Localizing Router Configuration Errors Using Minimal Correction Sets

Aaron Gember-Jacobson, Ruchit Shrestha, Xiaolin Sun
Colgate University
{agemberjacobson,rshrestha,xsun}@colgate.edu

ABSTRACT

Router configuration errors are unfortunately common and difficult to localize using current network verifiers. We introduce a novel configuration error localizer (CEL) that precisely identifies which configuration segments contribute to the violation of forwarding requirements. In particular, CEL generates a system of satisfiability modulo theories (SMT) constraints—which encode a network's configurations, control logic, and forwarding requirements—and uses a domain-specific minimal correction set (MCS) enumeration algorithm to identify problematic configuration segments. CEL efficiently locates several configuration errors in real university networks and identifies all routing-related and at least half of all ACL-related errors we introduce.

1 INTRODUCTION

Many networks use distributed routing protocols (e.g., BGP and OSPF) to compute forwarding rules. The computations are influenced by the network topology, link failure states, protocol-specific algorithms, and numerous inputs specified in router configurations—e.g., neighbor relationships, prefixes to advertise, link costs, route filters, etc. To ensure the network computes rules that satisfy forwarding requirements—e.g., reachability, waypoint traversals, path preferences—suitable inputs must be specified in router configurations.

Unfortunately, providing the correct inputs is complex, and configuration errors are common [31, 44]. This has motivated researchers to develop numerous systems to detect [2–4, 8, 13, 14, 16, 18, 19, 22–25, 28, 32, 34, 40, 42, 45], remediate [15, 38], and avoid [5, 10, 11, 37] configuration errors.

Error detection tools fall into four major categories. Configuration checkers [13, 22, 28] inspect configurations for deviations from common practice; however these tools ignore the network's forwarding requirements. Data plane verifiers [18, 23–25, 32] analyze forwarding rule snapshots to detect violations of forwarding requirements, but they ignore the network's configurations. Control plane verifiers [2–4, 14, 16, 19, 34, 40] derive from the configurations a model that encodes *all* possible forwarding rules, but only one control plane verifier (Batfish [14]) links problematic forwarding rules with the configuration statements that caused them, and

this verifier scales poorly [16]. Finally, provenance-based diagnosis tools [8, 42] provide explanations for the presence (or absence) of specific forwarding rules, but these explanations are based on observed events—e.g, route advertisements and link failures—rather than configurations. In summary, existing error detection tools are ill suited for identifying *which aspects of a network's configurations caused a forwarding requirement to be violated*. Thus, it is hard for engineers, or even automated repair tools [15], to correct the error(s).

Configuration synthesizers [5, 10, 11, 15, 37, 38] are designed to automatically create error-free configurations, supposedly avoiding the need to detect or remediate errors. However, existing synthesizers do not cover the full range of requirements networks may need to satisfy—e.g., preventing communication between specific subnets—or the range of features required to satisfy those requirements—e.g., packet filters. Consequently, portions of a network's configurations are often still written by-hand, introducing the possibility for errors to occur. Furthermore, to make synthesis tractable, some tools require engineers to provide configuration templates/hints, which may contain errors or be overly restrictive— e.g., we uncovered an error in the templates used by NetComplete [10] to syntheisze WAN configurations (§6.1).

In this paper, we introduce a novel *configuration error localizer* (CEL) that *precisely identifies which configuration segments contribute to the violation of one or more forwarding requirements*. CEL constructs a system of satisfiability modulo theories (SMT) constraints that encode a network's configurations, control plane logic, and forwarding requirements, and then checks if the constraints are satisfiable under all failure scenarios of interest. If the constraints are unsatisfiable under some failure scenario—indicating the forwarding requirements cannot always be satisfied by the current configurations—then CEL computes a *minimal correction set*—i.e., a subset of constraints whose removal from the problem would allow the remaining constraints to be satisfied. The constraints contained in the correction set are used to flag configuration statements that contribute to the violation of forwarding requirements.

Accurately identifying configuration errors using correction sets requires addressing several challenges: (1) Existing SMT-based network models [3, 10, 15] use constraints that

1

co-mingle configuration and control logic, which makes it difficult to separate configuration errors from software bugs. (2) Forwarding requirement violations may be caused by *missing* configuration statements (e.g., a missing filter rule), but most models only encode the configuration statements that are present; only a few models encode *plausible* configuration statements, based on manually-specified templates [10] or a limited set of protocol options [15]. (3) A forwarding requirement may be satisfied under some conditions (e.g., no link failures) and violated under others. (4) There may be multiple plausible explanations for a forwarding requirement violation. (5) Configurations may contain multiple errors whose impact varies across failure scenarios and forwarding requirements.

To overcome these challenges, CEL uses a carefully structured SMT formulation and domain-specific algorithms for exploring the space of correction sets. In particular, CEL separates the encoding of configuration from the encoding of control logic and explicitly encodes the absence of certain configuration statements—namely, those present elsewhere in the configurations. Then, CEL conducts a counterexample-guided exploration of the failure scenarios under which a forwarding requirement is violated and employs a domain-specific version of a state-of-the-art minimal correction set (MCS) enumeration algorithm [29] to efficiently compute multiple, configuration-focused MCSes. Finally, CEL ranks MCSes based on common network management practices [17, 26] to identify the most-likely configuration errors.

We implement CEL atop the Minesweeper network verifier [3], and we evaluate CEL using a combination of real and synthetic network configurations. Using CEL, we pinpoint several configuration errors in two real university networks as well as synthesized WAN configurations. Additionally, we show that CEL identifies all routing-related and at least half of all ACL-related errors we introduce, and achieves perfect precision for 40% of the scenarios. Finally, we show that CEL localizes errors in less than 15 seconds for half of the scenarios, and less than 10 minutes in the remaining scenarios, whereas network engineers took more than a hour in over half of cases they reported.

2 EXISTING TECHNIQUES

To better understand the gap in approaches for localizing configuration errors, we study the current techniques used by network engineers (§2.1) and review existing systems developed by researchers (§2.2).

2.1 Resources used by network engineers

To better understand how network engineers currently identify configuration errors, we surveyed 25 engineers.[1] We asked

(a) *Used*   **(b)** *Desired*
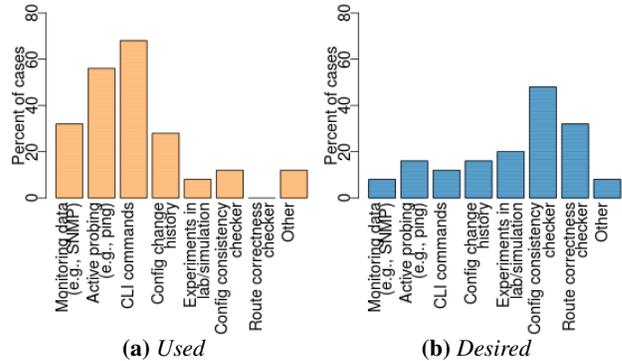
Figure 1: *Resources for identifying errors*

the engineers to consider a routing-related configuration error they recently corrected (excluding errors related to software bugs) and answer 9 questions about the error.

The engineers reported on configuration errors in enterprise campus (48%), service provider (32%), data center (16%), and Internet exchange point (4%) networks. The reported errors were approximately evenly divided between networks with less than 10 routers, 10 to 50 routers, and more than 50 routers. Almost half of the reported errors involved interface configurations. Engineers also reported on errors involving BGP, OSPF, IS-IS, MPLS, static routes, VPN, route filters, ACLs, and VLANs. Interestingly, one-third of the reported errors occurred in configuration segments that were automatically generated from templates or intents. This suggests that locating configuration errors is an important problem even with recent advances in configuration synthesis [5, 10, 11, 37].

Engineers reported using a wide range of resources to identify configuration errors (Figure 1a). Active probing (e.g., ping or traceroute) and router commands (e.g., `show route`) were the most commonly used resources, while tools that check for inconsistent/missing routes were not used by any engineers. Engineers used multiple resources in 72% of the cases. The most frequently used combination of resources was active probing and router commands (48% of cases).

Engineers were able to identify and correct errors reasonably quickly: almost half of the errors were identified in less than an hour, and 90% were identified in less than 12 hours.

Lastly, we asked what resources would have made it easier to identify the error (Figure 1b). Half of the engineers desired a tool that detects inconsistencies within/across device configurations, and one-third wanted a tool that checks for inconsistent/missing routes. Engineers also wanted tools that can detect missing route advertisements. Below, we discuss several existing tools that fit one of these categories.

2.2 Tools developed by researchers

Researchers have developed numerous tools for determining whether a network's configurations are correct and synthesizing correct (updates to) configurations. In this section, we
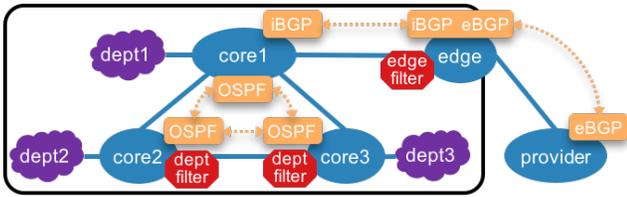
Figure 2: *Example campus network: blue ovals are routers; purple clouds are subnets; solid blue lines are links; orange rectangles are routing processes; dashed orange lines are routing adjacencies; red octagons are ACLs*

review five different categories of tools and illustrate the capabilities and limitations of the tools in each category using a small example network (Figure 2).

