

Pairing-Based Onion Routing^{*}

Aniket Kate, Greg Zaverucha, and Ian Goldberg

David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1
{akate,gzaveruc,iang}@cs.uwaterloo.ca

Abstract. This paper presents a novel use of pairing-based cryptography to improve circuit construction in onion routing anonymity networks. Instead of iteratively and interactively constructing circuits with a telescoping method, our approach builds a circuit with a single pass. The cornerstone of the improved protocol is a new pairing-based privacy-preserving non-interactive key exchange. Compared to previous single-pass designs, our algorithm provides practical forward secrecy and leads to a reduction in the required amount of authenticated directory information. In addition, it requires significantly less computation and communication than the telescoping mechanism used by Tor. These properties suggest that pairing-based onion routing is a practical way to allow anonymity networks to scale gracefully.

1 Introduction

The concept of onion routing [27] plays a key role in many efforts to provide anonymous communication. In the world of cryptographic protocols, bilinear pairings [9] have also had comparable impact. Their meeting is not surprising. This paper applies pairing-based cryptographic techniques—namely non-interactive key agreement—to the problem of session key establishment in anonymity networks based on onion routing. We show that this approach offers better performance, evidenced by reduced computational cost and fewer network communications. This improved performance is of particular interest to low-latency anonymity networks, as it increases responsiveness and network capacity. While using fewer resources for cryptography, we are careful to simultaneously meet the security goals provided by existing methods.

1.1 Our Contributions

This paper makes four primary contributions in the field of anonymous communication.

^{*} This is an extended version of our PET 2007 paper [15], including formal proofs of security and anonymity for our one-way and two-way anonymous key agreement protocols.

1. We define a privacy-preserving key agreement protocol using bilinear pairings in an identity-based infrastructure. We adapt it to achieve unilateral (one-way) anonymity with non-interactive key agreement and prove the security and anonymity of these protocols.
2. We then use our one-way anonymous key agreement protocol to build onion routing circuits for anonymity networks like Tor [7]. Our protocol constructs a circuit in a single pass and also provides a practical way to achieve forward secrecy.
3. The performance of our circuit construction protocol surpasses that of Tor, requiring significantly less computation and fewer network communications.
4. Our protocol does not require the public keys of onion routers to be authenticated. This reduces the load on directory servers and improves the scalability of anonymity networks.

The anonymous authentication scheme we present extends the non-interactive key agreement scheme of Sakai, Ohgishi, and Kasahara [29]. Previous work related to pairing-based key exchange, as well as to anonymity networks, is covered in Section 2. We describe the cryptographic protocols in Section 3, and an onion routing system built with a Boneh-Franklin identity-based infrastructure in Section 4. Some of the more practical issues in such a system are discussed in Section 5 and we compare our computational and communications costs to those of Tor in Section 6. Section 7 concludes the discussion and we analyze the security and anonymity of our protocols in the appendix.

2 Related Work

Over the years, a large number of anonymity networks have been proposed and some have been implemented. Common to many of them is *onion routing*, a technique whereby a message is wrapped in multiple layers of encryption, forming an *onion*. As the message is delivered via a number of intermediate *onion routers* (ORs), or *nodes*, each node decrypts one of the layers, and forwards the message to the next node. This idea goes back to Chaum [3] and has been used to build both low- and high-latency communication networks. Formalizations and security discussions of onion routing can be found in [2, 19, 22, 32].

A common realization of an onion routing system is to arrange a collection of nodes that will relay traffic for users of the system. Some examples are [5, 7, 10, 27, 28] (the related work section of [7] contains a thorough list). To date, the largest onion routing system is Tor, which has approximately 1000 onion routers and hundreds of thousands of users [33]. These numbers (and their growth) underscore the demand for anonymity online.

To use a network of onion routers, users randomly choose a path through the network and construct a *circuit*—a sequence of nodes which will route traffic. After the circuit is constructed, each of the nodes in the circuit shares a symmetric key with the user, which will be used to encrypt the layers of future onions. In the original Onion Routing project [14, 27, 32] (which was superseded by Tor) circuit construction was done as follows. The user created an onion where each

layer contained the symmetric key for one node and the location of the next node, all encrypted with the original node’s public key. Each node decrypts a layer, keeps the symmetric key and forwards the rest of the onion along to the next node. The main drawback of this approach is that it does not provide forward secrecy (as defined in [7]). Suppose a circuit is constructed from the user to the sequence of nodes $A \leftrightarrow B \leftrightarrow C$, and that A is malicious. If A records the traffic, and at a later time compromises B (at which point he learns the next hop is C), then compromises C , the complete route is known, and A learns who the user has communicated with.

A possible fix for this problem is to frequently change the public keys of each node. This limits the amount of time A has to compromise B and C , but requires that the users of the system frequently contact the directory server to retrieve authentic keys. Later systems constructed circuits incrementally and interactively (this process is sometimes called *telescoping*). The idea is to use the node’s public key only to initiate a communication during which a temporary session key is established via the Diffie-Hellman key exchange. Tor constructs circuits in this way, using the Tor authentication protocol (TAP). TAP is described and proven secure in previous work of the last author [13].

Trade-offs exist between the two methods of constructing circuits. Forward secrecy is the main advantage of telescoping, but telescoping also handles nodes that are not accepting connections; if the third node is down during the construction of a circuit, for example, the first two remain, and the user only needs to choose an alternate third. Information about the status and availability of nodes is therefore less important. The drawback of telescoping is the cost; establishing a circuit of length ℓ requires $O(\ell^2)$ network communications, and $O(\ell^2)$ symmetric encryptions/decryptions.

Overlier and Syverson [24] improve the efficiency of telescoping-based circuit construction using a half-certified Diffie-Hellman key exchange [21, Sec. 12.6]. They further define an efficient single-pass circuit construction and a few variants. The proposed variants offer different levels of forward secrecy, which is traded off against computation and communication. For example, their eventual forward secret variants use frequent rotation of nodes’ public keys, presenting the same issues as the first generation onion routing; their immediate forward secrecy variant uses the same amount of communication as the current Tor ($O(\ell^2)$), but less computation.

Privacy-preserving authentication schemes can be one- or two-way (also referred to as unilateral or bilateral). After one-way authentication between Anonymous and Bob, Anonymous has confirmed Bob’s identity and Bob learns nothing about Anonymous, except perhaps that he or she is a valid user of a particular system. In a two-way scheme, both users can confirm they are both valid users without learning who the other is.

The work of Okamoto and Okamoto [23] presents schemes for anonymous authentication and key agreement. In Rahman et. al. [26], an anonymous authentication protocol is presented as part of an anonymous communication system for mobile ad-hoc networks. The protocols in both papers are complex, and limited

motivation is given for design choices. Further, both papers neglect to discuss the security of their proposed protocols. The protocols we present in Section 3.2 are a great deal simpler than previous protocols. This allows them to be more easily understood, and simplifies the discussion of their security, which appears in Section 3.3.

Previous protocols (as well as ours) owe a lot to the non-interactive key exchange protocol of Sakai, Ohgishi and Kasahara [29]. In the next section, we will review their scheme after covering relevant background material.

3 Pairing-Based Key Agreement with User Anonymity

In one of the pioneering works of pairing-based cryptography, Sakai et al. suggested an identity-based, non-interactive key agreement scheme using bilinear pairings [29]. In this section, we extend this key agreement scheme. We replace the identities of the participants by pseudonyms and our new scheme provides unconditional anonymity to participating users.

3.1 Preliminaries

We briefly review bilinear pairings and the original non-interactive key agreement scheme of Sakai et al. For a detailed presentation of pairings and cryptographic applications thereof see Blake et al. [9] and references therein.

