

Reputation Systems for Anonymous Networks

Elli Androulaki, Seung Geol Choi, Steven M. Bellovin, and Tal Malkin

Department of Computer Science, Columbia University
{elli,sgchoi,smb,tal}@cs.columbia.edu

Abstract. We present a reputation scheme for a pseudonymous peer-to-peer (P2P) system in an anonymous network. Misbehavior is one of the biggest problems in pseudonymous P2P systems, where there is little incentive for proper behavior. In our scheme, using ecash for reputation points, the reputation of each user is closely related to his real identity rather than to his current pseudonym. Thus, our scheme allows an honest user to switch to a new pseudonym keeping his good reputation, while hindering a malicious user from erasing his trail of evil deeds with a new pseudonym.

1 Introduction

Pseudonymous System. Anonymity is a desirable attribute to users (or peers) who participate in peer-to-peer (P2P) system. A peer, representing himself via a pseudonym, is free from the burden of revealing his real identity when carrying out transactions with others. He can make his transactions unlinkable (i.e., hard to tell whether they come from the same peer) by using a different pseudonym in each transaction. Complete anonymity, however, is not desirable for the good of the whole community in the system: an honest peer has no choice but to suffer from repeated misbehaviors (e.g. sending an infected file to others) of a malicious peer, which lead to no consequences in this perfectly pseudonymous world.

Reputation System. We present a reputation system as a reasonable solution to the above problem. In our system, two peers, after carrying out a transaction, evaluate each other by giving (or not) a *reputation point*. Reputation points assigned to each peer sum up to create that peer's reputation value. In addition, reputation values are public, which helps peers to decide whether it is safe or not to interact with a particular peer (more exactly a pseudonym).

Identity Bound Reputation System. We stress that, in our system, the reputation value is bound to each peer. In existing reputation systems [16, 18], the reputation value is bound to each pseudonym. Consequently, a new pseudonym of a

malicious peer will have a neutral reputation, irrespective of his past evil deeds. Thus, honest peers may still suffer from future misbehavior. On the other side, honest users won't use a new pseudonym, in order to keep the reputation they have accumulated. Thus, they cannot fully enjoy anonymity and unlinkability. Motivated by this discussion, our goal in this paper is to design an identity bound reputation system, combining the advantages of anonymity and reputation.

Our Contribution. First, we formally define security for identity bound reputation systems (Section 3). As far as we are aware, this is the first such security definition. Our definition captures the following informal requirements:

- Each peer has a reputation which he cannot lie about or shed. In particular, though each peer generates as many one time pseudonyms as he needs for his transactions, all of them must share the same reputation. Also, our system is robust against a peer's deliberate attempts to increase his own reputation.
- Reputation are updated and demonstrated in a way that does not compromise anonymity. In particular, the system maintains unlinkability between the identity of a peer and his pseudonyms and unlinkability among pseudonyms of the same peer.

Our second contribution is the construction of a reputation scheme that satisfies the security definition. It is a nontrivial task to realize a secure identity bound reputation scheme, as the requirements of anonymity and reputation maintenance are (seemingly) conflicting. Here, we only briefly give basic ideas for the construction (see Section 2 for high level description of our scheme and Section 5 for the detail). To satisfy the first item, we need a central entity, *Bank*. Bank, aware of the identity of each peer, keeps reputation accounts by the peer, and is considered trusted to perform its functional operations — reputation updates etc. — correctly. Since we do not consider Bank trusted in terms of the anonymity requirements, we need to utilize a two-stage reputation deposit procedure. For the second item, we use the concept of *e-cash*. E-cash is well-suited to our system since it can be spent anonymously, even to Bank. We also use other primitives, such as anonymous credential system and blind signatures.

Organization. In Section 2 we provide a high level description of our scheme. In Section 3 we present our model, including security requirements. The building blocks used by our system are described in Section 4, followed by a detailed description of our system in Section 5. Related work and future directions are discussed in Sections 6 and 7 respectively.

2 System Considerations and High Level Approach

In this section we discuss system considerations and present a high level description of our scheme.


System Considerations and Assumptions. We assume that all communication takes place over an anonymous communication network, e.g., a Mixnet [8] or an Onion Router [24, 11]. We further assume that this network is, in fact, secure. While we are not minimizing the difficulty of achieving that — see, for example, [15] or [21] — we regard that problem as out of scope for this paper.

We also assume certain out-of-band properties that are necessary for correspondence to the real world. The most important such assumption is that there is some limit to the number of reputation points any party can hand out per unit time. While we don't specify how this limit is set, we tentatively assume that it costs real money to obtain such points to hand out. This might, for example, be the daily membership fee for participation in the P2P network. Note that the assumption corresponds quite well to the best-known existing reputation system, Ebay. One can only dispense reputation points there after making a purchase; that in turn requires payment of a fee to the auction site. Bhattacharjee and Goel have derived a model for what this fee should be [4]; they call the necessary property “inflation resistance”.

A last assumption is unbounded collusion. That is, any number of parties on this network may collude to break anonymity of some other party. We specifically include the bank in this assumption. We assume collusion because in most real environments, it is possible for one party to open multiple accounts on the system. It may cost more money, but it does achieve the goal. Since a bank employee can do the same, we assume that the bank is colluding, too, albeit perhaps in response to a court order. Even if we assume a foolproof system for restricting accounts to one per person, two or more people could communicate via a private back channel, thus effectively creating multiple accounts under control of a single entity.

On the other hand, the bank is trusted to behave honestly in its functional transactions, which involve maintenance of reputation levels and repcoins for each peer (see below). Thus, if the bank is misbehaving (possibly in coalition with other adversarial users), it can compromise the correctness of the system, but not the anonymity. It is possible to distribute the bank functionality among several parties in order to increase fault tolerance and reduce any trust assumptions, but we will not describe this here.

Protocol Overview. Bank keeps the record of each peer's reputation in the *reputation database*. As shown on the left of Figure 1, a peer U (via his pseudonym



- Reputation granting process (left): (1) U withdraws a wallet W (i.e., repcoins) from the Bank B . (2) U , via P_U , awards (i.e., spends) a repcoin (S, π) to M . (3) M , via P_M , deposits the repcoin (S, π) . (4) If the deposit is successful, P_M obtains from B a blind permission σ . Note that σ is blind to B and only visible to M . (5) M deposits σ , and B increases M 's reputation point.
- Reputation demonstration process (right): (1) M requests a credential for the group G_i . (2) If M has enough reputation count for G_i , B issues a credential $cred$ to M . (3) By using $cred$, P_M proves its membership of G_i to P_U .

