

Network-Wide Configuration Synthesis

Ahmed El-Hassany^(✉), Petar Tsankov, Laurent Vanbever,
and Martin Vechev

ETH Zürich, Zürich, Switzerland
eahmed@ethz.ch

Abstract. Computer networks are hard to manage. Given a set of high-level requirements (e.g., reachability, security), operators have to manually figure out the individual configuration of potentially hundreds of devices running complex distributed protocols so that they, collectively, compute a compatible forwarding state. Not surprisingly, operators often make mistakes which lead to downtimes.

To address this problem, we present a novel synthesis approach that automatically computes correct network configurations that comply with the operator’s requirements. We capture the behavior of existing routers along with the distributed protocols they run in stratified Datalog. Our key insight is to reduce the problem of finding correct input configurations to the task of synthesizing inputs for a stratified Datalog program.

To solve this synthesis task, we introduce a new algorithm that synthesizes inputs for stratified Datalog programs. This algorithm is applicable beyond the domain of networks.

We leverage our synthesis algorithm to construct the first network-wide configuration synthesis system, called SyNET, that support multiple interacting routing protocols (OSPF and BGP) and static routes. We show that our system is practical and can infer correct input configurations, in a reasonable amount time, for networks of realistic size (>50 routers) that forward packets for multiple traffic classes.

1 Introduction

Despite being mission-critical for most organizations, managing a network is surprisingly hard and brittle.

A key reason is that network operators have to manually come up with a configuration, which ensures that the underlying distributed protocols compute a forwarding state that satisfies the operator’s requirements.

Doing so requires operators to precisely understand: (i) the behavior of each distributed protocol; (ii) how the protocols interact with each other; and (iii) how each parameter in the configuration affects the distributed computation.

Because of this complexity, operators often make mistakes that can lead to severe network downtimes. As an illustration, Facebook (and Instagram) recently suffered from widespread issues for about an hour due to a misconfiguration [1]. In fact, studies show that most network downtimes are caused by humans, not equipment failures [2]. Such misconfigurations can have Internet-wide effects [3].

To prevent misconfigurations, researchers have developed tools that check if a given configuration is correct [4–7]. While useful, these works still require network operators to produce the configurations in the first place. Template-based approaches [8–11] along with vendor-agnostic abstractions [12–14] have been proposed to reduce the configuration burden. However, they still require operators to understand precisely the details of each protocol. Recently, Software-Defined Networks (SDNs) have emerged as another paradigm to manage networks by *programming* them from a central controller. Deploying SDN is, however, a major hurdle as it requires new network devices *and* management tools. Further, designing correct, robust and yet, scalable, SDN controllers is challenging [15–18]. Because of this, only a handful of networks are using SDN in production. As a result, configuring individual devices is by far the most widespread (and default) way to manage networks.

Problem Statement: Network-Wide Configuration Synthesis. Ideally, from a network operator perspective, one would like to solve what we refer to as the *Network-Wide Configuration Synthesis* problem: *Given a network specification \mathcal{N} , which defines the behavior of all routing protocols run by the routers, and a set \mathcal{R} of requirements on the network-wide forwarding state, discover a configuration \mathcal{C} such that the routers converge to a forwarding state compatible with \mathcal{R} .* That is, the operator simply provides the high-level requirements \mathcal{R} , and the configuration \mathcal{C} is obtained automatically.

Distributed vs. Static Routing. Relying as much as possible on distributed protocols to compute the forwarding state is critical to ensure network reliability and scalability. A simpler problem would be to statically configure the forwarding entries of each router via static routes (e.g. see [19, 20]). Relying solely on static routes is, however, undesirable for two reasons. First, they prevent routers from reacting locally upon failure. Second, they can be costly to update as routers often have a large number of static entries.

Key Challenges. Coming up with a solution to the network-wide synthesis problem is challenging for at least three reasons: (i) *Diversity*: protocols have different expressiveness in terms of the forwarding entries they compute. For instance, the Open Shortest Path First protocol (OSPF) can only direct traffic along shortest-paths, while the Border Gateway Protocol (BGP) can direct traffic along non-shortest paths. Conversely, BGP cannot forward traffic along multiple paths by default¹, while OSPF supports multi-path routing and is thus better suited for load-balancing traffic, a feature heavily used in practice. (ii) *Dependence*: distinct protocols often depend on one another, making it challenging to ensure that they collectively compute a compatible forwarding state. For instance, BGP depends on the network-wide intra-domain configuration; and (iii) *Feasibility*: the search space of configurations is massive and it is thus difficult to find one that leads to a forwarding state satisfying the requirements.

¹ While vendor-specific workarounds to make BGP multipath exist, these break the congruency between the control and data plane and could lead to correctness issues.

This Work. In this paper, we provide the first solution to the network-wide synthesis problem. Our approach is based on two steps. First, we use stratified Datalog to capture the behavior of the network, i.e. the distributed protocols ran by the routers together with any protocol dependencies. Datalog is indeed particularly well-suited for describing these protocols in a clear and declarative way. Here, the fixed point of a Datalog program represents the stable forwarding state of the network. Second, and a key insight of our work: we pose the network-wide synthesis problem as an instance of finding an input for a stratified Datalog program where the program’s fixed point satisfies a given property. That is, the network operator simply provides the high-level requirements \mathcal{R} on the forwarding state (i.e., which is the same as requiring the Datalog program’s fixed point to satisfy \mathcal{R}), and our synthesizer automatically finds an input \mathcal{C} to the Datalog program (i.e., which identifies the wanted network-wide configuration). We remark that our Datalog input synthesis algorithm is a general, independent contribution, and is applicable beyond networks.

Main Contributions. To summarize, our main contributions are:

- A formulation of the network-wide synthesis problem in terms of input synthesis for stratified Datalog (Sect. 2).
- The first input synthesis algorithm for stratified Datalog. This algorithm is of broader interest and is applicable beyond networks (Sect. 5).
- An instantiation and an end-to-end implementation of our input synthesis algorithm to the network-wide synthesis problem, along with network-specific optimizations, in a system called SyNET.
- An evaluation of SyNET on networks with multiple interacting widely-used protocols. In addition, we test the correctness of the generated configurations on an emulated network environment. Our results show that SyNET can automatically synthesize input configurations for networks of realistic size (>50 routers) carrying multiple traffic classes (Sect. 6).

2 Network-Wide Configuration Synthesis

We now illustrate our configuration synthesis approach on a simple example. We highlight how, given a network and a set of requirements, we can pose the synthesis problem as an instance of input synthesis for stratified Datalog.

2.1 Motivating Example

We consider the simple network topology, depicted in Fig. 1(b), composed of 4 routers denoted A , B , C and D . Routers B and C can reach the external network Ext , and router D is directly connected to two internal networks N1 and N2 . In the following, we use the term traffic class to refer to a set of packets (e.g. packets destined to N1) that are handled analogously according to the requirements. In practice, each traffic class may contain thousands of IP prefixes [21].