Bilinear Pairings. Consider two additive cyclic groups \mathbb{G} and $\hat{\mathbb{G}}$ and a multiplicative cyclic group \mathbb{G}_T , all of the same prime order n . A bilinear map e is a map $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$ with following properties.

1. **Bilinearity:** For all $P \in \mathbb{G}$, $Q \in \hat{\mathbb{G}}$ and $a, b \in \mathbb{Z}_n$, $e(aP, bQ) = e(P, Q)^{ab}$.
2. **Non-degeneracy:** The map does not send all pairs in $\mathbb{G} \times \hat{\mathbb{G}}$ to unity in \mathbb{G}_T .
3. **Computability:** There is an efficient algorithm to compute $e(P, Q)$ for any $P \in \mathbb{G}$ and $Q \in \hat{\mathbb{G}}$.

Our protocols, like many pairing-based cryptographic protocols, use a special form of bilinear map called a *symmetric pairing* which has $\mathbb{G} = \hat{\mathbb{G}}$. For such pairings $e(P, Q) = e(Q, P)$ for any $P, Q \in \mathbb{G}$. The modified Weil pairing over elliptic curve groups [34] is an example of a symmetric bilinear pairing. In the rest of the paper, unless otherwise specified, all bilinear pairings are symmetric.

The Bilinear Diffie-Hellman Assumption. The *Bilinear Diffie-Hellman* (BDH) problem is to compute $e(P, P)^{abc} \in \mathbb{G}_T$ given a generator P of \mathbb{G} and elements aP, bP, cP for $a, b, c \in \mathbb{Z}_n^*$. An equivalent formulation of the problem, due to the bilinearity of the map, is to compute $e(A, B)^c$ given a generator P of \mathbb{G} , and elements A, B and cP .

If there is no efficient algorithm to solve the BDH problem for $\langle \mathbb{G}, \mathbb{G}_T, e \rangle$, they are considered to satisfy the *BDH assumption*.

Boneh-Franklin Setup and Non-Interactive Key Agreement. In a Boneh-Franklin Identity-Based Encryption (BF-IBE) setup [1], a trusted authority, called a private key generator (PKG), generates private keys (d_i) for clients using their well-known identities (ID_i) and a master secret s . A client with identity ID_i receives the private key $d_i = sH(ID_i) \in \mathbb{G}$, where $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$ is a full-domain cryptographic hash function and \mathbb{G}^* denotes the set of all elements in \mathbb{G} except the identity.

Sakai et al. observed that, with such a setup, any two clients of the same PKG can compute a shared key using only the identity of the other participant and their own private keys. Only the two clients and the PKG can compute this key. For two clients with identities ID_A and ID_B , the shared key is given by $K_{AB} = e(Q_A, Q_B)^s = e(Q_A, d_B) = e(d_A, Q_B)$ where $Q_A = H(ID_A)$ and $Q_B = H(ID_B)$.

Dupont and Enge proved this protocol is secure in the random oracle model assuming the BDH problem in $\langle \mathbb{G}, \mathbb{G}_T, e \rangle$ is hard [8].

3.2 Anonymous Key Agreement

We observe that by replacing the identity hashes with pseudonyms generated by users, a key agreement protocol with unconditional anonymity is possible. In our protocol, a participant can confirm that the other participant is a client of the same PKG, but can not determine his identity. Each client can randomly generate many possible pseudonyms and the corresponding private keys.

Suppose Alice, with (identity, private key) pair (ID_A, d_A) , is seeking anonymity. She generates a random number r_A and creates the pseudonym and corresponding private key $(P_A = r_A Q_A = r_A H(ID_A), r_A d_A = s P_A)$. In a key agreement protocol, she sends the pseudonym P_A instead of her actual identity to another participating client, who may or may not be anonymous. For two participants (say Alice and Bob) with pseudonyms P_A and P_B , the shared session key is given as

$$K_{AB} = e(P_A, P_B)^s = e(Q_A, Q_B)^{r_A r_B s}$$

where r_A and r_B are random numbers generated respectively by Alice and Bob. If Bob does not wish to be anonymous, he can just use $r_B = 1$ instead of a random value, resulting in $P_B = Q_B$. If persistent pseudonymity is desired instead of anonymity, the random values can easily be reused.

Two participants can perform a session key agreement by exchanging pseudonyms. Further, two participants can also perform an authenticated key agreement by modifying any secure symmetric-key based mutual authentication protocol and simply replacing their identities by their pseudonyms.

One-Way Anonymous Key Agreement. Anonymous communication generally requires anonymity for just one of the participants; the other participant often works as a service provider and the anonymous participant needs to confirm her identity. In the key agreement protocol, the service provider uses her

actual identity rather than a pseudonym. Further, in this one-way anonymity setting two participants can agree on a session key in a non-interactive manner. A non-interactive scheme to achieve this is defined next.

Suppose Alice and Bob are clients of a PKG. As before, Alice has identity ID_A and private key $d_A = sQ_A = sH(ID_A)$. Alice wishes to remain anonymous to Bob, but she knows Bob's identity ID_B .

1. Alice computes $Q_B = H(ID_B)$. She chooses a random integer $r_A \in \mathbb{Z}_n^*$, generates the corresponding pseudonym $P_A = r_A Q_A$ and private key $r_A d_A = sP_A$, and computes the session key $K_{AB} = e(sP_A, Q_B) = e(Q_A, Q_B)^{sr_A}$. She sends her pseudonym P_A to Bob.
2. Bob, using P_A and his private key d_B , computes the session key $K_{AB} = e(P_A, d_B) = e(Q_A, Q_B)^{sr_A}$.

Note that in step 1, Alice can also include a message for Bob symmetrically encrypted with the session key; we will use this in Section 4. Note also that in practice, the session key is often derived from K_{AB} , and not K_{AB} itself.

Key Authentication and Confirmation. In most one-way anonymous communication situations, it is also required to authenticate the non-anonymous service provider. With the non-interactive protocols of this section, the key is implicitly authenticated; Alice is assured that only Bob can compute the key. If Alice must be sure Bob has in fact computed the key, explicit key confirmation can be achieved by incorporating any symmetric-key based challenge-response protocol.

3.3 Security and Anonymity

In this section, we discuss the security and anonymity of our key agreement schemes in the random oracle model. We make following claims:

Unconditional Anonymity: It is impossible for the other participant in a protocol run, the PKG or any third party to learn the identity of an anonymous participant in a protocol run.

Session Key Secrecy: It is infeasible for anyone other than the two participants or the PKG to determine a session key generated during a protocol run.

No Impersonation: It is infeasible for a malicious client of the PKG to impersonate another (non-anonymous) client in a protocol run. In the case of persistent pseudonymity, it is not feasible for a malicious entity to communicate using a different entity's pseudonym.

Next, we present informal proofs for each of our claims. For complete security proofs, we refer the reader to the appendix.

Unconditional Anonymity. For an anonymous client with identity ID_C , the pseudonym $P_C = r_C Q_C \in \mathbb{G}$ is the only parameter exchanged during the protocol that is derived from her identity. Because \mathbb{G} is a cyclic group of prime order, multiplying by the random r_C perfectly blinds the underlying identity. The anonymity set is restricted to the clients of a PKG, unless a random pair $(U, d_U) \in \mathbb{G}$ is made public. In the latter case, anyone can generate a pseudonym and participate in the protocol using (U, d_U) .

Session Key Secrecy. Dupont and Enge [8] prove the security of the key agreement scheme of Sakai et al. in the random oracle model. According to this proof, an attacker cannot compute the shared key if the BDH assumption holds on $\langle \mathbb{G}, \mathbb{G}_T, e \rangle$, and H is modelled by a random oracle. Our protocol simply modifies that of Sakai et al. to use $P_i = H'(ID_i)$ instead of $Q_i = H(ID_i)$, where $H'(x) = r_i \cdot H(x)$ for a random value r_i , so the proof of security in [8] is easily modified to suit our protocol.

