



Towards Model Checking Real-World Software-Defined Networks

Vasileios Klimis^(✉), George Parisis^{}, and Bernhard Reus^{}

University of Sussex, Brighton, UK
{v.klimis,g.paris,bernhard}@sussex.ac.uk

Abstract. In software-defined networks (SDN), a controller program is in charge of deploying diverse network functionality across a large number of switches, but this comes at a great risk: deploying buggy controller code could result in network and service disruption and security loopholes. The automatic detection of bugs or, even better, verification of their absence is thus most desirable, yet the size of the network and the complexity of the controller makes this a challenging undertaking. In this paper, we propose MOCS, a highly expressive, optimised SDN model that allows capturing subtle real-world bugs, in a reasonable amount of time. This is achieved by (1) analysing the model for possible partial order reductions, (2) statically pre-computing packet equivalence classes and (3) indexing packets and rules that exist in the model. We demonstrate its superiority compared to the state of the art in terms of expressivity, by providing examples of realistic bugs that a prototype implementation of MOCS in UPPAAL caught, and performance/scalability, by running examples on various sizes of network topologies, highlighting the importance of our abstractions and optimisations.

1 Introduction

Software-Defined Networking (SDN) [16] has brought about a paradigm shift in designing and operating computer networks. A logically centralised controller implements the control logic and ‘programs’ the data plane, which is defined by flow tables installed in network switches. SDN enables the rapid development of advanced and diverse network functionality; e.g. in designing next-generation inter-data centre traffic engineering [10], load balancing [19], firewalls [24], and Internet exchange points (IXPs) [15]. SDN has gained noticeable ground in the industry, with major vendors integrating OpenFlow [37], the de-facto SDN standard maintained by the Open Networking Forum, in their products. Operators deploy it at scale [27, 38]. SDN presents a unique opportunity for innovation and rapid development of complex network services by enabling all players, not just vendors, to develop and deploy control and data plane functionality in networks. This comes at a great risk; deploying buggy code at the controller could result in problematic flow entries at the data plane and, potentially, service disruption [13, 18, 47, 49] and security loopholes [7, 26]. Understanding and fixing such

bugs is far from trivial, given the distributed and concurrent nature of computer networks and the complexity of the control plane [44].

With the advent of SDN, a large body of research on verifying network properties has emerged [33]. Static network analysis approaches [2, 11, 30, 34, 45, 51] can only verify network properties on a given fixed network configuration but this may be changing very quickly (e.g. as in [1]). Another key limitation is the fact that they cannot reason about the controller program, which, itself, is responsible for the changes in the network configuration. Dynamic approaches, such as [23, 29, 31, 40, 48, 50], are able to reason about network properties as changes happen (i.e. as flow entries in switches' flow tables are being added and deleted), but they cannot reason about the controller program either. As a result, when a property violation is detected, there is no straightforward way to fix the bug in the controller code, as these systems are oblivious of the running code. Identifying bugs in large and complex deployments can be extremely challenging.

Formal verification methods that include the controller code in the model of the network can solve this important problem. Symbolic execution methods, such as [5, 8, 11, 12, 14, 28, 46], evaluate programs using symbolic variables accumulating path-conditions along the way that then can be solved logically. However, they suffer from the path explosion problem caused by loops and function calls which means verification does not scale to larger controller programs (bug finding still works but is limited). Model checking SDNs is a promising area even though only few studies have been undertaken [3, 8, 28, 35, 36, 43]. Networks and controller can be naturally modelled as transition systems. State explosion is always a problem but can be mitigated by using abstraction and optimisation techniques (i.e. partial order reductions). At the same time, modern model checkers [6, 9, 20, 21, 25] are very efficient.

NetsMC [28] uses a bespoke *symbolic* model checking algorithm for checking properties given a subset of computation tree logic that allows quantification only over all paths. As a result, this approach scales relatively well, but the requirement that only one packet can travel through the network at any time is very restrictive and ignores race conditions. NICE [8] employs model checking but only looks at a limited amount of input packets that are extracted through symbolically executing the controller code. As a result, it is a bug-finding tool only. The authors in [43] propose a model checking approach that can deal with dynamic controller updates and an arbitrary number of packets but require manually inserted non-interference lemmas that constrain the set of packets that can appear in the network. This significantly limits its applicability in realistic network deployments. Kuai [35] overcomes this limitation by introducing model-specific partial order reductions (PORs) that result in pruning the state space by avoiding redundant explorations. However, it has limitations explained at the end of this section.

In this paper, we take a step further towards the full realisation of model checking real-world SDNs by introducing MOCS (MOdel Checking for Software defined networks)¹, a highly expressive, optimised SDN model which we

¹ A release of MOCS is publicly available at <https://tinyurl.com/y95qtv5k>.

implemented in UPPAAL² [6]. MOCS, compared to the state of the art in model checking SDNs, can model network behaviour more realistically and verify larger deployments using fewer resources. The main contributions of this paper are:

Model Generality. The proposed network model is closer to the OpenFlow standard than previous models (e.g. [35]) to reflect commonly exhibited behaviour between the controller and network switches. More specifically, it allows for race conditions between control messages and includes a significant number of OpenFlow interactions, including barrier response messages. In our experimentation section, we present families of elusive bugs that can be efficiently captured by MOCS.

Model Checking Optimisations. To tackle the state explosion problem we propose context-dependent *partial order reductions* by considering the concrete control program and specification in question. We establish the soundness of the proposed optimisations. Moreover, we propose *state representation optimisations*, namely packet and rule indexing, identification of packet equivalence classes and bit packing, to improve performance. We evaluate the benefits from all proposed optimisations in Sect. 4.

Our model has been inspired by Kuai [35]. According to the contributions above, however, we consider MOCS to be a considerable improvement. We model more OpenFlow messages and interactions, enabling us to check for bugs that [35] cannot even express (see discussion in Sect. 4.2). Our context-dependent PORs systematically explore possibilities for optimisation. Our optimisation techniques still allow MOCS to run at least as efficiently as Kuai, often with even better performance.

2 Software-Defined Network Model

A key objective of our work is to enable the verification of network-wide properties in real-world SDNs. In order to fulfill this ambition, we present an extended network model to capture complex interactions between the SDN controller and the network. Below we describe the adopted network model, its state and transitions.

2.1 Formal Model Definition

The formal definition of the proposed SDN model is by means of an action-deterministic transition system. We parameterise the model by the underlying network topology λ and the controller program CP in use, as explained further below (Sect. 2.2).

Definition 1. *An SDN model is a 6-tuple $\mathcal{M}_{(\lambda, \text{CP})} = (S, s_0, A, \hookrightarrow, AP, L)$, where S is the set of all states the SDN may enter, s_0 the initial state, A the set of*

² UPPAAL has been chosen as future plans include extending the model to timed actions like e.g. timeouts. Note that the model can be implemented in any model checker.

actions which encode the events the network may engage in, $\hookrightarrow \subseteq S \times A \times S$ the transition relation describing which execution steps the system undergoes as it perform actions, AP a set of atomic propositions describing relevant state properties, and $L : S \rightarrow 2^{AP}$ is a labelling function, which relates to any state $s \in S$ a set $L(s) \in 2^{AP}$ of those atomic propositions that are true for s . Such an SDN model is composed of several smaller systems, which model network components (hosts, switches and the controller) that communicate via queues and, combined, give rise to the definition of \hookrightarrow . The states of an SDN transition system are 3-tuples (π, δ, γ) , where π represents the state of each host, δ the state of each switch, and γ the controller state. The components are explained in Sect. 2.2 and the transitions \hookrightarrow in Sect. 2.3.

Figure 1 illustrates a high-level view of OpenFlow interactions (left side), modelled actions and queues (right side).

