Journal Pre-proof

SERAPH: Towards secure and efficient multi-controller authentication with (t, n)-threshold signature in multi-domain SDWAN

Wendi Feng, Ke Liu, Shuo Sun, Bo Cheng, Wei Zhang

PII: S1084-8045(24)00097-3
DOI: https://doi.org/10.1016/j.jnca.2024.103920
Reference: YJNCA 103920

To appear in: *Journal of Network and Computer Applications*

Received date : 16 October 2023
Revised date : 25 March 2024
Accepted date : 16 May 2024

Please cite this article as: W. Feng, K. Liu, S. Sun et al., SERAPH: Towards secure and efficient multi-controller authentication with (t, n)-threshold signature in multi-domain SDWAN. *Journal of Network and Computer Applications* (2024), doi: https://doi.org/10.1016/j.jnca.2024.103920.

SERAPH: Towards Secure and Efficient Multi-Controller Authentication with (t, n)-Threshold Signature in Multi-domain SDWAN

Wendi Feng[a,b], Ke Liu[a], Shuo Sun[a], Bo Cheng[b], and Wei Zhang[a]

[a]*School of Computer Science, BISTU, 35 North 4th Ring Rd, 100101, Beijing, China*
[b]*State Key Laboratory of Networking and Switching Technology, BUPT, 10 Xitucheng RD, 100876, Beijing, China*

Abstract

The multi-controller scheme is widely adopted in Software-Defined Wide Area Networks (SDWANs), where a WAN is segmented into multiple domains, each controlled by one controller. These controllers communicate with each other in-band, necessitating authentication before exchanging control messages. However, relying solely on identification of a single node for authentication exposes the network to spoofing attacks, jeopardizing its security. To address this issue, we present SERAPH, an innovative (t, n)-threshold signature-based authentication scheme that verifies not only the node itself but also its "endorsement" nodes to establish its identity. We have investigated the best practice for defining the "endorsement" relationships concerning security and overheads, formulating the problem as an integer programming problem. We have demonstrated the polynomial-time hardness (NP-hardness) of the problem and proposed an efficient SERAPH algorithm. Through our rigorous simulation analysis, we show that SERAPH can provide comparative performance with Optimal and reduce time usage by over 90%.

1. Introduction

Software-defined Networking (SDN) [1] is a prevalent paradigm that extracts the *control plane* from the *data plane*, facilitating the control plane with a global view of the whole network. This new feature significantly benefits the network not only for finer-grained controlling and management but also for fast enforcement of network policies, making it easier and more efficient for traffic engineering for Wide Area Networks (WANs) engineering [2, 3, 4]. Thus, the SDN powered WAN becomes an *SDWAN*.

In large-scale commercial SDWAN networks, multiple controllers are employed for enhanced resiliency [5] and performance [6]. As a result, the network is divided into several domains, each under the control of one (logical) controller. By sharing control messages with other controllers, each controller has a comprehensive view of the entire network. Typically, these control messages are sent *in-band*, leveraging the underlying data plane network to relay them [3]. Hence, authentication is necessary before any controller communications to ensure network security.

The current authentication mechanisms for controllers primarily rely on the *Public Key Infrastructure* (PKI) approach, where a *Certificate Authority* (CA) is utilized to validate the identity of each controller. In most cases, SDWAN owners establish their own CA [7], but it can become a single point of vulnerability, as the entire network is compromised if the CA is hacked. Recent approaches have explored utilizing blockchain technology [8, 9], which creates a distributed database for storing authentication information. However, these approaches encounter challenges in terms of scalability and performance, as they require the participation of all controllers for authentication and can significantly impact the efficiency of network policy enforcement.

In this paper, we present SERAPH, a (t, n)-threshold signature-based approach. When authenticating a controller (*i.e.*, source node), SERAPH judiciously chooses few "endorsement" nodes to prove the authenticity of the source node. The intelligence behind is the (t, n)-threshold signature, where *siglets* (*i.e.*, signature pieces) of t nodes out of the total number of n nodes can be consolidated and verified [10]. However, employing all other nodes for "endorsement" can incur significant performance overhead, while using fewer can compromise the security as well. The question is: *how can we achieve both security and performance simultaneously?*

We define the challenge as the *"Endorsement" Mapping Problem* (EMP). Our observation is that

the "endorsement" relationship plays a crucial role in achieving the goals of controller authentication in a multi-domain SDWAN network. Consequently, we formulate the EMP problem by considering both authentication security and the performance overhead associated with "endorsement". Essentially, the problem aims to minimize the performance overhead while meeting specific security requirements. We mathematically formulate the EMP problem as an integer programming problem and prove its Non-polynomial hard (NP-hard) time complexity. To efficiently tackle the problem, we propose the SERAPH algorithm, which iteratively selects a node based on two criteria: the minimum communication overhead (distance) and the minimum probability of compromise. SERAPH repeats the process until the "endorsement" mapping satisfies the security requirement. Through extensive simulations, we demonstrate that SERAPH achieves comparable performance to the Optimal solution while significantly reducing the time required by over 90%.

To summarize, our contribution is three-fold:

- We identify the drawbacks of existing PKI-based point-to-point authentication schemes in multi-controller SDWAN and present the EMP problem and the SERAPH protocol.
- We mathematically formulate the EMP problem as an integer programming problem, prove its NP-hard time complexity, and propose an efficient heuristic algorithm called SERAPH.
- We conduct rigorous simulations to evaluate the performance of the SERAPH algorithm and prove it can achieve near-optimal performance with over 90% time-usage reduction.

The rest of the paper is structured as follows. Section 2 provides a concise overview of the essential background knowledge on SDWAN and the (t, n)-threshold signature mechanism, followed by a survey of related work. In Section 3, we illustrate the EMP problem with a concrete example and introduce our attack model. The detailed mechanism of the SERAPH protocol is described in Section 4, while in Section 5, we formally formulate the EMP problem. We further demonstrate the NP-hard time complexity of EMP and propose an efficient heuristic algorithm in Section 6. We conduct comprehensive simulations and analyze the simulation results in Section 7 to evaluate the performance of SERAPH, and we discuss SERAPH's various properties in Section 8. Finally, Section 9 concludes the paper.

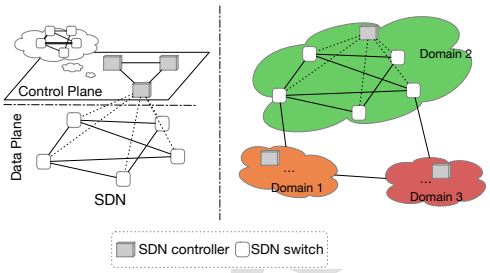

Figure 1: Two forms of multi-controller SDNs.

2. Background and Related Work

In this section, we brief the SDN and multi-controller SDN, followed by the drawbacks of existing controller authentication mechanisms.

2.1. SDN and Multi-Controller SDN

SDN, a revolutionary network paradigm, separates the *control plane* from the *data plane*. This approach enables the deployment of a (conceptually) centralized controller that provides a comprehensive overview of the network and directs the forwarding of packets in the data plane (consisting of controlled devices). Communication between the controller and SDN switches can be established through either *in-band* or *out-of-band* methods, utilizing SDN protocols such as OpenFlow [1]. In the in-band approach, control messages are transmitted over the links shared by the data plane. Conversely, the out-of-band method employs dedicated "controller – switch" links for communication.