No Impersonation. Suppose an adversarial client with ID_{adv}, d_{adv} wishes to impersonate a non-anonymous participant (say, Bob with ID_B) while communicating with an anonymous client with pseudonym P_A . The adversary would need to compute $K_{AB} = e(P_A, Q_B)^s$ given P_A, Q_B, Q_{Adv} and sQ_{Adv} . But this is just the BDH problem, so under the BDH assumption on $\langle \mathbb{G}, \mathbb{G}_T, e \rangle$, impersonation of other clients is infeasible.

Similarly, if the adversary wishes to communicate with Bob using the persistent pseudonym P_A of some other pseudonymous entity, it must compute $K_{AB} = e(P_A, Q_B)^s$ given P_A, Q_B, Q_{Adv} and sQ_{Adv} . Again, the adversary must solve the BDH problem.

3.4 Distributed PKG

The PKG in the BF-IBE framework, with the master key, has the power to decrypt all messages encrypted for clients. As our schemes use the same setup as BF-IBE, the PKG can compute a session key from the publicly available pseudonyms and the master key s . Due to this, compromise of the PKG is a single point of failure for security.

Boneh and Franklin suggest the use of a distributed PKG instead a single PKG to mitigate this problem. Their distributed PKG uses t out of m threshold cryptography [31], which involves distributing the master key information among m PKGs, such that any t of them, but no fewer, can compute the master key or generate a private key for a client. Their key distribution scheme uses a dealer who actually decides the master key and thus becomes a candidate for attack and can be a single point of failure. Instead, we suggest the use of a distributed key generation protocol such as that of Pedersen [25] or Gennaro et al. [12]. In these protocols, a master key is generated in a completely distributed way with each of m PKGs contributing a random share. The distributed design is additionally more robust; at any given time only t of the m PKGs must be online in order for a client to retrieve his private key.

3.5 Applications of Our Anonymity Schemes

Our anonymous key agreement schemes can be used to perform anonymous communication in any setting having a BF-IBE setup. In recent years, numerous BF-IBE based solutions have been suggested for various practical situations, such as ad-hoc networks. [4, 16, 30] Our anonymous key agreement schemes can be used in all of these setups without any extra effort. In this paper, we focus on a new pairing-based onion routing protocol which achieves forward secrecy and constructs circuits without telescoping. We describe this protocol in the next section.

4 Pairing-Based Onion Routing

Low-latency onion routing requires one-way anonymous key agreement and forward secrecy. In this section, we describe a new pairing-based onion routing protocol using the non-interactive key agreement scheme defined in Section 3.2.

Our onion routing protocol has a significant advantage over the original onion routing protocol [14] as well as the protocol used in Tor [7]; it provides a practical way to achieve forward secrecy without building circuits by telescoping. Though this is possible with the original onion routing protocol, that method involves regularly communicating authenticated copies of ORs' public keys to the system users; forward secrecy is achieved by periodically rotating these keys. This does not scale well; every time the public keys are changed *all* users must contact a directory server to retrieve the new authenticated keys. However, our onion routing protocol uses ORs' identities, which users can obtain or derive without repeatedly contacting a central server, thus providing practical forward secrecy without telescoping.

4.1 Design Goals and Threat Model

As our protocol only differs from existing onion routing protocols in the circuit construction phase, our threat model is that of Tor. For example, adversaries have complete control over some part (but not all) of the network, as well as control over some of the nodes themselves.

We aim at frustrating attackers from linking multiple communications to or from a single user. Like Tor, we do not try to develop a system secure against a global observer, which can in theory follow end-to-end traffic. Further, it should not be feasible for any node to determine the identity of any node in a circuit other than its two adjacent nodes. Finally, we require forward secrecy: after some amount of time, the session keys used to protect node identities and the contents of messages are irrecoverable, even if all participants in the network are subsequently compromised.

4.2 Pairing-Based Onion Routing Protocol

An onion routing protocol involves a service provider, a set of onion routers, and users. In our protocol, a user does not build the circuit incrementally via

telescoping, but rather in a single pass. The user chooses ℓ ORs from the available pool and generates separate pseudonyms for communicating with each of them. The user computes the corresponding session keys and uses them to construct a message with ℓ nested layers of encryption. This process uses the protocol given in Section 3.2 ℓ times.

The service provider works as the PKG for the ORs and provides private keys for their identities.

Forward Secrecy. There are two time-scale parameters in our protocol: the *master key validity period* (MKVP) and the *private key validity period* (PKVP). Both of these values relate to the forward secrecy of the system. The PKVP specifies how much exposure time a circuit has against compromises of the ORs that use it. That is, until the PKVP elapses, the ORs have enough information to collectively decrypt circuit construction onions sent during that PKVP. After each PKVP, ORs discard their current private keys and obtain new keys from the PKGs. This period can be short, perhaps on the order of an hour.

The MKVP specifies the circuit’s exposure time against compromises of the (distributed) PKG which reveal the master secret s . Because changing s involves the participation of all of the ORs as well as the PKGs, we suggest the MKVP be somewhat longer than the PKVP, perhaps on the order of a day. Remember that in the t of m distributed PKG, if at least $m - t + 1$ PKG members are honest and not compromised, no one will ever learn the value of a master secret.

Protocol Description. As discussed above, we propose the use of a distributed PKG, but for simplicity, our discussion will consider the PKG to be a single entity. Using a distributed PKG affects only the setup and key generation steps.

Setup: Given the security requirements, the PKG generates a digital signature key pair (for any secure digital signature scheme). It also generates a prime n , two groups \mathbb{G} (written additively) and \mathbb{G}_T (written multiplicatively) of order n and a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. Finally, the PKG chooses a full-domain cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$. The PKG publishes all of these values except its private signature key.

Key Generation: For each MKVP, the PKG generates a random master key $s \in \mathbb{Z}_n^*$ and a random $U \in \mathbb{G}$, and calculates sU . The PKG publishes a signed copy of (v_m, U, sU) , where v_m is a timestamp for the MKVP in question. This U is a common value to be shared by all users of the system.

For every valid OR with identity ID_i , and for every PKVP v that overlaps with the MKVP, the PKG generates the private key $d_{vi} = sH(v||ID_i)$. The PKG distributes these private keys, as well as a copy of the signed (v_m, U, sU) , to the appropriate ORs over a secure authenticated forward-secret channel. If an OR becomes compromised, the PKG can revoke it by simply no longer calculating its values of d_{vi} .

Note that this key distribution can be *batched*; that is, the PKG can precompute the master keys and private keys in advance (say a week at a

time), and deliver them to the ORs in batches of any size from one PKVP at a time on up. This batching reduces the amount of time the PKG has to be online, and does not sacrifice forward secrecy. On the other hand, large batches will delay the time until a revocation becomes effective.

User Setup: Once every MKVP v_m , each user must obtain a new signed tuple (v_m, U, sU) from any OR or from a public website. Once every PKVP v , the user computes the following pairing for each OR i and stores the results locally:

$$\gamma_{vi} = e(sU, Q_{vi}) = e(U, Q_{vi})^s \text{ where } Q_{vi} = H(v || \text{ID}_i)$$

Circuit Construction: During a PKVP v , a user U chooses a set of ORs (say A, B, \dots, N) and constructs a circuit $U \Leftrightarrow A \Leftrightarrow B \Leftrightarrow \dots \Leftrightarrow N$ with the following steps.