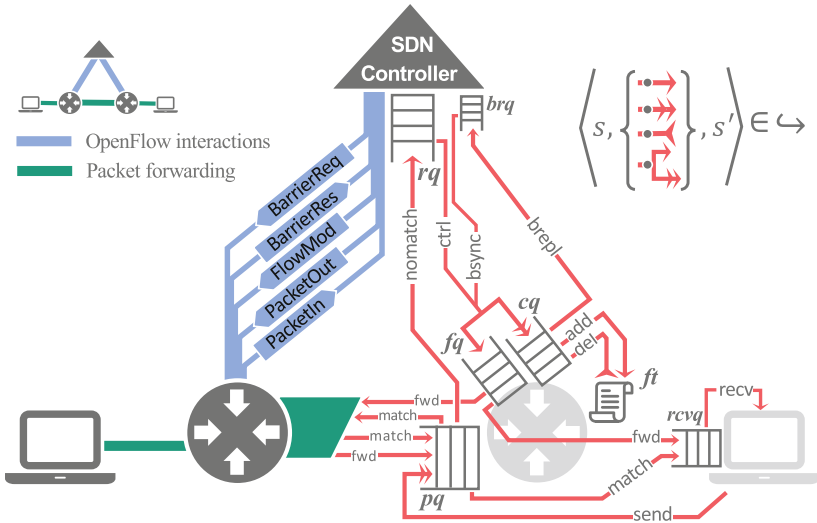


Fig. 1. A high-level view of OpenFlow interactions using OpenFlow specification terminology (left half) and the modelled actions (right half). A red solid-line arrow depicts an action which, when fired, (1) dequeues an item from the queue the arrow begins at, and (2) adds an item in the queue the arrowhead points to (or multiple items if the arrow is double-headed). Deleting an item from the target queue is denoted by a reverse arrowhead. A forked arrow denotes multiple targeted queues. (Color figure online)

2.2 SDN Model Components

Throughout we will use the common “dot-notation” (\dots) to refer to components of composite gadgets (tuples), e.g. queues of switches, or parts of the state. We use obvious names for the projections functions like $s.\delta.sw.pq$ for the packet

queue of the switch sw in state s . At times we will also use t_1 and t_2 for the first and second projection of tuple t .

Network Topology. A location (n, pt) is a pair of a node (host or switch) n and a port pt . We describe the network topology as a bijective function $\lambda : (Switches \cup Hosts) \times Ports \rightarrow (Switches \cup Hosts) \times Ports$ consisting of a set of directed edges $\langle (n, pt), (n', pt') \rangle$, where pt' is the input port of the switch or host n' that is connected to port pt at host or switch n . $Hosts$, $Switches$ and $Ports$ are the (finite) sets of all hosts, switches and ports in the network, respectively. The topology function is used when a packet needs to be forwarded in the network. The location of the next hop node is decided when a *send*, *match* or *fwd* action (all defined further below) is fired. Every SDN model is w.r.t. a fixed topology λ that does not change.

Packets. Packets are modelled as finite bit vectors and transferred in the network by being stored to the queues of the various network components. A *packet* $\in Packets$ (the set of all packets that can appear in the network) contains bits describing the proof-relevant header information and its location *loc*.

Hosts. Each *host* $\in Hosts$, has a packet queue (*rcvq*) and a finite set of ports which are connected to ports of other switches. A host can send a packet to one or more switches it is connected to (*send* action in Fig. 1) or receive a packet from its own *rcvq* (*recv* action in Fig. 1). Sending occurs repeatedly in a non-deterministic fashion which we model implicitly via the $(0, \infty)$ abstraction at switches' packet queues, as discussed further below.

Switches. Each *switch* $\in Switches$, has a flow table (*ft*), a packet queue (*pq*), a control queue (*cq*), a forwarding queue (*fq*) and one or more ports, through which it is connected to other switches and/or hosts. A flow table $ft \subseteq Rules$ is a set of forwarding rules (with $Rules$ being the set of all rules). Each one consists of a tuple $(priority, pattern, ports)$, where *priority* $\in \mathbb{N}$ determines the priority of the rule over others, *pattern* is a proposition over the proof-relevant header of a packet, and *ports* is a subset of the switch's ports. Switches match packets in their packet queues against rules (i.e. their respective *pattern*) in their flow table (*match* action in Fig. 1) and forward packets to a connected device (or final destination), accordingly. Packets that cannot be matched to any rule are sent to the controller's request queue (*rq*) (*nomatch* action in Fig. 1); in OpenFlow, this is done by sending a *PacketIn* message. The forwarding queue *fq* stores packets forwarded by the controller in *PacketOut* messages. The control queue stores messages sent by the controller in *FlowMod* and *BarrierReq* messages. *FlowMod* messages contain instructions to add or delete rules from the flow table (that trigger *add* and *del* actions in Fig. 1). *BarrierReq* messages contain barriers to synchronise the addition and removal of rules. MOCS conforms to the OpenFlow specifications and always execute instructions in an interleaved fashion obeying the ordering constraints imposed by barriers.

OpenFlow Controller. The controller is modelled as a finite state automaton embedded into the overall transition system. A controller program CP , as used to parametrise an SDN model, consists of $(CS, pktIn, barrierIn)$. It uses its own local state $cs \in CS$, where CS is the finite set of control program states. Incoming

PacketIn and *BarrierRes* messages from the SDN model are stored in separate queues (*rq* and *brq*, respectively) and trigger *ctrl* or *bsync* actions (see Fig. 1) which are then processed by the controller program in its current state. The controller’s corresponding handler, *pktIn* for *PacketIn* messages and *barrierIn* for *BarrierRes* messages, responds by potentially changing its local state and sending messages to a subset of *Switches*, as follows. A number of *PacketOut* messages (pairs of *pkt*, *ports*) can be sent to a subset of *Switches*. Such a message is stored in a switch’s forward queue and instructs it to forward packet *pkt* along the ports *ports*. The controller may also send any number of *FlowMod* and *BarrierReq* messages to the control queue of any subset of *Switches*. A *FlowMod* message may contain an *add* or *delete* rule modification instruction. These are executed in an arbitrary order by switches, and *barriers* are used to synchronise their execution. Barriers are sent by the controller in *BarrierReq* messages. OpenFlow requires that a response message (*BarrierRes*) is sent to the controller by a switch when a barrier is consumed from its control queue so that the controller can synchronise subsequent actions. Our model includes a *brepl* action that models the sending of a *BarrierRes* message from a switch to the controller’s barrier reply queue (*brq*), and a *bsync* action that enables the controller program to react to barrier responses.

Queues. All queues in the network are modelled as *finite* state. Packet queues *pq* for switches are modelled as multisets, and we adopt $(0, \infty)$ abstraction [41]; i.e. a packet is assumed to appear either zero or an arbitrary (unbounded) amount of times in the respective multiset. This means that once a packet has arrived at a switch or host, (infinitely) many other packets of the same kind repeatedly arrive at this switch or host. Switches’ forwarding queues *fq* are, by contrast, modelled as sets, therefore if multiple identical packets are sent by the controller to a switch, only one will be stored in the queue and eventually forwarded by the switch. The controller’s request *rq* and barrier reply queues *brq* are modelled as sets as well. Hosts’ receive queues *rcvq* are also modelled as sets. Controller queues *cq* at switches are modelled as a finite sequence of sets of control messages (representing add and remove rule instructions), interleaved by any number of barriers. As the number of barriers that can appear at any execution is finite, this sequence is finite.

2.3 Guarded Transitions

Here we provide a detailed breakdown of the transition relation $s \xrightarrow{\alpha(\vec{a})} s'$ for each action $\alpha(\vec{a}) \in A(s)$, where $A(s)$ the set of all enabled actions in s in the proposed model (see Fig. 1). Transitions are labelled by action names α with arguments \vec{a} . The transitions are only enabled in state s if s satisfies certain conditions called *guards* that can refer to the arguments \vec{a} . In guards, we make use of predicate *bestmatch*(*sw*, *r*, *pkt*) that expresses that *r* is the highest priority rule in *sw.ft* that matches *pkt*’s header. Below we list all possible actions with their respective guards.

***send*(*h*, *pt*, *pkt*).** Guard: *true*. This transition models packets arriving in the network in a non-deterministic fashion. When it is executed, *pkt* is added to

the packet queue of the network switch connected to the port pt of host h (or, formally, to $\lambda(h, pt)_1.pq$, where λ is the topology function described above). As described in Sect. 3.2, only relevant representatives of packets are actually sent by end-hosts. This transition is unguarded, therefore it is always enabled.

***recv*(h, pkt).** Guard: $pkt \in h.rcvq$. This transition models hosts receiving (and removing) packets from the network and is enabled if pkt is in h 's receive queue.

