Modelling Agent-Skipping Attacks in Message Forwarding Protocols

Zach Smith
Sjouke Mauw
zach@almou.se
sjouke.mauw@uni.lu
University of Luxembourg

Hugo Jonker
hugo.jonker@ou.nl
Open University of the Netherlands

Hyunwoo Lee
lee3816@purdue.edu
Purdue University

ABSTRACT

Message forwarding protocols are protocols in which a chain of agents handles transmission of a message. Each agent forwards the received message to the next agent in the chain. For example, TLS middleboxes act as intermediary agents in TLS, adding functionality such as filtering or compressing data.

In such protocols, an attacker may attempt to bypass one or more intermediary agents. Such an *agent-skipping attack* can the violate security requirements of the protocol.

Using the multiset rewriting model in the symbolic setting, we construct a comprehensive framework of such path protocols. In particular, we introduce a set of security goals related to *path integrity*: the notion that a message faithfully travels through participants in the order intended by the initiating agent. We perform a security analysis of several such protocols, highlighting key attacks on modern protocols.

KEYWORDS

security protocols; formal verification; multiparty protocols;

1 INTRODUCTION

The complexity of communication systems often necessitates the use of simplifying assumptions in order to enable functional models. For example, we typically treat internet communication as two-party protocols, focusing on the end-to-end security requirements. However, modern protocols often involve the addition of intermediate agents in order to enhance functionality. In these settings, messages are forwarded down a *path* of connected agents. In secure routing protocols [22, 30], agents might simply forward incoming messages to the next agent after inspecting the packet headers. However, some advanced protocols require agents to actively participate in some form. For example, the TLS protocol [14, 35] is used for an overwhelming majority of modern web communication. TLS users often use services that intercept messages, redirecting or modifying their contents, such as load balancers and firewalls. Such *TLS middleboxes* necessitate modifications to the TLS protocol in order to support multiple parties. Currently, the most common approach in handling this is known as Split TLS [20], in which the TLS session is "split" into a series of completely disjoint sessions between each pair of intermediate agents. Security concerns about Split TLS [13, 16, 45] have led to new solutions, such as mcTLS [32].

Mixnets and Onion Routing protocols also involve sending messages along a path. Such protocols originate from Chaum [8]. Modern protocols such as The Onion Router [40] and Sphinx [12] are built in a similar way. These protocols create layered messages ("onions"), where each layer contains information specific to one agent in the chain. This agent then peels off that layer and forwards the remaining onion to the next agent. A novel application of this is in the field of payment networks, e.g., the Lightning Network [34]. In this system, off-chain payment channels between pairs of agents replace traditional transactions. This leads to the concept of chained payments, in which an agent can send funds to a peer through a connected series of such channels, where each intermediate agent collects a small fee.

Several *agent-skipping attacks* have been found on such protocols [21, 27]. Intuitively, these attacks arise from the use of shortcuts: redirecting or modifying messages in order to bypass one or more agents in the chain. The basic structure of an attack is given in Figure 1.
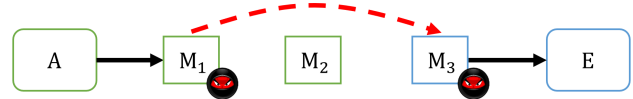

Figure 1: An agent-skipping attack: agents M_1 and M_3 collude to bypass M_2 using an out-of-band channel.

The impact of a skipping attack depends on the specific setting. For TLS, the principal agents want to ensure that their middleboxes are being respected. For example, it must not be possible to bypass a content filter, lest malicious injected code could reach an endpoint. For onion-style protocols, the sender wishes to preserve privacy by guaranteeing that the message travels only between the trusted intermediaries. Finally, for payment networks, intermediate agents must be assured that if they assist in a payment by forwarding a message, they are guaranteed to receive their transaction fees during the resolution stage. In order to protect against such attacks, protocols require a notion of *path integrity*. Participating parties must be sure that all messages follow the intended path. This paper is focused on building a simple and modular framework for verifying such a property.

Contributions. Our contributions are as follows:

- We highlight the threat posed by message skipping attacks on a wide domain of protocols, including an in-depth case study on the mbTLS extension to TLS.
- We introduce a symbolic framework, built upon the multiset rewriting model, that can be used to describe the structure of 'path-based' protocols.
- We give a formal definition of the *path integrity* security goal inside this framework.
- We provide a collection of models of a range of protocols from the literature using the Tamarin prover tool, showing the applicability of our framework.

Table 1: Protocols which rely on Path Integrity for security

Protocol Family	Example Protocols	Relevance of Path Integrity	Example Attack
Secure Routing	OPT [22], ICING [30], RPL [46]	Poisoning network topology	Tunelling [37]
Mixnet	Chaum [8], TOR [40], HOR-NET [9]	Privacy of path compromised	RAPTOR [39]
Payment Network	Lightning [34]	Transaction fee skimming	Wormhole [27]
Middlebox-Enabled TLS	mcTLS [32], mbTLS [31], maTLS [25], meTLS [26]	Bypass functionality-enhancing services	

2 BACKGROUND AND RELATED WORK

We begin by briefly highlighting examples of failures of Path Integrity across the literature, in Section 2.1. In Section 2.2, we discuss existing formalisms of related security goals. Finally, Section 2.3 contains a case study of the middlebox-enabled TLS extension mbTLS [31], to highlight a specific scenario where Path Integrity is relevant, and provide an intuitive definition.

2.1 Attacks on Path Integrity

Table 1 contains a list of protocol families of interest, as well as the consequences of skipping attacks.

Secure Routing. Secure Routing protocols form the simplest class of protocols where a notion of Path Integrity is a goal – with several different well-studied and comparable definitions, including Path Enforcement [7] or Path Compliance. In these protocols, a path is often set by using a set of path headers, along with some series of verification checks at intermediate nodes. In many cases, a key focus is the process of path selection and building an accurate view of the network topology. Recently, several attacks have been identified on such protocols for IoT devices [24], which have constrained computational capacity [46].

Mixnets. Mixnets are a variant of Secure Routing protocols which sacrifice the ability of individual nodes to validate the path in favour of enhanced privacy features. Each node is made aware only of the previous and next node in the path, often using a multiply-encrypted bundle containing the relevant headers. In this case, an attacker injecting additional nodes in the path can violate privacy goals, as they are able to infer the identity of the other intermediate agents. There are many network-based attacks on Mixnet protocols, for example, using traffic analysis. In the case where an adversary is granted additional privileges (for example, identity spoofing), interception attacks may exist [39].

Payment Networks. Payment Networks are Layer-2 cryptocurrency protocols for facilitating faster payments. The most notable is the Lightning Network [34] for Bitcoin. Agents set up long-term payment "contracts" using Bitcoin, and payments take place by re-negotiating this existing contract, allowing them to take place off-chain.

In order to send funds to a partner with whom you do not have an active contract, the Lightning Network allows chained payments: agents may publicly declare the presence of an open contract in order to allow others to route payments through it. Informally, if A wishes to pay C an amount n, he makes a deal with B such that A will pay B n if and only if B pays this same amount to C. B can then charge some small transaction fee for performing this service.

The atomicity of this transaction is protected by the use of a hashed challenge $h(x)$ (and associated preimage x). An agent can only retrieve their share of the funds if they can produce the preimage to the next agent in the path. Although the Lightning protocol is built on top of the Sphinx [12] Mixnet, the shared value $h(x)$ removes the privacy benefits – and worse, revealing the preimage x causes a Wormhole attack, as shown by Malavolta et al. [27]. This allows colluding parties to skim the transaction fees of agents between them in the path.