Multi-Controller SDNs. In cases where the network comprises a large number of forwarding devices or when the data plane is distributed far apart geographically, a single controller may not be sufficient to handle the requests effectively. Thus, Multi-Controller SDNs are employed. In this approach, multiple controllers are utilized, where each controller is responsible for a subset of SDN switches located in a specific region and exchanges topology information. Despite the distribution of controllers across different locations, the control plane remains logically centralized, achieving increased resilience to failure and enhanced performance. The multi-controller scheme is not only applied in large SD-WAN networks by dividing the network into multiple domains controlled by individual controllers with in-band message transmission but also in smaller local SDN networks, as depicted in Figure 1. This

paper focuses on multi-domain SDNs.

2.2. Controller Authentications

According to the de-facto standard of SDN protocols (*i.e.*, OpenFlow [11]), controller authentication is not mandatory in SDN. The specification only recommends configuring control channels between controllers and switches using the SSL/TLS protocol without explicitly handling authentications over controllers (due to the multi-controller paradigm is not specified). Consequently, controller authentication is typically achieved through Public Key Infrastructure (PKI), where a Certificate Authority (CA) issues certificates for controllers to serve as their identification. The switch verifies the received certificate before establishing a secure channel with the controller. Existing SDN frameworks such as Orion [12], NOX [13], POX [14], and Ryu [15] employ this approach for controller authentication.

2.3. Distributed Threshold Signature

The threshold signature scheme [10] allows for many-to-one signing and verifying, whereby a group of participants can generate their own partial signatures (*i.e.*, siglets), and the verifier can consolidate the individual siglets and verify the *global signature*. The scheme requires the number of participants to exceed a certain *threshold*, hence the name "threshold signature". The process consists of five steps as follows: **i) Parameter negotiation:** The scheme first determines the value of t, n, two large primes p, q, and each participant's polynomials. These parameters can be configured by the operator. **ii) Key calculation:** Individual public keys pk_i and private keys qk_i are calculated for each participant using the negotiated parameters, and the individual public keys are distributed. After receiving the individual public keys, the global public key pk can be constructed. **iii) Siglet generation:** Each participant calculates the siglet using their private key. **iv) Global signature generation and verification:** The verifier, which can be any participant, receives t siglets and calculates (consolidates) the global signature with them. The verifier is then able to verify the global signature using pk. **v) Signature sharing:** The global signature is shared among the participants who are required to present the signature for various purposes.

2.4. Related Work

This subsection surveys relevant work, which can be classified into three categories, detailed as follows.

2.4.1. Point-to-point Authentication

Point-to-point authentication mechanisms rely solely on the authenticating and authenticated nodes (i.e., source and destination). This mechanism is commonly used in multi-controller SDN schemes. For example, DISCO [16] and SPARC [17] utilize the Simple Authentication Security Layer (SASL) [18], which leverages an authentication ID issued by the system operator to authenticate controllers. Similarly, Onix [19] and Orion [6] employ remote procedure calls (RPCs) between controllers, employing SSL/TLS for security. Authentication in these cases is achieved using certificates issued by a central authority (CA). CIDC [20] and Hyperflow [21] are also examples of this kind. However, the point-to-point scheme requires a central management agent, such as the CA, which can introduce security vulnerabilities and become a single point of compromise and failure, thus impairing overall security enhancements.

2.4.2. Blockchain-based Distributed Authentication

Blockchain-based schemes replace the centralized CA with a distributed blockchain [22]. The blockchain can be seen as a decentralized database where blocks are connected through cryptographic hash functions. Each block contains a replica of the database, including a cryptographic hash of the previous block, a timestamp, and transaction data (records). In blockchains, records are immutable, ensuring that the stored identifications cannot be forged. BlockCtrl [8] and BlockREV [9] are examples of authentication schemes that utilize blockchain. While these blockchain-based schemes address the issue of a single point of failure associated with CAs, attackers can still attempt to spoof authentication information by intercepting the traffic and potentially compromising the network's security. Moreover, Some blockchain consensus mechanisms like Proof-of-Work (PoW) rely on computational power. An attacker with more than 51% of the computational resources can "rule" the consensus and compromise the system. Therefore, additional measures should be implemented to protect against such attacks and maintain the integrity of the authentication process.

2.4.3. Threshold Signature-based Schemes

Da Silva *et al.* [23] suggest attesting the authenticity of controller states using a signature generated with a threshold encryption mechanism. FastHand [24] addresses the limitations of traditional PKI-based authentication and employs a threshold signature

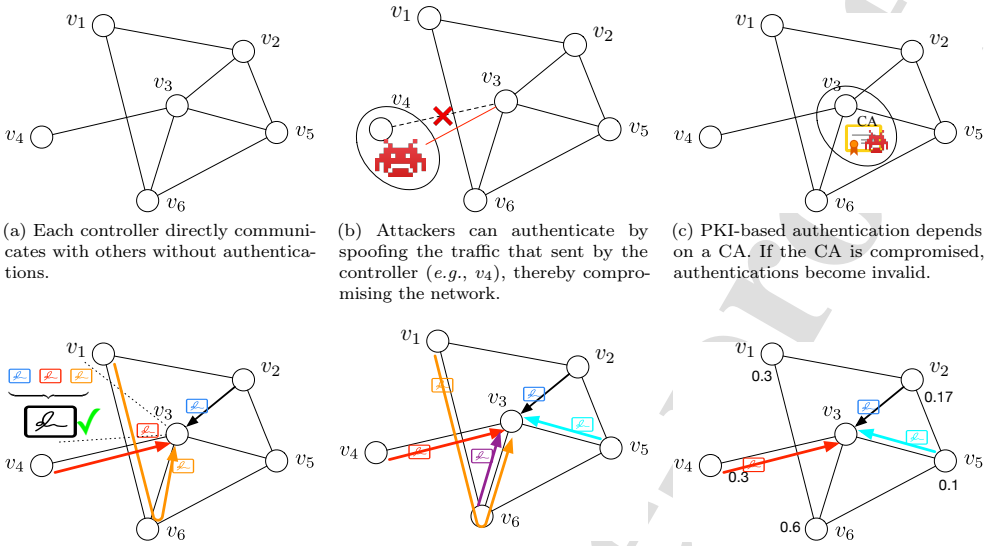

(a) Each controller directly communicates with others without authentications.

(b) Attackers can authenticate by spoofing the traffic that sent by the controller (*e.g.*, v_4), thereby compromising the network.

(c) PKI-based authentication depends on a CA. If the CA is compromised, authentications become invalid.

(d) (t, n)-threshold signature-based authentication enhances security with "endorsements" from others.

(e) Using all other nodes for "endorsement" is costly.

(f) Selecting suitable nodes for "endorsement".

Figure 2: Motivation examples. A multi-domain SDN network with six controllers, each node represents a controller for brevity.

mechanism for user equipment (UE) handover across 5G base stations. However, both schemes treat each node equally without considering best practices for selecting an appropriate number of suitable nodes to enhance security since proper consideration of node selection is essential to improve the overall security of the threshold signature mechanism (see Section 3.1).

3. Motivation and Attack Model

3.1. Motivation

In this subsection, we motivate the EMP problem through illustrative examples. To facilitate explanation, we employ a synthetic logical topology of an SDN network consisting of six nodes, and each represents controllers or domains with controllers in a multi-domain SDWAN setup (refer to Figure 2a). This simplified representation allows us to demonstrate the key concepts and challenges involved in achieving efficient controller authentication.

3.1.1. Needs for Controller Authentications

In the depicted multi-controller network with six nodes (as shown in Figure 2b), no security enhancement measures are applied. Consequently, if an end-host (*e.g.*a man-in-the-middle attacker) within the network or a domain decides to impersonate a controller, say v_2, the attacker can easily intercept the communication messages. By constructing packets to hijack the traffic, the attacker can then proceed to send misleading information to other controllers. As a result, the integrity of the network is compromised and susceptible to malicious activities.