***match*(sw, pkt, r).** Guard: $pkt \in sw.pq \wedge r \in sw.ft \wedge bestmatch(sw, r, pkt)$. This transition models matching and forwarding packet pkt to zero or more next hop nodes (hosts and switches), as a result of highest priority matching of rule r with pkt . The packet is then copied to the packet queues of the connected hosts and/or switches, by applying the topology function to the port numbers in the matched rule; i.e. $\lambda(sw, pt)_1.pq, \forall pt \in r.ports$. Dropping packets is modelled by having a special 'drop' port that can be included in rules. The location of the forwarded packet(s) is updated with the respective destination (switch/host, port) pair; i.e. $\lambda(sw, pt)$. Due to the $(0, \infty)$ abstraction, the packet is not removed from $sw.pq$.

***nomatch*(sw, pkt).** Guard: $pkt \in sw.pq \wedge \nexists r \in sw.ft . bestmatch(sw, r, pkt)$. This transition models forwarding a packet to the OpenFlow controller when a switch does not have a rule in its forwarding table that can be matched against the packet header. In this case, pkt is added to rq for processing. pkt is not removed from $sw.pq$ due to the supported $(0, \infty)$ abstraction.

***ctrl*(sw, pkt, cs).** Guard: $pkt \in controller.rq$. This transition models the execution of the packet handler by the controller when packet pkt that was previously sent by sw is available in rq . The controller's packet handler function $pktIn(sw, pkt, cs)$ is executed which, in turn (i) reads the current controller state cs and changes it according to the controller program, (ii) adds a number of rules, interleaved with any number of barriers, into the cq of zero or more switches, and (iii) adds zero or more forwarding messages, each one including a packet along with a set of ports, to the fq of zero or more switches.

***fwd*($sw, pkt, ports$).** Guard: $(pkt, ports) \in sw.fq$. This transition models forwarding packet pkt that was previously sent by the controller to sw 's forwarding queue $sw.fq$. In this case, pkt is removed from $sw.fq$ (which is modelled as a set), and added to the pq of a number of network nodes (switches and/or hosts), as defined by the topology function $\lambda(sw, pt)_1.pq, \forall pt \in ports$. The location of the forwarded packet(s) is updated with the respective destination (switch/host, port) pair; i.e. $\lambda(n, pt)$.

***FM*(sw, r),** where $FM \in \{add, del\}$. Guard: $(FM, r) \in head(sw.cq)$. These transitions model the addition and deletion, respectively, of a rule in the flow table of switch sw . They are enabled when one or more *add* and *del* control messages are in the set at the head of the switch's control queue. In this case, r is added to – or deleted from, respectively – $sw.ft$ and the control message is deleted from the set at the head of cq . If the set at the head of cq becomes empty it is removed. If then the next item in cq is a barrier, a *brepl* transition becomes enabled (see below).

***brepl*(*sw*, *xid*).** Guard: $b(xid) = head(sw.cq)$. This transition models a switch sending a barrier response message, upon consuming a barrier from the head of its control queue; i.e. if $b(xid)$ is the head of $sw.cq$, where $xid \in \mathbb{N}$ is an identifier for the barrier set by the controller, $b(xid)$ is removed and the barrier reply message $br(sw, xid)$ is added to the controller's brq .

***bsync*(*sw*, *xid*, *cs*).** Guard: $br(sw, xid) \in controller.brq$. This transition models the execution of the barrier response handler by the controller when a barrier response sent by switch sw is available in brq . In this case, $br(sw, xid)$ is removed from the brq , and the controller's barrier handler $barrierIn(sw, xid, cs)$ is executed which, in turn (i) reads the current controller state cs and changes it according to the controller program, (ii) adds a number of rules, interleaved with any number of barriers, into the cq of zero or more switches, and (iii) adds zero or more forwarding messages, each one including a packet along with a set of ports, to the fq of zero or more switches.

An Example Run. In Fig. 2, we illustrate a sequence of MOCS transitions through a simple packet forwarding example. The run starts with a *send* transition; packet p is copied to the packet queue of the switch in black. Initially, switches' flow tables are empty, therefore p is copied to the controller's request queue (*nomatch* transition); note that p remains in the packet queue of the switch in black due to the $(0, \infty)$ abstraction. The controller's packet handler is then called (*ctrl* transition) and, as a result, (1) p is copied to the forwarding queue of the switch in black, (2) rule r_1 is copied to the control queue of the switch in black, and (3) rule r_2 is copied to the control queue of the switch in white. Then, the switch in black forwards p to the packet queue of the switch in white (*fwd* transition). The switch in white installs r_2 in its flow table (*add* transition) and then matches p with the newly installed rule and forwards it to the receive queue of the host in white (*match* transition), which removes it from the network (*recv* transition).

2.4 Specification Language

In order to specify properties of packet flow in the network, we use LTL formulas without “next-step” operator \bigcirc ³, where atomic formulae denoting properties of states of the transition system, i.e. SDN network. In the case of safety properties, i.e. an invariant w.r.t. states, the $LTL_{\setminus\{\bigcirc\}}$ formula is of the form $\Box\varphi$, i.e. has only an outermost \Box temporal connective.

Let P denote unary predicates on packets which encode a property of a packet based on its fields. An atomic *state condition* (proposition) in AP is either of the following: (i) existence of a packet pkt located in a packet queue (pq) of a switch or in a receive queue ($rcvq$) of a host that satisfies P (we denote this by $\exists pkt \in n.pq . P(pkt)$ with $n \in Switches$, and $\exists pkt \in h.rcvq . P(pkt)$

³ This is the largest set of formulae supporting the partial order reductions used in Sect. 3, as stutter equivalence does not preserve the truth value of formulae with the \bigcirc .

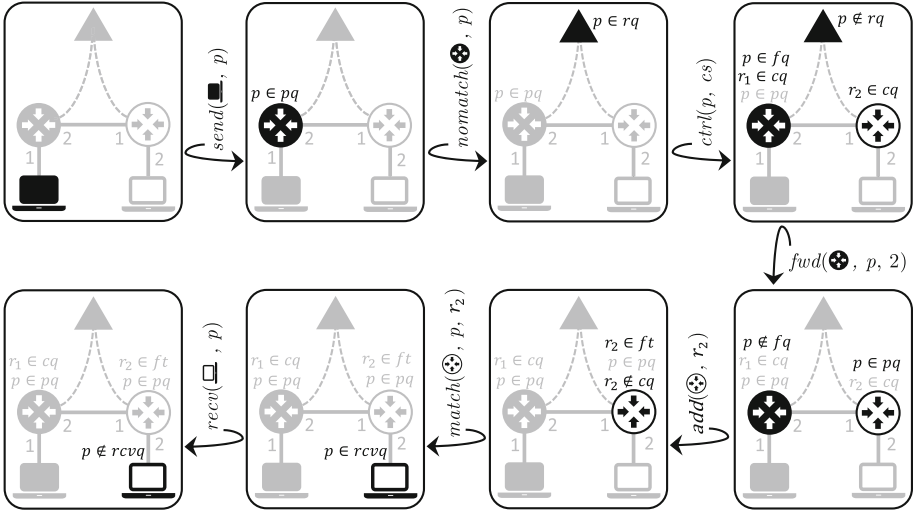


Fig. 2. Forwarding p from \blacksquare to \square . Non greyed-out icons are the ones whose state changes in the current transition.

with $h \in Hosts$ ⁴; (ii) the controller is in a specific *controller* state $q \in CS$, denoted by a unary predicate symbol $Q(q)$ which holds in system state $s \in S$ if $q = s.\gamma.cs$. The specification logic comprises first-order formula with equality on the finite domains of switches, hosts, rule priorities, and ports which are *state-independent* (and decidable).

For example, $\exists pkt \in sw.pq. P(pkt)$ represents the fact that the packet predicate $P(_)$ is true for at least one packet pkt in the pq of switch sw . For every atomic packet proposition $P(pkt)$, also its negation $\neg P(pkt)$ is an atomic proposition for the reason of simplifying syntactic checks of formulae in Table 1 in the next section. Note that universal quantification over packets in a queue is a derived notion. For instance, $\forall pkt \in n.pq. P(pkt)$ can be expressed as $\nexists pkt \in n.pq. \neg P(pkt)$. Universal and existential quantification over switches or hosts can be expressed by finite iterations of \wedge and \vee , respectively.