The example network uses the Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and access control lists (ACLs). edge receives routes from provider using eBGP and forwards routes to core1 using iBGP. core1 originates BGP routes for dept1/2/3. The BGP configuration on edge, provider, and core1 is bare-bones and does not contain any policies that filter or modify route advertisements. Routers core1, core2, and core3 run OSPF. core1 advertises a default route to core2 and core3, because external advertisements are not propagated to these two routers. The edgeFilter ACL defined on edge only allows traffic to core1 that is destined for dept1 or dept2. The deptFilter ACL defined on core3 (core2) only allows traffic to core2 (core3) that originated from dept2/3.

The network has two forwarding requirements: (FR1) all departments can access the Internet (in the absence of failures), and (FR2) all departments can communicate, even in the presence of a single link failure. However, these requirements are violated due to three different configuration errors:

CE1: edge's iBGP configuration is missing an option that causes a router to set itself as the next hop in routes forwarded to iBGP peers. Consequently, the routes edge forwards to core1 will contain provider as the next hop, which is not directly reachable from core1. This causes FR1 to be violated.

CE2: edgeFilter does not allow traffic destined for dept3 to be sent to core1. Consequently, external hosts cannot reach dept3, which violates FR1.

CE3: deptFilter does not allow traffic that originated from dept1. Thus, if the core1–core2 link fails, OSPF re-routes traffic from dept1 to dept2 through core3, and the traffic is dropped by core3, which violates FR2. A similar issue arises when the core1–core3 link fails.

Configuration checkers [13, 22, 28] statically analyze router configurations to detect anomalies. In particular, rcc [13] checks whether BGP configurations conform to common best

practice—e.g., private addresses are not advertised externally—while Minerals [28] and SelfStarter [22] check for inconsistencies between routers—e.g., a filter is defined on multiple routers but contains different rules on some routers.

Configuration checkers (specifically, rcc) can only detect CE1. No checker can detect CE2 or CE3. From the perspective of Minerals and ShapeShifter, the deptFilter ACLs on core2 and core3 contain the same rule and are applied in a consistent manner—namely, core2 filters incoming traffic from core3 and vice versa—so no problem exists with these ACLs. The edgeFilter ACL on edge is unique—it has a unique name, unique rules, and is applied in a unique context—so there is no point of comparison that can be used to check if this ACL contains errors. ACLs are not even considered by rcc, which is designed explicitly for checking BGP-related configuration.

Data plane verifiers [18, 23–25, 32] transform a snapshot of routers' forwarding information bases (FIBs) into algebraic functions [24], boolean constraints [32], or forwarding graphs [18, 23, 25, 45] and use algebraic analysis, satisfiability (SAT) solvers, or graph algorithms, respectively, to detect violations of engineer-specified forwarding requirements.

Data plane verifiers can immediately detect that FR1 is violated, because no entries for external networks will be installed in core1's FIB (due to CE1), and edge's FIB will contain a rule to discard traffic destined for dept3[2] (due to CE2). However, data plane verifiers will not detect a violation of FR2 (due to CE3) until the core1–core2 or core1–core3 link fails and the routers update their FIBs.

Although data plane verifiers can detect violations of the forwarding requirements, the verifiers do not identify which configuration segments led to the missing and erroneous FIB entries. FIBs are much easier to model than router configurations and control logic, so data plane verifiers only consider the former. Identifying the cause of problematic FIB entries is left as a manual exercise for network engineers. This diagnosis is further complicated by the fact that errors in FIBs may be the result of errors in configurations and/or bugs in the control plane software running on the routers. Determining which component to blame requires additional sleuthing.

Control plane verifiers [2, 3, 12, 14, 16, 34, 40] model the route advertisement and selection algorithms used by various protocols (e.g., BGP and OSPF) in order to compute routers' expected FIBs under various failure scenarios. The models—e.g., Datalog programs [14], weighted-directed graphs [2, 16], or systems of constraints [3, 40]—are based on router configurations, protocol standards, and vendor documentation. The models are used to check whether forwarding requirements are satisfied under all failure scenarios of interest.

---

[2]Routers store filtering rules separately from forwarding rules, but for simplicity we refer to them collectively as the FIB.

Control plane verifiers can detect all three configuration errors, and produce counterexamples with forwarding paths that demonstrate the forwarding requirement violations. However, akin to data plane verifiers, most control plane verifiers do not identify which segments of the configuration contributed to the problematic forwarding paths. Only one control plane verifier (Batfish [14]) is capable of identifying which configuration statements led to the problematic forwarding rules, and this tool scales poorly when exploring multiple failure scenarios of interest [16].

Provenance-based diagnosis tools [8, 42] focus on explaining *why* a FIB entry is present (or absent). However, these explanations are based on observed events—e.g., route advertisements and link failures—rather than routers' configurations. Consequently, the explanations are valuable for identifying control plane software bugs that lead to forwarding requirement violations, but the explanations do not help network engineers identify which configuration segments are responsible for forwarding requirement violations.

Configuration synthesizers [5, 10, 11, 15, 37, 38] are designed to automatically create error-free configurations, supposedly avoiding the need to localize errors. However, portions of configurations are often still written by-hand [17], introducing the possibility for errors to occur. Furthermore, some synthesizers require engineers to provide configuration templates, which may contain errors—e.g., we uncovered an error in the templates provided with NetComplete [10] (§6.1).

In summary, existing tools either *identify a limited set of configuration errors* or *fail to identify which configuration segments cause forwarding requirement violations*.

3 CORRECTION SETS

Given the aforementioned limitations of existing tools, we seek to develop a framework for accurately and efficiently identifying configuration segments that contribute to forwarding requirement violations. We view this task as a problem of identifying *incompatibilities* between router configurations and forwarding requirements. For example, in the network presented in §2.2, `core2`'s configuration contains an ACL that only permits traffic destined for `dept2/3`, whereas one of the forwarding requirements (`FR2`) is that all departments can communicate. This configuration segment and forwarding requirement are incompatible. Of course, incompatibilities are not always this obvious: e.g., the error in `edge`'s iBGP configuration (`CE1`) is subtle.

Our key insight is to identify incompatibilities using minimal correction sets (MCSes) derived from an SMT-based encoding of a network's configurations and forwarding requirements. Several prior works [3, 10, 11] have introduced SMT-based models that encode a network's configurations and control logic as a conjunction of logical constraints, $C \land L$.

(a) *With ACL and static routes*

(b) *With ACL, OSPF, and additional link*
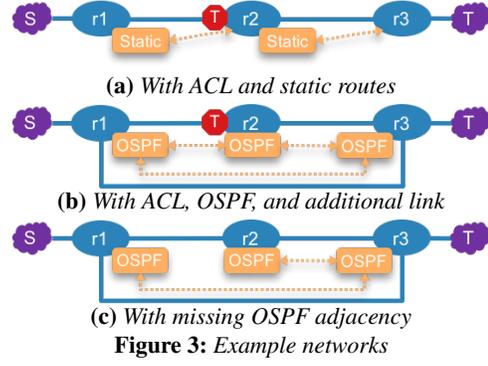
(c) *With missing OSPF adjacency*
Figure 3: *Example networks*

1   $reach_S^T$
2   $reach_S^T \iff reach_{r1}^T$
3   $reach_{r1}^T \iff fwd_{r1 \to T}^T \lor (fwd_{r1 \to r2}^T \land reach_{r2}^T)$
4   $fwd_{r1 \to T}^T \iff connected_{r1}^T \land allow_{r1 \to T}^T$
5   $fwd_{r1 \to r2}^T \iff \neg connected_{r1}^T \land static_{r1 \to r2}^T \land allow_{r1 \to r2}^T$
6   $static_{r1 \to r2}^T$
7   $static_{r2 \to r3}^T$
8   $connected_{r3}^T$
9   $\neg allow_{r1 \to r2}^T$

Figure 4: *Partial encoding of the forwarding requirement (red), control logic (yellow), and configuration (blue) for the example network in Figure 3a*

These constraints define the paths the network will compute for specific prefixes under various failures. Forwarding requirements of interest are appended as a conjunction of additional logical constraints, *R*. The satisfiability of these models can be checked using an SMT solver (e.g., Z3 [43]) in order to determine whether the configurations (can be modified to) satisfy all forwarding requirements.

For example, Figure 4 contains a simple, partial encoding of the forwarding requirement, control logic, and router configurations for the example network in Figure 3a. (The detailed encoding used by CEL is discussed in §4.) To satisfy the requirement that subnet *T* is reachable from *S* (line 1), *T* must be reachable from *r*1 (line 2). *T* is reachable from a router (e.g., *r*1) iff the router either forwards directly to *T* or to a neighboring router (e.g., *r*2) which can reach *T* (e.g., line 3). According to the control logic, a router (e.g., *r*1) forwards traffic directly to a subnet (e.g., *T*) iff the subnet is directly connected to the router and the traffic destined for the subnet is allowed by an ACL (e.g., line 4). Furthermore, a router (e.g., *r*1) forwards traffic destined for a subnet (e.g., *T*) to another router (e.g., *r*2) if a static route is configured, the subnet is not directly connected, and traffic destined for the subnet is allowed by an ACL (e.g., line 5). Finally, the encoding indicates which static routes, connected subnets, and ACLs are configured (e.g., lines 6–9). Checking these constraints

using an SMT solver reveals they are unsatisfiable—i.e., the forwarding requirement is violated.

