TANDEM: Securing Keys by Using a Central Server While Preserving Privacy

Wouter Lueks
SPRING Lab, EPFL
wouter.lueks@epfl.ch

Brinda Hampiholi
Radboud University
brinda@cs.ru.nl

Greg Alpár
Open University
Radboud University
g.alpar@cs.ru.nl

Carmela Troncoso
SPRING Lab, EPFL
carmela.troncoso@epfl.ch

Abstract—Users' devices, e.g., smartphones or laptops, are typically incapable of securely storing and processing cryptographic keys. We present TANDEM, a novel set of protocols for securing cryptographic keys with support from a central server. TANDEM uses *one-time-use key-share tokens* to, unlike traditional threshold-cryptographic solutions, preserve users' privacy with respect to a malicious central server. Additionally, TANDEM enables users to block their keys if they lose their shares, and it enables the server to limit how often an adversary can use an unblocked key. We prove TANDEM's security and privacy properties, and we empirically show that it causes little overhead using a proof of concept implementation. To illustrate TANDEM's advantages we use it to secure attribute-based credentials keys using a central server without hurting the privacy properties provided by the credential system.

Index Terms—privacy, threshold cryptography, central server, time-correlation attack

I. INTRODUCTION

The security of cryptographic schemes hinges on the security of the underlying keys. However, secure solutions to store and process keys on users' devices are hard to deploy in practice. Software-based approaches are extremely difficult to secure [1], [2], [3], [4], and secure hardware [5], [6], [7] might not be available on users' devices, not accessible to developers [8], [9], or harmful to usability [10].

As an alternative, users could use a secure central server to store their keys and perform cryptographic operations on their behalf, and to block their keys if their devices are compromised. The problem is that centralization introduces security and privacy concerns that are not an issue when keys are stored on the user's device. First, users must trust the central server to not impersonate them. Second, the central server is in a privileged position to learn private information about users from their interactions with other services. Brandão et al. illustrate these problems in the context of nation-scale brokered-identification systems [11]. They show how a central hub that acts as the broker between users, identity providers, and service providers can impersonate users, link users' transactions across different service providers, and also learn private identifiable information about users.

A natural solution to the impersonation problem is to involve the user in the storage and/or usage of the keys by using threshold cryptography. This approach additionally strengthens authentication security as the user needs a second factor: a key share. However, threshold cryptography does not address the privacy concerns associated with centralization. The central server learns the users' key-usage patterns and, as the time of access and use of the key are almost the same, it can use this information to deanonymize users' anonymous transactions, e.g., correlating interactions to public activities such as updates to a blockchain ledger [12], [13].

In this work, we present TANDEM, a set of protocols that augment threshold-cryptographic schemes to enable secure and privacy-preserving usage of key shares managed by a central server. To use a key, a user sends a *one-time-use key-share token* to the central server using an anonymous communication channel. This token contains a randomized version of the central server's key share for this user. The server uses this key share to run the threshold-cryptographic protocol *without* learning the user's identity.

The construction of key-share tokens permits to decouple the stages of obtaining and using the tokens, eliminating the possibility of time-correlation attacks. Furthermore, the one-time property enables two additional functionalities: the blocking of keys in case the user's key share is compromised, and the rate limiting of key usage to restrict how often an attacker can use an unblocked key. TANDEM provides these functionalities *without* the need to identify token owners.

TANDEM can be used to secure the keys of any cryptographic scheme (e.g., encryption, signature, or payments) for which a linearly randomizable threshold-cryptographic version of the scheme exists. As long as the threshold version is private, i.e., the scheme does not require information that identifies the user besides the key, TANDEM ensures that not even a malicious central server can learn with which user it is interacting. For example, Tandem can be applied to threshold variants of Schnorr [14] and RSA signatures [15], ElGamal-based [16], [17] and RSA decryption [15], as well as threshold-cryptographic versions of electronic cash schemes [18], [19] and attribute-based credential schemes [20], [21], [22], [23]. We note, however, that Tandem cannot be applied to existing threshold DSA schemes because they are multiplicative [24] or require identifying auxiliary information [25].

To demonstrate the potential of TANDEM we use it to secure keys in a threshold version of BBS+ attribute-based credentials (ABCs). ABCs [20], [21], [22], [23] protect users' anonymity during authentication on sensitive online services,

e.g., online health services[1]. Thus, enhancing them with a naïve centralized approach in which the central server could learn which sensitive services users access would defeat the very purpose of ABCs. Using TANDEM to secure her keys, the user can show her credential to the online service *without* the TANDEM server learning who is using the key, preserving the user's privacy even if the TANDEM server and the service provider collude. Moreover, the user never has the complete key in her device.

The anonymity provided by ABCs opens the door to malicious users abusing service providers. We also provide a simple modification to the threshold ABC schemes that enables service providers to confirm that TANDEM is used. Then, as long as all users use TANDEM, TANDEM can replace the complex ad-hoc cryptographic techniques to block users [26], [27] or limit key-usage [28].

We validate the practicality of the TANDEM protocols on a prototype C implementation. Using a key with TANDEM induces a $50-100$ ms overhead on the TANDEM server with respect to traditional threshold-cryptographic solutions, and only 5 ms overhead on the user. The cost for the server is manageable. On the user side, the overhead is negligible with respect to the delay imposed by the use of anonymous communications necessary for typical uses of TANDEM such as anonymous web-based authentication.

In summary, we make the following contributions:

✓ We introduce TANDEM; it enables the use of threshold-cryptographic protocols with a central server to secure cryptographic keys without this server learning what keys are used by whom. Additionally, TANDEM enables blocking and rate limiting of key usage.

✓ We provide a threshold version of an attribute-based credential system, and show how TANDEM can be used to augment its security. We show how the underlying constructions in TANDEM permit rate limiting and revocation of credentials *without* relying on complex purpose-built cryptographic techniques.

✓ We prove the security and privacy of TANDEM, and we use a prototype implementation to validate its practicality. All operations in TANDEM take less than one second, imposing a reasonable overhead on both server and users.

II. RELATED WORK

Exisiting solutions to protect cryptographic keys fall into two coarse categories, either *single-party* or *decentralized*. The former typically rely on secure hardware [29], [7], [6], [5] that securely stores and processes cryptographic keys within the secure environment. However, secure hardware is expensive, is not always available (e.g., in laptops) or not accessible by application developers, and is often not flexible enough to run advanced protocols.

Threshold cryptography aims to strengthen general cryptographic protocols by distributing the user's secret key

TABLE I: Comparison of generic and special-purpose TCPs with TCPs augmented with Tandem.

	Generic e.g., [14], [17]	Purpose [41], [21]	with TANDEM
Anonymous key usage	×	×	✓
Hide protocol data	×	✓	×[*]
Hide key-usage patterns	×	×	✓

[*] This property is irrelevant for TANDEM; see text.

among several parties. This approach was first proposed by Desmedt [30] and Boyd [31]. Several threshold encryption and signature schemes have been proposed since then [32], [33], [34], [14], [15], [35], [36], [37]. More recently, Atwater et al. [38] built a library to execute such protocols in users' personal devices. Other works have tackled more complicated protocols. For instance, Brands shows how to distribute the user's secret key in attribute-based credentials [21], and Keller et al. [39] show how to make threshold-cryptographic versions of zero-knowledge proofs.

Many works propose systems in which the user's secret key is shared between a user's device and a central server to protect the key and also to enable instant blocking of the key by the user [40], [41], [42], [43], [44], [45]. However, none of these schemes provide privacy for the user towards the central server. In all of these schemes users authenticate themselves to the server, making them susceptible to time-correlation attacks [12]. Camenisch et al. [41] attempt to ensure privacy to some extent in signature schemes by blinding the message being signed during the threshold protocol with the server. Yet, the server learns when and how often the user uses her signing key. Hence, users are still vulnerable to timing attacks. The scheme by Brands [21] protects against these attacks as long as the key-share holder is a smartcard, which cannot store a timed log of operations. However, if the smartcard is replaced by an online server that holds the key share, this server learns the key-usage patterns of users. Then, the cryptographic measures proposed by Brands alone cannot prevent time-correlation attacks.

TANDEM is designed to complement these threshold-cryptographic solutions to make them privacy friendly. We compare the privacy properties obtained when using TANDEM with those in previous proposals in Table I. We consider three privacy aspects: anonymity of the user when running the threshold protocol (i.e., need to authenticate); hide the data used (e.g., signed message) in the protocol from the server; and hide the usage pattern to avoid timing attacks. Generic schemes focus on achieving security of the secret key and thus provide no privacy. Special-purpose designs, have so far only focused on the protection of data involved in the protocol. TANDEM does not need to protect this data: the other two properties decouple the data from the user's identity.

Finally, similarly to TANDEM, password-hardening services [46], [47] use decentralization to increase security against brute-force attacks. These schemes introduce a hardening server that rate limits, or even blocks, requests from

---

[1] such as https://medical.mit.edu/services/mental-health-counseling and https://www.nhs.uk/Conditions/online-mental-health-services/Pages/introduction.aspx

the main authentication server. However, the cryptographic techniques behind these solutions cannot be directly applied to the problem tackled by TANDEM. First, they are designed for a particular task: securely verifying passwords, and adapting them to run other protocols is non-trivial. Second, in the password scenario the hardening server is only accessed by the authentication server. Therefore, there are no privacy concerns, and these techniques do not provide any privacy protection.

III. PROBLEM STATEMENT

We consider a scenario in which *users* are required to perform cryptographic operations to interact with a *service provider* (SP). Users use insecure devices, such as smartphones, tablets, or laptops, without secure hardware, to run the cryptographic protocols. To keep their keys safe, they use a centralized TANDEM *Server* (TS) to run threshold-cryptographic protocols (TCPs) in a distributed way. Users wish to keep their key-use pattern and their use of other services private with respect to the TS. We call an execution of the protocol between the user and the TS a *transaction*.

For simplicity, we assume that there is only one TANDEM server. However, we note secret-sharing the key with multiple TANDEM servers would increase security. In this case TANDEM would ensure that keys can be blocked and rate limited as long as at least one of the TANDEM servers is honest. Privacy is not affected by the number of servers.

A. TANDEM *Properties and Threat model*

PROPERTY 1 (Key security). TANDEM *protects the use of the user's key*. No entity other than the user is able to use the user's key. Even if the user's device is compromised, the user can maintain this property by *blocking the key* at the TANDEM server. Thereafter, the attacker cannot further use her key. We formalize this property in Game 1 in Section VI.

Any solution that recomputes the full user's key on the user's device, e.g., by deriving it from a user-entered passphrase, does not satisfy this key-security property. In such a solution, an attacker who compromises the user's device can observe the full key when it is used. Thereafter the attacker can use the key indefinitely, making blocking impossible.

PROPERTY 2 (Key rate-limiting). TANDEM *limits the rate of usage of keys*. Users can limit the number of times her key is used in a given interval of time. We call this interval an *epoch*. We formalize this property in Game 2 in Section VI.

The security and rate-limiting properties are related to the revocation and n-times-use concepts of attribute-based credentials [28], respectively. Yet, they are not the same. Revocation and n-times-use credentials trust the service providers to block credentials respectively to block a credential after n uses. Using TANDEM on the other hand, users need to trust only the TS, which *they* choose, to block and rate-limit keys. TANDEM can ensure this property for a large class of protocols, even if a system does not rely on credentials.

PROPERTY 3 (Key-use privacy). TANDEM *protects the privacy of key use in transactions*. The TANDEM server (TS) cannot distinguish between two users performing transactions. Even if the TS colludes with the service provider (SP) it cannot distinguish users (unless the SP could distinguish the users, in which case collusion leads to a trivial and unavoidable privacy breach). We formalize this property in Game 4 in Section VI.

We assume that the TANDEM server is honest with respect to security. That is, it follows the protocols so as to protect the security of users' keys (Property 1) and to ensure that keys are only used the allowed number of times (Property 2). Moreover, we trust the TANDEM server to be available, i.e., TANDEM does not protect from denial of service. However, the TANDEM server may be malicious with respect to privacy: It is interested in breaching the privacy of the users by trying to learn which keys and services they use (Property 3).

Why naïve solutions do not work. Consider an approach in which a user naïvely secret-shares her key with the TANDEM server. When she needs to run a threshold-cryptographic protocol, the user authenticates to the TS, the TS recognizes the user, retrieves its share of the user's key, and executes the TCP together with the user. This scheme offers key security (Property 1): The TS alone cannot use the user's key, and if an attacker compromises the user's device, the user can authenticate to the TS and request it to block her key. This scheme also provides key rate limitation (Property 2): the TS can observe when a user accesses her key. Hence, it can easily enforce a limit on the number of times the key is used. However, since the user is identified while using the key for a TCP, the scheme does not achieve key-use privacy (Property 3).

The lack of key-use privacy has further implications when the interactions between the user and the SP are anonymous (e.g., showing an anonymous credential). The SP can collude with the TS to learn the user's identity, exploiting the fact that there is a strong correlation between the time when the authenticated user interacts with the TS, and when the anonymous user interacts with the SP. Thus, for every anonymous transaction with the SP, the anonymity set of the user is reduced to the authenticated users interacting with the TS around the transaction time. This attack has been used in the early days of Tor to identify users and hidden services [48], [49]. The attack relies solely on time correlation between accesses. Therefore, the attack cannot be prevented by making the messages seen by the TS and the SP during the TCP cryptographically unlinkable [21].

There are two straightforward approaches to prevent time-correlation attacks: introducing delays and introducing dummy requests. These solutions are, however, difficult to use in practice. To significantly increase the anonymity set for users, operations may need to be delayed for a long time. This rules out applications that require short delays, such as showing an anonymous credential or performing a payment. Dummy traffic not only imposes an overhead on users and the TS, but it is widely known that generating dummy actions that
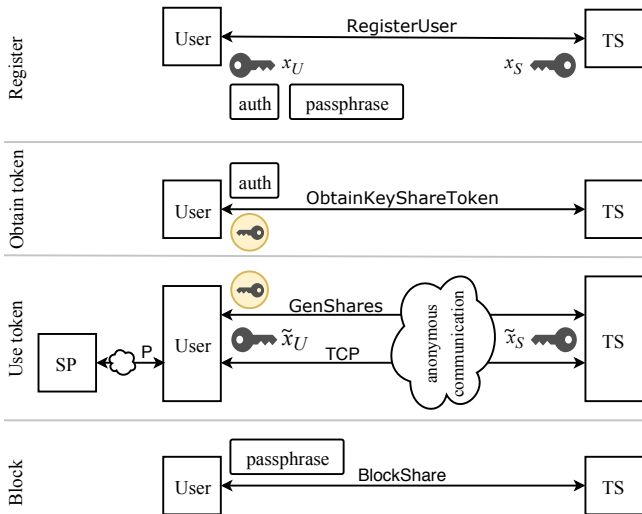
Fig. 1: TANDEM process: after registration, a user can authenticate herself and obtain key-share tokens, which can later be used anonymously to execute a TCP (the user and SP run protocol P). The user can block her keys at any time. Inputs are shown above the arrows, outputs below.

are indistinguishable from real activity is very difficult [50], [51]—especially because it is unrealistic that users would be always online so that their devices could produce such requests.

B. TANDEM *at a Glance*

We now provide a high-level overview of how users can use the TANDEM server to perform threshold-cryptographic protocols in a privacy-preserving way, see Fig. 1. We assume that users can use an anonymous communication channel [52], [53] to communicate with the TANDEM server (TS) and service providers (SPs) to protect their privacy at the network layer.

Registration. First, users register with the TS using the RegisterUser protocol. During registration, the user and TS jointly compute long-term shares x_U and x_S of a long-term key x appropriate for the threshold-cryptographic protocols they seek to run later. The user obtains credentials to authenticate when obtaining a token (e.g., a password) and also a means to block her keys (e.g., a passphrase); she stores the latter outside her device.