Middlebox-Enabled TLS. Unlike the other sets of protocols, a key focus in Middlebox-Enabled TLS protocols is that intermediate agents generally have full access to the message payload. As such, a key component of their design is in giving the endpoints control over who has access to the necessary encryption keys, to avoid the leaking of sensitive data. We will discuss these protocols in more detail in Section 2.3.

2.2 Existing Formalisations

Although the notion of Path Integrity is relevant to several domains, the majority of the literature around this topic is restricted to Secure Routing protocols, where it is the primary security goal.

Several formal treatments have been made for modelling such protocols [29]. This includes the work of Zhang et al. [47], who verify the OPT protocol [22] using an embedding of the LS^2 logic into the Coq [2] tool. More recently, Klenze et al. [23] built a more general approach using the theorem prover Isabelle/HOL. They use this approach to verify several secure routing protocols including ICING [30]. In this work we highlight the fact that this security goal is in fact relevant to several other protocol domains. Instead of using a general-purpose theorem prover, we design a framework that is directly compatible with the security tool Tamarin [28]. This framework allows us to specify protocols from all of the above domains, using the standard Dolev-Yao [15]. adversary in a symbolic setting, rather than the computational settings previously used.

For this work we assume that the intended message path is fixed during the setup of the protocol – i.e. the initiating agent has selected the full path. We do not make any assumptions about the knowledge of the intermediate agents Though the path selection process often presents unique problems [19, 33, 36], we consider it out of scope of this work, as it is often domain-specific.
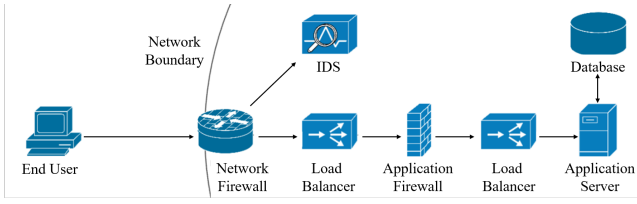
Figure 2: Example Network Architecture. Several intermediate applications operate between the Network Boundary and the Application Server.

2.3 TLS Middlebox Extensions

Transport Layer Security (TLS) [14, 35] is the de-facto standard for end-to-end security on the web. In some cases, endpoints may introduce TLS *middleboxes* – intermediate agents that attempt to add functionality to a TLS session, such as by filtering undesirable traffic or compressing data. In order to perform these functions, such middleboxes require full access to the (encrypted) message payloads, thus requiring that they are active participants in the TLS session.

Example Network. In order to motivate the discussion in this section, consider the example network infrastructure given in Figure 2. Requests to a web-facing service are passed through a series of middleboxes before reaching the application servers. Intermediaries include Load Balancers which control the flow of messages to multiple servers in parallel, and Application Firewalls.

Traditionally, enabling these middleboxes is achieved using a process known as SplitTLS [20], which requires endpoints to grant any middleboxes access to their certificate private keys (for a server), or to install a root certificate (for a client). This allows the middleboxes to impersonate the protected application server.

However, SplitTLS has been shown to degrade the security of a TLS session. This could be caused by TLS middleboxes that are either incorrectly or maliciously implemented [13, 16, 38, 42]. For example, a middlebox may support insecure or deprecated ciphersuites, allowing an adversary to perform a man-in-the-middle attack to degrade the security of the entire session.

As an alternative to SplitTLS, a series of middlebox-enabled TLS schemes have been proposed [25, 26, 31, 32]. These avoid the problems associated with sharing keys by instead increasing the visibility of these middleboxes and allowing the end user to confirm the authenticity of the application server directly.

Security Goals of TLS Middlebox Extensions. Middlebox-enabled TLS extensions aim to achieve three main goals:

- Allow middleboxes to identify themselves as active parts of a TLS session, rather than creating split sessions.
- As a result, avoid dangerous sharing of critical private keys (and their associated certificates) between multiple devices.
- Monitor and regulate the actions that middleboxes take.

These goals can be achieved in several ways. Generally, approaches involve some combination of the following strategies:

- Establishing a secure channel between the two endpoints as part of the initial handshake, in order to exchange session-critical data.

- Delegation of session keys from endpoints to middleboxes, rather than piecewise establishment between pairs of agents.
- Addition of a "modification log" to messages, showing which middleboxes have modified a message in-transit.

In this work, although we allow for dishonest middleboxes who release their keys, we assume that all honest agents are willing only to run the specific TLS variant in question, and follow it faithfully. Further, our analysis considers only the record phase of TLS, where application data is secured with keys established during the handshake phase of TLS. We assume the handshake was correctly performed. Indeed, our security claims follow from the order of agents that is believed to have been established during this handshake.

Although significant security analysis has been put into the core TLS protocol suite [3, 10], the security discussion of middlebox-enabled extensions is still somewhat limited. Some formal models have been created for accountable proxying, such as those of Bhargavan et al. [4, 5]. However, these focus more on end-to-end authentication guarantees in the presence of proxies, avoiding discussion of path integrity. They use a computational setting, rather than a symbolic model such as the one we present here. One middlebox-enabled TLS extension, maTLS [25], provides a definition of path integrity for their specific use-case, while we aim to produce a broad-scope definition.

The mbTLS Protocol. The mbTLS protocol [31] is a proposed middlebox-enabled TLS scheme. It differs from standard TLS in the following ways:

- During the mbTLS handshake between a client and a server, middleboxes report their presence to the endpoints with an additional handshake.
- At the end of the handshake phase, each middlebox is delegated two session keys for its associated sessions (i.e. one where it acts as a client, and one as a server) from the endpoint which has deployed them.
- During the record phase, messages between the endpoints are passed down the chain of middleboxes, which decrypt
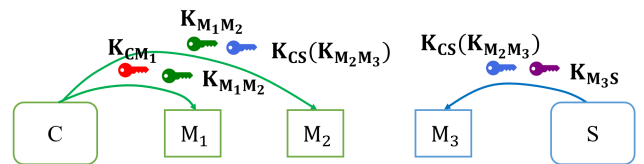
(a) The Client and Server delegate keys to the middleboxes that they have introduced for each of the two TLS (sub)sessions for the chained connection
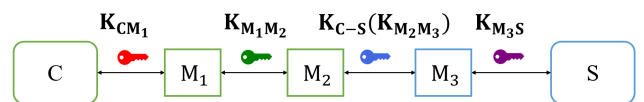
(b) Intended execution - messages flow down the pre-established path. Middleboxes use keys received from the handshake phase

Figure 3: Overview of the mbTLS protocol

then re-encrypt each message (whilst performing any of their functions on the message body)

Importantly, mbTLS only proposes significant alterations to the handshake phase of the TLS protocol. Once the session has been established, each segment in the path performs the TLS record phase protocol as usual. This is in contrast to e.g. mcTLS [32], in which agents who perform modifications re-encrypt the payload with a different key, or maTLS [25], where a MAC is appended to the message at each segment. An overview of the mbTLS scheme is given in Figure 3.

The mbTLS scheme assumes that middleboxes are running on hardware enclaves, such as the Intel SGX framework [1]. As a result, the authors assert that middleboxes can be seen as trusted agents for the purpose of security analysis. However, we argue that this is not a realistic security model. Software enclaves are designed to ensure that software is loaded without modifications after distribution, and is run in a secure environment. This does not provide any guarantees about security of the software itself.