When the constraints are unsatisfiable, we can compute a *minimal correction set* (MCS)—i.e., a subset of logical constraints whose removal from the problem causes the remaining constraints to become satisfiable. For example, if we remove the constraint that traffic for T is not allowed to be sent from $r1$ to $r2$ (line 9 in Figure 4), then the remaining constraints, including the forwarding requirement (line 1), are satisfiable. MCSes thus provide a way to identify incompatibilities between configuration segments (e.g., the ACL on $r2$ in Figure 3a) and forwarding requirements (e.g., FR2).

However, accurately identifying configuration errors using MCSes requires addressing several challenges, which we articulate below. §4 and §5 discuss how CEL addresses them.

3.1 Challenges

Challenge 1: limiting error localization to configurations.
The first challenge arises from the fact that both configurations and control logic impact whether a forwarding requirement is satisfied. For example, FIB entries partially depend on: (*i*) a router's logic for populating the global routing information base (RIB) from the RIBs of individual routing protocols, and (*ii*) ACLs defined in the router's configuration. Consequently, we may attribute forwarding requirement violations to a router's control logic or its configuration. For example, we showed above that the constraint representing the ACL on $r2$ (line 9 in Figure 4) is an MCS—i.e., removing this constraint allows all other constraints, including the forwarding requirement (line 1), to be satisfied. However, the constraint that $r1$ forwards directly to T iff T is directly connected (line 4) is also an MCS. The fundamental difference between these MCSes is that the former represents configuration while the latter represents control logic.

While configuration errors and software bugs are both common causes of network outages [30, 44], we focus on localizing the former, because: (*i*) configurations are fully under network engineers' control, whereas router software is often closed-source; and (*ii*) software fault localization has already been extensively studied [41]. Limiting our localization to configuration errors requires an encoding that separates configuration and control logic into separate constraints (§4.2) and an MCS extraction algorithm that only produces MCSes with configuration-related constraints (§5.2).

Challenge 2: partially violated forwarding requirements.
The second challenge arises from the fact that link and node failures may also impact whether a forwarding requirement is satisfied. For example, the network in Figure 3b has two loop-free physical paths from S to T. OSPF will choose the shorter path ($r1 \rightarrow r3$) when the $r1$–$r3$ link is active and the longer path ($r1 \rightarrow r2 \rightarrow r3$) when the link fails. However,

T is only reachable from S when the shorter path is used, because an ACL on $r2$ blocks traffic destined for T.

Using MCS-based error localization in this scenario is difficult due to the way existing SMT-based models handle failures. SMT-based models designed for verification [3] identify a single failure scenario (if any) under which a requirement is violated. In particular, the models include free variables representing link states and constraints encoding the negation of the forwarding requirement. The problem is satisfiable iff there exists a failure scenario under which the requirement is violated (e.g., the $r1$–$r3$ link fails), and unsatisfiable iff the requirement is fulfilled under every failure scenario (which is not the case for our example network). However, by definition, a satisfiable problem has no MCSes [33], so we cannot employ MCS-based error localization with this encoding of failures and forwarding requirements. If we remove the negation from the forwarding requirement, then the problem is satisfiable iff there exists a scenario under which the requirement is fulfilled (e.g., no link failures), and unsatisfiable iff the requirement is violated under every failure scenario (which is not the case for our example). So again there are no MCSes.

A different problem arises if we use an SMT-based model designed for synthesis [10]. These models avoiding introducing free variables to model link failure state by instead requiring backup (and primary) paths to be specified in the forwarding requirement. For example, we would need to reformulate our forwarding requirement to say that traffic from S to T is forwarded along any of $r1 \rightarrow r3, r1 \rightarrow r2 \rightarrow r3$. While this is feasible to do for a small network, computing such paths for a larger network is difficult [37].

To identify problematic configuration segments when forwarding requirements are violated under some, but not all, failure scenarios, we must employ a more sophisticated approach for considering failures (§5.1).

Challenge 3: identifying missing configuration statements.
While the aforementioned requirement violations were caused by the presence of an ACL, the *omission* of a configuration segment can also cause violations. For example, every router in Figure 3c runs OSPF and is physically connected to every other router, but OSPF is not configured on the interfaces that connect $r1$ and $r2$. Consequently, if the $r1$–$r3$ link fails, T will be unreachable from S. We expect such omissions occur in practice due to the lack of (full) automation in some networks (§2.1) and the large scope of some configuration updates [17].

Existing SMT-based models [3, 10] are unsuitable for locating the configuration error, because they only encode statements that are present in the configurations.[3] No constraints are created for absent configuration, so it is impossible to

---

[3]NetComplete, which is designed for configuration synthesis, accepts configuration templates containing holes, but only the holes that are explicitly present in the template are encoded in the model.

produce an MCS containing a constraint related to the configuration error. To ensure CEL can detect missing configuration segments that cause requirement violations, we must explicitly encode the absence of configuration (§4.3).

Challenge 4: ranking plausible errors. There are often many ways to satisfy a forwarding requirement. Thus, if a requirement is violated, there are often many MCSes: e.g., for 75% (45%) of the configuration errors we introduce into real university networks (§6.2), there are more than 20 (100) MCSes.

Requiring network engineers to sift through a large list of plausible errors places an undue burden on the engineers and undermines the motivation for CEL. Consequently, it is essential to rank errors based on their likelihood of being true errors. However, what constitutes a "true error" depends on engineers' network management practices [17]. To rank configuration error candidates, we must consider what design and operational practices a network follows (§5.3).

Challenge 5: identifying multiple configuration errors. The above examples only contain one configuration error, but real configurations often contain multiple errors [12, 14, 22, 28]. In the simplest case the errors are related, insofar as they impact the same set of requirements under the same set of failures. For example, if we added an ACL on $r2$ in the network in Figure 3c, then both the ACL and the absence of an OSPF relationship between $r1$ and $r2$ would cause T to be unreachable from S when the $r1$–$r3$ link fails. However, it is also possible for different requirement violations to be caused by different errors. For example, assume we instead configured an ACL on $r1$'s interface to $r3$. When the $r1$–$r3$ link is active, the ACL causes T to be unreachable from S; when the $r1$–$r3$ link fails, the absence of an OSPF relationship between $r1$ and $r2$ causes T to be unreachable. Ensuring all errors are identified requires a carefully designed MCS enumeration algorithm (§5)

§4 and §5 discuss how CEL addresses these challenges.

4 SMT-BASED NETWORK ENCODING

As discussed in §3.1, CEL requires an SMT-based network model that satisfies two important properties: (*i*) configuration is encoded separately from control logic, so we can limit our localization to configuration errors; and (*ii*) the absence of configuration is explicitly encoded, so we can identify omissions that contribute to forwarding requirement violations. Unfortunately, none of the existing SMT-based network models designed for control plane verification/synthesis [3, 10, 38, 40] satisfy these properties. However, Minesweeper [3] and NetComplete [10] accommodate a broad range of configuration constructs, so we use these models as a starting point, and adapt them to satisfy the aforementioned requirements.

4.1 Existing Models

Minesweeper and NetComplete encode a network's forwarding behavior as a system of logical constraints over symbolic route advertisements. Symbolic advertisements are created for each routing adjacency—i.e., pair of neighboring routers that are configured to exchange routes using a particular routing protocol. Each symbolic advertisement is composed of multiple symbolic variables that mirror the fields in actual route advertisements: e.g., prefix, path cost/length, and local preference. Constraints on these symbolic advertisements express: (*i*) the route advertisement import/export behavior of each router—which includes the application of route policies that filter and/or modify route advertisements, the forwarding of route advertisements, and the origination of route advertisements; (*ii*) the protocol-specific and cross-protocol route selection procedures within each router; and (*iii*) the consequent packet forwarding behavior. We explain these constraints in detail in Appendix A.

4.2 Encoding Configuration

As discussed earlier, configuration must be encoded separately from control logic in order to limit CEL's output to configuration errors. However, the constraints generated by Minesweeper and NetComplete co-mingle configuration and control logic. For example, the constraint that encodes the routes `core1` exports to `core2` (Figure 12b in Appendix A) includes the OSPF originated prefix for `dept1` (`1.0.1.0/24`) and link costs specified in `core1`'s configuration. Similarly, the constraint that encodes which packets `core2` forwards to `core3` (Figure 12c in Appendix A) includes the `deptFilter` ACL in `core2`'s configuration.

CEL separates configuration from control logic by: (*i*) replacing configuration-based expressions in the aforementioned route import/export, route selection, and packet forwarding constraints with symbolic configuration variables; and (*ii*) adding a constraint for each configuration variable that equates the variable with the expression specified in the current configurations. Figure 5 illustrates this encoding.