Obtain Token. Key-share tokens enable the user to anonymously use her key later (see below). To obtain a token, the user runs the ObtainKeyShareToken protocol with the TS. First, the user authenticates herself to the TS. Then, the user and the TS construct a one-time-use key-share token, containing a randomized version of the TS's key share x_S. At this stage, the TS can limit the number of tokens it provides the user, thus limiting how often the user can use her key.

Key-share tokens may seem similar to passwords: both unlock functionality. However, unlike passwords, key-share tokens can be verified and used *without* knowing the user's identity. Moreover, key-share tokens contain a randomized

key share $\tilde{x}_S$ essential for the TCP. Hence, TANDEM cannot be replaced by a password-hardening service [46], [47]. The randomized key shares contained in the tokens also distinguish the tokens from traditional eCash tokens [54], [18], [19].
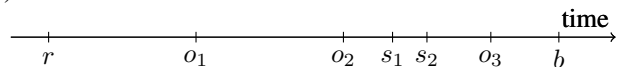
Using Keys. After obtaining tokens, a user can run threshold-cryptographic protocols with the TS. First, the user and the TS use the token to derive fresh shares $\tilde{x}_U$ and $\tilde{x}_S$ by running the GenShares protocol. The new share $\tilde{x}_S$ cannot be linked to the long-term share x_S, thus it does not reveal the user's identity to the TS. The user and the TS use the fresh shares $\tilde{x}_U$ and $\tilde{x}_S$ as input to the threshold-cryptographic protocol TCP, allowing the user to use her key in the cryptographic protocol P with the service provider.

The TS never communicates directly with service providers, but only via the user. Therefore, the use of TANDEM can remain invisible to the service provider, i.e., users can use TANDEM without the SP's knowledge or consent.

Blocking keys. Whenever a user wants to block her key, she requests the TS to block her key by using the BlockShare protocol with her blocking means (e.g., the passphrase). Thereafter, unused tokens become invalid, and no new tokens can be obtained (not even by an adversary that knows the authentication credential used to obtain tokens). Hence, the user's key cannot be used anymore.

Preventing Time Correlation. When obtaining tokens, the user is authenticated. Hence, to preserve privacy, the actions of obtaining and using tokens *must* be uncorrelated, i.e., tokens should not be obtained right before usage. To avoid correlation, the user can configure her device to obtain tokens at random times or at regular times (e.g., every night) such that tokens are always available. The user can authenticate herself to the device at those times, or automate the process by storing her authentication credential on the device. Note that the user's key can still be blocked at the TS if an attacker learns this authentication credential.

Here we show an example time line of registration (r), obtaining tokens (o_i), using tokens (s_i), and blocking the key (b) events:

This example illustrates that obtain and use events do not necessarily follow each other, but can be interleaved. As a result, the timing of these events needs *not* to be correlated. The token o_3, unused before the key is blocked by b, cannot be used after time b.

In Section V-B we explain why private information retrieval is not a suitable alternative to decouple obtaining and using of key shares, and how TANDEM outperforms generic alternatives based on secure multi-party computation.

IV. CRYPTOGRAPHIC PRELIMINARIES

Let ℓ be a security parameter. Throughout this paper, $\mathbb{G}$ is a cyclic group of prime order p (of 2ℓ bits) generated by g. We write $\mathbb{Z}_p$ for the integers modulo p. We use a cryptographic hash function $H : \{0, 1\}^* \to \mathbb{Z}_p$ that maps strings to integers

modulo p. We write $a \in_R A$ to denote that a is chosen uniformly at random from the set A. Furthermore, we write $[n]$ to denote the set $\{0, \ldots, n-1\}$.

For reference, Table II in Section V summarizes the notation used by TANDEM's building blocks, and Table III in Section V explains frequently-used symbols in TANDEM.

A. Cryptographic Building Blocks

TANDEM relies on a couple of cryptographic building blocks. First, we use an additive homomorphic encryption scheme given by the algorithms $\mathbf{E}_{pk}^+, \mathbf{D}_{sk}^+$ with plaintext space $\mathbb{Z}_N$ (i.e., integers modulo N) and space of randomizers $\mathcal{R}$. We write $c = \mathbf{E}_{pk}^+(m; \kappa)$ to denote the homomorphic encryption of the message $m \in \mathbb{Z}_N$ using randomness $\kappa \in \mathcal{R}$. The scheme is additively homomorphic, so

$$\mathbf{E}_{pk}^+(m_1; \kappa_1) \cdot \mathbf{E}_{pk}^+(m_2; \kappa_2) = \mathbf{E}_{pk}^+(m_1 + m_2 \pmod{N}; \kappa_1 \kappa_2).$$

Our proof of concept uses Joye and Libert's encryption scheme [55], but Paillier's scheme [56] would also work.

Second, TANDEM uses a CPA secure encryption scheme $\mathbf{E}_{pk_{id}}, \mathbf{D}_{sk_{id}}$ with plaintext space $\mathbb{G}$ such as ElGamal [57], that allows simple verifiable encryption.

Third, TANDEM uses two computationally hiding and binding commitment schemes. First, by $\mathsf{Commit}(m, r)$ we denote a commitment function that takes a message $m \in \mathbb{Z}_p$ and a randomizer $r \in \mathbb{Z}_p$. Analogously, we define $\mathsf{Commit}((m_1, \ldots, m_k), r)$ to commit to a tuple of messages. We instantiate this scheme using Pedersen's commitments [58], because it enables users to obtain a blind signature on the tuple $(m_1, \ldots, m_k)$. However, any other commitment scheme with these properties suffices as well.

Second, we denote by $\Delta = \mathsf{ExtCommit}(m, r)$ with $m \in \{0, 1\}^*, r \in \{0, 1\}^{2\ell}$ an extractable commitment scheme. That is, in our security reductions, we can extract the input m used to create a commitment Δ. For example, the instantiation $\mathsf{ExtCommit}(m, r) = H(m\|r)$ is extractable in the random oracle model.

B. Threshold-Cryptographic Protocols

In this paper, we focus on cryptographic protocols run between a user and a service provider, e.g., showing a credential to an SP or spending an electronic coin. The threshold-cryptographic version of such a protocol splits the user's key x and the user's side of the original protocol in two parts, run by different parties. Each party operates on a secret-share of the user's key. Security of the threshold-cryptographic protocol (TCP) ensures that a large enough subset of shares (two in the case of two parties) are required to complete the protocol.

We consider TCPs where the user's side of the protocol is distributed between the user and the TS. After registration, the user and the TS hold the key shares x_U and x_S of x, respectively. After running GenShares, the user and the TS hold the fresh key shares $\tilde{x}_U$ and $\tilde{x}_S$. They then run the TCP protocol, which we denote as:

$$\mathsf{P}(\mathsf{in}_{SP}) \leftrightarrow \mathsf{TCP.U}(\tilde{x}_U, \mathsf{in}_U) \leftrightarrow \mathsf{TCP.TS}(\tilde{x}_S), \qquad (1)$$

where the SP, the user and the TS respectively run the interactive programs P, TCP.U and TCP.TS. The user mediates all interactions between the service provider and the TS. The user and the SP take extra inputs needed for the execution of the target cryptographic protocol denoted as in_U and in_{SP}. For simplicity, we denote the complete protocol from (1) by $\mathsf{TCP}(\tilde{x}_U, \tilde{x}_S, \mathsf{in}_U, \mathsf{in}_{SP})$.

TANDEM can only enhance the privacy (Property 3) of certain TCPs. We formalize the condition that these TCPs should satisfy. To avoid that the TS can recognize the user based on the shares input to the TCP, we randomize the long-term secret shares. Thus, we require that TCPs enhanced with TANDEM still function with randomized key shares. In addition, our privacy-friendly GenShares protocol requires this randomization to be linear.

For simplicity, we assume that the user's secret $x \in \mathbb{Z}_p$ for some field $\mathbb{Z}_p$ of prime order p (e.g., corresponding to the group $\mathbb{G}$ we defined above). We note, however, that our constructions can be modified to settings with unknown order arising from RSA assumptions. Formally, we require the TCP to be linearly randomizable:

DEFINITION 1. Let $x_U, x_S \in \mathbb{Z}_p$ be secret shares of the user's secret x. Then, we say that the TCP is *linearly randomizable* if for all δ we have that (1) if $\mathsf{TCP}(x_U, x_S, \mathsf{in}_U, \mathsf{in}_{SP})$ completes successfully, then so does $\mathsf{TCP}(x_U - \delta, x_S + \delta, \mathsf{in}_U, \mathsf{in}_{SP})$, and (2) $x_S + \delta$ is independent from x_S.

The first condition implies that the original secret sharing (x_U, x_S) and the randomized secret sharing $(x_U - \delta, x_S + \delta)$ must share the same secret, whereas the second implies that the TS cannot recognize the user from the randomized secret share alone.

Security and privacy properties of TCPs. To ensure that a TCP with TANDEM satisfies the security properties (Property 1 and Property 2) we require that the TCP itself is secure. That is, if the TS no longer uses its share x_S to run its part of the TCP, then no malicious user can successfully complete the TCP with the SP. We formalize this notion in Game 3 in Section VI.

To ensure that a TCP with TANDEM satisfies the privacy property (Property 3) we require that the TCP itself offers privacy with respect to the TS respectively the TS and the SP. That is, if the TS runs its part of the TCP using a randomized key-share as input, then the TS respectively the TS and the SP cannot recognize the user. We formalize this notion in Game 5 in Section VI.

V. TANDEM

In this section, we present a construction that enables anonymous users to use their keys with the TS *without* the TS learning which key is being accessed. It uses homomorphic encryption to decouple the action of accessing the user's long-term key-share x_S at the TANDEM server from its subsequent use in the threshold-cryptographic protocols. Thus, it prevents time-correlation attacks.

Initially, the TS generates a private-public key pair (sk, pk) for an additively homomorphic encryption scheme (see Section IV). The TS publishes the public key pk. Upon registration with the TS, a user receives $\overline{x_S} = \mathbf{E}^+_{pk}(x_S)$—a homomorphic encryption of the TS's key-share x_S. Because the ciphertext $\overline{x_S}$ is encrypted against the TS' key, the user does not learn anything about the TS' share x_S.

When the user wants to *use* her key, she produces a randomized version of the TS' key-share x_S. To produce this randomization, she picks a large δ and computes $c = \overline{x_S} \cdot \mathbf{E}^+_{pk}(\delta) = \mathbf{E}^+_{pk}(x_S + \delta)$. On her side, she randomizes her key as $\tilde{x}_U = x_U - \delta \pmod{p}$. Then, she sends c to the TS via an anonymous channel. The TS decrypts c to recover its key for the threshold cryptographic protocol, $\tilde{x}_S = x_S + \delta \pmod{p}$. It is easy to see that a linear TCP with randomized shares completes successfully, because $\tilde{x}_S + \tilde{x}_U = x_U + x_S \pmod{p}$. Because c is randomized, the TS can no longer recognize its share x_S, effectively decoupling this action from the key-share generation.

In this approach, however, the TS cannot block or rate-limit keys. We present below a construction for one-time-use key-share tokens containing signed and randomized ciphertexts like c that enables blocking and rate-limiting while preserving users' privacy.

A. One-time-use Key-share Tokens

To construct a token the user picks a large δ and homomorphically computes $c = \overline{x_S} \cdot \mathbf{E}^+_{pk}(\delta)$, a randomized encryption of the TS' key share. Then, she sends a commitment to c to the TS, together with a proof that the committed c was constructed by additively randomizing $\overline{x_S}$. This proof is needed to enable secure blocking as we explain below. The user engages with the TS to obtain a blind signature σ on c. The signature σ is only known to the user at this stage. The user stores the token $\tau = (\sigma, c)$ and the randomizer δ.

To run a threshold-cryptographic protocol the user anonymously contacts the TS and sends her key-share token $\tau = (\sigma, c)$. The TS checks the signature and makes sure the token was not used before. Then, the TS recovers the randomized key-share $\tilde{x}_S = \mathbf{D}^+_{sk}(c) \pmod{p} = x_S + \delta \pmod{p}$ and uses it as the key for the threshold cryptographic protocol. The user, on the other hand, uses $\tilde{x}_U = x_U - \delta \pmod{p}$ as the key. As in the previous case, because c is fully randomized, the TS cannot leverage it to identify users. Moreover, as σ is a blind signature on c the TS cannot use σ or c to link the token creation to the token use.

When a user asks the TS to block her key, the TS no longer creates key-share tokens for this user (we explain how the TS blocks unspent tokens below). This prevents attackers from further running threshold-cryptographic protocols, even if they corrupt the user's device. For the blocking of keys to be effective, attackers must not be able to construct key-share tokens for a blocked user. Here is where the proof becomes handy that c is constructed as $\overline{x_S} \cdot \mathbf{E}^+_{pk}(\delta)$, where $\overline{x_S}$ belongs to the current user. Suppose that we omit the proof. Then, an attacker controlling an unblocked user can create tokens

TABLE II: Notation and cryptographic building blocks used by TANDEM.

Symbol	Interpretation
$[n]$	The set $\{0, \ldots, n-1\}$
ℓ	The security parameter
$\mathbb{G}, g, p$	Cyclic group $\mathbb{G} = \langle g \rangle$ of order p
Additively homomorphic encryption scheme	
$\mathbf{E}^+_{pk}(m; r)$	Encrypt message $m \in \mathbb{Z}_N$ with randomizer $r \in \mathcal{R}$
$\mathbf{D}^+_{sk}(c)$	Decrypt ciphertext c
N	Size of additive plaintext domain
$\mathcal{R}$	Space of randomizers
CPA secure verifiable encryption scheme	
$\mathbf{E}_{pk_{id}}(m)$	Encrypt message $m \in \mathbb{G}$
$\mathbf{D}_{sk_{id}}(c)$	Decrypt ciphertext c
Commitment schemes and hash function	
$\mathrm{Commit}(m, r)$	Commit to $m \in \mathbb{Z}_p$ (or a tuple of messages) with randomizer $r \in \mathbb{Z}_p$
$\mathrm{ExtCommit}(m, r)$	Commit to $m \in \{0,1\}^*$ with randomizer $r \in \{0,1\}^{2\ell}$
$H(s)$	Hash function from $s \in \{0,1\}^*$ to $\mathbb{Z}_p$

TABLE III: Notation in TANDEM protocols

Symbol	Interpretation
$\mathcal{D}$	Disclose subset in cut-and-choose construction
δ, δ_i	Randomizers of key shares
id	Token identifier
k	Token security parameter
ℓ_δ	Length of randomizers δ_i in bits
x	Long-term secret key for a user
pk, sk	Public-private key-pair of TS
pk_{id}, sk_{id}	Public-private key-pair of the user U
p	Order of the group $\mathbb{G}$
x_U	Long-term key share held by the user
x_S	Long-term key share held by the TS
$\overline{x_S}$	Homomorphic encryption of x_S
$\tilde{x}_U$	User's key share output by GenShares
$\tilde{x}_S$	TS' key share output by GenShares
ϵ	The current epoch
σ	Blind signature of the TS

for a corrupted blocked user. The attacker uses the unblocked user's account to make the TS blindly sign encrypted key shares for the blocked user. The attacker can use the resulting token to use the blocked user's key, defeating the purpose of TANDEM. Verifying which user's key share is embedded into the ciphertext blindly signed by the TS prevents the attack.

Finally, since tokens are one-use only, to restrict the number of times a user can use her key (rate-limit), the TS just signs a limited number of key-share tokens per-epoch per-user.

Registering Users. When a user first registers at the TS, the TS computes a key-share x_S for that user, and sends her an encrypted version $\overline{x_S} = \mathbf{E}^+_{pk}(x_S)$. To ensure that the TS cannot hide an identifier in higher-order bits of x_S that are not randomized by the user in the remainder of the protocol the TS proves that the plaintext x_S is in the correct range.

PROTOCOL 1. The RegisterUser protocol is run between a user and the TS, and proceeds as follows.