In order to be able to express that a condition holds when a certain event happened, we add to our propositions instances of *propositional dynamic logic* [17, 42]. Given an action $\alpha(\cdot) \in A$ and a proposition P that may refer to any variables in \vec{x} , $[\alpha(\vec{x})]P$ is also a proposition and $[\alpha(\vec{x})]P$ is true if, and only if, after firing transition $\alpha(\vec{a})$ (to get to the current state), P holds with the variables in \vec{x} bound to the corresponding values in the actual arguments \vec{a} . With the help of those basic modalities one can then also specify that more complex events occurred. For instance, dropping of a packet due to a *match* or *fwd* action can be expressed by $[match(sw, pkt, r)](r.fwd_port = \mathbf{drop}) \wedge [fwd(sw, pkt, pt)](pt = \mathbf{drop})$. Such predicates derived from modalities are used in [32] (extended version of this paper, with proofs and controller programs), Appendix B-CP5.

⁴ Note that these are *atomic* propositions despite the use of the existential quantifier notation.

The meaning of temporal LTL operators is standard depending on the trace of a transition sequence $s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots$. The trace $L(s_0)L(s_1)\dots L(s_i)\dots$ is defined as usual. For instance, trace $L(s_0)L(s_1)L(s_2)\dots$ satisfies invariant $\Box\varphi$ if each $L(s_i)$ implies φ .

3 Model Checking

In order to verify desired properties of an SDN, we use its model as described in Definition 1 and apply model checking. In the following we propose optimisations that significantly improve the performance of model checking.

3.1 Contextual Partial-Order Reduction

Partial order reduction (POR) [39] reduces the number of interleavings (traces) one has to check. Here is a reminder of the main result (see [4]) where we use a stronger condition than the regular (C4) to deal with cycles:

Theorem 1 (Correctness of POR). *Given a finite transition system $\mathcal{M} = (S, A, \hookrightarrow, s_0, AP, L)$ that is action-deterministic and without terminal states, let $A(s)$ denote the set of actions in A enabled in state $s \in S$. Let $\text{ample}(s) \subseteq A(s)$ be a set of actions for a state $s \in S$ that satisfies the following conditions:*

C1 (Non)emptiness condition: $\emptyset \neq \text{ample}(s) \subseteq A(s)$.

C2 Dependency condition: Let $s \xrightarrow{\alpha_1} s_1 \dots \xrightarrow{\alpha_n} s_n \xrightarrow{\beta} t$ be a run in \mathcal{M} . If $\beta \in A \setminus \text{ample}(s)$ depends on $\text{ample}(s)$, then $\alpha_i \in \text{ample}(s)$ for some $0 < i \leq n$, which means that in every path fragment of \mathcal{M} , β cannot appear before some transition from $\text{ample}(s)$ is executed.

C3 Invisibility condition: If $\text{ample}(s) \neq A(s)$ (i.e., state s is not fully expanded), then every $\alpha \in \text{ample}(s)$ is invisible.

C4 Every cycle in $\mathcal{M}^{\text{ample}}$ contains a state s such that $\text{ample}(s) = A(s)$.

where $\mathcal{M}^{\text{ample}} = (S_a, A, \twoheadrightarrow, s_0, AP, L_a)$ is the new, optimised, model defined as follows: let $S_a \subseteq S$ be the set of states reachable from the initial state s_0 under \twoheadrightarrow , let $L_a(s) = L(s)$ for all $s \in S_a$, and define $\twoheadrightarrow \subseteq S_a \times A \times S_a$ inductively by the rule

$$\frac{s \xrightarrow{\alpha} s'}{s \twoheadrightarrow s'} \quad \text{if } \alpha \in \text{ample}(s)$$

If $\text{ample}(s)$ satisfies conditions (C1)–(C4) as outlined above, then for each path in \mathcal{M} there exists a stutter-trace equivalent path in $\mathcal{M}^{\text{ample}}$, and vice-versa, denoted $\mathcal{M} \stackrel{\text{st}}{\equiv} \mathcal{M}^{\text{ample}}$.

The intuitive reason for this theorem to hold is the following: Assume an action sequence $\alpha_i \dots \alpha_{i+n} \beta$ that reaches the state s , and β is independent of $\{\alpha_i, \dots, \alpha_{i+n}\}$. Then, one can permute β with α_{i+n} through α_i successively n times. One can

therefore construct the sequence $\beta\alpha_i\dots\alpha_{i+n}$ that also reaches the state s . If this shift of β does not affect the labelling of the states with atomic propositions (β is called *invisible* in this case), then it is not detectable by the property to be shown and the permuted and the original sequence are equivalent w.r.t. the property and thus don't have to be checked both. One must, however, ensure, that in case of loops (infinite execution traces) the ample sets do not *preclude* some actions to be fired altogether, which is why one needs (C4).

The more actions that are both stutter and provably independent (also referred to as *safe actions* [22]) there are, the smaller the transition system, and the more efficient the model checking. One of our contributions is that we attempt to identify *as many safe actions as possible* to make PORs more widely applicable to our model.

The PORs in [35] consider only dependency and invisibility of *recv* and *barrier* actions, whereas we explore systematically all possibilities for applications of Theorem 1 to reduce the search space. When identifying safe actions, we consider (1) the actual controller program CP, (2) the topology λ and (3) the state formula φ to be shown invariant, which we call the *context* CTX of actions. It turns out that two actions may be dependent in a given context of abstraction while independent in another context, and similarly for invisibility, and we exploit this fact. The argument of the action thus becomes relevant as well.

Definition 2 (Safe Actions). *Given a context $\text{CTX} = (\text{CP}, \lambda, \varphi)$, and SDN model $\mathcal{M}_{(\lambda, \text{CP})} = (S, A, \hookrightarrow, s_0, AP, L)$, an action $\alpha(\cdot) \in A(s)$ is called ‘safe’ if it is independent of any other action in A and invisible for φ . We write safe actions $\check{\alpha}(\cdot)$.*

Definition 3 (Order-sensitive Controller Program). *A controller program CP is order-sensitive if there exists a state $s \in S$ and two actions α, β in $\{\text{ctrl}(\cdot), \text{bsync}(\cdot)\}$ such that $\alpha, \beta \in A(s)$ and $s \xrightarrow{\alpha} s_1 \xrightarrow{\beta} s_2$ and $s \xrightarrow{\beta} s_3 \xrightarrow{\alpha} s_4$ with $s_2 \neq s_4$.*

Definition 4. *Let φ be a state formula. An action $\alpha \in A$ is called ‘ φ -invariant’ if $s \models \varphi$ iff $\alpha(s) \models \varphi$ for all $s \in S$ with $\alpha \in A(s)$.*

Lemma 1. *For transition system $\mathcal{M}_{(\lambda, \text{CP})} = (S, A, \hookrightarrow, s_0, AP, L)$ and a formula $\varphi \in \text{LTL}_{\{\circ\}}$, $\alpha \in A$ is safe iff $\bigwedge_{i=1}^3 \text{Safe}_i(\alpha)$, where Safe_i , given in Table 1, are per-row.*

Proof. See [32] Appendix A.

Theorem 2 (POR instance for SDN). *Let $(\text{CP}, \lambda, \varphi)$ be a context such that $\mathcal{M}_{(\lambda, \text{CP})} = (S, A, \hookrightarrow, s_0, AP, L)$ is an SDN network model from Definition 1; and let safe actions be as in Definition 2. Further, let $\text{ample}(s)$ be defined by:*

$$\text{ample}(s) = \begin{cases} \{\alpha \in A(s) \mid \alpha \text{ safe}\} & \text{if } \{\alpha \in A(s) \mid \alpha \text{ safe}\} \neq \emptyset \\ A(s) & \text{otherwise} \end{cases}$$

Table 1. Safeness predicates

Action	Independence	Invisibility
$Safe_1(\alpha)$	$Safe_2(\alpha)$	$Safe_3(\alpha)$
$\alpha = ctrl(sw, pk, cs)$	CP is not order-sensitive	if $Q(q)$ occurs in φ , where $q \in CS$, then α is φ -invariant
$\alpha = bsync(sw, xid, cs)$	CP is not order-sensitive	if $Q(q)$ occurs in φ , where $q \in CS$, then α is φ -invariant
$\alpha = fwd(sw, pk, ports)$	\top	if $\exists pk \in b.q. P(pk)$ occurs in φ , for any $b \in \{sw\} \cup \{\lambda(sw, p)_1 \mid p \in ports\}$ and $q \in \{pq, recvq\}$, then α is φ -invariant
$\alpha = brepl(sw, xid)$	\top	\top
$\alpha = recv(h, pk)$	\top	if $\exists pk \in h.recvq. P(pk)$ occurs in φ , then α is φ -invariant

Then, *ample* satisfies the criteria of Theorem 1 and thus $\mathcal{M}_{(\lambda, CP)} \stackrel{st}{\equiv} \mathcal{M}_{(\lambda, CP)}^{ample}$ ⁵

Proof.