Fig. 1. Reputation granting and demonstration

P_U) can increase the reputation of a pseudonym P_M by giving a *repcoin*,¹ which is basically an e-coin. Bank manages the number of repcoins that each peer has using another database: *repcoin quota database*.

Note that M does not deposit the repcoin using his identity. This is for the sake of maintaining unlinkability between a pseudonym and a peer. If M directly deposited the repcoin, collusion of Bank and U would reveal that M and P_M are linked. In fact, this shows the difficulty of realizing a secure reputation scheme: it is not obtained by using an ecash scheme naively. To preserve unlinkability, we use a level of indirection. When P_M successfully deposits the repcoin, it gets a blind permission from Bank. The blind permission is basically a blind signature, which therefore does not contain any information about P_M . So, M can safely deposit the permission.

We chose to employ an anonymous credential system (see Section 4) to construct the reputation demonstration procedure (on the right side of Figure 1). The anonymous credential enables M , via his pseudonym P_M , to prove his membership in group G_i anonymously. Thus, unlinkability between M and P_M is maintained.

We also note that P_M , instead of revealing its exact reputation value, shows the membership of a group G_i . Demonstration of exact reputation value could allow an attacker who continuously queries for the reputation of many pseudonyms — without even needing to transact with them — to infer whether two pseudonyms

¹ If M wants to increase of reputation of P_U , they can carry out the same protocol with their roles reversed.

correspond to the same user. To make matters worse, with Bank's collaboration, pseudonyms can be linked to a limited number of identities that have the exact same reputation value with the pseudonym. On the other hand, grouping together identities which belong to the same reputation level, makes small changes in reputation accounts invisible to other pseudonyms. Bank can still see the changes that take place in peers' reputations, but cannot link them to specific pseudonyms any more. The reputation levels (i.e., groups G_i) are defined as a system parameter. Reputation levels are not necessarily required to be disjoint. One example would be that G_i contains peers who has more than 2^i different reputation values.

Punishing Misbehaving Parties. When modeling the security of the system, we aim to achieve our goals (such as anonymity, no lying about reputation level, no over-awarding reputations beyond the allowed limit, etc.) by rendering a break of the security computationally infeasible (modulo some standard cryptographic assumptions). However, some security breaches are impossible to completely prevent. For example, as long as there is no central party involved on-line in each transaction, a user can always award the same reppoint twice to different parties. As another example, if anonymity and unlinkability is to be preserved, a peer with a high reputation level can always give away all his data and secret keys to another peer, allowing the latter to claim and prove the high reputation as his own. In these cases, we build into our model an incentive structure (similar to previous work, e.g., [19]), whereby such security breaches would hurt the offender. In particular, for the first case above, we require that a double awarding of a reppoint would reveal the identity of the offender (which can then lead to consequences outside of our model). For the second case, we require that in order for Alice to empower Bob, who has a lower reputation level, to prove a reputation level as high as Alice's, Alice would have to effectively give Bob her master private key. This information may be quite sensitive, especially if the private key used within the reputation system is the same one used for a public-key infrastructure outside the system.

3 A Model for Anonymous Reputation Systems

In this section, we present our model for anonymous reputation systems. We first enumerate the types of entities and the operations considered in the system, followed by the security definition. The motivation and rationale for our model and choices were discussed in Section 2. We note that some of these definitions were inspired by previous work on other primitives, such as [6, 5].

3.1 Participating Entities

The entities in an anonymous reputation system are as follows.

- **Peers.** Peers are the regular users of a P2P network. A peer interacts with other peers via pseudonyms of his choice and can be either a User (buyer) or a Merchant in different transactions. Peers can award reputation points to other peers (through their pseudonyms), and can show their reputation level to other peers.
- **Bank.** Bank manages information with respect to each peer's reputation (where the information is tied to actual identities — public keys — of peers, not to pseudonyms). Specifically, it maintains three databases: the repcoin quota database (denoted D_{quota}), the reputation database (denoted D_{rep}), and the history database (denoted D_{hist}).

D_{quota} holds the amount of repcoins that each peer is allowed to award to other peers. When a peer withdraws a wallet of repcoins, the amount of his repcoin quota is decreased correspondingly. Bank also replenishes all the peer's account periodically, as per system parameters (for example, every day each peer can award at most 20 repcoins to others; see the discussion in Section 2). D_{rep} contains the amount of reputation points that each peer has earned by receiving repcoins from other peers. In order to prevent peers from double-awarding (awarding two peers with same-serial-numbered repcoins), D_{hist} holds all the repcoins that are deposited.

3.2 Operations

The operations supported in our system are listed below. When an operation is an interactive procedure (or a protocol consisting of multiple procedures) between two entities A and B , we denote it by $\langle O_A, O_B \rangle \leftarrow \text{Pro}(I_C)[A(I_A), B(I_B)]$, where Pro is the name of the procedure (or protocol). O_A (resp. O_B) is the private output of A (resp. B), I_C is the common input of both entities, and I_A (resp. I_B) is the private input of A (resp. B). We also note that depending on the setup, some operations may require additional global parameters (e.g., some common parameters for efficient zero-knowledge proofs, a modulus p , etc). Our system will need these additional parameters only when using underlying schemes that use such parameters, e.g., e-cash systems or anonymous credential systems. To simplify notation, we omit these potential global parameters from the inputs to all the operations.

- $(pk_B, sk_B) \leftarrow \text{Bkeygen}(1^k)$ is the key generation algorithm for Bank.
- $(pk_U, sk_U) \leftarrow \text{Ukeygen}(1^k)$ is the key generation algorithm for peers. We call pk_U the (master) public key of U , and sk_U the master secret key of U .
- $(P, si_P) \leftarrow \text{Pnymgen}(1^k)$ is the pseudonym generation algorithm for peers. The si_P is the secret information used to generate the pseudonym P .

