

Provably Secure and Practical Onion Routing

Michael Backes Saarland University and MPI-SWS Saarbrücken, Germany backes@cs.uni-saarland.de	Ian Goldberg University of Waterloo Waterloo, Canada iang@cs.uwaterloo.ca	Aniket Kate MPI-SWS Saarbrücken, Germany aniket@mpi-sws.org	Esfandiar Mohammadi Saarland University Saarbrücken, Germany mohammadi@cs.uni-saarland.de
--	--	--	--

Abstract—The onion routing network Tor is undoubtedly the most widely employed technology for anonymous web access. Although the underlying onion routing (OR) protocol appears satisfactory, a comprehensive analysis of its security guarantees is still lacking. This has also resulted in a significant gap between research work on OR protocols and existing OR anonymity analyses. In this work, we address both issues with onion routing by defining a provably secure OR protocol, which is practical for deployment in the next generation Tor network.

We start off by presenting a security definition (an ideal functionality) for the OR methodology in the universal composability (UC) framework. We then determine the exact security properties required for OR cryptographic primitives (onion construction and processing algorithms, and a key exchange protocol) to achieve a provably secure OR protocol. We show that the currently deployed onion algorithms with slightly strengthened integrity properties can be used in a provably secure OR construction. In the process, we identify the concept of predictably malleable symmetric encryptions, which might be of independent interest. On the other hand, we find the currently deployed key exchange protocol to be inefficient and difficult to analyze and instead show that a recent, significantly more efficient, key exchange protocol can be used in a provably secure OR construction.

In addition, our definition greatly simplifies the process of analyzing OR anonymity metrics. We define and prove forward secrecy for the OR protocol, and realize our (white-box) OR definition from an OR black-box model assumed in a recent anonymity analysis. This realization not only makes the analysis formally applicable to the OR protocol but also identifies the exact adversary and network assumptions made by the black box model.

Keywords—onion routing; security proof; universal composability

I. INTRODUCTION

Over the last few years the onion routing (OR) network Tor [28] has emerged as a successful technology for anonymous web browsing. It currently employs more than two thousand dedicated relays, and serves hundreds of thousands of users across the world. Despite its success, the existing Tor network still lacks a rigorous security analysis, as its security properties have neither been formalized cryptographically nor proven. (See [3], [11], [23] for previous attempts and their shortcomings.) In this paper, we define security for the third-generation OR protocol Tor, and construct a provably secure and practical OR protocol.

An OR network consists of a set of routers or OR nodes that relay traffic, a large set of users, and directory servers that provide routing information for the OR nodes to the users. A user (say Alice) constructs a *circuit* by choosing a small sequence of (usually three) OR nodes, where the chosen nodes route Alice’s traffic over the path formed. The crucial property of an OR protocol is that a node in a circuit can determine no circuit nodes other than its predecessor and its successor. Alice sends data over the constructed circuit by sending the first OR node a message wrapped in multiple layers of symmetric encryption (one layer per node), called an *onion*, using symmetric keys agreed upon during an initial *circuit construction* phase. Consequently, given a public-key infrastructure (PKI), cryptographic challenges in onion routing are to securely agree upon such symmetric keys, and then to use the symmetric keys to achieve confidentiality and integrity.

In the first generation onion routing [25], circuits are constructed in a single pass. However, the scalability issues while pursuing forward secrecy [7] in the single-pass construction prompted Dingledine, Mathewson and Syverson [9] to use a telescoping approach for the next-generation OR protocol Tor. In this telescoping approach, they employed a forward secret, *multi-pass* key agreement protocol called the Tor authentication protocol (TAP) to negotiate a symmetric session key between user Alice and a node. Goldberg [13] presented a security proof for individual instances of TAP. The security of TAP, however, does not automatically imply the security of the Tor protocol. (For a possible concurrent execution attack, see [30].) The Tor protocol constitutes a sequential execution of multiple TAP instances as well as onion construction and processing algorithms, and thus its security has to be analyzed in a composability setting.

In this direction, Camenisch and Lysyanskaya [3] defined an anonymous message transmission protocol in the universal composability (UC) framework, and presented a protocol construction that satisfies their definition. They motivated their choice of the UC framework for a security definition by its versatility as well as its appropriateness for capturing protocol compositions. However, Feigenbaum, Johnson and Syverson [11], [12] observe that the protocol definition presented by Camenisch and Lysyanskaya [3] does not correspond to the OR methodology, and a rigorous

security analysis of an OR protocol still remains an unsolved problem.

Studies on OR anonymity such as [11], [23], [26] assume simplified OR black-box models to perform an analysis of the anonymity guarantees of these models. Due to the complexity of an OR network’s interaction with the network and the adversary, such black-box models are not trivially realized by deployed OR networks, such as Tor. As a result, there is a gap between deployed OR protocols and anonymity analysis research that has to be filled.

A. Our Contributions

Our contribution is threefold. First, we present a security definition for the OR methodology as an ideal functionality \mathcal{F}_{OR} in the UC framework. This ideal functionality in particular gives appropriate considerations to the goals of various system entities. After that, we identify and characterize which cryptographic primitives constitute central building blocks of onion routing, and we give corresponding security definitions: a one-way authenticated key exchange (1W-AKE) primitive, and onion construction and processing algorithms. We then describe an OR protocol Π_{OR} that follows the current Tor specification and that relies on these building blocks as black boxes. We finally show that Π_{OR} is secure in the UC framework with respect to \mathcal{F}_{OR} , provided that these building blocks are instantiated with secure realizations (according to their respective security definitions).

Second, we present a practical OR protocol by instantiating Π_{OR} with the following OR modules: a 1W-AKE protocol ntor [14], employed onion construction and processing algorithms in Tor with a slightly enhanced integrity mechanism. We show that these instantiations fulfill the security definitions of the individual building blocks that we identified before. This yields the first practical and provably secure OR protocol that follows the Tor specification. As part of these proofs, we identify a novel security definition of symmetric encryption notion we show to be sufficient for showing Π_{OR} secure. This notion strictly lies between CPA-security and CCA-security and characterizes stateful deterministic countermode encryptions. We call this notion *predictably malleable encryptions*, which might be of an independent interest.

Third, we illustrate the applicability of the abstraction \mathcal{F}_{OR} by introducing the first cryptographic definition of forward circuit secrecy for onion routing, which might be of independent interest. We utilize the abstraction \mathcal{F}_{OR} and the UC composability theorem for proving that Π_{OR} satisfies forward circuit secrecy by means of a simple proof. As a second application, we close the gap between the OR black-box model, prevalently used in anonymity analyses [11], [12], [23], [26], and a cryptographic model (Π_{OR}) of onion routing. Again, we utilize our abstraction \mathcal{F}_{OR} and the UC composability theorem for proving that against local, static attackers the recent analysis of the OR black-box model [12]

also applies to our OR protocol Π_{OR} instantiated with secure core building blocks.

Compared to previous work [3], we construct an OR circuit interactively in multiple passes, whereas previous work did not consider circuit construction at all, and hence does not model the widely used Tor protocol. The previous approach, and even single-pass circuit construction in general, restricts the protocol to eventual forward secrecy, while a multi-pass circuit construction ensures forward secrecy immediately after the circuit is closed. Second, we show that their hop-to-hop integrity verification is not mandatory, and that an end-to-end integrity verification suffices for onion routing. Finally, they do not consider backward messages (from web-servers to Alice), and their onion wrapping and unwrapping algorithms also do not work in the backward direction.

There has also been work on universally composable Mix-Nets by Wikström [29]. That work has some similarities to our work, but it only considers Mix-Nets, e.g., it does not need to cope with circuits and sessions.

Another important approach for analyzing onion routing has been conducted by Feigenbaum, Johnson, and Syverson [10]. In contrast to our work, the authors analyze an I/O automaton that use idealized encryption, pre-shared keys, and assume that every party only constructs one circuit to one destination. Moreover, the result in that work only holds in the stand-alone model against a local attackers whereas our result holds in the UC model against global and partially global attackers. In particular, by the UC composability theorem our result even holds with arbitrary protocols surrounding and against an attacker that controls parts of the network.

Outline of the Paper. Section II provides background information relevant to onion routing, 1W-AKE, and the UC framework. Section III, presents our security definition for onion routing. Section IV, presents cryptographic definitions for predictably malleable encryptions and secure onion construction and processing algorithms. Section V, states that given a set of secure OR modules we can construct a secure OR protocol. Section VII utilizes our security definition to analyze some security and anonymity properties of onion routing. Finally, we discuss some further interesting directions in Section VIII. In this work, many proofs have been omitted due to space constraints, which can be found in the full version [1].

II. BACKGROUND

In this paper, we often omit the security parameter κ when calling an algorithm A ; i.e., we abbreviate $A(1^\kappa, x)$ by $A(x)$. We write $y \leftarrow A(x)$ for the assignment of the result of $A(x)$ to a variable y , and we write $y \stackrel{\$}{\leftarrow} S$ for the assignment of a uniformly chosen element from S to y . For a given security parameter κ , we assume a message space $M(\kappa)$ that

is disjoint from the set of onions. We assume a distinguished error message \perp ; in particular, \perp is not in the message space. For some algorithms, we write $Alg(a, b, c, [d])$ and mean that the argument d is optional. Finally, for stateful algorithms, we write $y \leftarrow A(x)$ but we actually mean $(y, s') \leftarrow A(x, s)$, where s' is used in the next invocation of A as a state, and s is the stored state from the previous invocation. We assume that for all algorithms $s \in \{0, 1^\kappa\}$.

A. Onion Routing Circuit Construction

In the original Onion Routing project [16], [17], [25], [27], circuits were constructed in a single pass. However, such a single-pass circuit construction does not provide *forward secrecy*: if an adversary corrupts a node and obtains the private key, the adversary can decrypt all of the node’s past communication. Although changing the public/private key pairs for all OR nodes after a predefined interval is a possible solution (*eventual forward secrecy*), this solution does not scale to realistic OR networks such as Tor, since at the start of each interval every user has to download a new set of public keys for all the nodes.

Pairing-based onion routing (PB-OR) [21] and certificate-less onion routing (CL-OR) [6] attempt to provide better single-pass constructions. However, both approaches do not yield satisfactory solutions: CL-OR suffers from the same scalability issues as the original OR protocol [20]; PB-OR requires a distributed private-key generator [19] that may lead to inefficiency in practice.

Another problem with the single-pass approach is its intrinsic restriction to *eventual forward secrecy* [22]; i.e., if the current private key is leaked, then past sessions remain secret only if their public and private keys have expired. A desirable property is that all past sessions that are closed remain secret even if the private key is leaked; such a property is called *immediate forward secrecy*.

In the current Tor protocol, circuits are constructed using a multi-pass approach that is based on TAP. The idea is to use the private key only for establishing a temporary session key in a key exchange protocol. Together with the private key, additional temporary (random) values are used for establishing the key such that knowing the private key does not suffice for reconstructing the session key. These temporary values are erased immediately after the session key has been computed. This technique achieves immediate forward secrecy in multi-pass constructions, which however was never formally defined or proven before.

The multi-pass approach incurs additional communication overhead. However, in practice, almost all Tor circuits are constructed for a circuit length of $\ell = 3$, which merely causes an overhead of six additional messages.¹ With this small overhead, the multi-pass circuit construction is the preferred choice in practice, due to its improved forward secrecy

¹The overhead reduces to four additional messages if we consider the “CREATE_FAST” option available in Tor.

guarantees. Consequently, for our OR security definition we consider a multi-pass circuit construction as in Tor.

B. One-Way Authenticated Key Exchange – 1W-AKE

In a multi-pass circuit construction, a session key is established via a Diffie–Hellman key exchange. However, the precise properties required of this protocol were not formalized until recently. Goldberg, Stebila and Ustaoglu [14] formalized the concept of 1W-AKE, presented an efficient instantiation, and described its utility towards onion routing. We review their work here and we refer the readers to [14] for a detailed description.

An authenticated key exchange (AKE) protocol establishes an authenticated and confidential communication channel between two parties. Although AKE protocols in general aim for key secrecy and mutual authentication, there are many practical scenarios such as onion routing where mutual authentication is undesirable. In such scenarios, two parties establish a private shared session key, but only one party authenticates to the other. In fact, as in Tor, the unauthenticated party may even want to preserve its anonymity. Their 1W-AKE protocol constitutes this precise primitive.

The 1W-AKE protocol consists of three procedures: *Initiate*, *Respond*, and *ComputeKey*. With procedure *Initiate*, Alice (or her onion proxy) generates and sends an authentication challenge to the server (an OR node). The OR node responds to the challenge by running the *Respond* procedure, and returning the authentication response. The onion proxy (OP) then runs the *ComputeKey* procedure over the received response to authenticate the OR node and compute the session key.

The security of a 1W-AKE is defined by means of a *challenger* that represents all honest parties. The attacker is then allowed to query this challenger. If the attacker is not able to distinguish a fresh session key from a randomly chosen session key, we say that the 1W-AKE is *secure*. This challenger is constructed in a way that security of the 1W-AKE implies one-way authentication of the responding party.

For the definition of *one-way anonymity* we introduce a proxy, called the anonymity challenger, that relays all messages from and to the usual 1W-AKE challenger except for a challenge party C . The attacker can choose two challenge parties, out of which the anonymity challenger randomly picks one, say i^* . Then, the anonymity challenger relays all messages that are sent to C to P_{i^*} (via the 1W-AKE challenger).