1) The user U and the TS generate secret shares $x_U \in_R \mathbb{Z}_p$ and $x_S \in_R \mathbb{Z}_p$, respectively. The user also generates a public-private key-pair (pk_{id}, sk_{id}) for encrypting token

identifiers and sends pk_{id} to the TS. The user needs the secret key sk_{id} to block unspent tokens if needed. We assume that the user stores sk_{id} externally so that it is available even after she loses her device. We propose that the user's device generates sk_{id} based on a high-entropy passphrase (such as a Diceware passphrase[2]), so that users can write down this string as a stand-in for sk_{id}.

2) The TS picks $\kappa \in_R \mathcal{R}$, computes $\overline{x_S} = \mathbf{E}_{pk}^+(x_S; \kappa)$ and sends $\overline{x_S}$ to the user. Moreover, the TS sends a range proof to the user that $\overline{x_S}$ is constructed correctly, i.e., that

$$\mathbf{D}_{sk}^+(\overline{x_S}) \in [0, p). \tag{2}$$

See Appendix B for how to instantiate this proof.

3) The TS records $(x_S, \overline{x_S}, pk_{id})$ for this user, and marks this user as active. The user stores $(x_U, \overline{x_S}, pk_{id})$ on her device, and stores sk_{id} externally.

Obtain a Key-share Token. First, the user needs to randomize the ciphertext $\overline{x_S}$. However, it seems difficult to prove directly, for example in zero-knowledge, that the randomized ciphertext produced by the user is of the correct form. Therefore, we use a standard cut-and-choose approach [59], [54] to allow the TS to check that the encrypted key share it is blindly signing is correct with overwhelming probability. The user constructs $2k$ randomized ciphertexts $c_i = \overline{x_S} \cdot \mathbf{E}_{pk}^+(\delta_i; \kappa_i)$, and sends commitments C_i to them to the TS. The TS then asks the user to open a subset $\mathcal{D}$ of cardinality k, so that the TS can verify that these k ciphertexts were correctly formed. Having checked all opened ciphertexts, the TS blindly signs the remaining k ciphertexts. By nature of the cut-and-choose protocol at least one of the remaining ciphertexts is a correct randomization of $\overline{x_S}$ with high probability.

Let $\ell_\delta = \lceil \log p \rceil + \ell + \log k + 2$ be the bit-length of the randomizer such as δ. This size ensures that the k unopened $x_S + \delta_i$ values statistically hide x_S. Furthermore, we require that $N > 3 \cdot 2^{\ell_\delta}$ to ensure no overflows occur.

In our security proofs, see Section VI, we show that an adversary cannot learn anything useful about x_S despite seeing $\overline{x_S}$ and having access to the TS. We reduce to the CPA security of the homomorphic encryption scheme to show that an adversary cannot use $\overline{x_S}$ to learn something about x_S. However, in the reduction to CPA security, we cannot decrypt ciphertexts. Yet, in the GenShares protocol, the TS *must* decrypt a randomized version of $\overline{x_S}$ to recover the randomized key share. To allow us to correctly answer GenShares queries *without* decrypting ciphertexts, the user additionally creates a commitment Δ_i to δ_i and κ_i. In our proof of security, we use the extractability of ExtCommit$(\cdot, \cdot)$ to extract δ_i from these commitments, thus allowing us to answer GenShares queries without actually decrypting.

Using an additively homomorphic CCA2 secure encryption scheme would obviate the need for the extractable commitments Δ_i, simplifying the scheme. Unfortunately, to the best of our knowledge no additively homomorphic CCA2 secure

scheme exists. The RCCA scheme by Canetti et al. [60] is not homomorphic, the schemes by Prabhakaran and Rosulek [61] are multiplicatively homomorphic, and the fully homomorphic scheme by Lai et al. [62] is not CCA2 secure.

PROTOCOL 2. The ObtainKeyShareToken protocol is run between a user and the TS.

1) The user recovers $(x_U, \overline{x_S}, pk_{id})$ from storage, and authenticates to the TS. The TS aborts if this user exceeded the rate-limit for the current epoch, was banned, or was blocked. Otherwise, the TS looks up the user's record $(x_S, \overline{x_S}, pk_{id})$.

2) The TS randomly chooses a subset $\mathcal{D} \subset \{1, \ldots, 2k\}$ of cardinality k of indices of ciphertexts that it will check at step 5. The TS commits to $\mathcal{D}$ by picking $\theta \in_R \{0,1\}^{2\ell}$ and sending $\Delta = \mathsf{ExtCommit}(\mathcal{D}, \theta)$ to the user.

3) The user picks randomizers $\delta_1, \ldots, \delta_{2k} \in \{0,1\}^{\ell_\delta}$ to randomize the encrypted secret share $\overline{x_S}$, randomizers $\kappa_1, \ldots, \kappa_{2k} \in \mathcal{R}$ to create ciphertexts, and randomizers $r_1, \ldots, r_{2k} \in \mathbb{Z}_p$ and $\xi_1, \ldots, \xi_{2k} \in \{0,1\}^{2\ell}$ for the commitments and sets:

$$\begin{aligned} c_i &= \overline{x_S} \cdot \mathbf{E}_{pk}^+(\delta_i; \kappa_i) \\ C_i &= \mathsf{Commit}(H(c_i), r_i) \\ \Delta_i &= \mathsf{ExtCommit}((\delta_i, \kappa_i), \xi_i), \end{aligned} \tag{3}$$

for $i = 1, \ldots, 2k$. Finally, she sends the commitments $C_1, \ldots, C_{2k}$ and $\Delta_1, \ldots, \Delta_{2k}$ to the TS. Note that the commitments C_i and Δ_i are computationally binding and hiding.

4) The TS opens the commitment Δ by sending the subset $\mathcal{D}$ and the randomizer θ to the user. The user checks that $\Delta = \mathsf{ExtCommit}(\mathcal{D}, \theta)$, and aborts if the check fails.

5) The user opens the requested commitments by sending $(c_i, \delta_i, \kappa_i, r_i, \xi_i)_{i \in \mathcal{D}}$ to the TS. The TS checks that all disclosed values are constructed as per equation (3) and that $\delta_i < 2^{\ell_\delta}$. If any check fails, the TS bans the user.

6) Next, the user generates a token identifier $id \in_R \mathbb{Z}_p$ at random. Let $\mathcal{H} = \{i_1, \ldots, i_k\} = \{1, \ldots, 2k\} \setminus \mathcal{D}$ be the set of indices of unopened commitments. For the blind signature the user picks $r \in_R \mathbb{Z}_p$ and creates a commitment

$$C = \mathsf{Commit}((id, \epsilon, H(c_{i_1}), \ldots, H(c_{i_k})), r)$$

to the unopened ciphertexts, the epoch ϵ and id. Then, she encrypts the token identifier id as $\overline{id} = \mathbf{E}_{pk_{id}}(id)$, and sends C and $\overline{id}$ to the TS. Finally, she proves in zero-knowledge to the TS that $\overline{id}$ encrypts the token identifier id in C against her own public key pk_{id} and that C commits to the unopened ciphertexts, i.e.,

$$\begin{aligned} PK\{((c_i, r_i, \eta_i)_{i \in \mathcal{H}}, id, r) \; : \; \overline{id} = \mathbf{E}_{pk_{id}}(id) \wedge \\ \forall i \in \mathcal{H} \, [C_i = \mathsf{Commit}(\eta_i, r_i)] \wedge \\ C = \mathsf{Commit}((id, \epsilon, \eta_{i_1}, \ldots, \eta_{i_k}), r)\}, \end{aligned}$$

where $\eta_i = H(c_i)$. The TS checks this proof.

[2]http://world.std.com/~reinhold/diceware.html

7

7) If any check fails, the TS bans the user and aborts the protocol. If all checks pass, the TS runs a blind signature protocol with the user on the commitment C so that the user obtains a signature σ on the tuple $(id, \epsilon, H(c_{i_1}), \ldots, H(c_{i_k}))$. The user stores $\tau = (\sigma, \epsilon, id, (c_i, \kappa_i, \delta_i)_{i \in \mathcal{H}})$. The TS stores $\overline{id}$.

The following lemma states that even if a user is malicious, at least one of the ciphertexts c_i must be correctly formed. (See Appendix A for the proof.)

Lemma 1. *Consider a token $\tau = (\sigma, \epsilon, id, (c_i, \kappa_i, \delta_i)_{i=1,\ldots,k})$ obtained using the above protocol by a (potentially malicious) user U with corresponding TS key-share x_S. Let $\Delta_1, \ldots, \Delta_k$ be the corresponding set of commitments used during the obtain step. Then, with probability $1 - 1/\binom{2k}{k}$ there exists an index i^*, and randomizers $\delta^* < 2^{\ell_\delta}$, κ^*, and ξ^* such that:*

$$c_{i^*} = \mathbf{E}_{pk}^+(x_S + \delta^*; \kappa^*)$$
$$\Delta_{i^*} = ExtCommit((\delta^*, \kappa^*), \xi^*).$$

Using a Key-share Token. When using a token, the user sends the tuple $(id, \epsilon, c_{i_1}, \ldots, c_{i_k})$ and the signature σ to the TS, and provides an index j of the ciphertext c_j that the TS should decrypt. The TS uses the corresponding plaintext as the key in the threshold-cryptographic protocol. We know from Lemma 1 that at least one index i^* exists such that c_{i^*} is correctly formed. Key-share tokens resemble Chaum et al.'s e-cash tokens [54]. For the e-cash tokens it suffices if some indices are correct, in TANDEM, however, the user chooses the index j, and we must thus ensure that c_j in particular is correct. To enable the TS to check this, the user also reveals the differences $\gamma_i = \delta_j - \delta_i$ for all $i = 1, \ldots, k$. If these differences are correct then because c_{i^*} is a randomization of x_S, so must be c_j.

PROTOCOL 3. The GenShares protocol is run between an anonymous user and the TS.

1) The user a token $\tau = (\sigma, \epsilon, id, (c_i, \kappa_i, \delta_i)_{i=1,\ldots,k})$ as input and connects to the TS via an anonymous channel. She sends $(id, \epsilon, c_1, \ldots, c_k)$ and the blind signature σ.
2) Next, the user finds j such that $\delta_j \geq \delta_i$ for all i and computes $\gamma_i = \delta_j - \delta_i \geq 0$ and $\nu_i = \kappa_j \cdot \kappa_i^{-1}$ such that

$$c_j = c_i \cdot \mathbf{E}_{pk}^+(\gamma_i; \nu_i) \qquad (4)$$

for $i = 1, \ldots, k$. Finally she sends j, and $\gamma_1, \ldots, \gamma_k$, $\nu_1, \ldots, \nu_k$ to the TS.
3) The TS verifies that the γ_is and ν_is satisfy equation (4), that σ is a correct signature on $(id, \epsilon, H(c_1), \ldots, H(c_k))$, token id was not blocked, ϵ corresponds to the current epoch, and that $\gamma_i < 2^{\ell_\delta}$. The TS aborts if any check fails.
4) The TS decrypts c_j to compute $\tilde{x}_S = \mathbf{D}_{sk}^+(c_j) \pmod{p}$.
5) The user calculates her key share $\tilde{x}_U$ as:

$$\tilde{x}_U \equiv x_U - \delta_j \pmod{p}$$

Using Lemma 1, we can show that the decrypted element c_j must also be of the right form. (See Appendix A for the proof.)

Lemma 2. *If revealed token $(id, \epsilon, c_1, \ldots, c_k,)$ with j, $\gamma_1, \ldots, \gamma_k$ and $\nu_1, \ldots, \nu_k$ satisfies equation (4), then with probability $1 - 1/\binom{2k}{k}$ there exists $\delta < 2^{\ell_\delta+1}$ such that*

$$\mathbf{D}_{sk}^+(c_j) = x_S + \delta$$

where x_S is the TS key-share for the corresponding user.

The range proof in registration is essential. The range proof in equation (2) in the RegisterUser protocol ensures that the plaintext $x_S = \mathbf{D}_{sk}^+(\overline{x_S})$ is small compared to the randomizers δ_i. As a result, the randomized ciphertexts c_i statistically hide x_S. It is not sufficient to skip the range proof and instead choose the randomizers δ_i from the full plaintext domain $[N]$ to hide x_S. Without the range proof, the TS can construct tokens that it can later recognize by exploiting the fact that a large x_S results in a reduction modulo N. More precisely, the TS can set x_S of its target user somewhat large, so that $x_S + \delta_j > N$ (with a non-negligible probability). The user believes that the TS derives $x_S + \delta_j \pmod{p}$ (because she believes no modular reduction took place) and compensates accordingly. However, the TS actually derives $\tilde{x}_S = (x_S + \delta_j \mod N) \pmod{p} = x_S + \delta_j - (N \mod p)$. To test if the current token is from its target user, the TS adds $(N \mod p)$ to $\tilde{x}_S$. If the guess was correct, the TCP completes correctly, otherwise the protocol fails. This allows the TS to detect specific users.

Blocking the Key. To block her key, the user runs the BlockShare protocol with TS to ensure no new key-share tokens are created for her, and that all her unspent tokens are blocked.

PROTOCOL 4. The BlockShare protocol is run by a user and the TS. The user takes as input her long-term key sk_{id} (which she recorded outside her device). The user authenticates to the TS (possibly using sk_{id}). The TS marks the user as blocked, so that it will no longer issue new tokens. Then they continue as follows to invalidate unspent tokens. The TS sends a list of all encrypted token identifiers $\overline{id}_1, \ldots, \overline{id}_t$ that the user obtained in this epoch. The user looks up a list of all spent token identifiers (see below). The user then uses sk_{id} to decrypt $\overline{id}_1, \ldots, \overline{id}_t$ and sends the decrypted token identifiers that have not yet been spent to the TS. The TS will then block all tokens with these identifiers.

Since we assume the TS is honest with respect to blocking, the TS accurately provides the list of encrypted token identifiers. In the ObtainKeyShareToken protocol, the user verifiably encrypts the token identifier id. As a result, even if the user's device is corrupted, the TS stores a correct encryption $\overline{id}$ of id, so the above procedure blocks all unspent tokens.

In the unlikely case that a user cannot recover the identifiers, the attacker can nevertheless use the TS only a limited number of times, as the attacker is still subject to the rate-limit.

List of spent tokens. The TS is malicious with respect to privacy. So, it might try to trick the user into revealing the

identifiers of tokens she has already spent (thus revealing that these tokens were hers). In particular, the TS is not trusted to provide an accurate list of spent tokens. Therefore, we propose that users externally store spent token identifiers, so that they have a reliable record. Alternatively, the TS can keep a verifiable log of spent tokens by appending spent token identifiers to a public append-only log (users must then verify that each spent token identifier is in fact added to the log). Users then use this log as a record of spent tokens. Finally, if epochs are short, and users are willing to risk revealing their actions in the current epoch, they can also use a list provided by the TS. If the TS cheats, users reveal at most their actions within the most recent epoch when they block their keys.

B. Alternative constructions

An alternative method to construct tokens could be to use an authenticated encryption scheme that the user and the TS evaluate using secure multi-party computation [63]. The server inputs its key share x_S while the user inputs the randomizer δ. The user's output is the authenticated encryption of $x_S + \delta$ for the TS's symmetric key which serves as token. To ensure that the TS cannot recognize this token, the protocol should resist malicious servers and the circuit should validate the TS's input (i.e., that they are always the same). Even though recent secure two-party computation protocols that are secure against a malicious server boast impressive performance [64], they still require at least one order of magnitude more computational power as well as more bandwidth than our custom scheme.

