Bing Shui

State Key Laboratory for Novel Software Technology Nanjing University Nanjing, Jiangsu, China bingshui@smail.nju.edu.cn Yufan Zhou State Key Laboratory for Novel Software Technology Nanjing University Nanjing, Jiangsu, China yufanzhou@smail.nju.edu.cn

Baowen Xu State Key Laboratory for Novel Software Technology Nanjing University Nanjing, Jiangsu, China bwxu@nju.edu.cn

Abstract

Routers, relying on routing protocols to determine how data packets travel across the Internet, serve as the backbone of modern networks. Vulnerable routing protocols can lead to serious consequences, including data leaks and network congestion. This work focuses on validating the implementation of a key class of routing protocols known as Interior Gateway Protocols (IGPs). Unlike communication protocols such as TCP/IP, which define structured data packets and state machines to facilitate communication, IGPs are designed to automatically manage the network topology. Thus, conventional techniques, which primarily focus on communication correctness, cannot be applied directly to IGPs. We propose ToDIFF, a differential validation technique to uncover IGP bugs in three steps: (1) it uses a network generation algorithm to create random yet valid IGP networks, (2) it applies a semantics-guided program synthesizer to generate equivalent topological programs, and (3) it simulates the network via the equivalent topological programs, with any discrepancies suggesting the presence of a potential bug. We have evaluated ToDIFF on the implementation of two common IGP protocols, OSPF and IS-IS. The results demonstrate that ToDIFF outperforms existing approaches. To date, our tool has successfully identified 26 bugs, all confirmed or fixed by developers.

CCS Concepts

• Security and privacy \rightarrow Software security engineering; Network security.

https://doi.org/10.1145/3719027.3744826

Qingkai Shi* State Key Laboratory for Novel Software Technology Nanjing University Nanjing, Jiangsu, China qingkaishi@nju.edu.cn

Keywords

Routing Protocols; Interior Gateway Protocols; Differential Validation; Topology Synthesis.

ACM Reference Format:

Bing Shui, Yufan Zhou, Jielun Wu, Baowen Xu, and Qingkai Shi. 2025. Validating Interior Gateway Routing Protocols via Equivalent Topology Synthesis. In Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25), October 13–17, 2025, Taipei, Taiwan. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3719027.3744826

1 Introduction

Unlike communication protocols such as TCP/IP, routing protocols aim to automatically manage network topology, establish routing tables, adapt to network changes, and ensure robust network connectivity. Given the critical role routing protocols play, any vulnerabilities in their implementations may lead to severe disruptions. For instance, in 2018, attackers exploited such a vulnerability to redirect traffic intended for Amazon's cloud service to a malicious IP address, causing users to lose at least \$150,000 [41]. Thus, it is crucial to validate the correctness of routing protocol implementations.

Existing Works. Many existing works can validate routing protocol implementations, but are subject to various limitations. First, if we treat routing protocol implementations as common software, traditional methods, such as symbolic analysis [4, 6, 48, 52, 57], fuzzing [18, 20, 24], and LLM-based analysis [51, 53] can be employed to detect general program defects like memory corruptions. However, network protocol bugs usually manifest as subtle deviations from specifications rather than obvious crashes, making their detection and analysis highly dependent on protocol-specific insights — an area where traditional methods often fall short.

In contrast, many works, such as network protocol fuzzing, e.g., [1, 29, 37, 44, 47, 49], differential analysis, e.g., [15, 16, 61], integrate protocol-specific knowledge. The former generates and mutates network packets within protocol constraints to explore execution states, effectively identifying memory safety issues and

Jielun Wu

State Key Laboratory for Novel Software Technology Nanjing University Nanjing, Jiangsu, China jielunwu@smail.nju.edu.cn

^{*}Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS '25, Taipei, Taiwan

^{© 2025} Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-1525-9/2025/10

packet-level errors. The latter heavily relies on multiple implementations of the same protocol and compares different implementations to detect bugs. However, both methods focus on detecting bugs related to communication among multiple parties but ignore problems in routing the communication messages. In other words, the approaches discussed above are ineffective in validating routing protocols, whose main functionality is not to exchange messages but to build a proper topology for routing.

Regarding the main functionality of routing protocols, a few techniques have been introduced, ranging from applying handcrafted topology test suites [10, 11, 13] to automated testing or modelchecking techniques [25, 30, 31, 34, 35, 50, 56, 59, 60]. Manually crafted test suites use predefined topology cases to validate specific routing functionalities but often suffer from limited coverage and miss edge cases. Although capable of generating diverse topological structures for testing or model checking, existing automated techniques still suffer from the notorious testing oracle problem. For example, they have to rely on non-trivial manual efforts to build testing oracles, i.e., manually compute the correct routing tables and compare them to the ones derived from routing protocols.

Our Approach. To automate the validation procedure for routing protocols and address the testing oracle problem, this work proposes ToDIFF, a differential testing technique to validate the implementation of Interior Gateway Protocols (IGPs), a critical category of routing protocols including OSPF [38] and IS-IS [5]. Our key observation is that the same network topology can be established via different but equivalent commands, which we refer to as topological programs. As such, ToDIFF systematically synthesizes multiple equivalent topological programs, which yield multiple network topologies. Since the synthesized topological programs are equivalent, any discrepancies in the outcome network topologies indicate potential bugs in IGP implementations.

More specifically, ToDIFF works in three steps. First, a target network with a random number of routers and a random topology is generated by a dedicated algorithm. Despite being random, we must follow certain constraints so that the outcome topology is valid. Second, ToDIFF synthesizes multiple and equivalent topological programs that are expected to yield the target network. To ensure equivalence and efficiency, we formally define the language semantics and apply a semantics-guided program synthesizer that is guaranteed to generate programs with bounded length. Third, by entering the topological programs into the IGP implementations, routers in the network will communicate with each other so that a correct network topology (e.g., router tables in each router) can be created at the end. Since it usually takes a long time for a topology to be built, ToDIFF includes a series of heuristic yet effective methods to facilitate the convergence of the topology building procedure, thereby improving the whole testing efficiency. Finally, we compare the topologies derived from the topological programs to find potential bugs in IGP implementations.

Contributions. Putting the three steps above together yields ToD-IFF, a new differential analyzer capable of generating diverse networks and uncovering functionality errors in IGP implementations via automated testing oracles. Our design significantly reduces the testing complexity and makes ToDIFF practical for different IGP protocols. In summary, we make the following contributions:

Bing Shui, Yufan Zhou, Jielun Wu, Baowen Xu, and Qingkai Shi



Figure 1: An OSPF network containing 4 routers and 3 areas.



Figure 2: An example of the topological program.

- We propose a novel differential testing framework to detect hidden bugs in IGP implementations.
 - It leverages a network generation algorithm that is capable of producing random yet valid routing topologies among a number of routers.
 - It features a program synthesizer that can generate equivalent topological programs with bounded length, under the guidance of first-order-logic formulated semantics.
 - It designs a series of heuristic methods to accelerate network convergence and differentiate topologies, thus improving the efficiency of testing.
- We implement our approach as a tool, namely ToDIFF, to validate the correctness of multiple IGP implementations. Our tool is efficient, as it can produce tens of topologies in one minute and complete hundreds of differential analyses in one hour. Our tool is also effective, as we have detected 26 zero-day bugs in the FRRouting project [12], a widely used and mature routing protocol suite for Linux and Unix platforms. In contrast, existing approaches fail to detect any of them. All the bugs detected have been confirmed or fixed by the developers. Notably, despite the extensive testing conducted on the protocol implementation, we uncovered some deeply hidden bugs that had remained undiscovered for over 20 years. ToDIFF is publicly available at https://todiff.github.io/.

2 Preliminaries

IGP is a class of essential routing protocols that enable efficient and secure routing within an autonomous system. By dynamically detecting and adapting to network topology changes, IGPs such as OSPF and ISIS help prevent routing loops and ensure reliable data transmission, enhancing overall network integrity. Without IGPs, networks would be more vulnerable to misconfigurations, unauthorized access, and traffic disruptions. Thus, it is a critical task to validate the correctness of IGP implementations. This section provides preliminary background on IGPs, using OSPF as an example. Other IGPs are similar to OSPF.

Routers. A network often contains many interconnected routers, denoted as r_i . Each router r_i has multiple interfaces eth_j, denoted as r_i -eth_j. Routers are connected via interfaces. In Figure 1, Router r_0 and Router r_1 are connected via the interfaces r_0 -eth₀ and r_1 -eth₀.

IP Address and Subnet Mask. For communication, each router interface is assigned an IP address and a subnet mask, both 32-bit integers, but often written as *a.b.c.d* where *a, b, c,* and *d* are the first, second, third, and fourth 8 bits of an integer, respectively. Two interfaces are in the same subnet if the subnet masks are the same and the bitwise-and operation of the IP address and the subnet mask yields identical values. For instance, 192.168.10.11/255.255.0.0 and 192.168.13.12/255.255.0.0 are in the same subnet because they have the same result of the bitwise-and operation, 192.168.0.0, and the same subnet mask, 255.255.0.0.

OSPF Network and Topological Program. OSPF, as a routing protocol, helps routers build routing tables such that data packets can travel from their source to the destination through the most optimal or feasible routes. We refer to a network using OSPF as an <u>OSPF network</u>. OSPF partitions a network into multiple areas, say area₀, area₁, area₂, ..., to simplify management and enhance efficiency. Each area contains a subset of interconnected routers. Figure 1 shows an OSPF network with four routers and three areas: area₀ consists of two routers, r₀ and r₁; area₁ consists of two routers, r₀ and r₂; and area₂ consists of two routers, r₁ and r₃.

To enable area partitioning and the computation of routing tables, each router in the network installs an OSPF implementation, which then reads a <u>topological program</u>, often provided by the network administrators and reflecting how they want to configure the network topology. Figure 2 shows a simple topological program for Router r_0 . Lines 1-2 and Lines 4-5 respectively specify the IP addresses of two interfaces, r_0 -eth₀ and r_0 -eth₁, as well as the areas, area₀ and area₁, they belong to. Line 7 specifies that area₁ is of a special type, namely NSSA, which is explained later.

According to the topological program, the network is automatically partitioned into areas by the OSPF implementation installed in each router. The routing information is also automatically computed at both the intra- and inter-area levels. Within an OSPF area, each router describes its known network topology in a data structure known as Link State Advertisement (LSA) and forwards LSAs to its neighbor routers. Upon receiving LSAs from neighbors, a router stores them in its link-state database and floods them to other neighbors. This process continues until all routers in the area have an identical view of the area's topology. In other words, each router finally knows how routers are connected in the area.

At the inter-area level, routers at the boundary of an area, known as Area Border Routers (ABRs), e.g., r_0 and r_1 in Figure 1, summarize and distribute information about one area's topology to other areas. As such, an area can have an overview of other area's topology to facilitate data packet transmission across different areas.

To sum up, an <u>OSPF network</u> consists of inter-connected routers with necessary topological information, such as the router types, area types, and other parameters such as the IP address and the cost of each router interface. An OSPF network can be described by a set of <u>topological programs</u> for each router in the network.