In the one-way anonymity experiment, the adversary can issue the following queries to the challenger C . All other queries are simply relayed to the 1W-AKE challenger. The session Ψ^* denotes the challenge session. The two queries are for activation and communication during the test session. We say that a 1W-AKE is *one-way anonymous* if the attacker

cannot guess which party has been guessed with more than $1/2 + \mu(\kappa)$ probability, where μ is a negligible function.

In terms of instantiation, Goldberg et al. showed that an AKE protocol suggested for Tor—the fourth protocol in [24]—can be attacked, leading to an adversary determining all of the user’s session keys. They then fixed the protocol (see Figure 13) and proved that the fixed protocol (ntor) satisfies the formal properties of 1W-AKE. In our OR analysis, we use their formal definition and their fixed protocol.

C. The UC Framework: An Overview

The UC framework is designed to enable a modular analysis of security protocols. In this framework, the security of a protocol is defined by comparing it with a setting in which all parties have a direct and private connection to a trusted machine that computes the desired functionality. As an example consider an authenticated channel between Alice and Bob with a passive attacker. In the real world Alice would call a protocol that signs the message m to be communicated and sends the signed message over the network such that Bob would verify the signature. In the setting with a trusted machine T , however, Alice sends the message m directly to T ; T notifies the attacker about m , and T directly sends m to Bob. This trusted machine is called the *ideal functionality*.

Security in the UC framework is defined as follows: a protocol π UC-realizes an ideal functionality \mathcal{F} if for all probabilistic poly-time (PPT) attackers \mathcal{A} there is a PPT simulator S such that no PPT machine can distinguish an interaction with π and \mathcal{A} from an interaction with \mathcal{F} and S . The distinguisher is connected to the protocol and the attacker (or the simulator).

In contrast to typical UC proofs, our attacker model considers a more fine-grained network topology. Typically, a global attacker is assumed in UC; however, as we also want to be able to argue about local attackers, we prove our result for partially global attackers, i.e., in particular also for completely global attackers. A network over which the attacker does not have full control is modelled by a network functionality $\mathcal{F}_{\text{NET}^q}$ in which the attacker can adaptively compromise up to q links between honest onion routers. This network functionality is a global setup assumption; consequently, we have to consider the generalized UC framework (GUC) by Canetti, Dodis, Pass, and Walfish [5].² Throughout this work, if we say that a protocol ρ UC realizes a protocol π we actually mean that ρ GUC realizes π . (For a thorough definition of GUC, we refer to [5].)

D. The OR Protocol

We describe an OR protocol Π_{OR} that follows the Tor specification [8]. We do not present the cryptographic algo-

²The authors show that composability also holds true in the presence of global functionalities as long as the environment has access to these functionalities, i.e., they are not simulated by the simulator.

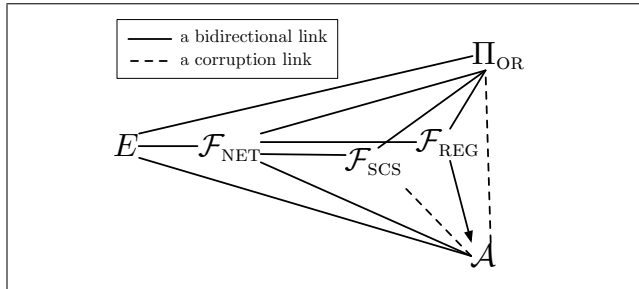


Figure 1. Overview of the set-up

rithms, e.g., wrapping and unwrapping onions, in this section but only present the skeleton of the protocol. A thorough characterization of these cryptographic algorithms follows in Section IV.

We describe our protocols using pseudocode and assume that a node maintains a state for every execution and responds (changes the state and/or sends a message) upon receiving a message as per its current state. In Figure 1, we give an overview of the setup that we consider.

As an attacker model we consider a partially global attacker in contrast to the global attacker that is typically used in UC analyses. For modelling a partially global attacker, we introduce an ideal functionality $\mathcal{F}_{\text{NET}^q}$ that allows the attacker to compromise at most q links.

There are two types of messages that the protocol generates and processes: the first type contains *input actions*, which carry inputs to the protocol from the user (Alice), and *output actions*, which carry outputs of the protocol to Alice. The second message type is a point-to-point *network message* (a cell in the OR literature), which is to be delivered by one protocol node to another. To enter a wait state, a thread may execute a command of the form **wait for** a network message.

With this methodology, we are able to effortlessly extract an OR protocol (Π_{OR}) from the Tor specification by categorizing actions based on the OR cell types (see Figure 2). For ease of exposition, we only consider Tor cells that are cryptographically important and relevant from the security definitional perspective. In particular, we consider create, created and destroy cells among control cells, and data, extend and extended cells among relay cells. We also include two input messages cc and send, where Alice uses cc to create OR circuits and uses send to send messages m over already-created circuits. We do not consider streams and the SOCKS interface in Tor as they are extraneous to the basic OR methodology. We unify instructions for an OP (onion proxy) node and an OR node for the simplicity of discussion. Moreover, for the sake of brevity, we restrict ourselves to messages $m \in M(\kappa)$ that fit exactly in one cell. It is straight-forward to extend our result to a protocol that accepts larger messages. The only difference is that the onion proxy and the exit node divide messages into smaller

```

upon an input (setup):
  Generate an asymmetric key pair  $(sk, pk) \leftarrow G$ .
  send a cell (register,  $P, pk$ ) to the  $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$  functionality
  wait for a cell (registered,  $\langle P_j, pk_j \rangle_{j=1}^{\mathcal{N}}$ ) from  $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ 
  output (ready,  $\mathcal{N} = \langle P_j \rangle_{j=1}^{\mathcal{N}}$ )
upon an input (cc,  $\mathcal{P} = \langle P, \langle P_j \rangle_{j=1}^{\ell} \rangle$ ):
  store  $\mathcal{P}$  and  $\mathcal{C} \leftarrow \langle P \rangle$ ; call  $\text{ExtendCircuit}(\mathcal{P}, \mathcal{C})$ 
upon an input (send,  $\mathcal{C} = \langle P \xleftrightarrow{\text{cid}_1} P_1 \xleftrightarrow{\text{cid}_2} \dots P_\ell \rangle, m$ ):
  if  $\text{Used}(\text{cid}_1) < \text{ttl}_{\mathcal{C}}$  then
    look up the keys  $(\langle k_j \rangle_{j=1}^{\ell})$  for  $\text{cid}_1$ 
     $O \leftarrow \text{WrOn}(m, \langle k_j \rangle_{j=1}^{\ell})$ ;  $\text{Used}(\text{cid}_1)++$ 
    send a cell  $(\text{cid}_1, \text{relay}, O)$  to  $P_1$  over  $\mathcal{F}_{\text{SCS}}$ 
  else
    call  $\text{DestroyCircuit}(\mathcal{C}, \text{cid}_1)$ ; output (destroyed,  $\mathcal{C}, m$ )
upon receiving a cell  $(\text{cid}, \text{create}, X)$  from  $P_i$  over  $\mathcal{F}_{\text{SCS}}$ :
   $\langle Y, k_{\text{new}} \rangle \leftarrow \text{Respond}(pk_P, sk_P, X)$ 
  store  $\mathcal{C} \leftarrow \langle P_i \xleftrightarrow{\text{cid}, k_{\text{new}}} P \rangle$ 
  send a cell  $(\text{cid}, \text{created}, Y, t)$  to  $P_i$  over  $\mathcal{F}_{\text{SCS}}$ 
upon receiving a cell  $(\text{cid}, \text{created}, Y, t)$  from  $P_i$  over  $\mathcal{F}_{\text{SCS}}$ :
  if  $\text{prev}(\text{cid}) = (P', \text{cid}', k')$  then
     $O \leftarrow \text{WrOn}(\langle \text{extended}, Y, t \rangle, k')$ 
    send a cell  $(\text{cid}', \text{relay}, O)$  to  $P'$  over  $\mathcal{F}_{\text{SCS}}$ 
  else if  $\text{prev}(\text{cid}) = \perp$  then
     $k_{\text{new}} \leftarrow \text{ComputeKey}(pk_i, Y, t)$ 
    update  $\mathcal{C}$  with  $k_{\text{new}}$ ; call  $\text{ExtendCircuit}(\mathcal{P}, \mathcal{C})$ 
upon receiving a cell  $(\text{cid}, \text{relay}, O)$  from  $P_i$  over  $\mathcal{F}_{\text{SCS}}$ :
  if  $\text{prev}(\text{cid}) = \perp$  then
    if  $\text{getkey}(\text{cid}) = (k_j)_{j=1}^{\ell}$  then
      (type,  $m$ ) or  $O \leftarrow \text{UnwrOn}(O, (k_j)_{j=1}^{\ell})$ 
       $(P', \text{cid}') \text{ or } \perp \leftarrow \text{next}(\text{cid})$ 
    else if  $\text{prev}(\text{cid}) = (P', \text{cid}', k')$  then
       $O \leftarrow \text{WrOn}(O, k')$  /* a backward onion */
  switch (type)
  case extend:
    get  $\langle P_{\text{next}}, X \rangle$  from  $m$ ;  $\text{cid}_{\text{next}} \xleftarrow{\$} \{0, 1\}^{\kappa}$ 
    update  $\mathcal{C} \leftarrow \langle P_i \xleftrightarrow{\text{cid}, k} P \xleftrightarrow{\text{cid}_{\text{next}}} P_{\text{next}} \rangle$ 
    send a cell  $(\text{cid}_{\text{next}}, \text{create}, X)$  to  $P_{\text{next}}$  over  $\mathcal{F}_{\text{SCS}}$ 
  case extended:
    get  $\langle Y, t \rangle$  from  $m$ ; get  $P_{\text{ex}}$  from  $(\mathcal{C}, \mathcal{P})$ 
     $k_{\text{ex}} \leftarrow \text{ComputeKey}(pk_{\text{ex}}, Y, t)$ 
    update  $\mathcal{C}$  with  $(k_{\text{ex}})$ ; call  $\text{ExtendCircuit}(\mathcal{P}, \mathcal{C})$ 
  case data:
    if  $(P = \text{OP})$  then output (received,  $\mathcal{C}, m$ )
    else if  $m = (S, m')$ 
      generate or lookup the unique  $\text{sid}$  for  $\text{cid}$ 
      send  $(P, S, \text{sid}, m')$  to  $\mathcal{F}_{\text{NET}^q}$ 
    case corrupted: /*corrupted onion*/
      call  $\text{DestroyCircuit}(\mathcal{C}, \text{cid})$ 
    case default: /*encrypted forward/backward onion*/
      send a cell  $(\text{cid}', \text{relay}, O)$  to  $P'$  over  $\mathcal{F}_{\text{SCS}}$ 
upon receiving a msg  $(\text{sid}, m)$  from  $\mathcal{F}_{\text{NET}^q}$ :
  get  $\mathcal{C} \leftarrow \langle P' \xleftrightarrow{\text{cid}, k} P \rangle$  for  $\text{sid}$ ;  $O \leftarrow \text{WrOn}(m, k)$ 
  send a cell  $(\text{cid}, \text{relay}, O)$  to  $P'$  over  $\mathcal{F}_{\text{SCS}}$ 
upon receiving a cell  $(\text{cid}, \text{destroy})$  from  $P_i$  over  $\mathcal{F}_{\text{SCS}}$ :
  call  $\text{DestroyCircuit}(\mathcal{C}, \text{cid})$ 

```

Figure 2. Π_{OR} : The OR Protocol for Party P pieces and recombine them in an appropriate way.

Function calls *Initiate*, *Respond* and *ComputeKey* cor-

```

 $\text{ExtendCircuit}(\mathcal{P} = \langle P_j \rangle_{j=1}^{\ell}, \mathcal{C} = \langle P \xleftrightarrow{\text{cid}_1, k_1} P_1 \xleftrightarrow{\text{cid}_2} \dots P_{\ell'} \rangle$ ):
  determine the next node  $P_{\ell'+1}$  from  $\mathcal{P}$  and  $\mathcal{C}$ 
  if  $P_{\ell'+1} = \perp$  then
    output (created,  $\langle P \xleftrightarrow{\text{cid}_1} P_1 \xleftrightarrow{\text{cid}_2} \dots P_{\ell'} \rangle$ )
  else
     $X \leftarrow \text{Initiate}(pk_{P_{\ell'+1}}, P_{\ell'+1})$ 
    if  $P_{\ell'+1} = P_1$  then
       $\text{cid}_1 \xleftarrow{\$} \{0, 1\}^{\kappa}$ 
      send a cell  $(\text{cid}_1, \text{create}, X)$  to  $P_1$  over  $\mathcal{F}_{\text{SCS}}$ 
    else
       $O \leftarrow \text{WrOn}(\{\text{extend}, P_{\ell'+1}, X\}, (k_j)_{j=1}^{\ell'})$ 
      send a cell  $(\text{cid}_1, \text{relay}, O)$  to  $P_1$  over  $\mathcal{F}_{\text{SCS}}$ 
 $\text{DestroyCircuit}(\mathcal{C}, \text{cid})$ :
  if  $\text{next}(\text{cid}) = (P_{\text{next}}, \text{cid}_{\text{next}})$  then
    send a cell  $(\text{cid}_{\text{next}}, \text{destroy})$  to  $P_{\text{next}}$  over  $\mathcal{F}_{\text{SCS}}$ 
  else if  $\text{prev}(\text{cid}) = (P_{\text{prev}}, \text{cid}_{\text{prev}})$  then
    send a cell  $(\text{cid}_{\text{prev}}, \text{destroy})$  to  $P_{\text{prev}}$  over  $\mathcal{F}_{\text{SCS}}$ 
  discard  $\mathcal{C}$  and all streams

```

Figure 3. Subroutines of Π_{OR} for Party P

respond to 1W-AKE function calls described in Section II-B. Function calls *WrOn* and *UnwrOn* correspond to the principal onion algorithms. *WrOn* creates a layered encryption of a payload (plaintext or onion) for given an ordered list of ℓ session keys for $\ell \geq 1$. *UnwrOn* removes ℓ layers of encryptions from an onion to output a plaintext or an onion given an input onion and a ordered list of ℓ session keys for $\ell \geq 1$. Moreover, onion algorithms also ensure end-to-end integrity. The cryptographic requirements for these onion algorithms are presented in Section IV-B.