- $\langle W, D'_{\text{quota}} \rangle / \langle \perp, \perp \rangle \leftarrow \text{RepCoinWithdraw} (pk_B, pk_U, n) [U(sk_U), B(sk_B, D_{\text{quota}})]$. A peer U tries to withdraw n repcoins (in the form of a wallet W) from Bank B . Bank, using D_{quota} , checks if U is eligible for withdrawal. If so, the withdrawal is carried out and D_{quota} is changed accordingly.
- $\langle (W', S, \pi), (S, \pi) \rangle / \langle \perp, \perp \rangle \leftarrow \text{Award} (P_U, P_M, pk_B) [U(si_{P_U}, W, pk_U, sk_U), M(si_{P_M})]$. A peer U (via P_U), using his wallet W , gives a repcoin (S, π) to M (via P_M). Here S is a serial number and π is the proof of a valid repcoin.
- $\langle \top, (D'_{\text{rep}}, D'_{\text{hist}}) \rangle / \langle \perp, \perp \rangle \leftarrow \text{RepCoinDeposit} (pk_B, S, \pi) [M(P_U, si_{P_U}, pk_U, sk_U), B(sk_B, D_{\text{rep}}, D_{\text{hist}})]$. A peer M deposits the repcoin into his reputation account. If the repcoin (S, π) is valid and not double-awarded, then the coin is stored in the history database D_{hist} , and the amount of reputation of pk_M in D_{rep} is increased by one.
- $(pk_U, \Pi_G) / \perp \leftarrow \text{Identify}(S, \pi_1, \pi_2)$. If a repcoin is double-awarded with (S, π_1) and (S, π_2) , Bank can find the peer who double-awarded the coin using this operation. Here, Π_G is a proof that pk_U double-awarded the repcoin with the serial number S .
- $\top / \perp \leftarrow \text{VerifyGuilt}(S, \Pi_G, pk_U)$ outputs \top if the peer U (represented by pk_U) indeed double-awarded the coin with the serial number S .
- $\langle C^l_U, \top \rangle / \langle \perp, \perp \rangle \leftarrow \text{RepCredRequest} (pk_B, pk_U, l) [U(sk_U), B(sk_B, D_{\text{rep}})]$. A peer U requests a credential that will enable U to prove to another peer that he has reputation level l . Bank B refers to D_{rep} , and if U has sufficient reputation it issues a credential C^l_U . (As discussed in Section 2, how exactly the reputation levels are defined is a system parameter).
- $\langle \top, \top \rangle / \langle \perp, \perp \rangle \leftarrow \text{ShowReputation} (P_{U_1}, P_{U_2}, pk_B, l) [U_1(sk_{U_1}, si_{P_{U_1}}, C^l_{U_1}), U_2(si_{P_{U_2}})]$. A peer U_1 (via P_{U_1}) proves to U_2 (via P_{U_2}) that he has reputation level l .

3.3 Security

In this section we define security for anonymous reputation systems.

Adversarial Model. We will consider two adversarial models, assuming the stronger one for the anonymity-related security properties (unlinkability and exculpability), and the weaker one for the reputation-handling properties (no over-awarding and reputation unforgeability).

For the weaker adversarial model, we assume Bank is *honest-but-curious*, that is, it follows the protocol specification correctly. All other peers may become malicious, and behave in arbitrary ways in the protocol. Adversarial parties may collude with each other, and as long as they are peers, they may decide to share any of their state or secret information with each other, and coordinate their actions; Bank may share the content of its maintained databases (D_{quota} , D_{rep} , and D_{hist}), but *not* Bank's secret keys (thus it is meaningful for Bank to be *honest-but-curious*, even when in coalition with other players).²

For the stronger adversarial model, we remove the honest-but-curious restriction on Bank: we assume all parties (including Bank) may be corrupted, collaborating with each other, and behaving arbitrarily.

Correctness.

² Note that if we allowed Bank to share its secret keys and to behave arbitrarily, it could issue more repcoins than allowed, generate reputation credentials that do not correspond to the correct reputation level, etc.

- If an honest peer U_1 , who has enough repcoins in his repcoin quota, runs RepCoinWithdraw with an honest Bank B , then neither will output an error message; if the peer U_1 , using the wallet (output of RepCoinWithdraw), runs Award with an honest peer U_2 (via his pseudonym), then U_2 accepts a repcoin (S, π) ; if the peer U_2 runs RepCoinDeposit with the honest Bank to deposit the repcoin (S, π) then U_2 's reputation in Bank will be increased by one.
- If an honest peer U_1 runs RepCredRequest with an honest Bank and a reputation level for which he is eligible, then U_1 gets a valid credential. For a valid credential C_U^l , its owner can always prove his reputation through ShowReputation(l, C_U^l, \dots) procedure.

Unlinkability.

For an adversary A who has corrupted certain parties including Bank, we say that a peer U appears consistent with a pseudonym P to A , if U and P 's owner are uncorrupted, and if the levels for which P successfully invoked ShowReputation are a subset of the levels for which U successfully invoked RepCredRequest. We now define the following two unlinkability properties:

Peer-Pseudonym Unlinkability. Consider an adversary who, having corrupted some parties including Bank, is participating in the system for some arbitrary sequence of operations executed by honest and by corrupted parties. Given a pseudonym P that does not belong to a corrupted party, the adversary can learn which peer owns P no better than guessing at random among all non-corrupted peers that appear consistent with P .

Pseudonym-Pseudonym Unlinkability. Consider an adversary who, having corrupted some peers (but not Bank), is participating in the system for some arbitrary sequence of operations executed by honest and corrupted parties. Given two pseudonyms P_1, P_2 that do not belong to corrupted parties, the adversary has no advantage in telling whether P_1, P_2 belong to the same peer or not. Next, consider an adversary who corrupted some peers and Bank as well. Then the above requirement should hold as long as there are *at least two* non-corrupted peers who appear consistent with both P_1 and P_2 (because if there is only one such uncorrupted peer, clearly both pseudonyms belong to the same one).

No Over-Awarding.

- No collection of peers should be able to award more repcoins than they withdrew. Suppose that n peers U_1, \dots, U_n collude together, and that the sum of the amount of repcoins allowed to them is N . Then, the number of different serial numbers of repcoins that can be awarded to other peers is at most N .
- Suppose that one or more colluding peers run the Award protocol with two pseudonyms P_{M_1} and P_{M_2} such that P_{M_1} gets (S, π_1) and P_{M_2} gets (S, π_2) . Then, we require that Identify(S, π_1, π_2) outputs a public key pk_U and a proof of guilt Π_G such that VerifyGuilt(pk_U, S, Π_G) accepts.