3.1.2. PKI-based Authentication Is Vulnerable

PKI-based authentication is a commonly employed mechanism for authenticating network entities. It relies on a central authority (CA) to issue certificates for these entities. However, this centralized approach poses a single point of vulnerability. As shown in Figure 2c, when the CA is compromised, an attacker can exploit this situation to grant a legitimate identity to a device under their control. This compromised device can then be utilized to compromise the network, enabling potential unauthorized access and malicious activities.

3.1.3. (t, n)-Threshold Signature Brings New Hope

In the (t, n)-threshold signature authentication scheme, the authorization of an entity by other enti-

ties requires "endorsement" from $t-1$ entities, thus significantly enhancing security. This process is illustrated in Figure 2d, where v_2 intends to communicate with v_3. A (t, n)-threshold signature-based authentication is initiated. Firstly, v_2 sends a signature, generated using its private key, to v_3. Subsequently, v_3 requests signatures from two other nodes (*e.g.*, v_1 and v_4) to "endorse" the communication. By consolidating the three signatures (fromv_1, v_2, and v_4), v_3 creates a *global signature*. Upon successfully verifying this signature using its own public key, v_3 can determine that v_2 is a *trustworthy partner*, thereby allowing further communication to proceed with confidence.

3.1.4. Choosing the "Endorsements" Is Non-trivial

Nevertheless, it is important to note that the (t, n)-threshold signature-based authentication can have a dual impact. While it enhances security, it also introduces performance overhead that needs to be taken into consideration.

Larger number of "endorsements" improves security but also results in increased overhead. This can be observed in the example depicted in Figure 2e. The overhead introduced includes the time taken to obtain each siglet (e.g., half round-trip time (RTT) between v_2 and v_3, and RTT among v_1 and v_3, v_4 and v_3), the overhead for consolidating the siglet into a global signature, and the overhead of verifying the global signature. It is clear that the overhead is directly related to the value of t, with larger t values leading to higher overhead. Additionally, the selection of nodes for "endorsement" also affects performance, as nodes located farther away contribute to longer RTT, further impacting the overhead.

In addition, the probability of each controller being compromised can vary depending on factors such as the number of hosts and operators. It is worth noting that to control the communication, attackers would need to compromise all "endorsement" nodes. This highlights the significance of carefully **selecting the appropriate nodes** for "endorsement" in the authentication process as an essential task.

Hence, **judiciously selecting t and the nodes with the consideration of efficiency and security is important.** As demonstrated in Figure 2f, it jointly considers security and efficiency, facilitating the network with security at a minimum performance overhead.

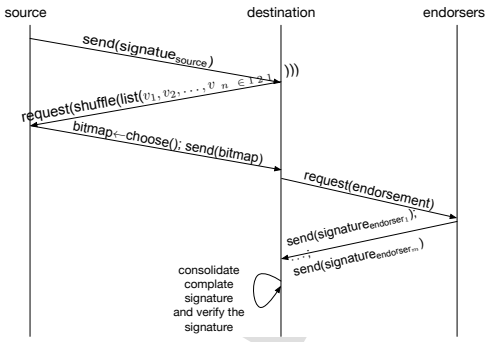

Figure 3: Protocol procedures and phases.

3.2. Attack Model

Our attack model is based on the traffic spoofing attack [25], where the attacker can intercept the traffic of a controller within the same domain. We make four assumptions in this attack model: **i)** Each attack is conducted by capturing the traffic only from the controller that resides in the same domain as the attacker's host. **ii)** The attacker is unaware of the total number of controllers in the network and which specific nodes are used for "endorsing" an authentication. **iii)** The attacker cannot compromise the host running the controller and thereby cannot gain access to the public/private keys used for generating and verifying signatures. **iv)** The attack can neither discover the bitmap structure (as described in Section 4) nor identify the "endorsing" nodes. While our attack model focuses on spoofing attacks, common authentication attacks like password guessing and dictionary attacks can be inherently defended because our system relies on a threshold signature scheme, not individual passwords. An attacker would need to compromise a minimum number of nodes (t) to forge a valid signature, making these traditional attacks ineffective.

4. The Seraph Authentication Protocol

This section outlines the Seraph protocol, taking into account the assumptions of our attack model as described in Section 3.2.

The initialization of the Seraph protocol is performed by the network operator, who configures the t value, "endorsement" mapping, and a polynomial required to establish the threshold signature scheme for authentication purposes, as described in Section 2.3. The "endorsement" mapping is derived

from the SERAPH algorithm, which is elaborated upon in Section 6.2.

Once the network begins, controllers exchange routing or control information among controller nodes. However, prior to this procedure, the SERAPH authentication protocol is essential for security measures. It consists of three phases: *authentication requesting*, *"endorsement" information passing*, and *joint signaturing*.

As depicted in Figure 3, when a controller (referred to as the "source") intends to transmit information to another entity (referred to as the "destination"), a series of steps should be followed: **i)** Firstly, the source sends its signature to the destination. **ii)** Upon receiving the signature, the destination acknowledges the authentication attempt and sends a shuffled list of node IDs back to the source. It is important to note that this list contains a greater number of nodes (ranging from 20% to 50%) than the actual topology to enhance security. **iii)** Subsequently, the source receives the shuffled list, selects the necessary "endorsement" nodes, and returns the sequence bitmap to the destination. In this bitmap, the i^{th} bit represents the usability of the i^{th} node in the received shuffled list. By not revealing the construction of the bitmap, the attacker is unable to obtain the required "endorsement" nodes. **iv)** The destination then receives the bitmap and requests signatures from the corresponding "endorsement" nodes. **v)** Finally, the destination consolidates the signatures obtained from the source and the "endorsement" into a single global signature, which is subsequently verified. If the verification process succeeds, the source is authenticated, and subsequent control messages can be transmitted.

5. Formulation

In this section, we begin by providing a mathematical description of the multi-controller network and introducing the metrics employed to evaluate its security and performance. Subsequently, we outline the constraints and objective functions for the EMP problem, formulating it as an optimization problem.

5.1. System Description

In this subsection, we present the mathematical formulation of the multi-controller network. For brevity, each domain is represented as a single node, and the detailed topology within each domain is omitted. The definitions for the notation used can be found in Table 1.

Table 1: Notation Definitions.

Nota.	Description
$\mathcal{V}$	Network node set. $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$.
$\mathcal{E}$	Edge set. $\mathcal{E} = \{e_1, e_2, \dots\}$.
n	Number of domains in the network.
T	Threshold set. $T = \{t_1, t_2, \dots, t_n\}$.
X	Matrix of the "endorsement" mapping. $X = \{x_{ij}\}, x_{ij} \in \{0, 1\}$.
$\mathcal{P}_i$	Set of (shortest) paths between node v_i and other nodes. $\mathcal{P}_i = \{\overrightarrow{p}_1, \overrightarrow{p}_2, \dots\}$.
$\mathbf{L}_i$	Set of lengths of paths in $\mathcal{P}_i$. $\mathbf{L}_i = \{\lvert\overrightarrow{p}_1\rvert, \lvert\overrightarrow{p}_2\rvert, \dots\} = \{l_1, l_2, \dots\}$.
$P_X(v_i)$	Probability of compromising the network under the "endorsement" mapping X.
p_i	Probability of compromising controller v_i.
$C_X^{v_{i_1}}(v_{i_2})$	Communication overhead for v_{i_1} authenticate at v_{i_2}.
$A_X^{v_{i_1}}(v_{i_2})$	Computational overhead for v_{i_1} authenticate at v_{i_2}.
α	Per-node signature consolidation time usage.
λ	Global signature verification time usage.