OSPF Areas. There are multiple special areas in an OSPF network. We use $area_0$ to denote the area that connects all other areas. An

CCS '25, October 13-17, 2025, Taipei, Taiwan

1. int 2. 3.	<pre>c ospf_area_nssa_unset() { auto *area = ospf_area_lookup_by_area_id(ospf, area_id); </pre>
4	area_NNSSATranslatorRole = OSPE NSSA ROLE CANDIDATE:
4.	alea-MSSATTAISIACOINCIE - OSPE_NSSA_NOEL_CANDIDATE,
5. 4	<pre>F 1f (area->NSSATranslatorState == USPF_NSSA_TRANSLATE_ENABLED)</pre>
6. +	<pre>ospf_asbr_status_update(ospf,ospf->redistribute);</pre>
7.	area->NSSATranslatorState = OSPF_NSSA_TRANSLATE_DISABLED;
8	

Figure 3: A bug in an OSPF implementation.

OSPF network ensures the existence of area₀. As demonstrated in Figure 1, area₀ connects area₁ and area₂.

An area could be a normal area, e.g., $area_0$ and $area_2$ in Figure 1, or of special types. For example, $area_1$ is of a special type known as NSSA, which does not forward complete topological information to other routers so as to reduce the workload.

OSPF LSAs. Routers maintain the topological information they have known in a data structure known as Link State Advertisement (LSA). As discussed before, routers exchange LSAs so as to have a complete view of the network topology. An LSA could be of special types. For instance, an external LSA (ExLSA) describes the topology of the network outside the OSPF network. In Figure 1, Router r_2 and Router r_3 forward ExLSA to other routers such that other routers have a view of the external world. Particularly, ExLSA traveling in a normal area is a Normal ExLSA, which is different from the one propagated in an NSSA area, i.e., NSSA ExLSA.

OSPF Routers. Routers in an OSPF network may be of different types. In addition to Area Border Routers (ABRs), which connect multiple areas, ASBR is the other important category of routers in an OSPF network. Usually, an ASBR, e.g., r_2 and r_3 in Figure 1, connects an OSPF network to the external world. However, some special routers are automatically designated as ASBRs to complete some special functionalities. For instance, in Figure 1, Router r_0 is set to ASBR such that it can translate the NSSA ExLSA into Normal ExLSA and forward the ExLSA to the normal area area₀.

3 Motivation and Overview

This section discusses a real-world bug detected by our tool, which shows the limitations of existing works and the merits of ToDIFF.

3.1 Motivation

Bug Example. Let us study a zero-day bug found by our tool in an OSPF implementation from the FRRouting protocol suite, a widely used and mature routing protocol suite for Linux and Unix platforms [12]. Unfortunately, the bug has been hidden for 22 years due to the lack of effective validation techniques.

Consider the OSPF network in Figure 1. The bug happens when area₂ is set to NSSA and then switched back to a normal area. Like Router r_0 in area₁, when area₂ is set to NSSA, Router r_1 becomes an ASBR so that it can translate NSSA ExLSA into Normal ExLSA. When area₂ is switched back to a normal area, the translation is no longer needed. Thus, Router r_1 is no longer an ASBR, either. However, the OSPF implementation fails to cancel its role as an ASBR.

If we look into the OSPF implementation in Figure 3, we can find that the task of canceling an NSSA area is in the function ospf_area_nssa_unset. The function misses Lines 5-6, which updates router to the correct role. This bug may lead to a series of network problems. For example, the bug forces ABR to generate CCS '25, October 13-17, 2025, Taipei, Taiwan

Bing Shui, Yufan Zhou, Jielun Wu, Baowen Xu, and Qingkai Shi



Figure 4: The overall workflow of ToDIFF.

unnecessary LSAs, which then spread widely throughout the network and increase the overhead of LSA propagation. In some cases, external network traffic may be misdirected to this false ASBR, resulting in a traffic black hole that drops all traffic destined for a specific destination, or even more severe network issues.

Limitation of Existing Works. General bug detection techniques such as symbolic analysis [4, 6, 48, 57] and fuzzing [18, 20] struggle to detect this bug because this bug relates to the OSPF specification rather than conventional vulnerabilities like memory corruptions. For example, missing Lines 5-6 in Figure 3 does not violate any memory safety property. Thus, the aforementioned techniques for general software bugs do not work.

Existing works such as network protocol fuzzers [1, 29, 37, 44], differential analyzers [15, 16, 61], and model checkers [2, 9, 39, 40] integrate protocol-specific knowledge to identify bugs. However, these methods primarily focus on issues related to network packet transmission or protocol state transition, treating the process of establishing network topology as a black box. As a result, they are incapable of identifying topology establishment errors, such as the one in the motivating example.

As discussed in Section 1, although there are a few approaches to testing routing protocols, they heavily depend on manual efforts to either construct different topologies [10, 11, 13] or build testing oracles [3, 21–23, 34, 56]. This labor-intensive procedure is errorprone and easy to miss scenarios that should be thoroughly tested. Consequently, this bug had been kept in the OSPF implementation for 22 years before we detected it.

3.2 ToDIFF in a Nutshell

Our approach, ToDIFF, has three key advantages over existing methods. First, our approach is specially designed for routing protocols and aims to detect bugs during topology establishment. Second, our approach does not rely on predefined topologies but randomly generates diverse network topologies for testing. Third, our approach does not rely on manual efforts to build testing oracles but applies differential testing to address the oracle problem.

Despite these advantages, the practical implementation of ToD-IFF must still overcome several network-related challenges. Next, following the workflow shown in Figure 4, we walk through each step of our approach, demonstrating how it identifies the bug in the motivating example and discussing our strategies to address these challenges. Again, we use OSPF as an example. Other IGPs follow similar methodologies.

Step 1: Generating Random yet Valid OSPF Network. The first step is to generate a random yet valid OSPF network. While randomness ensures the diversity of scenarios in which we test OSPF implementations, validity ensures that the generated OSPF



Figure 5: (a) A random yet invalid network. (b) A valid network generated by ToDIFF.

network can be used for testing purposes. The next step will synthesize multiple equivalent topological programs to specify this generated network.

Although it seems straightforward to generate a random graph where each node denotes a router and each edge represents the link between routers, such an approach could often produce invalid OSPF networks. For instance, recall that OSPF partitions a network into multiple areas, and all areas must be connected to area₀. If we randomly generate a graph like Figure 5(a) and randomly partition it into four areas, none of the four areas can play the role of area₀ because there must exist a disconnected area. In addition to the area constraint, there are also multiple other constraints a valid OSPF network must follow. Thus, the challenge below arises:

Challenge 1: How can we systematically generate random yet valid OSPF networks?

To satisfy the area constraint discussed above, we do not directly generate an OSPF network but randomly generate a list of OSPF areas, among which we randomly select one as $area_0$. All other areas are then connected to $area_0$ and are also randomly connected to each other. Figure 5(b) shows an example where the four areas are independently produced by a random graph generation algorithm, each with 4, 3, 3, and 5 routers. When connecting two areas, the area border routers are merged into one. Section 4 will detail how other constraints (e.g., IP address constraints) an OSPF network should follow are satisfied.

Step 2: Synthesizing Equivalent Topological Programs. Given the target network generated by Step 1, this step generates multiple equivalent topological programs describing the target network. For example, let us consider the OSPF network in our motivating example (see Figure 1). We can generate two topological programs for Router r_1 as shown in Figure 6. Figure 6(a) does not set the type of area₂. Thus, it is a normal area by default. In Figure 6(b), area₂ is designated as an NSSA area at Line 8, which is then canceled at Line 9. Thus, the semantics of the two topological programs are equivalent but may lead to different router behaviors, thereby helping us find bugs in the next step.

However, unlike the simple example, a real-world topological program can be highly complex, comprising a large variety of commands with different semantics. This complexity introduces a significant challenge:

Challenge 2: How can we generate equivalent topological programs consisting of a rich set of commands?

CCS '25, October 13-17, 2025, Taipei, Taiwan



Figure 6: An example of the equivalent topological programs.

To address the challenge, we describe the target OSPF network as a first-order logic constraint, denoted as ψ , and formally define the semantics of each command in the topological program via first-order logic, too. With the semantics in hand, we can generate a sequence of random commands, say c_1, c_2, \ldots, c_n , and compute the aggregate semantics, denoted as a constraint ϕ . With careful guidance, ToDIFF generates a topological program (i.e., a command sequence) with semantics ϕ such that *n* is less than a predefined constant and $\phi \equiv \psi$. In other words, the length of the synthesized topological program is bounded and it has equivalent semantics as the target OSPF network. More details can be found in Section 4.

Step 3: Determining Network Convergence and Differentiating Results. We enter the synthesized topological programs into the OSPF implementations (which are installed in the routers), routers will start to exchange topology information and build routing tables. Our goal is to compare the running results (e.g., the routers' status) of multiple equivalent topological programs. Any inconsistencies between the running results imply possible bugs in the OSPF implementations. Note that we often have to wait for a long time until the network converges, i.e., all routers have a complete view of the whole network and have successfully built the routing table. Comparing the intermediate results before network convergence is not meaningful. Thus, to improve the testing efficiency, we have to address the following challenge.

Challenge 3: How can we determine and accelerate the convergence of an OSPF network?

Since the OSPF specification does not provide a standard way to determine network convergence, we provide a few heuristic methods to determine and speed up network convergence by, for example, checking if all routers have successfully established connections to their neighbors. We provide more details in Section 4.

In the motivating example, after separately running the two equivalent topological programs generated in Step 2 and confirming the network converges according to the above rules, we observe that Router r_1 demonstrates inconsistent status — one indicates that Router r_1 is an ABR, while the other says it could be both ABR and ASBR. This inconsistency implies a bug in the OSPF implementation and helps us locate the buggy code in Figure 3.

4 Design: OSPF as an Example

This section formally details the three steps in ToDIFF: generating random yet valid networks (Section 4.1), synthesizing equivalent topological programs (Section 4.2), as well as determining network convergence and differentiating results (Section 4.3). Note that while this section uses OSPF as an example, the methodology is general for other IGPs.

4.1 Step 1: Network Generation

The first step of our approach is to generate a random yet valid network. This includes a physical network and configurations that adhere to Ethernet protocols and routing protocol specifications. A valid network ensures that routing functionalities can be correctly activated. Below, we use a graph model to define a valid network and detail the random network generation algorithm.

4.1.1 Valid Networks. Formally, we define a valid network running routing protocols as below.

Definition 4.1 (Network). A network is an undirected graph, $\mathbb{G} = (\mathbb{I}, \mathbb{L}, \mathbb{R})$, where $\mathbb{I} = \{ \text{eth}_0, \text{eth}_1, \ldots \}$ is a set of Ethernet interfaces and $\mathbb{L} \subseteq \mathbb{I} \times \mathbb{I}$ is a set of links between the interfaces. $\mathbb{R} = \{ \mathbf{r}_0, \mathbf{r}_1, \ldots \}$ denotes a set of routers in the network. A router $\mathbf{r}_i \in \mathbb{R}$ may contain multiple interfaces, denoted as $\nabla(\mathbf{r}_i) \subseteq \mathbb{I}$, and an interface can only belong to one router.