- Each repcoin that is accepted but not double-awarded in the Award protocol increases exactly one reputation point in the database D_{rep} irrespective of the beneficiary of the repcoin. However, we don't regard it as a breach of security when a peer M_1 received a repcoin but passed it to M_2 , who deposited it into his reputation account; in any event, this is just another form of collusion. Another justification is that the peer M_1 sacrifices one reputation point.

Exculpability.

This property is to protect the honest peer from any kind of framing attack against him. No coalition of peers, even with Bank, can forge a proof Π_G that $\text{VerifyGuilt}(pk_U, S, \Pi_G)$ accepts where pk_U is an honest peer U 's public key who did not double-award a repcoin with the serial number S .

Reputation Unforgeability.

- No coalition of peers, where l is the highest reputation level of any one of them, can show a reputation level higher than l for any of their pseudonyms. This implies as a special case that a single peer cannot forge his reputation.
- Consider a peer U with reputation level l , who owns a pseudonym P . Suppose that some coalition of peers has empowered U with the ability to prove that P has reputation level $l' > l$. Let Bad be the set of peers with reputation level at least l' among the coalition (note that by the previous requirement, there must be at least one peer in Bad). Then, it must be that U can learn the master secret key of a peer $U' \in \text{Bad}$.

4 Building Blocks of our Scheme

Anonymous Credential Systems. In anonymous credential systems — see, for example, [19, 6, 3] — there are three types of players: *users*, *organizations*, and *verifiers*. Users receive credentials, organizations grant and verify the credentials of users, and verifiers verify credentials of the users. Below are the supported procedures.

- $(pk_O, sk_O) \leftarrow \text{AC.OKeyGen}(1^k)$. Key generation algorithm for an organization. (pk_O, sk_O) denotes the key pair of the organization O .
- $(pk_U, sk_U) \leftarrow \text{AC.UKeyGen}(1^k)$. Key generation algorithm for a user. (pk_U, sk_U) denotes the key pair of the user U . Sometimes sk_U is called the master secret key of U .³

³ Anonymous credential systems do not typically require a specific form for the master public and secret keys, but assume it is inherited from some PKI, where users are motivated to keep their secret key secret. In other variations of anonymous credential systems (with all-or-nothing non-transferability) there is no master public key. Our scheme can be adapted to such systems as well.

- $\langle (N, \text{NSecr}_N), (N, \text{NLog}_N) \rangle \leftarrow \text{AC.FormNym}(pk_O) [U(sk_U), O(sk_O)]$. Nym⁴ generation protocol between U and O , where N is output nym, NSecr_N is secret information with respect to N , and NLog_N is the corresponding log on the organization side.
- $\langle \text{cred}_N, \text{CLog}_{\text{cred}_N} \rangle \leftarrow \text{AC.GrantCred}(N, pk_O) [U(pk_U, sk_U, \text{NSecr}_N), O(sk_O, \text{NLog}_N)]$. Credential granting protocol, where cred_N is a credential for the nym N , and $\text{CLog}_{\text{cred}_N}$ is the corresponding log on the organization side.
- $\langle \top, \top \rangle / \langle \perp, \perp \rangle \leftarrow \text{AC.VerifyCred}(pk_O) [U(N, \text{cred}_N), V]$. Credential verification protocol.
- $\langle \top, \top \rangle / \langle \perp, \perp \rangle \leftarrow \text{AC.VerifyCredOnNym}(N, pk_O, pk_{O_1}) [U(N_1, \text{cred}_{N_1}), O(\text{NLog}_N)]$. In this protocol, U proves to O that N is his valid nym issued by O and that cred_{N_1} on the nym N_1 issued by O_1 .

Secure anonymous credential systems satisfy the following conditions (see [19, 6, 3] for more details): (1) *Unique User for Each Nym*. Even though the identity of a user who owns a nym must remain unknown, the owner should be unique. (2) *Unlinkability of Nyms*. Nyms of a user are not linkable at any time with a probability better than random guessing. (3) *Unforgeability of Credentials*. A credential may not be issued to a user without the organization's cooperation. (4) *Consistency of Credentials*. It is not possible for different users to team up and show some of their credentials to an organization and obtain a credential for one of them that the user alone would not have gotten. (5) *Non-Transferability*. Whenever Alice discloses some information that allows Bob to use her credentials or nyms, she is effectively disclosing her master secret key to him.

E-Cash. An e-cash system consists of three types of players: the *bank*, *users* and *merchants*. Below are the supported procedures (see [5]).

- $(pk_B, sk_B) \leftarrow \text{EC.BKeyGen}(1^k)$ is the key generation algorithm for the bank.
- $(pk_U, sk_U) \leftarrow \text{EC.UKeyGen}(1^k)$ is the key generation algorithm for users.
- $\langle W, \top \rangle \leftarrow \text{EC.Withdraw}(pk_B, pk_U, n) [U(sk_U), B(sk_B)]$. The user U withdraws a wallet W of n coins from the bank.
- $\langle W', (S, \pi) \rangle \leftarrow \text{EC.Spend}(pk_M, pk_B, n) [U(W), M(sk_M)]$. The user U spends a coin by giving it to the merchant M . U gets the updated wallet W' , and M obtains a coin (S, π) where S is a serial number and π is a proof.
- $\langle \top / \perp, L' \rangle \leftarrow \text{EC.Deposit}(pk_M, pk_B) [M(sk_M, S, \pi), B(sk_B, L)]$. M deposits (S, π) into its account in the bank B . L' is the updated list of the spent coins (i.e., (S, π) is added to the list).
- $(pk_U, \Pi_G) \leftarrow \text{EC.Identify}(S, \pi_1, \pi_2)$. Given two coins with the same serial number, i.e., (S, π_1) and (S, π_2) , B finds the identity of the double-spender pk_U and the corresponding proof Π_G .
- $\top / \perp \leftarrow \text{EC.VerifyGuilt}(S, pk_U, \Pi_G)$. It verifies the proof Π_G that the user pk_U is guilty of double-spending coin S .