The network can be represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ is the set of network domains controllers. n represents the total number of domains, and v_i represents the i^{th} domain's controller[1] in the network. $\mathcal{E}$ is the links set. Each node in $\mathcal{V}$ can choose a number t as the minimum number of nodes (including "endorsement" nodes and the source) required for authenticating the source. We use $T = \{t_1, t_2, \dots, t_n\}$ as the set of threshold values for authenticating each node. Let X be the matrix of "endorsement" mapping, and $x_{ij} = 1$ represents authenticating node i requires node j for "endorsement". Let $\mathcal{P}_i = \{\overrightarrow{p}_1, \overrightarrow{p}_2, \dots\}$ be the set of (shortest) paths of v_i and other nodes, and the length of each path is represented as the cardinality of the path vector set $\lvert\overrightarrow{p}_k\rvert$. Hence, the length set of path set $\mathcal{P}_i$ is written as $\mathbf{L}_i = \{\lvert\overrightarrow{p}_1\rvert, \lvert\overrightarrow{p}_2\rvert, \dots\}$.

5.2. Metrics

5.2.1. Probability-based Security

Suppose an attacker can launch a spoofing attack by spoofing the controller of the same domain. To be authenticated, the attacker must obtain siglets of the "endorsement" nodes. However, since the construction of the bitmap (as mentioned in **Phase iv)**))

[1]We say i^{th} controller to represent the i^{th} domain's controller in the remainder of the paper.

cannot be recovered by the attacker, the probability of being compromised (*i.e.*, the probability that the attacker successfully passes the authentication) is composed of two main factors: the probability of successfully guessing the "endorsement" pattern and the probability that all "endorsement" nodes are compromised. Thus, it can be expressed as

$$P_X(v_{i_1}) = \frac{1}{C_{n-1}^{t_{i_1}} + \cdots + C_{n-1}^{n-1}} p_{v_{i_1}} \prod_{j=1}^{n} p_j^{x_{i_1 j}}, \quad (1)$$

where v_{i_1} is the source node, and $p_{v_{i_1}} \prod_{j=1}^{n} p_j^{x_{i_1 j}}$ indicates the probability of compromising the source node and all "endorsement" nodes.

5.2.2. Communication Overheads

As the threshold signature-based authentication necessitates additional assistance from the "endorsement" nodes, communication between the destination node and these nodes introduces authentication overheads. This can be expressed as the length of the corresponding point-to-point paths, which can be written as

$$C_X^{v_{i_1}}(v_{i_2}) = \mathbf{L}_{i_2} \cdot \overrightarrow{x_{i_1}}, \quad (2)$$

where $\mathbf{L}_{i_2}$ is the set of shortest lengths of paths between v_{i_2} and other nodes, and $C_X^{v_{i_1}}(v_{i_2})$ is the total length of the paths required for "endorsing" v_{i_1}'s authentication.

5.2.3. Signature Computational Overheads

When the destination node receives siglets from the source node and the "endorsement" nodes, it needs to consolidate these siglets into a single global signature and then verify it. The time required for verifying the global signature is constant, while the time taken to consolidate the siglets is proportional to the number of "endorsement" nodes. Thus, the overhead can be formulated as

$$A_X^{v_{i_1}}(v_{i_2}) = \alpha \sum_{j=1}^{n} x_{v_{i_1} j} + \lambda, \quad (3)$$

where α and λ are the constants that represent the calculation overhead of a per-node signature consolidation and signature verification, respectively.

5.3. Contraints

This subsection introduces the constraints considered in this work.

5.3.1. "Endorsement" Constraint

Our scheme utilizes threshold signatures to enhance security, and the authentication of each node necessitates "endorsement" from at least one additional node. Therefore, the endorsement nodes should neither be the source nor the destination. This can be expressed as follows

$$x_{i_1 i_1} = 0, \quad (4)$$

where $v_{i_1} \in \mathcal{V}$ is the source node.

5.3.2. Threshold Signature Constraint

Furthermore, the total number of nodes involved in "endorsing" the authentication should not exceed the number of nodes in the network. This can be formulated as

$$t_i \leq \sum_{j=1}^{n} x_{ij} \leq n - 2, t_i \geq 2, \quad (5)$$

where $n - 2$ represents the authenticating nodes (*i.e.*, source and destination nodes) should be omitted, and t_i is the threshold in the (t, n) threshold signature.

5.3.3. Overall Security Constraint

Each network may support different services or applications that require varying levels of security. For instance, a network supporting military industrial control is considered more security-sensitive than a home network [26]. Although application-level security requirements may vary, the fundamental security requirements for the physical WAN's domains rarely change. Hence, we utilize the level of security to assess whether a setup is suitable for safeguarding the system. Therefore, we can write

$$\log\left(p_i \prod_{j=1}^{n} p_j^{x_{ij}}\right) \leq \ell_{sec} \log\left(p_i \prod_{j=1}^{n} p_j\right), \quad (6)$$

where ℓ_{sec} is the level of security assigned by the network operator, and the level of compromising probability across nodes should not exceeds ℓ_{sec}. $p_i \prod_{j=1}^{n} p_j$ denotes that the possibility of compromising all nodes in the network when all are chosen as "endorsement" nodes. To diminish the impact of the size of different topologies, we apply logarithm calculations to the compromising probabilities.

5.4. Objective Function

The EMP problem focuses on securing authentication (security) while minimizing performance overhead (efficiency). As each network operator can determine the necessary security level based on network usage, our problem transforms into finding the authentication "endorsement" mapping that attains the minimum performance overhead for a given security level. Hence, we can represent the objective function as

$$\begin{aligned} obj &= \sum_{v_{i_1}\in\mathcal{V}} \sum_{v_{i_2}\in\mathcal{V}} \left(C_X^{v_{i_1}}(v_{i_2}) + A_X^{v_{i_1}}(v_{i_2})\right) \\ &= \sum_{v_{i_1}\in\mathcal{V}} \sum_{v_{i_2}\in\mathcal{V}} \left(\mathbf{L}_{i_2}\cdot\overrightarrow{x_{i_1}} + \alpha\sum_{j=1}^{n} x_{v_{i_1}j} + \lambda\right) \\ &= \sum_{v_{i_1}\in\mathcal{V}} \sum_{v_{i_2}\in\mathcal{V}} \sum_{j=1}^{n} x_{i_1 j}\left(l_{i_2 j} + \alpha\right) + (n-1)^2\lambda. \end{aligned} \tag{7}$$

5.5. Problem Formulation

As mentioned in the previous subsection, the EMP problem concerns finding a secure endorsement scheme X that attains minimum performance overhead (obj). This problem can be formulated as

$$\begin{aligned} \min_{x}\ & \sum_{v_{i_1}\in\mathcal{V}} \sum_{v_{i_2}\in\mathcal{V}} \sum_{j=1}^{n} x_{i_1 j}\left(l_{i_2 j} + \alpha\right) + (n-1)^2\lambda \\ \text{s.t.}\ & (4)(5)(6), \\ & x_{ij} \in \{0,1\}, \\ & \forall i, j, i_1, i_2 \in \mathcal{V}, i_2 \neq j, i_1 \neq i_2, \end{aligned} \tag{P}$$

where $\{x_{ij}\}$ are the designed binary integer variables, and $\mathbf{L}_i$ is the calculated path length vector. Symbols α, λ, and ℓ are constants indicating the overhead and security levels, respectively. Since the variable $\{x_{ij}\}$ is binary, Problem (P) is an *interger programming* problem.