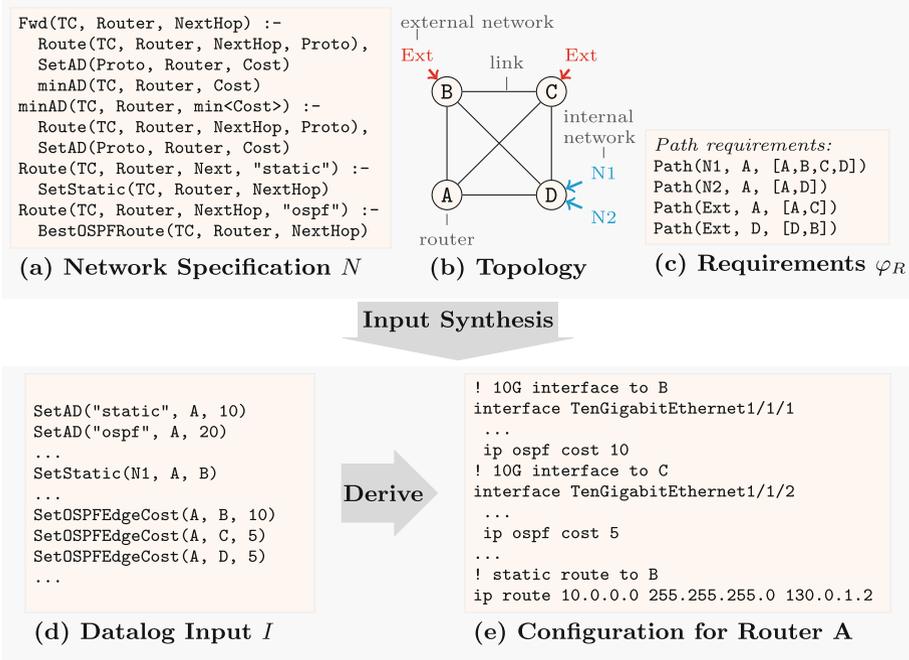


Fig. 1. Network-wide Configuration Synthesis. The input is: (a) declarative network specification N in stratified Datalog (b) network topology, and (c) routing requirements φ_R . The output is: (d) a Datalog input I that results in a forwarding state satisfying the requirements. Configurations (e) are derived from I .

Computation of Forwarding State. We first informally describe how each router’s forwarding entries are computed, assuming the configuration is provided.

Each router runs both, OSPF and BGP protocols, and in addition can also be configured with static routes. The computation of OSPF is based on finding least-cost paths to the internal destinations as well as to all routers in the network, where cost is the sum of the link weights defined in router configurations. The least-cost paths are then used to generate forwarding entries at each router to all internal destinations. In our example, these internal destinations are $N1$ and $N2$. In contrast, BGP computes forwarding entries to reach external destinations, Ext in our example. The computed forwarding entries define the next hop router for each destination. For example, BGP computes an entry at router A for Ext which forwards packets to a border router (i.e., either B or C). To decide which router the entry should forward to, each BGP router selects the egress point (i.e., border router) to reach an externally-learned prefix based on a preference value. This preference is (typically) defined in the configuration of each border router and propagated network-wide. If multiple routers announce the same preference for a prefix, internal BGP routers directs traffic to the closest egress point, according to the OSPF costs.

Once BGP and OSPF have finished computing their forwarding entries, each router takes these entries (along with those defined via static routes) and selects the OSPF-, BGP-, static route- produced forwarding entry with the highest preference (in networking terms, higher preference means lower administrative cost) defined in its local configuration. The union of all forwarding entries obtained at the routers is referred to as the forwarding state of the network.

Configuration Synthesis. Next, we illustrate the opposite direction (and one this work focuses on): given requirements φ_R , find a configuration which the protocols use to compute a forwarding state (as described above) that satisfies φ_R .

Let us consider the four path requirements given in Fig. 1(c). The first two state that A must forward packets for the traffic classes $N1$ and $N2$ along the paths $A \rightarrow B \rightarrow C \rightarrow D$ and $A \rightarrow D$, respectively. Note that these two requirements might reflect a security policy in the network or generated by a traffic engineering optimization tool [22, 23]. These two requirements cannot be enforced using OSPF alone. The reason is that, as discussed, OSPF works by selecting the least-cost path (by summing the weights on the links) and there is no assignment of weights to links which would lead to least-cost paths that exactly match the two path requirements.

Yet, the two requirements can be enforced by: (i) generating a static route-based forwarding entry at A to forward packets for $N1$ to B ; (ii) configuring link weights so paths $A \rightarrow D$ and $B \rightarrow C \rightarrow D$ have the lowest OSPF costs from A to D and, respectively, from B to D ; and (iii) on router A , configure a higher preference for forwarding entries based on static routes than OSPF forwarding entries. Because a static route forwarding entry is only generated for destination $N1$ (from (i)) and not $N2$, this means the entry for $N1$ will forward the traffic to router B while the entry for $N2$ will be the OSPF generated one (from (ii)).

The last two path requirements state that A and D must forward packets destined to the traffic class Ext to C and B , respectively. The two path requirements can be satisfied by: (i) setting identical BGP router preferences at the local configurations of B and C ; and (ii) configuring link weights so that paths $A \rightarrow C$ and $D \rightarrow B$ have the lowest costs from A to C and from D to B , respectively. In this way, BGP will use the results from the OSPF least-cost paths to compute its forwarding entries to Ext . This is an example where BGP interacts with OSPF and uses information from its computation.

The following is the final configuration produced by our synthesizer (the synthesizer is discussed in later sections):

- weight 10 is assigned to link $A \rightarrow B$,
- weight 5 is assigned to links $B \rightarrow C$, $C \rightarrow D$, and $A \rightarrow C$,
- weight 4 is assigned to link $D \rightarrow B$,
- weight 100 is assigned to the remaining links,
- a static route- based forwarding entry is defined at router A to forward traffic for $N1$ to B , and
- the router preference for all routers is set to 100.

In Fig. 1(e), we illustrate an excerpt of router A 's local configuration.

Phrasing the Problem as Inputs Synthesis for Stratified Datalog. A key insight of our work is to pose the question of finding a network configuration as an instance of input synthesis for stratified Datalog.

First, we declaratively specify the behavior of the network, i.e. the distributed protocols that the routers run, the protocol interactions, and the network topology, as a stratified Datalog program N . As requirements usually pertain to the stable forwarding state, the stratified Datalog encoding captures the stable state of these routing protocols as opposed to intermediate computation steps. Few relevant Datalog rules are given in Fig. 1(a); we detail this specification step in Sect. 4. The resulting Datalog program derives a predicate `Fwd` that defines the forwarding entries computed by all routers, where `Fwd(TC, Router, NextHop)` is derived if `Router` forwards packets for traffic class `TC` to router `NextHop`.

Second, we can directly express routing requirements as constraints over the predicate `Fwd`. We denote these constraints with φ_R in Fig. 1.

Finally, an input I to the Datalog program N identifies a network-wide configuration. We formalize the network-wide configuration synthesis problem as:

Definition 1. *The network-wide configuration synthesis problem is:*
Input A declarative network specification N and routing requirements φ_R .
Output A Datalog input I such that $\llbracket N \rrbracket_I \models \varphi_R$, if such an input exists; otherwise, return *unsat*.

In our definition, $\llbracket N \rrbracket_I$ denotes the fixed point of the Datalog program N for the input I , and $\llbracket N \rrbracket_I \models \varphi_R$ holds if this fixed point satisfies the constraints φ_R .

Synthesizing inputs for stratified Datalog is, however, a difficult (and, in general, undecidable) problem [24]. The problem is, however, decidable if we fix a finite set of values to bound the set of inputs. This is reasonable in the context of networks, where values represent finitely many routers, interfaces, and configuration parameters.

To address the problem, we introduce a new iterative synthesis algorithm that partitions the Datalog program P into strata P_1, \dots, P_n , finds an input I_i for each stratum P_i and then construct an input I for the Datalog program P . Each stratum P_i is a semi-positive Datalog program that enjoys the property that if a predicate is derived by the rules after some number of steps, then it must be contained in the fixed point. We describe this algorithm in Sect. 5.

3 Background: Stratified Datalog

We briefly overview the syntax and semantics of stratified Datalog.

Syntax. Datalog’s syntax is given in Fig. 2. We use \bar{r} , \bar{l} , and \bar{t} to denote zero or more rules, literals, and terms separated by commas, respectively. A Datalog program is *well-formed* if for any rule $a \leftarrow \bar{l}$, we have $\text{vars}(a) \subseteq \text{vars}(\bar{l})$, where $\text{vars}(\bar{l})$ returns the set of variables in \bar{l} .