Secure e-cash scheme satisfies the following condition: (1) *Correctness*. If an honest user runs EC.Withdraw with an honest bank, then neither will output

⁴ Usually, nym and pseudonym are used interchangeably. But to avoid confusion with the term pseudonym in our reputation scheme, we stick to the term nym in anonymous credential systems.

an error message. If an honest user runs EC.Spend with an honest merchant, then the merchant accepts the coin. (2) *Balance*. No collection of users and merchants can ever spend more coins than they withdrew. (3) *Identification of double-spenders*. Suppose the bank B is honest, and M_1 and M_2 are honest merchants who ran the EC.Spend protocol with the adversary whose public key is pk_U . Suppose the outputs of M_1 and M_2 are (S, π_1) and (S, π_2) respectively. This property guarantees that, with high probability, EC.Identify(S, π_1, π_2) outputs a key pk_U and proof Π_G such that EC.VerifyGuilt(S, pk_U, Π_G) accepts. (4) *Anonymity of users*. The bank, even when cooperating with any collection of malicious users and merchants, cannot learn anything about a user's spendings other than what is available from side information from the environment. (5) *Exculpability*. When S is a coin serial number not double-spent by user U with public key pk_U , the probability that EC.VerifyGuilt(S, Π_G, pk_U, n) accepts is negligible.

Blind Signatures. Blind signatures have two types of players: the *bank* and the *users*. A user requests the bank to generate a signature on a message m . Then the bank generates a signature without knowing the message m . Below are the supported procedures (see [14]).

- $(pk_B, sk_B) \leftarrow \text{BS.KeyGen}(1^k)$. Key-generation algorithm for the bank B .
- $\langle \top/\perp, \sigma/\perp \rangle \leftarrow \text{BS.Sign}(pk_B)[B(sk_B), U(m)]$. Signing protocol.
- $\top/\perp \leftarrow \text{BS.Verify}(m, \sigma, pk_B)$. Verification algorithm.

Secure blind signature scheme satisfies the following conditions: (1) *Unforgeability*. Only the bank who owns the secret key sk_B can generate valid signatures. (2) *Blindness*. The bank B does not learn any information about the message m on which it generates a signature σ .

5 Anonymous Identity-Bound Reputation System

In this section we describe a general scheme based on any implementation of the building blocks. See Appendix A for a specific instantiation of the scheme.

E-cash schemes will be used for the implementation of repcoins, blind signatures will be used in repcoin-withdraw and reputation-update procedures, and anonymous credential systems will be used for the reputation-demonstration procedures. As we shall see, while the first two are used in a relatively straightforward manner, the last one is used in a more complex way, since the reputation demonstration setting presents a new type of hurdle to overcome if unlinkability is to be achieved even against colluding bank and peers.

Underlying Protocols and Requirements. Our scheme will work with any implementation of these underlying primitives, as long as the master public and

secret keys for peers in our system are of the same form as those in the underlying e-cash scheme and anonymous credential system. That is, the key generation algorithms Ukeygen, EC.UKeyGen, and AC.Ukeygen are all the same.⁵

Our scheme will also require a zero knowledge proof of knowledge of both the master secret key corresponding to a master public key, and the secret information of a nym’s owner (which is given as an output of the AC.FormNym operation). Thus, when instantiating our scheme with specific primitives, it is useful to choose underlying primitives that admit efficient proofs of this form (as we do in the Appendix A).

Setup. We start with the setup procedure on Bank’s side.

- Bank B executes EC.BKeyGen procedure of e-cash scheme to create a digital signature key-pair (pk_B, sk_B) . This is the key-pair that will be used for creating the repcoins. Bank publishes pk_B .
- B executes BS.BkeyGen procedure of blind signatures scheme to create a blind signature key pair to be used in the Reputation Deposit procedure (pk_B^b, sk_B^b) . Bank publishes pk_B^b .
- B defines fixed reputation levels l_i , represented by a group G_i . These “reputation” groups — although managed by Bank — play a role similar to the one organizations play in anonymous credential systems. For each one of these groups, Bank runs AC.OKeyGen protocol to generate public-secret key pairs (pk_{G_i}, sk_{G_i}) . Bank also publishes pk_{G_i} s.
- B does the appropriate setup (if any) for the pseudonym generation. For example, this may involve selecting an appropriate algebraic group G_p .

On the peers’ side, each peer U_i invokes EC.UKeyGen to create a master public-secret keypair (pk_{U_i}, sk_{U_i}) .

Operations. As mentioned, we assume that messages are exchanged through perfectly secure channels. The system operations are realized as follows.

1. Generation of Pseudonyms. Each peer generates his own pseudonyms. There is no particular structure imposed on the pseudonyms, and they need not be certified or registered with Bank (or any other entity). The only requirement is that the pseudonym generation leaves the owner with some secret information (e.g., the random string used for the generation procedure), such that possession of this information proves ownership of the pseudonym. We will also need such a proof to be executed. Thus, in principle, we can simply use a random string r as the secret information and $P = f(r)$ as the pseudonym, where f is some

⁵ As discussed in Section 2, an important part our system setup is the assumption that peers are motivated to keep their master private key secret. For this reason, it is beneficial to have the master public and private keys be part of an external PKI which is used for other purposes (e.g., signing documents) outside our system.

one-way function, with an associated zero-knowledge proof of knowledge of the inverse of P . However, a more efficient solution is to let the pseudonym generation procedure to be a digital signature key generation, keeping the signing key as the secret information and the verification key as the pseudonym. Here, being able to produce valid signatures will prove ownership of the pseudonym, without a need for a zero-knowledge proof.

2. *RepCoin Withdrawal.* RepCoin Withdrawal takes place between Bank B and a peer U . Both U and B engage in EC.Withdraw procedure of a e-cash scheme. For simplicity purposes, we assume that a wallet W of n repcoins has been withdrawn. Since the only properties related to repcoins are anonymity of an honest withdrawer and repudiation of any double spender, the wallet can be like the one suggested in [5], or n separate digital coins withdrawn through any known e-cash scheme.