6. Solution

In this section, we first prove the NP-hard complexity of the EMP problem and then present an efficient heuristic algorithm – SERAPH.

6.1. Complexity Analysis

This subsection proves the NP-hardness by reducing a special case of the EMP problem to the *traveling salesman problem* (TSP).

Theorem 1. *For a special case where (1) each node requires exactly one node to "endorse" the authentication; (2) each node only "endorses" the authentication of one other node; and (3) the attacker cannot guess the "endorsement" nodes, the EMP problem is NP-hard.*

Proof 1. *We first introduce the traveling salesman problem (TSP). The TSP problem involves studying the problem of finding the optimal route for a salesman to traverse all the locations exactly once. A typical formulation of the TSP problem is as follows:*

$$\begin{aligned} \min_{x}\ & \sum_{i}\sum_{j} c_{ij}x_{ij} \\ s.t.\ & \sum_{j} x_{ij} = 1, \sum_{i} x_{ij} = 1 \\ & \sum_{i}\sum_{j} x_{ij} \leq |S| - 1, \\ & \forall S \subset V, 2 \leq |S| \leq n-2, \end{aligned} \tag{8}$$

where S is the set of all tours of the graph G. $\{x_{ij}\}$ indicates if the salesman visits city j immediately after city i, $x_{ij} = 1$; and 0 otherwise. It has been proven that the TSP problem is NP-hard [27].

We then prove under the aforementioned special case, Problem P and the TSP problem are equivalent. Given the spacial case in Theorem 1, the objective function can be derived as $\sum_i \sum_j l'_{ij}x_{ij}$, and Constraint (6) is eliminated (given Condition (3)). It is noted that Constraint (5) can be derived to $n \leq \sum_{i=1}^{n}\sum_{j=1}^{n} x_{ij} \leq n(n-2)$. Let $S' = \sum_{i=1}^{n}\sum_{j=1}^{n} x_{ij} - (n-2)$. Hence, the constraint can be written as $2 \leq S' \leq (n-1)(n-2)$. Consequently, Problem P can be reformulated as

$$\begin{aligned} \min_{x}\ & \sum_{i=1}^{n}\sum_{j=1}^{n} l'_{ij}x_{ij} \\ s.t.\ & \sum_{i=1}^{n} x_{ij} = 1, \sum_{j=1}^{n} x_{ij} = 1, \\ & 2 \leq S' \leq (n-1)(n-2), \end{aligned} \tag{P'}$$

where it is easy to prove that $(n-1)(n-2) \geq n-2$. Hence, $2 \leq S' \leq (n-1)(n-2)$ has a larger range. Problem (P') aims to minimize the overall transmission performance overhead. We can consider the "endorsement" mapping x_{ij} as the traveling sequence in the TSP problem and treat S' as the tour set S in TSP. We can prove that the solution with the minimum performance overhead exists if and only if

Algorithm 1: The SERAPH algorithm.

Input: $G = (V, E)$: The topology;
Input: $\{p_i\}$: Compromise probability vector.
Input: ℓ_{sec}: Configured security level.
Output: X: "Endorsement" mapping.

```
X = {0, ..., 0}; Calculate the node distance matrix D with the Dijkstra algorithm [28];
Derive list of vectors D', where the D'_i is vector of nodes sorted by distance between v_i in ascending order;
Sort each vector in D' by {p_i} in ascending order as D'';
for i ∈ [1, n] do
    X' ← X;
    for v_j ∈ D''_i do
        X'_ij = 1;
        if P_X'(v_i) < ℓ_sec then
            X ← X';
            break;
return X
```

the TSP problem has an optimal solution that uses the minimum cost to travel to all cities. The problem construction can be derived in polynomial time, but finding the optimal solution for Problem (P') is NP-hard due to the NP-hardness of the TSP problem.

Since Problem (P') is a special case of the TSP problem, we can conclude that

Theorem 2. *The "endorsement" mapping problem problem is NP-hard.*

6.2. The SERAPH Algorithm

The complexity of the EMP problem arises from the variations in the compromise probability of each node. Due to its NP-hard complexity, we propose an efficient heuristic algorithm to solve it.

As demonstrated in Algorithm 1, the intelligence behind Seraph lies in utilizing the *most secure and closest* nodes for authentication "endorsement". The SERAPH approach is a three-step iterative process described as follows.

(1) Initialization. SERAPH first computes the shortest paths for each pair of nodes, generating a list of "endorsement" candidates for each node. The list is then sorted in ascending order based on distance and in ascending order based on the probability of compromise.

(2) Security-first "endorsements". After initialization, SERAPH iteratively assigns nodes from the sorted candidate list to "endorse" the authentication and updates the endorsement mapping x_i. This ensures that the selected nodes are both the closest and the most secure.

(3) Constraint comparing. Finally, SERAPH calculates the compromise probability of the node (*i.e.*, $P_X(v_i)$) based on the "endorsement" mapping x_i. If the security level of x_i aligns with the network operator's configuration, SERAPH proceeds to the next node. Otherwise, node v_i requires additional "endorsement" nodes to be assigned from the candidate list.

7. Simulation

This section describes the simulation of SERAPH. We first present the simulation setup and then introduce the compared algorithms. Finally, we exhibit the simulation results and analyze them. **In summary, the simulation results indicate that SERAPH can not only achieve comparable performance to Optimal but also scale well with larger topologies.**

7.1. Simulation Setup

We utilize topologies from the Topology Zoo [29] in our simulations. It includes 262 real-world backbone network topologies provided in `gml` file format. Our simulation is implemented in Python, utilizing the popular `python-igraph` [30] graph library for reading `gml` files and performing shortest path calculations. In the simulation, each node in a topology represents a domain with one deployed controller. The compromise probability for each controller (node) is randomly generated within the range of [0.1, 0.9] following a uniform distribution. The setting is based on the statistics for various vulnerability databases [31, 32]. The configured security level of the whole network is set from 0.1 to 1.

7.2. Compared Algorithms

We compare three algorithms in the simulation described as follows.

- Single: this is the lower bound algorithm, where no "endorsement" authentications are required.
- Optimal: this is the optimal solution to the "endorsement" mapping problem problem.
- SecFirst: this solution aims at achieving the best security and is the upper bound.
- LowOverhead: this solution minimizes the "endorsement" overhead without considering security.
- SERAPH: this solution is detailed in Section 6.2.

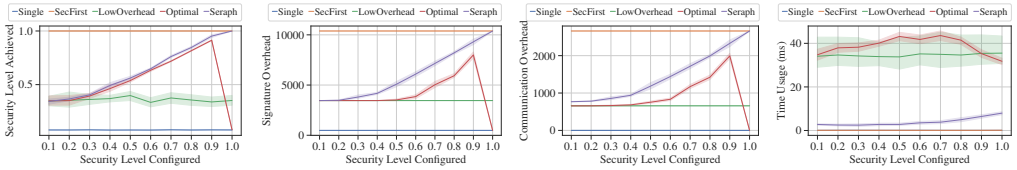

(a) Security ratio. The higher, the better.
(b) Signature overhead. The lower, the better.
(c) Communication overhead. The lower, the better.
(d) Time usage. The lower, the better.

Figure 4: Performance of the compared algorithms on the Abilene topology.

7.3. Simulation Results and Analysis

We compare the performance of the mentioned algorithms, evaluating *security ratios*, *signature overhead*, *communication overhead*, and *time usage*. Our results show that SERAPH achieves near-optimal performance on small topologies and scales efficiently to larger topologies with stricter security requirements, indicating significantly faster execution times for SERAPH.