A valid network has to satisfy the following constraints, namely, unique IP address and subnet consistency, given that $ip(eth_i)$ and mask (eth_i) represent the IP address and subnet mask, respectively. **Constraint 1: Unique IP Address.** This constraint requires that each interface has an IP address different from all others in a network $\mathbb{G} = (\mathbb{I}, \mathbb{L}, \mathbb{R})$:

$$\forall$$
 eth_i, eth_i \in I : ip(eth_i) \neq ip(eth_i).

Constraint 2: Subnet Consistency. This constraint requires that the interfaces at the two ends of a link are in the same subnetwork:

$$\forall (\mathsf{eth}_i, \mathsf{eth}_j) \in \mathbb{L} : \mathsf{ip}(\mathsf{eth}_i) \& \mathsf{mask}(\mathsf{eth}_i) = \mathsf{ip}(\mathsf{eth}_j) \& \mathsf{mask}(\mathsf{eth}_j) \\ \land \mathsf{mask}(\mathsf{eth}_i) = \mathsf{mask}(\mathsf{eth}_j),$$

where & denotes the bitwise-and operation.

When applying a specific routing protocol like OSPF, an interface will have more attributes, including but not limited to the following:

- area(eth_i): the area the interface belongs to (recall that OSPF partitions a network into multiple areas).
- area_type(eth_i): the type of the area the interface belongs to (recall that an area could be a normal area, an NSSA area, etc.).
- hello(eth_i): the time interval at which a router sends a Hello message to its neighbor via the interface. Routers in an OSPF network send Hello messages periodically to their neighbors to maintain connectivity.
- dead(eth_i): the time interval at which a router considers the neighbor to be down. Routers in an OSPF network consider their neighbor is down after this time interval without receiving a hello message from their neighbor.
- cost(eth_i): the cost measures the resource consumption to send data via this interface, which is used for traffic optimization.

CCS '25, October 13-17, 2025, Taipei, Taiwan

Bing Shui, Yufan Zhou, Jielun Wu, Baowen Xu, and Qingkai Shi

At the same time, additional constraints (discussed below) should be satisfied to keep an OSPF network valid. Note that, to ease discussion, we present the most important constraints.

Constraint 3: Area Consistency. This constraint requires that interfaces at the two ends of a link are in the same area:

 $\forall (\mathsf{eth}_i, \mathsf{eth}_i) \in \mathbb{L} : \mathsf{area}(\mathsf{eth}_i) = \mathsf{area}(\mathsf{eth}_i).$

Constraint 4: Area Connectivity. This constraint requires that interfaces in the same area should be interconnected:

$$area(eth_0) = area(eth_n) \Rightarrow \exists (eth_0, eth_1, \dots, eth_n), r \in \mathbb{R} :$$

$$area(eth_i) = area(eth_0)$$

$$\land \quad (eth_i, eth_{i+1} \in \nabla(r) \lor (eth_i, eth_{i+1}) \in \mathbb{L}).$$

Constraint 5: Existence of Area 0. Recall that a valid OSPF network ensures the existence of a special area, namely area₀:

 $\exists \operatorname{eth}_i \in \mathbb{I} : \operatorname{area}(\operatorname{eth}_i) = \operatorname{area}_0 \land \operatorname{area}_type(\operatorname{eth}_i) = \operatorname{normal}.$

Constraint 6: Connecting to Area 0. Recall that a valid OSPF network requires all areas other than area₀ to connect to area₀ via at least one router:

$$\forall \operatorname{area}_k \neq \operatorname{area}_0, \exists r \in \mathbb{R}, \operatorname{eth}_i, \operatorname{eth}_j \in \nabla(r) :$$
$$\operatorname{area}(\operatorname{eth}_i) = \operatorname{area}_k \wedge \operatorname{area}(\operatorname{eth}_j) = \operatorname{area}_0.$$

Constraint 7: Time Interval Consistency. Given two connected interfaces in an OSPF network, the hello and dead intervals of the two interfaces must be the same:

 $\forall (\mathsf{eth}_i, \mathsf{eth}_i) \in \mathbb{L} : \mathsf{hello}(\mathsf{eth}_i) = \mathsf{hello}(\mathsf{eth}_i)$ $\wedge \operatorname{dead}(\operatorname{eth}_i) = \operatorname{dead}(\operatorname{eth}_i).$

4.1.2 Random Network Generation. Algorithm 1 shows our approach to generating a random yet valid OSPF network satisfying the seven constraints discussed before. Basically, the algorithm can be split into three parts (Line 2, Lines 3-6, and Lines 7-8). The first part (Line 2) is to generate random OSPF areas. Each area is created by invoking the procedure in Lines 10-18. Lines 11-12 created random links, each connecting two interfaces, as illustrated in Figure 7(a), where each square is an interface. Lines 13-16 merge interfaces to form a router (recall that a router consists of multiple interfaces), as illustrated in Figure 7(b). In this procedure, we should ensure that all routers are interconnected, satisfying Constraints 3 and 4. Line 17 names the area as $area_k$ and specifies the type of area, where k is an input integer ranging from 0 to a random positive integer. When k = 0, area₀, which must be a normal area, is created to satisfy Constraint 5.

The second part (Lines 3-6) connects the areas by invoking the procedure at Line 19, as exemplified in Figure 7(c). Particularly, Lines 3-4 ensure that all areas other than area₀ are connected to area₀, satisfying Constraint 6. Lines 5-6 randomly connect other areas. The third part (Lines 7-8) is straightforward, which is to assign random attributes to routers' interfaces, satisfying Constraints 1, 2, and 7. For instance, if we assign an IP address to an interface, the IP address cannot be assigned to others.

Algorithm	1:	OSPF	Network	Generation.
-----------	----	------	---------	-------------

- procedure gen_ospf_network() $\mathbb{G} \leftarrow \{\mathbb{G}_i = (\mathbb{I}_i, \mathbb{L}_i, \mathbb{R}_i) : \mathbb{G}_i \leftarrow \text{gen_ospf_area}(i) \text{ for } i = 0, 1, 2, \dots\};\$
- for each $\mathbb{G}_i \in \mathbb{G}$ and $i \neq 0$ do
- 3 $[con_ospf_area(\mathbb{G}_i, \mathbb{G}_0);$ 4
- 5
- for random times do 6
- con_ospf_area (\mathbb{G}_i , \mathbb{G}_j) where *i* ≠ *j*;
- for each interface eth in $\mathbb G$ do assign ip(eth), mask(eth), hello(eth), dead(eth) with respect to Constraints 1, 2, and 7:

```
return G;
```

10 procedure gen_ospf_area(k)

- $\mathbb{I} \leftarrow a \text{ random set of interfaces;}$ 11
- $\mathbb{L} \leftarrow$ a random subset of $\mathbb{I} \times \mathbb{I}$ such that links do not share interfaces: 12 $\mathbb{R} \leftarrow \emptyset$:
- 13
- for random times until there's a path between any pair of interfaces do 14 $\mathbb{I}' \leftarrow$ a random subset of \mathbb{I} such that all interfaces in \mathbb{I}' are from 15
 - different links and have not been assigned to a router;

16 $\mathbb{R} \leftarrow \mathbb{R} \cup \{r\}; \nabla(r) \leftarrow \mathbb{I}';$

 $\forall eth \in \mathbb{I} : area(eth) \leftarrow area_k; area_type(eth) \leftarrow normal or nssa;$ 17

- return $\mathbb{G} = (\mathbb{I}, \mathbb{L}, \mathbb{R});$ 18
- 19 **procedure** con_ospf_area($\mathbb{G}_i, \mathbb{G}_j$)
- randomly pick two routers from \mathbb{G}_i and \mathbb{G}_i , respectively; 20
- 21 merge the two routers into one;



Figure 7: Example to show the network generation algorithm.

4.1.3 Extensions. We provide a simplified network model above to simplify the explanations of our approach. In practice, it can be easily extended to include network devices, such as switches and hosts. Unlike a link connecting two interfaces, a switch can be regarded as a special link connecting more than two interfaces, forming a subnet. As such, Constraint 2 should be rewritten as $\forall eth_i, eth_i \in subnet$: $ip(eth_i) \& mask(eth_i) = ip(eth_i) \& mask(eth_i) \land mask(eth_i)$ = $mask(eth_i)$. In addition to switches, hosts can be viewed as special routers that do not run OSPF.

It is also easy to extend our network model to other IGPs, such as IS-IS. Section 5 reports our experimental results on both OSPF and IS-IS, the two most popular IGPs.

4.2 Step 2: Topological Program Synthesis

Given a random yet valid network generated in Step 1, Step 2 generates multiple equivalent topological programs to specify this network, via a novel semantics-guided program synthesis. This subsection consists of two parts, explaining the syntax and semantics of each command in a topological program (Section 4.2.1) and the program synthesis algorithm (Section 4.2.2), respectively.

Program Command	:= :=	Command+ Router Interface
Router RouterSubCmd	:= :=	router ospf no? RouterSubCmd area <num> nssa network <ip>/<mask> area <num></num></mask></ip></num>
Interface InterfaceSubCmd	:= :=	<pre>interface <r_i-eth_j> no? InterfaceSubCmd ip address <ip>/<mask> ip ospf area <num> ip ospf cost <num> ip ospf hello-interval <num> ip ospf dead-interval <num></num></num></num></num></mask></ip></r_i-eth_j></pre>
<ip>, <mask>, <num></num></mask></ip>	€	$\{ 0, 1, 2, \dots \} \cup \{ \top \}$

Figure 8: Syntax of a Topological Program.

4.2.1 Syntax and Semantics of a Topological Program. Each IGP implementation provides certain commands to control the network topology. Although the commands may differ, they follow a similar syntax. Figure 8 illustrates an abstract syntax of topological programs for OSPF.

Basically, a topological program is a list of commands as demonstrated by the first rule in the grammar. To ease the explanation, the syntax is simplified to include a subset of important commands: either a router command or an interface command. A router command can control all interfaces and other information of a router while an interface command controls a single interface. In these commands, <ip>, <num>, and <mask> are all integers with a special value \top meaning an undefined value.

A router command starts with the keywords "router ospf" or "router ospf no", followed by a sub-command. Assume <num> = k. A sub-command in the simplified syntax could update the type of an area, i.e., area_k, to a special type, namely NSSA, or specify that all interfaces, eth_i, satisfying ip(eth_i) & <mask> = <ip> & <math><mask>, belongs to area_k. If a router command contains the keyword "no", it performs an inverse operation of the sub-command. For instance, Line 9 in Figure 6 is a router command that cancels area₂'s role as an NSSA. The detailed semantics of a router command are listed in Table 1 (see ID 1-4).

An interface command configures the interface eth_j of the router r_i . The corresponding sub-commands sets the IP address, area, cost, hello interval, and dead interval, respectively. A typical example is illustrated in Figure 6, where Lines 1-2 set the IP address of the interface r_1 -eth₀. Similar to a router command, if an interface command contains the keyword "no", it performs an inverse operation of the sub-command. The detailed semantics of an interface command are listed in Table 1 (see ID 5-14).