C1 The (non)emptiness condition is trivial since by definition of *ample*(s) it follows that *ample*(s) = \emptyset iff $A(s) = \emptyset$.

C2 By assumption $\beta \in A \setminus ample(s)$ depends on *ample*(s). But with our definition of *ample*(s) this is impossible as all actions in *ample*(s) are safe and by definition independent of all other actions.

C3 The validity of the invisibility condition is by definition of *ample* and safe actions.

C4 We now show that every cycle in $\mathcal{M}_{(\lambda, CP)}^{ample}$ contains a fully expanded state s , i.e. a state s such that *ample*(s) = $A(s)$. By definition of *ample*(s) in Theorem 2 it is equivalent to show that there is no cycle in $\mathcal{M}_{(\lambda, CP)}^{ample}$ consisting of safe actions only. We show this by contradiction, assuming such a cycle of only safe actions exists. There are five safe action types to consider: *ctrl*, *fwd*, *brepl*, *bsync* and *recv*. Distinguish two cases.

Case 1. A sequence of safe actions of same type. Let us consider the different safe actions:

- Let ρ an execution of $\mathcal{M}_{(\lambda, CP)}^{ample}$ which consists of only one type of *ctrl*-actions:

$$\rho = s_1 \xrightarrow{ctrl(pkt_1, cs_1)} s_2 \xrightarrow{ctrl(pkt_2, cs_2)} \dots s_{i-1} \xrightarrow{ctrl(pkt_{i-1}, cs_{i-1})} s_i$$

Suppose ρ is a cycle. According to the *ctrl* semantics, for each transition $s \xrightarrow{ctrl(pkt, cs)} s'$, where $s = (\pi, \delta, \gamma)$, $s' = (\pi', \delta', \gamma')$, it holds that $\gamma'.rq = \gamma.rq \setminus \{pkt\}$ as we use sets to represent *rq* buffers. Hence, for the execution ρ it holds $\gamma_i.rq = \gamma_1.rq \setminus \{pkt_1, pkt_2, \dots, pkt_{i-1}\}$ which implies that $s_1 \neq s_i$. Contradiction.

⁵ Stutter equivalence here implicitly is defined w.r.t. the atomic propositions appearing in φ , but this suffices as we are just interested in the validity of φ .

- Let ρ an execution which consists of only one type of *fwd*-actions: similar argument as above since *fq*-s are represented by sets and thus forward messages are removed from *fq*.
- Let ρ an execution which consists of only one type of *brepl*-actions: similar argument as above since control messages are removed from *cq*.
- Let ρ an execution which consists of only one type of *bsync*-actions: similar argument as above, as barrier reply messages are removed from *brq*-s that are represented by sets.
- Let ρ an execution which consists of only one type of *recv*-actions: similar argument as above, as packets are removed from *recv* buffers that are represented by sets.

Case 2. A sequence of different safe actions. Suppose there exists a cycle with mixed safe actions starting in s_1 and ending in s_i . Distinguish the following cases.

- There exists at least a *ctrl* and/or a *bsync* action in the cycle. According to the effects of safe transitions, the *ctrl* action will change to a state with smaller *rq* and the *bsync* will always switch to a state with smaller *brq*. It is important here that *ctrl* does not interfere with *bsync* regarding *rq*, *brq*, and no safe action of other type than *ctrl* and *bsync* accesses *rq* or *brq*. This implies that $s_1 \neq s_i$. Contradiction.
- Neither *ctrl*, nor *bsync* actions in the cycle.
 - There is a *fwd* and/or *brepl* in the cycle: *fwd* will always switch to a state with smaller *fq* and *brepl* will always switch to a state with smaller *cq* (*brepl* and *recv* do not interfere with *fwd*). This implies that $s_1 \neq s_i$. Contradiction.
 - There is neither *fwd* nor *brepl* in the cycle. This means that only *recv* is in the cycle which is already covered by the first case.

□

Due to the definition of the transition system via ample sets, each safe action is immediately executed after its enabling one. Therefore, one can merge every transition of a safe action with its precursory enabling one. Intuitively, the semantics of the merged action is defined as the successive execution of its constituent actions. This process can be repeated if there is a chain of safe actions; for instance, in the case of $s \xrightarrow{\text{nomatch}(sw, pkt)} s' \xrightarrow{\text{ctrl}(sw, pkt, cs)} s'' \xrightarrow{\text{fwd}(sw, pkt, ports)} s'''$ where each transition enables the next and the last two are assumed to be safe. These transitions can be merged into one, yielding a stutter equivalent trace as the intermediate states are invisible (w.r.t. the context and thus the property to be shown) by definition of safe actions.

3.2 State Representation

Efficient state representation is crucial for minimising MOCS's memory footprint and enabling it to scale up to relatively large network setups.

Packet and Rule Indexing. In MOCS, only a single instance of each packet and rule that can appear in the modelled network is kept in memory. An index is then used to associate queues and flow tables with packets and rules, with a single bit indicating their presence (or absence). This data structure is illustrated in Fig. 3. For a data packet, a value of 1 in the *pq* section of the entry indicates that infinite copies of it are stored in the packet queue of the respective switch. A value of 1 in the *fq* section indicates that a single copy of the packet is stored in the forward queue of the respective switch. A value of 1 in the *rq* section indicates that a copy of the packet sent by the respective switch (when a *nomatch* transition is fired) is stored in the controller’s request queue. For a rule, a value of 1 in the *ft* section indicates that the rule is installed in the respective switch’s flow table. A value of 1 in the *cq* section indicates that the rule is part of a *FlowMod* message in the respective switch’s control queue.

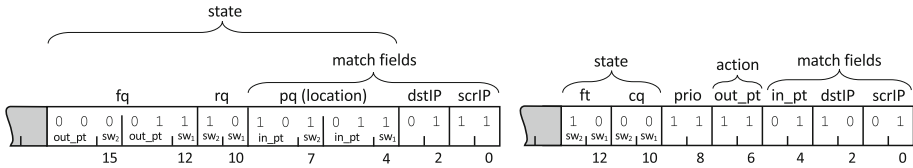


Fig. 3. Packet (left) and rule (right) indices

The proposed optimisation enables scaling up the network topology by minimising the required memory footprint. For every switch, MOCS only requires a few bits in each packet and rule entry in the index.

Discovering Equivalence Classes of Packets. Model checking with all possible packets, including all specified fields in the OpenFlow standard, would entail a huge state space that would render any approach unusable. Here, we propose the discovery of equivalence classes of packets that are then used for model checking. We first remove all fields that are not referenced in a statement or rule creation or deletion in the controller program. Then, we identify packet classes that would result in the same controller behaviour. Currently, as with the rest of literature, we focus on simple controller programs where such equivalence classes can be easily identified by analysing static constraints and rule manipulation in the controller program. We then generate one representative packet from each class and assign it to all network switches that are directly connected to end-hosts; i.e. modelling clients that can send an arbitrarily large number of packets in a non-deterministic fashion. We use the minimum possible number of bits to represent the identified equivalence classes. For example, if the controller program exerts different behaviour if the destination TCP port of a packet is 23 (i.e. destined to an SSH server) or not, we only use a 1-bit field to model this behaviour.

Bit Packing. We reduce the size of each recorded state by employing bit packing using the *int* type supported by UPPAAL, and bit-level operations for the entries in the packet and rule indices as well as for the packets and rules themselves.

4 Experimental Evaluation

In this section, we experimentally evaluate MOCS by comparing it with the state of the art, in terms of performance (verification throughput and memory footprint) and model expressivity. We have implemented MOCS in UPPAAL [6] as a network of parallel automata for the controller and network switches, which communicate asynchronously by writing/reading packets to/from queues that are part of the model discussed in Sect. 2. As discussed in Sect. 3, this is implemented by directly manipulating the packet and rule indices.