7.3.1. Performance Comparison on Abilene

We utilize the Abilene topology [33], which consists of 11 nodes and 14 links. Abilene, the former version of Internet2 [34], is extensively as a case studied in the networking community. Besides, due to the time-consuming nature of signature calculation and verification, which takes milliseconds to derive results, we preconfigure α (calculation overhead) and λ (signature verification overhead) to 9 ms and 4 ms, respectively. We believe the performance setting is sound as manifested in [35, 36]. Finally, we conduct 10 simulations for each run to minimize errors.

(1) Security. We employ a probability-based approach that calculates the ratio of the compromised probability of a specific solution X to the ratio of that in the most secure setting (*i.e.*, the SecFirst algorithm). This can be derived from Equation 6 and expressed as follows:

$$sec = \frac{\log\left(p_i \prod_{j=1}^{n} p_j^{x_{ij}}\right)}{\log\left(p_i \prod_{j=1}^{n} p_j\right)}, \tag{9}$$

where *sec* is the ratio, and we utilize logarithm to mitigate the impact of topology sizes (*i.e.*, the number of nodes in a topology). As shown in Figure 4a, Single only achieves an average of 0.079 compared to the most secure scheme, failing to meet all security requirements. LowOverhead fails to satisfy 60% of the cases as it does not consider any security factors. Additionally, the variation (see the error bar) in LowOverhead is larger than that of other algorithms due to the larger solution space it considers. SecFirst serves as our baseline in this evaluation, representing the most secure scheme. Despite incurring significant overheads (see Section 7.3.1**(2)**), SecFirst satisfies the security requirements even when they are not explicitly demanded. Furthermore, we observe that Optimal drops to 0 when the security requirement is set to 1.0. This is because Optimal fails to generate a feasible solution for the EMP problem. On the other hand, SERAPH can still fulfill the requirement thanks to its iterative approach. **Overall, SERAPH demonstrates comparable performance to Optimal while satisfying all security requirements.**

(2) Overheads.

We present the signature verification and communication overheads of the compared algorithms in Figures 4b-4c. SecFirst exhibits the highest overhead for both signature verification and communication due to its utilization of all nodes (excluding the source and destination) for "endorsement". This is anticipated since it pursues maximum security. In contrast, Single achieves the lowest overhead, with a 0 communication overhead as it solely authenticates with the source node itself. LowOverhead delivers a steady low-overhead performance in both signature verification and communication when compared to Optimal and SERAPH. This is attributable to the relaxation of the security constraint in LowOverhead. Compared to Optimal, SERAPH incurs up to 22.6% more overhead for signature verification and as high as 22.9% more overhead for communication, contributing less than 2 seconds for authentication. We consider the overhead to be acceptable for SD-WAN since the BGP protocol (a distributed protocol for multi-domain WAN) updates its routing information

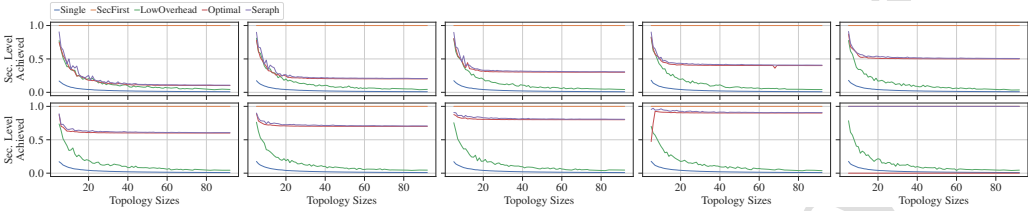

Figure 5: Security performance of different topologies under different security requirements. The higher, the better.

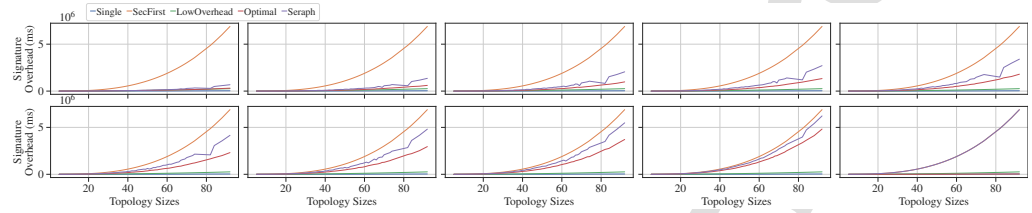

Figure 6: Signature verification overhead of different topologies under different security requirements. The lower, the better.

every 120 seconds [37]. Furthermore, as Optimal fails to generate a feasible solution for the 1.0 security requirement, the overhead drops to 0 in both figures.

(3) Time usage. We present the execution time comparisons of the algorithms. Single and SecFirst do not utilize any "endorsement" nodes, resulting in a fixed zero matrix solution and a 0 time usage. LowOverhead disregards constraints, enabling it to execute in a shorter duration but exhibiting greater variation due to the larger solution space mentioned in Section 7.3.1**(1)**. On the other hand, the Optimal algorithm achieves better security performance at the expense of longer time usage. SERAPH significantly reduces time expenses by up to 93.6% while attaining near-optimal performance on security and overhead metrics.

7.3.2. Performance Across Various Topology Sizes

To demonstrate the scalability of SERAPH, we conducted simulations across multiple topologies ranging from 5 nodes to 92 nodes from Topology Zoo. We used the same α and λ values as mentioned in Section 7.3.1. For each simulation, we ran it 10 times to reduce errors, with average values used to show performance. **In summary, SERAPH achieves comparable security performance to Optimal while attaining scalable to stricter constraints.**

(1) Security. We utilize the same security metric as shown in Equation 9. Figure 5 demonstrates the achieved security level under different topology sizes and security requirements (each displayed in each subplot). We observe that Single, SecFirst, and LowOverhead remain unchanged across different subplots as they are not affected by varying security requirements. Optimal and SERAPH exhibit comparable performance and can fulfill most security requirements (ranging from 0.1 to 0.8). However, under security requirements of 0.9 and 1.0, Optimal encounters difficulties in deriving feasible solutions (*e.g.*, 5-node topology under the 0.9 security requirement and all topologies under the 1.0 security requirement), whereas **SERAPH successfully generates corresponding solutions that satisfy all requirements.** Furthermore, we observe "L"-shape plots when the security requirement is below 0.9. This phenomenon arises from the fact that, with small topology sizes, the limited number of nodes leads to a significant increase or decrease in security with the "endorsement" of just one more or one less node. Additionally, we notice a peculiar occurrence where Optimal achieves a value close to 0.5 for the 5-node topology under the 0.9 security requirement, even though the value should either be 0 (indicating no feasible solution) or above 0.9. This discrepancy is due to averaging, where 5 out of 10 simulation rounds fail to derive feasible solutions.

(2) Overheads. As demonstrated in Figure 6 and Figure 7, Single once again achieves the mini-

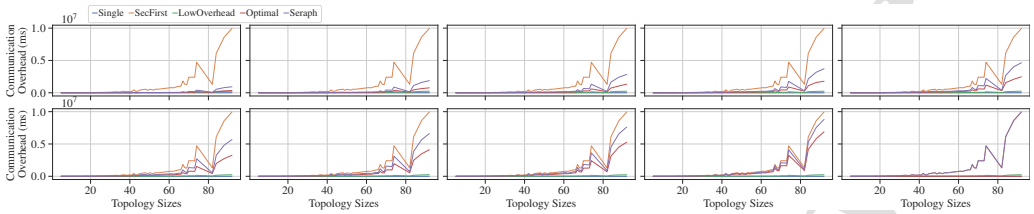

Figure 7: Communication overhead of different topologies under different security requirements. The lower, the better.