Tor uses a centralized approach to determine valid OR nodes and distribute their public keys. Every OR node has to be registered in so-called directory servers, where each registration is checked by an administrator. These directory servers then distribute the list of valid OR nodes and the respective public keys. We abstract the key registration procedure by assuming that the directory servers expect a fixed set of parties upon setup. Formally, we model these directory servers as an ideal functionality $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$, which is basically defined as by Canetti [4] except that $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ rejects all parties that are not in \mathcal{N} and only sends the public keys around once all parties in \mathcal{N} registered.³ Tor does not guarantee any anonymity once these directory servers are compromised. Therefore, we concentrate on the case in which these directory servers cannot be compromised.⁴ As in Tor, we assume that the list of valid OR nodes is given to the directory servers from outside, in our case from the environment. However, for the sake of simplicity we assume that the OR list is only synchronized initially.

³Technically, we also extend $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ such that upon each (register, sid, v)-message, $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ notifies the attacker. And only after the attacker confirmed this message, $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ registers v with P .

⁴Formally, this ideal functionality $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ does not accept compromise-requests from the attacker.

In detail, we slightly extend the functionality as follows. $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ initially receives a list of OR nodes from the environment, waits for each of these parties for a public key, and distributes the list of OR nodes and their public keys as (registered, $\langle P_j, pk_j \rangle_{j=1}^n$). Each OR node, on the other hand, initially computes its long-term keys (sk, pk) and registers the public part at $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$. Then, the node waits to receive the message (registered, $\langle P_j, pk_j \rangle_{j=1}^n$) from $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ before declaring that it is ready for use.⁵

OPs develop circuits incrementally, one hop at a time, using the *ExtendCircuit* function defined in Figure 3. To create a new circuit, an OP sends a create cell to the first node, after calling the *Initiate* function of 1W-AKE; the first node responds with a created cell after running the *Respond* function. The OP then runs the *ComputeKey* function. To extend a circuit past the first node, the OP sends an extend relay cell after calling the *Initiate* function, which instructs the last node in the circuit to send a create cell to extend the circuit.

Circuits are identified by circuit IDs ($cid \in \{0, 1\}^\kappa$) that associate two consecutive circuit nodes. We denote circuit at a node P_i using the terminology $\mathcal{C} = P_{i-1} \xleftrightarrow{cid_i, k_i} P_i \xleftrightarrow{cid_{i+1}} P_{i+1}$, which says that P_{i-1} and P_{i+1} are respectively the predecessor and successor of P_i in a circuit \mathcal{C} . k_i is a session key between P_i and the OP, while the absence of k_{i+1} indicates that a session key between P_{i+1} and the OP is not known to P_i ; analogously the absence of a circuit id cid in that notation means that only the first circuit id is known, as for OP, for example. Functions *prev* and *next* on cid correspondingly return information about the predecessor or successor of the current node with respect to cid ; e.g., $next(cid_i)$ returns (P_{i+1}, cid_{i+1}) and $next(cid_{i+1})$ returns \perp . The OP passes on to Alice $\langle P \xleftrightarrow{cid_1} P_1 \xleftrightarrow{\dots} P_\ell \rangle$.

Within a circuit, the OP and the exit node use relay cells created using *WrOn* to tunnel end-to-end commands and connections. The exit nodes use some additional mechanisms (abstracting the streams used in Tor) to synchronize communication between the network and a circuit \mathcal{C} . We represent that using *sid*. With this auxiliary synchronization, end-to-end communication between OP and the exit node happens with a *WrOn* call with multiple session keys and a series of *UnwrOn* calls with individual session keys in the forward direction, and a series of *WrOn* calls with individual session keys, and finally a *UnwrOn* call with multiple session keys in the backward direction. Communication in the forward direction is initiated by a send message by Alice to the OP, while communication in the backward direction is initiated by a network message to the exit node. Cells are exchanged between OR nodes over a secure and authenticated channels, e.g., a TLS connection. We abstract such a channel in the UC framework by a functionality \mathcal{F}_{SCS}

⁵The functionality $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$ additionally answers upon a request retrieve with the full list of participants $\langle P_j, pk_j \rangle_{j=1}^n$.

upon receiving a msg (compromise, $N_{\mathcal{A}}$) from \mathcal{A} :
 set $compromised(P) \leftarrow true$ for every $P \in N_{\mathcal{A}}$
 set $b \leftarrow \frac{|N_{\mathcal{A}}|}{|N_{\text{OR}}|}$
upon an input (send, S) from the environment for party U :
 with probability b^2 , send (sent, U, S) to \mathcal{A}
 with probability $(1-b)b$, send (sent, $-, S$) to \mathcal{A}
 with probability $b(1-b)$, send (sent, $U, -$) to \mathcal{A}
 with probability $(1-b)^2$, send (sent, $-, -$) to \mathcal{A}

Figure 4. Black-box OR Functionality \mathcal{B}_{OR} [12]

as proposed by Canetti [4] with the only difference that \mathcal{F}_{SCS} does not output the leakage to the attacker but to $\mathcal{F}_{\text{NET}^q}$, i.e., the network functionality.⁶ We write $storeX \leftarrow v$ for either introducing a new variable X with value v , or assign a value v to a variable X in case X was not previously defined.

To tear down a circuit completely, an OR or OP sends a destroy cell to the adjacent nodes on that circuit with appropriate cid using the *DestroyCircuit* function defined in Figure 3. Upon receiving an outgoing destroy cell, a node frees resources associated with the corresponding circuit. If it is not the end of the circuit, it sends a destroy cell to the next node in the circuit. Once a destroy cell has been processed, the node ignores all cells for the corresponding circuit. Note that if an integrity check fails during *UnwrOn*, the destroy cells are sent in the forward and backward directions in a similar way.

In the Tor the OP has a time limit (of ten minutes) for each established circuit; thereafter, the OP constructs a new circuit. However, the UC framework does not provide a notion of time. We model such a time limit in the UC framework by only allowing a circuit to transport at most a constant number (say $tll_{\mathcal{C}}$) of messages measured using the *used* function call. Afterwards, the OP discards the circuit and establishes a fresh circuit.

E. An OR Black Box Model

Anonymity in a low-latency OR network does not only depend upon the security of the onions but also upon the magnitudes and distributions of users and their destination servers. In the OR literature, considerable efforts have been put towards measuring the anonymity of onion routing [10]–[12], [23], [26].

Feigenbaum, Johnson, and Syverson used for an analysis of the anonymity properties of onion routing an ideal functionality \mathcal{B}_{OR} [12]. This functionality emulates an I/O-automata model for onion routing from [10], [11]. Figure 4 presents this functionality \mathcal{B}_{OR} .

Let N_{OR} be the set of onion routers, and let $N_{\mathcal{A}}$ of those be eavesdropped, where $b = |N_{\mathcal{A}}|/|N_{\text{OR}}|$ defines the fraction of compromised nodes. It takes as input from each user U the identity of a destination S . For every such connection between a user and a destination, the functionality may

⁶As leakage function l for \mathcal{F}_{SCS} , we choose $l(m) := |m|$.

reveal to the adversary the identity of the user ($\text{sent}, U, -$) (i.e., the first OR router is compromised), the identity of the destination ($\text{sent}, -, S, [m]$) (i.e., the exit node is compromised), both ($\text{sent}, U, S, [m]$) (i.e., the first OR router and the exit node are compromised) or only a notification that something has been sent ($\text{sent}, -, -$) (i.e., neither the first OR router nor the exit node is compromised).

We stress that this functionality only abstracts an OR network against local attackers. As the distribution of the four cases only depends on the first and the last router being compromised but not on the probability that the attacker controls sensitive links between honest parties, \mathcal{B}_{OR} only models OR against local adversaries. As an example consider, the case in which the attacker only wiretaps the connection between the exit node and the server. In this case, the attacker is able to determine which message has been sent to whom, i.e., the abstraction needs to leak ($\text{sent}, -, S, [m]$); however, the probability of this event is c , where c is the fraction of observed links between honest onion routers and users and servers. Therefore, \mathcal{B}_{OR} cannot be used as an abstraction for onion routing against partially global attackers.

We actually present \mathcal{B}_{OR} in two variants. In the first variant \mathcal{B}_{OR} does not send an actual message but only a notification. This variant has been analyzed by Feigenbaum, Johnson, and Syverson. We additionally consider the variant in which \mathcal{B}_{OR} sends a proper message m . We denote these two variants by marking the message m as optional, i.e., as $[m]$.

In order to justify these OR anonymity analyses that consider an OR network as a black box, it is important to ascertain that these black boxes indeed model onion routing. In particular, it is important under which adversary and network assumptions these black boxes model and securely abstract real-world OR networks. In this work, we show that the black box \mathcal{B}_{OR} can be UC-realized by a simplified version of the Tor network.

III. SECURITY DEFINITION OF OR

In this section, we first describe our system and adversary model for all protocols that we analyze (Section III-A). Thereafter, we present a composable security definition of OR by introducing an ideal functionality (abstraction) \mathcal{F}_{OR} in the UC framework (Section III-B).

Tor was designed to guarantee anonymity even against partially global attackers, i.e., attackers that do not only control compromised OR nodes but also a portion of the network. Previous work, however, only analyzed local, static attackers [10]–[12], such as the abstraction \mathcal{B}_{OR} presented in Figure 4. In contrast, we analyze onion routing against partially global attackers. As our resulting abstraction \mathcal{F}_{OR} has to faithfully reflect that an active attacker can hold back all onions that it observes, \mathcal{F}_{OR} is naturally more complex than \mathcal{B}_{OR} .

A. System and Adversary Model

We consider a fully connected network of n parties $N = \{P_1, \dots, P_n\}$. For simplicity of presentation, we consider all parties to be OR nodes that also can function as OPs to create circuits and send messages. It is also possible to use our formulation to model separate user OPs that only send and receive messages but do not relay onions.

Tor has not been designed to resist against global attackers. Such an attacker is too strong for many practical purposes as it can simply break the anonymity of an OR protocol by holding back all but one onion and tracing that one onion through the network. However, in contrast to previous work, we do not only consider local attackers, which do not control more than the compromised OR routers, but also partially global attackers that control a certain portion of the network. Analogous to the network functionality \mathcal{F}_{SYN} proposed by Canetti [4], we model the network as an ideal functionality $\mathcal{F}_{\text{NET}^q}$, which bounds the number of attacker-controlled links to $q \in [0, \binom{n}{2}]$. For attacker-controlled links the messages are forwarded to the attacker; otherwise, they are directly delivered. In Section VII we show that previous black-box analyses of onion routing against local attackers applies to our setting as well by choosing $q := 0$. Let S represent all possible destination servers $\{S_1, \dots, S_\Delta\}$ which reside in the network abstracted by a network functionality $\mathcal{F}_{\text{NET}^q}$.

We stress that the UC framework does not provide a notion of time; hence, the analysis of timing attacks, such as traffic analysis, is not in the scope of this work.

Adaptive Corruptions. Forward secrecy [7] is an important property for onion routing. In order to analyze this property, we allow adaptive corruptions of nodes by the attacker \mathcal{A} . Such an adaptive corruption is formalized by a message compromise, which is sent to the respective party. Upon such a compromise message the internal state of that party is deleted and a long-term secret key sk for the node is revealed to the attacker. \mathcal{A} can then impersonate the node in the future; however, \mathcal{A} cannot obtain the information about its ongoing sessions. We note that this restriction arises due to the currently available security proof techniques and the well-known selective opening problem with symmetric encryptions [18], and the restriction is not specific to our constructions [2], [15]. We could also restrict ourselves to a static adversary as in previous work [3]; however, that would make an analysis of forward secrecy impossible.