Throughout this section we will be using three examples of network controllers: (1) A *stateless firewall* ([32] Appendix B-CP1) requires the controller to install rules to network switches that enable them to decide whether to forward a packet towards its destination or not; this is done in a stateless fashion, i.e. without having to consider any previously seen packets. For example, a controller could configure switches to block all packets whose destination TCP port is 22 (i.e. destined to an SSH server). (2) A *stateful firewall* ([32] Appendix B-CP2) is similar to the stateless one but decisions can take into account previously seen packets. A classic example of this is to allow bi-directional communication between two end-hosts, when one host opens a TCP connection to the other. Then, traffic flowing from the other host back to the connection initiator should be allowed to go through the switches on the reverse path. (3) A *MAC learning application* ([32] Appendix B-CP3) enables the controller and switches to learn how to forward packets to their destinations (identified with respective MAC addresses). A switch sends a *PacketIn* message to the controller when it receives a packet that it does not know how to forward. By looking at this packet, the controller learns a mapping of a source switch (or host) to a port of the requesting switch. It then installs a rule (by sending a *FlowMod* message) that will allow that switch to forward packets back to the source switch (or host), and asks the requesting switch (by sending a *PacketOut* message) to flood the packet to all its ports except the one it received the packet from. This way, the controller eventually learns all mappings, and network switches receive rules that enable them to forward traffic to their neighbours for all destinations in the network.

4.1 Performance Comparison

We measure MOCS’s performance, and also compare it against Kuai [35]⁶ using the examples described above, and we investigate the behaviour of MOCS as we scale up the network (switches and clients/servers). We report three metrics:

⁶ Note that parts of Kuai’s source code are not publicly available, therefore we implemented it’s model in UPPAAL.

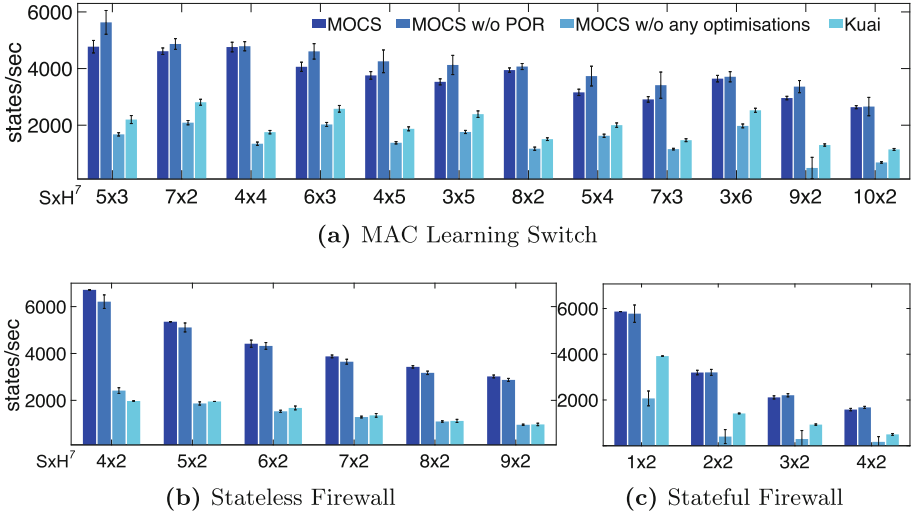


Fig. 4. Performance comparison – verification throughput

(1) *verification throughput* in visited states per second, (2) number of visited states, and (3) required memory. We have run all verification experiments on an 18-Core iMac pro, 2.3 GHz Intel Xeon W with 128 GB DDR4 memory.

Verification Throughput. We measure the verification throughput when running a single experiment at a time on one CPU core and report the average and standard deviation for the first 30 min of each run. In order to assess how MOCS’s different optimisations affect its performance, we report results for the following system variants: (1) MOCS, (2) MOCS without POR, (3) MOCS without any optimisations (neither POR, state representation), and (4) Kuai. Figure 4 shows the measured throughput (with error bars denoting standard deviation).

For the MAC learning and stateless firewall applications, we observe that MOCS performs significantly better than Kuai for all different network setups and sizes⁷, achieving at least double the throughput Kuai does. The throughput performance is much better for the stateful firewall, too. This is despite the fact that, for this application, Kuai employs the unrealistic optimisation where the *barrier* transition forces the immediate update of the forwarding state. In other words, MOCS is able to explore significantly more states and identify bugs that Kuai cannot (see Sect. 4.2).

The computational overhead induced by our proposed PORs is minimal. This overhead occurs when PORs require dynamic checks through the safety predicates described in Table 1. This is shown in Fig. 4a, where, in order to decide about the (in)visibility of $fwd(sw, pk, pt)$ actions, a lookup is performed in the history-array of packet pk , checking whether the bit which corresponds to switch sw' , which is connected with port pt of sw , is set. On the other hand, if a POR does not require any dynamic checks, no penalty is induced, as shown in Figs. 4b

⁷ $S \times H$ in Figs. 4 to 6 indicates the number of switches S and hosts H .

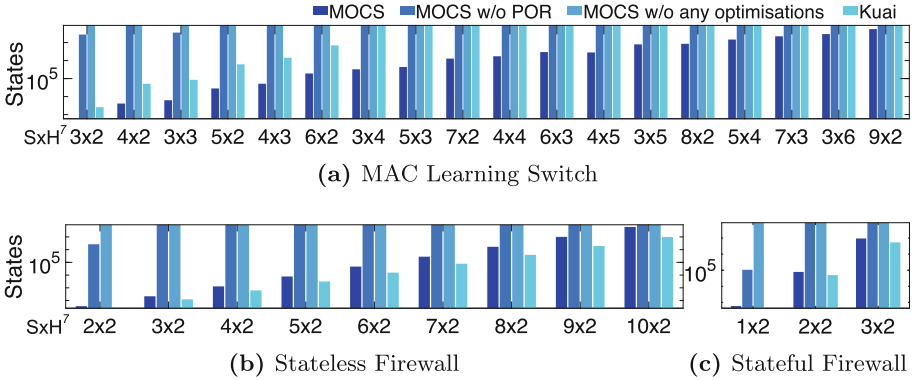


Fig. 5. Performance comparison – visited states (logarithmic scale)

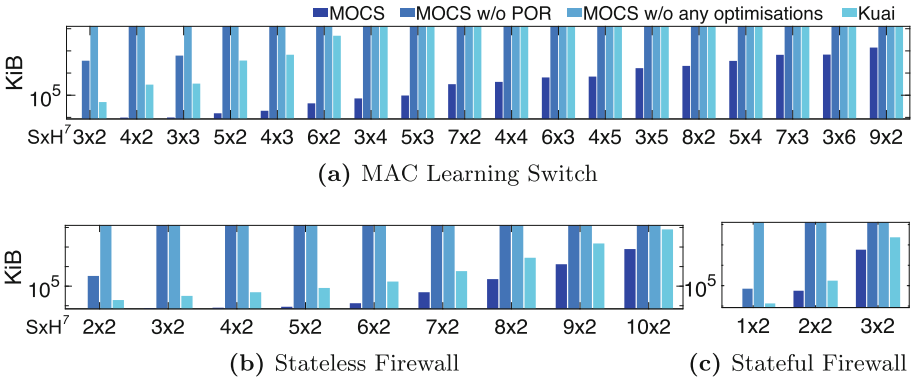


Fig. 6. Performance comparison – memory footprint (logarithmic scale)

and 4c, where the throughput when the PORs are disabled is almost identical to the case where PORs are enabled. This is because it has been statically established at a pre-analysis stage that all actions of a particular type are always safe for any argument/state. It is important to note that even when computational overhead is induced, PORs enable MOCS to scale up to larger networks because the number of visited states can be significantly reduced, as discussed below.

In order to assess the contribution of the state representation optimisation in MOCS’s performance, we measure the throughput when both PORs and state representation optimisations are disabled. It is clear that they contribute significantly to the overall throughput; without these the measured throughput was at least less than half the throughput when they were enabled.

Number of Visited States and Required Memory. Minimising the number of visited states and required memory is crucial for scaling up verification to larger networks. The proposed partial order reductions (Sect. 3.1) and identification of packet equivalent classes aim at the former, while packet/rule indexing

and bit packing aim at the latter (§3.2). In Fig. 5, we present the results for the various setups and network deployments discussed above. We stopped scaling up the network deployment for each setup when the verification process required more than 24 h or started swapping memory to disk. For these cases we killed the process and report a topped-up bar in Figs. 5 and 6.