Only expressions that are *solely based on configuration* are replaced with symbolic configuration variables. Expressions that correspond to control logic are left unchanged. For example, the prefix `core1` is configured to originate (`1.0.1.0/24`) is replaced with a symbolic variable ($cfg^{OSPF:originate}_{core1:1.0.1.0/24}$), because `core1`'s configuration lists each network advertised by OSPF. However, the comparison between the originated prefix and the destination prefix specified in the forwarding requirement is left in the export constraint (line 2 in Figure 5a), because a router always forwards packets according to their destination IP address. Similarly, the cost `core1` adds to advertisements sent to `core2` is replaced with a symbolic variable ($cfg^{OSPF:cost}_{core1 \to core2}$), because

1 **if** $\text{cfg}^{\text{OSPF:adjacency}}_{\text{core1}\rightarrow\text{core2}}$ **then**
2   **if** destination $\in \text{cfg}^{\text{OSPF:originate}}_{\text{core1:1.0.1.0/24}}$ **then**
3      $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{prefix} = \text{cfg}^{\text{OSPF:originate}}_{\text{core1:1.0.1.0/24}}$
4      $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{cost} = \text{cfg}^{\text{OSPF:cost}}_{\text{core1}\rightarrow\text{core2}}$
5   **else if** $\text{OSPF}_{\text{core1-best}}.\text{valid} \wedge \text{cfg}^{\text{OSPF:filter}}_{\text{core1:default}}$ **then**
6      $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{valid} = \text{True}$
7      $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{prefix} =$
         $\text{OSPF}_{\text{core1-best}}.\text{prefix}$
8      $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{cost} =$
         $\text{OSPF}_{\text{core1-best}}.\text{cost} + \text{cfg}^{\text{OSPF:cost}}_{\text{core1}\rightarrow\text{core2}}$
9   **else** $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{valid} = \text{False}$
10 **else** $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}.\text{valid} = \text{False}$

(a) *core1 export to core2*

1 $fwd_{\text{core2}\rightarrow\text{core3}} \Leftrightarrow \text{Overall}_{\text{core2-best}} =$
    $\text{OSPF}_{\text{core2}\leftarrow\text{core3}} \wedge \text{cfg}^{\text{OutACL}}_{\text{core2}\rightarrow\text{core3}}$

(b) *core2 forward to core3*

1 $\text{cfg}^{\text{OSPF:originate}}_{\text{core1:1.0.1.0/24}} = 1.0.1.0/24$
2 $\text{cfg}^{\text{OSPF:cost}}_{\text{core1}\rightarrow\text{core2}} = 1$
3 $\text{cfg}^{\text{ACL}}_{\text{core2:deptFilter}} = \text{source} \in 1.0.2.0/24 \vee \text{source} \in$
    $1.0.3.0/24$
4 $\text{cfg}^{\text{OutACL}}_{\text{core2}\rightarrow\text{core3}} = \text{cfg}^{\text{ACL}}_{\text{core2:deptFilter}}$
5 $\text{cfg}^{\text{OutACL}}_{\text{core2}\rightarrow\text{core1}} = \text{True}$

(c) *Configuration*

Figure 5: *Constraints with symbolic configuration variables*

`core1`'s configuration lists the OSPF cost for each interface, but the arithmetic expression is left in the export constraint (line 8 in Figure 5a), because OSPF costs are always additive. In contrast, the ACL-based expressions in forwarding constraints (e.g., line 2 in Figure 12c) are replaced with a symbolic configuration variable ($\text{cfg}^{\text{OutACL}}_{\text{core2}\rightarrow\text{core3}}$), because packet filters defined in a router's configuration specify both the packet fields and values to match.

It is important to note that CEL's use of symbolic configuration variables is fundamentally different from NetComplete's use of symbolic configuration variables. In particular, CEL replaces *all* configuration-based expressions with symbolic variables, whereas NetComplete only uses symbolic configuration variables for holes in configuration templates; concrete configuration values are still embedded directly in NetComplete's route import/export constraints. Furthermore, CEL fully constrains the value of symbolic configuration variables, whereas symbolic configuration variables in NetComplete are (partially) unconstrained free variables.

4.3 Encoding Absent Configuration

As illustrated above, the contents of constraints are based on the contents of routers' configurations: e.g., originated prefixes, link costs, and packets filters. However, as demonstrated

in §3.1, the absence of configuration is also meaningful and must be encoded in the network model.

Unfortunately, the space of potentially omitted configuration is extremely large: e.g., configurations from a large ISP contain 56 different top-level commands [7], and the command reference for the latest version of Cisco IOS contains hundreds of commands [1]. Including every possible configuration variable in the system of constraints would vastly increase the size of the problem and be prohibitively expense.

Fortunately, prior work on software fault localization has shown that erroneously omitted statements are often present (in a similar form) elsewhere in the code [39]. A review of configuration omission errors reported in prior works [12, 13, 28] and the high prevalence of configuration clones/templates [6, 22, 28] suggests omissions in network configurations are also likely present (in a similar form) elsewhere in the configurations. Consequently, we use the configuration statements present in the current configurations to guide the inclusion of potentially omitted configuration variables.

In particular, CEL determines which types of configuration variables (e.g., $\text{cfg}^{\text{OSPF:originate}}$, $\text{cfg}^{\text{OSPF:cost}}$, $\text{cfg}^{\text{OutACL}}$) appear in the network configurations, and includes all plausible instances of these variables in control logic (i.e., route import/export, route selection, and packet forwarding) constraints. For example, ACLs are configured on some interfaces in the example campus network in Figure 3. Thus, CEL adds a $\text{cfg}^{\text{OutACL}}$ variable to every packet forwarding constraint (e.g., Figure 5b), even if no ACL is currently configured on the corresponding interface. CEL also adds a configuration constraint for every $\text{cfg}^{\text{OutACL}}$ variable, which asserts the variable equals the currently configured ACL (e.g., line 4 in Figure 5c), or True (i.e., permit all traffic) if no ACL is configured for the interface (e.g., line 5). If a specific configuration option (e.g., $\text{cfg}^{\text{OSPF:filter}}$) is never used in the network, then no configuration variables of this type are introduced.

Two additional principles govern CEL's inclusion of configuration variables for potential omissions. First, we assume network engineers will not erroneously omit an entire routing protocol from a device's configuration. Hence, we only add omission variables for a protocol (e.g., $\text{cfg}^{\text{OSPF:cost}}$) if that protocol is currently enabled on the device. Second, we consider the semantics of each potentially omitted configuration option, and only introduce configuration variables for feasible instances. For example, CEL only adds $\text{cfg}^{\text{OSPF:orignate}}$ variables for directly connected prefixes that are not already configured to be announced by OSPF; announcing other prefixes (without enabling route redistribution) is infeasible.

5 COMPUTING MCSES

We now turn our attention to the task of computing MCSes as a means of identifying configuration errors that contribute

to forwarding requirement violations. As discussed in §3, an MCS is a subset of logical constraints whose removal from the problem allows the problem to be satisfied. More formally, let $\varphi = \{c_1, \ldots, c_{|C|}, l_1, \ldots, l_{|L|}, r_1, \ldots, r_{|R|}\}$ represent the conjunction (or union) of the constraints that encode a network's configurations (C), control logic (L), and forwarding requirements (R). An MCS is a subset of clauses $\gamma \subseteq \varphi$ such that $\varphi \setminus \gamma$ is satisfiable and $\forall x \in \gamma : (\varphi \setminus \gamma) \cup \{x\}$ is unsatisfiable. Our goal is to compute MCSes such that $\gamma \subseteq C$, because a router's control logic is typically outside of a network engineer's control (§3.1) and a network's forwarding requirements are fixed. Additionally, we seek to prioritize MCSes that are more likely to represent "true errors" (§3.1).

In this section, we discuss three key mechanisms we use to achieve this goal: (*i*) a counterexample-guided exploration of the failure scenarios under which a forwarding requirement is violated; (*ii*) a domain-specific MCS enumeration algorithm based on MARCO [29]; and (*iii*) an MCS ranking algorithm inspired by engineers' network management practices [17].

5.1 Handling link/node failures

As discussed in §3.1, existing SMT-based network models [3, 10] are ill-suited for localizing configuration errors that only arise under some failure scenarios. Minesweeper's encoding of link/node failures using free variables causes φ to be *satisfiable* (i.e., $\gamma = \varnothing$) if there is at least one failure scenario under which the forwarding requirement is (not) violated. NetComplete's encoding of primary/backup paths in the forwarding requirements scales poorly.

Consequently, CEL uses a *counterexample-guided exploration of failure scenarios* (Figure 6) to compute MCSes and localize errors. First, we augment each route import constraint (e.g., Figure 12a in Appendix A) with a conditional expression representing the availability of the corresponding link (e.g., $\neg\mathsf{failed}_{\mathsf{core1}\leftarrow\mathsf{core2}}$).[4] Then, we check whether $C \wedge L \wedge \neg R$ is satisfiable. If the problem is unsatisfiable, then the forwarding requirement is never violated (i.e., $\neg R = \mathsf{False}$) and the process terminates. Otherwise, the solver produces a satisfying solution which represents a counterexample: i.e., a scenario under which the forwarding requirement is violated. The satisfying solution includes concrete values for the free variables representing link failure states (e.g., $\mathsf{failed}_{\mathsf{core1}\leftarrow\mathsf{core2}}$). We produce an expression F that equates each failure state variable with its concrete value: e.g., $\mathsf{failed}_{\mathsf{core1}\leftarrow\mathsf{core2}} = \mathsf{True} \wedge \mathsf{failed}_{\mathsf{core1}\leftarrow\mathsf{core3}} = \mathsf{False} \wedge \ldots$

We construct a new system of constraints: $\varphi' = C \wedge L \wedge R \wedge F$. This system of constraints is unsatisfiable, because F represents a failure scenario in which the forwarding requirement

---

[4]It is trivial to extend CEL to also support node failures.
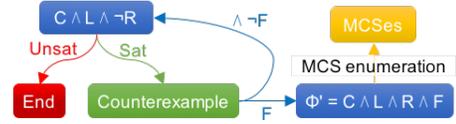
Figure 6: *Counterexample-guided exploration of failures*

is violated (i.e., $\neg R = \mathsf{True}$). Hence, we can use φ' to compute MCSes (§5.2) and localize errors that manifest under this failure scenario.

However, as demonstrated in §3.1, different configuration errors may manifest under different failure scenarios. To ensure we detect all of these errors, we must consider all failure states under which the forwarding requirement is violated. To achieve this, we conjunct $\neg F$ with $C \wedge L \wedge \neg R$ and check if the updated problem has a satisfying solution. Any solution will contain a different combination of link failures, since the original failure scenario is disallowed by the inclusion of $\neg F$ in the system of constraints. We repeat this process until no more satisfying solutions (i.e., counterexamples) exist.

Note that multiple failure scenarios may result in the same forwarding behavior. For example, if we modify the network in Figure 3c to include an ACL on $r3$'s incoming link from $r1$, traffic from S to T will be forwarded along the path $r1 \rightarrow r3$ and dropped at $r3$ when any of the following sets of links fail: $\varnothing$, $\{r1\text{–}r2\}$, $\{r2\text{–}r3\}$, or $\{r1\text{–}r2, r2\text{–}r3\}$. As an optimization, we can formulate φ' and compute MCSes using just one of these failure scenarios, since the network behaves equivalently, and hence manifests the same configuration error(s), under each of these failure scenarios.

5.2 Computing MCSes

We now present our domain-specific algorithms for computing MCSes for a set of unsatisfiable network constraints $\varphi' = \{c_1, \ldots, c_{|C|}, l_1, \ldots, l_{|L|}, r_1, \ldots, r_{|R|}, f_1, \ldots, f_{|F|}\}$. We begin with a simple algorithm for computing a single MCS. We then use this as a subroutine of an algorithm for enumerating all MCSes. Computing multiple MCSes is essential for exposing additional/alternative configuration errors. Both algorithms are based on a state-of-the-art MCS enumeration algorithm called MARCO [29].

Background. Before delving into our algorithms, we highlight two additional types of constraint sets that are relevant to MCSes. A *maximal satisfiable subset* (MSS) is a subset of constraints $\kappa \subseteq \varphi'$ such that κ is satisfiable and $\forall x \in \varphi' \setminus \kappa : \kappa \cup \{x\}$ is unsatisfiable. An MCS is the complement of an MSS: i.e., $\gamma = \varphi' \setminus \kappa$. A *minimal unsatisfiable subset* (MUS) is a subset of constraints $\mu \subseteq \varphi'$ such that μ is unsatisfiable and $\forall x \in \mu : \mu \setminus \{x\}$ is satisfiable. An MCS is a minimal hitting set of a constraint system's MUSes: i.e., an MCS contains at least one constraint from every MUS [9, 35].

8

Input: φ', a set of unsatisfiable constraints
1 $\kappa = \varnothing$
2 **for** $x \in \varphi'$ **do**
3      $\kappa = \kappa \cup \{x\}$
4   **if** κ is unsat **then**
5          $\kappa = \kappa \setminus \{x\}$
   Output: κ, an MSS

Figure 7: *Algorithm for computing a single MSS*

5.2.1 Computing a single MCS. We can compute an MCS by computing an MSS (κ) and taking its complement ($\gamma = \varphi' \setminus \kappa$). MARCO [29] employs a simple algorithm for computing an MSS (Figure 7): start with an empty set (line 1); iteratively add constraints to the set (lines 2–3); check the satisfiability of the constraint set after each addition (line 4); if adding a constraint causes the set of constraints to become unsatisfiable, then remove the most-recently-added constraint before adding the next constraint (lines 4–5).

Unfortunately, this simple algorithm does not guarantee that only configuration-related constraints are included in the MCS (i.e., $\gamma \subseteq C$). Thus, instead of starting with $\kappa = \varnothing$, we initialize κ to the set of all logic-, requirement-, and failure-related constraints (i.e., $\kappa = L \cup R \cup F$). This ensures the MSS always includes all non-configuration-related constraints (i.e., $\kappa \supseteq \varphi' \setminus C$), and the MCS, which is the complement of the MSS, only includes configuration-related constraints (i.e., $\gamma \subseteq C$). This avoids pointless correction sets with constraints that cannot be modified through configuration changes.

With this change, the time complexity for generating a single MCS for a system of constraints is linear in the number of configuration constraints. However, configuration errors are often (relatively) small [22, 28], so MCSes will usually be small; correspondingly, MSSes will usually include most configuration constraints. This motivates us to employ a divide-and-conquer approach to compute MCSes in $O(\log n)$: instead of adding configuration constraints to κ one at a time, we split the configuration constraints in half and add either half to κ. If κ becomes unsatisfiable, we recursively divide that half and try to add each of these smaller sets of configuration constraints. If κ is satisfiable after adding a group of configuration constraints, we do not recurse.

5.2.2 Enumerating all MCSes. To enumerate all MCSes, while preserving the restriction that MCSes only contain configuration-related constraints, we modify the MARCO algorithm [29]. MARCO methodically explores the power set lattice for a set of constraints (φ') to identify the frontier between the satisfiable and unsatisfiable regions. Figure 8 shows the lattice for a small example set of network constraints $\varphi' = \{c_1, c_2, l_1, l_2, r_1, f_1\}$. Every subset of φ' is either satisfiable (green) or unsatisfiable (red). MSSes and MUSes are local maxima (green circles) and minima (red circles), respectively,

along the frontier between the satisfiable and unsatisfiable regions. Since every subset (superset) of a (un)satisfiable set is (un)satisfiable, the set of all MSSes (MUSes) fully specify the frontier [9, 35].

MARCO efficiently identifies the frontier using the following algorithm: select an unexplored (i.e., uncolored) subset σ in the power set lattice; check the satisfiability of σ; if σ is (un)satisfiable grow (shrink) σ to an MSS (MUS) using the algorithm in Figure 7 (or a variant that removes one constraint at a time) with σ used as the initial set; mark the MSS (MUS) and all of its subsets (supersets) in the lattice as (un)satisfiable. The algorithm terminates when the satisfiability of all subsets is known (i.e., the lattice is fully colored).

However, as discussed above, we only care about MSSes that include all non-configuration related constraints (i.e., $\kappa \supseteq \varphi' \setminus C$). In other words, we only need to run MARCO on the sublattice whose bottom is $\varphi' \setminus C$ and whose top is φ' (e.g., the gray region in Figure 8). This has the beneficial side effect of substantially reducing the size of the region MARCO needs to explore.

5.3 Ranking MCSes

As we discussed in §3.1, there are often many MCSes: e.g., for 75% (45%) of the synthetic configuration errors we introduce in real university networks (§6.2), there are more than 20 (100) MCSes. The large number of MCSes stems from the multitude of ways a forwarding requirement may be satisfied [11] and/or the presence of multiple configuration errors [12, 14, 22, 28]. Requiring network engineers to sift through a large list of plausible errors places an undue burden on the engineers and undermines the motivation for CEL. Consequently, filtering/prioritizing MCSes is essential for CEL to be useful in practice.

Based on our prior observation that configuration errors are often (relatively) small [22, 28] (§5.2.1), we prioritize smaller MCSes. We also aggregate MCSes across all forwarding requirement violations and failure scenarios to eliminate duplicates and ensure we cover all errors. We show in §6.2 that this approach works well in practice.

6 IMPLEMENTATION & EVALUATION

We implemented CEL atop the Minesweeper network verifier in $\approx 6k$ lines of Java code. Our current implementation can localize errors to specific interface states (i.e., up/down), layer-3 adjacencies, OSPF and BGP routing adjacencies, originated prefixes, OSPF link costs, route filter definitions, uses of route filters, ACL definitions, and uses of ACLs. We plan to make our implementation open source.