B. Ideal Functionality

The presentation of the ideal functionality \mathcal{F}_{OR} is along the lines of the description OR protocol Π_{OR} from Section II-D. We continue to use the message-based state transitions from Π_{OR} , and consider sub-machines for all n nodes in the ideal functionality. To communicate with each other through messages and data structures, these sub-machines

```

upon an input (setup):
  draw a fresh handle  $h$ ; a set registered_flag  $\leftarrow true$ 
  store  $lookup(h) \leftarrow (dir, registered, \mathcal{N})$ 
  send  $(h, register, P)$  to  $\mathcal{A}$ 
  wait for a msg (dir, registered,  $\mathcal{N}$ ) via a handle
  output (ready,  $(P_j)_{j=1}^n = (ready, \mathcal{N})$ )
upon an input (cc,  $\mathcal{P} = \langle P, P_1, \dots, P_\ell \rangle$ ):
  store  $\mathcal{P}$  and  $\mathcal{C} \leftarrow \langle P \rangle$ ;  $ExtendCircuit(\mathcal{P}, \mathcal{C})$ 
upon an input (send,  $\mathcal{C} = \langle P \xleftrightarrow{cid_1} P_1 \iff \dots P_\ell \rangle, m$ ):
  if  $Used(cid_1) < ttl_{\mathcal{C}}$  then
     $Used(cid_1)++$ ;  $SendMessage(P_1, cid_1, relay, \langle data, m \rangle)$ 
  else
     $DestroyCircuit(\mathcal{C}, cid_1)$ ; output (destroyed,  $\mathcal{C}, m$ )
upon receiving a handle  $(P, P_{next}, h)$  from  $\mathcal{F}_{NET^q}$ :
  send  $(msg) \leftarrow lookup(h)$  to a receiving submachine  $P_{next}$ 
upon receiving a msg  $(P_i, cid, create)$  through a handle:
  store  $\mathcal{C} \leftarrow \langle P_i \xleftrightarrow{cid} P \rangle$ ;  $SendMessage(P_i, cid, create)$ 
upon receiving a msg  $(P_i, cid, created)$  through a handle:
  if  $prev(cid) = (P', cid')$  then
     $SendMessage(P', cid', relay, extended)$ 
  else if  $prev(cid) = \perp$  then
     $ExtendCircuit(\mathcal{P}, \mathcal{C})$ 
upon receiving a msg  $(P_i, cid, relay, O)$  through a handle:
  if  $prev(cid) = \perp$  then
    if  $next(cid) = \perp$  then
      get (type,  $m$ ) from  $O$ 
    else  $\{P', cid'\} \leftarrow next(cid)$ 
  else
     $(P', cid') \leftarrow prev(cid)$ 
  switch (type)
  case extend:
    get  $P_{next}$  from  $m$ ;  $cid_{next} \xleftarrow{\$} \{0, 1\}^\kappa$ 
    update  $\mathcal{C} \leftarrow \langle P_i \xleftrightarrow{cid} P \xleftrightarrow{cid_{next}} P_{next} \rangle$ 
     $SendMessage(P_{next}, cid_{next}, create)$ 
  case extended:
    update  $\mathcal{C}$  with  $P_{ex}$ ;  $ExtendCircuit(\mathcal{P}, \mathcal{C})$ 
  case data:
    if  $(P = OP)$  then output (received,  $\mathcal{C}, m$ )
    else if  $m = (S, m')$ 
      generate or lookup the unique  $sid$  for  $cid$ 
      send  $(P, S, sid, m')$  to  $\mathcal{F}_{NET^q}$ 
    case corrupted:  $!*corrupted onion*$ 
       $DestroyCircuit(\mathcal{C}, cid)$ 
    case default:  $!*encrypted forward/backward onion*$ 
       $SendMessage(P', cid', relay, O)$ 
upon receiving a msg  $(sid, m)$  from  $\mathcal{F}_{NET^q}$ :
  obtain  $\mathcal{C} = \langle P' \xleftrightarrow{cid} P \rangle$  for  $sid$ 
   $SendMessage(P', cid, relay, \langle data, m \rangle)$ 
upon receiving a msg  $(P_i, cid, destroy)$  through a handle:
   $DestroyCircuit(\mathcal{C}, cid)$ 
upon receiving a msg  $(P_i, P, h, [corrupt, T(\cdot)])$  from  $\mathcal{A}$ :
   $(message) \leftarrow lookup(h)$ 
  if corrupt = true then
     $message \leftarrow T(message)$ ; set  $corrupted(message) \leftarrow true$ 
    process  $message$  as if the receiving submachine was  $P$ 
upon receiving a msg (compromise,  $P$ ) from  $\mathcal{A}$ :
  set  $compromised(P) \leftarrow true$ 
  delete all local information at  $P$ 

```

Figure 5. The ideal functionality \mathcal{F}_{OR}^N (short \mathcal{F}_{OR}) for Party P

```

 $ExtendCircuit(\mathcal{P} = (P_j)_{j=1}^\ell, \mathcal{C} = \langle P \xleftrightarrow{cid_1} P_1 \iff \dots P_{\ell'} \rangle$ ):
  determine the next node  $P_{\ell'+1}$  from  $\mathcal{P}$  and  $\mathcal{C}$ 
  if  $P_{\ell'+1} = \perp$  then
    output (created,  $\mathcal{C}$ )
  else
    if  $P_{\ell'+1} = P_1$  then
       $cid_1 \xleftarrow{\$} \{0, 1\}^\kappa$ ;  $SendMessage(P_1, cid_1, create)$ 
    else
       $SendMessage(P_1, cid_1, relay, \{extend, P_{\ell'+1}\})$ 
 $DestroyCircuit(\mathcal{C}, cid)$ :
  if  $next(cid) = (P_{next}, cid_{next})$  then
     $SendMessage(P_{next}, cid_{next}, destroy)$ 
  else if  $prev(cid) = (P_{prev}, cid_{prev})$  then
     $SendMessage(P_{prev}, cid_{prev}, destroy)$ 
  discard  $\mathcal{C}$  and all streams
 $SendMessage(P_{next}, cid_{next}, cmd, [relay-type], [data])$ :
  create a  $msg$  for  $P_{next}$  from the input
  draw a fresh handle  $h$  and set  $lookup(h) \leftarrow msg$ 
  if  $compromised(P_{next}) = true$  then
     $P_{last}$  is the last node in the complete continuous
     $compromised$  path starting  $P_{next}$ 
    if  $(P_{last} = OP)$  or  $P_{last}$  is the exit node then
      send the entire msg to  $\mathcal{A}$ 
    else
      send  $\langle P, P_{next}, \dots, P_{last}, cid_{next}, cmd, h \rangle$  to  $\mathcal{A}$ 
  else
    send  $\langle P, P_{next}, h \rangle$  to the network

```

Figure 6. Subroutines of \mathcal{F}_{OR} for Party P

```

upon receiving a msg (observe,  $P, P_{next}$ ) from  $\mathcal{A}$ :
  set  $observedLink(P, P_{next}) \leftarrow true$ 
upon receiving a msg (compromise,  $S$ ) from  $\mathcal{A}$ :
  set  $compromised(S) \leftarrow true$ ; send  $\mathcal{A}$  all existing  $sid$ 
upon receiving a msg  $(P, P_{next}/S, m)$  from  $\mathcal{F}_{OR}$ :
  if  $P_{next}/S$  is a  $\mathcal{F}_{OR}$  node then
    if  $observedLink(P, P_{next}) = true$  then
      forward the msg  $(P, P_{next}, m)$  to  $\mathcal{A}$ 
    else
      reflect the msg  $(P, P_{next}, m)$  to  $\mathcal{F}_{OR}$ 
  else if  $P_{next}/S$  is a  $\mathcal{F}_{NET^q}$  server then
    if  $compromised(S) = true$  then
      forward the msg  $(P, S, m)$  to  $\mathcal{A}$ 
    else
      output  $(P, S, m)$ 
upon receiving a msg  $(P/S, P_{next}, m)$  from  $\mathcal{A}$ :
  forward the msg  $(P/S, P_{next}, m)$  to  $\mathcal{F}_{OR}$ 

```

Figure 7. The Network Functionality \mathcal{F}_{NET^q}

share a memory space in the functionality. The sub-machine pseudocode for the ideal functionality appears in Figure 5 and three subroutines are defined in Figure 6. As the similarity between pseudocodes for the OR protocol and the ideal functionality is obvious, rather than explaining the OR message flows again, we concentrate on the differences.

The only major difference between Π_{OR} and \mathcal{F}_{OR} is that cryptographic primitives such as message wrapping, unwrapping, and key exchange are absent in the ideal world; we do not have any keys in \mathcal{F}_{OR} , and the OR messages

$WrOn$ and $UnwrOn$ as well as the 1W-AKE messages $Initiate$, $Respond$, and $ComputeKey$ are absent.

The ideal functionality also abstracts the directory server and expects on the input/output interface of \mathcal{F}_{REG}^N (from the setting with Π_{OR}) an initial message with the list $\langle P_i \rangle_{i=1}^n$ of valid nodes. This initial message corresponds to the list of onion routers that have been approved by an administrator. We call the part of \mathcal{F}_{OR} that abstracts the directory servers dir . For the sake of brevity, we do not present the pseudocode of dir . Upon an initial message with a list $\langle P_i \rangle_{i=1}^n$ of valid nodes, dir waits for all nodes P_i ($i \in \{1, \dots, n\}$) for a message (register, P_i). Once all nodes registered, dir sends a message (registered, $\langle P_i \rangle_{i=1}^n$) with a list of valid and registered nodes to every party that registered, and to every party that sends a retrieve message to dir .

Messages from \mathcal{A} and \mathcal{F}_{NET^q} . In Figure 5 and Figure 7, we present the pseudocode for the attacker messages and the network functionality, respectively. For our basic analysis, we model an adversary that can control all communication links and servers in \mathcal{F}_{NET^q} , but cannot view or modify messages between parties due to the presence of the secure and authenticated channel. Therefore, sub-machines in the functionality store their messages in the shared memory, and create and send handles $\langle P, P_{next}, h \rangle$ for these messages \mathcal{F}_{NET^q} . The message length does not need to be leaked as we assume a fixed message size (for all $M(\kappa)$). Here, P is the sender, P_{next} is the receiver and h is a handle or a pointer to the message in the shared memory of the ideal functionality. In our analysis, all \mathcal{F}_{NET^q} messages flow to \mathcal{A} , which may choose to return these handles back to \mathcal{F}_{OR} through \mathcal{F}_{NET^q} at its own discretion. However, \mathcal{F}_{NET^q} also maintains a mechanism through *observedLink* flags for the non-global adversary \mathcal{A} . The adversary may also corrupt or replay the corresponding messages; however, these active attacks are always detected by the receiver due to the presence of a secure and authenticated channel between any two communicating parties and we need not model these corruptions.

The adversary can compromise a party P or server S by sending a compromise message to respectively \mathcal{F}_{OR} and \mathcal{F}_{NET^q} . For party P or server S , the respective functionality then sets the *compromised* tag to *true*. Furthermore, all input or network messages that are supposed to be visible to the compromised entity are forwarded to the adversary. In principle, the adversary runs that entity for the rest of the protocol and can send messages from that entity. In that case, it can also propagate corrupted messages which in Π_{OR} can only be detected during $UnwrOn$ calls at OP or the exit node. We model these corruptions using $corrupted(msg) = \{true, false\}$ status flags, where $corrupted(msg)$ status of messages is maintained across nodes until they reach end nodes. Furthermore, for every corrupted message, the adversary also provides a modification function $T(\cdot)$ as the

end nodes run by the adversary may continue execution even after observing a *corrupted* flag. In that case, $T(\cdot)$ captures the exact modification made by the adversary.

We stress that \mathcal{F}_{OR} does not need to reflect reroutings and circuit establishments initiated by the attacker, because the attacker learns, loosely speaking, no new information by rerouting onions.⁷ Similar to the previous work [3], a message is directly given to the adversary if all remaining nodes in a communication path are under adversary control.

IV. SECURE OR MODULES

We identify the core cryptographic primitives for a secure OR protocol. In this section, we present a cryptographic characterization of these core cryptographic primitives, which we call *secure OR modules*. Secure OR modules consist of two parts: first, secure onion algorithm, and second, a one-way authenticated key exchange primitive (1W-AKE), a notion recently introduced by Goldberg, Stebila, and Ustaoglu [14].

Onion algorithms typically use several layers of encryptions and possibly integrity mechanisms, such as message authentication codes. Previous attempts [3] for proving the security OR protocols use mechanisms to ensure hop-to-hop integrity, such as non-malleable encryption schemes. The widely-used Tor network, however, does not use hop-to-hop integrity but only end-to-end integrity. In the analysis of OR protocols with only end-to-end integrity guarantees, we also have to consider the cases in which the end node is compromised, thus no integrity check is performed at all. In order to cope with these cases, we identify a new notion of predictably malleable encryption schemes. Predictable malleability allows the attacker to change the ciphertexts but requires the resulting changes to the plaintext to be efficiently predictable given only the changes of the ciphertext. In Section IV-A we rigorously define the notion of *predictably malleable* encryption schemes.

Inspired by Section IV-A, we introduce in Section IV-B the notion of *secure onion algorithms*.

In the following definitions, we assume the PPT machines to actually be oracle machines. We write A^B to denote that A has oracle access to B .

A. Predictably Malleable Encryption

Simulation-based proofs often face their limits when dealing with malleable encryption schemes. The underlying problem is that malleability induces an essentially arbitrarily large number of possibilities to modify ciphertexts, and the simulator has no possibility to predict the resulting changes to the corresponding plaintext.

⁷More formally, the simulator can compute all responses for rerouting or such circuit establishments without requesting information from \mathcal{F}_{OR} because the simulator knows all long-term and session keys. The only information that the simulator does not have is the routing information, which the simulator gets in case of rerouting or circuit establishment.