For the MAC learning application, MOCS can scale up to larger network deployments compared to Kuai, which could not verify networks consisting of more than 2 hosts and 6 switches. For that network deployment, Kuai visited ~ 7 m states, whereas MOCS visited only ~ 193 k states. At the same time, Kuai required around 48 GBs of memory (7061 bytes/state) whereas MOCS needed ~ 43 MBs (228 bytes/state). Without the partial order reductions, MOCS can only verify tiny networks. The contribution of the proposed state representation optimisations is also crucial; in our experiments (results not shown due to lack of space), for the 6×2 network setups (the largest we could do without these optimisations), we observed a reduction in state space (due to the identification of packet equivalence classes) and memory footprint (due to packet/rule indexing and bit packing) from ~ 7 m to ~ 200 k states and from ~ 6 KB per state to ~ 230 B per state. For the stateless and stateful firewall applications, resp., MOCS performs equally well to Kuai with respect to scaling up.

4.2 Model Expressivity

The proposed model is significantly more expressive compared to Kuai as it allows for more asynchronous concurrency. To begin with, in MOCS, controller messages sent before a barrier request message can be interleaved with all other enabled actions, other than the control messages sent after the barrier. By contrast, Kuai always flushes all control messages until the last barrier in one go, masking a large number of interleavings and, potentially, buggy behaviour. Next, in MOCS *nomatch*, *ctrl* and *fwd* can be interleaved with other actions. In Kuai, it is enforced a mutual exclusion concurrency control policy through the *wait*-semaphore: whenever a *nomatch* occurs the mutex is locked and it is unlocked by the *fwd* action of the thread *nomatch-ctrl-fwd* which refers to the same packet; all other threads are forced to wait. Moreover, MOCS does not impose any limit on the size of the *rq* queue, in contrast to Kuai where only one packet can exist in it. In addition, Kuai does not support notifications from the data plane to the controller for completed operations as it does not support reply messages and as a result any bug related to the fact that the controller is not synced to data-plane state changes is hidden.⁸ Also, our specification language for states is more expressive than Kuai’s, as we can use any property in LTL without “next”, whereas Kuai only uses invariants with a single outermost \square .

The MOCS extensions, however, are conservative with respect to Kuai, that is we have the following theorem (without proof, which is straightforward):

⁸ There are further small extensions; for instance, in MOCS the controller can send multiple *PacketOut* messages (as OpenFlow prescribes).

Theorem 3 (MOCS Conservativity). *Let $\mathcal{M}_{(\lambda, \text{CP})} = (S, A, \hookrightarrow, s_0, AP, L)$ and $\mathcal{M}_{(\lambda, \text{CP})}^K = (S_K, A_K, \hookrightarrow_K, s_0, AP, L)$ the original SDN models of MOCS and Kuai, respectively, using the same topology and controller. Furthermore, let $\text{Traces}(\mathcal{M}_{(\lambda, \text{CP})})$ and $\text{Traces}(\mathcal{M}_{(\lambda, \text{CP})}^K)$ denote the set of all initial traces in these models, respectively. Then, $\text{Traces}(\mathcal{M}_{(\lambda, \text{CP})}^K) \subseteq \text{Traces}(\mathcal{M}_{(\lambda, \text{CP})})$.*

For each of the extensions mentioned above, we briefly describe an example (controller program and safety property) that expresses a bug that is impossible to occur in Kuai.

Control Message Reordering Bug. Let us consider a stateless firewall in Fig. 7a (controller is not shown), which is supposed to block incoming SSH packets from reaching the server (see [32] Appendix B-CP1). Formally, the safety property to be checked here is $\Box(\forall \text{pkt} \in S.\text{rcvq} . \neg \text{pkt}.\text{SSH})$. Initially, flow tables are empty. Switch A sends a *PacketIn* message to the controller when it receives the first packet from the client (as a result of a *nomatch* transition). The controller, in response to this request (and as a result of a *ctrl* transition), sends the following *FlowMod* messages to switch A ; rule $\mathbf{r1}$ has the highest priority and drops all SSH packets, rule $\mathbf{r2}$ sends all packets from port 1 to port 2, and rule $\mathbf{r3}$ sends all packets from port 2 to port 1. If the packet that triggered the transition above is an SSH one, the controller drops it, otherwise, it instructs (through a *PacketOut* message) A to forward the packet to S . A bug-free controller should ensure that $\mathbf{r1}$ is installed before any other rule, therefore it must send a barrier request after the *FlowMod* message that contains $\mathbf{r1}$. If, by mistake, the *FlowMod* message for $\mathbf{r2}$ is sent before the barrier request, A may install $\mathbf{r2}$ before $\mathbf{r1}$, which will result in violating the given property. MOCS is able to capture this buggy behaviour as its semantics allows control messages prior to the barrier to be processed in an interleaved manner.

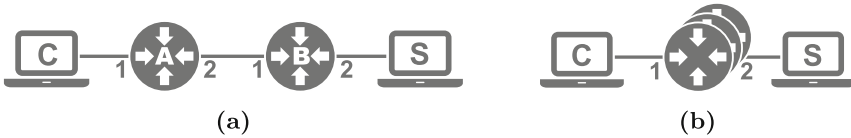


Fig. 7. Two networks with (a) two switches, and (b) n stateful firewall replicas

Wrong Nesting Level Bug. Consider a correct controller program that enforces that server S (Fig. 7a) is not accessible through SSH. Formally, the safety property to be checked here is $\Box(\forall \text{pkt} \in S.\text{rcvq} . \neg \text{pkt}.\text{SSH})$. For each incoming *PacketIn* message from switch A , it checks if the enclosed packet is an SSH one and destined to S . If not, it sends a *PacketOut* message instructing A to forward the packet to S . It also sends a *FlowMod* message to A with a rule that allows packets of the same protocol (not SSH) to reach S . In the opposite case (SSH), it checks (a Boolean flag) whether it had previously sent drop rules for SSH packets to the switches. If not, it sets flag to true, sends a *FlowMod* message with a rule

that drops SSH packets to A and drops the packet. Note that this inner block does not have an `else` statement.

A fairly common error is to write a statement at the wrong nesting level ([32] Appendix B-CP4). Such a mistake can be built into the above program by nesting the outer `else` branch in the inner `if` block, such that it is executed any time an SSH-packet is encountered but the SSH drop-rule has already been installed (i.e. flag `f` is true). Now, the SSH drop rule, once installed in switch A , disables immediately a potential $nomatch(A, p)$ with $p.ssh = true$ that would have sent packet p to the controller, but if it has not yet been installed, a second incoming SSH packet would lead to the execution of the `else` statement of the inner branch. This would violate the property defined above, as p will be forwarded to S ⁹.

MOCS can uncover this bug because of the correct modelling of the controller request queue and the asynchrony between the concurrent executions of control messages sent before a barrier. Otherwise, the second packet that triggers the execution of the wrong branch would not have appeared in the buffer before the first one had been dealt with by the controller. Furthermore, if all rules in messages up to a barrier were installed synchronously, the second packet would be dealt with correctly, so no bug could occur.

Inconsistent Update Bug. OpenFlow’s barrier and barrier reply mechanisms allow for updating multiple network switches in a way that enables *consistent packet processing*, i.e., a packet cannot see a partially updated network where only a subset of switches have changed their forwarding policy in response to this packet (or any other event), while others have not done so. MOCS is expressive enough to capture this behaviour and related bugs. In the topology shown in Fig. 7a, let us assume that, by default, switch B drops all packets destined to S . Any attempt to reach S through A are examined separately by the controller and, when granted access, a relevant rule is installed at both switches (e.g. allowing all packets from C destined to S for given source and destination ports). Updates must be consistent, therefore the packet cannot be forwarded by A and dropped by B . Both switches must have the new rules in place, before the packet is forwarded. To do so, the controller, ([32] Appendix B-CP5), upon receiving a *PacketIn* message from the client’s switch, sends the relevant rule to switch B (*FlowMod*) along with respective barrier (*BarrierReq*) and temporarily stores the packet that triggered this update. Only after receiving *BarrierRes* message from B , the controller will forward the previously stored packet back to A along with the relevant rule. This update is consistent and the packet is guaranteed to reach S . A (rather common) bug would be one where the controller installs the rules to both switches and at the same time forwards the packet to A . In this case, the packet may end up being dropped by B , if it arrives and gets processed before the relevant rule is installed, and therefore the invariant $\square([drop(pkt, sw)] . \neg(pkt.dest = S))$, where $[drop(pkt, sw)]$ is a quantifier that binds dropped packets (see definition in [32] Appendix B-CP5), would

⁹ Here, we assume that the controller looks up a static forwarding table before sending *PacketOut* messages to switches.

be violated. For this example, it is crucial that MOCS supports barrier response messages.