We evaluate CEL along three dimensions: (*i*) Can CEL locate real configuration errors? (*ii*) How accurate is CEL's error localization? (*iii*) How quickly can CEL locate errors?
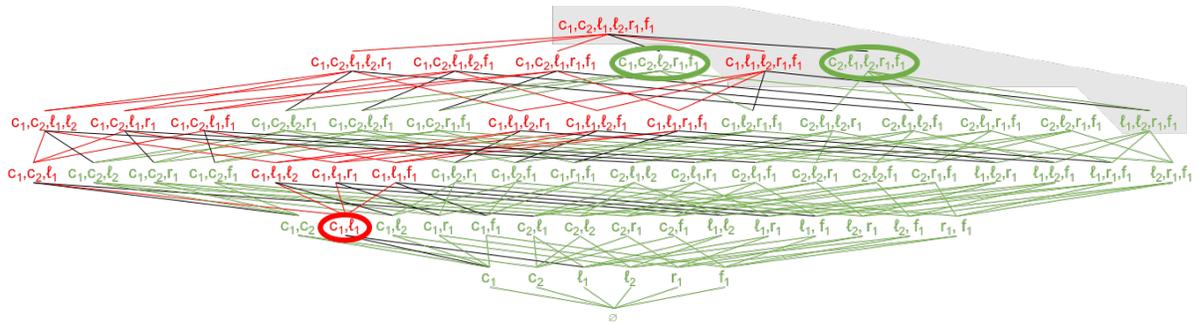
Figure 8: *Power set lattice for example network constraints; the sublattice CEL explores is highlighted in gray*

Configurations. We use real configurations from two universities networks and synthetic configurations for eight wide area networks (WANs). The real configurations come from the universities' core and distribution routers.[5] *UnivA* has 11 routers and supports 65k users; *UnivB* has 28 routers and supports 30k users. Both universities use OSPF for internal routes and iBGP for external routes; we exclude the latter from our analysis, because we do not know the networks' external routes. The synthetic configurations are modeled on eight real WAN [27] topologies that range in size from 34 to 159 routers. The configurations, which use eBGP, are synthesized using NetComplete [10]. All configurations are written in Cisco's IOS language.

Forwarding requirements. We do not know the complete forwarding requirements for the university networks, so we only consider: (*i*) reachability between subnets associated with different departments, (*ii*) protection of management subnets (*UnivA* only), and (*iii*) protection against source-spoofing (*UnivB* only). For the WANs, we only consider reachability between pairs of edge routers.

Experimental setup. We use a server with a 10-core 2.4 Ghz processor and 128GB of memory. Unless otherwise noted, we limit CEL's MCS enumeration time to 10 minutes.

6.1 Locating real configuration errors

We used CEL to check each network's compliance with the aforementioned forwarding requirements. Below, we highlight several types of violations we found and the corresponding configuration errors CEL located.

ACL not applied. One of *UnivA*'s forwarding requirements is to only allow certain, trusted sources to access device management VLANs. However, CEL reported that this requirement was violated for two of the device management VLANs. For each of these violations, CEL computed about a dozen MCSes, each of which contained configuration constraints related to VLAN interface state (up) or the absence of an ACL on some VLAN. This allowed us to quickly diagnose the problem

as a missing ACL on the device management VLANs. Configuration consistency checkers [4, 28] cannot easily detect this error, because there are no clear outliers: some VLANs should have the ACL applied, and some VLANs should not have the ACL applied; the correct behavior for a given VLAN can only be determined from *UnivA*'s forwarding requirements, which consistency checkers do not consider.

Incorrect ACL rules. One of *UnivB*'s forwarding requirements is to prevent source spoofing at VLAN granularity—i.e., a malicious host should not be allowed to impersonate a host within a different VLAN. However, CEL reported that this requirement was violated for two VLANs. For each of these violations, CEL computed two MCSes: one localized the error to a specific ACL definition and the other localized the error to the application of that ACL on a specific VLAN. Upon looking at the configurations, we immediately noticed that each of the identified ACLs was applied only to the VLAN for whom the forwarding requirement was violated, and the network address listed in the ACL did not match the address of the VLAN to which the ACL was applied. Although the ACL defined for each VLAN is based on a standard template, configuration consistency checkers [4, 28] cannot easily detect this error, because the parameter values are unique in each instantiation of this template and must match the address of the VLAN to which the ACL is applied.

Improperly configured backup paths. The *UnivA* network has redundant core and distribution routers and links, such that the network can tolerate three simultaneous link failures. However, CEL reported that departments connected to a particular pair of primary/backup routers would be unreachable if three links failed. CEL identified a single class of failure scenarios that caused the violation, and computed a single MCS that localized the problem to the absence of a layer-3 adjacency between a particular pair of VLAN interfaces on the backup distribution and core routers. Existing SMT-based network models [3, 10] which CEL is based on do not explicitly model layer 2, so we had to manually inspect the configuration based on CEL's output to locate the precise configuration error. Based on CEL's output, we were able to

---

[5]We ignore access switches, because they primarily operate at layer 2.

Type	Change → Impact
OmitNw	Remove BGP or OSPF `network` → router does not advertise a prefix
OmitNb	Remove BGP `neighbor` or OSPF `no passive` → router does not advertise routes to a neighbor
OmitAcl	Remove `ip access-group` from interface → all packets are allowed
OmitAclRule	Remove line from `access-list` → additional packets are blocked or allowed
ExtraAcl	Add `ip access-group` to interface → some packets are blocked

Table 1: *Synthetic errors injected into configurations*

(a) *Recall*   **(b)** *Precision*

Figure 9: *Accuracy by error type*

focus on configuration related to a specific VLAN on two specific routers and quickly determine that the distribution router's interface to the the core router was not configured to participate in the VLAN. In the future, we plan to extend CEL to explicitly model layer 2.

Improperly configured route filters. The sample WAN router configuration template provided with NetComplete [10] is designed to support reachabiliby between every pair of edge routers using a fixed primary or backup path. However, when the backup paths for different router pairs overlap, CEL reported that the edge routers are unreachable if their primary links fail. For each of these violations, CEL identified a single class of failure scenarios that caused the violation, and computed a single MCS that localized the problem to the match criteria used in the route filter on the overlapping router(s). Upon examining these route filters, we quickly noticed that all rules in the route filter used the same prefix list, which only contained the prefix for one of the edge routers, whereas different rules should have used different prefix lists. The authors of NetComplete accepted the bug fix we submitted. This error highlights the need for CEL even in the presence of tools that (partially) generate configurations.

6.2 Accuracy

Next, we evaluate CEL's accuracy. We focus on two metrics: *recall*—i.e., what fraction of a network's configuration errors are identified by CEL; and *precision*—i.e., what fraction of configuration segments identified by CEL contain errors. Unless otherwise noted, we only consider the smallest MCSes.

Errors. Since we do not have access to a list of known errors in the real configurations and the synthetic configurations are designed to be error-free, we introduce synthetic errors into the configurations. The errors we introduce (Table 1) are based on the real errors we found (§6.1) as well as common errors identified in previous measurement studies [31, 44]. We generate three sets of faulty configurations per-network per-error-type, with a few exceptions: we do not introduce *OmitNw* into *UnivB*, because its OSPF stanzas contains a single `network` statement covering all prefixes; and we do
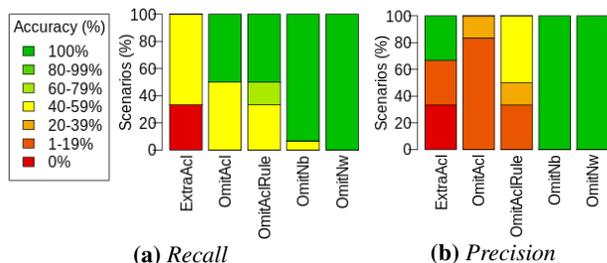
not introduce *OmitAcl*, *OmitAclRule*, or *ExtraAcl* into the synthetic WAN configurations, because they do not use ACLs.

For each type of error, we randomly select which statements we add/remove/modify. In some instances, we remove the same/related statements from multiple devices to avoid introducing errors that could be easily detected by configuration consistency checkers [4, 13, 28]: e.g., for *OmitNb* errors, we remove the adjacency-related configuration segments from both neighbors; for *OmitAclRule* errors, we remove the same line from from every device that contains a copy of the selected ACL; and for *UnivA*, we introduce the same error on a primary router and its corresponding backup router.

Impact of error type. Figure 9a shows CEL's recall for each type of error. CEL locates all *OmitNw* errors in all scenarios and all *OmitNb* errors in all but two scenarios. In contrast, for *OmitAclRule* and *OmitAcl* errors, CEL locates all of the errors for *UnivB* and about half of the errors for *UnivA*. As discussed above, *UnivA* replicates ACLs across devices, so we replicate ACL errors across devices. However, CEL only identifies the errors on some devices. Nonetheless, this is sufficient to draw network engineer's attention to the problem. (We elaborate on differences in accuracy between networks below.) CEL's recall is the worst for *ExtraAcl* errors, because CEL primarily enumerates MCSes containing interface state (up), routing adjacency, and layer-3 adjacency configuration segments. While disabling an interface or adjacency would satisfy the violated forwarding requirements, such changes would likely cause other forwarding requirements to be violated and not match engineers' management practices elsewhere in the network. Developing more sophisticated ranking heuristics that eschew these alternatives is an important area of future work.

Figure 9b shows CEL's precision for each type of error. For all *OmitNw* and *OmitNb* scenarios, CEL only flags the configuration segments we omitted. In contrast, for ACL-related errors, CEL's precision is generally less than 60%. The cause of this low precision is the same as we discussed above for *ExtraAcl* errors. This indicates that CEL is better suited for routing-related errors; improving CEL's accuracy for ACL errors and/or expanding the range of ACL errors other systems [22, 38] can detect is important future work.

