz, T. Yu, Y. Tobioka, S. Chaki, V. Sekar, S. Vyas, and Cmu. Buzz: Testing context-dependent policies in stateful networks buzz: Testing context-dependent policies in stateful networks. In *NSDI*, 2016.
10. A. Finkel and P. Schnoebelen. Well-structured transition systems everywhere! *Theoretical Computer Science*, 256(1-2):63–92, 2001.
11. C. Flanagan, S. N. Freund, S. Qadeer, and S. A. Seshia. Modular verification of multithreaded programs. *Theor. Comput. Sci.*, 338(1-3):153–183, 2005.

12. N. Foster, D. Kozen, M. Milano, A. Silva, and L. Thompson. A coalgebraic decision procedure for netkat. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 343–355, 2015.
13. J. Hoenicke, R. Majumdar, and A. Podelski. Thread modularity at many levels: a pearl in compositional verification. In *POPL*, pages 473–485, 2017.
14. P. Kazemian, M. Chang, H. Zeng, G. Varghese, N. McKeown, and S. Whyte. Real time network policy checking using header space analysis. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI '13)*, 2013.
15. P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI '12)*, 2012.
16. A. Khurshid, W. Zhou, M. Caesar, and B. Godfrey. Veriflow: verifying network-wide invariants in real time. *Computer Communication Review*, 42(4):467–472, 2012.
17. M. Kuzniar, P. Peresini, M. Canini, D. Venzano, and D. Kostic. A soft way for openflow switch interoperability testing. In *CoNEXT*, pages 265–276, 2012.
18. H. Mai, A. Khurshid, R. Agarwal, M. Caesar, B. Godfrey, and S. T. King. Debugging the Data Plane with Anteater. In *SIGCOMM*, 2011.
19. R. M. Marmorstein and P. Kearns. A tool for automated iptables firewall analysis. In *Usenix annual technical conference, Freenix Track*, pages 71–81, 2005.
20. A. Mayer, A. Wool, and E. Ziskind. Fang: A firewall analysis engine. In *Security and Privacy, 2000. S&P 2000. Proceedings. 2000 IEEE Symposium on*, pages 177–187. IEEE, 2000.
21. K. S. Namjoshi and R. J. Treffer. Uncovering symmetries in irregular process networks. In *Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings*, pages 496–514, 2013.
22. T. Nelson, C. Barratt, D. J. Dougherty, K. Fisler, and S. Krishnamurthi. The margrave tool for firewall analysis. In *LISA*, 2010.
23. T. Nelson, A. D. Ferguson, M. J. G. Scheer, and S. Krishnamurthi. Tierless programming and reasoning for software-defined networks. In *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2014, Seattle, WA, USA, April 2-4, 2014*, pages 519–531, 2014.
24. A. Panda, O. Lahav, K. J. Argyraki, M. Sagiv, and S. Shenker. Verifying reachability in networks with mutable datapaths. In *14th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2017, Boston, MA, USA, March 27-29, 2017*, pages 699–718, 2017.
25. G. D. Plotkin, N. Bjørner, N. P. Lopes, A. Rybalchenko, and G. Varghese. Scaling network verification using symmetry and surgery. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 69–83, 2016.
26. A. Pnueli, J. Xu, and L. Zuck. Liveness with $(0, 1, \text{infinity})$ -counter abstraction. In *Computer Aided Verification*, pages 93–111. Springer, 2002.
27. R. Potharaju and N. Jain. Demystifying the dark side of the middle: a field study of middle-box failures in datacenters. In *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23-25, 2013*, pages 9–22, 2013.
28. R. W. Ritchey and P. Ammann. Using model checking to analyze network vulnerabilities. In *Security and Privacy, 2000*.
29. A. W. Roscoe and C. A. R. Hoare. The laws of occam programming. *Theoretical Computer Science*, 60(2):177–229, 1988.
30. D. Sethi, S. Narayana, and S. Malik. Abstractions for model checking sdn controllers. In *FMCAD*, 2013.

31. J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else's problem: Network processing as a cloud service. In *SIGCOMM*, 2012.
32. A. Sivaraman, A. Cheung, M. Budiu, C. Kim, M. Alizadeh, H. Balakrishnan, G. Varghese, N. McKeown, and S. Licking. Packet transactions: High-level programming for line-rate switches. In *Proceedings of the ACM SIGCOMM 2016 Conference, Florianopolis, Brazil, August 22-26, 2016*, pages 15–28, 2016.
33. R. Skowrya, A. Lapets, A. Bestavros, and A. Kfoury. A verification platform for sdn-enabled applications. In *HiCoNS*, 2013.
34. R. Stoenescu, M. Popovici, L. Negreanu, and C. Raiciu. Scalable symbolic execution for modern networks. In *SIGCOMM*, 2016.
35. Y. Velner, K. Alpernas, A. Panda, A. Rabinovich, M. Sagiv, S. Shenker, and S. Shoham. Some complexity results for stateful network verification. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 811–830. Springer, 2016.