As such, maliciously designed or configured middlebox software is not protected against. This means that an attacker could write malware which poses as a service-providing middlebox (such as an ad-blocker), but instead fulfills some other purpose (such as leaking data). In addition, we note the existence of several attacks on trusted execution environments [6, 18, 43, 44], suggesting that even well-intending participants may accidentally leak secret data.

Middlebox Skipping Attack on mbTLS. Our analysis shows that the mbTLS protocol admits a skipping attack.

Looking back to Figure 2, suppose a malicious administrator has access to the two layers of load balancers. During the mbTLS handshake phase, the Application Firewall is registered as being an active participant in the session. However, once packets are sent during the record phase, the two load balancers communicate directly through a side channel, bypassing the application firewall. This idea is shown in Figure 4.
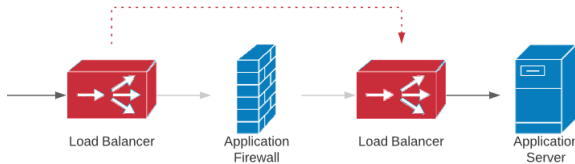

Figure 4: Firewall is bypassed by two collaborating agents

The application server is unable to differentiate between the firewall choosing not to modify a message, and the firewall never having received the message at all. In this way, malicious inputs could be targeted against the AS without the protection of the firewall.

Intuitively, the Application Server should have some guarantee that because it saw the firewall registered during the handshake phase, the firewall continues to be an activate participant in all record phase messages. This leads to the following requirement on message-forwarding protocols:

DEFINITION 2.1. *Path Integrity (Notion) Once a path has been established, if a message is received by one agent on a path, then all previous agents on the path should have also received it.*

In order to prevent an attack such as the one here, a simple approach is to add some form of read-receipt to messages, in the form of a MAC or signature from each middlebox in the path. Upon receipt of a message, endpoints can confirm that the path was followed faithfully by checking that the set of signatures has been correctly constructed. Indeed, the mcTLS [32] scheme makes use of "write keys" that privileged middleboxes use to update a write MAC, but possesses a similar vulnerability to mbTLS in that if no changes are made to the payload, the associated message also remains unchanged outside re-encryption.

3 MULTISET REWRITING THEORY

In this section, we introduce the framework that will be used to describe protocol execution and analysis for the rest of the paper. The language we use can be considered as a subset of that supported by the Tamarin prover tool [28, 41].

We employ a multiset rewriting system – a special form of term rewriting system. The state of the communication network is modelled by a collection of *facts*, with rewrite rules which add or remove from this collection, forming a labelled transition system.

3.1 Fundamentals

Our term rewriting system is built using terms from an order-sorted algebra (e.g. Goguen [17]). We define two top-level sorts msg and $Fact$, and subsorts pub, $fresh$, such that $pub < msg$ and $fresh < msg$.

Intuitively, msg terms represent any value that might be used on the communication network, while pub and $fresh$ terms represent public and freshly generated values, respectively. We write $x \colon y$ to indicate that the term x is of type y. We define a public term $'' \colon pub$, representing the empty string. Public terms will also be used to indicate agents' identities. Notationally, we write the following symbols to indicate the type of a term: $A \colon pub$, $x \colon fresh$, $m \colon msg$.

We allow for collections of function symbols $\Sigma_{msg^*, msg}$ and $\Sigma_{msg^*, Fact}$, which map a sequence of type msg (or its subsorts) to either another msg or a $Fact$. For simplicity we denote $\Sigma_{msg^*, msg}$ by Σ and $\Sigma_{msg^*, Fact}$ by F.

Atoms (i.e. undecomposable terms) can represent names (i.e. unassigned expressions), or variables (i.e. assigned values). A term is said to be *ground* if it contains no variables. A *substitution*, σ, is a (partial) function from variables to terms of the same sort (or a subsort). We say a substitution is *ground* if it maps to a ground term. A substitution is applied to a term by applying it to each subterm. Given a term t and a substitution σ, we say that the substitution σ *grounds* t if the resultant term $t\sigma$ is a ground term. We allow for an equational theory E over terms of type msg, or its subtypes. E is a collection of equations $lhs = rhs$. Two terms s and t are said to be equivalent modulo E if a series of equations can be applied to s or t (or both) such that the resulting terms s' and t' are equal. We will make use of a standard set of function symbols Σ, as well as an associated equational theory E, which is presented in Table 2.

We write $\{x\}_k$ to denote encryption when the kind is clear from the context. Similarly, we will sometimes omit the pair operators $\langle$ and $\rangle$ for readability.

Symbols	Equations
	Hashing
h /1	
	Pairing
$\langle\rangle$/2	
fst /1	$\mathrm{fst}(\langle x, y\rangle) = x$
snd /1	$\mathrm{snd}(\langle x, y\rangle) = y$
	Symmetric encryption
senc /2	
sdec /2	$\mathrm{sdec}(\mathrm{senc}(m, k), k) = m$
	Asymmetric encryption
pk /1	
aenc /2	
adec /2	$\mathrm{adec}(\mathrm{aenc}(m, \mathrm{pk}(k)), k) = m$
	Signatures
sign /2	
true /0	
verify /3	$\mathrm{verify}(\mathrm{sign}(m, k), m, \mathrm{pk}(k)) = \mathrm{true}$

Table 2: Collection of functions Σ and equational theory E

3.2 Protocol Specification

Thusfar, our discussion has been restricted to terms of type *msg*. We now divert our attention to *Facts*. Intuitively, while *msg* terms model the value of certain things (messages, agent names, encryption keys), terms of type *Fact* describe the state of the protocol execution itself. As such, facts are often parameterised by one or more message terms. We reserve the following fact symbols with corresponding intuition:

- Net /1: A message on the communication network
- K /1: Adversary knowledge of a term
- Pk /2, Ltk /2: Public and long-term keys
- ShKey /3: A shared encryption key

From now on we assume that we are working over multisets where all terms are of type *Fact*. A *State*, S, is a multiset where all of the terms are ground, and models the current execution state of a protocol. Each run of a protocol will begin with an empty state, which then transitions into future states through a series of *rules*.

A *rule* r is defined by a triplet of multisets $r : L \xrightarrow{E} R$. Given a state S, and a substitution σ, we can apply rule r if:

- σ is a grounding substitution for L and R, and
- $L\sigma \subset S$.

In this case, the state S' is produced by removing the submultiset of S equal to $L\sigma$, and replacing it with $R\sigma$. The elements of $E\sigma$ are known as the *event facts* of the rule.

For convenience, we will sometimes use the prefix ! in fact symbols to indicate that the fact is *persistent*. Persistent facts are never removed from a state as a result of rule execution . They represent reusable assets, such as encryption keys or adversary knowledge.

A simple example of a rule is given in `Dec_Fwd` below, in which an agent (whose name will instantiate the variable A), detects a message on the network encrypted with their public key, and decrypts it before forwarding on the result.

$$\mathtt{Dec\_Fwd} := \begin{bmatrix} \mathrm{Net}(\{m\}_{pk(k)}) \\ !\mathrm{Ltk}(A, k) \end{bmatrix} \xrightarrow{\mathrm{Fwd}(A, m)} \begin{bmatrix} \mathrm{Net}(m) \end{bmatrix}$$

When a rule is applied, the terms $E\sigma$ are appended to the *trace*, τ, an indelible ordered history of event markers. At the start of any execution, $\tau = \phi$, the empty trace. After the execution of a rule r, the resulting event facts are added along with a discrete time marker $\#t_i$. For example, an application of the rule above might append the fact $\mathrm{Fwd}(A, m)@t_1$ to the trace. Time markers are assumed to be ordered, unique, and increasing. However, they hold no values (they do not represent actual timestamps, only the order of events in an execution).