(a) *Recall*   **(b)** *Precision*
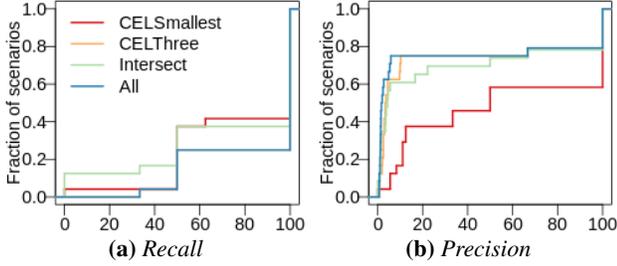
Figure 10: *Impact of ranking*

(a) *Constraints & Failures*   **(b)** *Generate a single MCS*

Figure 11: *Time to complete major tasks*

Impact of network design. Next, we examine CEL's accuracy for different networks. CEL achieves perfect recall and precision for all synthetic configurations, except for one *OmitNb* scenario in one WAN. CEL's recall is also perfect for all *UnivB* scenarios, but CEL achieves perfect recall for only one-third of the *UnivA* scenarios and 50% recall for most of the other *UnivA* scenarios. As we discussed above, *UnivA* replicates ACLs (and we replicate ACL errors) across devices, but CEL only locates the errors on about half of the devices. If we increase the time limit on MCS enumeration (to 1 hour), then CEL is able to explore more of the lattice (§5.2.2) and achieves higher recall.

CEL achieves 100% precision for about half of the *UnivA* scenarios and about one-third of the *UnivB* scenarios. In the remaining scenarios, which all contain ACL errors, the average precision is 25% and 19%, respectively. In summary, CEL's accuracy is partially impacted by network design, but the type of error CEL must locate has a greater impact.

Impact of ranking. Next, we examine how CEL's approach of aggregating MCSes from all violated forwarding requirements and ranking MCSes from smallest to largest (§5.3) impacts CEL's accuracy. We compare four different approaches: CEL's approach in which we only consider the smallest (i.e., highest ranked) MCSes (*CELSmallest*), CEL's approach in which consider the three smallest sizes of MCSes (*CELThree*), only considering MCSes that are common across all violated forwarding requirements (*Intersect*), and considering all MCSes CEL enumerates within 10 minutes (*All*). We limit our analysis to real university configurations, because all scenarios with synthetic WAN configurations have only one MCS.

We observe (Figure 10a) that *All* has perfect recall for 75% of the scenarios, whereas *CELSmallest* and *CELThree* have perfect recall for 60% of the scenarios. Furthermore, all approaches except *Intersect* have at least 50% recall for 96% of the scenarios. However, when we look at precision (Figure 10b), *CELSmallest* has the best overall precision, with at least 33% precision for 55% of the scenarios and perfect precision for 40% of the scenarios. In contrast, *All* has less than 6% precision for 75% of the scenarios and perfect precision for
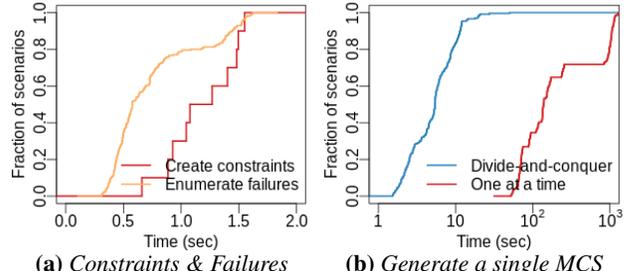
only 20% of the scenarios; the precision for *CELThree* is similar to *All*. In summary, *CELSmallest* significantly improves precision without a substantial decrease in recall.

6.3 Efficiency

Lastly, we evaluate CEL's efficiency. We measure the time CEL takes to complete each of its major tasks: constructing the system of constraints (§4), enumerating failure scenarios (§5.1), and enumerating MCSes (§5.2). We use the same configurations as §6.2. CEL can be applied to multiple forwarding requirements in parallel, so the time required for a single forwarding requirement is the limiting factor.

We observe (Figure 11a) that CEL can construct the system of constraints for the university and WAN networks in ≤1.5s. The time is strongly correlated (coeff = 0.70) with the number of routers in a network. Enumerating the failure scenarios under which a forwarding requirement is violated is also fast (≤1.6s) for the networks and requirements we study.

Enumerating MCSes is the most time consuming task. For half of the scenarios, it takes ≤12s to enumerate all MCSes, while for 40% of the scenarios, we reach the 10 minute time limit we set, and CEL does not enumerate all MCSes. However, as we showed in §6.2, CEL achieves reasonable accuracy even with partial enumeration of MCSes.

As shown in Figure 11b, our divide-and-conquer approach for computing MSSes (§5.2.1) decreases the time to compute a single MCS by approximately two orders of magnitude.

7 RELATED WORK

We discussed existing verification and synthesis tools in §2.2.

Fault localization has been extensively studied by the software engineering research community [41]. CEL's use of correction sets to identify faulty configuration statements is most closely related to Bug-Assist's use of a MaxSMT formulation to compute the maximal set of program statements that may cause a specific test case failure [20, 21]. CEL's aggregation of MCSes across failure scenarios and forwarding requirements (§5) is inspired by spectrum-based fault localization techniques [36].

8 CONCLUSION

We have presented CEL: a system that accurately localizes configuration errors by computing minimal correction sets (MCSes) for an SMT-based network model using domain-specific MCS enumeration algorithms. We showed that CEL is able to pinpoint errors in real university networks and synthesized WAN configurations and identify all routing-related and half of all ACL-related synthetic errors we introduce.

REFERENCES

[1] Command reference, cisco ios xe amsterdam 17.2.x (catalyst 9600 switches). https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9600/software/release/17-2/command_reference/b_172_9600_cr.html.

[2] A. Abhashkumar, A. Gember-Jacobson, and A. Akella. Tiramisu: Fast multilayer network verification. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2020.

[3] R. Beckett, A. Gupta, R. Mahajan, and D. Walker. A general approach to network configuration verification. In *SIGCOMM*, 2017.

[4] R. Beckett, A. Gupta, R. Mahajan, and D. Walker. Abstract interpretation of distributed network control planes. *Proceedings of the ACM on Programming Languages*, 4(POPL):1–27, 2019.

[5] R. Beckett, R. Mahajan, T. Millstein, J. Padhye, and D. Walker. Don't mind the gap: Bridging network-wide objectives and device-level configurations. In *SIGCOMM*, 2016.

[6] T. Benson, A. Akella, and D. Maltz. Unraveling the complexity of network management. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2009.

[7] D. Caldwell, S. Lee, and Y. Mandelbaum. Adaptive parsing of router configuration languages. In *IEEE Internet Network Management Workshop (INM)*, 2008.

[8] A. Chen, Y. Wu, A. Haeberlen, W. Zhou, and B. T. Loo. The good, the bad, and the differences: Better network diagnostics with differential provenance. In *SIGCOMM*, 2016.

[9] J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artif. Intell.*, 32(1):97–130, 1987.

[10] A. El-Hassany, P. Tsankov, L. Vanbever, and M. Vechev. Netcomplete: Practical network-wide configuration synthesis with autocompletion. In *NSDI*, 2018.

[11] A. El-Hassany, P. Tsankov, L. Vanbever, and M. T. Vechev. Network-wide configuration synthesis. In *Computer Aided Verification (CAV)*, 2017.

[12] S. K. Fayaz, T. Sharma, A. Fogel, R. Mahajan, T. D. Millstein, V. Sekar, and G. Varghese. Efficient network reachability analysis using a succinct control plane representation. In *Symposium on Operating Systems Design and Implementation (OSDI)*, 2016.

[13] N. Feamster and H. Balakrishnan. Detecting BGP configuration faults with static analysis. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2005.

[14] A. Fogel, S. Fung, L. Pedrosa, M. Walraed-Sullivan, R. Govindan, R. Mahajan, and T. Millstein. A general approach to network configuration analysis. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2015.

[15] A. Gember-Jacobson, A. Akella, R. Mahajan, and H. H. Liu. Automatically repairing network control planes using an abstract representation. In *Symposium on Operating Systems Principles (SOSP)*, 2017.

[16] A. Gember-Jacobson, R. Viswanathan, A. Akella, and R. Mahajan. Fast control plane analysis using an abstract representation. In *SIGCOMM*, 2016.

[17] A. Gember-Jacobson, W. Wu, X. Li, A. Akella, and R. Mahajan. Management plane analytics. In *Internet Measurement Conference (IMC)*, 2015.

[18] A. Horn, A. Kheradmand, and M. R. Prasad. Delta-net: Real-time network verification using atoms. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2017.

[19] K. Jayaraman, N. Bjørner, J. Padhye, A. Agrawal, A. Bhargava, P. C. Bissonnette, S. Foster, A. Helwer, M. Kasten, I. Lee, A. Namdhari, H. Niaz, A. Parkhi, H. Pinnamraju, A. Power, N. M. Raje, and P. Sharma. Validating datacenters at scale. In J. Wu and W. Hall, editors, *SIGCOMM*, 2019.

[20] M. Jose and R. Majumdar. Bug-assist: Assisting fault localization in ANSI-C programs. In *Computer Aided Verification (CAV)*, 2011.

[21] M. Jose and R. Majumdar. Cause clue clauses: error localization using maximum satisfiability. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2011.

[22] S. K. R. K., A. Tang, R. Beckett, K. Jayaraman, T. D. Millstein, Y. Tamir, and G. Varghese. Finding network misconfigurations by automatic template inference. In R. Bhagwan and G. Porter, editors, *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2020.