(Program) $P ::= \bar{r}$	(Literal) $l ::= a \mid \neg a$	(Variables) $X, Y \in Vars$
(Rule) $r ::= a \leftarrow \bar{l}$	(Predicates) $p, q \in Preds$	(Values) $v \in Vals$
(Atom) $a ::= p(\bar{t})$	(Term) $t ::= X \mid v$	

Fig. 2. Syntax of stratified Datalog

A predicate is called *extensional* if it appears only in the bodies of rules (right side of the rule), otherwise (if it appears at least once in a rule head) it is called *intensional*. We denote the sets of extensional and intensional predicates of a program P by $edb(P)$ and $idb(P)$, respectively.

A Datalog program P is *stratified* if its rules can be partitioned into strata P_1, \dots, P_n such that if a predicate p occurs in a positive (negative) literal in the body of a rule in P_i , then all rules with p in their heads are in a stratum P_j with $j \leq i$ ($j < i$). Stratification ensures that predicates that appear in negative literals are fully defined in lower strata.

We syntactically extend stratified Datalog with aggregate functions such as `min` and `max`. This extension is possible as stratified Datalog is equally expressive to Datalog with stratified aggregate functions; for details see [25].

Semantics. Let $\mathcal{A} = \{p(\bar{t}) \mid \bar{t} \subseteq Vals\}$ denote the set of all ground (i.e. variable-free) atoms. The complete lattice $(\mathcal{P}(\mathcal{A}), \subseteq, \cap, \cup, \emptyset, \mathcal{A})$ partially orders the set of interpretations $\mathcal{P}(\mathcal{A})$.

Given a substitution $\sigma \in Vars \rightarrow Vals$ mapping variables to values. Given an atom a , we will write $\sigma(a)$ for the ground atom obtained by replacing the variables in a according to σ ; e.g., $\sigma(p(X))$ returns the ground atom $p(\sigma(X))$. The consequence operator $T_P \in \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{A})$ for a program P is defined as

$$T_P(A) = A \cup \{\sigma(a) \mid a \leftarrow l_1 \dots l_n \in P, \forall l_i \in \bar{l}. A \vdash \sigma(l_i)\}$$

where $A \vdash \sigma(a)$ if $\sigma(a) \in A$ and $A \vdash \sigma(\neg a)$ if $\sigma(a) \notin A$.

An input for P is a set of ground atoms constructed using P 's extensional predicates. Let P be a program with strata P_1, \dots, P_n and I be an input for P . The model of P for I , denoted by $\llbracket P \rrbracket_I$, is M_n , where $M_0 = I$ and $M_i = \bigcap \{A \in \text{fp } T_{P_i} \mid A \subseteq M_{i-1}\}$ is the smallest fixed point of T_{P_i} that is greater than the lower stratum's model M_{i-1} .

4 Declarative Network Specification

In this section, we first describe how we declarative specify the behavior of the network as a Datalog program. Afterwards, we discuss how routing requirements are specified as constraints over the Datalog program's fixed point.

4.1 Specifying Networks

To faithfully capture a network's behavior, we model (i) the behavior of routing protocols and their interactions and (ii) the topology of the network.

Expressing Protocols in Stratified Datalog. We formalize individual routing protocols and how routers combine the forwarding entries computed by these protocols as a stratified Datalog program N . The Datalog program N derives the predicate $\text{Fwd}(\text{TC}, \text{Router}, \text{NextHop})$, which represents the network’s global forwarding state. In Fig. 1(a), for example, we show the relevant rules that define how the forwarding entries computed by OSPF are combined with those defined via static routes. The predicate $\text{Route}(\text{TC}, \text{Router}, \text{NextHop}, \text{Proto})$ captures the forwarding entries of OSPF and static routes. The top Datalog rule states that routers select, for each traffic class TC , the forwarding entry with the minimal administrative cost (minAD) calculated over all protocols via the second Datalog rule in Fig. 1(a). The bottom two rules define the predicate Route , which collects the forwarding entries defined via static routes and computed by OSPF. We remark that OSPF routes (represented by the predicate BestOSPFRoute) are defined through additional Datalog rules that capture the behavior of the OSPF protocol².

Network Topology. The network topology is also captured via Datalog rules in the program N . We model each router as a constant and use predicates to represent the topology. For example, the predicate $\text{SetLink}(R1, R2)$ represents that two routers $R1$ and $R2$ are connected via a link, and we add the Datalog rule $\text{SetLink}(R1, R2) \leftarrow \text{true}$ to define such a link.

4.2 Specifying Requirements

We specify the requirements as function-free first-order constraints over the predicate $\text{Fwd}(\text{TC}, \text{Router}, \text{NextHop})$, which defines the network’s forwarding state. We write $A \models \varphi$ to denote that a Datalog interpretation A satisfies φ . For illustration, we describe how common routing requirements can be specified:

- $\text{Path}(\text{TC}, R1, [R1, R2, \dots, Rn])$ (Path requirement): packets for traffic class TC must follow the path $R1, \dots, Rn$. These requirements are specified as a conjunction over the predicate Fwd .
- $\forall R1, R2. \text{Fwd}(\text{TC1}, R1, R2) \Rightarrow \neg \text{Fwd}(\text{TC2}, R1, R2)$ (Traffic isolation): the paths for two distinct traffic classes TC1 and TC2 do not share links in the same direction.
- $\text{Reach}(\text{TC}, R1, R2)$ (Reachability): packets for traffic class TC can reach router $R2$ from router $R1$. The predicate Reach is the transitive closure over the predicate Fwd (defined via Datalog rules).
- $\forall \text{TC}, R. (\neg \text{Reach}(\text{TC}, R, R))$ (Loop-freeness): generic requirement stipulating that the forwarding plane has no loops.

More complex requirements, such as way pointing, can be specified based on the core function-free first-order constraints provided by SyNET. Further, SyNET can be used as a backend for a high-level requirements language that is easier to use by a network operator.

² A detailed OSPF model can be found in the technical report [26].

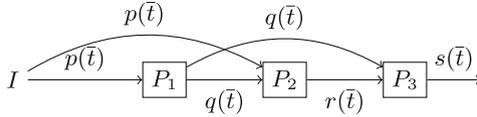


Fig. 3. A Datalog program with strata P_1 , P_2 , and P_3 , and flow of predicates between the strata.

4.3 Network-Wide Configurations

The input protocol configurations deployed at the network’s routers are represented as input *edb* predicates to the Datalog programs that formalize the protocols. For example, the local OSPF configuration for a router specifies the weights associated with the links connected to that router; this is represented by the *edb* predicate `SetOSPFEdgeCost(Router, NextHop, Weight)`.

A subset of the synthesized Datalog input for our motivating example is given in Fig. 1(d). Here, `SetAD` defines the administrative cost of static routes to be lower than that of OSPF (so static routes are preferred over forwarding entries computed by OSPF). The predicate `SetStatic(N1, A, B)`, which represents static routes, defines a static route for `N1` from `A` to `B`. The predicate `SetOSPFEdgeCost` defines the links’ weights.

5 Input Synthesis for Stratified Datalog

We now present a new iterative algorithm for synthesizing inputs for stratified Datalog. We first describe the high-level flow of the algorithm before presenting the details.

High-Level Flow. Consider the stratified Datalog program with strata P_1 , P_2 , and P_3 , depicted in Fig. 3. Incoming and outgoing edges of a stratum P_i indicate the *edb* predicates and, respectively, the *idb* predicates of that stratum. For example, the stratum P_3 takes as input predicates $q(\bar{t})$ and $r(\bar{t})$ and derives the predicate $s(\bar{t})$. Our iterative algorithm first synthesizes an input I_3 for P_3 which determines the predicates $q(\bar{t})$ and $r(\bar{t})$ that $P_1 \cup P_2$ must output. To synthesize such an input for a single stratum, we present an algorithm, called $\mathcal{S}_{SemiPos}$, that addresses the input synthesis problem for semi-positive Datalog programs [27, Chap. 15.2], i.e. Datalog programs where negation is restricted to *edb* predicates. After synthesizing an input I_3 for P_3 , our iterative algorithm synthesizes an input I_2 for P_2 such that the fixed-point $\llbracket P_2 \rrbracket_{I_2}$ produces the predicates $r(\bar{t})$ that are contained in the already synthesized input I_3 for P_3 . We note that this iterative process may require backtracking, in case no input for P_2 can produce the desired predicates $r(\bar{t})$ contained in I_3 . The algorithm terminates when it synthesizes inputs for all three strata.