<pre> upon (initialize) $k \leftarrow G(1^\eta)$ $s_d \leftarrow \varepsilon; s_e \leftarrow \varepsilon$ upon (encrypt, m) if $b = 0$ then $(c, s) \leftarrow E(0^{ m }, s_e, k)$ if $q(s_e) \neq \perp$ then $(d, u) \leftarrow q(s_e)$ $c \leftarrow d$ else if $b = 1$ then $(c, s) \leftarrow E(m, s_e, k)$ $q(s_e) \leftarrow (c, m)$ $s_e \leftarrow s; \text{respond } c$ </pre>	<pre> upon (decrypt, c) $(d, u) \leftarrow q(s_d)$ $T \leftarrow M(c, d)$ if $b = 0$ then if $q(s_d) = \perp$ then $(m, s) \leftarrow D(c, s_d, k)$ $q(s_d) \leftarrow (c, m)$ else if $q(s_d) \neq \perp$ then $m \leftarrow T(u)$ else if $b = 1$ then $(m, s) \leftarrow D(c, s_d, k)$ $s_d \leftarrow s; \text{respond } (m, T)$ </pre>
--	---

Figure 8. The IND-PM Challenger $\text{PM-Ch}_b^\varepsilon$

We characterize the property of predicting the changes to the plaintext merely given the modifications on the ciphertext. Along the lines of the IND-CCA definition for stateful encryption schemes, we define the notion of *predictably malleable* (IND-PM) encryption schemes.⁸ The attacker has access to an encryption and a decryption oracle, and either all encryption and decryption queries are honestly answered (the honest game) or all are faked (the faking game), i.e., $0^{|m|}$ is encrypted instead of a message m . In the faking game, the real messages are stored in some shared datastructure q , and upon a decryption query only look-ups in q are performed. The IND-PM challenger maintains a separate state, e.g., a counter, for encryption and decryption. These respective states are updated with each encryption decryption query.

In contrast to the IND-CCA challenger, the IND-PM challenger (see Figure 8) additionally stores the produced ciphertext together with the corresponding plaintext for each encryption query. Moreover, for each decryption call the challenger looks up the stored ciphertexts and messages. The honest decryption ignores the stored values and performs an honest decryption, but the faking decryption compares the stored ciphertext with the ciphertext from the query and tries to predict the modifications to the plaintext. Therefore, we require the existence of an efficiently computable algorithm M that outputs the description of an efficient transformation procedure T for the plaintext given the original ciphertext as well as the modified ciphertext.

Definition 1 (Predictable malleability): An encryption scheme $\mathcal{E} := (G, E, D)$ is IND-PM if there is a negligible function μ such that there is a deterministic polynomial-time algorithm M such that for all PPT attackers \mathcal{A}

$$\Pr[b' \stackrel{\$}{\leftarrow} \{0, 1\}, b \leftarrow \mathcal{A}(1^\kappa)^{\text{PM-Ch}_b^\varepsilon} : b = b'] \leq 1/2 + \mu(\kappa)$$

⁸The name predictable malleability is justified as it can be shown that every IND-CCA secure scheme is also IND-PM, and every IND-PM scheme in turn is IND-CPA secure. In Section VI-A, we present detCTR and state that it is IND-PM secure.

Moreover, we require that for all $m, c, s, k, k' \in \{0, 1\}^*$

$$\Pr[(c', s') \leftarrow E(m, k, s), (m', s'') \leftarrow D(c, k', s) : s' = s''] = 1$$

$\text{PM-Ch}_0^\varepsilon$ and $\text{PM-Ch}_1^\varepsilon$ are defined in Figure 8.

We stress that the definition implies a super-polynomial length for state-cycles; otherwise there is in the faking game at least one repeated state s for which the two encrypt queries output the same ciphertext for any two plaintexts.

In Section VI-A, we show that deterministic counter-mode is IND-PM.

B. Secure Onion Algorithms

We identify the onion wrapping (*Wrap*) and unwrapping (*Unwrap*) algorithms as central building blocks in onion routing. We identify four core properties of onion algorithms. The first property is *correctness*, i.e., if all parties behave honestly, the result is correct. The second property is the security of statefulness, coined *synchronicity*. It roughly states that whenever a wrapping and an unwrapping algorithms are applied to a message with unsynchronous states, the output is completely random. The third property is *end-to-end integrity*. The fourth property states that for all modifications to an onion the resulting changes in the ciphertext are predictable. We this property *predictable malleability*.

Onion Correctness. The first property of secure onion algorithms is *onion correctness*. It states that honest wrapping and unwrapping results in the same message. Moreover, the correctness states that whenever the unwrapping algorithm has a fake flag, it does not care about integrity, because for $m \in M(\kappa)$ the integrity measure is always added, as required by the end-to-end integrity. But for $m \notin M(\kappa)$ but of the right length, the wrapping is performed without an integrity measure. The fake flag then causes the unwrapping to ignore the missing integrity measure. Then, we also require that the state transition is independent from the message or the key.

Definition 2 (Onion correctness): Recall that $M(\kappa)$ is the message space for the security parameter κ . Let $\langle k_i \rangle_{i=1}^\ell$ be a sequence of randomly chosen bitstrings of length κ .

<p>Forward: $\Omega_f(m)$</p> $O_1 \leftarrow \text{Wrap}(m, \langle k_i \rangle_{i=1}^\ell)$ for $i = 1$ to ℓ do $O_{i+1} \leftarrow \text{Unwrap}(O_i, k_i)$ $x \leftarrow O_{\ell+1}$	<p>Backward: $\Omega_b(m)$</p> $O_\ell \leftarrow \text{Wrap}(m, k_\ell)$ for $i = \ell - 1$ to 1 do $O_i \leftarrow \text{Wrap}(O_{i+1}, k_i)$ $x \leftarrow \text{Unwrap}(O_1, \langle k_i \rangle_{i=1}^\ell)$
--	--

Let Ω'_f be the defined as Ω_f except that *Unwrap* additionally uses the fake flag. Analogously, Ω'_b is defined. We say that a pair of onion algorithms (*Wrap*, *Unwrap*) is *correct* if the following three conditions hold:

- (i) $\Pr[x \leftarrow \Omega_d(m) : x = m] = 1$ for $d \in \{f, b\}$ and $m \in M(\kappa)$.
- (ii) $\Pr[x \leftarrow \Omega'_d(m) : x = m] = 1$ for $d \in \{f, b\}$ and all $m \in M'(\kappa) := \{m' | \exists m'' \in M(\kappa). |m'| = |m''|\}$.

(iii) For all $m \in M'(\kappa)$, $k, k' \in \{0, 1\}^\kappa$ and $c, s \in \{0, 1\}^*$ such that c is a valid onion and s is a valid state

$$\Pr[(c', s') \leftarrow \text{WrOn}(m, k, s), \\ (m', s'') \leftarrow \text{UnwrOn}(c, k', s) : s' = s''] = 1$$

(iv) WrOn and UnwrOn are polynomial-time computable and randomized algorithms.

Synchronicity. The second property is synchronicity. In order to achieve replay resistance, we have to require that once the wrapping and unwrapping do not have synchronized states anymore, the output of the wrapping and unwrapping algorithms is indistinguishable from randomness.

Definition 3 (Synchronicity): For a machine \mathcal{A} , let $\Omega_{l, \mathcal{A}}$ and $\Omega_{r, \mathcal{A}}$ be defined as follows:

<p>Left: $\Omega_{l, \mathcal{A}}(\kappa)$</p> $(m_1, m_2, st) \leftarrow \mathcal{A}(1^\kappa)$ $k, s, s' \xleftarrow{\$} \{0, 1\}^\kappa$ $O \leftarrow \text{WrOn}(m_1, k, s)$ $O' \leftarrow \text{UnwrOn}(O, k, s')$ $b \leftarrow \mathcal{A}(O', st)$	<p>Right: $\Omega_{r, \mathcal{A}}(\kappa)$</p> $(m_1, m_2, st) \leftarrow \mathcal{A}(1^\kappa)$ $k, s, s' \xleftarrow{\$} \{0, 1\}^\kappa$ $O \leftarrow \text{WrOn}(m_2, k, s)$ $O' \leftarrow \text{UnwrOn}(O, k, s')$ $b \leftarrow \mathcal{A}(O', st)$
--	---

For all PPT machines \mathcal{A} the following is negligible in κ :

$$|\Pr[b \leftarrow \Omega_{l, \mathcal{A}}(\kappa) : b = 1] - \Pr[b \leftarrow \Omega_{r, \mathcal{A}}(\kappa) : b = 1]|$$

End-to-end integrity. The third property that we require is *end-to-end integrity*; i.e., the attacker is not able to produce an onion that successfully unwraps unless it compromises the exit node. For the following definition, we modify OS-Ch⁰ such that, along with the output of the attacker, also the state of the challenger is output. In turn, the resulting challenger OS-Ch^{0'} can optionally get a state s as input. In particular, $(a, s) \leftarrow A^B$ denotes in the following definition the pair of the outputs of A and B .

For the following definition we use the modified challenger OS-Ch^{0'}, which results from modifying OS-Ch⁰ such that along with the output of the attacker also the state of the challenger is output. The resulting challenger OS-Ch^{0'} can, moreover, optionally get a state s as input.

Definition 4 (End-to-end integrity): Let $S(O, cid)$ be the machine that sends a $(\text{destruct}, O)$ query to the challenger and outputs the response. Let $Q'(s)$ be the set of answers to construct queries from the challenger to the attacker. Let the last onion $O_{\ell'}$ of an onion O_1 be defined as follows:

Last(O_1):

for $i = 1$ to $\ell' - 1$ **do**
 $O_{i+1} \leftarrow \text{UnwrOn}(O_i)$

Let $Q(s) := \{\text{Last}(O_1) \mid O_1 \in Q'(s)\}$ be the set of last onions answers to the challenger. We say a set of onion algorithms has *end-to-end integrity* if for all PPT attackers \mathcal{A} the following is negligible in κ

$$\Pr[(O, s) \leftarrow \mathcal{A}(1^\kappa)^{\text{OS-Ch}^{0'}}, (m, s') \leftarrow S(O, cid)^{\text{OS-Ch}^{0'}(s)} \\ : m \in M(\kappa) \wedge P_{\ell'} \text{ is honest} \wedge O \notin Q(s')].$$

(setup, ℓ')
if $initiated = false$ **then**
for $i = 1$ to ℓ' **do**
 $k_i \xleftarrow{\$} \{0, 1\}^\kappa$; $cid_i \xleftarrow{\$} \{0, 1\}^\kappa$
 $initiated \leftarrow true$; store ℓ'
send cid

(compromise, i)
 $initiated \leftarrow false$; erase the circuit
 $compromised(i) \leftarrow true$; run setup;
send (cid_j, k_j) for all j with $compromised(j) = true$

(send, m)
 $O \leftarrow \text{WrOn}(m, \langle k_i \rangle_{i=1}^{\ell'})$
send O

(unwrap, O, cid)
look up the key k for cid
 $O' \leftarrow \text{UnwrOn}(O, k)$
send O'

(respond, m)
 $O \leftarrow \text{WrOn}(m, k_{\ell'})$
send O

(wrap, O, cid)
look up the key k for cid
 $O' \leftarrow \text{WrOn}(O, k)$
send O'

(destruct, O)
 $m \leftarrow \text{UnwrOn}(O, \langle k_i \rangle_{i=1}^{\ell'})$
send m

Figure 9. The Honest Onion Secrecy Challenger OS-Ch⁰: OS-Ch⁰ only answers for honest parties

Predictably Malleable Onion Secrecy. The fourth property that we require is *predictably malleable onion secrecy*, i.e., for every modification to a ciphertext the challenger is able to compute the resulting changes for the plaintext. This even has to hold for faked plaintexts.

In detail, we define a challenger OS-Ch⁰ that provides, a wrapping, a unwrapping and a send and a destruct oracle. In other words, the challenger provides the same oracles as in the onion routing protocol except that the challenger only provides one single session. We additionally define a faking challenger OS-Ch¹ that provides the same oracles but fakes all onions for which the attacker does not control the final node.

For OS-Ch¹, we define the maximal paths that the attacker knows from the circuit. A visible subpath of a circuit $(P_i, k_i, cid_i)_{i=1}^{\ell}$ from an honest onion proxy is a minimal subsequence of corrupted parties $(P_i)_{i=u}^s$ of $(P_i)_{i=1}^{\ell}$ such that P_{i-1} is honest and either $s = \ell$ or P_{s+1} is honest as well. The parties P_{i-1} and, if existent, P_{s+1} are called the guards of the visible subpath $(P_i)_{i=u}^s$. We store visible subpaths by the first $cid = cid_u$.

Figure 9 and 10 presents OS-Ch⁰, and OS-Ch¹, respec-

```

(setup,  $\ell'$ )
  do the same as OS-Ch0
  additionally draw a distinguished key  $k_S \leftarrow \{0, 1\}^\kappa$ 
(compromise,  $i$ )
  do the same as OS-Ch0


---


(send,  $m$ )
   $q(st_f^1) \leftarrow m$ ; look up the first visible subpath  $(cid_1, \langle k_i \rangle_{i=1}^j)$ 

  if  $j = \ell'$  then  $m' \leftarrow q(st_f^1)$ 
  else  $k_{j+1} \leftarrow k_S$ ;  $j \leftarrow j + 1$ ;  $m' \leftarrow 0^{|q(st_f^1)|}$ 
   $((O_i)_{i=0}^j, s') \leftarrow WrOn^j(m, \langle k_i \rangle_{i=1}^j, st_f^1)$ 
  update  $st_f^1 \leftarrow s'$ 
  store  $onions(cid_j) \leftarrow O_1$ ; send  $O_j$ 
(unwrap,  $O, cid_i$ )
  look up the forward v.s.  $\langle k_i \rangle_{i=u}^j$  for  $cid_i$ 
   $O' \leftarrow onions(cid_i)$ 
   $T \leftarrow M(O, O')$ ;  $q(st_f^i) \leftarrow T(q(st_f^i))$ 
  if  $j = \ell'$  then  $m \leftarrow q(st_f^i)$ 
  else  $k_{j+1} \leftarrow k_S$ ;  $j \leftarrow j + 1$ ;  $m \leftarrow 0^{|q(st_f^i)|}$ 
   $((O_i)_{i=u-1}^j, s') \leftarrow WrOn^{j-u+1}(m, \langle k_i \rangle_{i=u}^j, st_f^i)$ 
  update  $st_f^i \leftarrow s'$ 
  store  $onions(cid_j) \leftarrow O_u$ ; send  $O_j$ 
(respond,  $m$ )
   $q(st_b^{\ell'}) \leftarrow m$ ; look up the last visible subpath  $\langle k_i \rangle_{i=u}^{\ell'}$ 
  if  $u = 1$  then  $m \leftarrow q(st_b^{\ell'})$ 
  else  $k_{u-1} \leftarrow k_S$ ;  $u \leftarrow u - 1$ ;  $m \leftarrow 0^{|q(st_b^{\ell'})|}$ 
   $((O_i)_{i=u-1}^j, s') \leftarrow WrOn^{j-u+1}(m, \langle k_i \rangle_{i=u}^j, st_b^{\ell'})$ 
  update  $st_b^{\ell'} \leftarrow s'$ 
  store  $onions(cid_u) \leftarrow O_u$ ; send  $O_j$ 
(wrap,  $O, cid_i$ )
  look up the backward v.s.  $\langle k_i \rangle_{i=u}^j$  for  $cid_i$ 
   $O' \leftarrow onions(cid_i)$ ;  $T \leftarrow M(O, O')$ ;  $q(st_b^i) \leftarrow T(q(st_b^i))$ 
  get  $\langle k_i \rangle_{i=u}^j$  for  $cid$ 
  if  $u = 1$  then  $m \leftarrow q(st_b^i)$ 
  else  $k_{u-1} \leftarrow k_S$ ;  $u \leftarrow u - 1$ ;  $m \leftarrow 0^{|q(st_b^i)|}$ 
   $((O_i)_{i=u-1}^j, s') \leftarrow WrOn^{j-u+1}(m, \langle k_i \rangle_{i=u}^j, st_b^i)$ 
  update  $st_b^i \leftarrow s'$ 
  store  $onions(cid_u) \leftarrow O_u$ ; send  $O_j$ 
(destruct,  $O, cid$ )
   $m \leftarrow UnwrOn(k_1, st_b^1)$ ;  $O' \leftarrow onions(cid_1)$ 
   $T \leftarrow M(O, O')$ ;  $q(st_b^1) \leftarrow T(q(st_b^1))$ 
  if  $m \neq \perp$  then
    send  $q(st_b^1)$ 

```

Figure 10. The Faking Onion Secrecy Challenger OS-Ch¹: $((O_i)_{i=u-1}^j, s') \leftarrow WrOn^{j-u+1}(m, \langle k_i \rangle_{i=u}^j, st)$ is defined as $O_{u-1} \leftarrow m$; **for** $i = u$ **to** j **do** $(O_i, s') \leftarrow WrOn(O_{i-1}, k_{j+u-i}, st)$. OS-Ch¹ only answers for honest parties

tively.⁹

Definition 5 (Predictably malleable onion secrecy): Let onionAlg be a pair of algorithms $WrOn$ and $UnwrOn$. We say that the algorithms onionAlg satisfy *predictably*

⁹We stress that in Figure 10 the onion O_u denotes the onion from party P_j to party P_{j+1} .

malleable onion secrecy if there is a negligible function μ such that there is a efficiently computable function M such that for all PPT machines \mathcal{A} and sufficiently large κ

$$\Pr[b \stackrel{\$}{\leftarrow} \{0, 1\}, b' \leftarrow \mathcal{A}(1^\kappa)^{OS-Ch^b} : b = b'] \leq 1/2 + \mu(\kappa)$$

Definition 6 (Secure onion algorithms): A pair of onion algorithms $(WrOn, UnwrOn)$ is *secure* if it satisfies onion correctness, synchronicity, predictably malleable onion secrecy, and end-to-end integrity.

In Section VI-B, we show that the Tor algorithms are secure onion algorithms.

V. Π_{OR} UC-REALIZES \mathcal{F}_{OR}

In this section, we show that Π_{OR} can be securely abstracted as the ideal functionality \mathcal{F}_{OR} .

Recall that π securely realizes \mathcal{F} in the \mathcal{F}' -hybrid model if each party in the protocol π has a direct connection to \mathcal{F}' . \mathcal{F}_{REG}^N is the key registration, \mathcal{F}_{SCS} is the secure channel functionality, and \mathcal{F}_{NET^q} is the network functionality, where q is the upper bound on the corruptable parties. We prove our result in the $\mathcal{F}_{REG}^N, \mathcal{F}_{SCS}$ -hybrid model; i.e., our result holds for any key registration and secure channel protocol securely realizing \mathcal{F}_{REG}^N and \mathcal{F}_{SCS} , respectively. The network functionality \mathcal{F}_{NET^q} abstract partially global attacker and is a global functionality.¹⁰

Theorem 1: If Π_{OR} uses secure OR modules \mathcal{M} , then with the global functionality \mathcal{F}_{NET^q} the resulting protocol Π_{OR} in the $\mathcal{F}_{REG}^N, \mathcal{F}_{SCS}$ -hybrid model securely realizes the ideal functionality \mathcal{F}_{OR} for any q .

As a next step, we give a proof outline in order to highlight at which places we apply the required security properties or the secure OR modules. The full proof can be found in [1].

Proof outline: First, we show that the original scenario is indistinguishable from that in which a simulator computes $\Pi_{OR}, \mathcal{F}_{REG}^N, \mathcal{F}_{SCS}$, and \mathcal{A} . Then, we first modify the 1W-AKE primitive such that we use randomly chosen keys for honest pairs of parties instead of the computed ones. By the security of 1W-AKE the environment cannot tell the difference. Thereafter, we modify all honestly generated onions with honest exit-nodes such that all plaintexts are replaced by the constant-zero bitstring. By the onion secrecy and the synchronicity, we know that the environment can still not tell the difference. Finally, the simulator still internally runs $\mathcal{F}_{REG}^N, \mathcal{F}_{SCS}$, and \mathcal{A} but only uses the information that \mathcal{F}_{OR} offers. By an extensive case distinction of the definition of \mathcal{F}_{OR} and by applying the anonymity of the 1W-AKE primitive, we conclude that the environment also cannot tell this difference. ■

As our primitives are proven secure in the random oracle model (ROM), the main theorem uses the ROM.

¹⁰We stress that Π_{OR} (with any modules) is \mathcal{F}_{NET^q} -subroutine respecting; hence the GUC composition theorem holds.

```

 $G_c(1^\eta)$ 
  output  $k \xleftarrow{\$} G(1^\eta)$ 

 $E_c((x_1, \dots, x_t), (k, ctr)) = D_c((x_1, \dots, x_t), (k, ctr))$ 
  if  $ctr = \varepsilon$  then  $ctr = 0$ 
  output  $(\text{PRP}(s, k) \oplus x_1, \dots, \text{PRP}(s + t - 1, k) \oplus x_t, (k, ctr + t))$ 

```

Figure 11. The stateful deterministic counter-mode (detCTR) $\mathcal{E}_c = (G_c, E_c, D_c)$

Theorem 2: If pseudorandom permutations exist, there are secure OR modules (ntor, onionAlgs) such that the protocol Π_{OR} in the $\mathcal{F}_{\text{REG}}^N, \mathcal{F}_{\text{SCS}}, \mathcal{F}_{\text{NET}^q}$ -hybrid model using (ntor, onionAlgs) securely realizes in the ROM the ideal functionality \mathcal{F}_{OR} in the $\mathcal{F}_{\text{NET}^q}$ -hybrid model for any q .

Proof: If pseudorandom permutations exist Lemma 2 implies that secure onion algorithms exist. Lemma 3 shows that in the ROM 1W-AKE exist. Then, Theorem 1 implies the statement. ■

Note that we could not prove 1W-AKE security for the TAP protocol currently used in Tor as it uses a CCA-insecure version of the RSA encryption scheme.

VI. INSTANTIATING SECURE OR MODULES

We present a concrete instantiation of OR modules and show that this instantiation constitutes a set of secure OR modules. As onion algorithms we use the algorithms that are used in Tor with a strengthened integrity mechanism, and as 1W-AKE we use the recently proposed ntor protocol [14].

We prove that the onion algorithms of Tor constitute secure onion algorithms, as defined in Definition 6. The crucial part in that proof is to show that these onion algorithms are predictably malleable, i.e., for every modification of the ciphertext the changes in the resulting plaintext are predictable by merely comparing the modified ciphertext with the original ciphertext. We first show that the underlying encryption scheme, the deterministic counter-mode, is predictably malleable (Section VI-A). Thereafter, we show the security of Tor’s onion algorithms (Section VI-B).

In Section VI-C, we briefly present the ntor protocol and cite the result from Goldberg, Stebila, and Ustaoglu that ntor constitutes a 1W-AKE. The proofs of the lemmas in this section are postponed to the full version [1].

A. Deterministic Counter Mode and Predictable Malleability

We show that the deterministic counter-mode (detCTR) scheme is predictably malleable, as defined in Definition 1.

Lemma 1: If pseudorandom permutations exist, the deterministic counter mode (detCTR) with $\mathcal{E}_c = (G_c, E_c, D_c)$ as defined in Figure 11 predictably malleable.

```

 $WrOn_I(O, k)$ , for  $O \notin M(\kappa)$ 
   $O' \leftarrow Enc_{ctr}(O, k)$ ; return  $O'$ 

 $WrOn_I(m, k)$ , for  $m \in M(\kappa)$ 
   $(r, r') \leftarrow PRG(k)$ ;  $k_m \leftarrow Gen_m(r)$ 
   $k_e \leftarrow Gen_e(r')$ 
   $Mac(m, k_m)$ 
   $O' \leftarrow Enc_{ctr}(O, k_e)$ ; return  $O'$ 

 $WrOn_I(m, \langle k_i \rangle_{i=1}^\ell)$ , for  $m \in M(\kappa)$ 
   $O_2 \leftarrow WrOn_I(m, k_1)$ 
  for  $i = 2$  to  $\ell$  do
     $O_{i+1} \leftarrow WrOn_I(O_i, k_i)$ 
  return  $O_\ell$ 

 $UnwrOn_I(O, k)$ 
   $(r, r') \leftarrow PRG(k)$ ;  $k_m \leftarrow Gen_m(r)$ 
   $k_e \leftarrow Gen_e(r')$ 
   $O' \leftarrow Dec_{ctr}(O, k_e)$ 
  if  $O' = m || t$  and  $m \in M(\kappa)$  and  $V(m, t, k_m) = 1$  then
    return  $O''$ 
  else
     $O' \leftarrow Dec_{ctr}(O, k)$ ; return  $O'$ 

 $UnwrOn_I(O, k, \text{fake})$ 
   $O' \leftarrow Dec_{ctr}(O, k)$ ; return  $O'$ 

 $UnwrOn_I(O, \langle k_i \rangle_{i=1}^\ell)$ 
  for  $i = 1$  to  $\ell$  do
     $O_{i+1} \leftarrow WrOn_I(O_i, k_i)$ 
  return  $O_\ell$ 

```

Figure 12. The Onion Algorithms onionAlg

B. Security of Tor’s Onion Algorithms

Let $E := (Gen_e, Enc, Dec)$ be a stateful deterministic encryption scheme, and let $M := (Gen_m, Mac, V)$ be a deterministic MAC. Let PRG be a pseudo random generator such that for all $x \in \{0, 1\}^*$ $|PRG(x)| = 2 \cdot |x|$. We write $PRG(x)_1$ for the first half of $PRG(x)$ and $PRG(x)_2$ the second half. Moreover, for a randomized algorithm A , we write $A(x; r)$ for a call of $A(x)$ with the randomness r .

As a PRP candidate we use AES, as in Tor, and as a MAC use H-MAC with SHA-256. We use that in detCTR encrypting two blocks separately results in the same ciphertext as encrypting the pair of the blocks at once. Moreover, we assume that the output of H-MAC is exactly one block.

The correctness follows by construction. The synchronicity follows, because a PRP is used for the state. The end-to-end integrity directly follows from the SUF of the Mac. And the predictable malleability follows from the predictable malleability of the deterministic counter-mode.

Lemma 2: Let $\text{onionAlg} = \{UnwrOn_I, WrOn_I\}$. If pseudorandom permutations exist, onionAlg are secure onion algorithms.