[23] P. Kazemian, M. Chang, H. Zeng, G. Varghese, N. McKeown, and S. Whyte. Real time network policy checking using header space analysis. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2013.

[24] P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2012.

[25] A. Khurshid, X. Zou, W. Zhou, M. Caesar, and P. B. Godfrey. VeriFlow: Verifying network-wide invariants in real time. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2013.

[26] H. Kim, T. Benson, A. Akella, and N. Feamster. The evolution of network configuration: A tale of two campuses. In *Internet Measurement Conference (IMC)*, 2011.

[27] S. Knight, H. X. Nguyen, N. Falkner, R. A. Bowden, and M. Roughan. The internet topology zoo. *IEEE Journal on Selected Areas in Communications*, 29(9):1765–1775, 2011.

[28] F. Le, S. Lee, T. Wong, H. S. Kim, and D. Newcomb. Detecting network-wide and router-specific misconfigurations through data mining. *IEEE/ACM Trans. Netw.*, 17(1):66–79, 2009.

[29] M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. Fast, flexible MUS enumeration. *Constraints*, 21(2):223–250, 2016.

[30] H. H. Liu, X. Wu, W. Zhou, W. Chen, T. Wang, H. Xu, L. Zhou, Q. Ma, and M. Zhang. Automatic life cycle management of network configurations. In *SIGCOMM Workshop on Self-Driving Networks*, 2018.

[31] R. Mahajan, D. Wetherall, and T. E. Anderson. Understanding BGP misconfiguration. In *SIGCOMM*, 2002.

[32] H. Mai, A. Khurshid, R. Agarwal, M. Caesar, P. B. Godfrey, and S. T. King. Debugging the data plane with Anteater. In *SIGCOMM*, 2011.

[33] N. Narodytska, N. Bjørner, M. V. Marinescu, and M. Sagiv. Core-guided minimal correction set and core enumeration. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

[34] S. Prabhu, K.-Y. Chou, A. Kheradmand, B. Godfrey, and M. Caesar. Plankton: Scalable network configuration verification through model checking. In *Symposium on Networked Systems Design and Implementation (NSDI)*, 2019.

[35] R. Reiter. A theory of diagnosis from first principles. *Artif. Intell.*, 32(1):57–95, 1987.

[36] J. R. Ruthruff, E. Creswick, M. M. Burnett, C. R. Cook, S. Prabhakararao, M. F. II, and M. Main. End-user software visualizations for fault localization. In *ACM 2003 Symposium on Software Visualization*, 2003.

[37] K. Subramanian, L. D'Antoni, and A. Akella. Synthesis of fault-tolerant distributed router configurations. *POMACS*, 2(1):22:1–22:26, 2018.

[38] B. Tian, X. Zhang, E. Zhai, H. H. Liu, Q. Ye, C. Wang, X. Wu, Z. Ji, Y. Sang, M. Zhang, et al. Safely and automatically updating in-network acl configurations with intent language. In *SIGCOMM*, 2019.

[39] X. Wang. *Automatic Localization of Code Omission Faults*. PhD thesis, 2010.

[40] K. Weitz, D. Woos, E. Torlak, M. D. Ernst, A. Krishnamurthy, and Z. Tatlock. Scalable verification of border gateway protocol configurations with an SMT solver. In *ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2016.

[41] W. E. Wong, R. Gao, Y. Li, R. Abreu, and F. Wotawa. A survey on software fault localization. *IEEE Trans. Software Eng.*, 42(8):707–740, 2016.

[42] Y. Wu, M. Zhao, A. Haeberlen, W. Zhou, and B. T. Loo. Diagnosing missing events in distributed systems with negative provenance. In *SIGCOMM*, 2014.

[43] The z3 theorem prover. https://github.com/Z3Prover/z3.

[44] H. Zeng, P. Kazemian, G. Varghese, and N. McKeown. Automatic test packet generation. In *Conference on emerging Networking Experiments and Technologies (CoNEXT)*, 2012.

[45] P. Zhang, X. Liu, H. Yang, N. Kang, Z. Gu, and H. Li. APKeep: Realtime verification for real networks. In *NSDI*, 2020.

A   EXISTING NETWORK MODELS

As we briefly discussed in §4.1, Minesweeper and NetComplete encode a network's forwarding behavior as a system of logical constraints over symbolic route advertisements. Symbolic advertisements are created for each routing adjacency—i.e., pair of neighboring routers that are configured to exchange routes using a particular routing protocol (e.g., OSPF). Each symbolic advertisement is composed of multiple symbolic variables that mirror the fields in actual route advertisements: e.g., prefix, path cost/length, and local preference. Constraints on these symbolic advertisements express: (*i*) the route advertisement import/export behavior of each router, (*ii*) the route selection procedures within each router, and (*iii*) the consequent packet forwarding behavior. This appendix explains each of these constraints in detail using the example network in Figure 2.

Route import/export. The encoding includes import and export constraints for each direction of a routing adjacency. These constraints encode: (*i*) the application of route policies which filter and/or modify route advertisements, (*ii*) the forwarding of route advertisements, and (*iii*) the origination of route advertisements. For example, Figures 12a and 12b encode which OSPF advertisements `core1` imports from and exports to, respectively, `core2`. A symbolic advertisement representing an imported route (e.g., $\text{OSPF}_{\text{core1}\leftarrow\text{core2}}$) is valid iff an advertisement is exported by the adjacent router (line 1 in Figure 12a) and the advertisement is accepted by the inbound route filter (true by default). A symbolic advertisement representing an exported route (e.g., $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$) is valid iff the forwarding requirement's destination prefix falls within an originated prefix (line 1 in Figure 12b) or the best imported route (if any) is accepted by the outbound route filter (true by default) (line 4 in Figure 12b).

Route selection. Protocol-specific route selection algorithms are encoded using constraints that compare the imported advertisements (e.g., $\text{OSPF}_{\text{core1}\leftarrow\text{core2}}$ and $\text{OSPF}_{\text{core1}\leftarrow\text{core3}}$)

1 **if** $\text{OSPF}_{\text{core2}\rightarrow\text{core1}}$.valid **then** *// Accept all exported*
2 $\quad$ $\text{OSPF}_{\text{core1}\leftarrow\text{core2}}$.valid = True
3 $\quad$ $\text{OSPF}_{\text{core1}\leftarrow\text{core2}}$.prefix =
$\quad\quad$ $\text{OSPF}_{\text{core2}\rightarrow\text{core1}}$.prefix
4 $\quad$ $\text{OSPF}_{\text{core1}\leftarrow\text{core2}}$.cost = $\text{OSPF}_{\text{core2}\rightarrow\text{core1}}$.cost
5 **else** $\text{OSPF}_{\text{core1}\leftarrow\text{core2}}$.valid = False *// None imported*

(a) `core1` *import from* `core2`

1 **if** destination $\in$ 1.0.1.0/24 **then** *// Originate* `dept1`
2 $\quad$ $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$.prefix = 1.0.1.0/24
3 $\quad$ $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$.cost = 1
4 **else if** $\text{OSPF}_{\text{core1}-\text{best}}$.valid **then** *// Forward shortest*
5 $\quad$ $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$.valid = True
6 $\quad$ $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$.prefix = $\text{OSPF}_{\text{core1}-\text{best}}$.prefix
7 $\quad$ $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$.cost = $\text{OSPF}_{\text{core1}-\text{best}}$.cost + 1
8 **else** $\text{OSPF}_{\text{core1}\rightarrow\text{core2}}$.valid = False *// None exported*

(b) `core1` *export to* `core2`

1 $fwd_{\text{core2}\rightarrow\text{core3}} \Leftrightarrow \text{Overall}_{\text{core2}-\text{best}} =$
$\quad$ $\text{OSPF}_{\text{core2}\leftarrow\text{core3}}$
2 $\quad$ $\wedge$ (source $\in$ 1.0.2.0/24 $\vee$ source $\in$ 1.0.3.0/24)
$\quad\quad$ *// deptFilter*

(c) `core2` *forward to* `core3`

Figure 12: *Constraints encoding part of the behavior of the example network in Figure 2*

according to local preference, path length/cost, etc. Cross-protocol route selection algorithms are encoded using constraints that compare the best advertisements from each protocol (e.g., $\text{OSPF}_{\text{core1}-\text{best}}$ and $\text{BGP}_{\text{core1}-\text{best}}$) according to administrative distance. The highest-ranked symbolic advertisement (e.g., $\text{Overall}_{\text{core1}-\text{best}}$) represents an entry in the router's global RIB.

Packet forwarding. Minesweeper also models the contents of a router's FIB. Constraints on symbolic forwarding variables encode the application of: (*i*) routes from the global RIB, and (*ii*) packet filters defined in configurations. For example, Figure 12c encodes whether `core2` forwards to `core3`.

Forwarding requirements. A network's forwarding requirements are expressed in terms of constraints on symbolic route advertisements and/or symbolic forwarding entries in a manner similar to lines 1–3 in Figure 4. CEL considers one forwarding requirement at a time, but could easily be extended to handle multiple forwarding requirements at a time, as done by NetComplete.

Note that we have not discussed Minesweeper's encoding of link/node failure states (in route import/export constraints) or NetComplete's encoding of primary/backup paths (in forwarding requirements). We discuss CEL's handling of failures in §5.1.