In the following, we first present the algorithm $\mathcal{S}_{SemiPos}$ that is used to synthesize an input for a single stratum (which is a semi-positive program). Then, we present the general algorithm, called \mathcal{S}_{Strat} , that iteratively applies $\mathcal{S}_{SemiPos}$ for each stratum to synthesize inputs for stratified Datalog programs.

5.1 Input Synthesis for Semi-positive Datalog with SMT

The key idea is to reduce the input synthesis problem to satisfiability of SMT constraints: Given a semi-positive Datalog program P and a constraint φ , we encode the question $\exists I. \llbracket P \rrbracket_I \models \varphi$ into an SMT constraint ψ . If ψ is satisfiable, then from a model of ψ we can derive an input I such that $\llbracket P \rrbracket_I \models \varphi$.

SMT Encoding Challenges. Given a Datalog program P and a constraint φ , encoding the question $\exists I. \llbracket P \rrbracket_I \models \varphi$ with SMT constraints is non-trivial due to the mismatch between Datalog’s program fixed point semantics and the classical semantics of first-order logic. This means that simply taking the conjunction of all Datalog rules into an SMT solver does not solve our problem. For example, consider the following Datalog program P_{tc} :

$$\begin{aligned} tc(X, Y) &\leftarrow e(X, Y) \\ tc(X, Y) &\leftarrow tc(X, Z), tc(Z, Y) \end{aligned}$$

which computes the transitive closure of the predicate $e(X, Y)$. A naive way of encoding these Datalog rules with SMT constraints:

$$\begin{aligned} \forall X, Y. (e(X, Y) \Rightarrow tc(X, Y)) \\ \forall X, Y. ((\exists Z. tc(X, Z) \wedge tc(Z, Y)) \Rightarrow tc(X, Y)) \end{aligned}$$

and we denote the conjunction of these two SMT constraints as $\llbracket P_{tc} \rrbracket$. Now, suppose we have the fixed point constraint $\varphi_{tc} = (\neg e(v_0, v_2)) \wedge tc(v_0, v_2)$ and we want to generate an input I so that $\llbracket P_{tc} \rrbracket_I \models \varphi_{tc}$. A model that satisfies $\llbracket P_{tc} \rrbracket \wedge \varphi_{tc}$ is

$$\mathcal{M} = \{e(v_0, v_1), tc(v_0, v_1), tc(v_0, v_2)\}$$

The input derived from this model, obtained by projecting \mathcal{M} over the *edb* predicate e , is $I_{\mathcal{M}} = \{e(v_0, v_1)\}$. We get

$$\llbracket P_{tc} \rrbracket_{I_{\mathcal{M}}} = \{e(v_0, v_1), tc(v_0, v_1)\}$$

and so $\llbracket P_{tc} \rrbracket_{I_{\mathcal{M}}} \not\models \varphi_{tc}$, which is clearly not what is intended.

SMT Encoding. Our key insight is to split the constraint φ into a conjunction of positive and negative clauses, where a clause φ is positive (resp., negative) if $A \models \varphi$ implies that $A' \models \varphi$ for any interpretation $A' \supseteq A$ (resp., $A' \subseteq A$). We can then unroll recursive predicates to obtain a sound encoding for positive constraints, and we do not unroll them to get a sound encoding for negative constraints.

The encoding of a Datalog program P into an SMT constraint is defined in Fig. 4. The resulting SMT constraint is denoted by $\llbracket P \rrbracket_k$, where the parameter k defines the number of unroll steps. In the encoding we assume that (i) all terms in rules’ heads are variables and (ii) rules’ heads with the same predicate have identical variable names. Note that any Datalog program can be converted into this form using rectification [28] and variable renaming.

$$\begin{aligned}
 [P]_k &= \bigwedge_{p \in \text{idb}(P)} \text{ENCODE}(P, p) \wedge \text{UNROLL}(P, p, k) \\
 \text{ENCODE}(P, p) &= \bigwedge_{p(\bar{X}) \leftarrow \bar{l} \in P} \forall \bar{X}. ((\exists \bar{Y}. \bigwedge \bar{l}. \Rightarrow p(\bar{X})), \text{ where } \bar{Y} = \text{vars}(\bar{l}) \setminus \bar{X}) \\
 \text{UNROLL}(P, p, k) &= \bigwedge_{0 < i \leq k} \text{STEP}(P, p, i) \\
 \text{STEP}(P, p, i) &= \forall \bar{X}. (p_i(\bar{X}) \Leftrightarrow (\bigvee_{p(\bar{X}) \leftarrow \bar{l} \in P} \exists \bar{Y}. \tau(\bar{l}, i - 1))), \text{ where } \bar{Y} = \text{vars}(\bar{l}) \setminus \bar{X} \\
 \tau(\bar{l}, k) &= \begin{cases} \tau(l_1, k) \wedge \dots \wedge \tau(l_n, k) & \text{if } \bar{l} = l_1 \wedge \dots \wedge l_n \\ \neg \tau(p(\bar{t}), k) & \text{if } \bar{l} = \neg p(\bar{t}) \\ \text{false} & \text{if } \bar{l} = p(\bar{t}), p \in \text{idb}(P), k = 0 \\ p_k(\bar{t}) & \text{if } \bar{l} = p(\bar{t}), p \in \text{idb}(P), k > 0 \\ l & \text{otherwise} \end{cases}
 \end{aligned}$$

Fig. 4. Encoding a Datalog program P with constraints $[P]_k$

Function Encode. The constraint returned by $\text{ENCODE}(p, P)$ states that an atom $p(X)$ is derived if P has a rule that derives $p(\bar{X})$ and whose body evaluates to true. To capture Datalog’s semantics, the variables in $p(\bar{X})$ are universally quantified, while those in the rules’ bodies are existentially quantified. This constraint $\text{ENCODE}(p, P)$ is sound for negative requirements, but not for positive ones as it does not state that $p(\bar{X})$ is derived *only if* a rule body with $p(\bar{X})$ in the head evaluates to true.

Functions Unroll and Step. The constraint returned by $\text{STEP}(P, p, i)$ encodes whether an atom $p(X)$ is derived after i applications of P ’s rules; e.g., $p(X)$ ’s truth value after 3 steps is represented with the atom $p_3(X)$. Intuitively, $p(X)$ is true iff there is a rule that derives $p(X)$ and whose body evaluates to true using the atoms derived in previous iterations. Which atoms are derived in previous iterations is captured by the literal renaming function τ . Note that $\tau(l, 0)$ returns **false** for any *idb* literal l since all intensional predicates are initially **false**. Further, $\tau(l, k)$ returns l for any extensional literal l (the last case in Fig. 4) since their truth values do not change. Finally, the constraint returned by $\text{UNROLL}(P, p, k)$ conjoins $\text{STEP}(P, p, 0), \dots, \text{STEP}(P, p, k)$ to capture the derivation of $p(X)$ after k steps. This is sound for positive requirements, but not for negative ones since more $p(X)$ atoms may be derived after k steps.