5 Conclusion

We have shown that an OpenFlow compliant SDN model, with the right optimisations, can be model checked to discover subtle real-world bugs. We proved that MOCS can capture real-world bugs in a more complicated semantics without sacrificing performance.

But this is not the end of the line. One could automatically compute equivalence classes of packets that cover all behaviours (where we still computed manually). To what extent the size of the topology can be restricted to find bugs in a given controller is another interesting research question, as is the analysis of the number and length of interleavings necessary to detect certain bugs. In our examples, all bugs were found in less than a second.

References

1. Al-Fares, M., Radhakrishnan, S., Raghavan, B.: Hedera: dynamic flow scheduling for data center networks. In: NSDI (2010)
2. Al-Shaer, E., Al-Haj, S.: FlowChecker: configuration analysis and verification of federated OpenFlow infrastructures. In: SafeConfig (2010)
3. Albert, E., Gómez-Zamalloa, M., Rubio, A., Sammartino, M., Silva, A.: SDN-Actors: modeling and verification of SDN programs. In: Havelund, K., Peleska, J., Roscoe, B., de Vink, E. (eds.) FM 2018. LNCS, vol. 10951, pp. 550–567. Springer, Cham (2018)
4. Baier, C., Katoen, J.P.: Principles of Model Checking. The MIT Press, Cambridge (2008)
5. Ball, T., Bjørner, N., Gember, A., et al.: VeriCon: towards verifying controller programs in software-defined networks. In: PLDI (2014)
6. Behrmann, G., David, A., Larsen, K.G., et al.: Developing UPPAAL over 15 years. In: Practice and Experience, Software (2011)
7. Braga, R., Mota, E., Passito, A.: Lightweight DDoS flooding attack detection using NOX/OpenFlow. In: LCN (2010)
8. Canini, M., Venzano, D., Perešíni, P., et al.: A NICE way to test OpenFlow applications. In: NSDI (2012)
9. Cimatti, A., et al.: NuSMV 2: an OpenSource tool for symbolic model checking. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 359–364. Springer, Heidelberg (2002)
10. Curtis, A.R., Mogul, J.C., Tourrilhes, J., et al.: DevoFlow: scaling flow management for high-performance networks. In: SIGCOMM (2011)
11. Dobrescu, M., Argyraki, K.: Software dataplane verification. In: Communications of the ACM (2015)
12. El-Hassany, A., Tsankov, P., Vanbever, L., Vechev, M.: Network-wide configuration synthesis. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10427, pp. 261–281. Springer, Cham (2017)

13. Fayaz, S.K., Sharma, T., Fogel, A., et al.: Efficient network reachability analysis using a succinct control plane representation. In: OSDI (2016)
14. Fayaz, S.K., Yu, T., Tobioka, Y., et al.: BUZZ: testing context-dependent policies in stateful networks. In: NSDI (2016)
15. Feamster, N., Rexford, J., Shenker, S., et al.: SDX: A Software-defined Internet Exchange. Open Networking Summit (2013)
16. Feamster, N., Rexford, J., Zegura, E.: The road to SDN. SIGCOMM Comput. Commun. Rev. (2014)
17. Fischer, M.J., Ladner, R.E.: Propositional dynamic logic of regular programs. J. Comput. Syst. Sci. **18**, 194–211 (1979)
18. Fogel, A., Fung, S., Angeles, L., et al.: A general approach to network configuration analysis. In: NSDI (2015)
19. Handigol, N., Seetharaman, S., Flajslik, M., et al.: Plug-n-Serve: load-balancing web traffic using OpenFlow. In: SIGCOMM (2009)
20. Havelund, K., Pressburger, T.: Model checking JAVA programs using JAVA PathFinder. STTT **2**, 366–381 (2000)
21. Holzmann, G.J.: The model checker SPIN. IEEE Trans. Softw. Eng. **23**, 279–295 (1997)
22. Holzmann, G.J., Peled, D.: An improvement in formal verification. In: Hogrefe D., Leue S. (eds) Formal Description Techniques VII. IAICT, pp. 197–211. Springer, Boston, MA (1995)
23. Horn, A., Kheradmand, A., Prasad, M.R.: Delta-net: real-time network verification using atoms. In: NSDI (2017)
24. Hu, H., Ahn, G.J., Han, W., et al.: Towards a reliable SDN firewall. In: ONS (2014)
25. Jackson, D.: Alloy: a lightweight object modelling notation. ACM Trans. Softw. Eng. Methodol. **11**, 256–290 (2002)
26. Jafarian, J.H., Al-Shaer, E., Duan, Q.: OpenFlow random host mutation: transparent moving target defense using software defined networking. In: HotSDN (2012)
27. Jain, S., Zhu, M., Zolla, J., et al.: B4: experience with a globally-deployed software defined WAN. In: SIGCOMM (2013)
28. Jia, Y.: NetSMC: a symbolic model checker for stateful network verification. In: NSDI (2020)
29. Kazemian, P., Chang, M., Zeng, H., et al.: Real time network policy checking using header space analysis. In: NSDI (2013)
30. Kazemian, P., Varghese, G., McKeown, N.: Header space analysis: static checking for networks. In: NSDI (2012)
31. Khurshid, A., Zou, X., Zhou, W., et al.: VeriFlow: verifying network-wide invariants in real time. In: NSDI (2013)
32. Klimis, V., Parisi, G., Reus, B.: Towards model checking real-world software-defined networks (version with appendix). preprint [arXiv:2004.11988](https://arxiv.org/abs/2004.11988) (2020)
33. Li, Y., Yin, X., Wang, Z., et al.: A survey on network verification and testing with formal methods: approaches and challenges. IEEE Surv. Tutorials **21**, 940–969 (2019)
34. Mai, H., Khurshid, A., Agarwal, R., et al.: Debugging the data plane with ant eater. In: SIGCOMM (2011)
35. Majumdar, R., Deep Tetali, S., Wang, Z.: Kuai: a model checker for software-defined networks. In: FMCAD (2014)
36. McClurg, J., Hojjat, H., Černý, P., et al.: Efficient synthesis of network updates. In: PLDI (2015)
37. McKeown, N., Anderson, T., Balakrishnan, H., et al.: OpenFlow: enabling innovation in campus networks. SIGCOMM Comput. Commun. Rev. **38**, 69–74 (2008)

38. Patel, P., Bansal, D., Yuan, L., et al.: Ananta: cloud scale load balancing. SIGCOMM **43**, 207–218 (2013)
39. Peled, D.: All from one, one for all: on model checking using representatives. In: Courcoubetis, C. (ed.) CAV 1993. LNCS, vol. 697, pp. 409–423. Springer, Heidelberg (1993)
40. Plotkin, G.D., Bjørner, N., Lopes, N.P., et al.: Scaling network verification using symmetry and surgery. In: POPL (2016)
41. Pnueli, A., Xu, J., Zuck, L.: Liveness with $(0,1, \infty)$ - counter abstraction. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 107–122. Springer, Heidelberg (2002)
42. Pratt, V.R.: Semantical considerations on Floyd-Hoare logic. In: FOCS (1976)
43. Sethi, D., Narayana, S., Malik, S.: Abstractions for model checking SDN controllers. In: FMCAD (2013)
44. Shenker, S., Casado, M., Koponen, T., et al.: The future of networking, and the past of protocols. In: ONS (2011). <https://tinyurl.com/yxnuxobt>
45. Son, S., Shin, S., Yegneswaran, V., et al.: Model checking invariant security properties in OpenFlow. In: IEEE (2013)
46. Stoenescu, R., Popovici, M., Negreanu, L., et al.: SymNet: scalable symbolic execution for modern networks. In: SIGCOMM (2016)
47. Varghese, G.: Vision for network design automation and network verification. In: NetPL (Talk) (2018). <https://tinyurl.com/y2cnhvhf>
48. Yang, H., Lam, S.S.: Real-time verification of network properties using atomic predicates. IEEE/ACM Trans. Network. **24**, 887–900 (2016)
49. Zeng, H., Kazemian, P., Varghese, G., et al.: A survey on network troubleshooting. Technical report TR12-HPNG-061012, Stanford University (2012)
50. Zeng, H., Zhang, S., Ye, F., et al.: Libra: divide and conquer to verify forwarding tables in huge networks. In: NSDI (2014)
51. Zhang, S., Malik, S.: SAT based verification of network data planes. In: Van Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 496–505. Springer, Cham (2013)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