C. ntor: A 1W-AKE

Øverlier and Syverson [24] proposed a 1W-AKE for use in the next generation of the Tor protocol with improved efficiency. Goldberg, Stebila, and Ustaoglu found an authen-

Initiate(pk_Q, Q):

- 1) Generate an ephemeral key pair $(x, X \leftarrow g^x)$.
- 2) Set session id $\Psi_P \leftarrow H_{st}(X)$.
- 3) Update $st(\Psi_P) \leftarrow (\text{ntor}, Q, x, X)$.
- 4) Set $m_P \leftarrow (\text{ntor}, Q, X)$.
- 5) Output m_P .

Respond(pk_Q, sk_Q, X):

- 1) Verify that $X \in G^*$.
- 2) Generate an ephemeral key pair $(y, Y \leftarrow g^y)$.
- 3) Set session id $\Psi_Q \leftarrow H_{st}(Y)$.
- 4) Compute $(k', k) \leftarrow H(X^y, X^{sk_Q}, Q, X, Y, \text{ntor})$.
- 5) Compute $t_Q \leftarrow H_{mac}(k', Q, Y, X, \text{ntor}, \text{server})$.
- 6) Set $m_Q \leftarrow (\text{ntor}, Y, t_Q)$.
- 7) Set $out \leftarrow (k, \star, X, Y, pk_Q)$, where \star is the anonymous party symbol.
- 8) Delete y and output m_Q .

ComputeKey(pk_Q, Ψ_P, t_Q, Y):

- 1) Retrieve Q, x, X from $st(\Psi_P)$ if it exists.
- 2) Verify that $Y \in G^*$.
- 3) Compute $(k', k) \leftarrow H(Y^x, pk_Q^x, Q, X, Y, \text{ntor})$.
- 4) Verify $t_Q = H_{mac}(k', Q, Y, X, \text{ntor}, \text{server})$.
- 5) Delete $st(\Psi_P)$ and output k .

If any verification fails, the party erases all session-specific information and aborts the session.

Figure 13. The ntor protocol

tication flaw in this proposed protocol, fixed it, and proved the security of the fixed protocol [14]. We use this fixed protocol, called ntor, as a 1W-AKE.

The protocol ntor [14] is a 1W-AKE protocol between two parties P (client) and Q (server), where client P authenticates server Q . Let (pk_Q, sk_Q) be the static key pair for Q . We assume that P holds Q 's certificate (Q, pk_Q) . P initiates an ntor session by calling the *Initiate* function and sending the output message m_P to Q . Upon receiving a message m'_P , server Q calls the *Respond* function and sends the output message m_Q to P . Party P then calls the *ComputeKey* function with parameters from the received message m'_Q , and completes the ntor protocol. We assume a unique mapping between the session ids Ψ_P of the cid in Π_{OR} .

Lemma 3 (ntor is anonymous and secure [14]): The ntor protocol is a one-way anonymous and secure 1W-AKE protocol in the random oracle model (ROM).

VII. FORWARD SECRECY AND ANONYMITY ANALYSES

In this section, we show that our abstraction \mathcal{F}_{OR} allows for applying previous work on the anonymity analysis of onion routing to Π_{OR} . Moreover, we illustrate that \mathcal{F}_{OR} enables a rigorous analysis of forward secrecy of Π_{OR} .

In Section VII-A, we show that the analysis of Feigenbaum, Johnson, and Syverson [12] of Tor's anonymity properties in a black-box model can be applied to our protocol Π_{OR} . Feigenbaum, Johnson, and Syverson show their anonymity analysis an ideal functionality \mathcal{B}_{OR} (see Figure 4). By proving that the analysis of \mathcal{B}_{OR} applies to

upon the first input m
 send N_{OR} to \mathcal{F}_{REG}^N in Π ; send setup to Π
 wait for $(\text{ready}, \langle P_i \rangle_{i=1}^n)$; further process m

upon an input (send, $S, [m]$)
 draw P_1, \dots, P_ℓ at random from N_{OR}
 store (S, m_{dummy}) [or (S, m)] in the queue for $\langle P, P_1, \dots, P_\ell \rangle$
 send $(\text{cc}, \langle P, P_1, \dots, P_\ell \rangle)$ to Π

upon (created, $\langle P \xleftrightarrow{cid} P_1 \iff \dots P_\ell \rangle$) **from** Π
 look up (S, m) from the queue for $\langle P, P_1, \dots, P_\ell \rangle$
 send (send, $\langle P \xleftrightarrow{cid} P_1 \iff \dots P_\ell \rangle, (S, m)$) to Π

upon (received, \mathcal{C}, m) **from** Π
 do nothing /* \mathcal{B}_{OR} does not allow responses to messages*/

upon a message m from \mathcal{F}_{NET^0} to the environment
 do nothing /* \mathcal{B}_{OR} does not deliver messages*/

Figure 14. User-interface $U(\Pi)$ for party P

\mathcal{F}_{OR} , the UC composition theorem and Theorem 1 imply that the analysis applies to Π_{OR} as well.

In Section VII-B, we present the result that immediate forward secrecy for Π_{OR} holds, by merely by analyzing \mathcal{F}_{OR} . The proofs in this section are omitted due to space constraints but can be found in the full version [1].

A. OR Anonymity Analysis

Feigenbaum, Johnson and Syverson [12] analyzed the anonymity properties of OR networks. In their analysis, the authors abstracted an OR network against attackers that are local, static as a black-box functionality \mathcal{B}_{OR} . We reviewed their abstraction \mathcal{B}_{OR} in Section II-E. In this section, we show that the analysis of \mathcal{B}_{OR} is applicable to Π_{OR} against local, static attackers.

There is a slight mismatch in the user-interface of \mathcal{B}_{OR} and Π_{OR} . The main difference is that Π_{OR} expects separate commands for creating a circuit and sending a message whereas \mathcal{B}_{OR} only expects a command for sending a message. We construct for every party P a wrapper U for Π_{OR} that adjusts Π_{OR} 's user-interface. Recall that we consider two versions of \mathcal{B}_{OR} and U simultaneously: one version in which no message is sent and one version in which a message is sent (denoted as $[m]$).

Instead of Π_{OR} , U only expects one command: (send, $S, [m]$). We fix the length ℓ of the circuit.¹¹ Upon (send, $S, [m]$), $U(\Pi)$ draws the path $P_1, \dots, P_\ell \xleftarrow{\$} N_{OR}$ at random, sends a $(\text{cc}, \langle P, P_1, \dots, P_\ell \rangle)$ to Π , waits for the cid from Π , and sends a (send, cid, m) command, where m is

¹¹We fix the length for the sake of brevity. This choice is rather arbitrary. The analysis can be adjusted to the case in which the length is chosen from some efficiently computable distribution or specified by the environment for every message.

a dummy message if no message is specified. Moreover, in contrast to \mathcal{B}_{OR} the protocol Π_{OR} allows a response for a message m and therefore additionally sends a session id sid to a server.¹²

In addition to the differences in the user-interface, \mathcal{B}_{OR} assumes the weaker threat model of a local, static attacker whereas Π_{OR} assumes a partially global attacker. We formalize a local attacker by considering Π_{OR} in the $\mathcal{F}_{\text{NET}^0}$ -hybrid model, and connect the input/output interface of $\mathcal{F}_{\text{NET}^0}$ to the wrapper U as well. For considering a static attacker, we make the standard UC-assumption that every party only accepts compromise requests at the beginning of the protocol. Moreover, we also need to assume that \mathcal{B}_{OR} is defined for a fixed set of onion routers \mathcal{N} in the same way as $\mathcal{F}_{\text{OR}}^{\mathcal{N}}$.

Finally, our work culminates in the connection of previous work on black-box anonymity analyses of onion routing with our cryptographic model of onion routing.

Lemma 4 ($U(\Pi_{\text{OR}})$ UC realizes \mathcal{B}_{OR}): Let $U(\Pi_{\text{OR}})$ be defined as in Figure 14. If Π_{OR} uses secure OR modules, then $U(\Pi_{\text{OR}})$ in the $\mathcal{F}_{\text{NET}^0}$ -hybrid model UC realizes \mathcal{B}_{OR} against static attackers.

B. Forward Secrecy

Forward secrecy [7] in cryptographic constructions ensures that a session key derived from a set of long-term public and private keys will not be compromised once the session is over, even when one of the (long-term) private keys is compromised in the future. Forward secrecy in onion routing typically refers to the privacy of a user's circuit against an attacker that marches down the circuit compromising the nodes until he reaches the end and breaks the user's anonymity.

It is commonly believed that for achieving forward secrecy in OR protocols it is sufficient to securely erase the local circuit information once a circuit is closed, and to use a key exchange that provides forward secrecy. Π_{OR} uses such a mechanism for ensuring forward secrecy. Forward secrecy for OR, however, has never been proven, and not even rigorously defined.

In this section, we present a game-based definition for OR forward secrecy (Definition 10) and show that Π_{OR} satisfies our forward secrecy definition (Lemma 7). We require that a local attacker does not even learn anything about a closed circuit if he compromises all system nodes. The absence of knowledge about a circuit is formalized in the notion of *OR circuit secrecy* (Definition 10), a notion that might be of independent interest.

Recall that we formalize a local attacker by considering Π_{OR} in the $\mathcal{F}_{\text{NET}^0}$ -hybrid model, i.e., the attacker cannot

¹²It is also possible to modify Π_{OR} such that Π_{OR} does not accept responses and does not draw a session id sid . However, for the sake of brevity we slightly modify \mathcal{B}_{OR} .

```

upon the first input  $m$ 
  set  $N_{\mathcal{A}} := \emptyset$ ; send  $N_{\text{OR}}$  to  $\mathcal{F}_{\text{REG}}^{\mathcal{N}}$  in  $\mathcal{F}_{\text{OR}}$ 
  send setup to  $\mathcal{F}_{\text{OR}}$ ; wait for (ready,  $\langle P_i \rangle_{i=1}^n$ )
  further process  $m$ 

upon (compromise,  $P$ ) from  $\mathcal{A}$ 
  if all previous messages only were compromise messages
  then
    set  $N_{\mathcal{A}} := N_{\mathcal{A}} \cup \{P\}$ 
    forward (compromise,  $P$ ) to party  $P$  in  $U(\mathcal{F}_{\text{OR}})$ 

upon the first message  $m$  that is not compromise from  $\mathcal{A}$ 
  send (compromise,  $N_{\mathcal{A}}$ ) to  $\mathcal{B}_{\text{OR}}$ 
  further process  $m$ 

upon any other message  $m$  from  $\mathcal{A}$  to  $\mathcal{F}_{\text{NET}^0}$ 
  forward  $m$  to  $\mathcal{F}_{\text{NET}^0}$ 

upon any other message  $m$  from  $\mathcal{A}$  to  $U'(\mathcal{F}_{\text{OR}})$ 
  forward  $m$  to  $U'(\mathcal{F}_{\text{OR}})$ 

upon a message  $m$  from  $U'(\mathcal{F}_{\text{OR}})$  to the environment
  do nothing /*  $\mathcal{B}_{\text{OR}}$  already outputs the message */

upon (sent,  $U, S, [m]$ ) from  $\mathcal{B}_{\text{OR}}$ 
  choose  $P_1 \xleftarrow{\$} N_{\mathcal{A}}$  and  $P_\ell \xleftarrow{\$} N_{\mathcal{A}}$ 
  choose  $P_2, \dots, P_{\ell-1} \xleftarrow{\$} N_{\text{OR}}$ 
  send (send,  $\langle P_i \rangle_{i=1}^\ell, [m]$ ) to  $U'(\mathcal{F}_{\text{OR}})$ 

upon (sent,  $-, S, [m]$ ) from  $\mathcal{B}_{\text{OR}}$ 
  choose  $P_1 \xleftarrow{\$} N_{\text{OR}} \setminus N_{\mathcal{A}}$  and  $P_\ell \xleftarrow{\$} N_{\mathcal{A}}$ 
  choose  $P_2, \dots, P_{\ell-1} \xleftarrow{\$} N_{\text{OR}}$ 
  send (send,  $\langle P_i \rangle_{i=1}^\ell, [m]$ ) to  $U'(\mathcal{F}_{\text{OR}})$ 

upon (sent,  $U, -$ ) from  $\mathcal{B}_{\text{OR}}$ 
  choose  $P_1 \xleftarrow{\$} N_{\mathcal{A}}$  and  $P_\ell \xleftarrow{\$} N_{\text{OR}} \setminus N_{\mathcal{A}}$ 
  choose  $P_2, \dots, P_{\ell-1} \xleftarrow{\$} N_{\text{OR}}$ 
  send (send,  $\langle P_i \rangle_{i=1}^\ell, [m_{\text{dummy}}]$ ) to  $U'(\mathcal{F}_{\text{OR}})$ 

upon (sent,  $-, -$ ) from  $\mathcal{B}_{\text{OR}}$ 
  choose  $P_1 \xleftarrow{\$} N_{\text{OR}} \setminus N_{\mathcal{A}}$  and  $P_\ell \xleftarrow{\$} N_{\text{OR}} \setminus N_{\mathcal{A}}$ 
  choose  $P_2, \dots, P_{\ell-1} \xleftarrow{\$} N_{\text{OR}}$ 
  send (send,  $\langle P_i \rangle_{i=1}^\ell, [m_{\text{dummy}}]$ ) to  $U'(\mathcal{F}_{\text{OR}})$ 

```

Figure 15. The simulator $S_{\mathcal{A}}$: U' gets the path as input instead of drawing it at random

observe the link between any pair of nodes without compromising any of the two nodes.