We will freely quantify over time markers e.g. $\forall \#t_i$ when there is no ambiguity, and will make use of ordering of time markers (e.g. $\#t_i < \#t_j$).

We reserve the following special rule `Fresh`:

$$\mathtt{Fresh} := \begin{bmatrix} - \end{bmatrix} \longrightarrow \begin{bmatrix} \mathrm{Fr}(x : \textit{fresh}) \end{bmatrix}$$

This rule allows for the creation of freshly generated random variables, for example for use in creating encryption keys. We specially require that each execution of the `Fresh` rule is instantiated by a distinct, previously unused value for x, and that it is the only rule which can create Fr facts. Since the created Fr fact is linear, it is consumed if used by a later rule – this ensures that the same random value can never be generated twice.

A protocol, P, is given by $P = (R, \mathsf{F}, \Sigma, E)$, a tuple of rules, facts, functions and an equational theory. We will assume that R, F, Σ and E contain all the reserved elements indicated in this section (including the adversary rules below). We define $Traces(P)$ as the set of all (valid) traces that can be constructed as a result of executing the rules from R along with the associated material from the other protocol components.

3.3 Adversary Model

Our model uses the Dolev-Yao [15] adversary. It is also defined in terms of multiset rewriting rules, given in Figure 5. The Dolev-Yao adversary is capable of eavesdropping, modifying, and retransmitting messages, modelled by the `Block` and `Inject` rules. The `Adv_Pub` and `Adv_Fr` rules allow the adversary to deduce public and (previously unused) fresh terms, while $\mathtt{Fun_f}$ allows the adversary to derive new terms by applying function symbols to known terms. Finally, the various `Corrupt` rules model agents who are fully under the control of the adversary.

3.4 Security Properties

Security goals of a protocol are given in terms of first-order logic formulae on the set of traces of the protocol. Intuitively, they indicate that certain events can or cannot happen, or that they must occur in a certain order.

A security goal holds for a given protocol if all traces of the protocol satisfy it. A trace which violates the goal can be reconstructed into an attack on the protocol. As a simple example, consider the following protocol P in Alice-Bob notation:

$$A \to B : \quad \{x, y\}_{k_{AB}}$$
$$B \to A : \quad \{y, x\}_{k_{AB}}$$

$$\texttt{Inject} := \big[\ !K(x)\ \big] \rightarrow \big[\ \mathrm{Net}(x)\ \big]$$

$$\texttt{Block} := \big[\ \mathrm{Net}(x)\ \big] \xrightarrow{K(x)} \big[\ !K(x)\ \big]$$

$$\texttt{Adv\_Pub} := \big[\quad - \quad\big] \xrightarrow{\mathrm{K}(A)} \big[\ !K(A:pub)\ \big]$$

$$\texttt{Adv\_Fr} := \big[\ \mathrm{Fr}(x)\ \big] \xrightarrow{\mathrm{K}(x)} \big[\ !K(x)\ \big]$$

$$\texttt{Fun}_\mathtt{f} := \begin{bmatrix} !K(x_1) \\ \cdots \\ !K(x_2) \end{bmatrix} \xrightarrow{\mathrm{K}(f(x_1,\ldots,x_n))} \big[\ !K(f(x_1,\ldots,x_n))\ \big]$$

$$\texttt{Corrupt\_Ltk} := \big[\ !\mathrm{Ltk}(A,k)\ \big] \xrightarrow{\mathrm{Corrupt}(A)} \big[\ !K(k)\ \big]$$

$$\texttt{Corrupt\_L} := \big[\ !\mathrm{ShKey}(A,B,k)\ \big] \xrightarrow{\mathrm{Corrupt}(A)} \big[\ !K(k)\ \big]$$

$$\texttt{Corrupt\_R} := \big[\ !\mathrm{ShKey}(A,B,k)\ \big] \xrightarrow{\mathrm{Corrupt}(B)} \big[\ !K(k)\ \big]$$

Figure 5: Rules which define the Dolev-Yao adversary.

In this protocol, A sends a pair of encrypted terms to B, who reverses their order and sends them back. A reasonable security claim for this protocol might be "The value x is either unknown to the adversary, or one of A or B is corrupt". This could be expressed as:

$$\forall A, B, x, t_i\colon \mathrm{Secret}(A, B, x)@t_i \implies$$
$$\nexists \#t_a\colon \mathrm{K}(x)@t_a\ \vee\ \exists \#t_b\colon \mathrm{Corrupt}(A)@t_b\ \vee\ \exists \#t_c\colon \mathrm{Corrupt}(B)@t_c$$

Note that this definition implies that the adversary can *never* learn the value x. We can add event orderings (such as $\#t_a < \#t_i$) to refine these claims. Our claims are generally quantified over all traces $\tau \in \mathit{Traces}(P)$,

We also use the same syntax to restrict the set of traces of a protocol to be investigated. As an example, we introduce the Equal fact, denoting a test for equality, and consider only traces where this test is successful: $\forall x, y, \#t_i\colon \mathrm{Equal}(x, y)@t_i \implies x = y$.

Such a restriction ensures we only consider traces in which signatures are correctly verified – for example, a rule containing $\mathrm{Equal}(\mathrm{verify}(sig, msg, pkA), \mathrm{true})$ can only be executed in the case that the value sig is indeed a signature for msg.

4 MODELLING PATH PROTOCOLS

In this section we introduce the notion of a *path protocol*. We break down the structure of a path protocol into a set of phases, and describe each of these phases as a collection of generic rules.

4.1 Running Example

Consider the multi-party message forwarding protocol shown in Figure 6, which uses public key encryption.

Intuitively, the agent A wants an intermediate agent B to forward a message p to C. This is achieved by using nested encryption. The protocol is depicted with one forwarding agent B, but indeed it could be trivially extended for any number of forwarding agents.

It would be relatively straightforward to model this protocol for a fixed number of agents, by specifying the value of the message
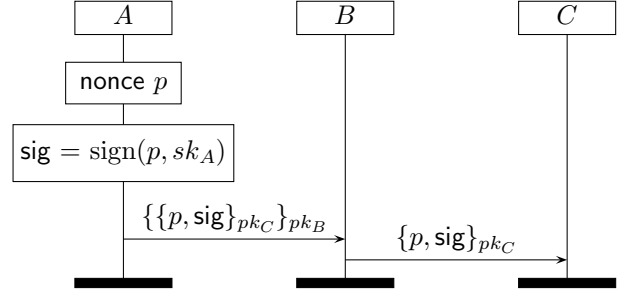

Figure 6: A simple message forwarding protocol.

at each step of execution. However, this design quickly becomes cumbersome as the number of agents grow. Moreover, it requires a separate model for each number of agents. Instead, we construct a single model which accounts for any number of agents. Figure 7 shows the set of rewriting rules which model this protocol.