In a topological program, the effects of a command c_i occurring before the other c_j could be covered by the effects of c_j . More complicated, the effects of a command could cover the partial effects of multiple commands. For instance, in Figure 9, Lines 2, 5, and 8 set the areas of the interfaces eth₀, eth₁, and eth₂ to area₁, area₀, and area₀, respectively. Line 10 affects all three interfaces — because all three IP addresses start with 177 — and changes the areas to area₂, area₂, and area₂, respectively. Line 11 only changes the areas of eth₁ and eth₂ to area₁ because their IP addresses start with 177.235.166 while the IP address of eth₀ does not. Finally, the areas of the three interfaces are area₂, area₁, and area₁, respectively. CCS '25, October 13-17, 2025, Taipei, Taiwan

<pre>1. interface r0-eth0 ip address 177.70.31.169/255.0.0.0 2. interface r0-eth0 ip ospf area 1 3</pre>	
4. interface r0-eth1 ip address 177.235.166.37/255.255.0.0	
5. interface r0-eth1 ip ospf area 0	
6.	
7. interface r0-eth2 ip address 177.235.166.38/255.255.0.0	
8. interface r0-eth2 ip ospf area 0	
9.	
10. router ospf network 177.0.0.0/255.0.0.0 area 2	
11. router ospf network 177.235.166.0/255.255.255.0 area 1	

Figure 9: Example to explain some complicated commands.

Al	Algorithm 2: Topological Program Synthesis for a Router.			
1 p	1 procedure synthesize_program($\psi \equiv \phi_1 \land \phi_2 \land \cdots \land \phi_n$)			
2	$\mathbb{U} \leftarrow \{\phi_1, \phi_2, \dots, \phi_n\};$			
3	$\mathbb{S} \leftarrow \emptyset;$			
4	$\mathbb{P} \leftarrow \text{empty list};$			
5	while U ≠ Ø do			
6	$c \leftarrow \text{select_command}(\mathbb{U}, \mathbb{S});$ /* Alg. 3 */			
7	add command c to the tail of \mathbb{P} ;			
8	return P;			

In practice, there are many complicated commands as above. Although the syntax used in this section is simplified to ease explanation, it covers sufficient and non-trivial cases we often encounter in practice. In the next part, we use the semantics of these commands to guide equivalent program synthesis.

4.2.2 Equivalent Program Synthesis. Given a valid network generated by Step 1, we now synthesize multiple equivalent topological programs for each router in the network. Each synthesized program should drive a router to the expected status in the network. The expected status of a router can be formulated as a conjunctive first-order logic formula, $\psi = \wedge_i \phi_i$, where each ϕ_i specifies the value of router interfaces' attributes. For instance, Step 1 may generate the network in Figure 1, where r_0 is expected to satisfy $\psi \equiv$ ip(eth₀) = 177.70.31.169 \wedge mask(eth₀) = 255.0.0.0 \wedge area(eth₀) = area₀ \wedge area type(eth₀) = normal \wedge ip(eth₁) = ... \wedge

As such, we send the target constraint ψ of a router to Algorithm 2 such that it synthesizes a random program, which can drive the router to the expected status denoted by ψ . Invoking Algorithm 2 multiple times yields multiple equivalent programs. Lines 2-3 of Algorithm 2 creates two sets, \mathbb{U} representing the target constraints we expect to satisfy but have not been satisfied, and \mathbb{S} is a set of constraints we expect to satisfy and have been satisfied. Clearly, the set \mathbb{U} contains all ϕ_i in the target constraint ψ . Line 4 initializes an empty command list to represent the synthesized topological program. Each iteration in the follow-up loop (Lines 5-7) produces one command until $\mathbb{U} = \emptyset$, meaning that all expected constraints have been satisfied.

Algorithm 3 presents the procedure of command selection invoked at Line 6 of Algorithm 2. This is the key part of our program synthesis, which guarantees the synthesis procedure can terminate by generating a program whose length is bounded.

Termination Guarantee. Line 2 of Algorithm 3 selects a random command, and Line 3 checks if this command is good enough via Algorithm 4. The criteria of a good command are discussed later. If

ID	Command c	Semantics [[c]]
1	router ospf area <num> nssa</num>	$\forall \text{eth}_i \in \nabla(\mathbf{r}_i), \text{area}(\text{eth}_i) = <\text{num}> : \text{area_type}(\text{eth}_i) = \text{nssa}$
2	router ospf no area <num> nssa</num>	$\forall eth_i \in \nabla(r_i), area(eth_i) = : area_type(eth_i) = normal (by default)$
3	router ospf network <ip>/<mask> area <num></num></mask></ip>	$\forall eth_i \in \nabla(r_i), ip(eth_i) \& mask(eth_i) = \langle ip \rangle \& \langle mask \rangle : area(eth_i) = \langle num \rangle$
4	router ospf no network <ip>/<mask> area <num></num></mask></ip>	$\forall eth_i \in \nabla(r_i), ip(eth_i) \& mask(eth_i) = \langle ip \rangle \& \langle mask \rangle : area(eth_i) = \top$
5	interface $\langle r_i - eth_j \rangle$ ip address $\langle ip \rangle / \langle mask \rangle$	$ip(eth_j) = \langle ip \rangle \land mask(eth_j) = \langle mask \rangle$
6	interface < r _{<i>i</i>} -eth _{<i>j</i>} > no ip address < i p>/ <mask></mask>	$ip(eth_j) = \top \land mask(eth_j) = \top$
7	interface <r<sub>i-eth_j> ip ospf area <num></num></r<sub>	$area(eth_j) = \langle num \rangle$
8	interface <r<sub>i-eth_j> no ip ospf area <num></num></r<sub>	$area(eth_j) = \top$
9	interface <r<sub>i-eth_j> ip ospf cost <num></num></r<sub>	$cost(eth_j) = $
10	interface <r<sub>i-eth_j> no ip ospf cost <num></num></r<sub>	$cost(eth_j) = 10$ (by default)
11	interface <r<sub>i-eth_j> ip ospf hello-interval <num></num></r<sub>	$hello(eth_j) = \langle num \rangle$
12	interface <r<sub>i-eth_j> no ip ospf hello-interval <num></num></r<sub>	$hello(eth_j) = 10$ (by default)
13	interface <r<sub>i-eth_j> ip ospf dead-interval <num></num></r<sub>	$dead(eth_j) = \langle num \rangle$
14	interface < r _{<i>i</i>} -eth _{<i>j</i>} > no ip ospf dead-interval < num >	$dead(eth_j) = 40$ (by default)

Table 1: Command Semantics for Router r_i.

Algorithm 3: Command Selection for Program Synthesis.

procedure select_command(U, S) 1 $c \leftarrow$ randomly generate a command; 2 **if** check_command(\mathbb{S} , c) = Good **then** 3 for each $\phi \in \mathbb{U}$ do 4 $| \quad \text{if } \llbracket c \rrbracket \Rightarrow \phi \text{ then } \mathbb{U} \leftarrow \mathbb{U} \setminus \{\phi\}; \mathbb{S} \leftarrow \mathbb{S} \cup \{\phi\};$ 5 for each $\phi \in \mathbb{S}$ do 6 $[if [\dot{c}] \land \phi \equiv false \text{ then } \mathbb{U} \leftarrow \mathbb{U} \cup \{\phi\}; \mathbb{S} \leftarrow \mathbb{S} \setminus \{\phi\};$ 7 return c; 8 else 9 $\phi \leftarrow \text{random constraint from } \mathbb{U};$ 10 11 $c \leftarrow$ generate one command to satisfy ϕ ; $\mathbb{U} \leftarrow \mathbb{U} \setminus \{\phi\}; \mathbb{S} \leftarrow \mathbb{S} \cup \{\phi\};$ 12 return c; 13

a command is good enough, we choose this command and update U and S at Line 4-8. In detail, Lines 4-5 move a constraint ϕ from U to S if ϕ is satisfied by the command, i.e., $[c] \Rightarrow \phi$. Recall that [c] denotes the constraint a command implies and is listed in Table 1. Similarly, Lines 6-7 move a constraint ϕ from \mathbb{S} to \mathbb{U} if ϕ is dissatisfied by the command, i.e., $[\![c]\!] \land \phi \equiv false$.

If a command is not good enough, Line 10 randomly selects a constraint from U. Lines 11-12 generate a command to satisfy this constraint and move it from \mathbb{U} to \mathbb{S} . Note that we can always generate a command to satisfy a given constraint. As an example, the second command in Figure 9 can always set the area of interface r_0 -eth₀ to area₁.

As for the criterion of a good enough command, Algorithm 4 provides a solution. In the algorithm, we check if the command *c* will dissatisfy a constraint ϕ in \mathbb{S} at Line 3. If so, the constraint ϕ will be moved from S to U, which is not an expected operation because the goal of our algorithm is to satisfy constraints in U (i.e., reduce the size of U). Thus, Line 3 also checks the other condition, i.e., $\#(\phi) > k$, where $\#(\phi)$ denotes the times of we move ϕ into \mathbb{U} and *k* is a predefined upper bound of $\#(\phi)$.

As such, the algorithm keeps a trend that constraints are moved from \mathbb{U} to \mathbb{S} . In other words, \mathbb{U} will gradually decrease in size until it becomes an empty set, jumping out of the loop in Algorithm 2 and terminating the algorithm.

Algorithm 4: Command Check for Program Synthesis.

procedure check_command(S, c)

for each $\phi \in \mathbb{S}$ do

if $[\![c]\!] \land \phi \equiv false and \#(\phi) > k$ then 3 **return** Bad; 4

return Good; 5

Example 4.2. Assume $\mathbb{U} = \{\phi_1, \phi_2\}$ and $\mathbb{S} = \{\}$ at the beginning of the program synthesis procedure, i.e., Algorithm 2. Our goal is to generate commands to satisfy ϕ_1 and ϕ_2 , thereby moving them from \mathbb{U} to \mathbb{S} . If we always generate random commands, e.g., c_1, c_2, c_3, \ldots , the constraint ϕ_1 could be satisfied by c_1, c_3, c_5, \ldots and dissatisfied by c_2, c_4, c_6, \ldots Consequently, ϕ_1 is constantly moved between the two sets, and there may be no chance for ϕ_2 to be moved out of \mathbb{U} .

In our algorithm, we limit the times, i.e., k, of moving ϕ from S to U. For instance, if k = 1, after generating at most three commands (one moving ϕ_1 from U to S, one moving from S to U, and one from U to S again), we cannot dissatisfy ϕ_1 any longer and, thus, have a chance to generate commands for ϕ_2 , thereby clearing the set \mathbb{U} and terminating the algorithm.

LEMMA 4.3. Algorithm 2 can always terminate.

Boundedness Guarantee. Recall the previous example and that k is a predefined constant limiting the times of moving a constraint ϕ from S to U. Thus, we generate at most 2k + 1 commands to move one constraint ϕ from U to S. That is to say, we generate at most $(2k + 1)|\mathbb{U}|$ commands to move all constraints from \mathbb{U} to \mathbb{S} , synthesizing a program of length at most $(2k + 1)|\mathbb{U}|$.

LEMMA 4.4. The length of a synthesized program is up to (2k +1) $|\mathbb{U}|$, where k is a predefined constant.

Putting the previous two lemmas together, we have the following theorem to conclude our program synthesis algorithm.

THEOREM 4.5. Algorithm 2 can always terminate by synthesizing a topological program whose length is up to $(2k + 1)|\mathbb{U}|$, where k is a predefined constant.

CCS '25, October 13-17, 2025, Taipei, Taiwan

4.3 Step 3: Differential Analysis

The third step is to input the synthesized topological programs into the routers, which then communicate with each other to exchange topological information until all routers compute complete routing tables. After the network converges, we compare the execution results of equivalent topological programs. Any inconsistency in the execution results indicates potential bugs.

4.3.1 Network Simulation. To facilitate the testing procedure, we do not use physical routers and networks. Instead, the routers, as well as the networks, are simulated by the open-sourced Mininet framework [27]. Mininet also allows us to install different routing protocol implementations into the simulated routers.

We then input the synthesized topological programs into the routers such that the protocol implementations can run the topological programs. The entire program is randomly split into several subsequences. At random intervals, a subsequence is entered until the entire program has been input. This is because a command may take some time to take effect. If commands are entered too quickly, earlier commands may be overwritten by later ones before they have had a chance to take effect.

4.3.2 Determining & Accelerating Network Convergence. After inputting the topological programs into the routers, it usually takes some time for the routers to exchange information and calculate routing data. We have to wait until this procedure is completed, i.e., the network converges – all commands in the topological programs have been executed, and all routers finish computing the routing table. Checking the intermediate status of the network is less meaningful for differential testing because a network may converge through different intermediate statuses.

Unfortunately, it is challenging to determine if a network converges because the specification of routing protocols, e.g., OSPF, does not explicitly define how to determine and speed up network convergence. To address this challenge, we use a few heuristic methods to estimate and accelerate convergence, as discussed below. The effectiveness of these methods is proven in our experiments, as discussed in Section 5.

Strategy 1: Neighbor Establishment. Like the TCP [14] handshake process, each router running OSPF maintains a state machine when it connects to another router. When the state machine reaches a final (a.k.a., terminated) state, the connection completes. Thus, to determine if a network converges, we check all routers' neighboring states to see if the states have been set to the final states.

Strategy 2: LSA Exchange. Recall that routers exchange LSAs to build routing tables. When a network converges, all LSA transmission queues in routers should be empty, meaning that a router no longer needs to exchange LSAs with others. Thus, we check if LSA transmission queues are empty to determine network convergence.

Strategy 3: Immediate Routing Calculation. By default, a router may calculate its routing table after some time to collect multiple LSAs from other routers. To speed up network convergence, we let all routers compute their routing table immediately after receiving an LSA from other routers.

Strategy 4: Time Interval Reduction. To establish a connection to neighboring routers and exchange LSAs, a router sends a packet,

Table 2: Differential Oracle in Five Groups.

Group	Fields	Description
1. Commands	-	Effective commands in a topological program
	IfAddress	IP address of the interface
	IfType	Interface type (e.g. broadcast, point-to-point, etc.)
Interfaces	IfArea	Area the interface belongs to
	IfCost	Cost of the interface
	NbState	State of a neighboring router
	NbPriority	Priority of the neighbor in the connection
Neighbors	NbRole	Role of a neighbor, e.g., ABR, ASBR
	NbLXRe	Length of neighbor's retransmission packet list
	LsaOption	Various options of a router
	LsaFlags	Router's role, e.g., ABR, ASBR
4. LSAs	LsaType	Types of the LSA
	LsaID	ID of the router sending the LSA
	RtDst	Destination address of the route
	RtCost	Cost to reach the destination address
5. Routes	RtNxtHop	Next router in the route to the destination
	RtIntf	The interface via which a packet should be sent out

e.g., a Hello/LSA packet, to its neighbors periodically, e.g., every n seconds. Since n > 5 by default in most cases, which is too long and slows down network convergence, we set n = 0.5 to speed up convergence as well as the whole differential testing procedure.

4.3.3 Differential Analysis. After running multiple equivalent topological programs, we collect the execution results from each router and compare the execution results derived from equivalent topological programs. In practice, there are a lot of execution results but not all of them can be used in differential analysis. For example, each router maintains the routing information as a database of LSAs and the age of each LSA (i.e., when an LSA is created). The ages of an LSA across different executions may differ, but the core contents of LSAs should be the same. Thus, in the differential analysis, we do not compare the LSA ages but other information that should be the same across different executions, which we refer to as the differential oracles.

Table 2 lists important differential oracles used in our approach, which can be put into five groups. First, a router in an OSPF network often outputs commands (from a topological program) that really take effect, i.e., the commands whose effects are not canceled by other commands. For instance, the command at Line 8 in Figure 6 is not included in the output commands because it is canceled by the command at Line 9. Thus, given two equivalent topological programs, the output effective commands should be the same. If comparing the final effective commands yields any inconsistency, it indicates some bugs in the protocol's implementation. Second, some information of an interface can play the role of differential oracles, including its IP address, type, area, cost, to name a few. Third, a router maintains the information for each of its neighbors. Thus, in the differential analysis, we can compare a router's neighbor information, including but not limited to the neighboring router's state, priority, type, etc. Fourth, recall that routers exchange LSAs (see Section 2) to compute a complete topology of the whole network. Each router records the LSAs it receives from other routers in a database. Thus, in the differential analysis, we can compare



Figure 10: Example to show the process of root cause analysis.

the LSAs recorded by routers, including LSA's sources, flags, types, etc. Last but not least, we compare the routing tables computed across multiple executions of equivalent topological programs. An entry in a routing table includes the destination a route, the cost of sending a packet to the destination, the next hop to send a packet, and many others.

4.3.4 Root Cause Analysis. Once ToDIFF detects a discrepancy between router fields (see Table 2) across equivalent topological programs, we perform a root cause analysis to identify bugs in the IGP implementation under test. Specifically, given two equivalent topological programs that lead to inconsistent router fields, we run both in parallel and perform single-step debugging using GDB [17], carefully monitoring their execution step by step until the internal program states begin to diverge. To narrow down the scope of debugging, we apply program slicing [54] to the inconsistent router fields reported by ToDIFF. This allows us to extract and track the relevant variables and control paths associated with the divergence, significantly reducing the amount of code that needs to be examined.

Example 4.6. Let us take the bug discussed in Section 3 as an example to illustrate the root cause analysis. Recall that we executed two equivalent topological programs in Figure 6 and observed a discrepancy in the LSA flag: one indicates that a router is both an ABR and an ASBR, while the other shows that it is only an ABR.

Since the struct field ospf->flags in the code captures the LSA flag of routers, we apply a program slicer to identify all statements and variables related to ospf->flags, including the function parameter redistribute and another struct field ospf->redistribute shown in Figure 10. When debugging, we track the values of these variables until their values diverge.

As illustrated in Figure 10, the LSA flag is set in the function asbr_status_update based on the value of ospf->redistribute. If this value is not zero, the ASBR bit in ospf->flags is set; otherwise, it is unset. When debugging, we find that the function asbr_status_update is invoked in different call sites: one topological program lets the function be invoked by init_ospf, where ospf->redistribute = 0; the other equivalent topological program lets the function be invoked by ospf_area_nssa_unset, where ospf->redistribute = 1. This difference explains the divergence in the LSA flag and identifies the root cause of the bug.

In our experience, the root cause analysis illustrated above typically allows us to identify root causes in 30 minutes. In the future, we plan to integrate existing automatic root cause analyzers, e.g., [58], into ToDIFF to fully automate the whole validation procedure.

4.4 Discussion

ToDIFF is fully automated except for a few manual efforts before and after the differential analysis. First, before applying ToDIFF, we must follow the official documents of IGPs to establish the constraints of an IGP network, as well as the semantics of IGP commands. Establishing these network constraints and command semantics allows us to build valid networks, thus avoiding the waste of resources exploring invalid networks and invalid topological programs in Steps 1 and 2. Note that the manual work above is a onetime effort and is common among modern automated techniques. For instance, automatic program analyzers, such as those for C/C++ programs, also require a one-time manual effort to model C/C++ semantics. Second, since root cause analysis is not the focus of this paper, we currently rely on a manual process to identify root causes after the differential analysis. In the future, we can replace it with automated root cause analyzers, e.g., [58].

5 Evaluation

We implement ToDIFF on top of Mininet [27], an open-source and widely-used network simulation framework. Mininet enables us to build a virtual network with multiple routers on a single server and run different IGP implementations. Rather than being a limitation, the use of a network simulated by Mininet brings two significant advantages. On the one hand, Mininet offers strong applicability as it creates realistic virtual networks by running real Linux kernel code, making it a widely adopted tool in both academia and industry for testing protocol implementations [8, 27]. As such, any approach that works in Mininet is also applicable to industry-grade routers and real-world deployment scenarios. Our method is therefore transferable to real environments without requiring fundamental modifications. On the other hand, the simulated network provides excellent scalability. Mininet supports the creation of large-scale virtual networks, which are often infeasible to construct or access in physical environments for testing purposes. This capability enables our method to be evaluated across a diverse range of topologies, demonstrating its scalability in testing complex protocol behaviors.

To show the efficacy of our approach, we conduct experiments with ToDIFF to address the following three research questions:

- RQ1. How efficient are the three steps of ToDIFF?
- **RQ2.** How effective is ToDIFF in detecting bugs, compared to state of the art?
- RQ3. What are the root causes of the discovered bugs?

Subjects. In the experiments, we use ToDIFF to test multiple implementations of OSPF and IS-IS from FRRouting [12]. We choose OSPF and IS-IS because they are the two most popular IGPs. We use implementations from FRRouting as it is an open-source routing protocol suite designed for Linux and Unix platforms and has been integrated into the software repositories of major Linux-based operating systems like Debian and CentOS. FRRouting is widely utilized in critical networking scenarios, including ISPs and SaaS infrastructure, and is trusted by technology giants such as NVIDIA and VMware. FRRouting has experienced robust growth in recent years, gradually becoming the standard implementation for popular routing protocols. To date, it has garnered 3.5K stars and 1.5K

Table 3: Details of OSPF and IS-IS Implementations.

IGP	Implementations	Size (KLoC)	# Files	# Commands
	frrouting-10.2	145	181	167
OSPF	frrouting-9.0	135	175	161
	frrouting-8.0	121	157	147
	frrouting-10.2	135	183	100
IS-IS	frrouting-9.0	126	177	100
	frrouting-8.0	108	158	74

forks on GitHub, with a continuous stream of contributions from hundreds of developers worldwide.

Particularly, the details of the OSPF and IS-IS implementations are listed in Table 3. We choose the three most recent major versions of the FRRouting project for testing. The OSPF and IS-IS implementations consist of hundreds of thousands of lines of code, with OSPF ranging from 121KLoC to 145KLoC and IS-IS from 108KLoC to 135KLoC. The project has a relatively complex structure, comprising hundreds of source code files, with OSPF involving 157 to 181 files and IS-IS from 158 to 183 files. The implementations support a rich set of commands to write topological programs and provide precise control over the protocols, with 147 to 167 kinds of commands for OSPF and 74 to 100 commands for IS-IS, respectively.

It is worth noting that although our evaluation is conducted on the open-source IGP implementation, FRRouting [12], the results are expected to generalize well to other routing protocol implementations. This is because, while we do not evaluate ToDiff on commercial implementations (e.g., Cisco, which are closed-source), these systems typically share similar architectural designs with FRRouting. For example, they use similar commands to configure and drive routers [7]. This architectural consistency ensures that ToDIFF can be effectively applied to both open-source and commercial implementations, and seamlessly extended to support these implementations with minimal adaptation effort.

Baselines. To show the effectiveness of ToDIFF, we compare ToD-IFF against existing validation techniques the FRRouting community is using. First, the FRRouting community actively handcrafts a rich test suite [10]. Currently, the test suite contains around 460 manually crafted test cases covering a wide range of network topologies. Second, we compare ToDIFF to OSS-Fuzz [19], which is a popular fuzzing framework developed by Google, having successfully detected 36,000 bugs across 1,000 projects, including the OSPF and IS-IS implementations in FRRouting.

Environment. All experiments were conducted on a server with the following configuration: 32 cores, 64 threads, 3.4 GHz CPU, and 256 GB of memory, running Ubuntu 20.04 and Mininet 2.3.0.

5.1 RQ1: Efficiency of ToDIFF

To answer RQ1, we measure the execution time of each step in our method and conduct a detailed analysis. The experiment is conducted under different scenarios where 1 to 15 routers are generated in a network. A network of 15 routers is sufficiently large to trigger deeply hidden bugs through our approach. We record the time taken by ToDIFF for its three main steps: network generation, topology synthesis, and differential IGP validation. The experiment is repeated ten times in different implementations, and the average time cost and number of generated commands are reported.





Figure 11: The time cost of the three steps and the number of generated commands for (a) OSPF and (b) IS-IS.



Figure 12: The time cost of Step 3 before and after applying our approach.

As plotted in Figure 11, (a) and (b) show the experiment results for OSPF and IS-IS, respectively. The X-axis denotes the number of routers in a generated network. The Y-axis, which is on a log scale, together with the bar chart, illustrates the time consumption of the three steps. The results demonstrate that our tool is highly efficient, with the total time cost of all three steps remaining under 3 minutes for both OSPF and IS-IS, with up to 15 routers. The testing time does not increase significantly with the network size, demonstrating the good scalability of our approach. For both OSPF and IS-IS, we complete Step 1, i.e., network generation, in less than one second, and Step 2, i.e., the synthesis of equivalent topological programs, within 2 seconds. Step 3, i.e., the network execution, can be completed in 100 seconds for OSPF and 170 seconds for IS-IS. The total testing time for IS-IS is longer than that of OSPF, because IS-IS inherently has a slower convergence speed than OSPF.

Since the time cost of Step 1 is negligible, we provide a detailed analysis for Step 2 and Step 3 below.

ID	IGP	Root Causes	ToDiff	FRRouting Test Suite	Google's OSS-Fuzz	Status
1	OSPF	Incorrect setting internal states about area structure	1	×	×	PR merged
2	OSPF	Incorrect setting internal states about nssa status	1	×	×	PR merged
3	OSPF	Incorrect setting internal states about area structure	1	×	×	PR merged
4	OSPF	Incorrect setting internal hello timer	1	×	×	PR merged
5	OSPF	Incorrect setting internal spf hold timer	1	×	×	PR merged
6	OSPF	Incorrect setting wait timer	1	×	×	Confirmed
7	OSPF	Incorrect setting dead multiplier timer	1	×	×	PR merged
8	OSPF	Incorrect parsing the commands of no VLINK	1	×	×	PR merged
9	OSPF	Incorrect parsing the commands of no AREANSSA	1	×	×	PR merged
10	OSPF	Incorrect parsing the commands of no AREANSSARANGE	1	×	×	PR merged
11	OSPF	Incorrect parsing the commands of no AREANSSACOST	1	×	×	PR merged
12	OSPF	Incorrect parsing the commands of no AREADEFAULTCOST	1	×	×	PR merged
13	OSPF	Incorrect parsing the commands of no OSPFWRITEMULTIPLIER	1	×	×	PR merged
14	OSPF	Incorrect parsing the commands of no DEFUALTINFORMATION	1	×	×	PR merged
15	OSPF	Incorrect parsing the commands of no DISTANCE	1	×	×	PR merged
16	OSPF	Incorrect parsing the commands of no AREARANGECOST	1	×	×	PR merged
17	OSPF	Incorrect parsing the commands of no AREASHORTCUT	1	×	×	PR merged
18	OSPF	Incorrect parsing the commands of no DEADINTERVAL	1	×	×	PR merged
19	OSPF	Incorrect parsing the commands of no ABRTYPE	1	×	×	PR merged
20	OSPF	Incorrect parsing the commands of no WRITEMULTIPLIER	\checkmark	×	×	PR merged
21	IS-IS	Incorrect setting interface's psnp-interval.	 Image: A second s	X	×	PR merged
22	IS-IS	Incorrect setting interface's csnp-inveral.	1	×	×	PR merged
23	IS-IS	Incorrect setting interface's priority.	1	×	×	PR merged
24	IS-IS	Incorrect setting interface's hello-inveral	1	×	×	PR merged
25	IS-IS	Incorrect setting interface's hello-multipllier	1	×	×	PR approved
26	IS-IS	Incorrect setting interface's circuit-type	1	×	×	PR approved

Table 4: Bugs Discovered by ToDIFF.

Detailed Analysis of Step 2. To show the scalability of our program synthesizer, we try to generate larger topological programs with up to 1,000 commands per router. For validation, the topological program comprising 1,000 commands is rich enough to cover a wide range of diverse network scenarios. The results show that this step takes less than 9 seconds to complete, attributable to our efficient program synthesis algorithm, which ensures good scalability for large-scale networks and complex topologies.

Detailed Analysis of Step 3. Step 3 is the most time-consuming step as we have to wait for a long time until the network converges before we can differentiate the results to find bugs. To reduce the waiting time, Section 4.3 presents a few heuristic methods to determine and speed up network convergence. In Figure 12, the dotted and blue line stands for the maximum time for a network to converge (the experiment uses our methods to determine network convergence but does not apply our method to accelerate convergence). We can observe that it takes up to about 175 seconds for a network to converge. It also means that if we do not provide any method discussed in Section 4.3 to determine and speed up network convergence, we have to always wait for at least 175 seconds for a network to converge, plotted by the solid and red line in the chart. With our strategies in Step 3, the average time waiting for a network to converge is significantly reduced by $1.6 \times$ to $4.7 \times$, as plotted by the dashed and green line in Figure 12.

5.2 RQ2: Effectiveness of ToDIFF

Code Coverage. We evaluate both the overall and topology-related code coverage. The former refers to the coverage across all source code of OSPF and ISIS. Meanwhile, since our approach focuses on identifying topology-related bugs, we also evaluate the topology-related coverage, which removes topology-irrelevant code.

Table 5: Overall / Topology-Related Code Coverage (%).

IGP	ToDiff	FRRouting Test Suite	Google's OSS-Fuzz
OSPF	62.3 / 85.1	51.8 / 32.2	15.7 / 1.1
ISIS	68.9 / 78.2	63.9 / 38.0	8.9 / 5.3

The coverage of ToDIFF is collected during our experiments, which involve 200 randomly-generated valid networks and 100 equivalent topological programs for each network. We run each topological program only once. For the FRRouting test suite, we executed all test cases related to the core functionality of OSPF and IS-IS. In the case of OSS-Fuzz, the campaign was run continuously for 24 hours. To mitigate randomness, we repeated each experiment 10 times and reported the average results.

The code coverage is listed in Table 5. ToDIFF consistently outperforms both the FRRouting test suite and OSS-Fuzz, achieving the highest overall and topology-related coverage. Unlike the FRRouting test suite, which depends on manually crafted test cases, ToDIFF automatically generates tests, enabling a more comprehensive exploration of network topologies. This automation results in a $2\times$ increase in topology-related coverage and is a key factor in ToDIFF 's ability to uncover 26 bugs that the FRRouting test suite fails to detect. On the other side, OSS-Fuzz predominantly generates invalid packets during mutation, which are often limited to exercising exception-handling paths and rarely reach the core logic of the protocol, resulting in significantly lower coverage.

The code not covered by ToDIFF is mainly related to user-defined protocol extensions, e.g., OSPF's opaque LSAs. As ToDIFF is designed to verify standard-compliant protocol functionality, such optional or implementation-specific features are intentionally excluded from its testing scope.

CCS '25, October 13-17, 2025, Taipei, Taiwan

1. DEFU	JN(ip_ospf_dead_interval, ip_ospf_dead_interval_cmd,
2	"ip ospf dead-interval minimal hello-multiplier (1-20)",
3.+	"ip ospf dead-interval minimal hello-multiplier (2-20)",
4.	•••••
	(a)
1. int 2.	<pre>ospf_vty_dead_interval_set(){</pre>
3.+ 4.+ 5.+	<pre>re = seconds - event_timer_remain_second(oi->t_wait); EVENT_OFF(oi->t_wait); if (re > 0) OSPF_ISM_TIMER_ON(oi->t_wait,);</pre>
6.	
	(b)

Figure 13: Case studies of discovered bugs.

Bug Detection Capability. To answer RQ2, we use ToDIFF to detect bugs in OSPF and IS-IS implementations listed in Table 3. As discussed before, the implementations are from high-quality open-source projects and have been frequently checked by mature bug detection tools. Nonetheless, ToDIFF can still detect 26 zero-day bugs listed in Table 4. Some of these bugs had even been hidden in the implementations for over 20 years, showing the high effectiveness of our approach. Meanwhile, many of the identified bugs are security-critical and exploitable. For example, the bug example discussed in Section 3 may redirect network traffic to the wrong routers, causing black holes, data leaks, and other severe consequences. As responsible researchers, we also upload patches to fix detected bugs. As shown in Table 3, "PR merged" means the patches have been merged into the code; "PR approved" means the patches have been approved by the developers and are pending merge; and "Confirmed" means developers confirm the validity of a bug report but are still working on fixing the bug. The links to these bug reports are publicly available at https://todiff.github.io/.

For bug detection, ToDIFF achieves 100% precision, i.e., none false positives are reported. This is attributed to the design that compares router attributes that must be consistent across equivalent topologies. Consequently, any observed discrepancy during differential testing signals the presence of a bug.

As shown in Table 4, the handcrafted test suite and OSS-Fuzz cannot detect any bugs we discover. On the one hand, the handcrafted test suite, albeit containing many test cases, cannot cover many cases that we can randomly generate. On the other hand, OSS-Fuzz fails to detect these bugs because it treats OSPF and IS-IS as common communication protocols and, thus, does not provide a special design to validate topology-related business logic in their implementations. These limitations highlight the advantages of our approach in effectively uncovering topology-related bugs.

As discussed in Section 1, although there are some automated approach to testing routing protocol implementations [25, 34, 35, 56], they are either too old and outdated or not publicly available, making them not directly comparable. As explained before, the key weaknesses of these existing methods are that they heavily depend on manual efforts to build testing oracles, whereas our approach applies differential testing to achieve fully automated validation.