Example. To illustrate the encoding, we translate the Datalog program:

$$\begin{aligned}
 tc(X, Y) &\leftarrow e(X, Y) \\
 tc(X, Y) &\leftarrow tc(X, Z), tc(Z, Y)
 \end{aligned}$$

which computes the transitive closure of the predicate $e(X, Y)$. This program has one *idb* predicate, tc . The function $\text{ENCODE}(P, tc)$ returns

$$\begin{aligned}
 &(\forall X, Y. e(X, Y) \Rightarrow tc(X, Y)) \\
 &\wedge (\forall X, Y. (\exists Z. tc(X, Z) \wedge tc(Z, Y)) \Rightarrow tc(X, Y))
 \end{aligned}$$

Algorithm 1. Algorithm $\mathcal{S}_{SemiPos}$ for semi-positive Datalog**Input:** Semi-positive Datalog program P and a constraint φ **Output:** An input I such that $\llbracket P \rrbracket_I \models \varphi$ or \perp

```

1 begin
2    $\varphi' \leftarrow \text{SIMPLIFY}(\varphi)$ 
3   for  $k \in [1..bound_k]$  do
4      $\varphi_k \leftarrow \text{REWRITE}(\varphi', k)$ 
5      $\psi \leftarrow [P]_k \wedge \varphi_k$ 
6     if  $\exists J. J \models \psi$  then
7        $I \leftarrow \{p(\bar{t}) \in J \mid p \in \text{edb}(P)\}$ , where  $J \models \psi$ 
8       return  $I$ 
9   return  $\perp$ 

```

We apply function $\text{UNROLL}(P, tc, 2)$ for $k = 2$, which after simplifications returns

$$\begin{aligned} &\forall X, Y. (tc_1(X, Y) \Leftrightarrow e(X, Y)) \\ &\forall X, Y. (tc_2(X, Y) \Leftrightarrow e(X, Y) \vee (\exists Z. tc_1(X, Z) \wedge tc_1(Z, Y))) \end{aligned}$$

In the constraints, the predicates tc_1 and tc_2 encode the derived predicates tc after 1 and, respectively, 2, derivation steps.

Algorithm. Algorithm $\mathcal{S}_{SemiPos}(P, \varphi)$, given in Algorithm 1, first calls function $\text{SIMPLIFY}(\varphi)$ that (i) instantiates any quantifiers in φ and (ii) transforms the result into a conjunction of clauses, where each clause is a disjunction of literals.

Then, the algorithm iteratively unrolls the Datalog rules, up to a pre-defined bound, called $bound_k$. In each step of the for-loop, the algorithm generates an SMT constraint that captures (i) which atoms are derived after k applications of P 's rules and (ii) which atoms are never derived by P . The resulting SMT constraint is denoted by $[P]_k$. The algorithm also rewrites the simplified constraint φ' using the function $\text{REWRITE}(\varphi', k)$ which recursively traverses conjunctions and disjunctions in the simplified constraint φ' and maps positive literals to the k -unrolled predicate $p_k(\bar{t})$ and negative literals to $\neg p(\bar{t})$:

$$\text{REWRITE}(\varphi, k) = \begin{cases} p_k(\bar{t}) & \text{if } \varphi = p(\bar{t}) \\ \neg p(\bar{t}) & \text{if } \varphi = \neg p(\bar{t}) \\ \text{REWRITE}(\varphi_1, k) \vee \dots \vee \text{REWRITE}(\varphi_n, k) & \text{if } \varphi = \varphi_1 \vee \dots \vee \varphi_n \\ \text{REWRITE}(\varphi_1, k) \wedge \dots \wedge \text{REWRITE}(\varphi_n, k) & \text{if } \varphi = \varphi_1 \wedge \dots \wedge \varphi_n \end{cases}$$

Note that since \vee and \wedge are monotone, negative literals constitute negative constraints and positive literals constitute positive constraints.

If the resulting constraint $[P]_k \wedge \psi_k$ is satisfiable, then an input is derived by projecting the interpretation I that satisfies the constraint over all edb predicates. Note that if there is an input I such that $\llbracket P \rrbracket_I \models \varphi$ and for which the fixed point $\llbracket P \rrbracket_I$ is reached in less than $bound_k$ steps, then $\mathcal{S}_{SemiPos}(P, \varphi)$ is guaranteed to return an input.

Algorithm 2. Input synthesis algorithm \mathcal{S}_{Strat} for stratified Datalog

Input: Stratified Datalog program $P = P_1 \cup \dots \cup P_n$, constraint φ over P_n
Output: An input I such that $\llbracket P \rrbracket_I \models \varphi$ or \perp

```

1 begin
2    $\mathcal{F}_1 \leftarrow \emptyset, \dots, \mathcal{F}_n \leftarrow \emptyset; I_1 \leftarrow \perp, \dots, I_n \leftarrow \perp; i \leftarrow n$ 
3   while  $i > 0$  do
4     if  $|\mathcal{F}_i| > bound_{\mathcal{F}}$  then
5        $\mathcal{F}_i \leftarrow \emptyset; \mathcal{F}_{i+1} \leftarrow \mathcal{F}_{i+1} \cup \{I_{i+1}\}$ 
6        $i \leftarrow i + 1; \quad // \text{ backtrack to higher stratum}$ 
7       continue
8      $\psi_{\mathcal{F}} \leftarrow \bigwedge_{I' \in \mathcal{F}_i} (\neg \bigwedge_{p \in edb(P_i)} ENCODEPRED(I', p))$ 
9     if  $i = n$  then
10       $\psi_i \leftarrow \varphi$ 
11     else
12       $\psi_i \leftarrow \bigwedge_{p \in \Delta_i} ENCODEPRED(I_{i+1} \cup \dots \cup I_n, p)$ 
13      where  $\Delta_i = (edb(P_i) \cup idb(P_i)) \cap (edb(P_{i+1}) \cup \dots \cup edb(P_n))$ 
14       $I_i = \mathcal{S}_{SemiPos}(P_i, \psi_i \wedge \psi_{\mathcal{F}})$ 
15      if  $I_i \neq \perp$  then
16         $i \leftarrow i - 1$ 
17      else
18        if  $i < n$  then
19           $\mathcal{F}_i \leftarrow \emptyset; \mathcal{F}_{i+1} \leftarrow \mathcal{F}_{i+1} \cup \{I_{i+1}\}$ 
20           $i \leftarrow i + 1 \quad // \text{ backtrack to higher stratum}$ 
21        else
22          return  $\perp$ 
23  return  $I = \{p(\bar{t}) \in I_1 \cup \dots \cup I_n \mid p \in edb(P)\}$ 

```

Theorem 1. Let P be a semi-positive Datalog program, φ a constraint. If $\mathcal{S}_{SemiPos}(P, \varphi) = I$ then $\llbracket P \rrbracket_I \models \varphi$.³

5.2 Iterative Input Synthesis for Stratified Datalog

Our iterative input synthesis algorithm for stratified Datalog, called \mathcal{S}_{Strat} , is given in Algorithm 2. We assume that the fixed point constraint φ is defined over predicates that appear in the highest stratum P_n ; this is without any loss of generality, as any constraint can be expressed using Datalog rules in the highest stratum, using a standard reduction to query satisfiability; cf. [24]. Starting with the highest stratum P_n , \mathcal{S}_{Strat} generates an input I_n for P_n such that $\llbracket P_n \rrbracket_{I_n} \models \varphi$. Then, it iteratively synthesizes an input for the lower strata P_{n-1}, \dots, P_1 using

³ The theorem's proof can be found in the technical report [26].

the algorithm $\mathcal{S}_{SemiPos}$. Finally, to construct an input for P , the algorithm combines the inputs synthesized for all strata and returns this.

Recall that the fixed point of a stratum P_i is given as input to the higher strata P_{i+1}, \dots, P_n . A key step when synthesizing an input I_i for P_i is thus to ensure that the *idb* predicates derived by P_i are identical to the *edb* predicates synthesized for the inputs I_{i+1}, \dots, I_n of the higher strata. Formally, let

$$\Delta_i = (edb(P_i) \cup idb(P_i)) \cap (edb(P_{i+1}) \cup \dots \cup edb(P_n))$$

We must ensure that $\{p(\bar{t}) \in \llbracket P_i \rrbracket_{I_i} \mid p \in \Delta_i\} = \{p(\bar{t}) \in I_{i+1} \cup \dots \cup I_n \mid p \in \Delta_i\}$.