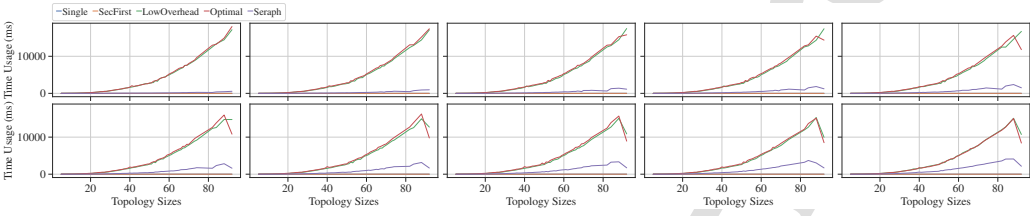

Figure 8: Time usage of different topologies under different security requirements. The lower, the better.

mum overhead across all requirements and topologies. The disparity between Optimal and SERAPH becomes more pronounced with increasing topology size. Generally, SERAPH contributes up to 23.3% more signature verification overhead and up to 21.8% more communication overhead. Moreover, communication overhead experiences irregularities in the Ulaknet topology [38] with 82 nodes. This can be attributed to Ulaknet's star-like topology, where three main nodes are interconnected while all other nodes connect to these three nodes. Turns out, the communication overhead is significantly reduced due to the maximum path length being limited to four.

(3) Time usage. As observed in Figure 8, LowOverhead and Optimal, which solve the integer programming problem, exhibit comparable time usage across all security requirement cases and topologies. Furthermore, an exponential-like increase can be observed as the size of the topologies increases. Conversely, SERAPH achieves linear performance improvement as the number of nodes increases, thanks to its iterative design. Additionally, it can reduce time usage by approximately 93% while achieving comparable security performance to Optimal.

8. Discussion

This section discusses miscellaneous properties of SERAPH. These properties exhibit SERAPH is practical, resilient, and performant in its applicable scenarios.

8.1. Practical Deployment

The SERAPH protocol is a security enhancement for the control plane of the SDWAN. Hence, it can best facilitate the *programmability* introduced by the SDN mechanism to implement the protocol as an application running on the control plane. This eliminates the need for hardware modifications or changes to the control plane operating system. Therefore, Deploying SERAPH becomes as straightforward as installing a new software program, significantly reducing deployment complexity and costs.

8.2. Resilient to Changes

SERAPH is designed for multi-domain SDWANs, where the network is logically segmented into domains, each managed by a controller. Within a domain, changes to the network topology or capacity are expected. However, the topology between domains, and by extension, the controllers managing those domains, tend to remain stable. Since SERAPH focuses on authentication between controllers in these interconnected domains, frequent topological changes within domains have minimal impact on SERAPH. In the rare instance where the overall network topology undergoes a significant shift, necessitating changes to the topology between domains, SERAPH can be adapted through software reconfiguration. Hence, SERAPH is resilient to changes.

8.3. Generality of the Dataset

We leverage the Topology Zoo topological dataset in our simulations to evaluate the performance of the SERAPH algorithm. We believe Topology Zoo is general, and the reason is three-fold. **i) Diverse network types.** Topology Zoo includes a variety of real-world topologies, encompassing star (*e.g.*, Padi), ring (*e.g.*, Viatel), tree (*e.g.*, AMRES), mesh (*e.g.*, GlobalCenter), and daisy chain (*e.g.*, Chinanet) structures. This diversity provides a strong foundation for evaluating SERAPH's adaptability to different network environments. **ii) Real-World relevance.** Unlike synthetic datasets, Topology Zoo provides a total number of 262 real-world network topologies. This ensures that the performance evaluation reflects SERAPH's effectiveness in realistic scenarios. **iii) Popularity.** The popularity of Topology Zoo is evident with over 1550 citations for the JSAC paper [29] on Google Scholar. This widespread use within the network research community strengthens the dataset's credibility as a benchmark for our simulations. In conclusion, Topology Zoo's diversity, real-world focus, and established reputation make it an ideal choice for evaluating Seraph's performance across a broad range of network conditions.

8.4. Applicable Scenarios

Our simulations demonstrate SERAPH's efficiency in multi-domain SDWAN environments. However, it is important to acknowledge that applications with extremely strict latency requirements might benefit from alternative authentication approaches due to the overhead. To this end, we plan to optimize the verification and solution derivation processes for the Seraph protocol and algorithm as part of our future work. This optimization aims at minimizing overhead while maintaining security, making SERAPH a compelling solution for a broader range of scenarios.

9. Conclusion and Future Work

This paper identifies the potential risk of spoofing attacks in the PKI-based multi-controller authentication schemes widely used in SDWAN. To enhance authentication security, we present an approach based on (t, n)-threshold signatures, which requires "endorsement" nodes to authenticate each node. We investigate the optimal method for setting the "endorsement" mapping, introduce the EMP problem, and formulate it as an integer programming problem. Additionally, we prove the NP-hardness of EMP and present the heuristic algorithm SERAPH, which efficiently solves it. Simulation results indicate that SERAPH achieves near-optimal performance with a 90% reduction in time usage. We hope our proposed (t, n)-threshold signature-based authentication paradigm will inspire the community to establish a more secure SDWAN network.

We envision two key areas for SERAPH's future development. Firstly, we will conduct more comprehensive simulations using diverse compromise probability distributions. This will strengthen our evaluation of SERAPH's robustness under various real-world vulnerability scenarios. Secondly, we plan to optimize the SERAPH protocol and algorithm to reduce overhead, including verification calculation overhead, communication overhead between nodes, and solution derivation time. This optimization will broaden the range of applications where SERAPH remains an efficient solution.

Acknowledgement

This work is supported in part by R&D Program of Beijing Municipal Education Commission under grant KM202311232005, National Key R&D Program of China under grant 2022YFC3320903, and Open Foundation of State key Laboratory of Networking and Switching Technology (Beijing University of Posts and Telecommunications) under grant SKLNST-2023-1-01. We acknowledge the anonymous reviewers for their insightful comments, which significantly improve the quality of the paper.

References

[1] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, J. Turner, Openflow: enabling innovation in campus networks, SIGCOMM Comput. Commun. Rev. 38 (2) (2008) 69–74.

[2] C.-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu, R. Alimi, C. Bhagat, S. Jain, J. Kaimal, S. Liang, K. Mendelev, et al., B4 and after: managing hierarchy, partitioning, and asymmetry for availability and scale in google's software-defined wan, in: Proceedings of the ACM SIGOMM 2018 Conference, 2018, pp. 74–87.

[3] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, R. Wattenhofer, Achieving high utilization with software-driven wan, in: Proceedings of the ACM SIGCOMM 2013 Conference, 2013, pp. 15–26.

[4] L. Poutievski, O. Mashayekhi, J. Ong, A. Singh, M. Tariq, R. Wang, J. Zhang, V. Beauregard, P. Conner, S. Gribble, et al., Jupiter evolving: transforming google's datacenter network via optical circuit switches and software-defined networking, in: Proceedings of the ACM SIGCOMM 2022 Conference, 2022, pp. 66–85.

[5] Z. Guo, S. Dou, S. Liu, W. Feng, W. Jiang, Y. Xu, Z.-L. Zhang, Maintaining control resiliency and flow programmability in software-defined wans during controller failures, IEEE/ACM Transactions on Networking 30 (3) (2022) 969–984.

[6] A. D. Ferguson, S. Gribble, C.-Y. Hong, C. Killian, W. Mohsin, H. Muehe, J. Ong, L. Poutievski, A. Singh, L. Vicisano, et al., Orion: Google's software-defined networking control plane, in: 18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21), 2021, pp. 83–98.