5.3 RQ3: Root Cause of Discovered Bugs

We analyzed the root causes of all the bugs we found (listed in Table 4) and categorized them into two groups. We also provided case studies and discussed the potential impacts of these bugs, typically including network congestion and routing black holes. **Group 1: Incorrect Parsing of Commands (13/26).** This type of error occurs when a protocol implementation incorrectly parses certain topological commands. These errors can cause the protocol implementation to reject valid or accept invalid commands, leading to misinterpretation of topological commands and, ultimately, incorrect network configurations. As illustrated below, such misconfigurations may lead to severe network problems like congestion.

Figure 13(a) shows a bug discovered by our tool in the OSPF protocol, which has been hidden for over 20 years. The code defines the command: ip ospf dead interval minimal hello-multiplier <num>, which is used to set the hello and dead timers of an interface. The correct value of <num> should range from 2 to 20, but the buggy code allows 1 to 20. In consequence, if a network administrator uses the command to set the hello timer to 1 second (i.e., <num> = 1), the dead timer is automatically set to 1/<num> = 1 second, too. In this case, OSPF's hello timer and dead timer are set to the same value, causing OSPF to continuously restart the neighbor establishment process and send related packets. This behavior can result in network congestion along the affected path.

Group 2: Incorrect Action Logic for Routers (13/26). These errors occur when a protocol implementation contains wrong logic to set variables that control router states, timers, and so on. Such errors can lead to incorrect protocol behavior, such as executing unintended actions or sending malformed packets. These issues can result in severe network problems like routing black holes. The bug discussed in the motivating example belongs to this category.

Figure 13(b) shows the other bug in the OSPF implementations. The code snippet handles the action logic for the command: ip ospf dead interval <num>, which sets the wait timer and is responsible for timing the transition from the waiting state to the connected state during the connection among routers. If a topological program uses a very large value for <num>, such as 1000, and then sets it back to a reasonably small value, the buggy code (without lines 3-5) fails to update the timer. As a result, the neighbor state remains stuck in the waiting state, thus preventing two routers from establishing neighbor relationships. This bug may result in an incomplete network topology and cause routing black holes or other potential network problems.

6 Related Work

Routing protocols differ from common (communication, management, or security) protocols as routing protocols concern network topologies. In contrast, the others focus on the communication among multiple parties and regard the network topology between two parties as black boxes. As discussed in Section 1, most existing works, e.g., network protocol fuzzing [1, 29, 32, 37, 42, 44, 45] or network differential analysis [15, 16, 61] are not designed for routing protocols and cannot detect errors happening at the time of topology establishment. In what follows, we discuss related works regarding the validation of routing protocols, but, in a word, to the best of our knowledge, ToDIFF is the first work that utilizes program synthesis to enable differential testing for routing protocols.

Validating Routing Protocol Specifications. Unlike our work, which assumes the protocol specification is correct and validates if an implementation correctly follows the specification, many existing work does not validate the implementation of routing protocols but their specifications. They create abstract models according to the protocol specification and then perform model checking or verification to check if certain properties may be violated in the model [26, 33, 36, 55]. Wibling [55] leveraged linear temporal logic to verify the correctness of connection establishment and message broadcasting for the LUNAR protocol. Khayou and Sarakbi [26] used abstract algebraic to model the EIGRP protocol [43], verifying if its routing calculations converge and adhere to the shortest-path distance constraints. Maag et al. [33] modeled the MANET protocol [46] using an extended finite state machines [28], thereby verifying if the protocol correctly establishes network topology. Maxa et al. [36] designed a secure protocol, SUAP, and analyzed its security features via model checking. The techniques used in these works are fundamentally different from our approach and outside the scope of this paper because we aim to validate specific routing protocol implementations instead of the specifications.

Validating Routing Protocol Implementations. Similar to our work, there are many testing techniques checking if a routing protocol implementation follows the protocol's specification. Using handcrafted test suites [10, 11, 13] is the most straightforward way to test a protocol's implementation. However, these test suites face coverage limitations and cannot thoroughly test protocol correctness across diverse network topologies. Compared to handcrafted test suites, our work makes it possible to generate diverse topologies with automated testing oracles to automate the testing procedure.

Automated methods can generate diverse test cases and can be categorized as active or passive. Active methods [21-23, 25, 34, 35, 56] use models such as graphs or finite state machines to actively generate test packets, which are then sent to the protocols' implementations for evaluation. Hao et al. [21] propose a probabilistic algorithm that randomly inserts connections into a network and calculates the correct routing tables as an oracle. They then test implementations using the generated networks and compare the results with the oracle. This approach requires significant manual effort to compute the testing oracle – the correct routing table for each protocol. In contrast, our method applies differential testing, which not only provides automated testing oracles but also analyzes both intermediate router status and the final routing tables. Helmy et al. [22] and Kasemsuwan et al. [25] mutate a set of prepared network topologies to generate additional test cases. These methods generate relatively limited topologies and also suffer from the testing oracle problem. In contrast, our approach is not restricted by the prepared networks and the testing oracles. With a random network generation algorithm and differential testing, our approach can synthesize diverse topologies with automated testing oracles, enabling broader test coverage. Maag and Zaidi [34], Malik [35], Wu et al. [56] have to transform the implementation of routing protocols into an abstract model, which is often a complex and error-prone process. Such abstraction is unnecessary for ToDIFF, thereby reducing analysis complexity and enhancing applicability.

Passive methods [30, 31, 50, 59, 60] do not actively generate networks or routing protocol packets but record and analyze realworld network traffic to detect bugs in routing protocol implementations. These passive approaches differ from ToDIFF, which is an active method, in the following aspects. First, we actively generate diverse networks and, thus, enable more comprehensive test coverage, whereas the passive approaches are restricted by observed network traffic. Second, we actively synthesize equivalent topological programs to address the testing oracle problem, while the passive approaches still rely on manual efforts to build testing oracles. This manual process is both labor-intensive and error-prone, restricting existing methods to a single protocol or a limited subset of protocol functionalities. Our approach addresses the oracle problem via differential analysis.

7 Conclusion

This work presents ToDIFF, a differential testing technique to detect hidden bugs within IGP implementations. ToDIFF first generates random yet valid networks, applies a semantics-guided and bounded program synthesizer to generate equivalent topological programs, and differentiates outcome topologies. To date, ToDIFF has identified 26 bugs across common IGPs, all confirmed or fixed.

References

- Jinsheng Ba, Marcel Böhme, Zahra Mirzamomen, and Abhik Roychoudhury. 2022. Stateful greybox fuzzing. In Proceedings of the 31st USENIX Security Symposium (SEC '22). USENIX Association, 3255–3272. https://www.usenix.org/conference/ usenixsecurity22/presentation/ba
- [2] Davide Benetti, Massimo Merro, and Luca Viganò. 2010. Model Checking Ad Hoc Network Routing Protocols: ARAN vs. endairA. In 2010 8th IEEE International Conference on Software Engineering and Formal Methods (SEFM '10). IEEE, 191–202. doi:10.1109/SEFM.2010.24
- [3] Sahana Bhosale and Ravindra Joshi. 2008. Conformance testing of OSPF protocol. In 2008 IET International Conference on Wireless, Mobile and Multimedia Networks (WoWMoM '08). IET, 42–47. https://ieeexplore.ieee.org/document/4470071
- [4] Cristian Cadar, Daniel Dunbar, and Dawson Engler. 2008. KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs. In Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation (OSDI '08). USENIX Association, 209–224. https://dl.acm.org/ doi/10.5555/1855741.1855756
- [5] R. Callon. 1990. Use of OSI IS-IS for routing in TCP/IP and dual environments. https://datatracker.ietf.org/doc/html/rfc1195.
- [6] Chia Yuan Cho, Vijay D Silva, and Dawn Song. 2013. BLITZ: compositional bounded model checking for real-world programs. In 2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE '13). IEEE, 136–146. doi:10.1109/ASE.2013.6693074
- [7] Cisco. 2016. IP Routing: OSPF Configuration Guide. https://www.cisco.com/c/en/ us/td/docs/ios-xml/ios/iproute_ospf/configuration/xe-16/iro-xe-16-book/irocfg.html.
- [8] Rogério Leão Santos de Oliveira, Christiane Marie Schweitzer, Ailton Akira Shinoda, and Ligia Rodrigues Prete. 2014. Using Mininet for emulation and prototyping Software-Defined Networks. In 2014 IEEE Colombian Conference on Communications and Computing (COLCOM '14). IEEE, 1–6. doi:10.1109/ColComCon. 2014.6860404
- [9] F. de Renesse and A.H. Aghvami. 2004. Formal verification of ad-hoc routing protocols using SPIN model checker. In *Proceedings of the 12th IEEE Mediterranean Electrotechnical Conference (MELECON '04)*. IEEE, 1177–1182. doi:10.1109/ MELCON.2004.1348275
- [10] FRR Developers. 2024. FRR RFC/Compliance Test. https://github.com/FRRouting/ frr/wiki/RFC_Compliance_Results/OSPF_extended_results.pdf.
- [11] FRR Developers. 2024. FRR Topotests. https://github.com/FRRouting/frr/tree/ master/tests/topotests.
- [12] FRR Developers. 2024. FRRouting Project. https://frrouting.org/.
- [13] Munet Developers. 2024. Munet. https://github.com/LabNConsulting/munet.
- [14] Wesley Eddy. 2022. Transmission Control Protocol (TCP). https://datatracker. ietf.org/doc/html/rfc9293.
- [15] Tiago Ferreira, Harrison Brewton, Loris D'Antoni, and Alexandra Silva. 2021. Prognosis: closed-box analysis of network protocol implementations. In Proceedings of the 2021 ACM SIGCOMM 2021 Conference (SIGCOMM '21). ACM, 762–774. doi:10.1145/3452296.3472938
- [16] Paul Fiterău-Broştean, Ramon Janssen, and Frits Vaandrager. 2016. Combining model learning and model checking to analyze TCP implementations. In Computer Aided Verification: 28th International Conference (CAV '16). 454–471. doi:10.1007/ 978-3-319-41540-6 25
- [17] GNU. 2025. GDB: The GNU Project Debugger. https://www.sourceware.org/gdb/.