Key Steps. The algorithm first partitions P into strata P_1, \dots, P_n . The strata can be computed using the predicates' dependency graph; see [27, Chap. 15.2]. For each stratum P_i , it maintains a set of inputs \mathcal{F}_i , which contains inputs for P_i for which the algorithm failed to synthesize inputs for the lower strata P_1, \dots, P_{i-1} . We call the sets \mathcal{F}_i *failed* inputs. All \mathcal{F}_i are initially empty.

In each iteration of the while loop, the algorithm attempts to generate an input I_i for stratum P_i . At line 4, the algorithm checks whether \mathcal{F}_i has exceeded a pre-defined bound $bound_{\mathcal{F}}$. If the bound is exceeded, it adds I_{i+1} to the failed inputs \mathcal{F}_{i+1} , re-initializes \mathcal{F}_i to the empty set, and backtracks to a higher stratum by incrementing i . This avoids exhaustively searching through all inputs to find an input compatible with those synthesized for the higher strata.

At line 8, the algorithm uses the helper function $\text{ENCODEPRED}(I', p)$. This function returns the constraint $\forall \bar{X}. (\bigvee_{p(\bar{t}) \in I'} \bar{X} = \bar{t}) \Leftrightarrow p(\bar{X})$, which is satisfied by an interpretation I iff I contains identical $p(\bar{t})$ predicates as those in I' . That is, if $I \models \text{ENCODEPRED}(I', p)$ then for any $p(\bar{t})$ we have $p(\bar{t}) \in I$ iff $p(\bar{t}) \in I'$. Therefore, the constraint $\psi_{\mathcal{F}}$ constructed at line 8 is satisfied by an input I_i iff $I_i \notin \mathcal{F}_i$, which avoids synthesizing inputs from the set of failed inputs.

The constraint ψ_i in the algorithm constrains the fixed point of P_i . For the highest stratum P_n , ψ_i is set to the constraint φ given as input to the algorithm. For the remaining strata P_i , ψ_i is satisfied iff the fixed point of P_i is compatible with the synthesized inputs for the higher strata P_{i+1}, \dots, P_n . In addition to constraining P_i 's *idb* predicates, we also constrain the input *edb* predicates. This is necessary to eagerly constrain the inputs.

At line 14, the algorithm invokes $\mathcal{S}_{SemiPos}$ to generate an input I_i such that $\llbracket P_i \rrbracket_{I_i} \models \varphi_i \wedge \psi_{\mathcal{F}}$. The algorithm proceeds to the lower stratum if such an input is found ($I \neq \perp$); otherwise, if $i < n$ the algorithm backtracks to the higher stratum by increasing i and updating the sets \mathcal{F}_{i+1} , and if $i = n$ it returns \perp .

Finally, the while-loop terminates when the inputs of all strata have been generated. The algorithm constructs and returns the input I for P .

Theorem 2. *Let P be a stratified Datalog program with strata P_1, \dots, P_n , and φ a constraint over predicates in P_n . If $\mathcal{S}_{Strat}(P, \varphi) = I$ then $\llbracket P \rrbracket_I \models \varphi$.⁴*

⁴ The theorem's proof can be found in the technical report [26].

6 Implementation and Evaluation

In this section we first describe SyNET, and end-to-end implementation of our input synthesis algorithm applied to the network-wide synthesis problem. We then turn to our evaluation of SyNET on practical topologies and requirements.

6.1 Implementation

SyNET is implemented in Python and automatically encodes stratified Datalog programs specified in the LogicBlox language [29] into SMT constraints specified in the SMT-LIB v2 format [30]. It uses the Python API of Z3 [31] to check whether the generated SMT constraints are satisfiable and to obtain a model.

SyNET supports routers that run both, OSPF and BGP protocols, and that can be configured with static routes. SyNET uses natural splitting for protocols: external routes are handled by BGP, while internal routes are handled by IGP protocols (OSPF and static, where static routes are preferred over OSPF). We have partitioned the Datalog rules that capture these protocols and their dependencies into 8 strata. SyNET relies on additional SMT constraints to ensure the well-formedness of the OSPF, BGP, and static route configurations output by our synthesizer. For most topologies and requirements, the Datalog program reaches a fixed point within 20 iterations, and so we fixed the unroll and backtracking bounds ($bound_k$ and $bound_{\mathcal{F}}$) to 20.

SyNET is vendor agnostic with respect to the synthesized configurations. A simple script can be used to convert the output of SyNET into any vendor specific configuration format and then deploy them in production routers. Indeed, to test the correctness of SyNET, we implemented a small script to convert the input synthesized by SyNET to Cisco router configurations.

SyNET supports two key optimizations that improve its performance. The first optimization is *partial evaluation*: SyNET partially-evaluates Datalog rules with predicates whose truth values are known apriori. For example, all `SetLink` predicates are known and can be eliminated. This reduces the number of variables in the rules and, in turn, in the generated SMT constraints. The second optimization is *network-specific constraints*: we have configured SyNET with generic constraints, which are true for all forwarding states, and with protocol-specific constraints, i.e. constraints that hold for any input to a particular protocol. An example constraint is: “No packet is forwarded out of the router if the destination network is directly connected to the router”. These constraints are not specific to particular requirements or topology. They are thus defined one time and can be used to synthesize configurations for any requirements and networks.

6.2 Experiments

To investigate SyNET’s performance and scalability, we experimented with different: (i) topologies, (ii) requirements; and (iii) protocol combinations. Further to

test correctness, we ran all synthesized configurations on an emulated environment of Cisco routers [32] and we verified that the forwarding paths computed match the requirements for each experiment.

Network Topologies. We used network topologies that have between 4 and 64 routers. The 4-router network is our overview example where we considered the same requirements as those described in Sect. 2. The 9-router network is Internet2 (see Fig. 5), a US-based network that connects several major universities and research institutes. The remaining networks are $n \times n$ grids.

Routing Requirements. For each router and each traffic class, we generate a routing requirement that defines where the packets for that traffic class must be forwarded to. We consider 1, 5, and 10 traffic classes. For a topology with n routers and m traffic classes, we thus generate $n \times m$ requirements.

For topologies with multiple traffic classes, we add one external network announced by two randomly selected routers. We add requirements to enforce that all packets destined to the external networks are forwarded to one of the two routers. This models a scenario where the operator is planning maintenance downtime for one of the two routers. Further, to show that SyNET synthesizes configurations with partially defined input and protocol dependencies, we assume the local BGP preferences are fixed by the network operator and thus SyNET has to synthesize correct OSPF costs to meet the BGP requirements.

Protocols. We consider three different combinations of protocols: (i) static routes; (ii) OSPF and static routes; and (iii) OSPF, BGP, and static routes. The protocol combinations (i) and (ii) ignore requirements for external networks since only BGP computes routes for them.

Experimental Setup. We run SyNET on a machine with 128 GB of RAM and a modern 12-core dual-processors running at 2.3 GHz.

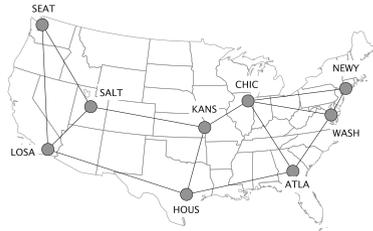


Fig. 5. Internet2 topology

Results. The synthesis times for the different networks and protocol combinations are shown in Table 1 (averaged over 10 runs). SyNET synthesizes the overview example’s configuration described in Sect. 2 in 10 s. For the largest network (64 routers) and number of traffic classes (10 classes), SyNET synthesizes a configuration for static routes (protocol combination (i)) in less than 1 h, and for the combination of static routes and OSPF, SyNET takes less than 22 h. When using both OSPF and BGP protocols along with static routes, for all network topologies SyNET synthesizes configurations for 1 and 5 traffic classes within 8 h; for 10 traffic classes, SyNET times out after 24 h for the largest topologies with 49 and 64 routers.