Another simple alternative construction is to let users retrieve $\overline{x_S} = \mathsf{Enc}(x_S)$ using private information retrieval (PIR) via an anonymous channel—the user must still hide her identity. Then, users randomize $\overline{x_S}$ similarly to our construction, and the TS decrypts the ciphertext to recover $\overline{x_S} + \delta$, which it then uses in the TCP. To enable blocking of keys, the TS needs to frequently refresh its encryption keys, effectively invalidating previously retrieved ciphertexts $\overline{x_S}$. This simple protocol, however, has serious drawbacks. First, blocking is only enforced upon key refreshing, thus the timespan when compromised keys can be used depends on the refreshing schedule of the TS. Second, because the encryption of x_S for the current period can be randomized as often as the user wants (and the use of PIR precludes record-keeping), this scheme cannot provide rate-limiting. Third, because the TS acts as a decryption oracle for a homomorphic encryption scheme, which is only CPA secure, proving security in this setting requires very strong and non-standard assumptions.

VI. Security and Privacy of Tandem

In this section we formalize the security and privacy properties offered by TANDEM. We refer to the appendix for the complete security and privacy proofs.

A. Security of TANDEM

We capture the security of TANDEM using the following game. It models that if the user's key is compromised (e.g.,

her device is stolen), the user can block the use of her key, provided that the TANDEM server remains honest.

GAME 1. The TANDEM *security game* is between a challenger controlling the TS and the SP, and an adversary controlling up to n users. The adversary's goal is to complete a threshold-cryptographic protocol for a blocked user.

Setup phase The challenger sets up the TS and the SP. The challenger runs RegisterUser with the adversary for each of the n users the adversary controls.

Query phase During the query phase, the adversary can ask the TS to run the RegisterUser, ObtainKeyShareToken and BlockShare protocols with users controlled by the adversary. Moreover, the adversary can make RunTCP queries to the challenger. In response, the TS first runs the GenShares protocol with the user (controlled by the adversary), followed by a run of the TCP protocol.

Selection phase At some point the adversary outputs the identifier of a blocked user U^* on which it wants to be challenged later. To allow the challenger to confirm that all unspent tokens are blocked (to prevent trivial wins), the adversary also outputs the long term secret sk_{id} of user U^*. The challenger checks sk_{id} against the recorded public key pk_{id} and then blocks all tokens of user U^* using sk_{id}. The adversary loses if sk_{id} is not correct.

Second query phase The adversary can keep asking the TS to run RegisterUser, ObtainKeyShareToken, BlockShare protocols. The adversary can also make RunTCP queries as before (however, following the protocols the TS will not allow ObtainKeyShareToken queries of blocked users).

Challenge phase Finally, upon request of the adversary, the challenger acts as SP in the TCP protocol. At the same time, the adversary may still make queries and run protocols as in the previous phase. The adversary wins if it successfully completes the TCP with the SP on behalf of the blocked user U^*. To prevent trivial wins, this TCP protocol must be completable only by user U^*.[3]

In this game, all users are automatically corrupted right from the moment they start the registration protocol. This models the notion that users can even be blocked if the adversary is present right from the start, and also implies that honest users—which are only corrupted later—can still be blocked.

GAME 2. The TANDEM *rate-limiting* game is identical to the TANDEM security game, except that in the selection phase the adversary outputs a rate-limited user (i.e., a user who is not allowed to obtain more tokens in this epoch).

Of course, to have security using TANDEM, the TCP itself must be secure. Hence, we require that even if a malicious user has interacted many times with the TS, she cannot use her key

---

[3]One option is that the TCP protocol identifies the user. So for example, for the attribute-based credential TCPs, this means that the showing protocol must disclose an attribute that identifies U^*. Another option is to make sure that user U^* is the only user who can successfully complete the protocol, e.g., by revoking all other credentials.

when she does not have access to the TS. We formalize this using the following game.

GAME 3. The *TCP security game* is between a challenger controlling the TS and the SP and the adversary controlling a malicious user.

Setup phase During the setup phase, the adversary generates $x_U \in_R \mathbb{Z}_p$, whereas the TS, controlled by the challenger, generates $x_S \in_R \mathbb{Z}_p$.

Query phase In the query phase, the adversary can make $\mathsf{TCP}(\delta)$ queries to request that the TS runs $\mathsf{TCP.TS}(x_S + \delta)$ with the user. The adversary is responsible for running $\mathsf{TCP.U}$. Optionally, the adversary-controlled user can communicate with the challenger-controlled SP running $\mathsf{P}()$ as well.

Challenge phase In the challenge phase, the adversary is not allowed to make TCP queries. Instead, it interacts solely with the challenger-controlled SP running $\mathsf{P}()$. The adversary wins if the SP accepts.

Theorem 1. *No PPT adversary can win the* TANDEM *security game or the* TANDEM *rate-limiting game with non-negligible probability, provided that the TCP is secure (i.e., no PPT adversary can win the TCP security game).*

Proof sketch. We prove the security of the scheme by reducing it to the TCP security property. First, we show how to run $\mathsf{GenShares}$ without decrypting ciphertexts. During the $\mathsf{ObtainKeyShareToken}$ protocol, we model hash functions as random oracles to allow us to extract the token identifier id from the proof of knowledge in step 6, and the unopened randomizers δ_i from the extractable commitments Δ_i in step 3. Hence, during $\mathsf{GenShares}$ we can identify the user, and thus the corresponding key share x_S, as well as the randomizer δ_j (with overwhelming probability, using Lemma 2).

Knowing x_S and δ_j we no longer need to decrypt ciphertexts to run $\mathsf{GenShares}$, therefore, we can use the CPA security of the homomorphic encryption scheme to replace the initial ciphertext $\overline{x_S} = \mathbf{E}_{pk}^+(x_S)$ for the challenge user by $\overline{x_S} = \mathbf{E}_{pk}^+(0)$ an encryption of 0. During $\mathsf{GenShares}$ we add x_S to compensate. (To enable the reduction to CPA, we simulate the range proof in step 2 of $\mathsf{RegisterUser}$.)

Finally, we answer all queries for the challenge user using the TCP security oracle. Hence, a break of the TANDEM security game results in a break of the TCP security game. □

See Appendix C for the full proof.

B. Privacy of TANDEM

The following game models that TANDEM provides privacy for users. A malicious TANDEM server cannot distinguish between two honest users performing a transaction using the TANDEM server even if it colludes with the service provider, provided that the service provider on its own cannot distinguish transactions by these two honest users. The following privacy game asks the TS to recognize users for which it earlier issued a key-share token.

GAME 4. The TANDEM *privacy game with colluding SP* is between a challenger, who controls two honest users U_0 and U_1, and an adversary $\mathcal{A}$ who controls the TS and the SP.

Setup phase The adversary $\mathcal{A}$ outputs the number of key-share tokens n_T each honest user should obtain. The adversary is responsible for setting up the SP and the TS, i.e., it should publish a public key pk. Next, the honest users U_0 and U_1 interact with the adversary-controlled TS to obtain n_T key-share tokens each. First, U_0 runs $\mathsf{ObtainKeyShareToken}$ n_T times to obtain tokens $\tau_{0,1}, \ldots, \tau_{0,n_T}$. Then, U_1 runs the obtain protocol n_T times to obtain tokens $\tau_{1,1}, \ldots, \tau_{1,n_T}$.

Query phase During the query phase, the adversary can make $\mathsf{RunTCP}(U_i, j, \mathsf{in}_U)$ queries to request that user U_i uses token $\tau_{i,j}$ and then runs the TCP with input in_U. If $i \in \{0, 1\}$ and user U_i did not use token $\tau_{i,j}$ before, then user U_i, controlled by the challenger, first runs $\mathsf{GenShares}$ with the TS using token $\tau_{i,j}$ and then runs $\mathsf{TCP.U}(\mathsf{in}_U)$ with the TS and the SP (running $\mathsf{TCP.TS}$ and P respectively).

Challenge phase At some point, the adversary outputs a pair of token indices (i_0, i_1) for user U_0 and U_1 respectively on which it wants to be challenged. Let $\tau_0 = \tau_{0,i_0}$ and $\tau_1 = \tau_{1,i_1}$ be the corresponding tokens. The adversary loses if either token τ_0 or τ_1 has been used before. Then, the challenger picks a bit $b \in \{0, 1\}$ and proceeds as if the adversary first made a $\mathsf{RunTCP}(U_b, \tau_b)$ query and then a $\mathsf{RunTCP}(U_{1-b}, \tau_{1-b})$ query.

Guess phase The adversary outputs a guess b' of b. The adversary wins if $b' = b$.

The privacy game models the fact that there is *no time correlation* between when tokens are obtained by a user, and when they are spent by a user. At the same time, the adversary has full control over the TS and the SP, so this game also models the fact that the TS and the SP *can* correlate events that they see.

Since the SP is controlled by the adversary, the TCP must ensure privacy with respect to the SP and the TS, if all that the TS sees are randomized secret shares. We formalize this in the following game.

GAME 5. The *TCP privacy game with colluding SP* is between a challenger controlling two honest users U_0 and U_1 and an adversary $\mathcal{A}$, controlling the TS and the SP.

Setup The adversary publishes the TS public key and is responsible for setting up the SP. The challenger sets up its users. First, user U_0 generates $x_{0,U} \in_R \mathbb{Z}_p$ while the TS generates $x_{0,S} \in_R \mathbb{Z}_p$, then U_1 and TS similarly generate $x_{1,U}$ and $x_{1,S}$. Finally, the TS sends $x_{0,S}$ and $x_{1,S}$ to users U_0 and U_1 respectively.

Queries Adversary $\mathcal{A}$ can make $\mathsf{RunTCP}(i, \mathsf{in}_U)$ queries, to request U_i to run the TCP protocol using input in_U with the TS and the SP (both controlled by $\mathcal{A}$). In the first step, the user picks $\delta \in_R \mathbb{Z}_p$ and sends the randomized secret-share $\tilde{x}_S = x_{i,S} + \delta \pmod{p}$ to the TS. The user

then sets $\tilde{x}_U = x_{i,U} - \delta$ and runs $\mathsf{TCP.U}(\tilde{x}_U, \mathsf{in}_U)$ with the TS and the SP running $\mathsf{TCP.TS}$ and P respectively.

Challenge Adversary $\mathcal{A}$ outputs an input in_U. Challenger picks a bit $b \in_R \{0, 1\}$. Then the challenger acts as if $\mathcal{A}$ first made a $\mathsf{RunTCP}(b, \mathsf{in}_U)$ query, and then a $\mathsf{RunTCP}(1 - b, \mathsf{in}_U)$ query.

Guess The adversary outputs a guess b' for b, $\mathcal{A}$ wins if $b = b'$.

TANDEM also offers privacy for users against the TS alone. That is, even if the service provider can identify users, the TS cannot observe their key-usage patterns as long it does not collude with the service provider. (If the SP can identify users, then so can a coalition of the TS and the SP, so we exclude this case to prevent a trivial win.) We model this situation as a variant of the previous two games.

DEFINITION 2. The TANDEM *privacy game with honest SP* and the *TCP privacy game with honest SP* are as in Game 4 and Game 5 above, however, the challenger controls the SP. The adversary can interact with the SP as a normal user.

Theorem 2. *No PPT adversary can win the* TANDEM *privacy game with colluding SP (respectively the* TANDEM *privacy game with honest SP) with probability non-negligibly better than* $1/2$*, provided that the* TCP *is privacy-friendly (i.e., no PPT adversary can win the TCP privacy game with colluding SP respectively the TCP privacy game with honest SP).*

Proof sketch. We first argue that we can remove all identifying information from the key-share tokens of the challenge users. First, we extract the server's key-shares $x_{0,S}$ and $x_{1,S}$ from the proof of knowledge in step 2 of the RegisterUser protocol. Then we simulate the proof of knowledge in step 6 of ObtainKeyShareToken, replace the ciphertext $\overline{id}$ by the encryption of zero (using the CPA security of the ElGamal encryption scheme), extract the subset $\mathcal{D}$ so that we can send random commitments C_i, Δ_i for $i \notin \mathcal{D}$ (because the commitment schemes are computationally hiding), and finally, we set the unrevealed ciphertexts $c_i = \mathbf{E}_{pk}^+(\delta_i, \kappa_i)$ for $i \notin \mathcal{D}$. None of these changes are detectable by the adversary during ObtainKeyShareToken.

To use tokens as requested by the adversary, the challenge users add $x_{i,S}$ to their long-term share x_U to compensate for the changes made so that GenShares completes successfully. Moreover, because the randomizers δ_i statistically hide x_S, the adversary cannot detect the final change to the ciphertexts during GenShares.

Therefore, by the blindness of the signature scheme, we can swap tokens between the users and still simulate protocols perfectly. Therefore, any adversary that can then still distinguish users must break the security of the TCP privacy game (with a colluding SP or with a honest SP). To extract secrets from the proofs and commitments, and to make the final reduction to TCP privacy, we model hash functions as random oracles. □

See Appendix D for the full proof.

VII. SECURING PROTOCOLS WITH TANDEM

Recall that to use TANDEM to protect the private key in a cryptographic scheme we must convert the protocols into linearly randomizable threshold-cryptographic protocols.

For this composition to be secure, the threshold-cryptographic protocols must satisfy the natural security definition (see Game 3). For this composition to be private, i.e., so that the TANDEM server alone respectively by colluding with the SP cannot identify the user, the threshold-cryptographic protocols must additionally be private (see Game 5).

Many traditional threshold-cryptographic schemes already satisfy these requirements. Threshold variants of Schnorr [14] and RSA signatures [15] as well as ElGamal-based [16], [17] and RSA encryption [15] schemes rely on Shamir secret-sharing and are thus linearly randomizable. Moreover, the threshold protocols are private, i.e., the server-side protocols for signing and decrypting, respectively, operate solely on the secret-share and the common input, the message or ciphertext.

Threshold-cryptographic versions of electronic cash schemes [18], [19] and attribute-based credential (ABC) schemes [20], [21], [22], [23] can also be used with TANDEM. For some of these, the threshold-cryptographic versions already exist [21]. For the others, the threshold-cryptographic versions of the zero-knowledge proofs on which these schemes are based must be created. As an example, we now show how convert the BBS+ ABC scheme [20] into a TANDEM-suitable threshold-cryptographic scheme.

A. The use-case of ABCs

Attribute-based credentials can be conceptualized as digital equivalents to classic documents like passports, driver's license, student cards, etc. The owner of a credential can selectively disclose any subset of attributes to a service provider in such a way that the the validity of the disclosed attributes can be validated. In many ABC systems credentials are unlinkable, that is, users are anonymous within the set of users having the same disclosed attributes.

To bind credentials to a user, and to ensure that only the owner can operate with them, credentials contain the user's secret key. Typically, all credentials of a user contain the same secret key. When credentials are stored on insecure platforms such as smart phones or personal computers TANDEM can be used to strengthen the security of the secret key. This ensures that valuable credentials cannot be abused, and can be blocked, while preserving users' privacy.

To use ABCs with TANDEM we need to convert the protocols for issuing and verifying credentials into threshold-cryptographic alternatives that are secure, private, and linearly randomizable. During issuance, the issuer (taking the place of the service provider in Section III) provides the user with a credential bound to the user's secret key. The issuer does not learn the user's secret key. During verification, a user authenticates to a service provider by selectively disclosing attributes from her ABCs.

In typical ABC schemes, these two protocols rely heavily on zero-knowledge proofs over the user's secret key. In the

TS	User	Issuer
$\tilde{x}_S \in \mathbb{Z}_p, B_0$	$\tilde{x}_U \in \mathbb{Z}_p, s', g, B, B_0, B_1$	sk_I

TS		User		Issuer
$\hat{x}_S \in_R \mathbb{Z}_p$		$\hat{x}_U, \hat{s}' \in_R \mathbb{Z}_p$		
$u_S = B_0^{\hat{x}_S}$	$\xrightarrow{\quad u_S \quad}$	$u_U = B_0^{\hat{x}_U}$		
		$\tilde{U} = B^{\hat{s}'} \cdot u_U \cdot u_S$	$\xrightarrow{\quad \tilde{U} \quad}$	
	$\xleftarrow{\quad c \quad}$		$\xleftarrow{\quad c \quad}$	$c \in_R \mathbb{Z}_p$
$r_S = \hat{x}_S + c \cdot \tilde{x}_S$	$\xrightarrow{\quad r_S \quad}$	$r_{s'} = \hat{s}' + c \cdot s'$		
		$r_U = \hat{x}_U + c \cdot \tilde{x}_U$		
		$r = r_U + r_S$	$\xrightarrow{\quad r_{s'}, r \quad}$	$\tilde{U} \overset{?}{=} U^{-c} \cdot B^{r_{s'}} \cdot B_0^r$

Fig. 2: Full details of the proof of knowledge of the user's commitment $U = B^{s'} B_0^{x_U} B_0^{x_S}$ in the BBS+ TCP issuance protocol. The TANDEM server only knows $\tilde{x}_S$ and the user knows $\tilde{x}_U$ and the randomness s' (recall $\tilde{x}_S$ and $\tilde{x}_U$ are the respective outputs of the GenShares protocol). The TS effectively creates a zero-knowledge proof of knowing $\tilde{x}_S$.

remainder of this section, we show how these non-threshold protocols for BBS+ credentials [20] can be converted to threshold-cryptographic versions suitable for TANDEM.

BBS+ credentials are anonymous credentials built from BBS+ signatures [20]. BBS+ signatures operate in a pairing setting and rely on discrete-logarithm based assumptions. Let $(\mathbb{G}_1, \mathbb{G}_2)$ be a bilinear group pair, both of prime order p, generated by g and h respectively. The pairing is given by $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ where $\mathbb{G}_T$, also of order p, is generated by $\hat{e}(g, h)$. Let l be the number of attributes. In the BBS+ credential scheme, an issuer randomly chooses generators $B, B_0, .., B_l \in_R \mathbb{G}_1$, picks a private key $sk_I \in_R \mathbb{Z}_p$, and computes $w = h^{sk_I}$. The issuer's public key is $pk_I = (w, B, B_0, .., B_l)$.

Obtaining a credential. Attribute-based credentials contain the user's secret key as an attribute. For simplicity, we describe the TANDEM BBS+ issuance and showing protocols below with two attributes: the secret key x and an issuer-determined attribute a_1. To obtain a credential, the user (and the TANDEM server) run the following TCP version of the issuance protocol with the issuer. The issuance protocol is run jointly by the user, TS, and an issuer. Let $\tilde{x}_U$ and $\tilde{x}_S$ be the two shares of the user's secret key $x = \tilde{x}_U + \tilde{x}_S$ that are held by the user and the TS respectively after running GenShares. The user first commits to her secret key x, to allow the issuer to blindly sign it. As we share the user's secret key between the user and the TS, they both have to participate in creating the commitment. First, the user sends B_0 to the TS so that it can compute $B_0^{\tilde{x}_S}$ before sending back to the user. Then the user and the TS create a commitment $U = B^{s'} B_0^{\tilde{x}_U} B_0^{\tilde{x}_S} = B^{s'} B_0^x$ where $s' \in_R \mathbb{Z}_p$. To prove to the issuer that U is well-formed, the user and the TS construct the proof

$$PK\{(x, s') : U = B^{s'} B_0^x\}. \tag{5}$$

Fig. 2 shows how to construct this proof. If this proof of knowledge verifies, the issuer randomly generates $s'', e \in_R \mathbb{Z}_p$ and calculates

$$A = \left(gB^{s''} U B_1^{a_1}\right)^{\frac{1}{e+sk_I}} \quad \in \mathbb{G}_1$$

and the tuple (A, e, s'') to the user. The user calculates $s = s' + s''$ and stores the credential $\sigma = (A, e, s)$.

Showing a credential. After the issuance protocol, the user can show the credential to a service provider to get access to a service or a resource. Again we convert the showing protocol into a TCP that uses the TANDEM server.

In the showing protocol, the user proves the possession of a credential $\sigma = (A, e, s)$ over her key x and the attribute a_1. To show she possesses such a credential, while hiding her key and disclosing her attribute, she proves in zero-knowledge that

$$\hat{e}(A, h^e w) = \hat{e}(gB^s B_0^x B_1^{a_1}, h). \tag{6}$$

To prove the validity of this equation in zero-knowledge, without revealing any of the values A, e, s (that would make the user linkable), we follow the approach by Au et al. [20]. Let g_1, g_2 be two extra generators in $\mathbb{G}_1$. First, the user creates a commitment $C_1 = A g_2^{r_1}$ to A, where r_1 is a randomizer chosen from $\mathbb{Z}_p$ by the user. The user then commits to her randomizer as well using $C_2 = g_1^{r_1} g_2^{r_2}$ where $r_2 \in_R \mathbb{Z}_p$. The user sends these commitments to the service provider. These commitments perfectly hide the value of A. Finally, she and the TS engage in the following zero-knowledge proof with the service provider:

$$PK\Big\{(r_1, r_2, \alpha_1, \alpha_2, e, x, s) : C_2 = g_1^{r_1} g_2^{r_2} \wedge C_2^e = g_1^{\alpha_1} g_2^{\alpha_2} \wedge$$
$$\hat{e}(C_1, w) \cdot \hat{e}(C_1, h)^e = \hat{e}(g, h)\hat{e}(B, h)^s \hat{\mathbf{e}}(\mathbf{B_0}, \mathbf{h})^x \cdot$$
$$\hat{e}(B_1, h)^{a_1} \hat{e}(g_2, w)^{r_1} \hat{e}(g_2, h)^{\alpha_1}\Big\}$$

to prove that she indeed posseses the signature over the hidden and the disclosed attributes and that equation (6) is satisfied. In the proof, $\alpha_1 = er_1$ and $\alpha_2 = er_2$. The user can easily generate the proofs for the first two conjuncts. Only the third conjuct contains the user's secret key x of which the user only has a secret share. Thus, the user has to contact the TS to construct this part of the proof. This proof is just a proof of representation, like in equation 5, albeit a bit more complex. As a result, a very similar construction as in Fig. 2 allows the user and the TS to jointly compute this proof.

Security and privacy of the TCPs. These TCPs satisfy the TCP security and privacy notions defined in Section VI. For security (see Game 3), note that the TS computes zero-knowledge proofs of knowing $\tilde{x}_S$. A malicious user learns nothing about $\tilde{x}_S$ (thus nor x_S) as a result of the zero-knowledge property. Hence, the TCP showing and issuance protocols satisfy the TCP security property.

For privacy (see Game 5), the TS operates on a fully randomized key $\tilde{x}_S$, so the TS cannot distinguish users based on the key if the SP is honest. The indistinguishability property of the credential scheme guarantees that the TS cannot distinguish users based on the resulting showing proof by colluding with the SP either. Thus, the TCP showing protocol is private for both honest and colluding SPs.

B. Rate-limiting in ABCs

Anonymous users can use the cover of privacy to misbehave, negatively impacting the system. ABC systems are not exempt from such misbehavior. Suppose, for example, that a user shares her "I am older than 18" credential with many under-aged users who do not hold such a credential. Then, those under-aged users can incorrectly convince service providers that they are over 18 years of age. If this happens often, service providers can no longer rely on these credentials to verify that a user is older than 18.

To limit such misbehavior, ABCs could benefit from rate-limiting. One method to limit abuse is to rate-limit credentials by ensuring that credentials can only be used a limited number of times. For instance, solutions such as n-times anonymous credentials [28] use custom cryptographic techniques to construct a special type of ABC that can be used only a limited number of times.

TANDEM can achieve a similar type of rate-limiting *without* modifying the underlying cryptographic construction of ABCs. To rate-limit use of a system, the TS enforces a per-user and per-epoch limit q on the number of tokens it issues per user and per epoch. As a result, no credential can be shown more than q times per epoch. This approach limits *all* credentials associated to a user's key. If desired, TANDEM can equally be applied on a per-credential basis.

This rate-limiting strategy *requires* that all users use TANDEM. However, recall that the SPs (issuers and verifiers) cannot detect the use of TANDEM, allowing users to forego sharing their keys with the TS, thus avoiding the rate limit. To enable the TS to enforce a rate-limit on all credentials, issuers must only issue credentials on keys that are shared with the TS.

A small change to the threshold-cryptographic version of the issuance protocol enables the issuer to confirm that users use TANDEM. To signal its involvement, the TS signs its proof (u_S, c, r_S) and sends the signature σ to the user. The user forwards the messages u_S, r_S and σ from the TS to the issuer together with its own messages u_U, $r_{s'}$ and r_U. The issuer, rather than the user, combines the proofs and verifies them. Moreover, the issuer checks the signature σ. If the signature

and proofs are correct, then the user's key was shared with the TS and the issuer signs the credential.

VIII. PERFORMANCE EVALUATION

In this section, we evaluate TANDEM's computational and bandwidth cost. We use an attribute-based credential instantiation as a case study and compare the performance when using no key protection, traditional threshold-cryptographic schemes, and TANDEM.

TANDEM consists of four protocols: RegisterUser, ObtainKeyShareToken, GenShares, and BlockShare. We implemented in C the time-critical protocols, ObtainKeyShareToken and GenShares.[4] We used Pedersen commitments [58] as commitment scheme, and BBS+ credentials [20] to construct the blind signature. We use the RELIC cryptographic library to implement them [65].[5] We use a recent implementation [66] of Joye and Libert's additive homomorphic encryption scheme [55] for our protocols. We set the modulus size to 2048 bits and the plaintext space to 394 bits, such that $N > 3 \cdot 2^{\ell_\delta}$ for $k \leq 85$. With this setting, encrypting a single 394 bits plaintext takes 0.9 ms whereas it takes 24.2 ms to decrypt a ciphertext. We also experimented with an optimized implementation[6] of Paillier's encryption scheme [56], but our experiments show that Joye and Libert's scheme gives better performance. Finally, we use ElGamal's encryption scheme [57] to encrypt token identifiers.

We empirically measure performance on a single core of an Intel i7-7700 running at 3.6 GHz. Smart phones and tablets generally have slower processors. Yet, we believe that given our measurements, TANDEM's performance would be practical on these devices.

Obtaining a token. We first justify our choices for the value of the security parameters k in our experiments. Our security analysis shows that an attacker can break TANDEM's security property by constructing a key-share token for a blocked user with probability $\binom{2k}{k}^{-1}$. Hence, $k = 42$ gives 80 bits of security, and $k = 66$ gives 128 bits security. However, ObtainKeyShareToken is an *interactive* protocol. The success probability of an attacker is limited by how often the TS lets the attacker try to construct a malicious token rather than by the adversary's computational power. Because the TS bans users trying to construct malicious tokens, one can choose a smaller k in practice. In a system with 100 000 users, $k = 20$ ensures that the probability that an attacker (corrupting all users) can at least *once* use any blocked key is less than 10^{-6}.

Fig. 3 shows the computing time (without communication) for the ObtainKeyShareToken protocol at the user (black) and server (blue) for different values of the parameter k. The homomorphic encryption scheme—creating the ciphertexts (user), and checking a subset of these (TS)—dominates the computational cost. Our experiments reveal that the timing

---

[4]The code will be available upon publication.

[5]We set up RELIC to use a BLS curve over a 381 bits field. This setup ensures 128 bits security, while the group order remains 255 bits.

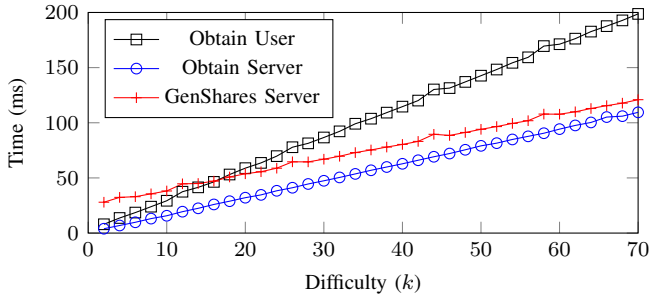[6]https://github.com/mcornejo/libpaillier

Fig. 3: ObtainKeyShareToken protocol computing time at the user (black) and the server side (blue), and GenShares protocol computation time at the server side (red) for increasing difficulty levels k.

TABLE IV: Comparison of computational cost and properties between not using a TCP, using a traditional TCP, and using a TCP with TANDEM ($k = 20$).

	No TCP	Traditional TCP	TCP + TANDEM
Obtain Token			
User	-	-	59 ms
Server	-	-	32 ms
Run Protocol			
User	5 ms	5 ms	5 + 4 ms
Server	-	1 ms	1 + 54 ms
Key blocking	✗	✓	✓
Rate limiting	✗	✓	✓
Privacy	✓	✗	✓

variance accross executions is negligible. The bandwidth cost is low. For a security level of $k = 20$, the user sends about 26 KiB and receives less than 200 bytes.

Using the key. On the user side running GenShares is very cheap. Even for $k = 60$ the user requires less than 5 ms. In terms of bandwidth the user just needs to send 12 KiB for $k = 20$ and 36 KiB for $k = 60$. We show the server's computational cost for recovering the TS key-share from the token in Fig. 3. For a reasonable security level of $k = 20$, the server computational overhead is around 50 ms. The sending of the token in the GenShares protocol can be combined with the request to start the TCP resulting in no extra latency on top of the delay incurred by the Tor network [53] (1–2 s to send a small and receive amount of data on a fresh circuit[7]). Note that the circuit creation and GenShares can be run preemptively, thereby reducing the user-perceived delay.

Given the above measurements, a modern 4-core server can participate in approximately 50 TCPs per second (not counting the cost of the application-dependent TCP itself), i.e., serve 3 000 users per minute, requiring about 15 Mbit/s incoming bandwidth.

RegisterUser and BlockShare. These protocols are run few times (only upon registration and for blocking) and are thus not critical for scalability. We estimate the cost for RegisterUser to be well below a second for both the user and the TS (given its similarity with the ObtainKeyShareToken and GenShares protocols, and that the cost of the range proof is around 500 ms). In BlockShare the ElGamal decryption the token identifiers dominates the run-time cost, we estimate it to be in the order of seconds for thousands of tokens.

Comparison. Table IV compares the computational cost of creating a single BBS+ disclosure proof with 5 hidden attributes without key protection, using a traditional TCP, and using a TANDEM-augmented TCP configured with $k = 20$.

Without a TCP, the credential showing is very fast and, as there is no party involved in the use of the key, the showing of the credential is perfectly private. However, it is not possible to perform key blocking nor limit the key-usage unless specific credentials (e.g. [28]) are used. When introducing a traditional TCP, the overhead is minimal (only 1 ms at the server side) and key blocking and rate limiting are possible, at the cost of privacy. TANDEM provides the three properties. Without taking into account the ObtainKeyShareToken operation that happens offline, the user's overhead is negligible (4 ms), and well below a second (54 ms) for the server. In all cases, the cost of Tandem's cryptographic operations are very small compared to Tor's network cost.

IX. CONCLUSION

Protecting cryptographic keys is imperative to maintain the security of cryptographic protocols. As users' devices are most of the time insecure, the community has turned to threshold-cryptographic protocols to strengthen the security of keys. When run with a central server, however, these protocols raise privacy concerns. In this paper, we have proposed TANDEM, a provably secure scheme that, when composed with threshold-cryptographic protocols, provides privacy-preserving access to the keys. TANDEM also enables users to block her keys and rate-limit their usage, in ways that previous work could not handle. Our proof-of-concept implementation of TANDEM shows that for reasonable security parameters TANDEM's protocols run in less than 60 ms, hence being suitable for use in practice.

TANDEM is particularly suited for privacy-friendly applications such as eCash and ABCs because it retains their inherent privacy properties. Yet, TANDEM can be used to strengthen a wide variety of primitives, including signature and encryption schemes, as long as they can be transformed into linearly-randomizable threshold protocols. Using attribute-based credentials we have shown that deriving such a threshold protocol can be done with standard techniques, and that thereafter adding TANDEM is straightforward.

REFERENCES

[1] A. Hern, "Stagefright: new Android vulnerability dubbed 'heartbleed for mobile'," *The Guardian*, 2015. [Online]. Available: https://www.theguardian.com/technology/2015/jul/28/stagefright-android-vulnerability-heartbleed-mobile

[2] V. van der Veen, Y. Fratantonio, M. Lindorfer, D. Gruss, C. Maurice, G. Vigna, H. Bos, K. Razavi, and C. Giuffrida, "Drammer: Deterministic Rowhammer Attacks on Mobile Platforms," in *CCS*, 2016.

[7]As reported by https://metrics.torproject.org/torperf.html, visited July 6, 2018.

[3] Kim Zetter, WIRED magazine, "How the top 5 PC makers open your laptop to hackers," https://www.wired.com/2016/05/2036876/, 2016.

[4] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown," *ArXiv e-prints*, Jan 2018.

[5] J. Ekberg, K. Kostiainen, and N. Asokan, "The Untapped Potential of Trusted Execution Environments on Mobile Devices," in *S&P*, 2014.

[6] C. Marforio, N. Karapanos, C. Soriente, K. Kostiainen, and S. Capkun, "Secure enrollment and practical migration for mobile trusted execution environments," in *SPSM*, 2013.

[7] R. S. Sandhu and X. Zhang, "Peer-to-peer access control architecture using trusted computing technology," in *SACMAT*, 2005.

[8] B. McGillion, T. Dettenborn, T. Nyman, and N. Asokan, "Open-TEE - An Open Virtual Trusted Execution Environment," in *TrustCom*, 2015.

[9] Android security website, "Developing third party applications with Trusty TEE," https://source.android.com/security/trusty/#third-party_trusty_applications, 2017.

[10] S. Das, A. Dingman, and L. J. Camp, "Why Johnny Doesn't Use Two Factor A Two-Phase Usability Study of the FIDO U2FSecurity Key," in *FC*, 2018.

[11] L. T. A. N. Brandão, N. Christin, G. Danezis, and anonymous, "Toward mending two nation-scale brokered identification systems," *PoPETs*, 2015.

[12] H. A. Jawaheri, M. A. Sabah, Y. Boshmaf, and A. Erbad, "When A small leak sinks A great ship: Deanonymizing tor hidden service users through bitcoin transactions analysis," *CoRR*, vol. abs/1801.07501, 2018.

[13] S. Goldfeder, H. A. Kalodner, D. Reisman, and A. Narayanan, "When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies," *PoPETs*, 2018.

[14] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, "Secure Distributed Key Generation for Discrete-Log Based Cryptosystems," *J. Cryptology*, 2007.

[15] V. Shoup, "Practical threshold signatures," in *EUROCRYPT*, 2000.

[16] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," in *CRYPTO*, 1984.

[17] V. Shoup and R. Gennaro, "Securing Threshold Cryptosystems against Chosen Ciphertext Attack," *J. Cryptology*, 2002.

[18] J. Camenisch, S. Hohenberger, and A. Lysyanskaya, "Compact e-cash," in *EUROCRYPT*, 2005.

[19] I. Miers, C. Garman, M. Green, and A. D. Rubin, "Zerocoin: Anonymous distributed e-cash from bitcoin," in *S&P*, 2013.

[20] M. H. Au, W. Susilo, and Y. Mu, "Constant-Size Dynamic k-TAA," in *SCN*, 2006.

[21] S. A. Brands, *Rethinking public key infrastructures and digital certificates: building in privacy*. MIT Press, 2000.

[22] J. Camenisch and E. V. Herreweghen, "Design and Implementation of the *Idemix* Anonymous Credential System," in *CCS*, 2002.

[23] J. Camenisch and A. Lysyanskaya, "A Signature Scheme with Efficient Protocols," in *SCN*, 2002.

[24] P. D. MacKenzie and M. K. Reiter, "Two-party generation of DSA signatures," *Int. J. Inf. Sec.*, 2004.

[25] R. Gennaro, S. Goldfeder, and A. Narayanan, "Threshold-optimal DSA/ECDSA signatures and an application to bitcoin wallet security," in *ACNS*, 2016.

[26] M. H. Au, P. P. Tsang, and A. Kapadia, "PEREA: Practical TTP-free Revocation of Repeatedly Misbehaving Anonymous Users," *TISSEC*, 2011.

[27] P. P. Tsang, M. H. Au, A. Kapadia, and S. W. Smith, "BLAC: Revoking Repeatedly Misbehaving Anonymous Users without Relying on TTPs," *TISSEC*, 2010.

[28] J. Camenisch, S. Hohenberger, M. Kohlweiss, A. Lysyanskaya, and M. Meyerovich, "How to Win the Clone Wars: Efficient Periodic n-times Anonymous Authentication," in *CCS*, 2006.

[29] C. J. Mitchell, "What is trusted computing?" in *Trusted Computing*, 2005, vol. 6.

[30] Y. Desmedt, "Society and Group Oriented Cryptography: A New Concept," in *CRYPTO*, 1987.

[31] C. Boyd, "Digital multisignatures," *Cryptography and Coding*, 1989.

[32] Y. Desmedt and Y. Frankel, "Shared Generation of Authenticators and Signatures (Extended Abstract)," in *CRYPTO*, 1991.

[33] T. Rabin, "A Simplified Approach to Threshold and Proactive RSA," in *CRYPTO*, 1998.

[34] R. Gennaro, T. Rabin, S. Jarecki, and H. Krawczyk, "Robust and Efficient Sharing of RSA Functions," *J. of Cryptology*, 2000.

[35] J. F. Almansa, I. Damgård, and J. B. Nielsen, "Simplified Threshold-RSA with Adaptive and Proactive Security," in *EUROCRYPT*, 2006.

[36] R. Peeters, S. Nikova, and B. Preneel, "Practical RSA threshold decryption for things that think," in *WISSec 2008*, 2008.

[37] C. Hazay, G. L. Mikkelsen, T. Rabin, and T. Toft, "Efficient RSA Key Generation and Threshold Paillier in the Two-Party Setting," in *CT-RSA*, 2012.

[38] E. Atwater and U. Hengartner, "Shatter: Using Threshold Cryptography to Protect Single Users with Multiple Devices," in *WISEC*, 2016.

[39] M. Keller, G. L. Mikkelsen, and A. Rupp, "Efficient Threshold Zero-Knowledge with Applications to User-Centric Protocols," in *ICITS*, 2012.

[40] P. D. MacKenzie and M. K. Reiter, "Networked Cryptographic Devices Resilient to Capture," in *S&P*, 2001.

[41] J. Camenisch, A. Lehmann, G. Neven, and K. Samelin, "Virtual Smart Cards: How to Sign with a Password and a Server," in *SCN*, 2016.

[42] D. Boneh, X. Ding, G. Tsudik, and C. Wong, "A Method for Fast Revocation of Public Key Certificates and Security Capabilities," in *USENIX*, 2001.

[43] D. Boneh, X. Ding, and G. Tsudik, "Fine-grained control of security capabilities," *TOIT*, 2004.

[44] B. Libert and J. Quisquater, "Efficient revocation and threshold pairing-based cryptosystems," in *PODC*, 2003.

[45] A. Buldas, A. Jürgenson, A. Kalu, and M. Oruaas, "Server-supported RSA signatures for mobile devices," in *ESORICS*, 2017.

[46] A. Everspaugh, R. Chatterjee, S. Scott, A. Juels, and T. Ristenpart, "The pythia PRF service," in *USENIX*, 2015.

[47] J. Camenisch, A. Lehmann, and G. Neven, "Optimal distributed password verification," in *CCS*, 2015.

[48] T. G. Abbott, K. J. Lai, M. R. Lieberman, and E. C. Price, "Browser-Based Attacks on Tor," in *PETS*, 2007.

[49] L. Øverlier and P. F. Syverson, "Locating Hidden Servers," in *S&P*, 2006.

[50] E. Balsa, C. Troncoso, and C. Díaz, "OB-PWS: Obfuscation-Based Private Web Search," in *S&P*, 2012.

[51] R. Chow and P. Golle, "Faking Contextual Data for Fun, Profit, and Privacy," in *WPES*, 2009.

[52] A. M. Piotrowska, J. Hayes, T. Elahi, S. Meiser, and G. Danezis, "The Loopix Anonymity System," in *USENIX*, 2017.

[53] R. Dingledine, N. Mathewson, and P. F. Syverson, "Tor: The second-generation onion router," in *USENIX*, 2004.

[54] D. Chaum, A. Fiat, and M. Naor, "Untraceable electronic cash," in *CRYPTO '88*, 1988.

[55] M. Joye and B. Libert, "Efficient Cryptosystems from 2^k-th Power Residue Symbols," in *EUROCRYPT*, 2013.

[56] P. Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," in *EUROCRYPT*, 1999.

[57] T. ElGamal, "A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms," *Trans. Inf. Theory*, vol. 31, no. 4, 1985.

[58] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *CRYPTO*, 1991.

[59] G. Brassard, D. Chaum, and C. Crépeau, "Minimum disclosure proofs of knowledge," *J. Comput. Syst. Sci.*, 1988.

[60] R. Canetti, H. Krawczyk, and J. B. Nielsen, "Relaxing chosen-ciphertext security," in *CRYPTO*, 2003.

[61] M. Prabhakaran and M. Rosulek, "Homomorphic encryption with CCA security," in *ICALP*, 2008.

[62] J. Lai, R. H. Deng, C. Ma, K. Sakurai, and J. Weng, "CCA-Secure Keyed-Fully Homomorphic Encryption," in *PKC*, 2016.

[63] A. C. Yao, "How to Generate and Exchange Secrets (Extended Abstract)," in *FOCS*, 1986.

[64] X. Wang, S. Ranellucci, and J. Katz, "Authenticated garbling and efficient maliciously secure two-party computation," in *SIGSAC*, 2017.

[65] D. F. Aranha and C. P. L. Gouvêa, "RELIC is an Efficient Library for Cryptography," https://github.com/relic-toolkit/relic.

[66] M. Barbosa, D. Catalano, and D. Fiore, "Labeled Homomorphic Encryption: Scalable and Privacy-Preserving Processing of Outsourced Data," in *ESORICS*, 2017.

[67] M. Bellare and S. Goldwasser, "Verifiable partial key escrow," in *CCS*, 1997.

[68] A. Juels, M. Luby, and R. Ostrovsky, "Security of blind digital signatures (extended abstract)," in *CRYPTO*, 1997.

[69] D. Schröder and D. Unruh, "Security of blind signatures revisited," *J. of Cryptol.*, 2017.

APPENDIX A
PROOFS OF LEMMAS

Proof of Lemma 1. Whenever a ciphertext c_i is selected by the TS for opening, the TS checks that it and the corresponding randomizers κ_i, δ_i, ξ_i, and r_i are as in equation (3) and that $\delta_i < 2^{\ell_\delta}$, and hence as stated in the theorem.

Since the TS checks k tuples, every adversary needs to include at least k correct tuples in its set of $2k$ tuples. If no index i^* exists for the remaining tuples, then all k of them were incorrectly formed. The probability that none of these k bad tuples were selected during the cut and choose protocol is $1/\binom{2k}{k}$. $\qquad\square$

Proof of Lemma 2. From Lemma 1 we know that with probability $1 - 1/\binom{2k}{k}$ there exists i^* and δ^*, x_S such that

$$\mathbf{D}^+_{sk}(c_{i^*}) = x_S + \delta^*$$

Let $c_j = \mathbf{E}^+_{pk}(\alpha)$. From equation (4) we know that:

$$c_j = c_{i^*} \cdot \mathbf{E}^+_{pk}(\gamma_{i^*}; \kappa_{i^*})$$

By decrypting we find that $\alpha = x_S + \delta^* + \gamma_{i^*} \pmod{N}$. Moreover, $\delta^* < 2^{\ell_\delta}$ (by Lemma 1), $x_S < p < 2^{\ell_\delta}$ (by construction) and $\gamma_{i^*} < 2^{\ell_\delta}$ as checked by the TS. Since $\ell_\delta = \lceil \log p \rceil + \ell + \log k + 2$ and $N > 3 \cdot 2^{\ell_\delta}$, we have that $\alpha = x_S + \delta^* + \gamma_{i^*}$ as integers, and thus c_j is a proper randomization, with randomizer $\delta^* + \gamma_{i^*} < 2^{\ell_\delta+1}$, of x_S as well. $\qquad\square$

APPENDIX B
CONSTRUCTING CORRECTNESS PROOF OF $\overline{x_S}$

In this section we describe the details of the range proof of $\mathbf{D}^+_{sk}(\overline{x_S})$ in the RegisterUser protocol. The range proof ensures that the TS cannot recognize anonymous users by constructing specially crafted versions of $\overline{x_S}$ as explained earlier. When using a homomorphic encryption scheme that supports zero-knowledge proofs, such as Paillier's encryption scheme, we can use standard techniques, see for example the bitwise technique by Bellare and Goldwasser [67], to prove that $\mathbf{D}^+_{sk}(\overline{x_S})$ is at most 2ℓ bits (which is a sufficient proxy for p in our schemes).

In our implementation, however, we use Joye and Libert's encryption scheme which does not readily allow zero-knowledge proofs. Therefore, we instantiate the range proof using a construction that consists of two parts.

I) The TS constructs a commitment C to x_S using a commitment scheme whose message space is at least as big as the plaintext space of the encryption scheme. The TS then uses a traditional zero-knowledge proof to show that the value x_S committed in C is smaller than p.

II) Next, the TS uses a cut-and-choose technique to show that C commits to $\mathbf{D}^+_{sk}(\overline{x_S}) = x_S$.

The details are as follows. The user and TS take $\overline{x_S}$ as input. The TS takes as private input x_S and the randomizer κ used to construct $\overline{x_S}$. Let $\overline{\mathbb{G}}$ be a cyclic group of order $\overline{p}$ generated by $\overline{g}$ such that $\overline{p} > N$ (recall, N is the size of the plaintext domain of the homomorphic encryption scheme). Let $\overline{h}$ be another generator of $\overline{\mathbb{G}}$ such that the discrete logarithm of $\overline{h}$ with respect to $\overline{g}$ is unknown. We use this group to create a commitment scheme with a large message space.

The details of the first step are as follows. Part I is represented by step 1, whereas part II is represented by the cut-and-choose technique in steps 2 – 7. If at any step a verification fails, the protocol is aborted. The cut-and-choose technique is very similar to the construction we use in the ObtainKeyShareToken and GenShares protocols. Let k be the difficulty level of the cut-and-choose protocol.

1) The TS creates a non-interactive proof that the commitment C contains key-share x_S of the correct size:

$$PK\{(x_S, r) : C = \overline{g}^{x_S}\overline{h}^r \wedge x_S \in [0, p)\}. \quad (7)$$

and sends this proof to the user. This proof can be implemented using a standard technique like the bitwise commitment technique of Bellare and Goldwasser [67]. The user checks the correctness of the proof.

2) The user randomly chooses a subset $\mathcal{D} \subset \{1, \ldots, 2k\}$ of cardinality k. She commits to $\mathcal{D}$ by picking $\theta \in_R \{0,1\}^\ell$ and sending $\Delta = \mathsf{ExtCommit}(\mathcal{D}, \theta)$ to the TS.

3) The TS picks randomizers $\delta_1, \ldots, \delta_{2k} \in_R \{0,1\}^{\ell_\delta}$ and $\kappa_1, \ldots, \kappa_{2k} \in_R \mathcal{R}$ to construct ciphertexts, and $r, r_1, \ldots, r_{2k} \in \mathbb{Z}_{\overline{p}}$ to create commitments. Then, the TS computes a commitment $C = \overline{g}^{x_S}\overline{h}^r$ and sets:

$$\begin{aligned} c_i &= \mathbf{E}^+_{pk}(\delta_i; \kappa_i) \\ C_i &= \overline{g}^{\delta_i}\overline{h}^{r_i} \end{aligned} \quad (8)$$

for $i = 1, \ldots, 2k$. Finally, the TS sends the ciphertexts $c_1, \ldots, c_{2k}$ and commitments $C, C_1, \ldots, C_{2k}$ to the user. The commitments are computationally binding and information theoretically hiding. (Contrary to the ObtainKeyShareToken protocol, the TS can safely send the ciphertexts, because the user cannot decrypt them.)

4) The user sends the subset $\mathcal{D}$ and the commitment randomizer θ to the TS.

5) If $\Delta = \mathsf{ExtCommit}(\mathcal{D}, \theta)$, then the TS sends $(\delta_i, \kappa_i, r_i)_{i \in \mathcal{D}}$ to the user (otherwise, it aborts). The user verifies that the values c_i, C_i for $i \in \mathcal{D}$ satisfy equation (8). Moreover, the user checks that $\delta_i < 2^{\ell_\delta}$ for $i \in \mathcal{D}$.

6) Next, the TS computes

$$\gamma_i = \delta_i - x_S, \quad \rho_i = r_i - r, \quad \nu_i = \kappa_i \kappa^{-1}$$

for $i \notin \mathcal{D}$, and sends them to the user.

7) Finally, the user checks that

$$\begin{aligned} c_i &= \overline{x_S} \cdot \mathbf{E}^+_{pk}(\gamma_i; \nu_i) \\ C_i &= C \cdot \overline{g}^{\gamma_i}\overline{h}^{\rho_i} \end{aligned} \quad (9)$$

and that $0 \le \gamma_i < 2^{\ell_\delta}$ for $i \notin \mathcal{D}$, and accepts the proof if all verifications are correct.

Lemma 3. *If the user does not reject in the above protocol, then with probability $1 - 1/\binom{2k}{k}$ we have that $\mathbf{D}^+_{sk}(\overline{x_S}) \in [0, p)$ as required.*

Proof. From the zero-knowledge proof in step 1, we know that the TS knows an opening α', r' of $C = \overline{g}^{\alpha'}\overline{h}^{r'}$ such that $0 \leq \alpha' < p$. We complete the proof by showing that $\alpha' = \mathbf{D}_{sk}^+(\overline{x_S})$.

We continue as per Lemma 1 and Lemma 2. We restate them here for completeness. First, along the lines of Lemma 1, with probability $1 - 1/\binom{2k}{k}$ there exists an index i^* such that the TS knows an opening δ^*, r^* such that:

$$\delta^* = \mathbf{D}_{sk}^+(c_{i^*}) < 2^{\ell_\delta}$$
$$C_{i^*} = \overline{g}^{\delta^*}\overline{h}^{r^*}. \tag{10}$$

The user checks that the TS knows an opening for the k pairs that are opened by the TS in step 4. So, the TS must include at least k pairs for which it knows a correct opening. Suppose, for contradiction, that the index i^* does not exist, i.e., that the remaining k pairs are incorrect or cannot be opened by the TS. Since the protocol completed, the user did not detect foul play. This situation can only occur if the TS correctly guesses the set $\mathcal{D}$ in advance. Since the TS does not learn anything about $\mathcal{D}$ before step 3, the probability that none of the remaining pairs is correct is $1/\binom{2k}{k}$, as required.

Assume now that this index i^* as required above exists. We use this to show that C commits to $\mathbf{D}_{sk}^+(c)$, i.e., that $\alpha' = \mathbf{D}_{sk}^+(c)$. From equation (9) we know that:

$$C_{i^*} = C \cdot \overline{g}^{\gamma_{i^*}}\overline{h}^{\rho_{i^*}}$$

so, by using equation (10) and equating exponents, we find that $\delta^* = \alpha' + \gamma_{i^*} \pmod{\overline{p}}$. We know from the zero-knowledge proof that $\alpha' < p$ and by direct inspection that $\gamma < 2^{\ell_\delta}$ therefore, the equality holds over the integers as well, and we have

$$\delta^* = \alpha' + \gamma_{i^*} < 2^{\ell_\delta+1} < N. \tag{11}$$

From equation (9) we also know that:

$$c_{i^*} = \overline{x_S} \cdot \mathbf{E}_{pk}^+(\gamma_{i^*}; \nu_{i^*})$$

By decrypting and using equation (10) we find that:

$$\delta^* = \mathbf{D}_{sk}^+(\overline{x_S} \cdot \mathbf{E}_{pk}^+(\gamma_{i^*}; \nu_{i^*})) = \mathbf{D}_{sk}^+(\overline{x_S}) + \gamma_{i^*} \pmod{N}.$$

Substituting δ^* from equation (11) and substracting γ_{i^*} shows that $\alpha' = \mathbf{D}_{sk}^+(\overline{x_S}) \pmod{N}$, and therefore, by size of α' and $\mathbf{D}_{sk}^+(\overline{x_S}) < N$, that $\alpha' = \mathbf{D}_{sk}^+(\overline{x_S})$ as required. $\square$

In the security proof, we replace $\overline{x_S}$ with the encryption of 0, so that the adversary who has corrupted a user learns nothing about x_S (except what is revealed as a result of the threshold-cryptographic protocol). The following lemma states that we can do so, without the adversary detecting this change.

Lemma 4. *TS can simulate the correctness proof given above such that $\overline{x_S} = \mathbf{E}_{pk}^+(0)$, provided that the encryption scheme is CPA secure and the commitment scheme ExtCommit$(\cdot, \cdot)$ is extractable. This simulation does not require any knowledge of how $\overline{x_S}$ was created.*

This proof uses a sequence of games that interpolates between the situation where the RegisterUser protocol is executed normally, and the situation, where $\overline{x_S}$ is an encryption of 0. This game is as in the security game: the adversary can make RegisterUser, ObtainKeyShareToken, GenShares, and BlockShare queries. It's task is to determine if $\overline{x_S}$ is as in the original protocol, or $\overline{x_S} = \mathbf{E}_{pk}^+(0)$. In particular:

- Game 0. In Game 0, $\overline{x_S}$ is constructed as per the protocol.
- Game 1. We proceed as in Game 0, but simulate the cut-and-choose proof in steps $2 - 7$ by extracting $\mathcal{D}$.
- Game 2. As in Game 1, but simulate the zero-knowledge proof in step 1 of the protocol.
- Game 3. As in Game 2, but replace the commitment C by a random commitment.
- Game 4. As in Game 3, but replace $\overline{x_S}$ with an encryption of 0.

We show that each pair of consecutive games is indistinguishable to a polynomially-bounded adversary. Hence, no adversary can distinguish Game 0 from Game 4, thus proving the lemma.

Proof of Lemma 4. We first show how to simulate the cut-and-choose proof in steps $2 - 7$. The adverary sends a commitment Δ to the TS in step 1. We use the extractability of ExtCommit$(\cdot, \cdot)$ to recover $\mathcal{D}$ from Δ (for example, using the random oracle model if it is implemented using a hash-function).

We change how TS acts in step 3. Let $\mathcal{D} \subset \{1, \ldots, 2k\}$ be the subset of cardinality k extracted from Δ. For all $i \in \mathcal{D}$ the TS sets c_i and C_i as per equation (8). For other elements, i.e., for $i \in \{1, \ldots, 2k\} \setminus \mathcal{D}$, the TS generates $\gamma \in_R \{0, \ldots, 2^{\ell_\delta} - 1\}, \rho \in_R \mathbb{Z}_{\overline{p}}, \nu \in_R \mathcal{R}$ and sets c_i and C_i as per equation (9).

In step 4, the adversary reveals $\mathcal{D}$ and θ. If $\Delta = \text{ExtCommit}(\mathcal{D}, \theta)$ then with overwhelming probability, we correctly extracted $\mathcal{D}$. If we correctly extracted $\mathcal{D}$, the TS can open the tuples for $i \in \mathcal{D}$ in step 5 and return γ_i, ρ_i, ν_i for the other elements. Both satisfy the adversary's checks in steps 5 and 7.

Games 0 is indistinguishable from Game 1. The simulated proof can go wrong for two reasons. One, we can fail to extract the disclose set $\mathcal{D}$, but this can only happen with negligible probability. Second, the distribution of γ_is for $i \notin \mathcal{D}$ is not completely correct, however, the size of δ ensures that this difference is statistically hidden from the adversary. So, from the point of view of the adversary, Games 0 and 1 are indistinguishable.

In Game 2 we simulate the zero-knowledge proof in step 1. By construction of the simulator of this proof, the adversary cannot detect this change.

As a result of the changes we made in Game 1, the answers of TS do not depend on the opening of C. So, in Game 3 the TS can generate a random commitment $C \in_R \overline{\mathbb{G}}$. Since Pedersen's commitment scheme is information-theoretically hiding, the adversary cannot detect this change.

In Game 4, the TS sends $\overline{x_S} = \mathbf{E}_{pk}^+(0)$ to the user instead of an encryption of the key-share x_S. As a result of the changes

we made in Game 1, the TS can still complete the remaining part of the protocol.

We claim that the adversary $\mathcal{A}$ cannot distinguish Games 3 and 4. Suppose to the contrary that $\mathcal{A}$ *can* distinguish Games 3 and 4. We then show that $\mathcal{A}$ can break the CPA security of the homomorphic encryption scheme.

To do so, we build an adversary $\mathcal{B}$ against the CPA security of the encryption scheme. Recall that $\mathcal{B}$ can make a challenge query on two messages m_0 and m_1. In our case, $\mathcal{B}$ picks $m_0 = x_S$ and $m_1 = 0$. Then, its challenger returns a ciphertext $c_* = \mathbf{E}_{pk}^{+}(pk, m_b)$ for some bit $b \in_R \{0, 1\}$. Adversary $\mathcal{B}$ needs to guess b.

In RegisterUser queries for the challenge user U^*, adversary $\mathcal{B}$ (which acts as a challenger to $\mathcal{A}$) uses $\overline{x_S} = c_*$. Clearly, if $b = 0$, then $\mathcal{B}$ perfectly simulates Game 3. If $b = 1$, it perfectly simulates Game 4. Therefore, if $\mathcal{A}$ can distinguish between Games 3 and 4, it can break the CPA security of the encryption scheme. □

APPENDIX C
SECURITY PROOF

In the security proof, the challenger controls the TS and the adversary tries to attack a user. The security proof is a sequence of games. In the final game, the challenger simulates the game using only the TCP oracle of the TCP security game, without knowing the corresponding TS' key-share x_S. As a result, any adversary that manages to use the blocked key of that user must therefore break the security of the underlying threshold-cryptographic protocol.

We use the following sequence of games:

- Game 0: We play the game as described in the TANDEM Security game, see Game 1 on page 9.
- Game 1: We change the definition of GenShares. The challenger simulates the workings of TS but does not decrypt any ciphertext. Instead, the TS uses the extractability of ExtCommit$(\cdot, \cdot)$ and the Δ_is (from the corresponding ObtainKeyShareToken protocol) to compute the plaintext corresponding to c_j (without decrypting), which it uses as $\tilde{x}_S$. Finally, the TS constructs the proof of knowledge of $\tilde{x}_S$ as before.
- Game 2: We guess the challenge user U^* and we change the definition of RegisterUser for this user: we replace $\overline{x_S} = \mathbf{E}_{pk}^{+}(x_S)$ by $\overline{x_S} = \mathbf{E}_{pk}^{+}(0)$.
- Game 3: For all non-challenge users we answer GenShares queries as in the previous game. For U^* the TS simulates the TCP following GenShares using the TCP security oracle (without knowing x_S of U^*).

We then prove the following:

- The adversary cannot distinguish Game 0 from Game 1. We prove that as long as one of the tuples is as it should be—and Lemma 1 shows that this is the case with high probability—then we correctly recover the plaintext of c_j and thus the TS extracts the correct $\tilde{x}_S$, and therefore the TCP is correct as well.
- The adversary cannot distinguish Game 1 from Game 2. We no longer decrypt ciphertexts. Hence, we can use

the CPA security of the encryption scheme to show that the adversary cannot distinguish Game 1 from Game 2. More formally, we build a distinguisher that interpolates between Games 1 and 2. The distinguisher makes a query for $m_0 = x_S$ and $m_1 = 0$ to its CPA challenger, and uses the answer as $\overline{x_S}$. Lemma 4 shows the adversary cannot detect this change to RegisterUser. If the CPA challenger returned an encryption of x_S then the distinguisher perfectly simulates Game 1, otherwise it simulates Game 2. We can still answer GenShares queries correctly, since we no longer need to decrypt any ciphertexts.

- The adversary cannot distinguish Game 2 from Game 3 because the TCP oracle simulates the same protocol.
- Finally, if we have an adversary that can win Game 3, then it breaks the security of the TCP because by construction the challenger has no new tokens for the challenge user U^* (because she is blocked or rate-limited) in the challenge phase.

Proof of Theorem 1. This proof follows the sequence of games highlighted above. Let U^* be the challenge user. We guess this user. If the guess turns out to be incorrect, we repeat the reduction with a new guess.

In Game 1 we change how the TS responds to RunTCP queries, in particular, we change GenShares for the challenge user U^*. The TS (controlled by the challenger) no longer decrypts the ciphertext c_j revealed in a token, but instead directly recovers the plaintext using the Δ_i values. The TS then continues as before.

To enable the TS to answer RunTCP queries without decrypting, the TS stores some extra values whenever $\mathcal{A}$ runs the ObtainKeyShareToken protocol. Whenever the TS blindly signs a token, it extracts, id, the token's identifier (normally, the TS cannot learn this value). The challenger uses the extractability of ExtCommit$(\cdot, \cdot)$ to find inputs $\delta'_{i_1}, \ldots, \delta'_{i_m}$ and $\kappa'_{i_1}, \ldots, \kappa'_{i_m}$ used to create the unopened commitments $\Delta_{i_1}, \ldots, \Delta_{i_m}$. (The adversary might cheat so that not all Δ_is are true commitments.) By Lemma 1, $m \geq 1$, and there exists i^* such that the extracted inputs are correct, i.e., $\delta'_{i^*} = \delta_{i^*}$ and $\kappa'_{i^*} = \kappa_{i^*}$. The challenger records the tuple $(id, U, (i_1, \delta'_{i_1}, \kappa'_{i_1}), \ldots, (i_m, \delta'_{i_m}, \kappa'_{i_m}))$ for later use.

We now show how to answer RunTCP queries without needing to decrypt the ciphertexts. The TS initially follows the GenShares protocol. At the start of the protocol, $\mathcal{A}$ sends a token $(id, c_1, \ldots, c_k)$ to the TS (run by the challenger) together with a (blind) signature on it produced by the TS. Moreover, $\mathcal{A}$ provides $\gamma_1, \ldots, \gamma_k$ and $\nu_1, \ldots, \nu_k$. The TS then checks that these values are correct. If not, it aborts. So far, the challenger follows the protocol.

Now, we start to deviate from the protocol. The challenger looks up the corresponding tuple $(id, U, (i_1, \delta'_{i_1}, \kappa'_{i_1}), \ldots, (i_m, \delta'_{i_m}, \kappa'_{i_m}))$ from tokens it issued. Let $\overline{x_S}$ be the encrypted key share for this user. We use the values $\delta'_{i_1}, \ldots, \delta'_{i_m}$ and $\kappa'_{i_1}, \ldots, \kappa'_{i_m}$ to find the plaintext of one of $c_{i_1}, \ldots, c_{i_m}$ and then use this to compute the plaintext of c_j.

For $i \in i_1, \ldots, i_m$ test if:

$$c_i = \overline{x_S} \cdot \mathbf{E}_{pk}^+(\delta_i', \kappa_i')$$

Let $(i^*, \delta_{i^*}', \kappa_{i^*}')$ be the tuple that satisfies this equation. By Lemma 1 we know that there must exist an index i^* such that:

$$c_{i^*} = \overline{x_S} \cdot \mathbf{E}_{pk}^+(\delta_{i^*}, \kappa_{i^*}),$$
$$\Delta_{i^*} = \mathsf{ExtCommit}((\delta_{i^*}, \kappa_{i^*}), \xi^*),$$

so this procedure does indeed find such a tuple $(i^*, \delta_{i^*}', \kappa_{i^*}')$. The plaintext of c_{i^*} thus is $x_S + \delta_{i^*}'$. If $i^* = j$ we are done, and $\tilde{x}_S = x_S + \delta_{i^*} \pmod{p}$. Otherwise, the plaintext of c_j is $x_S + \delta_{i^*} + \gamma_{i^*}$ and therefore $\tilde{x}_S = x_S + \delta_{i^*}' + \gamma_{i^*}' \pmod{p}$.

Now that the challenger has derived $\tilde{x}_S$ it continues with the TCP as normal. This shows how we can answer RunTCP queries without needing to decrypt the ciphertexts.

Games 0 and 1 cannot be distinguished by the adversary. During ObtainKeyShareToken queries, the TS extracts the token identifier id using rewinding, so this is not detected by the adversary. By Lemma 1 the index i^* exists with overwhelming probability, so the responses of the TS are completely identical for the RunTCP queries made by the adversary.

In Game 2, we do not send $\overline{x_S} = \mathbf{E}_{pk}^+(x_{U^*,S})$ to the adversary when it makes RegisterUser queries for the challenge user U^*. Instead, we send $\overline{x_S} = \mathbf{E}_{pk}^+(0)$. During RunTCP queries, we first extract the plaintext of c_j as above, and then add $x_{U^*,S}$. The fact that the TS does not need to decrypt c_j to answer RunTCP queries together with Lemma 4 shows that the adversary cannot detect this change.

In Game 3, we again change how we answer RunTCP queries for the challenge user U^*. In particular, we will answer this query without using the corresponding key-share $x_{U^*,S}$. Instead, we use the challenge oracle for the TCP security in the query phase. We proceed as before, to find the plaintext δ of c_j when running GenShares. However, now we use the TCP challenge oracle to run the TCP by making a $\mathsf{TCP}(\delta \mod p)$ query. The TANDEM security challenger relays the messages to the adversary $\mathcal{A}$. After the selection phase, we advance the TCP security challenger to the challenge phase. Moreover, the challenge user U^* cannot obtain new tokens (because U^* is either blocked or rate-limited), and all old tokens have been invalidated, so we no longer need access to the TCP oracle to answer queries. Finally, in the challenge phase, we relay the messages to the TCP challenger. Then, if adversary $\mathcal{A}$ wins Game 3, it breaks the TCP security of the underlying TCP. Since we assumed this cannot happen, the TANDEM scheme is secure as well. The only difference between Game 2 and Game 3 is that we use the TCP oracle to run the TCP. However, since the TCP oracle uses to correct randomized key, this change is indistinguishable to the adversary. $\qquad\square$

APPENDIX D
PRIVACY PROOF

In our privacy proof, we reduce an attacker against privacy to an attacker on the underlying blind signature scheme (which we instantiate using the BBS+ credential scheme). In terms of attribute-based credentials, this game is precisely the issuer-unlinkability game. This game is the standard blind-signature game [68], [69].

GAME 6. The *blind-signature game* is between a challenger controlling an honest user U and an adversary controlling the signer.

Setup At the start of the game, $\mathcal{A}$ publishes the public key of the signer and outputs all other necessary public parameters.

Challenge At some point, the adversary outputs two messages m_0 and m_1 on which it wants to be challenged. The challenger picks a bit $b \in_R \{0,1\}$ and proceeds as follows. (1) User U engages with the signer to obtain a signature on m_b. Let σ_b be the corresponding signature. (2) User U engages with the signer to obtain a signature on m_{1-b}. Let σ_{1-b} be the corresponding signature. If either signing protocol fails, set $\sigma_0 = \sigma_1 = \bot$. Finally, the challenger sends (m_0, σ_0) and (m_1, σ_1) to the adversary.

Guess Finally, the adversary outputs a guess b' of b. The adversary wins if $b' = b$.

If no adversary can win this game then the signer can recognize neither the signature nor the message.

The computationally hiding commitments in the Obtain-KeyShareToken protocol ensure that the TS learns nothing about the unrevealed ciphertexts c_i which it then blindly signs—again without learning anything about the message. So, when the user runs GenShares and thereby reveals these ciphertexts, they cannot be directly correlated to a corresponding run of ObtainKeyShareToken. Moreover, the plaintext corresponding to the ciphertexts c_i are fully randomized, so that these too do not reveal anything about the user with which the TS is currently interacting.

The privacy proof folows a sequence of games. Throughout we use a guess i_0, i_1 for the challenge tokens. If this guess turns out to be incorrect when the adversary makes it challenge query, we abort and try again. We first use a sequence of games to show that we can remove identifying information from the ObtainKeyShareToken protocol.

- Game 0 is the Tandem privacy game, see Game 4 on page 10.
- In Game 1, we extract the TS key-shares $x_{0,S}$ and $x_{1,S}$ for users U_0 and U_1 from the TS' proof of knowledge in step I of the RegisterUser protocol, see Appendix B.
- In Game 2, we forge the user's zero-knowledge proof of correct construction of C, the commitment to the token identifier id and the randomized ciphertexts, at the end of ObtainKeyShareToken protocol.
- In Game 3, we replace the ciphertext $\overline{id}$ with the encryption of 0. The CPA security of the encryption scheme ensures that the adversary cannot detect this change.
- In Game 4, we use the extractability of $\mathsf{ExtCommit}(\cdot, \cdot)$ to forge the user's cut-and-choose proof in the Obtain-KeyShareToken protocol, and send random commitments C_i, Δ_i for $i \notin \mathcal{D}$. However, we honestly construct C as per the protocol.

- In Game 5, for user U_i and the challenge token, we set $c_i = \mathbf{E}_{pk}^+(x_{i,S} + \delta_i, \kappa_i)$ for $i \notin \mathcal{D}$, rather than using $\overline{x_S}$. We commit to c_i for $i \notin \mathcal{D}$ as usual. Lemma 3 shows that with high probability we still follow the protocol correctly.
- In Game 6, we omit $x_{i,S}$ altogether in the construction of the unrevealed c_i, that is, we set:

$$c_i = \mathbf{E}_{pk}^+(\delta_i, \kappa_i) \tag{12}$$

for all $i \notin \mathcal{D}$, and use these values to construct C. When answering RunTCP queries, user i adds $x_{i,S}$, which we extract during the RegisterUser protocol, to its long-term secret-share x_U to compensate for this change. The size of the randomizers δ_i ensure that the TS cannot detect this change.

We are now in the situation where the tokens held by user 0 and 1 are exchangeable. We use this to show that no adversary can distinguish situations $b = 0$ and $b = 1$. We use a sequence of games to interpolate between the two situations. We start from Game 6.

- In Game A, the challenger uses $b = 0$ but otherwise proceeds as in Game 6.
- In Game B, the challenger swaps the signatures of the challenge tokens of users U_0 and U_1. By the blind signature game, the adversary cannot detect this change.
- In Game C, the challenger also swaps the users U_0 and U_1 in the challenge phase. As a result, it perfectly simulates $b = 1$ in Game 6. The privacy property of the threshold cryptographic protocol (with colluding respectively honest SP) ensures that the adversary cannot detect this change.

Since these steps are indistinguishable, no adversary can distinguish the situations $b = 0$ and $b = 1$ in Game 6, and by indistinguishability again, neither can any adversary distinguish these two in the original privacy game.

Proof of Theorem 2. Throughout this proof, we use a guess for the challenge tokens i_0 and i_1 of users U_0 and U_1 respectively. If this guess turns out to be wrong in the challenge step, we abort and try again.

In Game 1, the challenger extracts $x_{0,S}$ and $x_{1,S}$ for users U_0 and U_1. In particular, the challenger runs the knowledge extractor on the proof of knowledge of the RegisterUser protocol, see Equation 7, for each of the users. Since the extractor uses rewinding, the adversary does not detect this.

In Game 2, the challenger forges the proof of knowledge of correctness of the commitment C at the end of the Obtain-KeyShareToken protocol for the challenge tokens i_0 and i_1 of users U_0, U_1 respectively. By simulatability of this proof, the adversary cannot detect this change.

In Game 3, the challenger replaces the encryption of the token identifier id for the challenge tokens i_0 and i_1 with the encryption of the value 0. The proof of knowledge of correct encryption is already forged in the previous game. A reduction to the CPA security of the encryption scheme shows that an

adversary that can distinguish Games 2 and 3 can break the CPA security of the encryption scheme.

In Game 4, the challenger extracts the subset $\mathcal{D}$ from the commitment Δ as soon as it receives it. For the two challenge tokens, the challenger (acting as the user) now proceeds as follows. It computes C_i, Δ_i for $i \in \mathcal{D}$ as per the protocol. However, for $i \notin \mathcal{D}$ it lets the unrevealed commitments C_i and Δ_i commit to random values. The proof of knowledge that C commits to the same values as C_i is already forged since a previous step. Because the commitment scheme is information theoretically hiding, the adversary cannot detect this change.

Despite the changes we made, the final token that is stored by the user is exactly the same as in the original ObtainKeyShareToken protocol. In Game 5 we compute the values c_i for user U_j and $i \notin \mathcal{D}$ as $c_i = \mathbf{E}_{pk}^+(x_{j,S} + \delta_i, \kappa_i)$ (recall, we extracted $x_{j,S}$ in the RegisterUser phase) instead of $c_i = \overline{x_S} \cdot \mathbf{E}_{pk}^+(\delta_i, \kappa_i)$. Lemma 3 shows that with overwhelming probability $\mathbf{D}_{sk}^+(\overline{x_S})$ equals the value $x_{j,S}$ we extracted in the RegisterUser protocol, so this change does not modify the adversary's view.

In Game 5, the user U_j computes

$$c_i = \mathbf{E}_{pk}^+(x_{j,S} + \delta_i, \kappa_i)$$

In Game 6, we remove the $x_{j,S}$ component from this equation, and instead just compute

$$c_i = \mathbf{E}_{pk}^+(\delta_i, \kappa_i) \tag{13}$$

for the challenge tokens. To compensate for the fact that $x_{j,S}$ is no longer included, the users adds $x_{j,S}$ to x_U. As a result, the threshold cryptographic protocol still completes as before.

The size of the domain from which the δ_is are drawn, ensures that the adversary cannot detect this change. More formally, the user sends c_is, γ_is, and ν_is. However, the last two sets are redundant, they can be computed directly based on the c_is. As a result, we can focus on $\delta_i = \mathbf{D}_{sk}^+(c_i)$. By the size of the domain δ_is and the size of $x_{j,S}$ tuples $(\delta_1, \ldots, \delta_k)$ and $(x_{j,S} + \delta_1, \ldots, x_{j,S} + \delta_k)$ are statistically indistinguishable. As a result no adversary can distinguish Games 5 and 6.

We now show that no adversary can win Game 6. We again use a sequence of games, but now interpolate between Game A, where the challenger uses $b = 0$ in Game 6, and Game C, where the challenger uses $b = 1$ in Game 6. We construct the intermediate Game B, where user U_0 uses the token i_1 of user U_1 and vice versa. Since the challenge tokens in Game 6 (and thus in Games A, B, and C) do not depend on the user, the threshold-cryptographic protocols complete correctly as in Game 6.

We first show that Games A and B are indistinguishable. Suppose to the contrary that $\mathcal{A}$ can distinguish Games A and B. We show that we can use $\mathcal{A}$ to build an adversary $\mathcal{B}$ that breaks the blindness property of the signature scheme. In the blind signature game, $\mathcal{B}$ gets oracle access to two users that request a blind signature on one message each. Adversary $\mathcal{B}$ acts as the challenger towards $\mathcal{A}$ in Game 6. At the start of the game $\mathcal{B}$ generates two messages, corresponding to key-share

tokens, for which users U_0 and U_1 need a blind signature. It creates:

$$m_0 = (id, H(c_1), \ldots, H(c_k)) \tag{14}$$
$$m_1 = (id', H(c'_1), \ldots, H(c'_k)), \tag{15}$$

where the values in the tuples are as in Game 6. Adversary $\mathcal{B}$ sends m_0, m_1 to its blind signature challenger.

During the ObtainKeyShareToken protocols for the challenge tokens, $\mathcal{B}$ simulates its users as follows. When user U_0 is running the blind signature protocol to create the challenge token τ_0, $\mathcal{B}$ uses the its challenger of the blind signature game to act as the user. When U_1 runs the blind signature protocol to create token τ_1, $\mathcal{B}$ again uses its blind signature game challenger. Finally, the blind signature challenger outputs two signatures σ_0 and σ_1 on messages m_0 and m_1 respectively. Adversary $\mathcal{B}$ uses σ_0 to construct the key-share token for user U_0, and uses σ_1 to construct the key-share token for user U_1. If $b = 0$ in the blind-signature game, $\mathcal{B}$s challenge user first blindly signed m_0, so $\mathcal{B}$ perfectly simulates Game A. If $b = 1$ in the blind-signature game, then $\mathcal{B}$ perfectly simulates Game B. Hence, any distinguisher between Games A and B breaks the blindness property of the blind signature scheme.

We now show that if the TCP scheme is private (with a colluding respectively honest SP), no adversary can distinguish between Games B and C. Suppose to the contrary that adversary $\mathcal{A}$ can distinguish Game B from Game C. We show that we can use $\mathcal{A}$ to build an adversary $\mathcal{B}$ that breaks the privacy property of the TCP scheme. Adversary $\mathcal{B}$ simulates users U_0 and U_1 towards $\mathcal{A}$. The RegisterUser and ObtainKeyShareToken protocols do not involve the users' secrets, so $\mathcal{B}$ computes them directly. We now show how to answer RunTCP queries.

Whenever $\mathcal{A}$ makes a RunTCP(U_i, j, in_U) query, $\mathcal{B}$ makes a RunTCP(i, in_U) query of its challenger. Distinguisher $\mathcal{B}$'s challenger replies with the TS' key-share $\tilde{x}_S$. Let $\tau = (\sigma, \epsilon, id, (c_l, \kappa_l, \delta_l)_{l=1,\ldots,k})$ be the jth token of user U_i. Normally, this token dictates a TS key-share unequal to $\tilde{x}_S$, but we can use the random oracle and change the token to ensure that the TS will recover $\tilde{x}_S$. To do so, adversary $\mathcal{B}$ picks $\delta'_1, \ldots, \delta'_k \in_R \{0, \ldots, p2^\ell\}$. Let δ'_m be the largest, then we slightly increase this value (by at most p) so that $\delta'_m = \tilde{x}_S \pmod{p}$. (With overwhelming probability this modified δ'_m is less than 32^{ℓ_δ}; if not, we try again.) Then, we pick $\kappa'_1, \ldots, \kappa'_k \in_R \mathcal{R}$ and set $c'_l = \mathbf{E}^+_{pk}(\delta'_l; \kappa'_l)$. Adversary $\mathcal{B}$ updates the random oracle to ensure that $H(c'_i) = H(c_i)$, i.e., the new pairs have the same hash values as the original pairs. Next, $\mathcal{B}$ uses token $\tau' = (\sigma, id, (c'_l, \kappa'_i, \delta'_l)_{l=1,\ldots,k})$ to run GenShares with the TS. The changes to the random oracle ensure that this token is valid. Moreover, the changes to the random oracle succeed with high probability since at no point in the games does the TS learn the inputs to these hash-functions. The TS will derive the correct secret share $\tilde{x}_S$ from τ'. So it runs the correct TCP protocol with the requested user which is simulated by $\mathcal{B}$'s challenger.

To answer $\mathcal{A}$'s challenge queries, $\mathcal{B}$ again uses his challenger and proceeds as above to answer the queries. If $b = 0$

in the TCP privacy game, then $\mathcal{B}$'s first run of RunTCP uses user U_0's key, so $\mathcal{B}$ simulates Game B. Otherwise, if $b = 1$, then $\mathcal{B}$ simulates Game C. So, any adversary $\mathcal{A}$ that can distinguish Games B and C breaks the privacy property of the TCP scheme. This completes the privacy proof. $\square$