- [18] Patrice Godefroid, Michael Y. Levin, and David Molnar. 2012. SAGE: whitebox fuzzing for security testing. *Queue* 10, 1 (2012), 20–27. doi:10.1145/2090147. 2094081
- [19] Google. 2025. OSS-Fuzz. https://google.github.io/oss-fuzz/.
- [20] Istvan Haller, Asia Slowinska, Matthias Neugschwandtner, and Herbert Bos. 2013. Dowsing for overflows: a guided fuzzer to find buffer boundary violations. In Proceedings of the 22nd USENIX Conference on Security (SEC'13). USENIX Association, 49–64. https://dl.acm.org/doi/10.5555/2534766.2534772
- [21] Ruibing Hao, David Lee, Rakesh K. Sinha, and Dario Vlah. 2000. Testing IP routing protocols – from probabilistic algorithms to a software tool. In Formal Methods for Distributed System Development Joint International Conference on Formal Description Techniques for Distributed Systems and Communication Protocols (FORTE/PSTV '00). Springer US, 249–264. doi:10.1007/978-0-387-35533-7_16
- [22] A. Helmy and D. Estrin. 1998. Simulation-based 'STRESS' testing case study: a multicast routing protocol. In Proceedings. Sixth International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS '98). IEEE, 36–43. doi:10.1109/MASCOT.1998.693672
- [23] Ahmed Helmy, Deborah Estrin, and Sandeep K. S. Gupta. 1998. Fault-oriented test generation for multicast routing protocol design. In Proceedings of the FIP TC6 WG6.1 Joint International Conference on Formal Description Techniques for Distributed Systems and Communication Protocols and Protocol Specification, Testing and Verification (FORTE XI / PSTV XVIII '98). Kluwer, B.V., 93–109. doi:10.1007/978-0-387-35394-4_6
- [24] Heqing Huang, Peisen Yao, Rongxin Wu, Qingkai Shi, and Charles Zhang. 2020. Pangolin: incremental hybrid fuzzing with polyhedral path abstraction. In 2020 IEEE Symposium on Security and Privacy (SP '20). IEEE, 1613–1627. doi:10.1109/ SP40000.2020.00063
- [25] Poonyavee Kasemsuwan and Vasaka Visoottiviseth. 2017. OSV: OSPF vulnerability checking tool. In 2017 14th International Joint Conference on Computer Science and Software Engineering (JCSSE '17). IEEE, 1–6. doi:10.1109/JCSSE.2017.8025919
- [26] Hussein Khayou and Bakr Sarakbi. 2017. A validation model for non-lexical routing protocols. J. Netw. Comput. Appl. 98, C (2017), 58-64. doi:10.1016/j.jnca. 2017.09.006
- [27] Bob Lantz, Brandon Heller, and Nick McKeown. 2010. A network in a laptop: rapid prototyping for software-defined networks. In Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets '24). ACM, 1–6. doi:10. 1145/1868447.1868466
- [28] D. Lee and M. Yannakakis. 1996. Principles and methods of testing finite state machines – a survey. Proc. IEEE 84, 8 (1996), 1090–1123. doi:10.1109/5.533956
- [29] Junqiang Li, Senyi Li, Gang Sun, Ting Chen, and Hongfang Yu. 2022. SNPSFuzzer: A fast greybox fuzzer for stateful network protocols using snapshots. *IEEE Transactions on Information Forensics and Security* 17 (2022), 2673–2687. doi:10. 1109/TIFS.2022.3192991
- [30] J. Liu, Yougu Yuan, D.M. Nicol, R.S. Gray, C.C. Newport, D. Kotz, and L.F. Perrone. 2004. Simulation validation using direct execution of wireless ad-hoc routing protocols. In 18th Workshop on Parallel and Distributed Simulation (PADS '04). IEEE, 7–16. doi:10.1109/PADS.2004.1301280
- [31] Jason Liu, Yougu Yuan, David M Nicol, Robert S Gray, Calvin C Newport, David Kotz, and Luiz Felipe Perrone. 2005. Empirical validation of wireless models in simulations of ad hoc routing protocols. *Simulation* 81, 4 (2005), 307–323. doi:10.1177/0037549705055017
- [32] Zhengxiong Luo, Junze Yu, Feilong Zuo, Jianzhong Liu, Yu Jiang, Ting Chen, Abhik Roychoudhury, and Jiaguang Sun. 2023. Bleem: packet sequence oriented fuzzing for protocol implementations. In 32nd USENIX Security Symposium (SEC '23). USENIX Association, 4481–4498. https://dl.acm.org/doi/10.5555/3620237. 3620488
- [33] Stephane Maag, Cyril Grepet, and Ana Cavalli. 2008. A formal validation methodology for MANET routing protocols based on nodes' self similarity. *Comput. Commun.* 31, 4 (2008), 827–841. doi:10.1016/j.comcom.2007.10.031
- [34] Stéphane Maag and Fatiha Zaidi. 2006. Testing methodology for an ad hoc routing protocol. In Proceedings of the ACM International Workshop on Performance Monitoring, Measurement, and Evaluation of Heterogeneous Wireless and Wired Networks (PM2HW2N '06). ACM, 48–55. doi:10.1145/1163653.1163663
- [35] Saif Ur Rehman Malik. 2014. Using formal methods to validate the usage, protocols, and feasibility in large scale computing systems. Ph. D. Dissertation. North Dakota State University. https://core.ac.uk/download/pdf/211302484.pdf#page=123
- [36] Jean-Aimé Maxa, Mohamed Slim Ben Mahmoud, and Nicolas Larrieu. 2016. Extended verification of secure UAANET routing protocol. In 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC '16). IEEE, 1–16. doi:10.1109/DASC. 2016.7777970
- [37] Ruijie Meng, Martin Mirchev, Marcel Böhme, and Abhik Roychoudhury. 2024. Large language model guided protocol fuzzing. In *Proceedings of the 31st Annual Network and Distributed System Security Symposium (NDSS '24)*. IEEE. doi:10. 14722/ndss.2024.24556
- [38] John Moy. 1998. OSPF Version 2. https://datatracker.ietf.org/doc/html/rfc2328.
- [39] Madanlal Musuvathi and Dawson R. Engler. 2004. Model checking large network protocol implementations. In Proceedings of the 1st Conference on Symposium

on Networked Systems Design and Implementation - Volume 1 (NSDI'04). USENIX Association, 12. https://dl.acm.org/doi/10.5555/1251175.1251187

- [40] Madanlal Musuvathi, David Y. W. Park, Andy Chou, Dawson R. Engler, and David L. Dill. 2003. CMC: a pragmatic approach to model checking real code. SIGOPS Oper. Syst. Rev. 36, SI (2003), 75–88. doi:10.1145/844128.844136
- [41] Ameet Naik. 2018. Anatomy of a BGP Hijack on Amazon's Route 53 DNS Service. https://www.thousandeyes.com/blog/amazon-route-53-dns-and-bgp-hijack.
- [42] Roberto Natella. 2022. Stateafl: greybox fuzzing for stateful network servers. Empirical Software Engineering 27, 7 (2022), 191. doi:10.1007/s10664-022-10233-3
 [43] Ivan Pepelnjak. 1999. EIGRP network design solutions. Cisco press. https:
- //dl.acm.org/doi/10.5555/519477
 [44] Van-Thuan Pham, Marcel Böhme, and Abhik Roychoudhury. 2020. AFLNET: A greybox fuzzer for network protocols. In 2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST '20). IEEE, 460–465. doi:10. 1109/ICST46399.2020.00062
- [45] Shisong Qin, Fan Hu, Zheyu Ma, Bodong Zhao, Tingting Yin, and Chao Zhang. 2023. Nsfuzz: towards efficient and state-aware network service fuzzing. ACM Transactions on Software Engineering and Methodology 32, 6 (2023), 1–26. doi:10. 1145/3580598
- [46] E.M. Royer and Chai-Keong Toh. 1999. A review of current routing protocols for ad hoc mobile wireless networks. *IEEE Personal Communications* 6, 2 (1999), 46–55. doi:10.1109/98.760423
- [47] Qingkai Shi, Junyang Shao, Yapeng Ye, Mingwei Zheng, and Xiangyu Zhang. 2023. Lifting network protocol implementation to precise format specification with security applications. In Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23). ACM, 1287–1301. doi:10.1145/ 3576915.3616614
- [48] Qingkai Shi, Xiao Xiao, Rongxin Wu, Jinguo Zhou, Gang Fan, and Charles Zhang. 2018. Pinpoint: fast and precise sparse value flow analysis for million lines of code. In Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'18). ACM, 693–706. doi:10.1145/ 3192366.3192418
- [49] Qingkai Shi, Xiangzhe Xu, and Xiangyu Zhang. 2023. Extracting protocol format as state machine via controlled static loop analysis. In Proceedings of the 32nd USENIX Conference on Security Symposium (SEC '23). USENIX Association, 7019– 7036. doi:10.5555/3620237.3620630
- [50] Hasan Ural, Zhi Xu, and Fan Zhang. 2007. An improved approach to passive testing of FSM-based systems. In Second International Workshop on Automation of Software Test (AST '07). IEEE, 6–6. doi:10.1109/AST.2007.1
- [51] Nalin Wadhwa, Jui Pradhan, Atharv Sonwane, Surya Prakash Sahu, Nagarajan Natarajan, Aditya Kanade, Suresh Parthasarathy, and Sriram Rajamani. 2024. CORE: Resolving Code Quality Issues using LLMs. Proc. ACM Softw. Eng. 1, 36 (2024), 789 – 811. doi:10.1145/3643762
- [52] Chengpeng Wang, Wenyang Wang, Peisen Yao, Qingkai Shi, Jinguo Zhou, Xiao Xiao, and Charles Zhang. 2023. Anchor: Fast and Precise Value-flow Analysis for Containers via Memory Orientation. ACM Trans. Softw. Eng. Methodol. 32, 3 (2023), 1–39. doi:10.1145/3565800
- [53] Chengpeng Wang, Wuqi Zhang, Zian Su, Xiangzhe Xu, Xiaoheng Xie, and Xiangyu Zhang. 2024. LLMDFA: Analyzing Dataflow in Code with Large Language Models. In Advances in Neural Information Processing Systems (NEURIPS '24). Curran Associates, Inc., 131545–131574.
- [54] Mark Weiser. 1984. Program slicing. IEEE Transactions on Software Engineering 10, 4 (1984), 352–357. doi:10.1109/TSE.1984.5010248
- [55] Oskar Wibling. 2005. Ad hoc routing protocol validation. Ph.D. Dissertation. Uppsala University. https://uu.diva-portal.org/smash/get/diva2:117170/ FULLTEXT01.pdf
- [56] Jianping Wu, Zhongjie Li, and Xia Yin. 2003. Towards modeling and testing of IP routing protocols. In Proceedings of the 15th IFIP International Conference on Testing of Communicating Systems (TestCom'03). Springer-Verlag, 49–62. https: //dl.acm.org/doi/10.5555/1764575.1764582
- [57] Yichen Xie and Alex Aiken. 2005. Scalable error detection using boolean satisfiability. In Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'05). ACM, 351–363. doi:10.1145/1040305.1040334
- [58] Carter Yagemann, Simon P. Chung, Brendan Saltaformaggio, and Wenke Lee. 2021. Automated bug hunting with data-driven symbolic root cause analysis. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21). ACM, 320–336. doi:10.1145/3460120.3485363
- [59] Yixin Zhao, Jianping Ju, and Yin Xia. 2002. From active to passive—progress in testing Internet routing protocols. *Journal of Computer Science and Technology* 17, 3 (2002), 264–283. doi:10.1007/0-306-47003-9_7
- [60] Yixin Zhao, Xia Yin, and Jianping Wu. 2001. Online Test System, an application of passive testing in routing protocols test. In Proceedings of 9th International Conference on Networks (ICON'01). IEEE, 190–195. doi:10.1109/ICON.2001.962339
- [61] Mingwei Zheng, Qingkai Shi, Xuwei Liu, Xiangzhe Xu, Le Yu, Congyu Liu, Guannan Wei, and Xiangyu Zhang. 2024. ParDiff: practical static differential analysis of network protocol parsers. Proc. ACM Program. Lang. 8, OOPSLA1 (2024), 1208 – 1234. doi:10.1145/3649854