Interpretation. Our results show that SyNET scales to real-world networks. Indeed, a longitudinal analysis of more than 260 production networks [33] revealed that 56% of them have less than 32 routers. SyNET would synthesize

Table 1. SyNET’s synthesis times (averaged over 10 runs) for different number of routers, protocol combinations, and traffic classes in the requirements.

Protocol	# Routers	1 Traffic class		5 Traffic classes		10 Traffic classes	
		Avg	Std	Avg	Std	Avg	Std
Static	9	1.3 s	(0.5)	2.0 s	(0.1)	2.8 s	(0.4)
	9 (Internet2)	1.3 s	(0.5)	2.0 s	(0.0)	4.0 s	(0.8)
	16	5.9 s	(0.3)	7.8 s	(0.4)	11.2 s	(0.4)
	25	32.0 s	(0.6)	37.0 s	(0.6)	46.1 s	(0.9)
	36	2 m 49.7 s	(3.0)	3 m 1.5 s	(4.5)	3 m 27.0 s	(4.4)
	49	12 m 29.2 s	(7.0)	13 m 02.3 s	(10.6)	14 m 10.7 s	(15.0)
	64	46 m 36.2 s	(49.0)	47 m 23.8 s	(27.2)	49 m 22.2 s	(39.3)
OSPF+Static	9	9.4 s	(0.5)	19.8 s	(0.4)	39.9 s	(0.5)
	9 (Internet2)	9.0 s	(1.4)	21.3 s	(1.2)	49.3 s	(0.5)
	16	43.5 s	(0.7)	1 m 19.8 s	(0.6)	4 m 5.8 s	(1.6)
	25	2 m 55.2 s	(6.1)	7 m 3.8 s	(9.9)	15 m 56.4 s	(38.1)
	36	10 m 00.5 s	(9.5)	23 m 58.9 s	(22.5)	1 h 11 m 38.2 s	(127.5)
	49	24 m 11.6 s	(43.5)	1 h 30 m 00.3 s	(89.6)	5 h 22 m 55.8 s	(421.2)
	64	2 h 22 m 13.2 s	(209.9)	5 h 42 m 58.9 s	(619.4)	21 h 13 m 16.0 s	(1986.7)
BGP+OSPF+Static	9	15.3 s	(0.5)	27.7 s	(0.5)	1 m 0.5 s	(2.6)
	9 (Internet2)	13.3 s	(0.9)	22.7 s	(0.9)	1 m 19.7 s	(0.5)
	16	56.0 s	(1.6)	2 m 24.7 s	(0.9)	8 m 29.0 s	(10.7)
	25	3 m 56.3 s	(3.1)	8 m 46.3 s	(5.3)	40 m 09.3 s	(99.2)
	36	14 m 14.0 s	(15.0)	43 m 38.0 s	(5.7)	2 h 35 m 11.7 s	(197.7)
	49	1 h 23 m 20.7 s	(211.1)	2 h 15 m 18.0 s	(12.8)	timeout (>24h)	
	64	1 h 46 m 35.0 s	(165.8)	7 h 24 m 51.3 s	(519.2)	timeout (>24h)	

configurations for such networks within one hour. SyNET also already supports a reasonable amount of traffic classes. According to a study on real-world enterprise and WAN networks [21], even large networks with 100,000s of IP prefixes in their forwarding tables usually see less than 15 traffic classes in total.

While SyNET can take more than 24 h to synthesize a configuration for the largest networks (with all protocols activated and 10 traffic classes), we believe that this time can be reduced through divide-and-conquer. Real networks tend to be hierarchically organized around few regions (to ensure the scalability of the protocols [34]) whose configurations can be synthesized independently. We plan to explore the synthesis of such hierarchical configurations in future work.

7 Related Work

Analysis of Datalog Programs. Datalog has been successfully used to declaratively specify variety of static analyzers [35, 36]. It has been also used to verify network-wide configurations for protocols such as OSPF and BGP [4]. Recent work [37] has extended Datalog to operate with richer classes of lattice structures. Further, the μZ tool [38] extends the Z3 SMT solver with support for fixed points. The focus of all these works is on computing the fixed point of a program P for a given input I and then checking a property φ on the fixed point. That is, they check whether $\llbracket P \rrbracket_I \models \varphi$. All of these works assume that the input is provided a priori. In contrast, our procedure discovers an input that produces a fixed point satisfying a given (user-provided) property on the fixed point.

The algorithm presented in [36] can be used to check whether certain tuples are not derived for a given set of inputs. Given a Datalog program P (without

negation in the literals), a set Q of tuples, and a set \mathcal{I} of inputs, the algorithm computes the set $Q \setminus \bigcap \{ \llbracket P \rrbracket_I \mid I \in \mathcal{I} \}$. This algorithm cannot address our problem because it does not support stratified Datalog programs, which are not monotone. While their encoding can be used to synthesize inputs for each stratum of a stratified Datalog program, it supports only negative properties, which require that certain tuples are not derived. Our approach is thus more general than [36] and can be used in their application domain.

The FORMULA system [39,40] can synthesize inputs for non-recursive Datalog programs, as it supports non-recursive Horn clauses with stratified negation (even though [41] which uses FORMULA shows examples of recursive Horn clauses w/o negation). Handling recursion with stratified negation is nontrivial as bounded unrolling is unsound if applied to all strata together. Note that virtually all network specifications require recursive rules, which our system supports.

Symbolic Analysis and Synthesis. Our algorithm is similar in spirit to symbolic (or concolic) execution, which is used to automatically generate inputs for programs that violate a given assertion (e.g. division by zero); see [42–44] for an overview. These approaches unroll loops up to a bound and find inputs by calling an SMT solver on the symbolic path. While we also find inputs for a symbolic formula, the entire setting, techniques and algorithms, are all different from the standard symbolic execution setting.

Counter-example guided synthesis approaches are also related [45]. Typically, the goal of synthesis is to discover a program, while in our case the program is given and we synthesize an input for it. There is a connection, however, as a program can be represented as a vector of bits. Most such approaches have a single counter-example generator (i.e., the oracle), while we use a sequence of oracles. It would be interesting to investigate domains where such layered oracle counter-example generation can benefit and improve the efficiency of synthesis.

Network Configuration Synthesis. Propane [46] and Genesis [19] also produce network-wide configurations out of routing requirements. Unlike our approach, however, Propane only supports BGP and Genesis only supports static routes. In contrast to our system, Propane and Genesis support failure-resilience requirements. While we could directly capture such requirements by quantifying over links, this would make synthesis more expensive. A more efficient way to handle such requirements would be to synthesize a failure-resilient forwarding plane using a system like Genesis [19], and to then feed this as input to our synthesizer to get a network-wide configuration. In contrast to these approaches, our system is more general: one can directly extend it with additional routing protocols, by specifying them in stratified Datalog, and synthesize configurations for *any combination* of routing protocols.

ConfigAssure [47] is a general system that takes as input requirements in first-order constraints and outputs a configuration conforming to the requirements. The fixed point computation performed by routing protocols cannot be captured using the formalism used in ConfigAssure. Therefore, ConfigAssure cannot be used to specify networks and, in turn, to synthesize protocol configurations for networks.

8 Conclusion

We formulated the network-wide configuration synthesis problem as a problem of finding inputs of a Datalog program, and presented a new input synthesis algorithm to solve this challenge. Our algorithm is based on decomposing the Datalog rules into strata and iteratively synthesizing inputs for the individual strata using off-the-shelf SMT solvers. We implemented our approach in a system called SyNET and showed that it scales to realistic network size using any combination of OSPF, BGP and static routes.