Definition 7 (Local attackers): We say that we consider a protocol Π against local attackers if we consider Π in the $\mathcal{F}_{\text{NET}^0}$ -hybrid model.

The definition of circuit secrecy compares a pair of circuits and requires that the attacker cannot tell which one has been used. Of course, we can only compare two circuits that are not trivially distinguishable. The following notion of *visibly coinciding circuits* excludes trivially distinguishable pairs of circuits. Recall that a visible subpath of a circuit is

<p>CS-Ch_b^Π: (setup) from \mathcal{A} if $initial = \perp$ then send (setup) to Π $challenge \leftarrow false$ $initial \leftarrow true$</p>	<p>CS-Ch_b^Π: $(cc, \mathcal{P}^0, \mathcal{P}^1, P)$ from \mathcal{A} if $challenge = true$ then if \mathcal{P}^0 and \mathcal{P}^1 visibly coincide then forward (cc, \mathcal{P}^b, P) to Π</p>
<p>CS-Ch_b^Π: (compromise, P) from \mathcal{A} if $challenge = false$ then store that P is compromised forward (compromise, P) to Π</p>	<p>CS-Ch_b^Π: for every other message m from \mathcal{A} forward m to Π</p> <p>CS-Ch_b^Π: for every message m from Π if $challenge = true$ and $m = (created, \langle P \xleftrightarrow{cid} P_1 \iff \dots \iff P_{\ell'} \rangle, P)$ then store P for cid forward m to \mathcal{A}</p>
<p>CS-Ch_b^Π: (close_initial) from \mathcal{A} $challenge \leftarrow true$</p>	

Figure 16. OR Circuit Secrecy Game

<p>FS-Ch_b^Π behaves exactly like CS-Ch_b^Π except for the following message:</p>
<p>FS-Ch_b^Π: upon (close_challenge) from \mathcal{A} if $challenge = true$ then $challenge \leftarrow false$ for every circuit cid created in the challenge phase do look up onion proxy P for cid; send $(cid, destroy, P)$ to Π</p>

Figure 17. OR Forward Secrecy Challenger: FS-Ch_b^Π

a maximal contiguous subsequence of compromised nodes.

Definition 8 (Visibly coinciding circuits): A subsequence $\langle P_j \rangle_{j=u}^s$ of a circuit $\langle P_i \rangle_{i=1}^\ell$ is an *extended visible subpath* if $\langle P_j \rangle_{j=u+1}^{s-1}$ is a visible subpath or $s = \ell$ and $\langle P_j \rangle_{j=u+1}^s$ is a visible subpath.

We say that two circuits $\mathcal{P}^0 = \langle P_i^0 \rangle_{i=0}^{\ell^0}$, $\mathcal{P}^1 = \langle P_i^1 \rangle_{i=0}^{\ell^1}$ are trivially distinguishable if the following three conditions hold:

- (i) the onion proxies P_0^0, P_0^1 are not compromised,
- (ii) the sequences of extended visible subpaths of \mathcal{P}^0 and \mathcal{P}^1 are the same, and
- (iii) the exit nodes of \mathcal{P}^0 and \mathcal{P}^1 are the same, i.e., $P_{\ell^0}^0 = P_{\ell^1}^1$.

For the definition of circuit secrecy of a protocol Π , we define a challenger that communicates with the protocol Π and the attacker. The challenger C_b is parametric in $b \in \{0, 1\}$. C_b forwards all requests from the attacker to the protocol except for the cc commands. Upon a cc command C_b expects a pair $\mathcal{P}^0, \mathcal{P}^1$ of node sequences, checks whether \mathcal{P}^0 and \mathcal{P}^1 are visibly coinciding circuits, chooses \mathcal{P}^b , and forwards (cc, \mathcal{P}^b) to the protocol Π . We require that the attacker does not learn anything about visibly coinciding circuits.

A protocol can be represented without loss of generality as an interactive Turing machine that internally runs every sin-

gle protocol party as a submachine, forwards each messages for a party P to that submachine, and sends every message from that submachine to the respective communication partner. We assume that upon a message (setup), a protocol responds with a list of self-generated party identifiers. The protocol expects for every message from the communication partner a party identifier and reroutes the message to the corresponding submachine. In the following definition, we use this notion of a *protocol*.

Definition 9: Let Π be a protocol and CS-Ch be defined as in Figure 16. An OR protocol has circuit secrecy if there is a negligible function μ such that the following holds for all PPT attackers \mathcal{A} and sufficiently large κ

$$\Pr[b \xleftarrow{\$} \{0, 1\}, b' \leftarrow \mathcal{A}(\kappa)^{CS-Ch_b^\Pi(\kappa)} : b = b'] \leq 1/2 + \mu(\kappa)$$

Forward secrecy requires that even if all nodes are compromised after closing all challenge circuits the attacker cannot learn anything about the challenge circuits.

Definition 10: Let Π be a protocol and FS-Ch be defined as in Figure 17. An OR protocol has circuit secrecy if there is a negligible function μ such that the following holds for all PPT attackers \mathcal{A} and sufficiently large κ

$$\Pr[b \xleftarrow{\$} \{0, 1\}, b' \leftarrow \mathcal{A}(\kappa)^{FS-Ch_b^\Pi(\kappa)} : b = b'] \leq 1/2 + \mu(\kappa)$$

Lemma 5: \mathcal{F}_{OR} against local attackers satisfies OR circuit secrecy (see Definition 9).

The protocol as introduced in Section II-D presents Π_{OR} as one (sub-)machine for every protocol party. Equivalently, Π_{OR} can be represented as one interactive Turing machine that runs all parties as submachines, upon a message (setup) from the communication partner, sends (setup) to every party, and sends an answer with a list of party identifiers to the communication partner. In the following definition, Π_{OR} is represented as one interactive Turing machine that internally runs all protocol parties.

Lemma 6: Π_{OR} instantiated with secure OR modules against local attackers satisfies OR circuit secrecy (see Definition 9).

It is easy to see that in \mathcal{F}_{OR} , once a circuit is closed, all information related to the circuit at the uncompromised nodes is deleted. Therefore, forward secrecy for \mathcal{F}_{OR} is obvious from the circuit secrecy in Lemma 6. Hence, the following lemma immediately follows.

Lemma 7: Π_{OR} instantiated with secure OR modules against local attackers satisfies OR forward secrecy (see Definition 10).

VIII. FUTURE WORK

For future work an interesting direction could be to incorporate hidden services into the UC security analysis. We already designed the abstraction in a way that allows for a modular extension of the UC proof to a hidden service functionality. Moreover, our work offers a framework for

the analysis of other desirable OR properties, such as circuit position secrecy.

It is well known that the UC framework lacks a notion of time; consequently any UC security analysis neglects timing attacks, in particular traffic analysis. A composable security analysis that also covers, e.g., traffic analysis, is an interesting task for future work.

Acknowledgements: This work is partially supported by CISP, NSERC, and Mprime. We thank Gregory Zaverucha for helpful preliminary discussions, and Aaron Johnson, Paul Syverson, Dominique Unruh and the anonymous reviewers for their valuable comments on earlier drafts of the paper. We also thank Sebastian Meiser for many insightful discussions about and ideas for predictable malleability.

REFERENCES

- [1] M. Backes, I. Goldberg, A. Kate, and E. Mohammadi, “Provably secure and practical onion routing,” IACR Cryptology ePrint Archive, Report 2011/308, 2012.
- [2] V. Boyko, P. D. MacKenzie, and S. Patel, “Provably Secure Password-Authenticated Key Exchange Using Diffie-Hellman,” in *EUROCRYPT*, 2000, pp. 156–171.
- [3] J. Camenisch and A. Lysyanskaya, “A formal treatment of onion routing,” in *Advances in Cryptology—CRYPTO 2005*, 2005, pp. 169–187.
- [4] R. Canetti, “Universally Composable Security: A New Paradigm for Cryptographic Protocols,” in *Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001, pp. 136–145.
- [5] R. Canetti, Y. Dodis, R. Pass, and S. Walfish, “Universally composable security with global setup,” in *Proc. 4th Theory of Cryptography Conference (TCC)*, 2007, pp. 61–85.
- [6] D. Catalano, D. Fiore, and R. Gennaro, “Certificateless onion routing,” in *Proc. 16th ACM Conference on Computer and Communication Security (CCS)*, 2009, pp. 151–160.
- [7] W. Diffie, P. C. van Oorschot, and M. J. Wiener, “Authentication and Authenticated Key Exchanges,” *Des. Codes Cryptography*, vol. 2, no. 2, pp. 107–125, 1992.
- [8] R. Dingledine and N. Mathewson, “Tor Protocol Specification,” https://gitweb.torproject.org/torspec.git?a=blob_plain;hb=HEAD;f=tor-spec.txt, 2008, accessed Nov 2011.
- [9] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The Second-Generation Onion Router,” in *Proc. 13th USENIX Security Symposium (USENIX)*, 2004, pp. 303–320.
- [10] J. Feigenbaum, A. Johnson, and P. F. Syverson, “A model of onion routing with provable anonymity,” in *Proc. 11th Conference on Financial Cryptography and Data Security (FC)*, 2007, pp. 57–71.
- [11] J. Feigenbaum, A. Johnson, and P. F. Syverson, “Probabilistic analysis of onion routing in a black-box model,” in *Proc. 6th ACM Workshop on Privacy in the Electronic Society (WPES)*, 2007, pp. 1–10.
- [12] J. Feigenbaum, A. Johnson, and P. F. Syverson, “Probabilistic Analysis of Onion Routing in a Black-box Model, Tech. Rep. arXiv:1111.2520, 2011, <http://arxiv.org/abs/1111.2520v1>.
- [13] I. Goldberg, “On the Security of the Tor Authentication Protocol,” in *Proc. 6th Workshop on Privacy Enhancing Technologies*, 2006, pp. 316–331.
- [14] I. Goldberg, D. Stebila, and B. Ustaoglu, “Anonymity and one-way authentication in key exchange protocols,” *Designs, Codes and Cryptography*, pp. 1–25, 2012, proposal for Tor: <https://gitweb.torproject.org/torspec.git/blob/HEAD:/proposals/ideas/xxx-ntor-handshake.txt>.
- [15] O. Goldreich and Y. Lindell, “Session-Key Generation Using Human Passwords Only,” in *CRYPTO*, 2001, pp. 408–432.
- [16] D. M. Goldschlag, M. Reed, and P. Syverson, “Hiding Routing Information,” in *Proc. 1st Workshop on Information Hiding*, 1996, pp. 137–150.
- [17] D. M. Goldschlag, M. G. Reed, and P. F. Syverson, “Onion Routing,” *Commun. ACM*, vol. 42, no. 2, pp. 39–41, 1999.
- [18] D. Hofheinz, “Possibility and Impossibility Results for Selective Decommitments,” *J. Cryptology*, vol. 24, no. 3, pp. 470–516, 2011.
- [19] A. Kate and I. Goldberg, “Distributed Private-Key Generators for Identity-Based Cryptography,” in *Proc. 7th Conference on Security and Cryptography for Networks (SCN)*, 2010, pp. 436–453.
- [20] A. Kate and I. Goldberg, “Using Sphinx to Improve Onion Routing Circuit Construction,” in *Proc. 14th Conference on Financial Cryptography and Data Security (FC)*, 2010, pp. 359–366.
- [21] A. Kate, G. M. Zaverucha, and I. Goldberg, “Pairing-Based Onion Routing,” in *Proc. 7th Privacy Enhancing Technologies Symposium (PETS)*, 2007, pp. 95–112.
- [22] A. Kate, G. M. Zaverucha, and I. Goldberg, “Pairing-Based Onion Routing with Improved Forward Secrecy,” *ACM Trans. Inf. Syst. Secur.*, vol. 13, no. 4, p. 29, 2010.
- [23] S. Mauw, J. Verschuren, and E. P. de Vink, “A Formalization of Anonymity and Onion Routing,” in *Proc. 9th European Symposium on Research in Computer Security (ESORICS)*, 2004, pp. 109–124.
- [24] L. Øverlier and P. Syverson, “Improving efficiency and simplicity of tor circuit establishment and hidden services,” in *Proc. 7th Privacy Enhancing Technologies Symposium (PETS)*, 2007, pp. 134–152.
- [25] M. Reed, P. Syverson, and D. Goldschlag, “Anonymous Connections and Onion Routing,” *IEEE J-SAC*, vol. 16, no. 4, pp. 482–494, 1998.
- [26] V. Shmatikov, “Probabilistic analysis of an anonymity system,” *Journal of Computer Security*, vol. 12, no. 3-4, pp. 355–377, 2004.
- [27] P. Syverson, G. Tsudik, M. Reed, and C. Landwehr, “Towards an Analysis of Onion Routing Security,” in *Proc. Workshop on Design Issues in Anonymity and Unobservability (WDIAU)*, 2000, pp. 96–114.
- [28] “The Tor Project,” <https://www.torproject.org/>, 2003, accessed Nov 2011.
- [29] D. Wikström, “A universally composable mix-net,” in *Proc. of the 1st Theory of Cryptography Conference (TCC)*, 2004, pp. 317–335.
- [30] Y. Zhang, “Effective attacks in the tor authentication protocol,” in *NSS '09*, 2009, pp. 81–86.