4.2 Modelling Multi-Step Messages

Our model in Section 3 considers messages to be sent between pairs of agents. We now build a set of rewriting rules which allow us to specify protocols that use these multi-step messages. Intuitively, such messages are constructed by "wrapping" layers of encryption on top of each other. This approach will allow us to define *path protocols*, which we break down into a series of phases.

DEFINITION 4.1 (PATH PROTOCOLS (NOTION)). *A Path Protocol is a protocol in which rules (other than those modelling adversary capabilities) can be categorised as belonging to one of the following phases:*

- **Setup Phase**: *A preliminary phase in which agents' encryption keys are established*
- **Construction Phase**: *An initial agent creates a message from a payload, by configuring it to pass through a series of intermediate agents*
- **Forwarding Phase**: *Each intermediate agent receives, repackages and forwards the message*
- **Receive Phase**: *The intended recipient receives the final message and retrieves the payload.*

Throughout the course of our discussion, we will generally use the name p to refer to the *payload* – the intended value for the final recipient. The name m will be used to refer to *messages* sent between agents – including things like encryption (which we do model) and header or miscellaneous data (which we do not). We make the assumption that for each individual session, the payload will always contain some session data or randomness that makes it unique, and thus suitable as a session identifier.

Over the course of this section, we build a framework of generic rules which is sufficient to cover each of the individual phases in Definition 4.1. This resulting set of rules can be used to describe a large majority of path protocols.

Setup Phase. Our execution model begins with the empty multiset. In order for the protocol to begin, agents must be instantiated and assigned asymmetric and shared encryption keys. Our models make use of the `Gen_ShKey` and `Gen_Ltk` rules for generating encryption

Setup Phase

$$\texttt{Gen\_Ltk} := \big[\ \mathrm{Fr}(ltk)\ \big] \longrightarrow \begin{bmatrix} !\mathrm{Ltk}(A, ltk) \\ !\mathrm{Pk}(A, \mathrm{pk}(\cancel{ltk})) \\ \mathrm{Net}(\mathrm{pk}(ltk)) \end{bmatrix}$$

Construction Phase

$$\texttt{Create} := \begin{bmatrix} \mathrm{Fr}(p) \\ !\mathrm{Pk}(E, pkE) \\ !\mathrm{Ltk}(A, ltkA) \end{bmatrix} \xrightarrow[\mathrm{StartBuild}(A,p)]{\mathrm{Add}(p,E,\{p,m,''\})} \big[\ \mathrm{Build}(p, E, m)\ \big]$$

$$m = \{p, \mathrm{sign}(p, ltkA)\}_{pkE}$$

$$\texttt{Wrap} := \begin{bmatrix} \mathrm{Build}(p, M_i, m) \\ !\mathrm{Pk}(M_j, pk) \end{bmatrix} \xrightarrow{\mathrm{Add}(p,E,m,\{m\}_{pk})} \big[\ \mathrm{Build}(p, M_j, \{m\}_{pk})\ \big]$$

$$\texttt{Send} := \big[\ \mathrm{Build}(p, M_1, m)\ \big] \longrightarrow \big[\ \mathrm{Net}(m)\ \big]$$

Forwarding Phase

$$\texttt{Unwrap} := \begin{bmatrix} \mathrm{Net}(\{m\}_{\mathrm{pk}(ltk)}) \\ !\mathrm{Ltk}(M_i, ltk) \end{bmatrix} \xrightarrow{\mathrm{Forward}(M_i, \{m\}_{\mathrm{pk}(ltk)}, m)} \big[\ \mathrm{Net}(m)\ \big]$$

Receive Phase

$$\texttt{Receive} := \begin{bmatrix} \mathrm{Net}(\{p, sig\}_{\mathrm{pk}(ltkE)}) \\ !\mathrm{Ltk}(E, ltkE)) \\ !\mathrm{Pk}(A, pkA) \end{bmatrix} \xrightarrow[\mathrm{Equal}(\mathrm{verify}(sig,p,pkA),\mathrm{true})]{\mathrm{Forward}(E,\{p,sig\}_{\mathrm{pk}(ltkE)},'')} \big[\ -\ \big]$$

Figure 7: Full set of rewriting rules for the example protocol given in Figure 6.

keys, specified as follows:

$$\texttt{Gen\_ShKey} := \big[\ \mathrm{Fr}(k)\ \big] \longrightarrow \big[\ !\mathrm{ShKey}(A, B, k)\ \big]$$

$$\texttt{Gen\_Ltk} := \big[\ \mathrm{Fr}(ltk)\ \big] \longrightarrow \begin{bmatrix} !\mathrm{Ltk}(A, ltk) \\ !\mathrm{Pk}(A, \mathrm{pk}(ltk)) \\ \mathrm{Net}(\mathrm{pk}(ltk)) \end{bmatrix}$$

We also allow for the corruption of agents created during the setup phase. We will assume that the initiating agent in each session is honest, but allow all other agents to be under full adversarial control.

Construction Phase. The Construction Phase represents the beginning of a session of the protocol. We break this down into three rules: a `Create` rule which determines the payload, a `Wrap` rule which modifies the message to pass through an intermediate agent, and a `Send` rule, in which the message is released onto the network.

Implicit in the application of these rules is the notion of a *path order*. Intuitively, this is a relation that models the (intended) order in which the message will pass between protocol participants. The construction phase that takes place in each run of the protocol will define the path order for that execution.

DEFINITION 4.2 (PATH ORDER). *A Path Order is a total order $<_\pi$ on a (finite) set of public terms. We call the minimal element A of $<_\pi$ the initial agent, the maximal element E the final agent, and all other elements M_i intermediate agents.*

We allow for a different path to be chosen for each session. We assume that our protocols are designed such that paths are non-repeating (i.e. the path order is always well-defined). A path order

need not include all agent identifiers, and so the path can be of any length. The idea of a path order allows us to present an informal notion for the Path Integrity security goal that we will define in the next section.

DEFINITION 4.3 (PATH INTEGRITY (REVISITED)). *A protocol satisfies path integrity if for every session, for all M_i and M_j such that $M_i <_\pi M_j$, if M_j has forwarded the message, then M_i has also forwarded the message.*

In order to formalise this notion, we will make use of *expected messages*. We define the linear fact $\mathrm{Build}(p, M, msg)$, which represents the message as it is being constructed: the payload p is used as a unique path identifier, while the second and third terms indicate the current agent being considered as well as the current value of the message (for example, as successive layers of cryptography are applied). For more complex protocols, additional parameters may be added to the Build fact to track state. The event fact Add represents that an agent has been added to the path. It is parameterised by the path identifier, the agent who has been added to the path, and how the initiating party anticipates the message will be altered as it passes through them (for example, through de- or re-encryption). The StartBuild event fact is used to mark the beginning of the protocol execution.

These facts are used by three rules during this phase. In the first rule, the initiating agent determines the payload to be sent to the other endpoint. The second rule is repeatedly applied to add new intermediate agents to the path in order, each time replacing the current Build fact with a new version containing any required changes to the message. Finally, the last rule sends the message on to the network:

$$\begin{bmatrix} \mathrm{Fr}(p) \\ !\mathrm{Pk}(E, pkE) \end{bmatrix} \xrightarrow[\mathrm{StartBuild}(A,p)]{\mathrm{Add}(p,E,f(p),'')} \big[\ \mathrm{Build}(p, E, f(p))\ \big]$$

$$\begin{bmatrix} \mathrm{Build}(p, M_i, m) \\ !\mathrm{Pk}(M_j, pkJ) \end{bmatrix} \xrightarrow{\mathrm{Add}(p,M_i,m,g(m))} \big[\ \mathrm{Build}(p, M_j, g(m))\ \big]$$

$$\big[\ \mathrm{Build}(p, M_1, m)\ \big] \longrightarrow \big[\ \mathrm{Net}(m)\ \big]$$

We use anonymous functions f and g to depict how the message is changed – these are instantiated for each specification based on the protocol in question. For our example protocol, f is pairing with a signature, and g is asymmetric encryption using a public key.

The primary requirement of our "wrapping" function g (and later on the associated "unwrapping" function) is that the general structure of the message is preserved as it is transmitted between agents. In our example, each agent expects to receive (and send) a message that consists of a single term encrypted by a public key, $\{m\}_{pk}$. In a more complex situation, this message may contain components such as a public term (containing the identity of the next agent in the path), or a list of signatures. We make this requirement on the protocol specification level – an agent may not be able to verify themselves that part of the encrypted body of their message has this structure. We require this structural symmetry between the four following message types sent during protocol execution:

- The outwards packet sent by A
- The final packet received by E
- Each packet going into a forwarding agent M_i
- Each packet going out of a forwarding agent M_i

The ordering of Add facts in any given trace establishes a path order $<_\pi$: If the Add fact for M_i is added to the trace before that of M_j, we have that $M_j <_\pi M_i$.

Forwarding Phase. The forwarding phase of the protocol occurs as intermediate agents transceive the message We model this with the use of the `Unwrap` rule, in which each intermediate agent forwards the message. The exact nature of this forwarding is dependent on the protocol – it may involve de- or re-encryption, or reading information about how to route the forwarded message.

$$\begin{bmatrix} \text{Net}(m) \\ !\text{Ltk}(M_i, ltk) \end{bmatrix} \xrightarrow{\text{Forward}(M_i, m, f(m))} [\, \text{Net}(f(m)) \,]$$

We introduce the Forward fact to denote that the agent has forwarded the message, including the values it has changed from and to. In a faithful execution of the protocol, the parameters of these facts should agree with those in the Add facts created in the Construction phase.

Receive Phase. The last step of a protocol is upon the successful receipt of the payload by the endpoint.

$$\begin{bmatrix} \text{Net}(f(p)) \\ !\text{Ltk}(E, ltkE)) \end{bmatrix} \xrightarrow{\text{Forward}(E, f(p), \text{``''})} [\quad - \quad]$$

The `Receive` rule may include validation of the final payload. For example, in the example protocol, the final agent must ensure that the attached signature matches the payload. More advanced validation is discussed in the following section.

5 SECURITY GOALS FOR PATH PROTOCOLS

We introduce two security goals to cover the range of protocols built in the framework from the previous section. The first, *Path Integrity* (Sec. 5.1), covers the simplest case of message forwarding protocols. The second, *Verification-dependent Path Integrity* (Sec. 5.2) makes Path Integrity conditional on the receiving party validating the

---

$$\forall A, M_i, M_j, p_{ID}, f_i, t_i, f_j, t_j,$$
$$\#ta_i, \#ta_j, \#tk_i, \#ts.$$

1.	$(\not\exists \#ta_c.\ \text{Corrupt}(A)@ta_c)$
2.	$\wedge \text{StartBuild}(A, p_{ID})@ts$
3.	$\wedge \text{Add}(p_{ID}, M_i, f_i, t_i)@ta_i$
4.	$\wedge \text{Add}(p_{ID}, M_j, f_j, t_j)@ta_j$
5.	$\wedge(\#ts < \#ta_i < \#ta_j)$
6.	$\wedge \big(\text{Forward}(M_i, f_i, t_i)@tk_i \big)$
7.	$\implies$
8.	$\exists \#tk_j.\ (\#tk_j < \#tk_i)$
9.	$\wedge \big($
10.	$\big(\text{Forward}(M_j, f_j, t_j)@tk_j \big)$
11.	$\vee \big(\exists\ \#tc_j.\ \text{Corrupt}(M_j)@tc_j \wedge$
12.	$\text{K}(\langle f_j, t_j \rangle)@tk_j \big)$
	$\big)$

Figure 8: Statement of the Path Integrity security goal

---

received message. Note that these security goals are perpendicular to many existing security goals, e.g. those relating to secrecy or synchronisation between agents. For example, if an attacker is able to confuse an agent into believing that they are performing a different role in the protocol (e.g. that they are an endpoint when they should be forwarding the message), other attacks may arise. The security goals we present here form an extension of the classical security goals for this protocol domain, not a replacement.

Intuition. The intuition behind the structure of our security goals is as follows. Given a specific protocol session, we assume that a path order $<_\pi$ has been defined by a sequence of events. We then examine the case that some agent M_i has successfully forwarded the message. Our goal is satisfied if there is (fundamentally) only one way to fill the gap between these events: that each intermediate agent M_j such that $M_j <_\pi M_i$ has also forwarded the message. These claims are typically verified by performing backwards reasoning – checking all ways of reconstructing each partial trace.

5.1 Path Integrity

We begin by formalising Definition 4.3, the idea of *Path Integrity*. This goal represents the initiating agent's belief that a sent message will indeed travel through the list of intermediate agents in the intended path in the correct order. Intuitively, this requires a correspondence between the order in which agents were named in applications of the `Wrap` and `Unwrap` rules.

DEFINITION 5.1. *Path Integrity*
We say that a protocol P satisfies Path Integrity if and only if all traces $\tau \in Traces(P)$ satisfy the property displayed in Figure 8.

The intuition of the definition in Figure 8 is as follows: (1) Suppose A is an honest agent, who (2) starts a session ID p_{ID}, (3) adding agents M_I and (4) M_j to the path, such that (5) $M_j <_\pi M_i$, and (6) M_i has successfully forwarded the message, (7) then (8) at some earlier time, (9) either: (10) M_j forwarded the message, (11) or M_j is corrupt and (12) the adversary had the necessary knowledge to forward the message.

Note that in case the agent is corrupt, the adversary may forward the message. In this a situation, the path order is still preserved, as the message was still correctly forwarded.

5.2 Verification-Dependent Path Integrity

The security property defined in the previous section can be seen as *on-the-fly* security: it ensures that path integrity holds even while a message is in-flight. This is common for onion-style protocols, which use layered encryption such that the message can only be decrypted in a specific order. However, for many protocols, this requirement might be too strict – we may want to loosen it to only consider *completed* sessions. This approach is typical in middlebox-enabled TLS protocols [25, 26, 32], where an additional *verification phase* exists to ensure that a session has executed completely. We extend our framework to model such protocols and their security.

Verification Phase. We assume that the final message received by the endpoint can be broken into two main parts: one containing the session payload (or some function thereof), and another which includes validation data, such as appended MACs.

We modify the `Receive` rule to include a $\text{Check}(E, f(p), m)$ fact, separating these two components. Successive applications of a `Verify` rule then check the validation data added by each intermediate agent in turn. A final rule runs after all verification steps to confirm successful completion.

$$\begin{bmatrix} \text{Check}(E, f(p), m) \\ !\text{Pk}(M_i, pkM_i)) \end{bmatrix} \longrightarrow \begin{bmatrix} \text{Check}(E, g(p), l(m)) \end{bmatrix}$$

$$\begin{bmatrix} \text{Check}(E, p, \text{``}) \\ !\text{Pk}(A, pkA)) \end{bmatrix} \xrightarrow{\text{Complete}(E, f(p))} \begin{bmatrix} - \end{bmatrix}$$

As in similar rules, anonymous functions f, g, l are used to denote the changing values of the payload and validation portions of the message as it is decomposed between verification steps. These functions are instantiated for individual protocols.

The security definition for Verification-Dependent Path Integrity differs only in the addition of an event marking that verification was successful at the end of the protocol's execution. We extend the final rule of the protocol with a `Complete` event fact, representing that the verification process has successfully completed. This event fact is parameterised by the path identifier, allowing it to be readily associated with the corresponding `StartBuild` fact. The corresponding security claim differs only in that we include the existence of this fact in the premise of the implication.

6 EXPERIMENTS

To demonstrate the applicability of our framework, we consider multiple protocols from the literature. We perform a security survey, building implementations using the Tamarin [28] prover tool[1].

We split our analysis into three main families, as follows: **Middlebox-Enabled TLS**, **Mixnets** and **Payment Networks**. We consider onion-style protocols (such as those demonstrated by Chaum [8]) as part of the Mixnet family – although there can be some differences on the network layer (such as by batching messages), the message structure is often very similar. Our reasons for choosing these specific protocols are as follows:

- **Middlebox-Enabled TLS**. mcTLS [32] uses session keys shared between multiple parties based on their permissions (read, write), rather than their location in the path. mbTLS [31] uses unique session keys for each pair of adjacent agents. maTLS [25] furthers this scheme with chained signatures (forming a *modification log*).
- **Payment Networks**. The Lightning Network [34] uses per-hop payloads to conceal routing data, with a term $h(x)$ that is shared between agents. We break the protocol into two phases – the setup phase in which HTLCs are established, and the unlock phase in which funds are released by sharing the inverse hash x between participating agents.
- **Mixnets**. The original models by Chaum [8] use chained asymmetric encryption. The TOR [40] ecosystem establishes connections using symmetric encryption (with public keys only for key-establishment) - paths are defined by a connection ID. HORNET [9] reduces the use of state compared

---

[1]Our Tamarin models, as well as protocol diagrams showing our level of detail, are available from https://github.com/path-integrity-analysis/path-integrity

Table 3: Security of protocols considered in our analysis

Protocol Name	Path Integrity
Middlebox-Enabled TLS	
- mcTLS [32]	$\times$
- mbTLS [31]	$\times$
- maTLS [25]	$\checkmark^*$
Payment Network	
- Lightning (setup phase) [34]	$\checkmark$
- Lightning (unlock phase) [34]	$\times$
Mixnet	
- Chaum [8]	$\checkmark$
- TOR – Establishment [40]	$\checkmark$
- TOR – Data Exchange [40]	$\checkmark$
- HORNET [9]	$\checkmark$

to TOR, by including routing data as part of the message in place of a connection ID, using a construction based on Sphinx [12].

The results of our analysis can be found in Table 3. We give results for Path Integrity (as per Definition 5.1). For each protocol we indicate if the security goal is met ($\checkmark$) or violated ($\times$). An asterisk (*) indicates the goal is dependent on a Verification phase (as in Verification-Dependent Path Integrity).

7 CONCLUSION

In this paper we have considered multi-party protocols in which messages are forwarded from one endpoint to another through a series of intermediate agents. We gave several formalisations of path integrity inside a multiset rewriting model used by the automated verification tool Tamarin.

We applied these definitions in a comprehensive security survey, demonstrating a novel attack on mbTLS, and allowing for simple, automated security proofs for payment networks on the Bitcoin Lightning infrastructure and other scenarios. These attacks demonstrate that path integrity is an important security property that is not covered by traditional authentication goals. Though the impact of a message skipping attack can vary based on the domain, the consequences can often be significant.

Future Work.
There are several avenues for future work. Our approach focused on ensuring that agents behave in the correct order as defined by a pre-set path. However, our analysis does not consider traditional security goals, such as those for secrecy or integrity. Although there has been some work into approaches for the extension of these security goals into multiparty settings [11], this is often on a mutually pairwise basis. Instead, considerations of the differing levels of participation of agents may lead to more precise security definitions.

In this work we modelled the Lightning Payment protocol as two separate one-way protocols. However, our model can be extended in order to enable analysis as a singular protocol. This gives rise to a notion of *symmetry* – in the case where the initial path of the message is not known or fixed, but we wish to ensure that the return

journey is identical to the forward journey. This may be relevant in several fault-tolerant versions of payment network protocols, where a payment is split into many small atomic transactions that are independently routed.

REFERENCES

[1] Ittai Anati, Shay Gueron, Simon Johnson, and Vincent Scarlata. 2013. Innovative technology for CPU based attestation and sealing. In *Proceedings of the 2nd international workshop on hardware and architectural support for security and privacy*, Vol. 13. ACM New York, NY, USA.

[2] Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Jean-Christophe Filliatre, Eduardo Gimenez, Hugo Herbelin, Gerard Huet, Cesar Munoz, Chetan Murthy, et al. 1997. The Coq proof assistant reference manual: Version 6.1. (1997).

[3] Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. 2017. Verified models and reference implementations for the TLS 1.3 standard candidate. In *2017 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 483–502.

[4] Karthikeyan Bhargavan, Ioana Boureanu, Antoine Delignat-Lavaud, Pierre-Alain Fouque, and Cristina Onete. 2018. A formal treatment of accountable proxying over TLS. In *2018 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 799–816.

[5] Karthikeyan Bhargavan, Ioana Boureanu, Pierre-Alain Fouque, Cristina Onete, and Benjamin Richard. 2017. Content delivery over TLS: a cryptographic analysis of keyless SSL. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 1–16.

[6] Ferdinand Brasser, Urs Müller, Alexandra Dmitrienko, Kari Kostiainen, Srdjan Capkun, and Ahmad-Reza Sadeghi. 2017. Software grand exposure: SGX cache attacks are practical. In *11th USENIX Workshop on Offensive Technologies (WOOT'17)*.

[7] Kai Bu, Yutian Yang, Avery Laird, Jiaqing Luo, Yingjiu Li, and Kui Ren. 2018. What's (not) validating network paths: A survey. *arXiv preprint arXiv:1804.03385* (2018).

[8] David L Chaum. 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* 24, 2 (1981), 84–90.

[9] Chen Chen, Daniele E Asoni, David Barrera, George Danezis, and Adrain Perrig. 2015. HORNET: High-speed onion routing at the network layer. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1441–1454.

[10] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. 2017. A comprehensive symbolic analysis of TLS 1.3. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1773–1788.

[11] C.J.F. Cremers and S. Mauw. 2006. *Generalizing Needham-Schroeder-Lowe for Multi-Party Authentication*. CS-Report 06/04. Department of Mathematics and Computing Science, Eindhoven University of Technology.

[12] George Danezis and Ian Goldberg. 2009. Sphinx: A compact and provably secure mix format. In *Security and Privacy, 2009 30th IEEE Symposium on*. IEEE, 269–282.

[13] Xavier de Carné de Carnavalet and Mohammad Mannan. 2016. Killed by proxy: Analyzing client-end TLS interception software. In *Network and Distributed System Security Symposium*.

[14] Tim Dierks and Eric Rescorla. 2008. *The transport layer security (TLS) protocol version 1.2*. Technical Report. NIST.

[15] Danny Dolev and Andrew Chi-Chih Yao. 1983. On the security of public key protocols. *IEEE Trans. Information Theory* 29, 2 (1983), 198–207.

[16] Zakir Durumeric, Zane Ma, Drew Springall, Richard Barnes, Nick Sullivan, Elie Bursztein, Michael Bailey, J. Alex Halderman, and Vern Paxson. 2017. The security impact of HTTPS interception. In *Network and Distributed Systems Symposium*.

[17] Joseph A Goguen and José Meseguer. 1992. Order-sorted algebra I: Equational deduction for multiple inheritance, overloading, exceptions and partial operations. *Theoretical Computer Science* 105, 2 (1992), 217–273.

[18] Marcus Hähnel, Weidong Cui, and Marcus Peinado. 2017. High-resolution side channels for untrusted operating systems. In *2017 USENIX Annual Technical Conference (USENIX ATC'17)*. 299–312.

[19] Yih-Chun Hu, Adrian Perrig, and David B Johnson. 2005. Ariadne: A secure on-demand routing protocol for ad hoc networks. *Wireless networks* 11, 1-2 (2005), 21–38.

[20] Jeff Jarmoc and DSCT Unit. 2012. SSL/TLS interception proxies and transitive trust. *Black Hat Europe* (2012).

[21] Chris Karlof and David Wagner. 2003. Secure routing in wireless sensor networks: Attacks and countermeasures. *Ad hoc networks* 1, 2-3 (2003), 293–315.

[22] Tiffany Hyun-Jin Kim, Cristina Basescu, Limin Jia, Soo Bum Lee, Yih-Chun Hu, and Adrian Perrig. 2014. Lightweight source authentication and path validation. In *Proceedings of the 2014 ACM conference on SIGCOMM*. 271–282.

[23] Tobias Klenze, David Basin, and Christoph Sprenger. 2021. Formal Verification of Secure Forwarding Protocols. In *34th IEEE Computer Security Foundations Symposium (CSF 2021)*.

[24] Arvind Kumar, Rakesh Matam, and Shailendra Shukla. 2016. Impact of packet dropping attacks on RPL. In *2016 Fourth International Conference on Parallel, Distributed and Grid Computing (PDGC)*. IEEE, 694–698.

[25] Hyunwoo Lee, Zach Smith, Junghwan Lim, Gyeongjae Choi, Selin Chun, Taejoong Chung, and Ted Taekyoung Kwon. 2019. maTLS: How to Make TLS middlebox-aware?. In *NDSS*.

[26] Jie Li, Rongmao Chen, Jinshu Su, Xinyi Huang, and Xiaofeng Wang. 2019. ME-TLS: Middlebox-Enhanced TLS for Internet-of-things Devices. *IEEE Internet of Things Journal* (2019).

[27] Giulio Malavolta, Pedro Moreno-Sanchez, Clara Schneidewind, Aniket Kate, and Matteo Maffei. 2019. Anonymous Multi-Hop Locks for Blockchain Scalability and Interoperability. In *NDSS*.

[28] Simon Meier, Benedikt Schmidt, Cas Cremers, and David Basin. 2013. The TAMARIN prover for the symbolic analysis of security protocols. In *International Conference on Computer Aided Verification*. Springer, 696–701.

[29] Sebastian Nanz and Chris Hankin. 2006. A framework for security analysis of mobile wireless networks. *Theoretical Computer Science* 367, 1-2 (2006), 203–227.

[30] Jad Naous, Michael Walfish, Antonio Nicolosi, David Mazières, Michael Miller, and Arun Seehra. 2011. Verifying and enforcing network paths with ICING. In *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*. 1–12.

[31] David Naylor, Richard Li, Christos Gkantsidis, Thomas Karagiannis, and Peter Steenkiste. 2017. And Then There Were More: Secure Communication for More Than Two Parties. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*. ACM, 88–100.

[32] David Naylor, Kyle Schomp, Matteo Varvello, Ilias Leontiadis, Jeremy Blackburn, Diego López, Konstantina Papagiannaki, Pablo Rodriguez Rodriguez, and Peter Steenkiste. 2015. Multi-context TLS (mcTLS): Enabling secure in-network functionality in TLS. In *ACM SIGCOMM Computer Communication Review*, Vol. 45. ACM, 199–212.

[33] Dmytro Piatkivskyi and Mariusz Nowostawski. 2018. Split Payments in Payment Networks. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*. Springer, 67–75.

[34] Joseph Poon and Thaddeus Dryja. 2016. The bitcoin lightning network: Scalable off-chain instant payments. *See https://lightning. network/lightning-network-paper. pdf* (2016).

[35] Eric Rescorla. 2018. *The transport layer security (TLS) protocol version 1.3*. Technical Report. NIST.

[36] Elias Rohrer, Jann-Frederik Laß, and Florian Tschorsch. 2017. Towards a Concurrent and Distributed Route Selection for Payment Channel Networks. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*. Springer, 411–419.

[37] Kimaya Sanzgiri, Bridget Dahill, Brian Neil Levine, Clay Shields, and Elizabeth M Belding-Royer. 2002. A secure routing protocol for ad hoc networks. In *10th IEEE International Conference on Network Protocols, 2002. Proceedings*. IEEE, 78–87.

[38] Timothy J. Seppala. 2015. New Lenovo PCs shipped with factory-installed adware. (2015). https://www.engadget.com/2015/02/19/lenovo-superfish-adware-preinstalled/

[39] Yixin Sun, Anne Edmundson, Laurent Vanbever, Oscar Li, Jennifer Rexford, Mung Chiang, and Prateek Mittal. 2015. {RAPTOR}: Routing attacks on privacy in tor. In *24th {USENIX} Security Symposium ({USENIX} Security 15)*. 271–286.

[40] P. Syverson, R. Dingledine, and N. Mathewson. 2004. Tor: The second generation onion router. In *Usenix Security*.

[41] The Tamarin Team. 2018. *The Tamarin User Manual*. https://tamarin-prover.github.io/manual/tex/tamarin-manual.pdf

[42] Giorgos Tsirantonakis, Panagiotis Ilia, Sotiris Ioannidis, Elias Athanasopoulos, and Michalis Polychronakis. 2018. A Large-scale Analysis of Content Modification by Open HTTP Proxies. In *Network and Distributed System Security Symposium (NDSS)*.

[43] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx. 2018. Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution. In *27th USENIX Security Symposium (USENIX Security'18)*. 991–1008.

[44] Jo Van Bulck, Nico Weichbrodt, Rüdiger Kapitza, Frank Piessens, and Raoul Strackx. 2017. Telling your secrets without page faults: Stealthy page table-based attacks on enclaved execution. In *26th USENIX Security Symposium (USENIX Security'17)*. 1041–1056.

[45] Louis Waked, Mohammad Mannan, and Amr Youssef. 2018. To Intercept or Not to Intercept: Analyzing TLS Interception in Network Appliances. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*. ACM, 399–412.

[46] Thomas Watteyne, Antonella Molinaro, Maria Grazia Richichi, and Mischa Dohler. 2010. From manet to ietf roll standardization: A paradigm shift in wsn routing protocols. *IEEE Communications Surveys & Tutorials* 13, 4 (2010), 688–707.

[47] Fuyuan Zhang, Limin Jia, Cristina Basescu, Tiffany Hyun-Jin Kim, Yih-Chun Hu, and Adrian Perrig. 2014. Mechanized network origin and path authenticity proofs. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and*

Communications Security. 346–357.