1. For each OR i in the circuit, the user generates a random integer $r_i \in \mathbb{Z}_n^*$ and computes the pseudonym $P_{U_i} = r_i U$ and the value $\gamma_{vi}^{r_i} = e(U, Q_{vi})^{sr_i}$. From $\gamma_{vi}^{r_i}$ two session keys are derived: a forward session key K_{U_i} and a backward session key K_{iU} . Finally, the following onion is built and sent to A , the first OR in the circuit:

$$r_A U, \{B, r_B U, \{\dots \{N, r_N U, \{\emptyset\}_{K_{UN}}\} \dots\}_{K_{UB}}\}_{K_{UA}}$$

Here $\{\dots\}_{K_{U_i}}$ is symmetric-key encryption and \emptyset is an empty message which informs N that it is the exit node.

2. After receiving the onion, the OR with identity ID_i uses the received $r_i U$ and its currently valid private key d_{vi} to compute $e(r_i U, d_{vi}) = e(U, Q_i)^{r_i s} = \gamma_{vi}^{r_i}$. It derives the forward session key K_{U_i} and the backward session key K_{iU} . It decrypts the outermost onion layer $\{\dots\}_{K_{U_i}}$ to obtain the user's next pseudonym, the nested ciphertext, and the identity of the next node in the circuit. The OR then forwards the pseudonym and ciphertext to the next node. To avoid replay attacks, it also stores pseudonyms (see Section 5). The process ends when an OR (N in this case) gets \emptyset .
3. The exit node N sends a confirmation message encrypted with the backward session key $\{Confirm\}_{K_{NU}}$ to the previous OR in the circuit. Each OR encrypts the confirmation with its backward session key and sends it to the previous node, until the ciphertext reaches the user. The user decrypts the ciphertext layers to verify the confirmation.
4. If the user does not receive the confirmation in a specified time, she selects a different set of ORs and repeats the protocol.

The circuit construction is further illustrated in Figure 1, where a user builds a three-node circuit.

Anonymous Communication: After the circuit is constructed, communication proceeds in the same manner as Tor. The user sends onions through the

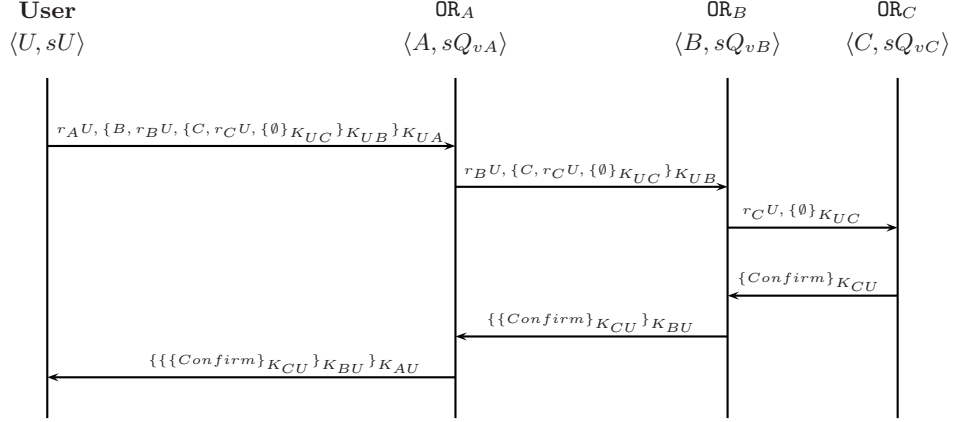


Fig. 1. A user builds a circuit with three ORs.

circuit with each layer encrypted with the forward keys K_{U_i} , and each hop decrypts one layer. Replies are encrypted at each hop with the backward key K_{iU} , and the user decrypts the received onion.

Note that as an optimization, one or more messages can be bundled inside the original circuit construction onion, in place of \emptyset .

4.3 Security Analysis

Camenisch and Lysyanskaya [2] formally define the requirements of a secure onion routing construction in the universal composability (UC) framework in terms of onion-correctness, onion-integrity and onion-security. We observe that our circuit construction trivially achieves onion-correctness and onion-integrity.

They also define a secure (in the UC framework) onion routing circuit construction using any IND-CCA-2 public key encryption scheme. As an encryption in our circuit construction ($r_{U_i}, \{\dots\}_{K_{U_i}}$) is a generalization of the **BasicIdent** scheme ($r_{U_i}, \{\dots\} \oplus \{K_{U_i}\}$) from BF-IBE [1], it is IND-CPA secure. Along similar lines to [1], we can use a technique due to Fujisaki-Okamoto [11] to convert our scheme to an IND-CCA-2 construction: $(r_{U_i}, \{\sigma\}_{K_{U_i}}, \{\dots\}_{H'(\sigma)})$ for a random binary string σ and a cryptographic hash H' . Therefore, a combination of this IND-CCA-2 encryption and the Camenisch-Lysyanskaya circuit construction is secure in the UC framework. But such a circuit construction is less efficient than ours, and we consider proving onion-security for our circuit construction defined in Section 4.2 to be important future work.

4.4 Advantages Over First-Generation Onion Routing

As discussed earlier, it is possible to achieve forward secrecy in first-generation onion routing by periodically replacing the public-private key pairs of the ORs.

Following the change, the service provider publishes signed copies of the new OR public keys after getting authentic copies from the ORs. However, this requires all users to regularly obtain fresh authenticated public key information for all ORs.

In contrast, with our system, each user only needs to obtain the single authenticated value (v_m, U, sU) , and only once every MKVP. The user can then calculate the required γ_{vi} values on her own until the end of that period, thus reducing the load on the service provider. This load is further reduced by having the service provider never communicate directly with users at all, but only with the ORs.

As a consequence, our pairing-based onion routing is a more practical solution for low-latency anonymous communication.

4.5 Advantages Over Telescoping in Tor

The Tor network, in practice, uses the telescoping approach based on the Diffie-Hellman key exchange to form an anonymity circuit. We find the following advantages for our protocol over the telescoping approach.

- Although our above-defined protocol requires occasional private key generation for ORs to achieve forward secrecy, it saves communication cost at every circuit construction by avoiding telescoping. We discuss our communication and computational advantages in Section 6.4.
- The absence of telescoping in our protocol provides flexibility to the user to modify a circuit on the fly. For example, suppose a user U has constructed a circuit $(U \Leftrightarrow A \Leftrightarrow B \Leftrightarrow \dots \Leftrightarrow K \Leftrightarrow \dots \Leftrightarrow N)$. In our protocol, she can bundle instructions to immediately replace K with K' in the next message, while keeping the remaining circuit intact. Her circuit would then be $(U \Leftrightarrow A \Leftrightarrow B \Leftrightarrow \dots \Leftrightarrow K' \Leftrightarrow \dots \Leftrightarrow N)$.

4.6 Issues with the Proposed Scheme

The certifying authorities in the Tor system need to be less trusted than the PKG in our scheme. With a short PKVP and MKVP (compared to the key replacement period in Tor), our PKGs (any t of them) need to be online with greater reliability. Further, if fewer than t are available, the whole system is paralysed after the current batch.

It is also possible for t malicious PKGs to passively listen to all of the traffic as they can compute private keys for all ORs. A geographically and politically distributed implementation of m PKGs certainly reduces this possibility.

To passively decrypt an OR's messages, an adversary of the Tor system must know the OR's private key, as well as the current Diffie-Hellman key (established for each circuit). In our scheme, as it is non-interactive, an adversary who knows only the OR's private key can decrypt all of the messages for that OR. This may be an acceptable trade-off, considering the advantages gained from the non-interactive protocol.

5 Systems Issues

In this section, we describe how components of an onion routing system such as Tor would behave in a pairing-based setting. To implement pairings, we must choose groups where pairings are known, and are efficiently computable. Once these groups are fixed we can estimate the computational cost required to construct a circuit. The next section will compare the cost of our scheme to the cost of setting up a circuit in Tor.

PKG. As discussed in Section 3.4, the PKG should be distributed across servers run by independent parties. To provide robustness, a “ t of m ” secret sharing scheme may be employed; this would mean that an OR need only contact t of m “pieces” of the PKG to learn its complete private key. Naturally, private key information must always be communicated over a secure channel. We note that end users of the system will have no reason to contact the PKG; the PKG only communicates with ORs, and sends one private key (an element of \mathbb{G}) per PKVP to each. The load on the PKG should therefore be quite manageable. For added protection from attack, the PKG could even situate itself as a “hidden service” [7, §5], so that only known ORs could even connect to it, and no one would know where many of the pieces were located.

Channel Security. The security and forward secrecy depends on the channel between the PKG and the OR used to compute the private key. With a non-distributed PKG, an attacker can compromise an OR’s private key by compromising this channel. The distributed PKG provides robustness here as well, since the attacker must subvert t secure channels to reconstruct the private key from the shares.

Onion Router Identities. Users calculate γ_{vi} based on each router’s identity ID_i . This identity can be as simple as a port number and a hostname or IP address. In that case, the BF-IBE setup ensures that if a user knows how to contact an OR, she automatically knows its public key.

The value γ_{vi} is also based on the current PKVP v . To avoid requiring tight synchronization between the clocks of ORs and users, ORs should keep their private keys d_{vi} around for a short time after the official end of the PKVP, but must securely discard them after that.

Replay Prevention. To avoid attacks where adversaries replay old circuit construction onions, ORs should store the pseudonyms they receive for the duration of a PKVP and drop onions which re-use a pseudonym. After circuit construction, replay attacks can be prevented with existing methods (see [6] for an example).

Directory Servers. Directory servers can be used to provide signed information about the list of available ORs to the users of the system. The directory servers in Tor, for example, provide a list of the ORs along with their public keys, status, capabilities and policies. In our pairing-based setting, of course, the public keys are unnecessary.

6 Performance

In this section, we consider the cost of creating a circuit from a user through ℓ onion routers. We estimate the computational cost, and count the number of AES-encrypted network communications. We compare the performance of our system to that of Tor.

6.1 Security Levels and Parameter Sizes

Before comparing the costs of the cryptography in both schemes we determine the parameter sizes required to provide the same level of security currently provided by Tor.

Tor uses public key parameters to provide security at the 80-bit level [13]. The discrete log problem is in a 1024-bit field, and the RSA problem is also at the 1024-bit level. The symmetric parameters provide significantly more security, by using AES with a 128-bit key.

We must choose appropriate groups \mathbb{G} and \mathbb{G}_T over which our pairing will be defined, in order to offer similar strength. The current favourite choice is the group of torsion points of an elliptic curve group over a finite field, with either the Weil or Tate pairing. To achieve an 80-bit security level, the elliptic curve discrete log problem an attacker faces must be in a group of at least 160 bits. Due to the reduction of Menezes, Okamoto and Vanstone [20], we must also ensure that discrete logs are intractable in the target group, \mathbb{G}_T . In our case, $\mathbb{G}_T = \mathbb{F}_{p^k}$, where k is the embedding degree of our curve taken over \mathbb{F}_p . We must then choose our curve E , a prime p , and embedding degree k such that $E(\mathbb{F}_p)$ has a cyclic subgroup of prime order $n \approx 2^{160}$, and p^k is around 2^{1024} . This can be achieved in a variety of ways, but two common choices are $k = 2, p \approx 2^{512}$ and $k = 6, p \approx 2^{171}$. Pairing implementations with both sets of parameters are available in the PBC library [18]. Efficiency studies suggest that $k = 2$ and the Tate pairing can offer better performance at this security level [17], so we make that choice.

6.2 Cost of Building a Circuit with Tor

Tor builds circuits by telescoping. A user Uriel chooses a Tor node (say Alice), and establishes a secure channel using an encrypted Diffie-Hellman exchange. She then picks a second node, Bob, and over this secure channel, establishes a new secure channel to Bob with another (end-to-end) encrypted Diffie-Hellman exchange. She proceeds in this manner until the circuit is of some desired length

ℓ . For details, see the Tor specification [6]. Note that Uriel cannot use the same Diffie-Hellman parameters with different nodes, lest those nodes be able to determine that the same user was communicating with each of them.

Each Diffie-Hellman exchange requires Uriel to perform two modular exponentiations with 1024-bit moduli and 320-bit exponents. Likewise, each server also performs two of these exponentiations. Uriel RSA encrypts the Diffie-Hellman parameter she sends the server, and the server decrypts it. The AES and hashing operations involved have negligible costs compared to these.

Uriel’s circuit construction to Alice takes two messages: one from Uriel to Alice, and one from Alice to Uriel. When Uriel extends this circuit to Bob (via Alice), there are four additional messages: Uriel to Alice, Alice to Bob, Bob to Alice, and Alice to Uriel. Continuing in this way, we see that the total number of messages required for Tor to construct a circuit of length ℓ is $\ell(\ell + 1)$. Note that each of these messages needs to be encrypted and decrypted at each hop.

6.3 Cost of Building a Circuit with Pairing-Based Onion Routing

In order to create a circuit of length ℓ with our scheme, the user Uriel must choose ℓ random elements r_i of \mathbb{Z}_n^* . As above, Uriel should not reuse these values. She then computes $r_S U$ and $\gamma_S^{r_S}$, and derives the forward and backward keys K_{US} and K_{SU} from $\gamma_S^{r_S}$, for each server S in the circuit. Each server computes $e(r_S U, d_S) = \gamma_S^{r_S}$ for its current private key d_S and derives K_{US} and K_{SU} .

Uriel creates one message, as in Figure 1, and sends it to the first server in the chain. This server decrypts a layer and sends the result to the second server in the chain, and so on, for a total of ℓ hop-by-hop encrypted messages. At the end of the chain, the last server replies with a confirmation message that travels back through the chain, producing ℓ more messages, for a total of 2ℓ .

6.4 Comparison and Discussion

We summarize the results of the previous two sections in Table 1. We count the number of “bignum” operations for each of the client and the servers, both for Tor and for our pairing-based onion routing protocol. We ignore the comparatively negligible computational costs of AES operations and hashing.

For each bignum operation, we include a benchmark timing. These timings were gathered on a 3.0 GHz Pentium D desktop using the PBC pairing-based cryptography library [18]. We can see that the total computation time to construct a circuit of length ℓ using our method is 61% less on the client side and 49% less on the server side as compared to using Tor. In addition, our method uses only a linear number of AES-encrypted messages, while Tor uses a quadratic number.

7 Conclusion

We have presented a new pairing-based approach for circuit construction in onion routing anonymity networks. We first extended the protocol of Sakai et al. [29]

Operation	Time	Tor		PB-OR	
		client	each server	client	each server
Pairing	2.9 ms	0	0	0	1
RSA decryption	2.7 ms	0	1	0	0
Modular exponentiation	1.5 ms	2ℓ	2	0	0
Multiplication in \mathbb{G}	1.0 ms	0	0	ℓ	0
Exponentiation in \mathbb{G}_T	0.2 ms	0	0	ℓ	0
RSA encryption	0.1 ms	ℓ	0	0	0
Total time (ms)		3.1ℓ	5.7	1.2ℓ	2.9
Total AES-encrypted messages		$\ell(\ell + 1)$		2ℓ	

Table 1. Comparison of costs of setting up a circuit of length ℓ . The values in the Tor column are based on the Tor specification [6]. PB-OR is our pairing-based onion routing scheme.

to allow for one-way or two-way anonymous or pseudonymous key agreement. We then used this extension to produce a new circuit construction protocol for onion routing networks. Our new pairing-based protocol creates circuits in a single pass, and also provides forward secrecy.

This protocol uses significantly less computation and communication than the corresponding protocol in Tor, and reduces the load on the network support infrastructure. These improvements can be used to enhance the scalability of low-latency anonymity networks.

Acknowledgements. We would like to thank the Natural Sciences and Engineering Research Council of Canada for supporting this research with a Discovery Grant (Kate, Goldberg) and a PGS-D postgraduate scholarship (Zaverucha). We thank the anonymous reviewers for their constructive feedback, which improved the quality of this paper. We would also like to thank Sk. Md. Mizanur Rahman for providing us with an advance copy of the proceedings version of [26].

References

1. D. Boneh and M. Franklin. Identity-Based Encryption from the Weil Pairing. In *Advances in Cryptology—CRYPTO 2001, Lecture Notes in Computer Science 2139*, pages 213–229. Springer-Verlag, August 2001.
2. J. Camenisch and A. Lysyanskaya. A Formal Treatment of Onion Routing. In *Advances in Cryptology—CRYPTO 2005, Lecture Notes in Computer Science 3621*, pages 169–187. Springer-Verlag, August 2005.
3. D. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, 4(2):84–88, February 1981.
4. H. Chien and R. Lin. Identity-based Key Agreement Protocol for Mobile Ad-hoc Networks Using Bilinear Pairing. In *IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC’06)*, pages 520–529, 2006.

5. W. Dai. PipeNet 1.1. Post to Cypherpunks mailing list, November 1998.
6. R. Dingledine and N. Mathewson. The Tor Protocol Specification. <http://tor.eff.org/svn/trunk/doc/spec/tor-spec.txt>. Accessed February 2007.
7. R. Dingledine, N. Mathewson, and P. Syverson. Tor: The Second-Generation Onion Router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.
8. R. Dupont and A. Enge. Provably secure non-interactive key distribution based on pairings. *Discrete Applied Mathematics*, 154(2):270–276, 2006.
9. I. Blake (Editor). *Advances in Elliptic Curve Cryptography*. Number 317 in London Mathematical Society Lecture Note Series. Cambridge University Press, 2005.
10. M. J. Freedman and R. Morris. Tarzan: A Peer-to-Peer Anonymizing Network Layer. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS 2002)*, Washington, DC, November 2002.
11. Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *Advances in Cryptology—CRYPTO 2001*, pages 537–554, 1999.
12. R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure Distributed Key Generation for Discrete-Log Based Cryptosystems. *Journal of Cryptology*, 20(1):51–83, 2007.
13. I. Goldberg. On the Security of the Tor Authentication Protocol. In *Proceedings of the Sixth Workshop on Privacy Enhancing Technologies (PET 2006)*, *Lecture Notes in Computer Science 4258*, pages 316–331. Springer-Verlag, June 2006.
14. D. Goldschlag, M. Reed, and P. Syverson. Hiding Routing Information. In *Proceedings of Information Hiding: First International Workshop, Lecture Notes in Computer Science 1174*, pages 137–150. Springer-Verlag, May 1996.
15. A. Kate, G. M. Zaverucha, and I. Goldberg. Pairing-Based Onion Routing. In *7th Privacy Enhancing Technologies Symposium*, June 2007.
16. A. Khalili, J. Katz, and W. Arbaugh. Toward Secure Key Distribution in Truly Ad-Hoc Networks. In *IEEE Workshop on Security and Assurance in Ad-Hoc Networks 2003*, pages 342–346, 2003.
17. N. Koblitz and A. Menezes. Pairing-Based Cryptography at High Security Levels. In *Tenth IMA International Conference on Cryptography and Coding, Lecture Notes in Computer Science 3796*, pages 13–36. Springer-Verlag, December 2005.
18. B. Lynn. PBC Library – The Pairing-Based Cryptography Library. <http://crypto.stanford.edu/pbc/>. Accessed February 2007.
19. S. Mauw, J. Verschuren, and E. de Vink. A Formalization of Anonymity and Onion Routing. In *ESORICS 2004, Lecture Notes in Computer Science 3193*, pages 109–124. Springer-Verlag, September 2004.
20. A. Menezes, T. Okamoto, and S. Vanstone. Reducing Elliptic Curve Logarithms to Logarithms in a Finite Field. In *STOC '91: Proc. of the twenty-third annual ACM Symposium on Theory of Computing*, pages 80–89, 1991.
21. A. Menezes, P. Van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1st edition, 1997.
22. B. Möller. Provably Secure Public-Key Encryption for Length-Preserving Chaumian Mixes. In *CT-RSA 2003, Lecture Notes in Computer Science 2612*. Springer-Verlag, April 2003.
23. E. Okamoto and T. Okamoto. Cryptosystems Based on Elliptic Curve Pairing. In *Modeling Decisions for Artificial Intelligence—MDAI 2005, Lecture Notes in Computer Science 3558*, pages 13–23. Springer-Verlag, July 2005.

24. L. Øverlier and P. Syverson. Improving efficiency and simplicity of Tor circuit establishment and hidden services. In *Proceedings of the 7th Privacy Enhancing Technologies Symposium (these proceedings)*, 2007.
25. T. Pedersen. A Threshold Cryptosystem without a Trusted Party. In *Advances in Cryptology—Eurocrypt 1991, Lecture Notes in Computer Science 547*, pages 522–526. Springer-Verlag, 1991.
26. S. Rahman, A. Inomata, T. Okamoto, M. Mambo, and E. Okamoto. Anonymous Secure Communication in Wireless Mobile Ad-hoc Networks. In *First International Conference on Ubiquitous Convergence Technology (ICUCT2006)*, December 2006.
27. M. Reed, P. Syverson, and D. Goldschlag. Anonymous Connections and Onion Routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, May 1998.
28. M. Rennhard and B. Plattner. Introducing MorphMix: Peer-to-Peer based Anonymous Internet Usage with Collusion Detection. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2002)*, Washington, DC, USA, November 2002.
29. R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems based on pairing. In *Symposium on Cryptography and Information Security (SCIS 2000)*, 2000.
30. A. Seth and S. Keshav. Practical Security for Disconnected Nodes. In *IEEE ICNP Workshop on Secure Network Protocols, 2005 (NPSec)*, pages 31–36, 2005.
31. A. Shamir. How to Share a Secret. *Commun. ACM*, 22(11):612–613, 1979.
32. P. Syverson, G. Tsudik, M. Reed, and C. Landwehr. Towards an Analysis of Onion Routing Security. In *Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability, Lecture Notes in Computer Science 2009*, pages 96–114. Springer-Verlag, July 2000.
33. The Tor Project. Tor: anonymity online. <http://tor.eff.org/>. Accessed February 2007.
34. E. Verheul. Evidence that XTR Is More Secure than Supersingular Elliptic Curve Cryptosystems. In *Advances in Cryptology—Eurocrypt 2001, Lecture Notes in Computer Science 2045*, pages 195–210, 2001.

A Security and Anonymity Analysis

In this section, we discuss the security and anonymity of our one-way and two-way anonymous key agreement protocols in a BF-IBE setup. We prove the unconditional anonymity, session key secrecy and no impersonation properties for our protocols.

A.1 Unconditional Anonymity

We prove that it is impossible for an adversary \mathcal{A} to learn the identity of an anonymous participant in a protocol run. Note that \mathcal{A} can be the other participant in a protocol run, the PKG for the system or any third party. To prove unconditional anonymity, consider the following game between an adversary and a challenger.

Setup. The adversary \mathcal{A} publishes the system parameters: a cyclic additive group \mathbb{G} of prime order n and a hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$.