References

1. Ryall, J.: Facebook, Tinder, Instagram suffer widespread issues. <http://mashable.com/2015/01/27/facebook-tinder-instagram-issues/>
2. Juniper Networks. What's Behind Network Downtime? Proactive Steps to Reduce Human Error and Improve Availability of Networks. Technical report, May 2008
3. BGPmon. Internet prefixes monitoring. <http://www.bgpmon.net/blog/>
4. Fogel, A., Fung, S., Pedrosa, L., Walraed-Sullivan, M., Govindan, R., Mahajan, R., Millstein, T.: A general approach to network configuration analysis. In: NSDI (2015)
5. Feamster, N., Balakrishnan, H.: Detecting BGP configuration faults with static analysis. In: NSDI (2005)
6. Nelson, T., Barratt, C., Dougherty, D.J., Fidler, K., Krishnamurthi, S.: The margrave tool for firewall analysis. In: LISA (2010)
7. Yuan, L., Chen, H., Mai, J., Chuah, C.-N., Su, Z., Mohapatra, P.: FIREMAN: a toolkit for firewall modeling and analysis. In: S&P (2006)
8. Vanbever, L., Quoitin, B., Bonaventure, O.: A hierarchical model for BGP routing policies. In: ACM SIGCOMM PRESTO (2009)
9. Chen, X., Mao, M., Van der Merwe, J.: PACMAN: a platform for automated and controlled network operations and configuration management. In: CoNEXT (2009)
10. Enck, W., Moyer, T., McDaniel, P., Sen, S., Sebos, P., Spoerel, S., Greenberg, A., Sung, Y.-W.E., Rao, S., Aiello, W.: Configuration management at massive scale: system design and experience. *IEEE J. Sel. Areas Commun.* (2009)
11. Gottlieb, J., Greenberg, A., Rexford, J., Wang, J.: Automated provisioning of BGP customers. *IEEE Netw.* **17**, 44–55 (2003)
12. Alaettinoglu, C., Villamizar, C., Gerich, E., Kessens, D., Meyer, D., Bates, T., Karrenberg, D., Terpstra, M.: Routing Policy Specification Language. RFC 2622
13. Bjorklund, M.: YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF). RFC 6020
14. Enns, R., et al.: Network Configuration Protocol (NETCONF). RFC 4741
15. El-Hassany, A., Miserez, J., Bielik, P., Vanbever, L., Vechev, M.: SDNRacer: concurrency analysis for SDNs. In: PLDI (2016)
16. Canini, M., Venzano, D., Peresini, P., Kostic, D., Rexford, J., et al.: A NICE way to test OpenFlow applications. In: NSDI (2012)
17. Scott, C., Wundsam, A., Raghavan, B., Panda, A., Or, A., Lai, J., Huang, E., Liu, Z., El-Hassany, A., Whitlock, S., Acharya, H.B., Zarifis, K., Shenker, S.: Troubleshooting Blackbox SDN control software with minimal causal sequences. In: ACM SIGCOMM (2014)

18. Ball, T., Bjørner, N., Gember, A., Itzhaky, S., Karbyshev, A., Sagiv, M., Schapira, M., Valadarsky, A.: VeriCon: towards verifying controller programs in software-defined networks. In: PLDI (2014)
19. Subramanian, K., D’Antoni, L., Akella, A.: Genesis: synthesizing forwarding tables in multi-tenant networks. In: POPL (2017)
20. Kang, N., Liu, Z., Rexford, J., Walker, D.: Optimizing the “One Big Switch” abstraction in software-defined networks. In: CoNEXT (2013)
21. Benson, T., Akella, A., Maltz, D.A.: Mining policies from enterprise network configuration. In: IMC (2009)
22. Awduche, D., et al.: Overview and Principles of Internet Traffic Engineering. RFC3272
23. Fortz, B., Rexford, J., Thorup, M.: Traffic engineering with traditional IP routing protocols. *IEEE Commun. Mag.* **40**, 118–124 (2002)
24. Halevy, A.Y., Mumick, I.S., Sagiv, Y., Shmueli, O.: Static analysis in datalog extensions. *J. ACM* **48**, 971–1012 (2001)
25. Mumick, I.S., Shmueli, O.: How expressive is stratified aggregation? *Ann. Math. Artif. Intell.* **15**, 407–435 (1995)
26. El-Hassany, A., Tsankov, P., Vanbever, L., Vechev, M.T.: Network-wide configuration synthesis. CoRR, abs/1611.02537 (2016). <http://arxiv.org/abs/1611.02537>
27. Abiteboul, S., Hull, R., Vianu, V. (eds.): *Foundations of Databases: The Logical Level* (1995)
28. Ullman, J.D.: *Principles of Database and Knowledge-Base Systems*. Computer Science Press, New York (1989)
29. <https://logicblox.com/content/docs4/coreference/html/index.html>
30. Barrett, C., et al.: *The SMT-LIB Standard: Version 2.0* (2010)
31. De Moura, L., Bjørner, N.: Z3: an efficient SMT solver. In: TACAS (2008)
32. Graphical Network Simulator-3 (GNS3). <https://www.gns3.com/>
33. Knight, S., Nguyen, H.X., Falkner, N., Bowden, R.A., Roughan, M.: The internet topology zoo. *IEEE J. Sel. Areas Commun.* **29**, 1765–1775 (2011)
34. Doyle, J., Carroll, J.: *Routing TCP/IP, vol. 1*. Cisco Press, Indianapolis (2005)
35. Smaragdakis, Y., Bravenboer, M.: Using datalog for fast and easy program analysis. In: de Moor, O., Gottlob, G., Furche, T., Sellers, A. (eds.) *Datalog Reloaded*. LNCS, vol. 6702, pp. 245–251. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-24206-9_14](https://doi.org/10.1007/978-3-642-24206-9_14)
36. Zhang, X., Mangal, R., Grigore, R., Naik, M., Yang, H.: On abstraction refinement for program analyses in datalog. In: PLDI (2014)
37. Madsen, M., Yee, M.-H., Lhoták, O.: From datalog to flix: a declarative language for fixed points on lattices. In: PLDI (2016)
38. Hoder, K., Bjørner, N., De Moura, L.: μZ : an efficient engine for fixed points with constraints. In: CAV (2011)
39. Jackson, E.K., Sztipanovits, J.: Towards a formal foundation for domain specific modeling languages. In: EMSOFT (2006)
40. Jackson, E.K., Schulte, W.: Model generation for horn logic with stratified negation. In: Suzuki, K., Higashino, T., Yasumoto, K., El-Fakih, K. (eds.) *FORTE 2008*. LNCS, vol. 5048, pp. 1–20. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-68855-6_1](https://doi.org/10.1007/978-3-540-68855-6_1)
41. Jackson, E.K., Kang, E., Dahlweid, M., Seifert, D., Santen, T.: Components, platforms and possibilities: towards generic automation for MDA. In: EMSOFT (2010)
42. Cadar, C., Sen, K.: Symbolic execution for software testing: three decades later. *Commun. ACM* **26**, 82–90 (2013)

43. Kroening, D., Tautschnig, M.: CBMC – C bounded model checker. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. TACAS 2014, vol. 8413, pp. 389–391. Springer, Heidelberg (2014). doi:[10.1007/978-3-642-54862-8_26](https://doi.org/10.1007/978-3-642-54862-8_26)
44. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 168–176. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-24730-2_15](https://doi.org/10.1007/978-3-540-24730-2_15)
45. Solar-Lezama, A., Tancau, L., Bodik, R., Seshia, S., Saraswat, V.: Combinatorial Sketching for Finite Programs. In: ASPLOS (2006)
46. Beckett, R., Mahajan, R., Millstein, T., Padhye, J., Walker, D.: Don't mind the gap: bridging network-wide objectives and device-level configurations. In: SIGCOMM (2016)
47. Narain, S., Levin, G., Malik, S., Kaul, V.: Declarative infrastructure configuration synthesis and debugging. *J. Netw. Syst. Manag.* **16**, 235–258 (2008)