Identity-Based Steganography and Its Applications to Censorship Resistance

Tim Ruffing[1], Jonas Schneider[1], and Aniket Kate[2]

[1] Saarland University, Saarbrücken, Germany
{tim.ruffing,s9joscne}@stud.uni-saarland.de
[2] MMCI, Saarland University, Saarbrücken, Germany
aniket@mmci.uni-saarland.de

Abstract. Public-key steganography has been proposed for several censorship-resistance systems. However, distribution of the employed public keys presents an availability, scalability, and security challenge in many of those. To mitigate this problem, we introduce the notion of identity-based steganography. In particular, we define identity-based steganographic tagging (IBST), which allows a sender to produce a steganographic tag for a recipient's identity such that the tag can only be recognized by the intended recipient using her (identity-based) private key. We instantiate our definition by presenting an efficient IBST scheme, provably secure under the bilinear decisional Diffie-Hellman assumption. We find IBST to be particularly useful in censorship-resistance systems when the censors are able to impede distribution of cryptographic keys or break forward security by compromising system agents. As two representative applications of IBST, we present an efficient and dynamic solution for the Collage covert communication system and a scalable approach to make the Telex censorship-resistance solution forward secure.

Keywords: Steganography, Censorship Resistance, Key Distribution, Identity-Based Cryptography, Collage, Telex

1 Introduction

Censorship resistance systems [1, 2, 3, 4] typically employ steganography to establish their covert channels. For both symmetric-key and public-key steganography, the distribution of a steganographic key (stego-key) constitutes an important bootstrapping problem in these systems: A simple attack for the censor is to block the key distribution. The problem is further aggravated for systems aiming for forward security, which is typically achieved by frequently changing to new keys and deleting old keys.

In an identity-based encryption (IBE) scheme, as proposed in [5] and brought to practicality in [6], messages are encrypted for an identity of an intended recipient, e.g., her email or IP address, or even the domain name of her webserver. The recipient decrypts a ciphertext using a private key associated to her identity, which she has to obtain from a trusted entity, a private key generator (PKG),

holding a master key. Identity-based cryptography (IBC) mitigates the key distribution problem as only a single master public key associated with the master key needs to be distributed to clients.

Contribution. In this work, we observe that the efficient key distribution of IBC is useful for censorship-resistance systems, especially when private keys have to be distributed only outside the censored area. Therefore, we combine steganography and IBC to introduce the concept of *identity-based steganography*. Using identity-based steganography in censorship-resistance systems, essentially no key distribution in the censored area is required with the exception of a single master public key, which can be bundled with the software. In general, identity-based steganography is helpful whenever a censor is able to block mass key distribution and whenever a client publishes steganographic messages optimistically, i.e., hoping that someone outside the censored area will react to them.

In particular, we define the notion of an identity-based steganographic tagging (IBST) scheme. It allows a sender A, given only a single public key and the identity of a recipient B to produce a steganographic tag that only B can recognize, and without possession of the corresponding private key the tag is indistinguishable from a random bitstring. We also give a formal security definition that captures the undetectability of tags. Building upon the IBE scheme by Boneh and Franklin [6] and steganographic techniques from Telex [7], we instantiate our definition with a concrete IBST construction and prove its security in the random oracle model under the bilinear decisional Diffie-Hellman (BDDH) assumption.

Finally, we demonstrate the utility of our construction to censorship-resistance systems: First, we exemplify the usefulness of identity-based steganography to Collage [8]. Here, besides obtaining eventual forward security, we achieve better security and flexibility without hindering the bootstrapping process. Second, we use our construction to provide eventual forward security for arbitrarily short time intervals also in Telex [7], which improves upon the original design in terms of scalability.

2 Background

Public-Key Steganography. Anderson [9] is the first to consider steganography in the public-key world, however only with informal security arguments. Security definitions are introduced by Ahn and Hopper [10] along with the first provably secure scheme for public-key steganography. Backes and Cachin [11] extend public-key steganography to adaptive chosen-covertext attacks; their constructions however only fulfill an intermediate security notion. Following this work, the construction of Hopper [12] achieves the full security notion defined in [11]. A multi-recipient variant of public-key steganography is proposed and instantiated by Fazio, Nicolosi, and Perera [13].

Identity-based Encryption. Boneh and Franklin [6] give the first IBE scheme to reach practical efficiency. Our scheme is built upon BasicIdent, an intermediate

construction in their work. A common application for identity-based encryption is the implementation of forward security, by taking the current time interval as an identity. This technique is also relevant for the applications we propose.

Censorship Resistance. Steganography is often used for covert communication and censorship-resistance technology [1, 2, 3, 4]. The Telex censorship-resistance system e.g., proposed by Wustrow et al. [7], relies on public-key steganography to enable users to request a proxy service. The user establishes a specially crafted TLS connection to an innocent-looking website and hopes that a router along the path, belonging to a trusted Internet Service Provider (ISP) in possession of a private key, can detect the connection as a service request. Several other censorship-resistance systems follow a similarly opportunistic paradigm. Among them are MIAB (Message in a Bottle) [14] and Collage [8]. They allow users to post steganographic messages on blogs and websites with user-generated content, respectively, such that the intended recipients can fetch them from there. One of our example applications improves the flexibility of Collage.

The prototype of Flash Proxy, the censorship-resistance system introduced by Fifield et al. [15], makes use of IBE to achieve a primitive similar to identity-based steganographic tagging defined in this work. They do not define the primitive or prove its properties. As opposed to our work, their primitive includes the transmission of messages and thus requires a full IBE encryption scheme. Our construction avoids that in order to allow very short tags, e.g., to be able to use them with covertexts such as random nonces.

3 Identity-Based Steganographic Tagging

Given a public key of a recipient, steganographic tagging allows to produce *tags* that look like random covertexts for everybody who is not in possession of the corresponding private key. Assume a censored user publishes a tag in the hope that the intended recipient will find it and react to it, typically by establishing a connection to the user. For instance, in the case of Collage [8], an uncensored party retrieves a picture from Flickr in which the the tag and possibly a message are embedded. Additionally, the tag can be used to establish a shared secret, e.g., a symmetric stego-key, between sender and recipient.

In this section, we extend steganographic tagging to the identity-based setting and give a concrete instantiation.

3.1 Definition

An identity-based steganographic tagging scheme (IBST) provides algorithms to produce a tag and a secret using a master public key and an identity. The intended recipient can detect a tag and extract the corresponding shared secret using the private key corresponding to her identity.

Additionally, a random innocent message (a *covertext*) should be a tag only with negligible probability. We denote this property by *rareness*. It ensures that normal covertext messages are not recognized as tags.

The following definition captures steganographic tagging in the identity-based setting. Regarding the notion of steganography, we follow the approach of [11].

Definition 1 (Identity-Based Steganographic Tagging). *Let $\mathcal{C}_\lambda$ be a distribution on covertexts. An* identity-based steganographic tagging scheme (IBST) *is a tuple of algorithms $\mathcal{T} = (\mathsf{Setup}, \mathsf{PrivateKey}, \mathsf{Tag}, \mathsf{Detect})$, where*

$\mathsf{Setup}(1^\lambda)$ *takes as input the security parameter 1^λ and returns a key pair (mpk, msk) with a master public key mpk and a master private key msk. Furthermore, it publishes parameters params that are assumed to be implicitly known to all following algorithms.*

$\mathsf{Tag}(mpk, id)$ *takes as input the master public key mpk and an $id \in \mathcal{ID}$ and returns a pair (τ, s) of the so called tag $\tau \in \mathcal{C}_\lambda$ and a corresponding shared secret s.*

$\mathsf{PrivateKey}(msk, id)$ *takes as input the master private key msk together with an $id \in \mathcal{ID}$. It returns the private key sk_{id} for id.*

$\mathsf{Detect}(sk_{id}, \tau)$ *takes as input the private key sk_{id} belonging to an $id \in \mathcal{ID}$ and a candidate tag $\tau \in \mathcal{C}_\lambda$. If τ is a tag, it returns the shared secret s corresponding to the tag. Otherwise, it returns $\perp$.*

An IBST scheme must satisfy the following properties:

Correctness: *Let $id \in \mathcal{ID}$ and a key pair (mpk, msk) output by $\mathsf{Setup}(1^\lambda)$ be given. Furthermore, let $sk_{id} \leftarrow \mathsf{PrivateKey}(msk, id)$. For all pairs (τ, s) output by $\mathsf{Tag}(mpk, id)$, it should hold that*

$$\mathsf{Detect}(sk_{id}, \tau) = s \ .$$

Rareness: *For all $id \in \mathcal{ID}$ and (mpk, msk) output by $\mathsf{Setup}(1^\lambda)$, the density $\gamma(\lambda)$ defined as follows must be a negligible function:*

$$\gamma(\lambda) := \Pr[\mathsf{Detect}(sk_{id}, c) \neq \perp; c \leftarrow \mathcal{C}_\lambda, sk_{id} \leftarrow \mathsf{PrivateKey}(msk, id)] \ .$$

The security notion of IBST captures two essential security properties: First, tags cannot be distinguished from random covertexts without possession of the private key. In other words, they are not detectable for the censor. Second, no adversary can associate shared secrets with the corresponding tags. This ensures the confidentiality of the exchanged secret. Our definition allows an adversary to query private keys, adaptively and for arbitrary identities, as long as it is not the identity on which she chooses to be challenged, analogously to the IND-ID-CPA game for identity-based encryption.

Definition 2 (Indistinguishability Game). *The game $\mathsf{ID\text{-}TAG}^b_{\mathcal{T},\mathcal{A}}(\lambda)$ that is parameterized over a bit $b \in \{0, 1\}$, a polynomially bounded adversary $\mathcal{A}$ and an IBST $\mathcal{T} = (\mathsf{Setup}, \mathsf{PrivateKey}, \mathsf{Tag}, \mathsf{Detect})$ is defined as follows:*

Setup: *The challenger takes the security parameter λ and runs $\mathsf{Setup}(1^\lambda)$ to compute (mpk, msk). It gives mpk to the adversary and keeps msk.*

Phase 1: *The adversary issues (adaptively) a number of extraction queries* $(id_1, \ldots, id_{q_1})$*, which the challenger answers with the corresponding private keys* $(sk_{id_1}, \ldots, sk_{id_{q_1}})$ *after generating them through* $\mathsf{PrivateKey}(msk, \cdot)$.

Challenge: *At some point, the adversary declares the first phase to be over and outputs a challenge identity* id^* *with* $id^* \neq id_i$ *for all* $1 \leq i \leq q_1$*. The challenger now sets* $(\tau_0, s) \leftarrow \mathsf{Tag}(mpk, id^*)$ *and picks a random covertext* $\tau_1 \leftarrow \mathcal{C}_\lambda$*. It sends* (τ_b, s) *to the adversary.*

Phase 2: *The adversary issues additional queries* $(id_{q_1+1}, \ldots, id_{q_2})$ *where* $id_i \neq id^*$ *for all* $q_1 < i \leq q_2$*. The challenger responds as in phase 1.*

Guess: *Finally, the adversary outputs a guess bit* $b' \in \{0, 1\}$*, which is also the output of the game.*

The advantage of $\mathcal{A}$ *is defined as*

$$\mathbf{Adv}^{\mathsf{ID\text{-}TAG}}_{\mathcal{T},\mathcal{A}}(\lambda) := \left|\Pr[\mathsf{ID\text{-}TAG}^0_{\mathcal{T},\mathcal{A}}(\lambda) = 0] - \mathsf{ID\text{-}TAG}^1_{\mathcal{T},\mathcal{A}}(\lambda) = 0]\right| \ .$$

We say that an identity-based steganographic tagging scheme has indistinguishable tags and shared secrets *if for every polynomially bounded adversary* $\mathcal{A}$*, we have that* $\mathbf{Adv}^{\mathsf{ID\text{-}TAG}}_{\mathcal{T},\mathcal{A}}(\lambda)$ *is negligible.*

3.2 Construction

We now instantiate our definition of an IBST scheme with a construction based on the BasicIdent IBE scheme as defined by Boneh and Franklin [6]. Indistinguishability is achieved similarly as in [7]. Since the tag contains a random group element on an elliptic curve, it is crucial to encode this element such that its representation is indistinguishable from a bitstring drawn uniformly at random. Given the x- and y-coordinate of a group element, it can be checked if the element is on the curve. Thus only the x-coordinate is part of the tag. As the group order is selected close to a power of two, almost every bitstring represents an x-coordinate that is either on the curve or on its quadratic twist.[3] If elements are randomly chosen from either of those groups, their representations are pseudorandom [7].

We construct our pairing-based stego-tagging scheme $\mathcal{T}$ as follows. In our case, the covertext distribution is just the uniform distribution on bitstrings of a fixed length.

$\mathsf{Setup}(1^\lambda)$: Similar to the BasicIdent IBE, choose groups $\mathbb{G}_0$, $\mathbb{G}_1$, $\hat{\mathbb{G}}_0$, $\hat{\mathbb{G}}_1$ of prime order $q(\lambda)$ and two bilinear maps $e_b : \mathbb{G}_b \times \mathbb{G}_b \to \hat{\mathbb{G}}_b$ for $b \in \{0, 1\}$ such that $\mathbb{G}_1 = \mathbb{G}_0^{-1}$ is a quadratic twist of $\mathbb{G}_0$. We require the BDDH assumption [6] to hold with respect to e_0 and e_1. Furthermore, we assume that $2^{\ell_q} - q < \sqrt{q}$, where ℓ_q is the bit-length of q. Let $\phi : \{0, 1\}^{\ell_q} \times \{0, 1\}^{\ell_q} \to \{0, 1\}^{\ell_q}$ be the point multiplication on representations of x-coordinates. Choose random generators $P_0 \leftarrow \mathbb{G}_0, P_1 \leftarrow \mathbb{G}_1$ and random $s_0, s_1 \leftarrow \mathbb{Z}_q^*$. Let $\ell_h := \lambda - \ell_q$ and choose cryptographic hash functions $G_b^{\text{tag}} : \hat{\mathbb{G}}_b \to \{0, 1\}^{\ell_h}, G_b^{\text{key}} : \hat{\mathbb{G}}_b \to \{0, 1\}^\lambda$, and $H_b : \{0, 1\}^* \to \mathbb{G}_b^*$ for $b \in \{0, 1\}$.

[3] The idea to use curves related by twisting to achieve pseudorandomness came up first in [16].

Publish $params := (q, \mathbb{G}_0, \mathbb{G}_1, \hat{\mathbb{G}}_0, \hat{\mathbb{G}}_1, e_0, e_1, G_0^{\text{tag}}, G_1^{\text{tag}}, G_0^{\text{key}}, G_1^{\text{key}}, H_0, H_1)$ and the master public key $mpk := (P_0, s_0P_0, P_1, s_1P_1)$. Keep the master private key $msk := (s_0, s_1)$.

$\mathsf{PrivateKey}(msk, id)$: Compute $Q_0^{id} := H_0(id)$ and $Q_1^{id} := H_1(id)$ and return $sk_{id} := (s_0Q_0^{id}, s_1Q_1^{id})$.

$\mathsf{Tag}(mpk, id)$: Parse mpk as $(P_0, P_0^{pub}, P_1, P_1^{pub})$. Pick a random bit $b \leftarrow \{0,1\}$ and $r \leftarrow \mathbb{Z}_q^*$. Let $Q_b^{id} := H_b(id)$ and compute $p := e_b(rP_b^{pub}, Q_b^{id})$.
Set $\tau := \phi(r, P_b)||G_b^{\text{tag}}(p)$ and $s := G_b^{\text{key}}(p)$. Return (s, τ).

$\mathsf{Detect}(\tau, sk_{id})$: Parse sk_{id} as (K_0, K_1) and τ as $x||h$ with $|x| = \ell_q$ and $|h| = \ell_h$. Let X be the point with the x-coordinate represented by x; X lies either on $\mathbb{G}_0$ or on the twist $\mathbb{G}_1$. Say we have $X \in \mathbb{G}_b$. Check whether

$$h = G_b^{\text{tag}}(e_b(X, K_b))$$

holds. If the checks succeeds, output the shared secret $G_b^{\text{key}}(e_b(X, K_b))$, otherwise output $\perp$.

Note that it is also possible to compute some pairings efficiently in case only the y-coordinate of one argument is given, without explicitly computing the missing y-coordinate [17].

In Appendix A, we prove correctness and rareness as well as indistinguishability of tags for our construction. The latter can be reduced to the infeasibility of the bilinear decisional Diffie-Hellman problem [6] and is formalized by the following theorem, which holds in the random oracle model.

Theorem 1 (Indistinguishability of Our Construction). *Suppose there is an* $\mathsf{ID\text{-}TAG}$ *adversary* $\mathcal{A}_{\mathcal{T}}$ *that has advantage* $\mathbf{Adv}_{\mathcal{T},\mathcal{A}}^{\mathsf{ID\text{-}TAG}} = \epsilon(\lambda)$ *against the IBST* $\mathcal{T}$ *as defined in Section 3.2. Then there is an adversary* $\mathcal{B}_{BDDH}$ *that has advantage at least* $\epsilon'(\lambda) \approx \epsilon(\lambda)/e(1 + q_E)$ *against the BDDH assumption. Its running time is* $O(time(\mathcal{A}_{\mathcal{T}}))$.

The proof of this theorem is given in Appendix A.2.

3.3 Identity-Based Steganographic Encryption

We observe that IBST can also be used to construct identity-based steganographic encryption in a hybrid manner. For simplicity, assume that the covertext distribution is the uniform distribution on bitstrings of a fixed length, like for our construction in Section 3.2, i.e., we aim for pseudorandom ciphertexts. To obtain an IBE scheme with this property, we can use a symmetric stego-system to encrypt a message and concatenate a tag (produced by IBST) to encapsulate the symmetric stego-key. Additionally, we can embed the pseudorandom ciphertext into covertext messages to allow for other covertext distributions, as done for instance in [9] and [11] for public-key steganography.

Conversely, given IBE with pseudorandom ciphertexts, IBST can be constructed in a black-box way: Tagging is implemented by encrypting $0^\lambda||s$, where

0^λ is used to ensure rareness of the tags and s is the shared secret, as noted by Fifield et al. [15, p. 17].[4] While outside the scope of this work, we believe that identity-based steganographic encryption is applicable to censorship-resistance systems whenever a message shall be transported (and not only a shared key).

3.4 Distributed PKG

A common concern with identity-based primitives is that the PKG introduces a key escrow and a single point of failure problem: If the PKG becomes unavailable for some reason, it cannot provide private keys anymore and the whole system does not work, until the PKG is available again. However, there are solutions readily available to distribute PKG over several entities [18] for the employed Boneh-Franklin IBE setup, which can mitigate both problems.

4 Applications

Our IBST construction is useful for covert communication in any scenario where a well-formed recipient identity is known to the sender. It becomes particularly useful in scenarios where one public (or symmetric) stego-key cannot be used by all recipients for security reasons and the recipient's public key cannot easily be communicated to the sender. Thus, IBST can be applied in many censorship resistance and covert communication scenarios. In this section, we discuss its utility to two representative censorship-resistance systems: Collage [8] and Telex [7].

4.1 Application to Collage

Collage Overview. Collage [8] is a covert messaging system. It offers covert communication through postings on websites that allow user-generated content (e.g., Flickr, Twitter, or blogs). Collage assumes that the sender and the recipient of a message share a database of *tasks*, which helps the sender and the recipient to agree on a *rendezvous point*, i.e., a location where a stegotext should be stored. An example for a typical sender task is to embed a message steganographically into an image, and upload the image on Flickr under a keyword "flower". The corresponding recipient task is downloading images with this keyword and trying to stego-decrypt each of those.

Collage with IBST. In Collage, identity-based steganography can be used if a client inside the censored area would like to initialize a connection to a recipient outside this area, say a web proxy. The censored client selects a random task from its database and uses this task also as an identity string for IBST; e.g., in a scenario where the task refers to website Flickr and keyword flower, the client would not only use the task to determine the location (Flickr) where

[4] Additionally, we require the natural property that a random covertext decrypts to $0^\lambda||x$ for some x only with negligible probability.

the data should be stored but also use {`flickr`, `flower`} as identity string for the tagging algorithm. As the produced tag is pseudorandom, it can be steganographically embedded into an image. The recipient proxy will monitor Flickr for a few keywords assigned to it, and try to detect the tags using the private keys for the corresponding identities. A query from the sender can be added using identity-based steganographic encryption as described in Section 3.3.

The most important advantage of using identity-based steganography in Collage is that the task database can be made dynamic without sacrificing efficiency or trading off security properties. Burnett, Feamster, and Vempala [8] propose a list of most popular tags on Flickr as a means to establishing a task database without any direct communication between sender and recipient. This lack of communication in their system however introduces a key distribution problem with both public-key and symmetric-key steganography. It can be solved efficiently by using identity-based steganography because only the master public key has to be conveyed to the senders.

Since the number of identities is not fixed a priori in identity-based steganography, the IBST-based Collage solution scales well with an increasing number of proxies and tasks. Moreover, as the proxy knows only the private keys for the keywords it is responsible for, the impact of a compromise of one proxy is limited by design. We stress that distribution of the private keys is not a problem because both the proxies and the PKG are placed outside the censored area.

Forward Security. The security of many protocols relies on the secrecy of long-term keys. Since in practice, these keys can get compromised, it is a natural desire to ensure that the security of finished protocol sessions cannot be undermined. A cryptographic protocol is forward secure if an attacker that manages to compromise a protocol party is not able to obtain a session secret used in the *previous* protocol sessions [19]. Here, "previous" can either refer to any completed session (immediate forward security), or to a session that was finished before the key has been rotated for the last time (eventual forward security).

Consider again the example that a censored client wishes to establish a connection to an uncensored proxy, and assume we use IBST as described above. In identity-based cryptography, (eventual) forward security can be achieved by dividing time into intervals and concatenating the current interval to the identity string. Since we can apply this method directly to the IBST-based Collage in our example, we obtain forward security with potentially short time intervals. The PKG hands out private keys (one for each rendezvous point) to the proxies for a time period t, shortly before t begins. Once t is over, every proxy securely erases the corresponding private keys such that an attacker compromising this proxy cannot obtain the private keys for past periods $t' < t$. Notice that our solution provides forward security without introducing any additional communication between client and proxy.

Even if the task database does not entirely depend on public information and some agreement on tasks has to take place, the efficiency gain in terms of communication for key distribution is considerable; e.g., for T forward security time intervals and (say) the k most popular hash-tags on Flickr, kT public keys

have to be distributed with standard public key steganography instead of only one public key required in our identity-based solution.

4.2 Application to Telex

Telex Overview. In the Telex system [7], a censored user equipped with a Telex client initiates a TLS connection to some unsuspicious website and embeds a steganographic tag in the nonce used for TLS key agreement. The nonce appears random to the censor; however, at an ISP outside the censored area the tag is detected by a router (Telex station), which possesses the Telex private key. Upon detection, the Telex client and station collaborate to allow the Telex station to act as a benevolent man-in-the-middle in the TLS connection between the client and the unsuspicious website, and establish a new connection with a censored website, which the user wishes to visit. The new connection shares attributes with the previous one and lets the censor believe that the user is still connected to the unsuspicious website. Under the decisional Diffie-Hellman (DDH) assumption [20], the censor is not able to distiguish a tagged TLS connection from a normal one.

Forward Security in Telex. In Telex, the client is unsure about the identity of the involved Telex station, and all Telex stations share the same private key for efficiency and availability reasons. In this setting, however, any compromised station would jeopardize the security of all current and previous communication, and thus forward security is a well desired property in Telex. The Telex design addresses this by bundling a rather large number of public keys with the software.

We do not find this solution to be perfect due the following reason: Obtaining the (new) key material can become a problem for a client who acquires a software version for which the public keys have been expired due to a compromise of the Telex server. For such a client, it would be ideal if the key material is small (a few hundred bytes) such that it can be provided covertly and in an asynchronous manner, e.g., through an online image.

As we argue in Appendix B, immediate forward security is impossible for systems like Telex. However, using our IBST scheme, we are able to obtain eventual forward security with short time intervals while avoiding the aforementioned key management problem. For this purpose, we replace the stego-tagging scheme used in Telex by our IBST scheme and use time periods as identities, similar as proposed for Collage.

With this solution, the client can use an old software version as long as the IBST master key remains the same. Even when the master key is rotated, only the latest master public key (a few hundred bytes) has to be conveyed to the client; this can fit in almost all covert channels.

5 Conclusions and Future Work

In this paper, we have developed the notion of IBST, defined its security, and instantiated an IBST scheme with applicability to censorship-resistance systems

such as Collage and Telex. We have observed that IBST is particular useful if a client inside a censored area wishes to establish a connection to a recipient outside this area. Additionally, our construction can often provide forward security in censorship-resistance systems.

It would be interesting to extend the IBST security definition to a notion similar to adaptive chosen-covertext attacks [11]. Moreover, we leave a proper selection of parameters for a scheme that could be used in Telex for future work.

Acknowledgment. This work is supported by the German Universities Excellence Initiative.

References

[1] A. Houmansadr, G. T. Nguyen, M. Caesar, and N. Borisov. "Cirripede: circumvention infrastructure using router redirection with plausible deniability." In: *Proceedings of the 18th ACM conference on Computer and communications security*. CCS'11. ACM, 2011, pp. 187–200.

[2] A. Baliga, J. Kilian, and L. Iftode. "A web based covert file system." In: *Proceedings of the 11th USENIX workshop on Hot topics in operating systems*. HOTOS'07. USENIX, 2007, 12:1–12:6.

[3] Z. Weinberg, J. Wang, V. Yegneswaran, L. Briesemeister, S. Cheung, F. Wang, and D. Boneh. "StegoTorus: a camouflage proxy for the Tor anonymity system." In: *Proceedings of the 2012 ACM conference on Computer and communications security*. CCS'12. ACM, 2012, pp. 109–120.

[4] H. Mohajeri Moghaddam, B. Li, M. Derakhshani, and I. Goldberg. "SkypeMorph: protocol obfuscation for Tor bridges." In: *Proceedings of the 2012 ACM conference on Computer and communications security*. CCS'12. New York, NY, USA: ACM, 2012, pp. 97–108.

[5] A. Shamir. "Identity-based cryptosystems and signature schemes." In: *Advances in Cryptology – CRYPTO '84*. New York, NY, USA: Springer, 1985, pp. 47–53.

[6] D. Boneh and M. Franklin. "Identity-based encryption from the Weil pairing." In: *SIAM Journal on Computing* 32.3 (Mar. 2003), pp. 586–615.

[7] E. Wustrow, S. Wolchok, I. Goldberg, and J. A. Halderman. "Telex: anticensorship in the network infrastructure." In: *Proceedings of the 20th USENIX conference on Security*. Berkeley, CA, USA: USENIX, 2011.

[8] S. Burnett, N. Feamster, and S. Vempala. "Chipping away at censorship firewalls with user-generated content." In: *Proceedings of the 19th USENIX conference on Security*. USENIX, 2010, p. 29.

[9] R. Anderson. "Stretching the limits of steganography." In: *Information Hiding*. Springer, 1996, pp. 39–48.

[10] L. Ahn and N. J. Hopper. "Public-key steganography." In: *Advances in Cryptology - EUROCRYPT 2004*. Vol. 3027. Lecture Notes in Computer Science. Springer, 2004, pp. 323–341.

[11] M. Backes and C. Cachin. "Public-key steganography with active attacks." In: *Proceedings of the 2nd Theory of Cryptography Conference*. Vol. 3378. TCC'05. Springer, 2005, pp. 210–226.
[12] N. Hopper. "On steganographic chosen covertext security." In: *Proceedings of the 32nd international conference on Automata, Languages and Programming*. ICALP'05. Springer, 2005, pp. 311–323.
[13] N. Fazio, A. R. Nicolosi, and I. M. Perera. *Broadcast steganography*. 2013. IACR Cryptology ePrint Archive: 2013/087.
[14] L. Invernizzi, C. Kruegel, and G. Vigna. "Message in a bottle: sailing past censorship." In: Hot Topics in Privacy Enhancing Technologies. 2012.
[15] D. Fifield, N. Hardison, J. Ellithorpe, E. Stark, D. Boneh, R. Dingledine, and P. Porras. "Evading censorship with browser-based proxies." In: *Proceedings of the 12th international conference on Privacy Enhancing Technologies*. PETS'12. Springer, 2012, pp. 239–258.
[16] B. S. Kaliski Jr. "A pseudo-random bit generator based on elliptic logarithms." In: *Advances in Cryptology – CRYPTO' 86*. Vol. 263. Lecture Notes in Computer Science. Springer, 1987, pp. 84–103.
[17] S. D. Galbraith and X. Lin. "Computing pairings using x-coordinates only." In: *Designs, Codes and Cryptography* 50.3 (Jan. 2009), pp. 305–324.
[18] A. Kate and I. Goldberg. "Distributed private-key generators for identity-based cryptography." In: *Security and Cryptography for Networks*. Springer, 2010, pp. 436–453.
[19] W. Diffie, P. C. Oorschot, and M. J. Wiener. "Authentication and authenticated key exchanges." In: *Designs, Codes and Cryptography* 2.2 (1992), pp. 107–125.
[20] A. Joux and K. Nguyen. "Separating decision Diffie-Hellman from computational Diffie-Hellman in cryptographic groups." In: *Journal of Cryptology* 16.4 (2003), pp. 239–247.
[21] D. Galindo. "Boneh-Franklin identity based encryption revisited." In: *Proceedings of the 32nd international conference on Automata, Languages and Programming*. ICALP'05. Springer, 2005, pp. 791–802.

A Correctness, Rareness and Indistinguishability Proofs

We have to show correctness, rareness (Definition 1), and indistinguishability (Definition 2) of our construction.

A.1 Correctness

Let $id \in \mathcal{ID}$ and a key pair $(mpk, msk) = ((P_0, s_0P_0, P_1, s_1P_1), (s_0, s_1))$ output by $\mathsf{Setup}(1^\lambda)$ be given. Further, let $sk_{id} = (s_0Q_0^{id}, s_1Q_1^{id}) \leftarrow \mathsf{PrivateKey}(msk, id)$. Let (τ, s) be output by $\mathsf{Tag}(mpk, id)$. For some $b \in \{0,1\}$ and $r \in \mathbb{Z}_p^*$,

$$(\tau, s) = (\phi(r, P_b)||G_b^{\text{tag}}(p), G_b^{\text{key}}(p)), \quad \text{where} \quad p = e_b(rP_b^{pub}, Q_b^{id}) \ .$$

It follows that

$$\begin{aligned} G_b^{\text{tag}}(e_b(rP_b, s_bQ_b^{id})) &= G_b^{\text{tag}}(e_b(P_b, Q_b^{id})^{rs_b}) \\ &= G_b^{\text{tag}}(e_b(rs_bP_b, Q_b^{id})) \\ &= G_b^{\text{tag}}(e_b(rP_b^{pub}, Q_b^{id})) \\ &= G_b^{\text{tag}}(p) \ . \end{aligned}$$

and thus

$$\mathsf{Detect}(sk_{id}, \tau) = G_b^{\text{key}}(e_b(rP_b, s_bQ_b^{id})) = G_b^{\text{key}}(p) = s \ . \qquad \square$$

A.2 Rareness

For rareness, let again $id \in \mathcal{ID}$ and (mpk, msk) output by $\mathsf{Setup}(1^\lambda)$ be given. The density $\gamma(\lambda)$ behaves as follows:

$$\begin{aligned} \gamma(\lambda) &= \Pr[\mathsf{Detect}((K_0, K_1), c) \neq \perp; x||y \leftarrow \mathcal{C}_\lambda; (K_0, K_1) \leftarrow \mathsf{PrivateKey}(msk, id)] \\ &= \Pr[y = G_0^{\text{tag}}(e_0(X, K_0)) \vee y = G_1^{\text{tag}}(e_1(X, K_1)) \\ &\qquad c = x||y \leftarrow \mathcal{C}_\lambda; (K_0, K_1) \leftarrow \mathsf{PrivateKey}(msk, id)] \\ &\overset{(1)}{\leq} 2\Pr[y = G_0^{\text{tag}}(e_0(X, K_0)); \\ &\qquad c = x||y \leftarrow \mathcal{C}_\lambda; (K_0, K_1) \leftarrow \mathsf{PrivateKey}(msk, id)] \\ &\overset{(2)}{=} 2\Pr[y = G_0^{\text{tag}}(r); y \leftarrow \{0,1\}^{\ell_h}, r \leftarrow \hat{\mathbb{G}}_0^*] \\ &\overset{(3)}{=} 2^{-\ell_h+1} \approx 0 \ , \end{aligned}$$

where

(1) is a union bound for the events $y = G_0^{\text{tag}}(e_0(X, K_0))$ and $y = G_1^{\text{tag}}(e_1(X, K_1))$, which have the same probability,

(2) follows from the fact that $e_0(X, K_0) = e_0(rP_0, s_0Q_0^{id})$ for some $r \in \mathbb{Z}_p^*$ is a random element of $\hat{\mathbb{G}}_0^*$ and

(3) follows from modeling G_0^{tag} and G_1^{tag} as random oracles. $\square$

A.3 Indistinguishability

The following lemma captures the security of our construction. The proof relies on a reduction to the semantic security (IND-ID-CPA) of the BasicIdent IBE scheme due to Boneh and Franklin [6].

Lemma 1. *Suppose there is an* ID-TAG *adversary* $\mathcal{A}_\mathcal{T}$ *that has an advantage* $\mathbf{Adv}_{\mathcal{T},\mathcal{A}}^{\text{IND-ID-CPA}} = \epsilon(\lambda)$ *against the IBST* $\mathcal{T}$ *as defined in Section 3.2. Then there is an* IND-ID-CPA *adversary* $\mathcal{B}_{\mathsf{BasicIdent}}$ *that has advantage at least* $\epsilon'(\lambda) \approx \epsilon(\lambda)$ *against* BasicIdent. *Its running time is* $O(time(\mathcal{A}_\mathcal{T}))$.

Proof. The proof considers a sequence of games that differ only in the challenge given to the adversary. In the following, let τ_0, τ_1, s be defined as in Definition 2. Additionally, the challengers in the games draw bitstrings u_k of length k uniformly at random. We describe the games by the challenge given to the adversary:

Game 1: $(\tau_0, s) = \big(\phi\,(r, P_b)\,||\,G_b^{\text{tag}}(\alpha),\, G_b^{\text{key}}(\alpha)\big)$, where $\alpha = e_b\big(rs_bP_b, H_b(id)\big)$
This is the game $\mathsf{ID\text{-}TAG}^0_{\mathcal{T},\mathcal{A}}$.

Game 2:
$$(\tau_0, u_\lambda) = \left(\phi\,(r, P_b)\,||\,G_b^{\text{tag}}\left(e_b\big(rsP_b, H_b(id)\big)\right), u_\lambda\right)$$

Game 3: $\left(\phi\,(r, P_b)\,||\,u_{\ell_h}, u_\lambda\right)$

This is a hybrid game where the hash part of the tag is replaced by a random bitstring.

Game 4: $(\tau_1, u_\lambda) = (u'_\lambda, u_\lambda) = (u_{\ell_q}\,||\,u_{\ell_h}, u_\lambda)$

(Recall that $\tau_1 \leftarrow \mathcal{C} = \{0,1\}^\lambda$ is just a random covertext.)

Game 4a: $(\tau_1, s) = \left(u_\lambda, G_b^{\text{key}}\left(e_b\big(rs_bP_b, H_b(id)\big)\right)\right)$

This is the game $\mathsf{ID\text{-}TAG}^1_{\mathcal{T},\mathcal{A}}$.

By assumption, $\mathbf{Adv}^{\mathsf{ID\text{-}TAG}}_{A}(\lambda) = \big|\Pr[\mathsf{ID\text{-}TAG}^0_{\mathcal{T},\mathcal{A}}(\lambda) = 0] - \mathsf{ID\text{-}TAG}^1_{\mathcal{T},\mathcal{A}}(\lambda) = 0]\big|$ is non-negligible. In other words, $\mathcal{A}$ can distinguish Game 1 and Game 4a. Thus, $\mathcal{A}$ has a non-negligible probability in distinguishing two subsequent of our games. We distinguish the respective cases.

Game 1 $\not\approx$ Game 2 First, suppose that this advantage is between Game 1 and Game 2; let $\mathcal{D}$ be an adversary that distinguishes these games. In Game 1, the second element of the challenge pair is $G_b^{\text{tag}}(e_b(rs_bP_b, H_b(id)))$, while it is a uniformly chosen bitstring in Game 2. As G_b^{tag} is a random oracle, the games are perfectly indistinguishable as long as the adversary $\mathcal{D}$ does not query G_b^{tag} on $e_b(rs_bP_b, H_b(id))$. For case it does, we construct an adversary $\mathcal{B}$ against the simultaneous IND-ID-CPA security of two BasicIdent variants with the domain groups $\mathbb{G}_0$, $\mathbb{G}_1$, respectively. We denote the hash functions they use by $H_0^{\mathsf{BI}}, G_0^{\mathsf{BI}}$ as well as $H_1^{\mathsf{BI}}, G_1^{\mathsf{BI}}$. Under abuse of notation, we denote the combined game also by IND-ID-CPA if it is clear from the context which one we refer to.

$\mathcal{B}$ receives two master public keys P_0^{pub}, P_1^{pub} for BasicIdent and sends the pair (P_0^{pub}, P_1^{pub}) as master public key for $\mathcal{T}$ to $\mathcal{D}$. In phase 1 and phase 2, $\mathcal{B}$ relays all extraction oracle queries from $\mathcal{D}$ to the corresponding oracles of its two challengers and answers to $\mathcal{D}$ with (sk_0, sk_1), where sk_b is the answer from the corresponding oracle. Moreover, $\mathcal{D}$'s queries to $G_{b'}^{\text{tag}}$ are relayed to $G_{b'}^{\mathsf{BI}}$, those to $H_{b'}$ are relayed to $H_{b'}^{\mathsf{BI}}$ for all $b' \in \{0,1\}$. $\mathcal{B}$ keeps a list of all queries to $G_{b'}^{\text{key}}$.

When $\mathcal{D}$ sends a challenge identity id, $\mathcal{B}$ chooses two challenge messages $m_0, m_1 \leftarrow \{0,1\}^{\ell_h}$ uniformly at random and sends (m_0, m_1, id) to one of the challengers. $\mathcal{B}$ decides uniformly at random which challenger to play against.

Let $c = (c_1, c_2)$ denote the received challenge ciphertext. $\mathcal{B}$ sends the challenge $(c_1||c_2, u_\lambda)$ to $\mathcal{D}$.

When $\mathcal{D}$ returns a bit $b_\mathcal{D}$, $\mathcal{B}$ checks if $\mathcal{D}$ has queried $G^{\text{key}}_{b'}$ on a value x such that $G^{\text{tag}}_{b'}(x) \oplus c_2 = m_{\tilde{b}}$ for some $\tilde{b}, b' \in \{0,1\}$. If yes, $\mathcal{B}$ returns $\tilde{b}$ to its chosen challenger. In that case, as $G^{\text{key}}_{b'}(x) = G^{\text{key}}_{b'}(e_b(rs_bP_b, H_b(id)))$ implies $x = e_b(rs_bP_b, H_b(id))$ with overwhelming probability,[5] we have $G^{\text{tag}}_{b'}(x) \oplus c_2 = G^{\text{tag}}_{b'}(x) \oplus G^{\text{tag}}_{b'}(x) \oplus m_{\tilde{b}} = m_{\tilde{b}}$. That is, $\mathcal{B}$ has decrypted the challenge ciphertext and returned the correct bit $\tilde{b}$.

If no query x fulfills the mentioned property, $\mathcal{B}$ returns $b_\mathcal{D}$. Altogether, we know that $\mathcal{B}$ wins one of the IND-ID-CPA games with overwhelming probability, whenever $\mathcal{D}$ manages to query G^{key}_0 or G^{key}_1 on $e_b(rs_bP_b, H_b(id))$. One can verify that the simulation is perfect up that point. If $\mathcal{D}$ does not issue such a query, Game 1 and Game 2 look clearly identically. It follows

$$\mathbf{Adv}^{\text{IND-ID-CPA}}_{\text{BasicIdent},\mathcal{B}}(\lambda) \geq \left| \Pr_{\text{Game 1}(\lambda)}[\mathcal{D} \text{ outputs } 0] - \Pr_{\text{Game 2}(\lambda)}[\mathcal{D} \text{ outputs } 0] \right| - \delta(\lambda) \ ,$$

which is non-negligible for a negligible function δ because of the case condition.

Game 2 $\not\approx$ Game 3 In this case, we construct again an adversary $\mathcal{B}$ against the two IND-ID-CPA games with groups $\mathbb{G}_0$ and $\mathbb{G}_1$. $\mathcal{B}$ relays all oracle and extraction queries from the distinguisher $\mathcal{D}$, including the answers, exactly as in the previous case. Additionally like in the previous case, $\mathcal{B}$ chooses in the challenge phase one of the IND-ID-CPA challengers uniformly at random and sends it (id, m_0, m_1), where $m_0 = 0^{\ell_h}$ and $m_1 \leftarrow \{0,1\}^{\ell_h}$. It sends the challenge $(c_0||c_1, u_\lambda)$ to $\mathcal{D}$, where (c_0, c_1) is the received challenge ciphertext and $u_\lambda \leftarrow \{0,1\}^\lambda$. When $\mathcal{D}$ outputs its guess $b_\mathcal{D}$, $\mathcal{B}$ relays this guess to its challenger.

Observe that $\mathcal{B}$ simulates Game 2 perfectly if the input bit of the chosen IND-ID-CPA challenger is 0. This is because $h \oplus m_0 = h \oplus 0^{\ell_h} = h$ for the second part $h := e_b(rs_bP_b, H_b(id))$ of a real tag.

On the other hand, if the challenge bit is 1, $\mathcal{B}$ simulates Game 3 perfectly: Since m_1 is chosen at random, the second part of the challenge tag $h \oplus m_1$ is completely independent of all other information that $\mathcal{D}$ has.

Thus, $\mathcal{B}$'s guess in the IND-ID-CPA game is correct whenever $\mathcal{D}$ correctly distinguishes between Game 2 and Game 3 and we have

$$\mathbf{Adv}^{\text{IND-ID-CPA}}_{\text{BasicIdent},\mathcal{B}}(\lambda) \geq \left| \Pr_{\text{Game 2}(\lambda)}[\mathcal{D} \text{ outputs } 0] - \Pr_{\text{Game 3}(\lambda)}[\mathcal{D} \text{ outputs } 0] \right| \ ,$$

which is non-negligible because of the case condition.

Game 3 $\not\approx$ Game 4 As everything except the second part of the candidate tag is chosen independently at random, this is a contradiction to the fact that x-coordinates of random group elements output by ϕ are indistinguishable from

[5] It is possible that $\mathcal{D}$ has found a collision in the random oracles G^{tag}_0 and G^{tag}_1. This happens only with negligible probability, say $\delta(\lambda)$.

random bitstrings of size ℓ_q. Like in Telex, that holds, because we have chosen q and the group $\mathbb{G}_0$ and its twist $\mathbb{G}_1$. See the appendix in [7] for a proof.

Game 4 $\not\approx$ Game 4a This is impossible because these games are identical from the point of view of the adversary. To see that, recall that G_0^{key} and G_1^{key} are random oracles and the adversary is given no information about r, which is chosen at random.

Summing up, there is a combined adversary $\mathcal{B}$ that breaks the $\mathsf{IND\text{-}ID\text{-}CPA}_{\mathsf{BasicIdent}}$ game either with $\mathbb{G}_0$ or with $\mathbb{G}_1$ and has advantage $\mathbf{Adv}^{\mathsf{IND\text{-}ID\text{-}CPA}}_{\mathsf{BasicIdent},\mathcal{B}} \approx 2\,\mathbf{Adv}^{\mathsf{ID\text{-}TAG}}_{\mathcal{T},\mathcal{A}}$ and running time $O(time(\mathcal{A}))$. □

To prove our main result, we require two results regarding the security of BasicIdent, which is proven by Boneh and Franklin [6] using an intermediate public-key encryption scheme BasicPub secure under the computational bilinear Diffie-Hellman assumption. Galindo [21] shows a tighter security reduction of BasicPub but against the stronger decisional variant of the assumption. By relying on these results and the one we just proved, it is straightforward to prove our main theorem. It reduces the indistinguishability of our construction to the BDDH assumption in the random oracle model.

Proof (of Theorem 1, page 6). The statement follows immediately from the combination of Lemma 1, Result 7 in [21, Appendix B], and Lemma 4.2 in [6]. □

B Impossibility of Immediate Forward Security in Telex

A natural approach to achieve immediate forward security in a protocol is as follows: For each session, the protocol parties P_i draw individual random secrets s_i and perform a key exchange to derive a shared secret s. Once such a secret is not necessary anymore for a party (no later than at the end of the session), this party can securely erase s_i and all derived secrets. Afterwards, the corresponding session secrets can not be derived anymore.

Consider a protocol, or a phase of such, with some party P whose output does not depend on any randomness. P must be able to participate in the protocol phase, using as input only the received messages and its long-term key.

Thus, an eavesdropping attacker that has been able to read and store all messages on the network possesses all the necessary information to compute the whole view of P in past sessions once she manages to corrupt P. This shows that such protocol phases cannot achieve immediate forward security.

Telex has such a protocol phase: For practicality, it is critical that the Telex station can recognize connection requests after the first message from the client. The reason is that the Telex station should only interfere, i.e., change data flowing from the unsuspicious website to the client, with TLS connections for which the client almost certainly wishes to establish a Telex connection. Otherwise, Telex would break TLS on a wide scale, because the station would try to act as a man-in-the-middle in normal TLS sessions. That would be detected, and thus the session would abort. Altogether, the station cannot send data to the client before deciding to treat a connection as a Telex connection.