Challenge. \mathcal{A} chooses two identity strings ID_A and ID_B and sends them to the challenger. The challenger computes $Q_A = H(ID_A)$ and $Q_B = H(ID_B)$. He then uniformly at random chooses $r \in \mathbb{Z}_n^*$ and $b \in \{0, 1\}$,

1. if $b = 0$, computes a pseudonym $P = rQ_A$ or
2. if $b = 1$, computes a pseudonym $P = rQ_B$

and sends P to \mathcal{A} .

Guess. \mathcal{A} wins the game, if she can guess the correct value of b with probability significantly greater than $1/2$.

As \mathbb{G} is a cyclic prime order group, both Q_A and Q_B are generators of \mathbb{G} . For the uniform random element $r \in \mathbb{Z}_n^*$, the pseudonym P equal to rQ_A or rQ_B is also a uniform random element of \mathbb{G}^* . Therefore, an attacker cannot determine which of the two ways the challenger generated P and consequently cannot guess the value of b with probability greater than $1/2$ to win this game. The inability of the attacker to win this game for system parameters generated by her, even with unbounded computation power, proves our unconditional anonymity claim.

A.2 Session Key Secrecy

Here, we prove that it is infeasible for anyone other than the two participants or the PKG to determine a session key generated during a protocol run of the one-way or two-way anonymous key agreement, under the BDH assumption.

Consider the following game to prove key secrecy in the one-way anonymous case.

Setup. The challenger generates groups \mathbb{G} and \mathbb{G}_T of prime order n , a cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$, a symmetric bilinear pairing $e : \mathbb{G} \times \mathbb{G} \mapsto \mathbb{G}_T$ and a master secret $s \in \mathbb{Z}_n^*$.

Extraction Queries. The adversary \mathcal{A}_1 issues q extraction queries for identities ID_1, ID_2, \dots, ID_q to the challenger. The challenger queries H to compute the corresponding private keys $sH(ID_1), sH(ID_2), \dots, sH(ID_q)$ and sends them back to \mathcal{A}_1 .

Challenge. Once \mathcal{A}_1 informs the challenger that it has collected enough information, the challenger picks an element $P_A \in \mathbb{G}^*$ and sends it to \mathcal{A}_1 .

Guess. \mathcal{A}_1 outputs a binary string (an identity) ID_B and $K_{AB} \in \mathbb{G}_T$.

The attacker's advantage can be defined as

$$\text{Adv}(\mathcal{A}_1) = \text{Prob}(e(P_A, H(ID_B))^s = K_{AB})$$

We say \mathcal{A}_1 (t_1, ϵ_1)-wins the game, if it runs in time at most t_1 and has advantage ϵ_1 .

Now, consider the following game to prove key secrecy in the two-way anonymous case.

Setup. The challenger generates groups \mathbb{G} and \mathbb{G}_T of order n , a cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$, a symmetric bilinear pairing $e : \mathbb{G} \times \mathbb{G} \mapsto \mathbb{G}_T$ and a master secret $s \in \mathbb{Z}_n^*$.

Extraction Queries. The adversary \mathcal{A}_2 issues q extraction queries for identities $ID_1, ID_2, \dots, ID_q \in \mathbb{G}$ to the challenger. The challenger queries H to compute the corresponding private keys $sH(ID_1), sH(ID_2), \dots, sH(ID_q)$ and sends them back to \mathcal{A}_2 .

Challenge. Once \mathcal{A}_2 informs the challenger that it has collected enough information, the challenger picks two elements P_A and P_B in \mathbb{G}^* and sends them to \mathcal{A}_2 .

Guess. \mathcal{A}_2 outputs $K_{AB} \in \mathbb{G}_T$.

The attacker's advantage can be defined as

$$\text{Adv}(\mathcal{A}_2) = \text{Prob}(e(P_A, P_B)^s = K_{AB})$$

We say \mathcal{A}_2 (t_2, ϵ_2) -wins a game, if it runs in time at most t_2 and has advantage ϵ_2 .

Suppose that there is an adversary \mathcal{A}_1 who (t_1, ϵ_1) -wins the one-way anonymous game and an adversary \mathcal{A}_2 who (t_2, ϵ_2) -wins the two-way anonymous game. Note that these adversaries always exist, for appropriate choices of t_i and ϵ_i . We now show that an algorithm \mathcal{B} can make use of \mathcal{A}_1 or \mathcal{A}_2 to solve a random instance of the BDH problem.

Theorem 1. *Let the hash function H be modelled by a random oracle. Suppose there exist adversaries \mathcal{A}_1 and \mathcal{A}_2 such that the adversary \mathcal{A}_1 (t_1, ϵ_1) -wins the one-way anonymous protocol security game and the adversary \mathcal{A}_2 (t_2, ϵ_2) -wins the two-way anonymous protocol security game. Then there exists an algorithm \mathcal{B} which solves the BDH problem*

- using the adversary \mathcal{A}_1 with probability $\frac{\epsilon_1}{e(1+q)}$ in time $t_1 + wq + t_T + t_{inv}$ or
- using the adversary \mathcal{A}_2 with probability ϵ_2 in time $t_2 + wq$.

Here e is the base of natural logarithms, w is a small constant, q is an upper bound on the number of extraction queries performed by an adversary, t_T is the time required for exponentiation in \mathbb{G}_T and t_{inv} is the time required to invert an element of \mathbb{Z}_n^* .

Proof. Let $(P, aP, bP, cP) \in \mathbb{G}$ be a random and uniformly distributed instance of the BDH problem, which algorithm \mathcal{B} receives as input. To find the solution $e(P, P)^{abc}$, \mathcal{B} simulates the challenger for \mathcal{A}_1 or \mathcal{A}_2 . This means that \mathcal{B} must simulate the random oracle H and answer the private key extraction queries by \mathcal{A}_1 or \mathcal{A}_2 . As the steps for H -queries and extraction queries are same in both \mathcal{A}_1 or \mathcal{A}_2 , we denote both of them by \mathcal{A} .

H -queries. At any time, \mathcal{A} can query the random oracle H . To respond to these queries, \mathcal{B} maintains an initially empty list L of tuples $(X, Q, h, \beta) \in \{0, 1\}^* \times \mathbb{G}^* \times \mathbb{Z}_n^* \times \{0, 1\}$. When \mathcal{A} queries for the hash value of some bit-string X_i , algorithm \mathcal{B} responds as follows:

1. If L contains a tuple (X_i, Q_i, h_i, β_i) , \mathcal{B} responds by sending Q_i .

2. Otherwise, \mathcal{B} generates at random $\beta_i \in \{0, 1\}$, so that $\text{Prob}(\beta_i = 0) = \delta$, where δ depends on B 's choice for the attacker (\mathcal{A}_1 or \mathcal{A}_2) and will be determined below.
3. Algorithm \mathcal{B} picks a random $h_i \in \mathbb{Z}_n^*$. If $\beta_i = 0$, set $Q_i = h_i P$, else set $Q_i = h_i(bP)$. Note that either way, Q_i is uniformly random in \mathbb{G}^* and independent of \mathcal{A} 's current view.
4. Finally, algorithm \mathcal{B} adds the tuple (X_i, Q_i, h_i, β_i) to the list L and responds with Q_i .

Extraction queries. \mathcal{A} can ask for extraction queries for identity strings. For an input string ID_i for private key extraction, \mathcal{B} responds as follows:

1. Algorithm \mathcal{B} runs the above H -query algorithm for input $X_i = \text{ID}_i$ to obtain $(\text{ID}_i, Q_i, h_i, \beta_i)$.
2. If $\beta_i = 1$ then \mathcal{B} reports failure.
3. Otherwise, \mathcal{B} computes the private key $h_i(cP) = cQ_i$ and sends it to algorithm \mathcal{A} .

Challenge. After completing the extraction queries, B challenges \mathcal{A}_1 with $P_A = aP$ or \mathcal{A}_2 with $P_A = aP$ and $P_B = bP$.

Guess. \mathcal{A}_1 outputs $(\text{ID}_B, K_{AB}) \in \{0, 1\}^* \times \mathbb{G}_T$ or \mathcal{A}_2 outputs $K_{AB} \in \mathbb{G}_T$. In case of the adversary \mathcal{A}_2 , \mathcal{B} outputs $\sigma = K_{AB}$ as its guess for the solution to the BDH problem. For the adversary \mathcal{A}_1 , algorithm \mathcal{B} performs following steps:

1. \mathcal{B} obtains the tuple $(\text{ID}_B, Q_B, h_B, \beta_B)$ from the list L . Absence of the tuple $(\text{ID}_B, Q_B, h_B, \beta_B)$ in the list L indicates that \mathcal{A}_1 did not ask the random oracle for $H(\text{ID}_B)$. As the probability of the adversary's success in this case is negligible¹, we safely assume the presence of the tuple $(\text{ID}_B, Q_B, h_B, c_B)$.
2. If $\beta_B = 1$, \mathcal{B} outputs $\sigma = K_{AB}^{h_B^{-1}}$ as its guess for the BDH instance.
3. If $\beta_B = 0$, \mathcal{B} reports failure.

Suppose that \mathcal{B} does not report failure and outputs σ while using \mathcal{A}_1 . As $\beta_B = 1$, $H(\text{ID}_B) = h_B(bP)$ and with probability ϵ_1 , $\sigma = e(aP, h_B(bP))^{ch_B^{-1}} = e(P, P)^{abc}$. Therefore \mathcal{B} will guess correctly with probability ϵ_1 , when it does not abort. The probability that \mathcal{B} does not abort while extracting a single private key query is δ ; for q queries, the probability is δ^q . The probability that \mathcal{B} does not abort while guessing the BDH solution is $1 - \delta$. Therefore, the overall probability of non-abortion is $\delta^q(1 - \delta)$. Maximizing this probability, the optimal value can be obtained at $\delta = \frac{q}{1+q}$ and by choosing the value of δ optimally, the overall probability of non-abortion is $\frac{q^q}{(1+q)^{q+1}}$. Therefore, \mathcal{B} outputs the correct solution to the BDH instance with probability at least $\frac{\epsilon_1 q^q}{(1+q)^{q+1}} \geq \frac{1}{e(1+q)}$ as $(1 - \frac{1}{q+1})^q \geq 1/e$. The solution is computed in time $t_1 + wq + t_T + t_{inv}$, where t_1 is the time

¹ If \mathcal{A}_1 does not query the random oracle for $H(\text{ID}_B)$, the probability it can guess this value is negligible. As there is a bijection between x and $e(P_A, x)$ for a given P_A , the probability that \mathcal{A}_1 can output $K_{AB} = e(P_A, H(\text{ID}_B))$ is also negligible. Thus, \mathcal{A}_1 's advantage in this case is negligible.

required by \mathcal{A} , w is the time required to answer an extraction query (generate a random element r and compute the r -th multiple of cP), q is an upper bound on the number of such queries, and $t_T + t_{inv}$ is the time to invert an element of \mathbb{Z}_n^* and to compute an exponentiation of $K_{AB} \in \mathbb{G}_T$ in the guessing phase.

For adversary \mathcal{A}_2 , we simply set the value of δ to 1. Suppose that \mathcal{B} does not report failure and outputs σ while using \mathcal{A}_2 . With probability ϵ_2 , $\sigma = e(aP, bP)^c = e(P, P)^{abc}$, which is the correct solution to the BDH problem. The solution is computed in time $t_2 + wq$. \square

Note that it is possible to prove the security of the two-way anonymous key agreement protocol without random oracles, if we do not consider the query extraction phase. Assume that only one identity hash and private key pair (U, sU) is publicly available and each user uses the same pair to generate a pseudonym and corresponding private key. Given an adversary \mathcal{A} to (t, ϵ) -compute $K_{AB} = e(P_A, P_B)^s$ when challenged by P_A and P_B , a random instance (P, aP, bP, cP) of the BDH problem can be solved in time t with probability ϵ by publishing (P, cP) as the publicly available identity hash and private key and challenging \mathcal{A} with $P_A = aP$ and $P_B = bP$.

A.3 No Impersonation

We claim that it is infeasible for a malicious client of the PKG to impersonate another (non-anonymous) client in a protocol run. To successfully impersonate a non-anonymous participant ID_N in our one-way anonymous key agreement protocol, given a pseudonym and ID_N , an adversary needs to determine the corresponding session key. Now, we present an adversary game for non-anonymous participant impersonation, which is the same as the key secrecy game of the one-way anonymous key agreement.

Setup. The challenger generates groups \mathbb{G} and \mathbb{G}_T of prime order n , a cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$, a symmetric bilinear pairing $e : \mathbb{G} \times \mathbb{G} \mapsto \mathbb{G}_T$ and a master secret $s \in \mathbb{Z}_n^*$.

Extraction Queries. The adversary \mathcal{A}_1 issues q extraction queries for identities ID_1, ID_2, \dots, ID_q to the challenger. The challenger queries H to compute the corresponding private keys $sH(ID_1), sH(ID_2), \dots, sH(ID_q)$ and sends them back to \mathcal{A}_1 .

Challenge. Once \mathcal{A}_1 informs the challenger that it has collected enough information, the challenger picks an element $P_A \in \mathbb{G}^*$ and sends it to \mathcal{A}_1 .

Guess. \mathcal{A}_1 outputs a binary string (an identity) ID_B and $K_{AB} \in \mathbb{G}_T$.

The attacker's advantage can be defined as

$$\text{Adv}(\mathcal{A}_1) = \text{Prob}(e(P_A, H(ID_B))^s = K_{AB})$$

Consequently, the corresponding theorem and proof are same as those for key secrecy in the one-way anonymous key agreement protocol (see Section A.2).

In the case of persistent pseudonymity, we claim that it is not feasible for a malicious entity to communicate using a different entity's pseudonym. Here, the

malicious entity needs to find the shared secret key for a persistent pseudonym generated and used by some other anonymous entity and an arbitrary identity or pseudonym for which it does not know the private key. In the one-way anonymous communication protocol, the corresponding adversary game remains the same as impersonation of the non-anonymous entity, and in the two-way anonymous case, the game is the same as the one used to prove key secrecy. A combined adversary game can be defined as follows.

Setup. The challenger generates groups \mathbb{G} and \mathbb{G}_T of order n , a cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}^*$, a symmetric bilinear pairing $e : \mathbb{G} \times \mathbb{G} \mapsto \mathbb{G}_T$ and a master secret $s \in \mathbb{Z}_n^*$.

Extraction Queries. The adversary \mathcal{A} issues q extraction queries for identities $\text{ID}_1, \text{ID}_2, \dots, \text{ID}_q \in \mathbb{G}$ to the challenger. The challenger queries H to compute the corresponding private keys $sH(\text{ID}_1), sH(\text{ID}_2), \dots, sH(\text{ID}_q)$ and sends them back to \mathcal{A} .

Challenge. Once \mathcal{A} informs the challenger that it has collected enough information, the challenger picks a pseudonym $P_A \in \mathbb{G}^*$ and an identity string ID_B or a pseudonym $P_B \in \mathbb{G}^*$ and sends them to \mathcal{A} .

Guess. \mathcal{A} outputs $K_{AB} \in \mathbb{G}_T$.

The attacker's advantage can be defined as

$$\text{Adv}(\mathcal{A}) = \text{Prob}(e(P_A, P_B)^s = K_{AB} \quad \text{or} \quad e(P_A, H(\text{ID}_B))^s = K_{AB})$$

It is easy to observe that a theorem and proof for this game are same as those used to prove key secrecy in Section A.2.