3. *Reputation Award.* This procedure is executed between two pseudonyms, one (i.e., P_U) belonging to a peer U and one (i.e., P_M) belonging to a peer M . Both engage in EC.Spend protocol of a e-cash scheme. However, this protocol takes place strictly between the two pseudonyms P_U and P_M instead of involving the actual identities U and M . Thus, P_U gives a repcoin to P_M , where no information about identities of the parties involved is revealed.

4. *Reputation Update.* This protocol is invoked when a peer M wants to increase his reputation based on the repcoins that his pseudonyms have received since the last time he updated his reputation record. As previously discussed, maintaining unlinkability between a pseudonym and its owner is a crucial feature of our system. Towards this end, a single interaction for update (with a merchant presenting himself to Bank either as a peer or as a pseudonym) will not work, as we explain below.

Assume peer M wants to deposit a repcoin he received as P_M from pseudonym P_U of User U . Note that no one except M knows who is the owner of P_M . Given the fact that U knows the exact form of the repcoin he gave to M , if M tried to deposit the repcoin by presenting himself as M to Bank, a collusion of Bank and U would reveal that M is the owner of P_M . Trying to solve this by letting M “rerandomize” the repcoin in some way before depositing it presents problems for enforcing the no over-awarding requirement. On the other hand, if Reputation Update procedure was done by the pseudonym P_M of M , there would be a problem in persuading the Bank to update M ’s record without revealing that M is the owner of P_M .

Therefore, our Reputation Update protocol has two stages. First, P_M contacts Bank and gets a blind permission from it that shows a repcoin has been deposited and is valid. Second, M deposits that blind permission. In particular, the following procedure takes place:

4.1 Obtaining Blind Permission. Peer M executes EC.Deposit procedure of e-cash scheme using his pseudonym P_M , but here the actual deposit does not happen. Rather, if Bank B accepts the repcoin, M gets from B a blind signature on a random message. That is, P_M sends to B a repcoin that it has received. If B accepts the coin as valid, P_M chooses a random message C and gets a blind signature of C : σ_B^b . We call (C, σ_B^b) a *blind permission*.

4.2 Deposit of the Blind Permission. M sends B the permission (C, σ_B^b) . Then, B checks if the tuple is fresh and increases the reputation of M .

5. *Reputation Demonstration.* This protocol is invoked when one peer wants to demonstrate his reputation to another peer, both interacting strictly through their pseudonyms. We will utilize predefined groups G_i corresponding to reputation levels l_i , which are managed by Bank. For a peer U who wants, via P_U , to prove his reputation level l_i to a pseudonym P_V of a peer-verifier V , the protocol proceeds as follows:

- If he has not done it before, U contacts the bank to register in the group G_i that corresponds to the desired reputation level l_i . U interacts with G_i (Bank) by invoking AC.FormNym protocol of an anonymous credential system, in order to generate a nym $N_U^{l_i}$ for U under that group.⁶ (U can generate as many nyms as he wants.)
- U contacts G_i , providing its master public pk_U key and a zero knowledge proof of knowledge π that he possesses the corresponding master secret key sk_U . U also presents $N_U^{l_i}$ and a zero-knowledge proof π_N that it has been created correctly and he is the owner.
- G_i checks that U is valid and that his reputation is indeed in that group (or higher), and executes AC.GrantCred to generate a credential $C_N^{l_i}$ for $N_U^{l_i}$.
- U interacts with the verifier P_V under his pseudonym P_U . P_U proves by executing AC.VerifyCred that he possesses a credential from group G_i . Specifically, P_U proves that its owner has registered under a nym to G_i and has acquired — through that nym — a credential of membership.

5.1 Security

The following theorem states the correctness and security of our general scheme. For lack of space, we refer the reader to our technical report [1] for proofs.

⁶ Recall that there is a big difference between pseudonyms and nyms. As discussed before, Pseudonyms are public-secret key-pairs, used as means to preserve peers' anonymity when involved in transactions. A nym of a peer will be associated with a particular reputation group. Bank, as the manager of the reputation groups, will be able to link the nyms with the peer identities (master public key). In contrast, unlinkability of peers and pseudonyms is maintained, as per our security definitions.

Theorem 1. *If the underlying primitives (anonymous credential system, e-cash system, and blind signatures) are secure, then our scheme satisfies correctness, peer-pseudonym unlinkability, pseudonym-pseudonym unlinkability, no over-awarding, exculpability, and reputation unforgeability.*

5.2 Practical Issues

In the absence of a concrete implementation, it is hard to make concrete statements about practical issues. Furthermore, our main result is a framework which can accommodate different algorithms. That said, there are at least two areas that deserve further attention, performance and system security.

In general, our protocol is neither real-time nor high-performance. We are not proposing per-packet operations; most of what we do is per-user or per-purchase. As such, real-time performance is not critical; in particular, there are no bottleneck nodes.

A full performance analysis is given in [1]. Here, we note that all of our primitive operations are $O(1)$ in system size. That is, there are no design elements in our scheme whose performance degrades as the size of the system increases. Similarly, no operation takes more than a few messages; all are $O(k + w)$ in message size, where k and w are security parameters. More details for our specific instantiation are given in Appendix A.

In addition to the anonymous peer-to-peer communication necessary for the underlying application, there is now a new communications path: from each party to the bank. Parties who are engaging in our protocol will need to contact the bank. This provides another channel that might be detected by, say, the attacks described in [15]. Indeed, there may exist a sort of “meta-intersection attack” [9]: the peer-to-peer traffic alone may not be suspicious, but it when coupled with conversations with the bank might be sufficient for identification.

A second area for security concern is CPU consumption. Our scheme (see Appendix A) requires public key operations; these are CPU-intensive. An attacker who has identified a candidate participant in real-time might be able to connect to it — we are, after all, talking about peer-to-peer systems — and measure how long its own communications take. The obvious defense is to make sure that any given operation takes constant time; in turn, this likely means pre-configuring each peer node with a maximum number of concurrent connections supported.

6 Related Work

A number of papers have addressed the issue of reputation and privacy.

There are many papers on reputation systems for peer-to-peer networks. Most focus on building distributed reputation systems, rather than worrying about privacy; [12] is typical.

The difficulty of building systems like this is outlined by Dingedine, Mathewson, and Syverson [10]. They present a number of similar systems and show why bolting on reputation is hard.

A typical approach is typified by [26], who incorporate privacy into their scheme. However, their system does not provide unlinkability. It also requires a trusted “observer” module for full functionality.

The work by Kinatader et al. [16, 18] is close to ours. The system in [16] differs from ours in two notable ways. First, its reputations are linkable. Indeed, they see this as a virtue, in that recommendations can be weighted depending on the reputation of the recommender. Second, they assume a trusted hardware module (i.e., a TPM chip) on every endpoint. In [18], they describe a more general system based on UniTEC [17]. Reputation statements are signed by a pseudonym’s private key. Unlinkability is achieved by switching public keys. Apparently, the UniTEC layer can share reputations between different pseudonyms, but the authors do not explain how this is done. Presumably, this is handled by bookkeeping at that layer. More seriously, although they assert that a trusted module is desirable but not necessary, they do not explain how that could work, and in particular how they can prevent cheating.

Pavlov et al. [22] present a system, based on secret-sharing, which has many of the same properties as ours. However, it depends on locating “witnesses”, other parties with knowledge of the target’s reputation. In a sufficiently-large community with a low density of interaction, this may be difficult. Furthermore, it does not provide unlinkability; witness testify about a known party’s past behavior.

Another work related to ours is Voss [25] and Steinbrecher [23]. In both of the systems, users interact with each other through pseudonyms, and reputation is strongly connected to identities. In fact, in [25] reputation points are implemented as coins, which may have positive or negative value. However, in both cases, Trusted Third Parties⁷ are required to ensure unlinkability between identities and pseudonyms.

Approaches other than reputation systems have also been presented to deal with misbehaving users in anonymous or pseudonymous systems. Belenkiy et al. [2] make use of endorsed e-cash to achieve fair and anonymous two-party protocol wherein parties buy or barter blocks of data. Whereas e-cash stands for reputation in our scheme, e-cash stands for actual money in their scheme; a peer uses e-cash to buy data from other peers. Johnson et al. [13] focus on protecting

⁷ In [23] TTP appear in the form of designated identity providers.

a service in Tor from a malicious user without blocking all the exit Tor nodes. In particular, they present a protocol where misbehaving anonymous users are blacklisted by servers.

7 Future Directions

A few interesting open problems remain.

First, our current scheme uses unit coins for reputation. That is, all reputation credits are worth the same amount. It would be nice to permit variable values; we suspect that this is easy.

More seriously, we do not have negative feedback. There is a vast difference between knowing that a seller has performed well on m transactions and knowing that that seller has performed well on m out of n . The difficulty is forcing the seller to commit to depositing a coin indicating bad behavior; most sellers know when they have done something wrong. In the technical report [1], we developed a partial solution. The scheme does not satisfy the complete unlinkability requirement stipulated in our definition, as Bank knows the number of transactions a peer had interacted in as a seller (modulo this information being leaked, all anonymity requirements are preserved).

Finally, we would like to get rid of the bank, which in our scheme is trusted to maintain reputation balances correctly (though not trusted from the privacy perspective). A fully decentralized scheme would eliminate single points of failure, and would be more in keeping with a widespread, anonymous, peer-to-peer network. Note that this would require two significant changes: using a digital cash scheme that does not require a central bank, and devising some other mechanism for inflation resistance.

Acknowledgment

We are grateful to Moti Yung for useful discussions regarding this work. We would like to thank Patrick Tsang, Apu Kapadia, and anonymous referees for helpful comments on the paper.

References

1. Elli Androulaki, Seung Geol Choi, Steven M. Bellovin, and Tal Malkin. Reputation systems for anonymous networks. Technical Report CUCS-029-07, Computer Science Dept., Columbia University, 2007. <http://www.cs.columbia.edu/research/publications>.
2. Mira Belenkiy, Melissa Chase, C. Christopher Erway, John Jannotti, Alptekin Küpçü, Anna Lysyanskaya, and Eric Rachlin. Making p2p accountable without losing privacy. In *WPES*, pages 31–40, 2007.

3. Mira Belenkiy, Melissa Chase, Markulf Kohlweiss, and Anna Lysyanskaya. P-signatures and noninteractive anonymous credentials. In *TCC*, pages 356–374, 2008.
4. Rajat Bhattacharjee and Ashish Goel. Avoiding ballot stuffing in ebay-like reputation systems. In *P2PECON*, pages 133–137, 2005.
5. Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Compact e-cash. In *EUROCRYPT*, pages 302–321, 2005.
6. Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT*, pages 93–118, 2001.
7. Jan Camenisch and Markus Stadler. Efficient group signature schemes for large groups. In *CRYPTO*, pages 410–424, 1997.
8. David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–88, 1981.
9. George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *Information Hiding*, pages 293–308, 2004.
10. Roger Dingledine, Nick Mathewson, and Paul Syverson. Reputation in p2p anonymity systems. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.
11. Roger Dingledine, Nick Mathewson, and Paul F. Syverson. Tor: The second-generation onion router. In *USENIX Security Symposium*, pages 303–320, 2004.
12. Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *NOSSDAV*, 2003.
13. Peter C. Johnson, Apu Kapadia, Patrick P. Tsang, and Sean W. Smith. Nymble: Anonymous ip-address blocking. In *Privacy Enhancing Technologies*, pages 113–133, 2007.
14. Ari Juels, Michael Luby, and Rafail Ostrovsky. Security of blind digital signatures (extended abstract). In *CRYPTO*, pages 150–164, 1997.
15. Dogan Kesdogan, Dakshi Agrawal, Vinh Pham, and Dieter Rautenbach. Fundamental limits on the anonymity provided by the mix technique. In *S&P*, pages 86–99, 2006.
16. Michael Kinateder and Siani Pearson. A privacy-enhanced peer-to-peer reputation system. In *EC-Web*, pages 206–215, 2003.
17. Michael Kinateder and Kurt Rothermel. Architecture and algorithms for a distributed reputation system. In *iTrust*, pages 1–16, 2003.
18. Michael Kinateder, Ralf Terdic, and Kurt Rothermel. Strong pseudonymous communication for peer-to-peer reputation systems. In *SAC*, pages 1570–1576, 2005.
19. Anna Lysyanskaya, Ronald Rivest, Amit Sahai, and Stefan Wolf. Pseudonym systems. In *SAC*, pages 184–199, 1999.
20. Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In *CRYPTO*, pages 31–53, 1992.
21. Lasse Øverlier and Paul F. Syverson. Locating hidden servers. In *S&P*, pages 100–114, 2006.
22. Elan Pavlov, Jeffrey S. Rosenschein, and Zvi Topol. Supporting privacy in decentralized additive reputation systems. In *iTrust*, pages 108–119, 2004.
23. Sandra Steinbrecher. Design options for privacy-respecting reputation systems within centralised internet communities. In *SEC*, pages 123–134, 2006.
24. Paul F. Syverson, David M. Goldschlag, and Michael G. Reed. Anonymous connections and onion routing. In *IEEE Symposium on Security and Privacy*, pages 44–54, 1997.
25. Marco Voss. Privacy preserving online reputation systems. In *International Information Security Workshops*, pages 245–260, 2004.
26. Marco Voss, Andreas Heinemann, and Max Muhlhauser. A privacy preserving reputation system for mobile information dissemination networks. In *SECURECOMM*, pages 171–181, 2005.

A An Example of Scheme Instantiation

In this section we give a specific instantiation of our scheme, where we make use of the anonymous credential system by Camenisch and Lysyanskaya [6] (denoted by CL), the e-cash scheme by Camenisch et al. [5] (denoted by CHL), and the blind signature scheme by Okamoto [20] (denoted by Ok). We do so in order to present a concrete and efficient construction (we include the efficiency analysis, relying on that of the underlying primitives, with each of the operations).

Setup(1^k).

Bank B does the setup as follows:

- B executes CHL.BKeygen(1^k) to generate an e-cash key pair (pk_B^{ec}, sk_B^{ec}) , and publishes $pk_B^{ec} = (g_{ec}, \hat{g}_{ec}, \tilde{g}_{ec})$.
- B executes Ok.KeyGen(1^k) to generate a blind signature key pair (pk_B^{bs}, sk_B^{bs}) and publishes pk_B^{bs} .
- For each reputation group G_i ($1 \leq i \leq k$), B executes CL.OKeYGen(1^k) to generate the anonymous credential system key pair $(pk_B^{ac_i}, sk_B^{ac_i})$ for G_i , and publishes $pk_B^{ac_i} = (n_{ac_i}, a_{ac_i}, b_{ac_i}, d_{ac_i}, g_{ac_i}, h_{ac_i})$.
- B creates a cyclic group $G_p = \langle g_p \rangle$ of order $p = \Theta(2^k)$ where the DDH assumption holds. This algebraic group is used for pseudonym generation on the peer's side.

On the peers' side, each peer U executes CHL.UKeyGen(1^k) to obtain $(pk_U, sk_U) = (g_{ec}^{x_U}, x_U)$, and publishes pk_U . Note that x_U will be used as the master secret key of U in the anonymous credential system (and this discrete-log based key is a reasonable choice for a more general PKI key as well).

Operations.

1. Generation of Pseudonyms. Each peer generates his pseudonyms locally using \mathbb{G}_p . Specifically, he chooses a random number $r_i \in \mathbb{Z}_p$ and compute $g_p^{r_i}$. The value $g_p^{r_i}$ is considered a pseudonym P_U^i of peer U .

2. RepCoin Withdrawal. A peer U executes CHL.Withdraw with Bank, and obtains a wallet W of 2^w repcoins. This procedure takes $O(1)$ exponentiations and $O(1)$ rounds.

3. Reputation Award. A pseudonym P_U gives a repcoin to P_M by executing CHL.Spend with P_M . This procedure also takes $O(1)$ exponentiations and $O(1)$ rounds.

4. Reputation Update.

4.1 Obtaining Blind Permission. A pseudonym P_M and Bank B participate in CHL.Deposit protocol, which takes $O(1)$ exponentiations and $O(1)$

rounds. If CHL.Deposit accepts, P_M acquires the blind permission $\sigma_B^{\text{bs}} = \text{Ok.Sig}n(sk_B^{\text{bs}}, r_{\text{perm}})$ where r_{perm} is a random message. Obtaining the blind permission takes $O(1)$ exponentiations and $O(1)$ rounds.

4.2 Deposit of the Blind Permission. M (the owner of P_M) sends σ_B^{bs} to B . B checks if the permission $(r_{\text{perm}}, \sigma_B^{\text{bs}})$ is fresh; if so, it increases M 's reputation value. This procedure takes $O(1)$ exponentiations and $O(1)$ rounds.

5. *Reputation Demonstration.* Suppose that a pseudonym P_U asks P_M to demonstrate its reputation level, and that M (the owner of P_M) wants to show to P_U that it belongs to G_i , i.e., his reputation is at least at level l_i .

- Obtaining a nym under G_i . M contacts Bank B and executes CL.FormNym with respect to G_i ⁸. Let $N_M^{l_i}$ be the nym that M obtained from this procedure. Note that $N_M^{l_i}$ is of the form: $g_{\text{ac}_i}^{x_U} \cdot h_{\text{ac}_i}^r$. This takes $O(1)$ exponentiations and $O(1)$ rounds.
- Obtaining a credential for G_i . M contacts B , and he sends B the message $(pk_M, N_M^{l_i})$. Then, M executes with B a zero-knowledge proof of knowledge

$$PK\{(\alpha, \beta) : pk_M = g_{\text{ec}}^\alpha, N_M^{l_i} = g_{\text{ac}_i}^\alpha \cdot h_{h_i}^\beta\}^9$$

This takes $O(1)$ exponentiations and $O(1)$ rounds.

Now, B verifies the proof. If the proof is verified so that M is eligible for a credential of the group G_i , B executes the CL.GrantCred (protocol4) with respect to G_i . Let C_{l_i} be the output credential. This takes $O(1)$ exponentiations and $O(1)$ rounds.

- Showing reputation using the credential. P_M contacts P_U and executes CL.VerifyCred (protocol3) with respect to G_i to prove that owner of P_M has a credential for the group G_i . This takes $O(1)$ exponentiations and $O(1)$ rounds.

⁸ We use both protocol1 and protocol6 of [6] instead of just protocol1 to ensure the non-transferability of credentials.

⁹ This proof can be parsed as “I know the exponent α and β that was used in generating pk_M and $N_M^{l_i}$ ”. See [7, 6] for more detail. The proof can be regarded as an authentication procedure.