[7] B. Osborn, J. McWilliams, B. Beyer, M. Saltonstall, Beyondcorp: Design to deployment at google, The USENIX ;login: Journal 41 (2016) 28–34.

[8] H. Yang, Y. Liang, Q. Yao, S. Guo, A. Yu, J. Zhang, Blockchain-based secure distributed control for software defined optical networking, China Communications 16 (6) (2019) 42–54.

[9] P. Li, S. Guo, J. Wu, Q. Zhao, Blockrev: Blockchain-enabled multi-controller rule enforcement verification in sdn, Security and Communication Networks 2022 (2022).

[10] A. Shamir, How to share a secret, Communications of the ACM 22 (11) (1979) 612–613.

[11] Open Networking Foundation, Openflow switch specification, `https://opennetworking.org/wp-content/uploads/2014/10/openflow-switch-v1.5.1.pdf` (2022).

[12] Y. Fu, J. Bi, K. Gao, Z. Chen, J. Wu, B. Hao, Orion: A hybrid hierarchical control plane of software-defined networking for large-scale networks, in: 2014 IEEE 22nd International Conference on Network Protocols, IEEE, 2014, pp. 569–576.

[13] N. Gude, T. Koponen, J. Pettit, B. Pfaff, M. Casado, N. McKeown, S. Shenker, Nox: towards an operating system for networks, SIGCOMM Comput. Commun. Rev. 38 (3) (2008) 105–110.

[14] S. Kaur, J. Singh, N. S. Ghumman, Network programmability using pox controller, in: ICCCS International conference on communication, computing & systems, IEEE, Vol. 138, sn, 2014, p. 70.

[15] Ryu SDN Framework Community, Ryu sdn framework, `https://ryu-sdn.org` (2023).

[16] K. Phemius, M. Bouet, J. Leguay, Disco: Distributed multi-domain sdn controllers, in: 2014 IEEE network operations and management symposium (NOMS), IEEE, 2014, pp. 1–4.

[17] M. Li, X. Wang, H. Tong, T. Liu, Y. Tian, Sparc: Towards a scalable distributed control plane architecture for protocol-oblivious sdn networks, in: 2019 28th International Conference on Computer Communication and Networks (ICCCN), IEEE, 2019, pp. 1–9.

[18] J. Myers, Simple authentication and security layer (sasl), IETF RFC (1997).

[19] T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, et al., Onix: A distributed control platform for large-scale production networks, in: 9th USENIX Symposium on Operating Systems Design and Implementation (OSDI 10), 2010.

[20] F. Benamrane, R. Benaini, et al., An east-west interface for distributed sdn control plane: Implementation and evaluation, Computers & Electrical Engineering 57 (2017) 162–175.

[21] A. Tootoonchian, Y. Ganjali, Hyperflow: A distributed control plane for openflow, in: Proceedings of the 2010 internet network management conference on Research on enterprise networking, Vol. 3, 2010, pp. 10–5555.

[22] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, H. Wang, Blockchain challenges and opportunities: A survey, International Journal of Web and Grid Services 14 (4) (2018) 352–375.

[23] J. P. da Silva, E. Alchieri, J. Bordim, L. Costa, A secure and distributed control plane for software defined networks, in: International Conference on Advanced Information Networking and Applications, Springer, 2020, pp. 994–1006.

[24] P. K. Roy, P. Sahu, A. Bhattacharya, Fasthand: A fast handover authentication protocol for densely deployed small-cell networks, Journal of Network and Computer Applications 205 (2022) 103435.

[25] J. K. Tugnait, Pilot spoofing attack detection and countermeasure, IEEE Transactions on Communications 66 (5) (2018) 2093–2106.

[26] W. Feng, C. Liu, B. Cheng, J. Chen, Z. Wan, An end-host-importance-aware secure service-enabled hybrid sdn deployment, IEEE Transactions on Network and Service Management 20 (2) (2023) 2056–2070. `doi:10.1109/TNSM.2022.3208695`.

[27] M. Jünger, G. Reinelt, G. Rinaldi, The traveling salesman problem, Handbooks in operations research and management science 7 (1995) 225–330.

[28] D. B. Johnson, A note on dijkstra's shortest path algorithm, Journal of the ACM (JACM) 20 (3) (1973) 385–388.

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

[30] G. Csardi, T. Nepusz, et al., The igraph software package for complex network research, InterJournal, Complex Systems 1695 (5) (2006) 1–9.

[31] NATIONAL VULNERABILITY DATABASE, Nvd–nvd dashboard, `https://nvd.nist.gov/general/nvd-dashboard` (2024).

[32] Chinese National Vulnerability Database, Cnnvd vulnerability reports, `https://www.cnnvd.org.cn/home/report` (2024).

[33] Wikipedia, Abilene network, `https://en.wikipedia.org/wiki/Abilene_Network` (2023).

[34] Internet2, About internet2, `https://internet2.edu/community/about-us/` (2023).

[35] C. Stathakopoulou, C. Cachin, Threshold signatures for blockchain systems, Tech. rep., Swiss Federal Institute of Technology (2017).

[36] H. Yu, H. Wang, Elliptic curve threshold signature scheme for blockchain, Journal of Information Security and Applications 70 (2022) 103345.

[37] C. Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T. C. Schmidt, M. Wahlisch, Bgp beacons, network tomography, and bayesian computation to locate route flap damping, in: Proceedings of the ACM Internet Measurement Conference, 2020, pp. 492–505.

[38] Topology Zoo, The ulaknet topology, `http://www.topology-zoo.org/maps/Ulaknet.jpg` (2023).

1. We identify the drawbacks of existing PKI- based point-to-point authentication schemes in multi-controller SDWAN and present a (t,n)-threshold signature-based authentication scheme.
2. We identify the Endorsement" Mapping Problem (EMP) problem in the proposed scheme, formulate the problem as an integer programming problem, prove its NP-hard time complexity, and propose an efficient heuristic algorithm called SERAPH.
3. We conduct rigorous simulation to evaluate the performance of the SERAPH algorithm and prove it can achieve near-optimal performance with over 90% time-usage reduction.

Wendi Feng is an associate professor in the School of Computer Science at Beijing Information Science and Technology University. He is also affiliated with the State Key Laboratory of Network and Switching Technology at Beijing University of Posts and Telecommunications, where he obtained his Ph.D. degree under the guidance of Prof. Junliang Chen. From 2018 to 2020, was co-supervised by Prof. Zhi-Li Zhang at the University of Minnesota - Twin Cities. His research interests encompass mobile computing, cloud computing, computer networking, software-defined networks, and network function virtualization. Wendi serves as a reviewer for prestigious journals such as JNCA and TNSM.

Ke Liu received her B.E. from Beijing Information Science and Technology University in 2020. She is currently a master student advised by Prof. Wei Zhang and Wendi Feng at Beijing Information Science and Technology University. Her research interests include cryptographic application, privacy protection, and software-defined network.

Shuo Sun received his B.E. from Beijing Information Science and Technology University in 2021. He is currently a master student advised by Prof. Wei Zhang and Wendi Feng at Beijing Information Science and Technology University. His research interests include monolithic OS kernels and network systems.

Bo Cheng received the Ph.D. degree in computer science from the University of Electronics Science and Technology of China, Chengdu, China, in 2006. He is a Professor with the State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications, Beijing, China. His research interests include service computing, the Internet of Things, and multimedia communications.

Wei Zhang received both his B.E. and Ph.D. degree from Tsinghua University. He is currently a full Professor with School of Computer Science at Beijing Information Science and Technology University. His current research interests include big data storage and security, software and hardware co-design.

Declaration of interests

☒ The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

☐ The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: